arXiv:2210.00158v1 [math.CO] 1 Oct 2022

Local and global expansion in random geometric graphs

Siqi Liu* Sidhanth Mohanty† Tselil Schramm‡ Elizabeth Yang§

February 27, 2023

Abstract

Consider a random geometric 2-dimensional simplicial complex X sampled as follows: first, sample n vectors $\boldsymbol{u}_1,\dots,\boldsymbol{u}_n$ uniformly at random on $\mathbb{S}^{d-1}$; then, for each triple $i,j,k \in [n]$, add $\{i,j,k\}$ and all of its subsets to X if and only if $\langle \boldsymbol{u}_i,\boldsymbol{u}_j\rangle \geqslant \tau, \langle \boldsymbol{u}_i,\boldsymbol{u}_k\rangle \geqslant \tau$, and $\langle \boldsymbol{u}_j,\boldsymbol{u}_k\rangle \geqslant \tau$. We prove that for every $\varepsilon > 0$, there exists a choice of $d = \Theta(\log n)$ and $\tau = \tau(\varepsilon,d)$ so that with high probability, X is a high-dimensional expander of average degree n^{ε} in which each 1-link has spectral gap bounded away from $\frac{1}{2}$.

To our knowledge, this is the first demonstration of a natural distribution over 2-dimensional expanders of arbitrarily small polynomial average degree and spectral link expansion better than $\frac{1}{2}$. All previously known constructions are algebraic. This distribution also furnishes an example of simplicial complexes for which the trickle-down theorem is nearly tight.

En route, we prove general bounds on the spectral expansion of random induced subgraphs of arbitrary vertex transitive graphs, which may be of independent interest. For example, one consequence is an almost-sharp bound on the second eigenvalue of random n-vertex geometric graphs on $\mathbb{S}^{d-1}$, which was previously unknown for most n,d pairs.

*UC Berkeley. sliu18@berkeley.edu. Supported in part by the Berkeley Haas Blockchain Initiative and a donation from the Ethereum Foundation.

†UC Berkeley. sidhanthm@cs.berkeley.edu. Supported by a Google PhD Fellowship.

‡Stanford University. tselil@stanford.edu. Supported by an NSF CAREER award # 2143246.

§UC Berkeley. elizabeth_yang@berkeley.edu. Supported by the NSF GRFP under Grant No. DGE 1752814.

Contents

1 Introduction **1**
- 1.1 Our results . . . 2
- 1.2 Spectra of random restrictions . . . 3
- 1.3 Related work . . . 4
- 1.4 Discussion and open questions . . . 6
- 1.5 Overview of the proof . . . 8

2 Preliminaries **13**
- 2.1 Linear algebra . . . 13
- 2.2 Probability . . . 14
- 2.3 The uniform distribution over the unit sphere . . . 15

3 The second eigenvalue of random restrictions **16**

4 Analyzing the discrete walk with Brownian motion **22**
- 4.1 Concentration of spherical Brownian Motion within a cap . . . 31

5 The second eigenvalue of links **32**
- 5.1 Spectral norm bound for centered links . . . 34

6 The second eigenvalue of the shell walk **40**
- 6.1 Coupling for the shell walk . . . 41
- 6.2 Relating typical rows and bounding outlier columns . . . 43

7 2-dimensional expansion of the random geometric complex **47**

8 Tightness of the tricking-down theorem **49**

A Azuma-Hoeffding for continuous processes **56**

B Bounding the expected contribution of outlier shells **56**

C Computations for evaluating ratios of typical entries **58**

1 Introduction

A graph G is called a spectral λ-expander if the second eigenvalue of its normalized adjacency matrix, $\lambda_2(G)$, is at most λ. More generally, a sequence of graphs of increasing size $(G_n)_{n\in\mathbb{N}}$ is said to be a family of (1-dimensional) spectral λ-expanders if $\lambda_2(G_n) \leqslant \lambda < 1$ as $n \to \infty$, and importantly, this implies that no vertex cut of G_n has sub-constant sparsity. Expanders are an indispensible tool in theoretical computer science and mathematics, underlying advances in pseudorandomness, coding theory, routing algorithms, and more (e.g. [INW94, SS96, Pin73], see also the survey [HLW06]); similarly, the phenomenon of expansion has enabled the analysis of approximation algorithms, probabilistically checkable proofs, embeddability of metric spaces (e.g. [ABS15, Din07, LLR95]), as well as numerous results in number theory, group theory, and other areas of pure mathematics (see e.g. the survey [Lub12]).

Sparse expander graphs were first shown to exist via the probabilistic method [KB93, Pin73]. In fact, even the most extreme version of expansion is ubiquitous: The best possible spectral expansion for a d-regular graph is the Ramanujan bound $\lambda = \frac{2\sqrt{d-1}}{d}$, and this bound is achieved (up to an additive $o_n(1)$) by a random d-regular graph with high probability [Alo86, Nil91, Fri93, Fri08].[1] Even an Erdős–Rényi graph $\mathsf{G}(n,p)$ forms a "decorated expander" for any $p > \frac{1}{n}$, which is to say the graph is an expander when one omits small isolated connected components and dangling trees [FR08, BKW14]. The sampling of these graphs may be derandomized [BL06, MOP20], so that in addition to several known explicit constructions (see [Mar73, GG81, LPS88, Mar88], and others) we have a wealth of algorithmic constructions of expander graphs to use in applications.

Higher-dimensional spectral expansion is a generalization of expansion to simplicial complexes. For simplicity, we will for the moment limit ourselves to 2-dimensional spectral expansion, which can be stated easily in terms of simple graphs. A graph $G = (V,E)$ is said to be a 2-dimensional spectral expander if G itself is an expander, and further for every vertex $v \in V$, the induced graph on v's neighbors $G[N(v)]$ (called the "link" of v) is a λ-expander for $\lambda < \frac{1}{2}$. The significance of $\frac{1}{2}$ is that when the local expansion is $\lambda < \frac{1}{2}$, this is enough to trigger a "trickling down" phenomenon that ensures that G in its entirety is an expander! Hence, higher-dimensional expanders have the remarkable property that global expansion is witnessed by local expansion. This local-to-global phenomenon has led to a number of recent breakthroughs in theoretical computer science: objects inspired by high-dimensional expanders are crucial in explicit constructions of locally testable codes [DEL$^+$22, LH22] and quantum LDPC codes [PK22, LZ22, DHLV22], and the local-to-global phenomenon has been essential in analyzing Markov chains for a wide variety of sampling problems [ALOV19].

The simplest example of a 2-dimensional expander is the complete complex, based on the complete graph K_n. Sparse examples are known as well (e.g. [CSŻ03, Li04, LSV05a, LSV05b, KKL14]), though at first their existence may seem remarkable: such graphs must be globally sparse, and yet the $O(1)$-sized local neighborhood of every vertex must be densely connected to ensure sufficient expansion. This is a delicate balance, and indeed given the state of our knowledge today the phenomenon of sparse high-dimensional expansion seems "rare," in sharp contrast with the ubiquity of 1-dimensional expansion. Only a few sparse constructions are currently known, and many of these constructions are algebraic, inheriting their expansion properties from the groups used to define them (as discussed further in Section 1.3).

[1] Though even in [Alo86] it is conjectured that random graphs achieve the Ramanujan bound, the first graphs proven achieve this bound were explicit algebraic constructions [LPS88, Mar88].

A prominent open problem in the area is to identify natural distributions over sparse higher-dimensional expanders [Lub18, Lin, Lub]; this would be highly beneficial, both for a deeper mathematical understanding and for applications in algorithms and complexity. The simplest distributions immediately fail: random d-regular graphs are locally treelike, and so with high probability $G[N(v)]$ will be an independent set (with $\lambda = 1$) for most $v \in V$. The same is true for an Erdős-Rényi graph $G(n,p)$ when $p \ll \frac{1}{\sqrt{n}}$. Though a number of distributions have been shown to have some higher-dimensional expansion properties [LM06, FGL+12, Con19, CTZ20, LMY20, Gol21], they all fall short in some sense: either they are quite dense (degree $\Omega(\sqrt{n})$) or fail to satisfy the spectral conditon $\lambda < \frac{1}{2}$. In this work, our primary question is the following:

Are there natural, high-entropy distributions over 2-dimensional expanders of average degree $\ll \sqrt{n}$?

We answer this question in the affirmative. We prove that for any $\varepsilon > 0$ and any large enough $n \in \mathbb{N}$, there exists a choice of $d \in \mathbb{N}$ such that a random n-vertex geometric graph on $\mathbb{S}^{d-1}$ with average degree n^{ε} is a 2-dimensional expander with high probability.

1.1 Our results

In order to state our results, we first give some formal definitions.

Definition 1.1 (Simplicial complex)**.** A *k-dimensional simplicial complex* X is a downward-closed collection of subsets of size at most $k+1$ over some ground set X_0, with a downward-closed weight function w.[2] Any $S \in X$ is called a $(|S|-1)$-*face*, and the restriction of X to sets of size at most $\ell+1 \leqslant k$ is called the *ℓ-skeleton* of X. The *degree* of $v \in X_0$ is the number of top-level faces that contain it.

For example, the set of all cliques of size at most $k+1$ in a graph G, where the weight of each clique is proportional to the number of $(k+1)$-cliques it occurs in, defines a k-dimensional simplicial complex.

Definition 1.2 (Link)**.** Let X be a simplicial complex. For any face S, the *link* of S in X is the simplicial complex X_S with weight function w_S, consisting of all sets in X which contain S, minus S:

$$X_S = \{T \setminus S \mid T \supseteq S, T \in X\}, \qquad w_S(T \setminus S) = w(T) \quad \forall T \in X_S$$

For example, in the simplicial complex whose highest order faces are the triangles in a graph G, the link of a vertex v is the induced graph on the neighbors of v with its isolated vertices removed.

We are interested in simplicial complexes where the links expand enough to trigger a "local-to-global phenomenon" via the trickling-down theorem, stated below in the 2-dimensional case.[3]

Theorem 1.3 (Trickling-down theorem [Opp18])**.** *Let X be a 2-dimensional simplicial complex. If its 1-skeleton is connected, and the second eigenvalue of every link's random walk matrix is at most λ, then the second absolute eigenvalue of the random walk matrix of the 1-skeleton of X is at most $\frac{\lambda}{1-\lambda}$.*

This theorem explains the significance of $\lambda = \frac{1}{2}$, since when $\lambda < \frac{1}{2}$, local expansion "trickles down" to imply global expansion. We will show that random geometric graphs, in a carefully-chosen parameter regime, have sufficient link expansion.

[2] Recall that X is called downward-closed if $S \subseteq T$ and $T \in X$ imply $S \in X$, and w is called downward-closed if weights are assigned to maximal faces, and for each non-maximal $S \in X$, we recursively define $w(S) = \sum_{x \in X_0} w(S \cup \{x\})$.

[3] The trickling-down theorem also generalizes to higher dimensions: sufficiently strong local spectral expansion of only the highest-order links implies global spectral expansion.

Definition 1.4 (Random geometric graph)**.** A *random geometric graph* $\boldsymbol{G} \sim \mathsf{Geo}_d(n,p)$ is sampled as follows: for each $i \in [n]$, a vector $\boldsymbol{u}_i$ is drawn independently from the uniform distribution over $\mathbb{S}^{d-1}$ and identified with vertex i. Then, each edge $\{i,j\}$ is included if and only if $\langle \boldsymbol{u}_i, \boldsymbol{u}_j \rangle \geqslant \tau$ where $\tau = \tau(p,d)$ is chosen so that $\mathbf{Pr}_{\mathsf{Geo}_d(n,p)}[(i,j) \in \boldsymbol{G}] = p$.

Definition 1.5 (Random geometric complex)**.** The *random geometric k-complex* $\mathsf{Geo}_d^{(k)}(n,p)$ is the distribution defined by sampling $\boldsymbol{G} \sim \mathsf{Geo}_d(n,p)$ and taking the downward-closure of the complex whose k-faces are the cliques of size $(k+1)$ in $\boldsymbol{G}$.

Our main result proves that there are conditions under which random geometric 2-complexes of degree n^ε are high-dimensional expanders enjoying the trickling-down phenomenon:

Theorem 1.6. *For every $0 < \varepsilon < 1$, there exist constants C_ε and $\delta = \exp(-O(1/\varepsilon))$ such that when $\boldsymbol{H} \sim \mathsf{Geo}_d^{(2)}(n, n^{-1+\varepsilon})$ for $d = C_\varepsilon \log n$, with high probability every vertex link of $\boldsymbol{H}$ is a $(\frac{1}{2} - \delta)$-expander, and hence its 1-skeleton is a $(1 - \frac{4\delta}{1+2\delta})$-expander.*

Along the way, we also analyze the spectrum of $\boldsymbol{G} \sim \mathsf{Geo}_d(n,p)$ directly and obtain sharper control of its second eigenvalue in a more general setting, giving bounds on the spectral norm of random geometric graphs in the full high-dimensional ($d \to_n \infty$) regime. To our knowledge, previous results in this vein are only for $d \sim n^{1/k}$ for fixed integers k [EK10, CS13, DV13, Bor13, FM19, LY22].

Theorem 1.7. *Let $\boldsymbol{G} \sim \mathsf{Geo}_d(n,p)$ and $\tau := \tau(p,d)$. Then with high probability $\boldsymbol{G}$ is a μ-expander, where*

$$\mu := (1 + o(1)) \cdot \max\left\{ (1 + o_{d\tau^2}(1)) \cdot \tau, \frac{\log^4 n}{\sqrt{pn}} \right\},$$

where $o_{d\tau^2}(1)$ denotes a function that goes to 0 as $d \cdot \tau(p,d)^2 \to \infty$.

In Section 8 we show that an eigenvalue close to τ is achieved (for some p,d), so Theorem 1.7 is close to sharp. Since in Theorem 1.6 we show that the vertex links of $\boldsymbol{G}$ have eigenvalue $\lambda \leqslant \frac{\tau}{1+\tau}$, this implies that the trickling-down theorem is tight

Proposition 1.8 (Trickling-down theorem is tight)**.** *For each $\lambda \in (0, \frac{1}{2}]$ and $\eta > 0$ there exists a 2-dimensional expander in which all vertex link eigenvalues are at most λ for which the 1-skeleton is connected with eigenvalue at least $\frac{\lambda}{1-\lambda} - \eta$.*

1.2 Spectra of random restrictions

Theorem 1.7 (and morally Theorem 1.6) is a consequence of a more general theorem that we prove concerning the spectral properties of random restrictions of graphs. We describe this result here, both because it may be of independent interest, and because it may help demystify Theorem 1.6.

Random restriction is a procedure for approximating a large graph X by a smaller graph $\boldsymbol{G}$: one selects a random subset of vertices $\boldsymbol{S}$, and then takes $\boldsymbol{G}$ to be the induced graph $X[\boldsymbol{S}]$. The random restriction $\boldsymbol{G}$ is now a smaller (and often sparser) approximation to X; this idea has been useful in a number of contexts in theoretical computer science (e.g. [GGR98, AdlVKK03, BHHS11, LRS15, HKP+17]). The core question is: to what extent do random restrictions actually inherit properties of the original graph? We will show that if random walks on X mix rapidly enough, then random restrictions inherit the *spectral* properties of the original graph.

To see the relevance of this result in our context, notice that a random geometric graph on the sphere is a random restriction of the (infinite) graph with vertex set $\mathbb{S}^{d-1}$ and edge set $\{(u,v) \mid \langle u,v\rangle \geqslant \tau\}$. Theorem 1.6 is then a consequence of the fact that the sphere is *itself* a 2-dimensional expander.

We state the theorem precisely below.

Definition 1.9 (Random restriction)**.** Suppose X is a (possibly infinite) graph, and that the simple random walk on X has unique stationary distribution ρ. We define an *n-vertex random restriction* of X to be a graph $\boldsymbol{G} \sim \mathsf{RR}_n(X)$ sampled by sampling n vertices independently according to ρ, $\boldsymbol{S} \sim \rho^{\otimes n}$, then taking $\boldsymbol{G} = X[\boldsymbol{S}]$ to be the graph induced on those vertices.

We show that if the average degree in $\boldsymbol{G}$ is not too small, $\lambda_2(\boldsymbol{G})$ reflects the rapid mixing of the random walk on X.

Theorem 1.10. *Let X be a (possibly infinite) vertex-transitive graph on which the associated simple random walk has a unique stationary distribution ρ, and let $p = \mathbf{Pr}_{\boldsymbol{G}\sim\mathsf{RR}_n(X)}[(i,j) \in E(\boldsymbol{G})]$ be the marginal edge probability of a n-vertex random restriction of X. Suppose there exist $C \geqslant 1$ and $\lambda \in [(np)^{-1/2}, 1]$ such that for any $k \in \mathbb{N}$, k-step walks on X satisfy the following mixing property: for any distribution α over $V(X)$,*

$$\mathrm{d}_{\mathrm{TV}}\left(X^k\alpha, \rho\right) \leqslant C \cdot \lambda^k,$$

where X^k denotes the k-step random walk operator on X, and furthermore suppose $pn \gg C^6 \log^4 n$. Then for any constant $\gamma > 0$,

$$\mathop{\mathbf{Pr}}_{\boldsymbol{G}\sim\mathsf{RR}_n(X)}\left[\left|\lambda_2(\widehat{A}_{\boldsymbol{G}})\right|, \left|\lambda_n(\widehat{A}_{\boldsymbol{G}})\right| \leqslant (1+o(1)) \cdot \max\left(\lambda, \frac{\log^4 n}{\sqrt{pn}}\right)\right] \geqslant 1 - n^{-\gamma},$$

where $\widehat{A}_{\boldsymbol{G}}$ is the (normalized) adjacency matrix of $\boldsymbol{G}$.

Remark 1.11. It is likely that some of the conditions of Theorem 1.10 could be weakened. The decay of total variation could plausibly be replaced with a (much weaker) assumption about the spectral gap of X; this would not impact our results for $\mathbb{S}^{d-1}$, but may be useful in other applications. Transitivity is assumed mostly to make the proof of Theorem 1.10 go through at this level of generality; to prove Theorem 1.6 we re-prove a version of Theorem 1.10 for the specific non-transitive case where X is a link of a vector in the sphere (a spherical cap).

1.3 Related work

We give a brief overview of related work. While so far we have focused on a spectral notion of high-dimensional expanders (HDX), there are two additional notions: coboundary and cosystolic expansion. These are meant to generalize the Cheeger constant, a cut-based measure of graph expansion.

Distributions over high-dimensional expanders. The existence of natural distributions over sparse HDXs has been a question of interest since sparse HDX were first shown to exist (and this was highlighted as an important open problem in e.g. [Lub18, Lub]).

The early work of Linial and Meshulam [LM06] considered the distribution over 2-dimensional complexes in which all edges $\binom{[n]}{2}$ are included, and each triangle is included independently with probability p; they identified the phase transition at p for coboundary connectivity for this distribution (see also the follow-ups [BHK11, MW09, LP16]). This distribution has the drawback that the 1-skeleton of these complexes is K_n, and so the resulting complex is far from sparse.

In [FGL+12], the authors show that a union of d random partitions of $[n]$ into sets of size $k+1$ with high probability produces a *geometric* expander [Gro10], which is a notion of expansion which measures how much the faces must intersect when the complex is embedded into $\mathbb{R}^k$. The resulting complexes have disconnected links when $d \ll \sqrt{n}$, and so they fail to be spectral HDXs.

The work of [LMY20] introduces a distribution over spectral expanders with expansion exactly $\frac{1}{2}$ by taking a tensor product of a random graph and a HDX; the authors show that down-up walks on these expanders mix rapidly, and [Gol21] introduces a reweighing of these complexes which yields improved mixing time bounds. However, the links in these complexes fail to satisfy $\lambda < \frac{1}{2}$, and so fall outside of the range of the trickling-down theorem. The same drawback applies to [Con19, CTZ20]: they show that up-down walks mix on random polylogarithmic-degree graphs given by subsampling a random set of generators of a Cayley graph. However, these graphs do not satisfy the conditions of the trickling-down theorem.

Explicit constructions. One of the first constructions of sparse high-dimensional spectral expanders was the Ramanujan complex of [CSŻ03, Li04, LSV05a, LSV05b], which generalize the Ramanujan expander graphs of [LPS88]. Not only are these spectral expanders, but [KKL14, EK16] also show that they are co-systolic expanders. These Ramanujan complexes are algebraic by nature, constructed from the Cayley graphs of $\mathrm{PSL}_d(\mathbb{F}_q)$. Other algebraic constructions include that of [KO18]; the authors analyze the expansion properties of coset complexes for various matrix groups. They achieve sparse spectral expanders, with local expansion arbitrarily close to 0. More recently, [OP22] extend the coset complex construction to the more general family of Chevalley groups.

A few combinatorial constructions for HDX are also known. [CLP20] prove that objects called (a,b)-expanders are two-dimensional spectral expanders; they give a graph-product-inspired construction of a family of such expanders, and show that other known complexes [CSŻ03, Li04, LSV05a, LSV05b, KO18] are also (a,b)-expanders. Their work is extended by [FI20] to higher dimensions.

Applications of HDX. The local-to-global phenomenon in HDX has already been useful in many settings. [DK17] use high-dimensional spectral expanders to construct "agreement expanders," whose links give rise to local agreement tests: given "shards" of a function that pass a large fraction of the local agreement tests, the authors can conclude the presence of a "global" function g that stitches the shards together. In coding theory, the locally testable codes of [DEL+22] and quantum LDPCs of [PK22, LZ22, DHLV22] utilize a common simplicial-complex-like structure called the square Cayley complex, whose local-to-global properties are essential in the analysis of these codes.

The local-to-global phenomenon also implies that "down-up" walks on the associated simplicial complex mix (as made formal in [AL20]). A k-down-up walk is supported k-faces of the simplicial complex, and transitions occur by dropping down into a random $(k-1)$-face, then transitioning up to a random k-face (one can also define the "up-down" walk analagously). This local-to-global analysis has recently been influential in the study of mixing times of Markov chains. Several well-studied Markov chains can be recast as the k-down-up random walk of a carefully designed simplicial complex. One notable example is the matroid basis exchange walk, which is an algorithm for sampling independent sets of a matroid (e.g. spanning trees in the graphical matroid). [ALOV19] were able to obtain an improved mixing time bound for the basis exchange walk–a significant breakthrough that, due to the local-to-global property, was achieved through the analysis of simple, "local" view of the matroids.

Random geometric graphs and random kernel matrices. Random restrictions of metric spaces such as $\mathbb{S}^{d-1}$ and $[-1,1]^d$ are well-studied in the fixed-dimensional regime, where $d = O(1)$ and

$n \to \infty$ (see the survey of Penrose [Pen03]). In our work we are interested in the high-dimensional setting, where $d \to \infty$ with n. The high-dimensional setting was first studied only recently, initiated by [DGLU11, BDER16], and many mysteries remain in this young area of study.

Our Theorem 1.7 is related to the study of *kernel random matrices*: random $n \times n$ matrices whose (i,j)-th entry is given by $f_d\left(\langle \boldsymbol{u}_i, \boldsymbol{u}_j \rangle\right)$, for $f_d : \mathbb{R} \to \mathbb{R}$ and $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_n$ sampled independently from some distribution over $\mathbb{R}^d$. The special case of $\boldsymbol{u}_i \sim \mathrm{Unif}(\mathbb{S}^{d-1})$ and $f_d(x) = \mathbf{1}[x \geqslant \tau(p,d)]$ yields the adjacency matrix of $\mathsf{Geo}_d(n,p)$. A line of work initiated by [KG00] studies the spectrum of kernel random matrices [EK10, CS13, DV13, Bor13, FM19], and the most recent work [LY22] characterizes the limiting empirical spectral distribution when $d = \Theta(n^{1/k})$ for k a fixed constant and f can be "reasonably" approximated by polynomials (in a sense that is flexible enough to capture the indicator $f_d(x) = \mathbf{1}[x \geqslant \tau(p,d)]$). In comparison with our results, they characterize the entire empirical spectral distribution, but we do not need to restrict $d \sim n^{1/k}$ for integer k, which is crucial for our applications.

1.4 Discussion and open questions

Sparser high-dimensional expanders from random restrictions? As hinted in Section 1.2 the random geometric complex fits in the broader framework of *random restrictions* of simplicial complexes: starting with a dense high-dimensional expander X, we sample a subset of vertices $\boldsymbol{S}$ of X to produce the sparser induced complex $X[\boldsymbol{S}]$.

We have shown in Theorem 1.10 that $X[\boldsymbol{S}]$ (to some extent) inherits the spectral properties of X itself, and we've leveraged this to show that for any polynomial average degree, one can produce a 2-dimensional expander by taking a random restriction of X the sphere in a particular dimension and with a particular connectivity distance. We hope that Theorem 1.10 (or a strengthening thereof, see Remark 1.11) might help us identify additional natural distributions over sparser and/or higher-dimensional complexes. More specifically,

> *Is there a simplicial complex X whose random restrictions yield high-dimensional expanders whose links have eigenvalue $< \frac{1}{2}$, of sub-polynomial or polylogarithmic degree*[4]*?*

As a starting point, we remark that geometric graphs on the unit sphere work because the corresponding X itself has link expansion better than $\frac{1}{2}$, witnessing that $\mathbb{S}^{d-1}$ itself is an expander. Some simpler-to-analyze metric spaces do not have this property; for example:

Shortest path metric in a graph. Starting with a connected, locally-treelike d-regular graph G, consider the geometric graph given by connecting pairs of vertices at distance $\leqslant 2$ in G. The triangle complex on the resulting graph has links which are connected, and further the 1-skeleton is an expander if G is an expander.[5] However, the links cannot identify whether G is an expander, and so link expansion cannot be better than $\frac{1}{2}$. To see why, consider a first case where G is a random d-regular graph (expanding), and a second case where G consists of two random d-regular graphs connected by a bridge (non-expanding); because G is locally treelike in both cases, the links in these two cases will be identical.

[4] We note that constant average degree would likely require additional work; this is not just because of the polylogarithmic factors appearing in the statement of Theorem 1.10, but because in a random restriction the degree distribution of each vertex is $\mathsf{Binom}(n,p)$ and so when $p = \Theta(1/n)$ one will have isolated vertices; this is the same as the phenomenon wherein Erdős-Rényi graphs of degree $O(1)$ are not expanders until one restricts to the giant component.

[5] Technically, we require this of a reweighting of G where each edge is weighted according to the number of triangles it participates in; concentration phenomena ensure that the expansion of G and this weighted graph are similar.

The d-dimensional torus with ℓ_∞ metric. Consider the geometric graph X on the d-dimensional torus $[-R,R]^d$ with "wraparound" (so that $-R$ is identified with R), in which we connect u,v if $\|u-v\|_\infty \leqslant \tau$. This space is simple to analyze because of its product structure; the expansion is dictated by the ratio of τ to the side length R, worsening as R grows relative to τ.[6] Each link is simply the box $[-\tau,\tau]^d$ with the same ℓ_∞-edge condition, regardless of the value of R. Since R dictates the global expansion, the link expansion cannot be better than $\frac{1}{2}$.

By way of contrast, it is not possible to plant the links of the geometric graph on $\mathbb{S}^{d-1}$ in a nonexpanding graph: for instance, it is possible to determine the radius R of a sphere of unknown scale given only a link in its geometric graph.

We also remark that Theorem 1.10 could be used to obtain expanders of dimension $k>2$; indeed, it seems that this is within reach even using $\mathsf{Geo}_d^{(k)}(n,p)$. A direct approach, in the case of the sphere, is to perform the conditioning from Section 5 not only for spherical caps, but for intersections of $k-1$ spherical caps as well; perhaps there is a more elegant alternative approach?

How faithfully do random geometric graphs discretize continuous manifolds? One interpretation of Theorem 1.7 is that the random geometric graph $\mathsf{Geo}_d(n,p)$ offers a good approximation (in spectral norm) for the corresponding metric on $\mathbb{S}^{d-1}$ when pn is large enough relative to d. A natural question is to extend this to other properties of $\mathbb{S}^{d-1}$; for example, do random geometric graphs offer a good approximation on the rest of the spectrum? Numerical experiments suggest the following (informal) conjecture.

Conjecture 1.12. *For $\boldsymbol{G}\sim\mathsf{Geo}_d(n,p)$, the spectrum of the normalized adjacency matrix $A_{\boldsymbol{G}}$ breaks into a "bulk" portion and an "outlier" portion where every bulk eigenvalue is at most $O(1/\sqrt{pn})$ in magnitude, and every outlier eigenvalue is "close" to an eigenvalue of the graph on the sphere with an edge between every u,v with $\langle u,v\rangle \geqslant \tau(p,d)$.*

A proof of the above conjecture, and an investigation of whether an analogous phenomenon holds on general manifolds, would be very interesting.

Spectral algorithms for random geometric graphs. Here, we have given some of the first analyses of the spectral radius of random geometric graphs on the sphere. One appeal of random geometric graphs on the sphere, or in Gaussian space, is that they offer a more natural model for networks arising from data than, e.g., Erdős-Rényi graphs. The idea is that in modern networks, we often think of each node as being representable by a latent feature vector, with nearby nodes having similar features. Hence, geometric graphs are promising as an alternative testbed for rigorous analysis of algorithms. Yet currently, they have not been studied much in such a context, in part because of the absence of tools for their analysis.

A natural question is whether one could build on our work to analyze spectral clustering algorithms in "random geometric block model" graphs.

Question 1.13. Suppose $\boldsymbol{G}\sim\mathsf{RR}_n(\frac{1}{2}\mathcal{N}(0,\Sigma_1)+\frac{1}{2}\mathcal{N}(\mu,\Sigma_2))$; that is, n points $\boldsymbol{u}_1,\dots,\boldsymbol{u}_n$ are sampled from the uniform mixture over the d-dimensional Gaussian distributions $\mathcal{N}(0,\Sigma_1)$ and $\mathcal{N}(\mu,\Sigma_2)$, then $(i,j)\in E(\boldsymbol{G})$ if and only if $\|\boldsymbol{u}_i-\boldsymbol{u}_j\|\leqslant\varepsilon$. Does spectral clustering recover the component membership of the datapoints?

[6] This can be seen by analogy to the spectrum of a d-tensor-power of the discrete $R\tau^{-1}$-cycle.

This question is a more accurate representation of clustering problems arising from real data than, say, the question of applying spectral clustering to recover cluster memberships in the stochastic block model; it would be interesting to understand the conditions (on n, d, δ, μ, and Σ_1, Σ_2) which guarantee that spectral clustering succeeds.

1.5 Overview of the proof

We now explain how we prove our main theorem, Theorem 1.6, which states that for a complex sampled from $\boldsymbol{H} \sim \mathsf{Geo}_d^{(2)}(n,p)$ for $p = n^{-1+\varepsilon}$ with $0 < \varepsilon < 1$ and $d = C_\varepsilon \log n$, with high probability every link of $\boldsymbol{H}$ is a $\left(\frac{1}{2} - \delta\right)$-expander for some $\delta = \exp(-O(1/\varepsilon))$, and its 1-skeleton is a $\left(1 - \frac{4\delta}{1+2\delta}\right)$-expander. By the trickling-down theorem, it suffices for us to prove:

1. All n vertices' corresponding links in $\boldsymbol{H}$ are $\left(\frac{1}{2} - \delta\right)$-expanders with high probability.
2. The 1-skeleton of $\boldsymbol{H}$ is connected with high probability.

To show Item 2, it is enough to show that some reweighting of the 1-skeleton expands; Item 1 implies that every edge (i,j) must participate in at least one triangle (otherwise the link would contain isolated vertices), so the unweighted 1-skeleton is just the adjacency matrix of an unweighted graph from $\mathsf{Geo}_d(n,p)$. En route to proving Item 1 we'll prove that unweighted random geometric graphs expand, by this logic yielding Item 2 a consequence.

Analyzing link expansion. We establish Item 1 by showing that that each of the n links is a $\left(\frac{1}{2} - \delta\right)$-expander with probability $1 - o(1/n)$, then applying a union bound. We can think of sampling the link of vertex i_w in $\boldsymbol{H}$ by first choosing the number of neighbors $\boldsymbol{r} \sim \mathsf{Binom}(n-1,p)$, then sampling $\boldsymbol{r}$ points $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{\boldsymbol{r}}$ independently and uniformly from a measure-p cap in $\mathbb{S}^{d-1}$ centered at some point w (corresponding to the vector of the link vertex i_w), placing an edge between every i,j such that $\langle \boldsymbol{v}_i, \boldsymbol{v}_j \rangle \geqslant \tau(p,d)$. Finally, we remove any isolated vertices; here, we'll show that the graph expands with high probability before removing these isolated vertices, which implies that no isolated vertices have to be removed. For the remainder of the overview, let $\tau = \tau(p,d)$. We'll show that:

Theorem 1.14 (Informal version of Theorem 5.1)**.** *Let $\boldsymbol{G}$ be the link of some point $w \sim \mathbb{S}^{d-1}$ induced by $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m \sim \mathrm{cap}_p(w)$. Then with high probability $\boldsymbol{G}$ is a μ-expander where*

$$\mu := (1 + o(1)) \cdot \max\left\{ \frac{\tau}{\tau+1}, \frac{\log^4 m}{\sqrt{qm}} \right\} + o_d(1).$$

Here $q = \mathbf{Pr}_{u,v \sim \mathbb{S}^{d-2}}\left[\langle u, v \rangle \geqslant \frac{\tau}{\tau+1}\right]$.

Links are essentially random geometric graphs in one lower dimension. Since most of the measure of the cap lies close to its boundary, intuitively the link is distributed *almost* like a random geometric graph with points drawn independently from the cap boundary, i.e. the shell $\mathrm{shell}_p(w) := \{x : \langle x, w \rangle = \tau\}$. Our proof of Theorem 1.14 must pay attention to the fluctuations in $\langle \boldsymbol{v}_i, w \rangle - \tau$, but to simplify our current discussion we assume each link is in fact a random geometric graph on $\mathrm{shell}_p(w)$, and address the fluctuations later in the overview.

Observe that a uniformly random $\boldsymbol{v}$ from $\mathrm{shell}_p(w)$ is distributed as $\tau \cdot w + \sqrt{1-\tau^2} \cdot \boldsymbol{u}$ where $\boldsymbol{u}$ is a uniformly random unit vector orthogonal to w. Using this decomposition, we see that $\langle \boldsymbol{v}_i, \boldsymbol{v}_j \rangle \geqslant \tau$

if and only if $\langle \boldsymbol{u}_i, \boldsymbol{u}_j \rangle \geqslant \frac{\tau}{1+\tau}$. Thus, under our simplifying assumption, the link is distributed exactly like a random geometric graph on $\mathbb{S}^{d-2}$ with inner product threshold $\frac{\tau}{1+\tau}$. Hence (up to the difference between $\mathrm{cap}_p(w)$ and $\mathrm{shell}_p(w)$), to understand link expansion we can study the second eigenvalue of a random geometric graph on the sphere.

