Title

Interactions in random structures

Author

Hiesmayr, Ella Veronika

Publication Date

2024

Peer reviewed|Thesis/dissertation

Interactions in random structures

by

Ella Veronika Hiesmayr

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Statistics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Steven N. Evans, Co-chair
Professor Shirshendu Ganguly, Co-chair
Professor Fraydoun Rezakhanlou
Associate Professor Nikhil Srivastava

Summer 2024

Interactions in random structures

Abstract

Interactions in random structures

by

Ella Veronika Hiesmayr

Doctor of Philosophy in Statistics

University of California, Berkeley

Professor Steven N. Evans, Co-chair

Professor Shirshendu Ganguly, Co-chair

Interaction plays an important role in probability. When analyzing random structures, a lot of understanding is to be gained from the relationship between different aspects of an object, and the influence its different substructures have on each other. In this thesis we will explore this theme through three projects.

The first project describes the edge of the spectrum of sparse Erdős-Rényi graphs. In those graphs the spectral edge is typically determined by the neighborhood of the vertices with the highest degrees. Our main results are a description of the largest eigenvalue in terms of local geometric features of the graph, as well as the localization of the corresponding eigenvectors in the balls around high-degree vertices. Many crucial elements of the proof consist in showing that these small parts of the graph don't interact with each other and the rest of the graph. Our analysis of the largest eigenvalue and eigenvector of rooted trees using continued fractions could potentially be useful in other contexts.

The second project combines two well-known probabilistic objects – sparse Erdős-Rényi random graphs and random matrices – to obtain weighted random graphs. We once more study the spectral edge, but this time from a large deviation perspective, i.e. we focus on an extremely unlikely event. Depending on the tail of the weights, the graph and the weights interact differently to produce a largest eigenvalue that is atypically large or small. In the light-tailed case moderately large edge-weights on large stars turn out to be the most competitive structure, and in the heavy-tailed case very large weights on small clique. We provide large deviation probabilities, as well as a law of large numbers for the largest eigenvalue in both cases. Surprisingly the large deviation probabilities are universal for light-tailed weights, and identical to those of unweighted graphs. Our analysis also led to a linear algebraic result relating entry-wise matrix norms to the operator norm of a given matrix.

The third project is more explicitly about interactions: we define a multi-type birth death process, in which the evolution of a lineage depends on the empirical distribution of the other lineages present in the system. The motivation for this project comes from germinal centers, where antibodies are optimized for a specific immune response. This process has properties, like a carrying capacity and frequency-dependent selection, that classical models cannot reproduce. Our simulations suggest that by introducing this interaction between cells we obtain a model that is closer to observations. We prove that the processes effectively decouple in the limit and provide an implicit description of the limiting flow. We crucially use ideas from propagation of chaos to show that the interaction leads to a specific limiting flow, but disappears in the limit.

Contents

List of Figures

Acknowledgments

Just as the properties of the mathematical objects in this thesis emerge through an interplay between their substructures, a combination of many kinds of support was crucial to the making of this thesis.

From a mathematical perspective I am first and foremost very grateful to my advisors Steve Evans and Shirshendu Ganguly for the topics and ideas they introduced me to, and for the time they shared with me. With Steve I embarked on a subject entirely new to me, and I benefitted a lot from his patience while learning the principles of stochastic calculus. Throughout our meetings I gained a lot of experience in connecting abstract ideas with each other as well as with applications. In my work with Shirshendu I learned how important it is to switch back and forth between a high-level and detailed understanding, and that a guiding idea, paired with the willingness to revise it, is crucial to making progress towards solving a problem.

Secondly I want to thank my collaborators William DeWitt, Sebastian Hummel, Theo McKenzie and Kyeongsik Nam for the discussions, the dead-ended ones being just as eye-opening as the successful ones. I also want to thank the other members of my committee, Fraydoun Rezakhanlou and Nikhil Srivastava for their feedback and advice, as well as my other professors, in particular Jim Pitman, Alan Hammond, and Alistair Sinclair, for what they taught me. The probability student community at Berkeley was an amazing venue to discover and think through ideas and concepts. Milind and Adam, thank you for making sure we did not entirely drift off into our individual directions. Without Ümit Işlak, who was my advisor during my master, and believed in me, I would never have started this journey.

Of course many other aspects of life beyond mathematics play a role when doing a PhD. The statistics department staff made me feel very welcome and supported throughout my time in Berkeley. I especially want to thank La Shana for supporting our individual needs and for being a calming presence during the stressful moments. The statistics graduate students offered incredible support to each other from the start to the end. In particular I want to thank everyone in the SGSA for their work to create social and academic possibilities for connection. Dan's open-mindedness and his willingness to offer advice and support was greatly felt. Working alongside Hoon, Karissa and Yaxuan in our office was always uplifting, especially when going through a more difficult stretch. I thank my cohort for working and winding down together, especially during the first years. Our experience as graduate students is shaped in many ways by our employment and I am thankful to our union for defending our rights.

The PhD was not always easy, and I am grateful that my therapists helped me live and grow through my experiences. My life in Berkeley would not have been the same without my friends here. Getting a break from the demands of the PhD while dancing, hiking, climbing, talking, eating and swimming was just as important as the time spent working. Adrian, Andy, Alex, Damir, Irem, Jake, Jess, Jimmy, Kao, Lalu, Olivia, Peyten, Tuğba, and Walter, I am so grateful I met you. My friends abroad, through whom I stayed connected to life outside of the PhD and Berkeley, were also very helpful for maintaining a healthy perspective. I thank especially Anna, Carmen and Ece for the fun and long conversations. Finally I thank my family for being there, for helping me through crises, and in particular my sister for making me laugh when I felt lost.

Chapter 1

Introduction

Most objects we encounter in real life are in interaction with other objects of the same kind, thus forming a network. For instance, people are always in interaction with other people, be it physically, virtually, emotionally, economically or intellectually. On the other hand, networks can be considered to be made of different components that interact: there is of course the graph that simply records whether there is a connection or not, but each node of the graph, and each connection between one or several nodes can have its own properties. Furthermore the graph can be embedded in an environment, for instance nodes can be located in a metric space. A social network can thus be considered to be a combination of the individual properties of the people in the group, of the connections of different types and strengths between them, as well as the biological, societal and economic setting they are in. All of these parts interact and play a role in co-creating the properties of the network.

By virtue of being completely abstract, mathematics has a big advantage compared to other disciplines: it is generally much easier to study objects in isolation. We can for instance simply study a graph where all nodes and connections are the same, without needing to think about their individual properties or their environment. We can also study a group of objects with the assumption that these are independent of each other, and thus not interacting. This is one reason why mathematics can generate precise statements in a way that other disciplines cannot.

Once objects are decently understood in isolation, it can nevertheless be fruitful to combine them and study how their interaction creates an object that is richer than the sum of the isolated parts. From a mathematical perspective it is interesting to see if new properties emerge from this combination. These new objects are especially important if there is desire for mathematics to be useful for understanding real world phenomena despite its abstract nature. If the mathematical objects we study are complex enough they might shed some light into what we observe in reality. From this perspective there is a balance to strike between mathematical objects being too simple to be useful and too complicated to be studied in a meaningful way.

Randomness is another important way of introducing complexity. By assuming that the elements of the structure we study are all samples from the same distribution, we obtain a system with variability, while still having uniformity on a higher level. This allows us to derive properties that most realizations of the system will have, and quantify just how unlikely unusual behavior is. Interactions between different elements are particularly interesting in this context, as typical behavior in one aspect can interact with atypical behavior of another, or atypical behaviors of both can conspire to create unexpected effects.

In this thesis we explore the topic of interaction in random structures in three different ways. First, we see an example where we need to identify relevant local structures in a graph and isolate them to get a picture of the whole. In this case the global picture, which includes all connections in the graph, obstructs rather than clarifies understanding. A lot of work in this case goes into showing that interaction between different structures is so weak that it can essentially be ignored. Second, we merge two well-known structures, sparse Erdős-Rényi graphs and random matrices, to obtain sparse weighted graphs, and study how their dynamics interact. Here we focus on unlikely events, and our reasoning is guided by the idea that unlikely events in weighted graphs occur due to unlikely events in the graph, in the weight matrix, or a collaboration between the two. Which of these mechanisms dominates depends on how unlikely each of the different events is. The third part studies interacting branching processes. The descendants of a single organism are typically represented by a tree. In isolation we have a good understanding of these genealogical trees, but in practice these trees evolve in the presence of other trees, which can substantially change their behavior. We define a model where a set of branching processes interacts in a simple enough way so that we can still derive important properties of the system. We now describe these three parts in more detail.

The first topic of this thesis, which corresponds to the first two parts mentioned above, are properties of a specific type of random graph, namely Erdős-Rényi graphs. This model is easily defined: for a fixed set of vertices labeled $\{1, \ldots, N\}$, connect any two independently with probability p. A typical node in real-world graphs has constant degree, even if the number of nodes increases. To mirror this property, we take $p = \frac{d}{N}$ with d constant or varying slowly with N. While this model is too simple to exhibit many properties of real-world networks, it lends itself to many questions that are easily posed, but hard to answer. This allows for the development of mathematical tools that can be valuable when studying other random graph models that are closer to observed ones.

For random graphs, deriving the exact distribution of most statistics quickly becomes overly complex, so the analysis is often focused on the asymptotic regime, where the number of nodes tends to infinity. Given a certain statistic of the random graph, one can then either study the typical behavior, i.e. what happens with high probability as the size of the graph increases, or atypical events, and the structural changes they imply, in other words, which changes must the graph undergo for these rare events to occur.

One specific perspective one can take on random graphs is to look at the spectral properties of a matrix representation. A graph can be represented by several different matrices. In this thesis we use its adjacency matrix, i.e. an $N \times N$ symmetric matrix whose entries are one whenever two vertices are connected, and zero otherwise. One generally differentiates between the edge and the bulk of the spectrum. The latter is concerned with the region where most eigenvalues lie, while the former focuses on the largest and smallest eigenvalues and their eigenvectors. Those quantities capture some important features of the graph: they can for instance describe the spread of a disease on the graph. While the bulk of the eigenvalues of a random matrix tends to exhibit more universal behavior, as is exemplified in Wigner's semi-circle law, the effect of the precise distribution of the matrix entries is generally more visible for the eigenvalues at the edge. We will study the edge of the spectrum from two perspectives in this thesis.

The first contribution of this thesis is a description of the typical behavior of the largest eigenvalues and eigenvectors of Erdős-Rényi graphs with close to constant average degree. The goal here is to find geometric properties of the graph, like degrees, subgraphs, and neighborhoods, to characterize the spectral edge. Regarding the eigenvectors, a crucial question is whether they are localized or delocalized, in other words, how many vertices concentrate most of the mass of the eigenvector, and how those vertices are related.

It is known that the largest eigenvalue of sparse Erdős-Rényi graphs, both in the typical [KS03] and the large deviation case [BBG21], is equal to the square root of the largest degree. The proofs of these statements rely on the idea that the eigenvectors of the largest eigenvalues are localized around vertices with largest degrees. To prove this, an elementary and important ingredient is that the eigenvalue of a star, which is a graph consisting only of a central vertex of degree d and its neighbors, has eigenvalue $\sqrt{d}$.

This idea has been formalized and confirmed in a series of works, [ADK21a, ADK22, ADK23b, ADK23a], where it is shown that in sparse graphs the largest eigenvalues are determined by the neighborhoods of vertices with highest degrees and that the corresponding eigenvectors are localized around one of those vertices. The general approach developed in those papers is to look at the spectrum of truncated balls around high-degree vertices, which are essentially disjoint trees, and can thus be analyzed more easily, and then show that those essentially do not interact with the rest of the graph.

However, the case where $p = \frac{d}{N}$, with d constant, a problem mentioned by Alice Guionnet in her plenary lecture at the European Congress of Mathematics [Gui21], remained open. The main difficulty in this case is that the degrees of the neighbors of large degree vertices, as well as the growth of the spheres around them, are less concentrated. With Theo McKenzie, we used the general framework mentioned above, but expressed the eigenvector of the truncated balls as a continued fraction, rather than finding an explicit approximation for it. By working directly with the true eigenvector, larger fluctuations can be tolerated, which allows us to derive the above mentioned properties for this sparser case as well [HM23].

The second perspective we take on the spectral edge differs in two ways. Firstly, we no longer consider the typical behavior, and rather focus on large deviations, in other words, the event where the largest eigenvalue is atypically small or large. Secondly, we now add independent weights to the graph. Thus each edge is assigned an independent random weight whose distribution has Weibull shape: its tails are of the form $e^{-t^{\alpha}}$. When $0 < \alpha < 2$, then we consider these to be heavy tails, and when $\alpha > 2$ the tails are considered to be light. The case $\alpha = 2$ essentially corresponds to Gaussian edge weights and was treated in [GN22]. Conceptually, "$\alpha = \infty$" corresponds to the case without edge weights, for which large deviations results were obtained in [BBG21]. Those two previous results guided us in our investigation, as the edge spectrum of graphs without Gaussian weights is governed by entirely different structures than the graph without weights. While, as mentioned above, high-degree vertices are the relevant structure in the latter case, small cliques with very large edge weights turned out to be more competitive when Gaussian weights are added to the model.

Together with my advisor Shirshendu Ganguly and Kyeongsik Nam [GHN24], we derived lower and upper tail large deviation results for heavy- and light-tailed weights. Based on our proofs, it seems like graphs with heavy-tailed edge weights behave similarly to graphs with Gaussian tails, in that small cliques with very large weights are driving the large deviations. Nevertheless there is a crucial difference: while the size of the clique is bounded for $\alpha < 2$, no matter the size of the deviation, the size of the optimal clique goes to infinity in the Gaussian case. It is important to note that the two papers [BBG21] and [GN22] actually prove structural results, while the structural conjectures we make for both light- and heavy-tailed weights are based on our proof strategy.

For light tails we were guided by the idea that high-degree vertices remain the relevant structure. Here too there is a crucial difference, compared to the unweighted case, as the vertices with maximum degree might no longer be competitive. The reason behind this is that there are too few of them. Dependent on α, the group of vertices with large but not quite maximum degree has collectively a better chance at accumulating large weights on their edges. To our surprise the large deviation probability is universal for light tails, in the sense that it does not depend on the precise weight distribution, which is given by the parameter α. Indeed, it is identical to the one for unweighted graphs. Such a universality phenomenon is unexpected in the large deviation regime.

The second main topic of this thesis is an interacting multi-type birth-death process, which I studied with my advisor Steven Evans, William S. DeWitt and Sebastian Hummel [DEHH24]. The initial motivation for this project was antibody maturation in germinal centers, where cells are optimized to bind to a virus that just entered a body. There are a lot of discrepancies between observations of this process and the existing mathematical models. For instance such centers have a carrying capacity, which means that the growth of the process slows down as the number of cells increases.

To better match the observations, we define a system of multi-type birth-death processes. For each branching process in the system, its birth, death and mutation rates depend strongly on its own state and weakly on the state of the other processes present in the system. This is modeling the process in a germinal cell, where we start with an initial set of cells, which then each reproduce, die and mutate. The rates at which each cell does this depends on its affinity to bind to the virus that just entered the immune system, relative to the binding affinity of all the other cells present in the germinal center.

Since interactions make this process substantially more complicated, we restrict ourselves to mean-field interactions, which means that we can focus on one process that interacts with the other ones through the empirical measure. Besides an implicit description of the flow of the empirical measure when the number of initial particles tends to infinity, we also show that any finite number of branching processes effectively decouple in the limit, which means that while the interactions lead to a specific limiting marginal distribution, the interactions are no longer locally visible.

Structure of this thesis

The rest of this thesis consists of six chapters, the first five focusing on random graphs and the last chapter introducing our multi-type birth death process. We start by introducing known and new results that will be necessary to prove our main theorems. In Chapter 2 we first introduce a new result relating the spectral norm of a matrix to its entry-wise L^p-(quasi)norm, which generalizes the classical Motzkin-Strauss theorem [MS65]. We then state a few basic spectral properties of graphs.

Structural properties of sparse Erdős-Rényi graphs are gathered in Chapter 3. We start by citing results about the profile of the large degrees in the graph. In other words we quantify how large the maximum degree typically is, and prove or state results on the number of vertices whose degree is equal to a fraction of that maximum degree. We then study the degree profile closer to the maximum degree more closely, by estimating the number of vertices that deviate from the maximum degree by an additive constant. We also prove results about the local neighborhoods of vertices with large degree, and end the chapter with some results on the existence of a clique and the connectivity structure of highly sub-critical graphs.

Chapter 4, the last introductory chapter, focuses on distributions that are relevant in our proofs, which are the Binomial and the Poisson distribution. After some results that quantify the asymptotic equivalence of the $\text{Binom}(N, d/N)$ and the $\text{Pois}(d)$ distribution, we state tail bounds for both of these, that we will regularly use. We then show that in Erdős-Rényi graphs, this equivalence can not only be applied to each individual degree in the graph, but also holds true for the largest degrees and their neighborhoods. We end the chapter by providing bounds on the tails of sums of independent random variables. In both projects such quantities will show up in our expression for the spectrum.

The remaining three chapters are each devoted to one project. Each of those chapters starts with an introduction of the problem, followed by the main theorems, relevant previous results, and a sketch of the proofs, before providing those in full detail.

In Chapter 5 we state and prove a precise description of the edge spectrum of Erdős-Rényi graphs with close to constant average degree. Importantly we can relate each eigenvalue at the edge of the spectrum to a high-degree vertex and express it up to a small error only using its degree and the number of neighbors of its neighbors. Our other main result shows that the corresponding eigenvectors are localized in a small ball around the high-degree vertices.

We continue by studying the spectrum of weighted Erdős-Rényi graphs in Chapter 6. We derive upper and lower tail large deviation probabilities for the largest eigenvalue of Erdős-Rényi graphs with constant average degree for both heavy- and light-tailed edge weights. As a consequence of these results we also get a law of large numbers for the largest eigenvalue.

In the final chapter we introduce our multi-type birth death process model with mean field interactions, and derive some of its asymptotic properties. We start by showing that its empirical measure converges to a deterministic flow of measures, and because of the mean field interactions, we show that this also corresponds to the marginal distribution of an individual process. Finally we provide some simulations that demonstrate how a simple variety of our process exhibits some of the properties we were hoping to model.

When proofs were either simple, similar to previously published proofs, or less clearly related to the overall topic of this thesis, we moved them to the appendix.

Chapter 2

Spectral properties of graphs

In this chapter we prove and state results about the spectrum of the adjacency matrix of graphs. Finding the spectrum and eigenvectors of the adjacency matrix of graphs is an ubiquitous problem in combinatorics and spectral theory, with important applications to computer science and mathematical physics, see [Chu97, KS97, Alo98] for general overviews on their applicability. None of the results in this section are probabilistic, but we will use them to analyze our random graphs later on.

Let us first introduce some notation. For any graph G, denote by $V = V(G)$ and $E = E(G)$ the set of vertices and edges in G respectively. For any vertex $v \in V(G)$, define $d(v)$ to be the degree of a vertex v, and let $d_1(G)$ be the maximum degree of G. For any graph $G = (V, E)$ with a vertex set $V = [N] := \{1, 2, \cdots, N\}$, we write each *undirected* edge joining two vertices i and j with $i < j$ as (i, j). We use the notation $i \sim j$ for two vertices i, j if i and j are connected by an edge. We denote by A_G, or A, when G is clear from context, the adjacency matrix of G. A is the symmetric matrix that satisfies $a_{ii} = 0$, and for $i \neq j$, $a_{ij} = 1$ if $i \sim j$ and 0 otherwise.

We also use the notation $G = (V, E, A)$ to denote a network with an underlying graph $G = (V, E)$ having A, a real $n \times n$ symmetric matrix, as its weight matrix. We abuse the notation, but only slightly, as a graph can be considered as a network where all edges have weight 1. In other words, $a_{ij} = a_{ji}$ is the weight of an edge (i, j), given that $(i, j) \in E$, and 0 otherwise.

For a symmetric matrix Z, we denote its eigenvalue in non-increasing order by $\lambda_1(Z) \geq \lambda_2(Z) \geq \cdots \geq \lambda_n(Z)$. When it is clear from the context we will suppress Z in the notation. We will also sometimes use the notation $\lambda_{\max}(Z)$ to denote the maximum eigenvalue of Z.

For a vector $v \in \mathbb{R}^n$ we denote by $\|v\|$ its Euclidean norm. For a matrix A in $\mathbb{R}^{m \times n}$, we denote by $\|A\|$ the operator norm, i.e. $\|A\| = \sup_{v \in \mathbb{R}^n} \frac{\|Av\|_2}{\|v\|_2}$.

Here and throughout, when writing $x \ll y$, we mean that $x = o(y)$, $x \lesssim y$ means $x = O(y)$, and $x \asymp y$ means $x = \Theta(y)$.

$\mathbf{1}_X, \mathbf{1}(X)$ denote the indicator on the event X occurring.

We start with a key ingredient in our proofs for heavy-tailed weights in the form of a new deterministic bound on the largest eigenvalue in terms of the 'entry-wise' L^p-(quasi)norm of the matrix, generalizing the classical Motzkin-Straus theorem [MS65] corresponding to $p = 2$ case. The idea of proof in Section 6.1 explains why this bound is crucial in the case of heavy-tailed weights. We then state some basic spectral properties of graphs, some of which were originally stated for unweighted graphs, but generalize easily.

The first section of this chapter comes from [GHN24], while the second section contains results from both [GHN24] and [HM23], as well as results from other papers.

2.1 Spectral norm and L^p-(quasi)norm

For $p > 0$, we denote by $\|A\|_p$ the entry-wise L^p-(quasi)norm[1] of the symmetric matrix A:

$$\|A\|_p := \Big(\sum_{1\le i,j\le n} |a_{ij}|^p \Big)^{1/p}.$$

To state our bound for the largest eigenvalue of the symmetric matrix A in terms of $\|A\|_p$, we first recall the following auxiliary function which appeared in the statement of Theorem 6.3: For $\theta > 1$ and any integer $k \ge 2$, let

$$\phi_\theta(k) := \sup_{f=(f_1,\cdots,f_k):\|f\|_1=1} \sum_{i,j\in[k],i\ne j} |f_i|^\theta |f_j|^\theta. \tag{2.1}$$

We will assume without loss of generality that the vector f appearing above is non-negative.

Proposition 2.1. Suppose that $1 < p < 2$ and let $k \ge 2$ be an integer. Then, for any network $G = (V, E, A)$ such that the maximum size of clique contained in G is k,

$$\lambda_1(A) \le \phi_{\frac{p}{2(p-1)}}(k)^{\frac{p-1}{p}} \|A\|_p\,. \tag{2.2}$$

In the case $0 < p \le 1$, for any network $G = (V, E, A)$,

$$\lambda_1(A) \le 2^{-\frac{1}{p}} \|A\|_p\,. \tag{2.3}$$

Before proving this proposition, we state some useful lemmas. First, the next lemma identifies the structure that leads to the equality in the above expressions. We will need those characterizations when planting the structures that lead to an atypically large eigenvalue in the heavy-tailed edge-weights case.

[1] A quasinorm satisfies the norm axioms except that the triangle inequality is replaced by $\|x + y\| \le K(\|x\| + \|y\|)$ for some $K > 1$.

Lemma 2.1. Assume that $1 < p < 2$. Then, for any integer $k \geq 2$, there exist $k_1, k_2 \geq 0$ with $k_1 + k_2 \leq k$ and $x, y \geq 0$ such that if $G = (V, E)$ is a clique with $V = [k]$ and $A = (a_{ij})_{i,j\in[k]}$ is a block matrix given by

$$a_{ij} = \begin{cases} x^2 & i \neq j, i, j \in \{1, \cdots, k_1\} =: V_1, \\ y^2 & i \neq j, i, j \in \{k_1 + 1, \cdots, k_1 + k_2\} =: V_2, \\ xy & i \in V_1, j \in V_2 \text{ or } i \in V_2, j \in V_1, \\ 0 & \text{otherwise}, \end{cases} \tag{2.4}$$

then A satisfies the equality in (2.2).
In the case $0 < p \leq 1$, the equality in (2.3) holds when A is the adjacency matrix of a clique of size 2, i.e. a graph consisting of a single edge.

Note that this matrix is not unique, as for any tuple (k_1, k_2, x, y) and any constant c, the tuple (k_1, k_2, cx, cy) also satisfies the equality.

We defer the proof of this lemma to the end of this section. To prove Proposition 2.1, we rely on an alternative characterization of $\phi_\theta(k)$, which is given in the following lemma. It turns out that $\phi_\theta(k)$ is equal to the supremum of the same objective function (which we call $\widehat{\phi_\theta}(k)$) over all graphs G whose *maximum clique size* is k.

Lemma 2.2. For $\theta > 1$ and an integer $k \geq 2$, define

$$\widehat{\phi_\theta}(k) := \sup_{G=(V,E)} \sup_{f=(f_1,\cdots,f_{|V|}):\|f\|_1=1} \sum_{i,j\in[|V|], i\sim j} |f_i|^\theta |f_j|^\theta. \tag{2.5}$$

Here, the first supremum is taken over all graphs G whose maximum clique size is k. Recall that in the summation, $i \sim j$ means that vertices i and j are connected by an edge. Then, we have that

$$\widehat{\phi_\theta}(k) = \phi_\theta(k).$$

We defer the proof to the end of this section. Given this lemma, one can conclude the proof of Proposition 2.1.

Proof of Proposition 2.1. Let $V = [n]$ and a_{ij}s denote the edge-weights. By the variational characterization of the largest eigenvalue,

$$\lambda_1(A) = \sup_{\|f\|_2=1} \sum_{i\sim j} a_{ij} f_i f_j. \tag{2.6}$$

We now consider the different ranges of p. In the case $1 < p < 2$, we apply Hölder's inequality to bound the above quantity, whereas in the case $0 < p \leq 1$, we simply use the monotomicity of ℓ^p norms.

Case 1: $1 < p < 2$. Setting $q = \frac{p}{p-1} > 2$ to be the conjugate of p, by Hölder's inequality,

$$\sum_{i\sim j} a_{ij} f_i f_j \leq \Big(\sum_{i\sim j} |a_{ij}|^p\Big)^{\frac{1}{p}} \Big(\sum_{i\sim j} |f_i|^q |f_j|^q\Big)^{\frac{1}{q}}.$$

By the definition of $\widehat{\phi}_{\frac{q}{2}}$ and since $\phi_{\frac{q}{2}} = \widehat{\phi}_{\frac{q}{2}}$ (see Lemma 2.2), for any vector f such that $\|f\|_2 = 1$,

$$\sum_{i\sim j} |f_i|^q |f_j|^q \leq \widehat{\phi}_{\frac{q}{2}}(k) = \phi_{\frac{q}{2}}(k)$$

(note that in (2.5), supremum is taken over $\|f\|_1 = 1$). Therefore, we have

$$\lambda_1(A) \leq \phi_{\frac{q}{2}}(k)^{\frac{1}{q}} \|A\|_p = \phi_{\frac{p}{2(p-1)}}(k)^{\frac{p-1}{p}} \|A\|_p .$$

Case 2: $0 < p \leq 1$. Since $|f_i f_j| \leq \frac{1}{2}$ for any $i \neq j$, whenever $\|f\|_2 = 1$, by the monotonicity of ℓ^p norms,

$$\lambda_1(A) = 2 \sup_{\|f\|_2=1} \sum_{i<j, i\sim j} a_{ij} f_i f_j \leq \sum_{i<j, i\sim j} |a_{ij}| \leq \Big(\sum_{i<j, i\sim j} |a_{ij}|^p\Big)^{\frac{1}{p}} = 2^{-\frac{1}{p}} \Big(\sum_{i\sim j} |a_{ij}|^p\Big)^{\frac{1}{p}}.$$

□

We now establish some useful properties of the function ϕ_θ.

Lemma 2.3. Let $\theta > 1$. Then,

1. For each $k \geq 2$, there exists a k-dimensional vector f of the form $(x, \cdots, x, y, \cdots, y, 0, \cdots, 0)$ which attains the maximum of $\phi_\theta(k)$ in (2.1). In other words, there exist $k_1, k_2 \geq 0$ with $k_1 + k_2 \leq k$ and $x, y \geq 0$ such that $k_1 x + k_2 y = 1$ and
$$\phi_\theta(k) = \sum_{i,j\in[k], i\neq j} f_i^\theta f_j^\theta$$
holds with $f_1 = \cdots = f_{k_1} = x$, $f_{k_1+1} = \cdots = f_{k_1+k_2} = y$ and $f_{k_1+k_2+1} = \cdots = f_k = 0$.

2. For any $k \geq 2$,
$$\phi_\theta(k) \leq \Big(\frac{2\theta-2}{2\theta-1}\Big)^{2\theta-2} - \Big(\frac{2\theta-2}{2\theta-1}\Big)^{2\theta-1}. \tag{2.7}$$

3. For any $k \leq \frac{2\theta-1}{2\theta-2}$,
$$\phi_\theta(k) = \frac{1}{k^{2\theta-2}} - \frac{1}{k^{2\theta-1}}. \tag{2.8}$$
In addition,
$$\phi_\theta(2) = \frac{1}{2^{2\theta-1}}. \tag{2.9}$$

4. $\phi_\theta(k)$ is non-decreasing and becomes constant for large enough k.

We give a brief interpretation of this lemma. The statement (1) implies that for any $k \geq 2$, maximum of $\phi_\theta(k)$ in (2.1) is attained at the vector f with at most two distinct non-zero elements. (2) states a general upper bound for the function $\phi_\theta(k)$. (3) states that for $k \leq \frac{2\theta-1}{2\theta-2}$, the maximum of $\phi_\theta(k)$ is attained at the k-dimensional vector $(\frac{1}{k}, \cdots, \frac{1}{k})$. The fact that $\phi_\theta(k)$ becomes constant for large k, stated in (4), is a crucial ingredient in our analysis.

Proof of Lemma 2.3.

Proof of (1)**.** We assume, without loss of generality, that the supremum is taken over all k-tuples $(f_1, \cdots, f_k)$ with $\sum_{i=1}^k f_i = 1$ and $f_i \geq 0$. Since the collection of such k-tuples is a compact set, the function $(f_1, \cdots, f_k) \mapsto \sum_{i,j\in[k],i\neq j} f_i^\theta f_j^\theta$ over this set attains its maximum. Note that there may be several k-tuples which attain the maximum, and in this case we arbitrarily choose one of them. We further assume, without loss of generality, that for some integer $1 \leq \ell \leq k$, the maximum is attained in the interior of the $(\ell-1)$-dimensional simplex $f_1 + \cdots + f_\ell = 1$ with $f_{\ell+1} = \cdots = f_k = 0$ (i.e. $0 < f_1, \cdots, f_\ell < 1$). By the Lagrange multiplier theorem [Ste15, Chapter 14.8], which states that the gradient of the objective function at a local extreme point is a scalar multiple of the gradient of the constraint function, applied to our maximization problem on this simplex, setting $s := f_1^\theta + \cdots + f_\ell^\theta$, we have

$$(s - f_1^\theta) f_1^{\theta-1} = \cdots = (s - f_\ell^\theta) f_\ell^{\theta-1}.$$

Defining $g_i := f_i^{\theta-1}$ and $\bar{\theta} := \frac{2\theta-1}{\theta-1} > 1$, this implies that the quantities $\frac{g_i^{\bar{\theta}} - g_j^{\bar{\theta}}}{g_i - g_j}$ are all equal to s for any $1 \leq i < j \leq \ell$ for which $g_i \neq g_j$. By the mean value theorem applied to the convex function $x \mapsto x^{\bar{\theta}}$, one can deduce that there are at most two distinct (non-zero) values that g_is (and thus f_is) for $1 \leq i \leq \ell$ can take. This concludes the proof of (1).

Proof of (2)**.** Applying Hölder's inequality and using that $\sum_{i=1}^k f_i = 1$,

$$\sum_{i=1}^k f_i^\theta \leq \left(\sum_{i=1}^k f_i\right)^{\frac{\theta}{2\theta-1}} \left(\sum_{i=1}^k f_i^{2\theta}\right)^{\frac{\theta-1}{2\theta-1}} = \left(\sum_{i=1}^k f_i^{2\theta}\right)^{\frac{\theta-1}{2\theta-1}}. \tag{2.10}$$

Thus, setting $r := \sum_{i=1}^k f_i^{2\theta}$, we have

$$\sum_{i,j\in[k],i\neq j} f_i^\theta f_j^\theta = \left(\sum_{i=1}^k f_i^\theta\right)^2 - \sum_{i=1}^k f_i^{2\theta} \leq r^{\frac{2\theta-2}{2\theta-1}} - r. \tag{2.11}$$

The function $t \mapsto t^{\frac{2\theta-2}{2\theta-1}} - t$ is increasing on $\left(0, \left(\frac{2\theta-2}{2\theta-1}\right)^{2\theta-1}\right)$ and decreasing on $\left(\left(\frac{2\theta-2}{2\theta-1}\right)^{2\theta-1}, \infty\right)$. Thus,

$$\sum_{i,j\in[k],i\neq j} f_i^\theta f_j^\theta \leq \left(\frac{2\theta-2}{2\theta-1}\right)^{2\theta-2} - \left(\frac{2\theta-2}{2\theta-1}\right)^{2\theta-1}. \tag{2.12}$$

Note that the equality above may not be attained in general. In fact, if the equality is attained, then by the equality condition in the Hölder's inequality (2.10), $f_1 = \cdots = f_m = \frac{1}{m}$ and $f_{m+1} = \cdots = f_k = 0$ for some integer $1 \le m \le k$. Also, in order that (2.12) becomes an equality, $\sum_{i=1}^k f_i^{2\theta} = r = \left(\frac{2\theta-2}{2\theta-1}\right)^{2\theta-1}$. Since $\sum_{i=1}^k f_i^{2\theta} = m \cdot (\frac{1}{m})^{2\theta} = (\frac{1}{m})^{2\theta-1}$, the equality in (2.12) is possible only when $\frac{2\theta-1}{2\theta-2}$ is a positive integer.

Proof of (3)**.** We first prove the first part of the statement. By Hölder's inequality,

$$1 = \sum_{i=1}^{k} f_i \le \Big(\sum_{i=1}^{k} f_i^{2\theta}\Big)^{\frac{1}{2\theta}} k^{1-\frac{1}{2\theta}}. \tag{2.13}$$

Hence, if $k \le \frac{2\theta-1}{2\theta-2}$, then $r = \sum_{i=1}^k f_i^{2\theta} \ge \left(\frac{1}{k}\right)^{2\theta-1} \ge \left(\frac{2\theta-2}{2\theta-1}\right)^{2\theta-1}$. Thus, recalling that the function $t \mapsto t^{\frac{2\theta-2}{2\theta-1}} - t$ is decreasing on $\left(\left(\frac{2\theta-2}{2\theta-1}\right)^{2\theta-1}, \infty\right)$ and using $r \ge (\frac{1}{k})^{2\theta-1}$,

$$\sum_{i,j\in[k], i\neq j} f_i^\theta f_j^\theta \overset{(2.11)}{\le} r^{\frac{2\theta-2}{2\theta-1}} - r \le \frac{1}{k^{2\theta-2}} - \frac{1}{k^{2\theta-1}}.$$

One can also deduce that the maximum of $\phi_\theta(k)$ is attained when $f_1 = \cdots = f_k = \frac{1}{k}$. To see this, by the above inequality, if f attains the maximum, then $r = (\frac{1}{k})^{2\theta-1}$, which implies that f satisfies the equality in (2.13). By the equality condition in the Hölder's inequality, all the f_is are same and thus $f_1 = \cdots = f_k = \frac{1}{k}$.

The second statement (2.9) follows from the fact that $|xy| \le \frac{1}{4}$ whenever $|x| + |y| = 1$.

Proof of (4)**.** It is straightforward to observe that $\phi_\theta(k)$ is non-decreasing in k since as k increases, the supremum is taken over a larger class of vectors f. We now show that $\phi_\theta(k)$ becomes constant for large enough k. By considering $f_1 = \cdots = f_k = \frac{1}{k}$, we have

$$\phi_\theta(k) \ge \frac{1}{k^{2\theta-2}} - \frac{1}{k^{2\theta-1}}.$$

Since $\phi_\theta(k)$ is non-decreasing in k and the function $t \mapsto \frac{1}{t^{2\theta-2}} - \frac{1}{t^{2\theta-1}}$ is decreasing for any large enough t, we have that for large enough k,

$$\phi_\theta(k) \ge \sup_n \Big(\frac{1}{n^{2\theta-2}} - \frac{1}{n^{2\theta-1}}\Big).$$

If the equality holds for all large enough k, we are done. Otherwise, there is $K_0 > 0$ such that

$$\phi_\theta(k) > \sup_n \Big(\frac{1}{n^{2\theta-2}} - \frac{1}{n^{2\theta-1}}\Big) \tag{2.14}$$

for all $k \ge K_0$. We show that this implies that the the support of the maximizer f of $\phi_\theta(k)$ in (2.1) is uniformly bounded in k.

By the statement (1) of this lemma, for any k, the maximum of the objective function $\phi_\theta(k)$ is attained at some vector $f = (f_i)_i$ with $f_1 = \cdots = f_{k_1} = x$, $f_{k_1+1} = \cdots = f_{k_1+k_2} = y$, and $f_{k_1+k_2+1} = \cdots = f_k = 0$ for some $k_1, k_2 \geq 0$ with $k_1 + k_2 \leq k$ and $x, y \geq 0$. Thus, if the support of f is not bounded in k (without loss of generality we assume $k_2 \to \infty$), then for sufficiently small $\iota > 0$,

$$\begin{aligned}\phi_\theta(k) = 2\binom{k_1}{2}x^{2\theta} + 2\binom{k_2}{2}y^{2\theta} + 2k_1k_2x^\theta y^\theta &\leq 2\binom{k_1}{2}\frac{1}{k_1^{2\theta}} + \Big[2\binom{k_2}{2}\frac{1}{k_2^{2\theta}} + 2k_2\frac{1}{k_2^\theta}\Big] \\ &\leq \frac{1}{k_1^{2\theta-2}} - \frac{1}{k_1^{2\theta-1}} + \iota \\ &\leq \sup_n\Big(\frac{1}{n^{2\theta-2}} - \frac{1}{n^{2\theta-1}}\Big) + \iota \overset{(2.14)}{<} \phi_\theta(k),\end{aligned}$$

where we used $k_1x^\theta \leq k_1x \leq 1, y \leq \frac{1}{k_2}$ in the first inequality and $k_2 \to \infty, \theta > 1$ in the second inequality. The final RHS bound yields a contradiction. Hence, we conclude that the support of the vector f which maximizes the objective function $\phi_\theta(k)$ is uniformly bounded in k. This implies that $\phi_\theta(k)$ becomes constant for large k. □

Using this lemma, one can establish Lemma 2.1 which provides the equality condition of the inequalities in Proposition 2.1.

Proof of Lemma 2.1. Let us first consider the case $1 < p < 2$. Let $q > 2$ be the conjugate of p. By Lemma 2.3, there exist $k_1, k_2 \geq 0$ with $k_1 + k_2 \leq k$ and $x, y \geq 0$ such that the vector $f = (f_1, \cdots, f_{[k]})$ defined by

$$f_i = \begin{cases} x & i \in \{1, \cdots, k_1\} =: V_1, \\ y & i \in \{k_1+1, \cdots, k_1+k_2\} =: V_2, \\ 0 & \text{otherwise} \end{cases}$$

satisfies $\left\|f\right\|_1 = 1$ and

$$\phi_{\frac{q}{2}}(k) = \sum_{i,j\in[k], i\neq j} f_i^{\frac{q}{2}} f_j^{\frac{q}{2}}.$$

Now define the vector $\tilde{f} = \left(\tilde{f}_1, \cdots, \tilde{f}_k\right)$ by setting $\tilde{f}_i = \sqrt{f_i}$ so that $\|\tilde{f}\|_2 = \left\|f\right\|_1 = 1$. Next, define the $k \times k$ matrix $A = (a_{ij})_{i,j\in[k]}$ by

$$a_{ij} := \begin{cases} x^{\frac{1}{p-1}} & i \neq j, i, j \in V_1, \\ y^{\frac{1}{p-1}} & i \neq j, i, j \in V_2, \\ x^{\frac{1}{2(p-1)}} y^{\frac{1}{2(p-1)}} & i \in V_1, j \in V_2 \text{ or } i \in V_2, j \in V_1, \\ 0 & \text{otherwise.} \end{cases}$$

Note that we defined A so that $a_{ij}^p = f_i^{\frac{q}{2}} f_j^{\frac{q}{2}} = \tilde{f}_i^q \tilde{f}_j^q$ for $i \neq j$. This implies that

$$\|A\|_p = \left(\sum_{i,j\in[k],i\neq j} \tilde{f}_i^q \tilde{f}_j^q \right)^{\frac{1}{p}} = \left(\sum_{i,j\in[k],i\neq j} f_i^{\frac{q}{2}} f_j^{\frac{q}{2}} \right)^{\frac{1}{p}} = \phi_{\frac{q}{2}}(k)^{\frac{1}{p}} = \phi_{\frac{p}{2(p-1)}}(k)^{\frac{1}{p}} \tag{2.15}$$

and

$$\sum_{i,j\in[k],i\neq j} a_{ij}\tilde{f}_i\tilde{f}_j = \sum_{i,j\in[k],i\neq j} \tilde{f}_i^q \tilde{f}_j^q = \sum_{i,j\in[k],i\neq j} f_i^{\frac{q}{2}} f_j^{\frac{q}{2}} = \phi_{\frac{q}{2}}(k) = \phi_{\frac{p}{2(p-1)}}(k). \tag{2.16}$$

By the variational characterization of the largest eigenvalue (recall that $\|\tilde{f}\|_2 = 1$) and Proposition 2.1,

$$\phi_{\frac{p}{2(p-1)}}(k) \overset{(2.16)}{=} \sum_{i,j\in[k],i\neq j} a_{ij}\tilde{f}_i\tilde{f}_j \leq \lambda_1(A) \leq \phi_{\frac{p}{2(p-1)}}(k)^{\frac{p-1}{p}} \|A\|_p \overset{(2.15)}{=} \phi_{\frac{p}{2(p-1)}}(k),$$

which establishes that equality holds in (2.2).

Now, let us consider the case $0 < p \leq 1$. Note the largest eigenvalue of a network with a single edge is nothing other than the edge-weight which we call a. Since $\|A\|_p = 2^{\frac{1}{p}}a$ for the 2×2 matrix $A = \begin{pmatrix} 0 & a \\ a & 0 \end{pmatrix}$, which corresponds to a single edge with weight a, we obtain the equality in (2.3). □

We now establish Lemma 2.2, which claims that our two characterizations of ϕ_θ and $\widehat{\phi}_\theta$ are equivalent.

Proof of Lemma 2.2. Assuming that G is not a clique of size k, we can choose two vertices v_1 and v_2 that are not connected by an edge in G. Without loss of generality, we assume that $\sum_{i\sim v_1} f_i^\theta \geq \sum_{j\sim v_2} f_j^\theta$. Since

$$\sum_{i\sim j} f_i^\theta f_j^\theta = \Big(\sum_{i\sim v_1} f_i^\theta\Big) f_{v_1}^\theta + \Big(\sum_{j\sim v_2} f_j^\theta\Big) f_{v_2}^\theta + \sum_{i,j\neq v_1,v_2,i\sim j} f_i^\theta f_j^\theta,$$

the objective function does not decrease when we move the weight from v_2 to v_1 by replacing $f = (\cdots, f_{v_1}, \cdots, f_{v_2}, \cdots)$ by $f^{(1)} = (\cdots, f_{v_1} + f_{v_2}, \cdots, 0, \cdots)$. This follows from the fact that for $\theta > 1$, if $a \geq b \geq 0$, then $\max_{x+y=s,x,y\geq 0} ax^\theta + by^\theta = a \cdot s^\theta + b \cdot 0^\theta = as^\theta$.

After removing the zero at v_2, we obtain a new vector $f^{(1)}$ on the new graph G_1 obtained by a deletion of the vertex v_2 and the edges incident on it. We repeat this procedure to get a sequence of vectors $f^{(1)}, \cdots, f^{(m)}$ and graphs $G_1, \cdots, G_m$ such that G_{i+1} is obtained by the removal of some vertex w_{i+1} together with the edges incident on it in G_i. This procedure can be repeated until G_m becomes a clique, showing that the maximum of $\widehat{\phi}_\theta(k)$ is attained at a clique of size k. □

We finish this section by introducing a technical lemma that will be useful when estimating the largest eigenvalue of tree-like networks.

Lemma 2.4. Suppose that G is a tree with a vertex set $[n]$. Let $s, \xi > 0$ and $\theta \geq 1$ be constants. Then, for any vector $f = (f_1, \cdots, f_n)$ with $\sum_{i=1}^n f_i = s$ and $0 \leq f_i \leq \xi$ for all i,

$$\sum_{i<j, i\sim j} f_i^\theta f_j^\theta \leq \begin{cases} \frac{1}{4}s^{2\theta} & \text{if } s < 2\xi, \\ \xi^\theta (s-\xi)^\theta & \text{if } s \geq 2\xi. \end{cases} \tag{2.17}$$

In particular, if $0 < \xi \leq \frac{1}{2}$, $\theta \geq 1$, $\sum_{i=1}^n f_i^2 = 1$ and $0 \leq f_i \leq \xi$, then

$$\sum_{i<j, i\sim j} f_i^{2\theta} f_j^{2\theta} \leq \xi^{2\theta}. \tag{2.18}$$

Note that the above estimates still hold even when G is a forest (i.e. vertex-disjoint union of trees), since f_is are all non-negative and one can easily construct a tree on the same vertex set as G with the latter as a subgraph and apply the above result.

Proof. Let $\rho := \arg\max_i f_i$ (there may be several vertices which attain the maximum, and in this case we choose any of them) and regard G as a tree rooted at ρ. Then, since every edge can be seen as connecting a vertex i to its unique parent p_i,

$$\sum_{i<j, i\sim j} f_i^\theta f_j^\theta = \sum_{(i,j)\in E(G)} f_i^\theta f_j^\theta = \sum_{i\neq\rho} f_{p_i}^\theta f_i^\theta \leq \sum_{i\neq\rho} f_\rho^\theta f_i^\theta \leq f_\rho^\theta (s - f_\rho)^\theta,$$

where the last inequality follows from the fact that $x_1^\theta + \cdots + x_m^\theta \leq (x_1 + \cdots + x_m)^\theta$ whenever $\theta \geq 1$ and $x_1, \cdots, x_m \geq 0$, which itself is a straightforward consequence of convexity of the function $x \to x^\theta$ on the positive real line. Finally, since the function $x \mapsto x(s-x)$ is increasing on $[0, \frac{s}{2}]$ and $f_\rho \leq \xi$, (2.17) follows.

(2.18) is a direct consequence of (2.17), by replacing f_i with f_i^2 and setting $s = 1$. □

2.2 Basic spectral properties of graphs

We end this chapter by introducing basic but crucial spectral properties of general weighted graphs. The first two statements of the following lemma appear as [KS03, Proposition 3.1] for unweighted graphs, but it is straightforward to see that a weighted version holds as well. The last statement follows from the variational characterization of the largest eigenvalue, since for any subset of vertices W and a subnetwork A_W induced by these vertices,

$$\lambda_1(A_W) = \sup_{\text{unit vector } v \text{ supported on } W} v^\intercal A v \leq \lambda_1(A).$$

Lemma 2.5. Let $G = (V, E, A)$ be any network. Suppose that $G_1, \cdots, G_k$ are subgraphs of G and let $A_1, \cdots, A_k$ be the corresponding networks. If $E(G) = \cup_{i=1}^{k} E(G_i)$, then $\lambda_1(A) \leq \sum_{i=1}^{k} \lambda_1(A_i)$. If in addition, the graphs $G_1, \cdots, G_k$ are vertex disjoint, then $\lambda_1(A) = \max_{i=1,\cdots,k} \lambda_1(A_i)$. Moreover, for any network A_1 induced by a subset of the vertices of G, we have that $\lambda_1(A) \geq \lambda_1(A_1)$.

The next lemma characterizes the largest eigenvalue of the weighted star graph in terms of the Frobenius norm of its weight matrix.

Lemma 2.6. If $G = (V, E, A)$ is a weighted star graph with $s + 1$ vertices and edge-weights $w_1, \cdots, w_s$, then $\lambda_1(A) = \sqrt{\sum_{j=1}^{s} w_j^2}$.

Proof. Note that A is of the form

$$A = \begin{pmatrix} 0 & w_1 & \dots & w_s \\ w_1 & 0 & & \\ \vdots & & \ddots & \\ w_s & & & \end{pmatrix}.$$

Setting $u := (1, 0, \dots, 0)$ and $v := \left(\sum_i w_i^2\right)^{-\frac{1}{2}} (0, w_1, \dots, w_s)$, one can write $A = \left(\sum_i w_i^2\right)^{\frac{1}{2}} (uv^T + vu^T)$. Since u and v are unit vectors satisfying $\langle u, v \rangle = 0$, the spectrum of $uv^T + vu^T$ counted with multiplicity is exactly $\{1, -1, 0, \dots, 0\}$, which implies the result. □

The next lemma bounds the largest eigenvalue of tree in terms of its maximum degree.

Lemma 2.7. [Kes59] If T is an unweighted forest with maximum degree bounded by Δ, then $\lambda_{\max}(A_T) \leq 2\sqrt{\Delta - 1}$.

The result below gives an upper bound for the largest eigenvalue of symmetric matrices.

Lemma 2.8. For any symmetric matrix $A = (a_{ij})_{i,j}$ whose diagonal entries are all zero,

$$\lambda_1(A) \geq \max_{i \neq j} |a_{ij}|. \tag{2.19}$$

Proof. Let $|a_{k\ell}|$ be the maximal value, i.e. $|a_{k\ell}| = \max_{i \neq j} |a_{ij}|$. Let the vector $\mathbf{v} = (v_i)_i$ be defined by $v_k = \frac{1}{\sqrt{2}}$, $v_\ell = \frac{\operatorname{sgn}(a_{k\ell})}{\sqrt{2}}$ and $v_i = 0$ otherwise, then $\|\mathbf{v}\|_2 = 1$ and $\mathbf{v}^T A \mathbf{v} = |a_{k\ell}|$. The result now follows by the variational formulation for $\lambda_1(A)$ applied to the vector $\mathbf{v}$. □

We conclude this section by citing a result that quantifies the proximity of true eigenvalues and eigenvectors to approximate ones.

Lemma 2.9 ([ADK23b], Lemma 4.10)**.** Consider a self-adjoint matrix M and $\Delta, \epsilon > 0$ satisfying $5\epsilon \leq \Delta$. For $\lambda \in \mathbb{R}$, assume M has a unique eigenvalue μ in the interval $[\lambda - \Delta, \lambda + \Delta]$ with eigenvector $\mathbf{w}$. If there is a normalized vector $\mathbf{v}$ such that $\|(M - \lambda)\mathbf{v}\| \leq \epsilon$, then

$$\mu - \lambda = \langle \mathbf{v}, (M - \lambda)\mathbf{v} \rangle + O\left(\frac{\epsilon^2}{\Delta}\right), \|\mathbf{w} - \mathbf{v}\| = O\left(\frac{\epsilon}{\Delta}\right)$$

In particular this lemma will be used in the following form when describing the spectrum of sparse Erdős-Rényi graphs in detail.

Lemma 2.10. For $i \geq 2$, let A be the adjacency matrix of the ball of radius i around a vertex x of degree α, such that $B_i(x)$ is an unweighted tree, $\frac{|S_2(x)|}{\alpha} \leq s \ll \alpha$ and the degree of each vertex in $B_i(x) \backslash \{x\}$ is at most $t \leq \frac{\alpha}{5}$.

Then the maximum eigenvalue μ of A satisfies

$$\mu = \sqrt{\alpha} + O\left(\frac{s}{\sqrt{\alpha}}\right).$$

Proof. We take as our test vector $\mathbf{w}$ the eigenvector corresponding to the star graph consisting of the central vertex x and its neighbors. Thus $\mathbf{w}|_x = \frac{1}{\sqrt{2}}$ and for $y \sim x$, $\mathbf{w}|_y = \frac{1}{\sqrt{2\alpha}}$. Since $B_i(x)$ is a tree, each $z \in S_2(x)$ has exactly one neighbor in $S_1(x)$, so $(A\mathbf{w})|_z = \frac{1}{\sqrt{2\alpha}}$. Moreover the number of non-zero entries in $A\mathbf{w} - \sqrt{\alpha}\mathbf{w}$ is $|S_2(x)| \leq \alpha s$.

The above implies that

$$\|A\mathbf{w} - \sqrt{\alpha}\mathbf{w}\| \leq \sqrt{\alpha s \frac{1}{2\alpha}} = \sqrt{\frac{s}{2}},$$

which corresponds to ε in Lemma 2.9.

To utilize Lemma 2.9 we require Δ such that A has a unique eigenvalue in $[\lambda - \Delta, \lambda + \Delta]$. For this we use eigenvalue interlacing: after deleting the row and column of A corresponding to x, the matrix is the adjacency matrix of a forest with degree at most t. By the spectral radius of a tree from Lemma 2.7, the maximum eigenvalue of this submatrix is at most $2\sqrt{t}$. Thus A has at most one eigenvalue in the interval $\left[2.1\sqrt{t}, 2\sqrt{\alpha} - 2.1\sqrt{t}\right]$, namely, if any, its maximum eigenvalue. Thus we can take $\Delta = \sqrt{\alpha} - 2.1\sqrt{t} \geq \left(1 - 2.1/\sqrt{5}\right)\sqrt{\alpha}$ in Lemma 2.9. The estimates on the errors now simply follow from plugging in our values for ε and Δ, since $\varepsilon = \sqrt{\frac{s}{2}} \ll \Delta$ by assumption. □

Chapter 3

Structure of sparse Erdős-Rényi graphs

One of the oldest and most studied random graph models is the Erdős-Rényi random graph $\mathcal{G}_{N,p}$, where there is an edge between any two of N vertices independently with probability p. How this graph looks like changes a lot according to the dependence of p on N, for instance with high probability the graph is connected if $p \gg \frac{\log N}{N}$ and disconnected if $p \ll \frac{\log N}{N}$ [ER60]. See the monograph of Guionnet for an overview of known results and the state of the field for this model [Gui21].

In terms of many applications, the regime of p of most interest is the sparse regime (i.e. $p \to 0$ as $N \to \infty$). In particular, the constant average degree regime of sparsity $p = \frac{d}{N}$ ($d > 0$ is a constant), i.e., when the typical number of connections of a single vertex tends to stay constant, arises naturally in several models in statistical mechanics (see [DM10] for a comprehensive treatment of statistical physics models on such sparse graphs).

Before embarking on the study of the spectrum of sparse Erdős-Rényi graphs, we now give some background on the structure of sparse Erdős-Rényi graphs, which will be useful later. The results below apply to $\mathcal{G}\left(N, \frac{d}{N}\right)$. The dependence of d on N will be specified in each result or section.

The combinatorial aspects of an Erdős-Rényi graph are governed by binomial distributions. Therefore, by $\mathrm{Bin}(k; N, p)$ we denote the probability that a binomial random variable with N trials and success probability p is equal to k.

3.1 Profile of large degrees

All new results in this section are from the paper [GHN24].

We first record that the maximum degree is almost deterministic when d is small.

Lemma 3.1 ([Bol01], Theorem 3.7)**.** Define $\mu_k = N\mathrm{Bin}(k; N-1, \frac{d}{N})$ to be the expected number of vertices of degree k in the graph. Now define

$$\mathfrak{u} := \arg\min_{k \in \mathbb{Z}} \left\{ \max\left\{ \mu_k, \mu_k^{-1} \right\} \right\}. \tag{3.1}$$

Then if $d = o(\log N)$, with high probability the maximum degree is in $\{\mathfrak{u}-1, \mathfrak{u}\}$.

In order to calculate $\mathfrak{u}$, note that by the Stirling approximation, having $\mu_k \approx 1$ implies

$$\Big(1 + o_N(1)\Big) \log N - \mathfrak{u} \log \mathfrak{u} + \mathfrak{u} - d + \mathfrak{u} \log d - \frac{1}{2} \log(2\pi\mathfrak{u}) = 0.$$

Therefore, in our regime of d,

$$\mathfrak{u} = \Big(1 + o_N(1)\Big) \frac{\log N}{\log\log N - \log d}. \tag{3.2}$$

and $\mathfrak{u} = \Theta\Big(\frac{\log N}{\log\log N}\Big)$.

In this part, all subsequent results assume that d is constant. Regarding large deviations of the largest degrees of $\mathcal{G}\Big(N, \frac{d}{N}\Big)$ we have the following result.

Proposition 3.1. [BBG21, Proposition 1.3] Let us denote by d_s, the s-th largest degree of $G = \mathcal{G}_{N,\frac{d}{N}}$ with d constant. Then, setting

$$t_N := \frac{\log N}{\log\log N},$$

we have, for any $\delta_1, \ldots, \delta_p \geq 0$,

$$\lim_{n\to\infty} \frac{-\log \mathbb{P}\left(d_1 \geq (1+\delta_1)t_N, \ldots, d_r \geq (1+\delta_r)t_N\right)}{\log N} = \sum_{s=1}^{r} \delta_s. \tag{3.3}$$

For our result on the large deviations of weighted Erdős-Rényi graphs, we need a more precise description of the degree distribution. An important input in our arguments will be that, for any constant $\kappa > 0$:

1. For any fixed $0 < \gamma < 1$, with high probability, there exist $N^{1-\gamma-\kappa}$ vertices having at least $\gamma \frac{\log N}{\log\log N}$ neighbors with no edges between each other. This is captured by Proposition 3.2.

2. For a suitable discretization $\{\gamma_i\}_{i=1,2,\cdots}$ of $(0, 1)$, with high probability, the number of vertices of degree between $\gamma_i \frac{\log N}{\log\log N}$ and $\gamma_{i+1} \frac{\log N}{\log\log N}$ is at most $N^{1-\gamma_i+\kappa}$ for all $i = 1, 2, \cdots$. This is formalized in Proposition 3.4.

In order to establish (1), we first estimate the probability that a vertex has a large degree in $\mathcal{G}_{N,\frac{d}{N}}$. For a subset $L \subseteq V = V(G)$, we denote by $d_L(v)$ the number of vertices in L connected to v. Throughout the thesis to simplify the notation, for $\gamma \geq 0$, define

$$g(\gamma) := \Big\lceil \gamma \frac{\log N}{\log\log N} \Big\rceil. \tag{3.4}$$

By a well-known binomial tail estimate (see Lemma 4.5 in the Appendix), for any vertex v,

$$\mathbb{P}(d(v) \geq g(\gamma)) = N^{-\gamma+o(1)}. \tag{3.5}$$

In the next lemma, we state a simple extension of (3.5), i.e. $d(v)$ replaced with $d_L(v)$ for general subsets $L \subseteq V$. Although a straightforward consequence of a binomial tail estimate, we provide a proof, which consists of straightforward but tedious algebra, for the sake of completeness.

Lemma 3.2. For $0 < \rho < 1$, let L be any subset of V of size $\lfloor \rho N \rfloor$. Then for any vertex $v \in L^c$ and $\gamma > 0$,

$$\mathbb{P}\left(d_L(v) \geq g(\gamma)\right) = N^{-\gamma+o(1)}. \tag{3.6}$$

Note that this probability does not depend on the parameter ρ, which shows that as long as $|L|$ is of order N, the probability is of the same order.

Proof. Since $d_L(v)$ is distributed as $\text{Bin}\left(\lfloor \rho N \rfloor, \frac{d}{N}\right)$, by the mentioned bound in Lemma 4.5, setting $\theta := \frac{1}{\lfloor \rho N \rfloor} g(\gamma) = \frac{1}{\lfloor \rho N \rfloor} \lceil \gamma \frac{\log N}{\log \log N} \rceil$,

$$\frac{1}{\sqrt{8\lfloor \rho N \rfloor \theta(1-\theta)}} e^{-\lfloor \rho N \rfloor I_{\frac{d}{N}}(\theta)} \leq \mathbb{P}\left(d_L(v) \geq g(\gamma)\right) \leq e^{-\lfloor \rho N \rfloor I_{\frac{d}{N}}(\theta)}.$$

Since $\frac{d}{N} = o\left(\frac{1}{\lfloor \rho N \rfloor} \lceil \gamma \frac{\log N}{\log \log N} \rceil\right)$, by the relative entropy estimate (see Lemma 4.4),

$$I_{\frac{d}{N}}(\theta) = (1 + o(1)) \frac{1}{\lfloor \rho N \rfloor} \left\lceil \gamma \frac{\log N}{\log \log N} \right\rceil \log\left(\frac{n}{d} \frac{1}{\lfloor \rho N \rfloor} \left\lceil \gamma \frac{\log N}{\log \log N} \right\rceil\right).$$

Thus, $\lfloor \rho N \rfloor I_{\frac{d}{N}}(\theta) = \lfloor \rho N \rfloor \frac{(1+o(1))\gamma \log N}{\lfloor \rho N \rfloor} = (\gamma + o(1)) \log N$.

The correction term in the lower bound, namely $\frac{1}{\sqrt{8\lfloor \rho N \rfloor \theta(1-\theta)}}$, is $N^{o(1)}$, which implies the matching bound $N^{-\gamma+o(1)}$.

□

We now proceed to estimate the number of such high-degree vertices satisfying additional useful properties.

Proposition 3.2. For $0 < \gamma, \rho < 1$, let $\mathcal{A}_{\gamma,\rho}$ be the event that there exist $m := \lceil N^{1-\gamma-\rho} \rceil$ vertices $v_1, \cdots, v_m$ and m subsets $W_1, \cdots, W_m \subseteq V$ of size $g(\gamma)$ satisfying the following properties:

1. Vertices $v_1, \cdots, v_m$ and elements in $W_1, \cdots, W_m$ are all distinct.
2. For each i, the vertex v_i is connected to all the elements in W_i.

3. For each i, there are no edges within W_i.

Then,

$$\mathbb{P}(\mathcal{A}_{\gamma,\rho}) \geq 1 - e^{-N^{1-\gamma-\rho+o(1)}}.$$

Proof. Let us partition the set of vertices into two subsets $S := \{s_1, \ldots, s_{\lceil \frac{N}{2} \rceil}\}$, which are the potential centers of the stars, and $L := S^c = \{\ell_1, \ldots, \ell_{\lfloor \frac{N}{2} \rfloor}\}$, which will be the potential leaves of the stars. The ordering on these two sets of vertices is arbitrary and only necessary so that the following algorithm is well-defined. Now we sequentially reveal the neighbors of vertex s_1 in L, by first checking whether ℓ_1 is its neighbor, and so on. Then,

1. We either obtain $g(\gamma)$ neighbors of s_1 before all edges from s_1 to vertices in L are revealed, or

2. There are less than $g(\gamma)$ neighbors of s_1 in L.

In the first case, we *mark* s_1 and define L_1 to be the collection of the first $g(\gamma)$ revealed vertices connected to s_1. In the second case, we do not mark s_1 and set $L_1 = \varnothing$.

Assume that we implemented the above process up to the k-th vertex s_k in S and obtained subsets $L_1, \cdots, L_k \subseteq L$. We then proceed similarly for s_{k+1}, but we only reveal edges from s_{k+1} to vertices in $L \setminus \cup_{i=1}^{k} L_i$. This guarantees that L_is are all disjoint. As before, we mark s_{k+1} and define L_{k+1} to be the collection of the first $g(\gamma)$ revealed vertices connected to s_{k+1} in the former case, and set $L_{k+1} = \varnothing$ in the latter case. We stop this process either once $\lceil N^{1-\gamma-\rho} \rceil$ vertices in S are marked, in which case we consider the process to be *successful*, or once we revealed edges to vertices in L for all vertices in S.

Let $\mathcal{B}$ be the event that the this revealing process is successful. We now show that this event happens with high probability. Since we discard exactly $g(\gamma)$ vertices in L for each marked vertex, at each k-th step, the set $L \setminus \cup_{i=1}^{k-1} L_i$ contains at least $\left\lfloor \frac{N}{2} \right\rfloor - N^{1-\gamma-\rho} g(\gamma) \geq \frac{N}{4}$ vertices, as long as n is large enough. Since the edges we reveal at each step are independent of any edges that have been revealed before, by Lemma 3.2 the probability that s_k has at least $g(\gamma)$ neighbors in $L \setminus \cup_{i=1}^{k-1} L_i$ is $N^{-\gamma+w}$ for some $w = o(1)$. Hence

$$\mathbb{P}(\mathcal{B}) \geq \mathbb{P}\left(\text{Binom}\left(\left\lceil \frac{N}{2} \right\rceil, N^{-\gamma+w}\right) \geq N^{1-\gamma-\rho}\right),$$

and thus, by Lemma 4.5,

$$\mathbb{P}(\mathcal{B}) \geq 1 - \exp\Big(-\left\lceil \frac{N}{2} \right\rceil I_{n^{-\gamma+w}}\Big(\frac{n^{1-\gamma-\rho}}{\lceil \frac{n}{2} \rceil}\Big)\Big).$$

Since $\frac{N^{1-\gamma-\rho}}{\lceil \frac{N}{2} \rceil} \leq \frac{1}{2} N^{-\gamma+w}$ for large N, by Lemma 4.3, there exists a constant $c > 0$ such that

$$I_{N^{-\gamma+w}}\left(\frac{N^{1-\gamma-\rho}}{\lceil \frac{N}{2} \rceil}\right) \geq cN^{-\gamma+w} = N^{-\gamma+o(1)}.$$

Thus,

$$\mathbb{P}(\mathcal{B}) \geq 1 - e^{-\lceil \frac{N}{2} \rceil N^{-\gamma+o(1)}} \geq 1 - e^{-N^{1-\gamma+o(1)}}.$$

Conditioned on the event $\mathcal{B}$, let us enumerate the marked vertices by $v_1, \cdots, v_{\lceil N^{1-\gamma-\rho} \rceil}$ and the collection of $g(\gamma)$ neighbors that we revealed by $W_1, \cdots, W_{\lceil N^{1-\gamma-\rho} \rceil}$ respectively. We call a vertex v_i *good* if there are no edges within W_i. Since having edges within W_i is independent of the revealing process, for large enough n,

$$\mathbb{P}(v_i \text{ is good}|\mathcal{B}) = \left(1 - \frac{d}{N}\right)^{\binom{g(\gamma)}{2}} \geq \frac{1}{2}. \tag{3.7}$$

Since W_is for $i = 1, \ldots, \lceil N^{1-\gamma-\rho} \rceil$ are disjoint, by independence, the number of good vertices stochastically dominates $\text{Binom}\left(\lceil N^{1-\gamma-\rho} \rceil, \frac{1}{2}\right)$. Thus, again by Lemma 4.5, for some constant $c' > 0$,

$$\mathbb{P}\Big(\text{There exist at least } \frac{1}{4} N^{1-\gamma-\rho} \text{ good vertices}\Big|\mathcal{B}\Big) \geq 1 - e^{-c' N^{1-\gamma-\rho}}.$$

Hence, putting things together yields that the probability that there exist at least $\frac{1}{4} N^{1-\gamma-\rho}$ good marked vertices is at least

$$\left(1 - e^{-N^{1-\gamma+o(1)}}\right)\left(1 - e^{-c' N^{1-\gamma-\rho}}\right) \geq 1 - e^{-N^{1-\gamma-\rho+o(1)}}.$$

Since one can absorb the factor $\frac{1}{4}$ into the exponent of N by adjusting the parameter $\rho > 0$, we are done. □

Now we focus on the unlikely appearance of vertices of degree close to $\gamma \frac{\log N}{\log \log N}$ for $\gamma > 1$. By (3.3), the probability of the existence of such vertex is $N^{1-\gamma+o(1)}$. In the next proposition, we improve this statement by further requiring the absence of edges between the neighbors of such vertex.

