

Convergence of Kikuchi matrices to Γ -independent and q -Gaussian limits

Afonso S. Bandeira* Dmitriy Kunisky† Petar Nizic-Nikolac‡

Lucas Pesenti§ Robert Wang¶

October 2, 2026

Abstract

Kikuchi matrices are a family of structured matrices that were introduced to study problems involving tensors and hypergraphs. We show that, as the ambient dimension grows, dense random Kikuchi matrices have a limit described by a system of Γ -independent semicircular elements. This characterizes their limiting spectral distribution and yields improved bounds on their spectral norm, a key quantity in the analysis of algorithms for Tensor PCA. Finally, we show that, in an appropriate double limit, independent Kikuchi matrices converge to the q -Gaussian system, another central object in noncommutative probability.

1 Introduction

Kikuchi matrices are a family of structured matrices that were recently introduced to reduce problems involving higher-degree polynomials, hypergraphs, or tensors to more tractable quadratic, graph, or matrix problems. For many applications, the bottleneck in this approach lies in precisely estimating the spectral norm of Kikuchi matrices as the ambient dimension grows. In this paper, we study the *limit* of Kikuchi matrices, focusing on the dense random Kikuchi matrices that arise in the analysis of algorithms for tensor PCA [WEM19, Has20, BCSvH26, Man24, KX26].

Given integers n, ℓ, r satisfying $1 \leq r \leq \ell \leq n - r$, the *Kikuchi matrix* $Y_n = Y_n(\ell, r)$ is the random matrix with entries

$$Y_n[U, V] := \begin{cases} \frac{1}{\sqrt{d_n}} g_{U \Delta V} & \text{if } |U \Delta V| = 2r \\ 0 & \text{otherwise} \end{cases} \quad \text{for } U, V \in \binom{[n]}{\ell}, \quad (1)$$

where g_I are independent standard Gaussian random variables for all $I \subseteq [n]$ of size $|I| = 2r$. Here, $U \Delta V = (U \cup V) \setminus (U \cap V)$ denotes the symmetric difference between sets, and

$$d_n := d_n(\ell, r) = \binom{\ell}{r} \binom{n - \ell}{r} \quad (2)$$

*Department of Mathematics, ETH Zürich. bandeira@math.ethz.ch

†Department of Applied Mathematics & Statistics, Johns Hopkins University. kunisky@jhu.edu

‡Department of Mathematics, ETH Zürich. petar.nizic-nikolac@ifor.math.ethz.ch

§Department of Mathematics, ETH Zürich. lpesenti@ethz.ch

¶Cheriton School of Computer Science, University of Waterloo. robert.wang2@uwaterloo.ca

is the normalization that makes the expected squared ℓ_2 -norm of each row of the Kikuchi matrix equal to 1. Our goal is to describe the limit as $n \rightarrow \infty$ of the spectrum and norm of Y_n for fixed ℓ, r .

Henceforth and unless stated otherwise, asymptotics are with respect to $n \rightarrow \infty$ for fixed ℓ, r . Similarly, “with high probability” should be understood as with probability tending to 1 as $n \rightarrow \infty$.

Tensor PCA. One motivation for understanding the spectrum of the Kikuchi matrix is the *tensor PCA* problem [MR14]. In one of its idealized forms, we aim to recover a vector of spins $x \in \{-1, 1\}^n$ from order- $2r$ tensor measurements. For each set $I \subseteq [n]$ of size $|I| = 2r$, we observe

$$\mathcal{T}_I = \lambda_n \prod_{i \in I} x_i + g_I,$$

where the g_I are i.i.d. standard Gaussians and $\lambda_n > 0$ is the signal strength. To apply the “Kikuchi matrix method”, given these observations, for any $r \leq \ell \leq n - r$, one can form the matrix indexed by $U, V \in \binom{[n]}{\ell}$ with entries

$$M_n[U, V] := \begin{cases} \mathcal{T}_{U \Delta V} & \text{if } |U \Delta V| = 2r, \\ 0 & \text{otherwise.} \end{cases}$$

When λ_n is large enough, M_n has an outlier eigenvector that can be used to recover x up to a global sign. To understand the spectrum of M_n , we may assume without loss of generality that x is the all-ones vector. Then $M_n = \lambda_n A_n + \sqrt{d_n} Y_n$, where A_n is the adjacency matrix of a regular graph (the Johnson graph) whose degree is d_n defined in (2), and Y_n is the Kikuchi matrix defined in (1). We therefore expect the signal to be visible in the spectrum of M_n whenever $\lambda_n d_n > \sqrt{d_n} \|Y_n\|$; the analysis reduces to estimating the norm of Y_n .¹ We point the reader to [BCSvH26] for a sharp derivation of the detection threshold of tensor PCA using this construction when $\ell < \frac{3}{2}r$.²

Prior work. We first review previous bounds on the norm of Kikuchi matrices. A simple application of the classical noncommutative Khintchine inequality [LP91] gives, with high probability,

$$\|Y_n\| \lesssim \sqrt{\log \binom{n}{\ell}} \lesssim \sqrt{\ell \log n}. \quad (3)$$

This bound is used, for instance, in the first application of Kikuchi matrices to tensor PCA [WEM19]. However, in many applications of this inequality, the logarithmic factor is known to be an artifact of the analysis and can be removed with more careful or application-specific calculations. To address this issue, a series of works [BBvH23, BCSvH26, Ban26] developed improved noncommutative Khintchine inequalities based on ideas from free probability. In particular, when specialized to Kikuchi matrices, these inequalities show that, in the regime $r \leq \ell < 2r$, the spectrum of Y_n is asymptotically semicircular and $\|Y_n\| = 2 + o(1)$ with high probability [BCSvH26]. However, these techniques break down once $\ell \geq 2r$, a regime in which the Kikuchi matrix does not asymptotically have a semicircular distribution of eigenvalues. This was shown formally by Kothari and Xu [KX26]: using the trace method and a

¹Here we allow for a constant factor of inaccuracy in understanding the size of λ_n necessary to create an outlier; understanding the exact threshold would raise the more complicated question of Baik-Ben Arous-Péché (BBP) transitions with a Kikuchi noise matrix, which we do not consider here.

²The asymptotics of (2) suggested an intriguing phenomenon for Tensor PCA: the detection threshold decreases as $\ell \rightarrow \infty$. This was the subject of [BKMR25, Conjecture 9], subsequently resolved in [KX26]. As a direct consequence of Theorem 1.6 below, we further improve the dependence on r to at most a polynomial factor away from optimal.

delicate charging argument, they established that instead we have $\|Y_n\| \asymp_r \sqrt{\ell}$ in a large regime of $\ell \leq n^{\delta(r)}$, up to constants depending only on r .

In this paper, we initiate a program to determine the *exact* asymptotic spectrum and spectral norm of Kikuchi matrices, for arbitrary fixed ℓ and r .

Weak and strong convergence. The limiting object associated with a single Kikuchi matrix is a bounded linear operator $A : \mathcal{H} \rightarrow \mathcal{H}$ on an infinite-dimensional Hilbert space $\mathcal{H}$. We refer the reader to Section 2.1 for our notation for operators on Hilbert spaces. Our results are formulated more generally for tuples of independent Kikuchi matrices, in terms of *weak* and *strong* convergence.

Definition 1.1. Fix $k \geq 1$. Let $(\mathcal{H}_n)_{n \geq 1}$ and $\mathcal{H}$ be Hilbert spaces, with states τ_n on $\mathcal{B}(\mathcal{H}_n)$ and τ on $\mathcal{B}(\mathcal{H})$. For each $n \geq 1$, let $\mathbf{A}_n = (A_{n,1}, \dots, A_{n,k})$ be a tuple of random self-adjoint elements of $\mathcal{B}(\mathcal{H}_n)$, and let $\mathbf{A} = (A_1, \dots, A_k)$ be a tuple of self-adjoint elements of $\mathcal{B}(\mathcal{H})$. We say that

1. $\mathbf{A}_n$ converges to $\mathbf{A}$ *weakly* if

$$\tau_n(P(\mathbf{A}_n)) \xrightarrow{n \rightarrow \infty} \tau(P(\mathbf{A})) \quad \text{in probability,}$$

for every noncommutative polynomial $P \in \mathbb{C}\langle X_1, \dots, X_k \rangle$.

2. $\mathbf{A}_n$ converges to $\mathbf{A}$ *strongly* if $\mathbf{A}_n$ converges to $\mathbf{A}$ weakly, and

$$\|P(\mathbf{A}_n)\| \xrightarrow{n \rightarrow \infty} \|P(\mathbf{A})\| \quad \text{in probability,}$$

for every noncommutative polynomial $P \in \mathbb{C}\langle X_1, \dots, X_k \rangle$.

In the special case where $\mathcal{H}_n$ is finite-dimensional, the random operators $A_{n,i}$ are random matrices, and we take $\tau_n = \text{tr}$ to be the normalized trace (the usual trace divided by the dimension of the matrix), while $\|\cdot\|$ is the usual spectral norm. For a single matrix, weak convergence corresponds to convergence of the spectral distribution, while strong convergence additionally captures convergence of the spectral norm. Powerful new techniques for establishing strong convergence have emerged in recent years [HT05, CGTvH26, CGvH26, vH26].

Γ -independent systems. Next, we define the candidate limiting operator. It is built from a family of operators with semicircular distributions whose joint distribution satisfies a notion of independence known as Γ -*independence*. This notion generalizes both free and classical independence. The construction below is due to Bożejko and Speicher [BS94], with a simplified formulation due to Oliveira Santos, Tripathi, and Youssef [OTY25].

Definition 1.2. Let $\Gamma = (V, E)$ be an undirected graph. The set of words $\bigsqcup_{k \geq 0} V^k$ on the alphabet V has a monoid structure given by concatenation. We denote by $w \cdot w'$ the concatenation of two words w, w' . The *trace monoid* $\mathcal{T} = \mathcal{T}(\Gamma)$ is the quotient monoid obtained by identifying two words whenever one can be obtained from the other by repeatedly interchanging consecutive letters that are adjacent in Γ . We denote by $\emptyset$ the element of the trace monoid corresponding to the empty word.

For example, if Γ is the graph on $V = \{a, b, c\}$ with a single edge between a and b , then we have $a \cdot b \cdot c = b \cdot a \cdot c$ in the trace monoid, whereas $a \cdot b \cdot c \neq a \cdot c \cdot b$. At the two extremes, the trace monoid is the free monoid when Γ has no edges, and the free commutative monoid when Γ is complete.

Definition 1.3. Let $\mathcal{T}$ be the trace monoid associated with the graph $\Gamma = (V, E)$. Consider the complex Hilbert space $\ell^2(\mathcal{T})$ with canonical basis $\{\delta_w : w \in \mathcal{T}\}$. We equip $\mathcal{B}(\ell^2(\mathcal{T}))$ with the state $\tau(L) := \langle \delta_\emptyset, L\delta_\emptyset \rangle$. For each $u \in V$, let $L_u \in \mathcal{B}(\ell^2(\mathcal{T}))$ be the *creation operator*, which is the linear map such that $L_u(\delta_w) = \delta_{w \cdot u}$. We call $(L_u + L_u^*)_{u \in V}$ a Γ -independent semicircular system.

Here, L_u^* denotes the adjoint of the creation operator L_u . On a basis element, the effect of L_u^* is to remove an occurrence of u on the right of the word, if possible, and to map to zero otherwise:

$$L_u^*(\delta_w) = \begin{cases} \delta_{w'} & \text{if } w = w' \cdot u \text{ for some } w' \in \mathcal{T}, \\ 0 & \text{otherwise.} \end{cases}$$

The factorization $w = w' \cdot u$ is understood as an equality in the trace monoid (in particular, such a w' , when it exists, is unique). The adjoint of a creation operator is often called an *annihilation operator*.

In a Γ -independent semicircular system $(S_u)_{u \in V}$, each S_u has a semicircular distribution, in the sense that, for each $p \geq 1$, $\tau(S_u^{2p})$ is the p -th Catalan number and $\tau(S_u^{2p+1}) = 0$. Furthermore, S_u and S_v are classically independent when u and v are adjacent in Γ , and freely independent otherwise. The generalized notion of independence satisfied by these operators is called Γ -independence [SW16, CGvH26].³ See Lemma 2.8 for a general formula for computing the joint moments of a Γ -independent semicircular system.