Remark 1.15 (Requiring $d = \Theta(\log n)$)**.** In light of Theorem 1.14 (and even the heuristic discussion above), it turns out that $d = \Theta(\log n)$ is the only regime for which the links can be connected while the 1-skeleton has average degree $\ll \sqrt{n}$. To see this, we consider the relationship between p, τ, and d; we have that

$$p = \mathop{\mathbf{Pr}}_{\boldsymbol{v},\boldsymbol{v}' \sim \mathbb{S}^{d-1}} \left[\langle \boldsymbol{v}, \boldsymbol{v}' \rangle \geqslant \tau\right] = \Theta\left(\tfrac{1}{\tau d}\right) \cdot \left(1 - \tau^2\right)^{\frac{d-1}{2}} \approx \exp(-d\tau^2/2). \tag{1}$$

See Lemma 2.8 for a formal argument.[7] Note that the arguments above in conjunction with (1) imply that the probability that two vertices within the link are connected is also roughly

$$q = \mathop{\mathbf{Pr}}_{\boldsymbol{u},\boldsymbol{u}' \sim \mathbb{S}^{d-2}} \left[\langle \boldsymbol{u}, \boldsymbol{u}' \rangle \geqslant \tfrac{\tau}{1+\tau}\right] = \Theta\left(\tfrac{1}{\tau d}\right) \cdot \left(1 - \tfrac{\tau^2}{(1+\tau)^2}\right)^{\frac{d-2}{2}},$$

since the link is like a random geometric graph on $\mathrm{shell}_p(w)$.

Connectivity within the links in conjunction with sparsity now requires us to have $d \in \Theta(\log n)$: The number of vertices inside each link concentrates around $m = np$, so the average degree inside the link is $qm \approx qpn$; we must have the average link degree $qpn \geqslant 1$, otherwise the link is likely disconnected. Now, if $\tau = o(1)$, then $\tau \approx \frac{\tau}{1+\tau}$ and $p \approx q$, so $qpn \geqslant 1 \implies p^2 n \gtrapprox 1 \implies p \gtrapprox n^{-1/2}$, ruling out a 1-skeleton with average degree $\ll \sqrt{n}$. Hence we need $\tau = \Omega(1)$. Given that $\tau = \Omega(1)$, (1) implies that to have the average 1-skeleton degree $\sqrt{n} \geqslant pn \geqslant 1$ we need $d \in \Theta(\log n)$.

Spectral expansion in random geometric graphs. We now explain how to prove near-sharp second eigenvalue bounds for random geometric graphs.

Theorem (Restatement of Theorem 1.7)**.** *Let $\boldsymbol{G} \sim \mathsf{Geo}_d(n,p)$ and $\tau := \tau(p,d)$. Then with high probability $\boldsymbol{G}$ is a μ-expander, where*

$$\mu := (1 + o(1)) \cdot \max\left\{(1 + o_{d\tau^2}(1)) \cdot \tau, \frac{\log^4 n}{\sqrt{pn}}\right\},$$

where $o_{d\tau^2}(1)$ denotes a function that goes to 0 as $d \cdot \tau(p,d)^2 \to \infty$.

As mentioned above, Theorem 1.7 is a consequence of the more general Theorem 1.10 about the second eigenvalue of random restrictions of vertex-transitive graphs, and the inner product threshold $\tau = \tau(p,d)$ appears as the mixing rate of the random walk on $\mathbb{S}^{d-1}$ where a step originating at v walks to a random vector in $\mathrm{cap}_p(v)$. Via standard concentration arguments applied to the vertex degrees, to prove the above it suffices to bound $\|A_{\boldsymbol{G}} - \mathbf{E} A_{\boldsymbol{G}}\| \leqslant \mu \cdot pn$, where $A_{\boldsymbol{G}}$ is the (unnormalized) adjacency matrix of $\boldsymbol{G}$. We'll focus on the regime where $pn \gg \mathrm{polylog} n$, so that $\mu \approx \tau$.

Trace method for random geometric graphs. To bound $\|A_{\boldsymbol{G}} - \mathbf{E} A_{\boldsymbol{G}}\|$, we employ the trace method, bounding the expected trace of a power of $A_{\boldsymbol{G}} - \mathbf{E} A_{\boldsymbol{G}}$. This is sufficient for the following

[7] Heuristically, it makes sense that $p = \mathbf{Pr}[\langle \boldsymbol{v}, \boldsymbol{v}' \rangle \geqslant \tau] \approx \exp(-\Theta(\tau^2 d))$, because $\langle \boldsymbol{v}, \boldsymbol{v}' \rangle$ is approximately $\mathcal{N}(0, \frac{1}{d})$.

reason: for convenience, let $\overline{A}_{\boldsymbol{G}} = A_{\boldsymbol{G}} - \mathbf{E}A_{\boldsymbol{G}}$, and let ℓ be any non-negative, even integer. Since ℓ is even,

$$\left\|\overline{A}_{\boldsymbol{G}}\right\|^{\ell} = \left\|\overline{A}_{\boldsymbol{G}}^{\ell}\right\| \leqslant \operatorname{tr}\left(\overline{A}_{\boldsymbol{G}}^{\ell}\right),$$

And so applying Markov's inequality,

$$\mathbf{Pr}\left(\left\|\overline{A}_{\boldsymbol{G}}\right\| \geqslant e^{\varepsilon}\left(\mathbf{E}\operatorname{tr}\left(\overline{A}_{\boldsymbol{G}}^{\ell}\right)\right)^{1/\ell}\right) = \mathbf{Pr}\left(\left\|\overline{A}_{\boldsymbol{G}}\right\|^{\ell} \geqslant e^{\varepsilon\ell}\,\mathbf{E}\operatorname{tr}\left(\overline{A}_{\boldsymbol{G}}^{\ell}\right)\right) \leqslant \exp(-\varepsilon\ell).$$

Thus, our goal reduces to bounding the expectation of $\operatorname{tr}(\overline{A}_{\boldsymbol{G}}^{\ell})$ for a sufficiently large even ℓ; in particular, if we choose $\ell \gg \log n$, then since $\overline{A}_{\boldsymbol{G}}$ has n eigenvalues, $\operatorname{tr}(\overline{A}_{\boldsymbol{G}}^{\ell})^{1/\ell}$ is a good "soft-max" proxy for $\|\overline{A}_{\boldsymbol{G}}\|$, and we will obtain high-probability bounds.

We now explain why properties of random walks on $\mathbb{S}^{d-1}$ naturally arise when applying the trace method. Concretely, $\operatorname{tr}(\overline{A}_{\boldsymbol{G}}^{\ell})$ is a sum over products of entries of $\overline{A}_{\boldsymbol{G}}$ corresponding to closed walks of length ℓ in the complete graph K_n on n vertices:

$$\operatorname{tr}\left(\overline{A}_{\boldsymbol{G}}^{\ell}\right) = \sum_{i_0,\ldots,i_{\ell-1}\in[n]} \prod_{t=0}^{\ell-1} (\overline{A}_{\boldsymbol{G}})_{i_t i_{t+1 \bmod \ell}},$$

The walk $i_0, i_2, \ldots, i_{\ell-1}, i_0$ can be represented as a directed graph. When we take the expectation, the symmetry of the distribution means that all sequences $i_0, \ldots, i_{\ell-1}$ which result in the same graph (up to relabeling) give the same value. That is, letting $\mathcal{W}_\ell$ be the set of all such graphs, and for each $W \in \mathcal{W}_\ell$ letting N_W be the number of ways it can arise in the sum above,

$$\mathbf{E}\operatorname{tr}\left(\overline{A}_{\boldsymbol{G}}^{\ell}\right) = \sum_{W\in\mathcal{W}_\ell} N_W \cdot \mathbf{E}\prod_{(i,j)\in W} (\overline{A}_{\boldsymbol{G}})_{ij}. \tag{2}$$

To bound this sum, we must bound the expectation contributed by each $W \in \mathcal{W}_\ell$. For the sake of this overview we will consider only the case when $W = C_\ell$, the cycle on ℓ vertices, as it requires less accounting than the other cases; however it is reasonable to restrict our attention to this case for now, as bounding it already demonstrates our main ideas, and because this term roughly dominates the sum with $N_{C_\ell} \gg N_{W'}$ for all other $W' \in \mathcal{W}_\ell$ at $\ell = \operatorname{polylog} n$ and $pn \gg \operatorname{polylog} n$.[8]

We now bound the expectation for the case $W = C_\ell$; readers uninterested in the finer details may skip to the conclusion in (4). We expand the product using that $(\overline{A}_{\boldsymbol{G}})_{ij} = \boldsymbol{A}_{ij} - p$ (since $\mathbf{E}[\boldsymbol{A}_{ij}] = p$):

$$\mathbf{E}\prod_{i=1}^{\ell}(\boldsymbol{A}_{i,i+1} - p) = \sum_{T\subseteq[\ell]} (-p)^{\ell-|T|}\,\mathbf{E}\prod_{i\in T}\boldsymbol{A}_{i,i+1} = \sum_{T\subseteq[\ell]} (-p)^{\ell-|T|}\,\mathbf{Pr}[\{(i,i+1) : i \in T\} \text{ is subgraph of } \boldsymbol{G}]. \tag{3}$$

and thus our focus is to understand subgraph probabilities in a random geometric graph. It is not too hard to see that when the edges specified by T form a forest, its subgraph probability is $p^{|T|}$, identical to its counterpart in an Erdős–Rényi graph; the nontrivial correlations introduced by the geometry only play a role when T has cycles. Hence, the sum (3) simplifies,

$$\mathbf{E}\prod_{i=1}^{\ell}(\overline{A}_{\boldsymbol{G}})_{i,i+1} = \sum_{T\subsetneq[\ell]} (-p)^{\ell-|T|}p^{|T|} + \mathbf{Pr}[C_\ell \text{ is subgraph of } \boldsymbol{G}] = \mathbf{Pr}[C_\ell \text{ is subgraph of } \boldsymbol{G}] - p^{\ell}, \tag{4}$$

where we used that the binomial sum is equal to $(p-p)^\ell = 0$.

Hence it remains to estimate the subgraph probability of a length-ℓ cycle. We will now see how subgraph probabilities are related to the mixing rate of a random walk on $\mathbb{S}^{d-1}$.

[8] Briefly, this is because whenever $i_0, \ldots, i_{\ell-1}$ are all distinct elements of $[n]$, the resulting walk's graph is a cycle, and when $\ell = \operatorname{polylog} n$, ℓ indices sampled at random from $[n]$ are all distinct with high probability.

Subgraph probability of a cycle in a random geometric graph. For the cycle $C_\ell = 0, 1, \ldots, \ell - 1, 0$, by Bayes' rule:

$$\mathbf{Pr}[C_\ell \in \boldsymbol{G}] = \prod_{i=0}^{\ell-1} \mathbf{Pr}[(i, i+1) \in \boldsymbol{G} \mid \forall j < i, (j, j+1) \in \boldsymbol{G}] = p^{\ell-1} \cdot \mathbf{Pr}[(\ell-1, 0) \in \boldsymbol{G} \mid 0, 1, \ldots \ell - 1 \in \boldsymbol{G}],$$

since in all but the step $i + 1 = \ell$, the graph in question is a forest. Identifying each i with a point $\boldsymbol{x}_i$ on $\mathbb{S}^{d-1}$, for any choice of $\boldsymbol{x}_0$ the above probability can equivalently be written as

$$p^{\ell-1} \cdot \mathbf{Pr}[\langle \boldsymbol{x}_{\ell-1}, \boldsymbol{x}_0 \rangle \geqslant \tau \mid \langle \boldsymbol{x}_i, \boldsymbol{x}_{i+1} \rangle \geqslant \tau : 0 \leqslant i \leqslant \ell - 2].$$

Denoting with P the transition kernel of the random walk we alluded to earlier, where in one step we walk from a point x to a uniformly random point in $\mathrm{cap}_p(x)$, we can write the distribution of $\boldsymbol{x}_\ell \mid \{\boldsymbol{x}_0, \langle \boldsymbol{x}_i, \boldsymbol{x}_{i+1} \rangle \geqslant \tau : 0 \leqslant i \leqslant \ell - 2\}$ as $P^{\ell-1}\delta_{\boldsymbol{x}_0}$ where $\delta_{\boldsymbol{x}_0}$ refers to the point mass probability distribution supported at $\boldsymbol{x}_0$. In turn, we can write the subgraph probability as:

$$p^{\ell-1} \cdot \mathop{\mathbf{Pr}}_{\boldsymbol{x}_{\ell-1} \sim P^{\ell-1}\delta_{\boldsymbol{x}_0}} \left[\boldsymbol{x}_{\ell-1} \in \mathrm{cap}_p(\boldsymbol{x}_0)\right].$$

If $\boldsymbol{x}_{\ell-1}$ were sampled from the uniform distribution ρ on $\mathbb{S}^{d-1}$ then the probability of landing in $\mathrm{cap}_p(\boldsymbol{x}_0)$ would be p, which lets us upper bound the subgraph probability by:

$$p^{\ell-1} \cdot \left(p + \mathrm{d_{TV}}\left(P^{\ell-1}\delta_{\boldsymbol{x}_0}, \rho\right)\right).$$

The terms for more complicated subgraphs $W' \in \mathscr{W}_\ell$ also similarly depend on the mixing properties of P via subgraph probabilities. Our next goal then is to understand the mixing properties of P.

Remark 1.16. To prove Theorem 1.10 about random restrictions, the same strategy is used to relate subgraph probabilities with mixing rate of the random walk on the original graph we start with.

Mixing properties of P. We show that the walk over $\mathbb{S}^{d-1}$ with transition kernel P contracts the TV distance by coupling this discrete walk with the continuous Brownian motion U_t over $\mathbb{S}^{d-1}$. Then via a known log-Sobolev inequality for Brownian motion on spheres, we can prove the following contraction property for P.

Theorem 1.17 (Informal version of Theorem 4.6)**.** *For any probability measure α over $\mathbb{S}^{d-1}$ and integer $k \geqslant 0$,*

$$\mathrm{d_{TV}}\left(P_p^k \alpha, \rho\right) \leqslant \left((1 + o_{d\tau^2}(1)) \cdot \tau\right)^{k-1} \cdot \sqrt{\frac{1}{2} \log \frac{1}{p}},$$

where P_p denotes the transition kernel in which every $x \in \mathbb{S}^{d-1}$ walks to a uniformly random point in the measure-p cap around it and $o_{d\tau^2}(1)$ denotes a function that goes to 0 as $d\tau^2 \to \infty$.

We leave the details to Section 4, but in brief, the reason we are able to execute this coupling is that the probability mass in $P\delta_{\boldsymbol{x}_0}$ concentrates around $\mathrm{shell}_{=\tau}(\boldsymbol{x}_0)$, and most of the $(\frac{1}{d-1} \log \frac{1}{\tau})$-step Brownian motion starting from $\boldsymbol{x}_0$ concentrates at $\mathrm{shell}_{=\tau}(\boldsymbol{x}_0)$, so when $t = \frac{1}{d-1} \log \frac{1}{\tau}$ the operators P and U_t have similar action.

We can now apply Theorem 1.17 to bound $\mathrm{d_{TV}}\left(P^{\ell-1}\delta_{\boldsymbol{x}_0}, \rho\right)$ with $\alpha = \delta_{\boldsymbol{x}_0}$ and $k = \ell - 1$:

$$\mathrm{d_{TV}}\left(P^{\ell-1}\delta_{\boldsymbol{x}_0}, \rho\right) \leqslant ((1 + o(1))\tau)^{\ell-2} \sqrt{\tfrac{1}{2} \log \tfrac{1}{p}}.$$

Spectral norm of random geometric graph. We now return to bounding the expected trace of $\overline{A}_{\boldsymbol{G}}^{\ell}$; putting together the above, we have the bound

$$\mathbf{E}\prod_{(i,j)\in C_\ell}(\overline{A}_{\boldsymbol{G}})_{ij} \leqslant \mathbf{Pr}[C_\ell \in \boldsymbol{G}] - p^\ell \leqslant p^{\ell-1}\left(p + \mathrm{d_{TV}}\left(P^{\ell-1}\delta_{\boldsymbol{x}_0}, \rho\right)\right) - p^\ell \leqslant p^{\ell-1}((1+o(1))\tau)^{\ell-2}\sqrt{\tfrac{1}{2}\log\tfrac{1}{p}}.$$

The coefficient N_{C_ℓ} in front of the $W = C_\ell$ term in (2) is the number of sequences $i_1,\dots,i_\ell \in [n]$ which yield an ℓ-cycle graph; this happens if and only if all of the indices are distinct, so $N_{C_\ell} = \ell! \cdot \binom{n}{\ell} \leqslant n^\ell$. Hence the contribution of the ℓ-cycle to the sum is at most $((1+o(1))np\tau)^{\ell-2}\cdot\mathrm{poly}(n)$ when $p > 1/n$. By a careful accounting similar to the above for all graphs $W \in \mathscr{W}_\ell$, one can show that in the parameter regime $pn \gg \mathrm{poly}\log(n)$ and $\ell = \mathrm{poly}\log n$, the term $W = C_\ell$ contains $(1-o(1))$ of the total value of this sum, so we obtain the bound

$$\left[\mathbf{E}\,\mathrm{tr}(\overline{A}_{\boldsymbol{G}}^{\ell})\right]^{1/\ell} \leqslant \left((1+o(1))\cdot((1+o(1))np\tau)^{\ell-2}\cdot\mathrm{poly}(n)\right)^{1/\ell} = (1+o(1))np\tau,$$

when we choose $\ell = \omega(\log n)$. Applying Markov's inequality we conclude that $\|\overline{A}_{\boldsymbol{G}}\| \leqslant (1+o(1))np\tau$ with high probability, and normalizing by the degrees (which concentrate well around np) we conclude our upper bound of τ in Theorem 1.7.

Adapting the spectral norm bound to links. Up until now, we have pretended that the link of i_w is a random geometric graph, where the vertices are identified with vectors in $\mathrm{shell}_{=\tau}(w)$, rather than $\mathrm{cap}_{\geqslant\tau}(w)$. While it is true that most of the probability mass in $\mathrm{cap}_{\geqslant\tau}(w)$ is close to the boundary, some $\frac{1}{\mathrm{poly}(m)}$-fraction of the vertices j in the link will have $\langle \boldsymbol{v}_j, w\rangle = \boldsymbol{\kappa}_j > (1+\delta)\tau$ for some $\delta > 0$. And within the link, these vertices will have higher expected degree: for $\boldsymbol{v}_i, \boldsymbol{v}_j$ having $\langle \boldsymbol{v}_i, w\rangle = \boldsymbol{\kappa}_i$ and $\langle \boldsymbol{v}_j, w\rangle = \boldsymbol{\kappa}_j$, following a similar calculation to the one above,

$$q_{ij} \coloneqq \mathbf{Pr}[i \sim j] = \mathop{\mathbf{Pr}}_{\boldsymbol{u}_i,\boldsymbol{u}_j\sim\mathbb{S}^{d-2}}\left[\langle \boldsymbol{u}_i, \boldsymbol{u}_j\rangle \geqslant \frac{\tau - \boldsymbol{\kappa}_i\boldsymbol{\kappa}_j}{\sqrt{(1-\boldsymbol{\kappa}_i^2)(1-\boldsymbol{\kappa}_j^2)}}\right] \tag{5}$$

And this quantity is $\gg q = \mathbf{Pr}_{\boldsymbol{u}_i,\boldsymbol{u}_j\sim\mathbb{S}^{d-2}}[\langle \boldsymbol{u}_i, \boldsymbol{u}_j\rangle \geqslant \frac{\tau}{1+\tau}]$ when $\boldsymbol{\kappa}_i > (1+\delta)\tau$ and $\boldsymbol{\kappa}_j \geqslant \tau$. Hence, vertex degrees are not as well concentrated within each link as they are (around pn) in the entire graph $\boldsymbol{H}$.

As a result, if we let $\boldsymbol{G}_w$ now stand for the link and $A_{\boldsymbol{G}_w}$ now stand for the adjacency matrix of the link, it is no longer the case that $\|A_{\boldsymbol{G}_w} - \mathbf{E}A_{\boldsymbol{G}_w}\|$ is small: $\mathbf{E}A_{\boldsymbol{G}_w}$ still has every entry equal to q, but the top eigenvector of $A_{\boldsymbol{G}}$ will not be close to the all-1 vector.

To contend with this, we analyze the spectral norm of $A_{\boldsymbol{G}}$ *conditioned* on the shells that the points in $\mathrm{cap}_p(w)$ are in. Letting $\boldsymbol{\kappa} \in [\tau, 1]^m$ be such that $\boldsymbol{\kappa}_i = \langle \boldsymbol{v}_i, w\rangle$, vertex degrees concentrate in $\boldsymbol{G}_w$ conditioned on $\boldsymbol{\kappa}$, and we can readily bound the spectral norm of $\overline{A}_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa} = A_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa} - \mathbf{E}[A_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa}]$.

The analysis of the spectral norm of $\overline{A}_{\boldsymbol{G}_w}$ is then not so different from that of $\overline{A}_{\boldsymbol{G}}$ for $\boldsymbol{G}$ a random geometric graph; the main difference is that now, instead of working with the walk P in which we walk from $\boldsymbol{u}_i$ to a random point in $\mathrm{cap}_{\geqslant\tau}(\boldsymbol{u}_i)$, at each step of the walk we must adjust the volume of the cap: when considering the probability that the edge i,j is present, we apply the operator $P_{q_{ij}}$ for $q_{ij}(\boldsymbol{\kappa}_i, \boldsymbol{\kappa}_j)$ as defined in (5), which walks from $\boldsymbol{u}_i$ to a random point in $\mathrm{cap}_{q_{ij}}(\boldsymbol{u}_i)$. This requires some additional accounting, but one can show that the slowest mixing occurs when $\boldsymbol{\kappa}_i = \boldsymbol{\kappa}_j = \tau$ and $q_{ij} = \frac{\tau}{1+\tau}$, from which we obtain the desired bound on $\|\overline{A}_{\boldsymbol{G}_w} \mid \kappa\|$. For details, see Section 5.

One additional complication is that $\mathbf{E}A_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa}$ is not a rank-1 matrix, so bounding $\|\overline{A}_{\boldsymbol{G}_w} \mid \kappa\|$ does not directly imply a bound on the second eigenvalue of $A_{\boldsymbol{G}_w}$. However, it turns out that $\mathbf{E}A_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa}$

is sufficiently close to a rank-1 matrix $R_{\boldsymbol{G}_w}$ (the matrix whose (i,j)th entry is the product of the expected degrees conditioned on $\boldsymbol{\kappa}$) that we can apply the triangle inequality:

$$\left\|(A_{\boldsymbol{G}_w} - R_{\boldsymbol{G}_w}) \mid \boldsymbol{\kappa}\right\| \leqslant \left\|A_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa} - \mathbf{E}[A_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa}]\right\| + \left\|\mathbf{E}[A_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa}] - R_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa}\right\|,$$

the first term we bound using the trace method as described above. The second term we bound via more-or-less direct calculation: because all but an $o(1)$ fraction of $\boldsymbol{\kappa}_i \approx \tau$, when ignoring an $o(1)$ fraction of rows and columns, the rows of $\mathbf{E}\, A_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa}$ are almost constant multiples of each other, and further these $o(1)$ fraction of rows and columns represent an $o(1)$ fraction of the total absolute value of $\mathbf{E}[A_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa}]$. (This is because the high-degree vertices in $\boldsymbol{G}_w$ represent an $o(1)$ fraction of the total edges in $\boldsymbol{G}_w$.) Now, thinking of $\mathbf{E}[A_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa}]$ as a transition operator of a Markov chain, we are able to use this to argue that the Markov chain mixes so rapidly that $\mathbf{E}[A_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa}]$ must be close to $R_{\boldsymbol{G}_w} \mid \boldsymbol{\kappa}$, yielding the desired bound. For details, see Section 6.

Organization

In Section 2 we give some technical preliminaries. In Section 3 we use the trace method prove the spectral norm bound for random restrictions of arbitrary graphs, Theorem 1.10. To apply Theorem 1.10 to bound the spectrum of $\mathsf{Geo}_d(n,p)$ (and also to ultimately prove that $\mathsf{Geo}_d^{(2)}(n,p)$ is a 2-dimensional expander), we must prove the total variation decay condition for a random walk on the sphere with steps consisting of jumps to a random point in a spherical cap. We do this in Section 4 by relating this walk with discrete jumps to Brownian motion on $\mathbb{S}^{d-1}$. The links of vertices in $\mathsf{Geo}_d^{(2)}(n,p)$ do not conform to the requirements of Theorem 1.10 because they are random restrictions of graphs which are not vertex-transitive, and so in Section 5 and Section 6 we prove a version of Theorem 1.10 specialized to these links; Section 5 contains the trace method and Section 6 addresses the fact that the top eigenvector is not proportional to $\vec{1}$. Finally we put the pieces together in Section 7 to prove Theorem 1.6. In Section 8, we show that the trickling-down theorem is tight.

2 Preliminaries

Notation. For a self-adjoint matrix M, we denote its eigenvalues in decreasing order as $\lambda_1(M) \geqslant \ldots \geqslant \lambda_n(M)$, the absolute values of its eigenvalues as $|\lambda|_1(M) \geqslant \ldots \geqslant |\lambda|_n(M)$, and $\lambda_{\max}(M)$ and $|\lambda|_{\max}(M)$ to denote $\lambda_1(M)$ and $|\lambda|_1(M)$ respectively. Given a sequence of matrices $M_1,\ldots,M_T$ we use $\prod_{i=1}^{T} M_i$ to denote the matrix $M_T \cdot M_{T-1} \cdots M_1$.

For a graph G, we use $V(G)$ to refer to its vertex set and $E(G)$ to refer to its edge set. For a vertex $v \in V(G)$, we use $N(v)$ to denote the set of neighbors of v.

For a probability distribution $\mathscr{D}$, we use $\Phi_{\mathscr{D}}(x)$ to denote the CDF of $\mathscr{D}$ at x, and $\overline{\Phi}_{\mathscr{D}}(x) \coloneqq 1 - \Phi_{\mathscr{D}}(x)$ to denote the tail of $\mathscr{D}$ at x. For any point x, we use δ_x to denote the delta distribution at x .

2.1 Linear algebra

The following articulates how one gets a handle on the second eigenvalue of a matrix after subtracting a rank-1 term, which will be used in Section 3 and Section 5.

Fact 2.1. *For any $n \times n$ symmetric matrix M and rank-1 PSD matrix R, $|\lambda|_2(M) \leqslant \|M - R\|$.*

Proof. By Cauchy's interlacing theorem, $\lambda_2(M) \leqslant \lambda_1(M-R) \leqslant \|M-R\|$ and $-\lambda_n(M) \leqslant -\lambda_n(M-R) \leqslant \|M-R\|$. The desired inequality is then true since $|\lambda|_2(M) \leqslant \max\{\lambda_2(M), -\lambda_n(M)\}$. □

Establishing second eigenvalue bounds in Section 3 and Section 5 also involves bounding the spectral norm of some matrices via the "trace method" articulated below.

Claim 2.2 (Trace Method)**.** Let $\boldsymbol{M}$ be a symmetric (random) matrix. Then for any even integer $\ell \geqslant 0$,

$$\mathbf{Pr}\left[\|\boldsymbol{M}\| \geqslant e^{\varepsilon} \cdot \mathbf{E}\left[\operatorname{tr}\left((\boldsymbol{M})^{\ell}\right)\right]^{1/\ell}\right] \leqslant \exp(-\varepsilon\ell).$$

Proof. By Markov's inequality, $\mathbf{Pr}[\|\boldsymbol{M}\| \geqslant t] \leqslant t^{-\ell}\mathbf{E}\left(\|\boldsymbol{M}\|^{\ell}\right)$. The claim then follows because for any self-adjoint matrix M, $\lambda_{\max}\left(M^{\ell}\right) \leqslant \operatorname{tr}(M^{\ell})$ when ℓ is even. □

We will also require the following bound on the spectrum of a matrix, which is a special case of the Gershgorin circle theorem.

Claim 2.3 (Row sum bound)**.** For any matrix M, $|\lambda|_{\max}(M) \leqslant \max_i \|M[i,*]\|_1$.

Proof. Let v be the eigenvector achieving $\lambda = |\lambda|_{\max}(M)$. Then letting k be the index maximizing $|v_k|$, we have

$$|\lambda v_k| = |(Mv)_k| = \left|\sum_j M_{kj} v_j\right| \leqslant |v_k| \sum_j |M_{kj}| \leqslant |v_k| \max_i \|M[i,*]\|_1,$$

and dividing through by $|v_k|$ gives the conclusion. □

2.2 Probability

Definition 2.4. The *total variation distance* between probability distributions μ and ν is defined as:

$$\mathrm{d}_{\mathrm{TV}}(\mu,\nu) := \max_{\mathscr{E}} \left|\mu(\mathscr{E}) - \nu(\mathscr{E})\right|.$$

Fact 2.5. *When ρ is a nonnegative measure such that μ and ν are absolutely continuous with respect to ρ, then:*

$$\mathrm{d}_{\mathrm{TV}}(\mu,\nu) = \frac{1}{2}\int \left|\frac{d\mu}{d\rho}(x) - \frac{d\nu}{d\rho}(x)\right| d\rho(x) = \int \left(\frac{d\mu}{d\rho}(x) - \frac{d\nu}{d\rho}(x)\right) \cdot \mathbf{1}\left[\frac{d\mu}{d\rho}(x) > \frac{d\nu}{d\rho}(x)\right] d\rho(x).$$

When μ and ν are supported on $[n]$, then:

$$\mathrm{d}_{\mathrm{TV}}(\mu,\nu) = \frac{1}{2}\|\mu-\nu\|_1 = \sum_{i=1}^{n} (\mu(i)-\nu(i)) \cdot \mathbf{1}[\mu(i) > \nu(i)]$$

where μ and ν are the vectors of probabilities.

We describe a Markov chain via its transition operator P where $P(i,j)$ denotes the probability of transitioning from state i to state j.

We call the joint distribution $\omega(\mu,\nu)$ a coupling between two distributions μ and ν if $\mu = \omega(\cdot,\nu)$ and $\mu = \omega(\mu,\cdot)$. In other words, the marginals of ω correspond to μ and ν.

Fact 2.6. *Let x and y be two arbitrary states in a Markov chain over state space Ω with transition operator P, and sample $X \sim P(x,\cdot)$ and $Y \sim P(y,\cdot)$, where $P(z,\cdot)$ denotes the distribution over Ω given by a single step of the walk starting from state z. Then, there exists a coupling of X and Y such that $X = Y$ with probability $1-\varepsilon$ if and only if $\mathrm{d}_{\mathrm{TV}}(P(x,\cdot),P(y,\cdot)) \leqslant \varepsilon$.*

2.3 The uniform distribution over the unit sphere

We use ρ to denote the uniform distribution on $\mathbb{S}^{d-1}$.

Let $v \in \mathbb{S}^{d-1}$ and $\boldsymbol{w} \sim \rho$. Then the distribution $\mathsf{D}_{\mathsf{ip}}(d)$ of $\langle \boldsymbol{w}, v\rangle$ is invariant under the choice of v, is supported on $[-1,1]$ and has probability density function:

$$\psi_d(x) = \frac{\Gamma\left(\frac{d}{2}\right)}{\Gamma\left(\frac{d-1}{2}\right)\sqrt{\pi}} \cdot \left(1-x^2\right)^{(d-3)/2}.$$

Henceforth, we use Z_d to denote the normalizing constant $\frac{\Gamma\left(\frac{d}{2}\right)}{\Gamma\left(\frac{d-1}{2}\right)\sqrt{\pi}}$.

Fact 2.7. $Z_d \leqslant O(\sqrt{d})$.

In addition, we will rely heavily on the following sharp estimate of the tail of $\mathsf{D}_{\mathsf{ip}}(d)$.

Lemma 2.8. *Let $\Phi_{\mathsf{D}_{\mathsf{ip}}(d)}(t) \coloneqq \mathbf{Pr}_{X\sim\mathsf{D}_{\mathsf{ip}}(d)}[X \geqslant t]$. Then, when $t \geqslant 0$:*

$$\frac{Z_d}{t(d-1)} \cdot \left(1-t^2\right)^{(d-1)/2} \cdot \left(1 - \frac{4\log\left(1+d\cdot t^2\right)}{d\cdot t^2}\right) \leqslant \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}(t) \leqslant \frac{Z_d}{t(d-1)} \cdot \left(1-t^2\right)^{(d-1)/2}.$$

Proof. It suffices to upper and lower bound $\int_t^1 (1-x^2)^{(d-3)/2}$. We first obtain an upper bound.

$$\begin{aligned}
\int_t^1 \left(1-x^2\right)^{(d-3)/2} dx &= \frac{1}{t}\int_t^1 t\left(1-x^2\right)^{(d-3)/2} dx \\
&\leqslant \frac{1}{t}\int_t^1 x\left(1-x^2\right)^{(d-3)/2} dx \\
&= -\frac{1}{t(d-1)} \cdot \left(1-x^2\right)^{(d-1)/2}\Big|_t^1 \\
&= \frac{1}{t(d-1)} \cdot \left(1-t^2\right)^{(d-1)/2}
\end{aligned}$$

Now we prove the lower bound. For any $\varepsilon > 0$ such that $t\cdot\sqrt{1-\varepsilon+\frac{\varepsilon}{t^2}} \leqslant 1$, and defining $\delta \coloneqq \frac{\varepsilon}{t^2} - \varepsilon$, we have the following.

$$\begin{aligned}
\int_t^1 (1-x^2)^{(d-3)/2} dx &\geqslant \frac{1}{t\sqrt{1+\delta}} \int_t^{t\sqrt{1+\delta}} \left(t\sqrt{1+\delta}\right)\left(1-x^2\right)^{(d-3)/2} dx \\
&\geqslant \frac{1-\delta}{t}\int_t^{t\sqrt{1+\delta}} x\left(1-x^2\right)^{(d-3)/2} dx \\
&= -\frac{1-\delta}{t(d-1)} \cdot \left(1-x^2\right)^{(d-1)/2}\Big|_t^{t\sqrt{1+\delta}} \\
&= \frac{1-\delta}{t(d-1)} \cdot \left(1-t^2\right)^{(d-1)/2} \cdot \left(1-(1-\varepsilon)^{(d-1)/2}\right)
\end{aligned}$$

where the second inequality uses $\frac{1}{\sqrt{1+\delta}} \geqslant 1-\delta$ and the last equality uses $1-t^2(1+\delta) = (1-t^2)(1-\varepsilon)$. Choosing $\varepsilon = \frac{2\log(1+dt^2)}{d-1}$ yields:

$$\int_t^1 \left(1-x^2\right)^{(d-3)/2} \geqslant \frac{1}{t(d-1)} \cdot \left(1-t^2\right)^{(d-1)/2} \cdot \left(1 - \frac{4\log\left(1+dt^2\right)}{dt^2}\right). \qquad \square$$

We use $\mathsf{D}_{\mathsf{ip}}(d)|_{\geqslant\tau}$ to represent $\mathsf{D}_{\mathsf{ip}}(d)$ conditioned on lying in $[\tau,1]$.