Proposition 3.3. For $\gamma > 1$, let $\mathcal{A}'_\gamma$ be the event that there exists a vertex v and a subset $W \subseteq V$ of size $g(\gamma) = \left\lceil \gamma \frac{\log N}{\log \log N} \right\rceil$ satisfying the following properties:

1. The vertex v is connected to all elements in W.
2. There are no edges within W.

Then,

$$\mathbb{P}\left(\mathcal{A}'_\gamma\right) = N^{1-\gamma+o(1)}.$$

Proof. Since the upper bound immediately follows from (3.3) with $r = 1$ and $\delta_1 = \gamma - 1$, we only prove the lower bound. We again partition the vertices into the subsets $S := \left\{s_1, \ldots, s_{\lceil \frac{N}{2} \rceil}\right\}$ and $L := S^c = \left\{\ell_1, \ldots, \ell_{\lfloor \frac{N}{2} \rfloor}\right\}$. Then, by Lemma 3.2, for each $s_k \in S$,

$$\mathbb{P}\left(d_L(s_k) \geq g(\gamma)\right) = N^{-\gamma+o(1)}.$$

Let $\mathcal{B}$ be the event that there exists a vertex $s_k \in S$ such that $d_L(s_k) \geq g(\gamma)$. Since $|S| = \lceil \frac{N}{2} \rceil$,

$$\mathbb{P}(\mathcal{B}) \geq 1 - (1 - N^{-\gamma-o(1)})^{\lceil \frac{N}{2} \rceil} \geq 1 - e^{-N^{1-\gamma+o(1)}} \geq N^{1-\gamma+o(1)}, \tag{3.8}$$

where we used the fact that $1 - e^{-x} \geq \frac{x}{2}$ for small $x > 0$. Given the event $\mathcal{B}$, let us take any subset $W \subseteq L$ of size $g(\gamma)$ consisting of neighbors of s_k. Then, as in (3.7) in the previous proof, using the independence between different edges,

$$\mathbb{P}\left(\text{There are no edges within } W | \mathcal{B}\right) \geq \frac{1}{2}. \tag{3.9}$$

Therefore, the statement follows by multiplying (3.8) and (3.9). □

The next result concerns the degree profile of high-degree vertices of G. To this end we define, for $\gamma \geq 0$,

$$D_\gamma := \left\{v \in V : d(v) \geq g(\gamma)\right\}. \tag{3.10}$$

By (3.5), if $0 \leq \gamma < 1$, then the expected number of elements in D_γ is of order $N^{1-\gamma}$. In Proposition 3.2, we established that with high probability, $|D_\gamma| \geq N^{1-\gamma-\rho}$ for any $\rho > 0$, which corresponds to a bound on the lower tail of $|D_\gamma|$ for $0 \leq \gamma < 1$.

Now, we establish (2) mentioned in the beginning of Section 3.1. To accomplish this, we start by estimating the moments of $|D_\gamma|$.

Lemma 3.3. For any $0 \leq \gamma < 1$, $\varepsilon > 0$ and a positive integer j, for sufficiently large n,

$$\mathbb{E}\left[\left|D_\gamma\right|^j\right] \leq N^{j-j\gamma+j\varepsilon}. \tag{3.11}$$

Proof. First note that the bound is trivial for $\gamma = 0$, since $|D_\gamma|^j \leq |V|^j \leq N^j$ for any integer j. As before, let us arbitrarily label all vertices and partition the set of vertices $V = \{v_1, \cdots, v_M\}$ into two subsets $S := \{v_1, \ldots, v_j\}$ and $L := V \backslash S$. This definition makes $d_L(v_1), \cdots, d_L(v_j)$ independent (recall that $d_L(v)$ denotes the number of vertices in L connected to v), which we will crucially use in the following moment calculations. First note that

$$\mathbb{E}\left[\left|D_\gamma\right|^j\right] = \mathbb{E}\left[\left(\sum_{i=1}^{N} \mathbf{1}\left(d(v_i) \geq g(\gamma)\right)\right)^j\right]$$

$$\leq \sum_{k=1}^{j} \binom{N}{k} k^j \mathbb{P}\left(d(v_1), \ldots, d(v_k) \geq g(\gamma)\right). \tag{3.12}$$

The above can be seen by expanding the sum and noting that each term is a product of j indicators. Given $v_1, \ldots, v_k$ with $k \leq j$, the number of summands leading to the indicator involving exactly these vertices is at most k^j, since each of the j factors can be one of $v_1, \ldots, v_k$. Further, exchangeability of the vertex degrees leads to the $\binom{N}{k}$ factors, yielding the inequality.

We now proceed to estimate the joint probabilities of interest. Since each vertex $v_i \in S$ can have at most $j-1$ edges into S, for any $1 \leq k \leq j$, for large enough n,

$$\begin{aligned}\mathbb{P}\left(d(v_1), \ldots, d(v_k) \geq g(\gamma)\right) &\leq \mathbb{P}\left(d_L(v_1), \ldots, d_L(v_k) \geq g(\gamma) - j + 1\right) \\ &\leq \mathbb{P}\Big(d_L(v_1), \ldots, d_L(v_k) \geq g\Big(\gamma - \frac{\varepsilon}{2}\Big)\Big) \leq \left(\frac{1}{N^{\gamma - \frac{\varepsilon}{2} + o(1)}}\right)^k,\end{aligned}$$

where we used the independence of $d_L(v_1), \cdots, d_L(v_k)$ and a tail probability estimate for the vertex degree (Lemma 3.2) in the last inequality. Applying this estimate to each term in (3.12), for all $1 \leq k \leq j$,

$$\binom{N}{k} k^j \mathbb{P}\left(d(v_1), \ldots, d(v_k) \geq g(\gamma)\right) \leq \frac{1}{k!} k^j N^{k\left(1-\gamma+\frac{\varepsilon}{2}\right)+o(1)} \leq N^{j(1-\gamma+\frac{\varepsilon}{2})+o(1)},$$

since $1-\gamma+\frac{\varepsilon}{2} > 0$. Therefore, (3.12) is bounded by $j \cdot n^{j(1-\gamma+\frac{\varepsilon}{2})+o(1)} \leq N^{j(1-\gamma+\varepsilon)}$ for large N, which concludes the proof. □

The moment bound yields the following.

Proposition 3.4. For any $0 < \kappa < 1$, let m be an integer such that $m\kappa < 1 \leq (m+1)\kappa$. Then, for any $\mu > 0$ and sufficiently large n,

$$\mathbb{P}\Big(|D_{i\kappa}| \leq N^{1-i\kappa+\kappa} \text{ for all } i = 0, 1, \cdots, m\Big) \geq 1 - N^{-\mu\kappa}. \tag{3.13}$$

Proof. By a union bound and the Markov's inequality combined with Lemma 3.3, for any $\varepsilon > 0$ and a positive integer j,

$$\begin{aligned}\mathbb{P}\left(|D_{i\kappa}| \geq N^{1-i\kappa+\kappa} \text{ for some } i = 0, 1, \cdots, m\right) &\leq \sum_{i=0}^{m} \mathbb{P}\left(|D_{i\kappa}| \geq N^{1-i\kappa+\kappa}\right) \\ &\leq \sum_{i=0}^{m} \frac{\mathbb{E}\left[|D_{i\kappa}|^j\right]}{N^{j-ji\kappa+j\kappa}} \leq (m+1) N^{j\varepsilon - j\kappa}.\end{aligned}$$

Taking $\varepsilon = \frac{\kappa}{2}$, and noticing that $m < \frac{1}{\kappa}$, the above expression is bounded by $(m+1)N^{-j\frac{\kappa}{2}} \leq \left(\frac{1}{\kappa}+1\right)N^{-j\frac{\kappa}{2}}$. Since j is an arbitrary integer, this concludes the proof. □

Note that the previous lemma and proposition could be stated with an asymptotic notation rather than the additional constants ε and μ, but the current phrasing will make it easier to use the results in our main proofs.

3.2 Degree profile near the maximum degree

For our precise description of the largest eigenvalues and eigenvectors of the unweighted graph G we follow [ADK23b] by analyzing the spectral contribution of high degree vertices, which we separate into three regimes. After defining these regimes we analyze their sizes. These results are different from the previous ones in that they contain estimates of the number of vertices whose degree is very close to the largest degree in the graph, i.e. of order $\mathfrak{u}$, with a small additive correction term. They all appear in the paper [HM23].

Once more, we first need to define some notation. For a vertex $x \in [N]$ and $i \geq 0$, we denote by $B_i(x)$ the ball of radius i around x, rooted at x. Moreover, we define $S_i(x)$, the sphere of radius i around x, to be the set of vertices y such that the shortest path from x to y is of length i. In other words $S_i(x)$ are all vertices that are not in $B_{i-1}(x)$, and that are connected to x by a path of length i.

Given a root vertex x, we define a partial ordering on vertices by writing $u \leq v$ if there is a shortest path from x to v that goes through u. We also write for $y \in [N]$,

$$N_y := \left| \left\{ u \in [N] : u \geq y, u \sim y \right\} \right| \tag{3.14}$$

as the number of children of y in the rooted graph. Similarly, for a rooted or unrooted graph, for a vertex $y \in [N]$ we define $\Gamma_y = \{z \in [N] : z \sim y\}$ to be the neighborhood of y.

The following parameters will be used to approximate the largest eigenvalues.

Definition 3.1. We denote by

1. $\alpha_x := |\Gamma_x|$, the degree of the vertex x
2. $\beta_x := \sum_{y \sim x} N_y$, the number of vertices in $S_2(x)$, which we also call the size of the 2-neighborhood,
3. $\beta_x^{(1,1)} = \sum_{y_2 \in S_2(x)} N_{y_2}$, the number of vertices in $S_3(x)$, and
4. $\beta_x^{(2)} = \sum_{y \sim x} N_y^2$.

As we will see, the largest eigenvalues are mostly determined by these four statistics. In our regime, the last two statistics (as well as all others) are well concentrated enough that we can write an accurate enough formula for the eigenvalues based on only α_x, β_x, for the vertices x with the largest degrees.

We define the following sets of high degree vertices in our graph. For $m \geq 0$, let

$$\mathcal{X}_m := \left\{ x \in [N] : \alpha_x \geq \mathfrak{u} - m \right\}.$$

Using this notation we can define the regimes into which we split up the high-degree vertices:

Definition 3.1. We differentiate:

1. the *fine regime*: $\mathcal{W} := \mathcal{X}_{\mathfrak{u}^{1/4}}$,
2. the *intermediate regime*: $\mathcal{V} := \mathcal{X}_{\mathfrak{u}^{2/3}}$, and
3. the *rough regime*: $\mathcal{U} := \mathcal{X}_{\mathfrak{u}/2}$.

The precise thresholds are not significant and rather of technical nature, our analysis would continue to work if we replaced these thresholds with $\mathfrak{u} - \mathfrak{u}^{c_1}, \mathfrak{u} - \mathfrak{u}^{c_2}, c_3\mathfrak{u}$ for some constants satisfying $0 < c_1 < 1/2, 1/2 < c_2 < 1$ and $0 < c_3 < 1$.

For large N, we have $\mathcal{W} \subset \mathcal{V} \subset \mathcal{U}$. First we bound the sizes of these sets. The upper bounds will let us perform union bounds, whereas the lower bound on $|\mathcal{X}_m|$ tells us that all of our highest degree vertices have almost the same degree.

Lemma 3.4. Let $\log^{-1/15} N \leq d \leq \log^{1/40} N$. For $m \geq 0$, with probability $1 - O\left(\left(\frac{\mathfrak{u}}{d}\right)^{-(m+1/2)}\right)$,

$$|\mathcal{X}_m| \leq \frac{3}{2}\left(\frac{\mathfrak{u}}{d}\right)^{m+1/2}. \tag{3.15}$$

Moreover, for $1 \leq m \leq \mathfrak{u}^c$ with $c < 1/2$, with probability $1 - O\left(\left(\frac{\mathfrak{u}}{d}\right)^{-(m-1/2)}\right)$,

$$|\mathcal{X}_m| \geq \frac{1}{2}\left(\frac{\mathfrak{u}}{d}\right)^{m-1/2}. \tag{3.16}$$

Proof. By Lemma 3.11 in [Bol01], we know that $\mathrm{Var}(|\mathcal{X}_m|) = O(\mathbb{E}[|\mathcal{X}_m|])$. By Chebyshev's inequality, we have that with probability $1 - O\left(\frac{1}{\mathbb{E}[|\mathcal{X}_m|]}\right)$,

$$\frac{1}{2}\mathbb{E}[|\mathcal{X}_m|] \leq |\mathcal{X}_m| \leq \frac{3}{2}\mathbb{E}[|\mathcal{X}_m|]. \tag{3.17}$$

Therefore, it is sufficient to show that $\mathbb{E}[|\mathcal{X}_m|]$ and $\mathbb{E}[|\mathcal{U}|]$ satisfy the above bounds, and that each is $\omega_N(1)$. Recall that $\mu_k := N\mathbb{P}(\text{ a vertex is of degree } k\text{ }) = N\mathrm{Bin}(k; N-1, p)$. Thus for any $k \in \mathbb{N}$,

$$\frac{\mu_{k+1}}{\mu_k} = \frac{\mathrm{Bin}(k+1; N-1, p)}{\mathrm{Bin}(k; N-1, p)} = \frac{(N-k-1)p}{(k+1)(1-p)}. \tag{3.18}$$

By the definition of $\mathfrak{u}$, and the fact that μ_k monotonically decreases in k,

$$\begin{aligned} \mu_{\mathfrak{u}}^{-1} \leq \mu_{\mathfrak{u}-1} &= \frac{N\mathfrak{u}\left(1 - \frac{d}{N}\right)}{d(N-\mathfrak{u})}\mu_{\mathfrak{u}}, \\ \mu_{\mathfrak{u}} \leq \mu_{\mathfrak{u}+1}^{-1} &= \frac{N(\mathfrak{u}+1)\left(1 - \frac{d}{N}\right)}{d(N-\mathfrak{u}-1)}\mu_{\mathfrak{u}}^{-1} \end{aligned}$$

Therefore

$$\Big(1 - o_N(1)\Big)\sqrt{\frac{d}{\mathfrak{u}}} \le \mu_{\mathfrak{u}} \le \Big(1 + o_N(1)\Big)\sqrt{\frac{\mathfrak{u}}{d}}.$$

We have by (3.18), for $m \le \mathfrak{u}$,

$$\begin{aligned}\mu_{\mathfrak{u}-m} &= \mu_{\mathfrak{u}}\left(\frac{N-d}{d}\right)^m \prod_{i=1}^{m} \frac{\mathfrak{u}-i+1}{N-\mathfrak{u}+i-1} \\ &= \Big(1 + o_N(1)\Big)\mu_{\mathfrak{u}}\frac{\mathfrak{u}^m}{d^m}\prod_{i=1}^{m}(1 - \frac{i-1}{\mathfrak{u}}).\end{aligned}$$

An upper bound on this is $(1 + o_N(1))(\frac{\mathfrak{u}}{d})^{m+1/2}$. Summing over all $0 \le n \le m$ gives (3.15). Assuming that $m \le \mathfrak{u}^c$ for $c < 1/2$,

$$\Big(1 + o_N(1)\Big)\mu_{\mathfrak{u}}\frac{\mathfrak{u}^m}{d^m}\prod_{i=1}^{m}\left(1 - \frac{i-1}{\mathfrak{u}}\right) \ge \Big(1 + o_N(1)\Big)\mu_{\mathfrak{u}}\frac{\mathfrak{u}^m}{d^m}e^{-m^2/\mathfrak{u}} \ge \Big(1 + o_N(1)\Big)\mu_{\mathfrak{u}}\frac{\mathfrak{u}^m}{d^m},$$

giving (3.16).

□

Corollary 3.1. For our regimes this implies that with high probability

$$\begin{aligned}|\mathcal{W}| &\ll e^{u^c} \text{ for any } c > \frac{1}{4} \\ |\mathcal{V}| &\ll e^{u^c} \text{ for any } c > \frac{2}{3} \\ |\mathcal{U}| &\ll N^c \text{ for any } c > \frac{1}{2}.\end{aligned}$$

Remark 3.1. Note that in the proof we derive bounds on the expected values of the sizes of these sets, which immediately imply bounds on the probability that a given vertex falls into one of the sets, since for any set $\mathcal{T}$, $\mathbb{E}[|\mathcal{T}|] = N\mathbb{P}(\text{ vertex } 1 \in \mathcal{T})$.

3.3 Local neighborhoods

When analyzing the largest eigenvalues and eigenvectors of G we will need a precise understanding of the local neighborhoods around the high degree vertices. In this section we assume that $\log N^{-1/15} \le d \le \log N^{1/40}$. All results in this section are from the paper [HM23].

Structure around the vertices in the fine and intermediate regime

We first define an event under which the balls around vertices in $\mathcal{V}$ have a nice structure. The approximate eigenvalue and eigenvalues will be defined using $B_r(x)$ for some fixed r. The following structural results are for slightly larger balls so that we can also bound the error coming from truncating the balls and guarantee that there is no intersection with balls of radius 3 around vertices from $\mathcal{U}$.

Definition 3.2. Define $\Omega_{3.2}$ to be the event that the following are true.

1. For all $x \neq y \in \mathcal{V}$, $B_{r+3}(x) \cap B_{r+3}(y) = \emptyset$.
2. For all $x \in \mathcal{V}$, $B_{r+3}(x)$ is a tree.
3. For $1 \leq i \leq r$ and every vertex $x \in \mathcal{V}$,
$$\left||S_i(x)| - d^{i-1}\alpha_x\right| = O\left(d^{i-3/2} + 1\right) \mathfrak{u}^{\frac{7}{8}}$$
Moreover, for every vertex $x \in \mathcal{W}$,
$$\left||S_i(x)| - d^{i-1}\alpha_x\right| \leq O\left(d^{i-3/2} + 1\right) \mathfrak{u}^{\frac{2}{3}}.$$
4. For $x \in \mathcal{V}$, every $y \in B_{r+3}(x)\backslash\{x\}$, satisfies $N_y \leq \mathfrak{u}^{3/4}$.
Moreover, for $x \in \mathcal{W}$, every $y \in B_{r+3}(x)\backslash\{x\}$ satisfies $N_y \leq \mathfrak{u}^{1/3}$.
5. For every $x \in \mathcal{V}$,
$$\left|\sum_{y \in S_1(x)} N_y^2 - \left(d^2 + d\right)\alpha_x\right| \leq O\left(\mathfrak{u}^{3/2}\right).$$
Moreover, for every $x \in \mathcal{W}$,
$$\left|\sum_{y \in S_1(x)} N_y^2 - \left(d^2 + d\right)\alpha_x\right| \leq O\left(\mathfrak{u}^{2/3}\right).$$

Note that statement 3 in 3.2 implies that $|S_2(x)| \leq 2d\alpha_x$ for our regime of d. If d was smaller than $\log^{-c} N$, for some $c > \frac{1}{4}$, this would no longer be true.

Because of the sparsity of the graph and concentration of independent binomials, we can show that the event $\Omega_{3.2}$ almost always occurs. Considering statements similar to most of these bounds have appeared previously, in, e.g. [ADK23b], we defer the proof of the following lemma to the appendix in Section A.1.

Lemma 3.5. The event $\Omega_{3.2}$ occurs with high probability.

Structure around the vertices in the rough regime

For the vertices in the rough regime we use a different approach to construct approximate eigenvectors, because the growth of the spheres and the maximum degree in the balls around them cannot be bounded as tightly as for vertices in $\mathcal{V}$. We now state two weaker structural lemmas around vertices in $\mathcal{U}$. Firstly we have weaker bounds on the neighborhood growth and the fluctuations of the degrees of the neighbors. The following lemma has a similar proof as Lemma 3.5, and we defer the proof to Section A.1 of the appendix.

Lemma 3.6. We have that with high probability for any vertex $x \in \mathcal{U}$ simultaneously and any $i \leq 3$, the sphere $S_i(x)$ at distance i from x satisfies

$$|S_i(x)| = O\left((d + \log\log N)^{i-1}\mathfrak{u}\right).$$

Moreover,

$$\sum_{y\in S_1(x)} \left(N_y - d\right)^2 \leq O\left((\log N)^2\right).$$

Next, we show the balls around vertices in $\mathcal{U}$ are close to disjoint trees: with high probability the neighborhoods around vertices in the rough regime are almost trees and contain few disjoint paths that contain other vertices from the rough regime. This result basically corresponds to Lemma 5.5 and Lemma 7.3 in [ADK21b], albeit for a different regime of d. The proof is also very similar and is deferred to Section A.1 of the appendix.

Lemma 3.7. Let $\mathcal{U}_\eta = \{x \in [N] : \alpha_x \geq \eta\mathfrak{u}\}$, and s be some positive integer, then with high probability for some constants C_1 and C_2 that only depend on η, simultaneously for all $x \in \mathcal{U}_\eta$,

1. $|E(B_s(x))| < |V(B_s(x))| - 1 + C_1$ and
2. $B_s(x)$ contains less than C_2 edge disjoint paths in $B_s(x)$ containing other vertices from $\mathcal{U}_\eta$.

Note that $\mathcal{U} = \mathcal{U}_{\frac{1}{2}}$ and that it is enough to take constants $C_1 = 2$ and $C_2 > \frac{2}{\eta}$.

We now construct a "pruned" graph in which the neighborhoods of vertices in $\mathcal{U}$ are disjoint trees. The construction works in the same manner as in Lemma 7.2 of [ADK21b] and we use it to prove a statement similarly to Proposition 6.19 in [ADK23b]. Once more the proof can be found in the appendix.

Lemma 3.8. Recall that we denote by G the random graph sampled from $\mathcal{G}\left(N, \frac{d}{N}\right)$. With high probability, there is a subgraph $\hat{G} \subset G$ such that for all vertices $x \in \mathcal{U}$,

1. Balls of radius 3 around x in $\hat{G}$, which we denote by $\hat{B}_3(x)$, are disjoint;

2. The subgraphs induced by $\hat{B}_3(x)$ are trees;
3. The maximum degree of $G - \hat{G}$ is bounded;
4. For $i \le 3$, the spheres $\hat{S}_i(x)$ in the pruned graph $\hat{G}$ satisfy
$$\left|\hat{S}_i(x)\right| = O\left((d + \log\log N)^{i-1}\mathfrak{u}\right);$$
5. We have that
$$\sum_{y \in \hat{S}_1(x)} \left(\hat{N}_y(x) - \frac{\hat{\beta}}{\hat{\alpha}}\right)^2 \le O\left((\log N)^2\right).$$

3.4 Probability of the existence of a clique

As will be explained in the idea of proof section, atypically high degree stars (with high edge-weights on them) are the driving mechanism behind the largest eigenvalue in the case of light-tailed weights. In the case of heavy-tailed weights on the other hand, a similar role is played by cliques. In this short section, we state a bound from [GHN24] for the probability that $G \sim \mathcal{G}_{N,\frac{d}{N}}$, with d constant, contains a clique of size k. While the lower bound can be seen for instance in [GN22, Lemma 4.3], the upper bound follows from the first moment method: there are $\binom{M}{k} \le N^k$ possible cliques of size k, and each of them requires the existence of $\binom{k}{2}$ edges, with each has a probability d/N of appearing.

Lemma 3.9. For any integer $k \ge 3$, there exists a constant $C = C(k, d) > 0$ such that
$$CN^{-\binom{k}{2}+k} \le \mathbb{P}(G \text{ contains a clique of size } k) \le d^{\binom{k}{2}} N^{-\binom{k}{2}+k}.$$

3.5 Connectivity structures of highly sub-critical Erdős-Rényi graph

As a key step in our proofs about weighteds graphs, we will decompose the underlying graph X into graphs with low and high edge-weights respectively. Because of the threshold we choose, the latter graph turns out to be a highly subcritical graph $\mathcal{G}_{N,q}$ with
$$q \le \frac{d'}{N(\log N)^{\varepsilon}} \tag{3.19}$$
for some constants $\varepsilon, d' > 0$. In this section, we record some properties of such graphs, namely that all connected components look like trees and that their sizes are well-controlled.

Throughout this section, we assume that the edge density q satisfies (3.19). First, we have a bound on the largest degree denoted by $d_1(\mathcal{G}_{N,q})$, as a direct consequence of (3.3).

Lemma 3.10. For $\delta_1 > 0$, define the event

$$\mathcal{D}_{\delta_1} := \left\{ d_1(\mathcal{G}_{N,q}) \le (1+\delta_1) \frac{\log N}{\log \log N} \right\}. \tag{3.20}$$

Then,

$$\liminf_{N\to\infty} \frac{-\log \mathbb{P}\left(\mathcal{D}^c_{\delta_1}\right)}{\log N} \ge \delta_1.$$

Proof. Note that $\mathcal{D}_{\delta_1}$ is a decreasing event and that $\mathcal{G}_{N,q}$ is stochastically dominated by $\mathcal{G}_{N,\frac{d}{N}}$. Hence, by (3.3), we obtain the result. □

Next, we state a quantitative bound on the size of the largest connected component. The next two results can be obtained from the cited lemmas by noting that the sparsity considered in [GN22] is of the form $\frac{d'}{N(\log N)^{\varepsilon/2}}$, while it is $\frac{d'}{N(\log N)^{\varepsilon}}$ in our case.

Lemma 3.11 ([GN22, Lemma 5.4])**.** Let $C_1, \cdots, C_L$ denote the connected components of $\mathcal{G}_{N,q}$. For $\delta_2 > 0$, define the event

$$\mathcal{C}_{\varepsilon,\delta_2} := \left\{ \max_{i=1,\dots,L} |C_i| \le (1+\delta_2) \frac{1}{\varepsilon} \frac{\log N}{\log \log N} \right\}.$$

Then,

$$\liminf_{N\to\infty} \frac{-\log \mathbb{P}\left(\mathcal{C}^c_{\varepsilon,\delta_2}\right)}{\log N} \ge \delta_2. \tag{3.21}$$

Note that in this lemma, we replaced δ_2 in [GN22, Lemma 5.4] by $2\delta_2$.

To conclude this section we state two results about the structure of the connected components. The first one quantifies how similar all connected components are to trees, in the sense that they have a small number of tree-excess edges. The second one concerns the event that all connected components are trees.

Lemma 3.12 ([GN22, Lemma 5.6])**.** Let $C_1, \cdots, C_L$ denote the connected components of $\mathcal{G}_{N,q}$. For $\delta_3 > 0$, define the event

$$\mathcal{E}_{\delta_3} := \left\{ \max_{i=1,\dots,L} \left\{ |E(C_i)| - |V(C_i)| \right\} \le \delta_3 \right\}. \tag{3.22}$$

Then,

$$\liminf_{N\to\infty} \frac{-\log \mathbb{P}\left(\mathcal{E}^c_{\delta_3}\right)}{\log N} \ge \delta_3. \tag{3.23}$$

In addition, define the event

$$\mathcal{T} := \Big\{ |E(C_i)| = |V(C_i)| - 1, \ \forall i = 1, \cdots, L \Big\}.$$

In other words, $\mathcal{T}$ is the event that all the connected components of $\mathcal{G}_{N,q}$ are trees. Then, there is a constant $C > 0$ that depends on d, such that

$$\mathbb{P}\Big(\mathcal{T}^c\Big) \leq \frac{C}{(\log N)^{2\varepsilon}}. \tag{3.24}$$

Chapter 4

Distributional properties of sparse Erdős-Rényi graphs

Many statistics of Erdős-Rényi graphs follow the binomial distribution. In particular the degree of a single vertex is Binom $(N-1, d/N)$, and when d is small and N large, this can be well approximated by a Pois(d) distribution. Moreover any two degrees only depend on each other very weakly, so we expect that in some ways, the whole collection of degrees should behave similarly to a collection of independent Poisson random variables of a similar size. The third section contains a precise statement of that form with regards to the maximal degrees and their neighborhoods. We start by collecting a few results regarding the Binomial and the Poisson distribution in the first section, as well as comparisons between them in the second. Those will be used repeatedly in this and other chapters. We end this chapter with a section that contains some results about the tails of sums of independent random variables, which show up for different reasons in the two projects about Erdős-Rényi graphs. The proof of the approximations in the first two and the last sections are given in the appendix in Section A.2.

The new results in this section are aggregated from [GHN24] and [HM23].

4.1 Distributional comparisons

Lemma 4.1. If $X \sim \text{Binom}(n, p)$ and $Y \sim \text{Pois}(np)$, and if $k, np \leq \sqrt{n}$, then

$$\mathbb{P}(X = k) = \left(1 + O\left(\frac{k^2 + (np)^2 + 1}{n}\right)\right)\mathbb{P}(Y = k).$$

This implies that the tails are also the same up to a small error.

Corollary 4.1. If $X \sim \text{Binom}(n, p)$ and $Y \sim \text{Pois}(np)$, and if $k \leq \sqrt{n}$, $np \leq n^{1/2-c}$ for some fixed constant c, then

$$\mathbb{P}(X \geq k) = \left(1 + O\left(\frac{k^2 + (np)^2 + 1}{n}\right)\right)\mathbb{P}(Y \geq k) + O\left(\left(ep\sqrt{n}\right)^{\sqrt{n}}\right).$$

4.2 Distributional identities

Poisson tails

Having established these comparisons, we use very tight bounds on the Poisson tail. Tao gives such a tight bound on his blog [Tao22], where he notes that forms of this bound are given previously [Gly87, Tal95]. In the post, Tao gives the proof of the upper bound and leaves the proof of the lower bound to the reader. We prove both sides in the appendix in Section A.2.

Lemma 4.2. For $X \sim \text{Pois}(\lambda)$ and $\delta \geq \frac{1}{\sqrt{\lambda}}$, for sufficiently large λ

$$\mathbb{P}\Big(X \geq \lambda(1+\delta)\Big) \leq \frac{e^{-\lambda h(\delta)}}{\sqrt{\lambda \min\{\delta, \delta^2\}}},$$

where $h(\delta) = (\delta + 1)\log(\delta + 1) - \delta$.

Moreover, if $\lambda(1 + \delta)$ is an integer, then there is a universal constant $c_{4.2}$ such that for sufficiently large λ

$$\mathbb{P}\Big(X \geq \lambda(1+\delta)\Big) \geq c_{4.2}\frac{e^{-\lambda h(\delta)}}{\sqrt{\lambda \min\{\delta, \delta^2\}}}.$$

The integrality assumption is necessary as for very large δ, the difference in probability between $\mathbb{P}(X \geq \lambda(1 + \delta))$ and $\mathbb{P}(X \geq \lambda(1 + \delta) + 1)$ is large enough that for very small $c > 0$, this two sided bound could not possibly hold for $\lceil \lambda(1+\delta) \rceil$ and $\lfloor \lambda(1+\delta) \rfloor + c$ simultaneously. For its use in our proofs, the integrality assumption is irrelevant, as we will only need the lower bound in the small δ regime.

Corollary 4.2. For $X \sim \textit{Pois}(\lambda)$, if $\lambda\delta^3 = o_\lambda(1)$ and $\delta \geq \frac{1}{\sqrt{\lambda}}$, then

$$\Big(1 - o_\lambda(1)\Big)c_{4.2}\frac{e^{-\frac{\lambda\delta^2}{2}}}{\delta\sqrt{\lambda}} \leq \mathbb{P}\Big(X \geq \lambda(1+\delta)\Big) \leq \Big(1 + o_\lambda(1)\Big)\frac{e^{-\frac{\lambda\delta^2}{2}}}{\delta\sqrt{\lambda}}. \tag{4.1}$$

Binomial tails

Although we will mostly approximate the degrees by Poisson random variables it is sometimes more convenient to work with the precise distribution. To this end we state here some classical tail bounds we will use, that rely on the relative entropy.

We denote by I_p the relative entropy functional

$$I_p(q) := q \log \frac{q}{p} + (1-q) \log \frac{1-q}{1-p}. \tag{4.2}$$

Lemma 4.3. There is a universal constant $c > 0$ such that for any $0 < p < 1$,

$$I_p\left(\frac{p}{2}\right) \geq cp. \tag{4.3}$$

This implies that $I_{1-p}\left(1 - \frac{p}{2}\right) \geq cp$ and for any $0 < q \leq \frac{p}{2}$, $I_p(q) \geq cp$.

Lemma 4.4 (Lemma 3.3 in [LZ17])**.** If $0 < p \ll q$ (i.e. $\frac{q}{p} \to \infty$ as $p, q \to 0$) and $q \leq 1 - p$ then

$$I_p(p+q) = (1 + o(1)) q \log\left(\frac{q}{p}\right).$$

The next lemma provides a sharp bound on the tail probability of the Binomial distribution.

Lemma 4.5 (Lemma 4.7.2 in [Ash65])**.** For $m \in \mathbb{N}$ and $0 < q < 1$, let X be a random variable that has the distribution $\text{Binom}(m, q)$. Then, for any $q < \theta < 1$,

$$\frac{1}{\sqrt{8m\theta(1-\theta)}} e^{-mI_q(\theta)} \leq \mathbb{P}\left(X \geq \theta m\right) \leq e^{-mI_q(\theta)}. \tag{4.4}$$

By applying the above to the random variable $m - X$ and using that $I_q(\theta) = I_{1-q}(1-\theta)$ we additionally have that, for $0 < \theta < q < 1$,

$$\frac{1}{\sqrt{8m\theta(1-\theta)}} e^{-mI_q(\theta)} \leq \mathbb{P}\left(X \leq \theta m\right) \leq e^{-mI_q(\theta)}. \tag{4.5}$$

As we will be interested in vertices of large degree in our graph, and the degrees follow a binomial distribution, we will repeatedly use tail bounds such as the following.

Lemma 4.6. Let $m = n + o(n)$ and $p = \frac{d}{n}$, and define $X \sim \text{Bin}(m, p)$ then, for $\tau = o(n)$, and $\tau > \frac{md}{n}$, it holds for some constant c, that for n large enough,

$$\mathbb{P}\left(X \geq \tau\right) \leq e^{-\tau \log(\tau) + c\tau \max\{\log d, 1\}}. \tag{4.6}$$

We will use a weaker, simpler version of this bound for low probability events.

Lemma 4.7 ([JRL11] Theorem 2.1)**.** If $X \sim \mathit{Binom}(N, p)$, and $\lambda = Np$, then for $t \geq 0$

$$\mathbb{P}(X - \lambda \geq t) \leq \exp\left(-\frac{t^2}{2\lambda + 2t/3}\right), \quad \mathbb{P}(X - \lambda \leq -t) \leq \exp\left(-\frac{t^2}{2\lambda}\right).$$

4.3 Poisson approximation in sparse Erdős-Rényi graphs

In our analysis of the spectral edge of unweighted Erdős-Rényi graphs we will crucially use that we can approximate high-degree vertices and their local neighborhoods using Poisson random variables. The first result is similar to [ADK23b] Lemma 7.1, but proven in a somewhat different way. It shows that if we consider a set of fixed vertices that is not too big, their degrees and the size of their 2-neighborhoods are close to the distribution of Poisson random variables.

Lemma 4.8. For $\log^{-1/9} N \leq d \leq \log^{1/40} N$, let G be a graph generated from the Erdős-Rényi graph distribution $\mathcal{G}\left(N, \frac{d}{N}\right)$. Moreover, for $k \leq e^{log^{2/3} N}$, consider vertices $z_1, \dots z_k \in [N]$, along with $\frac{1}{2}\mathfrak{u} \leq v_1, \dots, v_k \leq 2\mathfrak{u}$ and $w_1, \dots w_k$ such that $1 \leq w_i \leq dv_i + \mathfrak{u}^{7/8}$ for $1 \leq i \leq k$. Then define i.i.d. $X_1, X_2, \dots X_k \sim Pois(d)$ and independent $Y_{v_1} \sim Pois(dv_1), Y_{v_2} \sim Pois(dv_2), \dots, Y_{v_k} \sim Pois(dv_k)$. If A is the event that there are no intersections between the balls of radius 1 around the vertices $z_1, \dots, z_k$, and no edges from $S_1(z_i)$ to $S_1(z_j)$ for any i, j (including i=j), then

$$\mathbb{P}\left(\bigcap_{i=1}^{k}\{\alpha_{z_i} = v_i, \beta_{z_i} = w_i\} \cap A\right) = \left(1 + N^{-1+o_N(1)}\right)\prod_{i=1}^{k}\mathbb{P}(X_i = v_i)\mathbb{P}(Y_{v_i} = w_i) \tag{4.7}$$

Proof. Let $Z = \{z_1, \dots, z_k\}$, with A_Z being the adjacency matrix of Z. Recall that Γ_{z_i} denotes the neighbors of a vertex z_i.

We analyse the event that $\cap_i\{\alpha_{z_i} = v_i\}$ and that there are no edges between any z_i, as well as no intersection between the neighborhoods of the z_i, sequentially.

That $A_Z = 0$ happens with probability $(1 - d/N)^{\binom{k}{2}}$. Then we first need to choose exactly v_1 vertices among $[N]\backslash Z$ connected to z_1. Subsequently we need to choose exactly v_2 vertices among $[N]\backslash(Z \cup \Gamma_{z_1})$ and make sure that there are no edges between z_2 and Γ_{z_1}, and so on. Note that this way the edges we consider at each step are independent of the previously considered events and moreover the number of edges between z_i and $[N]\backslash(Z \cup \cup_{j=1}^{i-1}\Gamma_j)$ is binomially distributed with parameters $N - k - \sum_{j=1}^{i-1} v_i$ and d/N. This gives

$$\mathbb{P}\left(\bigcap_{i=1}^{k}\left\{\alpha_{z_i} = v_i\right\}\bigcap\left\{A_Z = 0\right\}\bigcap\left\{\cap_{i=1}^{k}\Gamma_{z_i} = \varnothing\right\}\right)$$
$$= \left[\prod_{i=1}^{k}\mathbb{P}\left(\text{Binom}\left(N - k - \sum_{j=1}^{i-1} v_j, \frac{d}{N}\right) = v_i\right)\left(1 - \frac{d}{N}\right)^{\sum_{j=1}^{i-1} v_i}\right]\left(1 - \frac{d}{N}\right)^{\binom{k}{2}}.$$

We now use Lemma 4.1 to approximate the binomial probabilities, the bound on the error terms follows from our assumptions on v_i, k and the bounds on d from 5.1.

$$\mathbb{P}\left(\text{Binom}\left(N-k-\sum_{j=1}^{i-1} v_j, \frac{d}{N}\right)=v_i\right)=\left(1+\tilde{O}\left(N^{-1}\right)\right)\mathbb{P}\left(\text{Pois}\left(d-\frac{d\left(k+\sum_{j=1}^{i-1} v_j\right)}{N}\right)=v_i\right)$$
$$=\left(1+N^{-1+o_N(1)}\right)\mathbb{P}\left(\text{Pois}(d)=v_i\right).$$

Moreover $\left(1-\frac{d}{N}\right)^{\sum_{j=1}^{i-1} v_i}$ as well as $\left(1-\frac{d}{N}\right)^{\binom{k}{2}}$ can also be written as $1+N^{-1+o_N(1)}$. There are $N^{o_N(1)}$ such error terms, which all together implies that

$$\mathbb{P}\left(\bigcap_{i=1}^{k}\left\{\alpha_{z_i}=v_i\right\}\bigcap\left\{A_Z=0\right\}\bigcap\left\{\bigcap_{i=1}^{k}\Gamma_{z_i}=\varnothing\right\}\right)=\left(1+N^{-1+o_N(1)}\right)\mathbb{P}\Big(\text{Pois}(d)=v_i\Big). \qquad (4.8)$$

We now condition on the event $\bigcap_{i=1}^{k}\left\{\alpha_{z_i}=v_i\right\}\bigcap\left\{A_Z=0\right\}\bigcap\left\{\bigcap_{i=1}^{k}\Gamma_{z_i}=\varnothing\right\}$ and similarly analyse $\{\beta_{z_i}=w_i\}$. Note that the event A does not require that the $S_2(z_i)$ are all disjoint which makes the analysis slightly simpler. The number of edges from $S_1(z_i)$ to $[N]\backslash(Z\cup\Gamma_Z)$ is Binomial with parameters $v_i\left(N-k-\sum_{j=1}^{k} v_j\right)$ and $\frac{d}{N}$, and those random variables are independent since we condition on $A_Z=0$ and $\bigcap_{i=1}^{k}\Gamma_{z_i}=\varnothing$. Finally, the probability that there are no edges within and across any Γ_{z_i} is equal to $(1-d/N)^{\sum_{i\neq j} v_i v_j+\sum_i\binom{v_i}{2}}$.

Thus, defining A_Γ to be the adjacency matrix of $\Gamma_Z=\bigcup_{i=1}^{k}\Gamma_{z_i}$, we get

$$\mathbb{P}\left(\bigcap_{i=1}^{k}\left\{\beta_{z_i}=w_i\right\}\bigcap\left\{A_\Gamma=0\right\}\Big|\bigcap_{i=1}^{k}\left\{\alpha_{z_i}=v_i\right\}\bigcap\left\{A_Z=0\right\}\bigcap\left\{\bigcap_{i=1}^{k}\Gamma_{z_i}=\varnothing\right\}\right)$$
$$=\left[\prod_{i=1}^{k}\mathbb{P}\left(\text{Binom}\left(v_i\left(N-k-\sum_{j=1}^{k} v_j\right),\frac{d}{N}\right)=w_i\right)\right]\left(1-\frac{d}{N}\right)^{\sum_{i,j=1,i\neq j}^{k} v_i v_j+\sum_{i=1}^{k}\binom{v_i}{2}}$$

The last term can as before be written as $1+N^{-1+o_N(1)}$. When $v_i=0$, we can immediately replace the Binomial random variables by Poisson random variables with parameter 0, since they are both constant 0. For $v_i>0$, we once more use the Poisson approximation from Lemma 4.1, which together with our bounds on v_i, w_i, k and d, gives

$$\mathbb{P}\left(\text{Binom}\left(v_i\left(N-k-\sum_{j=1}^{k} v_j\right),\frac{d}{N}\right)=w_i\right)$$
$$=\left(1+\tilde{O}\left(N^{-1}\right)\right)\mathbb{P}\left(\text{Pois}\left(v_i d\left(1-\frac{k+\sum_{j=1}^{k} v_j}{N}\right)\right)=w_i\right)$$
$$=\left(1+N^{-1+o_N(1)}\right)\mathbb{P}\Big(\text{Pois}\left(v_i d\right)=w_i\Big)$$

Combining this with (4.8), implies the result. □

Our next result basically shows that what is true for a fixed number of vertices is also true for the vertices with the largest degrees. Crucially, the number of vertices we are allowed to consider in Lemma 4.8 is larger than the number of vertices with very high degree we typically have in our graph.

We start by stating a lemma from [ADK23b] that expresses the distribution of a point process in terms of the point probabilities. Define the parameter $(Z_x)_{x\in[N]}$ to be an exchangeable family of random variables in a measurable space $\mathcal{Z}$. The point process Φ is defined to be equal to $\sum_{x\in[N]}\delta_{Z_x}$. We then introduce for $F\subset\mathcal{Z}^k$ and the point process Φ, the correlation measure

$$q_\Phi(F) := N(N-1)\cdots(N-k+1)\mathbb{P}\Big((Z_1,\dots Z_k)\in F\Big).$$

Lemma 4.9 ([ADK23b], Lemma 7.8)**.** For $n,m\in\mathbb{N}$ and disjoint, measurable $I_1,\dots I_n$

$$\mathbb{P}\left(\Phi(I_1)=k_1,\dots,\Phi(I_n)=k_n\right)=\frac{1}{k_1!\cdots k_n!}\sum_{\substack{\ell_1,\dots,\ell_n\\ \sum\ell_i\le m}}\frac{(-1)^{\sum_i\ell_i}}{\ell_1!\cdots\ell_n!}q_\Phi\left(I_1^{k_1+\ell_1}\times\cdots\times I_n^{k_n+\ell_n}\right)$$
$$+O\left(\frac{1}{k_1!\cdots k_n!}\sum_{\substack{\ell_1,\dots,\ell_n\\ \sum\ell_i=m+1}}\frac{(-1)^{\sum_i\ell_i}}{\ell_1!\cdots\ell_n!}q_\Phi\left(I_1^{k_1+\ell_1}\times\cdots\times I_n^{k_n+\ell_n}\right)\right). \quad (4.9)$$

We use this lemma to show that the edge of the process of the pairs (α_x,β_x), is close to the maxima of a set of Poisson random variables.

Proposition 4.1. Let $\Phi := \left\{(\alpha_x,\beta_x) : x\in[N], \alpha_x > \mathfrak{u}-2\log^{1/8}N\right\}$ denote the point process consisting of the largest degrees in G, with their two-neighborhoods. Let $P := \{(X_i,Y_i)\}_{i\in[N]}$ denote a set of i.i.d. random variables with $X_i\sim\text{Pois}(d)$ and $Y_i|X_i = x\sim\text{Pois}(dx)$. Then set $\Psi' := \left\{(X_i,Y_i) : i\in[N], \mathfrak{u}-2\log^{1/8}N\le X_i\le\mathfrak{u}, 0\le Y_i\le uX_i+u^{7/8}\right\}$. It holds that

$$d_{TV}(\Phi,\Psi') = o_N(1).$$

Proof. We begin by restricting Φ to points (x,y) such that $\mathfrak{u}-2\log^{1/8}N\le x\le u$ and $0\le y\le dx+u^{7/8}$, and to the event that all neighborhoods of vertices with such degrees are disjoint and tree-like. More precisely we define $\mathcal{A}(x,y)$ to be the event that for the ordered pair $(x,y)\in\mathbb{N}^2$, $\mathfrak{u}-2\log^{1/8}N\le x\le\mathfrak{u}$ and $0\le y\le dx+\mathfrak{u}^{7/8}$, and let $\mathcal{B}$ be the event that the neighborhoods around all such vertices $v\in[N]$ are disjoint tree-like. Now we set $\Phi' := \left\{(\alpha_x,\beta_x)\mathbf{1}_{\mathcal{A}(\alpha_x,\beta_x)\cap\mathcal{B}} : x\in[N]\right\}$. Then $\Phi'=\Phi$ with high probability by Lemma 3.5. Thus it is enough to show that $d_{TV}(\Phi',\Psi') = o_N(1)$.

We enumerate all possible location of points for these point processes, in other words, all potential ordered pairs (x,y) such that $\mathfrak{u}-2\log^{1/8}N\le x\le\mathfrak{u}$ and $0\le y\le dx+\mathfrak{u}^{7/8}$. Therefore, if $n_x := 2\log^{1/8}N+1$, $n_y = d\mathfrak{u}+\mathfrak{u}^{7/8}+1$, there are at most $n := n_xn_y$ possibilities.

Define $|\Phi'((x,y))|$ to be the number of points the point process Φ' has at (x,y). We consider an event E, which for some set K_E of vectors in $\mathbb{N}^n$, is defined as follows. We write $\mathcal{E}(X)$ here to mean that the event X occurs.

$$E := \bigsqcup_{\mathbf{k}\in K_E} \mathcal{E}\left(\left|\Phi'((x_1,y_1))\right| = k_1, \dots, \left|\Phi'((x_{n_x},y_{n_y}))\right| = k_n\right).$$

By Lemma 4.9,

$$\begin{aligned}
&\Pr\left(\left|\Phi'((x_1,y_1))\right| = k_1, \dots, \left|\Phi'((x_{n_x},y_{n_y}))\right| = k_n\right) \\
&\quad = \frac{1}{k_1!\cdots k_n!} \sum_{\substack{\ell_1,\dots,\ell_n \\ \sum l_i \le e^{\log^{2/3} N}/2-1}} \frac{(-1)^{\sum_i \ell_i}}{\ell_1!\cdots \ell_n!} q_{\Phi'}\left(I_1^{k_1+\ell_1}\times\cdots\times I_n^{k_n+\ell_n}\right) \\
&\qquad + O\left(\frac{1}{k_1!\cdots k_n!} \sum_{\substack{\ell_1,\dots,\ell_n \\ \sum \ell_i = e^{\log^{2/3} N}/2}} q_{\Phi'}\left(I_1^{k_1+\ell_1}\times\cdots\times I_n^{k_n+\ell_n}\right)\right) \qquad (4.10)
\end{aligned}$$

where I_i is the lattice point (x_i,y_i).

We use this threshold for $\sum k_i$, as by Lemma 3.4 and Markov's inequality, with probability $1-e^{-\Omega(\log^{2/3} N)}$, there are at most $e^{\log^{2/3} N}/2-1$ vertices with degree larger than $\mathfrak{u}-2\log^{1/8} N$, which implies that we only need to consider vectors such that $k := \sum_{i=1}^n k_i \le e^{\log^{2/3} N}/2-1$.

By Lemma 4.8,

$$\begin{aligned}
&\sum_{\mathbf{k}\in K_E} \frac{1}{k_1!\cdots k_n!} \sum_{\substack{\ell_1,\dots,\ell_n \\ \sum l_i \le e^{\log^{2/3} N}/2-1}} \frac{(-1)^{\sum_i \ell_i}}{\ell_1!\cdots \ell_n!} q_{\Phi'}\left(I_1^{k_1+\ell_1}\times\cdots\times I_n^{k_n+\ell_n}\right) \\
&= \left(1+N^{-1+o_N(1)}\right) \sum_{\mathbf{k}\in K_E} \frac{1}{k_1!\cdots k_n!} \sum_{\substack{\ell_1,\dots,\ell_n \\ \sum l_i \le e^{\log^{2/3} N}/2-1}} N^{k+\ell} \frac{(-1)^{\sum_i \ell_i}}{\ell_1!\cdots \ell_n!} \\
&\quad \cdot \prod_{i=1}^n \left(\mathbb{P}(Pois(d)=x_i)\mathbb{P}(Pois(dx_i)=y_i)\right)^{k_i+\ell_i}.
\end{aligned}$$

Also,

$$\begin{aligned}
&\frac{1}{k_1!\cdots k_n!} \sum_{\substack{\ell_1,\dots,\ell_n \\ \sum \ell_i = e^{\log^{2/3} N}/2}} q_{\Phi'}\left(I_1^{k_1+\ell_1}\times\cdots\times I_n^{k_n+\ell_n}\right) \\
&\qquad = O\left(\frac{1}{k_1!\cdots k_n!} \sum_{\substack{\ell_1,\dots,\ell_n \\ \sum \ell_i = e^{\log^{2/3} N}/2}} \frac{1}{(\ell/n)!^n} N^{k+\ell} \mathbb{P}\left(\mathrm{Pois}(d) = \mathfrak{u} - 2\log^{1/8} N\right)^{k+\ell}\right).
\end{aligned}$$

where $\ell = \sum_i \ell_i = e^{\log^{2/3} N}/2$. By the definition of the Poisson,

$$\frac{1}{k_1! \cdots k_n!} \sum_{\substack{\ell_1,\dots,\ell_n \\ \sum \ell_i = e^{\log^{2/3} N}/2}} \frac{1}{((\ell/n)!)^n} N^{k+\ell} \mathbb{P}\left(Pois(d) = \mathfrak{u} - 2\log^{1/8} N\right)^{k+\ell} = e^{-\Omega\left(e^{\log^{2/3} N}\right)}.$$

By using Lemma 4.9 once again,

$$\frac{1}{k_1! \cdots k_n!} \sum_{\sum \ell_1 \dots, \ell_n \le e^{\log^{2/3} N}/2 - 1} N^{k+\ell} \frac{(-1)^{\sum_i \ell_i}}{\ell_1! \cdots \ell_n!} \prod_{i=1}^{n} \left(\mathbb{P}(Pois(d) = x_i)\mathbb{P}(Pois(dx_i) = y_i)\right)^{k_i + \ell_i}$$
$$= \left(1 + N^{-1+o_N(1)}\right) \Pr\left(\Psi'(I_1) = k_1 + \ell_1, \dots, \Psi'(I_n) = k_n + \ell_n, \Psi'(I_n) = k_n\right) + e^{-\Omega\left(e^{\log^{2/3} N}\right)}.$$

We now wish to pass from this error to total variation distance. The total number of possibilities of k for $\sum_{i=1}^n k_i \le e^{\log^{2/3} N}/2$ is given by the balls and bins paradigm as $\sum_{k=0}^{e^{\log^{\frac{2}{3}} N}/2} \binom{n+k-1}{k-1} \le e^{\log^3 N}$. Therefore, the error for any event is at most

$$d_{TV}\left((\alpha_v, \beta_v)\mathbf{1}_{\mathcal{A}(\alpha_v,\beta_v)\cap\mathcal{B}}, (X, Y_X)\mathbf{1}_{\mathcal{A}(X,Y_X)}\right) = e^{\log^3 N} e^{-\Omega\left(e^{\log^{2/3} N}\right)} + N^{-1+o_N(1)} = o_N(1). \tag{4.11}$$

□

The above two lemmas are crucial as we can now essentially treat the relevant pairs (α_x, β_x) as independent Poisson random variables.

4.4 Tails of sums

Sums of squares of binomials

When approximating the largest eigenvalues of $\mathcal{G}\left(N, \frac{d}{N}\right)$ using local neighborhood statistics of high degree vertices, we will need bound the sums of squares of the degrees in neighborhoods of those high degree vertices. Therefore we require estimates for the sum of distributions with heavy Weibull tails. Such a bound follows from (4.7) and the tail results in [BMdlP23]. Justification for this generalization is given in the appendix in Section A.2.

Lemma 4.10. For any $n > 0$ and $d = o(n^{1/3})$, consider n independent i.i.d. samples $X_1, \dots X_n \sim \text{Binom}(N, \frac{d}{N})$. There is some constant $c_{4.10} > 0$ such that for any $t > n^{2/3}$,

$$\mathbb{P}\left(\left|\sum_{i=1}^{n} X_i^2 - \mathbb{E}\left[\sum_{i=1}^{n} X_i^2\right]\right| > t\right) \le 2n \exp\left(-\frac{c_{4.10}}{d^3 + 1}\sqrt{t}\right).$$

Moreover, if $t > 2(d^2+1)n^{2/3}$,

$$\mathbb{P}\left(\left|\sum_{i=1}^n X_i^2 - \mathbb{E}\left[\sum_{i=1}^n X_i^2\right]\right| > t\right) \le 2n \exp\left(-c_{4.10}\sqrt{t}\right).$$

Sums of random variables with Weibull shape

Similar bounds are also needed when considering the spectrum of weighted graphs. In the light-tailed case the sum of squares of the weights appears because it corresponds to the eigenvalue of a star. In the heavy-tailed case the sum of the αth power of the weights is introduced through our results on the relationship between the α-(quasi)norm of a matrix and its spectrum from section 2.1.

In this section, we state two key lemmas about the tail of a sum of i.i.d. Weibull random variables and their conditioned version. Throughout this section, we assume that $\{Y_i\}_{i=1,2,\dots}$ are i.i.d. random variables such that for $t > 0$,

$$\frac{C_1}{2}e^{-t^\alpha} \le \mathbb{P}\Big(Y_i \ge t\Big) \le \frac{C_2}{2}e^{-t^\alpha} \quad \text{and} \quad \frac{C_1}{2}e^{-t^\alpha} \le \mathbb{P}\Big(Y_i \le -t\Big) \le \frac{C_2}{2}e^{-t^\alpha}. \tag{4.12}$$

This implies that

$$C_1 e^{-t^\alpha} \le \mathbb{P}(|Y_i| \ge t) \le C_2 e^{-t^\alpha}.$$

Also, for our applications, as will appear several times in chapter 6 we define, for $\varepsilon > 0$, $\tilde{Y}_i$ as the random variable Y_i conditioned to be greater than $(\varepsilon \log\log N)^{\frac{1}{\alpha}}$ in absolute value.

Tails of sums of light-tailed random variables

First, we introduce a tail bound for the sum of squares of i.i.d. Weibull random variables having a lighter tail than the Gaussian distribution, i.e. $\alpha > 2$. Recall that $\lambda_\alpha^{\text{light}}$ denotes the typical value of the largest eigenvalue, defined in (6.3):

$$\lambda_\alpha^{\text{light}} = \left(\frac{2}{\alpha}\right)^{\frac{1}{\alpha}} \left(1 - \frac{2}{\alpha}\right)^{\frac{1}{2}-\frac{1}{\alpha}} \frac{(\log N)^{\frac{1}{2}}}{(\log\log N)^{\frac{1}{2}-\frac{1}{\alpha}}}.$$

Lemma 4.11. Assume that $\alpha > 2$. Then, for any $t > k \ge 2$,

$$C_1^k e^{-t^{\frac{\alpha}{2}} k^{1-\frac{\alpha}{2}}} \le \mathbb{P}\left(Y_1^2 + \cdots + Y_k^2 \ge t\right) \le C_2^k \left(\frac{2et}{k}\right)^k e^{-(t-k)^{\frac{\alpha}{2}} k^{1-\frac{\alpha}{2}}}. \tag{4.13}$$

In particular, assume that $t = d^2\left(\lambda_\alpha^{\text{light}}\right)^2 + o\left(\frac{\log N}{(\log\log N)^{1-\frac{2}{\alpha}}}\right)$ and $k = b\frac{\log N}{\log\log N} + o\left(\frac{\log N}{\log\log N}\right)$ for some constants $b, d > 0$. Then,

$$\lim_{N\to\infty} -\frac{\log \mathbb{P}\left(Y_1^2 + \cdots + Y_k^2 \ge t\right)}{\log N} = d^\alpha \frac{2}{\alpha - 2}\left(1 - \frac{2}{\alpha}\right)^{\frac{\alpha}{2}} b^{1-\frac{\alpha}{2}} \tag{4.14}$$

and

$$\lim_{N\to\infty} -\frac{\log \mathbb{P}\left(\tilde{Y}_1^2 + \cdots + \tilde{Y}_k^2 \geq t\right)}{\log N} \geq d^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}} b^{1-\frac{\alpha}{2}} - b\varepsilon. \tag{4.15}$$

(recall that $\tilde{Y}_i$ is a conditioned version of Y_i).

Finally, in the case $t = d^2\left(\lambda_\alpha^{\text{light}}\right)^2 + o\left(\frac{\log N}{\log\log N}\right)$ and $k = o(1)\frac{\log N}{\log\log N}$, we have

$$\lim_{N\to\infty} -\frac{\log \mathbb{P}\left(Y_1^2 + \cdots + Y_k^2 \geq t\right)}{\log N} = \lim_{N\to\infty} -\frac{\log \mathbb{P}\left(\tilde{Y}_1^2 + \cdots + \tilde{Y}_k^2 \geq t\right)}{\log N} = \infty. \tag{4.16}$$

The particular choices of t and k considered in (4.14)-(4.16) appear in our applications in Section 6.2.

Tails of sums of the αth-power of random variables

The next crucial result is about the tail estimate for the i.i.d. sum of αth-power of *conditioned* Weibull random variables for any $\alpha > 0$.

Lemma 4.12. Suppose that $\alpha, \varepsilon > 0$. Then, there exists a constant $C > 0$ depending only on C_1, C_2 such that the following holds: For any $L > m$,

$$\mathbb{P}\left(|\tilde{Y}_1|^\alpha + \cdots + |\tilde{Y}_m|^\alpha \geq L\right) \leq C^m e^{-L} e^m \left(\frac{L}{m}\right)^m e^{\varepsilon m \log\log N}. \tag{4.17}$$

In particular, assume that $m \leq b\frac{\log N}{\log\log N} + c$ and $L = a\log N$ for some constants $a, b, c > 0$. Then,

$$\mathbb{P}\left(|\tilde{Y}_1|^\alpha + \cdots + |\tilde{Y}_m|^\alpha \geq a\log N\right) \leq N^{-a+\varepsilon b+o(1)}. \tag{4.18}$$

Remark 4.1. The above lemmas 4.11 and 4.12, with the same argument, hold as well for a slightly more general class of Weibull random variables satisfying

$$C_1 t^{-c_1} e^{-\eta t^\alpha} \leq \mathbb{P}(|Y_i| > t) \leq C_2 t^{-c_2} e^{-\eta t^\alpha}$$

with some constants $c_1, c_2 \geq 0$. This can be used to generalize our results as indicated in Remark 6.2.

Chapter 5

Description of the edge spectrum of sparse Erdős-Rényi graphs

Given the preliminary results from the previous three chapters, we now focus on the behavior of the eigenvalues of $\mathcal{G}\left(N, \frac{d}{N}\right)$ near the edge of the spectrum. This spectral edge has received much attention, as it has its own specific applications. For example, the edge governs the mixing rate of Markov chains, and graph partitioning, as shown in [HLW06], albeit using the Laplacian rather than the adjacency matrix.

In this chapter we will focus on properties of the extreme eigenvalues and eigenvectors that happen with high probability, and all results mentioned in this introduction, unless otherwise specified, are also to be interpreted as such. The typical behavior of the extreme eigenvalues and eigenvectors in the Erdős-Rényi model is known to go through various phase transitions. When $d \gg N^{1/3}$ these edge eigenvalues have Tracy-Widom fluctuations, similar to the fluctuations of the eigenvalues of a GOE matrix [Sos99, EYY12, EKYY13, LS18]. When $N^{\epsilon} \leq d \ll N^{1/3}$, for some fixed $\epsilon > 0$, the top eigenvalues lose GOE behavior and edge eigenvalues become Gaussian distributed [HLY20, HK21].

For sparser Erdős-Rényi graphs, Krivelevich and Sudakov showed using a graph decomposition that eigenvalues are governed by the highest degree vertices [KS03]. More precisely, with $\mathfrak{u}$ denoting the largest degree we expect to occur in the graph, Krivelevich and Sudakov showed that the largest eigenvalue of an Erdős-Rényi graph is typically $(1 + o_N(1)) \max\{d, \sqrt{\mathfrak{u}}\}$. For $\log^{-1} N \ll d \ll \log N$, $\mathfrak{u} = \Theta\left(\frac{\log N}{\log\log N}\right)$. This shows that there is a phase transition in the largest eigenvalue at $d = \sqrt{\frac{\log N}{\log\log N}}$.

In fact, we begin to see the local affect of high degree vertices in the spectrum at $d \asymp \log N$. This is well known to be the threshold for connectivity (as was shown in the original work of Erdős and Rényi [ER60]), but is also the threshold for large fluctuations in the degree sequence, as opposed to greater concentration seen for larger d. Specifically, Benaych-Georges, Bordenave, and Knowles, as well as Latała, van Handel, and Youssef showed that when $d \gg \log N$, edge eigenvalues converge to the edge of the support of the asymptotic eigenvalue distribution [LvHY18, BGBK20]. However Benaych-Georges, Bordenave, and Knowles also showed that when $d \ll \log N$, roughly, edge eigenvalues are "governed" by the largest degree vertices of the adjacency matrix [BGBK19], with the specific threshold later given independently by Alt, Ducatez, and Knowles, as well as Tikhomirov and Youssef [ADK21b, TY21].

Alt, Ducatez, and Knowles further studied this problem, and managed to obtain impressively detailed results. Through the works of [ADK21a, ADK22, ADK23b, ADK23a], the authors show a transition between the occurrence of delocalized eigenvectors in the bulk of the spectrum, and localized eigenvectors near the edge for $d \lesssim \log N$ (the specific bounds on d vary paper to paper, but all results are for this sparse regime).

We focus specifically on [ADK23b]. In this result, Alt, Ducatez, and Knowles show that the largest eigenvalues of the graph are determined by two combinatorial statistics around the high degree vertices. As was shown previously through [KS03] and [BGBK19], the primary term is the degree of the high degree vertex, which we denote here by α_x. To gain the necessary levels of accuracy, they also track the secondary term β_x, which is the number of vertices of distance exactly 2 from a high degree vertex x in the graph.[1] Note that this notation slightly differs from the one in [ADK23b], where α_x and β_x are normalized by d. Reinterpreting their result, they show the following.

Theorem 5.1 ([ADK23b]). For $\zeta > 4$ and sufficiently small constant $\xi > 0$, assume that $(\log\log N)^\zeta \leq d \leq \left(\frac{1}{\log(4)-1} - (\log N)^{-\xi}\right)\log N$. For $K := d^{1/2-2/\zeta-16\xi}$ there are some $\delta, \epsilon > 0$ such that the first K eigenvalues are of the form

$$\frac{\alpha_x}{\sqrt{\alpha_x - \frac{\beta_x}{2\alpha_x}\left(\frac{\alpha_x}{d} + \frac{\beta_x}{\alpha_x d}\right) + \frac{\beta_x}{2\alpha_x}\sqrt{\left(\frac{\alpha_x}{d} + \frac{\beta_x}{d\alpha_x}\right)^2 - 4\frac{\alpha_x}{d}}}} + O(d^{-\epsilon}\mathfrak{u}^{-1}) \tag{5.1}$$

for K vertices of degree at least $\mathfrak{u} - \frac{d^{-\delta}}{\log \mathfrak{u}}$.

This is done by, given α_x and β_x, making an educated guess for the structure of the eigenvector. To use this approximate eigenvector it is crucial that the statistics of the local neighborhoods of high degree vertices are concentrated, in particular that the degrees of the vertices in the neighborhood is reasonably close to d, which is approximately their expected value. According to (5.1), to be able to properly estimate eigenvalues past the typical gap in the spectrum at the edge of $\Theta(\mathfrak{u}^{-1})$, it must be the case that $d \gg 1$.

