

A non-backtracking method for long matrix and tensor completion

Ludovic Stephan

École Polytechnique Fédérale de Lausanne (EPFL), IdePHICS Lab

LUDOVIC.STEPHAN@EPFL.CH

Yizhe Zhu

Department of Mathematics, University of California Irvine

YIZHE.ZHU@UCI.EDU

Editors: Shipra Agrawal and Aaron Roth

Abstract

We consider the problem of low-rank rectangular matrix completion in the regime where the matrix M of size $n \times m$ is “long”, i.e., the aspect ratio m/n diverges to infinity. Such matrices are of particular interest in the study of tensor completion, where they arise from the unfolding of a low-rank tensor. In the case where the sampling probability is $\frac{d}{\sqrt{mn}}$, we propose a new spectral algorithm for recovering the singular values and left singular vectors of the original matrix M based on a variant of the standard non-backtracking operator of a suitably defined bipartite weighted random graph, which we call a *non-backtracking wedge operator*. When d is above a Kesten-Stigum-type sampling threshold, our algorithm recovers a correlated version of the singular value decomposition of M with quantifiable error bounds. This is the first result in the regime of bounded d for weak recovery and the first for weak consistency when $d \rightarrow \infty$ arbitrarily slowly without any polylog factors. As an application, for low-CP-rank orthogonal k -tensor completion, we efficiently achieve weak recovery with sample size $O(n^{k/2})$ and weak consistency with sample size $\omega(n^{k/2})$. A similar result is obtained for low-multilinear-rank tensor completion with $O(n^{k/2})$ many samples.

Keywords: Tensor completion, long matrix completion, non-backtracking operator

1. Introduction

Matrix completion is the problem of reconstructing a matrix from a (usually random) subset of entries, leveraging prior structural knowledge such as rank and incoherence. When the two dimensions n, m of the matrix M are comparable, this problem has been widely studied in the past decades, with a fairly general pattern emerging. Roughly speaking, if the ground truth matrix is low-rank and eigenvectors are sufficiently incoherent (or delocalized), sampling $\Omega(n \log n)$ many entries uniformly at random with a certain efficient optimization procedure, one can exactly recover the ground truth matrix (Keshavan et al., 2009, 2010; Candes and Tao, 2010; Candes and Plan, 2010; Recht, 2011; Candes and Recht, 2012; Jain et al., 2013). This is not the case when the sample size is only $O(n)$, where full completion is provably impossible under uniform sampling. However, an approximation of the ground truth matrix is still possible (Heiman et al., 2014; Gamarnik et al., 2017; Brito et al., 2022) with efficient algorithms, and more refined estimates of the eigenvectors were studied in (Bordenave et al., 2022b).

In recent years, attention has shifted to “long” matrix models where the row and column sizes of the matrix are not proportional. This is the case, for example, in the spiked long matrix and tensor models (Montanari and Wu, 2022; Ding et al., 2020; Ben Arous et al., 2023; Montanari and Richard, 2014). This situation also occurs in bipartite graph clustering and community detection, where the goal is to recover the community structure in the smaller vertex set (Florescu and Perkins, 2016;

(Ndaoud et al., 2021; Cai et al., 2021; Braun and Tyagi, 2023; Braun, 2023) for a bipartite stochastic block model, and in subspace recovery problems (Cai et al., 2021; Zhou et al., 2021).

Our work is closely related to tensor completion, which is the higher-order analog of matrix completion. Indeed, tensor problems are highly related to their long matrix counterparts through the unfolding operation for an order- k tensor of size n^k :

$$\text{unfold}_{a,b} : \mathbb{R}^{n^k} \rightarrow \mathbb{R}^{n^a \times n^b},$$

where $a + b = k$. For example, an order-3 tensor of size n yields an unfolded matrix of size $n \times n^2$; in general, uneven unfoldings (i.e., with $a \neq b$) of a tensor give rise to long rectangular matrices. In those tensor problems, the ground truth is assumed to have a low CP-rank, which implies that the unfolded matrix is low-rank. In low-rank k -tensor completion, the best known polynomial-time algorithms require $\tilde{O}(n^{k/2})$ revealed entries (Jain and Oh, 2014; Montanari and Sun, 2018; Xia and Yuan, 2019; Cai et al., 2022), while information theoretically $\tilde{O}(n)$ suffices (Ghadermarzy et al., 2019; Harris and Zhu, 2021; Harris et al., 2023). In (Barak and Moitra, 2022), the authors showed lower bounds for a family of algorithms based on the sum-of-squares hierarchy for order-3 noisy tensor completion, where $\Omega(n^{3/2})$ many samples are needed, by connecting it to the literature on refuting random 3-SAT. They also conjectured such a sample size is needed for any polynomial-time algorithm. This is aligned with other tensor estimation problems where a diverging statistical-to-computational gap is expected (Montanari and Richard, 2014; Wein et al., 2019; Jagannath et al., 2020; Arous et al., 2020; Dudeja and Hsu, 2024; Auddy and Yuan, 2022). However, so far no known work has reached the exact threshold $\Omega(n^{3/2})$, only approaching it up to polylog factors, in a similar fashion to the matrix completion phase diagram described above. For exact recovery results on tensor completion, all the methods in the literature are either convex optimization or spectral initialization with iterative optimization steps. Our method achieves weak consistency without any log factors, which could improve the sample size for spectral initialization in the work (Jain and Oh, 2014; Cai et al., 2022; Xia and Yuan, 2019; Xia et al., 2021).

The applications we mentioned above are related to the same question:

Can we improve the sample complexity for estimating a one-sided structure in a long matrix?

Specifically, we explore the use of the *non-backtracking operator* in this context. The non-backtracking operator is an important object in the study of spectral graph theory (Bass, 1992; Terras, 2010). Recently, it has been a powerful tool for understanding the spectrum of a very sparse random matrix (Bordenave et al., 2018; Stephan and Massoulié, 2022; Bordenave, 2020; Bordenave and Collins, 2019; Brito et al., 2022; Benaych-Georges et al., 2020; Alt et al., 2021b,a; Dumitriu and Zhu, 2022; Stephan and Zhu, 2022) and for designing efficient algorithms in community detection and matrix completion (Krzakala et al., 2013b; Bordenave et al., 2018; Stephan and Massoulié, 2022; Bordenave et al., 2022b; Stephan and Zhu, 2022; Zhu and Zhu, 2023). In the bipartite graph setting, the non-backtracking spectrum was analyzed in (Bruto et al., 2022; Dumitriu and Zhu, 2022, 2021) when the two sides have a bounded aspect ratio. For the spectrum of the adjacency matrix of a random bipartite graph with a diverging aspect ratio, only a few results were obtained in (Zhu, 2023; Dumitriu and Zhu, 2023) for random bipartite biregular graphs, and in (Guruswami et al., 2022) for the adjacency matrix of bipartite random graphs with random sign flips.

In this work, we aim to address this gap by defining a non-backtracking operator suitable for sparse matrix estimation with a diverging aspect ratio, which we call a *non-backtracking wedge matrix*. Our work sheds light on the fundamental limits and practical algorithms for sparse long matrix estimation with a focus on one-sided structure recovery. To the best of our knowledge, this is

the first result for a long $n \times m$ low-rank matrix completion with sampling probability $p = \frac{d}{\sqrt{mn}}$ in the regime of bounded d for weak recovery and the first result for weak consistency when $d \rightarrow \infty$ arbitrarily slowly without any polylog factors. These two types of recovery guarantees are well studied in the literature of community detection (Abbe, 2018). In fact, a direct adaptation of the bipartite non-backtracking operator we define here leads to a simpler spectral algorithm for the detection of the bipartite stochastic block model studied in (Florescu and Perkins, 2016). Below, We summarize our results informally. See Theorems 2 and 3 for the precise statements.

Theorem 1 (long matrix completion, informal) *With $O(\sqrt{mn})$ many samples, if the left singular vectors of a long $n \times m$ matrix ($m \gg n$) are sufficiently delocalized, and the singular values are above a specific Kesten-Stigum threshold, then the non-backtracking wedge operator can recover the ground truth singular values and find correlated estimates of the left singular vectors.*

The *Kesten-Stigum threshold* in Theorem 1 is analogous to a similar threshold for community detection in the stochastic block model, which characterizes a threshold of the signal-to-noise ratio for the existence of polynomial time algorithms to achieve weak recovery; see (Abbe, 2018) for more details. We discuss the Kesten-Stigum bound in matrix completion in Section 2.4.

We then apply our main results to the low-rank tensor completion problem under an orthogonality assumption on the r components in a tensor decomposition. This is the first efficient algorithm to achieve weak recovery of a low-rank tensor with $\Omega(n^{k/2})$ samples without any extra polylog(n) factors. In particular, an explicit $n^{k/2}$ threshold is obtained for rank-1 tensor completion when the entries are ± 1 . The weak recovery results can serve as a good spectral initialization step for other optimization procedures (Montanari and Sun, 2018; Jain and Oh, 2014; Cai et al., 2022). Unlike the common approach of unfolding a tensor to a matrix as square as possible (i.e., applying $\text{unfold}_{\lfloor k/2 \rfloor, \lceil k/2 \rceil}$ to T) in (Mu et al., 2014; Montanari and Sun, 2018), we unfold T to an $n \times n^{k-1}$ matrix in the most unbalanced way, and this allows us to directly estimate individual components of the tensor decomposition with sample complexity $O(n^{k/2})$. Note that simply applying the results in (Bordenave et al., 2022b) to $\text{unfold}_{\lfloor k/2 \rfloor, \lceil k/2 \rceil}(T)$ will have sample complexity $O(n^{\lceil k/2 \rceil})$, and one cannot directly estimate any individual component in T . Therefore, our algorithm has a better computational complexity for tensor completion for all integer values k . The long matrix completion result also has a direct application in low-multilinear-rank tensor completion. Detailed statements and discussion on tensor completion are given in Sections 2.5 and 2.6.

Although there is a strong connection between the detection problem for matrix completion and community detection in the stochastic block model (Angelini et al., 2015; Stephan and Massoulié, 2022; Bordenave et al., 2018, 2022b), the statistical-to-computational gap in the hypergraph stochastic block model behaves very differently from tensor completion (Pal and Zhu, 2021; Stephan and Zhu, 2022; Gu and Polyanskiy, 2023). In hypergraph community detection, efficient spectral methods can achieve weak recovery with the sample complexity down to the information-theoretical threshold up to a constant factor in all cases. However, in tensor completion, there is a conjectured diverging gap between the information-theoretical sample complexity and sample complexity for efficient algorithms (Barak and Moitra, 2022).

The non-backtracking wedge operator we defined in Section 2.1 is tailored for the situation where the underlying sparse bipartite graph has a diverging aspect ratio. It is similar in spirit to the non-backtracking operator for hypergraphs (Angelini et al., 2015; Liu et al., 2022; Stephan and Zhu, 2022), and the nomadic walk operator in (Mohanty et al., 2020; Guruswami et al., 2022) for (signed) random bipartite biregular graphs. The main difference compared to (Guruswami et al.,

2022; Mohanty et al., 2020) is that we work with a non-centered random matrix model without degree homogeneity, and we only count one-sided non-backtracking walks. In addition, a detailed eigenvector analysis is provided while (Guruswami et al., 2022; Mohanty et al., 2020) are only about extreme eigenvalues. In terms of tensor estimation with path-counting arguments, another relevant work is the estimation of a spiked tensor model under heavy-tail noise (Ding et al., 2020), where the authors use $n \times n$ symmetric matrices to count the number of self-avoiding walks of length $c \log n$. However, those matrices are usually more expensive to compute and require a new parameter (the length of the paths considered) that needs tuning in practice.

2. Setting and main results

2.1. Detailed setting for long matrix completion

We consider a rectangular matrix M of size $n \times m$, with aspect ratio $\alpha = m/n$. We write the singular value decomposition of M as

$$M = \sum_{i=1}^n \nu_i \phi_i \psi_i^\top,$$

where the ν_i are the singular values of M , and the $(\phi_i)_{i \in [r]}$ (resp. the $(\psi_i)_{i \in [r]}$) are an orthonormal family of its left (resp. right) singular vectors. We order the singular values of M in decreasing order, that is $\nu_1 \geq \dots \geq \nu_n$. Our observed matrix is of the form

$$A = \frac{\sqrt{mn}}{d} (X \circ M), \quad (1)$$

where X is an i.i.d Bernoulli matrix:

$$X_{ij} \stackrel{i.i.d}{\sim} \text{Ber} \left(\frac{d}{\sqrt{mn}} \right).$$

We will use the shortcut $p = d/\sqrt{mn}$ as the sampling probability. The scaling of A is chosen so that the expected value of A is exactly M . Our goal is to weakly recover the eigendecomposition:

$$MM^\top = \sum_{i=1}^n \nu_i^2 \phi_i \phi_i^\top.$$

The non-backtracking wedge matrix We can view the matrix A as the weighted biadjacency matrix of a bipartite graph $G = (V_1 \cup V_2, E)$ with $|V_1| = n$, $|V_2| = m$. We call V_1 a *left* vertex set and V_2 a *right* vertex set. A path of length 2 in G denoted by (e_1, e_2, e_3) is called a *wedge* if $e_1, e_3 \in V_1, e_2 \in V_2$, and $e_1 \neq e_3$. Define the set of *oriented wedges* \mathbf{E} as

$$\mathbf{E} = \{(e_1, e_2, e_3) \in V_1 \times V_2 \times V_1 : \{e_1, e_2\} \in E, \{e_2, e_3\} \in E, e_1 \neq e_3\}. \quad (2)$$

The *non-backtracking wedge* matrix B associated to G is defined on $\mathbf{E} \times \mathbf{E}$ by

$$B_{ef} = \mathbf{1}_{e \rightarrow f} A_{f_1 f_2} A_{f_3 f_2}, \quad \text{where} \quad \mathbf{1}_{e \rightarrow f} = \mathbf{1}_{e_3 = f_1} \mathbf{1}_{e_2 \neq f_2}. \quad (3)$$




Figure 1: Example of a non-backtracking path between two wedges (e_1, e_2, e_3) and (f_1, f_2, f_3) .

Equivalently, $\mathbf{1}_{e \rightarrow f}$ is one if and only if $(e_1, e_2, e_3, f_2, f_3)$ is a non-backtracking walk of length 4 in G . See Figure 1 for an example. Note that the non-backtracking wedge matrix B is not symmetric: its index set is the oriented wedge set defined in (2), and $B_{ef} \neq B_{fe}$ in (3).

For comparison, we recall the definition of the *non-backtracking operator* for weighted graphs as follows. Let A be a weighted adjacency matrix of a graph $G = (V, E)$. Define the oriented edge set \mathbf{E} for G as $\mathbf{E} = \{(i, j) : \{i, j\} \in E\}$. The *non-backtracking operator* B associated with A is a non-Hermitian operator of size $|\mathbf{E}| \times |\mathbf{E}|$. For any $(u, v), (x, y) \in \mathbf{E}$, B is defined as

$$B_{(u,v),(x,y)} = \begin{cases} A_{xy} & v = x, u \neq y, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

In the case of an $n \times m$ bipartite weighted graph defined in (1), the non-backtracking wedge operator only counts non-backtracking walks starting from a vertex in V_1 while the non-backtracking operator defined in (4) counts non-backtracking walks from both V_1 and V_2 .

Intuition for the non-backtracking wedge matrix It has been shown that the non-backtracking operator, as a non-hermitian matrix, outperforms the standard (weighted) adjacency matrix when the sample size is very small in many graph clustering and matrix completion problems (Bordenave et al., 2018, 2022b; Stephan and Massoulié, 2022; Stephan and Zhu, 2022). The key difference is that the spectrum of the non-backtracking operator is more informative and stable when the underlying graph is very sparse with a constant average degree. It is also a general phenomenon that the spectrum of non-hermitian random matrices is less sensitive to a row or column with a large ℓ_2 -norm (Coste, 2023; Bordenave et al., 2022a).

In terms of the sample size, the non-backtracking matrix in (4) was used in (Bordenave et al., 2022b; Stephan and Massoulié, 2022) to achieve the optimal scale $O(n)$ for square matrix completion. However, for a low-rank tensor in a matrix form of size $n \times n^2$, the work (Bordenave et al., 2022b) only achieves weak recovery with sample size $O(n^2)$, while with the non-backtracking wedge operator, we only need $O(n^{1.5})$. The standard way in (Montanari and Sun, 2018; Cai et al., 2022) to obtain $\tilde{O}(n^{1.5})$ sample size is to apply a spectral method on the hollowed Gram matrix $Z = AA^T - \text{diag}(AA^T)$. For $i \neq j$, we have

$$Z_{ij} = \sum_{\text{wedge } i \rightarrow k \rightarrow j} A_{ik} A_{kj}.$$

Hence, the spectrum of Z depends on wedge walk counts. Our new non-backtracking operator only counts the wedge walks that are *non-backtracking*, which avoids the localization effect from high-degree vertices (Benaych-Georges et al., 2019) in the weighted graph corresponding to Z . This new operator introduces a tool to “regularize” the spectrum of a heavy-tailed long random matrix without any tuning parameter, which might be useful for other applications.

2.2. Quantifying the convergence

The alignment between the spectral decomposition of B and the SVD of A will be governed by a wealth of different quantities. For ease of understanding, we divide those parameters into two categories, depending on their role in the convergence.

Threshold parameters These parameters describe which parts of the SVD of M will be reflected in the spectral decomposition of B .

(i) **Variance parameter:**

$$\rho = \|Q\|, \quad \text{where } Q = \sqrt{mn}(M \circ M). \quad (5)$$

(ii) **Amplitude parameter:**

$$L = \sqrt{mn} \max_{x \in [n], y \in [m]} |M_{xy}|.$$

(iii) **Detection thresholds**

$$\theta = \max(\theta_1, \theta_2),$$

where the two parameters θ_1 and θ_2 are defined as

$$\theta_1 = \sqrt{\frac{\rho}{d}} \quad \text{and} \quad \theta_2 = \frac{L}{d}.$$

(iv) **Detection rank:**

$$r_0 = \max\{i \in [n] : \nu_i > \theta\}.$$

Equivalently, it is the number of singular values of M above the detection threshold θ .

(v) **Inverse signal-to-noise ratios (inverse SNRs):** for $i \in [r_0]$,

$$\tau_i = \frac{\theta}{\nu_i} \in (0, 1).$$

The largest of those inverse SNRs, τ_{r_0} , will be denoted by τ .

Complexity parameters Those parameters quantify how intrinsically “easy” the matrix M is to recover. In particular, contrary to the previous ones, they are all invariant under any rescaling of M .

(i) **Rescaled amplitude parameter:**

$$K = \frac{L^2}{\rho}. \quad (6)$$

(ii) **Aspect log-ratio:**

$$\eta = \frac{\log(m)}{\log(n)} - 1. \quad (7)$$

Equivalently, it is the solution of the equation $m = n^{1+\eta}$. When M is obtained from the (a, b) unfolding of a tensor, we get $\eta = \frac{b-a}{a}$.

(iii) **Rank:**

$$r = \text{rank}(M) = \max\{i \in [n] : \nu_i \neq 0\}.$$

(iv) **Delocalization parameter:**

$$\kappa = \sqrt{n} \max_i \|\phi_i\|_\infty.$$

Note that $\kappa \geq 1$ by definition. We require no such conditions on the right singular vectors ψ_i .

Simplifying assumptions We also make a few simplifying assumptions throughout the paper; although our results are still valid when those assumptions aren't satisfied, the proofs become much messier. Recall the aspect ratio is defined by $\alpha = \frac{m}{n} \geq 1$. We first assume that

$$\liminf_n d > 1 \quad \text{and} \quad d \leq n. \quad (8)$$

We also assume that

$$\eta \geq \frac{2 \log(d) + 2 \log \log_d(n)}{\log(n)}, \quad \text{or equivalently} \quad \alpha \geq d^2 (\log_d n)^2. \quad (9)$$

We note that if $d > n$, our results become trivially true (see the definition of ℓ in Theorem 2, and the bounds thereafter).

2.3. Main results for long matrix completion

Our first result shows that the eigenvalues of B exhibit a structure similar to those observed in (Stephan and Massoulié, 2022; Bordenave et al., 2022b): namely, most eigenvalues of B lie in a circle of known radius, with a few outlier eigenvalues correlated with the singular values of M .

Theorem 2 (Singular value estimation) *Let*

$$\ell = \lfloor \varepsilon \log_d(n) \rfloor, \quad \text{with} \quad \varepsilon = \frac{(\eta/2) \wedge 1}{25}.$$

There exist constants $C_0, n_0 \geq 1$ that depend polynomially on $r, K, \kappa, d, \log n$ such that, if $n \geq n_0$ and $C_0 \tau^{2\ell} \leq 1$, the following holds on an event with probability at least $1 - cd^4 n^{-1/4}$:

(i) *There exists an ordering $\lambda_1, \dots, \lambda_{r_0}$ of the top r_0 eigenvalues of B such that for any $i \in [r_0]$,*

$$|\lambda_i - \nu_i^2| \leq C_0 \nu_i^2 \tau_i^{2\ell}.$$

(ii) *All other eigenvalues of B have modulus at most $C_0^{1/\ell} \theta^2$.*

Our second result concerns quantifying the overlaps between the top eigenvectors of B and the singular vectors of M . For $i \in [r_0]$, we define the relative eigengap

$$\delta_{i,\ell} = \min_{\nu_j \neq \nu_i} \left| 1 - \left(\frac{\nu_j}{\nu_i} \right)^{2\ell} \right|. \quad (10)$$

Let ξ_i^R, ξ_i^L a right (resp. left) unit eigenvector of B associated with λ_i , and define

$$\zeta_i^R(x) := \sum_{e \in \mathbf{E}: e_1=x} A_{e_1 e_2} A_{e_3 e_2} \xi_i^R(e) \quad \text{and} \quad \zeta_i^L(x) := \sum_{e \in \mathbf{E}: e_1=x} \xi_i^L(e). \quad (11)$$

Theorem 3 (Left singular vector estimation) *Under the same setting and conditions as Theorem 2, with probability at least $1 - cd^4 n^{-1/4}$, for $i \in [r_0]$, there exists a unit left singular vector ϕ_i' of M associated with ν_i such that*

$$\left| \langle \zeta_i^R, \phi_i' \rangle - \frac{1}{\sqrt{\gamma_i}} \right| \leq \frac{C_0 \tau_i^{2\ell}}{\delta_{i,\ell}} \quad \text{and} \quad \left| \langle \zeta_i^L, \phi_i' \rangle - \frac{1}{\sqrt{\gamma_i}} \right| \leq \frac{C_0 \tau_i^{2\ell}}{\delta_{i,\ell}}, \quad (12)$$

where γ_i is defined as

$$\gamma_i = \left\langle \mathbf{1}, \left(I - \frac{QQ^*}{\nu_i^4 d^2} \right)^{-1} (\phi_i \circ \phi_i) \right\rangle. \quad (13)$$

2.4. Discussion

Comparison to (Bordenave et al., 2022b) In the seminal work of (Bordenave et al., 2022b), the authors provide a spectral algorithm to weakly recover a low-rank rectangular matrix of size $m \times n$ with $O(m + n)$ many samples, where a Kesten-Stigum bound for estimation singular values and both left and right eigenvectors are obtained. In our low-rank long matrix setting ($m \gg n$), their sample complexity scales as $O(m)$. In contrast, our non-backtracking wedge operator is able to estimate singular values and left singular vectors of a long matrix with sample size $O(\sqrt{mn})$. This improvement in sample complexity to recover one-sided structure is useful in tensor completion; see Section 2.5 for more details.