Definition 2.9. For a vector y, we use $\mathrm{cap}_p(y)$ and $\mathrm{cap}_{\geqslant\tau(p)}(y)$ interchangeably to denote the measure-p spherical cap around y:

$$\mathrm{cap}_p(y) = \mathrm{cap}_{\geqslant\tau(p)}(y) := \left\{u : \langle u, y\rangle \geqslant \tau(p), u \in \mathbb{S}^{d-1}\right\}.$$

We use $\underline{\mathrm{cap}}_p(y)$ and $\underline{\mathrm{cap}}_{\geqslant\tau(p)}(y)$ to denote the uniform measure over the set $\mathrm{cap}_p(y)$. We denote the boundary of $\mathrm{cap}_p(y)$ by $\mathrm{shell}_p(y)$ or $\mathrm{shell}_{=\tau(p)}(y)$. That is,

$$\mathrm{shell}_p(y) := \left\{u : \langle u, y\rangle = \tau(p), u \in \mathbb{S}^{d-1}\right\}.$$

3 The second eigenvalue of random restrictions

In this section we prove Theorem 1.10. Let X be a (possibly infinite) vertex-transitive graph with a unique stationary measure ρ. Let $\boldsymbol{G} \sim \mathsf{RR}_n(X)$ be a random restriction of X as defined in Definition 1.9, and let $p = \mathbf{Pr}_{\boldsymbol{G}\sim\mathsf{RR}_n(X)}[(i,j) \in E(\boldsymbol{G})]$ be the marginal edge probability in $\boldsymbol{G}$. Suppose furthermore that

$$\exists C, \lambda \text{ with } C \geqslant 1 \text{ and } \frac{1}{\sqrt{pn}} \leqslant \lambda \leqslant 1 \quad \text{s.t. for any distribution } \alpha \text{ on } V(X), \quad \mathrm{d_{TV}}\left(X^k\alpha, \rho\right) \leqslant C\lambda^k. \tag{6}$$

We overload notation and use X to denote the transition operator for the simple random walk on X, and for $H \subseteq V(X)$ we also use H to denote the indicator vector of the set H.

We denote its adjacency matrix by $A_{\boldsymbol{G}}$, the diagonal degree matrix by $D_{\boldsymbol{G}}$, the centered adjacency matrix by $\overline{A}_{\boldsymbol{G}} = A_{\boldsymbol{G}} - \mathbf{E}A_{\boldsymbol{G}}$, and the normalized adjacency matrix by $\widehat{A}_{\boldsymbol{G}} = D_{\boldsymbol{G}}^{-1/2}A_{\boldsymbol{G}}D_{\boldsymbol{G}}^{-1/2}$. Then we'll show the following.

Theorem 3.1. *As long as $pn \gg C^6\log^8 n$, for any constant $\gamma > 0$, with probability at least $1 - n^{-\gamma}$,*

$$|\lambda|_2\left(\widehat{A}_{\boldsymbol{G}}\right) \leqslant (1+o(1))\cdot\max\left(\lambda, \frac{\log^4 n}{\sqrt{pn}}\right).$$

Proof. By Fact 2.1, for any rank-1 PSD matrix R, $|\lambda|_2\left(\widehat{A}_{\boldsymbol{G}}\right) \leqslant \left\|\widehat{A}_{\boldsymbol{G}} - R\right\|$. Thus we turn our attention to bounding $\left\|\widehat{A}_{\boldsymbol{G}} - R\right\|$ for appropriately chosen R. Setting $R_{\boldsymbol{G}} = pD_{\boldsymbol{G}}^{-1/2}JD_{\boldsymbol{G}}^{-1/2}$ where J is the all-ones matrix and using submultiplicativity of the operator norm, we see:

$$\left\|\widehat{A}_{\boldsymbol{G}} - R_{\boldsymbol{G}}\right\| \leqslant \left\|D_{\boldsymbol{G}}^{-1/2}\right\|^2 \cdot \|A_{\boldsymbol{G}} - pJ\|.$$

Now, observe that $\left\|D_{\boldsymbol{G}}^{-1/2}\right\|^2 = \left\|D_{\boldsymbol{G}}^{-1}\right\|$. To bound this quantity, we'll use the concentration of the vertex degrees (the entries of the diagonal of $D_{\boldsymbol{G}}$). For every vertex, the marginal distribution of the degree is $\mathsf{Binom}(n,p)$. So by Hoeffding's inequality and the union bound, when $pn \gg \log^8 n$, for any fixed $\gamma > 0$, $|(D_{\boldsymbol{G}})_{ii} - pn| \leqslant \sqrt{pn\log^2 n}$ for all $i \in [n]$ with probability at least $1 - n^{\gamma}$. So with probability at least $1 - n^{-\gamma}$, $D_{\boldsymbol{G}}^{-1} = \frac{1}{pn}I + \Delta$ for Δ a diagonal matrix with entries with absolute value of order $\sqrt{\log^2 n/(pn)^3}$. Thus, $\left\|D_{\boldsymbol{G}}^{-1}\right\| \leqslant \frac{1}{pn}\cdot\left(1 + \frac{\log n}{\sqrt{pn}}\right)$.

Next, $\|A_{\boldsymbol{G}} - pJ\| \leqslant \left\|\overline{A}_{\boldsymbol{G}}\right\| + p$, where recall $\overline{A}_{\boldsymbol{G}} = A_{\boldsymbol{G}} - \mathbf{E}A_{\boldsymbol{G}}$. We will show:

$$\left\|\overline{A}_{\boldsymbol{G}}\right\| \leqslant (1+o(1))\cdot\max\{\lambda pn, \sqrt{pn}\log^4 n\}.$$

Putting these bounds together gives:

$$|\lambda|_2\left(\widehat{A}_{\boldsymbol{G}}\right) \le (1+o(1))\cdot\max\left\{\lambda, \frac{\log^4 n}{\sqrt{pn}}\right\}.$$

Finally, we devote the rest of the proof to bounding $\left\|\overline{A}_{\boldsymbol{G}}\right\|$. By Claim 2.2, it suffices to bound $\mathbf{E}\,\mathrm{tr}((\overline{A}_{\boldsymbol{G}})^\ell)$ for a large enough even ℓ.

For an $n\times n$ matrix M, $\mathrm{tr}(M^\ell)$ can be written as a sum over length-ℓ closed walks on the complete graph $\mathcal{K}_n$, with each walk W weighted according to $\prod_{(i,j)\in W} M_{ij}$. The exchangeability of entries in $\overline{A}_{\boldsymbol{G}}$ means that the walks can be partitioned into equivalence classes based on their topology as graphs, where the members of each class contribute identically to the summation.

Definition 3.2. We use $\mathcal{W}_\ell$ to denote the collection of length-ℓ walks in $\mathcal{K}_n$, the complete graph on n vertices. For $W\in\mathcal{W}_\ell$, we use $G(W)=(V(W),E(W))$ to denote the simple graph induced by edges walked on in W. We let the *multiplicity* of e in W, $m(e)$, be the number of times e occurs in W.

We can then write:

$$\mathbf{E}\,\mathrm{tr}\left((A_{\boldsymbol{G}}-\mathbf{E}A_{\boldsymbol{G}})^\ell\right) = \sum_{W\in\mathcal{W}_\ell}\mathbf{E}\prod_{e\in E(W)}(\mathbf{1}[e\in\boldsymbol{G}]-p)^{m(e)} \tag{7}$$

We now focus on understanding each term of the above summand in terms of the properties of $G(W)$. Our first step is to handle leaves.

Definition 3.3. We use $G_2(W)=(V_2(W),E_2(W))$ to denote the 2-*core* of $G(W)$, the graph obtained by recursively deleting degree-1 vertices from $G(W)$. We denote the graph induced on the edges deleted in this process as $G_1(W)$.

Observation 3.4. We have $G(W)=G_1(W)\cup G_2(W)$. Further, every vertex in $G_2(W)$ has degree at least 2, and $G_1(W)$ is a forest where each connected component has at most one vertex in $G_2(W)$.

Notice that if F is a forest, then $\mathbf{Pr}[F\in\boldsymbol{G}]=p^{|E(F)|}$, and further if F is a forest sharing at most one vertex with a graph H, then the events $\{H\in\boldsymbol{G}\}$ and $\{F\in\boldsymbol{G}\}$ are independent. Hence, with the above decomposition in hand, we can "peel off" the one-core and for any $W\in\mathcal{W}_\ell$ we can write:

$$\begin{aligned}
(7) &= \mathop{\mathbf{E}}_{\boldsymbol{u}_i\,i\in V_2(W)}\ \mathop{\mathbf{E}}_{\boldsymbol{u}_j\,j\in V_1(W)\setminus V_2(W)}\ \prod_{e\in E(W)}(\mathbf{1}[e\in\boldsymbol{G}]-p)^{m(e)}\\
&= \prod_{e\in E_1(W)}\mathbf{E}\left((\mathbf{1}[e\in\boldsymbol{G}]-p)^{m(e)}\right)\cdot\mathop{\mathbf{E}}_{\boldsymbol{u}_i\,i\in V_2(W)}\prod_{e\in E_2(W)}(\mathbf{1}[e\in\boldsymbol{G}]-p)^{m(e)}\\
&= \prod_{e\in E_1(W)}\mathbf{E}\left(\mathbf{1}[e\in\boldsymbol{G}]\left((1-p)^{m(e)}-(-p)^{m(e)}\right)+(-p)^{m(e)}\right)\cdot\mathop{\mathbf{E}}_{\boldsymbol{u}_i\,i\in V_2(W)}\prod_{e\in E_2(W)}(\mathbf{1}[e\in\boldsymbol{G}]-p)^{m(e)}\\
&\le \left|\prod_{e\in E_1(W)}\left(p(1-p)^{m(e)}+(1-p)(-p)^{m(e)}\right)\right|\cdot\left|\mathop{\mathbf{E}}_{\boldsymbol{u}_i:i\in V_2(W)}\prod_{e\in E_2(W)}(\mathbf{1}[e\in\boldsymbol{G}]-p)^{m(e)}\right|,
\end{aligned}\tag{8}$$

where in the third line we've used that $(\mathbf{1}[e\in\boldsymbol{G}]-p)^k=\mathbf{1}[e\in\boldsymbol{G}]((1-p)^k-(-p)^k)+(-p)^k$.

It now remains to handle the 2-core $G_2(W)$. To simplify the expression, we'll exploit the following fact: if J is a subset of vertices in $G_2(W)$, conditional on an assignment of $\boldsymbol{u}_i$ for all $i\in J$, the existence of edges in regions of $G_2(W)$ separated by J are independent. We'll take advantage of this fact by splitting $G_2(W)$ into regions separated by the set of vertices in $G_2(W)$ of degree at least 3, leaving us to bound a collection of paths and cycles.

Definition 3.5 (Junction vertices)**.** We use $J(W)$ to denote the set of *junction vertices* of $G_2(W)$, which are vertices with degree-$\geqslant 3$ in $G_2(W)$, or in the case that $G_2(W)$ only has vertices of degree-2, we choose an arbitrary vertex ot $G_2(W)$ and add it to $J(W)$. We use $G_J(W) = (J(W), E_J(W))$ to denote the *junction graph* of $G_2(W)$, which is a multigraph obtained by starting with $G_2(W)$ and contracting to an edge all walks $\gamma = u_0 \dots u_t$ satisfying the following conditions:

1. u_0 and u_t are (possibly identical) junction vertices,
2. $u_1, \dots, u_{t-1}$ are distinct vertices with degree-2 in $G_2(W)$.

For an edge $f \in E_J(W)$, we use $\gamma(f) = u_0, \dots, u_t$ to identify the walk from which f arose in $G_2(W)$, $s(f)$ to denote the "start" vertex u_0 of $\gamma(f)$, and $t(f)$ to denote the "terminal" vertex u_t of $\gamma(f)$.

Then we can bound the contribution of the 2-core in terms of the contribution of the walk $\gamma(f)$ corresponding to each edge f in the junction graph:

$$\begin{aligned}
\left| \mathop{\mathbf{E}}_{\substack{\boldsymbol{u}_i \\ i \in V_2(W)}} \prod_{e \in E_2(W)} (\mathbf{1}[e \in \boldsymbol{G}] - p)^{m(e)} \right| &= \left| \mathop{\mathbf{E}}_{\boldsymbol{u}_i : i \in J(W)} \mathop{\mathbf{E}}_{\boldsymbol{u}_i : i \notin J(W)} \prod_{e \in E_2(W)} (\mathbf{1}[e \in \boldsymbol{G}] - p)^{m(e)} \right| \\
&= \left| \mathop{\mathbf{E}}_{\boldsymbol{u}_i : i \in J(W)} \prod_{f \in E_{J(W)}} \mathop{\mathbf{E}}_{\boldsymbol{u}_i : i \in \gamma(f) \setminus J(W)} \prod_{e \in \gamma(f)} (\mathbf{1}[e \in \boldsymbol{G}] - p)^{m(e)} \right| \\
&\leqslant \mathop{\mathbf{E}}_{\boldsymbol{u}_i : i \in J(W)} \prod_{f \in E_{J(W)}} \left| \mathop{\mathbf{E}}_{\boldsymbol{u}_i : i \in \gamma(f) \setminus J(W)} \prod_{e \in \gamma(f)} (\mathbf{1}[e \in \boldsymbol{G}] - p)^{m(e)} \right|. \qquad (9)
\end{aligned}$$

We now focus on understanding the innermost expected value, the expectation over the internal vertices along a path, conditioned on the endpoints. Again using $(\mathbf{1}[e \in \boldsymbol{G}] - p)^k = \mathbf{1}[e \in \boldsymbol{G}]((1-p)^k - (-p)^k) + (-p)^k$,

$$\begin{aligned}
&\left| \mathop{\mathbf{E}}_{\boldsymbol{u}_i : i \in \gamma(f) \setminus J(W)} \prod_{e \in \gamma(f)} (\mathbf{1}[e \in \boldsymbol{G}] - p)^{m(e)} \right| \\
&\quad = \left| \mathop{\mathbf{E}}_{\boldsymbol{u}_i : i \in \gamma(f) \setminus J(W)} \prod_{e \in \gamma(f)} \left(\mathbf{1}[e \in \boldsymbol{G}] \cdot \left((1-p)^{m(e)} - (-p)^{m(e)} \right) + (-p)^{m(e)} \right) \right|, \\
&\quad = \left| \sum_{T \subseteq \gamma(f)} \mathop{\mathbf{E}}_{\boldsymbol{u}_i : i \in \gamma(f) \setminus J(W)} \prod_{e \in T} \mathbf{1}[e \in \boldsymbol{G}] \cdot \left((1-p)^{m(e)} - (-p)^{m(e)} \right) \prod_{e \in \gamma(f) \setminus T} (-p)^{m(e)} \right|
\end{aligned}$$

Now, using the independence of edges in a forest, we can bound terms where $T \neq \gamma(f)$ simply, and the term $T = \gamma(f)$ in terms of the probability that a $|\gamma(f)|$-length walk in X starting at $\boldsymbol{u}_{s(f)}$ ends at $\boldsymbol{u}_{t(f)}$ (which is where properties of the random walk in X will enter into the bound):

$$\begin{aligned}
&= \left| \sum_{\substack{T \subseteq \gamma(f) \\ T \neq \gamma(f)}} \prod_{e \in T} p \cdot \left((1-p)^{m(e)} - (-p)^{m(e)} \right) \cdot \prod_{e \in \gamma(f) \setminus T} (-p)^{m(e)} \right. \\
&\qquad \left. + \prod_{e \in \gamma(f)} \left((1-p)^{m(e)} - (-p)^{m(e)} \right) \cdot p^{|\gamma(f)|-1} \cdot \left\langle N(\boldsymbol{u}_{s(f)}), X^{|\gamma(f)|-1} \delta_{\boldsymbol{u}_{t(f)}} \right\rangle \right|,
\end{aligned}$$

where $N(\boldsymbol{u}_{s(f)})$ is the neighborhood of $\boldsymbol{u}_{s(f)}$ in X, and $\delta_{\boldsymbol{u}_{t(f)}}$ is the point mass at $\boldsymbol{u}_{t(f)}$. Now adding and subtracting $\prod_{e \in \gamma(f)} \left(\left((1-p)^{m(e)} - (-p)^{m(e)} \right) + (-p)^{m(e)} \right) \cdot p^{|\gamma(f)|}$, we complete the first summation and from the triangle inequality we obtain the bound

$$\leqslant \left| \prod_{e \in \gamma(f)} \left(p(1-p)^{m(e)} + (1-p)(-p)^{m(e)} \right) \right|$$

$$+\left|\prod_{e\in\gamma(f)}\left((1-p)^{m(e)}-(-p)^{m(e)}\right)\cdot p^{|\gamma(f)|-1}\cdot\left(\left\langle N(\boldsymbol{u}_{s(f)}),X^{|\gamma(f)|-1}\delta_{\boldsymbol{u}_{t(f)}}\right\rangle-p\right)\right|. \tag{10}$$

We bound (10) based on the graphical properties of $\gamma(f)$.

Definition 3.6. We say an edge e is a *singleton edge* if $m(e)=1$ and a *duplicative edge* otherwise.

If $\gamma(f)$ contains any singleton edges, then the first term of (10) is 0; otherwise it is bounded by

$$\prod_{e\in\gamma(f)}(p(1-p)^2+(1-p)p^2)\leqslant\prod_{e\in\gamma(f)}p(1-p)\leqslant p^{|\gamma(f)|}.$$

The second term can always be bounded by

$$\prod_{e\in\gamma(f)}\left((1-p)^{m(e)}+p^{m(e)}\right)\cdot p^{|\gamma(f)|-1}\cdot\left|\left\langle N(\boldsymbol{u}_{s(f)}),X^{|\gamma(f)|-1}\delta_{\boldsymbol{u}_{t(f)}}\right\rangle-p\right|\leqslant p^{|\gamma(f)|-1}\cdot\left|\left\langle N(\boldsymbol{u}_{s(f)}),X^{|\gamma(f)|-1}\delta_{\boldsymbol{u}_{t(f)}}\right\rangle-p\right|.$$

Using $D_J(W)$ to denote the collection of edges f in G_J such that $\gamma(f)$ contains no singleton edges, and $S_J(W)$ to use the collection of edges f in G_J such that $\gamma(f)$ contains a singleton edge, and plugging the above bounds into (9) tells us:

$$\begin{aligned}(9)\leqslant\mathop{\mathbf{E}}_{\substack{\boldsymbol{u}_i\\ i\in J(W)}}&\prod_{f\in D_J(W)}p^{|\gamma(f)|-1}\cdot\left(\left|\left\langle N(\boldsymbol{u}_{s(f)}),X^{|\gamma(f)|-1}\delta_{\boldsymbol{u}_{t(f)}}\right\rangle-p\right|+p\right)\cdot\\&\prod_{f\in S_J(W)}p^{|\gamma(f)|-1}\cdot\left|\left\langle N(\boldsymbol{u}_{s(f)}),X^{|\gamma(f)|-1}\delta_{\boldsymbol{u}_{t(f)}}\right\rangle-p\right|.\end{aligned}$$

If $G_J(W)$ were a tree, we could recursively take the expectation over leaf vertices to bound the quantity above, as we did to get rid of G_1. However, it is not a tree, so we'll pick an arbitrary spanning tree $T_J(W)$ of $G_J(W)$, and bound edges outside of the spanning tree directly. For $f\in E_J(W)\setminus T_J(W)$, we use Assumption 6 to conclude that $\mathrm{d_{TV}}\left(X^{|\gamma(f)|-1}\delta_{\boldsymbol{u}_{t(f)}},\rho\right)\leqslant C\lambda^{|\gamma(f)|-1}$, which thus implies that

$$\left|\left\langle N(\boldsymbol{u}_{s(f)}),X^{|\gamma(f)|-1}\delta_{\boldsymbol{u}_{t(f)}}\right\rangle-p\right|\leqslant C\lambda^{|\gamma(f)|-1}, \tag{11}$$

because $\langle N(\boldsymbol{u}_{s(f)}),X^{|\gamma(f)|-1}\delta_{\boldsymbol{u}_{t(f)}}\rangle$ represents the probability that a point sampled at random from the measure $X^{|\gamma(f)|-1}\delta_{\boldsymbol{u}_{t(f)}}$ lands in $N(\boldsymbol{u}_{s(f)})$, which is a set of measure p under ρ. We now prove the following by induction.

Claim 3.7. We have the following bound on the contribution of $f\in T_J(W)$:

$$\begin{aligned}\mathop{\mathbf{E}}_{\boldsymbol{u}_i:i\in J(W)}\prod_{f\in T_J(W)}p^{|\gamma(f)|-1}\cdot\left(\left|\left\langle N(\boldsymbol{u}_{s(f)}),X^{|\gamma(f)|-1}\delta_{\boldsymbol{u}_{t(f)}}\right\rangle-p\right|+p\cdot\mathbf{1}[f\in D_J(W)]\right)\\\leqslant\prod_{f\in T_J(W)}p^{|\gamma(f)|}\cdot\left(2C\lambda^{|\gamma(f)|}+\mathbf{1}[f\in D_J(W)]\right).\end{aligned}$$

Proof. We fix an order for $i\in J(W)$, $i_0,\ldots,i_t$ such that i_j is a leaf in $T_J^{(j)}(W)$, the graph obtained by taking $T_J(W)$ and deleting $i_{j+1},\ldots,i_t$. We use f_j to denote the unique edge incident to i_j in $T_J^{(j)}(W)$. Then if we define

$$a_j\coloneqq\mathop{\mathbf{E}}_{\boldsymbol{u}_i:i\in V(T_J^{(j)}(W))}\prod_{f\in T_J^{(j)}(W)}p^{|\gamma(f)|-1}\cdot\left(\left|\left\langle N(\boldsymbol{u}_{s(f)}),X^{|\gamma(f)|-1}\delta_{\boldsymbol{u}_{t(f)}}\right\rangle-p\right|+p\cdot\mathbf{1}[f\in D_J(W)]\right)$$

Because f_j is independent of $f_{j'}$ for $j' < j$ we can write:

$$a_j := \mathop{\mathbf{E}}_{\boldsymbol{u}_{i_0}} \cdots \mathop{\mathbf{E}}_{\boldsymbol{u}_{i_{j-1}}} \prod_{f \in T_J^{(j-1)}(W)} p^{|\gamma(f)|-1} \cdot \left(\left| \left\langle N(\boldsymbol{u}_{s(f)}), X^{|\gamma(f)|-1} \delta_{\boldsymbol{u}_{t(f)}} \right\rangle - p \right| + p \cdot \mathbf{1}[f \in D_J(W)] \right) \cdot$$
$$\mathop{\mathbf{E}}_{\boldsymbol{u}_{i_j}} p^{|\gamma(f_j)|-1} \cdot \left(\left| \left\langle N(\boldsymbol{u}_{s(f_j)}), X^{|\gamma(f_j)|-1} \delta_{\boldsymbol{u}_{t(f_j)}} \right\rangle - p \right| + p \cdot \mathbf{1}[f_j \in D_J(W)] \right)$$

Without loss of generality we can assume $i_j = t(f_j)$, and because $N(\boldsymbol{u}_{s(f_j)}) = pX\delta_{\boldsymbol{u}_{s(f_j)}}$,

$$\begin{aligned}
\mathop{\mathbf{E}}_{\boldsymbol{u}_{i_j}} \left| \left\langle N(\boldsymbol{u}_{s(f_j)}), X^{|\gamma(f_j)|-1} \delta_{\boldsymbol{u}_{t(f_j)}} \right\rangle - p \right| &= \mathop{\mathbf{E}}_{\boldsymbol{u}_{i_j}} \left| \left\langle N(\boldsymbol{u}_{s(f_j)}), X^{|\gamma(f_j)|-1} \delta_{\boldsymbol{u}_{i_j}} \right\rangle - p \right| \\
&= p \mathop{\mathbf{E}}_{\boldsymbol{u}_{i_j}} \left| \left\langle X^{|\gamma(f_j)|} \delta_{\boldsymbol{u}_{s(f_j)}}, \delta_{\boldsymbol{u}_{i_j}} \right\rangle - 1 \right| \\
&= 2p \cdot \mathrm{d_{TV}} \left(X^{|\gamma(f_j)|} \delta_{\boldsymbol{u}_{s(f_j)}}, \rho \right) \\
&\leqslant p \cdot 2C\lambda^{|\gamma(f_j)|}.
\end{aligned}$$

This gives us the inequality:

$$\alpha_j \leqslant \alpha_{j-1} \cdot p \left(2C\lambda^{|\gamma(f_j)|} + \mathbf{1}[f_j \in D_J(W)] \right).$$

The above inequality combined with the fact that $\alpha_0 = 1$ yields the claim. □

We use $e(W)$ to denote $|E(W)|$ and $\mathrm{sing}(W)$ to denote the number of singleton edges in $G_2(W)$,[9]. For any graph H we use $\mathrm{exc}(H)$ to denote the *excess* of H, which is $|E(H)| - |V(H)| + 1$, the number of edges H has over a tree.

Observation 3.8. $\mathrm{exc}(G(W)) = \mathrm{exc}(G_2(W)) = \mathrm{exc}(G_J(W))$. Thus, we denote this quantity as $\mathrm{exc}(W)$.

Observation 3.9. $|E_J(W)| \leqslant 3\mathrm{exc}(W)$.

Proof. We use Observation 3.8 to write:

$$2\mathrm{exc}(W) - 2 = 2|E_J(W)| - 2|V_J(W)| = \sum_{v \in V_J} (\deg_{G(J)}(v) - 2) \geqslant |V_J(W)| - 1,$$

where the degree a self-loop incurs on a vertex is 2, and the -1 on the right-hand side is to capture the possibility that $|J(W)| = 1$ when $G_2(W)$ has no degree-3 vertices. Adding $\mathrm{exc}(W)$ to both sides gives:

$$3\mathrm{exc}(W) \geqslant |E_J(W)|.$$ □

Using the bound on the non-tree edges from (11) and Claim 3.7, we get:

$$(9) \leqslant \prod_{f \in T_J(W)} p^{|\gamma(f)|} \cdot \left(2C\lambda^{|\gamma(f)|} + \mathbf{1}[f \in D_J(W)] \right) \cdot \prod_{f \in E_J(W) \setminus T_J(W)} p^{|\gamma(f)|-1} \cdot \left(C\lambda^{|\gamma(f)|-1} + p \cdot \mathbf{1}[f \in D_J(W)] \right)$$

[9] Note $\mathrm{sing}(W)$ is the same as the number of singleton edges in $G(W)$ since $G_1(W)$ cannot have singleton edges, as it is the multigraph induced by a closed walk of length ℓ.

Now, we bound separately the contribution of singleton and duplicative edges. For each $f \in S_J(W)$, we pull out a factor of $(p\lambda)^{|\gamma(f)|}2C$ if the edge was in the tree, and a factor $(p\lambda)^{|\gamma(f)|-1}C$ if the edge was not in the tree; this fully accounts for the contributions of singleton edges. For each $f \in D_J(W)$, we upper bound its contribution by $p^{|\gamma(f)|}3C$ if the edge was in the tree, and a factor $p^{|\gamma(f)|-1}3C$ otherwise; this is potentially loose because we don't keep the factors of λ, but it is a valid upper bound because $C \geqslant 1$ and $p, \lambda \leqslant 1$. We thus have a factor of p from $|E_2(W)| - \mathrm{exc}(W)$ edges, a factor of λ from $\mathrm{sing}(W) - \mathrm{exc}(W)$ edges, and a factor of at most $3C$ from each edge in $E_J(W)$. Summarizing,

$$\leqslant p^{|E_2(W)|-\mathrm{exc}(W)} \lambda^{\mathrm{sing}(W)-\mathrm{exc}(W)} \cdot (3C)^{|E_J(W)|},$$

and by Observation 3.9, the above is bounded by:

$$\leqslant p^{|E_2(W)|-\mathrm{exc}(W)} \lambda^{\mathrm{sing}(W)-\mathrm{exc}(W)} \cdot (3C)^{3\mathrm{exc}(W)}.$$

Since $m(e) \geqslant 2$ for every edge in $e \in E_1$ (otherwise the walk cannot be closed), by an analysis identical to that of the first term of (10), we have:

$$(8) \leqslant p^{|E_1(W)|} \cdot p^{|E_2(W)|-\mathrm{exc}(W)} \lambda^{\mathrm{sing}(W)-\mathrm{exc}(W)} \cdot (3C)^{3\mathrm{exc}(W)} = p^{e(W)-\mathrm{exc}(W)} \lambda^{\mathrm{sing}(W)} \left(\frac{27C^3}{\lambda}\right)^{\mathrm{exc}(W)}.$$

Finally, we can bound the trace power (7) as follows.

$$\begin{aligned}
(7) &\leqslant \sum_{W \in \mathscr{W}_\ell} p^{e(W)-\mathrm{exc}(W)} \lambda^{\mathrm{sing}(W)} \left(\frac{27C^3}{\lambda}\right)^{\mathrm{exc}(W)} \\
&= \sum_{a=1}^{\ell} \sum_{b=1}^{\ell} \sum_{c=1}^{\ell} \sum_{\substack{W \in \mathscr{W}_\ell \\ e(W)=a, \mathrm{sing}(W)=b, \mathrm{exc}(W)=c}} p^{a-c} \lambda^b \left(\frac{27C^3}{\lambda}\right)^c \\
&= \sum_{a=1}^{\ell} \sum_{b=1}^{\ell} \sum_{c=1}^{\ell} p^{a-c} \lambda^b \left(\frac{27C^3}{\lambda}\right)^c \cdot |\{W \in \mathscr{W}_\ell : e(W) = a, \mathrm{sing}(W) = b, \mathrm{exc}(W) = c\}| \qquad (12)
\end{aligned}$$

To finish bounding the trace power, it remains to count length-ℓ closed walks with a specified number of edges, excess edges, and singleton edges.

Claim 3.10. The number of walks W such that $e(W) = a$, $\mathrm{sing}(W) = b$, and $\mathrm{exc}(W) = c$ is at most:

$$n^{a-c+1} \cdot \ell^{2(\ell-b)} \cdot \ell^{2c}.$$

Proof. Observe that W has $a - c + 1$ vertices. Then the following information about W is sufficient to reconstruct it:

- The labels of the visited vertices in $[n]$ in the order in which they are visited. *There are at most n^{a-c+1} labelings.*
- The timestamps when the edge walked on is not a singleton edge. *There are at most $\ell^{\ell-b}$ possibilities.*
- The timestamps when W takes a step uv such that the edge $\{u, v\}$ has not been previously covered by W, but v has been previously visited, along with the timestamp of when v was visited for the first time. *There are c such steps, and hence there are at most ℓ^{2c} possibilities.*

- The timestamps when W takes a step uv such that the edge $\{u,v\}$ has been previously covered by W along with the timestamp of when $\{u,v\}$ was covered the first time. *There are at most* $\frac{\ell-b}{2}$ *such steps, and hence there are at most* $\ell^{\ell-b}$ *possibilities.*

Putting the above bounds together completes the proof. □

Observation 3.11. Any walk with b singleton edges and c excess edges has at most $\frac{\ell+b}{2}$ edges.

Proof. Each nonsingleton edge must be visited at least twice. There are at most $\ell-b$ non-singleton steps. So, there are at most $\frac{\ell-b}{2}$ nonsingleton edges, and the total number of edges is at most $\frac{\ell+b}{2}$. □

Now we can continue bounding the trace power.

$$\begin{aligned}(12) &\leqslant \sum_{a=1}^{\ell}\sum_{b=1}^{\ell}\sum_{c=1}^{\ell} p^{a-c}\lambda^b\left(\frac{27C^3}{\lambda}\right)^c \cdot n^{a-c+1}\cdot \ell^{3(\ell-b)}\cdot \ell^{2c}\\ &= n\sum_{a=1}^{\ell}\sum_{b=1}^{\ell}\sum_{c=1}^{\ell}(pn)^a\lambda^b\left(\frac{27C^3\ell^2}{\lambda pn}\right)^c\cdot \ell^{2(\ell-b)}\\ &\leqslant n\ell\cdot\max\left\{1,\left(\frac{27C^3\ell^2}{\lambda pn}\right)^{\ell}\right\}\cdot\sum_{a=1}^{\ell}\sum_{b=1}^{\ell}(\lambda pn)^b\cdot(pn)^{a-b}\cdot\ell^{2(\ell-b)}\end{aligned}$$

By Observation 3.11 and the assumption on λ from Assumption 6, we can bound the total edges a and hence the below.

$$\begin{aligned}&\leqslant n\ell\cdot\max\left\{1,\left(\frac{27C^3\ell^2}{\sqrt{pn}}\right)^{\ell}\right\}\cdot\sum_{a=1}^{\ell}\sum_{b=1}^{\ell}(\lambda pn)^b\cdot(pn)^{\frac{\ell-b}{2}}\cdot\ell^{2(\ell-b)}\\ &\leqslant n\ell^3\cdot\max\left\{1,\left(\frac{27C^3\ell^2}{\sqrt{pn}}\right)^{\ell}\right\}\max\left\{(\lambda pn)^{\ell},(pn\ell^4)^{\ell/2}\right\}\end{aligned}$$

By Claim 2.2,

$$\mathbf{Pr}\left[\|\overline{A}_{\boldsymbol{G}}\|\geqslant e^{\varepsilon}\cdot\left(n^{1/\ell}\ell^{3/\ell}\right)\max\left\{1,\frac{27C^3\ell^2}{\sqrt{pn}}\right\}\max\left\{\lambda pn,\sqrt{pn}\ell^2\right\}\right]\leqslant\exp(-\varepsilon\ell),$$

and choosing $\ell=\log^2 n$, $\varepsilon=\log\log n/\log n$, for any constant γ, we get:

$$\|A_{\boldsymbol{G}}-\mathbf{E}A_{\boldsymbol{G}}\|\leqslant(1+o(1))\cdot\left(1+\frac{27C^3\log^4 n}{\sqrt{pn}}\right)\cdot\max\left\{\lambda pn,\sqrt{pn}\log^4 n\right\}\tag{13}$$

with probability at least $1-n^{-\gamma}$. Now, by the assumption of the theorem, $pn\gg C^6\log^8 n$, so $1+\frac{27C^3\log^4 n}{\sqrt{pn}}=1+o(1)$. □

4 Analyzing the discrete walk with Brownian motion

In this section we quantify the extent to which convolving a measure α over $\mathbb{S}^{d-1}$ with a spherical cap of measure p brings α closer to uniform, provided that α satisfies a certain monotonicity property. We now define this monotonicity property, establishing a couple of additional definitions along the way.