5.1 Main results

Guionnet posed the question of deducing the behavior for d constant [Gui21]. The difficulty in this regime is that important statistics, like the growth of the spheres around

[1] We have translated their parameters into the unnormalized versions we use in our proof.

high-degree vertices and the maximum degrees in those spheres, are less concentrated. We use new techniques to show that in fact, the same behavior holds in the constant d regime, and continues until the average degree is subconstant. All the results in this chapter are from the paper [HM23].

Theorem 5.2. Consider a $G \sim \mathcal{G}(N, \frac{d}{N})$ graph with $\log^{-1/9} N \le d \le \log^{1/40} N$. With high probability, for each of the $e^{\log^{1/8} N}$ largest eigenvalues λ of the adjacency matrix, there is some vertex x such that

$$\lambda = \sqrt{\alpha_x + \frac{\beta_x}{\alpha_x} + \frac{d^2 + d}{\alpha_x}} + O\left(\left(d^{3/2} + 1\right) \mathfrak{u}^{-11/6}\right).$$

Moreover, for $k \le e^{\log^{1/8} N}$, the vertex x corresponding to the kth largest eigenvalue is the kth vertex in the lexicographic ordering (α_x, β_x).

As we will show, these high degree vertices are spaced throughout the graph, and have almost independent local statistics. Therefore, similar to [ADK23b], the distribution of the highest eigenvalues is described by a Poisson point process with density given by the probability of existence of an (α, β) pair, where α is close to maximal among all vertices.

To this end, define the discrete intensity measure $\rho : \mathbb{R} \to \mathbb{R}$,

$$\rho\left(\frac{s}{\mathfrak{u}}\right) := N \sum_{\ell=0}^{2\log^{1/8} N} \left(\frac{e^{-d} d^{\mathfrak{u}-\ell}}{(\mathfrak{u}-\ell)!} \frac{e^{-d(\mathfrak{u}-\ell)} (d(\mathfrak{u}-\ell))^{(s-\mathfrak{u}+\ell)(\mathfrak{u}-\ell)}}{((s-\mathfrak{u}+\ell)(\mathfrak{u}-\ell))!} \mathbf{1}_{\langle s(u-\ell)\rangle = 0} \mathbf{1}_{(s-u+\ell)(u-\ell) \le d(\mathfrak{u}-\ell)+\mathfrak{u}^{7/8}} \right).$$

where $\mathbf{1}_{\langle x\rangle=0}$ is the indicator that x is a whole number. ρ induces a Poisson point process Ψ. The meaning of ρ is that it is the intensity measure of $\alpha + \frac{\beta}{\alpha}$ if $\alpha \sim Pois(d)$ and $\beta \sim Pois(d\alpha)$, restricted to $\alpha \in \left[\mathfrak{u} - 2\log^{1/8} N, \mathfrak{u}\right], \beta \in \left[0, d\alpha + \mathfrak{u}^{7/8}\right]$. As we will see, Theorem 5.2 implies Ψ approximates the density of λ_x^2 at the edge of the spectrum.

Formally, we we will consider proximity in Lévy-Prokhorov distance, which is a metrization of the weak topology. Namely we define, for two Borel measures ν_1, ν_2 on $\mathbb{R}$,

$$D(\nu_1, \nu_2) = \inf\left\{\epsilon > 0 \ : \ \forall A \in \mathfrak{B}, \nu_1(A) \le \nu_2(A_\epsilon) + \epsilon \text{ and } \nu_2(A) \le \nu_1(A_\epsilon) + \epsilon\right\}$$

where $\mathfrak{B}$ is the set of Borel measurable sets in $\mathbb{R}$ and A_ϵ is the neighborhood of radius ϵ around A.

Moreover, for $K > 0$, define $\kappa(K)$ as

$$\kappa(K) = \inf\left\{s \in \mathbb{R} : \rho([s, \infty)) \le K\right\}. \tag{5.2}$$

Theorem 5.3. Set $K = e^{\log^{1/8} N}$, and recall the definition of $\kappa(K)$ from (5.2). Consider the density function

$$\Phi := \sum_{\lambda \in spec(A)} \delta_{\mathfrak{u}\left(\lambda^2 - \frac{d^2+d}{\mathfrak{u}}\right)}.$$

Then for d as defined in Theorem 5.2,

$$\lim_{N\to\infty} D\left(\Phi\mathbf{1}_{[\kappa(K),\infty)}, \Psi\mathbf{1}_{[\kappa(K),\infty)}\right) \to 0.$$

We multiplied both point processes by $\mathfrak{u}$ to emphasize that the approximation of λ^2 in Theorem 5.2 is $o(1/\mathfrak{u})$. As shown in Theorem 5.2, the largest eigenvalues are approximately the square root of a rational function on local statistics, therefore it is simpler to consider the point process of λ^2, as it leads to a nicer expression. Much of our analysis will be on the squared eigenvalue λ^2.

Theorem 5.3 implies the fluctuations of the top eigenvalue. Similar to if there is increasing degree as in [ADK23b], fluctuations are determined at two scales. If the expected number of vertices of degree $\mathfrak{u}$ (that is $\mu_{\mathfrak{u}}$) is constant, then the maximum degree has nonzero variance and will dominate the fluctuation of the maximum eigenvalue. Therefore, in this case the top eigenvalue will fluctuate at a large scale, as its first order fluctuation is a shifted Bernoulli, based on whether there is a vertex of degree $\mathfrak{u}$ or not. If the expected number of vertices of degree $\mathfrak{u}$ is subconstant or superconstant, then the largest degree of the graph becomes deterministic, and the fluctuations are determined by β. As β is then distributed as a Poisson, the fluctuations become much smaller and become those of the maximum of Poissons.

The fact that these eigenvalues are almost completely determined by local neighborhoods is intrinsically linked to the fact that they decay exponentially around a fixed vertex. We show the following, which implies eigenvectors are close to as localized as possible.

Theorem 5.4. Define $B_r(x)$ to be the ball of radius r around x in G. For $c \le 1/15$, consider $\log^{-c} N \le d \le \log^{1/40} N$, and $K_2 = \log^{o_N(1)} N$. Moreover, we fix $r' \ge 1$, and if $c > 0$, we add the requirement that $r' \le 1/(3c)$. With high probability, the eigenvectors $\mathbf{v}$ corresponding to the K_2 largest eigenvalues are *exponentially localized* in the sense that for each $\mathbf{v}$ there is some vertex x such that

$$\mathbf{v}|_x = \frac{1}{\sqrt{2}} + O\left(\left(1 + d^{-1/2}\right)\mathfrak{u}^{-1/3}\right)$$

and for $1 \le i \le r'$,

$$\|\mathbf{v}|_{S_i(x)}\| = \left(\frac{d}{\alpha}\right)^{(i-1)/2} \frac{1}{\sqrt{2}}\left(1 + O\left(\left(1 + d^{-1/2} + d^{-i+1}\right)\mathfrak{u}^{-1/3}\right)\right)$$

and

$$\|\mathbf{v}|_{[N]\setminus B_i(x)}\| = \left(\frac{d}{\alpha}\right)^{i/2} \frac{1}{\sqrt{2}}\left(1 + O\left(\left(1 + d^{-1/2} + d^{-r'+1}\right)\mathfrak{u}^{-1/3}\right)\right).$$

Our desire to keep our result as general as possible has resulted in this long expression for our error. There are multiple error terms and for $d \asymp 1$, the ratio of d, i, and $\mathfrak{u}$ governs which type of error will dominate.

We consider a significantly smaller number of eigenvalues in this theorem than in Theorem 5.2 as in order to show eigenvector localization, we must quantify the gaps of the eigenvalues induced by Theorem 5.2 and Theorem 5.3, whereas the previous theorems do not require such gaps. We do this only for the region given above, as our focus is the spectral edge, but most likely further analysis could extend this to further eigenvalues.

Using the regimes as defined in Definition 3.1, we can prove a theorem about the approximate diagonalization of the adjacency matrix A, that we will use to prove the other main theorems above.We will decompose our matrix using a unitary transform U made clear later. Here $D_{\mathcal{W}}, D_{\mathcal{V}\setminus\mathcal{W}}, D_{\mathcal{U}\setminus\mathcal{V}}$ are diagonal operators associated with the balls surrounding vertices in $\mathcal{U}$. A summary of the results concerning this decomposition is as follows. This can be compared to Proposition 3.1 in [ADK23a], and it also bears resemblance to the more directly combinatorial decompositions of other sparse matrix results [KS03, BBG21].

Theorem 5.5. With high probability, there is a unitary transformation $U : \mathbb{R}^N \to \mathbb{R}^N$ such that we can write

$$A = U \begin{pmatrix} D_{\mathcal{W}} & 0 & 0 & E^*_{\mathcal{W}} \\ 0 & D_{\mathcal{V}\setminus\mathcal{W}} & 0 & E^*_{\mathcal{V}\setminus\mathcal{W}} \\ 0 & 0 & D_{\mathcal{U}\setminus\mathcal{V}} + \mathcal{E}_{\mathcal{U}\setminus\mathcal{V}} & E^*_{\mathcal{U}\setminus\mathcal{V}} \\ E_{\mathcal{W}} & E_{\mathcal{V}\setminus\mathcal{W}} & E_{\mathcal{U}\setminus\mathcal{V}} & \mathcal{X} \end{pmatrix} U^* \tag{5.3}$$

where (5.3) satisfies the following:

1. The *fine regime* operator $D_{\mathcal{W}}$ is diagonal, of dimension $2|\mathcal{W}|$, and has at least $e^{\log^{1/8} N}$ eigenvalues of value at least $\sqrt{\mathfrak{u}} - O\left(\mathfrak{u}^{-3/8}\right)$.

2. The *intermediate regime* operator $D_{\mathcal{V}\setminus\mathcal{W}}$ is diagonal, of dimension $2|\mathcal{V}\setminus\mathcal{W}|$, and $\|D_{\mathcal{V}\setminus\mathcal{W}}\| \leq \sqrt{\mathfrak{u}} - \Omega\left(\mathfrak{u}^{-1/4}\right)$.

3. The *rough regime* operator $D_{\mathcal{U}\setminus\mathcal{V}}$ is diagonal and of dimension $2|\mathcal{U}\setminus\mathcal{V}|$ and satisfies $\|D_{\mathcal{U}\setminus\mathcal{V}}\| \leq \sqrt{\mathfrak{u}} - \mathfrak{u}^{1/6 - o_N(1)}$.

4. The *bulk* operator $\mathcal{X}$ satisfies $\|\mathcal{X}\| = \left(\frac{1}{\sqrt{2}} + o_N(1)\right)\sqrt{\mathfrak{u}}$.

5. The error terms satisfy $\|E_{\mathcal{W}}\|, \|E_{\mathcal{V}\setminus\mathcal{W}}\| = O\left((d^r + 1)\mathfrak{u}^{-r/2+1}\right)$, as well as $\|\mathcal{E}_{\mathcal{U}\setminus\mathcal{V}}\| + \|E_{\mathcal{U}\setminus\mathcal{V}}\| = O\left((\log\log N)^2\right)$.

These results, along with results concerning the structure of the eigenvectors associated with the operator surrounding vertices of $\mathcal{W}$, imply that the edge eigenvectors and eigenvalues come from the operator associated with $D_{\mathcal{W}}$. This in turn will give the theorems from Section 5.1. Thus the next few sections are dedicated to showing this decomposition.

Negative eigenvalues

The discussion above concerns the largest eigenvalues of the adjacency matrix, however, by the exact same analysis we can consider the most negative eigenvalues. Under the high probability assumption that the neighborhood of every high degree vertex is a tree, by the bipartite nature of a tree, every positive eigenvalue of a neighborhood of a tree has a corresponding negative eigenvalue that is of the same magnitude and has the same localization properties. Therefore, Theorems 5.2, 5.3, and 5.4 all apply to the most negative eigenvalues as well.

Extension of results

We believe that by increasing the analysis from our given set of local statistics to higher moments, our methods can be used to give even more accurate formulae for the largest eigenvalues based on the degree sequence of the highest degree vertices. Such an argument could show separation of the largest $\log^a N$ eigenvalues for any fixed $a \geq 0$, giving a more specific (and more complicated) Poisson point process and showing eigenvector localization for all these $\log^a N$ eigenvectors. However, for simplicity of the argument, in this work we only consider $K = \log^{o_N(1)} N$.

Using this same argument, we may be able to improve the necessary lower bound on d. Some estimates and concentration results required us to lower bound d by $\log^{-c} N$ for some $c < 1$, and error bounds simplify given our concrete assumption on d, but there are also important structural consideration for smaller d. If $d = \log^{-c} N$ for $c > 0$, then the connected components surrounding high degree vertices can be of small radius, which means that neighborhoods of high degree vertices could be identical, leading to the eigenvalues of those neighborhoods being identical, and thus the eigenvectors would no longer be localized around one high-degree vertex. This implies that we cannot remove the dependence of d on r in Theorem 5.4.

On the other hand, including more terms in our expression for λ_x could improve the lower bound on d for Theorem 5.2 and Theorem 5.3. However, in our opinion we will need new methods to achieve the threshold of $d = e^{-(\log\log N)^2}$ appearing in [KS03] and [BBG21]. Such a threshold is natural as for $d \leq e^{-(\log\log N)^2}$, all connected components are of size $(1 + o(1))\mathfrak{u}$, making localization and independence trivial.

Similarly, by using slightly tighter bounds on the probabilities of some tail events, we expect we can improve the upper bound on d in Theorem 5.2. However, there is a natural barrier at $\log^{1/2} N$, both in the exponential rate of decay of eigenvectors, and the application of the method in [KS03]. We are not motivated to optimize our technique for the upper bound considering results are already known for larger d from [ADK23b].

Related Work

There are many results concerning the bulk of the spectrum of random matrices. Focusing specifically on sparse Erdős-Rényi graphs, Khorunzhy, Shcherbina, and Vengerovsky, then Zakharevich, analyzed the moments of the limiting distribution of Wigner matrices of general models that include constant degree Erdős-Rényi graphs in order to study the limiting measure of the spectral distribution [KSV04, Zak06]. Benaych-Georges, Guionnet, and Male give a central limit theorem for linear statistics of a model that includes constant degree Erdős-Renyi graphs [BGGM14].

These random matrices have been studied as a model for quantum physics, specifically Hamiltonians of disordered systems. We see similar eigenvector localization in the edge of the spectrum in the Anderson model (see [And58]), where vertices on an integer lattice are given random potential, and we study the spectrum of the resulting Schrödinger operator. Eigenvectors near the edge of the spectrum are known to be localized for various models (e.g. [GMP77, FS83, AM93, DS20]) whereas there has been less progress on the structure of eigenvectors in the bulk.

Lévy matrices, a model of Wigner matrices where entries are sampled from distributions with heavy tails, have also proved to be a useful model for studying eigenvector localization. Similar to sparse Erdős-Rényi, there is a transition from delocalized eigenvectors in the bulk to localized eigenvectors at the edge [BG13, BG17, ALM21, ALY21, ABL22]. Moreover, similar to the sparse Erdős-Rényi model, eigenvalues near the edge of the spectrum in sparse Lévy models are known to converge to a Poisson point process [Sos04, ABP09].

Idea of Proof

We follow the framework of the proof of [ADK23b]. Our first goal is to show that the largest eigenvalues are determined by the local geometry of the highest degree vertices, i.e. by truncated balls around them. Once we show this, the Poisson point process and eigenvector structure follow from the randomness of the graph. We show this determination by classifying vertices by degree, and associating an eigenvector with each high degree vertex. These are the eigenvectors with eigenvalues of largest magnitude. In order to analyze these eigenvectors, they are approximated with the eigenvector of an infinite tree with much symmetry in [ADK23b]. The main issue with generalizing this approximation to gaphs with average constant degree is that in this regime, the fluctuations of statistics of the balls surrounding individual entries become too large to estimate the eigenvector based on the inputs α_x and β_x only.

The formula in [ADK23b] already is quite technical, so we avoid directly coming up with a more involved equation that gives a more accurate explicit approximation. Instead, we analyze the properties of the true eigenvector of a neighborhood directly. The advantage of such a method is that the only error we generate in this analysis is from truncating at level r. Therefore, if we can show the eigenvector is localized away from level r, then the error from this truncation can be drastically smaller. The disadvantage is that, considering we have no symmetry in our tree, we initially have no information about the eigenvector or eigenvalue, even whether it is localized or not.

The neighborhoods of the highest degree vertices are typically tree-like, and the central vertex has much higher degree than all others. Therefore, by analyzing the eigenvector equation at each vertex, we create a system of linear equations for the eigenvector and eigenvalue. The knowledge that the neighborhood is a tree and that the central degree is much larger than all other degrees is sufficient to show eigenvector localization (see Lemma 5.2) and a formula for the eigenvalue generated from a recursive equation (see (5.6) and Lemma 5.3). Moreover, because the central vertex has much higher degree, we can show this equation can be truncated up to small error with a rational function of local statistics, giving near exact dependence. Because of the general nature and simplicity of these lemmas, we believe they could be applicable elsewhere.

When we fully write out the equation for the eigenvalue, the first two terms are α_x and $\frac{\beta_x}{\alpha_x}$, which are used in [ADK23b] to completely determine the eigenvector. We show, explicitly giving the next few terms, that the equation for the eigenvalue beyond α, β concentrates. Specifically, although the fluctuation of statistics increases as the average degree decreases, the dependence on the fluctuating statistics decreases at a quicker rate (see Lemma 5.4). In fact, in our regime, the eigenvalue decays quickly enough that it implies the lexicographic ordering of Theorem 5.2 (see Lemma 5.5).

Given the dependence of the eigenvalue and eigenvectors on statistics of local neighborhoods, we translate this into statements about the overall graph. By standard perturbation theoretic arguments, this reduces to showing these local statistics are well separated for vertices corresponding to the edge of the spectrum. The requisite statistics are binomially distributed, which are approximately Poisson. Therefore it becomes useful to give precise tail bounds on the Poisson distribution. Tao recently gave such a tight, two-sided bound on his blog [Tao22], which is sufficient to show that (α_x, β_x) that are close to lexicographically maximizing are well separated (see Section 5.5).

Once we have proper control over these high degree vertices, we need to control the contribution of the rest of the spectrum. To do this, for vertices of still somewhat large degree, we proceed as per [ADK23b] and take a localized test vector supported on a pruned graph, where edges are pruned in such a way that high degree vertices are separated and neighborhoods are tree-like (see Section 5.3). For all other vertices, we use the work of Krivelevich and Sudakov, who prove a convenient decomposition of edges of the graph into different components [KS03]. The only one of these components that can significantly contribute to large eigenvalues is a subgraph of disjoint stars, meaning it is much simpler to bound the spectral radius of the operator away from the largest eigenvectors of the highest degree vertices and avoids using the Ihara-Bass argument of [ADK23b] (see Lemma 5.7).

Overview of the chapter

The next two sections analyze the regimes defined in 3.2. In Section 5.2, we analyze the eigenvector and eigenvalue of the largest eigenvalue of the ball of radius r surrounding the highest degree vertices, and we show that it is localized and well approximated by a formula involving only α, β, (see the beginning of Section 3.2 for the definition of these parameters) and d. In Section 5.3, we use a test vector to show that eigenvectors corresponding to vertices of high, but not too high degree, do not have large eigenvalues. In Section 5.4, we analyze the bulk using the decomposition from [KS03], and we show the given block decomposition is such that the highest degree vertices dominate. This gives theorems 5.5 and 5.2. In Section 5.5, we use this formula to prove the Poisson process, approximation and in Section 5.6 we show this implies Theorem 5.4, our result about the eigenvector localization.

Parameters

Although our choice of parameters is mentioned for all the main results we state them here a condensed form, for easier reference.

Definition 5.1 (Choice of parameters)**.** We fix $c > 0$ and we take the average degree $d(N)$ as any function such that for Theorem 5.2 and Theorem 5.3, $\log^{-1/9} N \leq d \leq \log^{1/40} N$, and for Theorem 5.4, $\log^{-c} N \leq d \leq \log^{1/40} N$ for $c \leq 1/15$.

Our analysis will be based on considering balls of radius r around the highest degree vertices. Most results are true for sufficiently large r, but in fact, it is enough to take $r = 5$ in order to prove Theorems 5.2 and 5.3. For Theorem 5.4, we need a slightly larger radius. So for the rest of this chapter it is sufficient to take $r := \max\{5, 2r'\}$, where r' is the parameter from Theorem 5.4.

5.2 Fine and Intermediate Regime

In this section we start by using the structural results about the balls around vertices in $\mathcal{V}$ from Section 3.3 to derive a recursion for the largest eigenvalue and eigenvector of the balls around the vertices in that regime. We also derive a first approximation of the largest eigenvalue for all vertices in $\mathcal{V}$. In the second part we then use these ingredients to prove the exponential decay of this eigenvector for all vertices in $\mathcal{V}$. In the last part we then derive a more precise expression for the eigenvalue of vertices in $\mathcal{W}$.

Notation

We want to emphasize that in this chapter we depart from our previous notation of denoting the largest eigenvalues of the graph by $\lambda_1 \geq \cdots \geq \lambda_N$. Instead we will generally use λ_x to denote the maximum eigenvalue of the truncated ball around the vertex $x \in [N]$.

General Structure

In this section, we create a test vector and test eigenvalue for the neighborhoods of vertices in $\mathcal{V}$. In order to do this we assume that the structural properties from Definition 3.2 hold. We know that this happens with high probability from Lemma 3.5. We also know from Lemma 3.1, that with high probability the maximum degree is bounded by $\mathfrak{u}$. Therefore, we will consistently use these structural properties.

Assumption 5.1. For the rest of Section 5.2, we assume that $\Omega_{3.2}$ occurs. Moreover we assume the high probability event that the maximum degree in G is in $\{\mathfrak{u}-1, \mathfrak{u}\}$.

Under this event we know enough about the structure of the neighborhoods around vertices in the intermediate regime, to analyze its top eigenvalue and eigenvector, which we use as the test eigenvector and eigenvalue.

Definition 5.2. For $x \in \mathcal{V}$, we define λ_x to be the top eigenvalue of $A_{B_r(x)}$. Moreover, define $\mathbf{w}_+(x)$ to be the eigenvector corresponding to λ_x and $\mathbf{w}_-(x)$ to be the eigenvector of the most negative eigenvalue of $A_{B_r(x)}$. We use the notation $\mathbf{w}_\pm(x)$, when a statement is true for both $\mathbf{w}_+(x)$ and $\mathbf{w}_-(x)$. Depending on the context, we will also use $\mathbf{w}_\pm(x)$ to denote the above eigenvector padded with 0's to make it a vector in $\mathbb{R}^N$.

Under $\Omega_{3.2}$, $B_r(x)$ is a tree, therefore it is bipartite. Therefore the most negative eigenvalue is $-\lambda_x$, and if $y \in B_r(x)$, then

$$\mathbf{w}_-(x)|_y = (-1)^{d(x,y)}\mathbf{w}_+(x)|_y.$$

Moreover, as $B_r(x)$ is a tree, its eigenvectors and eigenvalues satisfy a nice recursion. Consider an eigenvector $\mathbf{w}$ of $A_{B_r(x)}$ with eigenvalue λ. We consider the eigenvector equation at x:

$$\lambda \mathbf{w}|_x = \sum_{y\sim x} \mathbf{w}|_y$$

which, if $\mathbf{w}|_x \neq 0$, can be rewritten as

$$\lambda = \sum_{y\sim x} \frac{\mathbf{w}|_y}{\mathbf{w}|_x}. \tag{5.4}$$

More generally, we have for $v \sim u, v \geq u$, i.e. for v a child of u in the tree rooted at x, if $\mathbf{w}|_v \neq 0$,

$$\lambda \mathbf{w}|_v = \mathbf{w}|_u + \sum_{y_1\sim v, y_1\geq v} \mathbf{w}|_{y_1} \Rightarrow \mathbf{w}|_v = \frac{1}{\lambda - \sum_{y_1\sim v, y_1 \geq v} \frac{\mathbf{w}|_{y_1}}{\mathbf{w}|_v}} \mathbf{w}|_u. \tag{5.5}$$

Plugging in (5.5) into (5.4) for every $y_1 \sim x$ gives

$$\lambda = \sum_{y_1 \sim x} \frac{\frac{1}{\lambda - \sum_{y_2 \sim y_1, y_2 \geq y} \frac{\mathbf{w}|_{y_2}}{\mathbf{w}|_{y_1}}} \mathbf{w}|_x}{\mathbf{w}|_x} = \sum_{y_1 \sim x} \frac{1}{\lambda - \sum_{y_2 \sim y_1, y_2 \geq y_1} \frac{\mathbf{w}|_{y_2}}{\mathbf{w}|_{y_1}}}.$$

Repeating this process for all vertices gives that

$$\lambda = \sum_{y_1 \sim x} \frac{1}{\lambda - \sum_{y_2 \sim y_1, y_2 \geq y_1} \frac{1}{\lambda - \sum_{y_3 \sim y_2, y_3 \geq y_2} \frac{1}{\lambda - \sum_{y_4 \sim y_3, y_4 \geq y_3} \cdots}}}, \tag{5.6}$$

where the right hand side is a continued fraction of at most r levels.

Since $A_{B_r(x)}$ is a connected graph, its adjacency matrix is irreducible and this implies by the Perron-Frobenius theorem, that the top eigenvector $\mathbf{w}|_+(x)$ of $B_r(x)$ is the only positive eigenvector, implying in particular that (5.6) does not contain any 0 denominators. This means that we can use (5.6) for our definition of λ_x. To further examine this, we require an initial two sided bound, based on $B_r(x)$ being close to a star graph. This will be enough to bound the contribution of balls around vertices in $\mathcal{V} \setminus \mathcal{W}$, and for vertices in $\mathcal{W}$, we will eventually bootstrap it into a tighter bound in Lemma 5.4.

Lemma 5.1. For any vertex $x \in \mathcal{V}$,

$$\alpha_x \leq \lambda_x^2 \leq \alpha_x + O(d).$$

Proof. As the spectral radius of a star is the square root of the degree of the central vertex, $\lambda_x^2 \geq \alpha_x$. For the upper bound, we apply Lemma 2.10. By the definition of $\Omega_{3.2}$, (3) we can take $s(n) = 2d$, and by the definition of $\Omega_{3.2}$, (4) we can take $t(n) = \text{u}^{3/4}$. Lemma 2.10 thus implies that $\lambda_x = \sqrt{\alpha_x} + O\left(\frac{d}{\sqrt{\alpha_x}}\right)$ implying that $\lambda_x^2 = \alpha_x + O(d)$.

□

Eigenvector structure

For easier readability, we will suppress x in the notation for the rest of the section, so we write $\alpha := \alpha_x$, $\lambda := \lambda_x$, $\mathbf{w}_\pm = \mathbf{w}_\pm(x)$, and $S_1 = S_i(x)$. Moreover we define $N_{y^*} = \max_{y \in B_r(x) \setminus \{x\}} N_y$.

We now prove that the entries of the top eigenvector decay exponentially with the distance from the root for any tree where the root is of much higher degree than all other vertices.

Lemma 5.2. For any finite tree with a fixed root vertex x of degree α and all other vertices of degree at most $N_{y^*} \ll \alpha$, for any vertex u such that $u \sim v$, $u \leq v$, it holds that for the largest eigenvalue λ with eigenvector $\mathbf{w}$,

$$\mathbf{w}_+|_u = \left(1 + O\left(\frac{N_{y^*}}{\lambda^2}\right)\right) \lambda \mathbf{w}_+|_v. \tag{5.7}$$

Therefore by Lemma 5.1 and Lemma 3.2, we have the following. If $x \in \mathcal{V}$, then for $u, v \in B_r(x)$ such that $u \sim v$ and $u \leq v$,

$$\mathbf{w}_+|_u = \left(1 + O\left(\mathfrak{u}^{-1/4}\right)\right)\lambda \mathbf{w}_+|_v$$

and if $x \in \mathcal{W}$, then for $u \in B_r(x)$,

$$\mathbf{w}_+|_u = \left(1 + O\left(\mathfrak{u}^{-2/3}\right)\right)\lambda \mathbf{w}_+|_v.$$

Proof. By the eigenvector equation (5.5), we must have $\mathbf{w}_+|_u \leq \lambda \mathbf{w}_+|_v$. For a lower bound on $\mathbf{w}_+|_u$, we proceed by induction on the distance from x, starting from the leaves in $B_r(x)$ (note that these are not necessarily leaves in G). Any leaf v only has one neighbor u, making this base case trivial, as there is only one neighbor and $\mathbf{w}_+|_u = \lambda \mathbf{w}_+|_v$.

Now, assume (5.7) is true for all $z \geq u$. Then, applying (5.5) to v, we get

$$\begin{aligned}
\mathbf{w}_+|_u &= \left(\lambda - \sum_{y \sim v, y \geq v} \frac{\mathbf{w}_+|_y}{\mathbf{w}_+|_v}\right)\mathbf{w}_+|_v \\
&\geq \left(\lambda - \frac{N_{y^*}}{\lambda - O\left(\frac{N_{y^*}}{\lambda}\right)}\right)\mathbf{w}_+|_v \\
&\geq \left(1 - O\left(\frac{N_{y^*}}{\lambda^2}\right)\right)\lambda \mathbf{w}_+|_v, \qquad (5.8)
\end{aligned}$$

where we used the inductive hypothesis to get the first inequality and then used that by Lemma 3.5 the rooted trees around λ satisfy $N_{y*} \ll \lambda^2$. □

Such a tight bound implies exponential decay on various levels. We can now bound the error from approximating these eigenvectors using the truncation.

Proposition 5.1. We define

$$\Lambda := \sum_{x \in \mathcal{V}} \left(\lambda_x \mathbf{w}_+(x)\mathbf{w}_+(x)^* - \lambda_x \mathbf{w}_-(x)\mathbf{w}_-(x)^*\right).$$

Then

$$\max_{x \in \mathcal{V}, \sigma \in \{\pm 1\}} \left\|(A - \Lambda)\mathbf{w}_\sigma(x)\right\| = O\left(\left(d^{r/2} + 1\right)\mathfrak{u}^{-(r-1)/2}\right).$$

Proof. For any $x \in \mathcal{V}, \sigma \in \{\pm 1\}$, $\mathbf{w}_\sigma(x)$ satisfies $(A - \Lambda)\mathbf{w}_\sigma(x) = (A - A_{B_r(x)})\mathbf{w}_\sigma(x)$. The only nonzero entries of this vector are supported on $S_{r+1}(x)$. This holds because the only rows of $(A - A_{B_r(x)})$ that have non-zero entries corresponding to $B_r(x)$ are vertices in $S_{r+1}(x)$, since $B_{r+1}(x)$ are disjoint trees by $\Omega_{3.2}$. By Lemma 5.1 and Lemma 5.2, each entry in $\mathbf{w}_\sigma(x)$ corresponding to vertices in $S_r(x)$ has value at most $(1 + o_N(1))\mathfrak{u}^{-r/2}$. Moreover, under $\Omega_{3.2}$, the number of vertices in the $r+1$ level is $d^r \alpha_x + O(d^{r-1/2} + 1)\mathfrak{u}^{\frac{7}{8}} \leq O((d^r + 1)\mathfrak{u})$. Therefore, $\|(A - \Lambda)\mathbf{w}_{\sigma(x)}\| = O((d^{r/2} + 1)\mathfrak{u}^{-r/2+1/2})$. □

This error term gives us sufficient information about $\mathcal{V} \backslash \mathcal{W}$, and we can now focus only on the fine regime $\mathcal{W}$. The following proposition gives bounds on the total mass the eigenvector assigns to each sphere.

Proposition 5.2. For all $x \in \mathcal{W}$, the eigenvector $\mathbf{w}_+$ satisfies

$$\mathbf{w}_+|_x = \frac{1}{\sqrt{2}} + O\left(\left(1 + d^{-1/2}\right) \mathfrak{u}^{-1/3}\right). \tag{5.9}$$

and for $1 \leq i \leq r$,

$$\|\mathbf{w}_+|_{S_i}\| = \left(\frac{d}{\alpha}\right)^{(i-1)/2} \frac{1}{\sqrt{2}} \left(1 + O\left(\left(1 + d^{-1/2} + d^{-i+1}\right) \mathfrak{u}^{-1/3}\right)\right) \tag{5.10}$$

and

$$\|\mathbf{w}_+|_{[N]\backslash B_i}\| = \sqrt{\frac{1}{1 - \frac{d}{\alpha}}} \left(\frac{d}{\alpha}\right)^{i/2} \frac{1}{\sqrt{2}} \left(1 + O\left(\left(1 + d^{-1/2} + d^{-r+1}\right) \mathfrak{u}^{-1/3}\right)\right). \tag{5.11}$$

Proof. In this proof we repeatedly use the approximation of λ from Lemma 5.1 in order to replace λ by $\alpha^{\frac{1}{2}}$ or vice-versa up to some small multiplicative error.

First note that Lemma 5.2 implies that for each $v \in S_i$,

$$\mathbf{w}_+|_v = \left(1 + O\left(\mathfrak{u}^{-2/3}\right)\right) \lambda^{-i} (\mathbf{w}_+|_x).$$

Therefore, as we know by Lemma 3.5 that $|S_i| = d^{i-1}\alpha + O\left(\left(d^{i-3/2} + 1\right) \mathfrak{u}^{2/3}\right)$,

$$\|\mathbf{w}_+|_{S_i}\| \;=\; d^{(i-1)/2} \alpha^{1/2} \left(1 + O\left(\left(d^{-1/2} + d^{-i+1}\right) \mathfrak{u}^{-1/3}\right)\right) \lambda^{-i} \mathbf{w}_+|_x. \tag{5.12}$$

Note that for this approximation to hold and for the error term to go to 0, we require $d \geq \mathfrak{u}^{-\frac{1}{3r}}$ The complicated error term is necessary as different terms could maximize for different regimes of d.

Specifically, this means that for the normalized vector $\mathbf{w}_+$,

$$\begin{aligned} (\mathbf{w}_+|_x)^2 + \|\mathbf{w}_+|_{S_1}\|^2 \;&\geq\; 1 - O\left(\frac{d}{\mathfrak{u}} \left(1 + \left(d^{-1/2} + d^{-1}\right)\right) \mathfrak{u}^{-1/3})\right) \\ &\geq\; 1 - O\left(\frac{d}{\mathfrak{u}}\right). \end{aligned}$$

Using this together with (5.12) for $\mathbf{w}_+|_{S_1}$ and solving for $\mathbf{w}_+|_x$ gives

$$\begin{aligned} (\mathbf{w}_+|_x)^2 &= \frac{1}{2} + O\left(\left(1 + d^{-1/2}\right) \mathfrak{u}^{-1/3} + d\mathfrak{u}^{-1}\right) \\ \mathbf{w}_+|_x &= \frac{1}{\sqrt{2}} + O\left(\left(1 + d^{-1/2}\right) \mathfrak{u}^{-1/3}\right). \end{aligned}$$

Combining (5.12) and (5.9),

$$\begin{aligned}\|\mathbf{w}_+|_{S_i}\| &= d^{(i-1)/2}\alpha^{1/2}\left(1+O\left(\left(d^{-1/2}+d^{-i+1}\right)\mathfrak{u}^{-1/3}\right)\right)\lambda^{-i}\left(\frac{1}{\sqrt{2}}+O\left(\left(1+d^{-1/2}\right)\mathfrak{u}^{-1/3}\right)\right)\\ &= d^{(i-1)/2}\alpha^{-(i-1)/2}\frac{1}{\sqrt{2}}\left(1+O\left(\left(1+d^{-1/2}+d^{-i+1}\right)\mathfrak{u}^{-1/3}\right)\right).\end{aligned}$$

Similarly, for (5.11), we have by (5.12),

$$\begin{aligned}\left\|\mathbf{w}_+|_{[N]\setminus B_i}\right\|^2 &= \sum_{j=i+1}^{r}\|\mathbf{w}_+|_{S_j}\|^2\\ &= \frac{1}{1-\frac{d}{\alpha}}d^i\alpha^{-i}\frac{1}{2}\left(1+O\left(\left(1+d^{-1/2}+d^{-r+1}\right)\mathfrak{u}^{-1/3}\right)\right).\end{aligned}$$

□

Note that Proposition 5.2 implies that almost all mass of the vector $\mathbf{w}_+$ is on x and S_1.

Eigenvalue structure

Along with the eigenvector, we further analyze the eigenvalue. To do this, we expand (5.6) as an infinite sum which we get by repeatedly using the expansion $\frac{1}{\lambda-q}=\sum_{k=0}^{\infty}\frac{q^k}{\lambda^{k+1}}$.

We will bound the eigenvalue λ through the moments of the degree sequence surrounding $x\in\mathcal{W}$. We recall the definitions of $\beta^{(2)}$ and $\beta^{(1,1)}$ from Definition 3.1. Note that $\beta^{(1,1)}=|S_3|$ and $\beta^{(2)}=\sum_{y\sim x}N_y^2$.

Lemma 5.3. Under $\Omega_{3.2}$, for any $x\in\mathcal{W}$,

$$\lambda^2=\alpha+\lambda^{-2}\beta+\lambda^{-4}\left(\beta^{(2)}+\beta^{(1,1)}\right)+O\left(\left(d^2+d\right)\mathfrak{u}^{-5/3}\right). \tag{5.13}$$

Proof. We can rewrite (5.6) as

$$\lambda^2=\sum_{y_1\sim x}\frac{1}{1-\frac{1}{\lambda^2}\sum_{y_2\sim y_1,y_2\geq y_1}\frac{1}{1-\frac{1}{\lambda^2}\sum_{y_3\sim y_2,y_3\geq y_2}\frac{1}{1-\frac{1}{\lambda^2}\sum_{y_4\sim y_3,y_4\geq y_3}\cdots}}}.$$

and expand this as

$$\lambda^2=\sum_{y_1\sim x}\sum_{k_1=0}^{\infty}\left(\frac{1}{\lambda^2}\sum_{y_2\sim y_1,y_2\geq y_1}\frac{1}{1-\frac{1}{\lambda^2}\sum_{y_3\sim y_2,y_3\geq y_2}\frac{1}{1-\frac{1}{\lambda^2}\sum_{y_4\sim y_3,y_4\geq y_3}\cdots}}\right)^{k_1}.$$

and once again as

$$\lambda^2 = \sum_{y_1 \sim x} \sum_{k_1=0}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1} \sum_{k_2=0}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} \frac{1}{1 - \frac{1}{\lambda^2} \sum_{y_4 \sim y_3, y_4 \geq y_3} \cdots} \right)^{k_2} \right)^{k_1}. \tag{5.14}$$

Of course we could repeat this process, but this is sufficient accuracy for our purposes.

We do the same analysis as before, starting at the innermost level, corresponding to the leaves, and then inducting our way back up the tree. For any vertex v, we can write a recursive equation by defining

$$F(v) := \frac{1}{1 - \frac{1}{\lambda^2} \sum_{y \sim v, y \geq v} F(y)}$$

which gives that

$$\lambda^2 = \sum_{y_1 \sim x} \sum_{k_1=0}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1} \sum_{k_2=0}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} F(y_3) \right)^{k_2} \right)^{k_1}.$$

We now estimate F_v for $v \in B_r(x)$. For any leaf v, as there are no $y \geq v$, $F(v) = 1$. For the rest, we use induction. Recall that N_y^* is the maximum degree in $B_r(x) \backslash \{x\}$, and that $N_{y^*} \ll \lambda^2$ under the event $\Omega_{3.2}$. Assume that for all $y \sim v, y \geq v$, $F(y) = 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right)$. Then

$$F(v) = \frac{1}{1 - \frac{1}{\lambda^2} \sum_{y \sim v, y \geq v} F(y)} \leq \frac{1}{1 - \frac{N_{y^*}}{\lambda^2}\left(1 + O\left(\frac{N_{y^*}}{\lambda^2}\right)\right)} = 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right).$$

Therefore (5.14) becomes

$$\lambda^2 = \sum_{y_1 \sim x} \sum_{k_1=0}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1} \sum_{k_2=0}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right) \right)^{k_2} \right)^{k_1}.$$

We want to show this expansion relies only on the terms in Definition 3.1. Therefore we rewrite the first few terms as

$$\begin{aligned}
\alpha &= \sum_{y_1 \sim x} 1 \\
\lambda^{-2}\beta &= \sum_{y_1 \sim x} \left(\frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right) \right)^0 \right)^1 \\
\lambda^{-4}\beta^{(1,1)} &= \sum_{y_1 \sim x} \left(\frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} 1 \right)^1 \right)^1.
\end{aligned}$$

The contribution of $\beta^{(2)}$ is more complicated, but as $\beta^{(2)} = \sum_{y_1 \sim x} N_y^2$, we can write

$$\lambda^{-4}\beta^{(2)} = \sum_{y_1 \sim x} \left(\frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right)\right)^0 \right)^2$$

meaning we have

$$\begin{aligned}
&\sum_{y_1 \sim x} \left(\frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1} \sum_{k_2=0}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right)\right)^{k_2} \right)^2 - \lambda^{-4}\beta^{(2)} \\
&\quad = \sum_{y_1 \sim x} \left(\frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1} \sum_{k_2=1}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right)\right)^{k_2} \right)^2 \\
&\qquad + 2\lambda^{-4} \sum_{y_1 \sim x} N_y \sum_{y_2 \sim y_1, y_2 \geq y_1} \sum_{k_2=1}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right)\right)^{k_2}.
\end{aligned}$$

Therefore subtracting these terms from (5.14) gives

$$\begin{aligned}
&\lambda^2 - \alpha - \lambda^{-2}\beta - \lambda^{-4}\beta^{(2)} - \lambda^{-4}\beta^{(1,1)} \\
&= \sum_{y_1 \sim x} \sum_{k_1=3}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1} \sum_{k_2=0}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right)\right)^{k_2} \right)^{k_1} \\
&\quad + \sum_{y_1 \sim x} \frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1} \sum_{k_2=2}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right)\right)^{k_2} \\
&\quad + \sum_{y_1 \sim x} \frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1} \frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} O\left(\frac{N_{y^*}}{\lambda^2}\right) \\
&\quad + \sum_{y_1 \sim x} \left(\frac{1}{\lambda^2} \sum_{y_2 \sim y_1, y_2 \geq y_1} \sum_{k_2=1}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right)\right)^{k_2} \right)^2 \\
&\quad + 2\lambda^{-4} \sum_{y_1 \sim x} N_y \sum_{y_2 \sim y_1, y_2 \geq y_1} \sum_{k_2=1}^{\infty} \left(\frac{1}{\lambda^2} \sum_{y_3 \sim y_2, y_3 \geq y_2} 1 + O\left(\frac{N_{y^*}}{\lambda^2}\right)\right)^{k_2}.
\end{aligned}$$

Each of our error terms now has coefficient λ^{-6} or smaller. Therefore, by the same exponential decay and the fact that $\lambda^{-6} \leq \alpha^{-3}$ and $N_{y^*} \ll \lambda^2$, each of the five sums can be written as

$$O\left(\alpha^{-3} N_{y^*} \left(\beta^{(1,1)} + \beta^{(2)}\right)\right). \tag{5.15}$$

By the assumptions of $\Omega_{3.2}$, namely the bound on the maximal degree and the approximations of $\beta^{(1,1)}$ and $\beta^{(2)}$, as well as our bounds for d, for $x \in \mathcal{W}$

$$\begin{aligned} O\left(\alpha^{-3} N_{y^*}\left(\beta^{(1,1)}+\beta^{(2)}\right)\right) &= O\left(\mathfrak{u}^{-3}\mathfrak{u}^{1/3}\left(\beta^{(2)}+\beta^{(1,1)}\right)\right) \\ &= O\left(\mathfrak{u}^{-8/3}\left(\left(d^2+d\right)\mathfrak{u}+\mathfrak{u}^{2/3}+d^2\mathfrak{u}+\left(d^{3/2}+1\right)\mathfrak{u}^{2/3}\right)\right) \\ &= O\left(\left(d^2+d\right)\mathfrak{u}^{-5/3}\right). \end{aligned}$$

□

We now solve for λ in this equation. Note that while $\lambda^2 = \alpha(1+o(1))$ by Lemma 5.1, this by itself does not give a precise enough approximation. Instead we bootstrap this initial approximation using the structural properties from $\Omega_{3.2}$.

Lemma 5.4. For $x \in \mathcal{W}$,

$$\lambda^2 = \alpha + \frac{\beta}{\alpha} + \frac{\beta^{(1,1)}+\beta^{(2)}}{\alpha^2} - \frac{(\beta)^2}{\alpha^3} + O\left(\left(d^2+d\right)\mathfrak{u}^{-5/3}\right). \tag{5.16}$$

Proof. We rewrite (5.13) as

$$\lambda^2 = \alpha + \frac{1}{\alpha+(\lambda^2-\alpha)}\beta + \frac{1}{(\alpha+(\lambda^2-\alpha))^2}\beta^{(1,1)} + \frac{1}{(\alpha+(\lambda^2-\alpha))^2}\beta^{(2)} + O\left(\left(d^2+d\right)\mathfrak{u}^{-5/3}\right).$$

Moreover, by Lemma 5.1, $\lambda^2-\alpha = O(d)$, so we can approximate

$$\frac{1}{\alpha+(\lambda^2-\alpha)} = \frac{1}{\alpha}\frac{1}{1+\frac{(\lambda^2-\alpha)}{\alpha}} = \frac{1}{\alpha}\left(1+O\left(\frac{d}{\mathfrak{u}}\right)\right).$$

Therefore

$$\begin{aligned} \lambda^2 &= \alpha + \frac{1}{\alpha}\beta\left(1+O\left(\frac{d}{\mathfrak{u}}\right)\right) + \frac{1}{\alpha^2}\left(\beta^{(1,1)}+\beta^{(2)}\right)\left(1+O\left(\frac{d}{\mathfrak{u}}\right)\right) + O\left(\left(d^2+d\right)\mathfrak{u}^{-5/3}\right) \\ &= \alpha + \frac{\beta}{\alpha} + O\left(\frac{d^2+d}{\mathfrak{u}}\right). \end{aligned}$$

This implies that $\frac{1}{\lambda^2} = \frac{1}{\alpha+\frac{\beta}{\alpha}+\left(\lambda^2-\alpha+\frac{\beta}{\alpha}\right)} = \frac{1}{\alpha-\frac{\beta}{\alpha}}\left(1+O\left(\frac{d^2+d}{\mathfrak{u}^2}\right)\right)$. Plugging this more precise approximation for λ^2 once again into (5.13) we get

$$\begin{aligned} \lambda^2 = \alpha &+ \frac{1}{\alpha+\frac{\beta}{\alpha}}\beta\left(1+O\left(\frac{d^2+d}{\mathfrak{u}^2}\right)\right) + \frac{1}{(\alpha+\frac{\beta}{\alpha})^2}(\beta^{(1,1)}+\beta^{(2)})\left(1+O\left(\frac{d^2+d}{\mathfrak{u}^2}\right)\right) \\ &+ O\left(\left(d^2+d\right)\mathfrak{u}^{-5/3}\right) \end{aligned}$$

$$= \alpha + \frac{1}{\alpha + \frac{\beta}{\alpha}}\beta + \frac{1}{\left(\alpha + \frac{\beta}{\alpha}\right)^2}\left(\beta^{(1,1)} + \beta^{(2)}\right) + O\left(\left(d^2 + d\right)\mathfrak{u}^{-5/3}\right).$$

Due to our bounds on d from 5.1, we can expand $\frac{1}{\alpha+\frac{\beta}{\alpha}} = \frac{1}{\alpha} - \frac{\beta}{\alpha^3} + O\left(\frac{d^2}{\mathfrak{u}^3}\right)$, which gives us

$$\lambda^2 \;=\; \alpha + \frac{\beta}{\alpha} + \frac{\beta^{(1,1)} + \beta^{(2)}}{\alpha^2} - \frac{(\beta)^2}{\alpha^3} + O\left(\left(d^2 + d\right)\mathfrak{u}^{-5/3}\right)$$

as desired. □

Given this much tighter approximation of λ_x for $x \in \mathcal{W}$, we can now show that the order of the eigenvalues corresponding to $B_r(x)$ for $x \in \mathcal{W}$ is the same as the lexicographic order of α_x and β_x.

Lemma 5.5. For two vertices $u, v \in \mathcal{W}$, if $\alpha_u \geq \alpha_v$, $\lambda_u^2 - \lambda_v^2 \geq \alpha_u - \alpha_v + O\left(\left(d^{1/2} + 1\right)\mathfrak{u}^{-1/3}\right)$. Moreover, if $\alpha_u = \alpha_v$, and $\beta_u \geq \beta_v$, then $\lambda_u^2 - \lambda_v^2 \geq \frac{\beta_u - \beta_v}{\mathfrak{u}} + O\left(\left(1 + d^{3/2}\right)\mathfrak{u}^{-4/3}\right)$. Therefore, for $x \in \mathcal{W}$, λ_x^2 are ordered according to the lexicographic ordering of (α_x, β_x).

Proof. By (5.16),

$$\lambda_u^2 - \lambda_v^2 \geq \alpha_u - \alpha_v - \left|\frac{\beta_u}{\alpha_u} - \frac{\beta_v}{\alpha_v}\right| + \left|\frac{\beta_u^{(1,1)}}{\alpha_u^2} - \frac{\beta_v^{(1,1)}}{\alpha_v^2}\right| + \left|\frac{\beta_u^{(2)}}{\alpha_u^2} - \frac{\beta_v^{(2)}}{\alpha_v^2}\right| + \left|\frac{(\beta_u)^2}{\alpha_u^3} - \frac{(\beta_v)^2}{\alpha_v^3}\right| + O\left(\left(d^2 + d\right)\mathfrak{u}^{-5/3}\right).$$

Using the definition of $\Omega_{3.2}$ (3-5), namely the concentration of $\beta, \beta^{(1,1)}$ and $\beta^{(2)}$, this implies

$$\lambda_u^2 - \lambda_v^2 \geq \alpha_u - \alpha_v + O\left(\left(d^{1/2} + 1\right)\mathfrak{u}^{-1/3}\right).$$

Now, assume that $\alpha_u = \alpha_v$. Then by (5.16),

$$\lambda_u^2 - \lambda_v^2 \;\geq\; \frac{\beta_u - \beta_v}{\alpha_u} - \left|\frac{\beta_u^{(1,1)}}{\alpha_u^2} - \frac{\beta_v^{(1,1)}}{\alpha_v^2}\right| + \left|\frac{\beta_u^{(2)}}{\alpha_u^2} - \frac{\beta_v^{(2)}}{\alpha_v^2}\right| + \left|\frac{(\beta_u)^2}{\alpha_u^3} - \frac{(\beta_v)^2}{\alpha_v^3}\right| + O\left(\left(d^2 + d\right)\mathfrak{u}^{-5/3}\right).$$

Once again, by the definition of $\Omega_{3.2}$ (3-5), this implies

$$\lambda_u^2 - \lambda_v^2 \;\geq\; \frac{\beta_u - \beta_v}{\mathfrak{u}} + O\left(\left(d^{3/2} + 1\right)\mathfrak{u}^{-4/3}\right).$$

To see that this induces a lexicographic ordering, if $\alpha_u \neq \alpha_v$, then $\alpha_u - \alpha_v \geq 1 \gg (d^{1/2}+1)\mathfrak{u}^{-1/3}$. Similarly, if $\alpha_u = \alpha_v$, but $\beta_u \neq \beta_v$, then $\frac{\beta_u - \beta_v}{\mathfrak{u}} \geq \frac{1}{\mathfrak{u}} \gg (1 + d^{3/2})\mathfrak{u}^{-4/3}$, by our assumptions on d from Definition 5.1. □

5.3 Rough Regime

In this section we construct approximate eigenvectors corresponding to small balls around vertices in $\mathcal{U}\backslash\mathcal{V}$, and we derive a less precise approximation for the eigenvalues corresponding to those balls. The approach used in this section is very similar to the approach in Section 6.4 of [ADK23b]. Note a result like Lemma 5.1, which we will eventually use to show that eigenvalues from vertices in $\mathcal{V}\backslash\mathcal{W}$ cannot compete with the largest ones from vertices in $\mathcal{W}$, cannot directly be derived for all vertices in $\mathcal{U}$. The main obstructions are that the growth of the spheres and the maximum degree in the balls cannot be bounded as tightly as for vertices in $\mathcal{V}$.

More precisely our goal for this section is to show the following.

Lemma 5.6. For $x \in \mathcal{U}$, we can create a set of unit vectors $\mathbf{w}_\sigma(x)$ with $\sigma \in \{\pm 1\}$ such that

1. For $u, v \in \mathcal{U}, u \neq v$, and $\sigma_1, \sigma_2 \in \{\pm 1\}$ we have $supp(\mathbf{w}_{\sigma_1}(u)) \cap supp(\mathbf{w}_{\sigma_2}(v)) = \emptyset$.

2. We have $\left\|A\mathbf{w}_\sigma(x) - \sigma\lambda_x \mathbf{w}_\sigma(x)\right\| = O(\log\log N)$, where $\lambda_x = \sqrt{\alpha_x + \frac{\beta_x}{\alpha_x}}$.

Proof. Define $\hat{\alpha}_x, \hat{\beta}_x$ to be the parameters of the pruned graph $\hat{G}$. We use the same test vector as [ADK23b], restated with our parameters this is the unit vector

$$\mathbf{w}_\sigma(x) := \frac{1}{\sqrt{2}}\left(\frac{\sqrt{\hat{\alpha}_x}}{\sqrt{\hat{\alpha}_x + \frac{\hat{\beta}_x}{\hat{\alpha}_x}}}\mathbf{1}_x + \sigma\frac{1}{\sqrt{\hat{\alpha}_x}}\mathbf{1}_{\hat{S}_1(x)} + \frac{1}{\sqrt{\hat{\alpha}_x(\hat{\alpha}_x + \frac{\hat{\beta}_x}{\hat{\alpha}_x})}}\mathbf{1}_{\hat{S}_2(x)}\right). \tag{5.17}$$

The first statement of the Lemma now follows by Lemma 3.8.

We now fix x and σ and drop them from our notation for better readability. To prove the second statement we define $\hat{\lambda} = \sqrt{\hat{\alpha} + \frac{\hat{\beta}}{\hat{\alpha}}}$ and $\hat{A} = A_{\hat{G}}$ and use a triangle inequality to bound

$$\|A\mathbf{w} - \sigma\lambda\mathbf{w}\| \leq \left\|A\mathbf{w} - \hat{A}\mathbf{w}\right\| + \left\|\hat{A}\mathbf{w} - \sigma\hat{\lambda}\mathbf{w}\right\| + \left\|\hat{\lambda}\mathbf{w} - \lambda\mathbf{w}\right\|.$$

The first term on the right hand side is at most constant, as by Lemma 3.8 the maximum degree of $G - \hat{G}$ is bounded by a constant and since the maximum row sum is an upper bound for the maximum eigenvalue of a positive symmetric matrix.

Similarly the last term is bounded since by Lemma 3.8, $\hat{\alpha}$ differs from α by at most a constant, and β and $\hat{beta}$ can both be bounded by $(d + \log\log N)u$. This implies that

$$|\lambda - \hat{\lambda}| = \sqrt{\alpha + \frac{\beta}{\alpha}} - \sqrt{\hat{\alpha} + \frac{\hat{\beta}}{\hat{\alpha}}} = \sqrt{\alpha}\sqrt{1 + \frac{\beta}{\alpha^2}} - \sqrt{\alpha}\sqrt{1 + \frac{\hat{\alpha} - \alpha}{\alpha} + \frac{\hat{\beta}}{\alpha\hat{\alpha}}} \ll 1.$$

The second term can be computed as

$$\sqrt{2}\left(\hat{A}\mathbf{w} - \sigma\hat{\lambda}\mathbf{w}\right)$$

$$= \left(\sigma\sqrt{\hat{\alpha}} - \sigma\sqrt{\hat{\alpha}}\right)\mathbf{1}_x + \sum_{y\in\hat{S}_1}\left(\frac{\sqrt{\hat{\alpha}}}{\hat{\lambda}} + \frac{\hat{N}_y}{\sqrt{\hat{\alpha}}\hat{\lambda}} - \frac{\hat{\lambda}}{\sqrt{\hat{\alpha}}}\right)\mathbf{1}_y + \sum_{y\in\hat{S}_2}\left(\frac{\sigma}{\sqrt{\hat{\alpha}}} - \frac{\sigma}{\sqrt{\hat{\alpha}}}\right)\mathbf{1}_y + \sum_{y\in\hat{S}_3}\frac{1}{\sqrt{\hat{\alpha}}\hat{\lambda}}\mathbf{1}_y.$$

Therefore, using Lemma 3.8 and the lower bound on α_x for $x \in \mathcal{U}$, we get

$$\begin{aligned} 2\left\|\hat{A}\mathbf{w} - \sigma\hat{\lambda}\mathbf{w}\right\|^2 &\leq \frac{1}{\hat{\alpha}\left(\hat{\alpha} + \frac{\hat{\beta}}{\hat{\alpha}}\right)}\sum_{y\in\hat{S}_1}\left(\hat{N}_y - \frac{\hat{\beta}}{\hat{\alpha}}\right)^2 + \left|\hat{S}_3\right|\frac{1}{\hat{\alpha}\left(\hat{\alpha} + \frac{\hat{\beta}}{\hat{\alpha}}\right)} \\ &\leq O\left((\log\log N)^2 + \frac{(d + \log\log N)^2}{\mathfrak{u}}\right). \end{aligned}$$

Putting these three bounds together and using our expression (3.2) for $\mathfrak{u}$ and the bounds in 5.1 for d, we get that,

$$\|A\mathbf{w} - \sigma\lambda\mathbf{w}\| = O(\log\log N).$$

□

5.4 Block Decomposition

This section is devoted to proving Theorem 5.5, for which we use results from Sections 5.2 and 5.3. For this we first bound the contribution of the remainder of the matrix, i.e. from vectors that are orthogonal to the largest eigenvectors of small balls around the high-degree vertices. We end this section by using the approximate diagonalization to prove Theorem 5.2.

Bulk vectors

We now prove that there is no contribution from any other vector. To do this, we use the decomposition of Krivelevich and Sudakov [KS03]. This lets us reduce to only considering stars around high degree vertices. Here, we state a structure theorem that combines elements of the proof of Theorem 1.1 and Lemma 2.2 in [KS03]. To do this, recall that Γ_x are all vertices adjacent to x and consider the sets of vertices

$$\begin{aligned} \mathcal{Y}_1 &:= \left\{x \in [N] : \alpha_x \geq \mathfrak{u}^{3/4}\right\} \\ \mathcal{Y}_2 &:= \left\{x \in [N] : \Gamma_x \cap \mathcal{Y}_1 \neq \varnothing\right\} \end{aligned}$$

Proposition 5.3 ([KS03])**.** For $G \sim \mathcal{G}(N, \frac{d}{N})$ graph, if $d = o(\log^{1/2} N)$, then with high probability, there is a subgraph $\mathcal{H} \subset G$ such that

1. $\mathcal{H}$ is contained in the bipartite subgraph induced by $\mathcal{Y}_1$ and $\mathcal{Y}_2$,
2. $\mathcal{H}$ is a union of stars on disjoint vertices,

3. $\|A_{G\setminus\mathcal{H}}\| = O\left(d + \mathfrak{u}^{7/16}\right)$.

Note that $\mathcal{H}$ is the graph $G_6 - H$ in [KS03]. This is strong enough to show that no other vector interferes in the largest eigenvalues.

Define $U_{\mathcal{U}}$ to be the space spanned by $\mathbf{w}_\pm(x)$ as defined in Definition 5.2 for $x \in \mathcal{V}$, and $\mathbf{w}_\pm(x)$ as defined in Lemma 5.6 for $x \in \mathcal{U}\setminus\mathcal{V}$.

Lemma 5.7. For any vector $\mathbf{v} \in \mathbb{R}^N$ that satisfies $\|v| = 1$ and $\mathbf{v} \perp U_{\mathcal{U}}$,

$$\langle \mathbf{v}, A\mathbf{v}\rangle \leq \left(1 + o_N(1)\right)\frac{1}{\sqrt{2}}\sqrt{\mathfrak{u}}. \tag{5.18}$$

Proof. By Proposition 5.3, we have that

$$\langle \mathbf{v}, A\mathbf{v}\rangle \leq \max_{x\in\mathcal{Y}_1}\langle \mathbf{v}, A_{\tilde{B}_1(x)}\mathbf{v}\rangle + O\left(d + \mathfrak{u}^{7/16}\right)$$

where $\tilde{B}_1(x)$ is the ball of radius 1 around x in $\mathcal{H}$ from Proposition 5.3 and $A_{\tilde{B}_1(x)}$ is the adjacency matrix of the graph on the vertices $[N]$ induced by the ball $\tilde{B}_1(x)$.

Therefore we split into cases based on the degree of x. For $x \in \mathcal{V}$, we know that $\mathbf{v} \perp \frac{1}{\sqrt{2}}\left(\mathbf{w}_+(x) + \mathbf{w}_-(x)\right)$. Therefore

$$\langle \mathbf{v}, \mathbf{1}_x\rangle = \langle \mathbf{v}, \frac{1}{\sqrt{2}}\left(\mathbf{w}_+(x)+\mathbf{w}_-(x)\right)\rangle + \langle \mathbf{v}, \mathbf{1}_x - \frac{1}{\sqrt{2}}\left(\mathbf{w}_+(x)+\mathbf{w}_-(x)\right)\rangle \leq \left\|\mathbf{1}_x - \frac{1}{\sqrt{2}}\left(\mathbf{w}_+(x) + \mathbf{w}_-(x)\right)\right\|.$$

By Proposition 5.2,

$$\left\|\mathbf{1}_x - \frac{1}{\sqrt{2}}\left(\mathbf{w}_+(x) + \mathbf{w}_-(x)\right)\right\| = O\left(\mathfrak{u}^{-1/3}\right).$$

We then have

$$\langle \mathbf{v}, A_{\tilde{B}_1(x)}\mathbf{v}\rangle \leq 2\left|\langle \mathbf{v}, \mathbf{1}_x\rangle\right|\sum_{y\sim x}\left|\langle \mathbf{v}, \mathbf{1}_y\rangle\right| = O\left(\mathfrak{u}^{-1/3}\cdot\sqrt{\mathfrak{u}}\right).$$

Similarly, for $x \in \mathcal{U}\setminus\mathcal{V}$, we have $\langle \mathbf{v}, \mathbf{1}_x\rangle \leq \left\|\mathbf{1}_x - \frac{1}{\sqrt{2}}(\mathbf{w}_+(x) + \mathbf{w}_-(x))\right\|$. By the definition of the eigenvector in (5.17), and using properties from Lemma 3.8, we have that

$$\left\|\mathbf{1}_x - \frac{1}{\sqrt{2}}(\mathbf{w}_+(x) + \mathbf{w}_-(x))\right\| = O\left(\sqrt{\left[\frac{\sqrt{\hat{\alpha} + \frac{\hat{\beta}}{\hat{\alpha}}} - \sqrt{\hat{\alpha}}}{\sqrt{\hat{\alpha} + \frac{\hat{\beta}}{\hat{\alpha}}}}\right]^2 + \hat{\beta}\frac{1}{\hat{\alpha}\left(\hat{\alpha} + \frac{\hat{\beta}}{\hat{\alpha}}\right)}}\right)$$
$$= O\left(\sqrt{\frac{\beta^2}{\alpha^3} + \frac{\beta}{\alpha^2}}\right) = O\left(\frac{d + \log\log N}{\mathfrak{u}}\right).$$

Therefore, by the same argument as before

$$\langle \mathbf{v}, A_{\tilde{B}_1(x)}\mathbf{v}\rangle = O\left(\frac{d+\log\log N}{\sqrt{\mathfrak{u}}}\right) = o_N(1)\,.$$

For any vertex $x \in \mathcal{Y}_1 \backslash \mathcal{U}$, the maximum degree is $\mathfrak{u}/2$ and the spectral norm is given by the spectral radius of a star graph, namely for any vector $\mathbf{v}$ such that $\|\mathbf{v}\| = 1$,

$$\langle \mathbf{v}, A_{\tilde{B}_1(x)}\mathbf{v}\rangle \leq \sqrt{\frac{\mathfrak{u}}{2}}.$$

Combining these cases gives the result. □

Proof of the Structure Theorem

We now have all the ingredients to prove the structure theorem.

Proof of Theorem 5.5. We can now fully define the block decomposition from (5.3),

$$A = U\begin{pmatrix} D_{\mathcal{W}} & 0 & 0 & E^*_{\mathcal{W}} \\ 0 & D_{\mathcal{V}\backslash\mathcal{W}} & 0 & E^*_{\mathcal{V}\backslash\mathcal{W}} \\ 0 & 0 & D_{\mathcal{U}\backslash\mathcal{V}} + \mathcal{E}_{\mathcal{U}\backslash\mathcal{V}} & E^*_{\mathcal{U}\backslash\mathcal{V}} \\ E_{\mathcal{W}} & E_{\mathcal{V}\backslash\mathcal{W}} & E_{\mathcal{U}\backslash\mathcal{V}} & \mathcal{X} \end{pmatrix} U^*$$

We first define the unitary matrix U. We set the first $2|\mathcal{W}|$ columns of U to vectors $\mathbf{w}_\pm(x)$ for $x \in \mathcal{W}$, and denote this part of the matrix by $U_{\mathcal{W}}$, then we set the next $2|\mathcal{V}\backslash\mathcal{W}|$ columns to $\mathbf{w}_\pm(x)$ for $x \in \mathcal{V}\backslash\mathcal{W}$ and denote this part of the matrix by $U_{\mathcal{V}\backslash\mathcal{W}}$, for $\mathbf{w}_\pm(x)$ as defined in Definition 5.2. The next $2|\mathcal{U}\backslash\mathcal{V}|$ columns are the vectors $\mathbf{w}_\pm(x)$ for $x \in \mathcal{U}\backslash\mathcal{V}$ as defined in Lemma 5.6, and we denote this part of the matrix by $U_{\mathcal{U}\backslash\mathcal{V}}$. We denote these three parts of the matrix together by $U_{\mathcal{U}}$. We then complete U arbitrarily with a basis of the rest of $\mathbb{R}^N$, namely $U_{\mathcal{U}^\perp} \subset \mathbb{R}^N$.