Overlap convergence and weak consistency In full generality, the overlaps γ_i in Theorem 3 have a complicated expression. However, this can be simplified under a homogeneity assumption, namely

$$D_x := \sum_{y \in V_1, z \in V_2} Q_{xz} Q_{yz}$$

is independent of x . In this case, the Perron-Frobenius theorem implies that $D_x = \rho^2$, and hence

$$\langle \zeta_i^R, \phi_i \rangle = \sqrt{1 - \tau_i^4} + o(1), \quad \langle \zeta_i^L, \phi_i \rangle = \sqrt{1 - \tau_i^4} + o(1).$$

This implies a *weak recovery* of the singular vectors ϕ_i . In general, as $d \rightarrow \infty$, we have $(\gamma_i - 1) \sim C/d^2$, and hence (12) implies the eigenvectors of B are asymptotically aligned with the left singular vectors of M . Under the additional assumption that all singular values of M are distinct, the top r eigenvectors of B are asymptotically aligned with all left singular vectors of M , and eigenvalues are consistent estimators for $\nu_i, i \in [r]$: this is a form of *weak consistency* for matrix completion.

Optimality of the bounds and thresholds We have made no effort towards getting closed-form exponents in the polynomials C_0 and n_0 , which would probably not be sharp. An explicit bound on C_0, n_0 in the proportional case for matrix completion is worked out in (Bordenave et al., 2022b, Theorem 5.4). Similarly, although we do expect (when d, η are fixed) that the rate of convergence is indeed of the form n^{-c} , we didn't try to optimize the value of such c as much as possible.

Our work shows a threshold behavior for long matrix completion: one can weakly recover the components in MM^\top whose signal-to-noise ratio is above the threshold. On the other hand, the optimality of θ is still an open and interesting question since if this threshold can be lowered, more information about M can be recovered. The first threshold θ_1 is inherent to the problem and can be found in many similar problems (e.g., the *Kesten-Stigum threshold* in community detection (Krzakala et al., 2013a; Bordenave et al., 2018; Mossel et al., 2015, 2018)): it is related to an intrinsic property on the existence of a pseudo-eigenvector in a Galton-Watson tree with Poisson offspring distribution of mean d , see Section E for details. But whether θ_2 is needed or is an artifact of proof is still unclear. We note that in a similar problem regarding the configuration model (Coste, 2021), such a seemingly spurious threshold turns out to be the sharp one (albeit only for one eigenvalue). Note that as d grows, θ_2 decays faster than θ_1 , therefore $\theta = \theta_1$ for sufficiently large d .

Recovering right singular vectors We draw attention to the fact that our result is almost completely independent of the right singular vectors ψ_i of M , both on the level of assumptions (except for their influence on K) and results. We believe this is not an artifact of proof: indeed, most of the right vertices (a proportion $1 - \alpha^{-1}$ of them) are isolated, and it is impossible to recover any

signal from those. This is a completely different behavior compared with (Bordenave et al., 2022b), where the authors show that when α is bounded, both left and right singular vectors of M can be recovered. However, it matches recent results from (Montanari and Wu, 2022) in the framework of long matrix reconstruction showing different reconstruction thresholds for the left and right singular vectors. Our work also shows that one can effectively estimate the rank of a long matrix of size $n \times m$ with $O(\sqrt{mn})$ random samples when the rank $r = O(1)$. When m, n are proportional, a similar result was obtained in (Saade et al., 2015) based on the Bethe-Hessian matrix.

Parameter scaling All of the results in Theorems 2 and 3 are non-asymptotic, and with more care in the proof, we would be able to give explicit expressions for C_0, n_0 (see, e.g., (Bordenave et al., 2022b)). However, the given bounds are meaningful only when $\ell \gg 1$, and $C_0 \tau^{2\ell} \ll 1$. These can easily be seen to be satisfied whenever

$$K, \kappa, r = n^{o(1)}, \quad \text{and} \quad \log(d)^2 = o(\log(n)).$$

2.5. Application in low-CP-rank tensor completion

Throughout this part, our goal will be to recover a tensor T of order $k \geq 3$, which has a CP decomposition (see, e.g., (Kolda and Bader, 2009)) with r components:

$$T = \sum_{i=1}^r \nu_i w_i^{(1)} \otimes w_i^{(2)} \otimes \cdots \otimes w_i^{(k)}, \quad (14)$$

where $\|w_i^{(j)}\|_2 = 1$ for all $i \in [r], j \in [k]$. The smallest integer r for the existence of the decomposition (14) is called the *CP-rank* of T , and we will call it the rank of T for simplicity in this section.

It is known that for a general low-rank tensor T , finding its CP-decomposition is NP-hard (Hillar and Lim, 2013). Even with the complete information of the ground truth T , identifying each component in $\{w_i^{(j)}, i \in [r], j \in [k]\}$ is not feasible. Since we consider asymmetric tensor completion, it will be useful to consider a general version of unfolding:

Definition 4 Let \mathcal{S}, \mathcal{T} be a partition of $[k]$ in two subsets. The $(\mathcal{S}, \mathcal{T})$ -unfolding of T ; denoted as $\text{unfold}_{\mathcal{S}, \mathcal{T}}(T)$, is a matrix $M \in \mathbb{R}^{n^{\mathcal{S}} \times n^{\mathcal{T}}}$ defined as follows: for any multi-index $(i_s)_{s \in [k]}$,

$$M_{(i_s)_{s \in \mathcal{S}}, (i_s)_{s \in \mathcal{T}}} = T_{(i_s)_{s \in [k]}}.$$

In particular, when T is as in (14), we have

$$\text{unfold}_{\mathcal{S}, \mathcal{T}}(T) = \sum_{i=1}^r \nu_i \left(\bigotimes_{j \in \mathcal{S}} w_j \right) \left(\bigotimes_{j \in \mathcal{T}} w_j \right)^*. \quad (15)$$

For simplicity, we define the mode- i unfolding of T as $\text{unfold}_i(T) := \text{unfold}_{\{i\}, [k] \setminus \{i\}}(T)$.

However, we make an important remark: Equation (15) is not necessarily the SVD of the tensor unfolding. Indeed, the CP decomposition does not always have orthogonal components, and hence, the vectors in (15) are not necessarily orthogonal. Generally, for a given tensor T , the spectra of the unfolded matrices are not the same for different unfoldings, and there is a hierarchy of unfolded matrix norms (Wang et al., 2017). In the following, we will denote by \tilde{T} an observed tensor from T with sampling probability $p = \frac{d}{n^{k/2}}$. For an unfolding mode i , we let $M^{(i)} = \text{unfold}_i(\tilde{T})$, and $B^{(i)}$ be its associated non-backtracking wedge matrix.

2.5.1. THE RANK-1 CASE

When $r = 1$, we write

$$T = w^{(1)} \otimes \dots \otimes w^{(k)}, \quad (16)$$

where we choose $\nu_1 = 1$ in (14) for convenience. Then we have

$$\text{unfold}_i(T) := \text{unfold}_{\{i\}, [k] \setminus \{i\}}(T) = w^{(i)}(w^{(1)} \otimes \dots \otimes w^{(i-1)} \otimes w^{(i+1)} \otimes \dots \otimes w^{(k)})^\top.$$

We now write down the corresponding model parameters for $M^{(i)}$. It can be checked that they are the same for all i , independent of the choice for unfolding, except for κ :

$$\begin{aligned} \theta_1 &= \sqrt{\frac{n^{k/2} \prod_{j=1}^k \|w^{(j)}\|_4^2}{d}}, & \theta_2 &= \frac{n^{k/2} \prod_{j=1}^k \|w^{(j)}\|_\infty}{d} \\ K &= n^{k/2} \prod_{j=1}^k \frac{\|w^{(j)}\|_\infty^2}{\|w^{(j)}\|_4^2}, & \kappa^{(i)} &= \sqrt{n} \|w^{(i)}\|_\infty \leq \sqrt{n} \max_i \|w^{(i)}\|_\infty =: \kappa_{\max}. \end{aligned} \quad (17)$$

Note that we always have $\theta_1^2 \leq \theta_2$. Applying Theorems 2 and 3 to $\text{unfold}_i(T)$, we obtain the following corollary:

Corollary 5 (Rank-1 tensor completion) *Let $k \geq 3$, and define ℓ as*

$$\ell = \lfloor c_1 \log_d(n) \rfloor, \quad c_1 = \frac{(k/2 - 1) \wedge 1}{25}.$$

For any constant $\varepsilon > 0$, assume that

$$d > (1 + \varepsilon) n^{k/2} \prod_{j=1}^k \|w^{(j)}\|_\infty. \quad (18)$$

Then there exists $C_0, n_0 \geq 1$ that depend polynomially on $K, \kappa_{\max}, d, \log n$ such that for any $i \in [k]$, $n \geq n_0$ and $C_0 \theta^{2\ell} \leq 1$, the following holds on an event with probability at least $1 - cd^4 n^{-1/4}$:

(i) *The top eigenvalue $\lambda^{(i)}$ of $B^{(i)}$ satisfies*

$$|\lambda^{(i)} - 1| \leq C_0 \theta^{2\ell},$$

where $\theta = \max(\theta_1, \theta_2)$ with θ_1, θ_2 in (17).

(ii) *All other eigenvalues of $B^{(i)}$ have modulus at most $C_0^{1/\ell} \theta^2$.*

(iii) *There exist associated left and right eigenvectors $\xi^{(i),L}, \xi^{(i),R}$ of $B^{(i)}$ such that*

$$\left| \langle \zeta^{(i),L}, w^{(i)} \rangle - \sqrt{1 - \theta^2} \right| \leq C_0 \theta^{2\ell} \quad \text{and} \quad \left| \langle \zeta^{(i),R}, w^{(i)} \rangle - \sqrt{1 - \theta^2} \right| \leq C_0 \theta^{2\ell}, \quad (19)$$

where $\zeta^{(i),L}, \zeta^{(i),R}$ are defined in (11).

Hence, successively applying tensor unfolding k times to the observed tensor is enough to estimate all of the components of T . When all components w_1, \dots, w_k are vectors in $\frac{1}{\sqrt{n}} \{\pm 1\}^n$, with $(1 + \varepsilon) n^{k/2}$ samples, we have

$$\langle \zeta^{(i),L}, w^{(i)} \rangle = \frac{\sqrt{(1 + \varepsilon)^2 - 1}}{1 + \varepsilon} + O(n^{-c}).$$

2.5.2. LOW-RANK ORTHOGONAL COMPONENTS ESTIMATION

Our long matrix completion results aim to estimate individual left singular vectors above a Kesten-Stigum-type threshold. However, in the tensor case, without orthogonal assumption on the vectors $[w_i^{(1)}, \dots, w_i^{(k)}]$, such a reduction to the estimation of left singular vectors in long matrix completion is impossible. Therefore, we consider the same setting as in (Jain and Oh, 2014; Potechin and Steurer, 2017), and assume that T can be written as

$$T = \sum_{j=1}^r \nu_j w_j^{(1)} \otimes w_j^{(2)} \otimes \dots \otimes w_j^{(k)}, \quad (20)$$

where for each $i \in [k]$, $\{w_1^{(i)}, \dots, w_r^{(i)}\}$ forms an orthonormal basis, and $\nu_1 \geq \nu_2 \geq \dots \geq \nu_r > 0$. In contrast with the rank-one case, we don't have access to an orthogonal CP-decomposition of $T \circ T$, and hence most parameters used in Theorem 2 depend on the unfolding mode i ; we shall denote them by $\rho^{(i)}$, $Q^{(i)}$, and so on. We can still provide the following estimates of $\rho^{(i)}$ and $L^{(i)}$:

$$n^{k/2} \|T\|_2 \leq \rho^{(i)} \leq n^{k/2} \|T\|_F, \quad \text{and} \quad L^{(i)} = L := n^{k/2} \|T\|_\infty.$$

Corollary 6 (Low-rank orthogonal tensor completion) *Let T be a k -tensor with an orthogonal decomposition (20) and $k \geq 3$, and define ℓ as*

$$\ell = \lfloor c_1 \log_d(n) \rfloor, \quad c_1 = \frac{(k/2 - 1) \wedge 1}{25}.$$

For any $i \in [k]$, there exists $C_0^{(i)}, n_0^{(i)}$, that depend polynomially on $K^{(i)}, \kappa^{(i)}, r^{(i)}, d^{(i)}, \log(n)$, such that, if $n \geq n_0$ and $C_0^{(i)} (\tau_j^{(i)})^{2\ell} \leq 1$, the following holds with probability at least $1 - cd^4 n^{-1/4}$:

(i) *There exists an ordering $\lambda_1^{(i)}, \dots, \lambda_{r_0}^{(i)}$ of the top r_0 eigenvalues of $B^{(i)}$ such that for any $j \in [r_0^{(i)}]$,*

$$|\lambda_j - \nu_j^2| \leq C_0^{(i)} \nu_j^2 (\tau_j^{(i)})^{2\ell}.$$

(ii) *All other eigenvalues of $B^{(i)}$ have modulus at most $(C_0^{(i)})^{1/\ell} (\theta^{(i)})^2$.*

(iii) *For $j \in [r_0^{(i)}]$, there exist associated left and right eigenvectors ξ_j^L, ξ_j^R of $B^{(i)}$ such that*

$$\left| \langle \zeta_j^{(i),R}, w_j^{(i)} \rangle - \frac{1}{\sqrt{\gamma_j^{(i)}}} \right| \leq \frac{C_0^{(i)} (\tau_j^{(i)})^{2\ell}}{\delta_{j,\ell}}, \quad \left| \langle \zeta_j^{(i),L}, w_j^{(i)} \rangle - \frac{1}{\sqrt{\gamma_j^{(i)}}} \right| \leq \frac{C_0^{(i)} (\tau_j^{(i)})^{2\ell}}{\delta_{j,\ell}},$$

where γ_j is defined as

$$\gamma_j^{(i)} = \left\langle \mathbf{1}, \left(I - \frac{Q^{(i)} Q^{(i)*}}{\nu_j^4 d^2} \right)^{-1} (w_j^{(i)} \circ w_j^{(i)}) \right\rangle,$$

$\delta_{j,\ell}$ is defined in (10), and $\zeta_j^{(i),R}, \zeta_j^{(i),L}$ are defined in (11).

We obtain a weak recovery of an orthogonal tensor decomposition with sample size $O(n^{k/2})$, which removes the $\text{polylog}(n)$ factors in (Jain and Oh, 2014; Potechin and Steurer, 2017). Same as in the discussion for long matrix completion in Section 2.4, under the additional assumption that $\{\nu_i, i \in [r]\}$ are separated, with sample size $\omega(n^{k/2})$, weak consistency for T can be achieved:

Theorem 7 *Let T have an orthogonal decomposition as in (20), such that $\delta_{i,\ell} \geq c$ for some constant $c > 0$. For $i \in [k], j \in [r]$, define the estimator*

$$\hat{T} = \sum_{j=1}^r \hat{\nu}_j \hat{w}_j^{(1)} \otimes \cdots \otimes \hat{w}_j^{(k)} \quad \text{for} \quad \hat{w}_j^{(i)} = \zeta_j^{(i),R} \quad \text{and} \quad \hat{\nu}_j = \sqrt{\lambda_i^{(1)}}.$$

Then, with probability at least $1 - cd^4 n^{-1/4}$, for large enough d ,

$$\frac{\|T - \hat{T}\|_F}{\|T\|_F} \leq \frac{\sqrt{kr} K \tau}{d} + C_1 \left(\frac{\tau}{d}\right)^{\ell-1}, \quad \text{where} \quad \tau = \nu_r^{-2} \max_i \rho^{(i)}.$$

Theorem 7 is proved in Appendix I. To the best of our knowledge, this is the first result that shows weak recovery (in the sense $\|T - \hat{T}\|_F = o(1)$) in the regime $d \rightarrow \infty$, as opposed to $d \gtrsim \log(n)^c$.

2.5.3. DISCUSSION ON TENSOR COMPLETION

Comparison to nearly square unfolding Corollaries 5 and 6 provide an efficient algorithm that outputs an estimate for each component $w^{(i)}$. This directly implies that one can have a correlated estimation of T with sample size $O(n^{k/2})$. This improves the sample complexity $O(n^{\lceil k/2 \rceil})$ when directly applying the matrix completion algorithm in (Bordenave et al., 2022b) to $\text{unfold}_{\lfloor k/2 \rfloor, \lceil k/2 \rceil}(T)$ when k is odd. This mode- i unfolding approach is more computationally efficient than the nearly square unfolding (Montanari and Sun, 2018) for component-wise estimation. For example, in the rank-1 case, applying the matrix completion algorithm to $\text{unfold}_{\lfloor k/2 \rfloor, \lceil k/2 \rceil}(T)$, one can find a correlated estimation of $w^{(1)} \otimes \cdots \otimes w^{(\lfloor k/2 \rfloor)}$. However, a further recursive unfolding is needed to recover individual components. The advantage of single-mode unfolding was also discussed in (Ben Arous et al., 2023) for the spiked tensor model.

Optimal sample complexity for rank-one completion We note that although our method does reach the $O(n^{k/2})$ scaling for rank-one tensor completion, this is not the optimal sample complexity. Indeed, for a boolean rank-one tensor, completion is equivalent to k -XORSAT, and hence (Creignou et al., 2003) implies the following theorem:

Theorem 8 *Assume that $T = x^{\otimes k}$ for a given vector $x \in \mathbb{R}^n$ such that $\|x\|_2 = 1$. Then there exists a constant $c_k < 1$ such that if $p \geq c_k n^{-(k-1)}$, there exists an algorithm that outputs an estimator \hat{x} for x satisfying*

$$\frac{\langle \hat{x}, x \rangle}{\|\hat{x}\|_2} \geq \frac{1}{\sqrt{n} \|x\|_\infty}.$$

Hence, for the specific case of rank-one tensors with bounded entries, a near-optimal algorithm can reach a sample complexity of $O(n)$. However, the mapping to a XORSAT problem disappears for tensors of rank $r > 1$, as well as with the addition of any type of noise (see, e.g., (Barak and Moitra, 2022)). Therefore, the conjectured $n^{k/2}$ threshold in (Barak and Moitra, 2022) for general low-rank tensor completion does not apply to this specific noiseless example. The proof of Theorem 8 is given in Appendix J.

Subspace recovery for general low-rank tensors Without the orthogonal decomposition assumption of the tensor T , one cannot directly associate the singular vectors of the unfolded matrix to the components of T . In this general case, the best one can hope for is to recover the subspace spanned by the spikes from T . Although our result is of independent interest in long matrix completion, it can also be used directly in non-orthogonal tensor completion. First, when $d \rightarrow \infty$, the results of (Montanari and Sun, 2018) imply that a good reconstruction of the left singular space of M translates directly to weak consistency for tensor reconstruction. On the other hand, several algorithms such as the ones in (Montanari and Sun, 2018; Liu and Moitra, 2020; Cai et al., 2022) make use of a first spectral estimate of the CP factors of the tensor T before a refined optimization step; a popular choice for such an estimate is usually a truncated SVD of A . We conjecture that even for fixed d when the SVD of A contains no information about the one of M , the embedded eigenvectors of B can be used to perform an informed initialization of those algorithms.

2.6. Application in low-multilinear-rank tensor completion

Besides CP-decomposition, another popular choice for tensor decomposition is the Tucker decomposition (Tucker, 1966), which is also called higher-order singular value decomposition (HOSVD) (De Lathauwer et al., 2000). For an order k -tensor T , we write its HOSVD as

$$T = \sum_{q_1=1}^{r_1} \sum_{q_2=1}^{r_2} \cdots \sum_{q_k=1}^{r_k} s_{q_1, \dots, q_k} \phi_{q_1}^{(1)} \otimes \cdots \otimes \phi_{q_k}^{(k)}, \quad (21)$$

where (r_1, \dots, r_k) is called the *multilinear rank* of T , $U_j = [\phi_1^{(j)}, \dots, \phi_{r_j}^{(j)}] \in \mathbb{R}^{n \times r_j}$ is a matrix with orthonormal columns, and $S \in \mathbb{R}^{r_1 \times \dots \times r_k}$ is the core tensor. One can show that the columns of $U^{(j)}$ are the left singular vectors of the mode- j unfolding of T denoted by $\text{unfold}_j(T) \in \mathbb{R}^{n \times n^{k-1}}$. See (Kolda and Bader, 2009) for more details. Tensor completion with low-multilinear-rank was studied in (Kressner et al., 2014; Yuan and Zhang, 2017; Montanari and Sun, 2018). Recovering an HOSVD under additive noise was considered in (Lebeau et al., 2024).

We can apply the result for long matrix completion in Section 2.3 and to $\text{unfold}_j(T)$ to recover the mode- j components of the HOSVD. We denote the SVD of $\text{unfold}_j(T)$ as

$$\text{unfold}_j(T) = \sum_{i=1}^{r_j} \nu_i^{(j)} \phi_i^{(j)} \psi_i^{(j)}.$$

Let $B^{(j)}$ be the non-backtracking wedge operator for the observed matrix of $\text{unfold}_j(T)$ after sampling each entry with probability $p = \frac{d}{n^{k/2}}$. The following corollary follows directly from Theorem 3.

Corollary 9 (Low-multilinear-rank tensor completion) *Let T be a tensor defined in (21). With the same notations as in Section 2.2 and Theorem 3 applied to $\text{unfold}_j(T)$ and $B^{(j)}$, there exist constants $C_0, n_0 \geq 1$ that depend polynomially on $r, K, \kappa, d, \log n$ such that, if $n \geq n_0$ and $C_0 \tau_i^{2\ell} \leq 1$, the following holds on an event with probability at least $1 - cd^4 n^{-1/4}$: there exists a unit left singular vector $\phi_i^{\prime(j)}$ of $\text{unfold}_j(T)$ associated with $\nu_i^{(j)}$ such that*

$$\left| \langle \zeta_i^R, \phi_i^{\prime(j)} \rangle - \frac{1}{\sqrt{\gamma_i}} \right| \leq \frac{C_0 \tau_i^{2\ell}}{\delta_{i,\ell}} \quad \text{and} \quad \left| \langle \zeta_i^L, \phi_i^{\prime(j)} \rangle - \frac{1}{\sqrt{\gamma_i}} \right| \leq \frac{C_0 \tau_i^{2\ell}}{\delta_{i,\ell}}.$$

By applying Corollary 9 to different unfoldings, we can achieve weak recovery for the orthonormal components $U^{(j)}$, $1 \leq j \leq k$, with $O(n^{k/2})$ many samples. Our method, however, does not provide estimations for the core tensor S , the same as the work on low-multilinear-rank tensor approximation (Lebeau et al., 2024).