The graphs Γ of interest in this paper are the Kneser graphs.

Definition 1.4. For $\ell, r \geq 1$, the *Kneser graph* $K(\ell, r)$ is the graph on the vertex set $\binom{[\ell]}{r}$, with an edge between a and b when $a \cap b = \emptyset$. For $k \geq 1$, let $K^{(k)}(\ell, r)$ denote the graph with vertex set $[k] \times \binom{[\ell]}{r}$, where (i, a) and (j, b) are adjacent when $a \cap b = \emptyset$.

1.1 Our results

The Γ -independent limit. In this paper, we identify the limit of Kikuchi matrices: it is given by a Γ -independent semicircular system, where Γ is a Kneser graph. Our first main result establishes weak convergence of tuples of independent Kikuchi matrices.

Theorem 1.5. Fix $k \geq 1$ and $1 \leq r \leq \ell$. Let $\{S_{i,a} : i \in [k], a \in \binom{[\ell]}{r}\}$ be a $K^{(k)}(\ell, r)$ -independent semicircular system. Define

$$\Sigma_i = \Sigma_i(\ell, r) := \binom{\ell}{r}^{-1/2} \sum_{a \in \binom{[\ell]}{r}} S_{i,a} \quad \forall i \in [k]. \quad (4)$$

For each $n \geq \ell + r$, let $\mathbf{Y}_n = (Y_{n,1}, \dots, Y_{n,k})$ be a k -tuple of independent copies of the Kikuchi matrix $Y_n(\ell, r)$ defined in (1). Then $\mathbf{Y}_n$ converges weakly as $n \rightarrow \infty$ to $\Sigma = (\Sigma_1, \dots, \Sigma_k)$.

This determines the asymptotic spectral distribution of Kikuchi matrices. Its moments have a simple explicit form; see Lemma 2.8. As a special case, when $k = 1$ and $\ell < 2r$, the Kneser graph $K(\ell, r)$ has no edges, so the limit is semicircular, recovering the result of [BCSvH26].

The basic reason to expect the Kneser graph to be involved is as follows. In the combinatorics of tracial moments of a Kikuchi matrix, we sum over closed walks of a certain length on the *Kikuchi*

³The same notion has also been referred to as ε -free independence [CC21], Λ -freeness [Mio04], or Γ -right-angled independence [MT26] in prior works.

graph given by the pattern of non-zero entries. Such walks can be viewed as moving around on $\binom{[n]}{\ell}$ by replacing r indices of a subset at a time. The leading-order contributions to the tracial moments come from walks in which each Gaussian variable (i.e., each total set of $2r$ indices affected by a step) occurs exactly twice, and such that the first step of each such pair introduces r new indices that have not yet been encountered on the walk. Further, these replacements can be viewed as occurring in ℓ labeled “slots” of an element of $\binom{[n]}{\ell}$, whose labels persist across replacements of indices. In total, we then have a sequence of replacements of r indices, ordered in a cycle, each replacement occurring in some r slots out of ℓ , where each replacement is matched to another that occurs on the same set of slots. The main point is that two matches can only “cross” along the cycle if they occur on disjoint sets of slots. This is precisely the adjacency relation in the Kneser graph, whereby the associated Γ -independent system has moments matching the leading-order contributions to tracial moments. See also Section 1.2 for another interpretation of the limit in terms of intrinsic freeness.

Given Theorem 1.5, a natural follow-up question is:

*Under which assumptions on ℓ, r can weak convergence be upgraded to strong convergence?
In particular, when does $\|Y_n\| \rightarrow \|\Sigma\|$?*

We are unable to answer this question in full, but we prove several partial results. Note that weak convergence implies $\|Y_n\| \geq \|\Sigma\| - o(1)$ with high probability, so the question is whether Y_n can asymptotically have outlier eigenvalues lying beyond the support of its limiting spectral distribution.

To put our results in context, we first estimate $\|\Sigma\|$ as a function of ℓ and r . Counting the contribution of words adding $p \approx \ell/r^2$ pairwise disjoint subsets of $[\ell]$ of size r and removing them in an arbitrary order gives the lower bound

$$\|\Sigma\| \geq \tau(\Sigma^{2p})^{\frac{1}{2p}} \geq \left(p! \prod_{i=0}^{p-1} \frac{\binom{\ell-ir}{r}}{\binom{\ell}{r}} \right)^{\frac{1}{2p}} \gtrsim \frac{\sqrt{\ell}}{r}. \quad (5)$$

We later show (Proposition 6.7) that in fact, $\|\Sigma\| \asymp \max\{1, \sqrt{\ell}/r\}$. This is compatible with the recent result of Kothari and Xu [KX26], which confirms that $\|Y_n\| \asymp_r \sqrt{\ell}$ with high probability in the regime where r is constant. Our second main theorem is a refinement of [KX26] with an explicit dependence on r and an improved allowed range of $\ell = \ell(n)$.

Theorem 1.6. *If $\ell \lesssim n/(r^2 \log^2 n)$, then $\|Y_n\| \lesssim \sqrt{\ell}/r$ with high probability.*

Together with the lower bound (5) and weak convergence, this determines the limit of $\|Y_n\|$ for fixed ℓ, r up to a factor $O(\sqrt{r})$.

Next, we show that the norm convergence question can be settled entirely for a slight variant of Kikuchi matrices. When $\ell = r = 1$, we note that Y_n is (up to its diagonal) just a normalized real symmetric matrix drawn from the *Gaussian orthogonal ensemble* (GOE); we define the *directed Kikuchi matrix* to be the corresponding generalization of the complex Hermitian *Gaussian unitary ensemble* (GUE). Formally, given integers n, ℓ, r satisfying $r \leq \ell \leq n - r$, the directed Kikuchi matrix $X_n = X_n(\ell, r)$ is the random matrix with complex entries

$$X_n[U, V] := \begin{cases} \frac{1}{\sqrt{d_n}} g_{V \setminus U, U \setminus V} & \text{if } |U \triangle V| = 2r \\ 0 & \text{otherwise} \end{cases} \quad \text{for } U, V \in \binom{[n]}{\ell}.$$

Here, $g_{H,T}$ are standard complex Gaussians for all disjoint $H, T \subseteq [n]$ of size $|H| = |T| = r$, where $g_{H,T}$ and $g_{H',T'}$ are independent for every $\{H, T\} \neq \{H', T'\}$, and $g_{H,T} = \overline{g_{T,H}}$ (complex conjugate). The normalization $d_n = d_n(\ell, r)$ is defined in (2).

Directed Kikuchi matrices have the same weak limit as Kikuchi matrices (see Theorem 3.2). Our third main result shows that, for a single directed Kikuchi matrix, this weak convergence can be upgraded to convergence of the norm.

Theorem 1.7. *For fixed $1 \leq r \leq \ell$, we have $\|X_n\| \rightarrow \|\Sigma\|$ in probability as $n \rightarrow \infty$.*

The proof of Theorem 1.7 generalizes to ℓ growing with n , in the sense that $\|X_n\| \leq (1 + o(1))\|\Sigma\|$ continues holding for constant r and $\ell(n) \leq n^\delta$ for some small constant $\delta > 0$.

The q -Gaussian limit. After taking the limit $n \rightarrow \infty$, a further limit $\ell \rightarrow \infty$ with $r = r(\ell)$ connects Kikuchi matrices with another central object in noncommutative probability: the q -Gaussian system [FB70, BS91, BS92, Spe92, BKS97]. To define this object, observe that the creation operators in Definition 1.3 satisfy

$$L_u^* L_v = \Gamma_{uv} L_v L_u^* + \mathbb{1}\{u = v\} \cdot I, \quad (6)$$

where I is the identity operator and Γ is the adjacency matrix of the graph Γ . The relations (6) for arbitrary fixed matrices Γ (not necessarily $\{0, 1\}$ -valued) are known as the Γ -Toeplitz relations. The q -Gaussian system with k variables is obtained analogously by constructing creation operators satisfying the Toeplitz relations with $\Gamma_{uv} := q$ for all $u, v \in [k]$, for some constant $q \in [0, 1]$. See Definition 2.6 for a formal definition.

Our fourth main result shows that tuples of independent Kikuchi matrices converge to a q -Gaussian system in an appropriate double limit.

Theorem 1.8. *Fix $k \geq 1$ and $\alpha > 0$. Let $\mathbf{X}_n$ and $\mathbf{Y}_n$ be k -tuples of independent directed Kikuchi and Kikuchi matrices, respectively. Then in the double limit $n \rightarrow \infty$, followed by $\ell \rightarrow \infty$ with $r = r(\ell)$ such that $r^2/\ell \rightarrow \alpha$, both $\mathbf{X}_n$ and $\mathbf{Y}_n$ converge weakly to the $e^{-\alpha}$ -Gaussian system with k variables. Furthermore, for $k = 1$, X_n converges strongly to the same limit.*

In light of the multivariate central limit theorem for Γ -independent random variables of [COY24], the weak convergence statement in Theorem 1.8 is equivalent to showing that the blowups of the Kneser graphs converge to the constant graphon. We conjecture that the strong convergence statement for the directed Kikuchi matrix extends to any $k \geq 1$ and $\alpha > 0$ (for $\alpha = 0$, the linear operators of the corresponding “1-Gaussian system” are unbounded, so strong convergence does not make sense).

1.2 Connections and broader context

Intrinsic Γ -independence. The Kikuchi matrix is a Gaussian matrix series: in the notation of (1),

$$Y_n = \sum_{I \in \binom{[n]}{2r}} g_I M_I, \quad \text{where } M_I[U, V] := \begin{cases} \frac{1}{\sqrt{d_n}} & \text{if } U \triangle V = I, \\ 0 & \text{otherwise.} \end{cases}$$

Given this representation, the intrinsic freeness approach of [BBvH23] suggests comparing the spectrum of Y_n to that of the operator obtained by replacing the independent Gaussians with freely independent semicircular elements:

$$Y_{\text{free}} := \sum_{I \in \binom{[n]}{2r}} S_I \otimes M_I, \quad \{S_I\} \text{ is a freely independent semicircular system.} \quad (7)$$

However, as discussed above, the spectrum of Y_{free} does not asymptotically match that of Y_n once $\ell \geq 2r$. Instead, our results suggest that the appropriate operator to compare Y_n to is

$$Y_\Gamma := \sum_{I \in \binom{[n]}{2r}} S_I \otimes M_I, \quad \{S_I\} \text{ is a } K(n, 2r)\text{-independent semicircular system.} \quad (8)$$

Indeed, for every ℓ, n, r , the operator Y_Γ has exactly the same moments as the normalized sum of the Γ -independent semicircular system associated with the graph product $\Gamma = K(\ell, r) \times K(n - \ell, r)$.⁴ As $n \rightarrow \infty$ for fixed ℓ, r , the second factor $K(n - \ell, r)$ approaches a complete graph, and therefore Y_Γ converges weakly to the normalized sum of the $K(\ell, r)$ -independent semicircular system – precisely the weak limit of Y_n identified in Theorem 1.5.

Thus, the example of Kikuchi matrices motivates the development of a theory of “intrinsic Γ -independence”. As a first step in this direction, the parallel work [BKN⁺26b] derives general upper bounds on the norms of operators of the form (8), in the spirit of what Lechner’s formula achieves for (7) (see Proposition 6.6).

Hypergraph Moore bound. The dense Kikuchi matrix (1) has a natural sparse generalization: Given any $2r$ -uniform hypergraph $\mathcal{H}$ on n vertices, define $Y_{\mathcal{H}}$ by keeping only the entries (U, V) of Y such that $U \triangle V$ is a hyperedge of $\mathcal{H}$ (and setting all other entries to 0). Bounding the norm of these random sparse Kikuchi matrices $Y_{\mathcal{H}}$ was an instrumental step in the resolution, up to logarithmic factors, of Feige’s conjecture on the optimal girth-density tradeoff of hypergraphs, also known as the hypergraph Moore bound [GKM22, HKM23]. One of the sources of the spurious logarithmic factor is the application of bounds such as (3).