It is implied by the definition of U that the diagonal matrices $D_{\mathcal{W}}$ and $D_{\mathcal{V}\backslash\mathcal{W}}$ have entries $\sigma\lambda_x$ on the diagonal, i.e. the eigenvalue of the truncated balls corresponding to each $w_\sigma(x)$. The diagonal operator $D_{\mathcal{U}\backslash\mathcal{V}}$, is defined to have entries $\sigma\lambda_x$, from Lemma 5.6.

0's exist in the requisite places as we can assume by $\Omega_{3.2}$ that balls of vertices in $\mathcal{V}$ are disjoint, and for each $x \in \mathcal{V}$, the maximum degree of a vertex in $B_{r+3}(x)\backslash x$ is $\mathfrak{u}^{3/4}$, implying that there are no intersections with balls of radius 3 around vertices in $\mathcal{U}\backslash\mathcal{V}$. By Lemma 3.4, with high probability, the $e^{\log^{1/8} N}$ vertices of largest degree have degree at least $\mathfrak{u} - 2\log^{1/8} N$. Therefore the eigenvalues corresponding to these vertices have value at least $\sqrt{\mathfrak{u} - 2\log^{1/8} N} = \sqrt{\mathfrak{u}} - O\left(\log^{-3/8} N\right)$ by Lemma 5.1.

To get a bound on $D_{\mathcal{V}\backslash\mathcal{W}}$, we use the upper bound from Lemma 5.1. This gives that for any vertex $x \in \mathcal{V}$, and for the range of d defined in 5.1,

$$\lambda_x \leq \sqrt{\mathfrak{u} - \mathfrak{u}^{\frac{1}{4}} + O(d)} = \sqrt{\mathfrak{u}} - \frac{\mathfrak{u}^{-\frac{1}{2}}}{2} + O\left(\frac{d}{\sqrt{\mathfrak{u}}}\right) \leq \sqrt{\mathfrak{u}} - \Theta\left(\mathfrak{u}^{-\frac{1}{2}}\right).$$

By Lemma 5.1, for any $x \in \mathcal{V}$, $\sigma \in \{\pm 1\}$, $\|(A-\Lambda)\mathbf{w}_\sigma(x)\| = O\left(\left(d^{r/2}+1\right)\mathfrak{u}^{-(r-1)/2}\right)$. We will now show that this implies a bound on $\|E_W\|$: Using that $U^*_{\mathcal{W}}$ is a surjective projection of $\mathbb{R}^N$ onto $\mathbb{R}^{2|\mathcal{W}|}$ and $U_{\mathcal{U}^\perp}$ is an injective embedding of $\mathbb{R}^{N-2|\mathcal{U}|}$ onto $\mathbb{R}^N$, we can transform E_W, which maps $\mathbb{R}^{2|\mathcal{W}|}$ to $\mathbb{R}^{N-2|\mathcal{U}|}$, into an operator from $\mathbb{R}^N$ to $\mathbb{R}^N$, with the same spectral properties. Therefore, using additionally that outside of the choice of $\sigma \in \{\pm 1\}$, the supports of $\mathbf{w}_\sigma(x)$ are independent,

$$\begin{aligned}\left\|E_{\mathcal{W}}\right\| &= \left\|U_{\mathcal{U}^\perp}E_{\mathcal{W}}U^*_{\mathcal{W}}\right\| = \max_{\mathbf{v}\in span(U_{\mathcal{W}}),\|\mathbf{v}\|=1}\left\|U_{\mathcal{U}^\perp}E_{\mathcal{W}}U^*_{\mathcal{W}}\mathbf{v}\right\| \\ &\le \max_{x\in\mathcal{W},\sigma\in\{\pm1\}} 2\left\|U_{\mathcal{U}^\perp}E_{\mathcal{W}}U^*_{\mathcal{W}}\mathbf{w}_\sigma(x)\right\| = \max_{x\in\mathcal{W},\sigma\in\{\pm1\}} 2\left\|\left(A-U_{\mathcal{W}}D_{\mathcal{W}}U^*_{\mathcal{W}}\right)\mathbf{w}_\sigma(x)\right\| \\ &= \max_{x\in\mathcal{W},\sigma\in\{\pm1\}} 2\left\|(A-\Lambda)\,\mathbf{w}_\sigma(x)\right\| = O\left(\left(d^{r/2}+1\right)\mathfrak{u}^{-(r-1)/2}\right).\end{aligned}$$

Here we use the definition of Λ from Lemma 5.1. The same is true for $\|E_{\mathcal{V}\setminus\mathcal{W}}\|$.

Instead of bounding the operator norms of $E_{\mathcal{U}\setminus\mathcal{V}}$ and $\mathcal{E}_{\mathcal{U}\setminus\mathcal{V}}$ individually, we bound the operator norm of their concatenation, which will be an upper bound for both. Similarly to before we can write

$$\left\|\begin{bmatrix}\mathcal{E}_{\mathcal{U}\setminus\mathcal{V}}\\ E_{\mathcal{U}\setminus\mathcal{V}}\end{bmatrix}\right\| = \left\|\begin{bmatrix}U_{\mathcal{U}\setminus\mathcal{V}}U_{\mathcal{U}^\perp}\end{bmatrix}\begin{bmatrix}\mathcal{E}_{\mathcal{U}\setminus\mathcal{V}}\\ E_{\mathcal{U}\setminus\mathcal{V}}\end{bmatrix}U^*_{\mathcal{U}\setminus\mathcal{V}}\right\|.$$

By subsequently proceeding in the same way as for the error coming from the fine regime, Lemma 5.6, 2., gives the desired bound.

Finally for the bulk, $\|\mathcal{X}\| \le \left(\frac{1}{\sqrt{2}}+o_N(1)\right)\sqrt{\mathfrak{u}}$ by Lemma 5.7. □

This immediately gives the following.

Corollary 5.1.

$$\left\|\begin{pmatrix} D_{\mathcal{V}\setminus\mathcal{W}} & 0 & E^*_{\mathcal{V}\setminus\mathcal{W}} \\ 0 & D_{\mathcal{U}\setminus\mathcal{V}}+\mathcal{E}_{\mathcal{U}\setminus\mathcal{V}} & E^*_{\mathcal{U}\setminus\mathcal{V}} \\ E_{\mathcal{V}\setminus\mathcal{W}} & E_{\mathcal{U}\setminus\mathcal{V}} & \mathcal{X}\end{pmatrix}\right\| \le \sqrt{\mathfrak{u}} - \Theta\left(\mathfrak{u}^{-1/4}\right).$$

Proof. This norm is at most

$$\begin{aligned}&\max\left\{\left\|D_{\mathcal{V}\setminus\mathcal{W}}\right\|, \left\|\begin{pmatrix} D_{\mathcal{U}\setminus\mathcal{V}}+\mathcal{E}_{\mathcal{U}\setminus\mathcal{V}} & E^*_{\mathcal{U}\setminus\mathcal{V}} \\ E_{\mathcal{U}\setminus\mathcal{V}} & \mathcal{X}\end{pmatrix}\right\|\right\} + \left\|E_{\mathcal{V}\setminus\mathcal{W}}\right\| \\ &\qquad \le \max\left\{\left\|D_{\mathcal{V}\setminus\mathcal{W}}\right\|, \max\left\{\left\|D_{\mathcal{U}\setminus\mathcal{V}}+\mathcal{E}_{\mathcal{U}\setminus\mathcal{V}}\right\|, \left\|\mathcal{X}\right\|\right\} + \left\|E_{\mathcal{U}\setminus\mathcal{V}}\right\|\right\} + \left\|E_{\mathcal{V}\setminus\mathcal{W}}\right\|. \quad (5.19)\end{aligned}$$

The bound then follows from Theorem 5.5 and our bounds on d from 5.1. □

With this, we can show that the top eigenvalues correspond to $\mathcal{W}$.

Proposition 5.4. For every $k \leq e^{\log^{1/8} N}$, the kth largest eigenvalue λ of A corresponds to the kth largest lexicographic maximizer $x \in \mathcal{W}$ of (α_x, β_x) in that $\lambda = \lambda_x + O\left((d^r + 1)\,\mathfrak{u}^{-r+1}\right)$.

Proof. We start with the matrix

$$U \begin{pmatrix} D_{\mathcal{W}} & 0 & 0 & 0 \\ 0 & D_{\mathcal{V}\setminus\mathcal{W}} & 0 & E^*_{\mathcal{V}\setminus\mathcal{W}} \\ 0 & 0 & D_{\mathcal{U}\setminus\mathcal{V}} + \mathcal{E}_{\mathcal{U}\setminus\mathcal{V}} & E^*_{\mathcal{U}\setminus\mathcal{V}} \\ 0 & E_{\mathcal{V}\setminus\mathcal{W}} & E_{\mathcal{U}\setminus\mathcal{V}} & \mathcal{X} \end{pmatrix} U^*$$

We then make the transformation by performing the summation

$$U \begin{pmatrix} D_{\mathcal{W}} & 0 & 0 & 0 \\ 0 & D_{\mathcal{V}\setminus\mathcal{W}} & 0 & E^*_{\mathcal{V}\setminus\mathcal{W}} \\ 0 & 0 & D_{\mathcal{U}\setminus\mathcal{V}} + \mathcal{E}_{\mathcal{U}\setminus\mathcal{V}} & E^*_{\mathcal{U}\setminus\mathcal{V}} \\ 0 & E_{\mathcal{V}\setminus\mathcal{W}} & E_{\mathcal{U}\setminus\mathcal{V}} & \mathcal{X} \end{pmatrix} U^* + U \begin{pmatrix} 0 & 0 & 0 & E^*_{\mathcal{W}} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ E_{\mathcal{W}} & 0 & 0 & 0. \end{pmatrix} U^*$$

By perturbation theory, e.g. [Bau85, 7.1.1], [Bam20, Equation 23], each eigenvalue has changed by at most $O\left(\|E_{\mathcal{W}}\|^2\right) = O\left((d^r + 1)\mathfrak{u}^{-r+1}\right)$. Moreover, as $r \geq 5$, by Theorem 5.5 and Corollary 5.1, after the perturbation, nothing outside of $D_{\mathcal{W}}$ can correspond to one of the $e^{\log^{1/8} N}$ largest eigenvalues. Moreover, by Lemma 5.5, the ordering of eigenvalues must match the ordering in D_W, inducing the lexicographic ordering. □

Proof of the main eigenvalue theorem

Proof of Theorem 5.2. Consider the vertex x corresponding to one of the $e^{\log^{1/8} N}$ largest eigenvalues. By Lemma 5.4 and the concentration results from the definition of $\Omega_{3.2}$, (3-5), we have that

$$\lambda_x^2 = \alpha_x + \frac{\beta_x}{\alpha_x} + \frac{d^2}{\alpha_x} + \frac{d^2 + d}{\alpha_x} - \frac{d^2}{\alpha_x} + O\left(\frac{d^{3/2} + 1}{\mathfrak{u}^{4/3}}\right).$$

Therefore by Proposition 5.4, the true eigenvalue λ satisfies

$$\begin{aligned} \lambda &= \sqrt{\alpha_x + \frac{\beta_x}{\alpha_x} + \frac{d^2 + d}{\alpha_x} + O\left(\frac{d^{3/2} + 1}{\mathfrak{u}^{4/3}}\right)} + O\left((d^r + 1)\,\mathfrak{u}^{-(r-1)}\right). \\ &= \sqrt{\alpha_x + \frac{\beta_x}{\alpha_x} + \frac{d^2 + d}{\alpha_x}} + O\left(\left(d^{3/2} + 1\right)\mathfrak{u}^{-11/6} + (d^r + 1)\mathfrak{u}^{-(r-1)}\right) \\ &= \sqrt{\alpha_x + \frac{\beta_x}{\alpha_x} + \frac{d^2 + d}{\alpha_x}} + O\left(\left(d^{3/2} + 1\right)\mathfrak{u}^{-11/6}\right). \end{aligned}$$

as we have assumed $r \geq 5$. The lexicographic ordering follows immediately from Proposition 5.4.

□

5.5 Anticoncentration

In this section we prove that the distribution for λ^2 is anticoncentrated at the edge of the spectrum. For this we use the Poisson approximation for the largest degrees and their neighborhoods recorded in Proposition 4.1 and the representation of the eigenvalues from Theorem 5.2. We then proceed similarly to [ADK23b] to derive that this implies Theorem 5.3, which says that that the maximal eigenvalues are close to the maximal values of a Poisson process. This result is then used to show Lemma 5.9, which gives a lower bound on the distance between the size of the 2-spheres around vertices with maximal or almost maximal degree. In the final lemma of this section we show that this implies spacing of the largest eigenvalues.

Proof of Theorem 5.3. By Lemma 3.4, Theorem 5.2, and Lemma 5.5, with high probability, only vertices of degree between $\mathfrak{u}-2\log^{1/8} N$ and $\mathfrak{u}$ contribute to the top $e^{\log^{1/8} N}$ eigenvalues. Similarly, for every relevant vertex v, $\beta \leq d\alpha_v + \mathfrak{u}^{7/8}$ for every relevant vertex v with high probability by Lemma 3.5. In this region, Theorem 5.2 gives that $\lambda^2 = \alpha_v + \beta_v/\alpha_v + (d^2+d)/\alpha_v + O\left(\left(d^{3/2}+1\right)\mathfrak{u}^{-11/6}\right)$. Moreover, with high probability, neighborhoods of size r around such vertices v are disjoint and treelike by Lemma 3.5. Therefore, define $\mathcal{A}(x,y)$ to be the event that for the ordered pair $(x,y) \in \mathbb{N}^2$, $\mathfrak{u}-2\log^{1/8} N \leq x \leq \mathfrak{u}$ and $0 \leq y \leq dx+\mathfrak{u}^{7/8}$, and let $\mathcal{B}$ be the event that the neighborhoods around all vertices $v \in [N]$ such that $\mathcal{A}(\alpha_v, \beta_v) = 1$ are disjoint tree-like. If we define $\Lambda(\alpha,\beta) = \alpha+\beta/\alpha$, then with high probability Φ is contained in the intensity measure of $\left\{(\Lambda(\alpha_v,\beta_v) + \epsilon(d,N))\mathbf{1}_{\mathcal{A}(\alpha_v,\beta_v)\cap\mathcal{B}} : v \in [N]\right\}$, where $\epsilon = O\left(\frac{d^{3/2}+1}{\mathfrak{u}^{11/6}}\right)$. Moreover, Ψ has the same intensity measure as $\left\{\Lambda(X_i,Y_i)\mathbf{1}_{\mathcal{A}(X,Y_i)} : i \in [N]\right\}$, where the pairs (X_i, Y_i) are independent and satisfy $X_i \sim Pois(d)$ and $Y_i|X_i = x \sim Pois(dx)$.

It is then sufficient to show that,

$$d_{TV}\left(\left\{(\alpha_v,\beta_v)\mathbf{1}_{\mathcal{A}(\alpha_v,\beta_v)\cap\mathcal{B}} : v \in [N]\right\}, \left\{\Lambda(X_i,Y_i)\mathbf{1}_{\mathcal{A}(X,Y_i)} : i \in [N]\right\}\right) = o_N(1), \tag{5.20}$$

which is exactly the statement of Proposition 4.1. □

Now we can simply work with independent Poissons, for which the distribution of the maximum is easier to analyze. We start by determining an interval into which the maximizers of the Y_i, which approximate the β_x, fall.

Lemma 5.8. Consider any function $\zeta(N) = \omega_N(1)$ and $1 \leq K = \zeta^{o_N(1)}$. For fixed $m > 0$, $\mathfrak{a} = \Theta(\mathfrak{u})$, and i.i.d. $Y_1, \ldots, Y_\zeta \sim Pois(d\mathfrak{a})$, with probability $1 - O_N\left(\frac{1}{\sqrt{\log\zeta}}\right)$, the K largest values $Y_{(1)}, \ldots, Y_{(K)}$ are such that for every $1 \leq i \leq K$,

$$Y_{(i)} \in \left[d\mathfrak{a} + \sqrt{2d\mathfrak{a}\log\zeta} - \sqrt{d\mathfrak{a}}\frac{\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta}{\sqrt{2\log\zeta}}, d\mathfrak{a} + \sqrt{2d\mathfrak{a}\log\zeta}\right]$$

where $c_{4.2}$ is the constant from Lemma 4.2.

Proof. If $Y_{(K)}$ is less than some value T, then there are at least $\zeta - K + 1$ Y_i's less than T. Therefore,

$$\mathbb{P}(Y_{(K)} \geq T) \geq 1 - \binom{\zeta}{K-1}(1 - \mathbb{P}(Y_1 \geq T))^{\zeta - K + 1} \geq 1 - \binom{\zeta}{K-1} e^{-(\zeta - K + 1)\mathbb{P}(Y_1 \geq T)}$$
$$\geq 1 - e^{(K-1)\log\zeta - (\zeta - K + 1)\mathbb{P}(Y_1 \geq T)}.$$

To bound $\mathbb{P}(Y_1 \geq T)$ for $T = d\mathfrak{a} + \sqrt{2d\mathfrak{a}\log\zeta} - \frac{\sqrt{d\mathfrak{a}}\left(\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta\right)}{\sqrt{2\log\zeta}}$, we use the tail bound from Corollary 4.2, with

$$\delta := \frac{1}{\sqrt{d\mathfrak{a}}}\left(\sqrt{2\log\zeta} - \frac{\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta}{\sqrt{2\log\zeta}}\right).$$

As

$$d\mathfrak{a}\delta^2/2 = \log\zeta - \log K + \frac{1}{2}\log 2 + \log c_{4.2} - \frac{3}{2}\log\log\zeta$$
$$+ \frac{\left(\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta\right)^2}{4\log\zeta},$$

Corollary 4.2 gives that

$$\mathbb{P}\left(Y_1 \geq d\mathfrak{a} + \sqrt{d\mathfrak{a}}\left(\sqrt{2\log\zeta} - \frac{\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta}{\sqrt{2\log\zeta}}\right)\right)$$
$$\geq \frac{c_{4.2}}{\zeta}\frac{e^{\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta - \frac{\left(\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta\right)^2}{4\log\zeta}}}{\sqrt{2\log\zeta} - \frac{\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta}{\sqrt{2\log\zeta}}}$$
$$= (1 + o_N(1))\frac{K\log\zeta}{\zeta}.$$

Thus

$$1 - e^{(K-1)\log\zeta - (\zeta - K + 1)\mathbb{P}(Y_1 \geq T)} = 1 - \zeta^{-1 + o_N(1)}. \tag{5.21}$$

To prove the upper bound we proceed similarly, using that by a union bound

$$\mathbb{P}\Big(Y_{(1)} \geq T\Big) \leq \zeta\,\mathbb{P}\Big(Y_1 \geq T\Big). \tag{5.22}$$

Using once more the tail bound from Corollary 4.2 with $\delta = \sqrt{\frac{2\log\zeta}{d\mathfrak{a}}}$, we obtain

$$\mathbb{P}\Big(Y_1 \geq T\Big) \leq \frac{e^{-\log\zeta}}{\sqrt{2\log\zeta}},$$

which means that (5.22) can be upper bounded by $\frac{1}{\sqrt{2\log\zeta}}$.

□

The following lemma will imply spacing between eigenvalues.

Lemma 5.9. Fix $\mathfrak{a} \in \{\mathfrak{u}-1, \mathfrak{u}\}$. For any $K = \log^{o(1)} N$, with high probability the maximum $K+1$ values of $\beta_x^{(1)}$ of vertices with degree $\mathfrak{a}$ are separated by at least $\frac{(d\mathfrak{u})^{1/2}}{\log\left(\frac{\mathfrak{u}}{d}\right)^3 (K+1)^{3\log\log\log N}}$.

Proof. Denote by ζ the number of vertices of degree $\mathfrak{a}$. We will split into two cases, based on whether ζ is small relative to K. As we will see, if ζ is small, then we can bound the probability using the anticoncentration of the Poisson. If ζ is larger, we can shift our focus to the regime of Lemma 5.8. It is sufficient to split our cases according to $(K+1)^{\log\log\log(N)}$. Consider two vertices u, v such that $\alpha_u = \alpha_v = \mathfrak{a}$. If $\zeta \leq (K+1)^{\log\log\log(N)}$, then (4.11) in the proof of Theorem 5.3 implies that the distribution of the β_x's approaches the distribution of Poissons, so the probability that $\left|\beta_u^{(1)} - \beta_v^{(1)}\right| \leq \eta$ is at most $\frac{2\eta}{\sqrt{d\mathfrak{a}}} + \tilde{O}(N^{-1/2})$, considering the mode of a Poisson is at its mean, with probability at most $\frac{1}{\sqrt{d\mathfrak{a}}}$. Therefore the probability that any pair is within distance η is at most

$$\binom{\zeta}{2}\frac{2\eta}{\sqrt{d\mathfrak{a}}} \leq \frac{\eta(K+1)^{2\log\log\log N}}{\sqrt{d\mathfrak{a}}}$$

This converges to 0 for $\eta = \frac{(d\mathfrak{u})^{1/2}}{(K+1)^{3\log\log\log N}}$.

Otherwise, if $\zeta \geq (K+1)^{\log\log\log(N)}$, referring once more to (4.11) in the proof of Theorem 5.3 and Lemma 5.8, with high probability the K maximizers x of $\beta_x^{(1)}$ satisfy

$$\beta_x^{(1)} \in \left[d\mathfrak{a} + \sqrt{2d\mathfrak{a}\log\zeta} - \sqrt{d\mathfrak{a}}\frac{\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta}{\sqrt{2\log\zeta}}, d\mathfrak{a} + \sqrt{2(\log\zeta)d\mathfrak{a}}\right]. \tag{5.23}$$

To show the improvement in density, we consider the probability that $\beta_x^{(1)} = d\mathfrak{a} + t$, for $|t| = (1 + o(_N 1))\sqrt{2(\log\zeta)d\mathfrak{a}}$. We then have by the Stirling approximation,

$$\frac{e^{-d\mathfrak{a}}(d\mathfrak{a})^{d\mathfrak{a}+t}}{(d\mathfrak{a}+t)!} = (1 + o_N(1))\frac{e^t}{(1+\frac{t}{d\mathfrak{a}})^{d\mathfrak{a}+t}\sqrt{2\pi(d\mathfrak{a}+t)}}.$$

To approximate this, we have

$$\left(1 + \frac{t}{d\mathfrak{a}}\right)^{d\mathfrak{a}+t} = e^{\log\left(1+\frac{t}{d\mathfrak{a}}\right)(d\mathfrak{a}+t)} = e^{\left(\frac{t}{d\mathfrak{a}} - \frac{t^2}{2(d\mathfrak{a})^2} + O\left(\frac{t^3}{(d\mathfrak{a})^3}\right)\right)(d\mathfrak{a}+t)} = e^{t + \frac{t^2}{2d\mathfrak{a}} + O\left(\frac{t^3}{(d\mathfrak{a})^2}\right)}$$

In our window, $t \geq \sqrt{2(\log\zeta)d\mathfrak{a}} - \sqrt{d\mathfrak{a}}\frac{\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta}{\sqrt{2\log\zeta}}$, and we have

$$\frac{e^{-d\mathfrak{a}}(d\mathfrak{a})^{d\mathfrak{a}+t}}{(d\mathfrak{a}+t)!} = \frac{1}{e^{(1+o_N(1))t^2/(2d\mathfrak{a})}\sqrt{2\pi(d\mathfrak{a}+t)}} \leq \frac{e^{\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta}}{\zeta^{1-O\left(\sqrt{\frac{\log\zeta}{d\mathfrak{a}}}\right)}\sqrt{2\pi(d\mathfrak{a}+t)}}.$$

Here we must have $\sqrt{\frac{\log\zeta}{d\mathfrak{u}}} \to 0$, therefore, we use the assumption that $d \gg \frac{(\log\log N)^2}{\log N}$. The probability that there are at least two vertices in a window of length 2η around some $d\mathfrak{a} + t$ with $t = (1 + o_N(1))\sqrt{2d\mathfrak{a}\log\zeta}$ is therefore

$$\binom{\zeta}{2}\left(\frac{2\eta e^{\log K + \frac{1}{2}\log 2 - \log c_{4.2} + \frac{3}{2}\log\log\zeta}}{\zeta^{1-O\left(\sqrt{\frac{\log\zeta}{d\mathfrak{a}}}\right)}\sqrt{2\pi(d\mathfrak{a}+t)}}\right)^2 \le 4c_{4.2}^{-1}K^2\eta^2\log(\frac{\mathfrak{u}}{d})^3(d\mathfrak{u})^{-1} \tag{5.24}$$

for sufficiently large N, considering that with high probability $\zeta \le (\frac{\mathfrak{u}}{d})^{3/2}$ by Lemma 3.4.

To translate this into distance between $\beta_u^{(1)}$ and $\beta_v^{(1)}$, we cover the large interval corresponding to (5.23) with small intervals of length 2η, and centers spaced at distance η. To cover this large interval, we need at most $\eta^{-1}\sqrt{d\mathfrak{u}}$ small intervals. Therefore, union bounding the probability (5.24) gives that

$$\begin{aligned}
&\mathbb{P}\Big(\exists u, v \in [N] : \alpha_u = \alpha_v = \mathfrak{a}, \left|\beta_u^{(1)} - \beta_v^{(1)}\right| \le \eta\Big)\mathbf{1}\left(\zeta \ge (K+1)^{\log\log\log N}\right) \\
&\le 4c_{4.2}^{-1}K^2\eta^2\log\left(\frac{\mathfrak{u}}{d}\right)^3(d\mathfrak{u})^{-1}\eta^{-1}\sqrt{d\mathfrak{u}} \\
&\le 5c_{4.2}^{-1}K^2\eta\log\left(\frac{\mathfrak{u}}{d}\right)^3(d\mathfrak{u})^{-1/2}
\end{aligned}$$

This probability converges to 0 for $\eta = \frac{(d\mathfrak{u})^{1/2}}{\log\left(\frac{\mathfrak{u}}{d}\right)^3(K+1)^{3\log\log\log N}}$. □

Lemma 5.10. With high probability, for $u \neq v \in \mathcal{W}$ corresponding to the largest $K+1$ λ's, we have $|\lambda_u - \lambda_v| \ge \frac{d^{1/2}}{3\mathfrak{u}\log\left(\frac{\mathfrak{u}}{d}\right)^3(K+1)^{3\log\log\log N}}$.

Proof. By Lemma 5.5, this is immediately true if $\alpha_u \neq \alpha_v$. Therefore assume $\alpha_u = \alpha_v$. By Lemma 5.9, with high probability the $K+1$ maximizers of $\beta_u^{(1)}$ are spaced at distance $\frac{(d\mathfrak{u})^{1/2}}{\log\left(\frac{\mathfrak{u}}{d}\right)^3(K+1)^{3\log\log\log N}}$. Therefore, by Lemma 5.5, for $u \neq v$ and sufficiently large N, and as $d \gg \log^{-5/3} N$,

$$\begin{aligned}
|\lambda_u - \lambda_v| &= \frac{|\lambda_u^2 - \lambda_v^2|}{|\lambda_u + \lambda_v|} \\
&\ge \left(\frac{\frac{(d\mathfrak{u})^{1/2}}{\log\left(\frac{\mathfrak{u}}{d}\right)^3(K+1)^{3\log\log\log N}}}{\mathfrak{u}} + O\left(\left(1 + d^{3/2}\right)\mathfrak{u}^{-4/3}\right)\right)\frac{1}{2\sqrt{\mathfrak{u}} + O\left(\frac{d}{\sqrt{\mathfrak{u}}}\right)} \\
&\ge \frac{d^{1/2}}{3\mathfrak{u}\log\left(\frac{\mathfrak{u}}{d}\right)^3(K+1)^{3\log\log\log N}}
\end{aligned} \tag{5.25}$$

by the lower bound on d. □

5.6 Eigenvector Structure

Proposition 5.5. For $k \leq K = \log^{o_N(1)} N$, define x to be the vertex corresponding to the kth largest eigenvalue of A, as per Theorem 5.2. The eigenvector $\mathbf{v}$ of λ satisfies

$$\left\|\mathbf{v} - \mathbf{w}_+(x)\right\| = O\left(\mathfrak{u}^{-r/2+2}\right).$$

Proof. By Theorem 5.2, there is a correspondence between the top $K+1$ eigenvalues and eigenvectors of the matrix A and the top $K+1$ eigenvalues λ_x of the truncated balls around vertices $x \in \mathcal{W}$ together with their eigenvectors. Moreover, by Lemma 5.10, the difference between each pair of these $K+1$ eigenvalues is at least $\frac{d^{1/2}}{4\mathfrak{u}\log(\frac{\mathfrak{u}}{d})^3(K+1)^{3\log\log\log N}}$. Standard perturbation theory (see [GLO20] Theorem 2 and the remarks following it) gives that, if we fix the index $1 \leq i \leq K$,

$$\begin{aligned}\left\|\mathbf{v} - \mathbf{w}_+(x)\right\| &\leq \left\|E_W\right\| \cdot \left(\min_{j\neq i}\left|\lambda - \lambda_j\right|\right)^{-1} \\ &\leq O\left(\frac{\mathfrak{u}\log\left(\frac{\mathfrak{u}}{d}\right)^3(K+1)^{3\log\log\log N}}{d^{1/2}}(d^r+1)\mathfrak{u}^{-r/2+1/2}\right) \\ &= O\left(\mathfrak{u}^{-r/2+2}\right).\end{aligned}$$

by our assumptions on d from Definition 5.1 and K, and the bound on $\|E_W\|$ from Theorem 5.5. □

Proof of Theorem 5.4. By Proposition 5.2, Proposition 5.5, and the triangle inequality,

$$\begin{aligned}\mathbf{v}|_x &= \frac{1}{\sqrt{2}} + O\left(\left(1+d^{-1/2}\right)\mathfrak{u}^{-1/3} + \mathfrak{u}^{-r/2+2}\right) \\ &= \frac{1}{\sqrt{2}} + O\left(\left(1+d^{-1/2}\right)\mathfrak{u}^{-1/3}\right).\end{aligned}$$

as $r \geq 5$. Moreover, as $r \geq 2r'$, for $1 \leq i \leq r'$,

$$\begin{aligned}\left\|\mathbf{v}|_{S_i(x)}\right\| &= \left(\frac{d}{\alpha}\right)^{(i-1)/2}\frac{1}{\sqrt{2}}\left(1 + O\left(\left(1+d^{-1/2}+d^{-(i-1)}\right)\mathfrak{u}^{-1/3}\right)\right) + O\left(\mathfrak{u}^{-r/2+2}\right) \\ &= \left(\frac{d}{\alpha}\right)^{(i-1)/2}\frac{1}{\sqrt{2}}\left(1 + O\left(\left(1+d^{-1/2}+d^{-i+1}\right)\mathfrak{u}^{-1/3}\right)\right)\end{aligned}$$

as desired.

Similarly,

$$\left\|\mathbf{w}_+|_{[N]\setminus B_i(x)}\right\| = \left(\frac{d}{\alpha}\right)^{i/2}\left(1 + O\left(\left(1+d^{-1/2}+d^{-r+1}\right)\mathfrak{u}^{-1/3}\right)\right).$$

□

Chapter 6

The largest eigenvalue of weighted Erdős-Rényi graphs

Random graphs are used to model many real world phenomena, like electrical and social networks. Often for more effective encoding of physical phenomena, edges in a random network are equipped with random weights, which for instance could denote resistances in an electrical network. This leads one to consider Erdős-Rényi random graph $\mathcal{G}(N, p)$ with random weights assigned to the edges. Its adjacency matrix can be regarded as *sparse* or *diluted* random matrices, where each entry of a Wigner matrix is multiplied by an independent Bernoulli random variable with mean p.

Some important properties of the weighted random graph are captured by the spectrum of its adjacency matrix. For instance, the largest eigenvalue and the corresponding eigenvector can be used to measure the spread of diseases on graphs. Regarding this eigenvalue, a lot of research has been devoted towards studying its typical and atypical behavior: the former is concerned with the value and fluctuations that the largest eigenvalue has with high probability, and the latter is about the probability that the eigenvalue deviates significantly from its typical value. A particularly important question in this direction is how the spectral behavior depends on the precise matrix entry distribution, and we say that the *universality* phenomenon holds if there is no such dependency in some asymptotical sense.

An important distinction of sparse matrices from denser ones is that the universality of the spectral behavior breaks down in the sparse case and the spectrum depends rather crucially on the entry distribution. While the study of the bulk spectral properties of sparse random matrices has witnessed some activity, e.g., [BC15, BSV17], the precise edge statistics has still been mostly out of reach of the known methods which are primarily tailored to analyze denser graphs, see for instance [AGH21, CDG23].

The present chapter is based on the paper [GHN24] and is aimed at advancing our understanding in this direction, focusing on the large deviation properties. Towards this, we study the large deviation properties of the largest eigenvalue of the adjacency matrix of the Erdős-Rényi random graph $\mathcal{G}_{N,\frac{d}{N}}$ with general edge-weight distributions $(Y_{ij})_{1\le i<j\le N}$. The topic of large deviations of spectral observables of random matrices have attracted much interest over the last few decades leading to a considerable literature, some of which will be reviewed after the main results, in Section 6.1.

We will work in the particular setting where the edge-weights $(Y_{ij})_{1\le i<j\le N}$ have tails of the form

$$\mathbb{P}(|Y_{ij}| > t) \approx e^{-t^{\alpha}}, \qquad t > 1$$

for some $\alpha > 0$ (we will leave the notion of $\approx$ somewhat imprecise for now). It turns out that $\alpha = 2$ is critical in a sense which will become apparent once we state our main results which will address both the $\alpha > 2$ (light-tailed) and $\alpha < 2$ (heavy-tailed) cases.[1]

At this point, before embarking on stating our main results precisely in the next section, we choose to highlight some of their key features.

Perhaps the most interesting consequence of our main results is the surprising universality phenomenon for the large deviation of the largest eigenvalue in the light-tailed case $\alpha > 2$ where the deviation probabilities (ignoring smaller order terms) does not depend on α and is also identical to the one for Erdős-Rényi graphs without edge-weights, which essentially corresponds to "$\alpha = \infty$". Our results also yield new law of large numbers (LLN) results about the largest eigenvalue, which seem to be rather challenging to obtain using previous methods. Interestingly, it turns out that the LLN behavior exhibits a transition as well at $\alpha = 2$.

Let us now precisely define our model and state our main results. A short summary of the history of this problem will be provided subsequently.

6.1 Main results

All results below refer to $\mathcal{G}\left(N, \frac{d}{N}\right)$ with d constant.

Let $X = (X_{ij})_{i,j\in[N]}$ be the adjacency matrix of $\mathcal{G}_{N,p}$ and $Y = (Y_{ij})_{i,j\in[N]}$ be a standard (symmetric) Wigner matrix, that is we assume that $Y_{ji} = Y_{ij}$, $Y_{ii} = 0$ and $\{Y_{ij}\}_{1\le i<j\le N}$ are i.i.d random variables. The matrix of interest for us is $Z = X \odot Y$, i.e., $Z_{ij} = X_{ij}Y_{ij}$. This is a sparse random matrix which can be regarded as the adjacency matrix of a *weighted* random graph, whose underlying random graph is $\mathcal{G}_{N,\frac{d}{N}}$ with i.i.d. edge-weights coming from $\{Y_{ij}\}_{1\le i<j\le N}$. Throughout the chapter we interchangeably use the notation X both for the Erdős-Rényi graph $\mathcal{G}_{N,\frac{d}{N}}$ as well as its adjacency matrix.

Throughout the chapter matrix entries will be random variables with Weibull shape, which defined below essentially says that the tail probabilities are comparable to that of the Weibull distribution.

[1]The special case of the Gaussian distribution, which can be thought to be corresponding to $\alpha = 2$, up to polynomial prefactors, was studied previously in [GN22] by Ganguly and Nam; see Remark 6.2 for a further discussion in this direction.

Definition 6.1. A random variable W has *Weibull shape* with shape parameter $\alpha > 0$ if there exist constants $C_1, C_2 > 0$ such that for all $t > 1$,

$$\frac{C_1}{2}e^{-t^\alpha} \le \mathbb{P}\Big(W \ge t\Big) \le \frac{C_2}{2}e^{-t^\alpha} \quad \text{and} \quad \frac{C_1}{2}e^{-t^\alpha} \le \mathbb{P}\Big(W \le -t\Big) \le \frac{C_2}{2}e^{-t^\alpha}. \tag{6.1}$$

The definition ensures that

$$C_1 e^{-t^\alpha} \le \mathbb{P}\Big(|W| \ge t\Big) \le C_2 e^{-t^\alpha}, \tag{6.2}$$

which will be notationally convenient later.

The cause for assuming a symmetric tail behavior is that a much heavier lower tail causes the spectral norm of the random matrix to be governed by not the largest but the smallest eigenvalue which is a large negative number and hence the largest eigenvalue is not an interesting object of study anymore.

For brevity, throughout the chapter, we call a random variable with Weibull shape as a *Weibull random variable*.

We now state the main results of this chapter.

Light-tailed case, $\alpha > 2$.

Let

$$\lambda_\alpha^{\text{light}} := 2^{\frac{1}{\alpha}}\alpha^{-\frac{1}{2}}(\alpha-2)^{\frac{1}{2}-\frac{1}{\alpha}}\frac{(\log N)^{\frac{1}{2}}}{(\log\log N)^{\frac{1}{2}-\frac{1}{\alpha}}}. \tag{6.3}$$

We will shortly state (see Corollary 6.1) that the above is the typical value of $\lambda_1(Z)$ for $\alpha > 2$.

Theorem 6.1. For any $\delta > 0$,

$$\lim_{N\to\infty} -\frac{\log \mathbb{P}\left(\lambda_1(Z) \ge (1+\delta)\lambda_\alpha^{\text{light}}\right)}{\log N} = (1+\delta)^2 - 1.$$

Next, we establish the lower tail large deviation.

Theorem 6.2. For any $0 < \delta < 1$,

$$\lim_{N\to\infty} \frac{1}{\log N}\left(\log\log \frac{1}{\mathbb{P}\left(\lambda_1(Z) \le (1-\delta)\lambda_\alpha^{\text{light}}\right)}\right) = 1 - (1-\delta)^2.$$

A crucial observation is that the upper and lower tail large deviation results establish a sense of *universality* with the rate function not depending on α.

Further, as a corollary we have the following law of large numbers.

Corollary 6.1. We have

$$\lim_{N\to\infty} \frac{(\log\log N)^{\frac{1}{2}-\frac{1}{\alpha}}}{(\log N)^{\frac{1}{2}}} \lambda_1(Z) = 2^{\frac{1}{\alpha}} \alpha^{-\frac{1}{2}} (\alpha-2)^{\frac{1}{2}-\frac{1}{\alpha}}$$

in probability.

Note that the case of no edge-weights can be thought of as corresponding to "$\alpha = \infty$". Now while this was treated in [BBG21, Theorem 1.1] (see Theorem 6.5 stated later), in some sense the same can be deduced in a limiting sense from Theorem 6.1. In fact, noting that

$$\lim_{\alpha\to\infty} \lambda_\alpha^{\text{light}} = \lim_{\alpha\to\infty} \left[2^{\frac{1}{\alpha}} \alpha^{-\frac{1}{2}} (\alpha-2)^{\frac{1}{2}-\frac{1}{\alpha}} \frac{(\log N)^{\frac{1}{2}}}{(\log\log N)^{\frac{1}{2}-\frac{1}{\alpha}}} \right] = \frac{(\log N)^{\frac{1}{2}}}{(\log\log N)^{\frac{1}{2}}},$$

one obtains [BBG21, Theorem 1.1] by taking $\alpha \to \infty$ in Theorem 6.1. Note that the large deviation rate function turns out to be not only the same for all light-tailed distribution, i.e. whenever $\alpha > 2$, but also in the limit, when $\alpha = \infty$.

Heavy-tailed case, $\alpha < 2$

Counterpart to (6.3), we define

$$\lambda_\alpha^{\text{heavy}} := (\log N)^{\frac{1}{\alpha}}. \tag{6.4}$$

Stating the upper tail large deviation needs the following definition. For $\theta > 1$, let $\phi_\theta : \{2, 3, \cdots\} \to \mathbb{R}$ be defined by

$$\phi_\theta(k) := \sup_{v\in\mathbb{R}^k, v=(v_1,\cdots,v_k) \|v\|_1=1} \sum_{i,j\in[k], i\neq j} |v_i|^\theta |v_j|^\theta. \tag{6.5}$$

The statement involves further considering two sub-cases.

Theorem 6.3. Let $\delta > 0$.

1. In the case $1 < \alpha < 2$, let $\beta > 2$ be the conjugate of α (i.e. $\frac{1}{\alpha} + \frac{1}{\beta} = 1$). For an integer $k \geq 2$, define

$$\psi_{\alpha,\delta}(k) := \frac{k(k-3)}{2} + \frac{1}{2}(1+\delta)^\alpha \phi_{\beta/2}(k)^{1-\alpha}. \tag{6.6}$$

Then,

$$\lim_{N\to\infty} -\frac{\log \mathbb{P}\left(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha^{\text{heavy}}\right)}{\log N} = \min_{k=2,3,\cdots} \psi_{\alpha,\delta}(k). \tag{6.7}$$

2. In the case $0 < \alpha \leq 1$,

$$\lim_{N\to\infty} -\frac{\log \mathbb{P}\left(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha^{\text{heavy}}\right)}{\log N} = (1+\delta)^\alpha - 1. \tag{6.8}$$

The above theorem shows that once heavy-tailed edge-weights are induced on the *sparse* graph $\mathcal{G}_{N,\frac{d}{N}}$, the large deviation rate function for the largest eigenvalue exhibits a phase transition: the rate function is a *piecewise smooth* function whose behavior changes as $\arg\max_k \psi_{\alpha,\delta}(k)$ (i.e. size of the clique needed to have atypically large λ_1) varies. This is in a sharp contrast to the large deviation result for the standard (i.e. dense) Wigner matrices with heavy-tailed entries (i.e. edge-weights are induced on the complete graph) [Aug16], where it is proved that there exists a *smooth* function $I(\delta)$ such that for all $\delta > 0$,

$$\mathbb{P}(\lambda_1 > 2 + \delta) \approx e^{-I(\delta)N^{\alpha/2}}.$$

Remark 6.1. We do not have a closed expression for the quantity $\phi_\theta(k)$ defined in (6.5), unless $\theta = 1$ in which case we know $\phi_1(k) = \frac{k-1}{k}$ (by the classical Motzkin-Straus theorem [MS65]). While the definition of $\phi_\theta(\cdot)$, might seem somewhat unmotivated at this point, it appears quite naturally once we bound $\lambda_1(Z)$ in terms of the 'entrywise' L^p-(quasi)norm of Z. See Section 2.1 for details.

Next, we obtain the lower tail large deviation.

Theorem 6.4. For any $0 < \delta < 1$,

$$\lim_{N\to\infty} \frac{1}{\log N}\left(\log\log \frac{1}{\mathbb{P}\left(\lambda_1(Z) \le (1-\delta)\lambda_\alpha^{\text{heavy}}\right)}\right) = 1 - (1-\delta)^\alpha.$$

Note that unlike for light-tail edge weights, in this case the large deviation rate function *does* depend on α.

Finally, as a counterpart to Corollary 6.1, we have the following law of large numbers result.

Corollary 6.2. We have

$$\lim_{N\to\infty} \frac{\lambda_1(Z)}{(\log N)^{\frac{1}{\alpha}}} = 1$$

in probability.

Remark 6.2. While for the sake of simplicity as well as brevity we have chosen to work with distributions as in (6.1), certain generalizations are not hard to make. E.g., a simple rescaling shows that similar large deviation results hold for random variables possessing more general types of Weibull shape, where for $t > 1$,

$$\frac{C_1}{2}e^{-\eta t^\alpha} \le \mathbb{P}\Big(W \ge t\Big) \le \frac{C_2}{2}e^{-\eta t^\alpha} \quad \text{and} \quad \frac{C_1}{2}e^{-\eta t^\alpha} \le \mathbb{P}\Big(W \le -t\Big) \le \frac{C_2}{2}e^{-\eta t^\alpha} \tag{6.9}$$

for some additional scale parameter $\eta > 0$. A piece of straightforward algebra shows that the rate function for $\lambda_1(Z)$ does not change with η.

Further, a little additional work also allows one to consider an even more general class of distributions given by

$$\frac{C_1}{2}t^{-c_1}e^{-\eta t^\alpha} \le \mathbb{P}\Big(W \ge t\Big) \le \frac{C_2}{2}t^{-c_2}e^{-\eta t^\alpha} \quad \text{and} \quad \frac{C_1}{2}t^{-c_1}e^{-\eta t^\alpha} \le \mathbb{P}\Big(W \le -t\Big) \le \frac{C_2}{2}t^{-c_2}e^{-\eta t^\alpha}$$

for parameters $c_1, c_2 \ge 0$. This allows one to include polynomial pre-factors, which for $\alpha = 2$ covers the Gaussian case considered in [GN22].

We next include a brief overview of the literature on large deviations of spectral observables of random matrices.

Related results

Much effort has been devoted to proving whether spectral properties of random matrix ensembles exhibit certain universality features i.e. do not depend on the precise entry distribution as the matrix size goes to infinity. For instance, under some moment conditions, the appropriately scaled empirical distribution of the eigenvalues converges to the Wigner's semi-circle law [Wig58, Arn67] and the largest eigenvalue lies near the edge of the semi-circle distribution [Juh81, FK81, BY88]. On the other hand, large deviation behavior is generally far from universal making this an intriguing research direction which has also witnessed considerable activity. Large deviations for the empirical distribution and the largest eigenvalue were first derived for the Gaussian ensembles [BAG97, BADG01]. Progress stalled for a while, until a surprising recent result by Guionnet and Husson [GH20], where the large deviation universality was established for the largest eigenvalue of 'sharp' sub-Gaussian matrices. Their proof relied on the method of spherical integrals. Together with Augeri [AGH21], they also showed that for a more general class of sub-Gaussian matrices, the rate function is universal for small large deviations, but beyond that also depends on other properties of the moment generating function. Subsequently, continuing this line of work, Cook, Ducatez and Guionnet, in [CDG23], strengthened the result into a full large deviation principle for such sub-Gaussian Wigner matrices.

The behavior changes a lot when the entry distribution possesses tails *heavier* than Gaussian. Bordenave and Caputo [BC14] proved a large deviation result for the empirical distribution of Wigner matrices with stretch-exponential entries whose tails decay at the rate e^{-t^α} with $\alpha \in (0, 2)$. The proof relies on the observation that the atypical behavior is the result of a few atypically large entries. Building on the same idea, Augeri [Aug16] subsequently obtained a large deviation principle for the largest eigenvalue. In both cases, the large deviation speed as well as the rate function depend on α and thus on the precise tail behavior of the entries.

We next move on to the results on the spectral properties of the adjacency matrix of Erdős-Rényi random graphs $\mathcal{G}_{N,p}$. As mentioned in Chapter 5, the largest eigenvalue is typically linked to the average degree in the dense case, and to the largest degrees in the sparse case [KS03, ADK21b, TY21]. Beyond typical behavior, the study of the large deviation behavior for the spectrum of $\mathcal{G}_{N,p}$ becomes much more delicate and requires new methods. In the dense case (i.e. p fixed in N), Chatterjee and Varadhan [CV11, CV12] established a large deviation principle using the powerful graph limit theory and characterized the rate function in terms of a variational problem. This variational problem was later analyzed by Lubetzky and Zhao [LZ15]. It is important to note that the eigenvalues are normalized in a different way than in [BADG01, BAG97] here, and that the matrix entries are not assumed to be centered. The quantity of interest here are eigenvalues of the order of N, where $N \times N$ is the size of the matrix, of which there is typically just one when the entries are not centered. In contrast, in [BAG97, BADG01], roughly speaking, eigenvalues and deviations of order $\sqrt{N}$ were studied. In the sparse case $p \to 0$, the graph limit theory no longer applies. In the breakthrough work by Chatterjee and Dembo [CD16], a general framework of *nonlinear large deviations* was developed leading to a similar variational problem. This was later extended [CD20, Aug20] and analyzed in [BG20] to obtain large deviations for the largest eigenvalue in the sparsity regime $\frac{1}{\sqrt{N}} \ll p \ll 1$. Finally, the sparsity we will consider in this chapter, namely $p = \frac{d}{N}$ was covered in [BBG21]. It was shown that the large deviation behavior of edge eigenvalues is a consequence of the emergence of vertices of atypically large degree.

Finally, let us review the few existing results about the spectral behavior of the adjacency matrix of *weighted* Erdős-Rényi graphs, or in other words, sparse or diluted Wigner matrices. The typical largest eigenvalue of dense graphs (i.e. p fixed in N) with general edge-weights, under suitable moment conditions, belongs to the universality class of general Wigner matrices. In the sparser regime $p \to 0$, Khorunzhy [Kho01] proved that once $\frac{\log N}{N} \ll p \ll N^{\beta}$ for some $\beta > 0$, the largest eigenvalue is asymptotically $2\sqrt{Np}$ with high probability and does thus not depend on the precise distribution of the edge-weights. Recently, [BGBK20, TY21] treated the typical behavior of diluted Wigner matrices, where some of the above assumptions are relaxed. However all results mentioned before are in the regime $p \gg \frac{1}{N}$, which excludes the case $p = \frac{d}{N}$.

However, despite the above advances, significantly less is known about the spectral *large deviation* behavior in the setting of diluted Wigner matrices. Large deviations for the dense $\mathcal{G}_{N,p}$ with Gaussian edge-weights can be deduced from the aforementioned result for the sub-Gaussian Wigner matrices [AGH21]. This result together with [GH20] show that even for the same weight distribution, the large deviation behavior for the cases $p = 1$ and $p < 1$ (p is fixed) are different. No other regime of p, in particular when $p \to 0$, is covered by previous results [AGH21, GH20]. In this direction, a large deviation result for the largest eigenvalue, in the case $p = \frac{d}{N}$ with *Gaussian* edge-weights, was recently proved by Ganguly and Name in [GN22]. Finally, very recently, [Aug24] established a large deviation principle for the empirical spectral distribution of diluted Wigner matrices in the regime $\frac{\log N}{N} \ll p \ll 1$.

Recall that Theorem 6.1 implies a universal spectral large deviations behavior when $\alpha > 2$. While the large deviation result for the 'sharp' sub-Gaussian Wigner matrices [GH20] can be considered as a universality result for dense Wigner matrices, it seems

Theorem 6.1 is the first universality result of its kind in the sparse regime.

Idea of the proof

The proofs rely on identifying the correct geometric mechanisms responsible for the largest eigenvalue which we describe next.

Light-tailed weights

We start by considering the adjacency matrix X without edge-weights. In the influential work [KS03] it was shown that the largest eigenvalue is determined essentially by the star, incident on the vertex with the largest degree, which is typically $\frac{\log N}{\log \log N}$. Since the largest eigenvalue of a star of degree ℓ is equal to $\sqrt{\ell}$, we have $\lambda_1(X) \approx \frac{(\log N)^{\frac{1}{2}}}{(\log \log N)^{\frac{1}{2}}}$. It was subsequently established in [BBG21] that in this case, even in the large deviations regime, the largest eigenvalue is essentially determined by an atypically large degree vertex.

Now, let us consider the case when the edges are equipped with light-tailed edge-weights. Owing to the lightness of the tail, one might naturally expect that vertices with degree of order $\frac{\log N}{\log \log N}$ will continue to be the determining structure, but the existence of weights does add another element of randomness that can increase the largest eigenvalue. This calls for the need to balance the fact that there are more stars of lower degree, while high degree stars are more likely to have a big largest eigenvalue necessitating the estimation of the contributions from vertices of degree close to $\gamma \frac{\log N}{\log \log N}$ for $0 < \gamma < 1$.

Using a binomial tail estimate, one can deduce that the probability that a vertex has degree close to $\gamma \frac{\log N}{\log \log N}$ is roughly $N^{-\gamma}$ and hence there are approximately around $N^{1-\gamma}$ vertices of that degree. Because $\mathcal{G}_{N,\frac{d}{N}}$ is sparse, it is likely that a constant proportion of these vertices has distinct neighborhoods with no edges present within those neighborhoods and hence these high-degree vertices and their neighborhoods induce vertex-disjoint stars. The next step is to consider the contribution from the edge-weights. Here we need two ingredients. First, the fact that the largest eigenvalue of a weighted star is the square root of the sum of the squares of the weights. The second fact we use is that for light-tailed weights, the probabilistically optimal way to obtain a large squared sum of the weights is by all of them being uniformly large. This allows us to deduce that for any weighted star S of degree k,

$$\mathbb{P}(\lambda_1(S) > t) \approx e^{-t^{\alpha} k^{1-\frac{\alpha}{2}}}. \tag{6.10}$$

Using this, to find a typical value t of the largest eigenvalue of such collection of stars, the probability in (6.10) with $k = \gamma \frac{\log N}{\log \log N}$ should balance out the number of such stars which is roughly $n^{1-\gamma}$. A further optimization over γ indeed indicates that typical value of the largest eigenvalue is close to $\lambda_\alpha^{\text{light}}$ (defined in (6.3)).

The proof of the upper bound in Theorem 6.1 now relies on establishing the above heuristic even in the large deviations regime. A key subtle distinction arises from the fact that the large deviation of λ_1 may induce atypical behavior of both the degrees as well as the edge weights which has to be taken into consideration as well.

The proof of the lower bound is much simpler and consists of ensuring that there is a star of degree $\gamma_\delta \frac{\log N}{\log\log N}$, for some δ dependent constant γ_δ, with large enough weights on the edges.

Heavy-tailed weights

The results of [GN22] indicate that typically the largest eigenvalue of $\mathcal{G}_{N,\frac{d}{N}}$ with *Gaussian* weights is determined by the maximal edge-weight in absolute value. Since the latter increases as the tail of individual random variables becomes heavier, this suggests that the same mechanism persists when the edge-weights have heavier tails (i.e. $\alpha < 2$).

Note that there are on average $\frac{dn}{2}$ edges present in the graph $\mathcal{G}_{N,\frac{d}{N}}$. Thus, the probability that the largest entry is greater in absolute value than t is

$$1 - \mathbb{P}\Big(\max_{(i,j)\in E(X)} |Z_{ij}| < t\Big) \approx 1 - (1 - e^{-t^\alpha})^{\frac{dN}{2}} \approx \frac{dN}{2} e^{-t^\alpha}. \tag{6.11}$$

This shows that the typical value of the maximum entry (in absolute value) is $\lambda_\alpha^{\text{heavy}}$ defined in (6.4).

However, unlike typical behavior, it turns out that large deviations, i.e., $\{\lambda_1(Z) \geq (1+\delta)\lambda_\alpha^{\text{heavy}}\}$, is achieved by the emergence of a clique with high edge-weights on it. One can first estimate the probability that $\mathcal{G}_{N,\frac{d}{N}}$, contains a clique of size k for each integer $k \geq 2$. Next, one estimates the probability that the edge-weights on the clique are high. One of our main results is that we identify how to induce these high edge-weights in the most efficient way, which turns out to involve the variational problem (6.5).

The optimal clique size depends on both α and δ. In particular, it is 2 when $\alpha \leq 1$, i.e., large deviations is dictated by the existence of an atypically large edge-weight.

Again as before, while the proof of the upper bound involves several technical ingredients making the above heuristics precise, the lower bound is obtained rather quickly by planting a clique of an appropriate size with high edge weights.

We end this section with a brief overview of the variational problem described in (6.5) which we believe is of independent interest.

$\alpha-$norm generalization of the Motzkin-Straus theorem

When the edge-weights are given by heavy-tailed Weibull random variables (with shape parameter $0 < \alpha < 2$), we study the spectral behavior by relying on a new result relating the largest eigenvalue and the entrywise L^α-(quasi)norm of any symmetric matrix A, which we define by

$$\|A\|_\alpha := \Big(\sum_{i,j} |a_{ij}|^\alpha\Big)^{1/\alpha}.$$

In particular we show that for any $0 < \alpha < 2$ and any integer $k \geq 2$, there exists an explicit and *sharp* constant $C(\alpha, k) > 0$ such that for any network $G = (V, E, A)$ with maximum clique size k,

$$\lambda_1(A) \leq C(\alpha, k)\, \|A\|_\alpha\,. \tag{6.12}$$

The constant $C(\alpha, k)$ is expressed in terms of the function $\phi_\alpha(k)$ defined in (6.5) (see Proposition 2.1 for details). Moreover, it turns out that $C(\alpha, k)$ does not depend on k when $0 < \alpha \leq 1$. The special case $\alpha = 2$ had been studied earlier where the bound counterpart to (6.12) reads as

$$\lambda_1(A) \leq \sqrt{\frac{k-1}{k}}\, \|A\|_2\,, \tag{6.13}$$

and can be obtained as a straightforward consequence of the Motzkin-Straus theorem (see [GN22, Proposition 3.1]).

We finish this discussion with a brief comparison of the heavy tailed case to the Gaussian case, which corresponds to the case of Weibull random variables with $\alpha = 2$, albeit with additional polynomial pre-factors as described in Remark 6.2. For the latter, it is shown in [GN22] that if we set

$$\bar{\psi}_\delta(k) := \frac{k(k-3)}{2} + \frac{1+\delta}{2}\frac{k}{k-1} \tag{6.14}$$

for integers $k \geq 2$, then

$$\lim_{N\to\infty} -\frac{1}{\log N} \log \mathbb{P}\Big(\lambda_1(Z) \geq \sqrt{2(1+\delta)\log N}\Big) = \min_{k=2,3,\cdots} \bar{\psi}_\delta(k). \tag{6.15}$$

While, $\lim_{\delta\to\infty} \arg\min_{k\in\mathbb{N}_{\geq 2}} \bar{\psi}_\delta(k) = \infty$, it turns out that when $1 < \alpha < 2$, the function $\phi_{\beta/2}(k)$ that appears in (6.6) becomes constant for large k (see (4) in Lemma 2.3), which implies that $\arg\min_{k\in\mathbb{N}_{\geq 2}} \psi_{\alpha,\delta}(k)$ stays bounded in δ. Further, our proof indicates that the optimal way for the largest eigenvalue to be greater than $(1+\delta)\lambda_\alpha$ is to have a clique of size $\arg\min_{k\in\mathbb{N}_{\geq 2}} \psi_{\alpha,\delta}(k)$ (in particular, 2, which is just a single edge, when $0 < \alpha \leq 1$) in the random graph $X = \mathcal{G}_{N,\frac{d}{N}}$, and then have high valued edge-weights on this clique. A counterpart result in the Gaussian case was proven in [GN22], where it was shown that the corresponding clique size is given by $\arg\min_{k\in\mathbb{N}_{\geq 2}} \bar{\psi}_\delta(k)$. Thus, by the above discussion, the governing clique size for $\alpha < 2$ stays bounded as the deviation increases in contrast to the Gaussian case where the same goes to infinity.

In the two remaining sections of this chapter we prove the large deviation theorems for the cases of light- and heavy-tailed edge-weights in Section 6.2 and 6.3 respectively.

6.2 Light-tailed weights

We consider $\alpha > 2$ in this section and prove Theorems 6.1 and 6.2. For the reader's benefit let us recall that

$$\lambda_\alpha^{\text{light}} = 2^{\frac{1}{\alpha}} \alpha^{-\frac{1}{2}} (\alpha - 2)^{\frac{1}{2} - \frac{1}{\alpha}} \frac{(\log N)^{\frac{1}{2}}}{(\log \log N)^{\frac{1}{2} - \frac{1}{\alpha}}}. \tag{6.16}$$

For notational brevity, in this section we will denote $\lambda_\alpha^{\text{light}}$ simply by λ_α. To simplify the notation further, we set

$$B_\alpha := 2^{\frac{1}{\alpha}} \alpha^{-\frac{1}{2}} (\alpha - 2)^{\frac{1}{2} - \frac{1}{\alpha}}. \tag{6.17}$$

The upper tail

Let us recall the theorem that we will prove in this section.

Theorem 6.1. For any $\delta > 0$,

$$\lim_{N \to \infty} -\frac{\log \mathbb{P}\left(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha^{\text{light}}\right)}{\log N} = (1+\delta)^2 - 1.$$

The governing structure for the upper tail of $\lambda_1(Z)$ will turn out to be a star of degree $\lceil \gamma_\delta \frac{\log N}{\log \log N} \rceil$ with

$$\gamma_\delta := (1+\delta)^2 \left(1 - \frac{2}{\alpha}\right) \tag{6.18}$$

and high edge-weights on the edges. This stems from the following optimization problem. The probability that the maximum of the largest eigenvalue among all the typically present $n^{1-\gamma}$ stars of degree $\lceil \gamma \frac{\log N}{\log \log N} \rceil$ (see e.g., Lemma 3.3 and Proposition 3.4) is greater than $(1+\delta)\lambda_\alpha$ is maximized at $\gamma = \gamma_\delta$.

Lower bound for the upper tail

The strategy will change depending on whether γ_δ is less or greater than 1.

Case 1: $\gamma_\delta < 1$. For small enough $\rho > 0$, we condition on the event $\mathcal{A}_{\gamma_\delta, \rho}$ measurable with respect to X, defined in Proposition 3.2. Conditioned on that event, there exist $m := \left\lceil \frac{1}{4} N^{1-\gamma_\delta - \rho} \right\rceil$ vertices with $g(\gamma_\delta) = \left\lceil \gamma_\delta \frac{\log N}{\log \log N} \right\rceil$ disjoint neighbors with no edges between each neighbors. Denote by $S_1, \cdots, S_m$ the vertex-disjoint stars induced by these vertices and their $g(\gamma)$ neighbors. Then, by Lemma 2.5,

$$\lambda_1(Z) \geq \max_{k=1,\cdots,m} \lambda_1(S_k).$$

Thus, conditioned on the event $\mathcal{A}_{\gamma_\delta,\rho}$, by the characterization of the largest eigenvalue of stars in Lemma 2.6,

$$\mathbb{P}\Big(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha \mid X\Big) \geq \mathbb{P}\left(\max_{k=1,\dots,m} \sum_{(i,j)\in E(S_k)} Z_{ij}^2 \geq (1+\delta)^2\lambda_\alpha^2 \mid X\right) \tag{6.19}$$

(recall that (i,j) denotes the *undirected* edge joining vertices i and j with $i<j$). In the appendix, we derive a tail bound (4.14) for the sum of squares of Weibull random variables. Plugging in the bound with $d=1+\delta$ and $b=\gamma_\delta$, under the event $\mathcal{A}_{\gamma_\delta,\rho}$, for each $k=1,\cdots,m$,

$$\mathbb{P}\left(\sum_{(i,j)\in E(S_k)} Z_{ij}^2 \geq (1+\delta)^2\lambda_\alpha^2 \mid X\right) \geq N^{-(1+\delta)^\alpha \frac{2}{\alpha-2}(1-\frac{2}{\alpha})^{\frac{\alpha}{2}}\gamma_\delta^{1-\frac{\alpha}{2}}+o(1)} = N^{-(1+\delta)^2\frac{2}{\alpha}+o(1)},$$

where we used $\gamma_\delta = (1+\delta)^2\left(1-\frac{2}{\alpha}\right)$ in the last equality. Using the independence of edge-weights and recalling $m = \left\lceil \frac{1}{4}N^{1-\gamma_\delta-\rho}\right\rceil$, under the event $\mathcal{A}_{\gamma_\delta,\rho}$,

$$\begin{aligned}\mathbb{P}\left(\max_{k=1,\dots,m} \sum_{(i,j)\in E(S_k)} Z_{ij}^2 \geq (1+\delta)^2\lambda_\alpha^2 \mid X\right) &\geq 1-\left(1-N^{-(1+\delta)^2\frac{2}{\alpha}+o(1)}\right)^m \\ &\geq 1-e^{-N^{1-\gamma_\delta-\rho-(1+\delta)^2\frac{2}{\alpha}+o(1)}} \geq \frac{1}{2}N^{1-(1+\delta)^2-\rho+o(1)},\end{aligned}$$

where we used the fact that $1-e^{-x}\geq \frac{1}{2}x$ for small $x>0$ in the last inequality.

Therefore, applying this to (6.19) and using that $\mathbb{P}(\mathcal{A}_{\gamma_\delta,\rho})\geq\frac{1}{2}$ (see Proposition 3.2),

$$\mathbb{P}\Big(\lambda_1(Z)\geq(1+\delta)\lambda_\alpha\Big) \geq \mathbb{E}\left[\mathbb{P}\Big(\lambda_1(Z)\geq(1+\delta)\lambda_\alpha \mid X\Big)1_{\mathcal{A}_{\gamma_\delta,\rho}}\right] \geq \frac{1}{4}N^{1-(1+\delta)^2-\rho+o(1)}.$$

Due to the arbitrariness of $\rho>0$,

$$\limsup_{N\to\infty} -\frac{\log\mathbb{P}\Big(\lambda_1(Z)\geq(1+\delta)\lambda_\alpha\Big)}{\log N} \leq (1+\delta)^2-1.$$

Case 2: $\gamma_\delta \geq 1$. For any $\rho>0$, we condition on the event $\mathcal{A}'_{(1+\rho)\gamma_\delta}$ measurable with respect to X, defined in Proposition 3.3. Under this event there exists a vertex v with $\left\lceil(1+\rho)\gamma_\delta\frac{\log N}{\log\log N}\right\rceil$ neighbors with no edges between them. Denote by S the star induced by v and these neighbors. As before, using the tail bound (4.14) with $d=1+\delta$ and $b=(1+\rho)\gamma_\delta$, under the event $\mathcal{A}'_{(1+\rho)\gamma_\delta}$,

$$\mathbb{P}\Big(\lambda_1(Z)\geq(1+\delta)\lambda_\alpha \mid X\Big) \geq \mathbb{P}\left(\sum_{(i,j)\in E(S)} Z_{ij}^2 \geq (1+\delta)^2\lambda_\alpha^2 \mid X\right)$$

$$\geq N^{-(1+\delta)^{\alpha}\frac{2}{\alpha-2}(1-\frac{2}{\alpha})^{\frac{\alpha}{2}}((1+\rho)\gamma_\delta)^{1-\frac{\alpha}{2}}+o(1)} = N^{-(1+\delta)^2\frac{2}{\alpha}(1+\rho)^{1-\frac{2}{\alpha}}+o(1)}.$$

By Proposition 3.3,

$$\mathbb{P}\left(\mathcal{A}'_{(1+\rho)\gamma_\delta}\right) = N^{1-(1+\rho)\gamma_\delta+o(1)}.$$

Thus, as above,

$$\begin{aligned}\limsup_{N\to\infty} \frac{-\log \mathbb{P}\big(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha\big)}{\log N} &\leq -1 + (1+\rho)\gamma_\delta + (1+\delta)^2\frac{2}{\alpha}(1+\rho)^{1-\frac{2}{\alpha}} \\ &= -1 + (1+\delta)^2\Big[(1+\rho)\Big(1-\frac{2}{\alpha}\Big) + \frac{2}{\alpha}(1+\rho)^{1-\frac{2}{\alpha}}\Big].\end{aligned}$$

Since $\rho > 0$ is arbitrary, this implies the desired bound.