Organization of the appendix

Numerical simulations are provided in Appendix A. Appendix B is devoted to stating the main technical ingredient of Theorems 2 and 3. It consists of a detailed spectral structure of the matrix B (or, more precisely, of its power B^ℓ), with precise expressions given for the pseudo-eigenvectors of B and their relationships with each other. We also provide a sketch of the proof of how this technical result implies our main theorems. The rest of the sections are then devoted to proving Theorem 10. In short:

- (i) In Appendix C, we collect a few simple inequalities that will be useful in the proof.
- (ii) in Appendix D, we perform a local study of the neighborhoods of the weighted bipartite random graph G ; we show a local convergence (in a stronger sense than the one of Benjamini and Schramm (Benjamini and Schramm, 2001)) of G to an appropriately defined Galton-Watson tree T . The structure of this tree is tailored to our bipartite setting and is similar to the one of (Florescu and Perkins, 2016). In short, the tree T encodes the “local” properties of G , in the sense that most neighborhoods of G have the same distribution as the tree T .
- (iii) in Appendix E, we study this tree in-depth to extract information about specific processes on T . Those processes give rise to the pseudo-eigenvectors of B . The main ingredient of this Appendix is a reduction to the setting of (Stephan and Massoulié, 2022), which approximates the bipartite graph G with a randomly weighted (but non-bipartite) version.
- (iv) Appendices F and G are devoted to bounding the bulk eigenvalues of G . The method used – a so-called *tangle-free* decomposition and a variant of the trace method (Füredi and Komlós, 1981) – is similar to the one of (Bordenave et al., 2018, 2022b; Stephan and Massoulié, 2022; Stephan and Zhu, 2022). One of the major technical novelty is in the trace method proof of Theorem 2(ii). The B matrix counts non-backtracking wedges, but the appearance of different wedges in a random graph G is dependent. It is a big challenge to estimate the high-power trace of a random matrix with dependent entries compared to the standard moment method proof techniques in the random matrix theory literature, and such a technical challenge does not appear in (Bordenave et al., 2018, 2022b; Stephan and Massoulié, 2022). Additional proofs based on the trace method are included in Appendix H.

Finally, we include the proof of Theorem 7 in Appendix I, and the proof of Theorem 8 in Appendix J.

Acknowledgments

Y.Z. is partially supported by NSF-Simons Research Collaborations on the Mathematical and Scientific Foundations of Deep Learning and an AMS-Simons Travel Grant.

References

- Emmanuel Abbe. Community Detection and Stochastic Block Models: Recent Developments. *Journal of Machine Learning Research*, 18(177):1–86, 2018.
- Johannes Alt, Raphaël Ducatez, and Antti Knowles. Delocalization transition for critical Erdős–Rényi graphs. *Communications in Mathematical Physics*, 388(1):507–579, 2021a.
- Johannes Alt, Raphaël Ducatez, and Antti Knowles. Extremal eigenvalues of critical Erdős–Rényi graphs. *The Annals of Probability*, 49(3):1347–1401, 2021b.
- M. C. Angelini, F. Caltagirone, F. Krzakala, and L. Zdeborová. Spectral detection on sparse hypergraphs. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 66–73, September 2015. doi: 10.1109/ALLERTON.2015.7446987.
- Gérard Ben Arous, Reza Gheissari, and Aukosh Jagannath. Algorithmic thresholds for tensor PCA. *Annals of Probability*, 48(4):2052–2087, 2020.
- Arnab Auddy and Ming Yuan. On estimating rank-one spiked tensors in the presence of heavy tailed errors. *IEEE Transactions on Information Theory*, 68(12):8053–8075, 2022.
- Boaz Barak and Ankur Moitra. Noisy tensor completion via the sum-of-squares hierarchy. *Mathematical Programming*, 193(2):513–548, 2022.
- Hyman Bass. The Ihara-Selberg Zeta function of a tree lattice. *International Journal of Mathematics*, 03(06):717–797, December 1992. ISSN 0129-167X. doi: 10.1142/S0129167X92000357.
- F. L. Bauer and C. T. Fike. Norms and exclusion theorems. *Numerische Mathematik*, 2(1):137–141, December 1960. ISSN 0945-3245. doi: 10.1007/BF01386217.
- Gérard Ben Arous, Daniel Zhengyu Huang, and Jiaoyang Huang. Long random matrices and tensor unfolding. *The Annals of Applied Probability*, 33(6B):5753–5780, 2023.
- Florent Benaych-Georges, Charles Bordenave, and Antti Knowles. Largest eigenvalues of sparse inhomogeneous Erdős–Rényi graphs. *The Annals of Probability*, 47(3):1653–1676, May 2019. ISSN 0091-1798, 2168-894X. doi: 10.1214/18-AOP1293.
- Florent Benaych-Georges, Charles Bordenave, and Antti Knowles. Spectral radii of sparse random matrices. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 56(3):2141–2161, 2020.
- Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. *Electronic Journal of Probability*, 6:1–13, 2001.
- Charles Bordenave. A new proof of Friedman’s second eigenvalue theorem and its extension to random lifts. *Annales Scientifiques de l’École Normale Supérieure*, 4(6):1393–1439, 2020.
- Charles Bordenave and Benoît Collins. Eigenvalues of random lifts and polynomials of random permutation matrices. *Annals of Mathematics*, 190(3):811–875, 2019.

- Charles Bordenave, Marc Lelarge, and Laurent Massoulié. Nonbacktracking spectrum of random graphs: Community detection and nonregular ramanujan graphs. *Annals of probability*, 46(1): 1–71, 2018.
- Charles Bordenave, Djilil Chafaï, and David García-Zelada. Convergence of the spectral radius of a random matrix through its characteristic polynomial. *Probability Theory and Related Fields*, pages 1–19, 2022a.
- Charles Bordenave, Simon Coste, and Raj Rao Nadakuditi. Detection thresholds in very sparse matrix completion. *Foundations of Computational Mathematics*, pages 1–125, 2022b.
- Guillaume Braun. Strong consistency guarantees for clustering high-dimensional bipartite graphs with the spectral method. *arXiv preprint arXiv:2304.06903*, 2023.
- Guillaume Braun and Hemant Tyagi. Minimax optimal clustering of bipartite graphs with a generalized power method. *Information and Inference: A Journal of the IMA*, 12(3):iaad006, 2023.
- Gerandy Brito, Ioana Dumitriu, and Kameron Decker Harris. Spectral gap in random bipartite biregular graphs and applications. *Combinatorics, Probability and Computing*, 31(2):229–267, 2022.
- Erik Broman, Tim Brug, Wouter Kager, and Ronald Meester. Stochastic domination and weak convergence of conditioned bernoulli random vectors. *Alea*, 9, 09 2011.
- Changxiao Cai, Gen Li, Yuejie Chi, H Vincent Poor, and Yuxin Chen. Subspace estimation from unbalanced and incomplete data matrices: $\ell_{2,\infty}$ statistical guarantees. *The Annals of Statistics*, 49(2), 2021.
- Changxiao Cai, Gen Li, H Vincent Poor, and Yuxin Chen. Nonconvex low-rank tensor completion from noisy data. *Operations Research*, 70(2):1219–1237, 2022.
- E. J. Candes and T. Tao. The Power of Convex Relaxation: Near-Optimal Matrix Completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, May 2010. ISSN 1557-9654. doi: 10.1109/TIT.2010.2044061.
- Emmanuel Candes and Benjamin Recht. Exact matrix completion via convex optimization. *Communications of the ACM*, 55(6):111–119, 2012.
- Emmanuel J. Candes and Yaniv Plan. Matrix Completion With Noise. *Proceedings of the IEEE*, 98(6):925–936, June 2010. ISSN 1558-2256. doi: 10.1109/JPROC.2009.2035722.
- Simon Coste. The spectral gap of sparse random digraphs. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 57(2):644–684, 2021.
- Simon Coste. Sparse matrices: convergence of the characteristic polynomial seen from infinity. *Electronic Journal of Probability*, 28:1–40, 2023.
- Nadia Creignou and Hervé Daudé. Smooth and sharp thresholds for random k -xor-cnf satisfiability. *RAIRO-Theoretical Informatics and Applications-Informatique Théorique et Applications*, 37(2): 127–147, 2003.

- Nadia Creignou, Hervé Daudé, and Olivier Dubois. Approximating the satisfiability threshold for random k -XOR-formulas. *Combinatorics, Probability and Computing*, 12(2):113–126, March 2003.
- Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. A multilinear singular value decomposition. *SIAM journal on Matrix Analysis and Applications*, 21(4):1253–1278, 2000.
- Jingqiu Ding, Samuel Hopkins, and David Steurer. Estimating Rank-One Spikes from Heavy-Tailed Noise via Self-Avoiding Walks. In *Advances in Neural Information Processing Systems*, volume 33, pages 5576–5586, 2020.
- Rishabh Dudeja and Daniel Hsu. Statistical-computational trade-offs in tensor pca and related problems via communication complexity. *The Annals of Statistics*, 52(1):131–156, 2024.
- Ioana Dumitriu and Yizhe Zhu. Spectra of random regular hypergraphs. *The Electronic Journal of Combinatorics*, pages P3–36, 2021.
- Ioana Dumitriu and Yizhe Zhu. Extreme singular values of inhomogeneous sparse random rectangular matrices. *arXiv preprint arXiv:2209.12271*, 2022.
- Ioana Dumitriu and Yizhe Zhu. Global eigenvalue fluctuations of random biregular bipartite graphs. *Random Matrices: Theory and Applications*, page 2350004, 2023.
- Laura Florescu and Will Perkins. Spectral thresholds in the bipartite stochastic block model. In *Conference on Learning Theory*, pages 943–959, 2016.
- David Freedman. A remark on the difference between sampling with and without replacement. *Journal of the American Statistical Association*, 72(359):681–681, 1977.
- Zoltán Füredi and János Komlós. The eigenvalues of random symmetric matrices. *Combinatorica*, 1(3):233–241, 1981.
- David Gamarnik, Quan Li, and Hongyi Zhang. Matrix completion from $O(n)$ samples in linear time. In *Conference on Learning Theory*, pages 940–947. PMLR, 2017.
- Navid Ghadermarzy, Yaniv Plan, and Özgür Yilmaz. Near-optimal sample complexity for convex tensor completion. *Information and Inference: A Journal of the IMA*, 8(3):577–619, 2019.
- Yuzhou Gu and Yury Polyanskiy. Weak recovery threshold for the hypergraph stochastic block model. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 885–920. PMLR, 2023.
- Venkatesan Guruswami, Peter Manohar, and Jonathan Mosheiff. ℓ_p -spread and restricted isometry properties of sparse random matrices. In *37th Computational Complexity Conference (CCC 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
- Kameron Decker Harris and Yizhe Zhu. Deterministic tensor completion with hypergraph expanders. *SIAM Journal on Mathematics of Data Science*, 3(4):1117–1140, 2021.

- Kameron Decker Harris, Oscar López, Angus Read, and Yizhe Zhu. Spectral gap-based deterministic tensor completion. In *2023 International Conference on Sampling Theory and Applications (SampTA)*, pages 1–6. IEEE, 2023.
- Eyal Heiman, Gideon Schechtman, and Adi Shraibman. Deterministic algorithms for matrix completion. *Random Structures & Algorithms*, 45(2):306–317, 2014.
- Christopher J Hillar and Lek-Heng Lim. Most tensor problems are NP-hard. *Journal of the ACM (JACM)*, 60(6):45, 2013.
- Aukosh Jagannath, Patrick Lopatto, and Léo Miolane. Statistical thresholds for tensor PCA. *Annals of applied probability*, 30(4):1910–1933, 2020.
- Prateek Jain and Sewoong Oh. Provable tensor factorization with missing data. In *Advances in Neural Information Processing Systems*, pages 1431–1439, 2014.
- Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 665–674, 2013.
- Raghuveer Keshavan, Andrea Montanari, and Sewoong Oh. Matrix Completion from Noisy Entries. In Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 952–960. Curran Associates, Inc., 2009.
- Raghuveer H Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *IEEE transactions on information theory*, 56(6):2980–2998, 2010.
- Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- Daniel Kressner, Michael Steinlechner, and Bart Vandereycken. Low-rank tensor completion by riemannian optimization. *BIT Numerical Mathematics*, 54:447–468, 2014.
- Florent Krzakala, Cristopher Moore, Elchanan Mossel, Joe Neeman, Allan Sly, Lenka Zdeborová, and Pan Zhang. Spectral redemption in clustering sparse networks. *Proceedings of the National Academy of Sciences*, 110(52):20935–20940, December 2013a. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1312486110.
- Florent Krzakala, Cristopher Moore, Elchanan Mossel, Joe Neeman, Allan Sly, Lenka Zdeborová, and Pan Zhang. Spectral redemption in clustering sparse networks. *Proceedings of the National Academy of Sciences*, 110(52):20935–20940, 2013b.
- Hugo Lebeau, Florent Chatelain, and Romain Couillet. A random matrix approach to low-multilinear-rank tensor approximation. *arXiv preprint arXiv:2402.03169*, 2024.
- Allen Liu and Ankur Moitra. Tensor completion made practical. *Advances in Neural Information Processing Systems*, 33:18905–18916, 2020.

- Siqi Liu, Sidhanth Mohanty, and Prasad Raghavendra. On statistical inference when fixed points of belief propagation are unstable. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 395–405. IEEE, 2022.
- Sidhanth Mohanty, Ryan O’Donnell, and Pedro Paredes. The sdp value for random two-eigenvalue csps. In *37th International Symposium on Theoretical Aspects of Computer Science*, 2020.
- Andrea Montanari and Emile Richard. A statistical model for tensor PCA. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS’14*, pages 2897–2905, Cambridge, MA, USA, December 2014. MIT Press.
- Andrea Montanari and Nike Sun. Spectral algorithms for tensor completion. *Communications on Pure and Applied Mathematics*, 71(11):2381–2425, 2018.
- Andrea Montanari and Yuchen Wu. Fundamental limits of low-rank matrix estimation with diverging aspect ratios. *arXiv preprint arXiv:2211.00488*, 2022.
- Elchanan Mossel, Joe Neeman, and Allan Sly. Reconstruction and estimation in the planted partition model. *Probability Theory and Related Fields*, 162(3):431–461, 2015. ISSN 1432-2064. doi: 10.1007/s00440-014-0576-6.
- Elchanan Mossel, Joe Neeman, and Allan Sly. A Proof of the Block Model Threshold Conjecture. *Combinatorica*, 38(3):665–708, 2018. ISSN 1439-6912. doi: 10.1007/s00493-016-3238-8.
- Cun Mu, Bo Huang, John Wright, and Donald Goldfarb. Square deal: Lower bounds and improved relaxations for tensor recovery. In *International conference on machine learning*, pages 73–81. PMLR, 2014.
- Mohamed Ndaoud, Suzanne Sigalla, and Alexandre B Tsybakov. Improved clustering algorithms for the bipartite stochastic block model. *IEEE Transactions on Information Theory*, 68(3):1960–1975, 2021.
- Soumik Pal and Yizhe Zhu. Community detection in the sparse hypergraph stochastic block model. *Random Structures & Algorithms*, 59(3):407–463, 2021.
- Aaron Potechin and David Steurer. Exact tensor completion with sum-of-squares. In *Conference on Learning Theory*, pages 1619–1673. PMLR, 2017.
- Benjamin Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12(12), 2011.
- Alaa Saade, Florent Krzakala, and Lenka Zdeborová. Matrix completion from fewer entries: Spectral detectability and rank estimation. *Advances in Neural Information Processing Systems*, 28, 2015.
- Ludovic Stephan and Laurent Massoulié. Non-backtracking spectra of weighted inhomogeneous random graphs. *Mathematical Statistics and Learning*, 5(3):201–271, 2022.
- Ludovic Stephan and Yizhe Zhu. Sparse random hypergraphs: Non-backtracking spectra and community detection. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 567–575. IEEE, 2022.

- Audrey Terras. *Zeta functions of graphs: A stroll through the garden*, volume 128. Cambridge University Press, 2010.
- Ledyard R Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3): 279–311, 1966.
- Miaoyan Wang, Khanh Dao Duc, Jonathan Fischer, and Yun S Song. Operator norm inequalities between tensor unfoldings on the partition lattice. *Linear algebra and its applications*, 520: 44–66, 2017.
- Alexander S Wein, Ahmed El Alaoui, and Cristopher Moore. The Kikuchi hierarchy and tensor PCA. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1446–1468. IEEE, 2019.
- Dong Xia and Ming Yuan. On polynomial time methods for exact low-rank tensor completion. *Foundations of Computational Mathematics*, 19(6):1265–1313, 2019.
- Dong Xia, Ming Yuan, and Cun-Hui Zhang. Statistically optimal and computationally efficient low rank tensor completion from noisy entries. *The Annals of Statistics*, 49(1), 2021.
- Ming Yuan and Cun-Hui Zhang. Incoherent tensor norms and their applications in higher order tensor completion. *IEEE Transactions on Information Theory*, 63(10):6753–6766, 2017.
- Zhixin Zhou, Fan Zhou, Ping Li, and Cun-Hui Zhang. Rate-optimal subspace estimation on random graphs. *Advances in Neural Information Processing Systems*, 34:20283–20294, 2021.
- Xiangyi Zhu and Yizhe Zhu. Non-backtracking eigenvalues and eigenvectors of random regular graphs and hypergraphs. *arXiv preprint arXiv:2312.03300*, 2023.
- Yizhe Zhu. On the second eigenvalue of random bipartite biregular graphs. *Journal of Theoretical Probability*, 36(2):1269–1303, 2023.

Appendix A. Numerical experiments

To illustrate our results, we considered a rectangular matrix M of size $n \times m$, of the form

$$M = \frac{1}{\sqrt{mn}} \left(\nu_1 \mathbf{1}_n \mathbf{1}_m^\top + \nu_2 uv^\top \right) \quad (22)$$

where u and v are vectors with entries in $\{-1, 1\}$ orthogonal to $\mathbf{1}_n$ and $\mathbf{1}_m$, respectively. In this setting, the thresholds θ_1 and θ_2 are $\theta_1 = \sqrt{\frac{\nu_1^2 + \nu_2^2}{d}}$ and $\theta_2 = \frac{\nu_1 + \nu_2}{d}$. For all our experiments, we fixed $d = 3$, $\nu_1 = 1$, $\nu_2 = 0.9$, and $m = n^2$ to mimic the unfolding of a tensor of order 3.

The left part of Figure 2 shows the spectrum of the non-backtracking wedge matrix B constructed as in (3), with $n = 4000$. The structure shown in Theorem 2 is easily apparent: all but two of the eigenvalues are confined in a circle of radius close to θ^2 , with two real *outlier* eigenvalues around ν_1^2 and ν_2^2 , respectively. In the right part, we plot the reduced eigenvector ζ_2^L associated with λ_2 for $n = 30000$ against the theoretical vector u . We can check that $\frac{\langle \text{sign}(\zeta_2^L), u \rangle}{n} \approx 0.845$, which corresponds to recovering 92.3% of the entries of u . On the other hand, Figure 3 shows the same




Figure 2: **Left:** Spectrum of the non-backtracking wedge matrix of the matrix M in (22). The bulk eigenvalues are in orange, while the outliers λ_1 and λ_2 are in green. The circle of radius θ^2 and the theoretical locations at ν_1^2 and ν_2^2 are plotted in solid black and dashed light purple, respectively. **Right:** Reduced second eigenvector ζ_2^L of B (see (11) for the embedding), in green. The theoretical singular vector u is plotted in black. Both vectors are rescaled to have norm \sqrt{n} .




Figure 3: Spectrum density (left) and second eigenvector (right) of the matrix $Z = AA^T - \text{diag}(AA^T)$. The second eigenvector of Z is normalized and plotted against the vector u as in Figure 2.

plots for the matrix $Z = AA^T - \text{diag}(AA^T)$, widely used in tensor completion works (Montanari and Sun, 2018; Cai et al., 2022). The spectrum of Z is less well-behaved than the one of B , and in particular, there is no visible outlier. This is confirmed by plotting the second eigenvector ξ_2 of Z against u : the former is very localized, with no visible correlation with u . This time, the sign correlation between ξ_2 and u is $\frac{\langle \text{sign}(\xi_2), u \rangle}{n} \approx 0.01 \approx \frac{1.65}{\sqrt{n}}$, which is indistinguishable from the variance of a random guess.

Appendix B. Spectral structure of the matrix B

In this section, we state the main intermediary result used in our paper, which characterizes almost fully the structure of the top eigenvectors of B . Theorems 2 and 3 then stem from a perturbation analysis already done in (Stephan and Massoulié, 2022; Bordenave et al., 2022b; Stephan and Zhu, 2022), which we will sketch here.

B.1. Preliminaries

We begin with a few important definitions and properties. Define the *start* and *terminal* matrices $S \in \mathbb{R}^{V_1 \times E}$ and $T \in \mathbb{R}^{E \times V_1}$ by

$$S_{xe} = \mathbf{1}_{e_1=x} \quad \text{and} \quad T_{ex} = \mathbf{1}_{e_3=x}.$$

This implies that for any vector $\phi \in \mathbb{R}^{V_1}$, $[S^*\phi](e) = \phi_{e_1}$ and $[T\phi](e) = \phi_{e_3}$. We shall also need the so-called *edge inversion operator* J and a diagonal weight operator Δ , defined as

$$J_{e,f} = \mathbf{1}\{f = e^{-1}\} \quad \text{and} \quad \Delta_{e,e} = A_{e_1e_2}A_{e_3e_2},$$

where $e^{-1} = (e_3, e_2, e_1)$. As a shortcut, we will denote

$$S_\Delta = S\Delta, \quad T_\Delta = \Delta T, \quad J_\Delta = \Delta J = J\Delta \quad (23)$$

the weighted versions of S, T, J respectively; and for any vector $\chi \in \mathbb{R}^E$, we let

$$\check{\chi} = J_\Delta \chi. \quad (24)$$

Those matrices are related by the following identities:

$$S = T^*J, \quad T = S^*J, \quad (25)$$

$$S_\Delta T = ST_\Delta = AA^*, \quad B = TS_\Delta - J_\Delta. \quad (26)$$

In particular, those identities imply the following property, named the *parity-time symmetry* in (Bordenave et al., 2018):

$$J_\Delta B = B^* J_\Delta. \quad (27)$$

Since we aim to recover the left singular vectors of M , it will be useful to consider the “folded” equivalents of M and Q , namely *signal* and *variance* matrices:

$$P = MM^* = \sum_{i=1}^r \nu_i^2 \phi_i \phi_i^* \quad \text{and} \quad \Phi = QQ^* \quad (28)$$

Finally, we define the oriented equivalents

$$\chi_i = T\phi_i \quad (29)$$

of the left singular vectors of M .