Our initial motivation to study limits of Kikuchi matrices was to remove these remaining logarithmic factors to establish the hypergraph Moore bound. Indeed, we were able to show that a natural candidate limiting object does have the spectral properties that would be needed to establish the conjecture. This appears as an application of a method we developed to understand objects arising in non-commutative probability [BKN⁺26b]. However, we were not able to show the level of convergence needed to establish the hypergraph Moore bound via this route.

While this work was in progress, the hypergraph Moore bound was established in two independent works [BKN⁺26a, SH26], using different techniques.

Matrix models for noncommutative independence. We show that Γ -independent and q -Gaussian systems appear naturally as limits of Kikuchi matrices. Historically, significant attention has been devoted to the converse direction: starting from a prescribed set of noncommutative relations, can one construct “natural” sequences of random matrices that realize them asymptotically, in the sense of weak or strong convergence?

The question of constructing such a random matrix model for the q -Gaussian system has its origins in mathematical physics, going back to Frisch and Bourret in 1970 [FB70], and was solved in the weak sense in [Spe92, Šni01]. For Γ -independence, weak matrix models were constructed in [CC21] and strong matrix models in [MT26, CGvH26]. One example is the *tensor GUE* model. For every fixed k , the results of [GS14], combined with strong convergence of the GUE [HT05], yield a strong matrix model for q -Gaussian systems with k variables when q is smaller than a positive constant depending

⁴That is, the graph with vertex set $\binom{[\ell]}{r} \times \binom{[n-\ell]}{r}$, with an edge between (a, a') and (b, b') if and only if $a \cap b = a' \cap b' = \emptyset$.

on k . Very recently, strong matrix models were constructed in the range $q < \sqrt{2} - 1$, for any k [CC26]; the case of larger q remains open.⁵

In contrast, we obtain weak (and conjecturally strong) models for these systems from the opposite direction: they arise naturally in the analysis of efficient algorithms for tensor PCA.

1.3 Organization of the paper

We gather preliminaries for the technical sections in Section 2. Each subsequent section focuses on one of our main theorems and can be read largely independently: Section 3 proves Theorem 1.5; Section 4 proves Theorem 1.6; Section 5 proves Theorem 1.7; and Section 6 proves Theorem 1.8.

1.4 Acknowledgments and AI disclosure

LP acknowledges support from an ETH Zürich Postdoctoral Fellowship (Project No. 26-1 FEL-092). Part of this work was done while RW was an academic guest in the Department of Mathematics at ETH Zürich. RW's work on this project was supported by the Queen Elizabeth II Graduate Scholarship in Science and Technology (QEII-GSST). PNN and RW thank the Simons Institute for the Theory of Computing and the organizers of the Fall 2026 program on *Spectral Theory Beyond Graphs*.

Most of this project was developed before AI tools could help with mathematics as they can now. Nevertheless, more recently, we have used AI tools to help perform some computations, carry through some specific arguments, and help with references. The authors wrote the paper and take responsibility for all its (potential) errors.

2 Preliminaries

All graphs in this paper are simple and undirected. For a graph Γ , let $V(\Gamma)$, $E(\Gamma)$ denote its vertex and edge sets, respectively; we also write Γ for its adjacency matrix. We define

$$\text{Hom}(G, \Gamma) := \{\varphi : V(G) \rightarrow V(\Gamma) : \{\varphi(u), \varphi(v)\} \in E(\Gamma) \ \forall \{u, v\} \in E(G)\}$$

to be the set of graph homomorphisms from a graph G to Γ . Let $\mathcal{P}(p)$ be the set of *pairings* (perfect matchings) on $[p]$. We denote the edges of a pairing π by $ij \in E(\pi)$ when $1 \leq i < j \leq p$. For each $ij \in E(\pi)$, we call i an *opening index* and j a *closing index*. Finally, let

$$\mathcal{P}(\mathbf{w}) := \{\pi \in \mathcal{P}(p) : w_i = w_j \ \forall ij \in E(\pi)\} \quad \text{for any } \mathbf{w} = (w_1, \dots, w_p).$$

Pairings appear naturally when computing the moments of multivariate Gaussians:

Lemma 2.1 (Wick's formula). *If $g \in \mathbb{C}^p$ is a centered Gaussian vector, then*

$$\mathbb{E} \left[\prod_{i=1}^p g_i \right] = \sum_{\pi \in \mathcal{P}(p)} \prod_{ij \in E(\pi)} \mathbb{E}[g_i g_j].$$

⁵The construction in [CC26] has some similarities with ours. The authors consider random matrices converging strongly to a Γ_ℓ -independent semicircular system for each ℓ , where Γ_ℓ is a sequence of Erdős-Rényi random graphs. A further limit $\ell \rightarrow \infty$ combined with a diagonal argument yields their q -Gaussian matrix model.

2.1 State and operator norm

Let $\mathcal{B}(\mathcal{H})$ be the space of bounded linear maps $\mathcal{H} \rightarrow \mathcal{H}$ on a Hilbert space $\mathcal{H}$. Let I be the identity operator. We write A^* for the adjoint of the operator A , and $\|A\|$ for the operator norm of A .

The moments of an operator are specified by a *state* $\tau : \mathcal{B}(\mathcal{H}) \rightarrow \mathbb{C}$, which is a linear functional such that $\tau(I) = 1$ and $\tau(A^*A) \geq 0$ for all $A \in \mathcal{B}(\mathcal{H})$. The definition automatically implies

$$\tau((A^*A)^p)^{\frac{1}{2p}} \leq \|A\| \quad \forall p \geq 1. \quad (9)$$

Given self-adjoint operators $A_1, \dots, A_k \in \mathcal{B}(\mathcal{H})$, let $C^*(A_1, \dots, A_k)$ be the closure (in the operator norm topology) of the set of all noncommutative polynomials in $A_1, \dots, A_k$. We say that τ is *faithful* on $C^*(A_1, \dots, A_k)$ if $\tau(A^*A) \neq 0$ for all $A \in C^*(A_1, \dots, A_k) \setminus \{0\}$. Under that assumption, the norm of an operator can be recovered from its moments.

Lemma 2.2 ([NS06, Proposition 3.17]). *Let $A_1, \dots, A_k$ be self-adjoint operators. If τ is faithful on $C^*(A_1, \dots, A_k)$, then*

$$\lim_{p \rightarrow \infty} \tau((A^*A)^p)^{\frac{1}{2p}} = \|A\| \quad \forall A \in C^*(A_1, \dots, A_k).$$

Finally, we say that τ is *tracial* on $C^*(A_1, \dots, A_k)$ if $\tau(AB) = \tau(BA)$ for all $A, B \in C^*(A_1, \dots, A_k)$.

Lemma 2.3 (Hölder's inequality; see [PX03, Section 1]). *Let $A_1, \dots, A_k$ be self-adjoint operators. If τ is tracial on $C^*(A_1, \dots, A_k)$ and k is even, then*

$$|\tau(A_1 \cdots A_k)| \leq \tau(A_1^k)^{\frac{1}{k}} \cdots \tau(A_k^k)^{\frac{1}{k}}.$$

2.2 Γ -independent and q -Gaussian systems

The following theorem from [BS94] ensures the existence of operators satisfying the Γ -Toeplitz relations.

Theorem 2.4. *Let $\Gamma \in [-1, 1]^{k \times k}$ be a symmetric matrix such that $\Gamma_{ii} < 1$ for all $i \in [k]$. Then there exists a Hilbert space $\mathcal{H}$, a distinguished unit vector $x \in \mathcal{H}$ and operators $L_1, \dots, L_k \in \mathcal{B}(\mathcal{H})$ with the following properties. Letting $\tau(A) := \langle x, Ax \rangle$ and $S_i := L_i + L_i^*$, we have that*

1. $L_1, \dots, L_k$ satisfy the Γ -Toeplitz relations (6);
2. $L_i^*x = 0$ for each $i \in [k]$;
3. τ is faithful and tracial on $C^*(S_1, \dots, S_k)$.

Remark 2.5. Although the concrete Hilbert space representation need not be unique, conditions 1-2-3 of Theorem 2.4 uniquely determine the information needed for weak and strong convergence to $(S_1, \dots, S_k)$. Indeed, it is not hard to see that conditions 1 and 2 uniquely determine by induction the joint moments $\tau(P(S_1, \dots, S_k))$ for every noncommutative polynomial P . Condition 3 then implies by Lemma 2.2 that these moments also determine the norm of every such polynomial.

Definitions 1.2 and 1.3 follow a known construction for proving Theorem 2.4 in the special case where Γ is the adjacency matrix of a graph [BS94, OTY25]. In this case, $\mathcal{H}$ is the Hilbert space indexed by the elements of the trace monoid, $L_1, \dots, L_k$ are the corresponding creation operators, and x is the basis vector corresponding to the empty word.

Definition 2.6. Let $q \in [0, 1]$. A q -Gaussian system on k variables is a tuple $(S_1, \dots, S_k)$ satisfying the conclusion of Theorem 2.4 for $\Gamma_{ij} = q$ for all $i, j \in [k]$. Each S_i is called a q -Gaussian variable.

Finally, we compute the moments of Γ -independent and q -Gaussian systems explicitly. The *crossing graph*, G_π , of a pairing π is the graph whose vertices $V(G_\pi) = E(\pi)$ are edges of π , and two vertices in G_π are adjacent if the corresponding edges form a crossing in π . Two edges $uu', vv' \in E(\pi)$ are *crossing* if $u < v < u' < v'$ or $v < u < v' < u'$.

Lemma 2.7 (See [BS91, Proposition 2]). *Let $(S_1, \dots, S_k)$ be a q -Gaussian system. Then*

$$\tau(S_{w_1} \cdots S_{w_p}) = \sum_{\pi \in \mathcal{P}(\mathbf{w})} q^{|E(G_\pi)|} \quad \forall p \geq 1 \quad \forall \mathbf{w} \in [k]^p.$$

Lemma 2.8 (See [SW16, Theorem 5.2]). *Let $(S_u)_{u \in V(\Gamma)}$ be a Γ -independent semicircular system. Then for any $p \geq 1$ and any choice of complex coefficients $\{c_{i,u} : i \in [p], u \in V(\Gamma)\}$,*

$$\tau \left(\left(\sum_{u \in V(\Gamma)} c_{1,u} S_u \right) \cdots \left(\sum_{u \in V(\Gamma)} c_{p,u} S_u \right) \right) = \sum_{\substack{\pi \in \mathcal{P}(p) \\ \varphi \in \text{Hom}(G_\pi, \Gamma)}} \prod_{ij \in E(\pi)} c_{i, \varphi(ij)} c_{j, \varphi(ij)}.$$

In particular,

$$\tau \left(\left(\sum_{u \in V(\Gamma)} S_u \right)^p \right) = \sum_{\pi \in \mathcal{P}(p)} |\text{Hom}(G_\pi, \Gamma)|.$$

2.3 Kikuchi walks

In this subsection, we start introducing the common terminology for the technical sections. In particular, we define *Kikuchi walks*, which are the combinatorial objects that arise in the trace moment expansion of Kikuchi matrices.

Definition 2.9. A *directed hyperedge* is a pair

$$e = (H(e), T(e)) \in \binom{[n]}{r} \times \binom{[n]}{r} \quad \text{such that } H(e) \cap T(e) = \emptyset.$$

The sets $H(e)$ and $T(e)$ are called the *head* and *tail* of e , respectively. The *support* of e is $\text{supp}(e) := H(e) \cup T(e)$. The *reverse* of e is $e^* := (T(e), H(e))$. Let $\mathcal{E}_{n,r}$ be the set of all directed hyperedges.

Definition 2.10. A *Kikuchi walk* of length p is a tuple $\mathbf{e} = (e_1, \dots, e_p) \in \mathcal{E}_{n,r}^p$ for which

$$U_0 = [\ell], \quad U_i := (U_{i-1} \cup H(e_i)) \setminus T(e_i),$$

satisfy $U_p = U_0$, $T(e_i) \subseteq U_{i-1}$ and $H(e_i) \cap U_{i-1} = \emptyset$ for all $i \in [p]$.

Define the support sequence of a Kikuchi walk by $\text{supp}(\mathbf{e}) := (\text{supp}(e_i))_{i \in [p]}$. We call a Kikuchi walk *even* if every support appearing in its support sequence occurs an even number of times, and denote by $\mathcal{W}_p$ the set of even Kikuchi walks of length p . Note that even Kikuchi walks automatically have even length.