□

Remark 6.3. If $\gamma_\delta = 1$, then the precise behavior of the number of vertices of degree close to $\gamma_\delta \frac{\log N}{\log\log N} = \frac{\log N}{\log\log N}$ or the probability of the existence of such a vertex is somewhat delicate to track. Hence, we considered vertices with a slightly larger degree instead to exploit the large deviation bound for atypically large degrees from (3.3). This shortens the proof by not dealing with the case $\gamma_\delta = 1$ separately.

Upper bound of the upper tail

We proceed in a sequence of steps:

1. We first truncate the weights Y and then accordingly decompose Z into $Z^{(1)} + Z^{(2)}$:

$$Z^{(1)}_{ij} = X_{ij}Y^{(1)}_{ij} \quad \text{and} \quad Z^{(2)}_{ij} = X_{ij}Y^{(2)}_{ij}, \tag{6.20}$$

where

$$Y^{(1)}_{ij} = Y_{ij}\mathbf{1}_{|Y_{ij}|>(\varepsilon\log\log N)^{\frac{1}{\alpha}}} \quad \text{and} \quad Y^{(2)}_{ij} = Y_{ij}\mathbf{1}_{|Y_{ij}|\leq(\varepsilon\log\log N)^{\frac{1}{\alpha}}}. \tag{6.21}$$

Similarly, write $X = X^{(1)} + X^{(2)}$ with

$$X^{(1)}_{ij} = X_{ij}\mathbf{1}_{|Y_{ij}|>(\varepsilon\log\log N)^{\frac{1}{\alpha}}} \quad \text{and} \quad X^{(2)}_{ij} = X_{ij}\mathbf{1}_{|Y_{ij}|\leq(\varepsilon\log\log N)^{\frac{1}{\alpha}}}. \tag{6.22}$$

This particular threshold is chosen so that, as we will soon see, $Z^{(2)}$ is spectrally negligible.

By the tail decay of Weibull random variables, $X^{(1)}$ is distributed as $\mathcal{G}_{N,q}$ with

$$q \leq \frac{d'}{N(\log N)^{\varepsilon}} \tag{6.23}$$

for some constant $d' > 0$. Also, given $X^{(1)}$, the edge-weights on the network $Z^{(1)}$ can be regarded as i.i.d. Weibull random variables conditioned to be greater than $(\varepsilon\log\log N)^{\frac{1}{\alpha}}$ in absolute value.

2. We analyze the component structure of $X^{(1)}$, the underlying graph of the network $Z^{(1)}$. The sparsity of $X^{(1)}$ allows the results in Section 3.5 to be applicable. In particular, connected components of $X^{(1)}$ are tree-like (i.e. the number of tree-excess edges are small) and their sizes are relatively small with high probability.

3. We further decompose the network $Z^{(1)}$ into $Z_1^{(1)}$, consisting of vertex-disjoint weighted stars, and $Z_2^{(1)}$, whose degrees are well-controlled.

4. Using the results in Step (2) and the fact that the maximal degree in $Z_2^{(1)}$ is relatively small, we prove that $Z_2^{(1)}$ is spectrally negligible as well.

5. We analyze the spectral contribution of $Z_1^{(1)}$ by grouping stars according to their degrees. Since we have a complete characterization of the largest eigenvalue of a (weighted) star graph (see Lemma 2.6), one can explicitly compute the contribution from the collections of stars of a given degree. It turns out that the main contribution, which leads to the large deviation probability, comes from the stars of degree close to $\gamma_\delta \frac{\log N}{\log\log N}$ (see (6.18)).

Given the truncation in Step (1) above, we first estimate $\lambda_1(Z^{(2)})$.

Lemma 6.1. For any $\delta, \varepsilon > 0$,

$$\liminf_{N\to\infty} \frac{-\log \mathbb{P}\Big(\lambda_1\left(Z^{(2)}\right) \geq \varepsilon^{\frac{1}{\alpha}}(1+\delta)\frac{\lambda_\alpha}{B_\alpha}\Big)}{\log N} \geq (1+\delta)^2 - 1. \tag{6.24}$$

We start by stating the following theorem from [BBG21]. While the latter covers a varied range of values for p, we state it only in the sparsity regime considered in our results. Recall that we set $t_n = \frac{\log N}{\log\log N}$.

Theorem 6.5 ([BBG21, Thm. 1.1])**.** For any $\delta > 0$,

$$\lim_{N\to\infty} \frac{-\log \mathbb{P}(\lambda_1(X) \geq (1+\delta)t_N^{\frac{1}{2}})}{\log N} = (1+\delta)^2 - 1. \tag{6.25}$$

Proof of Lemma 6.1. Since $\left|Y_{ij}^{(2)}\right| \leq (\varepsilon \log\log N)^{\frac{1}{\alpha}}$ for all i, j,

$$\lambda_1(Z^{(2)}) \leq (\varepsilon \log\log N)^{\frac{1}{\alpha}} \lambda_1(X). \tag{6.26}$$

By Theorem 6.5 and recalling $\frac{\lambda_\alpha}{B_\alpha} = \frac{(\log N)^{\frac{1}{2}}}{(\log\log N)^{\frac{1}{2}-\frac{1}{\alpha}}}$ (see (6.16) and (6.17)), this immediately concludes the proof.

□

Now we analyze $Z^{(1)}$, the main spectral part for $\lambda_1(Z)$. Since the edge density of its underlying graph $X^{(1)} \overset{d}{\sim} \mathcal{G}_{N,q}$ satisfies (6.23), the results in Section 3.5 hold for $X^{(1)}$ as indicated in Step (2).

We now implement Step (3), i.e. we decompose $X^{(1)}$ into two parts, one of which consists of a *vertex-disjoint* union of stars while the other one has a relatively small maximum degree. This remainder graph and the union of stars are not necessarily vertex-disjoint. The latter part will be spectrally negligible and the dominating factor will be the former part. For this we rely on a result from [BBG21, Lemma 3.5], that we simplified slightly for our setting. We moreover state an explicit bound on the maximum degree, which can be deduced easily from the proof.

Lemma 6.2 ([BBG21, Lemma 3.5]). There exists an event $\mathcal{W}$ measurable with respect to $X^{(1)}$ that happens with probability at least [2] $1 - e^{-\omega(\log N)}$ under which $X^{(1)}$ can be decomposed into a graph $X_1^{(1)}$ which is a vertex-disjoint union of stars, and a graph $X_2^{(1)}$ whose maximum degree is bounded by $2\,(\log N)^{\frac{7}{16}}$ for large enough N.

The decomposition in [BBG21, Lemma 3.5] is stated for $\mathcal{G}_{N,p}$ with $p = O(\frac{1}{N})$ which is applicable for $X^{(1)}$ by (6.23).

From now on, we condition on the high probability event $\mathcal{W}$. Let $Z_1^{(1)}$ and $Z_2^{(1)}$ be the corresponding networks of $X_1^{(1)}$ and $X_2^{(1)}$ respectively. We will first focus on the spectral behavior of $Z_2^{(1)}$ by analyzing its underlying graph $X_2^{(1)}$. Then by Lemmas 3.11, 3.12 and 6.2, each connected component C_ℓ of $X_2^{(1)}$ satisfies the following properties with high probability for any $\delta_1, \delta_2 > 0$:

1. $d_1(C_\ell) \leq 2\,(\log N)^{\frac{7}{16}}$,
2. $|V(C_\ell)| \leq (1+\delta_1)\frac{1}{\varepsilon}\frac{\log N}{\log\log N}$,
3. $|E(C_\ell)| \leq |V(C_\ell)| + \delta_2$

(recall that for any graph G, $d_1(G)$ denotes the maximum degree of G).

We now state the following key general proposition, which claims that it is costly that any connected network satisfying the above three properties has the largest eigenvalue of order λ_α (recall that λ_α denotes the quantity which turns out to be the typical largest eigenvalue, as defined in (6.16)).

Proposition 6.1. Assume that $\alpha > 2$ and $\{u_N\}_{N\geq 1}$ be a sequence such that $u_N = o\left(\frac{\log N}{\log\log N}\right)$. For any $N \in \mathbb{N}$ and $\varepsilon, \delta_1, \delta_2 > 0$, let $\mathcal{G} := \mathcal{G}_{N,\varepsilon,\delta_1,\delta_2}$ be the set of connected networks $G = (V, E, A)$ ($A = (a_{ij})_{i,j\in V}$ denotes the weight matrix) such that

1. $d_1(G) \leq u_N$,

[2] Here, the quantity $x = \omega(\log N)$ means that $\lim_{n\to\infty} \frac{x}{\log N} = \infty$.

2. $|V| \leq (1+\delta_1)\frac{1}{\varepsilon}\frac{\log N}{\log\log N}$,

3. $|E| \leq |V| + \delta_2$.

Assume that the edge-weights are i.i.d. Weibull random variables with a shape parameter $\alpha > 2$ conditioned to be greater than $(\varepsilon \log\log N)^{\frac{1}{\alpha}}$ in absolute value. Then, for any constant $c > 0$,

$$\lim_{N\to\infty} \frac{-\log \sup_{G\in\mathcal{G}} \mathbb{P}\left(\lambda_1(A) \geq c\lambda_\alpha\right)}{\log N} = \infty.$$

Proof. The general strategy is to bound $\lambda_1(A)$ by expressing it in terms of the corresponding (random) top eigenvector and then analyzing the contributions from the high and low values of the entries separately. To make this precise, suppose that $V = [m]$, and let $f = (f_i)_{i\in[m]}$ with $\|f\|_2 = 1$ be any (random) eigenvector of A such that

$$\lambda_1(A) = f^T A f = \sum_{i,j=1}^{m} a_{ij} f_i f_j = 2\sum_{(i,j)\in E} a_{ij} f_i f_j \tag{6.27}$$

(recall that (i,j) denotes the *undirected* edge joining two vertices $i < j$).

For a constant $\xi \in \left(0, \frac{1}{2}\right)$ which will be chosen sufficiently small later, define

$$V_S := \{i \in [m] : |f_i| < |\xi|\}, \qquad V_L := \{i \in [m] : |f_i| \geq |\xi|\}, \tag{6.28}$$

where the indices stand for *small* and *large* respectively. Since $\sum_{i=1}^m f_i^2 = 1$, we have

$$|V_L| \leq \left\lfloor \frac{1}{\xi^2} \right\rfloor. \tag{6.29}$$

We also partition the set of edges into two parts, those that are incident on a vertex in V_L and the rest:

$$E_S := \{(i,j) \in E : i < j, i, j \in V_S\}, \qquad E_L := \{(i,j) \in E : i < j, i \in V_L \text{ or } j \in V_L\}. \tag{6.30}$$

We now decompose the summation in (6.27) into two parts λ_S and λ_L:

$$\lambda_1(A) = 2\sum_{(i,j)\in E} a_{ij} f_i f_j = 2\sum_{(i,j)\in E_S} a_{ij} f_i f_j + 2\sum_{(i,j)\in E_L} a_{ij} f_i f_j =: 2\lambda_S + 2\lambda_L. \tag{6.31}$$

The above is expressed in a way such that both λ_S and λ_L can be bounded by sums of i.i.d. random variables which will be convenient.

Thus, for any constant $0 \leq \tau \leq 1$ which will be chosen later,

$$\mathbb{P}(\lambda_1(A) \geq c\lambda_\alpha) \leq \mathbb{P}(2\lambda_S \geq \tau c\lambda_\alpha) + \mathbb{P}(2\lambda_L \geq (1-\tau)c\lambda_\alpha). \tag{6.32}$$

We now analyze these two probabilities separately.

Bounding λ_S. We apply the Cauchy-Schwarz inequality to $\sum_{(i,j)\in E_S} a_{ij} f_i f_j$, and then use a bound on $\sum_{(i,j)\in E_S} f_i^2 f_j^2$ which we now derive. Let T be a spanning tree of G (recall that G is connected) and $E(T)$ be the collection of edges in T. Since $|E(T)| = |V| - 1$, by our assumption on the number of tree-excess edges, $|E_S \setminus E(T)| \le |E| - |E(T)| \le \delta_2 + 1$. Hence,

$$\sum_{(i,j)\in E_S} f_i^2 f_j^2 = \sum_{(i,j)\in E_S \cap E(T)} f_i^2 f_j^2 + \sum_{(i,j)\in E_S \setminus E(T)} f_i^2 f_j^2 \le \xi^2 + (\delta_2 + 1)\xi^4 \le (2 + \delta_2)\xi^2, \tag{6.33}$$

where we used (2.18) with $\theta = 1$ to bound the first term (note that $|f_i| \le \xi < \frac{1}{2}$). Thus, setting

$$\tau := (2 + \delta_2)^{\frac{1}{4}} \xi^{\frac{1}{2}}, \tag{6.34}$$

by the Cauchy-Schwarz inequality together with the bound (6.33),

$$\lambda_S \le \Big(\sum_{(i,j)\in E_S} f_i^2 f_j^2 \Big)^{\frac{1}{2}} \Big(\sum_{(i,j)\in E_S} a_{ij}^2 \Big)^{\frac{1}{2}} \le \tau^2 \Big(\sum_{(i,j)\in E_S} a_{ij}^2 \Big)^{\frac{1}{2}} \le \tau^2 \Big(\sum_{(i,j)\in E} a_{ij}^2 \Big)^{\frac{1}{2}}.$$

By assumptions on the network size and the number of tree-excess edges, $\sum_{(i,j)\in E} a_{ij}^2$ is the sum of at most $\big\lfloor (1 + \delta_1) \frac{1}{\varepsilon} \frac{\log N}{\log \log N} + \delta_2 \big\rfloor$ many squares of Weibull random variables conditioned to be greater than $(\varepsilon \log \log N)^{\frac{1}{\alpha}}$ in absolute value. Hence, by the tail estimate for such sum of squares (the bound (4.15) with $d = \frac{c}{2\tau} = \frac{c}{2(2+\delta_2)^{\frac{1}{4}} \xi^{\frac{1}{2}}}$ and $b = \frac{1+\delta_1}{\varepsilon}$),

$$\begin{aligned} \liminf_{N\to\infty} \frac{-\log \mathbb{P}(2\lambda_S \ge \tau c \lambda_\alpha)}{\log N} &\ge \liminf_{N\to\infty} \frac{-\log \mathbb{P}\Big(\sum_{(i,j)\in E} a_{ij}^2 \ge \frac{c^2}{4\tau^2} \lambda_\alpha^2 \Big)}{\log N} \\ &\ge \frac{c^\alpha}{2^\alpha (2 + \delta_2)^{\frac{\alpha}{4}} \xi^{\frac{\alpha}{2}}} \frac{2}{\alpha - 2} \Big(1 - \frac{2}{\alpha}\Big)^{\frac{\alpha}{2}} \Big(\frac{1 + \delta_1}{\varepsilon}\Big)^{1 - \frac{\alpha}{2}} - (1 + \delta_1). \end{aligned} \tag{6.35}$$

Bounding λ_L. By the Cauchy-Schwarz inequality and the fact that

$$\sum_{(i,j)\in E_L} f_i^2 f_j^2 \le \Big(\sum_{i\in V} f_i^2 \Big) \Big(\sum_{j\in V} f_j^2 \Big) = 1,$$

we have

$$\lambda_L \le \Big(\sum_{(i,j)\in E_L} f_i^2 f_j^2 \Big)^{\frac{1}{2}} \Big(\sum_{(i,j)\in E_L} a_{ij}^2 \Big)^{\frac{1}{2}} \le \Big(\sum_{(i,j)\in E_L} a_{ij}^2 \Big)^{\frac{1}{2}}.$$

Since $|V_L| \le \left\lfloor \frac{1}{\xi^2} \right\rfloor$ by (6.29), the event $\{2\lambda_L \ge (1-\tau)c\lambda_\alpha\}$ implies the existence of a random subset $J \subseteq V$ with $|J| \le \left\lfloor \frac{1}{\xi^2} \right\rfloor$ such that

$$\sum_{(i,j)\in E_J} a_{ij}^2 \ge \frac{(1-\tau)^2 c^2}{4} \lambda_\alpha^2,$$

where $E_J := \{(i,j) \in E : i < j, i \in J \text{ or } j \in J\}$. For any deterministic subset J' with $|J'| \le \left\lfloor \frac{1}{\xi^2} \right\rfloor$, by our assumption on the maximum degree, $|E_{J'}| = o\left(\frac{1}{\xi^2} \frac{\log N}{\log\log N}\right)$. Thus, by the tail probability estimate for the sum of squares of Weibull random variables (the bound (4.16) with $d = \frac{(1-\tau)c}{2}$),

$$\lim_{N\to\infty} \frac{-\log \mathbb{P}\left(\sum_{(i,j)\in E_{J'}} a_{ij}^2 \ge \frac{(1-\tau)^2 c^2}{4} \lambda_\alpha^2\right)}{\log N} = \infty.$$

By the assumption on the component size, the cardinality of different values that a random subset J with $|J| \le \left\lfloor \frac{1}{\xi^2} \right\rfloor$ can take is bounded by $\left((1+\delta_1)\frac{1}{\varepsilon}\frac{\log N}{\log\log N}\right)^{\frac{1}{\xi^2}} = n^{o(1)}$. Thus, by a union bound,

$$\lim_{N\to\infty} \frac{-\log \mathbb{P}(2\lambda_L \ge (1-\tau)c\lambda_\alpha)}{\log N} = \infty. \tag{6.36}$$

Therefore, applying (6.35) and (6.36) to (6.32),

$$\liminf_{N\to\infty} \frac{-\log \mathbb{P}(\lambda_1(A) \ge c\lambda_\alpha)}{\log N} \ge \frac{c^\alpha}{2^\alpha (2+\delta_2)^{\frac{\alpha}{4}} \xi^{\frac{\alpha}{2}}} \frac{2}{\alpha - 2}\left(1 - \frac{2}{\alpha}\right)^{\frac{\alpha}{2}} \left(\frac{1+\delta_1}{\varepsilon}\right)^{1-\frac{\alpha}{2}} - (1+\delta_1).$$

Since $\xi > 0$ is arbitrary, the RHS above can be made arbitrarily large, concluding the proof. □

Since each connected component of $X_2^{(1)}$ satisfies the conditions in Proposition 6.1 by the discussion following Lemma 6.2 with high probability, and the number of components is bounded by n, a union bound completes Step (4). Therefore, it remains to analyze the spectral behavior of $Z_1^{(1)}$, a collection of disjoint stars. We will group these stars according to their sizes and then show that the main contribution comes from the group of stars with degrees close to $\gamma_\delta \frac{\log N}{\log\log N}$.

As a preparation, we now introduce some notations and a few lemmas. The first lemma concerns the spectral behaviour of a single weighted star. Recall from (3.4) that we set $g(\gamma) = \left\lceil \gamma \frac{\log N}{\log\log N} \right\rceil$ and let us define, for a star graph S, $d(S)$ to be the degree of the root vertex of S.

Lemma 6.3. Suppose that S is a weighted star graph such that $d(S) \le g(\gamma)$ for some $\gamma > 0$, with i.i.d. weights given by Weibull random variables with a shape parameter $\alpha > 2$ conditioned to be greater than $(\varepsilon \log\log N)^{\frac{1}{\alpha}}$ in absolute value. Then, for any $\rho > 0$,

$$\liminf_{N\to\infty} \frac{-\log \mathbb{P}(\lambda_1(S) \ge (1+\rho)\lambda_\alpha)}{\log N} \ge (1+\rho)^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}} \gamma^{1-\frac{\alpha}{2}} - \varepsilon\gamma.$$

Proof. Let $\left\{\tilde{Y}_i\right\}_{i=1,2,\cdots}$ be i.i.d. Weibull random variables with a shape parameter $\alpha > 2$ conditioned to be greater than $(\varepsilon \log\log N)^{\frac{1}{\alpha}}$ in absolute value. Since the largest eigenvalue of a weighted star is nothing other than the square root of the sum of squares of edge-weights (see Lemma 2.6),

$$\mathbb{P}\Big(\lambda_1(S) \ge (1+\rho)\lambda_\alpha\Big) \le \mathbb{P}\left(\tilde{Y}_1^2 + \cdots + \tilde{Y}_{g(\gamma)}^2 \ge (1+\rho)^2\lambda_\alpha^2\right).$$

By the tail bound (4.15) with $d = 1+\rho$ and $b = \gamma$, we are done.

□

Next, we estimate the spectral contribution from the group of stars with degree close to $g(\gamma)$. For this we first introduce some additional notations. Let $d\left(X^{(1)}, v\right)$ be the degree of v in the graph $X^{(1)}$, and for $\gamma \ge 0$, define

$$D_\gamma^{(1)} = \Big\{v \in V : d\Big(X^{(1)}, v\Big) \ge g(\gamma)\Big\}. \tag{6.37}$$

For small enough constant $\kappa > 0$ which will be chosen later, define m to be an integer such that $m\kappa < 1 \le (m+1)\kappa$. Then, define the event measurable with respect to $X^{(1)}$:

$$\mathcal{P}_\kappa := \Big\{\Big|D_{i\kappa}^{(1)}\Big| \le N^{1-i\kappa+\kappa} \text{ for all } i = 0, 1, \cdots, m\Big\}, \tag{6.38}$$

which guarantees that for the discretization $\{\kappa, 2\kappa, \ldots, m\kappa\}$ of the interval $(0,1)$, there are not unusually many vertices whose degrees fall into any bin of degree range given by the discretization.

Additionally we define the event measurable with respect to $X^{(1)}$:

$$\mathcal{L}_{\delta,\kappa} := \left\{\Big|D_{1+\kappa}^{(1)}\Big| \le \frac{(1+\delta)^2}{\kappa}\right\}, \tag{6.39}$$

which guarantees that there are uniformly bounded number of vertices of unusually large degree.

Using the estimate for the contribution of a single star (Lemma 6.3), we now prove a lemma that captures the contribution from the group of stars of degree close to $g(\gamma)$. Recall from Lemma 6.2 that $X_1^{(1)}$ is a vertex-disjoint union of stars, and $\kappa > 0$ is a given constant.

Lemma 6.4. Let $\mathcal{S}$ be the collection of stars in $X_1^{(1)}$. Moreover, define, for any $h > 0$ and $\gamma = i\kappa < 1$ (i is a non-negative integer) or $\gamma \geq 1 + \kappa$,

$$\lambda_{\max}(\gamma, h) := \max_{S \in \mathcal{S}, d(S) \in (g(\gamma), g(\gamma+h)]} \{\lambda_1(S)\} \tag{6.40}$$

if there is a star $S \in \mathcal{S}$ satisfying $d(S) \in (g(\gamma), g(\gamma + h)]$, and set $\lambda_{\max}(\gamma, h) := 0$ otherwise. Then, for any $\rho > 0$,

$$\begin{aligned} \liminf_{N\to\infty} & \frac{-\log \mathbb{E}\Big[\mathbb{P}\Big(\lambda_{\max}(\gamma, h) \geq (1+\rho)\lambda_\alpha \mid X^{(1)}\Big)\mathbf{1}_{\mathcal{P}_\kappa \cap \mathcal{L}_{\delta,\kappa}}\Big]}{\log N} \\ & \geq -f_{\alpha,\rho}(\gamma + h) - h - \kappa - \varepsilon(\gamma + h), \end{aligned} \tag{6.41}$$

where the function $f_{\alpha,\rho} : (0, \infty) \to \mathbb{R}$ is defined by

$$f_{\alpha,\rho}(x) := 1 - x - (1+\rho)^\alpha \frac{2}{\alpha - 2}\left(1 - \frac{2}{\alpha}\right)^{\frac{\alpha}{2}} x^{1-\frac{\alpha}{2}}. \tag{6.42}$$

In (6.41), the quantity $f_{\alpha,\rho}(\gamma + h)$ should be thought of as the dominant term with the rest being error terms.

Proof. The proof depends on whether $\gamma < 1$ or $\gamma \geq 1 + \kappa$. In the case $\gamma = i\kappa < 1$, the number of stars in $X_1^{(1)}$ of degree at least $g(\gamma)$ is bounded by $N^{1-\gamma+\kappa}$ under the event $\mathcal{P}_\kappa$. Hence, by union bound and Lemma 6.3 with $\gamma + h$ in place of γ,

$$\begin{aligned} \liminf_{N\to\infty} & \frac{-\log \mathbb{E}\Big[\mathbb{P}\Big(\lambda_{\max}(\gamma, h) \geq (1+\rho)\lambda_\alpha \mid X^{(1)}\Big)\mathbf{1}_{\mathcal{P}_\kappa}\Big]}{\log N} \\ & \geq -(1 - \gamma + \kappa) + (1+\rho)^\alpha \frac{2}{\alpha - 2}\left(1 - \frac{2}{\alpha}\right)^{\frac{\alpha}{2}} (\gamma + h)^{1-\frac{\alpha}{2}} - \varepsilon(\gamma + h) \\ & = -f_{\alpha,\rho}(\gamma + h) - h - \kappa - \varepsilon(\gamma + h). \end{aligned}$$

Let us now consider the second case $\gamma \geq 1 + \kappa$. Note that by Lemma 6.3, for any star S with $d(S) \leq g(\gamma + h)$,

$$\mathbb{P}(\lambda_1(S) \geq (1+\rho)\lambda_\alpha) \leq N^{-(1+\rho)^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}}(\gamma+h)^{1-\frac{\alpha}{2}} + \varepsilon(\gamma+h) + o(1)}.$$

Thus, since $\lambda_{\max}(\gamma, h) = 0$ if there is no star S in $\mathcal{S}$ satisfying $d(S) > g(\gamma)$, we have

$$\begin{aligned} & \mathbb{E}\Big[\mathbb{P}\Big(\lambda_{\max}(\gamma, h) \geq (1+\rho)\lambda_\alpha \mid X^{(1)}\Big)\mathbf{1}_{\mathcal{L}_{\delta,\kappa}}\Big] \\ & = \mathbb{E}\Big[\mathbb{P}\Big(\lambda_{\max}(\gamma, h) \geq (1+\rho)\lambda_\alpha \mid X^{(1)}\Big)\mathbf{1}_{\mathcal{L}_{\delta,\kappa}}\mathbf{1}_{\{\exists S \in \mathcal{S} : d(S) > g(\gamma)\}}\Big] \\ & \leq \frac{(1+\delta)^2}{\kappa} \cdot N^{-(1+\rho)^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}}(\gamma+h)^{1-\frac{\alpha}{2}} + \varepsilon(\gamma+h) + o(1)} \mathbb{P}(d_1(X^{(1)}) \geq g(\gamma)), \end{aligned}$$

where the last inequality follows from the union bound (under the event $\mathcal{L}_{\delta,\kappa}$, the number of stars in $X_1^{(1)}$ of degree at least $g(\gamma) \geq g(1+\kappa)$ is bounded by $\frac{(1+\delta)^2}{\kappa}$). Also, since $\gamma > 1$, by (3.3) and the fact that $X^{(1)}$ is distributed as $\mathcal{G}_{N,q}$ with $q \leq p = \frac{d}{N}$,

$$\mathbb{P}(d_1(X^{(1)}) \geq g(\gamma)) \leq N^{1-\gamma+o(1)}.$$

Therefore,

$$\begin{aligned}
&\liminf_{N\to\infty} \frac{-\log \mathbb{E}\Big[\mathbb{P}\Big(\lambda_{\max}(\gamma,h) \geq (1+\rho)\lambda_\alpha \mid X^{(1)}\Big)\mathbf{1}_{\mathcal{L}_{\delta,\kappa}}\Big]}{\log N} \\
&\quad \geq -(1-\gamma) + (1+\rho)^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}} (\gamma+h)^{1-\frac{\alpha}{2}} - \varepsilon(\gamma+h) \\
&\quad = -f_{\alpha,\rho}(\gamma+h) - h - \varepsilon(\gamma+h).
\end{aligned}$$

□

Having the expression for the contribution of any group of stars under the assumption that the underlying graph is reasonably nice, we now identify the group of stars for which this contribution is maximized. This is done in the following technical lemma by optimizing the value of the function $f_{\alpha,\rho}$. Note that $f_{\alpha,\rho}$ was originally defined for $\rho > 0$, but below we consider the wider range $\rho > -1$ for a later application.

Lemma 6.5. For $\alpha > 2$ and $\rho > -1$, recall the function $f_{\alpha,\rho} : (0,\infty) \to \mathbb{R}$ in (6.42):

$$f_{\alpha,\rho}(\gamma) = 1 - \gamma - (1+\rho)^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}} \gamma^{1-\frac{\alpha}{2}}.$$

Then,

$$\max_{\gamma>0} f_{\alpha,\rho}(\gamma) = 1-(1+\rho)^2 \quad \text{and} \quad \gamma_\rho := \arg\max_{\gamma>0} f_{\alpha,\rho}(\gamma) = (1+\rho)^2\left(1-\frac{2}{\alpha}\right). \tag{6.43}$$

Proof. For the sake of readability, we will drop the subscripts of $f_{\alpha,\rho}$ in the proof. Note that

$$\frac{d}{d\gamma} f(\gamma) = -1 + (1+\rho)^\alpha \left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}} \gamma^{-\frac{\alpha}{2}},$$

and thus f is maximized at $\gamma = (1+\rho)^2\left(1-\frac{2}{\alpha}\right)$. Plugging this back into $f(\gamma)$, we get $\max_{\gamma>0} f(\gamma) = 1-(1+\rho)^2$.

□

We are now ready to put all of this together to prove the upper bound of the upper tail.

Proof of the upper bound of the upper tail. By using the decomposition (6.20)-(6.22), we write $Z = Z^{(1)} + Z^{(2)}$. Since $\lambda_1(Z) \le \lambda_1\big(Z^{(1)}\big) + \lambda_1\big(Z^{(2)}\big)$ by Lemma 2.5, we have

$$\mathbb{P}(\lambda_1(Z) \ge (1+\delta)\lambda_\alpha) \le \mathbb{P}\Big(\lambda_1(Z^{(1)}) \ge (1+\delta)\Big(1 - \frac{\varepsilon^{\frac{1}{\alpha}}}{B_\alpha}\Big)\lambda_\alpha\Big) + \mathbb{P}\Big(\lambda_1(Z^{(2)}) \ge \varepsilon^{\frac{1}{\alpha}}(1+\delta)\frac{\lambda_\alpha}{B_\alpha}\Big). \tag{6.44}$$

By Lemma 6.1, the second term above can be bounded by

$$\mathbb{P}\Big(\lambda_1\Big(Z^{(2)}\Big) \ge \varepsilon^{\frac{1}{\alpha}}(1+\delta)\frac{\lambda_\alpha}{B_\alpha}\Big) \le N^{1-(1+\delta)^2+o(1)}. \tag{6.45}$$

Hence, it suffices to bound the probability

$$\mathbb{P}\Big(\lambda_1(Z^{(1)}) \ge (1+\delta)\Big(1 - \frac{\varepsilon^{\frac{1}{\alpha}}}{B_\alpha}\Big)\lambda_\alpha\Big). \tag{6.46}$$

Step 1. Given the previous results, we will work on the event ensuring:

1. Existence of the decomposition of $X^{(1)}$ into $X_1^{(1)}$ (vertex-disjoint union of stars) and $X_2^{(1)}$ (relatively small maximum degree).
2. All connected components of $X_2^{(1)}$ are relatively small and tree-like.
3. $X_1^{(1)}$ has a controlled number of stars of each given degree.

The first condition is achieved by the event $\mathcal{W}$ in Lemma 6.2, and the second one is fulfilled by the series of events in Section 3.5 (applied to $X^{(1)} \overset{\mathrm{d}}{\sim} \mathcal{G}_{N,q}$). For the last condition, we consider the events $\mathcal{P}_\kappa$ and $\mathcal{L}_{\delta,\kappa}$ in (6.38) and (6.39) respectively.
The events above make up the event $\mathcal{K}_0$ measurable with respect to $X^{(1)}$:

$$\mathcal{K}_0 := \mathcal{W} \cap \mathcal{D}_{(1+\delta)^2-1} \cap \mathcal{C}_{\varepsilon,(1+\delta)^2-1} \cap \mathcal{E}_{(1+\delta)^2-1} \cap \mathcal{P}_\kappa \cap \mathcal{L}_{\delta,\kappa}. \tag{6.47}$$

Using the previously proven or cited results, we have

$$\begin{aligned}
&\mathbb{P}(\mathcal{W}^c) \le e^{-\omega(\log N)} \quad \text{by Lemma 6.2,}\\
&\mathbb{P}(\mathcal{D}^c_{(1+\delta)^2-1}) \le N^{1-(1+\delta)^2+o(1)} \quad \text{by Lemma 3.10,}\\
&\mathbb{P}(\mathcal{C}^c_{\varepsilon,(1+\delta)^2-1}) \le N^{1-(1+\delta)^2+o(1)} \quad \text{by Lemma 3.11,}\\
&\mathbb{P}(\mathcal{E}^c_{(1+\delta)^2-1}) \le N^{1-(1+\delta)^2+o(1)} \quad \text{by Lemma 3.12,}\\
&\mathbb{P}(\mathcal{P}^c_\kappa) \le N^{-(1+\delta)^2} \quad \text{by Proposition 3.4 with } \mu = \frac{(1+\delta)^2}{\kappa},\\
&\mathbb{P}(\mathcal{L}^c_{\delta,\kappa}) \le N^{-(1+\delta)^2+o(1)} \quad \text{by (3.3).}
\end{aligned} \tag{6.48}$$

Note that Proposition 3.4 and (3.3) were proven for the random graph $\mathcal{G}_{N,\frac{d}{N}}$, whereas the events $\mathcal{P}_\kappa$ and $\mathcal{L}_{\delta,\kappa}$ are defined in terms of the sparser graph $X^{(1)} \overset{\mathrm{d}}{\sim} \mathcal{G}_{N,q}$. However, since these events are decreasing, the same bounds hold by monotonicity. Combining these implies

$$\mathbb{P}(\mathcal{K}_0^c) \le N^{1-(1+\delta)^2+o(1)}. \tag{6.49}$$

Since $\lambda_1(\mathrm{Z}^{(1)}) \le \lambda_1(\mathrm{Z}_1^{(1)}) + \lambda_1(\mathrm{Z}_2^{(1)})$, defining δ', for sufficiently small ε (depending on α and δ), as

$$(1+\delta)\Big(1-\frac{\varepsilon^{\frac{1}{\alpha}}}{B_\alpha}\Big) = (1+\delta') + \varepsilon(1+\delta), \tag{6.50}$$

we have

$$\begin{aligned}\mathbb{P}\left(\lambda_1\Big(\mathrm{Z}^{(1)}\Big) \ge (1+\delta)\left(1-\frac{\varepsilon^{\frac{1}{\alpha}}}{B_\alpha}\right)\lambda_\alpha\right) &\le \mathbb{E}\left[\mathbb{P}\left(\lambda_1\Big(\mathrm{Z}_1^{(1)}\Big) \ge (1+\delta')\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{K}_0}\right] \\ &+ \mathbb{E}\left[\mathbb{P}\left(\lambda_1\Big(\mathrm{Z}_2^{(1)}\Big) \ge \varepsilon(1+\delta)\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{K}_0}\right] + \mathbb{P}(\mathcal{K}_0^c).\end{aligned} \tag{6.51}$$

From now on, we estimate the quantities in the RHS above.

Step 2. Contribution from $\mathbf{Z_2^{(1)}}$. Under the event $\mathcal{C}_{\varepsilon,(1+\delta)^2-1} \cap \mathcal{E}_{(1+\delta)^2-1}$, and hence under the event $\mathcal{K}_0$, each connected component of $X^{(1)}$, and thus of its subgraph $X_2^{(1)}$, satisfies the conditions in Proposition 6.1 with $\delta_1 = \delta_2 = (1+\delta)^2 - 1$ (recall that the largest degree of $X_2^{(1)}$ is $o(\frac{\log N}{\log\log N})$). Since the number of connected components is bounded by N, by Proposition 6.1 combined with a union bound, whenever the event $\mathcal{K}_0$ holds,

$$\lim_{N\to\infty} \frac{-\log \mathbb{P}(\lambda_1(\mathrm{Z}_2^{(1)}) \ge \varepsilon(1+\delta)\lambda_\alpha \mid X^{(1)})}{\log N} = \infty. \tag{6.52}$$

Step 3. Contribution from $\mathbf{Z_1^{(1)}}$. Let M be the smallest integer such that $(1+\delta)^2 < M\kappa$. Note that this in particular implies that

$$M \le \frac{(1+\delta)^2}{\kappa} + 1. \tag{6.53}$$

Note that under the event $\mathcal{D}_{(1+\delta)^2-1}$, $d_1(X_1^{(1)}) \le d_1(X^{(1)}) \le (1+\delta)^2 \frac{\log N}{\log\log N}$. Thus, the degree of any star S in $X_1^{(1)}$ falls into in one of the following (not necessarily disjoint) categories:

1. $d(S) \le g(\kappa)$.
2. $d(S) \in (g(i\kappa), g((i+2)\kappa)]$ for $i = 1, \cdots, M-1$ and $i \ne m+1$ (recall that m is a unique integer such that $m\kappa < 1 \le (m+1)\kappa$),

Reiterating Remark 6.3, the reason why we exclude $i = m+1$ is essentially because we do not have a precise understanding on the number of vertices of degree close to $\frac{\log N}{\log\log N}$.

To bound the contribution of the first category, by Lemma 6.3 (with $\gamma = \kappa$ and $\rho = \delta'$) combined with a union bound, for sufficiently small $\kappa, \varepsilon > 0$,

$$\liminf_{N\to\infty} \frac{-\log \mathbb{E}\Big[\mathbb{P}\Big(\max_{\substack{S\in\mathcal{S}\\ d(S)\le g(\kappa)}} \{\lambda_1(S)\} \ge (1+\delta')\lambda_\alpha \mid X^{(1)}\Big)\mathbf{1}_{\mathcal{K}_0}\Big]}{\log N}$$
$$\ge -1 + (1+\delta')^\alpha \frac{2}{\alpha-2}\Big(1-\frac{2}{\alpha}\Big)^{\frac{\alpha}{2}} \kappa^{1-\frac{\alpha}{2}} - \varepsilon\kappa > (1+\delta)^2 - 1. \tag{6.54}$$

Note that the additional '-1' term in the middle quantity comes from a union bound (the number of stars is bounded by N), and the last inequality holds once $\kappa > 0$ is sufficiently small (recall that $\alpha > 2$).

For the group of stars in the second category, by Lemma 6.4 (with $\gamma = i\kappa$, $h = 2\kappa$ and $\rho = \delta'$), for each $i = 1, \cdots, M-1$ with $i \neq m+1$,

$$\liminf_{N\to\infty} \frac{-\log \mathbb{E}\Big[\mathbb{P}\Big(\max_{\substack{S\in\mathcal{S}\\ d(S)\in(g(i\kappa), g((i+2)\kappa)]}} \{\lambda_1(S)\} \ge (1+\delta')\lambda_\alpha \mid X^{(1)}\Big)\mathbf{1}_{\mathcal{K}_0}\Big]}{\log N}$$
$$\ge -f_{\alpha,\delta'}((i+2)\kappa) - 2\kappa - \kappa - \varepsilon(i+2)\kappa$$
$$\ge (1+\delta')^2 - 1 - 3\kappa - \varepsilon(M+1)\kappa \overset{(6.53)}{\ge} (1+\delta')^2 - 1 - 3\kappa - \varepsilon\Big(\frac{(1+\delta)^2}{\kappa} + 2\Big)\kappa =: L, \tag{6.55}$$

where we used Lemma 6.5 to bound $f_{\alpha,\delta'}((i+2)\kappa)$ in the second inequality.

Since the categories considered in (6.54) and (6.55) make up the total contribution of the stars in the network $Z^{(1)}$, by a union bound,

$$\liminf_{N\to\infty} \frac{-\log \mathbb{E}[\mathbb{P}(\lambda_1(Z_1^{(1)}) \ge (1+\delta')\lambda_\alpha) \mid X^{(1)})\mathbf{1}_{\mathcal{K}_0}]}{\log N} \ge \min\{(1+\delta)^2 - 1, L\}. \tag{6.56}$$

Applying the bounds (6.49), (6.52) and (6.56) to (6.51),

$$\liminf_{N\to\infty} \frac{-\log \mathbb{P}\Big(\lambda_1(Z^{(1)}) \ge (1+\delta)\Big(1-\frac{\varepsilon^{\frac{1}{\alpha}}}{B_\alpha}\Big)\lambda_\alpha\Big)}{\log N} \ge \min\{(1+\delta)^2 - 1, L\}.$$

Applying this with (6.45) to (6.44), we obtain

$$\liminf_{N\to\infty} \frac{-\log \mathbb{P}(\lambda_1(Z) \ge (1+\delta)\lambda_\alpha)}{\log N} \ge \min\{(1+\delta)^2 - 1, L\}.$$

Since $\lim_{\varepsilon\to 0} \delta' = \delta$ (see (6.50) for the definition of δ'), the quantity L defined in (6.55) becomes sufficiently close to $(1+\delta)^2 - 1$ for small enough $\varepsilon, \kappa > 0$, which completes the proof.

□

Remark 6.4. The above proof indicates that the following structural result holds: Conditioned on the upper tail event $\{\lambda_1(Z) \geq (1+\delta)\lambda_\alpha\}$, with high probability, X contains a star of size roughly $\gamma_\delta \frac{\log N}{\log\log N}$ with edge-weights greater than $(\frac{2}{\alpha-2}\log\log N)^{\frac{1}{\alpha}}$ in absolute value. Indeed, since $f_{\alpha,\rho}(\gamma)$ at $\gamma = \gamma_\rho$ is a strict maximum, by Lemma 6.4, the contribution from the stars of degree $g(\gamma)$ with $\gamma \notin (\gamma_\delta - \chi, \gamma_\delta + \chi)$ for $\chi > 0$ is negligible compared to that from the stars of degree $g(\gamma_\delta)$.

This, combined with Remark A.1 in the Appendix, which says that if the sum of squares of light-tailed random variables is large then these random variables tend to be uniformly large, implies that with high probability conditionally on the upper tail event, there exists a star of degree close to $\gamma_\delta \frac{\log N}{\log\log N}$ with edge-weights close to

$$\left(\frac{(1+\delta)^2\lambda_\alpha^2}{\gamma_\delta \frac{\log N}{\log\log N}}\right)^{\frac{1}{2}} = \Big(\frac{2}{\alpha-2}\log\log N\Big)^{\frac{1}{\alpha}}$$

in absolute value. In other words, the optimal size of the star increases in δ, whereas the edge-weights on the star, while atypically large, do not depend on the amount of deviation δ.

The lower tail

We now move on to prove a large deviation result for the lower tail, that we restate here for the reader's convenience.

Theorem 6.2. For any $0 < \delta < 1$,

$$\lim_{N\to\infty} \frac{1}{\log N}\left(\log\log \frac{1}{\mathbb{P}\left(\lambda_1(Z) \leq (1-\delta)\lambda_\alpha^{\text{light}}\right)}\right) = 1-(1-\delta)^2.$$

Analogous to the upper tail case, the governing structure in this case will turn out to be the collection of $N^{1-\gamma_\delta'}$ vertex-disjoint stars of degree close to $\gamma_\delta' \frac{\log N}{\log\log N}$ with

$$\gamma_\delta' := (1-\delta)^2\Big(1-\frac{2}{\alpha}\Big). \tag{6.57}$$

Note that γ_δ' is nothing but $\gamma_{-\delta}$ from (6.18).

Lower bound for the lower tail

Before embarking on the proof, we establish a lemma about the lower tail behavior of the maximum among the largest eigenvalue of $N^{1-\gamma+o(1)}$ weighted stars of degree close to $g(\gamma) = \left\lceil \gamma \frac{\log N}{\log\log N}\right\rceil$.

Lemma 6.6. Suppose that $\gamma > 0, h \geq 0$ and $\kappa \geq \gamma - 1$. Let $\mathcal{S}$ be a collection of at most $N^{1-\gamma+\kappa}$ vertex-disjoint weighted stars of size less than $g(\gamma + h)$. Assume that edge-weights are i.i.d. Weibull random variables with a shape parameter $\alpha > 2$ conditioned to be greater than $(\varepsilon \log\log N)^{\frac{1}{\alpha}}$ in absolute value. Then, for any $0 < \rho < 1$,

$$\limsup_{N\to\infty} \frac{1}{\log N}\left(\log\log \frac{1}{\mathbb{P}\left(\max_{S\in\mathcal{S}}\{\lambda_1(S)\} \leq (1-\rho)\lambda_\alpha\right)}\right) \leq f_{\alpha,-\rho}(\gamma+h)+\kappa+h+\varepsilon(\gamma+h), \quad (6.58)$$

where the function $f_{\alpha,-\rho}$ is as defined in (6.42):

$$f_{\alpha,-\rho}(x) = 1 - x - (1-\rho)^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}} x^{1-\frac{\alpha}{2}}.$$

Again, as before, above $\kappa + h + \varepsilon(\gamma + h)$ should be thought of as an error term.

Proof. We use the notation $\left\{\tilde{Y}_i\right\}_{i=1,2,\cdots}$ from the proof of Lemma 6.3. By Lemma 2.6,

$$\mathbb{P}\Big(\lambda_1(S) \geq (1-\rho)\lambda_\alpha\Big) \leq \mathbb{P}\left(\tilde{Y}_1^2 + \cdots + \tilde{Y}_{g(\gamma+h)}^2 \geq (1-\rho)^2\lambda_\alpha^2\right).$$

By the tail estimate (4.15) with $d = 1-\rho$ and $b = \gamma + h$, this probability is upper bounded by

$$N^{\varepsilon(\gamma+h)-(1-\rho)^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}}(\gamma+h)^{1-\frac{\alpha}{2}}+o(1)},$$

Thus, using that the number of stars in $\mathcal{S}$ is bounded by $N^{1-\gamma+\kappa}$, by the independence of edge-weights,

$$\begin{aligned}\mathbb{P}\left(\max_{S\in\mathcal{S}}\{\lambda_1(S)\} \leq (1-\rho)\lambda_\alpha\right) &\geq \left(1 - N^{\varepsilon(\gamma+h)-(1-\rho)^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}}(\gamma+h)^{1-\frac{\alpha}{2}}+o(1)}\right)^{N^{1-\gamma+\kappa}} \\ &\geq \exp\left(-N^{1-\gamma+\kappa+\varepsilon(\gamma+h)-(1-\rho)^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}}(\gamma+h)^{1-\frac{\alpha}{2}}+o(1)}\right) \\ &= \exp\left(-N^{f_{\alpha,-\rho}(\gamma+h)+\kappa+h+\varepsilon(\gamma+h)+o(1)}\right),\end{aligned}$$

where the second inequality follows since $1 - x > e^{-2x}$ for small $x > 0$ and the constant can be absorbed into $N^{o(1)}$. □

Proof of the lower bound of the lower tail.

Step 1. Using the decomposition (6.20)-(6.22), we write $Z = Z^{(1)} + Z^{(2)}$. First, we define the event measurable with respect to X:

$$\mathcal{B}_\delta := \left\{\lambda_1(X) \leq (1+\delta)\frac{(\log N)^{\frac{1}{2}}}{(\log\log N)^{\frac{1}{2}}}\right\}. \quad (6.59)$$

As before, conditioning on the event $\mathcal{W}$ defined in Lemma 6.2, allows us to decompose $Z^{(1)}$ into $Z_1^{(1)}$ (vertex-disjoint union of stars) and $Z_2^{(1)}$ (relatively small maximum degree). Let

$$\mathcal{R}_\kappa := \{\text{Maximum degree in } X^{(1)} \text{ is less than } g(1+\kappa)\}.$$

We now define an event similar to (6.47) by additionally excluding the existence of an atypically large degree vertex using the above event, by defining the event $\mathcal{K}_1$ which is measurable with respect to $\{X, X^{(1)}\}$:

$$\mathcal{K}_1 := \mathcal{B}_\delta \cap \mathcal{W} \cap \mathcal{C}_{\varepsilon,(1+\delta)^2-1} \cap \mathcal{E}_{(1+\delta)^2-1} \cap \mathcal{P}_\kappa \cap \mathcal{R}_\kappa. \tag{6.60}$$

By Theorem 6.5, $\lim_{N\to\infty} \mathbb{P}(\mathcal{B}_\delta) = 1$. Also, by (3.3), $\lim_{N\to\infty} \mathbb{P}(\mathcal{R}_\kappa) = 1$ (note that although (3.3) is stated for the random graph $\mathcal{G}_{N,\frac{d}{N}}$, since $\mathcal{R}_\kappa$ is a decreasing event and $X^{(1)}$ is sparser than $\mathcal{G}_{N,\frac{d}{N}}$, we still have this estimate). Together with the analysis in (6.48), we have

$$\lim_{N\to\infty} \mathbb{P}(\mathcal{K}_1) = 1. \tag{6.61}$$

Since $\lambda_1(Z) \le \lambda_1\big(Z_1^{(1)}\big) + \lambda_1\big(Z_2^{(1)}\big) + \lambda_1\big(Z^{(2)}\big)$, setting δ' via

$$1-\delta = (1-\delta') + (1+\delta)\left(\varepsilon + \frac{\varepsilon^{\frac{1}{\alpha}}}{B_\alpha}\right), \tag{6.62}$$

we have

$$\begin{aligned}\mathbb{P}\Big(\lambda_1(Z) \le (1-\delta)\lambda_\alpha\Big) \ge \mathbb{E}\Bigg[\ \mathbb{P}\Big(\ \lambda_1\big(Z_1^{(1)}\big) \le (1-\delta')\lambda_\alpha,\quad \lambda_1\big(Z_2^{(1)}\big) \le (1+\delta)\varepsilon\lambda_\alpha,\\ \lambda_1\big(Z^{(2)}\big) \le (1+\delta)\frac{\varepsilon^{\frac{1}{\alpha}}}{B_\alpha}\lambda_\alpha \ \Big|\ X, X^{(1)}\Big)\,\mathbf{1}_{\mathcal{K}_1}\ \Bigg].\end{aligned} \tag{6.63}$$

First of all, under the event $\mathcal{B}_\delta$, and hence under the event $\mathcal{K}_1$, by the same argument as in (6.26),

$$\lambda_1\big(Z^{(2)}\big) \le (1+\delta)\frac{\varepsilon^{\frac{1}{\alpha}}}{B_\alpha}\lambda_\alpha. \tag{6.64}$$

Furthermore, note that $Z_1^{(1)}, Z_2^{(1)}$ and X are conditionally independent given $X^{(1)}$. Thus by (6.64), under the event $\mathcal{K}_1$, the conditional probability inside the expectation in (6.63) is written as

$$\begin{aligned}&\mathbb{P}\left(\lambda_1\big(Z_1^{(1)}\big) \le (1-\delta')\lambda_\alpha, \lambda_1\big(Z_2^{(1)}\big) \le (1+\delta)\varepsilon\lambda_\alpha \mid X, X^{(1)}\right)\\ &= \mathbb{P}\left(\lambda_1\big(Z_1^{(1)}\big) \le (1-\delta')\lambda_\alpha, \lambda_1\big(Z_2^{(1)}\big) \le (1+\delta)\varepsilon\lambda_\alpha \mid X^{(1)}\right)\end{aligned}$$

$$= \mathbb{P}\left(\lambda_1\left(Z_1^{(1)}\right) \le (1-\delta')\lambda_\alpha \mid X^{(1)}\right)\mathbb{P}\left(\lambda_1\left(Z_2^{(1)}\right) \le (1+\delta)\varepsilon\lambda_\alpha \mid X^{(1)}\right). \tag{6.65}$$

Therefore, by (6.63)-(6.65),

$$\mathbb{P}\Big(\lambda_1(Z) \le (1-\delta)\lambda_\alpha\Big) \ge \mathbb{E}\Big[\mathbb{P}\left(\lambda_1\left(Z_1^{(1)}\right) \le (1-\delta')\lambda_\alpha \mid X^{(1)}\right)\mathbb{P}\left(\lambda_1\left(Z_2^{(1)}\right) \le (1+\delta)\varepsilon\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{K}_1}\Big]. \tag{6.66}$$

We now estimate the two conditional probabilities above.

Step 2. Contribution from $Z_2^{(1)}$. Note that, by definition, under the event $\mathcal{K}_1$, all components of $X_2^{(1)}$ as well as of $X^{(1)}$ satisfy the properties described in the events $\mathcal{C}_{\varepsilon,(1+\delta)^2-1}$ and $\mathcal{E}_{(1+\delta)^2-1}$. Hence by Proposition 6.1 together with a union bound (the number of connected components is bounded by N), for large enough n, under the event $\mathcal{C}_{\varepsilon,(1+\delta)^2-1} \cap \mathcal{E}_{(1+\delta)^2-1}$, and hence under the event $\mathcal{K}_1$,

$$\mathbb{P}\left(\lambda_1\left(Z_2^{(1)}\right) \le (1+\delta)\varepsilon\lambda_\alpha \mid X^{(1)}\right) \ge \frac{1}{2}. \tag{6.67}$$

Step 3. Contribution from $Z_1^{(1)}$. We proceed by considering groups of stars of similar degrees. Let $\mathcal{S}$ be the collection of stars in its underlying graph $X_1^{(1)}$ given the latter, by construction, is a vertex-disjoint union of stars. We define the events capturing the contributions of the small and large stars, by

$$\mathcal{J}_0 := \left\{ \max_{\substack{S\in\mathcal{S} \\ d(S)\le g(\kappa)}} \{\lambda_1(S)\} \le (1-\delta')\lambda_\alpha \right\},$$

$$\mathcal{J}_m := \left\{ \max_{\substack{S\in\mathcal{S} \\ d(S)\in(g(m\kappa),g((m+2)\kappa)]}} \{\lambda_1(S)\} \le (1-\delta')\lambda_\alpha \right\},$$

and the events capturing the contribution of the stars with intermediate degree, i.e. for $i = 1,\cdots,m-1$ (recall that m is an integer such that $m\kappa < 1 \le (m+1)\kappa$), by

$$\mathcal{J}_i := \left\{ \max_{\substack{S\in\mathcal{S} \\ d(S)\in(g(i\kappa),g((i+1)\kappa)]}} \{\lambda_1(S)\} \le (1-\delta')\lambda_\alpha \right\}.$$

Since $(m+2)\kappa > 1+\kappa$, under the event $\mathcal{R}_\kappa$ and thus under $\mathcal{K}_1$, there are no stars in $X_1^{(1)}$ of degree at least $g((m+2)\kappa)$. Thus, conditioned on $\mathcal{K}_1$, the event $\bigcap_{i=0}^m \mathcal{J}_i$ implies $\lambda_1\left(Z_1^{(1)}\right) \le (1-\delta')\lambda_\alpha$.

We will now lower bound the probabilities of the events $\mathcal{J}_i$, conditioned on $\mathcal{K}_1$. By Lemma 6.6 with $\rho = \delta', \gamma = \kappa$ and $h = 0$ (the number of stars in $X_1^{(1)}$ is bounded by n), under the event $\mathcal{K}_1$,

$$\mathbb{P}\left(\mathcal{J}_0 \middle| X^{(1)}\right) \ge \exp\left(-n^{f_{\alpha,-\delta'}(\kappa)+\kappa+\varepsilon\kappa+o(1)}\right).$$

We will lower bound the probability of the remaining events. Recall that under the event $\mathcal{P}_\kappa$, and thus under the event $\mathcal{K}_1$, for $\gamma = i\kappa$ with $i = 1, \cdots, m$, the number of stars in $X_1^{(1)}$ of degree at least $g(\gamma)$ is bounded by $N^{1-\gamma+\kappa}$. Thus, by Lemma 6.6 with $\rho = \delta', \gamma = i\kappa$ and $h = \kappa$ or 2κ, under the event $\mathcal{K}_1$,

$$\mathbb{P}\left(\mathcal{J}_i \middle| X^{(1)}\right) \geq \exp\left(-N^{f_{\alpha,-\delta'}((i+1)\kappa)+2\kappa+\varepsilon((i+1)\kappa)+o(1)}\right) \quad \text{and}$$
$$\mathbb{P}\left(\mathcal{J}_m \middle| X^{(1)}\right) \geq \exp\left(-N^{f_{\alpha,-\delta'}((m+2)\kappa)+3\kappa+\varepsilon((m+2)\kappa)+o(1)}\right).$$

Now note that the events $\mathcal{J}_i$ are conditionally independent given $X^{(1)}$. Since $m\kappa < 1$ and $f_{\alpha,-\delta'}(\gamma) \leq 1-(1-\delta')^2$ for any $\gamma > 0$ by Lemma 6.5, all exponents of N in the above lower bounds for $\mathbb{P}\left(\mathcal{J}_i \middle| X^{(1)}\right)$ with $i = 0, \cdots, m$ are less than

$$1-(1-\delta')^2 + 3\kappa + \varepsilon(1+2\kappa) + o(1).$$

Thus, whenever the event $\mathcal{K}_1$ holds,

$$\begin{aligned}
\mathbb{P}\left(\lambda_1\left(Z_1^{(1)}\right) \leq (1-\delta')\lambda_\alpha \middle| X^{(1)}\right) &\geq \mathbb{P}\left(\bigcap_{i=0}^{m} \mathcal{J}_i \middle| X^{(1)}\right) = \prod_{i=0}^{m} \mathbb{P}\left(\mathcal{J}_i \middle| X^{(1)}\right) \\
&\geq \left(\exp\left(-N^{1-(1-\delta')^2+3\kappa+\varepsilon(1+2\kappa)+o(1)}\right)\right)^{m+2} \geq \exp\Big(-\Big(\frac{1}{\kappa}+2\Big)N^{1-(1-\delta')^2+3\kappa+\varepsilon(1+2\kappa)+o(1)}\Big),
\end{aligned} \tag{6.68}$$

where we used $m \leq \frac{1}{\kappa}$ in the last inequality.

Therefore, applying (6.67) and (6.68) to (6.66),

$$\mathbb{P}\left(\lambda_1(Z) \leq (1-\delta)\lambda_\alpha\right) \geq \frac{1}{2}\exp\Big(-\Big(\frac{1}{\kappa}+2\Big)N^{1-(1-\delta')^2+3\kappa+\varepsilon(1+2\kappa)+o(1)}\Big)\mathbb{P}(\mathcal{K}_1).$$

Since $\mathbb{P}(\mathcal{K}_1) \geq \frac{1}{2}$ for large enough n (see (6.61)) and $\lim_{\varepsilon\to 0}\delta' = \delta$ (see (6.62) for the definition of δ'), by taking $\kappa, \varepsilon > 0$ small enough, we establish the desired bound. ☐

We now move on to the final part of our analysis of light-tailed weights.

Upper bound for the lower tail

We show it is unlikely that all stars induced by vertices of degree close to $g(\gamma'_\delta) = \lceil \gamma'_\delta \frac{\log N}{\log\log N} \rceil$, where $\gamma'_\delta = (1-\delta)^2\left(1-\frac{2}{\alpha}\right)$ was defined in (6.57), have a largest eigenvalue less than $(1-\delta)\lambda_\alpha$.

To prove this, we condition, for small enough $\rho > 0$, on the event $\mathcal{A}_{\gamma'_\delta,\rho}$ defined in Proposition 3.2, i.e. there exist $m := \left\lceil \frac{1}{4}N^{1-\gamma'_\delta-\rho} \right\rceil$ vertices having $g(\gamma'_\delta)$ disjoint neighbors with no edges between each neighbors. Let $S_1, \cdots, S_m$ be the vertex-disjoint stars induced by these vertices and their $g(\gamma'_\delta)$ neighbors. By Proposition 3.2,

$$\mathbb{P}\left(\mathcal{A}^c_{\gamma'_\delta,\rho}\right) \leq e^{-N^{1-\gamma'_\delta-\rho+o(1)}}. \tag{6.69}$$

Since by Lemma 2.5

$$\lambda_1^2(Z) \geq \max_{k=1,\cdots,m} \lambda_1^2(S_k) = \max_{k=1,\cdots,m} \sum_{(i,j)\in E(S_k)} Z_{ij}^2,$$

we have

$$\mathbb{P}\Big(\lambda_1(Z) \leq (1-\delta)\lambda_\alpha\Big) \leq \mathbb{E}\Big[\mathbb{P}\Big(\max_{k=1,\cdots,m} \sum_{(i,j)\in E(S_k)} Z_{ij}^2 \leq (1-\delta)^2\lambda_\alpha^2 \mid X\Big)\mathbf{1}_{\mathcal{A}_{\gamma'_\delta,\rho}}\Big] + \mathbb{P}(\mathcal{A}^c_{\gamma'_\delta,\rho}). \quad (6.70)$$

Using the tail estimate (4.14) with $d = 1-\delta$ and $b = \gamma'_\delta$, under the event $\mathcal{A}_{\gamma'_\delta,\rho}$,

$$\begin{aligned}
&\mathbb{P}\Big(\max_{k=1,\cdots,m} \sum_{(i,j)\in E(S_k)} Z_{ij}^2 \leq (1-\delta)^2\lambda_\alpha^2 \mid X\Big) \\
&\leq \Big(1 - N^{-(1-\delta)^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}}(\gamma'_\delta)^{1-\frac{\alpha}{2}}+o(1)}\Big)^m \\
&\leq \exp(-N^{1-\gamma'_\delta-\rho-(1-\delta)^\alpha \frac{2}{\alpha-2}\left(1-\frac{2}{\alpha}\right)^{\frac{\alpha}{2}}(\gamma'_\delta)^{1-\frac{\alpha}{2}}+o(1)}) \leq \exp(-N^{1-(1-\delta)^2-\rho+o(1)}), \quad (6.71)
\end{aligned}$$

where we used $\gamma'_\delta = (1-\delta)^2\left(1-\frac{2}{\alpha}\right)$ to simplify the exponent. Since $\gamma'_\delta = (1-\delta)^2\left(1-\frac{2}{\alpha}\right) < (1-\delta)^2$, (6.69) and (6.71) show that the dominant term in (6.70) is $e^{-N^{1-(1-\delta)^2-\rho+o(1)}}$. By taking $\rho > 0$ sufficiently small enough, we obtain the matching upper bound.

□

6.3 Heavy-tailed weights

In this section, we prove Theorems 6.3 and 6.4. As before, for notational brevity, we define $\lambda_\alpha := \lambda_\alpha^{\text{heavy}} = (\log N)^{\frac{1}{\alpha}}$.

The upper tail

Let us first recall the theorem that we will prove in this section. Recall that for $\theta > 1$ and the integer $k \geq 2$, we defined the following function

$$\phi_\theta(k) = \sup_{f=(f_1,\cdots,f_k):\|f\|_1=1} \sum_{i,j\in[k],i\neq j} |f_i|^\theta |f_j|^\theta.$$

Theorem 6.3. Let $\delta > 0$.

1. In the case $1 < \alpha < 2$, let $\beta > 2$ be the conjugate of α (i.e. $\frac{1}{\alpha} + \frac{1}{\beta} = 1$). For an integer $k \geq 2$, define

$$\psi_{\alpha,\delta}(k) := \frac{k(k-3)}{2} + \frac{1}{2}(1+\delta)^\alpha \phi_{\beta/2}(k)^{1-\alpha}. \quad (6.6)$$

Then,

$$\lim_{N\to\infty} -\frac{\log \mathbb{P}\left(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha^{\text{heavy}}\right)}{\log N} = \min_{k=2,3,\cdots} \psi_{\alpha,\delta}(k). \tag{6.7}$$

2. In the case $0 < \alpha \leq 1$,

$$\lim_{N\to\infty} -\frac{\log \mathbb{P}\left(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha^{\text{heavy}}\right)}{\log N} = (1+\delta)^\alpha - 1. \tag{6.8}$$

Lower bound for the upper tail

As mentioned in the idea of proof section, we lower bound the large deviation probability by having high edge-weights on a suitable size of clique. In the case when $\alpha < 1$, it turns out to suffice to only consider a clique of size 2 (i.e. an edge). When $1 < \alpha < 2$ on the other hand, we also need to consider the possibility of larger cliques appearing in X. The lower bound then follows by optimizing over the clique size.

We first note that using $\phi_{\beta/2}(2) = 2^{1-\beta}$ (see (2.9) in Lemma 2.3),

$$\psi_{\alpha,\delta}(2) = -1 + \frac{1}{2}(1+\delta)^\alpha \phi_{\beta/2}(2)^{1-\alpha} = -1 + \frac{1}{2}(1+\delta)^\alpha \Big(\frac{1}{2^{\beta-1}}\Big)^{1-\alpha} = (1+\delta)^\alpha - 1, \tag{6.72}$$

where we used the conjugacy relation $\frac{1}{\alpha} + \frac{1}{\beta} = 1$ in the last identity.

We establish the lower bound by separately proving

$$\limsup_{N\to\infty} -\frac{\log \mathbb{P}(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha)}{\log N} \leq \psi_{\alpha,\delta}(2) = (1+\delta)^\alpha - 1 \tag{6.73}$$

and for any $k \geq 3$,

$$\limsup_{N\to\infty} -\frac{\log \mathbb{P}(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha)}{\log N} \leq \psi_{\alpha,\delta}(k). \tag{6.74}$$

Single large edge-weight. We consider the scenario that there is a large edge-weight, which provides the bound (6.73). Let us define the event that the number edges in the random graph X is not unusually small:

$$\mathcal{M} := \Big\{|E(X)| \geq \frac{d(N-1)}{4}\Big\}. \tag{6.75}$$

Then,

$$\mathbb{P}\Big(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha\Big) \geq \mathbb{E}\Big[\mathbb{P}\Big(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha \mid X\Big)\mathbf{1}_{\mathcal{M}}\Big]$$

Since the number of missing edges $\binom{N}{2} - E(X)$ has a distribution $\mathrm{Binom}\left(\binom{N}{2}, 1 - \frac{d}{N}\right)$, by lemmas 4.3 and 4.5 about the relative entropy and binomial tail estimates respectively, there is a constant $c > 0$ such that

$$\mathbb{P}\left(\mathcal{M}^c\right) = \mathbb{P}\left(\binom{N}{2} - |E(X)| > \binom{N}{2} - \frac{d(N-1)}{4}\right) \leq e^{-\binom{N}{2} I_{1-\frac{d}{N}}\left(1-\frac{d}{2N}\right)} \leq e^{-cN}. \tag{6.76}$$

Also, under the event $\mathcal{M}$, by the independence of edge-weights,

$$\mathbb{P}\left(\max_{(i,j)\in E(X)} |Z_{ij}| \geq (1+\delta)\lambda_\alpha \mid X\right) \geq 1 - \left(1 - C_1 N^{-(1+\delta)^\alpha}\right)^{\frac{d(N-1)}{4}} \geq N^{1-(1+\delta)^\alpha + o(1)}.$$

Combining these two bounds, we obtain (6.73). This already concludes the proof in the case $0 < \alpha \leq 1$.