Definition 4.1. We say a distribution on $\mathbb{S}^{d-1}$ with relative density α is *symmetric about* $y \in \mathbb{S}^{d-1}$ if there exists a function $\ell_\alpha : [-1,1] \to \mathbb{R}$ such that $\alpha(z) = \ell_\alpha(\langle z, y\rangle)$. We note that ℓ_α is also the density α projected onto the line defined by y relative to the projection of the uniform distribution, so that

$$\ell_\alpha(t) = \frac{\int_{\mathbb{S}^{d-1}} \mathbf{1}[\langle z, y\rangle = t] \cdot \ell_\alpha(t)\, d\rho(z)}{\int_{\mathbb{S}^{d-1}} \mathbf{1}[\langle z, y\rangle = t]\, d\rho(z)} = \frac{\int_{\mathbb{S}^{d-1}} \mathbf{1}[\langle z, y\rangle = t] \cdot \alpha(z)\, d\rho(z)}{\int_{\mathbb{S}^{d-1}} \mathbf{1}[\langle z, y\rangle = t]\, d\rho(z)}.$$

Notice that $\ell_\rho = 1$.

Definition 4.2. A measure α over $\mathbb{S}^{d-1}$ which is symmetric about some $y \in \mathbb{S}^{d-1}$ is said to be *spherically monotone* if ℓ_α is monotone non-decreasing.

An alternate characterization of spherically monotone distributions is that their relative densities can be written as a non-negative combination of spherical caps. Recall that we use $\underline{\mathrm{cap}}_p(y)$ and $\underline{\mathrm{cap}}_{\geqslant \tau(p)}(y)$ interchangeably to denote the uniform measure over $\mathrm{cap}_p(y)$.

Claim 4.3. A density $\alpha : \mathbb{S}^{d-1} \to \mathbb{R}$ which is symmetric about $y \in \mathbb{S}^{d-1}$ is spherically monotone if and only if there is a distribution r on $[-1,1]$ such that:

$$\alpha = \int \underline{\mathrm{cap}}_{\geqslant \theta}\, dr(\theta).$$

We call the above way of writing α as the *cap decomposition* of α. Further, $\ell_\alpha = \int \ell_{\underline{\mathrm{cap}}_{\geqslant\theta}}\, dr(\theta)$.

We give the straightforward proof later. Notice that in writing the expression for ℓ_α we have replaced $\ell_{\underline{\mathrm{cap}}_{\geqslant\theta}(y)}$ with $\ell_{\underline{\mathrm{cap}}_{\geqslant\theta}}$; this is because $\ell_{\underline{\mathrm{cap}}_{\geqslant\theta}(y)}$ does not depend on y.

Definition 4.4. Given a measure μ over $\mathbb{S}^{d-1}$ which is symmetric about some $y \in \mathbb{S}^{d-1}$, its *spherical kernel* P_μ is the transition operator of the random walk on $\mathbb{S}^{d-1}$ where a single step, starting from $x \in \mathbb{S}^{d-1}$, samples $\boldsymbol{a} \sim \ell_\mu$ and then walks from x to a uniformly random $\boldsymbol{w} \in \mathbb{S}^{d-1}$ satisfying $\langle \boldsymbol{w}, x\rangle = \boldsymbol{a}$. Equivalently, the density of $P_\mu \alpha$ is $\mu * \alpha$ for $*$ denoting convolution.

Remark 4.5. For brevity, we will use P_p as a shorthand for $P_{\underline{\mathrm{cap}}_p}$.

The main result of this section, proved after developing some tools, is the following:

Theorem 4.6. *If a probability distribution α over $\mathbb{S}^{d-1}$ is symmetric and spherically monotone, then for any integer $k \geqslant 0$,*

$$\mathrm{d_{TV}}\left(P_p^k \alpha, \rho\right) \leqslant \left((1 + o_{d\tau^2}(1)) \cdot \tau\right)^k \cdot \sqrt{\tfrac{1}{2}\mathrm{D}(\alpha \| \rho)},$$

where $o_{d\tau^2}(1)$ denotes a function that goes to 0 as $d\tau^2 \to \infty$.

As an immediate corollary, we obtain the following version which can be used in conjunction with Theorem 1.10 to conclude a bound on the second eigenvalue of random geometric graphs.

Corollary 4.7. *For any probability distribution α over $\mathbb{S}^{d-1}$,*

$$\mathrm{d_{TV}}\left(P_p^k \alpha, \rho\right) \leqslant \left((1 + o_{d\tau^2}(1)) \cdot \tau\right)^{k-1} \cdot \sqrt{\frac{1}{2} \cdot \log \frac{1}{p}}.$$

Proof. We write α as a convex combination of (symmetric, spherically monotone) point masses δ_x. Then we apply Theorem 4.6 in conjunction with the triangle inequality and the fact that $P_p^k \delta_x = P_p^{k-1} \underline{\mathrm{cap}}_p(x)$ and $\mathrm{D}(\underline{\mathrm{cap}}_p(x) \| \rho) = \log \frac{1}{p}$. ∎

Our proof of Theorem 4.6 will relate the action of P_p to the action of the *Brownian motion kernel*.

Definition 4.8 (Brownian motion on $\mathbb{S}^{d-1}$)**.** Let $(\boldsymbol{B}_t)_{t\geqslant 0}$ be standard Brownian motion in $\mathbb{R}^d$. We define Brownian motion on $\mathbb{S}^{d-1}$ starting at some point $V_0 \in \mathbb{S}^{d-1}$ as the process $(\boldsymbol{V}_t)_{t\geqslant 0}$ via the following stochastic differential equation:

$$d\boldsymbol{V}_t = \sqrt{2}\left(\mathbb{1} - \boldsymbol{V}_t\boldsymbol{V}_t^\top\right) d\boldsymbol{B}_t - (d-1)\boldsymbol{V}_t\, dt.$$

Definition 4.9. For any $t \geqslant 0$, let the *time-t Brownian motion kernel* U_t be the transition operator of a random walk on $\mathbb{S}^{d-1}$ where a single step samples runs a time-t Brownian motion on the sphere. Equivalently, $U_t = P_{\beta_t}$ for β_t the (spherically symmetric) density of a t-step Brownian motion.

For any $y \in \mathbb{S}^{d-1}$, $P_p y$ is highly concentrated near the boundary of the cap of measure p around y. As we will show in Section 4.1, the same is true for $U_t y$; it is highly concentrated near the boundary of a cap of measure $q = q(t)$ around y. So, choosing $T > 0$ so that $q(T) \approx p$, we will argue that U_T and P_p have similar action on spherically monotone measures.

We can then take advantage of the contractive properties of U_T in order to prove that P_p is contractive. The Brownian motion kernel satisfies the following mixing condition (which can be obtained, e.g., as a corollary of [BGL+14, Theorem 5.2.1] and [DEKL14, Corollary 2]):

Theorem 4.10 (Mixing of Brownian motion on $\mathbb{S}^{d-1}$)**.** *For any probability distribution ϕ on $\mathbb{S}^{d-1}$,*

$$\mathrm{D}(U_t\phi \,\|\, \rho) \leqslant \exp(-2(d-1)t) \cdot \mathrm{D}(\phi \,\|\, \rho)$$

As a corollary of the above and Pinsker's inequality, for any $t > 0$ and measure α over $\mathbb{S}^{d-1}$,

$$2\left(\mathrm{d_{TV}}\left(U_t\alpha, \rho\right)\right)^2 \leqslant \mathrm{D}(U_t\alpha \,\|\, \rho) \leqslant \exp(-2(d-1)t) \cdot \mathrm{D}(\alpha \,\|\, \rho). \tag{14}$$

Armed with (14), we can pass to working exclusively with the 1-dimensional projection of the measures in question onto the direction y.

Claim 4.11. For any spherically symmetric distribution with relative density γ, $\mathrm{d_{TV}}\left(\gamma, \rho\right) = \mathrm{d_{TV}}\left(\ell_\gamma, \ell_\rho\right)$.

Proof. We express the total variation distance in terms of the ℓ_1 norm:

$$\begin{aligned} 2\mathrm{d_{TV}}\left(\gamma, \rho\right) = \int_{z\in\mathbb{S}^{d-1}} \left|\gamma(z) - 1\right| d\rho(z) &= \int_{z\in\mathbb{S}^{d-1}} \left|\ell_\gamma(\langle z, y\rangle) - 1\right| d\rho(z) \\ &= \int_{t\in[-1,1]} \left|\ell_\gamma(t) - 1\right| d\ell_\rho(t) = 2\mathrm{d_{TV}}\left(\ell_\gamma, \ell_\rho\right). \end{aligned}$$

□

Note that if α is spherically symmetric about y then so is $U_t\alpha$, by the rotational invariance of Brownian Motion on the sphere. Hence combining Claim 4.11 with (14), we have that

$$\mathrm{d_{TV}}\left(\ell_{U_t\alpha}, \ell_\rho\right) \leqslant \sqrt{\frac{1}{2} \cdot \exp(-(d-1)t) \cdot \mathrm{D}(\alpha\|\rho)}.$$

Now, we'll show that for a well-chosen $T > 0$, $\ell_{U_T\alpha}$ nearly stochastically dominates $\ell_{P_p\alpha}$, and that $P_p\alpha$ and $U_T\alpha$ are both spherically monotone, and that this furthermore implies that $\mathrm{d_{TV}}\left(\ell_{U_T\alpha}, \ell_\rho\right)$ and $\mathrm{d_{TV}}\left(\ell_{P_p\alpha}, \ell_\rho\right)$ are related. Specifically, we show the following lemmas:

Lemma 4.12. *If ν and μ are spherically monotone densities and $\ell_\nu \preceq_{\mathrm{st}} \ell_\mu$, then*[10]

$$\mathrm{d_{TV}}\left(\ell_\nu, \ell_\rho\right) \leqslant \mathrm{d_{TV}}\left(\ell_\mu, \ell_\rho\right).$$

We prove the lemma below, but intuitively, a spherically monotone distribution can be realized as a non-negative combination of spherical caps; the uniform distribution has all of its mass on the largest cap (of measure 1). If $\ell_\nu \preceq_{\mathrm{st}} \ell_\mu$, then the total probability mass within any radius θ of the mode of μ exceeds that of ν, witnessing a larger total variation distance.

Lemma 4.13. *Let μ, ν, α be spherically monotone densities over $\mathbb{S}^{d-1}$, with $\ell_\nu \preceq_{\mathrm{st}} \ell_\mu$. Then*

1. *$P_\mu \alpha$ is spherically monotone (as is $P_\nu \alpha$),*
2. *$\ell_{P_\alpha \nu} \preceq_{\mathrm{st}} \ell_{P_\alpha \mu}$, and*
3. *$\ell_{P_\nu \alpha} \preceq_{\mathrm{st}} \ell_{P_\mu \alpha}$.*

We will prove this lemma below as well; the crux of the proof of Part 1 is to realize that because α, μ are spherically monotone, they can be decomposed as a non-negative combination of spherical caps. Then, by linearity of P_μ and by the commutativity of convolution, Part 1 reduces to showing that the convolution of two spherical caps is spherically monotone (this is a statement that we find intuitive, and it is easy to verify by directly examining the expression for $\ell_{P_{\geqslant\theta}\underline{\mathrm{cap}}_{\geqslant\psi}}$). To show Part 2, we observe that by decomposing α in its cap decomposition, it is then enough to compare $\ell_{P_{\geqslant\theta}\nu}$ with $\ell_{P_{\geqslant\theta}\mu}$ for each θ. Here, when $\ell_\mu \succeq_{\mathrm{st}} \ell_\nu$, a straightforward coupling demonstrates that $\ell_{P_{\geqslant\theta}\mu} \succeq_{\mathrm{st}} \ell_{P_{\geqslant\theta}\nu}$. Part 3 is a consequence of Part 2 and commutativity of convolution.

Our aim is to now apply these lemmas with $\nu \approx \underline{\mathrm{cap}}_p(y)$ and $\mu = \beta_T$ (note that $P_\nu \alpha = P_p \alpha$ and $P_\mu \alpha = U_T \alpha$). We now verify that these densities meet the conditions above. The density $\underline{\mathrm{cap}}_p(y)$ is spherically monotone because it is the same as ρ conditoned on being closer to y; we now show that β_t is indeed spherically monotone.

Claim 4.14. The density of a time-t Brownian motion, β_t, is spherically monotone.

Proof. Since Brownian motion on $\mathbb{S}^{d-1}$ can be realized as a sequence of random steps within spherical caps of infintesimally small measure ds, the measure of a t-step Brownian motion starting from $y \in \mathbb{S}^{d-1}$ is achieved by iteratively applying $P_{\underline{\mathrm{cap}}_{ds}}$ to the point mass at y. The proof is then complete by noting that $\ell_{\underline{\mathrm{cap}}_p}$ is spherically monotone for every p, then applying Part 1 of Lemma 4.13. □

Next, we argue that for $T = T(p)$, there is some small δ for which $(1-\delta)\ell_{\underline{\mathrm{cap}}_p} + \delta \ell_\rho \preceq_{\mathrm{st}} \ell_{\beta_T}$; that is, the linear projection of the p-cap is almost stochastically dominated by the linear projection of Brownian motion run for the proper amount of time. In order to do this, we first establish that almost all of the probability mass of β_T is in a cap of radius close to p. In Section 4.1, we'll prove the following lemma:

Lemma 4.15. *Let $(\boldsymbol{V}_t)_{t\geqslant 0}$ be a Brownian motion on $\mathbb{S}^{d-1}$ starting at V_0. Then for any time $t \geqslant 0$,*

$$\mathbf{Pr}\left[\left|\langle V_0, \boldsymbol{V}_t\rangle - \exp\left(-(d-1)t\right)\right| \geqslant x\right] \leqslant 2\exp\left(-\tfrac{d-1}{2}\frac{x^2}{1-e^{-2(d-1)t}}\right).$$

[10] As will be apparent from the proof, one may replace ℓ_ν, ℓ_μ with any monotone non-decreasing densities on $[-1,1]$.

From this lemma, we can show that almost all of the mass of the cap decomposition of ℓ_{β_T} is contained inside a $(\geqslant \tau)$-cap:

Claim 4.16. Let $\nu > 0$, $T := \frac{1}{d-1}\left(\log\frac{1}{\nu} - 2\varepsilon\right)$, and $\varepsilon \in \left[0, \frac{1}{2}\log\frac{1}{\nu}\right]$. Then the total mass of ℓ_{β_T} outside of $\mathrm{cap}_{\geqslant(1+\varepsilon)\nu}(V_0)$ for V_0 the starting point of the Brownian motion is bounded:

$$\int_{-1}^{(1+\varepsilon)\tau} d\ell_{\beta_T}(x) \leqslant \delta(\varepsilon) := 2\exp\left(-\frac{(d-1)\varepsilon^2\nu^2}{2(1-\nu^2)}\right).$$

Proof. We let $(\boldsymbol{V}_t)_{t\geqslant 0}$ be a Brownian motion on the sphere, $\boldsymbol{A}_t = \langle \boldsymbol{V}_t, V_0\rangle$, and $\boldsymbol{A}_t = \exp(-(d-1)t) + \boldsymbol{R}_t$. At time T, we have

$$\begin{aligned}\boldsymbol{A}_T = \exp(-(d-1)\cdot T) + \boldsymbol{R}_T &= \nu\cdot\exp(2\varepsilon) + \boldsymbol{R}_T \\ &\geqslant \nu\cdot(1+2\varepsilon) + \boldsymbol{R}_T \geqslant \nu\cdot(1+2\varepsilon) + \boldsymbol{R}_T.\end{aligned}$$

The event that $\boldsymbol{A}_T \leqslant \nu\cdot(1+\varepsilon)$ implies $\boldsymbol{R}_T < -\varepsilon\nu$, so it suffices to upper bound the probability that $|\boldsymbol{R}_T| > \varepsilon\nu$. Applying Lemma 4.15,

$$\mathbf{Pr}[|\boldsymbol{R}_T| \geqslant \varepsilon\nu] \leqslant 2\exp\left(-\tfrac{d-1}{2}\frac{\varepsilon^2\nu^2}{1-e^{-2(d-1)\cdot T}}\right) = 2\exp\left(-\frac{\varepsilon^2\nu^2(d-1)}{2(1-\nu^2)}\right)$$ □

Now, we are ready to establish the stochastic domination of the combination.

Claim 4.17. Let $p \in \left(0, \frac{1}{2}\right)$ and $\nu = \tau(p) + \frac{4}{\sqrt{d}}$. For $T = \frac{1}{d-1}(\log\frac{1}{\nu} - 2\varepsilon)$ with $\varepsilon \in \left[\frac{5}{(d-1)\nu^2}, \frac{1}{2}\log\frac{1}{\nu}\right]$,

$$\ell_{\beta_T} \succeq_{\mathrm{st}} (1-2\delta(\varepsilon))\ell_{\underline{\mathrm{cap}}_p} + 2\delta(\varepsilon)\ell_\rho,$$

for $\delta(\varepsilon)$ as defined in the statement of Claim 4.16.

Proof. Using Claim 4.3, we write $\beta_T = \int_{-1}^{1} c_\theta \cdot \ell_{\underline{\mathrm{cap}}_{\geqslant\theta}} d\theta$, with $\int c_\theta d\theta = 1$. Let $\tau' \in [-1,1]$ be such that

$$\int_{-1}^{\tau'} c_\theta d\theta = 2\delta(\varepsilon), \quad \text{and} \quad \int_{\tau'}^{1} c_\theta d\theta = 1 - 2\delta(\varepsilon). \tag{15}$$

The proof strategy is to show that the conclusion follows if $\tau' \geqslant \tau(p)$, and then establish that inequality.

First observe that if α and $\{\gamma_x\}_{x\in X}$ are measures satisfying $\gamma_x \succeq_{\mathrm{st}} \alpha$ for all $x \in X$, then a convex combination $\int c_x\gamma_x dx \succeq_{\mathrm{st}} \alpha$ as well, from which the conclusion follows. Now, writing

$$\beta_T = \int_{-1}^{\tau'} c_\theta \cdot \ell_{\underline{\mathrm{cap}}_{\geqslant\theta}} d\theta + \int_{\tau'}^{1} c_\theta \cdot \ell_{\underline{\mathrm{cap}}_{\geqslant\theta}} d\theta,$$

we see that the first term on the right-hand-side stochastically dominates $2\delta(\varepsilon)\cdot\ell_{\underline{\mathrm{cap}}_{\geqslant -1}} = 2\delta(\varepsilon)\cdot\ell_\rho$ since for every $\theta \in [-1,\tau']$, $\theta \geqslant -1$ and therefore $\ell_{\underline{\mathrm{cap}}_{\geqslant\theta}} \succeq_{\mathrm{st}} \ell_{\underline{\mathrm{cap}}_{\geqslant -1}} = \ell_\rho$. By identical reasoning, the second term stochastically dominates $\ell_{\underline{\mathrm{cap}}_{\geqslant\tau(p)}}$ since for every $\theta \in [\tau', 1]$, $\theta \geqslant \tau(p)$ and therefore $\ell_{\underline{\mathrm{cap}}_{\geqslant\theta}} \succeq_{\mathrm{st}} \ell_{\underline{\mathrm{cap}}_{\geqslant\tau}} = \ell_{\underline{\mathrm{cap}}_p}$.

We now show that the τ' satisfying (15) is at least τ, for which it is sufficient to show $\tau' \geqslant \nu$. Let $\kappa = \int_{-1}^{\nu} c_\theta\, d\theta$; $\tau' \geqslant \nu$ is equivalent to showing that $\kappa \leqslant 2\delta(\varepsilon)$. Using Claim 4.16, we know that $\mathbf{Pr}_{\boldsymbol{v}\sim\beta_T}[\boldsymbol{v} \in \mathrm{cap}_{\geqslant(1+\varepsilon)\nu}(V_0)] \geqslant 1-\delta(\varepsilon)$.

$$1 - \delta(\varepsilon) \leqslant \mathop{\mathbf{Pr}}_{\boldsymbol{v}\sim\beta_T}\left[\boldsymbol{v} \in \mathrm{cap}_{\geqslant(1+\varepsilon)\nu}(V_0)\right]$$

$$
\begin{aligned}
&= \int_{-1}^{(1+\varepsilon)v} c_\theta \cdot \mathop{\mathbf{Pr}}_{\boldsymbol{x}\sim \ell_{\underline{\mathrm{cap}}_{\geqslant\theta}}} [\boldsymbol{x} \geqslant (1+\varepsilon)v]\, d\theta + \int_{(1+\varepsilon)v}^{1} c_\theta \, d\theta \\
&\leqslant \int_{-1}^{v} c_\theta \cdot \mathop{\mathbf{Pr}}_{\boldsymbol{x}\sim \ell_{\underline{\mathrm{cap}}_{\geqslant\theta}}} [\boldsymbol{x} \geqslant (1+\varepsilon)v]\, d\theta + \int_{v}^{1} c_\theta \, d\theta \\
&\leqslant \left(\max_{\theta\in[-1,v]} \mathop{\mathbf{Pr}}_{\boldsymbol{x}\sim \ell_{\underline{\mathrm{cap}}_{\geqslant\theta}}} [\boldsymbol{x} \geqslant (1+\varepsilon)v] \right) \cdot \kappa + \int_{v}^{1} c_\theta \, d\theta = \mathop{\mathbf{Pr}}_{x\sim \ell_{\underline{\mathrm{cap}}_{\geqslant v}}} [\boldsymbol{x} \geqslant (1+\varepsilon)v] \cdot \kappa + \int_{v}^{1} c_\theta \, d\theta.
\end{aligned}
$$

Using Lemma 2.8 and $v \geqslant 4/\sqrt{d}$,

$$
\begin{aligned}
\mathop{\mathbf{Pr}}_{\boldsymbol{x}\sim \ell_{\underline{\mathrm{cap}}_{\geqslant v}}} [\boldsymbol{x} \geqslant (1+\varepsilon)v] &= \frac{\rho(\mathrm{cap}_{\geqslant(1+\varepsilon)v})}{\rho(\mathrm{cap}_{\geqslant v})} \\
&\leqslant \frac{3v\left(1-((1+\varepsilon)v)^2\right)^{(d-1)/2}}{2v(1+\varepsilon)\left(1-v^2\right)^{(d-1)/2}} \\
&\leqslant \frac{3}{2}\cdot\left(\frac{1-(1+\varepsilon)^2v^2}{1-v^2}\right)^{(d-1)/2} \\
&= \frac{3}{2}\cdot\left(1-\frac{2\varepsilon v^2+\varepsilon^2v^2}{1-v^2}\right)^{(d-1)/2} \leqslant \frac{3}{2}\cdot\left(1-2\varepsilon v^2\right)^{(d-1)/2} \leqslant \frac{3}{1+(d-1)\varepsilon v^2}.
\end{aligned}
$$

The final quantity is smaller than $\frac{1}{2}$ given our lower bound on ε, and $\int_\tau^1 c_\theta\, d\theta = 1-\kappa$. Plugging into the above, we have that

$$
1-\delta(\varepsilon) \leqslant \frac{1}{2}\kappa + 1 - \kappa \implies \kappa \leqslant 2\delta(\varepsilon),
$$

which completes the proof. □

Finally, we will need the following claim to transfer the statement about the stochastic domination of a linear combination of $\ell_{P_p\alpha}$ and ℓ_ρ to just $\ell_{P_p\alpha}$:

Lemma 4.18. *Suppose μ and ν are spherically monotone distributions, then for any $\eta \in [0,1)$,*

$$
\mathrm{d_{TV}}\left(\ell_\mu, \ell_\rho\right) \leqslant \frac{1}{1-\eta}\mathrm{d_{TV}}\left((1-\eta)\ell_\mu + \eta\ell_\nu, \ell_\rho\right).
$$

Proof. Let $s \in [-1,1]$ be such that:

$$
\mathrm{d_{TV}}\left(\ell_\mu, \ell_\rho\right) = \int_s^1 (\ell_\mu(x)-1)\, d\rho_{1\mathrm{D}} = \int_{-1}^s (1-\ell_\mu(x))\, d\rho_{1\mathrm{D}},
$$

where $\rho_{1\mathrm{D}}$ is the density of the 1-dimensional projection of $\underline{\mathrm{cap}}_1(y)$. The choice of s satisfying the above is the one satisfying $\ell_\mu(s) = 1$. If $\ell_\nu(s) \geqslant 1$, by spherical monotonicity $\ell_\nu(x) \geqslant 1$ on $[s,1]$ and:

$$
\begin{aligned}
\mathrm{d_{TV}}\left((1-\eta)\ell_\mu + \eta\ell_\nu, \ell_\rho\right) &\geqslant \int_s^1 \left((1-\eta)(\ell_\mu(x)-1) + \eta(\ell_\nu(x)-1)\right) d\rho_{1\mathrm{D}} \\
&\geqslant (1-\eta)\int_s^1 (\ell_\mu(x)-1)\, d\rho_{1\mathrm{D}} = (1-\eta)\cdot\mathrm{d_{TV}}\left(\ell_\mu, \ell_\rho\right).
\end{aligned}
$$

On the other hand, if $\ell_\nu(s) \leqslant 1$, by an identical argument we know:

$$
\mathrm{d_{TV}}\left((1-\eta)\ell_\mu + \eta\ell_\nu, \ell_\rho\right) \geqslant (1-\eta)\int_{-1}^s (1-\ell_\mu(x))\, d\rho_{1\mathrm{D}} = (1-\eta)\cdot\mathrm{d_{TV}}\left(\ell_\mu, \ell_\rho\right).
$$

The desired statement follows from rearranging the above inequality. □

We are now ready to prove Theorem 4.6, following the reasoning above, in combination with induction on k, the number of applications of P_p. We state and prove a more refined version of Theorem 4.6 below.

Theorem 4.19. *If a probability distribution α over $\mathbb{S}^{d-1}$ is symmetric and spherically monotone, then for any integer $k \geqslant 0$ and for $\nu = \tau(p,d) + \frac{4}{\sqrt{d}}$,*

$$\mathrm{d_{TV}}\left(P_p^k \alpha, \rho\right) \leqslant \nu^k \left(\frac{\exp\left(\frac{4}{(d-1)^{1/4}\sqrt{\nu}}\right)}{\sqrt{1 - 2\exp(-\nu\sqrt{d-1})}} \right).$$

Note that when $\tau^2 d \to \infty$, the parenthesized term is $1 + o(1)$ and $\nu = \tau \cdot (1 + o(1))$.

Proof. Suppose $\tau(p) \geqslant 1 - 1/(d-1)^{1/4}$, then the statement is vacuously true. Thus, we assume from now on $\tau(p) < 1 - 1/(d-1)^{1/4}$.

Let $\nu = \tau(p) + 4/\sqrt{d}$, let $t = \frac{1}{d-1}\left(\log\frac{1}{\nu} - 2\varepsilon\right)$, and $\delta = 2\exp\left(-\frac{(d-1)\varepsilon^2\nu^2}{2(1-\nu^2)}\right)$ for $\varepsilon = \frac{\sqrt{2-2\nu^2}}{(d-1)^{1/4}\sqrt{\nu}}$; note that for d sufficiently large, $\varepsilon \in \left[\frac{5}{(d-1)\nu^2}, \frac{1}{2}\log\frac{1}{\nu}\right]$. For convenience's sake, define $P_{p,\delta} = (1-2\delta)P_p + 2\delta P_1$. We will prove that

$$\ell_{P_{p,\delta}^k \alpha} \preceq_{\mathrm{st}} \ell_{U_t^k \alpha}, \qquad \text{and} \qquad U_t^k \alpha,\ P_{p,\delta}^k \alpha \quad \text{are spherically monotone.} \tag{16}$$

Given this, the proof of the theorem will follow: by the linearity of the projection onto the line defined by y, and by the commutativity of convolution,

$$\ell_{P_{p,\delta}^k \alpha} = \sum_{j=0}^{k} (1-2\delta)^{k-j} (2\delta)^j \binom{k}{j} \ell_{P_p^{k-j} P_1^j \alpha},$$

So from Claim 4.11, Lemma 4.18, (16), and Lemma 4.12,

$$\mathrm{d_{TV}}\left(P_p^k \alpha, \rho\right) = \mathrm{d_{TV}}\left(\ell_{P_p^k \alpha}, \ell_\rho\right) \leqslant \frac{1}{(1-2\delta)^k} \mathrm{d_{TV}}\left(\ell_{P_{p,\delta}^k \alpha}, \ell_\rho\right) \leqslant \frac{1}{(1-2\delta)^k} \mathrm{d_{TV}}\left(\ell_{U_t^k \alpha}, \ell_\rho\right). \tag{17}$$

Then we can apply Claim 4.11 to get that

$$\mathrm{d_{TV}}\left(\ell_{U_t^k \alpha}, \ell_\rho\right) = \mathrm{d_{TV}}\left(U_t^k \alpha, \rho\right), \tag{18}$$

and finally using that $U_t^k = U_{k \cdot t}$ in conjunction with Lemma 4.10, we have that

$$\mathrm{d_{TV}}\left(U_t^k \alpha, \rho\right) = \mathrm{d_{TV}}\left(U_{k\cdot t} \alpha, \rho\right) \leqslant \sqrt{\frac{1}{2}\exp(-2(d-1)tk) \cdot \mathrm{D}(\alpha \| \rho)}, \tag{19}$$

So combining (17), (18), and (19), we have that

$$\mathrm{d_{TV}}\left(P_p^k \alpha, \rho\right) \leqslant \sqrt{\frac{1}{2(1-2\delta)^k}\exp(-2(d-1)tk) \cdot \mathrm{D}(\alpha \| \rho)}.$$

In our case, $\delta = \exp(-\nu\sqrt{d-1})$, $t = \frac{1}{d-1}\left(\log\frac{1}{\nu} - \frac{\sqrt{2-2\nu^2}}{(d-1)^{1/4}\sqrt{\nu}}\right)$, so combining these estimates,

$$\mathrm{d_{TV}}\left(P_p^k \alpha, \rho\right) \leqslant \nu^k \cdot \left(\frac{\exp\left(\frac{4}{(d-1)^{1/4}\sqrt{\nu}}\right)}{\sqrt{1 - 2\exp(-\nu\sqrt{d-1})}} \right)^k \cdot \sqrt{\tfrac{1}{2}\mathrm{D}(\alpha \| \rho)},$$

as desired.

Now we prove (16). The proof is by induction on k; when $k = 0$, there is nothing to prove. Suppose now that the statement holds true for k; we shall prove it for $k+1$. By Claim 4.14, the density of a time-t spherical Brownian motion β_t is spherically monotone about its starting point, and clearly, any convex combination of caps is spherically monotone. Hence we can apply Lemma 4.13, Part 1 in conjunction with the induction hypothesis to conclude that both $P_{p,\delta}^{k+1}\alpha = P_{p,\delta}(P_{p,\delta}^{k}\alpha)$ and $U_t^{k+1}\alpha = U_t(U_t^k\alpha)$ are spherically monotone, giving the second part of the induction hypothesis.

By our induction hypothesis $U_t^k\alpha$ and $P_{p,\delta}^k\alpha$ are spherically monotone with $\ell_{U_t^k\alpha} \succeq_{\mathrm{st}} \ell_{P_{p,\delta}^k\alpha}$, and so we can apply Lemma 4.13, Part 2 in conjunction with Claim 4.14 to conclude that

$$\ell_{U_t^{k+1}\alpha} = \ell_{P_{\beta_t}(U_t^k\alpha)} \succeq_{\mathrm{st}} \ell_{P_{\beta_t}(P_{p,\delta}^k\alpha)},$$

and then apply Lemma 4.13, Part 3 in conjunction with Claim 4.17 to conclude that

$$\ell_{P_{\beta_t}(P_{p,\delta}^k\alpha)} \succeq_{\mathrm{st}} \ell_{P_{p,\delta}(P_{p,\delta}^k\alpha)} = \ell_{P_{p,\delta}^k\alpha},$$

completing the proof. □

Now, we fill in the proofs of the lemmas from above.

Claim (Restatement of Claim 4.3)**.** A density $\alpha : \mathbb{S}^{d-1} \to \mathbb{R}$ which is symmetric about $y \in \mathbb{S}^{d-1}$ is spherically monotone if and only if there is a distribution r on $[-1,1]$ such that:

$$\alpha = \int \underline{\mathrm{cap}}_{\geqslant\theta}\, dr(\theta).$$

We call the above way of writing α as the *cap decomposition* of α. Further, $\ell_\alpha = \int \ell_{\underline{\mathrm{cap}}_{\geqslant\theta}}\, dr(\theta)$.

Proof of Claim 4.3. We first prove the "only if" direction. Since α is spherically symmetric about y, $\alpha(v) = \ell_\alpha(\langle v, y\rangle)$. Let $d\ell_\alpha$ be the distributional derivative of ℓ_α, and set $dr(\theta) = \rho(\mathrm{cap}_{\geqslant\theta}(y))\, d\ell_\alpha(\theta)$.

$$\int (\underline{\mathrm{cap}}_{\geqslant\theta}(y))(v)\, dr(\theta) = \int \frac{\mathbf{1}[\langle v, y\rangle \geqslant \theta]}{\rho(\mathrm{cap}_{\geqslant\theta}(y))} \cdot \rho(\mathrm{cap}_{\geqslant\theta}(y))\, d\ell_\alpha(\theta) = \int \mathbf{1}[\langle v, y\rangle \geqslant \theta]\, d\ell_\alpha(\theta) = \alpha(v).$$

To see that the measure dr indeed gives a probability distribution, first observe that $dr(\theta) \geqslant 0$ for every θ due to the monotonicity of ℓ_α, and next observe that

$$\begin{aligned} 1 = \int_{v\in\mathbb{S}^{d-1}} \alpha(v)\, d\rho(v) &= \int_{v\in\mathbb{S}^{d-1}} \int_{-1}^{1} (\underline{\mathrm{cap}}_{\geqslant\theta}(y))(v)\, dr(\theta)\, d\rho(v) \\ &= \int_{-1}^{1} \int_{v\in\mathbb{S}^{d-1}} (\underline{\mathrm{cap}}_{\geqslant\theta}(y))(v)\, d\rho(v)\, dr(\theta) = \int_{-1}^{1} dr(\theta). \end{aligned}$$

In summary, since r is a positive measure which integrates to 1, it is a probability distribution. The claim regarding ℓ_α follows because the line projection onto y is a linear operation.