To lighten notation, we will write U_i , $T_i := T(e_i)$, $H_i := H(e_i)$, where the dependence on $\mathbf{e}$ is implicit but should be clear from context.

Lemma 2.11. *The trace moments of the Kikuchi matrix Y_n and directed Kikuchi matrix X_n satisfy*

$$\mathbb{E} \operatorname{tr} Y_n^p = d_n^{-\frac{p}{2}} |\{(\mathbf{e}, \pi) \in \mathcal{W}_p \times \mathcal{P}(p) : \operatorname{supp}(\mathbf{e}_i) = \operatorname{supp}(\mathbf{e}_j) \ \forall ij \in E(\pi)\}|, \quad (10)$$

$$\mathbb{E} \operatorname{tr} X_n^p = d_n^{-\frac{p}{2}} |\{(\mathbf{e}, \pi) \in \mathcal{W}_p \times \mathcal{P}(p) : \mathbf{e}_i = \mathbf{e}_j^* \ \forall ij \in E(\pi)\}|. \quad (11)$$

Proof. This is a direct application of Wick's formula (Lemma 2.1), together with the invariance of X_n and Y_n by relabeling (which allows us to average out the choice of U_0). $\square$

We organize the Kikuchi walks that contribute to (10) and (11) in terms of the number of vertices that they introduce. For $i \in [p]$, let

$$R_i := H_i \setminus \left(\bigcup_{j=0}^{i-1} U_j \right)$$

be the set of vertices introduced for the first time at step i . For an even walk $\mathbf{e}$, every support appearing in $\operatorname{supp}(\mathbf{e})$ occurs at least twice, a support that has appeared previously cannot introduce any new vertex, and any step can introduce at most r vertices. Therefore, an even walk can introduce $\sum_{i=1}^p |R_i| \leq rp/2$ vertices in total. We call the integer $rp/2 - \sum_{i=1}^p |R_i|$ the *defect* of the walk $\mathbf{e}$.

2.4 Kikuchi transports

By Lemma 2.8, the combinatorial objects that contribute to the moments of the $K(\ell, r)$ -independent semicircular system are tuples (φ, π) where $\pi \in \mathcal{P}(p)$ and $\varphi \in \operatorname{Hom}(G_\pi, K(\ell, r))$. In words, φ is an assignment of elements of $\binom{\ell}{r}$ to every pair in π , with the property that these subsets are disjoint whenever two pairs cross. To map the paired Kikuchi walks from Lemma 2.11 to these predicted limiting objects, we introduce the notion of *Kikuchi transport*.

Definition 2.12. Let $m \geq 0$ and $(\mathbf{e}, \pi) \in \mathcal{W}_p \times \mathcal{P}(p)$. A *Kikuchi transport* of $(\mathbf{e}, \pi)$ is a sequence $\boldsymbol{\rho} = (\rho_0, \dots, \rho_p)$ of injective maps $\rho_i : U_i \rightarrow [\ell + m]$ with $\rho_0 = \operatorname{id}_{[\ell]}$, such that

- (Pairing compatibility) $\rho_{i-1}(T_i) = \rho_{j-1}(T_j)$ for any $ij \in E(\pi)$.
- (Crossing compatibility) $\rho_{i-1}(T_i) \cap \rho_{i'-1}(T_{i'}) = \emptyset$ for any $\{ij, i'j'\} \in E(G_\pi)$.

We call m the *excess size* of $\boldsymbol{\rho}$.

Equivalently, a Kikuchi transport is a sequence of embeddings such that the map

$$\begin{aligned} G_\pi &\longrightarrow K(\ell + m, r) \\ ij &\longmapsto \rho_{i-1}(T_i) = \rho_{j-1}(T_j) \end{aligned} \quad (12)$$

is a graph homomorphism. By Lemma 2.8, these are exactly the objects that moments of Kneser semicircular systems count (for $m = 0$).

When designing Kikuchi transports, we will try to make sure that they respect as much as possible the *canonical bijections* $\phi_t : U_t \rightarrow U_{t-1}$, which are the maps fixing $U_t \cap U_{t-1} = U_{t-1} \setminus T_t = U_t \setminus H_t$, and mapping the elements of $H_t = U_t \setminus U_{t-1}$ to the elements of $T_t = U_{t-1} \setminus U_t$ in increasing order.

Definition 2.13. Let $\boldsymbol{\rho}$ be a Kikuchi transport of $(\mathbf{e}, \pi)$. A pair $(t, u) \in [p] \times U_t$ is called a *switch* if $\rho_t(u) \neq \rho_{t-1}(\phi_t(u))$. The *cost* of $\boldsymbol{\rho}$ is the total number of switches (t, u) it incurs.

3 Weak Convergence of Kikuchi Matrices

In this section, we prove weak convergence of Kikuchi matrices (Theorem 1.5) and directed Kikuchi matrices to their Γ -independent limit as special cases of Theorem 3.2 below.

Definition 3.1. Let ℓ, r, n be integers such that $1 \leq r \leq \ell \leq n - r$. Let $d_n = d_n(\ell, r)$ be as in (2). A *generalized Kikuchi matrix* $Z_n = Z_n(\ell, r)$ is a matrix whose entries have the form

$$Z_n[U, V] := \begin{cases} \frac{1}{\sqrt{d_n}} g_{(V \setminus U, U \setminus V)} & \text{if } |U \triangle V| = 2r \\ 0 & \text{otherwise} \end{cases} \quad \forall U, V \in \binom{[n]}{\ell},$$

where $(g_e)_{e \in \mathcal{E}_{n,r}}$ is a centered (real or complex) Gaussian vector satisfying $g_{e^*} = \overline{g_e}$, $\mathbb{E} |g_e|^2 = 1$ and $\mathbb{E} [g_e g_f] = 0$ whenever $\text{supp}(e) \neq \text{supp}(f)$.

The Kikuchi matrix is obtained by sampling a collection of i.i.d. standard real Gaussians g'_S for $S \subseteq [n]$ of size $|S| = 2r$ and setting $g_e = g'_{\text{supp}(e)}$. The directed Kikuchi matrix is obtained by sampling $(g_e)_{e \in \mathcal{E}_{n,r}}$ as standard complex Gaussians, independently up to the symmetry $g_{e^*} = \overline{g_e}$.

Theorem 3.2. Fix $k \geq 1$ and $1 \leq r \leq \ell$. Let $\Sigma = (\Sigma_1, \dots, \Sigma_k)$ be as in (4). Any sequence of k -tuples of independent generalized Kikuchi matrices $\mathbf{Z}_n = (Z_{n,1}, \dots, Z_{n,k})$ converges weakly as $n \rightarrow \infty$ to Σ .

The proofs in this section and in Sections 4 and 5 are all based on counting Kikuchi walks and compatible pairings appearing in the expansion of the trace moments. The difference between the arguments lies in how the walks are counted. For weak convergence, the moment p is fixed as $n \rightarrow \infty$. In this regime, the dominant contributions come from the Kikuchi walks of defect 0. Since p is fixed, it then suffices to use crude bounds on the number of paired walks of higher defect (Lemma 3.7).

Lemma 3.3. An even Kikuchi walk of length p has defect 0 if and only if $\text{supp}(e)$ consists of $p/2$ supports, each occurring twice, and the first occurrence of each support introduces exactly r vertices.

Proof. For a defect-0 walk, the above argument that bounds the number of vertices introduced in the walk must hold with equality, so every support in $\text{supp}(e)$ occurs exactly twice and every first occurrence of a support introduces exactly r vertices. The converse statement is immediate. $\square$

Lemma 3.3 shows that a defect-0 walk e determines a unique pairing $\pi(e) \in \mathcal{P}(\text{supp}(e))$. This process pairs together reverse hyperedges:

Lemma 3.4. If e has defect 0, then $e_i = e_j^*$ for every $ij \in E(\pi(e))$.

Proof. We proceed by induction on the length p of e . The statement holds for $p = 0$.

Let $p \geq 2$ and suppose that the statement holds for $p - 2$. Let $j \in [p]$ be the smallest closing index, and let $ij \in E(\pi(e))$. For every $i < k < j$, e_k must be the first occurrence of some support by minimality of j , and by Lemma 3.3, the vertices in H_k are all new. This means that none of the vertices of T_i that have been removed at step i can be re-introduced before step j . Since $T_i \subseteq \text{supp}(e_j)$ and $|T_i| = |H_j| = r$, it must be that $T_i = H_j$. Since heads and tails are disjoint and $\text{supp}(e_i) = \text{supp}(e_j)$ by assumption, it follows automatically that $e_i = e_j^*$.

Finally, removing e_i and e_j from the walk yields a defect-0 walk of length $p - 2$ by Lemma 3.3. By induction, the property holds for any edge in $\pi(e)$. $\square$

Lemma 3.5. *Let p be even and $\pi \in \mathcal{P}(p)$. The number of even Kikuchi walks of length p , defect-0, and such that $\pi(\mathbf{e}) = \pi$ is*

$$\frac{(n - \ell) \cdots (n - \ell - \frac{rp}{2} + 1)}{(r!)^{\frac{p}{2}}} |\text{Hom}(G_\pi, K(\ell, r))|.$$

Lemma 3.6. *For any defect-0 even Kikuchi walk $\mathbf{e}$, $(\mathbf{e}, \pi(\mathbf{e}))$ has a Kikuchi transport of excess size 0 and cost 0.*

Proof of Lemma 3.5 from Lemma 3.6. The map (12) induced by the transport is a graph homomorphism $G_\pi \rightarrow K(\ell, r)$. Moreover, at each opening index, the head consists of r new vertices. Thus, the $p/2$ opening heads are pairwise disjoint subsets of $[n] \setminus [\ell]$ of size r ; the number of such objects is

$$\binom{n - \ell}{r} \cdot \binom{n - \ell - r}{r} \cdots \binom{n - \ell - r(\frac{p}{2} - 1)}{r} = \frac{(n - \ell) \cdots (n - \ell - \frac{rp}{2} + 1)}{(r!)^{\frac{p}{2}}}.$$

Conversely, a graph homomorphism together with the disjoint sets assigned to the opening heads uniquely determines the walk. Indeed, we can reconstruct the walk and a cost-0 transport of it simultaneously from left to right. Suppose that we have reconstructed $\rho_1, \dots, \rho_{i-1}$ and $U_1, \dots, U_{i-1}$. If i is a closing index, then the head and tail of e_i are determined from the matching opening by the mirroring property of Lemma 3.4. If i is an opening index with $ij \in E(\pi)$, then H_i is prescribed by the opening head data, and $T_i = \rho_{i-1}^{-1}(\varphi(ij))$. This determines U_i , and the cost-0 condition then uniquely determines ρ_i from ρ_{i-1} . $\square$

Proof of Lemma 3.6. Define recursively $\rho_0(u) = u$, $\rho_t(u) = \rho_{t-1}(\phi_t(u))$. Then ρ_t is automatically injective, has image $[\ell]$, and incurs 0 switches.