B.2. Main technical results

We are now in a position to state our main technical result. Let

$$\ell = \lfloor \varepsilon \log_d(n) \rfloor \quad \text{with} \quad \varepsilon = \frac{(\eta/2) \wedge 1}{25}.$$

The candidate right (resp. left) eigenvectors of B^ℓ are the following: for $i \in [r_0]$,

$$u_i = \frac{B^\ell \chi_i}{\nu_i^{2\ell}} \quad \text{and} \quad \hat{u}_i = \frac{(B^*)^\ell \tilde{\chi}_i}{\nu_i^{2\ell+2}} \quad (30)$$

with associated (pseudo)-eigenvalue $\nu_i^{2\ell}$. We shall collect those eigenvalues and eigenvectors in the matrices $\Sigma^\ell = \text{diag}(\nu_i^{2\ell}) \in \mathbb{R}^{r_0 \times r_0}$, $U, \hat{U} \in \mathbb{R}^{E \times r_0}$. Finally, we define for all $t \geq 0$ the matrix $\Gamma^{(t)} \in \mathbb{R}^{r_0 \times r_0}$ as

$$\Gamma_{ij}^{(t)} = \sum_{s=0}^t \frac{\langle \mathbf{1}, \Phi^s(\phi_i \circ \phi_j) \rangle}{(\nu_i \nu_j d)^{2s}}. \quad (31)$$

Then our results are as follows:

Theorem 10 *There exists an absolute constants $c, C_0 > 0$ such that for $n \geq C_0 K^{12}$, with probability at least $1 - cd^4 n^{-1/4}$, the following inequalities hold:*

$$\|U^*U - d^2\Gamma^{(\ell)}\| \leq C_1 n^{-\varepsilon/2} \tau^{4\ell}, \quad \|\hat{U}^*\hat{U} - (\Gamma^{(\ell+1)} - I_{r_0})\| \leq C_1 n^{-\varepsilon/2} \tau^{4\ell}, \quad (32)$$

$$\|U^*\hat{U} - I_{r_0}\| \leq C_1 n^{-\varepsilon/2} \tau^{4\ell}, \quad \|\hat{U}^*B^\ell U - \Sigma^\ell\| \leq C_1 n^{-\varepsilon/2} \tau^{4\ell} \theta^{2\ell}, \quad (33)$$

$$\|B^\ell P_{\text{im}(\hat{U})^\perp}\| \leq C_2 \theta^{2\ell}, \quad \|P_{\text{im}(U)^\perp} B^\ell\| \leq C_2 \theta^{2\ell}, \quad (34)$$

with $C_1 = cr\kappa^2 d^6 \log(n)^{7/2}$ and $C_2 = cr\kappa^2 d^6 K^{20} \log^{14}(n)$, where $P_{\text{im}(U)^\perp}$ is the projection operator onto the orthogonal complement of the vector space spanned by column vectors of U .

The statements (32) and (33) are proved in Section E.1, and (34) is proved in Section F.

B.3. Sketch of proof for Theorems 2 and 3

Theorems 2 and 3 stem from Theorem 10 via non-Hermitian perturbation arguments. Details can be found in (Stephan and Massoulié, 2022; Bordenave et al., 2022b). Following the steps in (Stephan and Massoulié, 2022, Section 5.2), the bounds in Theorem 10 imply that B is nearly diagonalized by U, \hat{U} , in the sense that for a constant C_3 depending polynomially on $r, K, \kappa, d, \log n$,

$$\|B^\ell - U^* \Sigma^\ell \hat{U}\| \leq C_3 \theta^{2\ell}$$

and U, \hat{U} are nearly orthogonal. In turn, an application of the Bauer-Fike Theorem (Bauer and Fike, 1960) yields that for $i \in [r_0]$,

$$\lambda_i^\ell = \nu_i^{2\ell} + C_4 \theta^{2\ell} = \nu_i^{2\ell} \left(1 + C_4 \left(\frac{\theta}{\nu_i} \right)^{2\ell} \right),$$

and taking the $1/\ell$ -th power on both sides, we get

$$\lambda_i = \nu_i^2 \left(1 + C_5 \left(\frac{\theta}{\nu_i} \right)^{2\ell} \right),$$

where C_4, C_5 depend polynomially on $r, K, \kappa, d, \log n$. All the other eigenvalues of B satisfy

$$|\lambda| \leq C_5^{1/\ell} \theta^2.$$

This corresponds to Theorem 2.

Now, let $\sigma = C_4 \theta^{2\ell}$. For $i \in [r_0]$, we define $\mathcal{M}_i = \{j \in [r_0] : \nu_j = \nu_i\}$, and the eigengap

$$\delta_i = \min_{j \neq \mathcal{M}_i} |\nu_i^{2\ell} - \nu_j^{2\ell}|.$$

We will assume in the following that $\delta_i \geq 2\sigma$; otherwise, the bounds we show are trivial. Recall the definition of u_i from (30), and the definition of χ_i from (29). By the proof of Theorem 8 in (Stephan and Massoulié, 2022), if ξ_i is an eigenvector of B associated to λ_i , there exists a vector \tilde{u}_i in $\text{vect}(u_j : j \in \mathcal{I})$ such that

$$\left\| \xi_i - \frac{\tilde{u}_i}{\|\tilde{u}_i\|} \right\| \leq \frac{3\sigma}{\delta_i - \sigma} \leq \frac{6\sigma}{\delta_i}.$$

From the definition of $u_i, i \in [r_0]$ in (30), by linearity, we can write

$$\tilde{u}_i = \frac{B^\ell T \tilde{\phi}_i}{\nu_i^{2\ell}},$$

where $\tilde{\phi}_i \in \text{vect}(\phi_j : j \in \mathcal{I})$ is a unit vector. For simplicity, we assume $\tilde{\phi}_i = \phi_i$ (and hence $\tilde{u}_i = u_i$); the proof holds, *mutatis mutandis*, in the general case.

Note that using the definition of S_Δ from (23), $S_\Delta S_\Delta^*$ is a diagonal matrix such that

$$(S_\Delta S_\Delta^*)_{xx} = \sum_{e \in \mathbf{E}: x=e_1} (A_{e_1 e_2} A_{e_3 e_2})^2 \leq \left(\frac{L}{d} \right)^4 \sum_{e_2 \in [m], e_3 \neq x} X_{x e_2} X_{e_3 e_2}.$$

Here $\sum_{e_2, e_3} X_{x e_2} X_{e_3 e_2}$ counts the number of wedges in G starting from x , which is bounded by $cd^2 \log n$ for all $x \in [n]$ for an absolute constant $c > 1$ with probability at least $1 - n^{-1}$ from Proposition 16. Therefore $\|S_\Delta\| \leq cK\theta^2 \sqrt{\log n}$ and

$$\left\| S_\Delta \xi_i - S_\Delta \frac{u_i}{\|u_i\|} \right\| \leq \frac{c\sigma K^2 \theta^2 \sqrt{\log n}}{\delta_i}.$$

On the other hand, using Lemma 22 and (52) in Lemma 23, and recalling matrix relations and definitions in (23), (24), (29), we have

$$\begin{aligned} \langle S_\Delta u_i, \phi_i \rangle &= \langle u_i, \check{\chi}_i \rangle = \nu_i^2 (1 + O(C_0 n^{-\varepsilon/2} \tau_i^{2\ell})) \\ \|S_\Delta u_i\|^2 &= \nu_i^4 \left(\Gamma_{ii}^{(\ell+1)} + O(C_0 n^{-\varepsilon/2} \tau_i^{4\ell}) \right) \end{aligned}$$

for a constant C_0 depending polynomially on $r, K, \kappa, d, \log n$. From (50) in Lemma 23,

$$\|u_i\|^2 = d^2 \Gamma_{ii}^{(\ell)} + O(C_0 n^{-\varepsilon/2} \tau_i^{4\ell}).$$

Since Φ is a positive matrix with spectral radius ρ^2 , from (31),

$$\Gamma_{ii}^{(\ell)} = \left\langle \mathbf{1}, \left(I - \frac{\Phi}{\nu_i^4 d^2} \right)^{-1} (\phi_i \circ \phi_i) \right\rangle + O(\tau_i^{4\ell}) = \gamma_i + O(\tau_i^{4\ell}).$$

Putting all of the previous bounds together, recalling the definition of ζ_i^R from (11), we do find

$$\left\langle \frac{\zeta_i}{\|\zeta_i\|}, \phi_i \right\rangle = \left\langle \frac{S_{\Delta} \xi_i}{\|S_{\Delta} \xi_i\|}, \phi_i \right\rangle = \frac{1}{\sqrt{\gamma_i}} + O\left(\frac{2K^2 d \sqrt{\log n} \sigma}{\delta_i}\right),$$

which is the statement of Theorem 3 upon diving both δ_i and σ by $\nu_i^{2\ell}$. The same holds for the left eigenvectors by replacing u_i by \hat{u}_i in the proof.

Appendix C. Parameter inequalities

In this section, we collect a few simple inequalities that will be useful in the proof. First, since

$$\rho = \|Q\| \leq \sqrt{nm} \max_{xy} |Q_{xy}|,$$

we have $K \geq 1$. Further, using the variational formula for singular values,

$$\rho \geq \frac{1}{\sqrt{mn}} \langle \mathbf{1}, Q\mathbf{1} \rangle = \frac{1}{\sqrt{nm}} \sum_{xy} Q_{xy} = \sum_{x,y} M_{xy}^2 = \|M\|_F^2 \geq \nu_1^2.$$

Hence

$$1 \leq K = \frac{L^2}{\rho} \leq \frac{L^2}{\nu_1^2}.$$

Next, we derive some entrywise bound on $\Phi := QQ^*$. For any x, y ,

$$\Phi_{xy} = (QQ^*)_{xy} = \sum_z Q_{xz} Q_{yz} \leq \frac{K^2 \rho^2}{n}. \quad (35)$$

Let (v_k) be a set of orthonormal left singular vectors for Q with singular values μ_1, \dots, μ_n . For $t \geq 1$,

$$(\Phi^t)_{xy} = [(QQ^*)^t]_{xy} = \sum_k \mu_k^{2t} v_k(x) v_k(y) \quad (36)$$

$$\leq \rho^{2t-2} \sum_k \mu_k^2 |v_k(x)| |v_k(y)| \quad (37)$$

$$\leq \rho^{2t-2} \sqrt{(\Phi)_{xx}} \sqrt{(\Phi)_{yy}} \leq \frac{K^2 \rho^{2t}}{n}, \quad (38)$$

where the last line follows from Cauchy-Schwartz inequality and (35). Similarly, let $\tilde{\Phi} = Q^*Q$, we have

$$(\tilde{\Phi}^t)_{xy} \leq \frac{K^2 \rho^{2t}}{m}. \quad (39)$$

And

$$(\Phi^t Q)_{xy} \leq \frac{K^2 \rho^{2t+1}}{\sqrt{nm}}, \quad (\tilde{\Phi}^t Q^*)_{xy} \leq \frac{K^2 \rho^{2t+1}}{\sqrt{nm}}. \quad (40)$$

Moreover, for any $t \geq 1$ and $v, w \in \mathbb{R}^n$,

$$\langle v, \Phi^t w \rangle = \sum_{xy} (\Phi^t)_{xy} v(x) w(y) \leq \frac{K^2 \rho^{2t}}{n} \|v\|_1 \|w\|_1. \quad (41)$$

Appendix D. Local study of G

D.1. Preliminaries

Rooted graphs and trees We begin with a rigorous definition for the objects considered in this section:

Definition 11 (Labeled rooted graphs) A labeled rooted graph is a triplet $g_\star = (g, o, \iota)$ that consists in:

- a graph $g = (V, E)$,
- a distinguished vertex (or root) $o \in V$,
- a labeling function $\iota : V \rightarrow \mathbb{N}$.

The set of labeled rooted graphs is denoted by \mathcal{G}_\star .

When the labeling function ι is clear, we will use the shortcut (g, o) for g_\star . In particular, this will be the case whenever $V \subseteq \mathbb{N}$ and $\iota(x) = x$. The notion of induced subgraph extends naturally to rooted and/or labeled graphs, provided that the induced subgraph does contain the root o .

An important object of study will be the bipartite Galton-Watson tree defined as follows:

Definition 12 (Bipartite Galton-Watson tree) The bipartite Galton-Watson tree (T, ρ) is a random tree (T, o, ι) defined as follows:

- the root of the tree (depth 0) has label $\iota(o) = \rho$,
- the number of children of each vertex $x \in T$ is independent, and is defined as follows:
 - if x has even depth $2h$, x has $\text{Poi}(d^2)$ children at depth $2h + 1$,
 - if x has odd depth $2h + 1$, x has exactly one child at depth $2h + 2$.
- each vertex x of even depth has a label $\iota(x) \sim \text{Unif}([n])$, and each vertex of odd depth has a label $\iota(x) \sim \text{Unif}([m])$.

Tangle-free neighbourhoods and graphs For a rooted graph (g, o) , we denote by $(g, o)_t$ the subgraph induced by the vertices within distance t from o . The boundary of this subgraph, spanned by the vertices at distance exactly t from o , will be denoted by $\partial(g, o)_t$. The notion of *tangle-freeness*, which will be crucial in our analysis, is defined as follows:

Definition 13 (Tangle-freeness) *A rooted graph (g, o) is said to be tangle-free if it contains at most one cycle, otherwise, it is called tangled. For a given $t \in \mathbb{N}$, a graph g is said to be tangle-free if for any vertex $x \in g$, the neighborhood $(g, x)_t$ is tangle-free.*

D.2. Bounding the neighborhood sizes

The main issue in bounding the neighborhood sizes is that a typical left vertex x has approximately $mp = d\sqrt{\alpha}$ neighbors, which diverges as $\alpha \rightarrow \infty$. However, most of those neighbors have degree one and thus do not matter in the spectrum of B ! We formalize this by defining the bipartite graph $\bar{G} = (V_1, \bar{V}_2, \bar{E})$ induced by V_1 and the non-leaf vertices in V_2 :

$$\bar{V}_2 := \{y \in V_2 : \deg(y) \geq 2\}.$$

We are then able to show the following proposition:

Proposition 14 *There exists constants $c_0, c_1 > 0$ such that for all $x \in [n]$ and $s > 0$,*

$$\mathbb{P}(\forall t \geq 0, \quad |\partial(\bar{G}, x)_{2t}| \leq sd^{2t} \text{ and } |\partial(\bar{G}, x)_{2t+1}| \leq sd^{2t+2}) \geq 1 - c_0e^{-c_1s}$$

Proof We define

$$\tilde{p} = \mathbb{P}(\text{Bin}(n, p) \geq 1) = 1 - (1 - p)^n, \quad \tilde{d}_1 = mp\tilde{p}, \quad \text{and} \quad \tilde{d}_2 = \tilde{d}_1(1 + np).$$

We shall actually show the following bound: for some constants c_0, c_1, c_2, c_3, c_4 , and $s > 1$,

$$\mathbb{P}\left(\forall t \geq 0, \quad |\partial(\bar{G}, x)_{2t}| \leq c_1s\tilde{d}_2^t \text{ and } |\partial(\bar{G}, x)_{2t+1}| \leq c_2s\tilde{d}_1\tilde{d}_2^t\right) \geq 1 - c_3e^{-c_4s} \quad (42)$$

We first show how this implies Proposition 14: for $s > 1$ the bound is vacuous as soon as $t > c \log_d(n)$ for some constant c , so we can always assume otherwise. Second, we have

$$\tilde{d}_1 = d^2 + O(d^3/\sqrt{\alpha}) \quad \text{and} \quad \tilde{d}_2 = d^2 + O(d^3/\sqrt{\alpha}),$$

and hence for the timescale considered for t we have $\tilde{d}_2^t \leq cd^{2t}$ under the assumption (9). Finally, we can adjust c_0 and c_1 such that $c_0e^{-c_1} \geq 1$ and hence the bound also becomes vacuous when $s \leq 1$.

We now move to prove (42). Define the following sequences:

$$\begin{aligned} \varepsilon_{2t} &= \max\left(\tilde{d}_1^{-1}\tilde{d}_2^{-t+1}(2t), \sqrt{np}\tilde{d}_1^{-1/2}\tilde{d}_2^{-(t-1)/2}\sqrt{2t}\right), \\ \varepsilon_{2t+1} &= \tilde{d}_2^{-t/2}\sqrt{2t+1} \quad \text{and} \quad f_t = \prod_{t'=1}^t (1 + \varepsilon_{t'}). \end{aligned}$$

Due to Assumption (8), by definition of ε_t, f_t , there exists some constants c_0, c_1, c_2 such that

$$c_0 \leq f_t \leq c_1 \quad \text{and} \quad \varepsilon_t \leq c_1.$$

We also define the following upper-bound sequence:

$$S_{2t} = f_{2t} \tilde{d}_2^t \quad \text{and} \quad S_{2t+1} = f_{2t+1} \tilde{d}_1 \tilde{d}_2^t,$$

which satisfies the following recursion equation:

$$S_{2t+1} = (1 + \varepsilon_{2t+1}) \tilde{d}_1 S_{2t} \quad \text{and} \quad S_{2t+2} = (1 + \varepsilon_{2t+2})(1 + np) S_{2t+1}.$$

We will also need the following lemma, which stems from a classical Chernoff bound:

Lemma 15 *Let (Y_1, \dots, Y_T) be i.i.d binomial random variables, with mean μ . Then, for any $u \geq 0$,*

$$\mathbb{P}\left(\sum_{i=1}^T Y_i \geq (1+u)\mu T\right) \leq e^{-T\mu\gamma(u)},$$

where $\gamma(u) = (1+u)\log(1+u) - u$. The function $\gamma(u)$ is strictly increasing, and satisfies the inequality

$$\gamma(u) \geq \min\left(\frac{u^2}{4}, u\right) \quad (43)$$

For a given vertex $x \in \partial(\bar{G}, x)_{2t}$, the number of neighbors of x in $\bar{V}_2 \setminus (\bar{G}, x)_{2t}$ is dominated by $\text{Bin}(m, p\tilde{p})$. For any $t \geq 0$, we can apply Lemma 15 with $u = \varepsilon_{2t+1}$ to get

$$\mathbb{P}(|\partial(\bar{G}, x)_{2t+1}| \geq sS_{2t+1} \mid |\partial(\bar{G}, x)_{2t}| \leq sS_{2t}) \leq \exp\left(-sf_{2t}\tilde{d}_1\tilde{d}_2^t\gamma(\varepsilon_{2t+1})\right).$$

Equation (43) implies that when $u \leq c_1$, $\gamma(u) \geq c_2u^2$ for some constant u^2 , and hence

$$\mathbb{P}(|\partial(\bar{G}, x)_{2t+1}| \geq sS_{2t+1} \mid |\partial(\bar{G}, x)_{2t}| \leq sS_{2t}) \leq \exp(-cs(2t+1)).$$

Now, the number of neighbors of a vertex in $\partial(\bar{G}, x)_{2t+1}$ in V_1 is dominated by a $X \sim \text{Bin}(n, p)$ random variable, conditioned on $X \geq 1$, which is itself dominated by a $1 + \text{Bin}(n, p)$ variable (see, e.g., (Broman et al., 2011, Proposition 1.1)). This time, we apply Lemma 15 by noticing that

$$\mathbb{P}\left(\sum_{i=1}^T (1 + Y_i) \geq (1+u)(1+\mu)T\right) = \mathbb{P}\left(\sum_{i=1}^T Y_i \geq \left(1+u+\frac{u}{\mu}\right)\mu T\right),$$

so that

$$\mathbb{P}(|\partial(\bar{G}, x)_{2t+2}| \geq sS_{2t+2} \mid |\partial(\bar{G}, x)_{2t+1}| \leq sS_{2t+1}) \leq \exp\left(-snpf_{2t+1}\tilde{d}_1\tilde{d}_2^t\gamma(\varepsilon_{2t+2}/np)\right).$$

We can now notice that by the bound on γ in (43) and definition of ε_{2t+2} , we have

$$np\tilde{d}_1\tilde{d}_2^t\gamma(\varepsilon_{2t+2}/np) \geq c(2t+2).$$

Hence, we have shown that for any $t \geq 0$,

$$\mathbb{P}(|\partial(\bar{G}, x)_{t+1}| \geq sS_{t+1} \mid |\partial(\bar{G}, x)_t| \leq sS_t) \leq \exp(-cs(t+1)).$$

This implies that

$$\mathbb{P}(\exists t \geq 0, |\partial(\bar{G}, x)_t| \geq S_t) \leq \sum_{t \geq 0} \exp(-cs(t+1)) = \frac{e^{-cs}}{1 - e^{-cs}},$$

from which (42) ensues since $s > 1$. ■

With the same arguments as in (Bordenave et al., 2018, Lemma 29), this tail bound implies the following proposition:

Proposition 16 *With probability at least $1 - 1/n$, we have for all $x \in [n]$,*

$$|(\bar{G}, x)_{2t}| \leq c \log(n) d^{2t} \quad \text{and} \quad |(\bar{G}, x)_{2t+1}| \leq c \log(n) d^{2t+2}.$$

Further, we have for all $x \in V_1$, $p \geq 1$,

$$\mathbb{E}[|(\bar{G}, x)_{2t}|^p]^{1/p} \leq cpd^{2t},$$

and

$$\mathbb{E}\left[\max_{x \in [n], t \geq 0} |(\bar{G}, x)_{2t}|^p\right] \leq (cp)^p + (c \log(n))^p.$$

D.3. Tree approximation

We now couple the neighborhoods of G with an appropriately defined tree process. Central to this section is the following breadth-first exploration process of a vertex $x \in [n]$:

- start with $A_0 = \{x\}$
- at step t , take any vertex $x_t \in A_t$ among those closest to x
 - if $x_t \in V_2$, we let N_t be the set of all neighbors of x_t in $V_1 \setminus \bigcup_{s \leq t} A_s$
 - if $x \in V_1$, N_t is the set of neighbors of x_t in $V_2 \setminus \bigcup_{s \leq t} A_s$ that have a neighbour in $V_1 \setminus \bigcup_{s \leq t} A_s$.
- update $A_{t+1} = A_t \cup N_t \setminus \{x_t\}$ and continue.

We denote by \mathcal{F}_t the filtration adapted to this process and call an edge $\{e_1, e_2\}$ a *discovered edge* if for some $t \geq 0$ we have $x_t = e_1$ and $e_2 \in N_t$. We begin with a result about the tree-like properties of neighborhoods in G :

Proposition 17 *For a given vertex x , and $h \geq 0$, we have*

$$\mathbb{P}((G, x)_{2h} \text{ contains a cycle}) \leq \frac{cd^{2h+2}}{n}.$$

Further, with probability at least $1 - cd^{4h+4}/n$, the graph G is *h -tangle-free*: for every vertex x , $(G, x)_{2h}$ contains at most one cycle.