Now we prove the converse. Suppose $\alpha = \int_{-1}^{1} \underline{\mathrm{cap}}_{\geqslant\theta}(y)\, dr(\theta)$. By linearity of projection onto the line defined by y, $\ell_\alpha = \int_{-1}^{1} \ell_{\underline{\mathrm{cap}}_{\geqslant\theta}}\, dr(\theta)$. Since $\ell_{\underline{\mathrm{cap}}_{\geqslant\theta}}$ is monotone for every θ, and a non-negative combination of monotone functions is monotone, ℓ_α is also monotone, concluding the proof. □

Lemma (Restatement of Lemma 4.12)**.** *If ν and μ are spherically monotone densities and $\ell_\nu \preceq_{\mathrm{st}} \ell_\mu$, then*[11]

$$\mathrm{d_{TV}}\left(\ell_\nu, \ell_\rho\right) \leqslant \mathrm{d_{TV}}\left(\ell_\mu, \ell_\rho\right).$$

[11] As will be apparent from the proof, one may replace ℓ_ν, ℓ_μ with any monotone non-decreasing densities on $[-1,1]$.

Proof of Lemma 4.12. First, observe that $\ell_\nu \succeq_{\text{st}} \ell_\rho$ and $\ell_\mu \succeq_{\text{st}} \ell_\rho$ by the assumption that μ, ν are spherically monotone. Thus, $\ell_\mu \succeq_{\text{st}} \ell_\nu \succeq_{\text{st}} \ell_\rho$. Further, if measures a, b on $[-1,1]$ satisfy $a \succeq_{\text{st}} b$, then their CDFs G_a and G_b satisfy $G_a(s) \leqslant G_b(s)$ for every s. Hence,

$$G_{\ell_\mu}(s) \leqslant G_{\ell_\nu}(s) \leqslant G_{\ell_\rho}(s) \quad \forall s \in [-1,1].$$

By definition of the total variation distance, for any non-decreasing density $\gamma : [-1,1] \to \mathbb{R}$,

$$\mathrm{d_{TV}}\left(\gamma, \ell_\rho\right) = \max_{s \in [-1,1]} G_{\ell_\rho}(s) - G_\gamma(s).$$

Thus,

$$\mathrm{d_{TV}}\left(\ell_\nu, \ell_\rho\right) = G_{\ell_\rho}(s^*) - G_{\ell_\nu}(s^*) \leqslant G_{\ell_\rho}(s^*) - G_{\ell_\mu}(s^*) \leqslant \mathrm{d_{TV}}\left(\ell_\mu, \ell_\rho\right),$$

which completes the proof. □

We'll now prove Lemma 4.13.

Lemma (Restatement of Lemma 4.13)**.** *Let μ, ν, α be spherically monotone densities over $\mathbb{S}^{d-1}$, with $\ell_\nu \preceq_{\text{st}} \ell_\mu$. Then*

1. *$P_\mu \alpha$ is spherically monotone (as is $P_\nu \alpha$),*
2. *$\ell_{P_\alpha \nu} \preceq_{\text{st}} \ell_{P_\alpha \mu}$, and*
3. *$\ell_{P_\nu \alpha} \preceq_{\text{st}} \ell_{P_\mu \alpha}$.*

Proof of Lemma 4.13. We first prove Part 1. We can write α and μ in terms of their *cap decompositions* as shown in Claim 4.3, $\alpha = \int_{-1}^{1} \underline{\text{cap}}_{\geqslant\theta}(y)\, dr(\theta)$ and $\mu = \int_{-1}^{1} \underline{\text{cap}}_{\geqslant\psi}(z)\, ds(\psi)$ for some $z \in \mathbb{S}^{d-1}$. P_μ is a linear operator, so $P_\mu \alpha = \int P_\mu \underline{\text{cap}}_{\geqslant\theta}(y)\, dr(\theta)$. Further, by the commutativity of convolution, $P_\mu \underline{\text{cap}}_{\geqslant\theta}(y) = P_{\geqslant\theta} \mu_y$, where μ_y denotes the version of μ centered at y. Hence,

$$P_\mu \alpha = \int P_\mu \underline{\text{cap}}_{\geqslant\theta}(y) dr(\theta) = \int P_{\geqslant\theta} \mu_y \, dr(\theta) = \int \int P_{\geqslant\theta} \underline{\text{cap}}_{\geqslant\psi}(y)\, ds(\psi)\, dr(\theta).$$

Each $P_{\geqslant\theta} \underline{\text{cap}}_{\geqslant\psi}(y)$ is clearly spherically symmetric about y. Since the projection onto the line defined by y is a linear operation, $\ell_{P_\mu \alpha} = \int \int \ell_{P_{\geqslant\theta} \underline{\text{cap}}_{\geqslant\psi}} ds(\psi) dr(\theta)$, and because a non-negative combination of monotone functions is monotone, it suffices to prove that for any $\theta, \psi \in [-1,1]$, $\ell_{P_{\geqslant\theta} \underline{\text{cap}}_{\geqslant\psi}}$ is monotone. By definition,

$$\begin{aligned}
\ell_{P_{\geqslant\theta} \underline{\text{cap}}_{\geqslant\psi}(y)}(t) &= \mathop{\mathbf{E}}_{\boldsymbol{v} \sim \rho} \left[\left(P_{\geqslant\theta} \underline{\text{cap}}_{\geqslant\psi}(y) \right)(\boldsymbol{v}) \mid \langle \boldsymbol{v}, y \rangle = t \right] \\
&= \mathop{\mathbf{E}}_{\boldsymbol{v} \sim \rho} \left[\mathop{\mathbf{E}}_{\boldsymbol{w} \sim \underline{\text{cap}}_{\geqslant\theta}(\boldsymbol{v})} \left[\left(\underline{\text{cap}}_{\geqslant\psi}(y) \right)(\boldsymbol{w}) \right] \mid \langle \boldsymbol{v}, y \rangle = t \right] \\
&= \mathop{\mathbf{E}}_{\boldsymbol{v} \sim \rho} \left[\mathop{\mathbf{E}}_{\boldsymbol{w} \sim \underline{\text{cap}}_{\geqslant\theta}(\boldsymbol{v})} \left[\frac{\mathbf{1}[\langle \boldsymbol{w}, y \rangle \geqslant \psi]}{\rho(\text{cap}_{\geqslant\psi})} \right] \mid \langle \boldsymbol{v}, y \rangle = t \right] = \frac{\mathbf{Pr}_{\boldsymbol{v}, \boldsymbol{w} \sim \rho} \left[\langle \boldsymbol{w}, y \rangle \geqslant \psi \mid \langle \boldsymbol{w}, \boldsymbol{v} \rangle \geqslant \theta, \langle \boldsymbol{v}, y \rangle = t \right]}{\mathbf{Pr}_{\boldsymbol{w} \sim \rho} \left[\langle \boldsymbol{w}, y \rangle \geqslant \psi \right]}.
\end{aligned}$$

This ratio is monotone increasing in t, completing the proof of (1).

Now we show Part 2. Claim 4.3 shows that by the spherical monotonicity of α, we can express α in its cap decomposition,

$$\alpha = \int_0^1 \underline{\text{cap}}_q \, dr(q),$$

and now by the linearity of convolution, $P_\alpha = \int_0^1 P_{\underline{\mathrm{cap}}_q}\, dr(q)$, and $P_\alpha \mu = \int P_q \mu\, dr(q)$, $P_\alpha \nu = \int P_q \nu dr(q)$. So, to show that $\ell_{P_\alpha \mu} \succeq_{\mathrm{st}} \ell_{P_\alpha \nu}$, it suffices to argue "slice-by-slice" that for every $q \in [0,1]$, $\ell_{P_q \mu} \succeq_{\mathrm{st}} \ell_{P_q \nu}$.

This follows from the following coupling argument: we sample $(\boldsymbol{x}, \boldsymbol{y})$ from $(\ell_{P_q\mu}, \ell_{P_q\nu})$ in a coupled manner as follows: first, sample $(\boldsymbol{a}_\mu, \boldsymbol{a}_\nu) \sim (\ell_\mu, \ell_\nu)$ in a coupled manner so that $\boldsymbol{a}_\mu \geqslant \boldsymbol{a}_\nu$; such a coupling is guaranteed because $\ell_\mu \succeq_{\mathrm{st}} \ell_\nu$. Next, choose $(\boldsymbol{v}_\mu, \boldsymbol{v}_\nu)$ at random in $\mathbb{S}^{d-1}$ conditioned on $\langle \boldsymbol{v}_\mu, y\rangle = \boldsymbol{a}_\mu$ and $\langle \boldsymbol{v}_\nu, y\rangle = \boldsymbol{a}_\nu$. Now, let $\boldsymbol{\theta}_\mu$ be the random variable $\langle y, \boldsymbol{u}_\mu\rangle$ for $\boldsymbol{u}_\mu \sim \underline{\mathrm{cap}}_q(\boldsymbol{v}_\mu)$, and $\boldsymbol{\theta}_\nu = \langle y, \boldsymbol{u}_\nu\rangle$ for $\boldsymbol{u}_\nu \sim \underline{\mathrm{cap}}_q(\boldsymbol{v}_\nu)$. Note that the marginal over $\boldsymbol{\theta}_\mu$ is $\ell_{P_q\mu}$ and the marginal over $\boldsymbol{\theta}_\nu$ is $\ell_{P_q\nu}$. The probability $\mathbf{Pr}[\theta_\mu > t]$ is proportional to the measure of the intersection of $\mathrm{cap}_{\geqslant t}(y)$ and $\mathrm{cap}_q(\boldsymbol{u}_\mu)$, and similarly the probability $\mathbf{Pr}[\theta_\nu > t]$ is proportional to the measure of the intersection of $\mathrm{cap}_{\geqslant t}(y)$ and $\mathrm{cap}_q(\boldsymbol{u}_\nu)$. By our choice of coupling, the angle between $\boldsymbol{u}_\mu$ and y is smaller than the angle between $\boldsymbol{u}_\nu$ and y, so for every $t \in [-1,1]$,

$$\mathbf{Pr}[\theta_\mu > t] \geqslant \mathbf{Pr}[\theta_\nu > t],$$

and hence we may couple $\boldsymbol{\theta}_\mu$ and $\boldsymbol{\theta}_\nu$ so that $\boldsymbol{\theta}_\mu \geqslant \boldsymbol{\theta}_\nu$ always. Taking $\boldsymbol{x} = \boldsymbol{\theta}_\mu$ and $\boldsymbol{y} = \boldsymbol{\theta}_\nu$ in this coupling gives our conclusion.

Finally, observe that by the commutativity of convolution, $P_\mu \alpha = P_\alpha \mu$ and $P_\nu \alpha = P_\alpha \nu$, and so Part 3 follows from Part 2. □

4.1 Concentration of spherical Brownian Motion within a cap

In this section, we study the concentration of Brownian Motion on $\mathbb{S}^{d-1}$ in the spherical cap around its starting point.

Lemma (Restatement of Lemma 4.15)**.** *Let $(\boldsymbol{V}_t)_{t\geqslant 0}$ be a Brownian motion on $\mathbb{S}^{d-1}$ starting at V_0. Then for any time $t \geqslant 0$,*

$$\mathbf{Pr}\left[\left|\langle V_0, \boldsymbol{V}_t\rangle - \exp(-(d-1)t)\right| \geqslant x\right] \leqslant 2\exp\left(-\frac{d-1}{2}\frac{x^2}{1-e^{-2(d-1)t}}\right).$$

Proof of Lemma 4.15. Letting $\boldsymbol{A}_t = \langle V_0, \boldsymbol{V}_t\rangle$ be the correlation of the motion at step t with the starting point, $(\boldsymbol{B}_t)_{t\geqslant 0}$ be standard Brownian motion on $\mathbb{R}^d$, $(\boldsymbol{B}'_t)_{t\geqslant 0}$ be standard Brownian motion on $\mathbb{R}$, and $\theta = d-1$,

$$\begin{aligned} d\boldsymbol{A}_t = \langle V_0, d\boldsymbol{V}_t\rangle &= -\theta \cdot \boldsymbol{A}_t\, dt + \sqrt{2}\langle V_0, \left(\mathbb{1} - \boldsymbol{V}_t\boldsymbol{V}_t^\top\right) d\boldsymbol{B}_t\rangle \\ &= -\theta \cdot \boldsymbol{A}_t\, dt + \sqrt{2}\left\langle \left(\mathbb{1} - \boldsymbol{V}_t\boldsymbol{V}_t^\top\right) V_0, d\boldsymbol{B}_t\right\rangle \\ &= -\theta \cdot \boldsymbol{A}_t\, dt + \sqrt{2}\sqrt{1-\boldsymbol{A}_t^2}\, d\boldsymbol{B}'_t \end{aligned}$$

The solution to the deterministic differential equation $dx_t = -\theta x_t$ with initial condition $x_0 = 1$ is $x_t = \exp(-\theta t)$. To this end, it's convenient to split $\boldsymbol{A}_t$ up into a deterministic and a random part:

$$\boldsymbol{A}_t = \exp(-\theta t) + \boldsymbol{R}_t,$$

with the initial condition $R_0 = 0$. Then via calculation,

$$d\boldsymbol{R}_t = -\theta \boldsymbol{R}_t dt + \sqrt{2}\sqrt{1-\boldsymbol{A}_t^2}\, d\boldsymbol{B}'_t. \tag{20}$$

We now relate $\boldsymbol{R}_t$ to a stochastic process without drift, as is done, for example, in the analysis of the Ornstein-Uhlenbeck process. Consider $\boldsymbol{R}_t \exp(\theta t)$. Note that

$$\begin{aligned} d(\boldsymbol{R}_t \exp(\theta t)) &= \exp(\theta t)\, d\boldsymbol{R}_t + \boldsymbol{R}_t \theta \exp(\theta t)\, dt \\ &= -\boldsymbol{R}_t \theta \exp(\theta t)\, dt + \sqrt{2}\exp(\theta t)\sqrt{1-\boldsymbol{A}_t^2}\, d\boldsymbol{B}'_t + \boldsymbol{R}_t \theta \exp(\theta t)\, dt \\ &= \sqrt{2}\exp(\theta t)\sqrt{1-\boldsymbol{A}_t^2}\, d\boldsymbol{B}'_t, \end{aligned}$$

a process without drift.

The following version of the Azuma–Hoeffding inequality will allow us to argue that this driftless process concentrates.

Lemma 4.20. *Let $(\boldsymbol{X}_t)_{t\geqslant 0} \subset \mathbb{R}$ be a stochastic process adapted to the filtration $\mathscr{F}_t$ with $\mathbf{E}[e^{r\, d\boldsymbol{X}_t} \mid \mathscr{F}_t] < \exp\left(r^2\sigma_t^2 dt\right)$, for all t, r. Then for all $s, x > 0$,*

$$\mathbf{Pr}[|\boldsymbol{X}_s - \boldsymbol{X}_0| \geqslant x] \leqslant 2\exp\left(\frac{-x^2}{4\int_0^s \sigma_t^2 dt}\right).$$

Versions of this lemma are known (c.f. [Dem96] and references therein), we include a proof in Appendix A for completeness.

We apply Lemma 4.20 to prove that

$$\mathbf{Pr}[|\boldsymbol{R}_s| \geqslant x] \leqslant 2\exp\left(-C' \cdot \frac{x^2 d\theta}{1-\exp(-\theta s)}\right).$$

Indeed, $\boldsymbol{R}_t \exp(\theta t)$ is a stochastic process without drift and satisfies that

$$\begin{aligned} \mathbf{E}[\exp(r\, d(\boldsymbol{R}_t \exp(\theta t)))] &= \mathbf{E}\left[\exp\left(\sqrt{2} r \exp(\theta t)\sqrt{1-\boldsymbol{A}_t^2}\, d\boldsymbol{B}'_t\right)\right] \\ &\leqslant \exp\left(r^2 \exp(2\theta t)\cdot\left(1-\boldsymbol{A}_t^2\right) dt\right) \\ &\leqslant \exp\left(r^2 \exp(2\theta t)\, dt\right), \end{aligned}$$

Since $\boldsymbol{A}_t$ is real-valued. So we can apply Lemma 4.20 to the process and derive that

$$\begin{aligned} \mathbf{Pr}[|\boldsymbol{R}_s| \geqslant x] &= \mathbf{Pr}[|\boldsymbol{R}_s \exp(\theta s)| \geqslant x \exp(\theta s)] \\ &\leqslant 2\exp\left(-\frac{x^2 \exp(2\theta s)}{4\int_0^s \exp(2\theta t) dt}\right) \\ &= 2\exp\left(-\frac{\theta x^2}{2\cdot(1-\exp(-2\theta s))}\right), \end{aligned}$$

and plugging in $\theta = d-1$ concludes the proof. □

5 The second eigenvalue of links

In this section we analyze the links of the random geometric complex. Each link is a random geometric graph in a cap centered around some $w \in \mathbb{S}^{d-1}$ on $\boldsymbol{m}$ vertices where $\boldsymbol{m} \sim \mathsf{Binom}(n, p)$. We are interested in obtaining a high probability bound on the second eigenvalue of $\widehat{A}_{\boldsymbol{G}} := D_{\boldsymbol{G}}^{-1/2} A_{\boldsymbol{G}} D_{\boldsymbol{G}}^{-1/2}$, the normalized adjacency matrix of link graph $\boldsymbol{G}$, where $A_{\boldsymbol{G}}$ and $D_{\boldsymbol{G}}$ denote its adjacency matrix and

diagonal degree matrix. Since the number of vertices $\boldsymbol{m}$ concentrates well in our setting, throughout this section we treat the number of vertices m as fixed and handle the variation in $\boldsymbol{m}$ in Section 7. We also specialize the parameters to the regime relevant in proving Theorem 1.6 in Section 7 — in particular, the relationship between n, p and d is such that $\lim_{n\to\infty}\tau(p,d)$ is a constant in $(0,1)$, np is a polynomially large function of n, and $d=\Omega(\log n)$.

Theorem 5.1. *Let $0<\tau<1$ be a constant. Let $\boldsymbol{v}_1,\dots,\boldsymbol{v}_m\sim \mathrm{cap}_{\geqslant\tau}(w)$ and $\boldsymbol{G}:=\mathrm{gg}_\tau(\boldsymbol{v}_1,\dots,\boldsymbol{v}_m)$. Then for $q:=\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}\left(\frac{\tau}{1+\tau}\right)$, suppose $qm\gg\log^8 m\cdot\log^{3/2}\frac{1}{q}\cdot\left(\frac{1+\tau}{\tau}\right)^3$ and $d\geqslant C\cdot\log m$ for any constant $C>0$, then for any constant $\gamma>0$,*

$$\mathbf{Pr}\left[|\lambda|_2\left(\widehat{A}_{\boldsymbol{G}}\right)>\frac{\tau}{1+\tau}+o_{d,m}(1)\right]\leqslant O\left(m^{-\gamma}\right).$$

To prove Theorem 5.1, by Fact 2.1 it suffices to bound $\left\|\widehat{A}_{\boldsymbol{G}}-R\right\|$ for any rank-1 PSD matrix R. For a given $\boldsymbol{G}$, the minimizing R for $\left\|\widehat{A}_{\boldsymbol{G}}-R\right\|$ is $R=R_{\boldsymbol{G}}=\frac{D_{\boldsymbol{G}}^{1/2}JD_{\boldsymbol{G}}^{1/2}}{\mathrm{tr}(D_{\boldsymbol{G}})}$ where J is the all-ones matrix.

One challenge in directly performing the trace method on $\widehat{A}_{\boldsymbol{G}}-R_{\boldsymbol{G}}$ is that the degree of any vertex i is a random variable that depends on the locations of all the vectors, and hence introduces extra correlations. In Section 3, this issue was resolved because the degrees concentrated very well, and hence $D_{\boldsymbol{G}}^{-1/2}$ and $D_{\boldsymbol{G}}^{1/2}$ were close to scalar multiples of identity. However, in the links the degrees of vertices in $\boldsymbol{G}$ no longer concentrate around a single value, and even the behavior of the *expected* degree of vertex i depends on which "shell" $\boldsymbol{v}_i$ is contained in around w, $\langle\boldsymbol{v}_i,w\rangle$. To better control the degrees, we will study the spectral norm of $\widehat{A}_{\boldsymbol{G}}-R_{\boldsymbol{G}}|\boldsymbol{\kappa}$ conditioned on the shells $\boldsymbol{\kappa}:=\{\boldsymbol{\kappa}_i:=\langle w,\boldsymbol{v}_i\rangle\}_{i=1}^m$.

Let $D_{\boldsymbol{\kappa}}\in\mathbb{R}^{m\times m}$ be the conditional expected diagonal degree matrix with $D_{\boldsymbol{\kappa}}[i,i]=\mathbf{E}[\deg_{\boldsymbol{G}}(i)\mid\boldsymbol{\kappa}]$. Then we define the new normalized matrix $\underline{A}_{\boldsymbol{G}}=D_{\boldsymbol{\kappa}}{}^{-1/2}A_{\boldsymbol{G}}D_{\boldsymbol{\kappa}}{}^{-1/2}$ and the new conditional rank-1 PSD matrix $R_{\boldsymbol{\kappa}}=\frac{D_{\boldsymbol{\kappa}}{}^{1/2}JD_{\boldsymbol{\kappa}}{}^{1/2}}{\mathrm{tr}(D_{\boldsymbol{\kappa}})}$. Then by optimality of $R_{\boldsymbol{G}}$:

$$\begin{aligned}\left\|\widehat{A}_{\boldsymbol{G}}-R_{\boldsymbol{G}}\right\|&\leqslant\left\|\widehat{A}_{\boldsymbol{G}}-\left(D_{\boldsymbol{G}}^{-1/2}D_{\boldsymbol{\kappa}}{}^{1/2}\right)R_{\boldsymbol{\kappa}}\left(D_{\boldsymbol{\kappa}}{}^{1/2}D_{\boldsymbol{G}}^{-1/2}\right)\right\|\\&=\left\|\left(D_{\boldsymbol{G}}^{-1/2}D_{\boldsymbol{\kappa}}{}^{1/2}\right)\left(\underline{A}_{\boldsymbol{G}}-R_{\boldsymbol{\kappa}}\right)\left(D_{\boldsymbol{\kappa}}{}^{1/2}D_{\boldsymbol{G}}^{-1/2}\right)\right\|\\&\leqslant\left\|\underline{A}_{\boldsymbol{G}}-R_{\boldsymbol{\kappa}}\right\|\cdot\left\|D_{\boldsymbol{G}}^{-1/2}D_{\boldsymbol{\kappa}}{}^{1/2}\right\|^2.\end{aligned}$$

Since $\left\|D_{\boldsymbol{G}}^{-1/2}D_{\boldsymbol{\kappa}}{}^{1/2}\right\|^2=\left\|D_{\boldsymbol{G}}^{-1}D_{\boldsymbol{\kappa}}\right\|$, this is equivalent to bounding

$$\left\|\underline{A}_{\boldsymbol{G}}-R_{\boldsymbol{\kappa}}\right\|\cdot\left\|D_{\boldsymbol{G}}^{-1}D_{\boldsymbol{\kappa}}\right\|\leqslant\left\|\underline{A}_{\boldsymbol{G}}-R_{\boldsymbol{\kappa}}\right\|\cdot\max_{i\in[m]}\frac{D_{\boldsymbol{\kappa}}[i,i]}{D_{\boldsymbol{G}}[i,i]}$$

Now, in the trace method it is convenient to work with $A_{\boldsymbol{G}}-\mathbf{E}[A_{\boldsymbol{G}}]\mid\boldsymbol{\kappa}$, which is not a rank-1 matrix. So, applying the triangle inequality,

$$\leqslant\left(\left\|\underline{A}_{\boldsymbol{G}}-\mathbf{E}[\underline{A}_{\boldsymbol{G}}\mid\boldsymbol{\kappa}]\right\|+\left\|\mathbf{E}[\underline{A}_{\boldsymbol{G}}\mid\boldsymbol{\kappa}]-R_{\boldsymbol{\kappa}}\right\|\right)\cdot\max_{i\in[m]}\frac{D_{\boldsymbol{\kappa}}[i,i]}{D_{\boldsymbol{G}}[i,i]},\tag{21}$$

It then suffices to bound $\left\|\mathbf{E}[\underline{A}_{\boldsymbol{G}}\mid\boldsymbol{\kappa}]-R_{\boldsymbol{\kappa}}\right\|$, $\max_{i\in[m]}\frac{D_{\boldsymbol{\kappa}}[i,i]}{D_{\boldsymbol{G}}[i,i]}$, and $\left\|\underline{A}_{\boldsymbol{G}}-\mathbf{E}[\underline{A}_{\boldsymbol{G}}\mid\boldsymbol{\kappa}]\right\|$ to complete the proof of Theorem 5.1.

In Section 6 we'll show that $\mathbf{E}[\underline{A}_{\boldsymbol{G}}\mid\boldsymbol{\kappa}]$ is close to $R_{\boldsymbol{G}}$ in spectral norm:

Lemma 5.2. *If $d\geqslant C\cdot\log m$ for some constant $C>0$ and the constant $\tau\in(0,1)$ satisfies $qm\gg\log^8 m$, then*

$$\left\|\mathbf{E}[\underline{A}_{\boldsymbol{G}}\mid\boldsymbol{\kappa}]-R_{\boldsymbol{\kappa}}\right\|\leqslant O\left(\sqrt{\frac{\log^2 d}{d}}\right)$$

with probability at least $1-o(m^{-\gamma})$ for any constant $\gamma>0$.

And the remainder of this section will be devoted to bounding the other two quantities, as follows:

Lemma 5.3. *For any* $0 < \alpha < 1$,

$$\max_{i\in[m]} \frac{D_\kappa[i,i]}{D_{\boldsymbol{G}}[i,i]} \leqslant \frac{1}{1-\alpha},$$

with probability at least $1 - m\cdot\exp\left(-\frac{\alpha^2 q(m-1)}{4}\right)$.

Lemma 5.4. *For any* κ, q *and* m *and* $qm \gg \log^8 m \cdot \log^{3/2}\frac{1}{q}\cdot\left(\frac{1+\tau}{\tau}\right)^3$

$$\left\|\underline{A}_{\boldsymbol{G}} - \mathbf{E}\left[\underline{A}_{\boldsymbol{G}} \mid \kappa\right]\right\| \leqslant (1+o_m(1))\cdot\frac{\tau}{1+\tau}.$$

In service of proving Lemma 5.2, Lemma 5.3 and Lemma 5.4, we need the following fact that arises in studying random geometric graphs with shifted edge-connectivity thresholds.

Definition 5.5. We define the bivariate function $T(x,y) := \frac{\tau - xy}{\sqrt{(1-x^2)(1-y^2)}}$ as the *shifted threshold* function, defined so that

$$\mathop{\mathbf{Pr}}_{\boldsymbol{x},\boldsymbol{y}\sim\mathbb{S}^{d-2}}[\langle \boldsymbol{x},\boldsymbol{y}\rangle \geqslant T(x,y)] = \mathop{\mathbf{Pr}}_{\boldsymbol{u},\boldsymbol{v}\sim\mathbb{S}^{d-1}}[\langle \boldsymbol{u},\boldsymbol{v}\rangle \geqslant \tau \mid \langle \boldsymbol{u},w\rangle = x, \langle \boldsymbol{v},w\rangle = y].$$

Claim 5.6. The shifted threshold function $T(x,y) := \frac{\tau - xy}{\sqrt{(1-x^2)(1-y^2)}}$ on the domain $x,y\in[\tau,1]$ is maximized when $x = y = \tau$, and achieves value $\frac{\tau}{1+\tau}$. Additionally $\partial_x T(x,y)$ and $\partial_y T(x,y)$ are both negative.

Proof. The derivatives $\partial_y T(x,y) = \frac{\tau}{\sqrt{1-x^2}}\cdot g(y) - \frac{x}{\sqrt{1-x^2}}\cdot h(y)$ and $\partial_x T(x,y) = \frac{\tau}{\sqrt{1-y^2}}\cdot g(x) - \frac{y}{\sqrt{1-y^2}}\cdot h(x)$, where $g(z) := \frac{z}{(1-z^2)^{3/2}}$ and $h(z) := \frac{1}{\sqrt{1-z^2}} + \frac{z^2}{(1-z^2)^{3/2}}$. Since $g(z) < h(z)$ for $z\in(0,1]$, then for $x,y\geqslant\tau$ we deduce that $\partial_y T, \partial_x T < 0$. Therefore, T achieves the maximum value $\frac{\tau-\tau^2}{1-\tau^2} = \frac{\tau}{1+\tau}$ when $x = y = \tau$. □

Now we prove Lemma 5.3.

Proof of Lemma 5.3. For any $\alpha\in(0,1)$, consider the event that $\max_{i\in[m]}\frac{D_\kappa[i,i]}{D_{\boldsymbol{G}}[i,i]} > \frac{1}{1-\alpha}$. We can bound the probability that this event happens by union bound and Bernstein's inequality:

$$\begin{aligned}\mathbf{Pr}[\exists i\in[m],\, D_{\boldsymbol{G}}[i,i] \leqslant (1-\alpha)D_\kappa[i,i]] &\leqslant \sum_{i=1}^m \mathbf{Pr}[D_{\boldsymbol{G}}[i,i] \leqslant (1-\alpha)D_\kappa[i,i]]\\ &\leqslant m\cdot\max_i \exp\left(-\frac{1}{2}\cdot\frac{\alpha^2 D_\kappa[i,i]^2}{(\alpha+1)D_\kappa[i,i]}\right)\\ &\leqslant m\cdot\max_i\exp\left(-\frac{\alpha^2 D_\kappa[i,i]}{4}\right)\end{aligned}$$

Observe that $D_\kappa[i,i] = \sum_{j\neq i}\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(\kappa_i,\kappa_j))$. By Claim 5.6, $T(\kappa_i,\kappa_j)\leqslant\frac{\tau}{1+\tau}$, so $\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}\left(T(\kappa_i,\kappa_j)\right)\geqslant \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}\left(\frac{\tau}{1+\tau}\right) = q$. Consequently, $D_\kappa[i,i]\geqslant q(m-1)$ from which the desired statement follows. □

5.1 Spectral norm bound for centered links

In the rest of the section, we prove Lemma 5.4 by bounding the expected trace $\mathbf{E}\left[\mathrm{tr}\left(\left(\underline{A}_{\boldsymbol{G}} - \mathbf{E}[\underline{A}_{\boldsymbol{G}}\mid\kappa]\right)^\ell\right)\right]$, for $\kappa\in[\tau,1]^m$ a fixed configuration of shells. The proof will be almost identical to the one in Section 3, but here we have to deal with the fact that the graph is not vertex-transitive.

Proof of Lemma 5.4. First observe:

$$\begin{aligned}\mathbf{E}\left[\operatorname{tr}\left(\left(\underline{A}_{\boldsymbol{G}}-\mathbf{E}[\underline{A}_{\boldsymbol{G}}\mid\kappa]\right)^{\ell}\right)\right]&=\mathbf{E}\left[\operatorname{tr}\left(\left(D_\kappa{}^{-1/2}A_{\boldsymbol{G}}D_\kappa{}^{-1/2}-D_\kappa{}^{-1/2}\mathbf{E}[A_{\boldsymbol{G}}|\kappa]D_\kappa{}^{-1/2}\right)^{\ell}\right)\right]\\&=\mathbf{E}\left[\operatorname{tr}\left(\left(D_\kappa{}^{-1}A_{\boldsymbol{G}}-D_\kappa{}^{-1}\mathbf{E}[A_{\boldsymbol{G}}|\kappa]\right)^{\ell}\right)\right].\end{aligned}$$

We rewrite the expression in terms of $D_\kappa{}^{-1}A_{\boldsymbol{G}}$ which approximates the transition matrix of the random walk on $\boldsymbol{G}$.[12] Next, we expand the expression in terms of walks in $\mathscr{K}_m$.

Following the convention of Section 3, we use $\mathscr{W}_\ell$ to denote the collection of length-ℓ walks in $\mathscr{K}_m$. For every $W\in\mathscr{W}_\ell$, use $G(W)=(V(W),E(W))$ to denote the multigraph obtained by the vertices and edges used in W. Use $m(e)$ to denote the number of times that an edge e appears in the walk W.

Definition 5.7. We also introduce the following notation. Let $d_W(\kappa):=\prod_{(i_t,i_{t+1})\in W}D_\kappa{}^{-1}[i_t,i_t]$ denote the normalization constant along the path W conditioned on the shells κ. Also define $p_e=\mathbf{E}[\mathbf{1}[e\in\boldsymbol{G}]\mid\kappa]$ to be the probability that an edge e exists conditioned on κ.