We show by induction on p that ρ is a Kikuchi transport. Let $ij \in E(\pi)$ be such that j is the smallest closing index. Then every index in $\{i + 1, \dots, j - 1\}$ is an opening index. None of these indices can remove a vertex of $H_i = T_j$: indeed, once removed, such a vertex could not be reintroduced before time j , since every intermediate head consists only of new vertices (which would contradict $T_j \subseteq U_{j-1}$). Therefore, the labels of H_i remain unchanged through the canonical bijections, and $\rho_{j-1}(T_j) = \rho_i(H_i) = \rho_{i-1}(T_i)$. Next, suppose that $\{ij, i'j'\} \in E(G_\pi)$ (note that $i' > i$ by minimality of j). By the above argument, we have $T_{i'} \cap H_i = \emptyset$ and $\rho_{i-1}(T_i) = \rho_i(H_i) = \rho_{i'-1}(H_i)$, so by injectivity of $\rho_{i'-1}$, we get $\rho_{i'-1}(T_{i'}) \cap \rho_{i-1}(T_i) = \emptyset$. This shows that the conditions in Definition 2.12 hold for ij ; removing these indices from the walk and the pairing yields a defect-0 walk of length $p - 2$, which implies the statement by induction. $\square$

Lemma 3.3 implies that defect-0 walks have exactly one pairing respecting their support. We robustify this bound as follows:

Lemma 3.7. *For any even Kikuchi walk $\mathbf{e}$ of length p and defect Δ , $|\mathcal{P}(\text{supp}(\mathbf{e}))| \leq p^{\Delta/r}$.*

Proof. Let m be the number of distinct supports in $\text{supp}(\mathbf{e})$. Since each directed hyperedge can introduce at most r new vertices and there are $rp/2 - \Delta$ new vertices in total, we have $p/2 - m \leq \Delta/r$. Furthermore, if the supports have multiplicities $\beta_1, \dots, \beta_m$,

$$|\mathcal{P}(\text{supp}(\mathbf{e}))| = \prod_{j=1}^m (\beta_j - 1)!! \leq \prod_{j=1}^m \beta_j^{\beta_j/2-1} \leq p^{p/2-m} \leq p^{\Delta/r},$$

which concludes the proof. $\square$

Let $\mathbf{Z} = (Z_1, \dots, Z_k)$ be a tuple of independent generalized Kikuchi matrices, where we drop the dependence on n to lighten notation. For $i \in [k]$, we denote by $(g_{i,e})_{e \in \mathcal{E}_{n,r}}$ the Gaussian vector associated to Z_i . Given a Kikuchi walk $\mathbf{e}$ of length p and $\mathbf{w} \in [k]^p$, define the weight

$$\text{wt}_{\mathbf{Z}}(\mathbf{w}, \mathbf{e}) := \mathbb{E} \left[\prod_{a=1}^p g_{w_a, e_a} \right] = \sum_{\pi \in \mathcal{P}(\mathbf{w})} \prod_{ab \in E(\pi)} \mathbb{E}[g_{w_a, e_a} g_{w_b, e_b}] , \quad (13)$$

where the identity follows from Wick's formula (Lemma 2.1) and independence of the matrices within the tuple. We make several simple observations about this definition of weight.

1. Only pairings $\pi \in \mathcal{P}(\text{supp}(\mathbf{e}))$ can have nonzero contribution to the sum on the right-hand side of (13). In particular, only even Kikuchi walks can have nonzero weight.
2. Since the Gaussians have variance 1, Cauchy-Schwarz implies that every term in the sum has magnitude at most 1. It then follows from Lemma 3.7 that the weight of any defect- Δ , length- p walk is bounded by $|\mathcal{P}(\text{supp}(\mathbf{e}))| \leq p^{\Delta/r}$.
3. By Lemma 3.3, if $\mathbf{e}$ has defect 0, then $\mathcal{P}(\text{supp}(\mathbf{e})) = \{\pi(\mathbf{e})\}$, so combining with the first observation and Lemma 3.4, we have $\text{wt}_{\mathbf{Z}}(\mathbf{w}, \mathbf{e}) = 1$ if $\pi(\mathbf{e}) \in \mathcal{P}(\mathbf{w})$, and 0 otherwise.

Proof of Theorem 3.2. Our definition of a Kikuchi walk assumes that its initial set is $[\ell]$. To reduce to this case, for each $U \in \binom{[n]}{\ell}$, fix an arbitrary permutation σ_U of $[n]$ satisfying $\sigma_U([\ell]) = U$, and let $\mathbf{Z}(U) = (Z_1(U), \dots, Z_k(U))$ be the tuple obtained by simultaneously permuting the rows and columns of every Z_i according to σ_U . Then $\mathbf{Z}(U)$ is again a tuple of independent generalized Kikuchi matrices.

For every $\mathbf{w} \in [k]^p$, we have

$$\mathbb{E} \text{tr} (Z_{w_1} \cdots Z_{w_p}) = \binom{n}{\ell}^{-1} \sum_{U \in \binom{[n]}{\ell}} d_n^{-p/2} \sum_{\mathbf{e} \in \mathcal{W}_p} \text{wt}_{\mathbf{Z}(U)}(\mathbf{w}, \mathbf{e}) . \quad (14)$$

By the first observation above, only even Kikuchi walks contribute to (14). Hence, when p is odd,

$$\mathbb{E} \text{tr} (Z_{w_1} \cdots Z_{w_p}) = 0 = \tau(\Sigma_{w_1} \cdots \Sigma_{w_p}) ,$$

where the second equality follows from Lemma 2.8. Suppose from now on that p is even.

An even Kikuchi walk of defect Δ introduces $rp/2 - \Delta$ vertices, so there are at most $O_{p,\ell,r}(n^{rp/2-\Delta})$ such walks. Since $|\text{wt}_{\mathbf{Z}(U)}(\mathbf{w}, \mathbf{e})| \leq p^{\Delta/r}$ by the second observation and $d_n^{p/2} = \Omega_{p,\ell,r}(n^{rp/2})$, the total contribution of positive-defect walks converges to 0, uniformly in U . By the third observation, defect-0 walks satisfy $\text{wt}_{\mathbf{Z}(U)}(\mathbf{w}, \mathbf{e}) = \mathbb{1}\{\pi(\mathbf{e}) \in \mathcal{P}(\mathbf{w})\}$. Therefore,

$$\begin{aligned} \mathbb{E} \text{tr} (Z_{w_1} \cdots Z_{w_p}) &= d_n^{-p/2} \sum_{\pi \in \mathcal{P}(\mathbf{w})} |\{\mathbf{e} \in \mathcal{W}_p : \mathbf{e} \text{ has defect 0 and } \pi(\mathbf{e}) = \pi\}| + o(1) \\ &= \frac{(n-\ell) \cdots (n-\ell - \frac{rp}{2} + 1)}{(r!)^{p/2} d_n^{p/2}} \sum_{\pi \in \mathcal{P}(\mathbf{w})} |\text{Hom}(G_\pi, K(\ell, r))| + o(1) \quad (\text{Lemma 3.5}) \\ &\xrightarrow{n \rightarrow \infty} \binom{\ell}{r}^{-p/2} \sum_{\pi \in \mathcal{P}(\mathbf{w})} |\text{Hom}(G_\pi, K(\ell, r))| = \tau(\Sigma_{w_1} \cdots \Sigma_{w_p}) . \quad (\text{Lemma 2.8}) \end{aligned}$$

This proves convergence of the trace moments in expectation. It remains to show that

$$\text{Var tr} (Z_{w_1} \cdots Z_{w_p}) \xrightarrow{n \rightarrow \infty} 0.$$

Expanding the expected square modulus by Wick's formula (as in (14)), the terms corresponding to pairings with no edge between the two copies of the length- p sequence cancel with the squared modulus of the expectation. Every remaining pairing contains an edge between the two copies, meaning that any pair of Kikuchi walks compatible with this pairing must share a vertex in the support. Such a pair of walks introduces at most $2\ell + rp - 1$ vertices. Hence, on the one hand the number of contributing pairs of walks is at most $O_{p,\ell,r}(n^{2\ell+rp-1})$. On the other hand, the normalization is

$$\binom{n}{\ell}^2 d_n^p = \Omega_{p,\ell,r}(n^{2\ell+rp}).$$

Since every term has magnitude at most 1 by the same argument as above, this implies that the variance is $O(1/n)$. Chebyshev's inequality therefore upgrades convergence in expectation to convergence in probability. $\square$

4 Norm Bound for Kikuchi Matrices

In this section, we prove the following bound on the norm of the Kikuchi matrix, which is tight up to a factor $O(\sqrt{r})$. Recall that the Kikuchi matrix $Y_n = Y_n(\ell, r)$ is defined in (1).

Theorem 4.1. *There exist constants $c, C > 0$ such that the following holds. For any integers $p, \ell, r, n \geq 1$ with $r \leq \ell$, $n \geq \ell + r$, and $p \leq c \min\{(n - \ell)/r^2, \sqrt{\ell(n - \ell)}/r\}$,*

$$(\mathbb{E} \text{tr} Y_n^{2p})^{\frac{1}{2p}} \leq C \sqrt{\frac{\ell}{r}}.$$

In particular, $\|Y_n\| \leq C \sqrt{\ell/r}$ with high probability when $\ell \leq cn/(r^2 \log^2 n)$.

Among all generalized Kikuchi matrices (Definition 3.1), the even trace moments are maximized for the undirected Kikuchi matrices, so Theorem 4.1 directly generalizes to arbitrary generalized Kikuchi matrices (and in particular to the directed Kikuchi matrices X_n).

The proof of Theorem 4.1 reduces to counting the number of objects appearing on the right-hand side of (10). The crude counting of higher-defect walks in Section 3 no longer suffices when p is growing with n .

Fix a tuple (e, π) such that $e \in \mathcal{W}_{2p}$ and $\pi \in \mathcal{P}(\text{supp}(e))$. Define

$$O_t := \{i < t : \exists j \geq t, ij \in E(\pi)\} \tag{15}$$

to be the set of opening indices whose pair is open immediately before index t . Call a closing index j with $ij \in E(\pi)$ *good* if $T_j \cap \bigcup_{k \in O_j \setminus \{i\}} \text{supp}(e_k) = \emptyset$, and *bad* otherwise. We call a pair (i, u) where i is an opening index and $u \in H_i \cap (\bigcup_{j < i} U_j)$ a *coincidence* of the walk. For a vertex u , we denote by $\delta(u)$ the number of times u appears as a coincidence. Note that $\sum_u \delta(u) = \Delta$.

Lemma 4.2. *If e has defect Δ , then at most Δ of the closing indices are bad.*

Proof. For a vertex u , the number of opening heads containing u equals the number of closing tails containing u . When j is a closing index, let $B_j := T_j \cap \bigcup_{k \in O_j \setminus \{i\}} \text{supp}(e_k)$, where i is the matching opening of j . If $u \in U_0$, then every occurrence of u in an opening head is a coincidence, so u belongs to at most $\delta(u)$ sets B_j . If $u \notin U_0$, then the number of opening heads containing u is $\delta(u) + 1$. However, the last occurrence of u is a closing tail and no other open pair at that point can contain u . Therefore, in each case we have $|\{j : u \in B_j\}| \leq \delta(u)$. Summing over u concludes the proof. $\square$

We say that $i \in O_t$ is *admissible* at time t if $|\text{supp}(e_i) \cap U_{t-1}| = r$. We say furthermore that it is *good* if $\text{supp}(e_i) \cap U_{t-1} \cap \text{supp}(e_j) = \emptyset$ for all $j \in O_t \setminus \{i\}$. When constructing a paired Kikuchi walk from left to right, admissible openings are the candidates that can be closed at time t . Good admissible openings create a good closing in this way.

Lemma 4.3. *The number of good admissible openings at any time is at most $\frac{\ell}{r}$.*

Proof. The candidate tails from the open pairs form a collection of pairwise disjoint subsets of U_{t-1} of size r . Such a collection has size at most $|U_{t-1}|/r = \ell/r$. $\square$

Lemma 4.4. *The number of admissible openings at any time is at most $\frac{\ell}{r} \left(1 + \frac{2\Delta}{\ell}\right)$.*

Proof. Let $\Phi_t := \sum_{i \in O_t} |\text{supp}(e_i) \cap U_{t-1}|$. On the one hand, an opening in O_t can be closed at step t only if its support intersects U_{t-1} in r elements, so the number of admissible choices is at most Φ_t/r . On the other hand, we claim that for $u \in U_{t-1}$, we have $|\{i \in O_t : u \in \text{supp}(e_i)\}| \leq 2\delta(u) + \mathbb{1}\{u \notin U_0\}$. Summing this inequality over u yields $\Phi_t \leq 2 \sum_{u \in U_{t-1}} \delta(u) + |U_{t-1} \setminus U_0| \leq 2\Delta + \ell$, as desired.