Proof Consider the exploration process outlined above, and let τ be the first time at which all the vertices in $(\bar{G}, x)_h$ have been revealed. By construction, τ is a stopping time for the filtration \mathcal{F}_t , and given \mathcal{F}_τ , the discovered edges of $(G, x)_h$ form a spanning tree. Hence, there are two possibilities for a cycle in $(G, x)_h$:

- there is an edge between a vertex of $(G, x)_h \cap V_1$ and $(G, x)_h \cap V_2$,
- there exists a vertex $z \in V_2 \setminus (G, x)_h$ and two vertices $x_1, x_2 \in (G, x)_h \cap V_1$ such that z is connected to both x_1 and x_2 .

The number of such events of the first type is stochastically dominated by a $\text{Bin}(|(G, x)_h|^2, p)$ variable and the second by a $\text{Bin}(|(G, x)_h|^2, mp^2)$. Further, given \mathcal{F}_τ , those two types of events are independent. Recall the following simple bounds:

$$\mathbb{P}(\text{Bin}(N, q) \geq 1) \leq Nq \quad \text{and} \quad \mathbb{P}(\text{Bin}(N, q) \geq 2) \leq N^2q^2.$$

Then,

$$\mathbb{P}((G, x)_h \text{ contains a cycle} \mid \mathcal{F}_\tau) \leq |(\bar{G}, x)_h|^2(mp^2 + p).$$

and similarly

$$\mathbb{P}((G, x)_h \text{ is tangled} \mid \mathcal{F}_\tau) \leq |(\bar{G}, x)_h|^4(m^2p^4 + p^2).$$

Taking expectations in both cases and using Proposition 16,

$$\mathbb{P}((G, x)_{2h} \text{ contains a cycle}) \leq \frac{cd^{2h+2}}{n} \quad \text{and} \quad \mathbb{P}((G, x)_{2h} \text{ is tangled}) \leq \frac{cd^{4h+4}}{n^2}.$$

The second statement of the proposition easily ensues from a union bound. \blacksquare

We now prove a coupling result between the neighborhoods of G and the Galton-Watson trees introduced in Definition 12. Recall that the total variation between two probability measures $\mathbb{P}_1, \mathbb{P}_2$ is defined as

$$d_{\text{TV}}(\mathbb{P}_1, \mathbb{P}_2) = \min_{\mathbb{P} \in \pi(\mathbb{P}_1, \mathbb{P}_2)} \mathbb{P}(X_1 \neq X_2),$$

where $\pi(\mathbb{P}_1, \mathbb{P}_2)$ is the set of all *couplings* of \mathbb{P}_1 and \mathbb{P}_2 , i.e. joint distributions of random variables (X_1, X_2) such that $X_1 \sim \mathbb{P}_1$ and $X_2 \sim \mathbb{P}_2$.

Proposition 18 *There exists a constant $c \geq 0$ such that for any $h \geq 0, x \in V_1$, we have*

$$d_{\text{TV}}(\mathcal{L}((G, x)_h), \mathcal{L}((T, x)_h)) \leq c \log(n)^2 \frac{d^{4h+3}}{n \wedge \sqrt{\alpha}}.$$

Proof We proceed in two steps: we first couple the unlabeled versions of the graphs and then move on to the labels. We shall use the following classical bounds for total variation distances:

$$d_{\text{TV}}(\text{Bin}(n, p), \text{Poi}(np)) \leq p \quad \text{and} \quad d_{\text{TV}}(\text{Poi}(\lambda), \text{Poi}(\lambda')) \leq |\lambda - \lambda'|. \quad (44)$$

Firstly, from Proposition 17, $(G, x)_h$ is a tree with probability $1 - cd^{2h+2}/n$, so we only need to couple the offspring distributions. Consider the exploration process $(x_t)_{t \geq 0}$ defined above; at each step t , we let n_t (resp. m_t) be the number of vertices in $V_1 \setminus \bigcup_{t' < t} A_{t'}$ (resp. $V_2 \setminus \bigcup_{t' < t} A_{t'}$). We aim to couple it with the same exploration process $(x'_t)_{t \geq 0}$ on (T, x) . Let \mathbb{P}_t (resp. \mathbb{Q}_t) be the offspring distribution of x_t (resp. x'_t). By definition, we have an explicit definition for \mathbb{P}_t and \mathbb{Q}_t :

- if x_t has even depth, $\mathbb{P}_t = \text{Bin}(m_t, p\tilde{p}_t)$ where $\tilde{p}_t = \mathbb{P}(\text{Bin}(n_t, p) \geq 1)$, and $\mathbb{Q}_t = \text{Poi}(d^2)$;
- if x_t has odd depth, $\mathbb{P}_t = \text{Bin}(n_t, p)$ conditioned on being at least 1, and $\mathbb{Q}_t = 1$ a.s.

We begin with the second case, as it is the simplest. As before, \mathbb{P}_t is dominated by a distribution equal to $1 + \text{Bin}(n, p)$, and hence

$$\mathbb{P}_t(|N_t| > 1) \leq \mathbb{P}(\text{Bin}(n, p) \geq 1) \leq np = d\alpha^{-1/2}.$$

We now move to the first case. We have

$$\tilde{p}_t = 1 - (1 - p)^{n_t} = pn_t + O(p^2 n^2) = O(np).$$

Using the bounds in eq. (44),

$$d_{\text{TV}}(\mathbb{P}_t, \mathbb{Q}_t) \leq p\tilde{p}_t + |m_t p\tilde{p}_t - d^2| \leq \frac{cd^2}{m} + |m_t n_t p^2 - d^2| + \frac{cd^3}{\sqrt{\alpha}}.$$

The middle term can be further decomposed by noticing that $nmp^2 = d^2$:

$$|m_t n_t p^2 - d^2| \leq |m_t - m|np^2 + m|n_t - n|p^2 \leq \frac{2d^2 |(G, x)_{2h}|}{n}$$

All in all, we showed that for any $t \geq 0$,

$$d_{\text{TV}}(\mathbb{P}_t, \mathbb{Q}_t) \leq cd^2 \left(\frac{|(G, x)_{2h}|}{n} \vee \frac{d}{\sqrt{\alpha}} \right). \quad (45)$$

Let $\tilde{\mathcal{L}}$ denote the unlabeled distribution of the graphs. By the high-probability bounds of Proposition 16, we obtain

$$\begin{aligned} d_{\text{TV}} \left(\tilde{\mathcal{L}}((G, x)_{2h}), \tilde{\mathcal{L}}((T, x)_{2h}) \right) &\leq \underbrace{\frac{cd^{2h+2}}{n}}_{(G, x)_h \text{ is not a tree}} + \underbrace{\frac{1}{n}}_{\text{Prop. 16 does not hold}} \\ &+ c \log(n)^2 d^{2h+2} \left(\frac{d^{2h}}{n} \vee d\alpha^{-1/2} \right) \\ &\leq c' \log(n)^2 \frac{d^{4h+3}}{n \wedge \sqrt{\alpha}}, \end{aligned}$$

where in the third part of the first inequality, we apply (45) with a union bound over all vertices in $(G, x)_{2h}$. This bound is lower than the one in Proposition 18 since $d > 1$.

Consider now the distribution of the labels. Given a coupling between the unlabeled versions of $(G, x)_{2h}$ and $(T, x)_{2h}$, the labels of the even (resp. odd) depth vertices of $(G, x)_{2h}$ are obtained by sampling without replacement from $[n]$ (resp. $[m]$), whereas those of $(T, x)_{2h}$ are sampled with replacement. From (Freedman, 1977), the total variation distance between sampling with and without replacement k elements from a population of N elements is bounded above by k^2/N , and hence

$$d_{\text{TV}}(\mathbb{P}, \mathbb{Q}) \leq 2 \frac{|(G, x)_{2h}|^2}{n},$$

where \mathbb{P}, \mathbb{Q} are the labels distributions for $(G, x)_{2h}, (T, x)_{2h}$, respectively. A final application of Proposition 16 completes the proof. \blacksquare

D.4. Quantitative bounds for local functionals

The bounds obtained in Proposition 18 can be viewed as a quantitative version of the local convergence of Benjamini and Schramm (Benjamini and Schramm, 2001), for possibly diverging depth. It is known that this coupling also implies weak convergence in the probability measure sense, i.e., the convergence of graph functionals. We now provide a quantitative version of this weak convergence, but only for a specific class of functionals, called *local* functionals:

Definition 19 (Local functionals) A functional $f : \mathcal{G}_* \rightarrow \mathbb{R}$ is said to be t -local if $f(G, o)$ is only a function of $(G, o)_t$.

Proposition 20 Let $f : \mathcal{G}_* \rightarrow \mathbb{R}$ be a $2h$ -local function for some $h \geq 0$, such that $f(g, o) \leq a|(g, o)_{2h}|^b$ for some a, b . Then with probability at least $1 - n^{-1}$,

$$\left| \sum_{x \in V_1} f(G, x) - \sum_{x \in [n]} \mathbb{E}[f(T, x)] \right| \leq ca \log(n)^{3/2+b} \frac{d^{2h(1+b)}}{\sqrt{n} \wedge \alpha^{1/4}}.$$

Proof Notice that the bounds on the neighborhood sizes match exactly the ones of (Bordenave et al., 2022b), and hence we can borrow the following concentration result from it (see the proof of Theorem 12.5 in (Bordenave et al., 2022b)): with probability at least $1 - n^{-1}$,

$$\left| \sum_{x \in V_1} f(G, x) - \sum_{x \in [n]} \mathbb{E}[f(G, x)] \right| \leq ca \log(n)^{3/2+b} d^{2h(1+b)} \sqrt{n}. \quad (46)$$

On the other hand, we have for any $x \in V_1$

$$|\mathbb{E}f(T, x) - \mathbb{E}f(G, x)| = \mathbb{E}[|f(T, x) - f(G, x)| \mathbf{1}_{\mathcal{E}(x)}],$$

where $\mathcal{E}(x)$ denotes the event of the coupling between $(G, x)_{2h}$ and $(T, x)_{2h}$ fails. From the Cauchy-Schwarz inequality,

$$\begin{aligned} |\mathbb{E}f(T, x) - \mathbb{E}f(G, x)| &\leq \sqrt{\mathbb{P}(\mathcal{E}(x))} \sqrt{\mathbb{E}[(f(T, x) - f(G, x))^2]} \\ &\leq \sqrt{d_{\text{TV}}(\mathcal{L}((G, x)_h), \mathcal{L}((T, x)_h))} \left(\sqrt{\mathbb{E}[|f(G, x)|^2]} + \sqrt{\mathbb{E}[|f(T, x)|^2]} \right). \end{aligned}$$

Proposition 18 bounds the first factor. It is easy to check that the bounds for $|(G, x)_{2h}|$ in Proposition 16 also apply to $|(T, x)_{2h}|$, which yields

$$|\mathbb{E}f(T, x) - \mathbb{E}f(G, x)| \leq ca \log(n) d^{2bh+1} \frac{d^{2h}}{\sqrt{n} \wedge \alpha^{1/4}}. \quad (47)$$

Combining (46) and (47) yields the desired bound. ■

Appendix E. From tree functionals to pseudo-eigenvectors

Since we now have a way to translate quantities from the Galton-Watson tree T to the graph G , we aim to apply these results to specific functionals. We begin with a few preliminaries: given a functional $f : \mathcal{G}_\star \rightarrow \mathbb{R}$, define its *Galton-Watson transform* \bar{f} as

$$\bar{f}(x) = \mathbb{E}[f(T, x)].$$

This definition is consistent since from Definition 12, the distribution of a Galton-Watson tree only depends on its root label x . We shall also need the following definition: given a matrix $W \in \mathbb{R}^{n \times m}$, and a functional $f : \mathcal{G}_\star \rightarrow \mathbb{R}$, define $\partial_W f$ as

$$\partial_W f(g, o) = \sum_{o', o'' \in \mathcal{N}_2(o)} W_{\iota(o)\iota(o')} W_{\iota(o'')\iota(o')} f(g \setminus \{o, o'\}, o'')$$

where o', o'' runs over the set of non-backtracking wedges starting from o . In particular, when g is a rooted tree, o' is uniquely determined by o'' , and the graph $(g \setminus \{o, o'\}, o'')$ is the subtree rooted at o'' .

The functionals we will consider are of the following form: for a vector ϕ ,

$$f_{\phi, t}(g, o) = \left(\frac{mn}{d^2}\right)^t \sum_{o_1, \dots, o_{2t}} \left(\prod_{s=0}^{t-1} M_{\iota(o_{2s}), \iota(o_{2s+1})} M_{\iota(o_{2s+2}), \iota(o_{2s+1})} \right) \phi_{\iota(o_{2t})} \quad (48)$$

where the sum runs over all non-backtracking paths of length $2t$ starting from o . Then, the following proposition holds:

Proposition 21 *For any $t \geq 0$, and $i, j \in [r]$,*

$$\begin{aligned} \overline{f_{\phi_i, t}} &= \nu_i^{2t} \phi_i, \\ \overline{f_{\phi_i, t} f_{\phi_j, t}} &= (\nu_i \nu_j)^{2t} \left(\sum_{s=0}^t \frac{\Phi^s}{(\nu_i \nu_j d)^{2s}} \right) (\phi_i \circ \phi_j) \end{aligned}$$

and if $F_{\phi_i, t} = (f_{\phi_i, t+1} - \nu_i^2 f_{\phi_i, t})^2$, then

$$\overline{F_{\phi_i, t}} = \frac{\Phi^t (\phi_i \circ \phi_i)}{d^{2t}}$$

Additionally, for any functional $f : \mathcal{G}_\star \rightarrow \mathbb{R}$, and any weight matrix W ,

$$\overline{\partial_W f} = \frac{d^2 (W W^*)}{mn} \bar{f}.$$

Proof We proceed by reduction to (Stephan and Massoulié, 2022). Let \tilde{T} be an ordinary Galton-Watson tree with offspring distribution $\text{Poi}(d^2)$, such that each vertex has a random label $\text{Unif}([n])$, and define the random weights

$$\tilde{W}_{xy} = \frac{mn}{d^2} M_{xZ} M_{yZ}, \quad Z \sim \text{Unif}([m])$$

The equivalent tree functional $\tilde{f}_{\phi,t}$ is defined as

$$\tilde{f}_{\phi,t}(g, o) = \sum_{o_1, \dots, o_t} \left(\prod_{s=0}^{t-1} \tilde{W}_{\iota(o_s), \iota(o_{s+1})} \right) \phi_{\iota(o_t)}.$$

It is easy to see that

$$f_{\phi,t}(T, x) \stackrel{d}{=} \tilde{f}_{\phi,t}(\tilde{T}, x).$$

On the other hand, the process \tilde{f} on (\tilde{T}, x) was already studied in depth in (Stephan and Massoulié, 2022). It corresponds to the signal matrix

$$\tilde{Q} = \frac{d^2 \mathbf{1}\mathbf{1}^\top}{n} \circ \left(\frac{mn}{d^2} \cdot \frac{MM^*}{m} \right) = P$$

and the variance matrix

$$\tilde{K} = \frac{d^2 \mathbf{1}\mathbf{1}^\top}{n} \circ \left(\frac{m^2 n^2 (M \circ M)(M \circ M)^*}{d^4 m} \right) = \frac{\Phi}{d^2}.$$

Proposition 21 then ensues from Propositions 6 and 7 in (Stephan and Massoulié, 2022). ■

E.1. Pseudo-eigenvectors

We now make use of Proposition 20 to translate the expectations on T to pseudo-eigenvectors of B . This corresponds to Equations (32)-(33) of Theorem 10. We shall prove stronger versions of those bounds in the following two lemmas and show how they imply the desired inequalities.

Lemma 22 *Let $h = c \log_d(n)$. With probability at least $1 - c' d^4 n^{12c-1}$, for any $t \leq 3h$, we have*

$$|\langle B^t \chi_i, \check{\chi}_j \rangle - \nu_i^{2t+2} \delta_{ij}| \leq c' \kappa^2 \theta_2^{2t+2} d^{12h+4} \frac{\log(n)^{5/2}}{\sqrt{n} \wedge \alpha^{1/4}}.$$

Proof Consider the following functional:

$$f(g, o) = \mathbf{1}_{(g,o)_t \text{ is tangle-free}} \phi_j(o) f_{\phi_i, t+1}(g, o),$$

where $f_{\phi_i, t}$ was defined in (48). It is easy to check that f is $(2t+2)$ -local, and when (g, o) is $3h$ -tangle-free (which happens with probability at least $1 - c' d^4 n^{12c-1}$ from Proposition 17), there are at most two non-backtracking paths between o and any vertex of (g, o) , hence from (48),

$$|f(g, o)| \leq 2 \frac{\kappa^2}{n} \theta_2^{2t+2} |(g, o)_{2t+2}|. \quad (49)$$

We are, therefore, in the setting of Proposition 20. It remains to compute the corresponding quantities on (G, x) and (T, x) . First, by definition,

$$\sum_{x \in [n]} f(G, x) = \langle B^t \chi_i, \check{\chi}_j \rangle,$$

and by Proposition 21 we have

$$\sum_{x \in [n]} f(T, x) = \sum_{x \in [n]} \phi_j(x) \nu_i^{2t+2} \phi_i(x) = \nu_i^{2t+2} \delta_{ij},$$

which completes the proof. \blacksquare

Recall the definition of $\Gamma_{ij}^{(t)}$ from (31). The following bounds hold:

Lemma 23 *Let $h = c \log_d(n)$. With probability at least $1 - c' d^4 n^{4c-1}$, for any $t \leq h$, we have*

$$\left| \langle B^t \chi_i, B^t \chi_j \rangle - \nu_i^{2t} \nu_j^{2t} d^{2t} \Gamma_{ij}^{(t)} \right| \leq c' d^6 \kappa^2 \log(n)^{7/2} \theta_2^{4t} \frac{d^{6t}}{\sqrt{n} \wedge \alpha^{1/4}}, \quad (50)$$

$$\left| \langle (B^*)^t \check{\chi}_i, (B^*)^t \check{\chi}_j \rangle - \nu_i^{2t+2} \nu_j^{2t+2} \left(\Gamma_{ij}^{(t+1)} - \delta_{ij} \right) \right| \leq c' d^6 \kappa^2 \log(n)^{7/2} \theta_2^{4t+4} \frac{d^{6t}}{\sqrt{n} \wedge \alpha^{1/4}}, \quad (51)$$

$$\left| \langle S_\Delta B^t \chi_i, S_\Delta B^t \chi_j \rangle - \nu_i^{2t+2} \nu_j^{2t+2} \Gamma_{ij}^{(t+1)} \right| \leq c' d^6 \kappa^2 \log(n)^{7/2} \theta_2^{4t+4} \frac{d^{6t}}{\sqrt{n} \wedge \alpha^{1/4}}. \quad (52)$$

Proof In the following, it will be useful to define the following functional, this time acting on wedge-rooted graphs:

$$\mathbf{f}_{\phi,t}(g, e) = \left(\frac{mn}{d^2} \right)^t \sum_{o_2=e_3, \dots, o_{2t}} \left(\prod_{s=1}^t M_{l(o_{2s}), l(o_{2s+1})} M_{l(o_{2s+2}), l(o_{2s+1})} \right) \phi_{l(o_{2t+2})}, \quad (53)$$

where e is a wedge and the sum runs over all tuples (o_i) such that $e_1, e_2, o_2, \dots, o_{2t+2}$ is a non-backtracking path of length $2t+2$. First, let

$$f(g, o) = \sum_{e: e_1=o} \mathbf{f}_{\phi_i,t}(g, e) \mathbf{f}_{\phi_j,t}(g, e).$$

Again, f is $(2t+2)$ -local, and

$$|f(g, o)| \leq 2 \frac{\kappa^2}{n} \theta_2^{4t} |(g, o)_{2t+2}|^2.$$

Additionally, it is easy to check that

$$\sum_{x \in V_1} f(G, x) = \langle B^t \chi_i, B^t \chi_j \rangle,$$

and that, letting W be the all-one matrix,

$$f(T, x) = \partial_W (f_{\phi_i,t} f_{\phi_j,t}),$$

hence using Proposition 21,

$$\bar{f} = \frac{d^2 \mathbf{1} \mathbf{1}^*}{n} (\nu_i \nu_j)^{2t} \left(\sum_{s=0}^t \frac{\Phi^s}{(\nu_i \nu_j d)^{2s}} \right) (\phi_i \circ \phi_j).$$

Taking the scalar product of the above equation with $\mathbf{1}$, and using Proposition 20, concludes the proof of (50).

For (51), the parity-time symmetry (27) implies that

$$\langle (B^*)^t \tilde{\chi}_i, (B^*)^t \tilde{\chi}_j \rangle = \langle \Delta B^t \chi_i, \Delta B^t \chi_j \rangle.$$

The corresponding graph functional is now given by

$$f(g, o) = \frac{mn}{d^2} \sum_{e: e_1=o} M_{e_1 e_2} M_{e_3 e_2} \mathbf{f}_{\phi_i, t}(g, e) \mathbf{f}_{\phi_j, t}(g, e),$$

which yields

$$\bar{f} = \Phi(\nu_i \nu_j)^{2t} \left(\sum_{s=0}^t \frac{\Phi^s}{(\nu_i \nu_j d)^{2s}} \right) (\phi_i \circ \phi_j) = (\nu_i \nu_j)^{2t+2} \left(\Gamma_{ij}^{(t+1)} - \delta_{ij} \right).$$

The rest of the proof proceeds as above. Finally, the last equation corresponds to

$$f(g, o) = f_{\phi_i, t}(g, o) f_{\phi_j, t}(g, o),$$

and proceeds identically. ■

Now we are ready to prove Equations (32)-(33) of Theorem 10.

Proof [Proof of (32) and (33)] Recall the definition of u_i, \hat{u}_i in (30). For any $r_0 \times r_0$ matrix M , we have

$$\|M\| \leq \|M\|_F \leq r_0 \max_{ij} |M_{ij}|. \quad (54)$$

The first statement in (32) follows from (54) and the entrywise bounds in (50) by taking $c = \varepsilon$ and $t = \ell$. The second statement in (32) follows in the same way due to (51). For any $i, j \in [r_0]$,

$$\langle u_i, \hat{u}_j \rangle = \frac{\langle B^\ell \chi_i, (B^*)^\ell \tilde{\chi}_j \rangle}{\nu_i^{2\ell} \nu_j^{2\ell+2}} = \frac{\langle B^{2\ell} \chi_i, \tilde{\chi}_j \rangle}{\nu_i^{2\ell} \nu_j^{2\ell+2}}.$$

Then the first statement of (33) follows from Lemma 22 by taking $c = \varepsilon$ and $t = 2\ell$. Similarly,

$$\langle \hat{u}_i, B^\ell u_j \rangle = \frac{\langle B^{3\ell} \chi_j, \tilde{\chi}_i \rangle}{\nu_i^{2\ell+2} \nu_j^{2\ell}}.$$

The second statement of (33) is proved by taking $t = 3\ell$ in Lemma 22. This completes the proof. ■

Finally, we show a result on a pseudo-eigenvector property of χ_i that shall be useful for the next step of the proof.