Then:

$$\mathbf{E}\left[\operatorname{tr}\left(\left(\underline{A}_{\boldsymbol{G}}-\mathbf{E}[\underline{A}_{\boldsymbol{G}}\mid\kappa]\right)^{\ell}\right)\mid\kappa\right]=\sum_{W\in\mathscr{W}_\ell}d_W(\kappa)\cdot\mathbf{E}\left[\prod_{e\in E(W)}(\mathbf{1}[e\in\boldsymbol{G}]-p_e)^{m(e)}\mid\kappa\right]\tag{22}$$

Next we apply the decomposition in Section 3 to $G(W)$ and obtain the 2-core graph $G_2(W)$ and the forest graph $G_1(W)$. Since conditioned on the vectors $\boldsymbol{v}_i\in V_2(W)$ the events $e\in\boldsymbol{G}$ are independent for all $e\in E_1(W)$, the expectation in (22) can be decomposed into two parts:

$$\begin{aligned}&\mathbf{E}\left[\prod_{e\in E(W)}(\mathbf{1}[e\in\boldsymbol{G}]-p_e)^{m(e)}\mid\kappa\right]\\&=\prod_{e\in E_1(W)}\mathbf{E}\left[(\mathbf{1}[e\in\boldsymbol{G}]-p_e)^{m(e)}\mid\kappa\right]\mathop{\mathbf{E}}_{\boldsymbol{v}_i\\ i\in V_2(W)}\left[\prod_{e\in E_2(W)}(\mathbf{1}[e\in\boldsymbol{G}]-p_e)^{m(e)}\mid\kappa\right]\end{aligned}\tag{23}$$

We bound the contribution from the edges in $E_2(W)$ by further spliting $G_2(W)$ into paths consisting of degree-2 vertices and the junction graph $G_J(W)=(J(W),E_J(W))$ as defined in Definition 3.5. As in Section 3 the key observation here is that conditioned on vertices in $J(W)$ the contributions from the paths of degree-2 vertices are all independent from each other:

$$\left|\mathop{\mathbf{E}}_{\boldsymbol{v}_i\\ i\in V_2(W)}\left[\prod_{e\in E_2(W)}(\mathbf{1}[e\in\boldsymbol{G}]-p_e)^{m(e)}\mid\kappa\right]\right|\leqslant\mathop{\mathbf{E}}_{\boldsymbol{v}_i\\ i\in J(W)}\prod_{f\in E_{J(W)}}\left|\mathop{\mathbf{E}}_{\boldsymbol{v}_i\\ i\in\gamma(f)\setminus J(W)}\left[\prod_{e\in\gamma(f)}(\mathbf{1}[e\in\boldsymbol{G}]-p_e)^{m(e)}\mid\kappa\right]\right|\tag{24}$$

Now, let $X_{\kappa,\kappa'}$ be the transition operator for the random step that walks from vector v in $\operatorname{shell}_{=\kappa}(w)$ to a uniformly random vector $\boldsymbol{v}'$ in $\operatorname{shell}_{=\kappa'}(w)\cap\operatorname{cap}_{\geqslant\tau}(v)$. Like in Section 3, we use $\gamma(f)=(f_0,f_1,\ldots,f_{\ell(f)})$ to identify the walk in $G_2(W)$ corresponding to the edge $f\in E_{J(W)}$. We denote the edge (f_i,f_{i+1}) with $\gamma_i(f)$. We simplify the contribution from each path $\gamma(f)$ where $f\in E_J(W)$ as follows:

$$\left|\mathop{\mathbf{E}}_{\boldsymbol{v}_i:i\in\gamma(f)\setminus J(W)}\left[\prod_{e\in\gamma(f)}(\mathbf{1}[e\in\boldsymbol{G}]-p_e)^{m(e)}\mid\kappa\right]\right|$$

[12] If D_κ were not the *expected* degree matrix but rather the exact degree matrix of $\boldsymbol{G}$, we would have a true transition matrix here.

$$
\begin{aligned}
&= \left| \mathop{\mathbf{E}}_{\boldsymbol{v}_i : i \in \gamma(f) \setminus J(W)} \left[\prod_{e \in \gamma(f)} \left(\mathbf{1}[e \in \boldsymbol{G}] \cdot \left((1-p_e)^{m(e)} - (-p_e)^{m(e)} \right) + (-p_e)^{m(e)} \right) \mid \kappa \right] \right| \\
&= \left| \sum_{T \subseteq \gamma(f)} \mathop{\mathbf{E}}_{\boldsymbol{v}_i : i \in \gamma(f) \setminus J(W)} \left[\prod_{e \in T} \mathbf{1}[e \in \boldsymbol{G}] \cdot \left((1-p_e)^{m(e)} - (-p_e)^{m(e)} \right) \prod_{e \in \gamma(f) \setminus T} (-p_e)^{m(e)} \mid \kappa \right] \right| \\
&\leqslant \left| \prod_{e \in \gamma(f)} \left(p_e (1-p_e)^{m(e)} + (1-p_e)(-p_e)^{m(e)} \right) \right| \\
&\quad + \left| \prod_{e \in \gamma(f)} \left((1-p_e)^{m(e)} - (-p_e)^{m(e)} \right) \cdot \prod_{i=0}^{\ell(f)-2} p_{\gamma_i(f)} \cdot \left(\left\langle \prod_{i=0}^{\ell(f)-2} X_{\kappa_{f_i}, \kappa_{f_{i+1}}} \delta_{\boldsymbol{v}_{f_0}}, \mathrm{cap}_{p_{\gamma_{\ell(f)-1}(f)}}(\boldsymbol{v}_{f_{\ell(f)}}) \right\rangle - p_{\gamma_{\ell(f)-1}(f)} \right) \right|
\end{aligned} \tag{25}
$$

Let $S_J(W) \subseteq E_J(W)$ be the set of edges f such $m(e) = 1$ for some $e \in \gamma(f)$, and $D_J(W) = E_J(W) \setminus S_J(W)$. For any $f \in S_J(W)$ the first term in (25) vanishes, while for any $f \in D_J(W)$

$$
\left| \prod_{e \in \gamma(f)} \left(p_e (1-p_e)^{m(e)} + (1-p_e)(-p_e)^{m(e)} \right) \right| \leqslant \prod_{e \in \gamma(f)} \left(p_e (1-p_e)^2 + (1-p_e) p_e^2 \right) \leqslant \prod_{e \in \gamma(f)} p_e .
$$

Therefore we can derive the following bound on the contribution from the 2-core graph.

$$
\begin{aligned}
(24) \leqslant \mathop{\mathbf{E}}_{\substack{\boldsymbol{v}_i \\ i \in J(W)}} \Bigg[& \prod_{f \in D_J(W)} \prod_{i=0}^{\ell(f)-2} p_{\gamma_i(f)} \cdot \left(\left| \left\langle \prod_{i=0}^{\ell(f)-2} X_{\kappa_{f_i}, \kappa_{f_{i+1}}} \delta_{\boldsymbol{v}_{f_0}}, \mathrm{cap}_{p_{\gamma_{\ell(f)-1}(f)}}(\boldsymbol{v}_{f_{\ell(f)}}) \right\rangle - p_{\gamma_{\ell(f)-1}(f)} \right| + p_{\gamma_{\ell(f)-1}(f)} \right) \cdot \\
& \prod_{f \in S_J(W)} \prod_{i=0}^{\ell(f)-2} p_{\gamma_i(f)} \cdot \left| \left\langle \prod_{i=0}^{\ell(f)-2} X_{\kappa_{f_i}, \kappa_{f_{i+1}}} \delta_{\boldsymbol{v}_{f_0}}, \mathrm{cap}_{p_{\gamma_{\ell(f)-1}(f)}}(\boldsymbol{v}_{f_{\ell(f)}}) \right\rangle - p_{\gamma_{\ell(f)-1}(f)} \right| \mid \kappa \Bigg]
\end{aligned}
$$

To bound the absolute value terms, we take an arbitrary spanning tree $T_J(W)$ of $G_J(W)$, and bound the absolute value differently depending on whether $f \in T_J(W)$ or not.

To bound this expectation, let $T_J(W)$ be a spanning tree of $G_J(W)$. For every edge not in $T_J(W)$, we apply a worst-case bound. To state this bound, we define $C := \sqrt{\frac{1}{2} \log \frac{1}{q}} \cdot \left(\frac{1+\tau}{\tau} \right)$ and $\lambda := \frac{\tau}{1+\tau}$.

Claim 5.8. For every shell configuration $\kappa \in [\tau, 1]^m$ and non-tree edge $f \in E_J(W) \setminus T_J(W)$, we have that

$$
\left| \left\langle \prod_{i=0}^{\ell(f)-2} X_{\kappa_{f_i}, \kappa_{f_{i+1}}} \delta_{\boldsymbol{v}_{f_0}}, \mathrm{cap}_{p_{\gamma_{\ell(f)-1}(f)}}(\boldsymbol{v}_{f_{\ell(f)}}) \right\rangle - p_{\gamma_{\ell(f)-1}(f)} \right| \leqslant C \cdot \lambda^{\ell(f)-1}
$$

Proof. To prove the claim, we first need to understand the random variable

$$
\left\langle \prod_{i=0}^{\ell(f)-2} X_{\kappa_{f_i}, \kappa_{f_{i+1}}} \delta_{\boldsymbol{v}_{f_0}}, \mathrm{cap}_{p_{\gamma_{\ell(f)-1}(f)}}(\boldsymbol{v}_{f_{\ell(f)}}) \right\rangle .
$$

Recall that at time step i the operator $X_{\kappa_{f_i}, \kappa_{f_{i+1}}}$ denotes the random step that takes a vector $\boldsymbol{v}_{f_i} = \kappa_{f_i} \cdot w + \sqrt{1 - \kappa_{f_i}^2} \cdot \boldsymbol{z}_i$ and outputs $\boldsymbol{v}_{f_{i+1}} := \kappa_{f_{i+1}} \cdot w + \sqrt{1 - \kappa_{f_{i+1}}^2} \cdot \boldsymbol{z}_{i+1}$ where $\boldsymbol{z}_{i+1}$ is a uniformly random unit vector orthogonal to w such that

$$
\langle \boldsymbol{v}_{f_i}, \boldsymbol{v}_{f_{i+1}} \rangle \geqslant \tau .
$$

This is equivalent to

$$
\kappa_{f_i} \kappa_{f_{i+1}} + \sqrt{(1 - \kappa_{f_i}^2)(1 - \kappa_{f_{i+1}}^2)} \cdot \langle \boldsymbol{z}_i, \boldsymbol{z}_{i+1} \rangle \geqslant \tau ,
$$

which can then be rearranged as

$$\langle \boldsymbol{z}_i, \boldsymbol{z}_{i+1} \rangle \geqslant T\big(\kappa_{f_i}\kappa_{f_{i+1}}\big) := \frac{\tau - \kappa_{f_i}\kappa_{f_{i+1}}}{\sqrt{(1-\kappa_{f_i}^2)(1-\kappa_{f_{i+1}}^2)}}.$$

In particular, we are choosing $\boldsymbol{z}_{i+1}$ in the $p_{\gamma_i(f)}$-cap of $\boldsymbol{z}_i$ within the $d-2$ dimensional unit sphere orthogonal to w. So the operator $\prod_{i=0}^{\ell(f)-2} X_{\kappa_{f_i},\kappa_{f_{i+1}}}$ can be decomposed into its action in the span of w and that in the space orthogonal to w. The action in the span of w conditioned on κ is deterministic. Orthogonal to w, it is the operator $\prod_{i=0}^{\ell(f)-2} P_{p_{\gamma_i(f)}}$ on $\mathbb{S}^{d-2}$. Thus the quantity we are interested in understanding is the same as

$$\left\langle \prod_{i=0}^{\ell(f)-2} P_{p_{\gamma_i(f)}} \delta_{\boldsymbol{z}_0}, \mathrm{cap}_{p_{\gamma_{\ell(f)-1}}}(\boldsymbol{z}_{\ell(f)}) \right\rangle.$$

Now, observe that:

$$\left| \left\langle \prod_{i=0}^{\ell(f)-2} P_{p_{\gamma_i(f)}} \delta_{\boldsymbol{z}_0}, \mathrm{cap}_{p_{\gamma_{\ell(f)-1}}}(\boldsymbol{z}_{\ell(f)}) \right\rangle - p_{\gamma_{\ell(f)-1}} \right| \leqslant \mathrm{d_{TV}}\left(\prod_{i=0}^{\ell(f)-2} P_{p_{\gamma_i(f)}} \delta_{\boldsymbol{z}_0}, \rho \right).$$

Recall that $p_{\gamma_i(f)} = \Phi_{\mathsf{D_{ip}}(d-1)}\big(\tau_{\kappa_i,\kappa_{i+1}}\big)$, which by Claim 5.6 is minimized when $\tau_{\kappa_i,\kappa_{i+1}} = \frac{\tau}{1+\tau}$, which means $p_{\gamma_i(f)} \geqslant q$. Thus, by Claim 4.11, Lemma 4.13, and Lemma 4.12, which make concrete the intuition that applying P_q should only mix slower than applying $P_{q'}$ for $q' \geqslant q$, we have:

$$\left| \left\langle \prod_{i=0}^{\ell(f)-2} P_{p_{\gamma_i(f)}} \delta_{\boldsymbol{z}_0}, \mathrm{cap}_{p_{\gamma_{\ell(f)-1}}}(\boldsymbol{z}_{\ell(f)}) \right\rangle - p_{\gamma_{\ell(f)-1}} \right| \leqslant \mathrm{d_{TV}}\left(P_q^{\ell(f)-1} \delta_{\boldsymbol{z}_0}, \rho \right) = \mathrm{d_{TV}}\left(\frac{1}{q} P_q^{\ell(f)-2} \mathrm{cap}_q \boldsymbol{z}_0, \rho \right).$$

Then by Theorem 4.6, the above is

$$\leqslant \sqrt{\frac{1}{2}\log\frac{1}{q}} \cdot \left(\frac{\tau}{1+\tau}\right)^{\ell(f)-2} = \sqrt{\frac{1}{2}\log\frac{1}{q}} \cdot \frac{1+\tau}{\tau} \cdot \left(\frac{\tau}{1+\tau}\right)^{\ell(f)-1} = C \cdot \lambda^{|\gamma(f)|-1},$$

which completes the proof. □

Next we bound the contribution of a tree edge $f \in T_J(W)$ using the following claim whose proof is identical to that of Claim 3.7.

Claim 5.9. For every shell vector κ and tree edge $f \in T_J(W)$, we have that

$$\mathop{\mathbf{E}}_{\substack{\boldsymbol{v}_i \\ i \in J(W)}} \prod_{f \in T_J(W)} \prod_{i=0}^{\ell(f)-2} p_{\gamma_i(f)} \cdot \left(\left| \left\langle \prod_{i=0}^{\ell(f)|-1} X_{\kappa_{f_i},\kappa_{f_{i+1}}} \delta_{\boldsymbol{v}_{f_0}}, \mathrm{cap}_{p_{\gamma_{\ell(f)-1}(f)}}(\boldsymbol{v}_{f_{\ell(f)}}) \right\rangle - p_{\gamma_{\ell(f)-1}(f)} \right| + p_{\gamma_{\ell(f)-1}(f)} \cdot \mathbf{1}[f \in D_J(W)] \right)$$
$$\leqslant \prod_{f \in T_J(W)} \prod_{i=0}^{\ell(f)-1} p_{\gamma_i(f)} \cdot \Big(2C\lambda^{\ell(f)} + \mathbf{1}[f \in D_J(W)] \Big).$$

Combining the two bounds for different edges in $G_J(W)$ to obtain the simplified bound for (24):

$$(24) \leqslant \prod_{f \in T_J(W)} \prod_{i=0}^{\ell(f)-1} p_{\gamma_i(f)} \Big(2C\lambda^{\ell(f)} + \mathbf{1}[f \in D_J(W)] \Big) \cdot$$
$$\cdot \prod_{f \in E_J(W) \setminus T_J(W)} \prod_{i=0}^{\ell(f)-2} p_{\gamma_i(f)} \cdot \Big(C\lambda^{\ell(f)-1} + p_{\gamma_{\ell(f)-1}(f)-1} \mathbf{1}[f \in D_J(W)] \Big)$$

We now recall some notation from Section 3. We use $e(W)$ to denote $|E(W)|$, $\mathrm{sing}(W)$ for the number of singleton edges in $G_2(W)$, and $\mathrm{exc}(G)$ for $|E(G)|-(|V(G)|-1)$, the number of edges G has more than a tree. The relations between these variable are already shown in Observation 3.9 and Claim 3.10. So here we directly apply these results to get that

$$\begin{aligned}(24) &\leqslant \prod_{e\in E_2(W)} p_e \cdot \prod_{f\in E_J(W)\setminus T_J(W)} p^{-1}_{\gamma_{\ell(f)-1}(f)} \cdot \lambda^{\mathrm{sing}(W)-\mathrm{exc}(W)} \cdot (3C)^{|E_J(W)|} \\ &\leqslant \prod_{e\in E_2(W)} p_e \cdot \prod_{f\in E_J(W)\setminus T_J(W)} p^{-1}_{\gamma_{\ell(f)-1}(f)} \cdot \lambda^{\mathrm{sing}(W)-\mathrm{exc}(W)} \cdot (3C)^{3\mathrm{exc}(W)} \quad \text{by Observation 3.9}\end{aligned}$$

Therefore

$$\begin{aligned}(23) &\leqslant \prod_{e\in E_1(W)} \mathbf{E}\left[(\mathbf{1}[e\in \boldsymbol{G}]-p_e)^{m(e)} \mid \kappa\right] \cdot \prod_{e\in E_2(W)} p_e \cdot \prod_{f\in E_J(W)\setminus T_J(W)} p^{-1}_{\gamma_{\ell(f)-1}(f)} \cdot \lambda^{\mathrm{sing}(W)-\mathrm{exc}(W)} \cdot (3C)^{3\mathrm{exc}(W)} \\ &\leqslant \prod_{e\in E_1(W)} p_e \cdot \prod_{e\in E_2(W)} p_e \cdot \prod_{f\in E_J(W)\setminus T_J(W)} p^{-1}_{\gamma_{\ell(f)-1}(f)} \cdot \lambda^{\mathrm{sing}(W)-\mathrm{exc}(W)} \cdot (3C)^{3\mathrm{exc}(W)} \\ &\leqslant \prod_{e\in E(W)} p_e \cdot \prod_{f\in E_J(W)\setminus T_J(W)} p^{-1}_{\gamma_{\ell(f)-1}(f)} \cdot \lambda^{\mathrm{sing}(W)} \left(\frac{27C^3}{\lambda}\right)^{\mathrm{exc}(W)} \qquad (26)\end{aligned}$$

Before finally bounding the trace power, we make the following observations.

Observation 5.10. As a consequence of Claim 5.6 for all $i\in[m]$, the expected degree of vertex i satisfies

$$D_\kappa[i,i] = \mathbf{E}[\deg_{\boldsymbol{G}}(i)\mid\kappa] \geqslant (m-1)\cdot q$$

We define Struc_ℓ to be the set of distinct unlabelled walks of length ℓ. Then as a corollary of Claim 3.10, we have

Corollary 5.11. *The number of unlabelled walks $U\in\mathsf{Struc}_\ell$ such that $e(U)=a, s(U)=b$, and $\mathrm{exc}(U)=c$ is at most:*

$$\ell^{3(\ell-b)}\cdot\ell^{2c}.$$

The result follows by observing that since U is unlabelled, we can remove the m^{a-c+1} term that counts the number of distinct labelings in Claim 3.10.

For an unlabeled walk U and labeled walk W, we say $W\sim U$ if W is a labeling of U in $[m]$.

Claim 5.12. For any unlabelled walk U we have that

$$\sum_{W\sim U} d_W(\kappa) \prod_{e\in E(W)} p_e \cdot \prod_{f\in E_J(W)\setminus T_J(W)} p^{-1}_{\gamma_{\ell(f)-1}(f)} \leqslant (m\cdot q)^{-\mathrm{exc}(U)-\frac{\ell-s(U)}{2}}$$

Proof. For each $W\sim U$ we use $i_1,\dots,i_a\in[m]$ to denote the label of each vertex in W in the order of visit. Then we construct the canonical spanning tree $T(W)$ by adding each directed edge in the order of W as long as the edge goes to an unvisited vertex. Use $\mathrm{Par}(i_j)$ to denote the parent of vertex i_j. Then the j-th edge of $T(W)$ is $(\mathrm{Par}(i_{j+1}), i_{j+1})$, and use $T(W)^{(j)}$ to denote the tree consisting of the first j edges of $T(W)$. Then i_{j+1} is always a leaf in $T(W)^{(j)}$.

$T(W)$ gives rise to a canonical spanning tree $T_J(W)$ in the contracted graph $G_J(W)$: an edge f is in $T_J(W)$ if and only if every edge in the path $\gamma(f)$ is in $T(W)$. From this fact we can deduce that

$$T(W) = E(W)\setminus\left\{\gamma_{\ell(f)-1}(f) : f\in E_J(W)\setminus T_J(W)\right\}.$$

Therefore, using Observation 5.10 we can take a loose upper bound on the contribution of edges outside of $T(W)$ and write

$$\begin{aligned}
&\sum_{W\sim U} d_W(\kappa) \prod_{e\in E(W)} p_e \cdot \prod_{f\in E_J(W)\setminus T_J(W)} p^{-1}_{\gamma_{\ell(f)-1}(f)} \\
&\leqslant \sum_{W\sim U} ((m-1)\cdot q)^{-\ell+(|V(U)|-1)} \prod_{(i,j)\in T(W)} \frac{p_{i,j}}{D_\kappa[i,i]} \\
&\leqslant ((m-1)\cdot q)^{-\ell+(|V(U)|-1)} \sum_{i_1,\dots,i_{a\in[m]}} \prod_{j=2}^{a} \frac{p_{\mathrm{Par}(i_j),i_j}}{D_\kappa[\mathrm{Par}(i_j),\mathrm{Par}(i_j)]} \qquad (27)
\end{aligned}$$

where $a = |V(U)|$. Next we show by induction on a that

$$\sum_{i_1,\dots,i_a} \prod_{j=2}^{a} \frac{p_{(\mathrm{Par}(i_j),i_j)}}{D_\kappa[\mathrm{Par}(i_j),\mathrm{Par}(i_j)]} = 1$$

The base case $a = 1$ is true by definition. Suppose this is true for $a-1$. Then:

$$\sum_{i_1,\dots,i_a} \prod_{j=2}^{a} \frac{p_{(\mathrm{Par}(i_j),i_j)}}{D_\kappa[\mathrm{Par}(i_j),\mathrm{Par}(i_j)]} = \sum_{i_1,\dots,i_{a-1}} \prod_{j=2}^{a-1} \frac{p_{(\mathrm{Par}(i_j),i_j)}}{D_\kappa[\mathrm{Par}(i_j),\mathrm{Par}(i_j)]} \cdot \sum_{i_a=1}^{m} \frac{p_{(\mathrm{Par}(i_a),i_a)}}{D_\kappa[\mathrm{Par}(i_a),\mathrm{Par}(i_a)]}$$

By definition $\sum_{i_a} \frac{p_{(\mathrm{Par}(i_a),i_a)}}{D_\kappa[\mathrm{Par}(i_a),\mathrm{Par}(i_a)]} = 1$, so we have:

$$= \sum_{i_1,\dots,i_{a-1}} \prod_{j=2}^{a-1} \frac{p_{(\mathrm{Par}(i_j),i_j)}}{D_\kappa[\mathrm{Par}(i_a),\mathrm{Par}(i_a)]} \cdot 1 = 1$$

Finally, observe that $\ell - (|V(U)|-1) \geqslant \mathrm{exc}(U) + \frac{\ell - s(U)}{2}$ is at least the number of steps that use a previously walked-on edge. The way to see this is to observe that the quantity $\ell - (|V(U)|-1)$ counts the number of steps to a previously visited vertex. Such a step can either (1) use an excess edge for the first time, of which there are $\mathrm{exc}(U)$ steps, or (2) use a previously walked-on edge, which must be at least half the steps that do not use a singleton edge, i.e. at least $\frac{\ell - s(U)}{2}$ steps. Thus we conclude that (27) $\leqslant ((m-1)\cdot q)^{-\ell+(|V(U)|-1)} \leqslant ((m-1)\cdot q)^{-\mathrm{exc}(U)-\frac{\ell-s(U)}{2}}$. □

Now we are finally already to bound the expected trace power. Plugging (26) into (22) gives:

$$\begin{aligned}
(22) &= \sum_{U\in\mathsf{Struc}_\ell} \sum_{W\sim U} d_W(\kappa) \prod_{e\in E(W)} p_e \cdot \prod_{f\in E_J(W)\setminus T_J(W)} p^{-1}_{\gamma_{\ell(f)-1}(f)} \cdot \lambda^{\mathrm{sing}(W)} \left(\frac{27C^3}{\lambda}\right)^{\mathrm{exc}(W)} \\
&\leqslant \sum_{U\in\mathsf{Struc}_\ell} ((m-1)\cdot q)^{-\mathrm{exc}(U)-\frac{\ell-s(U)}{2}} \cdot \lambda^{s(U)} \left(\frac{27C^3}{\lambda}\right)^{\mathrm{exc}(U)} && \text{by Claim 5.12} \\
&= \sum_{a=1}^{\ell}\sum_{b=1}^{\ell}\sum_{c=1}^{\ell} \sum_{\substack{U\in\mathsf{Struc}_\ell \\ e(U)=a, s(U)=b, \mathrm{exc}(U)=c}} ((m-1)\cdot q)^{-\frac{\ell-b}{2}} \cdot \lambda^b \left(\frac{27C^3}{\lambda q(m-1)}\right)^c \\
&= \sum_{a=1}^{\ell}\sum_{b=1}^{\ell}\sum_{c=1}^{\ell} ((m-1)\cdot q)^{-\frac{\ell-b}{2}} \cdot \lambda^b \left(\frac{27C^3}{\lambda q(m-1)}\right)^c \cdot \ell^{2(\ell-b)} \cdot \ell^{2c} && \text{by Claim 5.11} \\
&= \ell \sum_{b=1}^{\ell}\sum_{c=1}^{\ell} \left(\frac{\ell^2}{\sqrt{(m-1)\cdot q}}\right)^{\ell-b} \cdot \lambda^b \left(\frac{27C^3\ell^2}{\lambda q(m-1)}\right)^c \\
&= \ell^3 \max\left(1, \left(\frac{27C^3\ell^2}{\lambda q(m-1)}\right)^{\ell}\right) \cdot \max\left(\lambda^\ell, \left(\frac{\ell^2}{\sqrt{(m-1)\cdot q}}\right)^{\ell}\right)
\end{aligned}$$

By choosing $\ell = \log^2 m$, we can conclude that with probability at least $1 - m^{-\gamma}$,

$$\left\| \underline{A_G} - \mathbf{E}[\underline{A_G} \mid \kappa] \right\| \leqslant (1 + o(1)) \cdot \left(1 + \frac{27C^3 \log^4 m}{\lambda q m}\right) \cdot \max\left\{\lambda, \frac{\log^4 m}{\sqrt{qm}}\right\}$$

Since $qm \gg \log^8 m \cdot \log^{3/2} \frac{1}{q} \cdot \left(\frac{1+\tau}{\tau}\right)^3$, we have:

$$\left\| \underline{A_G} - \mathbf{E}[\underline{A_G} \mid \kappa] \right\| \leqslant (1 + o(1)) \cdot \lambda. \qquad \square$$

6 The second eigenvalue of the shell walk

The goal of this section is to prove Lemma 5.2, and in particular bound $\left\| \mathbf{E}[\underline{A_G} \mid \boldsymbol{\kappa}] - R_{\boldsymbol{\kappa}} \right\|$ where $\boldsymbol{\kappa} \sim (\mathsf{D}_{\mathsf{ip}}(d)_{\geqslant \tau})^{\otimes m}$ is a configuration of m shells, and we have conditioned on $\kappa_i = \langle w, \boldsymbol{v}_i \rangle$ for all $i \in [m]$.

To make the matrix more amenable to analysis via the coupling-based techniques we use here, we first observe that the spectral norm we are interested in bounding is equal to the largest eigenvalue of $\mathbf{E}[D_{\boldsymbol{\kappa}}{}^{-1} A_{\boldsymbol{G}} \mid \boldsymbol{\kappa}] - \vec{1}\boldsymbol{\pi}^\top$ where $\boldsymbol{\pi} := \frac{D_{\boldsymbol{\kappa}}}{\mathrm{tr}(D_{\boldsymbol{\kappa}})}\vec{1}$ is the stationary distribution of the Markov chain described by the transition matrix $\mathbf{E}[D_{\boldsymbol{\kappa}}{}^{-1} A_{\boldsymbol{G}} \mid \boldsymbol{\kappa}]$. Indeed:

$$\left\| \mathbf{E}[\underline{A_G} \mid \boldsymbol{\kappa}] - R_{\boldsymbol{\kappa}} \right\| = |\lambda|_{\max}\left(\mathbf{E}[\underline{A_G} \mid \boldsymbol{\kappa}] - R_{\boldsymbol{\kappa}}\right) = |\lambda|_{\max}\left(\mathbf{E}[D_{\boldsymbol{\kappa}}{}^{-1} A_{\boldsymbol{G}} \mid \boldsymbol{\kappa}] - 1\boldsymbol{\pi}^\top\right)$$

where the first equality uses symmetry of the matrix and the second equality uses the fact that the spectra of M and $D_{\boldsymbol{\kappa}}{}^{-1/2} M D_{\boldsymbol{\kappa}}{}^{1/2}$ are identical. For convenience, let $Q = \mathbf{E}[A_{\boldsymbol{G}} \mid \boldsymbol{\kappa}]$ and let $\underline{Q} = D_{\boldsymbol{\kappa}}{}^{-1} Q$. The following main result of this section implies Lemma 5.2.

Lemma 6.1. *There exists a constant $C > 0$ such that for any $d \geqslant C \log m$, any threshold $\tau \in (0,1)$ such that $q := \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}\left(\frac{\tau}{1+\tau}\right) \gg \log^8 m/m$, and any constant $\gamma > 0$, with probability at least $1 - o(m^{-\gamma})$ over the shells $\boldsymbol{\kappa} \sim (\mathsf{D}_{\mathsf{ip}}(d)_{\geqslant \tau})^{\otimes m}$,*

$$|\lambda|_{\max}\left(\underline{Q} - \vec{1}\boldsymbol{\pi}^\top\right) \leqslant O\left(\sqrt{\frac{\log^2 d}{d}}\right).$$

In service of proving Lemma 6.1, we show:

Lemma 6.2. *There exists a constant $C > 0$ such that for any $d \geqslant C \log m$, any threshold $\tau \in (0,1)$ such that $qm \gg \log^8 m$, and any constant $\gamma > 0$, with probability at least $1 - o(m^{-\gamma})$ over the shells $\boldsymbol{\kappa} \sim (\mathsf{D}_{\mathsf{ip}}(d)_{\geqslant \tau})^{\otimes m}$,*

$$\max_{i,j \in [n]} \left\| \left(\underline{Q}^2\right)_{i,*} - \left(\underline{Q}^2\right)_{j,*} \right\|_1 \leqslant O\left(\frac{\log^2 d}{d}\right),$$

where $\left(\underline{Q}^2\right)_{i,*}$ denotes the i-th row of the matrix $\underline{Q}^2$.

We show how to prove Lemma 6.1 using Lemma 6.2 and then dedicate the rest of the section to proving Lemma 6.2.

Proof of Lemma 6.1. First, observe that $|\lambda|_{\max}\left(\underline{Q} - \vec{1}\boldsymbol{\pi}^\top\right) = \sqrt{|\lambda|_{\max}\left(\underline{Q}^2 - \vec{1}\boldsymbol{\pi}^\top\right)}$. Via the row sum bound for the largest magnitude eigenvalue of a matrix (Claim 2.3), Lemma 6.2 and the fact that $\boldsymbol{\pi}$ is the stationary distribution of $\underline{Q}$:

$$|\lambda|_{\max}\left(\underline{Q}^2 - \vec{1}\boldsymbol{\pi}^\top\right) \leqslant \max_{i \in [n]} \left\| \left(\underline{Q}^2\right)_{i,*} - \boldsymbol{\pi}^\top \right\|_1$$

$$
\begin{aligned}
&= \max_{i\in[n]} \left\| \left(\underline{Q}^2\right)_{i,*} - \sum_{j\in[n]} \boldsymbol{\pi}_j \left(\underline{Q}^2\right)_{j,*} \right\|_1 \\
&= \max_{i\in[n]} \left\| \sum_{j\in[n]} \boldsymbol{\pi}_j \left(\left(\underline{Q}^2\right)_{i,*} - \left(\underline{Q}^2\right)_{j,*} \right) \right\|_1 \\
&\leqslant \max_{i,j\in[n]} \left\| \left(\underline{Q}^2\right)_{i,*} - \left(\underline{Q}^2\right)_{j,*} \right\|_1 .
\end{aligned}
$$

□

6.1 Coupling for the shell walk

In this section we give the proof of Lemma 6.2 assuming a few key lemmas. The proofs for the key lemmas are deferred to the next section.

6.1.1 A high-probability condition for $\boldsymbol{\kappa}$

To simplify the upcoming computations for Lemma 6.2, we will condition on the following high-probability event $\mathscr{E}_\gamma$ over the sample space of the shells $\boldsymbol{\kappa}$:

Definition 6.3. Let $\mathscr{E}_\gamma$ be the event that for all m shells $\boldsymbol{\kappa}_i \in \boldsymbol{\kappa}$ in the link, $\boldsymbol{\kappa}_i \leqslant \eta$, where

$$\eta = \tau\left(m^{-2\gamma-1} \cdot \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}(\tau), d\right).$$

Note that the outermost $\tau(\cdot)$ refers to the threshold function, rather than the value of the threshold such that $\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}(\tau) = p$.

Claim 6.4. The event $\mathscr{E}_\gamma$ occurs with probability at least $1 - m^{-2\gamma}$.

Proof. By definition, for any shell $\boldsymbol{\kappa}_i$: $\mathbf{Pr}_{\boldsymbol{\kappa}_i \sim \mathsf{D}_{\mathsf{ip}}(d)|_{\geqslant\tau}}[\boldsymbol{\kappa}_i \geqslant \eta] \leqslant m^{-2\gamma-1}$. Our conclusion follows from taking a union bound over all m shells. □

The conditioning on $\mathscr{E}_\gamma$ can be folded into high-probability guarantee over $\boldsymbol{\kappa}$ in Lemma 6.2. Thus, for the remainder of the section, we can assume that $\boldsymbol{\kappa}$ obeys event $\mathscr{E}_\gamma$. This will be especially relevant in the analysis of the outlier shells (Section 6.2.2).

Claim 6.5. If $d \geqslant C \log m$ for some constant $C > 0$, then $\eta \leqslant 1 - \varepsilon_\gamma$, where $\varepsilon_\gamma > 0$ is a constant depending only on γ.

Proof. Since τ is a constant bounded away from 1, and $d = \Omega(\log m)$, by the lower bound in Lemma 2.8, the quantity $m^{-2\gamma-1} \cdot \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}(\tau)$ is at least $\exp(-C_\gamma d)$ for some constant C_γ depending on γ. By the upper bound in Lemma 2.8, there is a constant $\varepsilon_\gamma > 0$ such that $\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}(1-\varepsilon_\gamma) \leqslant \exp(-C_\gamma d)$. Since $\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}$ is a decreasing function, $\eta \leqslant 1 - \varepsilon_\gamma$. □

6.1.2 "Typical" and "outlier" shells

In the proof of Lemma 6.2, we analyze the contributions of "typical" and "outlier" shells separately.

Definition 6.6. We say that a shell κ_i is "typical" if $\kappa_i \in [\tau, \tau(1+\alpha)]$, for $\alpha = \frac{36 \log d}{\tau^2(d-3)(1-\eta)}$.

Remark 6.7. α is chosen so that $\underline{Q}$, when restricted to typical rows and columns, will resemble a rank-1 matrix. For our eventual choices of d and m, the event that every shell is typical *does not* occur with high probability; we will inevitably need to deal with outlier shells.