Large edge-weights on a bigger clique. Next, we establish the lower bound (6.74) in the case $1 < \alpha < 2$. For any $k \geq 3$, using the result and notation from Lemma 2.1, take $k_1, k_2 \geq 0$ with $k_1 + k_2 \leq k$, $x, y \geq 0$ and the $k \times k$ matrix $A = (a_{ij})_{i,j\in[k]}$ given by

$$a_{ij} = \begin{cases} x^2 & i \neq j, i, j \in \{1, \cdots, k_1\} =: V_1, \\ y^2 & i \neq j, i, j \in \{k_1+1, \cdots, k_1+k_2\} =: V_2, \\ xy & i \in V_1, j \in V_2 \text{ or } i \in V_2, j \in V_1, \\ 0 & \text{otherwise,} \end{cases} \tag{6.77}$$

which achieves the equality in (2.2), i.e.

$$\lambda_1(A) = \phi_{\frac{\alpha}{2(\alpha-1)}}(k)^{\frac{\alpha-1}{\alpha}} \|A\|_\alpha = \phi_{\frac{\beta}{2}}(k)^{\frac{\alpha-1}{\alpha}} \|A\|_\alpha. \tag{6.78}$$

Conditioned on the event that X contains a clique of size k denoted by H, let $V(H) := \{v_1, \cdots, v_k\}$. Now consider the event that

$$Y_{v_i v_j} \geq \begin{cases} \frac{1}{\lambda_1(A)}(1+\delta)\lambda_\alpha x^2 & i \neq j, i, j \in V_1, \\ \frac{1}{\lambda_1(A)}(1+\delta)\lambda_\alpha y^2 & i \neq j, i, j \in V_2, \\ \frac{1}{\lambda_1(A)}(1+\delta)\lambda_\alpha xy & i \in V_1, j \in V_2 \text{ or } i \in V_2, j \in V_1, \\ 0 & \text{otherwise.} \end{cases} \tag{6.79}$$

By the distribution of the edge-weights as defined in (6.1), the conditional probability of this event is lower bounded by

$$\begin{aligned} N^{-(1+\delta)^\alpha \frac{1}{\lambda_1(A)^\alpha}\left(\binom{k_1}{2}x^{2\alpha} + \binom{k_2}{2}y^{2\alpha} + k_1 k_2 x^\alpha y^\alpha\right) + o(1)} &= N^{-(1+\delta)^\alpha \frac{\|A\|_\alpha^\alpha}{2\lambda_1(A)^\alpha} + o(1)} \\ &\overset{(6.78)}{=} N^{-\frac{1}{2}(1+\delta)^\alpha \phi_{\beta/2}(k)^{1-\alpha} + o(1)}. \end{aligned} \tag{6.80}$$

Also, under this event, we have $\lambda_1(Z) \geq \lambda_1(Z|_H) \geq (1+\delta)\lambda_\alpha$, since $Z|_H$ is entrywise greater than or equal to the matrix $\frac{(1+\delta)\lambda_\alpha}{\lambda_1(A)}A$, having non-negative entries, whose largest eigenvalue is $(1+\delta)\lambda_\alpha$. By Lemma 3.9 the probability that X contains a clique of size $k \geq 3$ is lower bounded by $CN^{-\binom{k}{2}+k}$. Therefore, combining this with (6.80), for any $k \geq 3$,

$$\limsup_{N\to\infty} -\frac{\log \mathbb{P}(\lambda_1(Z) \geq (1+\delta)\lambda_\alpha)}{\log N} \leq \binom{k}{2} - k + \frac{1}{2}(1+\delta)^\alpha \phi_{\beta/2}(k)^{1-\alpha} = \psi_{\alpha,\delta}(k). \tag{6.81}$$

□

Upper bound for the upper tail

As in the light-tailed case, we decompose $Z = Z^{(1)} + Z^{(2)}$ with a negligible part $Z^{(2)}$. However, the analysis of $Z^{(1)}$ will be significantly different since the governing structures will be distinct.

We first present a counterpart of Lemma 6.1. Noting that $t_N^{\frac{1}{\alpha}} \geq t_N^{\frac{1}{2}}$ for $\alpha < 2$, the proof is almost identical, so we omit it.

Lemma 6.7. For $\delta > 0$,

$$\liminf_{N\to\infty} \frac{-\log \mathbb{P}(\lambda_1(Z^{(2)}) \geq \varepsilon^{\frac{1}{\alpha}}(1+\delta)\lambda_\alpha)}{\log N} \geq (1+\delta)^2 - 1. \tag{6.82}$$

The results in Section 3.5 provide the structural properties of $Z^{(1)}$. The following key proposition, a counterpart of Proposition 6.1, states a bound on the largest eigenvalue of such networks. Recall that $\lambda_\alpha = (\log N)^{\frac{1}{\alpha}}$.

A key distinction between Proposition 6.1 and this proposition is that the former claims the smallness of $\lambda_1(A)$ under the condition that the maximum degree is $o(\frac{\log N}{\log\log N})$ and edge-weights are light. Whereas the latter provides the bound on $\lambda_1(A)$ for heavy-tailed weights under the weaker condition on the maximum degree, $O(\frac{\log N}{\log\log N})$.

Proposition 6.2. For any $k, N \in \mathbb{N}$ and constants $\varepsilon, c_1, c_2, c_3 > 0$, let $G = (V, E, A)$ be a connected network ($A = (a_{ij})_{i,j\in V}$ is a weight matrix) whose maximum clique size is k and which satisfies

1. $d_1(G) \leq c_1 \frac{\log N}{\log\log N}$,
2. $|V| \leq \frac{c_2}{\varepsilon}\frac{\log N}{\log\log N}$,
3. $|E| \leq |V| + c_3$.

Suppose that the edge-weights are given by i.i.d. Weibull random variables with a shape parameter $0 < \alpha < 2$ conditioned to be greater than $(\varepsilon \log\log N)^{\frac{1}{\alpha}}$ in absolute value. Let $\rho > 0$ and $0 < \xi < \frac{1}{2}$ be constants, then:

1. In the case $1 < \alpha < 2$, let β be the conjugate of α and set $\tau := (c_3+2)^{\frac{1}{2\beta}}\xi^{\frac{1}{2}}$. Then,

$$\mathbb{P}\left(\lambda_1(A) \geq \rho^{\frac{1}{\alpha}}\lambda_\alpha\right) \leq N^{-2^{-\alpha}\tau^{-\alpha}\rho+c_2+o(1)} + N^{-2^{-1}\phi_{\beta/2}(k)^{1-\alpha}(1-\tau)^{\alpha}\rho+c_1\xi^{-2}\varepsilon+o(1)}, \tag{6.83}$$

where $\phi_{\beta/2}$ is defined in (2.5).

2. In the case $0 < \alpha \leq 1$,

$$\mathbb{P}\left(\lambda_1(A) \geq \rho^{\frac{1}{\alpha}}\lambda_\alpha\right) \leq N^{-2^{-\alpha}\xi^{-\alpha}\rho+c_2+o(1)} + N^{-\rho(1-\xi)^{\alpha}+c_1\xi^{-2}\varepsilon+o(1)}. \tag{6.84}$$

The expression on the right hand side is technical but the constants ε, ξ will be suitably chosen so that the dominant behaviors in the case $1 < \alpha \leq 2$ and $0 < \alpha \leq 1$ are $N^{-2^{-1}\phi_{\beta/2}(k)^{1-\alpha}\rho}$ and $N^{-\rho}$ respectively. Note that in the case $1 < \alpha < 2$, since $\phi_{\beta/2}(k)$ is non-decreasing in k (see (4) in Lemma 2.3), the upper bound (6.83) gets worse as k increases. Whereas, the upper bound (6.84) in the case $0 < \alpha \leq 1$ does not depend on k.

Proof of Proposition 6.2. Let f be a (random) top eigenvector. For a constant $\xi \in (0, \frac{1}{2})$, define the subsets V_S, V_L, E_S, E_L as (6.28) and (6.30) in Proposition 6.1. Since $|V_L| \leq \frac{1}{\xi^2}$, by the condition (1),

$$|E_L| \leq \frac{c_1}{\xi^2}\frac{\log N}{\log\log N}. \tag{6.85}$$

As in (6.31), we write $\lambda_1(A) = 2\lambda_S + 2\lambda_L$. Then, for any $0 < \tau < 1$,

$$\mathbb{P}\left(\lambda_1(A) \geq \rho^{\frac{1}{\alpha}}\lambda_\alpha\right) \leq \mathbb{P}\left(2\lambda_S \geq \tau\rho^{\frac{1}{\alpha}}\lambda_\alpha\right) + \mathbb{P}\left(2\lambda_L \geq (1-\tau)\rho^{\frac{1}{\alpha}}\lambda_\alpha\right). \tag{6.86}$$

How we proceed from here depends on the value of α, but in both cases the parameter τ will be chosen in such way that it is costly for λ_S to be large. In the case $1 < \alpha < 2$ we apply Hölder's inequality to bound λ_S and λ_L, while in the case $0 < \alpha \leq 1$ we use the monotonicity of ℓ^p norms.

Case 1. $1 < \alpha < 2$. We first estimate λ_S. As in (6.33),

$$\sum_{(i,j)\in E_S} |f_i|^\beta |f_j|^\beta \leq \xi^\beta + (c_3+1)\xi^{2\beta} \leq (c_3+2)\xi^\beta,$$

where the first term ξ^β is obtained as an application of (2.18) with $\theta = \frac{\beta}{2} > 1$.

Note that by assumptions, $|E| \leq |V| + c_3 \leq \frac{c_2}{\varepsilon}\frac{\log N}{\log\log N} + c_3$. Hence, setting $\tau := (c_3+2)^{\frac{1}{2\beta}}\xi^{\frac{1}{2}}$, by Hölder's inequality,

$$\lambda_S \leq \Big(\sum_{(i,j)\in E_S} |f_i|^\beta |f_j|^\beta\Big)^{\frac{1}{\beta}}\Big(\sum_{(i,j)\in E_S} |a_{ij}|^\alpha\Big)^{\frac{1}{\alpha}} \leq \tau^2\Big(\sum_{(i,j)\in E_S} |a_{ij}|^\alpha\Big)^{\frac{1}{\alpha}} \leq \tau^2\Big(\sum_{(i,j)\in E} |a_{ij}|^\alpha\Big)^{\frac{1}{\alpha}}.$$

By a tail estimate for the sum of Weibull random variables (the bound (4.18) with $a = \frac{\rho}{2^\alpha \tau^\alpha}, b = \frac{c_2}{\varepsilon}, c = c_3$),

$$\mathbb{P}\left(2\lambda_S \geq \tau\rho^{\frac{1}{\alpha}}\lambda_\alpha\right) \leq \mathbb{P}\Big(\sum_{(i,j)\in E} |a_{ij}|^\alpha \geq \frac{\rho}{2^\alpha\tau^\alpha}\log N\Big) \leq N^{-2^{-\alpha}\tau^{-\alpha}\rho + c_2 + o(1)}. \tag{6.87}$$

Next, we estimate λ_L. By the definition of $\widehat{\phi}_{\beta/2}$ in (2.5) (where sup is taken over $\left\|f\right\|_1 = 1$ whereas our vector f satisfies $\left\|f\right\|_2 = 1$) and since $\widehat{\phi}_{\beta/2} = \phi_{\beta/2}$ (see Lemma 2.2),

$$\sum_{(i,j)\in E_L} |f_i|^\beta |f_j|^\beta \leq \sum_{(i,j)\in E} |f_i|^\beta |f_j|^\beta \leq \frac{1}{2}\widehat{\phi}_{\beta/2}(k) = \frac{1}{2}\phi_{\beta/2}(k).$$

Hence, by Hölder's inequality,

$$\lambda_L \leq \Big(\sum_{(i,j)\in E_L} |f_i|^\beta |f_j|^\beta\Big)^{\frac{1}{\beta}} \Big(\sum_{(i,j)\in E_L} |a_{ij}|^\alpha\Big)^{\frac{1}{\alpha}} \leq \Big(\frac{1}{2}\phi_{\beta/2}(k)\Big)^{\frac{1}{\beta}} \Big(\sum_{(i,j)\in E_L} |a_{ij}|^\alpha\Big)^{\frac{1}{\alpha}}. \tag{6.88}$$

Since the number of possible subsets that a random subset E_L can take is bounded by $|V|^{\left\lfloor \frac{1}{\xi^2} \right\rfloor} = n^{o(1)}$, by a union bound and the tail estimate (4.18) with

$$a = 2^{\frac{\alpha}{\beta}}\phi_{\beta/2}(k)^{-\frac{\alpha}{\beta}}(1-\tau)^\alpha \frac{\rho}{2^\alpha}, \quad b = \frac{c_1}{\xi^2}, \quad c = 0$$

(recall the bound of $|E_L|$ in (6.85)), we have

$$\begin{aligned}\mathbb{P}\left(2\lambda_L \geq (1-\tau)\rho^{\frac{1}{\alpha}}\lambda_\alpha\right) &\leq \mathbb{P}\Big(\sum_{(i,j)\in E_L} |a_{ij}|^\alpha \geq 2^{\frac{\alpha}{\beta}}\phi_{\beta/2}(k)^{-\frac{\alpha}{\beta}}(1-\tau)^\alpha\frac{\rho}{2^\alpha}\log N\Big) \\ &\leq N^{-2^{-1}\phi_{\beta/2}(k)^{1-\alpha}(1-\tau)^\alpha\rho + c_1\xi^{-2}\varepsilon + o(1)},\end{aligned} \tag{6.89}$$

where we used $2^{\frac{\alpha}{\beta}} \cdot 2^{-\alpha} = 2^{-1}$ and $-\frac{\alpha}{\beta} = 1 - \alpha$ by a conjugacy relation $\frac{1}{\alpha} + \frac{1}{\beta} = 1$.

Plugging the bounds (6.87) and (6.89) into (6.86) gives the desired bound.

Case 2. $0 < \alpha \leq 1$. In this case we use (6.86) with $\tau = \xi$. We first estimate λ_S. Since $|f_i f_j| \leq \xi^2$ when $|f_i|, |f_j| < \xi$, by the monotonicity of ℓ^p norms (note that $0 < \alpha \leq 1$),

$$\sum_{(i,j)\in E_S} a_{ij} f_i f_j \leq \xi^2 \sum_{(i,j)\in E_S} |a_{ij}| \leq \xi^2 \sum_{(i,j)\in E} |a_{ij}| \leq \xi^2 \Big(\sum_{(i,j)\in E} |a_{ij}|^\alpha\Big)^{\frac{1}{\alpha}}.$$

Thus, recalling $|E| \leq \frac{c_2}{\varepsilon}\frac{\log N}{\log\log N} + c_3$, by a tail estimate (4.18) with $a = \frac{\rho}{(2\xi)^\alpha}, b = \frac{c_2}{\varepsilon}, c = c_3$,

$$\mathbb{P}(2\lambda_S \geq \xi\rho^{\frac{1}{\alpha}}\lambda_\alpha) \leq \mathbb{P}\Big(\sum_{i<j,(i,j)\in E} |a_{ij}|^\alpha \geq \frac{\rho}{(2\xi)^\alpha}\log N\Big) \leq N^{-2^{-\alpha}\xi^{-\alpha}\rho + c_2 + o(1)}. \tag{6.90}$$

Next, we estimate λ_L. Since $|f_i f_j| \le \frac{1}{2}$ when $\|f\|_2 = 1$, again by the monotonicity of ℓ^p norms,

$$\sum_{(i,j)\in E_L} a_{ij} f_i f_j \le \frac{1}{2} \sum_{(i,j)\in E_L} |a_{ij}| \le \frac{1}{2}\Big(\sum_{(i,j)\in E_L} |a_{ij}|^\alpha \Big)^{\frac{1}{\alpha}}.$$

Thus, as in (6.89),

$$\mathbb{P}\left(2\lambda_L \ge (1-\xi)\rho^{\frac{1}{\alpha}}\lambda_\alpha\right) \le \mathbb{P}\Big(\sum_{(i,j)\in E_L} |a_{ij}|^\alpha \ge (1-\xi)^\alpha \rho \log N \Big) \le N^{-\rho(1-\xi)^\alpha + c_1\xi^{-2}\varepsilon + o(1)}. \tag{6.91}$$

Therefore, by plugging the bounds (6.90) and (6.91) into (6.86), the proof of (6.84) is concluded.

□

With all this preparation, we are now ready to prove the upper bound for the upper tail.

Proof of the upper bound of the upper tail. By a decomposition $Z = Z^{(1)} + Z^{(2)}$, setting δ' as

$$1+\delta = (1+\delta') + \varepsilon^{\frac{1}{\alpha}}(1+\delta), \tag{6.92}$$

we have

$$\mathbb{P}\Big(\lambda_1(Z) \ge (1+\delta)\lambda_\alpha\Big) \le \mathbb{P}\left(\lambda_1\Big(Z^{(1)}\Big) \ge (1+\delta')\lambda_\alpha\right) + \mathbb{P}\left(\lambda_1\Big(Z^{(2)}\Big) \ge \varepsilon^{\frac{1}{\alpha}}(1+\delta)\lambda_\alpha\right). \tag{6.93}$$

By Lemma 6.7,

$$\mathbb{P}\left(\lambda_1\Big(Z^{(2)}\Big) \ge \varepsilon^{\frac{1}{\alpha}}(1+\delta)\lambda_\alpha\right) \le N^{1-(1+\delta)^2+o(1)}, \tag{6.94}$$

which implies that $Z^{(2)}$ is spectrally negligible. Thus, it suffices to focus on the spectral behavior of $Z^{(1)}$. Let $C_1, \cdots, C_m$ be the connected components of $X^{(1)}$ and let $\mathcal{H}$ be the event defined by

$$\mathcal{H} := \Big\{ |\{\ell = 1, \cdots, m : C_\ell \text{ not tree}\}| < \log N \Big\}. \tag{6.95}$$

To bound the probability of $\mathcal{H}^c$, we use the fact that a graph with at least k connected components that are not trees, has at least k vertex-disjoint cycles, implying in particular the existence of k edge-disjoint cycles. Thus denoting by $\mathcal{T}^c$ the event of existence of a cycle and letting $\square$ to be the disjoint occurrence of events (see [Rei00] for a precise definition),

$$\mathcal{H}^c \subset \underbrace{\mathcal{T}^c \,\square\, \ldots \,\square\, \mathcal{T}^c}_{\lceil \log N \rceil \text{ times}}.$$

Then, by (3.24) in Lemma 3.12 and the Van-den Berg-Kesten (BK) inequality [Rei00], for large enough n,

$$\mathbb{P}(\mathcal{H}^c) \le C^{\log N}(\log N)^{-2\varepsilon \log N} \le (\log N)^{-\varepsilon \log N}.$$

Setting $\delta_0 := (1+\delta)^\alpha - 1$, define the event $\mathcal{F}_0$ measurable with respect to $X^{(1)}$ which guarantees that all components $C_1, \cdots, C_m$ satisfy the conditions of Proposition 6.2 and that there are only few components which are not trees:

$$\mathcal{F}_0 := \mathcal{D}_{2\delta_0} \cap \mathcal{C}_{\varepsilon,2\delta_0} \cap \mathcal{E}_{2\delta_0} \cap \mathcal{H},$$

where $\mathcal{D}_{2\delta_0}, \mathcal{C}_{\varepsilon,2\delta_0}$ and $\mathcal{E}_{2\delta_0}$ are the events defined in lemmas 3.10, 3.11 and 3.12 respectively (applied to the graph $X^{(1)}$). By the discussion above as well as the results in these lemmas, for large enough n,

$$\mathbb{P}(\mathcal{F}_0^c) \le N^{-2\delta_0 + o(1)}. \tag{6.96}$$

Conditioned on $X^{(1)}$, let $Z_\ell^{(1)}$ be the network $Z^{(1)}$ restricted to C_ℓ and denote by k_ℓ the size of the maximal clique in C_ℓ. We consider the case $1 < \alpha < 2$ and $0 < \alpha \le 1$ separately, since the maximum clique size turns out to be only relevant in the former case.

Case 1: $1 < \alpha < 2$. Let $\beta > 2$ be the conjugate of α. Under the event $\mathcal{F}_0$, in order to control $\lambda_1(Z_\ell^{(1)})$ for each $\ell = 1, \cdots, m$, we apply Proposition 6.2 with

$$c_1 = c_2 = 1 + 2\delta_0, \quad c_3 = 2\delta_0, \quad \xi := \varepsilon^{\frac{1}{4}}, \quad \tau = (2\delta_0 + 2)^{\frac{1}{2\beta}} \varepsilon^{\frac{1}{8}}, \quad \rho = (1+\delta')^\alpha \tag{6.97}$$

(recall that $\delta_0 = (1+\delta)^\alpha - 1$). Observing that for small enough $\varepsilon > 0$, the first term in the bound (6.83) is negligible compared to the second term,

$$\mathbb{P}\left(\lambda_1\left(Z_\ell^{(1)}\right) \ge (1+\delta')\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{F}_0} \le N^{-2^{-1}\phi_{\beta/2}(k_\ell)^{1-\alpha}(1-\tau)^\alpha(1+\delta')^\alpha + (1+2\delta_0)\varepsilon^{1/2} + o(1)}. \tag{6.98}$$

The argument for a component now depends on the size of its maximal clique, i.e. whether $k_\ell \ge 3$ or $k_\ell = 2$. For this define

$$I := \{\ell = 1, \cdots, m : k_\ell \ge 3\}, \quad J := \{\ell = 1, \cdots, m : k_\ell = 2\},$$

and let $\bar{k} := \max\{k_1, \cdots, k_m\}$. Since the maximum size of clique in any tree is equal to 2, under the event $\mathcal{H}$ defined in (6.95),

$$|I| \le \log N. \tag{6.99}$$

Then, by (6.98) and using the fact that $\phi_{\beta/2}(k_\ell) \le \phi_{\beta/2}(\bar{k})$ (recall that $\phi_{\beta/2}$ is non-decreasing, see (4) in Lemma 2.3), for $\ell \in I$,

$$\mathbb{P}\left(\lambda_1\left(Z_\ell^{(1)}\right) \ge (1+\delta')\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{F}_0} \le N^{-2^{-1}\phi_{\beta/2}(\bar{k})^{1-\alpha}(1-\tau)^\alpha(1+\delta')^\alpha + (1+2\delta_0)\varepsilon^{1/2} + o(1)}$$

$$\le N^{-2^{-1}\phi_{\beta/2}(\bar{k})^{1-\alpha}(1+\delta)^{\alpha}+\zeta_1+o(1)} \tag{6.100}$$

for some $\zeta_1 = \zeta_1(\varepsilon)$ with $\lim_{\varepsilon\to 0}\zeta_1 = 0$. Note that the last inequality follows from the fact that $\lim_{\varepsilon\to 0}\delta' = \delta$, $\lim_{\varepsilon\to 0}\tau = 0$ (recall the definition of δ' and τ in (6.92) and (6.97) respectively) and the uniform boundedness of the quantity $\phi_{\beta/2}(\bar{k})^{1-\alpha}$ (recall the monotonicity property of $\phi_{\beta/2}$ by Lemma 2.3 (4) and the fact $\alpha > 1$).

Similarly, for $\ell \in J$, by (6.98) again,

$$\mathbb{P}\left(\lambda_1\left(Z_\ell^{(1)}\right) \ge (1+\delta')\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{F}_0} \le N^{-2^{-1}\phi_{\beta/2}(2)^{1-\alpha}(1+\delta)^{\alpha}+\zeta_2+o(1)} \tag{6.101}$$

for some $\zeta_2 = \zeta_2(\varepsilon)$ with $\lim_{\varepsilon\to 0}\zeta_2 = 0$.

Note that since $X^{(1)}$ is distributed as $\mathcal{G}_{N,q}$ with $q \le p = \frac{d}{N}$, Lemma 3.9 implies that for any $k \ge 3$, the probability that $X^{(1)}$ contains a clique of size k is bounded by $d^{\binom{k}{2}}N^{-\binom{k}{2}+k}$. Thus, by (6.96), (6.100) and (6.101) combined with a union bound,

$$\begin{aligned}
&\mathbb{P}\left(\lambda_1\left(Z^{(1)}\right) \ge (1+\delta')\lambda_\alpha\right)\\
&\le C\log N \cdot \sum_{k=3}^{N} d^{\binom{k}{2}}N^{-\binom{k}{2}+k}\cdot N^{-2^{-1}\phi_{\beta/2}(k)^{1-\alpha}(1+\delta)^{\alpha}+\zeta_1+o(1)}\\
&+ CN\cdot N^{-2^{-1}\phi_{\beta/2}(2)^{1-\alpha}(1+\delta)^{\alpha}+\zeta_2+o(1)} + CN^{-2\delta_0+o(1)},
\end{aligned} \tag{6.102}$$

where the multiplicative factors $\log N$ and N arise from (6.99) and the fact $|J| \le N$ respectively. We analyze each term above.

Recalling the definition of $\psi_{\alpha,\delta}(k)$ in (6.6), the exponent of N in each summation is less than $-\min_{k\ge 3}\psi_{\alpha,\delta}(k)+\zeta_1+o(1)$. By a straightforward argument, one can deduce that the first term is bounded by $N^{-\min_{k\ge 3}\psi_{\alpha,\delta}(k)+\zeta_1+o(1)}$. In addition, by the first identity in (6.72), the second term is nothing other than $N^{-\psi_{\alpha,\delta}(2)+\zeta_2+o(1)}$. Furthermore, since $\psi_{\alpha,\delta}(2) \overset{(6.72)}{=} (1+\delta)^{\alpha}-1 = \delta_0$, the last term is bounded by $N^{-2\psi_{\alpha,\delta}(2)+o(1)}$.

Therefore, by combining the above bounds together, there exists $\zeta = \zeta(\varepsilon) > 0$ with $\lim_{\varepsilon\to 0}\zeta = 0$ such that the RHS in (6.102) is bounded by $N^{-\min_{k\ge 2}\psi_{\alpha,\delta}(k)+\zeta+o(1)}$. By taking $\varepsilon > 0$ sufficiently small, we conclude the proof.

Case 2: $0 < \alpha \le 1$. We apply Proposition 6.2 as before. The dominating term in the bound (6.84) is the second term for small enough $\varepsilon > 0$. For each $\ell = 1, \cdots, m$,

$$\mathbb{P}\left(\lambda_1\left(Z_\ell^{(1)}\right) \ge (1+\delta')\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{F}_0} \le N^{-(1+\delta')^{\alpha}(1-\varepsilon^{1/4})^{\alpha}+(1+2\delta_0)\varepsilon^{1/2}+o(1)} \le N^{-(1+\delta)^{\alpha}+\zeta'+o(1)} \tag{6.103}$$

for some $\zeta' = \zeta'(\varepsilon) > 0$ with $\lim_{\varepsilon\to 0}\zeta' = 0$. One can then conclude the proof by applying a union bound over at most n many connected components $C_1, \cdots, C_m$. □

Remark 6.5. Although we do not pursue proving this formally, with some additional work, one may be able to prove the following structure theorem: Let

$$\bar{k} := \arg\min_{k\in\mathbb{N}_{\ge 2}}\psi_{\alpha,\delta}(k).$$

Then, conditioned on the upper tail event $\{\lambda_1(Z) \geq (1+\delta)\lambda_\alpha\}$, with high probability, there exists a clique of size close to $\bar{k}$ in X with high edge-weights on it.

Also recall from the paragraph at the end of the idea og proof in Section 6.1 that the conditional structure given atypically large $\lambda_1(Z)$ is different in the Gaussian and heavy-tailed edge-weight cases. In the former case, the optimal size of the clique tends to infinity as the amount of deviation δ goes to infinity. whereas, in the latter case, it stays bounded.

The lower tail

We start by recalling the theorem we prove in this section.

Theorem 6.4. For any $0 < \delta < 1$,

$$\lim_{N\to\infty} \frac{1}{\log N}\left(\log\log \frac{1}{\mathbb{P}\left(\lambda_1(Z) \leq (1-\delta)\lambda_\alpha^{\text{heavy}}\right)}\right) = 1-(1-\delta)^\alpha.$$

Lower bound of the lower tail

We start by defining the X-measurable event $\mathcal{B}_1$ as

$$\mathcal{B}_1 := \left\{\lambda_1(\mathrm{X}) \leq 2\frac{(\log N)^{\frac{1}{2}}}{(\log\log N)^{\frac{1}{2}}}\right\}.$$

Recalling the event $\mathcal{H}$ introduced in (6.95), we define the event $\mathcal{F}_1$ measurable with respect to $\left\{X, X^{(1)}\right\}$:

$$\mathcal{F}_1 := \mathcal{D}_\delta \cap \mathcal{C}_{\varepsilon,\delta} \cap \mathcal{E}_\delta \cap \mathcal{H} \cap \mathcal{B}_1.$$

By Theorem 6.5, $\lim_{N\to\infty} \mathbb{P}(\mathcal{B}_1) = 1$. Combining this with a previous argument to derive (6.96), for large enough n,

$$\mathbb{P}\left(\mathcal{F}_1\right) \geq \frac{1}{2}. \tag{6.104}$$

Since $\lambda_1(Z) \leq \lambda_1(Z^{(1)}) + \lambda_1(Z^{(2)})$, setting

$$\delta'' := \delta + \varepsilon^{\frac{1}{\alpha}}, \tag{6.105}$$

we have

$$\mathbb{P}\left(\lambda_1(Z) \leq (1-\delta)\lambda_\alpha\right) \geq \mathbb{E}\left[\mathbb{P}\left(\lambda_1\left(Z^{(1)}\right) \leq (1-\delta'')\lambda_\alpha,\ \lambda_1\left(Z^{(2)}\right) \leq \varepsilon^{\frac{1}{\alpha}}\lambda_\alpha \mid X, X^{(1)}\right)\mathbf{1}_{\mathcal{F}_1}\right]. \tag{6.106}$$

Since $|Y^{(2)}_{ij}| \le (\varepsilon \log\log N)^{\frac{1}{\alpha}}$, under the event $\mathcal{B}_1$, hence under the event $\mathcal{F}_1$, for large enough n,

$$\lambda_1\left(Z^{(2)}\right) \le 2\frac{(\log N)^{\frac{1}{2}}}{(\log\log N)^{\frac{1}{2}}} \cdot (\varepsilon \log\log N)^{\frac{1}{\alpha}} \le \varepsilon^{\frac{1}{\alpha}}\lambda_\alpha$$

(recall that $\lambda_\alpha = (\log N)^{\frac{1}{\alpha}}$ and $\alpha < 2$). Thus, using the conditional independence of X and $Z^{(1)}$ given $X^{(1)}$, under the event $\mathcal{F}_1$,

$$\begin{aligned}&\mathbb{P}\left(\lambda_1\left(Z^{(1)}\right) \le (1-\delta'')\lambda_\alpha, \lambda_1\left(Z^{(2)}\right) \le \varepsilon^{\frac{1}{\alpha}}\lambda_\alpha \mid X, X^{(1)}\right)\\ &\quad= \mathbb{P}\left(\lambda_1\left(Z^{(1)}\right) \le (1-\delta'')\lambda_\alpha \mid X, X^{(1)}\right) = \mathbb{P}\left(\lambda_1\left(Z^{(1)}\right) \le (1-\delta'')\lambda_\alpha \mid X^{(1)}\right). \end{aligned} \tag{6.107}$$

Therefore, applying this to (6.106),

$$\mathbb{P}\Big(\lambda_1(Z) \le (1-\delta)\lambda_\alpha\Big) \ge \mathbb{E}\left[\mathbb{P}\left(\lambda_1\left(Z^{(1)}\right) \le (1-\delta'')\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{F}_1}\right]. \tag{6.108}$$

Let $C_1, \cdots, C_m$ be the connected components of $X^{(1)}$ and $Z^{(1)}_\ell$ be the restriction of $Z^{(1)}$ to C_ℓ for $\ell = 1, \cdots, m$. Let k_ℓ be the size of the maximal clique in C_ℓ.

By a similar reasoning as in the upper tail case, we separately consider $1 < \alpha < 2$ and $0 < \alpha \le 1$.

Case 1: $1 < \alpha < 2$. Let $\beta > 2$ be the conjugate of α. Under the event $\mathcal{F}_1$, in order to control $\lambda_1(Z^{(1)}_\ell)$ for each $\ell = 1, \cdots, m$, we apply Proposition 6.2 with

$$c_1 = c_2 = 1 + \delta, \quad c_3 = \delta, \quad \xi := \varepsilon^{\frac{1}{4}}, \quad \tau = (\delta+2)^{\frac{1}{2\beta}}\varepsilon^{\frac{1}{8}}, \quad \rho = (1-\delta'')^\alpha. \tag{6.109}$$

The dominating term in the bound (6.83) is the second term for small enough $\varepsilon > 0$. Thus,

$$\mathbb{P}\left(\lambda_1\left(Z^{(1)}_\ell\right) \ge (1-\delta'')\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{F}_1} \le N^{-2^{-1}\phi_{\beta/2}(k_\ell)^{1-\alpha}(1-\tau)^\alpha(1-\delta'')^\alpha + (1+\delta)\varepsilon^{1/2} + o(1)}. \tag{6.110}$$

As in the proof of upper bound for upper tails, we proceed differently for the components depending on the size of maximal cliques:

$$I := \{\ell = 1, \cdots, m : k_\ell \ge 3\}, \quad J := \{\ell = 1, \cdots, m : k_\ell = 2\},$$

and let $\bar{k} := \max\{k_1, \cdots, k_m\}$. To bound (6.110), we use the fact that there is a constant $c = c(\beta) > 0$ such that $\phi_{\beta/2}(k) \le c$ for all $k \ge 2$ (see (4) in Lemma 2.3). Thus, for $\ell \in I$,

$$\begin{aligned}\mathbb{P}\left(\lambda_1\left(Z^{(1)}_\ell\right) \ge (1-\delta'')\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{F}_1} &\le N^{-2^{-1}c^{1-\alpha}(1-\tau)^\alpha(1-\delta'')^\alpha + (1+\delta)\varepsilon^{1/2} + o(1)}\\ &\le N^{-2^{-1}c^{1-\alpha}(1-\delta)^\alpha + \zeta_1 + o(1)}\end{aligned} \tag{6.111}$$

for some $\zeta_1 = \zeta_1(\varepsilon)$ with $\lim_{\varepsilon\to 0}\zeta_1 = 0$. The last inequality follows from the fact that $\lim_{\varepsilon\to0}\delta'' = \delta$ and $\lim_{\varepsilon\to0}\tau = 0$ (see the definition of δ'' and τ in (6.105) and (6.109) respectively).

In addition, for $\ell \in J$, using the fact $\phi_{\beta/2}(2) = 2^{1-\beta}$ (see (2.9) in Lemma 2.3), by (6.110) again,

$$\mathbb{P}\left(\lambda_1\left(Z_\ell^{(1)}\right) \geq (1-\delta'')\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{F}_1} \leq N^{-(1-\tau)^\alpha(1-\delta'')^\alpha+(1+\delta)\varepsilon^{1/2}+o(1)} \leq N^{-(1-\delta)^\alpha+\zeta_2+o(1)} \tag{6.112}$$

for some $\zeta_2 = \zeta_2(\varepsilon)$ with $\lim_{\varepsilon\to 0}\zeta_2 = 0$, where we used $2^{-1}\cdot(2^{1-\beta})^{1-\alpha} = 1$ in the first inequality.

Thus, by (6.111) and (6.112) together with the independence of edge-weights across different components, under the event $\mathcal{F}_1$,

$$\begin{aligned}\mathbb{P}\left(\lambda_1\left(Z^{(1)}\right) \leq (1-\delta'')\lambda_\alpha \mid X^{(1)}\right) &\geq \left(1 - N^{-2^{-1}c^{1-\alpha}(1-\delta)^\alpha+\zeta_1+o(1)}\right)^{\log N}\left(1-N^{-(1-\delta)^\alpha+\zeta_2+o(1)}\right)^N \\ &\geq e^{-N^{1-(1-\delta)^\alpha+\zeta_2+o(1)}},\end{aligned}$$

where the powers $\log N$ and N come from the fact that $|I| \leq \log N$ and $|J| \leq N$ respectively. Therefore, applying this and (6.104) to (6.108),

$$\mathbb{P}\Big(\lambda_1(Z) \leq (1-\delta)\lambda_\alpha\Big) \geq \frac{1}{2}e^{-N^{1-(1-\delta)^\alpha+\zeta_2+o(1)}}.$$

By taking sufficiently small $\varepsilon > 0$, we conclude the proof.

Case 2: $0 < \alpha \leq 1$. We apply Proposition 6.2 as in the case $1 < \alpha < 2$. As mentioned after Proposition 6.2, the dominating term in the bound (6.84) is the second term. Hence, for each $\ell = 1, \cdots, m$,

$$\mathbb{P}\left(\lambda_1(Z_\ell^{(1)}) \geq (1-\delta'')\lambda_\alpha \mid X^{(1)}\right)\mathbf{1}_{\mathcal{F}_1} < N^{-(1-\delta'')^\alpha(1-\varepsilon^{1/4})^\alpha+(1+\delta)\varepsilon^{1/2}+o(1)} \leq N^{-(1-\delta)^\alpha+\zeta_3+o(1)}$$

for some $\zeta_3 = \zeta_3(\varepsilon)$ with $\lim_{\varepsilon\to 0}\zeta_3 = 0$. Thus, under the event $\mathcal{F}_1$,

$$\mathbb{P}\left(\lambda_1\left(Z^{(1)}\right) \leq (1-\delta'')\lambda_\alpha \mid X^{(1)}\right) \geq \left(1 - N^{-(1-\delta)^\alpha+\zeta_3+o(1)}\right)^n \geq e^{-N^{1-(1-\delta)^\alpha+\zeta_3+o(1)}}.$$

By the similar reasoning as before, we conclude the proof. □

Upper bound for the lower tail

Recall the event $\mathcal{M} = \left\{|E(X)| \geq \frac{d(N-1)}{4}\right\}$ defined in (6.75). Since $\lambda_1(Z) \geq \max_{(i,j)\in E(X)}|Z_{ij}| = \max_{(i,j)\in E(X)}|Y_{ij}|$ (see Lemma 2.8),

$$\mathbb{P}\Big(\lambda_1(Z) \leq (1-\delta)\lambda_\alpha\Big) \leq \mathbb{E}\left[\mathbb{P}\left(\max_{(i,j)\in E(X)}|Y_{ij}| \leq (1-\delta)\lambda_\alpha \mid X\right)\mathbf{1}_{\mathcal{M}}\right] + \mathbb{P}\left(\mathcal{M}^c\right).$$

Since $\mathbb{P}(|Y_{ij}| > (1-\delta)\lambda_\alpha) \geq C_1 N^{-(1-\delta)^\alpha}$ for any $i \neq j$, we have

$$\mathbb{P}\left(\max_{(i,j)\in E(X)}|Y_{ij}| \leq (1-\delta)\lambda_\alpha \mid X\right)\mathbf{1}_{\mathcal{M}} \leq \left(1 - C_1 N^{-(1-\delta)^\alpha}\right)^{\frac{d(N-1)}{4}} \leq e^{-N^{1-(1-\delta)^\alpha+o(1)}}.$$

Combining this with the bound $\mathbb{P}(\mathcal{M}^c) \leq e^{-cN}$ obtained in (6.76), we conclude the proof. □

Chapter 7

Mean field interacting multi-type birth-death processes

7.1 Introduction

The multi-type birth-death process

The *multi-type birth-death process* (MTBDP) is a continuous-time Markov chain generalizing the classical birth-death process [Fel68, Ken48] to a finite number of *types*. The state of the MTBDP counts the number of individuals (or *particles*) of each type while they undergo birth, death, and type transition events according to specified rates, which may be arbitrary functions of the current state and of time. If these rates are linear in the state, the MTBDP can be formulated as a branching process [Gri73]. If additionally, the rates for each type are proportional to the count of only that type, the MTBDP is said to be *simple*, and the rates can be specified particle-wise because particles do not interact. The general case of nonlinear rates has also been called a *multivariate competition process* [Reu61, Igl64], which, as noted by [HXC+18], is more restrictive than a multi-type branching process in that the latter allows for increments other than unity, and more general in that the latter is manifestly linear via its defining independence property.

Phylogenetic birth-death models

The MTBDP has facilitated the inference of diversification processes in biological systems, with applications ranging across scales of evolutionary time and biological organization. *Phylogenetic birth-death* models assume that a phylogenetic tree is generated by an MTBDP combined with a sampling process that censors subtrees that are not ancestral to any sampled leaves, so that histories are only partially observed. The diverse flavors of these models are reviewed and introduced with unified notation in [MLM+21]. Given a phylogeny, the inferential targets are the birth and death rates, as well as the type transition rates. Birth and death are variously interpreted as extinction and speciation rates in the context of macroevolutionary studies, or as transmission and recovery rates in the context of epidemiological or viral phylodynamic studies. The literature contains many variants of this modeling approach. Depending on the application, the birth and death rates may be assumed to be time-dependent, depend on particle type, or both.

To facilitate tractable likelihoods, phylogenetic birth-death models typically assume the restrictive non-interacting simple MTBDP, with particle-wise birth and death rates that depend only on particle type, and possibly on time. In this case, given a time-calibrated tree, the likelihood—defined via the conditional density of the tree assuming it has at least one sampled descendant—can be evaluated via tree message-passing computations. This message-passing structure can be seen to follow from elementary properties of branching processes, adapted to partial tree observation. The message functions [in work by NRS14, these are called *branch propagators*] are given by the solutions to master equations that marginalize over all possible unobserved subtrees subtending the branch, and are computed recursively via post-order traversal (from tree tips to root).

Biology involves interactions

Despite the robust computational development and wide usage of phylogenetic birth-death models for phylodynamic inference, their biological expressiveness is limited by the assumption that particles do not interact. Interactions may be essential to evolutionary dynamics. For example, environmental carrying capacity is a fundamental constraint on the long-term dynamics of any evolving population, and models of experimental microbial evolution generally allow for a transition from exponential growth to stationary phase as the population approaches capacity [BGPW19]. As another example, although the simple MTBDP facilitates modeling phenotypic selection via type-dependent birth and death rates, this does not capture *frequency-dependent selection*, where the fitness of a given type depends on the distribution of types in the population. In both of these examples, birth and death rates depend on the state of the population process, and this breaks the tree message-passing structure that phylogenetic birth-death models rely on.

As a motivating biological setting for the ideas to follow, we consider the somatic evolutionary process of *affinity maturation* of antibodies in micro-anatomical structures called *germinal centers* (GCs), which transiently form in lymph nodes during an adaptive immune response [reviewed in VM14, MEV16, SLW19, VN22, LLQ23]. In a GC, B cells—the cells that make antibodies—diversify and compete based on the ability of the antibodies they express to recognize a foreign *antigen* molecule. As GC B cells proliferate, they undergo targeted mutations in the genomic locus encoding the antibody protein that can modify its antigen binding affinity (they undergo type transitions). Via signaling from other GC cell types, the GC is able to monitor the binding phenotype of the B-cell population it contains, and provide survival signals to B cells with the highest-affinity antibodies (i.e., birth and death rates depend on type).

GCs have been studied extensively in mouse models that allow for experimental lineage tracing and manipulation of the B-cell population process. In particular, B cells can be *fate mapped* by genetically engineering them to express a fluorescent protein that marks them with a randomized color at the beginning of the GC evolutionary process [TMP+16, MSE+20, PJV21]. These initially random colors are non-randomly inherited by descendant cells, so a sample of the GC B-cell population at a future time can be partitioned into *lineages* of cells that share distinct common ancestors at the time of the initial color marking. Phylogenetic inference can then be used to reconstruct the evolutionary history of a GC B-cell lineage using the DNA sequences of the sampled B cells [DMV+18].

GC B cells compete for limited proliferative signaling based on the antigen binding affinity of their B-cell receptors, and the population distribution of binding affinities generally improves as affinity maturation unfolds, so a given binding phenotype may be high-fitness early in the process, but low-fitness later when the population distribution of affinity has improved. This invokes frequency-dependent selection, where the birth and death rates should depend on the population distribution of types. GCs are observed to reach a steady-state carrying capacity of several thousand cells, based on limited cell-mediated proliferative signaling, so carrying capacity is likely also important, meaning that birth and death rates should depend on the total population size.

Phylodynamic models have the potential to reveal how evolutionary dynamics is orchestrated in GCs to shape antibody repertoires and immune memories. However, phylogenetic birth-death models cannot accommodate key features of this system. [AMV+17] presented a simulation study using a birth-death model with competition to investigate features of the GC population process, but such agent-based simulations are not amenable to likelihood-based inference for partially observed histories. This motivates us to investigate a class of interacting MTBDPs that preserve tree-message passing for tractable likelihoods, and could thus be used in phylogenetic birth-death models.

Mean-field interactions between replica birth-death processes

Mean-field theories are a fundamental conceptual tool in the study of interacting particle systems. The ideas originated in statistical physics and quantum mechanics as a technique to reduce many-body problems—in fluids, condensed matter, and disordered systems—to effective one-body problems [see Par07, Kad09]. The theory was extensively developed in the context of general classes of stochastic processes, and has since been widely applied across many scientific domains [see CD22a, CD22b, for a review of theory and applications].

Motivated by the setting of GC evolutionary dynamics described above, with population-level interaction among many fate-mapped lineages, we set out to develop a mean-field model that couples the birth and death rates in a focal MTBDP (with D types) to the empirical distribution of states—i.e., the *mean-field*—over an exchangeable system of N replica MTBDPs. More concretely, this empirical distribution process is a stochastic process taking values in the space of probability measures on $\mathbb{N}_0^D$, where $\mathbb{N}_0$ denotes the non-negative integers: the mass assigned by this measure-valued process at time $t \geq 0$

to a vector $\mathbf{y} = (y_1, y_2, \ldots, y_D) \in \mathbb{N}_0^D$ is the proportion of replica processes (including the focal process) that have y_k individuals of type k for $1 \leq k \leq D$.

We prove that the empirical distribution process of the N replicas converges to a deterministic probability measure-valued flow as $N \to \infty$. Using the *propagation of chaos* theory [see CD22a, CD22b, Szn91, for surveys of this vast area and references to its many applications] we moreover show that in this limit, the replicas effectively decouple, and the focal process can instead be said to couple to a deterministic external field. This external field is *self-consistent* in the sense that, at any time $t \geq 0$, it is given by the very distribution of the state of the focal process. We calculate self-consistent fields by solving limiting nonlinear forward equations for the focal process. A key feature of this limit is that it restores message-passing likelihoods in the phylogenetic birth-death model setting, allowing for tractable phylodynamic models with interactions.

We note that there has been some work on mean-field models in the area of superprocesses (continuum analogs of branching processes) – see [Ove95, Ove96]. Finally, [Tha15] is tangentially related to our work in that it treats a particular question concerning mean-field interacting single-type birth-death processes. As the author of this paper observes regarding the literature about mean-field models and propagation of chaos, ". . . there are few results in discrete space."

Due to the difficulty of incorporating interactions in birth-death processes for inference applications, few results have been published so far, but we summarize some developments. [CMS14] developed techniques based on continued fraction representations of Laplace convolutions to calculate transition probabilities for general single-type birth-death processes, without state space truncation. [HXC+18] calculate transition probabilities for the *birth/birth-death* process—a restricted bivariate case where the death rate of one type vanishes, but rates may be otherwise nonlinear. [XGKMM15] use branching process approximations of birth-death processes and generating-function machinery [Wil05] for moment estimation. [GCPP21] study single-type branching processes with strong interactions, restricted to a regime in which duality methods can be used to characterize the stationary distribution.

Instead of the strong interactions considered in the above work, we introduce an MTBDP with mean-field interactions. This mean-field system restores (in the limit) the computational tractability of the non-interacting case. We establish the fairly general conditions under which this process is well-defined, demonstrate how to perform mean-field calculations in the context of a phylogenetic birth-death model, and provide an efficient software implementation. While we were motivated by evolutionary dynamics of antibodies in germinal centers, we also foresee applications to other somatic evolution settings, such as tumor evolution and developmental lineage tracing, and to experimental microbial evolution. While we have outline how to evaluate likelihoods for phylogenetic birth-death models with mean-field interaction, we leave inference on biological data for future work.

We finish with a note regarding possible extensions of this work. We suppress explicit time dependence in the particle-wise birth and death rates $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ for notational compactness, but all the results of §7.2 and §7.4 extend to the inhomogeneous case $\boldsymbol{\lambda}(t)$ and $\boldsymbol{\mu}(t)$ with suitable continuity assumptions in the time domain. We note, however, that our mean-field approach involves effective time-dependence in the rates even if the intrinsic rates are not explicitly time-dependent. This effective time-dependence arises from specifying a finite number of dynamical parameters (i.e., the rates and the interaction matrix W) that uniquely determine an effective field via the condition of self-consistency, Theorem 7.1.

Finally, we notice that our mean-field system of N interacting replica trees has a self-similarity property: if we consider a subset of N particles from one of the replicas at time $t > 0$, this looks like the starting configuration of a new N-system. This suggests that our mean-field model could also be used as an approximation for strong interactions within a single MTBDP. However, the appropriate notions of convergence and exchangeability are less clear in this case. The validity of a mean-field approximation for a single self-interacting MTBDP would seem to involve a delicate balance of quenched disorder from early times when the process is small, on the one hand, with the limiting mean-field interaction when the process is large, on the other hand. We save these questions for future work.

Structure of the chapter

The rest of this chapter, which corresponds to the paper [DEHH24], is structured as follows. In Section 7.2 we construct a system of MTBDPs that can model the properties and interactions between particles we have discussed so far. Moreover, this section contains our main theoretical results. Section 7.3 is dedicated to their proofs. In Section 7.4 we analyze a special case of an MTBDP system numerically.

7.2 Theoretical results: a mean-field interacting multi-type birth-death process with general rates

We start by describing a finite system of fairly general symmetrically interacting MTBDPs for which the interaction may be locally strong but is globally weak in the sense that different MTBDPs interact only via the empirical distribution of their states. Ultimately, we are interested in the joint law of a finite number of focal processes within an infinite system of such mean-field interacting MTBDPs. To this end, we establish that the process of the empirical distribution of families converges to a deterministic probability measure-valued flow. Any finite number of MTBDPs become asymptotically independent and identically distributed. In the limit, the law of any given focal process can be described by a time-inhomogeneous MTBDP. The time inhomogeneity comes from the deterministic probability measure-valued background flow that also describes the one-dimensional marginal distributions of the focal process.

One main contribution of our analysis compared to previous studies is that we allow for a quite general transition rate structure. The rate of a single MTBDP is only restricted to be of at most linear growth and Lipschitz continuous. To deal with the technical challenges that come with these general assumptions, we employ a localization technique and approximate the general system by one that has bounded transition rates. A key feature is that the system with bounded rates is in a certain sense close to the one with unbounded rates, uniformly in the system size.

Consider a finite set of types $\{1, \dots, D\} =: [D]$. The application we have in mind is that each type represents a certain affinity of B-cell receptors. We equivalently refer to the cells as particles, in line with the terminology used in the branching process literature. At the outset, let's envision a germinal center that initially contains a finite collection of $N \in \mathbb{N}$ such B-cells. The progeny process of each of the N founding cells in this GC can be modeled as an MTBDP. During the process of antibody affinity maturation, cells can divide into two daughters of the same type, mutate to one of the other $(D-1)$ affinity types, or die, according to specified rates. The interaction within lineages is (possibly) *strong*, whereas the interaction between the N lineages is *weak*. This means that the rates for the jth lineage depend on its state (locally-strong) and on the empirical distribution of MTBDP states over the N lineages (globally-weak). Note that this includes the special case of rates that depend on the global empirical type distribution aggregated over all N families in the GC. Initially, there are $N \in \mathbb{N}$ founding particles within the GC. A state of this system is then given by $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_N) \in (\mathbb{N}_0^D)^N$, where for $j \in [N]$ and $i \in [D]$, $z_{j,i}$ counts the number of type-i particles in the jth MTBDP. Let $\mathcal{M}_1(\mathbb{N}_0^D)$ denote the probability measures on $\mathbb{N}_0^D$. This space is embedded into the Banach space of finite signed measures on $\mathbb{N}_0^D$ equipped with the total variation norm. For $\nu \in \mathcal{M}_1(\mathbb{N}_0^D)$ and $\mathbf{y} \in \mathbb{N}_0^D$, let $\nu_{\{\mathbf{y}\}} := \nu(\{\mathbf{y}\})$. Then the total variation distance between $\nu, \nu' \in \mathcal{M}_1(\mathbb{N}_0^D)$ is $\|\nu - \nu'\|_{TV} := \frac{1}{2} \sum_{\mathbf{y} \in \mathbb{N}_0^D} |\nu_{\{\mathbf{y}\}} - \nu'_{\{\mathbf{y}\}}|$.

The *empirical distribution of MTBDP states* of $\mathbf{z} \in (\mathbb{N}_0^D)^N$ is

$$\pi_{\mathbf{z}} := \frac{1}{N} \sum_{j=1}^{N} \delta_{\mathbf{z}_j}.$$

For example, for $\mathbf{y} = (y_1, \dots, y_D) \in \mathbb{N}_0^D$ and $\mathbf{z} \in (\mathbb{N}_0^D)^N$, $\pi_{\mathbf{z}}(\{\mathbf{y}\})$ counts the relative frequency of lineages with composition $\mathbf{y}$, i.e. with y_i particles of type i, $i \in [D]$. Let $\mathcal{M}_{1,N}(\mathbb{N}_0^D) := \{\frac{1}{N} \sum_{j=1}^{N} \delta_{\mathbf{z}_j} \in \mathcal{M}_1(\mathbb{N}_0^D) : \mathbf{z} \in (\mathbb{N}_0^D)^N\}$, i.e. the probability measures that can arise as an empirical distribution of an MTBDP system with N initial particles.

Every successive change in the system affects only one particle at a time with a rate depending on the local state of its lineage, and the empirical distribution over the population of N lineages. That is, for $i \neq k \in [D]$, we have the following per lineage rates of various events

$$b^i : \mathbb{N}_0^D \times \mathcal{M}_1(\mathbb{N}_0^D) \to \mathbb{R}_+ \qquad (\textit{birth-rate of type-}i\textit{ particles}),$$
$$d^i : \mathbb{N}_0^D \times \mathcal{M}_1(\mathbb{N}_0^D) \to \mathbb{R}_+ \qquad (\textit{death-rate of type-}i\textit{ particles}),$$

$$m^{i,k} : \mathbb{N}_0^D \times \mathcal{M}_1(\mathbb{N}_0^D) \to \mathbb{R}_+ \qquad (\textit{mutation-rate from type-i to type-k particles}).$$

Throughout we assume that if $z_{j,i} = 0$, then $b^i(\mathbf{z}_j, \pi_{\mathbf{z}}) = d^i(\mathbf{z}_j, \pi_{\mathbf{z}}) = m^{i,k}(\mathbf{z}_j, \pi_{\mathbf{z}}) = 0$ for all $k \in [D]$. We stress that the rates do not depend on N.

We will assume that the rates per lineage grow at most linearly with the number of particles in the lineage and that the rates are Lipschitz continuous in the following sense. For $\mathbf{y} \in \mathbb{N}_0^D$, set $\mathbf{y}_\bullet := \sum_{i\in[D]} y_i \in \mathbb{N}_0$.

Assumption 7.1. (**A.1**) There exists a constant L such that for all $i, k \in [D]$, $i \neq k$, $\mathbf{y} \in \mathbb{N}_0^D$ and $\nu \in \mathcal{M}_1(\mathbb{N}_0^D)$,

$$b^i(\mathbf{y}, \nu) \leq L(\mathbf{y}_\bullet + 1), \quad d^i(\mathbf{y}, \nu) \leq L(\mathbf{y}_\bullet + 1), \quad m^{i,k}(\mathbf{y}, \nu) \leq L(\mathbf{y}_\bullet + 1).$$

(**A.2**) There exists a constant L such that for all $i, k \in [D]$, $i \neq k$, $\mathbf{y}, \mathbf{y}' \in \mathbb{N}_0^D$ and $\nu, \nu' \in \mathcal{M}_1(\mathbb{N}_0^D)$,

$$\begin{aligned} |b^i(\mathbf{y}, \nu) - b^i(\mathbf{y}', \nu')| &\leq L(|\mathbf{y} - \mathbf{y}'|_\bullet + \|\nu - \nu'\|_{TV}), \\ |d^i(\mathbf{y}, \nu) - d^i(\mathbf{y}', \nu')| &\leq L(|\mathbf{y} - \mathbf{y}'|_\bullet + \|\nu - \nu'\|_{TV}), \\ |m^{i,k}(\mathbf{y}, \nu) - m^{i,k}(\mathbf{y}', \nu')| &\leq L(|\mathbf{y} - \mathbf{y}'|_\bullet + \|\nu - \nu'\|_{TV}). \end{aligned}$$

Remark 7.1. To fully model features like carrying capacity constraints, an alternative would be to allow the rates to grow linearly with the mean of the measure, rather than bounding the contribution of the measure by a constant. In this case, the Lipschitz bounds would also depend on something like the Wasserstein-1 distances between the two measures involved. Proving similar results as ours under such assumptions is an open problem that we hope to return to in future work.

The system of MTBDPs is formally described through its infinitesimal generator, which requires some notation. To add and remove particles of type i in the jth MTBDP, we use $\mathbf{e}_{j,i} \in (\mathbb{N}_0^D)^N$, where $(\mathbf{e}_{j,i})_{k,\ell} = \mathbb{1}_k(i)\mathbb{1}_\ell(j)$. The domain of the generator is described using specific function spaces. We write $\bar{C}((\mathbb{N}_0^D)^N)$ for the space of (continuous) bounded functions on $(\mathbb{N}_0^D)^N$ and $\hat{C}((\mathbb{N}_0^D)^N)$ for the space of (continuous) bounded functions on $(\mathbb{N}_0^D)^N$ that vanish at infinity. Moreover, we write $C_c((\mathbb{N}_0^D)^N)$ for the space of compactly supported (finitely supported) (continuous) functions on $(\mathbb{N}_0^D)^N$. For $\mathbf{z} \in (\mathbb{N}_0^D)^N$, set $\mathbf{z}_{\bullet\bullet} := \sum_{j\in[N]} (\mathbf{z}_j)_\bullet = \sum_{j\in[N]} \sum_{i\in[D]} z_{j,i} \in \mathbb{N}_0$.

The generator A^N of the finite system of interacting MTBDPs acts on $f \in \hat{C}((\mathbb{N}_0^D)^N)$ via $A^N f(\mathbf{z}) := \sum_{j=1}^N (A_b^{N,j} + A_d^{N,j} + A_m^{N,j}) f(\mathbf{z})$ with

$$\begin{aligned} A_b^{N,j} f(\mathbf{z}) &:= \sum_{i=1}^D b^i(\mathbf{z}_j, \pi_{\mathbf{z}})[f(\mathbf{z} + \mathbf{e}_{j,i}) - f(\mathbf{z})] \\ A_d^{N,j} f(\mathbf{z}) &:= \sum_{i=1}^D d^i(\mathbf{z}_j, \pi_{\mathbf{z}})[f(\mathbf{z} - \mathbf{e}_{j,i}) - f(\mathbf{z})] \end{aligned}$$

$$A_m^{N,j} f(\mathbf{z}) := \sum_{i,k\in[D],\, i\neq k} m^{i,k}(\mathbf{z}_j, \pi_{\mathbf{z}})[f(\mathbf{z}+\mathbf{e}_{j,k}-\mathbf{e}_{j,i}) - f(\mathbf{z})].$$

Define

$$\Delta_N := \{f \in \hat{C}((\mathbb{N}_0^D)^N) :\ \mathbf{z} \mapsto \mathbf{z}_{\bullet\bullet} f(\mathbf{z}) \in \bar{C}((\mathbb{N}_0^D)^N),\ A^N f \in \hat{C}((\mathbb{N}_0^D)^N)\}.$$

Proposition 7.1 (Feller property for finite system)**.** The closure of $\{(f, A^N f) : f \in \Delta_N\}$ is single-valued and generates a Feller semigroup on $\hat{C}((\mathbb{N}_0^D)^N)$. Moreover, $C_c((\mathbb{N}_0^D)^N)$ is a core for this generator.

The proof of the proposition is in § 7.3. Write

$$\mathbf{Z}^N(t) := (\mathbf{Z}_1^N(t), \dots, \mathbf{Z}_N^N(t))$$

for a process with the semigroup guaranteed by Proposition 7.1 and set $\mathbf{Z}^N := (\mathbf{Z}^N(t))_{t\geq 0}$.

The system exhibits exchangeability among the MTBDPs due to the symmetries of the rates, provided that their initial distribution is also exchangeable. To formally establish this property, we utilize the *Markov mapping theorem*. As a result of this analysis, we also derive the Markovian nature of the system's empirical distribution process, subject to suitable initial conditions.

The empirical distribution of states in the system at time $t \geq 0$ is

$$\Pi^N(t) := \frac{1}{N}\sum_{j=1}^N \delta_{\mathbf{Z}_j^N(t)}.$$

The $\mathcal{M}_{1,N}(\mathbb{N}_0^D)$-valued empirical measure process is $\Pi^N = (\Pi^N(t))_{t\geq 0}$.

Its infinitesimal generator B^N acts on a subset of $\bar{C}(\mathcal{M}_{1,N}(\mathbb{N}_0^D))$, the bounded continuous functions on $\mathcal{M}_{1,N}(\mathbb{N}_0^D)$. More specifically, B^N acts on functions of the form

$$h(\nu) = \frac{1}{N!}\prod_{\substack{\mathbf{y}\in\mathbb{N}_0^D:\\ \mathbf{y}\in\operatorname{supp}(\nu)}} (N\nu(\{\mathbf{y}\}))! \sum_{\mathbf{z}\in(\mathbb{N}_0^D)^N:\pi_{\mathbf{z}}=\nu} f(\mathbf{z}),$$

with $f \in \Delta_N$, via $B^N h(\nu) := (B_b^N + B_d^N + B_m^N)h(\nu)$, where

$$B_b^N h(\nu) := \sum_{\mathbf{y}\in\mathbb{N}_0^D}\sum_{i=1}^D N\nu_{\{\mathbf{y}\}} b^i(\mathbf{y},\nu)\Big[h\Big(\nu + \frac{\delta_{\mathbf{y}+\mathbf{e}_i} - \delta_{\mathbf{y}}}{N}\Big) - h(\nu)\Big],$$

$$B_d^N h(\nu) := \sum_{\mathbf{y}\in\mathbb{N}_0^D}\sum_{i=1}^D N\nu_{\{\mathbf{y}\}} d^i(\mathbf{y},\nu)\Big[h\Big(\nu + \frac{\delta_{\mathbf{y}-\mathbf{e}_i} - \delta_{\mathbf{y}}}{N}\Big) - h(\nu)\Big],$$

$$B_m^N h(\nu) := \sum_{\mathbf{y}\in\mathbb{N}_0^D}\sum_{i,k\in[D],\, i\neq k} N\nu_{\{\mathbf{y}\}} m^{i,k}(\mathbf{y},\nu)\Big[h\Big(\nu + \frac{\delta_{\mathbf{y}+\mathbf{e}_k-\mathbf{e}_i} - \delta_{\mathbf{y}}}{N}\Big) - h(\nu)\Big],$$

with $\mathbf{e}_i$ the ith unit vector in $\mathbb{N}_0^D$.

To formally state the exchangeability of the system and the Markovian nature of Π^N, we require some notation. Let $\alpha^N(\nu, \mathbf{dz})$ be a kernel from $\mathcal{M}_{1,N}(\mathbb{N}_0^D)$ to $(\mathbb{N}_0^D)^N$ defined via

$$\alpha^N(\nu, \mathbf{dz}) := \frac{1}{N!} \prod_{\substack{\mathbf{y}\in\mathbb{N}_0^D: \\ \mathbf{y}\in\mathrm{supp}(\nu)}} (N\nu(\{\mathbf{y}\}))! \sum_{\mathbf{x}\in(\mathbb{N}_0^D)^N:\, \pi_{\mathbf{x}}=\nu} \delta_{\mathbf{x}}(\mathbf{z}),$$

i.e. $\alpha^N(\nu, \cdot)$ puts mass uniformly among all the system states $\mathbf{x} \in (\mathbb{N}_0^D)^N$ that are compatible with an empirical distribution ν. For $f \in \bar{C}((\mathbb{N}_0^D)^N)$, we write $\alpha^N f(\cdot) = \sum_{\mathbf{z}\in(\mathbb{N}_0^D)^N} f(\mathbf{z})\alpha^N(\cdot, \mathrm{d}\mathbf{z})$. (In particular, $\alpha^N f \in \bar{C}(\mathcal{M}_{1,N}(\mathbb{N}_0^D))$.)

Proposition 7.2 (Exchangeability)**.** Let $\nu^N \in \mathcal{M}_{1,N}(\mathbb{N}_0^D)$ and assume $\mathbf{Z}^N(0)$ has distribution $\alpha^N(\nu^N, \cdot)$. For all $t \geq 0$, $\mathbf{Z}^N(t) = (\mathbf{Z}_1^N(t), \dots, \mathbf{Z}_N^N(t))$ is exchangeable and Π^N is a Markov process with generator B^N.

In what follows, we consider the limit of large systems. For the germinal center application, this means that we assume the initial number of B-cells to be large. Any dependence of the rates on the *total mass* therefore is meant to be relative to the initial mass.

Our first main result describes the behavior of Π^N in the limit of large systems. We adopt the usual notation that if I is a closed subinterval of $\mathbb{R}_+$ and E is a metric space, then $D(I, E)$ is the Skorohod space of right-continuous, left-limited functions from I to E.