It remains to prove the above inequality. Let a, b, c, d denote the number of occurrences of u in an opening head, closing head, opening tail, and closing tail before time t . Since $u \in U_{t-1}$, we have $a + b - c - d = \mathbb{1}\{u \notin U_0\}$. Further, $a \leq \delta(u) + \mathbb{1}\{u \notin U_0\}$. Hence, the number of openings in O_t whose support contains u is $a + c - b - d = 2(a - d) - \mathbb{1}\{u \notin U_0\} \leq 2\delta(u) + \mathbb{1}\{u \notin U_0\}$. $\square$

Proof of Theorem 4.1. We start by fixing globally the defect of the walk $\Delta \geq 0$; the sequence $(\delta_1, \dots, \delta_p)$, where δ_i is the number of coincidences at the i -th opening index (at most $2^{p+\Delta}$ choices since $\sum_i \delta_i = \Delta$); and the positions of the opening indices and bad closings (at most 3^{2p} choices). Afterwards, we construct the pair $(\mathbf{e}, \pi)$ from left to right for $t = 1, \dots, 2p$. Assume that we have chosen $\text{supp}(e_1), \dots, \text{supp}(e_{t-1})$ and $\{ij \in E(\pi) : 1 \leq i < j < t\}$. Then there are two cases.

1. t is an opening index: The tail vertices can be identified as a subset of r elements from U_{t-1} . The coincidence vertices in the head are a subset of the $\leq 2pr$ previous vertices that have previously appeared in a tail in the walk. The non-coincidence vertices in the head are a subset of the $n - \ell$ vertices that have not yet appeared in the walk.
2. t is a closing index: It suffices to match t to an element of O_t . If t is a good (resp. bad) closing, there are $\frac{\ell}{r}$ (resp. $\frac{\ell}{r} \left(1 + \frac{2\Delta}{\ell}\right)$) admissible choices by Lemma 4.3 (resp. Lemma 4.4).

Overall, this counting applied to the right-hand side of (10) gives

$$\mathbb{E} \text{tr} Y_n^{2p} \leq \left(\frac{9}{d_n}\right)^p \sum_{\Delta \geq 0} \left[\left(\frac{\ell}{r}\right)^p \left(1 + \frac{2\Delta}{\ell}\right)^{\min\{p, \Delta\}} \sum_{\delta_1 + \dots + \delta_p = \Delta} \prod_{j=1}^p \binom{\ell}{r} \binom{2pr}{\delta_j} \binom{n - \ell}{r - \delta_j} \right].$$

The contribution of each opening index normalized by d_n can be upper bounded by

$$\frac{1}{d_n} \binom{\ell}{r} \binom{2pr}{\delta_j} \binom{n-\ell}{r-\delta_j} = \binom{2pr}{\delta_j} \frac{\binom{n-\ell}{r-\delta_j}}{\binom{n-\ell}{r}} \leq \left(\frac{2pr^2}{n-\ell-r+1} \right)^{\delta_j}.$$

The total contribution of the closing indices can be upper bounded by

$$\left(\frac{\ell}{r} \right)^p \left(1 + \frac{2\Delta}{\ell} \right)^{\min\{p, \Delta\}} \leq \left(\frac{\ell}{r} \right)^p \left(1 + \frac{2p}{\ell} \right)^\Delta.$$

Putting everything together,

$$\mathbb{E} \operatorname{tr} Y_n^{2p} \leq \left(\frac{18\ell}{r} \right)^p \sum_{\Delta \geq 0} \left(\frac{4pr^2}{n-\ell-r+1} \right)^\Delta \left(1 + \frac{2p}{\ell} \right)^\Delta \leq \left(\frac{C\ell}{r} \right)^p,$$

where the last inequality uses the fact that the series in Δ converges when p satisfies the assumption of Theorem 4.1. The concentration statement then follows from Markov's inequality. $\square$

5 Norm Convergence for Directed Kikuchi Matrices

In this section, we establish Theorem 1.7. In the previous section, we expanded the trace moments of the Kikuchi matrix Y_n in (10). The goal was to count the number of paired Kikuchi walks (e, π) , where π satisfies $\operatorname{supp}(e_i) = \operatorname{supp}(e_j)$ for every $ij \in E(\pi)$. For the directed Kikuchi matrix X_n , Wick's formula imposes a stronger “mirroring” property (Lemma 2.11). A *paired directed Kikuchi walk* of length p is a tuple (e, π) such that $e \in \mathcal{W}_p$ and $e_i = e_j^*$ for all $ij \in E(\pi)$. Then by (11), the p -th moment of X_n matches (up to an explicit normalization) the number of paired directed Kikuchi walks of length p .

Recall that the proof of weak convergence relies on the fact that defect-0 walks have transports of excess size 0 and cost 0 (Lemma 3.6). We now prove a generalization of this result.

Lemma 5.1. *Any defect- Δ paired directed Kikuchi walk has a Kikuchi transport of excess size Δ and cost at most 3Δ .*

Proof. Let (e, π) be a defect- Δ paired directed Kikuchi walk. Let $\rho_0(u) := u$. For each $t = 1, \dots, 2p-1$,

1. Suppose that $t+1$ is an opening index with closing j . For any $u \in T_{t+1}$ such that $\rho_{t-1}(\phi_t(u)) \in \rho_{i-1}(T_i)$ for some opening $i \leq t$ whose pair crosses $(t+1, j)$, let $\rho_t(u)$ be a fresh new label. For every other $u \in U_t$, let $\rho_t(u) := \rho_{t-1}(\phi_t(u))$.
2. Suppose that $t+1$ is a closing index with opening i . Then define

$$\rho_t(u) := \begin{cases} \rho_{i-1}(\phi_i(u)) & \text{if } u \in T_{t+1}, \\ \rho_{t-1}(\phi_t(u)) & \text{if } u \in U_t \setminus T_{t+1}. \end{cases}$$

Finally, we set $\rho_{2p}(u) := \rho_{2p-1}(\phi_{2p}(u))$. We show that this algorithm constructs a Kikuchi transport $\rho = (\rho_0, \dots, \rho_{2p})$ of excess size at most Δ and cost at most 3Δ . By construction, Steps 1 and 2 ensure that both compatibility conditions of Definition 2.12 are satisfied:

$$\rho_{i-1}(T_i) = \rho_{j-1}(T_j) \quad \forall ij \in E(\pi) \quad \text{and} \quad \rho_{i-1}(T_i) \cap \rho_{i'-1}(T_{i'}) = \emptyset \quad \forall \{ij, i'j'\} \in E(G_\pi).$$

For any label a , consider

$$O_t(a) := \{i \leq t : a \in \rho_{i-1}(T_i) \text{ and } \exists j > t, ij \in E(\pi)\}.$$

The crossing compatibility condition directly implies:

Claim 5.2. *For any t and a , out of all the openings in $O_t(a)$, the one that closes first is $\max O_t(a)$.*

Claim 5.3. *For any $t \geq 1$, the map ρ_t is injective. Furthermore, if $u \in U_t$ is such that $a := \rho_{t-1}(\phi_t(u))$ satisfies $O_t(a) \neq \emptyset$, then $u \in H_i$ and $a = \rho_{i-1}(\phi_i(u))$, where $i := \max O_t(a)$.*

Proof of Claim 5.3. We proceed by induction on t . The statements are immediate for $t = 0$, as $\rho_0 = \text{id}$ and $O_0(a) = \emptyset$ for all labels a . Now suppose that the claims hold for $0, \dots, t-1$. We first prove the “Furthermore” part. Let $u \in U_t$, $a = \rho_{t-1}(\phi_t(u))$, and suppose that $O_t(a) \neq \emptyset$. There are two cases:

- Suppose that t is an opening. We either have $O_t(a) = O_{t-1}(a)$ if $a \notin \rho_{t-1}(T_t)$, or $O_t(a) = O_{t-1}(a) \cup \{t\}$ if $a \in \rho_{t-1}(T_t)$. By injectivity of ρ_{t-1} , these two cases are equivalent respectively to $u \notin H_t$ and $u \in H_t$. In the former case, $a = \rho_{t-1}(u) = \rho_{t-2}(\phi_{t-1}(u))$, so the claim follows by induction. In the latter case, $t = \max O_t(a)$ and the claim is immediate.
- Suppose that t is a closing, with opening k . Assume first $u \notin H_t$. By the same reasoning as in the opening case, $a = \rho_{t-2}(\phi_{t-1}(u))$. We claim that $k \notin O_{t-1}(a)$: otherwise, Claim 5.2 shows that $k = \max O_{t-1}(a)$, which by induction would imply that $u \in H_k = T_t$, contradicting $u \in U_t$. Hence, $O_t(a) = O_{t-1}(a)$ and the claim follows by induction. Next, assume $u \in H_t$. Then by the construction in Step 2, $a = \rho_{k-1}(\phi_k(\phi_t(u))) = \rho_{k-1}(u)$ (using $\phi_k = \phi_t^{-1}$). By crossing compatibility, we must have $O_{k-1}(a) = O_t(a)$. In particular, $O_{k-1}(a) \neq \emptyset$ and therefore $a = \rho_{k-2}(\phi_{k-1}(u))$; the claim then follows from the induction hypothesis at step $k-1$.

The construction of ρ_t when $t+1$ is an opening clearly inherits injectivity from ρ_{t-1} . Suppose that $t+1$ is a closing index with opening i . The labels assigned by Step 2 to T_{t+1} are pairwise distinct, so it remains to show that none of these labels can occur outside of T_{t+1} . For any $a \in \rho_{i-1}(T_i)$, we have $i \in O_t(a)$, so $i = \max O_t(a)$ by Claim 5.2 and so any $u \in U_t$ such that $\rho_{t-1}(\phi_t(u)) \in \rho_{i-1}(T_i)$ must satisfy $u \in H_i = T_{t+1}$ by induction. This concludes the induction step. $\square$

For the transport size, charge every replacement of a label $a = \rho_{t-1}(\phi_t(u))$ in Step 1 to $(u, \max O_t(a))$. By Claim 5.2, the pair $ij \in E(\pi)$ containing $i := \max O_t(a)$ is the first pair in $O_t(a)$ to close. After its replacement, the label a disappears from the walk, and its first possible reappearance is at the closing of an opening in $O_t(a)$. Hence, a cannot reappear before step j , while after j , the pair ij is no longer open. In summary, every pair (u, i) can be charged at most once, and the pairs that are charged must satisfy that i is an opening index and $u \in H_i$ (by Claim 5.3). When $u \notin U_0$, the last occurrence of u in the walk is a closing, and its matching opening index cannot be charged. Thus, each u can be charged at most $\delta(u)$ times, and the total number of replacements in Step 1 is at most $\sum_u \delta(u) = \Delta$.

For the transport cost, by the above argument, the number of switches incurred by Step 1 is at most Δ . Let $B(u)$ be the number of switches of u incurred during Step 2. At every such switch occurring at a closing index $t+1$ with matching opening i , we have $u \in T_{t+1} = H_i$, so $B(u) \leq \delta(u) + 1$. Next, we claim that if $B(u) > 0$, then u must incur at least one switch in Step 1. Assuming this claim, the total number of switches due to Step 2 is at most $\sum_{u: B(u) > 0} (\delta(u) + 1) \leq 2\Delta$, as desired.

To prove the remaining claim, let t be the first iteration for which Step 2 incurs a switch at u . Let $t+1$ have opening index i , and let $a := \rho_{i-1}(\phi_i(u))$ be the new label. Consider the last iteration

$i \leq s < t$ for which $a \in \rho_{s-1}(U_{s-1}) \setminus \rho_s(U_s)$, and let $v \in U_s$ be defined by $a = \rho_{s-1}(\phi_s(v))$. Then $\max O_s(a) = i$ (otherwise, the maximizing pair would close before the pair of i and contradict the maximality of s), so by Claim 5.3, $a = \rho_{i-1}(\phi_i(v))$. By injectivity, this implies $v = u$. By minimality of t , the label of u at time s could not have been replaced at a Step 2, so it must be at a Step 1. $\square$

Lemma 5.4. *For any fixed $D_1, \dots, D_{2p} \in \binom{[\ell+\Delta]}{\ell}$,*

$$\sum_{\pi \in \mathcal{P}(2p)} |\{\varphi \in \text{Hom}(G_\pi, K(\ell + \Delta, r)) : \varphi(ij) \subseteq D_i \cap D_j \ \forall ij \in E(\pi)\}| \leq \binom{\ell}{r}^p \tau(\Sigma^{2p}).$$

Proof. Let $\{S_a : a \in \binom{[\ell+\Delta]}{r}\}$ be a $K(\ell + \Delta, r)$ -independent semicircular system, and

$$Q_i := \sum_{a \in \binom{D_i}{r}} S_a \quad \forall i \in [2p].$$

Each Q_i is a sum of a $K(\ell, r)$ -independent semicircular system, so $\tau(Q_i^{2p}) = \binom{\ell}{r}^p \tau(\Sigma^{2p})$. Further, by Lemma 2.8, the left-hand side of Lemma 5.4 is precisely $\tau(Q_1 \cdots Q_{2p})$. By Hölder's inequality (Lemma 2.3), we obtain

$$\tau(Q_1 \cdots Q_{2p}) \leq \prod_{i=1}^{2p} \tau(Q_i^{2p})^{\frac{1}{2p}} = \binom{\ell}{r}^p \tau(\Sigma^{2p}),$$

which concludes the proof. $\square$

Proof of Theorem 1.7. We count the triples $(\mathbf{e}, \pi, \boldsymbol{\rho})$ satisfying the conclusion of Lemma 5.1. By existence of the labeling, this upper bounds the number of paired directed Kikuchi walks (that the trace moments count exactly).