6.1.3 Total variation bound from similarity of typical rows and scarcity of outlier columns

To obtain the desired row-sum bound in Lemma 6.2, we will prove the following two lemmas about the matrix $\underline{Q}$. The first shows that outlier columns do not contribute much to the total row sum of any row:

Lemma 6.8. *For any $d \geq C \log m$ for some constant $C > 0$ and any $\tau \in [0,1]$ such that $qm \gg \log^8 m$, if $\kappa_i \leq \eta$,*

$$\sum_{k=1}^{m} \underline{Q}_{i,k} \cdot \mathbf{1}[k \text{ outlier}] \leq O\left(\frac{1}{d}\right)$$

with probability $1 - o(m^{-\log m})$.

The second shows that typical rows are similar at indices corresponding to typical columns:

Lemma 6.9. *For any dimension d and any threshold $\tau \in (0,1)$, if κ_i, κ_j correspond to typical shells, then for all ℓ such that κ_ℓ is typical,*

$$\underline{Q}_{i,\ell} \in \left(1 \pm O\left(\frac{\log^2 d}{d}\right)\right) \underline{Q}_{j,\ell}$$

These lemmas are both proven by direct calculation, and we leave their proofs to Section 6.2.2 and Section 6.2.3 respectively.

To illustrate these statements, we provide a schematic of the matrix Q below, organized into its typical and outlier rows and columns. Lemma 6.9 states that the sub-rows in area (I) of the matrix are all nearly equal to each other. Lemma 6.8 says that the sum of its entries in area (II) or area (IV) is a $O\left(\frac{1}{d}\right)$ fraction of the total row sum.

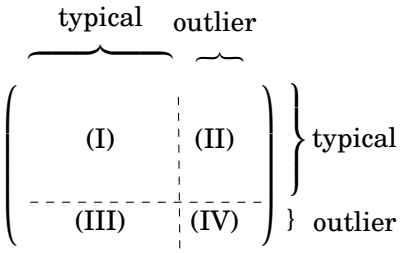

One straightforward corollary of Lemma 6.9 and Lemma 6.8 is that the ℓ_1 norms of the differences between any two *typical* rows of $\underline{Q}$ is at most $O\left(\frac{\log^2 d}{d}\right)$. More formally:

Corollary 6.10. *For any $d \geq C \log m$ for some constant $C > 0$ and any threshold $0 < \tau \leq 1$ such that $qm \gg \log^8 m$, let i, j be rows of $\underline{Q}$ corresponding to typical shells κ_i, κ_j. Then:*

$$\left\| \left(\underline{Q}\right)_{i,*} - \left(\underline{Q}\right)_{j,*} \right\|_1 \leq O\left(\frac{\log^2 d}{d}\right)$$

with probability $1 - o(m^{-\log m})$.

Proof. We split $\left\| \left(\underline{Q}\right)_{i,*} - \left(\underline{Q}\right)_{j,*} \right\|_1$ based on its contributions from typical columns and outlier columns.

$$\left\| \left(\underline{Q}\right)_{i,*} - \left(\underline{Q}\right)_{j,*} \right\|_1 = \sum_{\ell \text{ typical}} \left| \underline{Q}_{i,\ell} - \underline{Q}_{j,\ell} \right| + \sum_{\ell \text{ outlier}} \left| \underline{Q}_{i,\ell} - \underline{Q}_{j,\ell} \right|$$

Lemma 6.8 and the triangle inequality tell us that with probability $1-o(m^{-\log m})$:

$$\sum_{\ell \text{ outlier}} \left|\underline{Q}_{i,\ell} - \underline{Q}_{j,\ell}\right| \leqslant \sum_{\ell \text{ outlier}} \left|\underline{Q}_{i,\ell}\right| + \sum_{\ell \text{ outlier}} \left|\underline{Q}_{j,\ell}\right| \leqslant O\left(\frac{1}{d}\right)$$

Lemma 6.9 tells us that for some constant $C > 0$:

$$\begin{aligned}\sum_{\ell \text{ typical}} \left|\underline{Q}_{i,\ell} - \underline{Q}_{j,\ell}\right| &\leqslant \sum_{\ell \text{ typical}} \left[1 + \left(\frac{C\log^2 d}{d} - 1\right)\right] \underline{Q}_{j,\ell} \\ &= \frac{C\log^2 d}{d} \sum_{\ell \text{ typical}} \underline{Q}_{j,\ell} \leqslant \frac{C\log^2 d}{d}\end{aligned}$$

Combining the bounds on $\sum_{\ell \text{ typical}} \left|\underline{Q}_{i,\ell} - \underline{Q}_{j,\ell}\right|$ and $\sum_{\ell \text{ outlier}} \left|\underline{Q}_{i,\ell} - \underline{Q}_{j,\ell}\right|$ gives the desired result. □

We can furthermore use Lemma 6.8, Corollary 6.10, and a coupling argument, to prove Lemma 6.2.

Proof of Lemma 6.2. We may assume event $\mathscr{E}_\gamma$ (that all shells $\kappa_i \leqslant \eta$), and the outcomes of Lemma 6.8 and Corollary 6.10. The union of these three events occur with probability $1 - o(m^{-\gamma} + m^{-\log m})$.

Let $(\boldsymbol{X}_a^{(t)})_{t\geqslant 0}$ be the trajectory of Markov chain $\underline{Q}$ starting at vertex a. For any pair of vertices $i, j \in [n]$, we couple $\boldsymbol{X}_i^{(2)}$ and $\boldsymbol{X}_j^{(2)}$ such that they are equal with probability $1 - O\left(\frac{\log^2 d}{d}\right)$, and so by Fact 2.6:

$$\left\|\mathrm{pmf}\left(\boldsymbol{X}_i^{(2)}\right) - \mathrm{pmf}\left(\boldsymbol{X}_j^{(2)}\right)\right\|_{\mathrm{TV}} = \frac{1}{2}\left\|\left(\underline{Q}^2\right)_{i,*} - \left(\underline{Q}^2\right)_{j,*}\right\|_1 \leqslant O\left(\frac{\log^2 d}{d}\right)$$

from which the desired result follows.

We now exhibit such a coupling between $\boldsymbol{X}_i^{(2)}$ and $\boldsymbol{X}_j^{(2)}$. Observe that $\boldsymbol{X}_i^{(1)}$ and $\boldsymbol{X}_j^{(1)}$ are distributed according to $\left(\underline{Q}\right)_{i,*}$ and $\left(\underline{Q}\right)_{j,*}$. When κ_i and κ_j are both typical shells, we can couple $\boldsymbol{X}_i^{(1)}$ and $\boldsymbol{X}_j^{(1)}$ such that they are equal with probability $1 - O\left(\frac{\log^2 d}{d}\right)$ using Corollary 6.10 and Fact 2.6. As a result $\boldsymbol{X}_i^{(2)}$ and $\boldsymbol{X}_j^{(2)}$ can be coupled so that they agree with probability $1 - O\left(\frac{\log^2 d}{d}\right)$. When either κ_i or κ_j is an outlier shell, though $\boldsymbol{X}_i^{(1)}$ and $\boldsymbol{X}_j^{(1)}$ may have TV distance greater than $O\left(\frac{\log^2 d}{d}\right)$, by Lemma 6.8 for both random variables $1 - O\left(\frac{1}{d}\right)$-fraction of the probability mass is over the typical shells. Due to that, we can couple $\boldsymbol{X}_i^{(2)}$ and $\boldsymbol{X}_j^{(2)}$ with probability $\left(1 - O\left(\frac{1}{d}\right)\right) \cdot \left(1 - O\left(\frac{\log^2 d}{d}\right)\right)$ by Lemma 6.8. Thereby we complete the proof. □

6.2 Relating typical rows and bounding outlier columns

Our next step is to prove Lemma 6.8 and Lemma 6.9. Throughout this section, instead of working with $\underline{Q}$, we will work with $Q = \mathbf{E}[\boldsymbol{A}_{\boldsymbol{G}} \mid \kappa]$; it will be simpler to operate on the entries of Q and later relate them to $\underline{Q}$. We first characterize the entries of Q using the $\mathsf{D}_{\mathsf{ip}}(d-1)$ distribution.

6.2.1 The conditional expected adjacency matrix

For each pair of vertices $i, j \in [m]$, we have

$$Q_{i,j} = q_{\kappa_i}(\kappa_j) \coloneqq \mathop{\mathbf{Pr}}_{\boldsymbol{v}_i, \boldsymbol{v}_j \sim \rho_w}\left[\left\langle \boldsymbol{v}_i, \boldsymbol{v}_j\right\rangle \geqslant \tau \mid \left\langle \boldsymbol{v}_i, w\right\rangle = \kappa_i, \left\langle \boldsymbol{v}_j, w\right\rangle = \kappa_j\right].$$

Though $q_x(y)$ is symmetric in its inputs x, y, we choose this notation because we will often work with the function $q_x(\cdot)$, where the input is any value in $[\tau, 1]$.

Claim 6.11. The quantity $q_x(y)$ is exactly a tail probability of the $\mathsf{D}_{\mathsf{ip}}(d-1)$ distribution:

$$q_x(y) = \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,y))$$

Proof. Conditional on $\langle \boldsymbol{v}_i, w\rangle = x$ and $\langle \boldsymbol{v}_j, w\rangle = y$, $\boldsymbol{v}_i$ and $\boldsymbol{v}_j$ are distributed as $\boldsymbol{v}_i = x\cdot w + \sqrt{1-x^2}\cdot \boldsymbol{u}_i$ and $\boldsymbol{v}_j = y\cdot w + \sqrt{1-y^2}\cdot \boldsymbol{u}_j$, for $\boldsymbol{u}_i, \boldsymbol{u}_j$ uniformly random unit vectors orthogonal to w. Now, observe that the condition $\langle \boldsymbol{v}_i, \boldsymbol{v}_j\rangle \geqslant \tau$ is equivalent to $\langle \boldsymbol{u}_i, \boldsymbol{u}_j\rangle \geqslant \frac{\tau - xy}{\sqrt{(1-x^2)(1-y^2)}}$, and thus the desired statement follows since $\boldsymbol{u}_i$ and $\boldsymbol{u}_j$ are sampled from a space isometric to $\mathbb{S}^{d-2}$. □

6.2.2 The contribution of outlier columns

The goal of this section is to prove Lemma 6.8.

Lemma (Restatement of Lemma 6.8)**.** *For any $d \geqslant C\log m$ for some constant $C > 0$ and any $\tau \in [0,1]$ such that $qm \gg \log^8 m$, if $\kappa_i \leqslant \eta$,*

$$\sum_{k=1}^{m} Q_{i,k}\cdot \mathbf{1}[k \text{ outlier}] \leqslant O\left(\frac{1}{d}\right)$$

with probability $1 - o(m^{-\log m})$.

The lemma statement is equivalent to the following about Q: for all $i \in [m]$,

$$\frac{\sum_{k=1}^{m} Q_{i,k}\cdot \mathbf{1}[k \text{ outlier}]}{\sum_{k=1}^{m} Q_{i,k}} \leqslant O\left(\frac{1}{d}\right)$$

with probability $1 - o(m^{-\log m})$. Recalling that we use Z to denote the normalizing constant from Section 2.3, we can define:

$$N(x) \coloneqq \int_{\tau(1+\alpha)}^{1} q_x(y)\cdot Z^{-1}(1-y^2)^{(d-3)/2}dy = Z^{-1}\int_{\tau(1+\alpha)}^{1}(1-y^2)^{(d-3)/2}\cdot \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,y))\,dy$$
$$D(x) \coloneqq Z^{-1}\int_{\tau}^{1} q_x(y)\cdot (1-y^2)^{(d-3)/2}dy = Z^{-1}\int_{\tau}^{1}(1-y^2)^{(d-3)/2}\cdot \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,y))\,dy$$

By our definitions of $N(x)$ and $D(x)$, and recalling that we condition on $\mathscr{E}_\gamma$ (Definition 6.3) throughout this section,

$$N(\kappa_i) = \mathop{\mathbf{E}}_{\boldsymbol{\kappa}_\ell}[Q_{i,\ell}\cdot \mathbf{1}[\ell \text{ outlier}]],\ D(\kappa_i) = \mathop{\mathbf{E}}_{\boldsymbol{\kappa}_\ell}[Q_{i,\ell}]$$

The $Z^{-1}(1-y^2)^{(d-3)/2}$ expression in each integrand comes from the probability density over shells.

First, when $\kappa_i \leqslant \eta$, we establish that the ratio of the *expected* sum of outlier $Q_{i,k}$ and typical $Q_{i,k}$ is of the desired magnitude of $O\left(\frac{1}{d}\right)$.

Lemma 6.12. *For any $d \geqslant C\log m$ for some constant $C > 0$, any constant $\tau \in [0,1]$, and any $x \leqslant \eta$,*

$$\frac{N(x)}{D(x)} \leqslant O\left(\frac{1}{d}\right).$$

The proof is by computation and is provided in Appendix B. We are now ready to prove Lemma 6.8.

Proof of Lemma 6.8. For convenience, we use N and D as shorthand for $N(\kappa_i)$ and $D(\kappa_i)$. We compute a high probability lower bound for the numerator $\sum_{k=1}^{m} Q_{i,k} \cdot \mathbf{1}[k \text{ outlier}]$ and a high probability upper bound for the denominator $\sum_{k=1}^{m} Q_{i,k}$.

Concentration of the numerator: We first show that $\sum_{k=1}^{m} Q_{i,k} \cdot \mathbf{1}[k \text{ outlier}]$ concentrates around Nm. First, each $Q_{i,k} \leqslant 1$. Then, $\mathbf{Var}(Q_{i,k}) \leqslant \mathbf{E}[Q_{i,k}^2 \cdot \mathbf{1}[k \text{ outlier}]] \leqslant \mathbf{E}[Q_{i,k} \cdot \mathbf{1}[k \text{ outlier}]] = N$. Applying Bernstein's Inequality, we obtain:

$$\mathbf{Pr}\left[\sum_{k=1}^{m} Q_{i,k} \cdot \mathbf{1}[k \text{ outlier}] \geqslant Nm + \left(\sqrt{Nm}+1\right)\log^2 m\right] \leqslant \exp\left(-\frac{\left(\sqrt{Nm}+1\right)^2 \log^4 m}{\frac{1}{2}\cdot Nm + \frac{1}{3}\cdot\left(\sqrt{Nm}+1\right)\log^2 m}\right) \leqslant m^{-\log m}$$

Concentration of the denominator: We next show that $\sum_{k=1}^{m} Q_{i,k}$ concentrates around Dm. Using a similar bound on variance as above, and applying Bernstein's inequality again:

$$\mathbf{Pr}\left[\sum_{k=1}^{m} Q_{i,k} \leqslant Dm - \left(\sqrt{Dm}+1\right)\log m\right] \leqslant \exp\left(-\frac{\left(\sqrt{Dm}+1\right)^2 \log^2 m}{\frac{1}{2}\cdot Dm + \frac{1}{3}\cdot\left(\sqrt{Dm}+1\right)\log m}\right) \leqslant m^{-\log m}$$

Thus, with probability greater that $1 - 2m^{-\log m}$, the ratio $\frac{\sum_{k=1}^{n} Q_{i,k}\cdot \mathbf{1}[k \text{ outlier}]}{\sum_{k=1}^{m} Q_{i,k}}$ is at most

$$\frac{Nm + \left(\sqrt{Nm}+1\right)\log^2 m}{Dm - \left(\sqrt{Dm}+1\right)\log m}$$

We can upper bound this by $\leqslant O\left(\frac{1}{d}\right)$, as Lemma 6.12 tells us $\frac{N}{D} \leqslant O\left(\frac{1}{d}\right)$, and since $Dm \geqslant qm \gg \log^8 m$ the first terms in the ratio dominate. □

6.2.3 Relating typical rows

Our goal for this section is to prove:

Lemma (Restatement of Lemma 6.9)**.** *For any dimension d and any threshold $\tau \in (0,1)$, if κ_i, κ_j correspond to typical shells, then for all ℓ such that κ_ℓ is typical,*

$$\underline{Q}_{i,\ell} \in \left(1 \pm O\left(\frac{\log^2 d}{d}\right)\right)\underline{Q}_{j,\ell}$$

We will translate Lemma 6.9 into a statement about Q first. Let $Q_{i,*}$ and $Q_{j,*}$ be rows of Q corresponding to typical shells κ_i, κ_j. We will prove that $Q_{i,*}$ and $Q_{j,*}$, when restricted to typical columns, are nearly constant scalings of each other. Formally, we will prove:

Lemma 6.13. *For any dimension d and any threshold $\tau \in (0,1)$, let $\kappa_i, \kappa_j, \kappa_\ell$ be typical shells. Then,*

$$\frac{q_{\kappa_i}(\kappa_\ell)}{q_{\kappa_j}(\kappa_\ell)} \cdot \left(\frac{q_{\kappa_i}(\tau)}{q_{\kappa_j}(\tau)}\right)^{-1} \in 1 \pm O\left(\frac{\log^2 d}{d}\right)$$

In other words, this establishes that for any typical shells $\kappa_i, \kappa_j, \kappa_\ell$,

$$\frac{q_{\kappa_i}(\kappa_\ell)}{q_{\kappa_j}(\kappa_\ell)} \approx \frac{q_{\kappa_i}(\tau)}{q_{\kappa_j}(\tau)},$$

where the quantity on the right is a constant c_{ij} depending only on κ_i and κ_j (not κ_ℓ).

Proof of Lemma 6.9 using Lemma 6.13. By the definition of $\underline{Q}$,

$$\underline{Q}_{i,\ell} = \frac{Q_{i,\ell}}{\sum_{k=1}^m Q_{i,k}}, \underline{Q}_{j,\ell} = \frac{Q_{j,\ell}}{\sum_{k=1}^m Q_{j,k}}$$

It suffices to prove that $\frac{\underline{Q}_{i,\ell}}{\underline{Q}_{j,\ell}}$ is close to 1. Expanding $\frac{\underline{Q}_{i,\ell}}{\underline{Q}_{j,\ell}}$, we can upper bound:

$$\begin{aligned}
\frac{\underline{Q}_{i,\ell}}{\underline{Q}_{j,\ell}} &= \frac{Q_{i,\ell}}{\sum_{k=1}^m Q_{i,k}} \cdot \frac{\sum_{k=1}^m Q_{j,k}}{Q_{j,\ell}} \\
&= \frac{Q_{i,\ell}}{Q_{j,\ell}} \cdot \frac{\sum_{k=1}^m Q_{j,k}}{\sum_{k=1}^m Q_{i,k}} \\
&\leqslant \left(1 + \frac{C \log^2 d}{d}\right)\left(\frac{q_{\kappa_i}(\tau)}{q_{\kappa_j}(\tau)}\right) \cdot \frac{\sum_{k \text{ typical}} Q_{j,k} + \sum_{k \text{ outlier}} Q_{j,k}}{\sum_{k \text{ typical}} Q_{i,k} + \sum_{k \text{ outlier}} Q_{i,k}} \\
&\leqslant \left(1 + \frac{C \log^2 d}{d}\right)\left(\frac{q_{\kappa_i}(\tau)}{q_{\kappa_j}(\tau)}\right) \cdot \frac{\left(1 + \frac{C'}{d}\right) \sum_{k \text{ typical}} Q_{j,k}}{\sum_{k \text{ typical}} Q_{i,k}} \\
&\leqslant \left(1 + \frac{C \log^2 d}{d}\right)\left(\frac{q_{\kappa_i}(\tau)}{q_{\kappa_j}(\tau)}\right) \cdot \frac{\left(1 + \frac{C'}{d}\right)\left(1 + \frac{C \log^2 d}{d}\right) \sum_{k \text{ typical}} \left(\frac{q_{\kappa_j}(\tau)}{q_{\kappa_i}(\tau)}\right) \cdot Q_{i,k}}{\sum_{k \text{ typical}} Q_{i,k}} \\
&\leqslant 1 + \frac{C'' \log^2 d}{d}
\end{aligned}$$

The first inequality uses Lemma 6.13 to bound $\frac{Q_{i,\ell}}{Q_{j,\ell}}$. The second inequality uses the fact that the outlier entries of $Q_{i,*}$ and $Q_{j,*}$ only occupy an $O\left(\frac{1}{d}\right)$ fraction of the ℓ_1 norm of each row (Lemma 6.8). The third inequality again comes from an application of Lemma 6.13 to relate $Q_{i,k}$ to $Q_{j,k}$ when k is typical. The lower bound follows analogously. □

Proof of Lemma 6.13. By definition:

$$\frac{q_{\kappa_i}(\kappa_\ell)}{q_{\kappa_j}(\kappa_\ell)} \cdot \left(\frac{q_{\kappa_i}(\tau)}{q_{\kappa_j}(\tau)}\right)^{-1} = \frac{\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(\kappa_i, \kappa_\ell))}{\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}\left(T(\kappa_j, \kappa_\ell)\right)} \cdot \frac{\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}\left(T(\kappa_j, \tau)\right)}{\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(\kappa_i, \tau))} \tag{28}$$

Since $\kappa_i, \kappa_j, \kappa_\ell \in [\tau, \tau(1+\alpha)]$, by Claim 5.6, all terms of the form $T(x, y)$ in the above are lower bounded by $T(\tau(1+\alpha), \tau(1+\alpha))$, which is lower bounded by a constant for large enough d. Thus, by Lemma 2.8:

$$(28) = \left(1 \pm O\left(\frac{\log d}{d}\right)\right) \cdot \frac{T(\kappa_j, \kappa_\ell) \cdot T(\kappa_i, \tau)}{T(\kappa_i, \kappa_\ell) \cdot T(\kappa_j, \tau)} \cdot \left(\frac{A}{B}\right)^{(d-1)/2}$$

where $A := \left(1 - T(\kappa_i, \kappa_\ell)^2\right)\left(1 - T(\kappa_j, \tau)^2\right)$ and $B := \left(1 - T(\kappa_j, \kappa_\ell)^2\right)\left(1 - T(\kappa_i, \tau)^2\right)$. We show:

$$\left|\frac{T(\kappa_j, \kappa_\ell) \cdot T(\kappa_i, \tau)}{T(\kappa_i, \kappa_\ell) \cdot T(\kappa_j, \tau)} - 1\right| \leqslant O(\alpha^2) \tag{29}$$

$$\left|\frac{A}{B} - 1\right| \leqslant O(\alpha^2) \tag{30}$$

where (29) is proved in Claim C.1 and (30) is proved in Claim C.2. Consequently,

$$\frac{q_a(x)}{q_b(x)} \cdot \frac{q_b(\tau)}{q_a(\tau)} = \left(1 \pm O\left(\frac{\log d}{d}\right)\right) \cdot \left(1 \pm O(\alpha^2)\right) \cdot \left(1 \pm O(d\alpha^2)\right)$$

The term of order $d\alpha^2$ dominates, and because $\alpha = O(\frac{\log d}{d})$ we conclude the desired result. □

7 2-dimensional expansion of the random geometric complex

In this section we prove Theorem 1.6.

Theorem 7.1. *For every $0 < \varepsilon < 1$, $0 < \eta < 2\varepsilon$ and $d = \eta \log_{4/3} n$, if $\boldsymbol{H} \sim \mathsf{Geo}_d^{(2)}\left(n, n^{-1+\varepsilon}\right)$, then every link of $\boldsymbol{H}$ is a $\left(\frac{1}{2} - \delta\right)$-expander, and its 1-skeleton is a $\left(1 - \frac{4\delta}{1+2\delta}\right)$-expander with high probability where $\delta = \frac{1}{2} \cdot \frac{1-\sqrt{1-\exp\left(-2\log\frac{4}{3}\cdot(1-\varepsilon)/\eta\right)}}{1+\sqrt{1-\exp\left(-2\log\frac{4}{3}\cdot(1-\varepsilon)/\eta\right)}} - o_n(1)$.*

One of the ingredients in the proof of Theorem 7.1 is the spectral expansion of random geometric graphs, which is a corollary of Theorem 1.10 and Corollary 4.7:

Theorem (Restatement of Theorem 1.7)**.** *Let $\boldsymbol{G} \sim \mathsf{Geo}_d(n,p)$ and $\tau := \tau(p,d)$. Then with high probability $\boldsymbol{G}$ is a μ-expander, where*

$$\mu := (1+o(1)) \cdot \max\left\{(1+o_{d\tau^2}(1)) \cdot \tau, \frac{\log^4 n}{\sqrt{pn}}\right\},$$

where $o_{d\tau^2}(1)$ denotes a function that goes to 0 as $d \cdot \tau(p,d)^2 \to \infty$.

The second ingredient is a bound on the second eigenvalue of the links, proved in Section 5:

Theorem (Restatement of Theorem 5.1)**.** *Let $0 < \tau < 1$ be a constant. Let $\boldsymbol{v}_1, \dots, \boldsymbol{v}_m \sim \mathrm{cap}_{\geqslant\tau}(w)$ and $\boldsymbol{G} := \mathrm{gg}_\tau(\boldsymbol{v}_1, \dots, \boldsymbol{v}_m)$. Then for $q := \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}\left(\frac{\tau}{1+\tau}\right)$, suppose $qm \gg \log^8 m \cdot \log^{3/2} \frac{1}{q} \cdot \left(\frac{1+\tau}{\tau}\right)^3$ and $d \geqslant C \cdot \log m$ for any constant $C > 0$, then for any constant $\gamma > 0$,*

$$\mathbf{Pr}\left[|\lambda|_2\left(\widehat{A}_{\boldsymbol{G}}\right) > \frac{\tau}{1+\tau} + o_{d,m}(1)\right] \leqslant O\left(m^{-\gamma}\right).$$

Proof of Theorem 7.1. To show that the links expand, we apply Theorem 5.1 in combination with a union bound over all links. The second eigenvalue bound for the 1-skeleton is then proved using Theorem 1.3, the trickling-down theorem. Let $p = n^{-1+\varepsilon}$, $d = \eta \log_{4/3} n$ and $\tau = \tau(p,d)$.

Let $\boldsymbol{G} := \mathrm{gg}_\tau(\boldsymbol{v}_1, \dots, \boldsymbol{v}_n)$ to denote the geometric graph of the collection of vectors used to generate $\boldsymbol{H}$. The number of vertices that fall in the neighborhood of a vertex v within $\boldsymbol{G}$ is $\boldsymbol{m}_v \sim \mathsf{Binom}(n,p)$, and hence $\boldsymbol{m}_v \geqslant m := pn - 2\sqrt{pn \log n}$ except with probability $o(1/n)$. For the rest of the proof, we condition on the event that $\boldsymbol{m}_v \geqslant m$ for all v, which happens with probability $1 - o(1)$ by the union bound.

The link $\boldsymbol{H}_v$ of a vertex v is obtained by taking $\boldsymbol{G}_v$, the subgraph of $\boldsymbol{G}$ induced by the neighborhood of v, and then removing the isolate vertices. Note that the isolated vertices need to be removed since when sampling a random complex, we remove all edges that are not in any triangles. Our goal is to control the second eigenvalue of all the links in $\boldsymbol{H}$, and we do so by showing bounds on the second eigenvalue of $\boldsymbol{G}_v$ for all v. The second eigenvalue bounds show that with high probability, for all v, the graph $\boldsymbol{G}_v$ is connected, and hence has no isolated vertices. Consequently, $\boldsymbol{H}_v$ is in fact equal to $\boldsymbol{G}_v$ and the second eigenvalue bounds port over.

As a first step, we show that $\boldsymbol{G}_v$ satisfies the hypothesis of Theorem 5.1. In particular, we show for $q := \Phi_{\mathsf{D}_{\mathsf{ip}}(d-1)}\left(\frac{\tau}{1+\tau}\right)$

$$q \cdot m_v \gg \log^4 m_v \cdot \log^2 \frac{1}{q} \cdot \left(\frac{1+\tau}{\tau}\right)^4. \tag{31}$$

Using Lemma 2.8 and the fact that the tail function of a probability distribution is monotone decreasing, we can lower bound q:

$$q \geqslant \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}\left(\frac{1}{2}\right) \geqslant \frac{2Z_{d-1}}{d-2} \cdot \left(\frac{3}{4}\right)^{(d-2)/2} \cdot \left(1 - \frac{16\log d}{d-1}\right) \geqslant \Omega\left(\frac{1}{\sqrt{d}}\right) \cdot n^{-\eta/2},$$

where the first inequality holds since $\tau \in (0,1]$, and the last inequality holds by definitions of d and Z_d. We now lower bound τ by a constant. By Lemma 2.8:

$$\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}\left(\sqrt{1-\exp\left(-\frac{(1-\varepsilon)\log\frac{4}{3}}{\eta}\right)}\right) \geqslant \Omega\left(\frac{1}{\sqrt{d}}\right) \cdot \sqrt{p} \cdot \left(1 - O\left(\frac{\log d}{d}\right)\right) \geqslant p = \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}(\tau),$$

where the first inequality holds by definition of Z_d. Since $\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}$ is a decreasing function,

$$\tau \geqslant \sqrt{1-\exp\left(-\frac{(1-\varepsilon)\log\frac{4}{3}}{\eta}\right)}.$$

Consequently:

$$\log^4 m_v \cdot \log^2 \frac{1}{q} \cdot \left(\frac{1+\tau}{\tau}\right)^4 \leqslant \log^6 n.$$

On the other hand, $q m_v \geqslant \Omega\left(\frac{1}{\sqrt{d}}\right) \cdot n^{\varepsilon-\eta/2} \gg \log^6 n$, which establishes (31). By (31) and Theorem 5.1 with $\gamma = 2/\varepsilon$, with probability at least $1 - O(1/n^2)$:

$$|\lambda|_2\left(\widehat{A}_{\boldsymbol{G}_v}\right) \leqslant \frac{\tau}{1+\tau} + o_n(1).$$

By the union bound over all vertices, with probability $1 - O(1/n)$:

$$|\lambda|_2\left(\widehat{A}_{\boldsymbol{G}_v}\right) \leqslant \frac{\tau}{1+\tau} + o_n(1) \qquad \forall v \in [n].$$

Henceforth, we condition on the above. Since $\frac{\tau}{1+\tau} < 1$, for all $v \in [n]$, each $\boldsymbol{G}_v$ is connected and has no isolated vertices and hence $\boldsymbol{H}_v = \boldsymbol{G}_v$. Consequently

$$|\lambda|_2\left(\widehat{A}_{\boldsymbol{H}_v}\right) \leqslant \frac{\tau}{1+\tau} + o_n(1) \qquad \forall v \in [n].$$

Assuming the 1-skeleton $\boldsymbol{H}^{(1)}$ is connected, by the trickling-down theorem (Theorem 1.3) it satisfies:

$$|\lambda|_2\left(\widehat{A}_{\boldsymbol{H}^{(1)}}\right) \leqslant \frac{\frac{\tau}{1+\tau} + o_n(1)}{1 - \frac{\tau}{1+\tau} - o_n(1)} = \tau + o_n(1).$$

It remains to bound τ, $\tau/(1+\tau)$ and show $\boldsymbol{H}^{(1)}$ is connected. By Lemma 2.8, the following inequality must be satisfied:

$$p \leqslant \frac{Z_d}{\tau(d-1)} \cdot \left(1-\tau^2\right)^{(d-1)/2}.$$

Since the right hand side of the above is a decreasing function of τ and plugging in $\sqrt{1-\exp\left(-\frac{2(1-\varepsilon)\log\frac{4}{3}}{\eta}\right)}$ yields a value smaller than p, we know:

$$\tau \leqslant \sqrt{1-\exp\left(-\frac{2(1-\varepsilon)\log\frac{4}{3}}{\eta}\right)} = 1 - \frac{4\delta}{1+2\delta}. \tag{32}$$

The function $\tau/(1+\tau)$ is an increasing function and hence:

$$\frac{\tau}{1+\tau} \leqslant \frac{\sqrt{1-\exp\left(-\frac{2(1-\varepsilon)\log\frac{4}{3}}{\eta}\right)}}{1+\sqrt{1-\exp\left(-\frac{2(1-\varepsilon)\log\frac{4}{3}}{\eta}\right)}} + o_n(1) = \frac{1}{2} - \delta.$$

Finally, to show $\boldsymbol{H}^{(1)}$ is connected, it suffices to illustrate $\widetilde{\boldsymbol{H}}^{(1)}$, a reweighted version of $\boldsymbol{H}^{(1)}$, whose normalized adjacency matrix has a spectral gap. We use $\boldsymbol{G}$ as our reweighting of $\boldsymbol{H}^{(1)}$, which is valid since all edges in $\boldsymbol{G}$ occur in $\boldsymbol{H}^{(1)}$ with probability $1-o_n(1)$. Indeed, for every vertex v and neighbor w the vertex w has some neighbor w' within $\boldsymbol{G}_v$, which means $\{v,w,w'\}$ is a triangle in $\boldsymbol{H}$ causing $\{v,w\}$ to appear in $\boldsymbol{H}^{(1)}$. By our choice of parameters, the lower and upper bounds on τ shown in (7) and (32) respectively, and Theorem 1.7, we know $|\lambda|_2(A_{\boldsymbol{G}}) < \tau + o_n(1) < 1$, which implies $\boldsymbol{H}^{(1)}$ is connected, which completes our proof. □

8 Tightness of the tricking-down theorem

In this section we will show that the trickle-down theorem is tight:

Proposition (Restatement of Proposition 1.8)**.** *For each $\lambda \in (0, \frac{1}{2}]$ and $\eta > 0$ there exists a 2-dimensional expander in which all vertex link eigenvalues are at most λ for which the 1-skeleton is connected with eigenvalue at least $\frac{\lambda}{1-\lambda} - \eta$.*

We prove the proposition by showing that a random geometric graph's adjacency matrix (when weighted in a regular way) has second eigenvalue at least τ, and then prove that the random geometric complex indeed satisfies that regularity condition.