Theorem 7.1 (Convergence of empirical measure process)**.** Assume $\mathbf{Z}^N(0)$ has distribution $\alpha^N(\nu^N, \cdot)$ for $\nu^N \in \mathcal{M}_{1,N}(\mathbb{N}_0^D)$ satisfying $\nu^N \xrightarrow{N\to\infty} \nu \in \mathcal{M}_1(\mathbb{N}_0^D)$. Then there exists a unique solution to the initial value problem: $v(0) = \nu$ and for all $\mathbf{y} \in \mathbb{N}_0^D$,

$$\begin{aligned} v'_{\{\mathbf{y}\}}(t) = &- v_{\{\mathbf{y}\}}(t) \sum_{i=1}^{D} \Big(b^i(\mathbf{y}, v(t)) + d^i(\mathbf{y}, v(t)) + \sum_{k=1,\, k\neq i}^{D} m^{i,k}(\mathbf{y}, v(t))\Big) \\ &+ \sum_{i=1}^{D} \Big(v_{\{\mathbf{y}-\mathbf{e}_i\}}(t) b^i(\mathbf{y} - \mathbf{e}_i, v(t)) + v_{\{\mathbf{y}+\mathbf{e}_i\}}(t) d^i(\mathbf{y} + \mathbf{e}_i, v(t)) \\ &+ \sum_{k=1,\, k\neq i}^{D} v_{\{\mathbf{y}-\mathbf{e}_k+\mathbf{e}_i\}}(t)\, m^{i,k}(\mathbf{y} - \mathbf{e}_k + \mathbf{e}_i, v(t))\Big). \end{aligned} \tag{7.1}$$

(with the convention that for $\mathbf{y} \notin \mathbb{N}_0^D$, $v_{\{\mathbf{y}\}}(t) = b^i(\mathbf{y}, v(t)) = d^i(\mathbf{y}, v(t)) = m^{i,k}(\mathbf{y}, v(t)) = 0$). Moreover,

$$\Pi^N \xRightarrow{N\to\infty} v \tag{7.2}$$

(that is, the sequence of $D(\mathbb{R}_+, \mathcal{M}_{1,N}(\mathbb{N}_0^D))$-valued random elements $(\Pi^N)_{N\in\mathbb{N}}$ converges in distribution to the deterministic (continuous) function v).

Remark 7.2. From Theorem 7.1, the continuity of v, and the continuous mapping theorem, it follows that, under the conditions of Theorem 7.1,

$$\Pi^N(t) \xrightarrow{N\to\infty} v(t)$$

for every $t \geq 0$ (cf. Theorem 23.9 of [Kal21]). Hence, under the assumptions of Theorem 7.1, it follows from the theory of propagation of chaos, see Proposition 2.2 of [Szn91], that for every $k \in \mathbb{N}$ and $t \geq 0$ the $(\mathbb{N}_0^D)^k$-valued random vector $(\mathbf{Z}_1^N(t), \ldots, \mathbf{Z}_k^N(t))$ converges in distribution and that the limiting distribution is $v(t)^{\otimes k}$. We can do better than this, as the following second major result shows.

Corollary 7.1 (Convergence of focal processes)**.** Under the conditions of Theorem 7.1, for each $k \in \mathbb{N}$ there is convergence in distribution of the $D(\mathbb{R}_+, (\mathbb{N}_0^D)^k)$-valued sequence of random elements $\{(\mathbf{Z}_1^N, \ldots, \mathbf{Z}_k^N)\}_{N\in\mathbb{N}}$ to $(\mathbf{Z}_1^\infty, \ldots, \mathbf{Z}_k^\infty)$, where $\mathbf{Z}_1^\infty, \ldots, \mathbf{Z}_k^\infty$ are i.i.d. time-inhomogeneous MTBDPs with common initial distribution ν and the birth, death, and mutation rates of $\mathbf{Z}_j^\infty$, $1 \leq j \leq k$, at time $t \geq 0$ are given by $b^i(\mathbf{Z}_j(t), v(t))$, $d^i(\mathbf{Z}_j(t), v(t))$, and $m^{i,\ell}(\mathbf{Z}_j(t), v(t))$ for $i, \ell \in [D]$, $i \neq \ell$.

Idea of the proof

We begin by proving the properties of the finite system and its empirical measure process in Section 7.3. Instead of analyzing the limit of the finite system directly, we rely on a localized system, where we freeze processes once they reach a certain size. In Section 7.3 we show that the convergence of the localized system to the original system implies Theorem 7.1.

Since the localized system has a finite state space, we can rely on classical results to derive the limit of its empirical distribution. This is done in Section 7.3. To establish these properties, we define a finite system of independently evolving Yule-type processes that dominates and is coupled to the localized and standard version of the interacting MTBDP system 7.3. The simplicity of this pure birth process allows for easier estimates of its growth and fluctuations, which we can then relate back to the finite interacting systems. Moreover classical results imply the convergence of the empirical measure of this process, which is of use when proving tightness of the empirical measures of the interacting systems. We end the proof section by establishing Corollary 7.1 using the martingale property of the system and the fact that we can take the limit of the empirical distribution in the rates of the focal processes.

7.3 Proofs of the main results

Properties of the finite system

We initiate our analysis by proving the result concerning the generator of the finite system of MTBDPs.

Proof of Proposition 7.1. The proof consists of checking the conditions of Theorem 3.1 in Chapter 8 of [EK86]. The kernel that plays the role of the kernel $x \mapsto \lambda(x)\mu(x, \mathrm{d}y)$ in [EK86] is here

$$\mathbf{z} \mapsto \sum_{j=1}^{N} \sum_{i=1}^{D} \Big[b^i(\mathbf{z}_j, \pi_{\mathbf{z}})\delta_{\mathbf{z}+\mathbf{e}_{j,i}} + d^i(\mathbf{z}_j, \pi_{\mathbf{z}})\delta_{\mathbf{z}-\mathbf{e}_{j,i}} + \sum_{k\in[D], k\neq i} m^{i,k}(\mathbf{z}_j, \pi_{\mathbf{z}})\delta_{\mathbf{z}+\mathbf{e}_{j,k}-\mathbf{e}_{j,i}} \Big] \tag{7.3}$$

We will take the functions γ and η that appear in the statement of that result to both be $\mathbf{z} \mapsto \chi(\mathbf{z}) := (\mathbf{z}_{\bullet\bullet} \vee 1)$.

First note that $\mathbf{z} \mapsto 1/\chi(\mathbf{z}) \in \hat{C}((\mathbb{N}_0^D)^N)$, as required in [EK86]. Secondly,

$$\sup_{\mathbf{z}\in(\mathbb{N}_0^D)^N} \sum_{j=1}^{N} \sum_{i=1}^{D} \Big[b^i(\mathbf{z}_j, \pi_{\mathbf{z}}) + d^i(\mathbf{z}_j, \pi_{\mathbf{z}}) + \sum_{k\in[D], k\neq i} m^{i,k}(\mathbf{z}_j, \pi_{\mathbf{z}}) \Big] \Big/ \chi(\mathbf{z}) < \infty \tag{7.4}$$

by (**A.1**), and so hypothesis (3.2) of [EK86] is satisfied.

If $\mathbf{z}'$ is a point in the support of the measure on the right-hand size of (7.3), then $|\mathbf{z}_{\bullet\bullet} - \mathbf{z}'_{\bullet\bullet}| \leq 1$ and hence hypothesis (3.3) of [EK86] is satisfied.

Combining the bound (7.4) with the observation of the previous paragraph shows that hypotheses (3.4) and (3.5) of [EK86] hold, and this completes the proof. □

Next, we establish the exchangeability of the finite system and demonstrate the Markovianity of the empirical measure process.

Proof of Proposition 7.2. We first want to apply [Kur98, Corollary 3.5]. Note that for any $h \in \bar{C}(\mathcal{M}_{1,N}(\mathbb{N}_0^D))$ and $\pi_{\mathbf{z}} \in \mathcal{M}_{1,N}(\mathbb{N}_0^D)$, we have $\int h(\pi_{\mathbf{y}})\alpha^N(\pi_{\mathbf{z}}, \mathrm{d}\mathbf{y}) = h(\pi_{\mathbf{z}})$. Define

$$C^N = \left\{ \left(\textstyle\sum_{\mathbf{y}\in(\mathbb{N}_0^D)^N} f(\mathbf{y})\alpha^N(\cdot, \mathrm{d}\mathbf{y}), \sum_{\mathbf{y}\in(\mathbb{N}_0^D)^N} A^N f(\mathbf{y})\alpha^N(\cdot, \mathrm{d}\mathbf{y}) \right) : f \in \Delta_N \right\}. \tag{7.5}$$

We have to verify (the somewhat technical condition) that each solution of the extended forward equation of A^N corresponds to a solution of the martingale problem. Assume for now this is true. We show in Lemma 7.1 below that for $f(\mathbf{z}) = \prod_{j=1}^N g_j(\mathbf{z}_j)$ with $g_j \in \hat{C}(\mathbb{N}_0^D)$, $\alpha^N(A^N f)(\pi_{\mathbf{z}}) = B^N(\alpha^N f)(\pi_{\mathbf{z}})$. In particular, Π^N solves the C^N martingale problem. Thus, by Corollary 3.5 of [Kur98] (with $\gamma(\mathbf{z}) = \pi_{\mathbf{z}}$), Π^N is a Markov process. Moreover, by Theorem 4.1 of [Kur98], $\mathbf{Z}^N(t)$ is then exchangeable.

It remains to verify that each solution of an extended forward equation of A^N corresponds to a solution of the martingale problem. By Lemma 3.1 of [Kur98], it is enough to verify that A^N satisfies the conditions of Theorem 2.6 of [Kur98], that is, that A^N is a pre-generator and Hypothesis 2.4 of [Kur98] is satisfied. Because $(\mathbb{N}_0^D)^N$ is locally compact, the latter is satisfied by Remark 2.5 of [Kur98]. Another consequence of local compactness is that for A^N to be a pre-generator, it is enough to verify that A^N satisfies the positive maximum principle [Kur98, p.4], which is easily seen to be the case. □

The following lemma is a technical result used in the proof of Proposition 7.2.

Lemma 7.1. For $j \in [N]$, let $g_j \in \hat{C}(\mathbb{N}_0^D)$ and set $f(\mathbf{z}) = \prod_{j=1}^N g_j(\mathbf{z}_j^N)$. Then, for $\mathbf{z} \in (\mathbb{N}_0^D)^N$ and $\pi_{\mathbf{z}} \in \mathcal{M}_{1,N}(\mathbb{N}_0^D)$, we have

$$\alpha_N(A^N f)(\pi_{\mathbf{z}}) = B^N(\alpha_N f)(\pi_{\mathbf{z}}).$$

Proof. We will only show $\alpha_N(A_b^N f)(\pi_{\mathbf{z}}) = B_b^N(\alpha_N f)(\pi_{\mathbf{z}})$. That $\alpha_N A_d^N f(\pi_{\mathbf{z}}) = B_d^N(\alpha_N f)(\pi_{\mathbf{z}})$ and $\alpha_N A_m^N f(\pi_{\mathbf{z}}) = B_m^N(\alpha_N f)(\pi_{\mathbf{z}})$ can be proven in a similar vein. The result then follows from the linearity of A^N and B^N. We have

$$\begin{aligned}
&\alpha_N A_b^N f(\pi_{\mathbf{z}}) \\
&= \frac{1}{N!} \prod_{\substack{\mathbf{y}' \in \mathbb{N}_0^D: \\ \mathbf{y} \in \mathrm{supp}(\pi_{\mathbf{z}})}} (N\pi_{\mathbf{z}}(\{\mathbf{y}'\}))! \sum_{\substack{\mathbf{x} \in (\mathbb{N}_0^D)^N: \\ \pi_{\mathbf{x}} = \pi_{\mathbf{z}}}} \sum_{j=1}^N \sum_{i=1}^D b^i(\mathbf{x}^j, \pi_{\mathbf{z}})[f(\mathbf{x} + \mathbf{e}_{j,i}) - f(\mathbf{x})] \\
&= \frac{1}{N!} \prod_{\substack{\mathbf{y}' \in \mathbb{N}_0^D: \\ \mathbf{y} \in \mathrm{supp}(\pi_{\mathbf{z}})}} (N\pi_{\mathbf{z}}(\{\mathbf{y}'\}))! \sum_{\mathbf{y} \in \mathbb{N}_0^D} \sum_{i=1}^D \sum_{j=1}^N \sum_{\substack{\mathbf{x} \in (\mathbb{N}_0^D)^N: \\ \pi_{\mathbf{x}} = \pi_{\mathbf{z}}}} \mathbb{1}_{\mathbf{y}}(\mathbf{x}^j) b^i(\mathbf{x}^j, \pi_{\mathbf{z}})[f(\mathbf{x} + \mathbf{e}_{j,i}) - f(\mathbf{x})] \\
&= \sum_{\mathbf{y} \in \mathbb{N}_0^D} \sum_{i=1}^D N\pi_{\mathbf{z}}(\mathbf{y}) b^i(\mathbf{y}, \pi_{\mathbf{z}}) \Big[(\alpha_N f)\Big(\pi_{\mathbf{z}} + (\delta_{\mathbf{y}+\mathbf{e}_i} - \delta_{\mathbf{y}})/N\Big) - (\alpha_N f)(\pi_{\mathbf{z}})\Big] \\
&= B_b^N(\alpha_N f)(\pi_{\mathbf{z}}).
\end{aligned}$$

□

For the remainder of this section we will assume that the conditions of Theorem 7.1 hold; that is, (**A.1**) and (**A.2**) hold, and the sequence $\nu^N \in \mathcal{M}_{1,N}(\mathbb{N}_0^D)$, $N \in \mathbb{N}$, satisfies $\nu^N \xrightarrow{N\to\infty} \nu \in \mathcal{M}_1(\mathbb{N}_0^D)$.

A dominating system of Yule-type processes

It will be useful to compare $\mathbf{Z}^N$ to a system of asymptotically independent multi-type pure-birth-like processes that will have simultaneous births of different types. Even though these Markov processes are $\mathbb{N}_0^D$-valued, they inherit several useful properties from the classic Yule process. Let $\mathbf{R}^N = (\mathbf{R}_1^N, \ldots, \mathbf{R}_N^N)$ be the $(\mathbb{N}_0^D)^N$-valued Markov process transitioning from $(\mathbb{N}_0^D)^N \ni (\mathbf{r}_1, \ldots, \mathbf{r}_N) \to (\mathbf{r}_1, \ldots, \mathbf{r}_N) + (\mathbf{0}, \ldots, \mathbf{0}, \mathbf{1}, \mathbf{0}, \ldots, \mathbf{0})$ at rate $6LD^2(\mathbf{r}_j)_\bullet$, where $\mathbf{0} \in \mathbb{N}_0^D$ is the vector of all 0s and $\mathbf{1} \in \mathbb{N}_0^D$ is the vector of all 1s. Define $\rho : \mathbb{N}_0^D \to \mathbb{N}_0^D$ by $\rho(\mathbf{r}) := \mathbf{r}_\bullet \mathbf{1}$ (that is, if we think of $\mathbf{r}$ as a collection of particles of different types, then $\rho(\mathbf{r})$ replaces each particle by D particles where there is one particle of every one of the D types. Suppose that $\mathbf{R}^N(0)$ has distribution $\alpha^N(\nu^N, \cdot) \circ (\rho, \ldots, \rho)^{-1}$. It follows that each $(\mathbf{R}_j^N)_\bullet$ is an autonomous Markov process on $D\mathbb{N}_0$ that transitions from Dr to $Dr + D$ at rate

$6LD^3 r$. Consequently, $D^{-1}(\mathbf{R}_j^N)_\bullet$ is a Yule process that transitions from state x at rate $6LD^3 x$ (that is, the split rate per particle is $6LD^3$).

For $t \geq 0$ define $\Pi_{\mathbf{R}}^N(t) := \frac{1}{N}\sum_{j=1}^N \delta_{\mathbf{R}_j^N(t)}$ and set $\Pi_{\mathbf{R}}^N := (\Pi_{\mathbf{R}}^N(t))_{t\geq 0}$.

Lemma 7.2. i) For each $k \in \mathbb{N}$, there exist $\mathbf{R}_j^\infty$, $j \in [k]$, such that the sequence $\{(\mathbf{R}_1^N, \ldots, \mathbf{R}_k^N)\}_{N\in\mathbb{N}}$ converges in distribution to $(\mathbf{R}_1^\infty, \ldots, \mathbf{R}_k^\infty)$. The $\mathbf{R}_j^\infty$, $j \in [k]$ are i.i.d. Markov processes. Each one has initial distribution $\nu \circ \rho^{-1}$ and the same transition dynamics as the $\mathbf{R}_\ell^N$, $\ell \in [N]$, $N \in \mathbb{N}$, have in common.

ii) There is a unique solution $r := (r(t))_{t\geq 0}$ to the initial value problem: $r(0) = \nu \circ \rho^{-1}$ and for all $\mathbf{y} \in \mathbb{N}_0^D$,

$$r'_{\{\mathbf{y}\}}(t) = -\,6LD^2\mathbf{y}_\bullet r_{\{\mathbf{y}\}}(t) + 6LD^2(\mathbf{y}-\mathbf{1})_\bullet r_{\{\mathbf{y}-\mathbf{1}\}}(t), \tag{7.6}$$

where $r_{\{\mathbf{y}\}}(t) = 0$ for $\mathbf{y} \notin \mathbb{N}_0^D$.

iii) We may build $(\mathbf{R}^N)_{N\in\mathbb{N}}$ on a suitable probability space so that

$$\Pi_{\mathbf{R}}^N \xrightarrow{N\to\infty} r, \quad \text{a.s.}$$

Proof. i) Recall that the distribution of $\mathbf{R}^N(0)$ is $\alpha^N(\nu_0^N, \cdot) \circ (\rho, \ldots, \rho)^{-1}$. It suffices to show that the projection of this exchangeable probability measure onto the first k coordinates of $(\mathbb{N}_0^D)^N$ converges weakly to the product probability measure $(\nu \circ \rho^{-1})^{\otimes k}$ as $N \to \infty$. Moreover, from Proposition 2.2 of [Szn91] it suffices to check that the sequence of probability measures on $\mathcal{M}_1(\mathbb{N}_0^D)$ given by $(\alpha^N(\nu^N, \cdot) \circ (\rho, \ldots, \rho)^{-1}) \circ \pi^{-1}$, $N \in \mathbb{N}$, converges weakly to the unit point mass at the probability measure $\nu \circ \rho^{-1}$ (recall that $\mathbf{z} \mapsto \pi_{\mathbf{z}}$ is the map that takes $\mathbf{z} \in (\mathbb{N}_0^D)^N$ to $\frac{1}{N}\sum_{j=1}^N \delta_{\mathbf{z}_j} \in \mathcal{M}_{1,N}(\mathbb{N}_0^D)$). However, it is clear by construction that $(\alpha^N(\nu^N, \cdot) \circ (\rho, \ldots, \rho)^{-1}) \circ \pi^{-1}$ is simply the unit point mass at the probability measure $\nu^N \circ \rho^{-1}$.

ii) Note that (7.6) is just the Kolmogorov forward equations for a Markov process with transition dynamics the common transition dynamics of $\mathbf{R}_j^N$, $j \in [N]$, $N \in \mathbb{N}$, and initial distribution $\nu \circ \rho^{-1}$; that is, for a Markov process with the common distribution of $\mathbf{R}_j^\infty$, $N \in \mathbb{N}$. As we have remarked, such a Markov process is essentially a Yule process, and hence the Kolmogorov forward equations have a unique solution.

iii) From (i) and Proposition 2.2 of [Szn91] we have that the empirical measures on the path space $D(R_+, \mathcal{M}_1(\mathbb{N}_0^D))$ given by $\Sigma^N := \frac{1}{N}\sum_{j=1}^N \delta_{\mathbf{R}_j^N}$ converge in distribution to the point mass at the common distribution of the $\mathbf{R}_j^\infty$, $j \in \mathbb{N}$. By Skorohod's coupling, see Theorem 5.31 in [Kal21], it is possible to build random variables with the distributions of the Σ^N, $N \in \mathbb{N}$, on a suitable probability space so that Σ^N converges almost surely to the point mass at the common distribution of the $\mathbf{R}_j^\infty$, $j \in \mathbb{N}$. We may, of course, also assume that $\mathbf{R}_j^\infty$, $j \in \mathbb{N}$, are built on this probability space.

Fix $T > 0$. Let $\mathfrak{D}$ be a countable dense set in $[0, T]$ containing $\{0, T\}$. By the continuous mapping theorem, with probability one we have that for all $\mathbf{m} \in \mathbb{N}_0^D$

$$\Pi_{\mathbf{R}}^N(t)(\{\mathbf{y} : y_i \geq m_i,\ i \in [D]\}) \xrightarrow{N\to\infty} r(t)(\{\mathbf{y} : y_i \geq m_i,\ i \in [D]\}) \tag{7.7}$$

for all $t \in \mathfrak{D}$. By well-known results in real analysis, the monotonicity of the functions involved in the convergence in (7.7), plus the continuity of the right-hand side give firstly that the convergence holds for all $t \in [0, T]$ and secondly that the convergence is uniform. Consequently, almost surely $\Pi_{\mathbf{R}}^N(t)(\{\mathbf{y}\})$ converges uniformly to $r_{\{\mathbf{y}\}}(t)$ on $[0, T]$ for every $\mathbf{y} \in \mathbb{N}_0^D$.

Given any $\epsilon > 0$ we can choose K such that $\Pi_{\mathbf{R}}^N(T)(\{\mathbf{y} : \mathbf{y}_\bullet > K\}) \leq \epsilon$ for all N and $r(T)(\{\mathbf{y} : \mathbf{y}_\bullet > K\}) \leq \epsilon$. Therefore, using the monotonicity of $\Pi_{\mathbf{R}}^N(t)(\{\mathbf{y} : \mathbf{y}_\bullet > K\})$ and $r(t)(\{\mathbf{y} : \mathbf{y}_\bullet > K\})$ we have

$$\limsup_{N\to\infty} \sup_{t\in[0,T]} \|\Pi_{\mathbf{R}}^N(t) - r(t)\|_{\mathrm{TV}} \leq 2\epsilon.$$

Since T and ϵ are arbitrary, this completes the proof. □

For $\mathbf{z}, \mathbf{z}' \in (\mathbb{N}_0^D)^N$, we write $\mathbf{z} \leq \mathbf{z}'$ if $z_{j,i} \leq z'_{j,i}$ for all j, i.

Lemma 7.3 (Dominating pure-birth process coupling)**.** For each $N \in \mathbb{N}$ we can couple $\mathbf{Z}^N$ and $\mathbf{R}^N$ together so that almost surely $\mathbf{Z}^N(0) \leq \mathbf{R}^N(0)$ and almost surely for all $t \geq 0$, $j \in [N]$, and $i \in [D]$, $|Z_{j,i}^N(t) - Z_{j,i}^N(t-)| \leq R_{j,i}^N(t) - R_{j,i}^N(t-)$. In particular, almost surely for all $t \geq 0$, $\mathbf{Z}^N(t) \leq \mathbf{R}^N(t)$ and almost surely for all $0 \leq s < t$, $j \in [N]$, and $i \in [D]$, $|Z_{j,i}^N(t) - Z_{j,i}^N(s)| \leq R_{j,i}^N(t) - R_{j,i}^N(s)$.

Proof. First note that, because $(\mathbf{R}_1^N(0), \ldots, \mathbf{R}_N^N(0))$ has the same distribution as $(\rho(\mathbf{Z}_1^N(0)), \ldots, \rho(\mathbf{Z}_N^N(0)))$, it is certainly possible to couple $\mathbf{Z}^N(0)$ and $\mathbf{R}^N(0)$ together in the prescribed manner.

Next observe that the rate at which a given $\mathbf{Z}_j^N$ transitions to another state if $\mathbf{Z}_j^N$ is in state $\mathbf{z}_j \neq 0$ can be upper bounded using (**A.1**):

$$\sum_{i=1}^{D} \Big(b^i(\mathbf{z}_j, \pi_{\mathbf{z}}) + d^i(\mathbf{z}_j, \pi_{\mathbf{z}}) + \sum_{k\in[D], k\neq i} m^{i,k}(\mathbf{z}_j, \pi_{\mathbf{z}}) \Big)$$
$$\leq 3LD^2\Big(1 + (\mathbf{z}_j)_\bullet\Big) \leq 6LD^2(\mathbf{z}_j)_\bullet.$$

The same inequality holds trivially when $\mathbf{z}_j = 0$ by our assumption that in this case $b^i(\mathbf{z}_j, \pi_{\mathbf{z}}) = d^i(\mathbf{z}_j, \pi_{\mathbf{z}}) = m^{i,k}(\mathbf{z}_j, \pi_{\mathbf{z}}) = 0$ for $i, k \in [D]$, $i \neq k$.

Thus, we can couple $\mathbf{Z}^N$ to $\mathbf{R}^N$ by restricting the possible jump times of $\mathbf{Z}^N$ to the jump times of $\mathbf{R}^N$ and for τ a jump time of $\mathbf{R}^N$ such that the jump occurs in $\mathbf{R}_j^N$ for some $j \in [D]$, setting

$$\mathbf{Z}^N(\tau) = \begin{cases} \mathbf{Z}^N(\tau-) + \mathbf{e}_{j,i}, & \text{w.p. } \frac{b^i(\mathbf{Z}_j^N(\tau-),\Pi^N(\tau-))}{6LD^2(\mathbf{R}_j^N(\tau-))_\bullet}, \quad i \in [D] \\ \mathbf{Z}^N(\tau-) - \mathbf{e}_{j,i}, & \text{w.p. } \frac{d^i(\mathbf{Z}_j^N(\tau-),\Pi^N(\tau-))}{6LD^2(\mathbf{R}_j^N(\tau-))_\bullet}, \quad i \in [D] \\ \mathbf{Z}^N(\tau-) + \mathbf{e}_{j,k} - \mathbf{e}_{j,i}, & \text{w.p. } \frac{m^{i,k}(\mathbf{Z}_j^N(\tau-),\Pi^N(\tau-))}{6LD^2(\mathbf{R}_j^N(\tau-))_\bullet}, \quad i,k \in [D],\ i \neq k, \\ \mathbf{Z}^N(\tau-), & \text{otherwise.} \end{cases}$$

It is straightforward to check that then $\mathbf{Z}^N$ has the correct distribution and the other properties we want. □

From now on we assume $\mathbf{Z}^N$ is constructed on the basis of the coupling in Lemma 7.3.

For two probability measures $\nu, \nu' \in \mathcal{M}_1(\mathbb{N}_0^D)$, we say that ν' *stochastically dominates* ν if for every $\mathbf{m} \in \mathbb{N}_0^D$, $\nu(\{\mathbf{y} : y_i \geq m_i,\ i \in [D]\}) \leq \nu'(\{\mathbf{y} : y_i \geq m_i,\ i \in [D]\})$; we then write $\nu \preceq \nu'$.

Remark 7.3. The upper bound of $\mathbf{Z}^N$ in terms of $\mathbf{R}^N$ can be translated to a bound for their respective empirical measure processes. To this end, define $\Pi_{\mathbf{R}}^N(t) := \frac{1}{N}\sum_{j=1}^N \delta_{\mathbf{R}_j(t)}$ and set $\Pi_{\mathbf{R}}^N := (\Pi_{\mathbf{R}}^N(t))_{t\geq 0}$. Because of Lemma 7.3, we have $\Pi^N(t) \preceq \Pi_{\mathbf{R}}^N(t)$ for every $t \geq 0$. Moreover, $\Pi_{\mathbf{R}}^N$ is non-decreasing, i.e. for all $0 \leq s < t$, $\Pi_{\mathbf{R}}^N(s) \preceq \Pi_{\mathbf{R}}^N(t)$.

Proving convergence via localization

We employ a localization argument to establish Theorem 7.1. The core concept involves freezing families that reach a certain size $\kappa \in \mathbb{N}$. By utilizing classic methods, we can prove the convergence of the empirical distribution for a system undergoing such freezing.

Let $\mathbf{Z}^{N,\kappa} := (\mathbf{Z}_1^{N,\kappa}, \ldots, \mathbf{Z}_N^{N,\kappa})$ be the system of interacting MTBDPs that is coupled to $\mathbf{Z}^N$ by freezing lineages once they reach a state $\mathbf{y}$ where $\mathbf{y}_\bullet = \kappa$. Notably, the construction of $\mathbf{Z}^{N,\kappa}$ can therefore also be based on the system $\mathbf{R}^N$ in the manner of Lemma 7.3. Importantly, $\mathbf{Z}^{N,\kappa}(t) \leq \mathbf{R}^N(t)$ holds for all $t \geq 0$.

Let $b^{i,\kappa}$, $d^{i,\kappa}$, and $m^{i,k,\kappa}$ be the birth, death, and mutation rates of $\mathbf{Z}^{N,\kappa}$. For example, $b^{i,\kappa}(\mathbf{z},\nu) = b^i(\mathbf{z},\nu)\mathbb{1}(\mathbf{z}_\bullet < \kappa)$ We may think of $\mathbf{Z}^{N,\kappa}$ as a Markov process on the finite state space $(\bar{I}_\kappa)^N$ where $\bar{I}_\kappa := \{\mathbf{y} \in \mathbb{N}_0^D : \mathbf{y}_\bullet \leq \kappa\}$.

The generator $A^{N,\kappa}$ of $\mathbf{Z}^{N,\kappa}$ is then $A^{N,\kappa}f(\mathbf{z}) := \sum_{j=1}^N (A_b^{N,j,\kappa} + A_d^{N,j,\kappa} + A_m^{N,j,\kappa})f(\mathbf{z})$ for $f \in \hat{C}((\mathbb{N}_0^D)^N)$ with

$$A_b^{N,j,\kappa} f(\mathbf{z}) := \sum_{i=1}^D b^{i,\kappa}(\mathbf{z}_j, \pi_{\mathbf{z}})[f(\mathbf{z} + \mathbf{e}_{j,i}) - f(\mathbf{z})]$$

$$A_d^{N,j,\kappa} f(\mathbf{z}) := \sum_{i=1}^{D} d^{i,\kappa}(\mathbf{z}_j, \pi_{\mathbf{z}})[f(\mathbf{z} - \mathbf{e}_{j,i}) - f(\mathbf{z})]$$
$$A_m^{N,j,\kappa} f(\mathbf{z}) := \sum_{i,k \in [D],\, i \neq k} m^{i,k,\kappa}(\mathbf{z}_j, \pi_{\mathbf{z}})[f(\mathbf{z} + \mathbf{e}_{j,k} - \mathbf{e}_{j,i}) - f(\mathbf{z})].$$

Due to the state space of $\mathbf{Z}^{N,\kappa}$ being (essentially) finite, rendering it compact, we can now state the following proposition (see [EK86, Ch. 8.3.1]).

Proposition 7.3 (Feller property for the system of frozen processes)**.** The closure of $\{(f, A^{N,\kappa} f) : f \in C((\mathbb{N}_0^D)^N)\}$ is single-valued and generates a Feller semigroup on $C((\mathbb{N}_0^D)^N)$.

Also, in the system with frozen dynamics, the empirical distribution process is Markov. To be precise, define for $t \geq 0$, $\Pi^{N,\kappa}(t) := \frac{1}{N}\sum_{j=1}^{N} \delta_{\mathbf{Z}_j^{N,\kappa}(t)} \in \mathcal{M}_{1,N}(\mathbb{N}_0^D)$ and set $\Pi^{N,\kappa} := (\Pi^{N,\kappa}(t))_{t\geq 0}$. Its infinitesimal generator $B^{N,\kappa}$ is defined in the same way as B^N, but with the κ-frozen transition rates and modified domain (because the domain of $A^{N,\kappa}$ is different). The following holds via Proposition 7.2, since the rate functions of the frozen process satisfy (7.1).

Remark 7.4 (Exchangeability)**.** Let $\nu^N \in \mathcal{M}_{1,N}(\mathbb{N}_0^D)$ and assume $\mathbf{Z}^{N,\kappa}(0)$ has distribution $\alpha^N(\nu^N, \cdot)$. For all $t \geq 0$, $\mathbf{Z}^{N,\kappa}(t) = (\mathbf{Z}_1^{N,\kappa}(t), \ldots, \mathbf{Z}_N^{N,\kappa}(t))$ is exchangeable and $\Pi^{N,\kappa}$ is a Markov process with generator $B^{N,\kappa}$.

The proof of Theorem 7.1 revolves around three key propositions, all of which will be proved in § 7.3.

Proposition 7.4 (Approximation is uniform in system size)**.** For all $T > 0$ and for all $\kappa > 0$, there is $\varepsilon(\kappa, T)$ such that

$$\mathbb{E}\left[\sup_{t\in[0,T]} \|\Pi^{N,\kappa}(t) - \Pi^N(t)\|_{\mathrm{TV}}\right] \leq \varepsilon(\kappa, T)$$

and $\varepsilon(\kappa, T) \xrightarrow{\kappa\to\infty} 0$.

Proposition 7.5 (Convergence of empirical measure process in systems with freezing)**.** We have $\Pi^{N,\kappa} \xRightarrow{N\to\infty} v^\kappa$, where $v^\kappa = (v^\kappa(t))_{t\geq 0}$ is the unique solution to the initial value problem:

$v^\kappa(0) = \nu \in \mathcal{M}_1(\mathbb{N}_0^D)$, for $\mathbf{y} \in \mathbb{N}_0^D \setminus \bar{I}_\kappa$: $v^\kappa_{\{\mathbf{y}\}}(t) = \nu_{\{\mathbf{y}\}}(0)$; and for $\mathbf{y} \in \bar{I}_\kappa$:

$$\begin{aligned}(v^\kappa_{\{\mathbf{y}\}})'(t) = &- v^\kappa_{\{\mathbf{y}\}}(t) \sum_{i=1}^D \Big(b^{i,\kappa}(\mathbf{y}, v^\kappa(t)) + d^{i,\kappa}(\mathbf{y}, v^\kappa(t)) + \sum_{k=1,\, k\neq i}^D m^{i,k,\kappa}(\mathbf{y}, v^\kappa(t)) \Big) \\ &+ \sum_{i=1}^D \Big(v^\kappa_{\{\mathbf{y}-\mathbf{e}_i\}}(t) b^{i,\kappa}(\mathbf{y}-\mathbf{e}_i, v^\kappa(t)) + v^\kappa_{\{\mathbf{y}+\mathbf{e}_i\}}(t) d^{i,\kappa}(\mathbf{y}+\mathbf{e}_i, v^\kappa(t)) \\ &+ \sum_{k=1,\, k\neq i}^D v^\kappa_{\{\mathbf{y}-\mathbf{e}_k+\mathbf{e}_i\}}(t)\, m^{i,k,\kappa}(\mathbf{y}-\mathbf{e}_k+\mathbf{e}_i, v^\kappa(t)) \Big).\end{aligned} \tag{7.8}$$

Proposition 7.6 (Tightness of the empirical measure process)**.** The sequence $\{\Pi^N\}_{N\in\mathbb{N}}$ is tight.

We now prove Theorem 7.1.

Proof of Theorem 7.1. Fix $T > 0$. By Proposition 7.6, $(\Pi^N)_{N\in\mathbb{N}}$ is tight. Consider $(\Pi^{N_n})_{n\in\mathbb{N}}$ for a strictly increasing sequence $(N_n)_{n\in\mathbb{N}}$ in $\mathbb{N}$. There exists a weakly convergent subsequence $(\Pi^{N_{n_\ell}})_{\ell\in\mathbb{N}}$ and a càdlàg $\mathcal{M}_1(\mathbb{N}_0^D)$-valued process $\Pi^\star$ with $\Pi^{N_{n_\ell}} \overset{\ell\to\infty}{\Longrightarrow} \Pi^\star$.

On the one hand, by Proposition 7.4,

$$\mathbb{E}[\sup_{t\in[0,T]} \|\Pi^{N_{n_\ell},\kappa}(t) - \Pi^{N_{n_\ell}}(t)\|_{\mathrm{TV}}] \leq \varepsilon(\kappa, T).$$

On the other hand, by Proposition 7.5, $\Pi^{N_{n_\ell},\kappa} \overset{\ell\to\infty}{\Longrightarrow} v^\kappa$.

Let ρ be the following standard metric giving the Skorohod topology on the space $D([0,T], \mathcal{M}_1(\mathbb{N}_0^D))$ of càdlàg paths from $[0,T]$ to $\mathcal{M}_1(\mathbb{N}_0^D)$,

$$\rho(\mu, \nu) := \inf_{\lambda\in\Lambda} \left(\sup_{t\in[0,T]} |t - \lambda(t)| \vee \sup_{t\in[0,T]} \|\mu(t) - \nu \circ \lambda(t)\|_{\mathrm{TV}} \right),$$

where the infimum is over all continuous, increasing, bijections $\lambda : [0,T] \to [0,T]$. (cf. equation (12.13) of [Bil99]. Let W_1 be the Wasserstein–1 metric on the space of probability measures on $D([0,T], \mathcal{M}_1(\mathbb{N}_0^D))$ corresponding to ρ; that is,

$$W_1(P, Q) := \inf_R \int \rho(\mu, \nu)\, R(d\mu, d\nu),$$

where the infimum is over all probability measures R on the space $D([0,T], \mathcal{M}_1(\mathbb{N}_0^D)) \times D([0,T], \mathcal{M}_1(\mathbb{N}_0^D))$ that have respective marginals P and Q. Recall that W_1 metrizes weak convergence on the space of probability measures on $D([0,T], \mathcal{M}_1(\mathbb{N}_0^D))$ (see, for example, Theorem 6.9 of [V+09]). If Φ and Ψ are random elements of

$D([0,T],\mathcal{M}_1(\mathbb{N}_0^D))$, write $W_1(\Phi,\Psi)$ for the Wasserstein–1 distance between their respective distributions.

Observe that by setting $\Lambda(t) = t$, we get that $\rho(\mu,\nu) \leq \sup_{t\in[0,T]} \|\mu(t)-\nu(t)\|_{\mathrm{TV}}$, which implies that

$$W_1(\Phi,\Psi) \leq \inf_R \int \sup_{t\in[0,T]} \|\mu(t)-\nu(t)\|_{\mathrm{TV}} R(d\mu,d\nu).$$

If Φ and Ψ happen to be defined on the same probability space, we can choose as R the joint distribution of Φ and Ψ on that space to get

$$W_1(\Phi,\Psi) \leq \mathbb{E}\left[\sup_{t\in[0,T]} \|\Phi(t)-\Psi(t)\|_{\mathrm{TV}}\right].$$

Now,

$$\begin{aligned} &W_1\left(\Pi^\star, v^\kappa\right) \\ &\quad\leq W_1\left(\Pi^\star, \Pi^{N_{n_\ell}}\right) + \mathbb{E}\left[\sup_{t\in[0,T]} \|\Pi^{N_{n_\ell}}(t) - \Pi^{N_{n_\ell},\kappa}(t)\|_{TV}\right] \\ &\quad\quad + W_1\left(\Pi^{N_{n_\ell},\kappa}, v^\kappa\right). \end{aligned}$$

Taking $\ell \to \infty$ leads to the bound

$$W_1(\Pi^\star, v^\kappa) \leq \varepsilon(\kappa,T)$$

independent of the chosen subsequence (N_{n_ℓ}). Since $\varepsilon(\kappa,T) \to 0$ as $\kappa \to \infty$, we obtain $v^\kappa \overset{\kappa\to\infty}{\Longrightarrow} \Pi^\star$ and $\Pi^N \overset{N\to\infty}{\Longrightarrow} \Pi^\star$ upon taking $\kappa \to \infty$. In particular, $\Pi^\star = v$ of (7.1). □

Convergence of the dynamics under freezing

To establish the convergence of the system of MTBDPs that are frozen once they reach the set of frozen states parameterized by κ, we employ standard methods. In this regard, we rely on the following result, which is elaborated upon in [EK86, Ch. 4] concerning the notation used.

Theorem 7.2. [EK86, Corollary 4.8.16] Let (E,r) be complete and separable and $E_N \subset E$. Let $\mathcal{A} \subset \bar{C}(E) \times \bar{C}(E)$ and $v \in \mathcal{M}_1(E)$. Assume

1. the $D(\mathbb{R}_+,E)$ martingale problem for $(\mathcal{A},v)$ has at most one solution, and the closure of the linear span of $\mathcal{D}(\mathcal{A})$, the domain of $\mathcal{A}$, contains an algebra that separates points,

2. for each $N \in \mathbb{N}$, X_N is a progressively measurable process with measurable contraction semigroup $\{T_N(t)\}$, full generator $\hat{\mathcal{A}}_N$, and sample paths in $D(\mathbb{R}_+,E_N)$,

3. $\{X_N\}$ satisfies the compact containment condition; that is, for every $\eta > 0$ and $T > 0$ there is a compact set $\Gamma_{\eta,T} \subset E$ such that $\inf_N \mathbb{P}(X_N(t) \in \Gamma_{\eta,T} \text{ for } 0 \le t \le T) \ge 1 - \eta$,

4. for each $(f, g) \in \mathcal{A}$ and $T > 0$, there exists $(f_N, g_N) \in \hat{\mathcal{A}}_N$ and $G_N \subset E_N$ such that $\lim_{N\to\infty} \mathbb{P}(X_N(t) \in G_N,\ 0 \le t \le T) = 1$, $\sup_N \|f_N\| < \infty$, and $\lim_{N\to\infty} \sup_{x\in G_N} \|f(x) - f_N(x)\| = \lim_{N\to\infty} \sup_{x\in G_N} \|g(x) - g_N(x)\| = 0$,

5. $X_N(0) \xRightarrow{N\to\infty} v$.

Then, there exists a solution X of the $D(\mathbb{R}_+, E)$ martingale problem for $(\mathcal{A}, v)$ and $X_N \xRightarrow{N\to\infty} X$.

To apply Theorem 7.2, one of the things to check is that the sequence $\{\Pi^{N,\kappa}\}_{N\in\mathbb{N}}$ satisfies the compact containment condition. We will prove the following stronger result.

Lemma 7.4 (Compact containment)**.** The sequences $(\Pi^N)_{N\in\mathbb{N}}$ and $(\Pi^{N,\kappa})_{N\in\mathbb{N}}$ both satisfy the compact containment condition.

Proof. Fix $\eta, T > 0$. By how we have coupled together the construction of the processes involved, we have for any $t \in [0, T]$,

$$\begin{aligned} \Pi^N(t) \le \Pi^N_{\mathbf{R}}(t) \le \Pi^N_{\mathbf{R}}(T) \\ \Pi^{N,\kappa}(t) \le \Pi^N_{\mathbf{R}}(t) \le \Pi^N_{\mathbf{R}}(T). \end{aligned} \tag{7.9}$$

Recall that r is the solution to the Kolmogorov forward equation of a non-explosive Markov process that is essentially a Yule process.

Since $\Pi^N_{\mathbf{R}}(T) \xrightarrow{N\to\infty} r(T)$, by Lemma 7.2, we have that the collection of distributions of the sequence $\{\Pi^N_{\mathbf{R}}(T)\}_{N\in\mathbb{N}}$ is tight. Therefore, there exists a compact set $K_{\eta,T} \subseteq \mathcal{M}_1(\mathbb{N}_0^D)$ such that $\mathbb{P}(\Pi^N_{\mathbf{R}}(T) \in K_{\eta,T}) \ge 1 - \eta$ for all N.

It only remains to note that if K is a compact subset of $\mathcal{M}_1(\mathbb{N}_0^D)$, then so is the set $\bigcup_{\nu\in K}\{\mu \in \mathcal{M}_1(\mathbb{N}_0^D) : \mu \le \nu\}$ and then apply (7.9). □

We are now prepared to prove the convergence of the empirical measure process in a system with freezing.

Proof of Proposition 7.5. First, we note that the initial value problem can be reduced to a finite system of ODEs. Its right-hand side is Lipschitz continuous because the rates can be bounded using Assumption 7.1 and because $\mathbf{y} \in \bar{I}_\kappa$. The existence and uniqueness of a solution to this system follow from classic theory (e.g. [Dei77, Chapter 1]). Note that v^κ also solves (uniquely) the $(B^{N,\kappa}, \nu^\kappa)$ martingale problem, because the martingale problem and the ODE in this frozen (thus finite-dimensional) setting are equivalent [Kur11, Corollary 1.3].

We verify that the conditions of Theorem 7.2 are satisfied. To this end, take $E_N = \mathcal{M}_{1,N}(\bar{I}_\kappa)$ and $E = \mathcal{M}_1(\bar{I}_\kappa)$ in Theorem 7.2. The corresponding generators that we are interested in are $B^{N,\kappa}$, as defined before Remark 7.4, and $B^\kappa := B^\kappa_b + B^\kappa_d + B^\kappa_m$, where

$$B^\kappa_b h(\nu) = \sum_{\mathbf{y}\in\bar{I}_\kappa} \sum_{i=1}^{D} \nu_{\{\mathbf{y}\}} b^{i,\kappa}(\mathbf{y},\nu) \left[\frac{\partial h(\nu)}{\partial \nu_{\{\mathbf{y}+\mathbf{e}_i\}}} - \frac{\partial h(\nu)}{\partial \nu_{\{\mathbf{y}\}}} \right],$$

and B^κ_d and B^κ_m are obtained similarly by modifying the definitions of B^N_d and B^N_m before Proposition 7.2.

That there is at most one solution to the $(B^{N,\kappa}, \nu^\kappa)$ martingale problem follows from the discussion at the beginning of this proof. Moreover, we have that

$$\left\{ f(v) = \prod_{\mathbf{y}\in Y} g_{\mathbf{y}}(v_{\mathbf{y}}) \text{ with } g_{\mathbf{y}} \in \hat{C}^1(\mathbb{R}) \text{ and } Y \subset \mathbb{N}_0^D, |Y| < \infty \right\} \subset \bar{C}^1\left(\mathcal{M}_1(\bar{I}_\kappa)\right)$$

is an algebra that separates points. Thus, (1) holds. $\Pi^{N,\kappa}$ is an $\mathcal{M}_{1,N}(\bar{I}_\kappa)$-valued adapted, càdlàg Markov process and thus progressively measurable. Hence, (2) holds. Lemma 7.4 yields that (3) holds. For (4), fix $h \in \bar{C}^1(\mathcal{M}_1(\bar{I}_\kappa))$. Without loss of generality, we can assume that

$$h(\nu) = \tilde{h}(\nu_{\{\mathbf{y}^{(1)}\}}, \ldots, \nu_{\{\mathbf{y}^{(k)}\}})$$

for some $\tilde{h} \in \bar{C}^1([0,1]^k)$ with $k \in \mathbb{N}$, where $\mathbf{y}^{(1)}, \ldots, \mathbf{y}^{(k)} \in \bar{I}_\kappa$. We have to find a sequence $\{h^N\}$ of functions in the domain of the generator of $\Pi^{N,\kappa}$ that approximates h (recall its form from (7.5); but with $f \in C((\bar{I}_\kappa)^N)$ because the frozen system state space is compact). To this end, set

$$\tilde{f}^N(\mathbf{z}) = \tilde{h}(\pi_{\mathbf{z}}(\{\mathbf{y}^{(1)}\}), \ldots, \pi_{\mathbf{z}}(\{\mathbf{y}^{(k)}\}))$$

and

$$\begin{aligned} h^N(\nu) &= \frac{1}{N!} \prod_{\substack{\mathbf{x}\in\mathbb{N}_0^D: \\ \mathbf{x}\in\mathrm{supp}(\pi_{\mathbf{z}})}} (N\pi_{\mathbf{z}}(\{\mathbf{x}\}))! \sum_{\substack{\mathbf{z}\in(\mathbb{N}_0^D)^N: \\ \pi_{\mathbf{z}}=\nu}} \tilde{f}^N(\mathbf{z}) \\ &= \tilde{h}(\nu_{\{\mathbf{y}^{(1)}\}}, \ldots, \nu_{\{\mathbf{y}^{(k)}\}}). \end{aligned}$$

h^N is in the domain of $B^{N,\kappa}$. The only difference between h and h^N is that h^N is only defined on $\mathcal{M}_{1,N}(\bar{I}_\kappa)$, while the domain of h is $\mathcal{M}_1(\bar{I}_\kappa)$, and the two functions agree on $\mathcal{M}_{1,N}(\bar{I}_\kappa)$. This implies that

$$\sup_{\nu\in\mathcal{M}_{1,N}(\bar{I}_\kappa)} |h(\nu) - h^N(\nu)| = 0,$$

so in particular, the limit as $N \to \infty$ is 0. Since $\tilde{h}$ is bounded, also $\sup_N \|h^N\| < \infty$. Next, we show

$$\sup_{\nu\in\mathcal{M}_{1,N}(\mathbb{N}_0^D)} |B^{N,\kappa} h^N(\nu) - B^\kappa h(\nu)| \xrightarrow{N\to\infty} 0. \tag{7.10}$$

To this end, we start showing $\sup_{\nu\in\mathcal{M}_{1,N}(\bar{I}_\kappa)}|(B_b^{N,\kappa}h^N(\nu) - B_b^\kappa h(\nu))| \xrightarrow{N\to\infty} 0$. Note that by Taylor's formula and (**A.1**),

$$\begin{aligned}
&\left|B_b^{N,\kappa}h^N(\nu) - B_b^\kappa h(\nu)\right| \\
&= \sum_{\mathbf{y}\in\bar{I}_\kappa}\sum_{i=1}^{D} \nu_{\{\mathbf{y}\}} b^{i,\kappa}(\mathbf{y},\nu)\left[N\Big(h\Big(\nu + \frac{\delta_{\mathbf{y}+\mathbf{e}_i} - \delta_{\mathbf{y}}}{N}\Big) - h(\nu)\Big) - \frac{\partial h(\nu)}{\partial \nu_{\{\mathbf{y}+\mathbf{e}_i\}}} + \frac{\partial h(\nu)}{\partial \nu_{\{\mathbf{y}\}}}\right] \\
&\leq \sum_{\mathbf{y}\in\{\tilde{\mathbf{y}}^{(1)},\dots,\mathbf{y}^{(k)}\}} LD(\mathbf{y}_\bullet + 1)\cdot O(N^{-1}).
\end{aligned}$$

Since the right-hand side is independent of ν,

$$\sup_{\nu\in\mathcal{M}_{1,N}(\bar{I}_\kappa)} |B_b^{N,\kappa}h^N(\nu) - B_b^\kappa h(v)| \xrightarrow{N\to\infty} 0.$$

Analogously, it can be shown that $\sup_{\nu\in\mathcal{M}_{1,N}(\bar{I}_\kappa)}|B_d^{N,\kappa}h^N - B_d^\kappa h(\nu)| \xrightarrow{N\to\infty} 0$ and $\sup_{\nu\in\mathcal{M}_{1,N}(\bar{I}_\kappa)}|B_m^{N,\kappa}h^N - B_m^\kappa h(\nu)| \xrightarrow{N\to\infty} 0$. Then (4) follows by the triangle inequality. By assumption, (5) holds. In particular, we have checked (1)–(5) of Theorem 7.2 and thus the result follows. □

Next, we prove the bound on $\mathbb{E}[\sup_{t\in[0,T]}\|\Pi^{N,\kappa}(t) - \Pi^N(t)\|_{\mathrm{TV}}]$ that is uniform in N.

Proof of Proposition 7.4. Fix $T > 0$. For $u \leq \inf\{s \geq 0 : \mathbf{Z}_j^N(s)_\bullet = \kappa\}$, we have $\mathbf{Z}_j^N(u) = \mathbf{Z}_j^{N,\kappa}(u)$, $j \in [N]$. Thus, for $0 \leq t \leq T$, $|\{j \in N : \mathbf{Z}_j^N(t) \neq \mathbf{Z}_j^{N,\kappa}(t)\}| \leq |\{j \in [N] : \mathbf{R}_j^N(t) \notin \bar{I}_\kappa)\}| \leq |\{j \in [N] : \mathbf{R}_j^N(T) \notin \bar{I}_\kappa)\}|$. In words: the count of families that have different compositions under the original and frozen dynamics is bounded from above by the count of families in the dominating pure-birth-type process that exited $\bar{I}_\kappa$. Thus,

$$\begin{aligned}
&\mathbb{E}[\sup_{t\in[0,T]}\|\Pi^{N,\kappa}(t) - \Pi^N(t)\|_{\mathrm{TV}}] \\
&\quad\leq \mathbb{E}\left[\sup_{t\in[0,T]}\frac{1}{N}|\{j \in N : \mathbf{Z}_j^N(t) \neq \mathbf{Z}_j^{N,\kappa}(t)\}|\right] \\
&\quad\leq \mathbb{E}\left[\frac{1}{N}|\{j \in [N] : \mathbf{R}_j^N(T) \notin \bar{I}_\kappa)\}\right] \\
&\quad= \mathbb{P}(\mathbf{R}_1^N(T) \notin \bar{I}_\kappa).
\end{aligned}$$

We know, however, from Lemma 7.2(i) that the sequence $\{\mathbf{R}_1^N(T)\}_{N\in\mathbb{N}}$ is weakly convergent and hence tight, so $\varepsilon(\kappa, T) := \sup_{N\in\mathbb{N}} \mathbb{P}(\mathbf{R}_1^N(T) \notin \bar{I}_\kappa)$ has the desired properties.

□

We now address the tightness of the sequence $\{\Pi^N\}_{N\in\mathbb{N}}$.

Proof of Proposition 7.6. From Theorems 23.8 and 23.11 of [Kal21], it suffices to check the following.

1. For every $\eta > 0$ and $T > 0$ there is a compact set $\Gamma_{\eta,T} \subset \mathcal{M}_1(\mathbb{N}_0^D)$ such that $\inf_N \mathbb{P}(\Pi^N(t) \in \Gamma_{\eta,T} \text{ for } 0 \le t \le T) \ge 1 - \eta$.

2. For all $T > 0$,
$$\lim_{\theta \searrow 0} \sup_{N \in \mathbb{N}} \sup_{S \in \mathcal{S}_T^N} \sup_{0 \le u \le \theta} \mathbb{E}[\|\Pi^N(S+u) - \Pi^N(S)\|_{TV}] = 0,$$

where $\mathcal{S}_T^N$ is the set of all discrete $\sigma(\Pi^N)$-stopping times that are bounded by T.

Part 1 has been verified in Lemma 7.4.

For Part 2, note that a.s.

$$\begin{aligned}
&\|\Pi^N(S+u) - \Pi^N(S)\|_{TV} \\
&= \frac{1}{2}\frac{1}{N} \sum_{\mathbf{y} \in \mathbb{N}_0^D} \big||\{j \in [N] : \mathbf{Z}_j^N(S+u) = \mathbf{y}\}| - |\{j \in [N] : \mathbf{Z}_j^N(S) = \mathbf{y}\}|\big| \\
&\le \frac{1}{2}\frac{1}{N} \sum_{\mathbf{y} \in \mathbb{N}_0^D} \sum_{j=1}^N |\mathbb{1}_{\{\mathbf{Z}_j^N(S+u) = \mathbf{y}\}} - \mathbb{1}_{\{\mathbf{Z}_j^N(S) = \mathbf{y}\}}| \\
&\le \frac{1}{N} \sum_{j=1}^N \mathbb{1}_{\{\mathbf{Z}_j^N(S+u) \neq \mathbf{Z}_j^N(S)\}}.
\end{aligned}$$

In particular, using exchangeability, for $u \in [0, \theta]$ and $S \in \mathcal{S}_T^N$

$$\begin{aligned}
\mathbb{E}[\|\Pi^N(S+u) - \Pi^N(S)\|_{TV}] &\le \mathbb{P}(\mathbf{Z}_1^N(S+u) \neq \mathbf{Z}_1^N(S)) \\
&\le \mathbb{P}(\|\mathbf{Z}_1^N(S+u) - \mathbf{Z}_1^N(S)\|_1 \ge \varepsilon) \\
&\le \mathbb{P}(\mathbf{R}_1^N(S+u)_\bullet - \mathbf{R}_1^N(S)_\bullet \ge \varepsilon) \\
&\le \mathbb{P}(\mathbf{R}_1^N(S+\theta)_\bullet - \mathbf{R}_1^N(S)_\bullet \ge \varepsilon) \\
&\le \mathbb{P}(\mathbf{R}_1^N(T+\theta)_\bullet - \mathbf{R}_1^N(T)_\bullet \ge \varepsilon)
\end{aligned}$$

for any $\varepsilon > 0$. For all $N \in \mathbb{N}$, $\lim_{\theta \searrow 0}(\mathbf{R}_1^N(T+\theta)_\bullet - \mathbf{R}_1^N(T)_\bullet) = 0$ almost surely. Also, by Lemma 7.2(i), $\mathbf{R}_1^N(T+\theta)_\bullet - \mathbf{R}_1^N(T)_\bullet$ converges in distribution to $\mathbf{R}_1^\infty(T+\theta)_\bullet - \mathbf{R}_1^\infty(T)_\bullet$ and $\lim_{\theta \searrow 0}(\mathbf{R}_1^\infty(T+\theta)_\bullet - \mathbf{R}_1^\infty(T)_\bullet) = 0$ almost surely. Combining these observations gives Part 2. □

Finally, we address the convergence of the sequence $(\mathbf{Z}_1^N, \ldots, \mathbf{Z}_k^N)_{N \in \mathbb{N}}$ for each fixed $k \in \mathbb{N}$.

Lemma 7.5. For each $k \in \mathbb{N}$, the sequence $\{(\mathbf{Z}_1^N, \ldots, \mathbf{Z}_k^N)\}_{N \in \mathbb{N}}$ is tight.

Proof. From Lemma 7.3 we know that almost surely for all $0 \le s < t$, all $N \in \mathbb{N}$, $j \in [N]$, and $i \in [D]$, that $Z^N_{j,i}(t) \le R^N_{j,i}(t)$ and $|Z^N_{j,i}(t) - Z^N_{j,i}(s)| \le |R^N_{j,i}(t) - R^N_{j,i}(s)|$. It follows from these comparisons, the necessary and sufficient conditions for tightness in Theorem 7.2 in Chapter 4 of [EK86], and the convergence (hence tightness) of the sequence $\{(\mathbf{R}^N_1, \ldots, \mathbf{R}^N_k)\}_{N\in\mathbb{N}}$ established in Lemma 7.2, that the sequence $\{(\mathbf{Z}^N_1, \ldots, \mathbf{Z}^N_k)\}_{N\in\mathbb{N}}$ is tight. □

Proof of Corollary 7.1. For ease of notation, set $\mathbf{Z}^N_{[k]} := (\mathbf{Z}^N_1, \ldots, \mathbf{Z}^N_k)$, $N \in \mathbb{N}$. We know from Remark 7.2 that $(\mathbf{Z}^N_{[k]}(0))_{N\in\mathbb{N}}$ converges in distribution to a random element with distribution $\nu^{\otimes k}$.

By Lemma 7.5, the sequence $(\mathbf{Z}^N_{[k]})_{N\in\mathbb{N}}$ is tight.

Note for any function $f \in C_c((\mathbb{N}^D_0)^k)$ that

$$\begin{aligned}
&f(\mathbf{Z}^N_{[k]}(t)) \\
&- \int_0^t \sum_{j\in[k]} \Big[\sum_{i\in[D]} b^i(\mathbf{Z}^N_j(s), \Pi^N(s))(f(\mathbf{Z}^N_{[k]}(s) + e_{j,i}) - f(\mathbf{Z}^N_{[k]}(s))) \\
&+ \sum_{i\in[D]} d^i(\mathbf{Z}^N_j(s), \Pi^N(s))(f(\mathbf{Z}^N_{[k]}(s) - e_{j,i}) - f(\mathbf{Z}^N_{[k]}(s))) \\
&+ \sum_{i,\ell\in[D],\ell\neq i} m^{i,\ell}(\mathbf{Z}^N_j(s), \Pi^N(s))(f(\mathbf{Z}^N_{[k]}(s) + e_{j,\ell} - e_{j,i}) - f(\mathbf{Z}^N_{[k]}(s))) \Big] ds
\end{aligned}$$

is a martingale.

From Theorem 7.1, we have that $(\Pi^N)_{N\in\mathbb{N}}$ converges in probability to v, so any subsequential limit $\mathbf{Z}_{[k]} := (\mathbf{Z}^\infty_1, \ldots, \mathbf{Z}^\infty_k)$ is such that for any function $f \in C_c((\mathbb{N}^D_0)^k)$,

$$\begin{aligned}
&f(\mathbf{Z}^\infty_{[k]}(t)) \\
&- \int_0^t \sum_{j\in[k]} \Big[\sum_{i\in[D]} b^i(\mathbf{Z}^\infty_j(s), v(s))(f(\mathbf{Z}^\infty_{[k]}(s) + e_{j,i}) - f(\mathbf{Z}^\infty_{[k]}(s))) \\
&+ \sum_{i\in[D]} d^i(\mathbf{Z}^\infty_j(s), v(s))(f(\mathbf{Z}^\infty_{[k]}(s) - e_{j,i}) - f(\mathbf{Z}^\infty_{[k]}(s))) \\
&+ \sum_{i,\ell\in[D],\ell\neq i} m^{i,\ell}(\mathbf{Z}^\infty_j(s), v(s))(f(\mathbf{Z}^\infty_{[k]}(s) + e_{j,\ell} - e_{j,i}) - f(\mathbf{Z}^\infty_{[k]}(s))) \Big] ds
\end{aligned}$$

is a martingale (comparisons with $(\mathbf{R}^N)$ establish the necessary uniform integrability). This completes the proof. □

7.4 A computationally tractable special case: locally simple with moment-mediated interactions

In this section, we study a special case of the mean-field interacting MTBDP defined in §7.2 that is amenable to calculations in the context of a phylogenetic birth-death model. We specialize to a case with no local interactions, and with global interactions mediated by moments of the limiting transition probability $v(t)$ defined by (7.1). This class of processes is rich enough to model both carrying capacity and frequency-dependent selection, and does not add undue computational complexity.

Example 7.1 (Simple MTBDP with moment-mediated mean-field interactive death rates)**.** Consider the MTBDP with transition rates

$$b^i(\mathbf{y}, \nu) = y_i\lambda_i, \quad d^i(\mathbf{y}, \nu) = y_i\tilde{\mu}_i\left(\textstyle\sum_{\mathbf{y}\in\mathbb{N}_0^D} \mathbf{y}\nu_{\{\mathbf{y}\}}\right), \quad m^{i,j}(\mathbf{y}, \nu) = y_i\Gamma_{i,j},$$

$i, j \in [D]$, $j \neq i$, where $\boldsymbol{\lambda} \in \mathbb{R}^D_{\geq 0}$, $\Gamma \in \mathbb{R}^{D\times D}$ (with $\Gamma\mathbf{1} = \mathbf{0}$, and non-negative off-diagonal entries), and $\tilde{\boldsymbol{\mu}} \in C(\mathbb{R}^D_+, \mathbb{R}^D_+)$. Set $\mathbf{r}(t) := \sum_{\mathbf{y}\in\mathbb{N}_0^D} \mathbf{y}\, v_{\{\mathbf{y}\}}(t)$, where $(v(t))_{t\geq 0}$, the solution to (7.1), is the limit of the empirical distribution processes $(\Pi^N(t))_{t\geq 0}$. Then $\mathbf{r} = (\mathbf{r}(t))_{t\geq 0}$ is the first-moment process and it solves the finite, closed system,

$$\begin{aligned} r_i(t)' &= (\lambda_i - \tilde{\mu}_i(\mathbf{r}(t)))\, r_i(t) + \sum_{j=1}^{D} \Gamma_{ji} r_j(t), \quad i = 1, \ldots, D \\ \mathbf{r}(0) &= \mathbf{r}_0, \end{aligned} \tag{7.11}$$

where $\mathbf{r}_0$ is the expected initial state. Note that a solution of (7.1) has finite expectation, via Lemma 7.3.

Example 7.1 specializes the general mean-field interactions considered in Section 7.2 to a mean-field interaction mediated by the expected state vector (the first moment of the state distribution). In that case, the interaction field is the solution of the finite-dimensional nonlinear moment equation (7.11), so we can bypass solving the full infinite-dimensional nonlinear forward equation (7.1).

Example 7.2 (Linear moment interaction)**.** As a simple and biologically interpretable example of the special case of Example 7.1, we take $\tilde{\boldsymbol{\mu}}(\mathbf{r}(t)) = \boldsymbol{\mu} + W\mathbf{r}(t)$, where $\boldsymbol{\mu} \in \mathbb{R}^D_{\geq 0}$ is constant and the matrix $W \in \mathbb{R}^{D\times D}_{\geq 0}$ parameterizes the interaction. If W is the matrix of 1s, then each element of the D-vector $W\mathbf{r}(t)$ is the expected total population size of the focal process at time t, and death rates increase with this total size, enforcing a carrying capacity. Otherwise, the death rates are also sensitive to the expected relative frequency of each type. For example, if W is diagonally dominant, then the model includes negative frequency-dependent selection. Technically, to satisfy Assumption 7.1, we require that this linear term is truncated above some value of the expected total size. In practice, we take this cut-off to be very large, such that the numerical results below are not impacted.

Steady states induced by mean-field interaction

While the simple MTBDP displays only trivial steady states (or constant ones, if it is critical), the MTBDP with mean-field interaction admits more interesting behavior. Steady-state behavior can be examined by imposing a criticality condition on the self-consistent field. For Example 7.2, steady states $\mathbf{r}^* \in \mathbb{R}^D_{\geq 0}$ satisfy

$$\Big(\mathrm{diag}(\boldsymbol{\lambda} - \boldsymbol{\mu} - W\mathbf{r}^*) + \Gamma\Big)\mathbf{r}^* = 0. \tag{7.12}$$

Nontrivial solutions of this system of nonlinear algebraic equations for the critical field can be found numerically with standard root-finding methods, and are indeed steady states as long as the process is supercritical when the field vanishes. For results on steady-state solutions in strongly interacting MTBDPs, see [DDC18].

Numerical examples

For Example 7.2, the nonlinear moment equation (7.11) is of Ricatti type, with only quadratic nonlinearities. Figure 7.1 shows numerical results for the self-consistent field $\mathbf{r}$ of Example 7.2 with $D = 5$ types. The field in this case represents the vector of expected particle counts over the 5 types. These three examples model carrying capacity, negative frequency-dependent selection, and positive frequency-dependent selection, and all use the same rate parameters $\boldsymbol{\lambda}, \boldsymbol{\mu}, \Gamma$ (Figure 7.1A-C) but different W matrices. Without mean-field interaction ($W = 0$), this MTBDP is supercritical, and the expected particle counts grow exponentially (Figure 7.1D). One particle type has a higher birth rate than the others, so it grows faster. With a carrying capacity interaction (Figure 7.1E), the population reaches a stationary phase due to a mean-field interaction that increases death rates linearly with the expected population size. With negative frequency-dependent selection (via a diagonally dominant W), the types are more balanced (Figure 7.1F) because the death rate of a given type is suppressed only by growth of that type. With positive frequency-dependent selection (via a diagonally non-dominant W), the death rate of a given type is less suppressed by growth of that type than the others, leading to an enhancing effect on the type with the birth rate advantage (Figure 7.1G).