Fix the defect $\Delta \in \{0, \dots, rp\}$ and a set of switches $\sigma \subseteq [2p] \times [\ell + \Delta]^2$ from Definition 2.13, encoded as triples $(t, \rho_{t-1}(\phi_t(u)), \rho_t(u))$, of size $|\sigma| \leq 3\Delta$. There are at most $(2p(\ell + \Delta)^2 + 1)^{3\Delta}$ choices for this set. By construction, the switches determine the labels of the states $(\rho_i(U_i))_{i \in [2p]}$. Given that information, Lemma 5.4 applied to $D_i = \rho_{i-1}(U_{i-1})$ then implies that the total number of remaining possibilities for the pairing $\pi \in \mathcal{P}(2p)$ and the sequence $(\rho_{i-1}(T_i))_{i \in [2p]}$ is at most $\binom{\ell}{r}^p \tau(\Sigma^{2p})$. Finally, fix the number of coincidences δ_j at the j -th opening index.

We construct the maps ρ_i and the hyperedges e_i for $i = 1, \dots, 2p$ from left to right. At the j -th opening index $i \in [2p]$, the tail indices T_i can be obtained from $\rho_{i-1}(T_i)$ using injectivity of ρ_{i-1} . For the head H_i , we must choose δ_j vertices among the $\leq 2pr$ vertices that have appeared previously in a tail, and $r - \delta_j$ vertices among the $n - \ell$ vertices outside of the current state. At a closing index, the head and tail are uniquely determined from the matching opening index, using the mirroring property.⁶ The map ρ_i is then obtained from ρ_{i-1} using the set of switches. Overall,

$$\begin{aligned} \mathbb{E} \text{tr } X_n^{2p} &\leq \frac{\binom{\ell}{r}^p}{d_n^p} \tau(\Sigma^{2p}) \sum_{\Delta=0}^{rp} (2p(\ell + \Delta)^2 + 1)^{3\Delta} \sum_{\delta_1 + \dots + \delta_p = \Delta} \prod_{j=1}^p \binom{2pr}{\delta_j} \binom{n-\ell}{r-\delta_j} \\ &= \tau(\Sigma^{2p}) \sum_{\Delta=0}^{rp} (2p(\ell + \Delta)^2 + 1)^{3\Delta} \sum_{\delta_1 + \dots + \delta_p = \Delta} \prod_{j=1}^p \frac{\binom{2pr}{\delta_j} \binom{n-\ell}{r-\delta_j}}{\binom{n-\ell}{r}} \end{aligned}$$

⁶This is where we use the fact that the Kikuchi matrix is directed.

The combinatorial coefficient can be bounded by

$$\sum_{\delta_1 + \dots + \delta_p = \Delta} \prod_{j=1}^p \frac{\binom{2pr}{\delta_j} \binom{n-\ell}{r-\delta_j}}{\binom{n-\ell}{r}} \leq \sum_{\delta_1 + \dots + \delta_p = \Delta} \prod_{j=1}^p \left(\frac{4pr^2}{n-\ell} \right)^{\delta_j} \leq \left(\frac{4p^2 r^2}{n-\ell} \right)^{\Delta}.$$

Every term in the sum under the power Δ is then bounded by

$$(2p(\ell + \Delta)^2 + 1)^3 \cdot \left(\frac{4p^2 r^2}{n-\ell} \right) \lesssim \frac{p^5 \max\{\ell, rp\}^6 r^2}{n-\ell} \xrightarrow{n \rightarrow \infty} 0,$$

provided, for example, that $r \leq C$ and $\max\{p, \ell\} \leq n^c$ for constants $c, C > 0$. The desired upper bound then follows from Markov's inequality, and the matching lower bound follows from weak convergence. $\square$

6 Convergence to the q -Gaussian system

In this section, we prove Theorem 1.8. For a graph Γ on N vertices, we denote by $\lambda_1(\Gamma) \geq \dots \geq \lambda_N(\Gamma)$ the eigenvalues of its adjacency matrix. We start by computing the spectrum of the adjacency matrix of the Kneser graphs.

Lemma 6.1 (See [GR01, Theorem 9.4.3]). *For $\ell > 2r$, the adjacency matrix of $K(\ell, r)$ has eigenvalues*

$$(-1)^i \binom{\ell - r - i}{r - i} \quad i = 0, \dots, r,$$

and the eigenvalue for $i = 0$ has multiplicity 1.

We will use Lemma 6.1 to show that the Kneser graphs are quasirandom in the sense of Chung-Graham-Wilson [CGW89]. One equivalent condition for quasirandomness is the convergence of all homomorphism counts to those of an Erdős-Rényi random graph with matching density.

Lemma 6.2 (See [CGW89]). *Let $(\Gamma_\ell)_{\ell \geq 1}$ be a sequence of regular graphs whose edge density converges to $q \in [0, 1]$. Suppose that*

$$\max_{i \geq 2} |\lambda_i(\Gamma_\ell)| = o(|V(\Gamma_\ell)|).$$

Then for every fixed graph G ,

$$\frac{|\text{Hom}(G, \Gamma_\ell)|}{|V(\Gamma_\ell)|^{|V(G)|}} \xrightarrow{\ell \rightarrow \infty} q^{|E(G)|}.$$

We use the following convention in this section: denoting by K_2 the graph consisting of a single edge, the edge density of Γ is $|\text{Hom}(K_2, \Gamma)|/|V(\Gamma)|^2$.

Lemma 6.3 (See [BS92, Section 3]). *Let S be a q -Gaussian variable (Definition 2.6) for $q \in [0, 1]$. Then there exists a probability distribution μ_q with full support on*

$$\left[-\frac{2}{\sqrt{1-q}}, \frac{2}{\sqrt{1-q}} \right]$$

such that $\tau(S^p) = \int x^p d\mu_q(x)$ for all $p \geq 1$. In particular, we have $\|S\| = \frac{2}{\sqrt{1-q}}$.

6.1 Weak convergence

Lemma 6.4. *Let $\alpha > 0$ and Σ be as in Theorem 1.5. Suppose that $r = r(\ell)$ is such that $r^2/\ell \rightarrow \alpha$ as $\ell \rightarrow \infty$. Then Σ converges weakly as $\ell \rightarrow \infty$ to the $e^{-\alpha}$ -Gaussian system with k variables.*

Proof. We compute the limit of every monomial indexed by $\mathbf{w} \in [k]^p$. The moment formula for a Γ -independent semicircular system (Lemma 2.8) gives

$$\tau(\Sigma_{w_1} \cdots \Sigma_{w_p}) = \binom{\ell}{r}^{-p/2} \sum_{\pi \in \mathcal{P}(\mathbf{w})} |\text{Hom}(G_\pi, K(\ell, r))|.$$

The Kneser graph $K(\ell, r)$ is regular, with edge density

$$\frac{\binom{\ell-r}{r}}{\binom{\ell}{r}} = \frac{((\ell-r)!)^2}{(\ell-2r)! \ell!} \xrightarrow{\ell \rightarrow \infty} e^{-\alpha},$$

which can be verified by applying Stirling's formula. Moreover, Lemma 6.1 gives

$$\frac{1}{\binom{\ell}{r}} \max_{i \geq 2} |\lambda_i(K(\ell, r))| \xrightarrow{\ell \rightarrow \infty} 0.$$

Therefore, by Lemma 6.2, all homomorphism densities converge to those of an Erdős-Rényi random graph, i.e.,

$$\tau(\Sigma_{w_1} \cdots \Sigma_{w_p}) \xrightarrow{\ell \rightarrow \infty} \sum_{\pi \in \mathcal{P}(\mathbf{w})} e^{-\alpha |E(G_\pi)|}.$$

This is exactly the moment formula for the $e^{-\alpha}$ -Gaussian system (Lemma 2.7). $\square$

6.2 Norm convergence for quasirandom independence graphs

In this section, given a graph Γ , we let $(L_u)_{u \in V(\Gamma)}$ be the creation operators associated to the Γ -independent semicircular system, and define

$$\Sigma(\Gamma) := \frac{1}{\sqrt{|V(\Gamma)|}} \sum_{u \in V(\Gamma)} (L_u + L_u^*)$$

to be its normalized sum. We denote by $\omega(\Gamma)$ the maximum clique size of a graph Γ . The main result of this subsection is:

Proposition 6.5. *Let $(\Gamma_\ell)_{\ell \geq 1}$ be a sequence of regular graphs whose edge density converges to $q \in [0, 1)$. Suppose that*

$$\omega(\Gamma_\ell) \max\{0, \lambda_2(\Gamma_\ell)\} + |\lambda_{|V(\Gamma_\ell)|}(\Gamma_\ell)| = o(|V(\Gamma_\ell)|) \quad \text{as } \ell \rightarrow \infty. \quad (16)$$

Then

$$\|\Sigma(\Gamma_\ell)\| \xrightarrow{\ell \rightarrow \infty} \frac{2}{\sqrt{1-q}}.$$

Recall that by Lemma 6.3, the limit is precisely the edge of the support of the q -Gaussian distribution. The proof of Proposition 6.5 divides into a lower bound, which follows from weak convergence, and an upper bound. For the latter part, existing Khintchine inequalities for Γ -independent semicircular systems [CM26, OTY25] are not sufficiently sharp.

Instead, we exploit the fact that our graphs have small second eigenvalue and use the following bound:

Proposition 6.6 ([BKN⁺26b]). *For any graph Γ and $q \in [0, 1)$,*

$$\|\Sigma(\Gamma)\| \leq \frac{2}{\sqrt{1-q}} \sqrt{1 + \frac{\omega(\Gamma) \max\{0, \lambda_1(\Gamma - qJ)\}}{|V(\Gamma)|}},$$

where J denotes the all-ones matrix.

The proof of Proposition 6.6 can be found in the parallel work [BKN⁺26b, Corollary 5.3].