Lemma 24 *Let $h = c \log_d(n)$. With probability at least $1 - c' d^4 n^{4c-1}$, for any $t \leq h$,*

$$\|B^{t+1} \chi_i - \nu_i^2 B^t \chi_i\|^2 \leq K^2 d^2 \theta_1^{4t} + c' d^{12} \kappa^2 \theta_2^{4t} \log(n)^{7/2} \frac{d^{6t}}{\sqrt{n} \wedge \alpha^{1/4}}.$$

Proof Recalling the definition of f in (53), we define

$$f(g, o) = \sum_{e: e_1=o} (\mathbf{f}_{\phi_i, t+1}(g, e) - \nu_i^2 \mathbf{f}_{\phi_i, t}(g, e))^2$$

Then f is $(2t + 4)$ -local, and satisfies

$$|f(g, o)| \leq 4\kappa^2 \theta_2^{4t} |(g, o)_{2t+4}|^2.$$

On the other hand, by the same arguments as above, we have

$$\bar{f} = \frac{d^2 \mathbf{1} \mathbf{1}^*}{n} \frac{\Phi^t}{d^{2t}} (\phi_i \circ \phi_i),$$

and using (41)

$$\left| \sum_{x \in [n]} \bar{f}(x) \right| \leq K^2 d^2 \theta_1^{4t}.$$

Proposition 20 with a triangular inequality completes the proof. \blacksquare

Appendix F. Matrix expansion and norm bounds

We now move on to show (34). Let $\mathbf{E}_2(V)$ be the oriented wedge set on a complete bipartite graph with vertex sets V_1, V_2 . We can extend the definition of B to be an operator on $\mathbf{E}_2(V)$ such that for any $e = (e_1, e_2, e_3), f = (f_1, f_2, f_3) \in \mathbf{E}_2(V)$,

$$B_{ef} = \begin{cases} A_{f_1 f_2} A_{f_3 f_2} \mathbf{1}\{e \rightarrow f\} & \text{if } e, f \in \mathbf{E}_2, \\ 0 & \text{otherwise.} \end{cases} \quad (55)$$

For $k \geq 0$, we can define $\Gamma_{e,f}^{2k+2}$ to be the set of non-backtracking walks of length $(2k + 2)$ from e to f denoted by $(\gamma_0, \dots, \gamma_{2k+2})$. We then have

$$B_{ef}^k = \sum_{\gamma \in \Gamma_{ef}^{2k+2}} X_{e_1 e_2} X_{e_3 e_2} \prod_{s=1}^k A_{\gamma_{2s} \gamma_{2s+2}, \gamma_{2s+1}}.$$

A bipartite graph spanned by γ is given by all vertices and edges from γ . We say γ is a *tangle-free* path if the bipartite graph G spanned by γ contains at most one cycle. Otherwise, we call γ a tangled path. A bipartite graph G is called ℓ -tangle-free if for any $x \in V_1$, there is at most one cycle in the ℓ -neighborhood of x in G denoted by $(G, x)_\ell$.

For $k \geq 0$, let $F_{ef}^{2k+2} \subseteq \Gamma_{ef}^{2k+2}$ be the subset of all tangle-free paths of length $(2k + 2)$. If the bipartite graph G corresponding to the biadjacency matrix X is $(2\ell + 2)$ -tangle free, then for all $1 \leq k \leq \ell$, we must have $B^k = B^{(k)}$, with

$$B_{ef}^{(k)} = \sum_{\gamma \in F_{ef}^{2k+2}} X_{e_1 e_2} X_{e_3 e_2} \prod_{s=1}^k A_{\gamma_{2s} \gamma_{2s+2}, \gamma_{2s+1}}.$$

We introduce short-hand notations:

$$A_{e_1e_3,e_2} = A_{e_1e_2}A_{e_3e_2}, \quad X_{e_1e_3,e_2} = X_{e_1e_2}X_{e_3e_2}, \quad M_{e_1e_3,e_2} = M_{e_1e_2}M_{e_3e_2}.$$

And similarly,

$$Q_{e_1e_3,e_2} = Q_{e_1e_2}Q_{e_3e_2} = mnM_{e_1e_3,e_2}^2. \quad (56)$$

Define the corresponding centered random variables

$$\underline{A}_{e_1e_3,e_2} = A_{e_1e_2}A_{e_3e_2} - M_{e_1e_2}M_{e_3e_2}, \quad \underline{X}_{e_1e_3,e_2} = X_{e_1e_2}X_{e_3e_2} - \frac{d^2}{mn}.$$

We then define $\underline{B}^{(k)}$, a centered version of $B^{(k)}$, as

$$\underline{B}_{ef}^{(k)} = \sum_{\gamma \in F_{ef}^{2k+2}} \underline{X}_{e_1e_3,e_2} \prod_{s=1}^k \underline{A}_{\gamma_{2s}\gamma_{2s+2},\gamma_{2s+1}}.$$

We also define

$$B_{ef}^{(0)} = \mathbf{1}\{e = f\}X_{e_1e_3,e_2}, \quad \underline{B}_{ef}^{(0)} = \mathbf{1}\{e = f\}\underline{X}_{e_1e_3,e_2}.$$

The following telescoping sum formula holds for any real numbers $a_s, b_s, 0 \leq s \leq \ell$:

$$\prod_{s=0}^{\ell} a_s = \prod_{s=0}^{\ell} b_s + \sum_{t=0}^{\ell} \prod_{s=0}^{t-1} b_s (a_t - b_t) \prod_{s=t+1}^{\ell} a_s.$$

Separating the cases $t = 0$ in the sum, we can decompose $B^{(\ell)}$ as

$$\begin{aligned} B_{ef}^{(\ell)} &= \underline{B}_{ef}^{(\ell)} + \frac{d^2}{mn} \sum_{\gamma \in F_{ef}^{2\ell+2}} \prod_{s=1}^{\ell} A_{\gamma_{2s}\gamma_{2s+2},\gamma_{2s+1}} \\ &\quad + \sum_{t=1}^{\ell} \sum_{\gamma \in F_{ef}^{2\ell+2}} \underline{X}_{e_1e_3,e_2} \prod_{s=1}^{t-1} \underline{A}_{\gamma_{2s}\gamma_{2s+2},\gamma_{2s+1}} M_{\gamma_{2t}\gamma_{2t+2},\gamma_{2t+1}} \prod_{s=t+1}^{\ell} A_{\gamma_{2s}\gamma_{2s+2},\gamma_{2s+1}}. \end{aligned}$$

For $1 \leq t \leq \ell - 1$, define $F_{2t,ef}^{2\ell+2} \subset \Gamma_{ef}^{2\ell+2}$ the set of non-backtracking *tangled* paths $\gamma = (\gamma_0, \dots, \gamma_{2\ell+2})$ such that $(\gamma_0, \dots, \gamma_{2t}) \in F_{e,g}^{2t}$, $(\gamma_{2t+2}, \dots, \gamma_{2\ell+2}) \in F_{g',f}^{2\ell-2t}$ for some edges $g, g' \in \mathbf{E}_2(V)$. For $t = 0$, $F_{0,ef}^{2\ell+2} \subset \Gamma_{ef}^{2\ell+2}$ is the set of non-backtracking tangled paths $\gamma = (\gamma_0, \dots, \gamma_{2\ell+2})$ such that $(\gamma_0, \gamma_1) = (e_1, e_2)$, and $(\gamma_2, \dots, \gamma_{2\ell+2}) \in F_{gf}^{2\ell}$ for some $g \in \mathbf{E}_2(V)$. Necessarily, $g_1 = e_3$. Similarly, $F_{\ell,ef}^{2\ell+2} \subset \Gamma_{ef}^{2\ell+2}$ is the set of non-backtracking tangled paths $\gamma = (\gamma_0, \dots, \gamma_{2\ell+2})$ such that $(\gamma_0, \dots, \gamma_{2\ell}) \in F_{e,g}^{2\ell}$, for some $g \in \mathbf{E}_2$, and $(\gamma_{2\ell+1}, \gamma_{2\ell+2}) = (f_2, f_3)$. Necessarily, $g_3 = f_1$.

Next, we introduce three matrices $H, H^{(1)}$ and $H^{(2)}$ such that for $e, f \in \mathbf{E}_2(V)$,

$$H_{ef} = \frac{d^2}{mn} \mathbf{1}\{e \rightarrow f\} \quad (57)$$

$$H_{ef}^{(1)} = \mathbf{1}\{e \rightarrow f\} M_{f_1 f_3, f_2} \quad (58)$$

$$H_{ef}^{(2)} = \sum_{e \rightarrow g \rightarrow f} M_{e_3 f_1, g_2} A_{f_1 f_3, f_2}, \quad (59)$$

where the sum is over all $g \in \mathbf{E}_2$ such that $e \rightarrow g \rightarrow f$ is a non-backtracking walk of length 6. Then the following decomposition holds:

$$B^{(\ell)} = \underline{B}^{(\ell)} + HB^{(\ell-1)} + \sum_{t=1}^{\ell-1} \underline{B}^{(t-1)} H^{(2)} B^{(\ell-t-1)} + \underline{B}^{(\ell-1)} H^{(1)} - \sum_{t=0}^{\ell} R_t^{(\ell)},$$

where for $1 \leq t \leq \ell$,

$$(R_t^{(\ell)})_{ef} = \sum_{\gamma \in F_{2t, ef}^{2\ell+2}} X_{e_1 e_3, e_2} \prod_{s=1}^{t-1} A_{\gamma_{2s} \gamma_{2s+2}, \gamma_{2s+1}} M_{\gamma_{2t} \gamma_{2t+2}, \gamma_{2t+1}} \prod_{s=t+1}^{\ell} A_{\gamma_{2s} \gamma_{2s+2}, \gamma_{2s+1}}$$

$$(R_0^{(\ell)})_{ef} = \frac{d^2}{mn} \sum_{\gamma \in F_{0, ef}^{2\ell+2}} \prod_{s=1}^{\ell} A_{\gamma_{2s} \gamma_{2s+2}, \gamma_{2s+1}}.$$

From (23), the following holds:

$$(TPS_{\Delta})_{ef} = \sum_{i,j} T_{ei} P_{ij} S_{jf} \Delta_{ff} = P_{e_3 f_1} \Delta_{ff} = \sum_{u \in V_2} M_{e_3 u} M_{f_1 u} A_{f_1 f_2} A_{f_3 f_2}.$$

And from the singular value decomposition of P in (28), and the matrix relations from (25),

$$TPS_{\Delta} = \sum_{i=1}^n \nu_i^2 T \phi_i \phi_i^* S_{\Delta} = \sum_{i=1}^n \nu_i^2 \chi_i \check{\chi}_i^*. \quad (60)$$

Note that $H^{(2)}$ is close to TPS_{Δ}^* . We define \tilde{H} such that

$$H^{(2)} = \sum_{i=1}^r \nu_i^2 \chi_i \check{\chi}_i^* + \tilde{H}.$$

Therefore, the following decomposition holds:

$$B^{(\ell)} = \underline{B}^{(\ell)} + HB^{(\ell-1)} + \sum_{t=1}^{\ell-1} \sum_{k=1}^r \nu_k^2 \underline{B}^{(t-1)} \chi_k \check{\chi}_k^* B^{(\ell-t-1)}$$

$$+ \sum_{t=1}^{\ell-1} \underline{B}^{(t-1)} \tilde{H} B^{(\ell-t-1)} + \underline{B}^{(\ell-1)} H^{(1)} - \sum_{t=0}^{\ell} R_t^{(\ell)}. \quad (61)$$

From (58), deterministically,

$$\|H^{(1)}\| \leq \sqrt{\|H^{(1)}\|_1 \|H^{(1)}\|_{\infty}} \leq nm \|M\|_{\infty}^2 = L^2.$$

With the decomposition in (61), the following estimate thus holds:

Lemma 25 For any unit vector $x \in \mathbb{R}^{\mathbf{E}_2(V)}$:

$$\begin{aligned} \|B^{(\ell)}x\| &\leq \|\underline{B}^{(\ell)}\| + \|HB^{(\ell-1)}\| + \sum_{t=1}^{\ell-1} \sum_{i=1}^r \nu_i^2 \|\underline{B}^{(t-1)}\chi_i\| \left| \langle \check{\chi}_i, B^{(\ell-t-1)}x \rangle \right| \\ &\quad + \sum_{t=1}^{\ell-1} \|\underline{B}^{(t-1)}\tilde{H}B^{(\ell-t-1)}\| + L^2 \|\underline{B}^{(\ell-1)}\| + \sum_{t=0}^{\ell} \|R_t^{(\ell)}\|. \end{aligned}$$

The bulk estimates of (34) stem from Lemma 25 and the following proposition:

Proposition 26 Let χ be any vector among $\chi_1, \dots, \chi_r \in \mathbb{C}^{\mathbf{E}_2(V)}$. Let $\ell \leq \lfloor \log_d(n) \rfloor$. There exists a constant $C_1, c > 0$ such that for $n \geq C_1 K^{12}$, with probability at least $1 - cn^{-1/4}$, the following norm bounds hold for all $0 \leq k \leq \ell$:

$$\|\underline{B}^{(k)}\| \leq K^{11} d^{5/2} \log(n)^{10} \theta^{2k}, \quad (62)$$

$$\|R_k^{(\ell)}\| \leq K^{20} d^{3/2} \log(n)^{28} \frac{L^{2\ell}}{\sqrt{mn}}, \quad (63)$$

$$\|HB^{(k)}\| \leq d^5 \log(n)^8 \frac{L^{2k-2}}{mn}, \quad (64)$$

$$\|\underline{B}^{(k)}\chi\| \leq \kappa K^{11} d^{5/2} \log(n)^{12} \theta^{2k}. \quad (65)$$

And for all $1 \leq k \leq \ell - 1$,

$$\|\underline{B}^{(k-1)}\tilde{H}B^{(\ell-k-1)}\| \leq K^{11} \kappa^2 d^3 \log(n)^{18} \frac{\theta^{2(k-1)} L^{2(\ell-k)}}{mn}. \quad (66)$$

Before we prove Proposition 26 in the next section, we show how it implies eq. (34).

Proof [Proof of (34)] With the parity-time invariance from (27),

$$\langle \check{\chi}_i, B^t w \rangle = \langle B^t \chi_i, J_{\Delta} w \rangle.$$

and we have $\|J_{\Delta} w\| \leq \frac{L^2}{d^2}$. Since w is orthogonal to all $B^\ell \chi_i, i \in [r_0]$,

$$\begin{aligned} |\nu_i^{-2t} \langle \check{\chi}_i, B^t w \rangle| &= |\nu_i^{-2t} \langle \check{\chi}_i, B^t w \rangle - \nu_i^{-2\ell} \langle \check{\chi}_i, B^\ell w \rangle| \\ &= \left| \sum_{s=t}^{\ell-1} \nu_i^{-2s} \langle \check{\chi}_i, B^s w \rangle - \nu_i^{-2(s+1)} \langle \check{\chi}_i, B^{s+1} w \rangle \right| \\ &\leq \frac{L^2}{d^2} \sum_{s=t}^{\ell-1} \nu_i^{-2s-2} \|B^{s+1} \chi_i - \nu_i^2 B^s \chi_i\|. \end{aligned}$$

With Lemma 24,

$$\|B^{s+1} \chi_i - \nu_i^2 B^s \chi_i\| \leq K d \theta^{2s} + c \kappa d^6 \theta^{2s} \log(n)^{7/4} \frac{d^{3s}}{n^{1/4} \wedge \alpha^{1/8}}.$$

Therefore since $\nu_i \geq \theta$, for $i \in [r_0]$,

$$|\langle \check{\chi}_i, B^t w \rangle| \leq K d \nu_i^{2t} \sum_{s=t}^{\ell-1} \left(\frac{\theta}{\nu_i} \right)^{2s+2} + \frac{c \kappa d^6 \log^{7/4}(n)}{n^{1/4} \wedge \alpha^{1/8}} \nu_i^{2t} \sum_{s=t}^{\ell-1} \left(\frac{\theta}{\nu_i} \right)^{2s+2} d^{3s} \quad (67)$$

$$\leq \theta^{2t} \left(K d \log n + \frac{c \kappa d^6 \log^{9/4}(n) d^{3\ell}}{n^{1/4} \wedge \alpha^{1/8}} \right). \quad (68)$$

On the other hand, for $i > r_0$, $\nu_i \leq \theta$, from (50),

$$|\langle \check{\chi}_i, B^t w \rangle| \leq \frac{L^2}{d^2} \|B^t \chi_i\| \leq L \nu_i^{2t} \sqrt{\Gamma_{ii}^{(t)}} + c d^2 L \kappa \theta_2^{2t} \log(n)^{7/4} \frac{d^{3t}}{n^{1/4} \wedge \alpha^{1/8}}.$$

From (31) and (41),

$$\Gamma_{ii}^{(t)} \leq \sum_{s=0}^t \frac{K^2 \rho^{2s}}{(\nu_i^2 d)^{2s}} \leq K^2 \log(n) \theta_1^{4t} \nu_i^{-4t},$$

which implies for $i > r_0$,

$$|\langle \check{\chi}_i, B^t w \rangle| \leq \theta^{2t} \left(L K \log^{1/2}(n) + c d^2 L \kappa \log(n)^{7/4} \frac{d^{3t}}{n^{1/4} \wedge \alpha^{1/8}} \right). \quad (69)$$

Let w be any unit vector orthogonal to all $\hat{u}_i, i \in [r_0]$ defined in (30). By the tangle-free property in Proposition 17 and Lemma 25, with probability at least $1 - cn^{-1/4}$, we can apply the bounds in Proposition 26 as well as the two estimates (67) and (69) to conclude

$$\|B^\ell w\| \leq c r \kappa^2 d^6 K^{20} \log^{14}(n) \theta^{2\ell},$$

which is the first claim of (34). The proof of the second claim in (34) follows the same argument by considering the transpose of the decomposition in (61), as in the proof of (Stephan and Zhu, 2022, Lemma 26). \blacksquare

Appendix G. The trace method: proof of (62)

We now show Proposition 26. We shall only show eq. (62) in the main text; other inequality are dealt with using similar methods, and relegated to the appendix. For any integer $s \geq 0$,

$$\begin{aligned} \|\underline{B}^{(k)}\|^{2s} &\leq \text{tr} \left(\underline{B}^{(k)} \underline{B}^{(k)*} \right)^s = \sum_{(e_1, \dots, e_{2s})} \prod_{i=1}^s \underline{B}_{e_{2i-1}, e_{2i}}^{(k)} \underline{B}_{e_{2i+1}, e_{2i}}^{(k)} \\ &= \sum_{\gamma \in W_{k,s}} \prod_{i=1}^{2s} \underline{X}_{\gamma_{i,0} \gamma_{i,2} \gamma_{i,1}} \prod_{t=1}^k \underline{A}_{\gamma_{i,2t} \gamma_{i,2t+2} \gamma_{i,2t+1}}, \end{aligned} \quad (70)$$

where $W_{k,s}$ is the set of sequences of path $(\gamma_1, \dots, \gamma_{2s})$ such that each $\gamma_i = (\gamma_{i,0}, \dots, \gamma_{i,2k+2})$ is a non-backtracking tangle-free path of length $2k+2$, with boundary conditions that for all $i \in [s]$,

$$\begin{aligned} (\gamma_{2i,2k}, \gamma_{2i,2k+1}, \gamma_{2i,2k+2}) &= (\gamma_{2i-1,2k}, \gamma_{2i-1,2k+1}, \gamma_{2i-1,2k+2}), \\ (\gamma_{2i+1,0}, \gamma_{2i+1,1}, \gamma_{2i+1,2}) &= (\gamma_{2i,0}, \gamma_{2i,1}, \gamma_{2i,2}) \end{aligned} \quad (71)$$

with the convention that $\gamma_{2s+1} = \gamma_1$. For each $\gamma \in W_{k,s}$, we associate with an undirected bipartite graph $G_\gamma = (V_\gamma^1, V_\gamma^2, E_\gamma)$ of visited vertices and edges. From the boundary condition, G_γ is connected hence

$$|E_\gamma| - |V_\gamma| + 1 \geq 0.$$

Let $a(\gamma) = |E_\gamma|, v(\gamma) = |V_\gamma|$. We drop the dependence on γ for ease of notation.

Now we consider the expectation of (70). It also depends on the degree profile of vertices in V_γ^2 . There are no vertices of degree 1 in V_γ^2 by our construction of G_γ . Let T_γ be the set of distinct triplets visited by γ and let m_e be the number of times a triplet e is visited. If a triplet has an overlapped edge with another triplet, we call it a *bad triplet*; otherwise, we call it a *good triplet*. In addition, we call a vertex in V_γ^2 *bad* if it is a middle point of a certain bad triplet.