Lemma 8.1. *Let $\boldsymbol{G} \sim \mathsf{Geo}_d(n,p)$ generated by vectors $\boldsymbol{v}_1,\dots,\boldsymbol{v}_n$, and let W be the transition matrix of any time-reversible Markov chain on $\boldsymbol{G}$ with stationary distribution π. Then with high probability $\lambda_2(W) \geqslant \tau - o_n(1) - O(\mathrm{d_{TV}}(\pi, U_n)^2)$ where U_n is the uniform distribution on $[n]$.*

Proof. When $\mathrm{d_{TV}}(\pi,U_n) \geqslant 0.1$, the statement is vacuously true. Thus, we assume $\mathrm{d_{TV}}(\pi,U_n) < 0.1$ for the rest of this proof. We see that:

$$1-\lambda_2(W) = \min_{\substack{f:V(\boldsymbol{G})\to\mathbb{R}^d \\ f \text{ non-constant}}} \frac{\mathbf{E}_{x\sim_W y}\|f(x)-f(y)\|^2}{\mathbf{E}_{x,y\sim\pi}\|f(x)-f(y)\|^2} \leqslant \frac{\mathbf{E}_{x\sim_W y}\|\boldsymbol{v}_x-\boldsymbol{v}_y\|^2}{\mathbf{E}_{x,y\sim\pi}\|\boldsymbol{v}_x-\boldsymbol{v}_y\|^2} \leqslant \frac{2(1-\tau(p,d))}{\mathbf{E}_{x,y\sim\pi}\|\boldsymbol{v}_x-\boldsymbol{v}_y\|^2} \tag{33}$$

where the last inequality uses that for adjacent x,y, $\langle \boldsymbol{v}_x, \boldsymbol{v}_y\rangle \geqslant \tau(p,d)$. To lower bound the denominator, observe:

$$\begin{aligned}
\mathop{\mathbf{E}}_{x,y\sim\pi}\|\boldsymbol{v}_x-\boldsymbol{v}_y\|^2 &= \sum_{x,y\in[n]} \pi(x)\pi(y)\big(2-2\langle \boldsymbol{v}_x,\boldsymbol{v}_y\rangle\big) = 2\left(1-\sum_{x,y\in[n]}\langle \pi(x)\boldsymbol{v}_x, \pi(y)\boldsymbol{v}_y\rangle\right) \\
&= 2\left(1-\left\|\sum_{x\in[n]}\pi(x)\boldsymbol{v}_x\right\|^2\right) = 2\left(1-\left\|\sum_{x\in[n]}\frac{1}{n}\boldsymbol{v}_x + \sum_{x\in[n]}\left(\pi(x)-\frac{1}{n}\right)\boldsymbol{v}_x\right\|^2\right) \\
&\geqslant 2\left(1-\left\|\frac{1}{n}\sum_{x\in[n]}\boldsymbol{v}_x\right\|^2 - 4\left\|\frac{1}{n}\sum_{x\in[n]}\boldsymbol{v}_x\right\|\cdot \mathrm{d_{TV}}(\pi,U_n) - 4\mathrm{d_{TV}}(\pi,U_n)^2\right).
\end{aligned}$$

By standard concentration arguments, $\left\|\frac{1}{n}\sum_{x\in[n]}\boldsymbol{v}_x\right\|$ is $o_n(1)$ with high probability. Plugging in the lower bound into (33) tells us:

$$1-\lambda_2(W)\leqslant 1-\tau(p,d)+o_n(1)+O\big(\mathrm{d_{TV}}(\pi,U_n)^2\big),$$

which can be rearranged into the desired inequality. □

Armed with this lemma we can prove Proposition 1.8.

Proof of Proposition 1.8. Let $\tau=\frac{\lambda}{1-\lambda}$, which is in $(0,1)$ for $\lambda\in\left(0,\frac{1}{2}\right)$. Using the bounds from Lemma 2.8 we can choose n and $d=\Theta(\log n)$ such that for $p=\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}(\tau)$, we have $\frac{np^2}{2}\gg\operatorname{poly}\log n$.

Let $\boldsymbol{H}\sim\mathsf{Geo}_d^{(2)}(n,p)$. Since each link contains $\mathsf{Binom}(n-1,p)$ vertices, and $(n-1)p\gg\operatorname{poly}\log n$, every link has $(n-1)p(1\pm o_n(1))\geqslant m:=np/2$ vertices with probability $1-O(n^{-1})$. Also, $\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}\left(\frac{\tau}{1+\tau}\right)\geqslant\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}(\tau)$, so $m\cdot\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d)}\left(\frac{\tau}{1+\tau}\right)\geqslant\frac{np^2}{2}\gg\operatorname{poly}\log m$, and so the conditions of Theorem 5.1 are met so that by a union bound we can conclude that all links have second eigenvalue at most $\frac{\tau}{1+\tau}+o(1)=\lambda+o(1)$.

Simultaneously, for any pair of vertices the number of triangles they participate in are within a multiplicative factor of $1\pm\frac{\log n}{\sqrt{p^2n}}=1\pm o_n(1)$ of each other, as we argue in the next paragraph. Since the stationary distribution π of the random walk on $\boldsymbol{G}$ weighted according to $\boldsymbol{H}^{(1)}$, the 1-skeleton of $\boldsymbol{H}$, puts mass on vertex v proportional to the number of triangles v participates in, it must be the case that $\pi(v)=(1\pm o_n(1))\cdot\frac{1}{n}$. Consequently, $\mathrm{d_{TV}}(\pi,U_n)=o_n(1)$, and by Lemma 8.1, $\lambda_2(\boldsymbol{H}^{(1)})\geqslant\tau-o_n(1)=\frac{\lambda}{1-\lambda}-o_n(1)$.

We now show concentration for the number of triangles that contain a vertex. Indeed, the number of triangles that a vertex v participates in is equal to the number of edges in its link. Using, $\boldsymbol{m}_v$ to denote the number of vertices in the link of v, $\deg(u)$ to denote the degree of a vertex u within the link of v, and $\boldsymbol{\kappa}$ to denote the collection of shells that vertices in the link of v lie in, we have:

$$|E(\mathrm{Link}(v))|=\frac{1}{2}\sum_{i=1}^{\boldsymbol{m}_v}\deg_v(u).$$

Henceforth we condition on $\boldsymbol{m}_v$ achieving some value in $(1\pm o_n(1))p(n-1)$. The average degree of a vertex u within the link of v is at least $\frac{np^2}{2}$, and hence by Bernstein's inequality each $\deg_v(u)=(1\pm o_n(1))\mathbf{E}\big[\deg_v(u)|\boldsymbol{\kappa}_u\big]$ except with probability $O(n^{-4})$ since $\deg_v(u)|\boldsymbol{\kappa}_u$ is a sum of independent indicator random variables. The random variables $\mathbf{E}\big[\deg_v(u)|\boldsymbol{\kappa}_u\big]$ are independent and distributed as $p(\boldsymbol{\kappa}_u)\boldsymbol{m}_v$ where $p(\boldsymbol{\kappa}_u)$ is the probability that a uniformly random vector in $\mathrm{cap}_p(v)$ falls in $\mathrm{cap}_p(u)$ where $\langle u,v\rangle=\boldsymbol{\kappa}_u$. We can show with Bernstein's inequality that:

$$\sum_{u=1}^{\boldsymbol{m}_v}\mathbf{E}\big[\deg_v(u)|\boldsymbol{\kappa}_u\big]=(1\pm o(1))\mathbf{E}[p(\boldsymbol{\kappa}_u)]\boldsymbol{m}_v^2$$

except with probability $O(n^{-4})$. By the union bound, with probability $O(n^{-1})$ for all $v\in[n]$,

$$|E(\mathrm{Link}(v))|=\frac{1\pm o_n(1)}{2}\mathbf{E}[p(\boldsymbol{\kappa}_u)]\boldsymbol{m}_v^2,$$

which completes the proof. □

Acknowledgments

We thank Vishesh Jain for a game-changing pep talk on the power of the trace method. Thanks also to Ryan O'Donnell, Prasad Raghavendra, and Nikhil Srivastava for inspiring conversations about random graphs and high-dimensional expanders. We would like to thank Omar Alrabiah for feedback on an earlier version of the paper. S.M. and T.S. were visiting the "Computational Complexity of Statistical Inference" program at Simons Institute for the Theory of Computing when some of the work was conducted, and would like to thank the institute for their hospitality and support. T.S. thanks Samory Kpotufe for suggesting to a version of Question 1.13.

References

[ABS15] Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for unique games and related problems. *Journal of the ACM (JACM)*, 62(5):1–25, 2015. 1

[AdlVKK03] Noga Alon, W Fernandez de la Vega, Ravi Kannan, and Marek Karpinski. Random sampling and approximation of MAX-CSPs. *Journal of computer and system sciences*, 67(2):212–243, 2003. 3

[AL20] Vedat Levi Alev and Lap Chi Lau. Improved analysis of higher order random walks and applications. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1198–1211, 2020. 5

[Alo86] Noga Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986. 1

[ALOV19] Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials II: high-dimensional walks and an FPRAS for counting bases of a matroid. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 1–12, 2019. 1, 5

[BDER16] Sébastien Bubeck, Jian Ding, Ronen Eldan, and Miklós Z Rácz. Testing for high-dimensional geometry in random graphs. *Random Structures & Algorithms*, 49(3):503–532, 2016. 6

[BGL+14] Dominique Bakry, Ivan Gentil, Michel Ledoux, et al. *Analysis and geometry of Markov diffusion operators*, volume 103. Springer, 2014. 24

[BHHS11] Boaz Barak, Moritz Hardt, Thomas Holenstein, and David Steurer. Subsampling mathematical relaxations and average-case complexity. In *Proceedings of the twenty-second annual ACM-SIAM Symposium on Discrete Algorithms*, pages 512–531. SIAM, 2011. 3

[BHK11] Eric Babson, Christopher Hoffman, and Matthew Kahle. The fundamental group of random 2-complexes. *Journal of the American Mathematical Society*, 24(1):1–28, 2011. 4

[BKW14] Itai Benjamini, Gady Kozma, and Nicholas Wormald. The mixing time of the giant component of a random graph. *Random Structures & Algorithms*, 45(3):383–407, 2014. 1

[BL06] Yonatan Bilu and Nathan Linial. Lifts, discrepancy and nearly optimal spectral gap. *Combinatorica*, 26(5):495–519, 2006. 1

[Bor13] Charles Bordenave. On Euclidean random matrices in high dimension. *Electronic Communications in Probability*, 18:1–8, 2013. 3, 6

[CLP20] Michael Chapman, Nati Linial, and Yuval Peled. Expander graphs – both local and global. *Combinatorica*, 40(4):473–509, 2020. 5

[Con19] David Conlon. Hypergraph expanders from Cayley graphs. *Israel Journal of Mathematics*, 233(1):49–65, 2019. 2, 5

[CS13] Xiuyuan Cheng and Amit Singer. The spectrum of random inner-product kernel matrices. *Random Matrices: Theory and Applications*, 2(04):1350010, 2013. 3, 6

[CSŻ03] Donald I Cartwright, Patrick Solé, and Andrzej Żuk. Ramanujan geometries of type $\tilde{A}_n$. *Discrete mathematics*, 269(1-3):35–43, 2003. 1, 5

[CTZ20] David Conlon, Jonathan Tidor, and Yufei Zhao. Hypergraph expanders of all uniformities from Cayley graphs. *Proceedings of the London Mathematical Society*, 121(5):1311–1336, 2020. 2, 5

[DEKL14] Jean Dolbeault, Maria J Esteban, Michal Kowalczyk, and Michael Loss. Sharp interpolation inequalities on the sphere: new methods and consequences. In *Partial Differential Equations: Theory, Control and Approximation*, pages 225–242. Springer, 2014. 24

[DEL+22] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. Locally testable codes with constant rate, distance, and locality. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 357–374, 2022. 1, 5

[Dem96] Amir Dembo. Moderate deviations for martingales with bounded jumps. *Electronic Communications in Probability*, 1:11–17, 1996. 32

[DGLU11] Luc Devroye, András György, Gábor Lugosi, and Frederic Udina. High-dimensional random geometric graphs and their clique number. *Electronic Journal of Probability*, 16:2481–2508, 2011. 6

[DHLV22] Irit Dinur, Min-Hsiu Hsieh, Ting-Chun Lin, and Thomas Vidick. Good Quantum LDPC Codes with Linear Time Decoders. *arXiv preprint arXiv:2206.07750*, 2022. 1, 5

[Din07] Irit Dinur. The PCP theorem by gap amplification. *Journal of the ACM (JACM)*, 54(3):12–es, 2007. 1

[DK17] Irit Dinur and Tali Kaufman. High dimensional expanders imply agreement expanders. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 974–985. IEEE, 2017. 5

[DV13] Yen Do and Van Vu. The spectrum of random kernel matrices: universality results for rough and varying kernels. *Random Matrices: Theory and Applications*, 2(03):1350005, 2013. 3, 6

[EK10] Noureddine El Karoui. The spectrum of kernel random matrices. *The Annals of Statistics*, 38(1):1–50, 2010. 3, 6

[EK16] Shai Evra and Tali Kaufman. Bounded degree cosystolic expanders of every dimension. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 36–48, 2016. 5

[FGL+12] Jacob Fox, Mikhail Gromov, Vincent Lafforgue, Assaf Naor, and János Pach. Overlap properties of geometric expanders. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2012(671):49–83, 2012. 2, 5

[FI20] Ehud Friedgut and Yonatan Iluz. Hyper-regular graphs and high dimensional expanders. *arXiv preprint arXiv:2010.03829*, 2020. 5

[FM19] Zhou Fan and Andrea Montanari. The spectral norm of random inner-product kernel matrices. *Probability Theory and Related Fields*, 173(1):27–85, 2019. 3, 6

[FR08] Nikolaos Fountoulakis and Bruce A Reed. The evolution of the mixing rate of a simple random walk on the giant component of a random graph. *Random Structures & Algorithms*, 33(1):68–86, 2008. 1

[Fri93] Joel Friedman. Some geometric aspects of graphs and their eigenfunctions. *Duke Mathematical Journal*, 69(3):487–525, 1993. 1

[Fri08] Joel Friedman. *A proof of Alon's second eigenvalue conjecture and related problems*. American Mathematical Soc., 2008. 1

[GG81] Ofer Gabber and Zvi Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, 1981. 1

[GGR98] Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *Journal of the ACM (JACM)*, 45(4):653–750, 1998. 3

[Gol21] Louis Golowich. Improved product-based high-dimensional expanders. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021. 2, 5

[Gro10] Mikhail Gromov. Singularities, expanders and topology of maps. part 2: From combinatorics to topology via algebraic isoperimetry. *Geometric and Functional Analysis*, 20(2):416–526, 2010. 5

[HKP+17] Samuel B Hopkins, Pravesh K Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. The power of sum-of-squares for detecting hidden structures. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 720–731. IEEE, 2017. 3

[HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006. 1

[INW94] Russell Impagliazzo, Noam Nisan, and Avi Wigderson. Pseudorandomness for network algorithms. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 356–364, 1994. 1

[KB93] Andrey Kolmogorov and Ya M. Barzdin. On the realization of networks in three-dimensional space. In *Selected Works of AN Kolmogorov*, pages 194–202. Springer, 1993. 1

[KG00] Vladimir Koltchinskii and Evarist Giné. Random matrix approximation of spectra of integral operators. *Bernoulli*, pages 113–167, 2000. 6

[KKL14] Tali Kaufman, David Kazhdan, and Alexander Lubotzky. Ramanujan complexes and bounded degree topological expanders. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 484–493. IEEE, 2014. 1, 5

[KO18] Tali Kaufman and Izhar Oppenheim. Construction of new local spectral high dimensional expanders. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 773–786, 2018. 5

[LH22] Ting-Chun Lin and Min-Hsiu Hsieh. c^3-Local Testable Codes from Lossless Expanders. *arXiv preprint arXiv:2201.11369*, 2022. 1

[Li04] Wen-Ching Winnie Li. Ramanujan hypergraphs. *Geometric & Functional Analysis GAFA*, 14(2):380–399, 2004. 1, 5

[Lin] Nati Linial. Some geometric perspectives on combinatorics: High-dimensional, local and local-to-global II. Simons Institute big data bootcamp workshop, https://www.youtube.com/watch?v=Kt9SWmIPsH8. 2

[LLR95] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995. 1

[LM06] Nathan Linial and Roy Meshulam. Homological connectivity of random 2-complexes. *Combinatorica*, 26(4):475–487, 2006. 2, 4

[LMY20] Siqi Liu, Sidhanth Mohanty, and Elizabeth Yang. High-dimensional expanders from expanders. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020. 2, 5

[LP16] Nathan Linial and Yuval Peled. On the phase transition in random simplicial complexes. *Annals of mathematics*, pages 745–773, 2016. 4

[LPS88] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988. 1, 5

[LRS15] James R Lee, Prasad Raghavendra, and David Steurer. Lower bounds on the size of semidefinite programming relaxations. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 567–576, 2015. 3

[LSV05a] Alexander Lubotzky, Beth Samuels, and Uzi Vishne. Explicit constructions of ramanujan complexes of type ad. *European Journal of Combinatorics*, 26(6):965–993, 2005. 1, 5

[LSV05b] Alexander Lubotzky, Beth Samuels, and Uzi Vishne. Ramanujan complexes of typeã d. *Israel journal of Mathematics*, 149(1):267–299, 2005. 1, 5

[Lub] Alex Lubotzky. High-dimensional expanders and property testing. STOC workshop on advances in coding theory, https://youtu.be/P5vs4ARRQjU?t=2967. 2, 4

[Lub12] Alexander Lubotzky. Expander graphs in pure and applied mathematics. *Bulletin of the American Mathematical Society*, 49(1):113–162, 2012. 1

[Lub18] Alexander Lubotzky. High dimensional expanders. In *Proceedings of the International Congress of Mathematicians: Rio de Janeiro 2018*, pages 705–730. World Scientific, 2018. 2, 4

[LY22] Yue M Lu and Horng-Tzer Yau. An equivalence principle for the spectrum of random inner-product kernel matrices. *arXiv preprint arXiv:2205.06308*, 2022. 3, 6

[LZ22] Anthony Leverrier and Gilles Zémor. Quantum tanner codes. *arXiv preprint arXiv:2202.13641*, 2022. 1, 5

[Mar73] Grigorii Aleksandrovich Margulis. Explicit constructions of concentrators. *Problemy Peredachi Informatsii*, 9(4):71–80, 1973. 1

[Mar88] Grigory Margulis. Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators. *Problemy Peredachi Informatsii*, 24(1):51–60, 1988. 1

[MOP20] Sidhanth Mohanty, Ryan O'Donnell, and Pedro Paredes. Explicit near-Ramanujan graphs of every degree. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 510–523, 2020. 1

[MW09] Roy Meshulam and Nathan Wallach. Homological connectivity of random k-dimensional complexes. *Random Structures & Algorithms*, 34(3):408–417, 2009. 4

[Nil91] Alon Nilli. On the second eigenvalue of a graph. *Discrete Mathematics*, 91(2):207–210, 1991. 1

[OP22] Ryan O'Donnell and Kevin Pratt. High-Dimensional Expanders from Chevalley Groups. *arXiv preprint arXiv:2203.03705*, 2022. 5

[Opp18] Izhar Oppenheim. Local spectral expansion approach to high dimensional expanders part I: Descent of spectral gaps. *Discrete & Computational Geometry*, 59(2):293–330, 2018. 2

[Pen03] Mathew Penrose. *Random geometric graphs*, volume 5. OUP Oxford, 2003. 6

[Pin73] Mark S Pinsker. On the complexity of a concentrator. In *7th International Telegraffic Conference*, volume 4, pages 1–318. Citeseer, 1973. 1

[PK22] Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical ldpc codes. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 375–388, 2022. 1, 5

[SS96] Michael Sipser and Daniel A Spielman. Expander codes. *IEEE transactions on Information Theory*, 42(6):1710–1722, 1996. 1

A Azuma-Hoeffding for continuous processes

Lemma (Restatement of Lemma 4.20)**.** *Let $(\boldsymbol{X}_t)_{t\geqslant 0}\subset\mathbb{R}$ be a stochastic process adapted to the filtration $\mathscr{F}_t$ with $\mathbf{E}[e^{r\,d\boldsymbol{X}_t}\mid\mathscr{F}_t]<\exp\left(r^2\sigma_t^2dt\right)$, for all t,r. Then for all $s,x>0$,*

$$\mathbf{Pr}[|\boldsymbol{X}_s-\boldsymbol{X}_0|\geqslant x]\leqslant 2\exp\left(\frac{-x^2}{4\int_0^s\sigma_t^2dt}\right).$$

Proof of Lemma 4.20. The proof is the same as that of the standard Azuma-Hoeffding inequality. Without loss of generality assume $\boldsymbol{X}_0=0$. From Markov's inequality,

$$\mathbf{Pr}[|\boldsymbol{X}_s|\geqslant x]\leqslant e^{-\theta x}(\mathbf{E}[\exp(\theta\boldsymbol{X}_s)]+\mathbf{E}[\exp(-\theta\boldsymbol{X}_s)])$$

We have that $\boldsymbol{X}_{t+dt}=\boldsymbol{X}_t+d\boldsymbol{X}_t$, and so for any θ,

$$\begin{aligned}d\,\mathbf{E}[\exp(\theta\boldsymbol{X}_t)]&=\mathbf{E}[\exp(\theta\boldsymbol{X}_{t+dt})-\exp(\theta\boldsymbol{X}_t)]\\&=\mathbf{E}[\exp(\theta\boldsymbol{X}_t)]\cdot\mathbf{E}[\exp(\theta d\boldsymbol{X}_t)-1\mid\mathscr{F}_t]\\&\leqslant\mathbf{E}[\exp(\theta\boldsymbol{X}_t)]\cdot\left(\exp\left(\tfrac{\theta^2\sigma_t^2dt}{K}\right)-1\right)=\mathbf{E}[\exp(\theta\boldsymbol{X}_t)]\cdot\tfrac{\theta^2\sigma_t^2dt}{K},\end{aligned}$$

since the higher-order terms in the Taylor expansion of exp go to zero. Hence, we conclude that $d\log\mathbf{E}[\exp(\theta\boldsymbol{X}_t)]\leqslant\frac{\theta^2\sigma_t^2dt}{K}$, and thus $\log\mathbf{E}[\exp(\theta\boldsymbol{X}_s)]=\int_0^s d\log\mathbf{E}\exp(\theta\boldsymbol{X}_t)\leqslant\int_0^s\frac{\theta^2\sigma_t^2dt}{K}$ (where we used that $\boldsymbol{X}_0=0$). In conclusion,

$$\mathbf{E}[\exp(\theta\boldsymbol{X}_s)]\leqslant\exp\left(\frac{\theta^2\int_0^s\sigma_t^2dt}{K}\right),$$

and choosing $\theta=\frac{Kx}{2\int_0^s\sigma_t^2dt}$, and then repeating the argument with $-\boldsymbol{X}_s$, completes the proof. □

B Bounding the expected contribution of outlier shells

Lemma (Restatement of Lemma 6.12)**.** *For any $d\geqslant C\log m$ for some constant $C>0$, any constant $\tau\in[0,1]$, and any $x\leqslant\eta$,*

$$\frac{N(x)}{D(x)}\leqslant O\left(\frac{1}{d}\right).$$

Proof. For the function $T(x,y):=\frac{\tau-xy}{\sqrt{1-x^2}\sqrt{1-y^2}}$, we are interested in bounding

$$\frac{\int_{\tau(1+\alpha)}^1(1-y^2)^{(d-3)/2}\cdot\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,y))\,dy}{\int_{\tau}^1(1-y^2)^{(d-3)/2}\cdot\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,y))\,dy}\leqslant\frac{\int_{\tau(1+\alpha)}^1(1-y^2)^{(d-3)/2}\cdot\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,y))\,dy}{\int_{\tau}^{\tau(1+\alpha/2)}(1-y^2)^{(d-3)/2}\cdot\overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,y))\,dy}\tag{34}$$

For any $y \in [\tau(1+\alpha), 1]$ and $z \in [\tau, \tau(1+\alpha/2)]$, we show that the expression

$$\frac{(1-y^2)^{(d-3)/2} \cdot \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,y))}{(1-z^2)^{(d-3)/2} \cdot \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,z))} \tag{35}$$

is bounded by $\alpha\tau/d$. Consequently, (34) is bounded by $\frac{2}{d}$.

We now prove (35). Recall that by Claim 5.6, $T(x,y)$ is decreasing in x and y on $[\tau,1]\times[\tau,1]$. For fixed x, let y_* be chosen such that $T(x,y_*) = \frac{4}{\sqrt{d}}$. We split into cases depending on where y and z fall with respect to y_*.

Case 1: $y_* \leqslant z$. In this case, the numerator of (35) can be bounded by $(1-y^2)^{(d-3)/2}$ and the denominator can be lower bounded by $\beta \cdot (1-z^2)^{(d-3)/2}$ for some constant $\beta > 0$. Combined with the fact that the derivative of $1-y^2$ is $-2y$ which is at most -2τ and our choice of α, we have:

$$(35) \leqslant \frac{1}{\beta} \cdot \left(\frac{1-y^2}{1-z^2}\right)^{(d-3)/2} \leqslant \frac{(1-\alpha\tau^2)^{(d-3)/2}}{\beta} \leqslant \frac{\alpha\tau}{d}.$$

Case 2: $y_* \geqslant y$. In this case, via Lemma 2.8, the numerator of (35) can be bounded by

$$\left(1-y^2\right)^{(d-3)/2} \cdot \frac{Z_{d-1}}{T(x,y)(d-2)}\left(1-T(x,y)^2\right)^{(d-2)/2} = \left(\frac{1-y^2-\tau^2+2\tau xy}{1-x^2}\right)^{(d-3)/2} \cdot \frac{Z_{d-1}\sqrt{1-T(x,y)^2}}{T(x,y)(d-2)}$$

and similarly the denominator can be lower bounded by

$$\beta \cdot \left(\frac{1-z^2-\tau^2+2\tau xz}{1-x^2}\right)^{(d-3)/2} \cdot \frac{Z_{d-1}\sqrt{1-T(x,z)^2}}{T(x,z)(d-2)}$$

for some constant $\beta > 0$. Consequently,

$$(35) \leqslant \frac{1}{\beta} \cdot \left(\frac{1-y^2-\tau^2+2\tau xy}{1-z^2-\tau^2+2\tau xz}\right)^{(d-3)/2} \cdot \frac{T(x,z)\sqrt{1-T(x,y)^2}}{T(x,y)\sqrt{1-T(x,z)^2}} \leqslant O\left(\sqrt{d}\right) \cdot \left(\frac{1-y^2-\tau^2+2\tau xy}{1-z^2-\tau^2+2\tau xz}\right)^{(d-3)/2}$$

The derivative of the expression $G(z) := 1-z^2-\tau^2+2\tau xz$ is $-2z+2\tau x$, which is bounded by $-2\tau(1-x) \leqslant -2\tau(1-\eta)$ when $\tau \leqslant z$. By the derivative bound and our choice of α the above expression is at most:

$$O\left(\sqrt{d}\right) \cdot \left(\frac{3-2\tau(1-\eta)(y-z)}{3}\right)^{(d-3)/2} \leqslant O\left(\sqrt{d}\right) \cdot \left(1-\frac{(1-\eta)\alpha\tau^2}{3}\right)^{(d-3)/2} \leqslant \frac{\alpha\tau}{d}.$$

Case 3: $y \leqslant y_* \leqslant z$. In this case:

$$(35) = \frac{(1-y^2)^{(d-3)/2} \cdot \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,y))}{(1-y_*^2)^{(d-3)/2} \cdot \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,y_*))} \cdot \frac{(1-y_*^2)^{(d-3)/2} \cdot \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,y_*))}{(1-z^2)^{(d-3)/2} \cdot \overline{\Phi}_{\mathsf{D}_{\mathsf{ip}}(d-1)}(T(x,z))}.$$

By an identical calculation to the case where $y_* \geqslant y$, the first part of the above product is bounded by

$$O\left(\sqrt{d}\right) \cdot \left(1-\frac{2\tau(1-\eta)(y_*-y)}{3}\right)^{(d-3)/2}$$

and the second part is bounded by

$$O(1) \cdot (1-2(z-y_*)\tau)^{(d-3)/2}.$$

Either $y_* - y > \alpha\tau/4$ or $z - y_* > \alpha\tau/4$, and so

$$(35) \leqslant O\left(\sqrt{d}\right) \cdot \max\left\{1-\frac{(1-\eta)\alpha\tau^2}{6}, 1-\frac{\alpha\tau^2}{4}\right\}^{(d-3)/2} \leqslant \left(1-\frac{(1-\eta)\alpha\tau^2}{6}\right)^{(d-3)/2} \leqslant \frac{\alpha\tau}{d},$$

which completes the proof. □

C Computations for evaluating ratios of typical entries

Claim C.1. Let $\kappa_x = \tau(1+\delta_x)$, where α is the typicality threshold, and $x \in \{i,j,\ell\}$. Then,

$$\left|\frac{T(\kappa_j,\kappa_\ell)\cdot T(\kappa_i,\tau)}{T(\kappa_i,\kappa_\ell)\cdot T(\kappa_j,\tau)} - 1\right| \leqslant O(\alpha^2)$$

Proof. We first simplify $\frac{T(\kappa_j,\kappa_\ell)\cdot T(\kappa_i,\tau)}{T(\kappa_i,\kappa_\ell)\cdot T(\kappa_j,\tau)}$ by expanding the expression for $T(\cdot,\cdot)$ and cancelling terms:

$$\begin{aligned}
\frac{T(\kappa_j,\kappa_\ell)\cdot T(\kappa_i,\tau)}{T(\kappa_i,\kappa_\ell)\cdot T(\kappa_j,\tau)} &= \frac{(\tau-\kappa_j\kappa_\ell)(1-\kappa_i)}{(\tau-\kappa_i\kappa_\ell)(1-\kappa_j)} \\
\frac{T(\kappa_j,\kappa_\ell)\cdot T(\kappa_i,\tau)}{T(\kappa_i,\kappa_\ell)\cdot T(\kappa_j,\tau)} - 1 &= \frac{(\tau-\kappa_j\kappa_\ell)(1-\kappa_i)-(\tau-\kappa_i\kappa_\ell)(1-\kappa_j)}{(\tau-\kappa_i\kappa_\ell)(1-\kappa_j)}
\end{aligned}$$

The claim will follow from an upper bound on the numerator and a lower bound on the denominator.
Upper bound on numerator: The following calculation establishes that the magnitude of the numerator $|(\tau-\kappa_j\kappa_\ell)(1-\kappa_i)-(\tau-\kappa_i\kappa_\ell)(1-\kappa_j)|$ is $O(\alpha^2)$.

$$\begin{aligned}
|(\tau-\kappa_j\kappa_\ell)(1-\kappa_i)-(\tau-\kappa_i\kappa_\ell)(1-\kappa_j)| &= |(\tau-\tau\kappa_i-\kappa_j\kappa_\ell+\kappa_i\kappa_j\kappa_\ell)-(\tau-\tau\kappa_j-\kappa_i\kappa_\ell+\kappa_i\kappa_j\kappa_\ell)| \\
&= |\tau(\kappa_j-\kappa_i)-\kappa_\ell(\kappa_j-\kappa_i)| = |(\tau-\kappa_\ell)(\kappa_j-\kappa_i)| \\
&= \tau^2|1-(1+\delta_\ell)|\cdot|1+\delta_j-(1+\delta_i)| = \tau^2(\delta_\ell)|\delta_j-\delta_i| \\
&\leqslant \tau^2\alpha^2
\end{aligned}$$

where the last inequality uses $\delta_\ell \leqslant \alpha$ and $|\delta_j-\delta_i| \leqslant \alpha$.
Lower bound on denominator: The denominator $(\tau-\kappa_i\kappa_\ell)(1-\kappa_j)$ is lower bounded by a constant.

$$(\tau-\kappa_i\kappa_\ell)(1-\kappa_j) \geqslant (\tau-\tau^2(1+\alpha)^2)(1-\tau(1+\alpha))$$

As $\alpha \leqslant O\left(\frac{\log d}{d}\right)$, both $(\tau-\tau^2(1+\alpha)^2)$ and $(1-\tau(1+\alpha))$ are lower bounded by a constant. □

Claim C.2. Let $\kappa_x = \tau(1+\delta_x)$, where α is the typicality threshold, and $x \in \{i,j,\ell\}$, and define $A := \left(1-T(\kappa_i,\kappa_\ell)^2\right)\left(1-T(\kappa_j,\tau)^2\right)$ and $B := \left(1-T(\kappa_j,\kappa_\ell)^2\right)\left(1-T(\kappa_i,\tau)^2\right)$. Then,

$$\left|\frac{A}{B} - 1\right| \leqslant O(\alpha^2)$$

Proof. It suffices to lower bound B by a constant, and then prove that $|A-B| \leqslant O(\alpha^2)$. Lower bounding B by a constant is straightforward; $T(x,y)$ is maximized when $x=y=\tau$, achieving a value of $\frac{\tau}{1+\tau}$. Thus, $B \geqslant (1-\frac{\tau^2}{(1+\tau)^2})^2$, which is a constant.

To get a handle on $|A-B|$, we first expand the expressions:

$$\begin{aligned}
A-B &= [T(\kappa_i,\tau)^2 - T(\kappa_j,\tau)^2 - T(\kappa_i,\kappa_\ell)^2 + T(\kappa_j,\kappa_\ell)^2] + [T(\kappa_i,\kappa_\ell)^2T(\kappa_j,\tau)^2 - T(\kappa_j,\kappa_\ell)^2T(\kappa_i,\tau)^2] \\
&= \frac{1}{(1-\tau^2)(1-\kappa_i^2)(1-\kappa_j^2)(1-\kappa_\ell^2)}\cdot[\tau^2\cdot f_1(\delta_i,\delta_j,\delta_\ell) + \tau^4\cdot f_2(\delta_i,\delta_j,\delta_\ell)]
\end{aligned}$$

where

$$f_1(x,y,z) := (1-\tau(1+x))(1-\tau^2(1+y)^2)(1-\tau^2(1+z)^2) - (1-\tau(1+y))(1-\tau^2(1+x)^2)(1-\tau^2(1+z)^2)$$

$$-(1-\tau(1+x)(1+z))(1-\tau^2(1+y)^2)(1-\tau^2)+(1-\tau(1+y)(1+z))(1-\tau^2(1+x)^2)(1-\tau^2)$$

$$f_2(x,y,z)=(1-\tau(1+x)(1+z))(1-\tau(1+y))-(1-\tau(1+y)(1+z))(1-\tau(1+x))$$

The term $(1-\tau^2)(1-\kappa_i^2)(1-\kappa_j^2)(1-\kappa_\ell^2)$ is lower bounded by a constant for large enough d and so it suffices to bound $|f_1(\delta_i,\delta_j,\delta_\ell)|$ and $|f_2(\delta_i,\delta_j,\delta_\ell)|$ by $O(\alpha^2)$.

To do so, we use the fact that f_1 and f_2 is a polynomial of constant degree (in particular, of degree at most 5) whose degree-0 and degree-1 terms are 0, and whose remaining coefficients are bounded by a constant. The result then follows since $\delta_i,\delta_j,\delta_k \leqslant \alpha$.

It is easy to see that their coefficients are bounded by a constant. The fact that the degree-0 and degree-1 terms of f_k vanish for $k=1,2$ follows from the fact that $f_k(0,0,0)=0$ and that the univariate polynomials $f_k(x,0,0)$, $f_k(0,y,0)$, and $f_k(0,0,z)$ are identically 0. □