Proof of Proposition 6.5 from Proposition 6.6. The same argument as in Lemma 6.4 shows weak convergence: by Lemma 2.8,

$$\tau(\Sigma(\Gamma_\ell)^p) = \frac{1}{|V(\Gamma_\ell)|^{p/2}} \sum_{\pi \in \mathcal{P}(p)} |\text{Hom}(G_\pi, \Gamma_\ell)|.$$

The assumption (16) is stronger than that of Lemma 6.2, so we obtain

$$\tau(\Sigma(\Gamma_\ell)^p) \xrightarrow{\ell \rightarrow \infty} \sum_{\pi \in \mathcal{P}(p)} q^{|E(G_\pi)|} = \tau(S^p),$$

where S is a single q -Gaussian noncommutative random variable defined through Theorem 2.4 (see Lemma 2.7 for the associated moment formula). Weak convergence directly implies a lower bound on the norm. Indeed, by (9), for any fixed $p \geq 1$,

$$\|\Sigma(\Gamma_\ell)\| \geq \tau(\Sigma(\Gamma_\ell)^{2p})^{\frac{1}{2p}} \xrightarrow{\ell \rightarrow \infty} \tau(S^{2p})^{\frac{1}{2p}}.$$

By Lemmas 2.2 and 6.3, taking the $p \rightarrow \infty$ limit on the right-hand side yields

$$\liminf_{\ell \rightarrow \infty} \|\Sigma(\Gamma_\ell)\| \geq \frac{2}{\sqrt{1-q}}.$$

To prove the matching upper bound, we use Proposition 6.6. Let q_ℓ be the edge density of Γ_ℓ , so that q_ℓ converges to q . Since Γ_ℓ is regular, we have $\max\{0, \lambda_2(\Gamma_\ell)\} = \lambda_1(\Gamma_\ell - q_\ell J)$. Combining Proposition 6.6 with the assumption (16), we obtain

$$\limsup_{\ell \rightarrow \infty} \|\Sigma(\Gamma_\ell)\| \leq \frac{2}{\sqrt{1-q}}.$$

This concludes the proof. □

6.3 Applications to the Kneser graph

The graph $K(\ell, r)$ has $N_\ell := \binom{\ell}{r}$ vertices, and a clique is a collection of disjoint sets of $[\ell]$ of size r , so its maximum clique size is $\lfloor \ell/r \rfloor$. Combining this with the eigenvalues of the Kneser graph from Lemma 6.1,

$$\begin{aligned} \frac{\omega(K(\ell, r)) \max\{0, \lambda_2(K(\ell, r))\}}{N_\ell} &= \frac{\lfloor \ell/r \rfloor \binom{\ell-r-2}{r-2}}{\binom{\ell}{r}} \leq \frac{r-1}{\ell-1}, \\ \frac{|\lambda_{N_\ell}(K(\ell, r))|}{N_\ell} &= \frac{\binom{\ell-r-1}{r-1}}{\binom{\ell}{r}} \leq \frac{r}{\ell}. \end{aligned}$$

Proposition 6.7. $\|\Sigma(K(\ell, r))\| \asymp \max \left\{ 1, \frac{\sqrt{\ell}}{r} \right\}$

Proof. The edge density of $K(\ell, r)$ is $q := \binom{\ell-r}{r} / \binom{\ell}{r}$, from which one can verify that

$$\frac{2}{\sqrt{1-q}} \asymp \max \left\{ 1, \frac{\sqrt{\ell}}{r} \right\}.$$

Proposition 6.6 then implies the desired upper bound. For the lower bound, the normalization $\tau(\Sigma(K(\ell, r))^2) = 1$ gives the constant term, and the inequality (5) gives the other one. $\square$

Proof of Theorem 1.8. Lemma 6.4 combined with Theorem 3.2 directly implies the weak convergence assertion in Theorem 1.8. Theorem 1.7 shows that $\|X_n\| \rightarrow \|\Sigma(K(\ell, r))\|$ in probability as $n \rightarrow \infty$. In the regime $r^2/\ell \rightarrow \alpha$, condition (16) is verified and Proposition 6.5 implies that

$$\|\Sigma(K(\ell, r))\| \xrightarrow{\ell \rightarrow \infty} \frac{2}{\sqrt{1-e^{-\alpha}}},$$

so $\|X_n\|$ also converges in probability to that constant in the double limit $n \rightarrow \infty$, then $\ell \rightarrow \infty$. We generalize this to any polynomial P . Letting $\rho = 2/\sqrt{1-e^{-\alpha}}$ and S be a $e^{-\alpha}$ -Gaussian variable, Lemmas 6.3 and 2.2 give

$$\|P(S)\| = \lim_{p \rightarrow \infty} \tau((P(S)^*P(S))^p)^{\frac{1}{2p}} = \lim_{p \rightarrow \infty} \left(\int |P(x)|^{2p} d\mu_{e^{-\alpha}} \right)^{\frac{1}{2p}} = \max_{|x| \leq \rho} |P(x)|.$$

Since $\|P(X_n)\| \leq \max_{|x| \leq \|X_n\|} |P(x)|$, we obtain for every fixed $\varepsilon > 0$

$$\limsup_{\ell \rightarrow \infty} \limsup_{n \rightarrow \infty} \Pr(\|P(X_n)\| > \|P(S)\| + \varepsilon) = 0.$$

The matching lower bound follows from weak convergence: for every fixed $p \geq 1$,

$$\|P(X_n)\|^{2p} \geq \text{tr}((P(X_n)^*P(X_n))^p),$$

which converges in the double limit in probability to $\tau((P(S)^*P(S))^p)$. Taking the further limit $p \rightarrow \infty$ and using Lemma 2.2 concludes the proof of Theorem 1.8. $\square$

References

- [Ban26] Afonso S. Bandeira. Random matrices, intrinsic freeness, and sharp non-asymptotic inequalities. In *Proceedings of the International Congress of Mathematicians (ICM)*, pages 289–309, 2026. 2
- [BBvH23] Afonso S. Bandeira, March T. Boedihardjo, and Ramon van Handel. Matrix concentration inequalities and free probability. *Inventiones Mathematicae*, 234(1):419–487, 2023. 2, 6
- [BCSvH26] Afonso S. Bandeira, Giorgio Cipolloni, Dominik Schröder, and Ramon van Handel. Matrix concentration inequalities and free probability II. Two-sided bounds and applications. *Communications of the American Mathematical Society*, 6:896–946, 2026. 1, 2, 4
- [BKMR25] Afonso S. Bandeira, Anastasia Kireeva, Antoine Maillard, and Almut Rödter. Randomstrasse101: Open problems of 2024. *arXiv:2504.20539*, 2025. 2
- [BKN⁺26a] Afonso S. Bandeira, Dmitriy Kunisky, Petar Nizić-Nikolac, Lucas Pesenti, and Robert Wang. The hypergraph Moore bound. *arXiv:2607.14068v2*, 2026. 7

- [BKN⁺26b] Afonso S. Bandeira, Dmitriy Kunisky, Petar Nizić-Nikolac, Lucas Pesenti, and Robert Wang. Matrix concentration and equivalent operators on Fock spaces. *arXiv*, 2026. 7, 22
- [BKS97] Marek Bożejko, Burkhard Kümmerer, and Roland Speicher. q -Gaussian processes: Non-commutative and classical aspects. *Communications in Mathematical Physics*, 185:129–154, 1997. 6
- [BS91] Marek Bożejko and Roland Speicher. An example of a generalized Brownian motion. *Communications in Mathematical Physics*, 137(3):519–531, 1991. 6, 10
- [BS92] Marek Bożejko and Roland Speicher. An example of a generalized Brownian motion II. In *Quantum Probability and Related Topics*, volume 7, pages 67–77. World Scientific, 1992. 6, 20
- [BS94] Marek Bożejko and Roland Speicher. Completely positive maps on Coxeter groups, deformed commutation relations, and operator spaces. *Mathematische Annalen*, 300:97–120, 1994. 3, 9
- [CC21] Ian Charlesworth and Benoît Collins. Matrix models for ε -free independence. *Archiv der Mathematik*, 116:585–600, 2021. 4, 7
- [CC26] Martijn Caspers and Enli Chen. Strongly convergent matrix models for q -Gaussian algebras. *arXiv:2607.15100*, 2026. 8
- [CGTvH26] Chi-Fang Chen, Jorge Garza-Vargas, Joel A. Tropp, and Ramon van Handel. A new approach to strong convergence. *Annals of Mathematics*, 203(2):555–602, 2026. 3
- [CGvH26] Chi-Fang Chen, Jorge Garza-Vargas, and Ramon van Handel. A new approach to strong convergence II. The classical ensembles. *Geometric and Functional Analysis*, 36:831–890, 2026. 3, 4, 7
- [CGW89] Fan R. K. Chung, Ronald L. Graham, and Richard M. Wilson. Quasi-random graphs. *Combinatorica*, 9:345–362, 1989. 20
- [CM26] Benoît Collins and Akihiro Miyagawa. Operator-valued Khintchine inequality for ε -free semicircles. *Bulletin of the London Mathematical Society*, 58(1):e70229, 2026. 21
- [COY24] Guillaume Cébron, Patrick Oliveira Santos, and Pierre Youssef. Graphon-theoretic approach to central limit theorems for ε -independence. *arXiv:2411.13062*, 2024. 6
- [FB70] Uriel Frisch and Richard Bourret. Parastochastics. *Journal of Mathematical Physics*, 11:364–390, 1970. 6, 7
- [GKM22] Venkatesan Guruswami, Pravesh K. Kothari, and Peter Manohar. Algorithms and certificates for Boolean CSP refutation: “smoothed is no harder than random”. In *Symposium on Theory of Computing (STOC)*, pages 678–689, 2022. 7
- [GR01] Chris Godsil and Gordon Royle. *Algebraic Graph Theory*. Springer, 2001. 20
- [GS14] Alice Guionnet and Dimitri Shlyakhtenko. Free monotone transport. *Inventiones Mathematicae*, 197:613–661, 2014. 7
- [Has20] Matthew B. Hastings. Classical and quantum algorithms for tensor principal component analysis. *Quantum*, 4:237, 2020. 1
- [HKM23] Jun-Ting Hsieh, Pravesh K. Kothari, and Sidhanth Mohanty. A simple and sharper proof of the hypergraph Moore bound. In *Symposium on Discrete Algorithms (SODA)*, pages 2324–2344, 2023. 7
- [HT05] Uffe Haagerup and Steen Thorbjørnsen. A new application of random matrices: $\text{Ext}(C_{\text{red}}^*(F_2))$ is not a group. *Annals of Mathematics*, 162(2):711–775, 2005. 3, 7
- [KX26] Pravesh K. Kothari and Jeff Xu. Smooth trade-off for tensor PCA via sharp bounds for Kikuchi matrices. In *Symposium on Discrete Algorithms (SODA)*, pages 2617–2632, 2026. 1, 2, 5

- [LP91] Françoise Lust-Piquard and Gilles Pisier. Non commutative Khintchine and Paley inequalities. *Arkiv för Matematik*, 29:241–260, 1991. [2](#)
- [Man24] Peter Manohar. *New Spectral Techniques in Algorithms, Combinatorics, and Coding Theory: The Kikuchi Matrix Method*. PhD thesis, Carnegie Mellon University, 2024. [1](#)
- [Mło04] Wojciech Młotkowski. Λ -free probability. *Infinite Dimensional Analysis, Quantum Probability and Related Topics*, 7(1):27–41, 2004. [4](#)
- [MR14] Andrea Montanari and Emile Richard. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems (NIPS)*, volume 27, 2014. [2](#)
- [MT26] Michael Magee and Joe Thomas. Strongly convergent unitary representations of right-angled Artin groups. *Duke Mathematical Journal*, 175:1305–1343, 2026. [4](#), [7](#)
- [NS06] Alexandru Nica and Roland Speicher. *Lectures on the combinatorics of free probability*. Cambridge University Press, 2006. [9](#)
- [OTY25] Patrick Oliveira Santos, Raghavendra Tripathi, and Pierre Youssef. Khintchine inequalities, trace monoids and Turán-type problems. *arXiv:2506.02517*, 2025. [3](#), [9](#), [21](#)
- [PX03] Gilles Pisier and Quanhua Xu. Non-commutative L^p -spaces. In *Handbook of the Geometry of Banach Spaces*, volume 2, pages 1459–1517. North-Holland, 2003. [9](#)
- [SH26] Alexander Schmidhuber and Matthew B. Hastings. A spectral proof of the hypergraph Moore bound. *arXiv:2607.26028*, 2026. [7](#)
- [Śni01] Piotr Śniady. Gaussian random matrix models for q -deformed Gaussian variables. *Communications in Mathematical Physics*, 216:515–537, 2001. [7](#)
- [Spe92] Roland Speicher. A non-commutative central limit theorem. *Mathematische Zeitschrift*, 209:55–66, 1992. [6](#), [7](#)
- [SW16] Roland Speicher and Janusz Wysoczański. Mixtures of classical and free independence. *Archiv der Mathematik*, 107:445–453, 2016. [4](#), [10](#)
- [vH26] Ramon van Handel. Strong convergence: A short survey. In *Proceedings of the International Congress of Mathematicians (ICM)*, pages 145–165, 2026. [3](#)
- [WEM19] Alexander S. Wein, Ahmed El Alaoui, and Cristopher Moore. The Kikuchi hierarchy and tensor PCA. In *Symposium on Foundations of Computer Science (FOCS)*, pages 1446–1468, 2019. [1](#), [2](#)