Let k_e be the number of times that a triplet e is visited in γ , which carries the weight \underline{X}_e and denote $\tilde{m}_e = m_e - k_e$. Let $\bar{W}_{k,s}$ be the subset of $W_{k,s}$ such that each good triplet was visited at least twice in γ . We can now write the expectation of (70) as

$$\begin{aligned} & \sum_{\gamma \in \bar{W}_{k,s}} \mathbb{E} \prod_{e \in T_\gamma} \underline{X}_e^{k_e} \underline{A}_e^{\tilde{m}_e} \\ &= \sum_{\gamma \in \bar{W}_{k,s}} \mathbb{E} \prod_{e \in T_\gamma, \text{good}} \underline{X}_e^{k_e} \underline{A}_e^{\tilde{m}_e} \prod_{e \in T_\gamma, \text{bad}} \underline{X}_e^{k_e} \underline{A}_e^{\tilde{m}_e} \\ &\leq \sum_{\gamma \in \bar{W}_{k,s}} \prod_{e \in T_\gamma} \left(\frac{|M_e|}{p^2} \right)^{\tilde{m}_e} \cdot \left(\prod_{e \in T_\gamma, \text{good}} p^2 \right) \cdot \left| \mathbb{E} \prod_{e \in T_\gamma, \text{bad}} \underline{X}_e^{m_e} \right| \\ &= \sum_{\gamma \in \bar{W}_{k,s}} p^{-4ks} \prod_{e \in T_\gamma} |M_e|^{\tilde{m}_e} \cdot \left(\prod_{e \in T_\gamma, \text{good}} p^2 \right) \cdot \left| \mathbb{E} \prod_{e \in T_\gamma, \text{bad}} \underline{X}_e^{m_e} \right|, \end{aligned} \tag{72}$$

where in the last identity, we use that

$$\sum_{e \in T_\gamma} \tilde{m}_e = 2sk.$$

Notice that when $p \in [0, 1/2]$, deterministically, the following inequality holds:

$$|X_{e_1} X_{e_2} - p^2| \leq (1 + 4p) |(X_{e_1} - p)(X_{e_2} - p)|. \tag{73}$$

Let $T_\gamma(v)$ be the collection of bad triplets with the same middle point v . Let $d(v)$ be the number of distinct edges connected to v from bad triplets. For each $e \in E_\gamma$, with abuse of notation, let m_e be the number of visits of γ to the edge e . Since contributions of triplets with distinct middle points are independent, with (73), we can consider each contribution from $T_\gamma(v)$ as follows

$$\begin{aligned} \left| \mathbb{E} \prod_{e \in T_\gamma(v)} \underline{X}_e^{m_e} \right| &= \left| \mathbb{E} \prod_{e \in T_\gamma(v)} (X_{e_1 e_2} X_{e_3 e_2} - p^2)^{m_e} \right| \\ &\leq (1 + 4p)^{\sum_{e \in T_\gamma(v)} m_e} \prod_{e \in E_\gamma(v)} \mathbb{E} |X_e - p|^{m_e} \\ &\leq (1 + 4p)^{\sum_{e \in T_\gamma(v)} m_e} 2^{d(v)} p^{d(v)}. \end{aligned} \tag{74}$$

Notice that G_γ is a bipartite graph, and the imbalance between $|V_\gamma^1|$ and $|V_\gamma^2|$ is bounded by the genus plus the extra imbalance from the endpoints in each γ_i , so we have

$$\left| |V_\gamma^1| - |V_\gamma^2| \right| \leq a - v + 1 + 2s. \quad (75)$$

Let $v_i, v_{\geq i}$ be the number of vertices with degree i (at least i , respectively) in V_γ^2 . Then since there are no vertices of degree 1 in V_γ^2 ,

$$v_2 + v_{\geq 3} = |V_\gamma^2|, \quad \text{and} \quad 2v_2 + \sum_{i \geq 3} iv_i \leq a.$$

Therefore

$$\sum_{v \in V_\gamma^2, \text{bad}} d(v) \leq \sum_{i \geq 3} iv_i \leq \sum_{i \geq 3} 3(i-2)v_i \leq 3(a - 2|V_\gamma^2|) \leq 6(a - v + s + 1). \quad (76)$$

where the last inequality is from (75). Similarly,

$$v_{\geq 3} \leq a - 2|V_\gamma^2| \leq 2(a - v + 2s). \quad (77)$$

Since each γ_i is tangle-free, each bad vertex can be visited at most twice for each γ_i . Therefore each bad vertex can be visited at most $4s$ times by γ , which implies

$$\sum_{e \in T_\gamma, \text{bad}} m_e \leq 4s \cdot v_{\geq 3} \leq 8s(a - v + 2s). \quad (78)$$

From (74), we obtain

$$\left| \mathbb{E} \prod_{e \in T_\gamma, \text{bad}} \underline{X}_e^{m_e} \right| \leq (1+4p)^{8s(a-v+2s)} 2^{6(a-v+s+1)} p^{\sum_{v \in V_\gamma^2} d(v)}. \quad (79)$$

Since

$$\sum_{v \in V_\gamma^2, \text{good}} \deg(v) + \sum_{v \in V_\gamma^2} d(v) = a,$$

with (79), we can simplify (72) to be

$$\sum_{\gamma \in \overline{W}_{k,s}} \mathbb{E} \left[\prod_{e \in T_\gamma} \underline{X}_e^{k_e} \underline{A}_e^{\tilde{m}_e} \right] \leq \sum_{\gamma \in \overline{W}_{k,s}} (1+4p)^{8s(a-v+2s)} 2^{6(a-v+s+1)} p^{a-4ks} \prod_{e \in E_\gamma} |M_e|^{\tilde{m}_e}, \quad (80)$$

where for each $e \in E_\gamma$, with abuse of notation, \tilde{m}_e is defined as the number of visits of γ to the edge e excluding $(\gamma_{i,0}, \gamma_{i,1}), (\gamma_{i,1}, \gamma_{i,2}), i \in [2s]$.

Define $\mathcal{W}_{k,s}(v, a)$ to be the set of equivalence classes of paths in $\overline{W}_{k,s}$ with v distinct vertices and s distinct edges. The following lemma holds.

Lemma 27 *Let v, e be integers such that $a - v + 1 \geq 0$. Then*

$$|\mathcal{W}_{k,s}(v, a)| \leq (4(k+1)s)^{6s(a-v+1)+2s}.$$

Proof [Proof of Lemma 27] We can use the estimation from (Bordenave et al., 2018, Lemma 17) and notice that each graph is spanned by $2s$ concatenations of non-backtracking walks of length $2k + 2$. \blacksquare

We now bound the contribution of paths in each equivalence class.

Lemma 28 *Let $\gamma \in \overline{W}_{k,s}$ such that $|V_\gamma| = v, |E_\gamma| = a$. We have*

$$\sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\gamma'}} |M_e|^{\tilde{m}_e} \leq K^{6(a-v)+12s+2sk} (\sqrt{mn})^{-4sk+v} \rho^{2sk} \left(\sqrt{\frac{m}{n}} \right)^{a-v}.$$

Proof [Proof of Lemma 28] Let $\tilde{\gamma} = (\tilde{\gamma}_1, \dots, \tilde{\gamma}_{2s})$ be the sequence of paths from $\gamma \in \overline{W}_{k,s}$ such that $\tilde{\gamma}_i = (\gamma_{i,2}, \dots, \gamma_{i,2k+2})$, i.e., we remove the initial two steps (the first triplet) in each $\gamma_i, 1 \leq i \leq 2s$. From the boundary condition, the associated graph spanned by $\tilde{\gamma}$, denoted by $G_{\tilde{\gamma}}$, is connected and

$$|E_{\tilde{\gamma}}| = a, \quad |V_{\tilde{\gamma}}| = v.$$

Recall \tilde{m}_e is the number of visits of the edge $e \in E_{\tilde{\gamma}}$, we have

$$\sum_{e \in E_{\tilde{\gamma}}} \tilde{m}_e = 4ks.$$

Let H be the set of edges $e \in E_\gamma$ such that $\tilde{m}_e = 1$ and denote $h = |H|$. Note that by our construction, any edge in H is either one of the two initial or ending edges in $\tilde{\gamma}_i, i \in [2s]$, or from a bad triplet (note that it's impossible to have both edges in a triplet belong to H , then the contribution is zero by independence). Therefore with (76),

$$h \leq 4s + \sum_{v \in V_\gamma^2, \text{bad}} d(v) \leq 10s + 6(a - v + 1). \quad (81)$$

We also have

$$\sum_{e \in E_{\tilde{\gamma}} \setminus H} \tilde{m}_e = 4ks - h,$$

which implies

$$\begin{aligned} \sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\tilde{\gamma}'}} |M_e|^{\tilde{m}_e} &\leq \left(\sqrt{\frac{K\rho}{mn}} \right)^{4sk-2a+2h} \sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\tilde{\gamma}'} \setminus H} \frac{Q_e}{\sqrt{mn}}, \\ &= (K\rho)^{2sk+h-a} (\sqrt{mn})^{a-h-4sk} \sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\tilde{\gamma}'} \setminus H} Q_e, \end{aligned} \quad (82)$$

where Q_e is defined in (5), and we use (6) to get $|M_e| \leq \sqrt{\frac{K\rho}{mn}}$. Next, we consider an upper bound on

$$\sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\tilde{\gamma}'} \setminus H} Q_e.$$

Let $G'_{\tilde{\gamma}}$ be the graph spanned by all edges in $E_{\tilde{\gamma}} \setminus H$. Let $t_i, t_{\geq i}$ be the number of vertices in $V'_{\tilde{\gamma}}$ with a degree equal to i and at least i , respectively. We then have

$$\begin{aligned} v' &= t_1 + t_2 + t_{\geq 3} \geq v - h, \\ t_1 + 2t_2 + 3t_{\geq 3} &\leq \sum_{k \geq 1} kt_k = 2(a - h). \end{aligned} \quad (83)$$

The first inequality is because removing any edge in H can delete at most one vertex from $G_{\tilde{\gamma}}$. Since degree-1 vertices can only appear at the endpoints of each $\tilde{\gamma}_i, i \leq 2s$, which are $\gamma_{i,2}, \gamma_{i,2k+2}, i \in [2s]$, we must have $t_1 \leq 2s$. Together with (83), we obtain that

$$t_{\geq 3} \leq 2(a - v) + t_1 \leq 2(a - v) + 2s. \quad (84)$$

We reduce the graph $G'_{\tilde{\gamma}}$ to a multi-graph $\hat{G}_{\tilde{\gamma}}$ by gluing vertices of degree 2 while keeping the endpoints $\gamma_{i,2}, \gamma_{i,2k+2}, i \in [2s]$. Let \hat{v} be the number of vertices in $\hat{G}_{\tilde{\gamma}}$. Next, we define the edge set $\hat{E}_{\tilde{\gamma}}$ for $\hat{G}_{\tilde{\gamma}}$. We partition edges in $E'_{\tilde{\gamma}}$ into \hat{a} sequences of edges $\hat{e}_j = (e_{j_1}, \dots, e_{j_{q_j}})$ where $e_{j,t} = (x_{j,t-1}, x_{j,t}) \in E'_{\tilde{\gamma}}, x_{j,0}, x_{j,q_j} \in \hat{V}_{\tilde{\gamma}}$ and $x_{j_t} \notin \hat{V}_{\tilde{\gamma}}$ for $1 \leq t \leq q_j - 1$, and let $(x_{j,0}, x_{j,q_j})$ be an edge in $\hat{E}_{\tilde{\gamma}}$. This implies

$$\sum_{j=1}^{\hat{a}} q_j = a - h.$$

From the definition of $\hat{G}_{\tilde{\gamma}}$, the genus of the graph is preserved,

$$\hat{a} - \hat{v} = a' - v' = a - h - v' \geq a - v - h. \quad (85)$$

and since $a' = a - h, v' \geq v - h$,

$$\hat{a} - \hat{v} = a' - v' \leq a - v. \quad (86)$$

Moreover, since the number of degree-2 vertices in $\hat{G}_{\tilde{\gamma}}$ is bounded by the number of distinct endpoints in $\tilde{\gamma}_i, i \in [2s]$, we have from (84),

$$\hat{v} \leq t_1 + t_{\geq 3} + 2s \leq 2(a - v) + 6s. \quad (87)$$

On the other hand, removing edges in H does not increase the genus, so $a' - v' \leq a - v$. And with (87), it implies

$$\hat{a} = \hat{a} - \hat{v} + \hat{v} = a' - v' + \hat{v} \leq a - v + \hat{v} \leq 3(a - v) + 6s. \quad (88)$$

We classify the edges in $\hat{E}_{\tilde{\gamma}}$ according to endpoints $(x_{j,0}, x_{j,q_j})$ into 3 different groups \hat{E}_1, \hat{E}_2 , and \hat{E}_3 : (1) both are in $V_{\tilde{\gamma}}^1$, (2) both are in $V_{\tilde{\gamma}}^2$, (3) one from $V_{\tilde{\gamma}}^1$ and one from $V_{\tilde{\gamma}}^2$.

Let $y_1, \dots, y_{\hat{v}}$ be the elements in \hat{V}_γ , where the first \hat{v}_1 vertices are from \hat{V}_γ^1 , and let a_j, b_j be the indices such that $x_{j,0} = y_{a_j}, x_{j,q_j} = y_{b_j}$. Then

$$\begin{aligned}
 \sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\hat{\gamma}}'} Q_e &\leq \sum_{(y_1, \dots, y_{\hat{v}}) \in [n]^{\hat{v}_1} \times [m]^{\hat{v}_2}} \prod_{\hat{e}_j \in \hat{E}_1} (QQ^*)_{y_{a_j} y_{b_j}}^{q_j/2} \prod_{\hat{e}_j \in \hat{E}_2} (Q^*Q)_{y_{a_j} y_{b_j}}^{q_j/2} \prod_{\hat{e}_j \in \hat{E}_3} \frac{K^2 \rho^{q_j}}{\sqrt{nm}} \\
 &\leq \sum_{(y_1, \dots, y_{\hat{v}}) \in [n]^{\hat{v}_1} \times [m]^{\hat{v}_2}} \prod_{\hat{e}_j \in \hat{E}_1} \frac{K^2 \rho^{q_j}}{n} \prod_{\hat{e}_j \in \hat{E}_2} \frac{K^2 \rho^{q_j}}{m} \prod_{\hat{e}_j \in \hat{E}_3} \frac{K^2 \rho^{q_j}}{\sqrt{nm}} \\
 &= K^{2\hat{a}} \rho^{a-h} n^{\hat{v}_1 - |\hat{E}_1| - \frac{1}{2}|\hat{E}_3|} m^{\hat{v}_2 - |\hat{E}_2| - \frac{1}{2}|\hat{E}_3|} \\
 &= (K^2)^{\hat{a}} \rho^{a-h} (\sqrt{mn})^{\hat{v} - \hat{a}} \left(\sqrt{\frac{m}{n}} \right)^{\hat{v}_2 - |\hat{E}_2| - \hat{v}_1 + |\hat{E}_1|} \\
 &\leq (K^2)^{3(a-v) + 6s} \rho^{a-h} (\sqrt{mn})^{v+h-a} \left(\sqrt{\frac{m}{n}} \right)^{\hat{a} - \hat{v} - |\hat{E}_3| - 2(|\hat{E}_2| - \hat{v}_2)} \\
 &\leq (K^2)^{3(a-v) + 6s} \rho^{a-h} (\sqrt{mn})^{v+h-a} \left(\sqrt{\frac{m}{n}} \right)^{a-v - |\hat{E}_3| - 2|\hat{E}_2| + 2\hat{v}_2}
 \end{aligned}$$

where the first inequality is from (40), the second inequality is from (35) and (39), the third inequality is from (88) and (85), and the last one is from (86). Since all vertices in \hat{V}_γ^2 have degree at least 3,

$$-|\hat{E}_3| - 2|\hat{E}_2| + 2\hat{v}_2 = \sum_{v \in \hat{V}_\gamma^2} (2 - \deg(v)) \leq 0.$$

We get

$$\sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\hat{\gamma}} \setminus H} Q_e \leq (K^2)^{3(a-v) + 6s} \rho^{a-h} (\sqrt{mn})^{v+h-a} \left(\sqrt{\frac{m}{n}} \right)^{a-v}.$$

Therefore from (82),

$$\begin{aligned}
 \sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\hat{\gamma}}} |M_e|^{\tilde{m}_e} &\leq K^{6(a-v) + 12s + 2sk + h - a} (\sqrt{mn})^{-4sk + v} \rho^{2sk} \left(\sqrt{\frac{m}{n}} \right)^{a-v} \quad (89) \\
 &\leq K^{12(a-v+1) + 22s + 2sk - a} (\sqrt{mn})^{-4sk + v} \rho^{2sk} \left(\sqrt{\frac{m}{n}} \right)^{a-v},
 \end{aligned}$$

where the last inequality is from (81) and the fact that $K \geq 1$. This completes the proof of Lemma 28. \blacksquare

We take

$$k \leq \lfloor \log n \rfloor, \quad s = \left\lfloor \frac{\log \left(\frac{n}{8(dvK)^3 K^9} \right)}{8 \log \log n} \right\rfloor \quad (90)$$

and let $g = a - v + 1$. We have the first simple constraint on v, a as

$$0 \leq v - 1 \leq a \leq 4(k+1)s.$$

Moreover, for any edge in G_γ not included in any bad triplets must be visited at least twice to have a nonzero contribution in (70). And for any γ with a nonzero contribution in (70), the number of edges visited only once is bounded by $6(g+s)$ from (76). This implies $a \leq 2ks + 3g + 5s$.

From (70), (80) and (89),

$$\begin{aligned} \mathbb{E}\|\underline{B}^{(k)}\|^{2s} &\leq \sum_{a=1}^{4s(k+1)} \sum_{v=1}^{a+1} \mathbf{1}\{a \leq 2ks + 3g + 5s\} |\mathcal{W}_{k,s}(v, a)| \\ &\quad \cdot (1+4p)^{8s(a-v+2s)} 2^{6(a-v+s+1)} p^{a-4ks} \\ &\quad \cdot K^{12(a-v+1)+22s+2sk-a} (\sqrt{mn})^{-4sk+v} \rho^{2sk} \left(\sqrt{\frac{m}{n}}\right)^{a-v} \\ &\leq n(4(k+1)s)^{2s+2} (1+4p)^{16s^2} 2^{6s} K^{22s} \theta_1^{4sk} \\ &\quad \cdot \sum_{g=0}^{\infty} \sum_{a=1}^{(2ks+3g+5s) \wedge 4(k+1)s} \left(\frac{K}{d}\right)^{2sk-a} \left(\frac{2K^{12} (1+4p)^{8s} (4(k+1)s)^{6s}}{n}\right)^g \end{aligned}$$

We now consider two cases:

(1) When $K \geq d$, (91) can be further bounded by

$$n(4(k+1)s)^{2s+3} (1+4p)^{16s^2} 2^{6s} K^{22s} \theta_1^{4sk} \cdot \left(\frac{K}{d}\right)^{2sk} \sum_{g=0}^{\infty} \left(\frac{2K^{12} (1+4p)^{8s} (4(k+1)s)^{6s}}{n}\right)^g. \quad (91)$$

(2) When $K \leq d$, (91) is bounded by

$$n(4(k+1)s)^{2s+3} (1+4p)^{16s^2} 2^{6s} K^{22s} \theta_1^{4sk} (d/K)^{5s} \sum_{g=0}^{\infty} \left(\frac{2d^3 K^9 (1+4p)^{8s} (4(k+1)s)^{6s}}{n}\right)^g.$$

Therefore in both cases, (91) is bounded by

$$\begin{aligned} &n(4(k+1)s)^{2s+3} (1+4p)^{16s^2} 2^{6s} K^{22s} \theta_1^{4sk} \left(1 \vee \frac{d}{K}\right)^{5s} \\ &\quad \cdot \sum_{g=0}^{\infty} \left(\frac{2(d \vee K)^3 K^9 (1+4p)^{8s} (4(k+1)s)^{6s}}{n}\right)^g. \end{aligned}$$

Since $d \leq n^{1/12}$, from our choices of k and s in (90), for $n \geq C_1 K^{16}$, we have

$$(1+4p)^{8s} \leq 2, \quad \frac{(d \vee K)^3 K^9 (4(k+1)s)^{6s}}{n} \leq \frac{1}{8}, \quad n^{\frac{1}{2s}} \leq \log(n)^5.$$

which implies

$$\mathbb{E}\|\underline{B}^{(k)}\|^{2s} \leq 2n(4(k+1)s)^{2s+3} 2^{2s} K^{22s} \theta_1^{4sk} \left(1 \vee \frac{d}{K}\right)^{5s}.$$

Then by Markov's inequality and the assumption $C \geq 1$, with probability at least $1 - n^{-1/2}$, the following bound holds:

$$\|\underline{B}^{(k)}\| \leq K^{11} d^{5/2} \log(n)^{19/2} \theta^{2k}.$$

This finishes the proof for (62).

Appendix H. Additional trace methods

H.1. Proof of (63) on $R_k^{(\ell)}$

This is similar to the proof of (62). We only emphasize the main difference. For any $0 \leq k \leq \ell$,

$$\begin{aligned} \|R_k^{(\ell)}\|^{2s} &\leq \text{tr} \left[R_k^{(\ell)} R_k^{(\ell)*} \right]^s \\ &= \sum_{\gamma \in T_{\ell,s,k}} \prod_{i=1}^{2s} \underline{X}_{\gamma_{i,0}\gamma_{i,2}\gamma_{i,1}} \prod_{t=1}^{k-1} \underline{A}_{\gamma_{i,2t}\gamma_{i,2t+2},\gamma_{i,2t+1}} M_{\gamma_{i,2k}\gamma_{i,2k+2},\gamma_{i,2k+1}} \prod_{t=k+1}^{\ell} A_{\gamma_{i,2t}\gamma_{i,2t+2},\gamma_{i,2t+1}}. \end{aligned} \quad (92)$$

where $T_{\ell,s,k}$ is the set of all $(\gamma_1, \dots, \gamma_{2s})$ such that for all i ,

$$\gamma_i^1 = (\gamma_{i,0}, \gamma_{i,1}, \gamma_{i,2}, \dots, \gamma_{i,2k}) \quad \text{and} \quad \gamma_i^2 = (\gamma_{i,2k+2}, \dots, \gamma_{i,2\ell+1})$$

are non-backtracking tangle-free and γ_i is non-backtracking tangled with the same boundary condition as in (71).

For any $\gamma \in T_{\ell,s,k}$, denote $G_\gamma = (V_\gamma, E_\gamma)$ as the union of the graphs $G_{\gamma_i^j}$ for $i \in [2s]$ and $j = 1, 2$. Note that the edges $(\gamma_{i,2k}, \gamma_{i,2k+1}, \gamma_{i,2k+2})$ are not taken into account in G_γ . For any fixed $\gamma_i^1, \gamma_i^2, 1 \leq i \leq 2s$, there are at m^{2s} many ways to choose all $\gamma_{i,2k}, 1 \leq i \leq 2s$. Let $v = |V_\gamma|$ and $a = |E_\gamma|$. Since γ_i is tangled, each connected component in G_{γ_i} contains a cycle, therefore

$$v \leq a. \quad (93)$$

Define the set $\mathcal{T}_{\ell,s,k}(v, a)$ as the set of all equivalence classes of $T_{\ell,s,k}$ with given v and a . The following lemma from (Bordenave et al., 2018) holds.

Lemma 29 *Let $v \leq a$. Then $|\mathcal{T}_{\ell,s,k}(v, a)| \leq (8(\ell + 1)s)^{12s(a-v+1)+8s}$.*

Taking expectation of (92) yields

$$\mathbb{E} \|R_k^{(\ell)}\|^{2s} \leq \frac{L^{4s}}{(mn)^s} \sum_{\gamma \in \bar{T}_{\ell,s,k}} \mathbb{E} \prod_{i=1}^{2s} \underline{X}_{\gamma_{i,0}\gamma_{i,2}\gamma_{i,1}} \prod_{t=1}^{k-1} \underline{A}_{\gamma_{i,2t}\gamma_{i,2t+2},\gamma_{i,2t+1}} \prod_{t=k+1}^{\ell} A_{\gamma_{i,2t}\gamma_{i,2t+2},\gamma_{i,2t+1}}, \quad (94)$$

where $\bar{T}_{\ell,s,k}$ is the subset of $T_{\ell,s,k}$ where each $\gamma \in \bar{T}_{\ell,s,k}$ has nonzero contribution in the expectation. For each $\gamma \in \bar{T}_{\ell,s,k}$,

$$\mathbb{E} \prod_{i=1}^{2s} \underline{X}_{\gamma_{i,0}\gamma_{i,2}\gamma_{i,1}} \prod_{t=1}^{k-1} \underline{A}_{\gamma_{i,2t}\gamma_{i,2t+2},\gamma_{i,2t+1}} \prod_{t=k+1}^{\ell} A_{\gamma_{i,2t}\gamma_{i,2t+2},\gamma_{i,2t+1}} = \mathbb{E} \prod_{e \in T_\gamma} \underline{X}_e^{k_e} \underline{A}_e^{\tilde{m}_e} A_e^{t_e},$$

where k_e, \tilde{m}_e, t_e is the number of visits of γ to a triplet e that carry the weight $\underline{X}_e, \underline{A}_e$, and A_e , respectively and let $m_e = k_e + \tilde{m}_e$. Let T_γ be the set of all triplets from γ . We call a triplet *bad* if it overlaps another triplet. Separating good and bad triples yields

$$\begin{aligned} \mathbb{E} \prod_{e \in T_\gamma} \underline{X}_e^{k_e} \underline{A}_e^{\tilde{m}_e} A_e^{t_e} &= \prod_{e \in T_\gamma} \left(\frac{|M_e|}{p^2} \right)^{\tilde{m}_e + t_e} \left| \mathbb{E} \prod_{e \in T_\gamma \text{ good}} \underline{X}_e^{m_e} X_e^{t_e} \right| \cdot \left| \mathbb{E} \prod_{e \in T_\gamma \text{ bad}} \underline{X}_e^{m_e} X_e^{t_e} \right| \\ &\leq p^{-4(\ell-1)s} \prod_{e \in T_\gamma} |M_e|^{\tilde{m}_e + t_e} \left(\prod_{e \in T_\gamma \text{ good}} p^2 \right) \left| \mathbb{E} \prod_{e \in T_\gamma \text{ bad}} \underline{X}_e^{m_e} X_e^{t_e} \right|. \end{aligned} \quad (95)$$

Let $T_\gamma(v)$ be the collection of bad triplets with the same middle point v . Let $d(v)$ be the number of distinct edges connected to v from bad triplets. For each $e \in E_\gamma$, with abuse of notation, let m_e, t_e be the corresponding number of visits of γ to the edge e . Applying inequality (73), we find

$$\begin{aligned} \left| \mathbb{E} \prod_{e \in T_\gamma(v)} \underline{X}_e^{m_e} X_e^{t_e} \right| &\leq (1 + 4p)^{\sum_{e \in T_\gamma(v)} m_e} \prod_{e \in E_\gamma(v)} \mathbb{E} |X_e - p|^{m_e} X_e^{t_e} \\ &\leq (1 + 4p)^{\sum_{e \in T_\gamma(v)} m_e} 2^{d(v)} p^{d(v)}. \end{aligned} \quad (96)$$

Let $v_i, v_{\geq i}$ be the number of vertices with degree i (at least i , respectively) in V_γ^2 . Same as (76) and (77), we have

$$\begin{aligned} \sum_{v \in V_\gamma^2, \text{ bad}} d(v) &\leq 6(a - v + s + 1) \\ v_{\geq 3} &\leq 2(a - v + 2s). \end{aligned} \quad (97)$$

Since each γ_i^j is tangle-free, each bad vertex can be visited at most twice for each γ_i^j . Similar to (78), we obtain

$$\sum_{e \in T_\gamma, \text{ bad}} m_e \leq 8s \cdot v_{\geq 3} \leq 16s(a - v + 2s).$$

With the three inequalities above, (96) and (95) imply the following bound for any $\gamma \in \bar{T}_{\ell, s, k}$:

$$\mathbb{E} \prod_{e \in T_\gamma} \underline{X}_e^{k_e} \underline{A}_e^{\tilde{m}_e} A_e^{t_e} \leq (1 + 4p)^{16s(a-v+2s)} 2^{6(a-v+s+1)} p^{a-4(\ell-1)s} \prod_{e \in E_\gamma} |M_e|^{\tilde{m}_e + t_e}.$$

Next, we bound the contribution of each equivalence class.

Lemma 30 *Let $\gamma \in \bar{T}_{\ell, s, k}(v, a)$. We have*

$$\sum_{\gamma' : \gamma' \sim \gamma} \prod_{e \in E_{\gamma'}} |M_e|^{\tilde{m}_e + t_e} \leq K^{12(a-v+1)+2s(\ell-1)+40s-a} \rho^{2s(\ell-1)} (\sqrt{mn})^{v-4s(\ell-1)} \left(\sqrt{\frac{m}{n}} \right)^{a-v}. \quad (98)$$

Proof [Proof of Lemma 30] This is similar to the proof of Lemma 28. We only address the differences. Let $\tilde{\gamma} = (\tilde{\gamma}_1, \dots, \tilde{\gamma}_{2s})$ be the sequence of paths from $\gamma \in \overline{T}_{\ell, s, k}$ with the initial two steps (the first triplet) removed in each $\gamma_i, 1 \leq i \leq 2s$. From the boundary condition, the associated graph spanned by $\tilde{\gamma}$, denoted by $G_{\tilde{\gamma}}$, satisfies

$$|E_{\tilde{\gamma}}| = a, \quad |V_{\tilde{\gamma}}| = v.$$

Similarly, we have

$$\sum_{e \in E_{\tilde{\gamma}}} \tilde{m}_e + t_e = 4(\ell - 1)s.$$

Let H be the set of edges $e \in E_{\tilde{\gamma}}$ such that $\tilde{m}_e + t_e = 1$ and denote $h = |H|$. Therefore with (97),

$$h \leq 4s + \sum_{v \in V_{\tilde{\gamma}}^2, \text{bad}} d(v) \leq 10s + 6(a - v + 1). \quad (99)$$

We also have

$$\sum_{e \in E_{\tilde{\gamma}} \setminus H} \tilde{m}_e + t_e = 4(\ell - 1)s - h,$$

which implies

$$\sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\tilde{\gamma}}} |M_e|^{\tilde{m}_e + t_e} \leq (K\rho)^{2s(\ell-1)+h-a} (\sqrt{mn})^{a-h-4s(\ell-1)} \sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\tilde{\gamma}} \setminus H} Q_e, \quad (100)$$

Let $G'_{\tilde{\gamma}}$ be the graph spanned by all edges in $E_{\tilde{\gamma}} \setminus H$. Let $t_i, t_{\geq i}$ be the number of vertices in $V'_{\tilde{\gamma}}$ with a degree equal to i and at least i , respectively. We then have

$$\begin{aligned} v' &= t_1 + t_2 + t_{\geq 3} \geq v - h, \\ t_1 + 2t_2 + 3t_{\geq 3} &\leq \sum_{k \geq 1} kt_k = 2(a - h). \end{aligned} \quad (101)$$

Since degree-1 vertices can only appear at the endpoints of each $\tilde{\gamma}_i^j, i \leq 2s, j = 1, 2$, we must have

$$t_1 \leq 5s. \quad (102)$$

Together with (101), we obtain that

$$t_{\geq 3} \leq 2(a - v) + 5s. \quad (103)$$

We reduce the graph $G'_{\tilde{\gamma}}$ to a multi-graph $\hat{G}_{\tilde{\gamma}}$ by gluing vertices of degree 2 except for the endpoints of $\tilde{\gamma}_i^j, i \in [2s], j = 1, 2$. Let \hat{v} be the number of vertices in $\hat{G}_{\tilde{\gamma}}$. We partition edges in $E'_{\tilde{\gamma}}$ into \hat{a} sequences of edges as before. Similar to the proof of (62),

$$\sum_{j=1}^{\hat{a}} q_j = a - h, \quad \text{and} \quad a - v - h \leq \hat{a} - \hat{v} \leq a - v.$$

Since the number of degree-2 vertices in $\hat{G}_{\tilde{\gamma}}$ is bounded by the number of distinct endpoints in $\tilde{\gamma}_i^j, i \in [2s], j = 1, 2$, we have $t_2 \leq 5s$. And from (102) and (103),

$$\begin{aligned} \hat{v} &\leq t_1 + t_{\geq 3} + 5s \leq 2(a - v) + 15s, \\ \hat{a} &\leq 3(a - v) + 15s. \end{aligned} \quad (104)$$

Following the same steps in the proof of (62), we get

$$\sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\tilde{\gamma}} \setminus H} Q_e \leq (K^2)^{3(a-v)+15s} \rho^{a-h} (\sqrt{mn})^{v+h-a} \left(\sqrt{\frac{m}{n}} \right)^{a-v}. \quad (105)$$

Therefore from (99), (100), and (105),

$$\sum_{\gamma': \gamma' \sim \gamma} \prod_{e \in E_{\tilde{\gamma}}} |M_e|^{\tilde{m}_e + t_e} \leq K^{12(a-v+1)+2s(\ell-1)+40s-a} \rho^{2s(\ell-1)} (\sqrt{mn})^{v-4s(\ell-1)} \left(\sqrt{\frac{m}{n}} \right)^{a-v}.$$

This completes the proof of Lemma 30. ■

Now we take

$$\ell \leq \lfloor \log n \rfloor, \quad s = \left\lfloor \frac{\log \left(\frac{n}{8(d \vee K)^3 K^9} \right)}{24 \log \log n} \right\rfloor \quad (106)$$

and let $g = a - v + 1$. From (93), we have the first simple constraint on v, a as $0 \leq v \leq a \leq 4(\ell + 1)s$. Any edge in $\gamma_i^1, i \in [2s]$ not included in any bad triplets must be visited at least twice to have nonzero contribution in (94). The number of edges visited only once by $\gamma_i^1, i \in [2s]$ is bounded by $6(g + s)$ from (97). This implies

$$a \leq \frac{4ks - 6(g + s)}{2} + 6(g + s) + 2s(2\ell - 2k) = 2s(2\ell - k) + 3(g + s). \quad (107)$$

From (94), (98), (107), and Lemma 29,

$$\begin{aligned} \mathbb{E} \|R_k^{(\ell)}\|^{2s} &\leq \frac{L^{4s}}{(mn)^s} \sum_{a=1}^{4s(\ell+1)} \sum_{v=1}^a \mathbf{1}\{a \leq 2s(2\ell - k) + 3(g + s)\} (8(\ell + 1)s)^{12sg+8s} \\ &\quad \cdot (1 + 4p)^{16s(g+2s)} 2^{6(g+s)} p^{a-4(\ell-1)s} \\ &\quad \cdot K^{12g+2s(\ell-1)+40s-a} \rho^{2s(\ell-1)} (\sqrt{mn})^{v-4s(\ell-1)} \left(\sqrt{\frac{m}{n}} \right)^{a-v} \\ &\leq L^{4s} (mn)^{-s} n (8(\ell + 1)s)^{8s+1} (1 + 4p)^{32s^2} 2^{6s} K^{40s} \theta^{4s(\ell-1)} \left(1 \vee \frac{d}{K} \right)^{3s} \\ &\quad \cdot \sum_{g=0}^{\infty} \left(\frac{2^6 (d \vee K)^3 K^9 (1 + 4p)^{16s} (8(\ell + 1)s)^{12s}}{n} \right)^g. \end{aligned}$$

For $n \geq C_1 K^{16}$ with a constant C_1 depending only on C , from (106), we have

$$(1 + 4p)^{16s} \leq 2, \quad \frac{(d \vee K)^3 K^9 (8(\ell + 1)s)^{12s}}{n} \leq \frac{1}{8}, \quad n^{\frac{1}{2s}} \leq \log(n)^{13}.$$

which implies

$$\mathbb{E} \|R_k^{(\ell)}\|^{2s} \leq L^{4s} (mn)^{-s} n (8(\ell + 1)s)^{8s+1} 2^{2s} K^{40s} \theta^{4s(\ell-1)} \left(1 \vee \frac{d}{K}\right)^{3s}.$$

Recall $K, d \geq 1$ and $\theta \leq L$. Then by Markov's inequality, with probability at least $1 - n^{-1/2}$, the following bound holds:

$$\|R_k^{(\ell)}\| \leq \frac{L^{2\ell}}{\sqrt{mn}} K^{20} d^{3/2} \log(n)^{28}.$$

This finishes the proof for (63).

H.2. Proof of (64) on $HB^{(k)}$

Recall the definition of H from (57). We have

$$\begin{aligned} \|HB^{(k-1)}\|^{2s} &\leq \text{tr}[(HB^{(k-1)}(HB^{(k-1)})^*)^s] \\ &\leq \left(\frac{d^2}{mn}\right)^{2s} \sum_{\gamma \in W_{k,s}} \prod_{i=1}^s X_{\gamma_{2i-1,2}, \gamma_{2i-1,4}, \gamma_{2i-1,3}} \prod_{t=2}^{k-1} A_{\gamma_{2i-1,2t}, \gamma_{2i-1,2t+2}, \gamma_{2i-1,2t+1}} \\ &\quad \cdot \prod_{t=0}^{k-1} A_{\gamma_{2i,2t}, \gamma_{2i,2t+2}, \gamma_{2i,2t+1}} X_{\gamma_{2i,2k}, \gamma_{2i,2k+2}, \gamma_{2i,2k+1}}, \end{aligned}$$

where $W_{k,s}$ is the set of paths defined in (71). Let G_γ be defined as before, and $a = |E_\gamma|$, $v = |V_\gamma|$. The set of edges that carries weight from H has cardinality at most $2s$ by the boundary conditions. Hence we have

$$\begin{aligned} \mathbb{E} \|HB^{(k-1)}\|^{2s} &\leq \left(\frac{d^2}{mn}\right)^{2s} \sum_{\gamma \in W_{k,s}} \left(\frac{d}{\sqrt{mn}}\right)^{a-2s} \left(\frac{L^2}{d^2}\right)^{s(2k-4)} \\ &= \sum_{\gamma \in W_{k,s}} \left(\frac{d}{\sqrt{mn}}\right)^{a+2s} \left(\frac{L}{d}\right)^{4s(k-2)}. \end{aligned}$$

Recall the bound from Lemma 27 on $W_{k,s}(v, a)$. Since each equivalence class contains at most $n^{v_1} m^{v_2}$ many paths, we obtain

$$\mathbb{E} \|HB^{(k-1)}\|^{2s} \leq \sum_{a=1}^{4ks} \sum_{v=1}^{a+1} (4(k+1)s)^{6s(a-v+1)+2s} n^{v_1} m^{v_2} \left(\frac{d}{\sqrt{mn}}\right)^{a+2s} \left(\frac{L}{d}\right)^{4s(k-2)}.$$

The imbalance between v_1 and v_2 comes from the number of cycles in G_γ and the total number of components of the paths (since on one path, the imbalance from $v_1 - v_2$ increases by at most 1, and the first path of length k introduces an imbalance of 1). We have the following upper bound that

$$v_2 - v_1 \leq g - 1,$$

which implies

$$n^{v_1} m^{v_2} \leq (\sqrt{mn})^v \left(\sqrt{\frac{m}{n}} \right)^{g-1}. \quad (108)$$

Take $s = \left\lfloor \frac{\log n}{12 \log \log n} \right\rfloor$. With (108),

$$\begin{aligned} \mathbb{E} \|HB^{(k-1)}\|^{2s} &\leq (4(k+1)s)^{2s+1} L^{4ks-8s} d^{10s} n(\sqrt{mn})^{-2s} \sum_{g=0}^{\infty} \left(\frac{(4(k+1)s)^{6s}}{n} \right)^g \\ &\leq 2n(\sqrt{mn})^{-2s} (4(k+1)s)^{2s+1} L^{4ks-8s} d^{10s}. \end{aligned}$$

Applying Markov's inequality yields (64).

H.3. Proof of (65) and (66)

The proof of (65) is similar to the proof of (62), and one can adapt the proof of (Bordenave et al., 2022b, Equation (16.21)). We skip the proof.

For (66), first notice that $H_{ef}^{(2)}$ defined in (59) is equal to $(TPS^*)_{ef}$ defined in (60) except for $g \in \mathbf{E}_2$ such that (e, g, f) is a back-tracking path of length 6. This happens only when (1) $g_2 = e_2$, (2) $g_2 = f_2$, (3) $f_1 = e_3$. Therefore

$$\begin{aligned} -\tilde{H}_{ef} &= \left(M_{e_3 e_2} M_{f_1 e_2} + M_{e_3 f_2} M_{f_1 f_2} + \mathbf{1}\{f_1 = e_3\} \sum_{u \in V_2} M_{e_3 u}^2 \right) A_{f_1 f_3, f_2} \\ &:= M_{ef}^{(1)} + M_{ef}^{(2)} + M_{ef}^{(3)}. \end{aligned}$$

Then we can write

$$\|\underline{B}^{(k-1)} \tilde{H} B^{(\ell-k-1)}\| \leq \sum_{i=1}^3 \|\underline{B}^{(k-1)} M^{(i)} B^{(\ell-k-1)}\| \leq \sum_{i=1}^3 \|\underline{B}^{(k-1)}\| \|M^{(i)} B^{(\ell-k-1)}\|.$$

A similar argument of the bound (64) applied to each $i \in [3]$ finishes the proof.

Appendix I. Proof of Theorem 7

For simplicity, in this section, we assume that $\delta_{j,\ell}^{(i)}$ defined in (10) for each $\text{Unfold}_i(T)$ satisfies $\delta_{j,\ell}^{(i)} > c > 0$ for all i, j , and we place ourselves under the high-probability event of Corollary 6. Note that as $d \rightarrow \infty$, $L/d \ll \rho^{(i)}/\sqrt{d}$ for all i , and as such $\tau^{(i)} \leq \tau/\sqrt{d}$.

By a simple application of the triangle inequality,

$$\begin{aligned} \|T - \hat{T}\|_F &\leq \sum_{j=1}^r \left\| \nu_j \bigotimes_{i=1}^k w_j^{(i)} - \hat{\nu}_j \bigotimes_{i=1}^k \hat{w}_j^{(i)} \right\|_F \\ &\leq \sum_{j=1}^r |\nu_j - \hat{\nu}_j| + \nu_j \left\| \bigotimes_{i=1}^k w_j^{(i)} - \bigotimes_{i=1}^k \hat{w}_j^{(i)} \right\|_F \end{aligned}$$

The first term is easily bounded through Corollary 6:

$$|\nu_j - \hat{\nu}_j| \leq \frac{\nu_j^2 - \lambda_j^{(1)}}{2 \min(\nu_j, \hat{\nu}_j)} \leq C_1 \nu_j \left(\frac{\tau}{d}\right)^\ell,$$

and hence by the Cauchy-Schwarz inequality

$$\sum_{j=1}^r |\nu_j - \hat{\nu}_j| \leq C_2 \left(\frac{\tau}{d}\right)^\ell \|T\|_F.$$

For the second term, we also apply the Cauchy-Schwarz inequality to find

$$\begin{aligned} \sum_{j=1}^r \nu_j \left\| \bigotimes_{i=1}^k w_j^{(i)} - \bigotimes_{i=1}^k \hat{w}_j^{(i)} \right\|_F &\leq \|T\|_F \sqrt{\sum_{j=1}^r \left\| \bigotimes_{i=1}^k w_j^{(i)} - \bigotimes_{i=1}^k \hat{w}_j^{(i)} \right\|_F^2} \\ &= \|T\|_F \sqrt{\sum_{j=1}^r \left(2 - 2 \prod_{i=1}^k \langle w_j^{(i)}, \hat{w}_j^{(i)} \rangle \right)} \end{aligned}$$

From Corollary 6, we get

$$\langle w_j^{(i)}, \hat{w}_j^{(i)} \rangle \geq \frac{1}{\sqrt{\gamma_j^{(i)}}} - C_3 \left(\frac{\tau}{d}\right)^\ell,$$

where

$$\gamma_j^{(i)} = \left\langle \mathbf{1}, \left(I - \frac{Q^{(i)} Q^{(i)*}}{\nu_j^4 d^2} \right)^{-1} \left(w_j^{(i)} \circ w_j^{(i)} \right) \right\rangle \leq 1 + K^2 \frac{(\tau/d)^2}{1 - (\tau/d)^2}$$

where the last inequality comes from Equation (41) and $\langle \mathbf{1}, w_j^{(i)} \circ w_j^{(i)} \rangle = \|w_j^{(i)}\|^2 = 1$. Now, it is possible to check that for $a \geq 1$ and $x > 0$ small enough,

$$\frac{1}{\sqrt{1 + a \frac{x}{1-x}}} \geq 1 - ax$$

and applying this to the above bound yields

$$\langle w_j^{(i)}, \hat{w}_j^{(i)} \rangle \geq 1 - \left(\frac{K\tau}{d}\right)^2 - C_3 \left(\frac{\tau}{d}\right)^\ell.$$

Finally, for $a_1, \dots, a_k \geq 0$ such that $\sum a_i \leq 1$, we have $\prod(1 - a_i) \geq 1 - \sum a_i$, and hence

$$\sum_{j=1}^r \left(2 - 2 \prod_{i=1}^k \langle w_j^{(i)}, \hat{w}_j^{(i)} \rangle \right) \leq kr \left(\left(\frac{K\tau}{d}\right)^2 + C_3 \left(\frac{\tau}{d}\right)^\ell \right).$$

Theorem 7 ensues from the bound $\sqrt{a+b} \leq \sqrt{a} + \frac{b}{2\sqrt{a}}$.

As in Bordenave et al. (2022b), we could improve the bound of Theorem 7 by considering better estimators of the w_i , as well as an optimal shrinkage of the $\hat{\nu}_i$. However, this only gains constant factors over the “naive” ones, hence we chose to keep a simpler exposition.

Appendix J. Proof of Theorem 8

We first treat the case where $x \in \{\pm 1\}^n$. By the natural homeomorphism between $(\{\pm 1\}, \times)$ and (\mathbb{F}_2, \oplus) , each equation of the form $\hat{x}_{i_1} \dots \hat{x}_{i_k} = \tilde{T}_{i_1 \dots i_k}$ can be mapped to a linear equation in \mathbb{F}_2 . This is known as the k -XORSAT problem, and the optimal sample complexity can be found in (Creignou and Daudé, 2003):

Theorem 31 (Proposition 4.1 from (Creignou and Daudé, 2003)) *Let \mathcal{S} be a random k -XORSAT problem, where each clause of the form $y_{i_1} \oplus \dots \oplus y_{i_k} = 0$ is selected with probability p . Then, for $k \geq 3$, there exists a constant c_k such that if $p > c_k n^{1-k}$, then with high probability, the only solution to \mathcal{S} is $x = 0$.*

By considering $y_i = \hat{x}_i x_i$, when $p \geq c_k n^{1-k}$, the only solution to $\hat{x}_{i_1} \dots \hat{x}_{i_k} = \tilde{T}_{i_1 \dots i_k}$ for all i_1, \dots, i_k is $\hat{x} = x$. Note that this system can be solved through Gaussian elimination in \mathbb{F}_2 , which is a polynomial algorithm. Hence, the following holds:

Proposition 32 *If $p \geq c_k n^{1-k}$, and $T = x^{\otimes k}$ with $x \in \{\pm 1\}^n$, then there exists a polynomial-time algorithm that outputs an estimator \hat{x} such that $\hat{x} = x$ with high probability.*

Now, if $T = x^{\otimes k}$ for an arbitrary $x \in \mathbb{R}^n$, then $\text{sign}(T) = \text{sign}(x)^{\otimes k}$. Hence, applying Proposition 32 to $\text{sign}(T)$, we obtain an estimator $\hat{x} = \text{sign}(x)$. Therefore, we get

$$\langle \hat{x}, x \rangle = \langle \text{sign}(x), x \rangle = \|x\|_1.$$

Theorem 8 then ensues from the simple inequality

$$1 = \|x\|_2^2 = \sum x_i^2 \leq \|x\|_\infty \sum |x_i| = \|x\|_\infty \|x\|_1.$$