A Python implementation producing the results above is available at `https://github.com/WSDeWitt/mfbd`. This code is written in JAX [BFH+18] and relies on the Diffrax package [Kid21] for numerical ODE solutions. Specifically, to solve Riccati-type ODEs we use the Dormand-Prince 8/7 method [PD81]—a high-accuracy explicit Runge-Kutta solver—with Hermite interpolation for dense evaluation in the time domain. To adapt step sizes we use an I-controller [see HNW08, §II.4]. To solve the nonlinear algebraic equations for the critical field, we use root finding with automatic differentiation in JAXopt [BBC+21].

Regarding our approach for finding the solution in this specific example, we want to note that the moment-mediated interactions we study allow for direct solution of self-consistent fields via a nonlinear moment equation, which is amenable to standard numerical ODE techniques. Mean-field calculations in physical applications (typically on

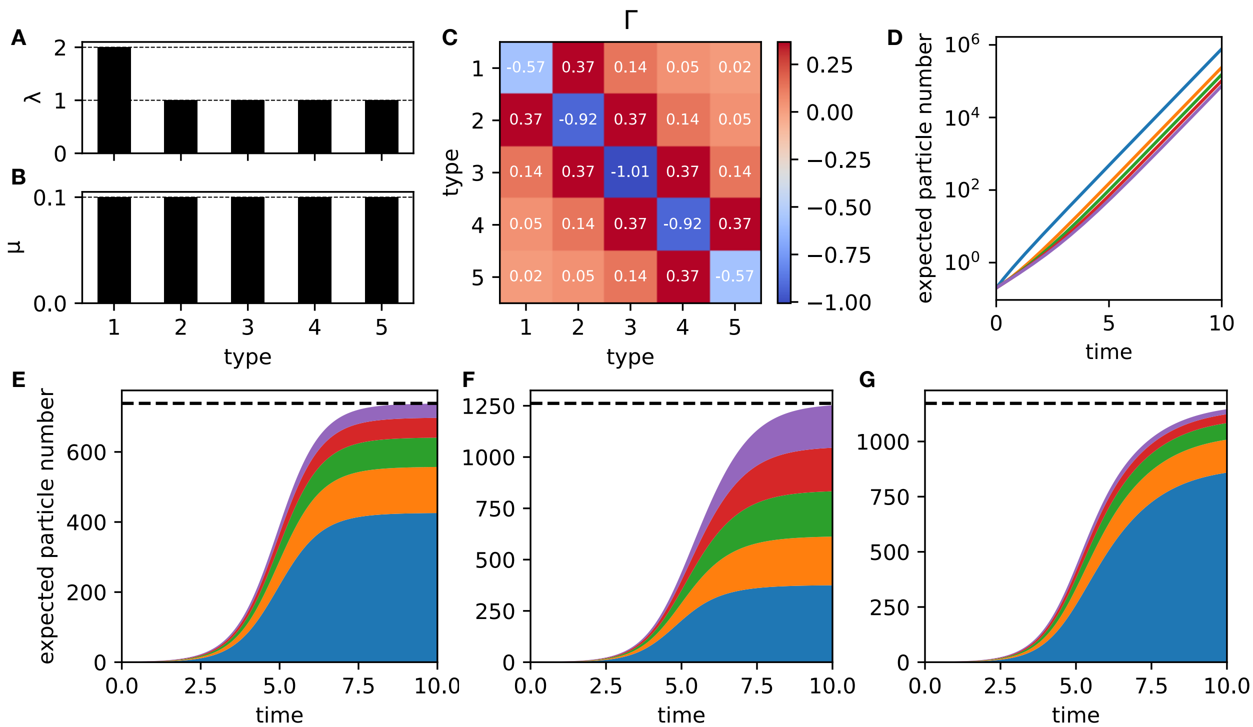

Figure 7.1: Numerical solutions of the self-consistent field $\mathbf{r}$ for the MTBDP with moment-mediated mean-field interaction (Example 7.2). In these examples, $D = 5$. (**A**.) Birth rates $\boldsymbol{\lambda}$, with type 1 elevated above the others. (**B**.) Death rate component $\boldsymbol{\mu}$, the same for all types. (**C**.) Type transition rate matrix Γ, of Toeplitz form, so that mutations between neighboring states are more likely. (**D**.) Expected particle count of each of the 5 types (colors) in the absence of any mean-field interaction, showing supercritical growth. **E**-F show stacked expected particle count of each of the 5 types, with various mean-field interactions of the form Example 7.2, with $\|W\|_F = 0.01$ in all cases. The dashed lines indicate the critical fields $\mathbf{r}^*$ computed by solving (7.12). (**E**.) Carrying capacity: $W \propto J$ (with J denoting the $D \times D$ matrix of 1s). (**F**.) Negative frequency-dependent selection: $W \propto I$. (**G**.) Positive frequency-dependent selection: $W \propto J - \frac{3}{5} I$.

continuous spaces with nonlinear PDEs) often rely on the *self-consistent field method*, which solves a sequence of linear systems with an external field that converges to a fixed point (for example, the Hartree-Fock, and density-functional theories for quantum many-body systems) [YS21, GKKR96]. Such methods tacitly assume a contractive mapping holds for this procedure, so that, by the Banach Fixed-Point Theorem, the field converges to a unique point. In practice, the method can suffer from slow convergence, non-convergence, or even divergence of the iterates, although there are several regularization techniques for controlling these issues. Our direct solution for the moment-mediated case avoids these issues.

Integrating mean-field interactions in phylogenetic birth-death models

Phylogenetic birth-death models augment the simple MTBDP with a sampling process that results in partially observed histories, and are considered as generative models for phylogenetic trees. They add two additional parameters: the sampling probability ρ gives the probability that any given particle at a specified final sampling time (the present) is sampled, and the fossilization probability σ gives the probability that a death event before the present is observed. The tree is then partially observed by pruning out all subtrees that are not ancestral to a sampled tip or fossil.

Computing likelihoods for rate parameters on phylogenetic trees requires marginalizing out all possible unobserved sub-histories, conditioned on the partially observed history. We briefly outline this calculation, augmented with mean-field interactions. We use notation like that of [KSVD16] and [BSVS18]. Given the parameters for the system in Example 7.1, and measuring time backward from the present sampling time, the probability density requires solving three coupled initial value problems (the standard case without mean-field interactions solves two systems).

First, the self-consistent fields $\mathbf{r}(t)$ are calculated as in Example 7.2 by solving a D-dimensional initial value problems (we reverse time such that the process starts at the tree root time $\tau > 0$, and ends at $t = 0$). Next, we need as an auxiliary calculation the probability $p_i(t)$ that a particle of type i at time t (before the present) will not be observed in the tree—that is, it will not be sampled and will not fossilize. These are given by the system of backward equations (of Riccati type)

$$\begin{aligned} p_i'(t) &= \lambda_i p_i(t)^2 - \left(\lambda_i + \mu_i + \sum_{j=1}^{D} W_{ij} r_j(t)\right) p_i(t) \\ &\quad + \sum_{j=1}^{D} \Gamma_{ij} p_j(t) + (1-\sigma)\left(\mu_i + \sum_{j=1}^{D} W_{ij} r_j(t)\right) \\ p_i(0) &= 1 - \rho, \end{aligned} \tag{7.13}$$

where $\mathbf{r}$ is given as in Example 7.1. These are solved on the interval $[0, \tau]$ where τ is the age of the root of the tree.

Finally, we compute the likelihood contribution for each of B tree branches $b = 1, \ldots, B \in \mathbb{N}$. Fixing some branch b with type i spanning the half-open interval $(t_1, t_2]$, let $q_i(t)$ denote its branch propagator, defined as the solution of the backward equation

$$q_i'(t) = \left(2\lambda_i p_i(t) + \Gamma_{ii} - \lambda_i - \mu_i - \sum_{j=1}^{D} W_{ij} r_j(t)\right) q_i(t)$$
$$q_i(t_1) = \begin{cases} \rho, & \text{if branch } b \text{ leads to a sample at } t_1 = 0 \\ \sigma\mu_i, & \text{if branch } b \text{ leads to a fossil at } t_1 > 0 \\ \lambda_i q_{\text{left}}(t_1) q_{\text{right}}(t_1), & \text{if branch } b \text{ splits at time } t_1 > 0 \\ \Gamma_{ij} q_j(t_1), & \text{if branch } b \text{ transitions to type } j \text{ at time } t_1 > 0 \end{cases} \tag{7.14}$$

where q_{left} and q_{right} denote the propagators of the left and right children of branch b. This system is coupled via the boundary conditions for each branch, and can be solved recursively by post-order tree traversal, yielding the tree likelihood accumulated at the root.

Standard phylogenetic birth-death models are recovered by setting $W = 0$, and only solving the **p** and **q** systems. By solving **p**, **q**, and **r** systems in the case $W \neq 0$, it is possible to compute tree likelihoods under phylogenetic birth-death processes that model interactions, while maintaining the efficient post-order calculation of likelihoods.

Bibliography

[ABL22] Amol Aggarwal, Charles Bordenave, and Patrick Lopatto. Mobility edge of Lévy matrices. *arXiv preprint arXiv:2210.09458*, 2022.

[ABP09] Antonio Auffinger, Gérard Ben Arous, and Sandrine Péché. Poisson convergence for the largest eigenvalues of heavy tailed random matrices. *Ann. Inst. Henri Poincaré Probab. Stat.*, 45(3):589–610, 2009.

[ADK21a] Johannes Alt, Raphael Ducatez, and Antti Knowles. Delocalization transition for critical Erdős–Rényi graphs. *Communications in Mathematical Physics*, 388(1):507–579, 2021.

[ADK21b] Johannes Alt, Raphael Ducatez, and Antti Knowles. Extremal eigenvalues of critical Erdős–Rényi graphs. *The Annals of Probability*, 49(3):1347–1401, 2021.

[ADK22] Johannes Alt, Raphael Ducatez, and Antti Knowles. The completely delocalized region of the Erdős–Rényi graph. *Electronic Communications in Probability*, 27, 2022.

[ADK23a] Johannes Alt, Raphael Ducatez, and Antti Knowles. Localized phase for the Erdős–Rényi graph. *arXiv preprint arXiv:2305.16294*, 2023.

[ADK23b] Johannes Alt, Raphael Ducatez, and Antti Knowles. Poisson statistics and localization at the spectral edge of sparse Erdős–Rényi graphs. *The Annals of Probability*, 51(1):277–358, 2023.

[AGH21] Fanny Augeri, Alice Guionnet, and Jonathan Husson. Large deviations for the largest eigenvalue of sub-Gaussian matrices. *Communications in Mathematical Physics*, 383(2):997–1050, 2021.

[ALM21] Amol Aggarwal, Patrick Lopatto, and Jake Marcinek. Eigenvector statistics of Lévy matrices. *Ann. Probab.*, 49(4):1778–1846, 2021.

[Alo98] Noga Alon. Spectral techniques in graph algorithms. In *LATIN'98: Theoretical Informatics: Third Latin American Symposium Campinas, Brazil, April 20–24, 1998 Proceedings 3*, pages 206–215. Springer, 1998.

[ALY21] Amol Aggarwal, Patrick Lopatto, and Horng-Tzer Yau. GOE statistics for Lévy matrices. *J. Eur. Math. Soc. (JEMS)*, 23(11):3707–3800, 2021.

[AM93] Michael Aizenman and Stanislav Molchanov. Localization at large disorder and at extreme energies: An elementary derivations. *Communications in Mathematical Physics*, 157:245–278, 1993.

[AMV+17] Assaf Amitai, Luka Mesin, Gabriel D Victora, Mehran Kardar, and Arup K Chakraborty. A population dynamics model for clonal diversity in a germinal center. *Front Microbiol.*, 8:1693, 2017.

[And58] Philip W Anderson. Absence of diffusion in certain random lattices. *Physical review*, 109(5):1492, 1958.

[Arn67] Ludwig Arnold. On the asymptotic distribution of the eigenvalues of random matrices. *Journal of Mathematical Analysis and Applications*, 20(2):262–268, 1967.

[Ash65] Robert B. Ash. *Information Theory*. Number 19 in Interscience tracts in pure and applied mathematics. Interscience, New York, 1965.

[Aug16] Fanny Augeri. Large deviations principle for the largest eigenvalue of Wigner matrices without Gaussian tails. *Electron. J. Probab.*, 21:49 pp., 2016.

[Aug20] Fanny Augeri. Nonlinear large deviation bounds with applications to Wigner matrices and sparse Erdős–Rényi graphs. *The Annals of Probability*, 48(5):2404–2448, 2020.

[Aug24] Fanny Augeri. Large deviations of the empirical spectral measure of supercritical sparse wigner matrices. *arXiv preprint arXiv:2401.11925*, 2024.

[BADG01] Gérard Ben Arous, Amir Dembo, and Alice Guionnet. Aging of spherical spin glasses. *Probability Theory and Related Fields*, 120(1):1–67, 2001.

[BAG97] Gérard Ben Arous and Alice Guionnet. Large deviations for Wigner's law and Voiculescu's non-commutative entropy. *Probability Theory and Related Fields*, 108(4):517–542, 1997.

[Bam20] Bassam Bamieh. A tutorial on matrix perturbation theory (using compact matrix notation). *arXiv preprint arXiv:2002.05001*, 2020.

[Bau85] Hellmut Baumgärtel. Analytic perturbation theory for matrices and operators. *Operator theory*, 15, 1985.

[BBC+21] Mathieu Blondel, Quentin Berthet, Marco Cuturi, Roy Frostig, Stephan Hoyer, Felipe Llinares-López, Fabian Pedregosa, and Jean-Philippe Vert. Efficient and modular implicit differentiation. *arXiv preprint arXiv:2105.15183*, 2021.

[BBG21] Bhaswar B Bhattacharya, Sohom Bhattacharya, and Shirshendu Ganguly. Spectral edge in sparse random graphs: Upper and lower tail large deviations. *The Annals of Probability*, 49(4):1847–1885, 2021.

[BC14] Charles Bordenave and Pietro Caputo. A large deviation principle for Wigner matrices without Gaussian tails. *Annals of Probability*, 42(6):2454–2496, 2014.

[BC15] Charles Bordenave and Pietro Caputo. Large deviations of empirical neighborhood distribution in sparse random graphs. *Probability Theory and Related Fields*, 163(1-2):149–222, 2015.

[BFH+18] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018.

[BG13] Charles Bordenave and Alice Guionnet. Localization and delocalization of eigenvectors for heavy-tailed random matrices. *Probability Theory and Related Fields*, 157(3-4):885–953, 2013.

[BG17] Charles Bordenave and Alice Guionnet. Delocalization at small energy for heavy-tailed random matrices. *Communications in Mathematical Physics*, 354:115–159, 2017.

[BG20] Bhaswar B Bhattacharya and Shirshendu Ganguly. Upper tails for edge eigenvalues of random graphs. *SIAM Journal on Discrete Mathematics*, 34(2):1069–1083, 2020.

[BGBK19] Florent Benaych-Georges, Charles Bordenave, and Antti Knowles. Largest eigenvalues of sparse inhomogeneous Erdős–Rényi graphs graphs. *Ann. Probab.*, 47(3):1653–1676, 2019.

[BGBK20] Florent Benaych-Georges, Charles Bordenave, and Antti Knowles. Spectral radii of sparse random matrices. *Ann. Inst. Henri Poincaré Probab. Stat.*, 56(3):2141–2161, 2020.

[BGGM14] Florent Benaych-Georges, Alice Guionnet, and Camille Male. Central limit theorems for linear statistics of heavy tailed random matrices. *Communications in Mathematical Physics*, 329(2):641–686, 2014.

[BGPW19] Ellen Baake, Adrián González Casanova, Sebastian Probst, and Anton Wakolbinger. Modelling and simulating Lenski's long-term evolution experiment. *Theor. Popul. Biol.*, 127:58–74, 2019.

[Bil99] Patrick Billingsley. *Convergence of probability measures*. John Wiley & Sons, Inc., New York, second edition, 1999.

[BMdlP23] Milad Bakhshizadeh, Arian Maleki, and Victor H de la Pena. Sharp concentration results for heavy-tailed distributions. *Information and Inference: A Journal of the IMA*, 12(3):iaad011, 2023.

[Bol01] Béla Bollobás. *Random Graphs*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2 edition, 2001.

[BSV17] Charles Bordenave, Arnab Sen, and Bálint Virág. Mean quantum percolation. *Journal of the European Mathematical Society*, 19(12):3679–3707, 2017.

[BSVS18] J. Barido-Sottani, T. G. Vaughan, and T. Stadler. Detection of HIV transmission clusters from phylogenetic trees using a multi-state birth–death model. *J. R. Soc. Interface*, 15(146):20180512, 2018.

[BY88] Z. D. Bai and Y. Q. Yin. Necessary and sufficient conditions for almost sure convergence of the largest eigenvalue of a Wigner matrix. *The Annals of Probability*, 16(4):1729 – 1741, 1988.

[CD16] Sourav Chatterjee and Amir Dembo. Nonlinear large deviations. *Adv. Math.*, 299:396–450, 2016.

[CD20] Nicholas Cook and Amir Dembo. Large deviations of subgraph counts for sparse erdős–rényi graphs. *Advances in Mathematics*, 373:107289, 2020.

[CD22a] Louis-Pierre Chaintron and Antoine Diez. Propagation of chaos: a review of models, methods and applications. I. Models and methods. *Kinet. Relat. Models*, 15(6):895–1015, 2022.

[CD22b] Louis-Pierre Chaintron and Antoine Diez. Propagation of chaos: a review of models, methods and applications. II. Applications. *Kinet. Relat. Models*, 15(6):1017–1173, 2022.

[CDG23] Nicholas A Cook, Raphael Ducatez, and Alice Guionnet. Full large deviation principles for the largest eigenvalue of sub-Gaussian Wigner matrices. *arXiv preprint arXiv:2302.14823*, 2023.

[Chu97] Fan RK Chung. *Spectral graph theory*, volume 92. American Mathematical Soc., 1997.

[CMS14] Forrest W. Crawford, Vladimir N. Minin, and Marc A. Suchard. Estimation for general birth-death processes. *J. Amer. Statist. Assoc.*, 109(506):730–747, 2014.

[CV11] Sourav Chatterjee and S. R. S. Varadhan. The large deviation principle for the Erdős-Rényi random graph. *European J. Combin.*, 32(7):1000–1017, 2011.

[CV12] Sourav Chatterjee and S. R. S. Varadhan. Large deviations for random matrices. *Comm. Stoch. Analysis*, 6(1):1–13, 2012.

[DDC18] Renaud Dessalles, Maria D'Orsogna, and Tom Chou. Exact steady-state distributions of multispecies birth–death–immigration processes: Effects of mutations and carrying capacity on diversity. *J. Stat. Phys.*, 173:182–221, 2018.

[DEHH24] William S. DeWitt, Steven N. Evans, Ella Hiesmayr, and Sebastian Hummel. Mean-field interacting multi-type birth-death processes with a view to applications in phylodynamics. *Theoretical Population Biology*, 159:1–12, 2024.

[Dei77] K. Deimling. *Ordinary Differential Equations In Banach Spaces*. Springer, 1977.

[DM10] Amir Dembo and Andrea Montanari. Gibbs measures and phase transitions on sparse random graphs. *Brazilian Journal of Probability and Statistics*, 24(2), 2010.

[DMV+18] William S DeWitt, Luka Mesin, Gabriel D Victora, Vladimir N Minin, and Frederick A Matsen IV. Using genotype abundance to improve phylogenetic inference. *Mol. Biol. Evol.*, 35(5):1253–1265, 2018.

[DS20] Jian Ding and Charles K Smart. Localization near the edge for the Anderson Bernoulli model on the two dimensional lattice. *Inventiones mathematicae*, 219:467–506, 2020.

[EK86] S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley Ser. Probab. Math. Stat. Wiley, Chichester, 1986.

[EKYY13] László Erdős, Antti Knowles, Horng-Tzer Yau, and Jun Yin. Spectral statistics of Erdős–Rényi graphs I: Local semicircle law. *Annals of Probability*, 41:2279—-2375, 2013.

[ER60] Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60, 1960.

[EYY12] László Erdős, Horng-Tzer Yau, and Jun Yin. Rigidity of eigenvalues of generalized wigner matrices. *Advances in Mathematics*, 229(3):1435–1515, 2012.

[Fel68] W. Feller. *An Introduction to Probability Theory and Its Applications, Volume 1*. Wiley, third edition, 1968.

[FK81] Zoltán Füredi and John Komlos. The eigenvalues of random symmetric matrices. *Combinatorica*, 1:233–241, 1981.

[FS83] Jürg Fröhlich and Thomas Spencer. Absence of diffusion in the Anderson tight binding model for large disorder or low energy. *Communications in Mathematical Physics*, 88(2):151–184, 1983.

[GCPP21] Adrián González Casanova, Juan Carlos Pardo, and José Luis Pérez. Branching processes with interactions: subcritical cooperative regime. *Adv. Appl. Probab.*, 53(1):251–278, 2021.

[GH20] Alice Guionnet and Jonathan Husson. Large deviations for the largest eigenvalue of Rademacher matrices. *The Annals of Probability*, 48(3):1436–1465, 2020.

[GHN24] Shirshendu Ganguly, Ella Hiesmayr, and Kyeongsik Nam. Spectral large deviations of sparse random matrices. *Journal of the London Mathematical Society*, 110(1):e12954, 2024.

[GKKR96] Antoine Georges, Gabriel Kotliar, Werner Krauth, and Marcelo J. Rozenberg. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Rev. Mod. Phys.*, 68:13–125, Jan 1996.

[GLO20] Anne Greenbaum, Ren-cang Li, and Michael L Overton. First-order perturbation theory for eigenvalues and eigenvectors. *SIAM review*, 62(2):463–482, 2020.

[Gly87] Peter W Glynn. Upper bounds on Poisson tail probabilities. *Operations Research Letters*, 6(1):9–14, 1987.

[GMP77] Ilya Goldsheid, Stanislav Molchanov, and Leonid Pastur. A random homogeneous schrödinger operator has a pure point spectrum. *Funct. Anal. Appl*, 11(1):1–10, 1977.

[GN22] Shirshendu Ganguly and Kyeongsik Nam. Large deviations for the largest eigenvalue of Gaussian networks with constant average degree. *Probability Theory and Related Fields*, 184(3):613–679, December 2022.

[Gri73] D. A. Griffiths. Multivariate Birth-and-Death Processes as Approximations to Epidemic Processes. *J. Appl. Probab.*, 10(1):15–26, 1973.

[Gui21] Alice Guionnet. Bernoulli random matrices. *arXiv preprint arXiv:2112.05506*, 2021.

[HK21] Yukun He and Antti Knowles. Fluctuations of extreme eigenvalues of sparse Erdős–Rényi graphs. *Probability Theory and Related Fields*, 180(3-4):985–1056, 2021.

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

[HLY20] Jiaoyang Huang, Benjamin Landon, and Horng-Tzer Yau. Transition from Tracy–Widom to Gaussian fluctuations of extremal eigenvalues of sparse Erdős–Rényi graphs. *Annals of Probability*, 2020.

[HM23] Ella Hiesmayr and Theo McKenzie. The Spectral Edge of Constant Degree Erdős–Rényi graphs, September 2023. arXiv:2309.11007 [math-ph].

[HNW08] E. Hairer, S.P. Nørsett, and G. Wanner. *Solving Ordinary Differential Equations I Nonstiff Problems*. Springer, Berlin, second revised edition edition, 2008.

[HXC+18] Lam Si Tung Ho, Jason Xu, Forrest W Crawford, Vladimir N Minin, and Marc A Suchard. Birth/birth-death processes and their computable transition probabilities with biological applications. *J. Math. Biol.*, 76:911–944, 2018.

[Igl64] Donald L. Iglehart. Multivariate Competition Processes. *Ann. Math. Statist.*, 35(1):350–361, 1964.

[JRL11] Svante Janson, Andrzej Rucinski, and Tomasz Luczak. *Random graphs*. John Wiley & Sons, 2011.

[Juh81] Ferenc Juhász. On the spectrum of a random graph. In János Bolyai Mathematical Society, editor, *Algebraic methods in graph theory (Coll. Math. Soc. J. Bolyai)*, volume 25, pages 313–316. North-Holland Publishing Co., Amsterdam; Budapest, 1981.

[Kad09] Leo P Kadanoff. More is the same; phase transitions and mean field theories. *J. Stat. Phys.*, 137:777–797, 2009.

[Kal21] Olav Kallenberg. *Foundations of Modern Probability*, volume 99 of *Probability Theory and Stochastic Modelling*. Springer, Cham, third edition, 2021.

[Ken48] David G. Kendall. On the Generalized "Birth-and-Death" Process. *Ann. Math. Statist.*, 19(1):1–15, 1948.

[Kes59] Harry Kesten. Symmetric random walks on groups. *Transactions of the American Mathematical Society*, 92(2):336–354, 1959.

[Kho01] A. Khorunzhy. Sparse random matrices: Spectral edge and statistics of rooted trees. *Advances in Applied Probability*, 33(1):124–140, 2001.

[Kid21] Patrick Kidger. *On Neural Differential Equations*. PhD thesis, University of Oxford, 2021.

[KS97] Tsampikos Kottos and Uzy Smilansky. Quantum chaos on graphs. *Physical review letters*, 79(24):4794, 1997.

[KS03] Michael Krivelevich and Benny Sudakov. The Largest Eigenvalue of Sparse Random Graphs. *Combinatorics, Probability and Computing*, 12(1):61–72, January 2003. Publisher: Cambridge University Press.

[KSV04] Oleksiy Khorunzhy, Mariya Shcherbina, and Valentin Vengerovsky. Eigenvalue distribution of large weighted random graphs. *Journal of Mathematical Physics*, 45(4):1648–1672, 2004.

[KSVD16] Denise Kühnert, Tanja Stadler, Timothy G. Vaughan, and Alexei J. Drummond. Phylodynamics with Migration: A Computational Framework to Quantify Population Structure from Genomic Data. *Mol. Biol. Evol.*, 33(8):2102–2116, 04 2016.

[Kur98] T. G. Kurtz. Martingale problems for conditional distributions of Markov processes. *Electron. J. Probab.*, 3:no. 9, 29 pp., 1998.

[Kur11] T. G. Kurtz. *Equivalence of Stochastic Equations and Martingale Problems*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[LLQ23] Xin Liu, Bo Liu, and Hai Qi. Germinal center reaction and output: recent advances. *Curr. Opin. Immunol.*, 82(102308):102308, April 2023.

[LS18] Ji Oon Lee and Kevin Schnelli. Local law and tracy–widom limit for sparse random matrices. *Probability Theory and Related Fields*, 171(1):543–616, 2018.

[LvHY18] Rafał Latała, Ramon van Handel, and Pierre Youssef. The dimension-free structure of nonhomogeneous random matrices. *Inventiones mathematicae*, 214:1031–1080, 2018.

[LZ15] Eyal Lubetzky and Yufei Zhao. On replica symmetry of large deviations in random graphs. *Random Structures Algorithms*, 47(1):109–146, 2015.

[LZ17] Eyal Lubetzky and Yufei Zhao. On the variational problem for upper tails in sparse random graphs. *Random Structures & Algorithms*, 50(3):420–436, 2017.

[MEV16] Luka Mesin, Jonatan Ersching, and Gabriel D Victora. Germinal center B cell dynamics. *Immunity*, 45(3):471–482, September 2016.

[MLM+21] Ailene MacPherson, Stilianos Louca, Angela McLaughlin, Jeffrey B Joy, and Matthew W Pennell. Unifying Phylogenetic Birth–Death Models in Epidemiology and Macroevolution. *Syst. Biol.*, 71(1):172–189, 06 2021.

[MS65] Theodore S Motzkin and Ernst G Straus. Maxima for graphs and a new proof of a theorem of Turán. *Canadian Journal of Mathematics*, 17:533–540, 1965.

[MSE+20] Luka Mesin, Ariën Schiepers, Jonatan Ersching, Alexandru Barbulescu, Cecília B. Cavazzoni, Alessandro Angelini, Takaharu Okada, Tomohiro Kurosaki, and Gabriel D. Victora. Restricted Clonality and Limited Germinal Center Reentry Characterize Memory B Cell Reactivation by Boosting. *Cell*, 180(1):92–106.e11, 2020.

[NRS14] R. A. Neher, C. A. Russell, and B. I. Shraiman. Predicting evolution from the shape of genealogical trees. *eLife*, 3:e03568, 2014.

[Ove95] L Overbeck. Superprocesses and McKean-Vlasov equations with creation of mass. Technical Report 428, University of California, Berkeley, 1995.

[Ove96] L. Overbeck. Nonlinear superprocesses. *Ann. Probab.*, 24(2):743–760, 1996.

[Par07] Giorgio Parisi. Course 3 Mean field theory of spin glasses: Statics and dynamics. In Jean-Philippe Bouchaud, Marc Mézard, and Jean Dalibard, editors, *Complex Systems*, volume 85 of *Les Houches*, pages 131–178. Elsevier, 2007.

[PD81] P. J Prince and J. R. Dormand. High order embedded Runge–Kutta formulae. *J. Comp. Appl. Math*, 7(1):67–75, 1981.

[PJV21] Juhee Pae, Johanne T Jacobsen, and Gabriel D Victora. Imaging the different timescales of germinal center selection. *Immunol. Rev.*, November 2021.

[Rei00] David Reimer. Proof of the van den Berg-Kesten conjecture. *Combinatorics, Probability and Computing*, 9(1):27–32, 2000.

[Reu61] GEH Reuter. Competition processes. In *Proc. 4th Berkeley Symp. Math. Statist. Prob*, volume 2, pages 421–430, 1961.

[SLW19] Mark J Shlomchik, Wei Luo, and Florian Weisel. Linking signaling and selection in the germinal center. *Immunol. Rev.*, 288(1):49–63, March 2019.

[Sos99] Alexander Soshnikov. Universality at the edge of the spectrum in Wigner random matrices. *Communications in Mathematical Physics*, 207:697–733, 1999.

[Sos04] Alexander Soshnikov. Poisson statistics for the largest eigenvalues of Wigner random matrices with heavy tails. *Electron. Comm. Probab.*, 9:82–91, 2004.

[Ste15] James Stewart. *Multivariable calculus*. Cengage learning, 2015.

[Szn91] Alain-Sol Sznitman. Topics in propagation of chaos. In *École d'Été de Probabilités de Saint-Flour XIX—1989*, volume 1464 of *Lecture Notes in Math.*, pages 165–251. Springer, Berlin, 1991.

[Tal95] Michel Talagrand. Concentration of measure and isoperimetric inequalities in product spaces. *Publications Mathématiques de l'Institut des Hautes Etudes Scientifiques*, 81:73–205, 1995.

[Tao22] Terence Tao. An Improvement to Bennett's Inequality for the Poisson Distribution, 2022. https://terrytao.wordpress.com/2022/12/13/an-improvement-to-bennetts-inequality-for-the-poisson-distribution/ Accessed 2023-06-01.

[Tha15] Marie-Noémie Thai. Birth and death process in mean field type interaction, 2015.

[TMP+16] Jeroen M. J. Tas, Luka Mesin, Giulia Pasqual, Sasha Targ, Johanne T. Jacobsen, Yasuko M. Mano, Casie S. Chen, Jean-Claude Weill, Claude-Agnès Reynaud, Edward P. Browne, Michael Meyer-Hermann, and Gabriel D. Victora. Visualizing antibody affinity maturation in germinal centers. *Science*, 351(6277):1048–1054, 2016.

[TY21] Konstantin Tikhomirov and Pierre Youssef. Outliers in spectrum of sparse wigner matrices. *Random Structures & Algorithms*, 58(3):517–605, 2021.

[V+09] Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.

[VM14] Gabriel D Victora and Luka Mesin. Clonal and cellular dynamics in germinal centers. *Curr. Opin. Immunol.*, 28:90–96, June 2014.

[VN22] Gabriel D Victora and Michel C Nussenzweig. Germinal centers. *Annu. Rev. Immunol.*, 40:413–442, 2022.

[Wig58] Eugene P. Wigner. On the distribution of the roots of certain symmetric matrices. *Annals of Mathematics*, 67(2):325–327, 1958.

[Wil05] Herbert S Wilf. *generatingfunctionology*. CRC press, third edition, 2005.

[XGKMM15] Jason Xu, Peter Guttorp, Midori Kato-Maeda, and Vladimir N. Minin. Likelihood-based inference for discretely observed birth–death-shift processes, with applications to evolution of mobile genetic elements. *Biometrics*, 71(4):1009–1021, 2015.

[YS21] Sarath Yasodharan and Rajesh Sundaresan. The four levels of fixed-points in mean-field models. In *2021 National Conference on Communications (NCC)*, pages 1–6. IEEE, 2021.

[Zak06] Inna Zakharevich. A generalization of Wigner's law. *Communications in Mathematical Physics*, 268:403–414, 2006.

Appendix A

Proofs

A.1 Structure surrounding vertices

Proof of Lemma 3.5, 1. We essentially follow the proof of Lemma 7.3 in [ADK21b]. If the balls of radius r around vertices of $\mathcal{V}$ are not disjoint there must be two vertices x and y in $\mathcal{V}$, connected by a simple path of length at least 1 (if the two vertices are connected by an edge) and at most $2r$. Let us denote the vertices on this path by $z_1, \dots, z_s$, where s can range from 0 to $2r-1$. If we define $z_0 = x$ and $z_{s+1} = y$, then we must have an edge from z_i to z_{i+1} for $i = 0, \dots, s$. Eventually we will take a union bound over all x, y and paths between them, but let us first compute the probability of a single such event.

For convenience of notation let $\tau = \mathfrak{u} - \mathfrak{u}^{2/3}$, the degree lower bound for the intermediate regime. For a given s as well as distinct vertices $x, y, z_1, \dots, z_s$ in $[N]$,

$$\begin{aligned}
&\mathbb{P}\left(x, y \in \mathcal{V}, (z_i, z_{i+1}) \in E(G) \text{ for } i \in \{0, \dots, s\}\right) \\
&\leq \mathbb{P}\left(\left|\Gamma_x \backslash \{y, z_1\}\right| \geq \tau - 2, \left|\Gamma_y \backslash \{x, z_s\}\right| \geq \tau - 2, (z_i, z_{i+1}) \in E(G) \text{ for } i \in \{0, \dots, s\}\right) \\
&= \mathbb{P}\left(\left|\Gamma_x \backslash \{y, z_1\}\right| \geq \tau - 2\right) \mathbb{P}\left(\left|\Gamma_y \backslash \{x, z_s\}\right| \geq \tau - 2\right) \left(\frac{d}{N}\right)^{s+1}
\end{aligned}$$

as now all events are independent and the path contains $s+1$ edges.

As $|\Gamma_x \backslash \{y, z_1\}|$ and $|\Gamma_y \backslash \{x, z_s\}|$ are distributed as $\mathrm{Bin}(N-2, d/N)$, we can use Lemma 4.6 to get

$$\mathbb{P}\left(|\Gamma_x \backslash \{y, z_1\}| \geq \tau - 2\right) \leq (1 + o_N(1)) e^{-\tau \log \tau + c\tau \max\{\log d, 1\}} \leq e^{O(\mathfrak{u}^{2/3}) - \log N}$$

by Lemma 3.4.

Thus by using a union bound, we can bound the probability that two balls of radius r around two vertices in $\mathcal{V}$ intersect, by considering that we can choose x, y in $\binom{N}{2}$ ways and then we need to choose the path between x and y, i.e. for some s between 0 and $2r$, we need s ordered vertices, for which there are $(N-2)_s$ ways. Combining this with the above bounds gives that the probability of two vertices in $\mathcal{V}$ being connected by such a path is bounded by

$$\binom{N}{2}\sum_{s=0}^{2r}(N-2)_s\left(\frac{d}{N}\right)^{s+1}e^{-(2+o_N(1))\log N} \leq N^2\sum_{s=0}^{2r}N^s\left(\frac{d}{N}\right)^{s+1}e^{-(1+o_N(1))2\log N}$$
$$\leq f(d,r)e^{-(1+o_N(1))\log N}$$

where we bounded $\sum_{s=0}^{2r} d^{s+1}$ by $f(d,r) := d\frac{d^{2r+1}-1}{d-1}$ when $d \neq 1$ and $f(d,r) = 2r+1$ when $d = 1$. Thus for any constant r the event holds with high probability. □

In order to prove that the balls around vertices in the fine regime are with high probability trees, we start by bounding the probability that the ball around a fixed vertex contains m excess edges, this result and its proof are almost identical to Lemma 5.5 in [ADK21b] but we chose to include them for sake of completeness.

Lemma A.1. For a vertex $x \in [N]$, any integer $C_1 \geq 1$ and any constant s it holds that

$$\mathbb{P}\Big(E(B_s(x)) \geq V(B_s(x)) - 1 + C_1 \,\Big|\, S_1(x)\Big) \leq Cd^{3C_1}\left(\frac{|S_1|}{N}\right)^{C_1},$$

where C is a constant that depends on the constants s and C_1.

Proof. We use the argument from Lemma 5.5 in [ADK21b], but obtain a slightly different bound that is better suited for our regime of d.

Let T be a spanning tree of $B_s(x)$. If $B_s(x)$ contains at least C_1 excess edges, there are C_1 edges in $B_s(x)$ not contained in T, denote those by E_E. Let V_E denote the vertices incident to those edges and E_P the edges on the unique paths in T from x to the vertices of V_E. Finally let V_P denote the vertices incident to edges in E_P. (See Figure 4 in [ADK21b] for an illustration.) We define H to be the graph with vertices $V_E \cup V_P$ and edges $E_E \cup E_P$.

Let $S_r^F(x)$ denote that sphere of radius r around x in the graph F. Then the graph H is a graph on the vertices $[N]$ that satisfies the following properties:

1. $x \in F$,
2. $S_1^F(x) \subseteq S_1^G(x)$,
3. $|S_1^F(x)| \geq 1$
4. $E(F) = V(F) - 1 + C_1$ and
5. $V(F) \leq 2C_1 r + 1$.

The last property holds since the edges E_E incident to at most $2C_1$ distinct vertices and the paths in T from x to those vertices are of length at most r, which implies that V_P contains at most $2C_1 r$ additional distinct vertices besides x.

Thus we can bound the probability that $E(B_s(x)) \geq V(B_s(x)) - 1 + C_1$ by the probability that $B_s(x)$ contains a subgraph F satisfying properties 1.-5. above. For a given $x \in \mathcal{V}$, we start by conditioning on $S_1(x)$. Then let $\mathbb{F}(x)$ denote the subgraphs satisfying properties 1.-5. Recalling that G denotes our Erdős-Rényi graph $\mathcal{G}\left(N, \frac{d}{N}\right)$, we can bound

$$\begin{aligned}\mathbb{P}\left(E\left(B_s(x)\right) \geq V\left(B_s(x)\right) - 1 + C_1 | S_1(x)\right) &\leq \mathbb{P}\left(\cup_{F\in\mathbb{F}(x)} F \subseteq G \middle| S_1(x)\right) \\ &\leq \sum_{F\in\mathbb{F}(x)} \mathbb{P}\left(F \subseteq G \middle| S_1(x)\right).\end{aligned}$$

Now note that conditioned on S_1 we can construct any graph $F \in \mathbb{F}(x)$ by first choosing $1 \leq s \leq C_1$ vertices from $S_1(x)$, then choosing $0 \leq t \leq 2C_1 r - s$ (we lose the +1 since x is always part of F) vertices from $[N]\backslash\mathcal{B}_1(x)$, and then building a tree with these $s + t + 1$ vertices such that the first s are neighbors of x and the remaining t vertices connect to that graph (but not to x), and then adding C_1 additional edges. We can bound the number of such graphs by the number of labeled trees on $s+t+1$ vertices (for which Cayley's formula gives that there are $(s + t + 1)^{s+t-1}$) times the number of ways of choosing C_1 edges, which can be bounded by $(s+t+1)^{2C_1}$. The probability that such a graph is contained in G is then equal to $\left(\frac{d}{N}\right)^{t+C_1}$ since the number of edges in F without those between x and vertices in $S_1(x)$ is equal to $t + C_1$. So continuing from above, we get

$$\begin{aligned}&\leq \sum_{s=1}^{C_1} \sum_{t=0}^{2C_1 r-s} \binom{|S_1|}{s} \binom{N - |S_1(x)| - 1}{t} (s + t + 1)^{s+t-1+2C_1} \left(\frac{d}{N}\right)^{t+C_1} \\
&\leq \sum_{s=1}^{C_1} \sum_{t=0}^{2C_1 r-s} \frac{|S_1|^s}{s!} \frac{N^t}{t!} (s + t + 1)^{s+t+2C_1-1} \left(\frac{d}{N}\right)^{t+C_1} \\
&\leq \frac{1}{N^{C_1}} \left(d(2C_1 r + 1)^2\right)^{C_1} \sum_{s=1}^{C_1} \frac{|S_1|^s}{s!} (2C_1 r + 1)^s \sum_{t=0}^{2C_1 r-s} \frac{1}{t!} (d(2C_1 r + 1))^t \\
&\leq \frac{1}{N^{C_1}} \left(d(2C_1 r + 1)^2\right)^{C_1} C_1 \left(|S_1|(2C_1 r + 1)\right)^{C_1} 2C_1 r \left((d(2C_1 r + 1))^2\right)^{2C_1 r} \\
&\leq \frac{1}{N^{C_1}} 2rC_1^2 d^{3C_1} (2C_1 r + 1)^{5C_1} |S_1|^{C_1} \\
&\leq Cd^{3C_1} \left(\frac{|S_1|}{N}\right)^{C_1},\end{aligned}$$

where C is a constant that depends on the constants r and C_1. □

Proof of Lemma 3.5, 2. Note that by a union bound over all vertices, we get

$$\begin{aligned}&\mathbb{P}\left(\exists x \in \mathcal{V} : \mathcal{B}_r(x) \text{ is not a tree}\right) \\
&\leq \sum_{x\in[N]} \mathbb{P}\left(\mathcal{B}_r(x) \text{ is not a tree}, x \in \mathcal{V}, \mathfrak{u} - \mathfrak{u}^{\frac{2}{3}} \leq \alpha_x < 2\frac{\log N}{\log\log N}\right) + \mathbb{P}\left(\alpha_x \geq 2\frac{\log N}{\log\log N}\right)\end{aligned}$$

$$
\begin{aligned}
&\leq \sum_{x\in[N]} \mathbb{E}\left[\mathbb{P}\left(\mathcal{B}_r(x) \text{ is not a tree}|S_1(x)\right)\mathbf{1}\left(\mathfrak{u}-\mathfrak{u}^{\frac{2}{3}} \leq \alpha_x < 2\frac{\log N}{\log\log N}\right)\right] + \mathbb{P}\left(d_x \geq 2\frac{\log N}{\log\log N}\right)\\
&\leq \sum_{x\in[N]} \mathbb{E}\left[Cd^{3C_1}\frac{|S_1|}{N}\mathbf{1}\left(\mathfrak{u}-\mathfrak{u}^{\frac{2}{3}} \leq \alpha_x < 2\frac{\log N}{\log\log N}\right)\right] + \mathbb{P}\left(d_x \geq 2\frac{\log N}{\log\log N}\right)\\
&\leq NCd^{3C_1}\frac{2\log N}{N}\frac{e^{d+\mathfrak{u}^{\frac{2}{3}}\log\mathfrak{u}}}{N} + N^{-\frac{1}{2}},
\end{aligned}
$$

where we apply the bound from Lemma A.1 with $C_1 = 1$, as well as the bound on $|\mathcal{V}|$ from 3.4 and then used Lemma 4.6 for the second term. □

Proof of Lemma 3.5, 3. We show this for $\mathcal{V}$. The proof for $\mathcal{W}$ is identical. For a vertex x, and constants C_i that will be set later, let us define the events

$$\mathcal{G}_i(x) := \left\{\left||S_i(x)| - d^{i-1}\alpha_x\right| \leq C_i\left(d^{i-\frac{3}{2}}+1\right)\mathfrak{u}^{\frac{7}{8}}\right\}$$

and $\mathcal{F}_i(x) := \bigcap_{j=1}^{i}\mathcal{G}_i(x)$. We will write $\mathcal{G}_i$ and $\mathcal{F}_i$ whenever it is clear from context which vertex they refer to. First note that under $\mathcal{F}_i(x)$, $|B_i(x)| \leq \sqrt{N}$.

Now fix a vertex x. $\mathcal{G}_1$ holds trivially by the definition of α_x.

For $i \geq 2$ we now first show that conditional on S_1 the probability that S_i is large given that S_{i-1} is small is small. More precisely we show that

$$\mathbb{P}\left(\mathcal{G}_i^c \cap \mathcal{F}_{i-1}|S_1\right) \leq 2\exp\left\{-\mathfrak{u}^{\frac{3}{4}}\right\} \tag{A.1}$$

To show the above equation first observe that conditioned on B_{i-1}, S_i consists of all the neighbors of vertices in S_{i-1} that are not in B_{i-1}. Thus, conditionally on B_{i-1}, $|S_i|$ is distributed as $\mathrm{Binom}(|S_{i-1}|(N-|B_{i-1}|), d/N)$.

Note that this implies that

$$\mathbb{E}\Big[|S_i|\,\Big|\,B_{i-1}\Big] = d|S_{i-1}| - d\frac{|S_{i-1}||B_{i-1}|}{N}. \tag{A.2}$$

Thus under the event $\mathcal{F}_i$, by Lemma 4.7, because $|B_i| \leq \sqrt{N}$,

$$
\begin{aligned}
&\mathbb{P}\left(\left||S_i| - \mathbb{E}\left[|S_i|\,\middle|\,B_{i-1}\right]\right| \geq \sqrt{d|S_{i-1}|}\mathfrak{u}^{\frac{3}{8}} + \mathfrak{u}^{\frac{7}{8}}\middle|B_{i-1}\right)\\
&\leq 2\exp\left\{-\frac{\left(\sqrt{d|S_{i-1}|}\mathfrak{u}^{\frac{3}{8}}+\mathfrak{u}^{\frac{7}{8}}\right)^2}{2\left(d|S_{i-1}| - d\frac{|S_{i-1}||B_{i-1}|}{N}\right) + \frac{2}{3}\left(\sqrt{d|S_{i-1}|}\mathfrak{u}^{\frac{3}{8}}+\mathfrak{u}^{\frac{7}{8}}\right)}\right\}\\
&\leq 2\exp\left\{-\frac{\left(\sqrt{d|S_{i-1}|}\mathfrak{u}^{\frac{3}{8}}+\mathfrak{u}^{\frac{7}{8}}\right)^2}{2\left(d|S_{i-1}| - d\right) + \frac{2}{3}\left(\sqrt{d|S_{i-1}|}\mathfrak{u}^{\frac{3}{8}}+\mathfrak{u}^{\frac{7}{8}}\right)}\right\}
\end{aligned}
$$

$$\leq 2\exp\left\{-C\mathfrak{u}^{\frac{3}{4}}\right\}$$

for some constant C that does not depend on i.

Now we need to transform the above inequality into the one that we are actually trying to prove. For this we need to estimate some quantities: Let us define $\delta_{i-1} = |S_{i-1}| - d^{i-2}\alpha_x$, note that under $\mathcal{F}_{i-1}$

$$d\delta_{i-1} \leq 2C_{i-1}(d^{i-\frac{3}{2}} + 1)u^{\frac{7}{8}}$$

and generally $d \leq \left(d^{i-\frac{3}{2}} + 1\right)\mathfrak{u}^{\frac{7}{8}}$.

For easier readability set $\varepsilon_i = C_i\left(d^{i-\frac{3}{2}} + 1\right)\mathfrak{u}^{\frac{7}{8}}$. Then

$$\begin{aligned}
\left||S_i| - d^{i-1}\alpha_x\right| \geq \varepsilon_i &\Rightarrow \left||S_i| - d|S_{i-1}|\right| \geq \varepsilon_i - d\delta_{i-1} \\
&\Rightarrow \left||S_i| - \mathbb{E}\left[|S_i| \,\middle|\, B_{i-1}\right]\right| \geq \varepsilon_i - d\delta_{i-1} - d \\
&\Rightarrow \left||S_i| - \mathbb{E}\left[|S_i| \,\middle|\, B_{i-1}\right]\right| \geq (C_i - 2C_{i-1} - 1)\left(d^{i-\frac{3}{2}} + 1\right)\mathfrak{u}^{\frac{7}{8}}
\end{aligned}$$

When $\mathcal{F}_{i-1}$ holds and $\alpha_x \leq 2\mathfrak{u}$,

$$\begin{aligned}
\sqrt{d|S_{i-1}|}u^{\frac{3}{8}} + \mathfrak{u}^{\frac{7}{8}} &\leq \sqrt{d\left(d^{i-2}\alpha_x + C_{i-1}\left(d^{i-\frac{5}{2}} + 1\right)\mathfrak{u}^{\frac{7}{8}}\right)}u^{\frac{3}{8}} + \mathfrak{u}^{\frac{7}{8}} \\
&\leq \left(\sqrt{2(C_{i-1} + 1)} + 1\right)\left(d^{i-\frac{3}{2}} + 1\right)u^{\frac{7}{8}}.
\end{aligned}$$

If we set C_i such that $C_i - 2C_{i-1} - 1 \geq \left(\sqrt{2(C_{i-1} + 1)} + 1\right)$, then whenever cF_{i-1} holds and $\alpha_x \leq 2\mathfrak{u}$,

$$\mathbb{P}\left(\left||S_i| - d^{i-1}\alpha_x\right| \geq C_i(d^{i-\frac{3}{2}} + 1)u^{\frac{7}{8}}\middle|B_{i-1}\right) \leq \mathbb{P}\left(\left||S_i| - \mathbb{E}[|S_i||B_{i-1}]\right| \geq \sqrt{d|S_{i-1}|}u^{\frac{3}{8}} + \mathfrak{u}^{\frac{7}{8}}\middle|B_{i-1}\right).$$

Finally we put all of this together in a union bound

$$\begin{aligned}
&\mathbb{P}\left(\exists x \in \mathcal{V} : \cup_{i=1}^{r+3}\mathcal{G}_i^c(x)\right) \\
&\leq N\mathbb{E}\left[\mathbb{P}\left(\cup_{i=1}^{r+3}\mathcal{G}_i^c(x)\middle|S_1(x)\right)\mathbf{1}(x \in \mathcal{V})\right] \\
&\leq N\mathbb{E}\left[\sum_{i=1}^{r+3}\mathbb{P}\left(\mathcal{G}_i^c(x) \cap \mathcal{F}_{i-1}(x)\middle|S_1(x)\right)\mathbf{1}\left(\mathfrak{u} - \mathfrak{u}^{\frac{2}{3}} \leq \alpha_x \leq 2\mathfrak{u}\right)\right] + N\mathbb{P}\left(\alpha_x > 2\mathfrak{u}\right) \\
&\leq N\mathbb{E}\left[\sum_{i=1}^{r+3}\mathbb{E}\left[\mathbb{P}\left(\left||S_i| - d^{i-1}\alpha_x\right| \geq \varepsilon_i\middle|B_{i-1}\right)\mathbf{1}\left(\mathcal{F}_{i-1}\right)\middle|S_1\right]\mathbf{1}\left(\mathfrak{u} - \mathfrak{u}^{\frac{2}{3}} \leq \alpha_x \leq 2\mathfrak{u}\right)\right] + N\mathbb{P}\left(\alpha_x > 2\mathfrak{u}\right) \\
&\leq N\mathbb{E}\left[\sum_{i=2}^{r+3}2\exp\left\{-C\mathfrak{u}^{\frac{3}{4}}\right\}\mathbf{1}\left(\mathfrak{u} - \mathfrak{u}^{\frac{2}{3}} \leq \alpha_x \leq 2\mathfrak{u}\right)\right] + N\mathbb{P}\left(\alpha_x > 2\mathfrak{u}\right) \\
&\leq N(r+3)e^{-C\mathfrak{u}^{\frac{3}{4}}}\mathbb{P}\left(\mathfrak{u} - \mathfrak{u}^{\frac{2}{3}} \leq \alpha_x\right) + N\mathbb{P}\left(\alpha_x > 2\mathfrak{u}\right)
\end{aligned}$$

$$\leq N(r+3)e^{-C\mathfrak{u}^{\frac{3}{4}}}\frac{1}{N}\frac{3}{2}e^{d+\mathfrak{u}^{\frac{2}{3}}\log\mathfrak{u}}\sqrt{\frac{\mathfrak{u}}{d}}+N\frac{1}{N^{\frac{3}{2}}}$$

where we used that $\mathfrak{u} = \Theta\left(\frac{\log N}{\log\log N}\right)$ by (3.2) and then applied Lemma 4.6 for the second term and the bound on $|\mathcal{V}|$ from Lemma 3.4 for the first term.

□

Proof of Lemma 3.5, 4. We will prove this for $\mathcal{V}$, the proof with $\mathcal{W}$ is identical. First note that by Lemma 4.7, for $X \sim \text{Binom}(N, d/N)$, $\mathbb{P}\left(X \geq \mathfrak{u}^{\frac{3}{4}}\right) \leq e^{-\Omega(\mathfrak{u}^{\frac{3}{4}})}$, since $d \leq (\log N)^{\frac{1}{5}}$ by 5.1. The basic idea now is that by Lemma 3.5 3, there are $O\left((r+3)d^{r+2}\mathfrak{u}\right)$ vertices in $B_{r+3}(x)$ and by Lemma 3.4 there are $e^{O\left(\mathfrak{u}^{\frac{2}{3}}\log\mathfrak{u}\right)}$ vertices in $\mathcal{V}$, so union bounding over all those vertices implies the result. Let us now make this precise.

We show this level by level. For all $y \in S_i(x)$, conditioned on B_i, the N_y are independent and distributed as $\text{Binom}(N - |B_i|, d/N)$, which is stochastically dominated by $\text{Binom}(N, d/N)$. Thus the probability that any $N_y \geq u^{\frac{3}{4}}$ is bounded by $e^{-\Omega\left(u^{\frac{3}{4}}\right)}$.

Putting everything together and using the notation $\mathcal{F}_i$ as defined in the previous proof, we first get

$$\mathbb{P}\left(\exists x \in \mathcal{V} : \exists y \in B_{r+3}(x) : N_y > \mathfrak{u}^{\frac{3}{4}}\right) \leq \mathbb{P}\left(\exists x \in \mathcal{V} : \exists y \in B_{r+3}(x) : N_y > \mathfrak{u}^{\frac{3}{4}}, \alpha_x \leq 2\mathfrak{u}\right) + \mathbb{P}\Big(\exists x : \alpha_x > 2\mathfrak{u}\Big)$$

and we know by 3.5, 3. that the latter event happens with low probability. The first term on the other hand we can bound by

$$\sum_{x\in[N]}\mathbb{P}\left(x \in \mathcal{V}, \cup_{i=1}^{r+3}\left\{\exists y \in S_i(x) : N_y > \mathfrak{u}^{\frac{3}{4}}\right\}, \alpha_x \leq 2\mathfrak{u}\right)$$
$$\leq \sum_{x\in[N]}\sum_{i=1}^{r+3}\mathbb{P}\left(\exists y \in S_i(x) : N_y > \mathfrak{u}^{\frac{3}{4}}, \mathcal{F}_i(x), x \in \mathcal{V}, \alpha_x \leq 2\mathfrak{u}\right) + \sum_{x\in[N]}\sum_{i=1}^{r+3}\mathbb{P}\left(\mathcal{F}_i^c(x), x \in \mathcal{V}, \alpha_x \leq 2\mathfrak{u}\right)$$

Now the latter term is small by the previous proof (note that we used a union bound there as well.) For the former term we proceed as follows:

$$\leq \sum_{x\in[N]}\sum_{i=1}^{r+3}\mathbb{E}\left[\mathbb{E}\left[\mathbf{1}\left(\exists y \in S_i(x) : N_y > \mathfrak{u}^{\frac{3}{4}}\right)\Big|B_i\right]\mathbf{1}_{\mathcal{F}_i}\mathbf{1}_{x\in\mathcal{V}}\mathbf{1}_{\alpha_x\leq 2\mathfrak{u}}\right]$$
$$\leq \sum_{x\in[N]}\sum_{i=1}^{r+3}e^{-\Omega(u^{\frac{3}{4}})}\mathbb{E}\left[|S_i|\mathbf{1}_{\mathcal{F}_i(x)}\mathbf{1}_{x\in\mathcal{V}}\mathbf{1}_{\alpha_x\leq 2\mathfrak{u}}\right]$$

$$\leq \sum_{x\in[N]} e^{-\Omega(u^{\frac{3}{4}})} \sum_{i=1}^{r+3} \mathbb{E}\left[\left(d^{i-1}\alpha + O\left(d^{i-\frac{3}{2}}\mathfrak{u}^{\frac{7}{8}} + \mathfrak{u}^{\frac{7}{8}}\right)\right)\mathbf{1}_{\alpha_x\leq 2\mathfrak{u}}\mathbf{1}_{x\in\mathcal{V}}\right]$$
$$\leq e^{-\Omega(u^{\frac{3}{4}})} O\left((r+3)\left(1+d^{r+2}\right)\mathfrak{u}\right)\frac{3}{2}e^{d+\mathfrak{u}^{\frac{2}{3}}\log\mathfrak{u}}\sqrt{\frac{\mathfrak{u}}{d}},$$

which is small as $N \to \infty$.

□

Proof of Lemma 3.5, 5. We prove this for $\mathcal{W}$, the proof for $\mathcal{V}$ is equivalent. Here we use Lemma 4.10 with $t = 2\mathfrak{u}^{2/3}$. The probability bound we obtain is

$$\exp\left(-\Omega\left(\mathfrak{u}^{-1/3}/(d^3+1)\right)\right).$$

Therefore, for our range of d, it is possible to union bound over all vertices in $|\mathcal{W}|$, as this gives us

$$\exp\left(\mathfrak{u}^{1/4}\right)\exp\left(-\Omega\left(\mathfrak{u}^{-1/3}/\left(d^3+1\right)\right)\right) = \exp\left(-\Omega\left(\mathfrak{u}^{-1/3}/\left(d^3+1\right)\right)\right)$$

by the bound in Lemma 3.4.

□

Proof of Lemma 3.6. We show that with high probability, for all vertices in $\mathcal{U}$,

$$\left|S_i\right| \leq 4^{i-1}\left(d + \log\log N - \log d\right)^{i-1}\mathfrak{u} \tag{A.3}$$

which implies the statement by our bounds on d from Definition 5.1.

The strategy is similar as in the proof of Lemma 3.5, 3: First note that by Lemma 4.7, under the event $\mathcal{F}_{i-1}$,

$$\mathbb{E}\left[|S_i|\,\middle|\,B_{i-1}\right] = d|S_{i-1}| - |B_{i-1}||S_{i-1}|\frac{d}{N} \leq d|S_{i-1}| - d$$

using this bound and the fact that $d \leq d + \log\log N - \log d$, we get that

$$\mathbb{P}\left(\left||S_i| - \mathbb{E}\left[|S_i|\,\middle|\,B_{i-1}\right]\right| \geq \left((d+\log\log N-\log d)\sqrt{2|S_{i-1}|u} + (d+\log\log N-\log d)\mathfrak{u}\right)\middle|B_{i-1}\right)$$
$$\leq 2\exp\left\{-\frac{\left((d+\log\log N-\log d)\sqrt{2|S_{i-1}|u} + (d+\log\log N-\log d)\mathfrak{u}\right)^2}{2\mathbb{E}\left[|S_i||B_{i-1}\right] + \frac{2}{3}\left((d+\log\log N-\log d)\sqrt{|S_{i-1}|u} + (d+\log\log N-\log d)\mathfrak{u}\right)}\right\}$$
$$\leq 2e^{-(d+\log\log N-\log d)\mathfrak{u}}$$
$$\leq 2N^{-1+o(1)}$$

by considering which term in the denominator is smaller and then using the approximation from (3.2) for $\mathfrak{u}$.

Now note that under the event $\mathcal{F}_{i-1}$,

$$\mathbb{E}\left[|S_i|\,\middle|\,B_{i-1}\right] \leq d4^{i-2}(d+\log\log N - \log d)^{i-2}\mathfrak{u} \leq 4^{i-2}(d+\log\log N - \log d)^{i-1}\mathfrak{u}$$

such that

$$\begin{aligned}
&|S_i| \geq 4^{i-1}(d+\log\log N - \log d)^{i-1}\mathfrak{u}\\
\Rightarrow &|S_i| - \mathbb{E}\left[|S_i|\middle|B_{i-1}\right] \geq 3\cdot 4^{i-2}(d+\log\log N - \log d)^{i-1}\mathfrak{u}\\
\Rightarrow &|S_i| - \mathbb{E}\left[|S_i|\middle|B_{i-1}\right] \geq (d+\log\log N - d)\sqrt{2|S_{i-1}|\mathfrak{u}} + (d+\log\log N - \log d)\mathfrak{u}.
\end{aligned}$$

This implies that

$$\mathbb{P}\left(|S_i| \geq 4^{i-2}(d+\log\log N - \log d)^{i-1}\mathfrak{u}|B_{i-1}\right) \leq 2N^{-1+o(1)}.$$

We now proceed as in the end of the proof of Lemma 3.5, 3, by using the bound on $\mathcal{U}$ from Lemma 3.4.

For the second statement of the Lemma, note that it is sufficient to bound $\sum_{y\sim x} N_y^2$. We once more use Lemma 4.10 as in the proof of Lemma 3.5, 5, setting $t = c_{4.10}^{-2}\log^2 N$ and using the bound on $|\mathcal{U}|$ from Corollary 3.1. □

Proof of Lemma 3.7. We use the arguments from the proofs of Lemma 5.5 and Lemma 7.3 in [ADK21b], but choose to include them here for sake of completeness.

For any $x \in [N]$ let $\mathcal{E}_x$ denote the event that there are at least C_1 excess edges in $B_s(x)$. Then by a union bound

$$\mathbb{P}\Big(\exists x \in \mathcal{U} : \mathcal{E}_x\Big) \leq \sum_{x\in[N]} \mathbb{P}\Big(\mathcal{E}_x, x\in\mathcal{U}, \alpha_x < 2\mathfrak{u}\Big) + \mathbb{P}\Big(\alpha_x \geq 2\mathfrak{u}\Big),$$

where the second summand can be bounded by $N^{-\frac{3}{2}}$ according to Lemma 4.6.

In order to bound the first term, we want to condition on $S_1(x)$, and then apply Lemma A.1 for this we write

$$\begin{aligned}
\mathbb{P}\Big(\mathcal{E}_x, x\in\mathcal{U}, \alpha_x < 2\mathfrak{u}\Big) &= \mathbb{E}\Big[\mathbb{P}(\mathcal{E}_x|S_1(x))\mathbf{1}(\{x\in\mathcal{U}, \alpha_x < 2\mathfrak{u}\})\Big]\\
&\leq \mathbb{E}\left[Cd^{3C_1}\left(\frac{|S_1|}{N}\right)^{C_1}\mathbf{1}\left(\eta\mathfrak{u} \leq \alpha_x \leq 2\mathfrak{u}\right)\right]\\
&\leq C\left(2d^3\right)^{C_1}\left(\frac{\log N}{N}\right)^{C_1}\frac{1}{N^{\frac{\eta}{2}}}
\end{aligned}$$

Thus by taking $C_1 \geq 2$ and then doing a union bound over all x we get the desired result.

For the second statement of the Lemma, we proceed as in the proof of Lemma 7.3 in [ADK21b] and write $\mathcal{I}_x$, the event that there are at least C_2 disjoint paths in $B_s(x)$ ending at vertices in $\mathcal{U}_\eta$, as a union over the specific paths:

$$\mathcal{I}_x = \bigcup_{\mathbf{y},\mathbf{z}} \Gamma^{(C_2)}_{\mathbf{y},\mathbf{z}},$$

where the union is taken over all vectors $\mathbf{y} = (y_1, \dots, y_{C_2})$ with distinct entries in $[N]\backslash\{x\}$ and the C_2-tuples $\mathbf{z}$ of disjoint vectors $(z^{(1)}, \dots, z^{(C_2)})$ of length $r_j \in \{0, \dots, s\}$ for $j \in [C_2]$, and

$$\Gamma_{\mathbf{y},\mathbf{z}} = \left\{ y_j \in \mathcal{U}_\eta, \left\{x, z_1^{(j)}\right\}, \left\{z_i^{(j)}, z_{i+1}^{(j)}\right\}, \left\{z_{r_j}^{(j)}, y^j\right\} \in E(G) \forall i \in [r_j - 1], j \in [k] \right\}.$$

For some fixed $\mathbf{y}$ and $\mathbf{z}$, and thus fixed set of $(r_1, \dots, r_{C_2})$, since all paths are disjoint, when we denote by $\mathcal{N}_x$ the neighborhood of a vertex x, and use the independence of the edges, we get that

$$\begin{aligned}\mathbb{P}\left(\Gamma_{\mathbf{y},\mathbf{z}}\right) \leq &\mathbb{P}\left(\left|\mathcal{N}_x \cap \left([N]\backslash\mathbf{y}\right)\right| \geq \eta \mathrm{u} - C_2\right) \\ &\prod_{j=1}^{C_2} \mathbb{P}\left(\left|\mathcal{N}_y \cap \left([N]\backslash\{x\} \cup \mathbf{y}\right)\right| \geq \eta \mathrm{u} - C_2 - 1\right) \\ &\left(\frac{d}{N}\right)^{\sum_{j=1}^{C_2} r_j + 1}.\end{aligned}$$

We now apply Lemma 4.2 and Corollary 4.1 to bound the remaining probabilities: since C_2 is constant all these probabilities will be bounded by

$$e^{-\eta u \log u + c\eta u} = e^{-(1+o(1))\eta \log N}.$$

This implies that the above probability is bounded by

$$\frac{d^{\sum_{j=1}^{C_2} r_j + 1}}{N^{\left(\sum_{j=1}^{C_2} r_j + 1\right) + \eta C_2 (1+o(1))}}.$$

To complete the union bound we need to count the number of terms, i.e. possible paths, for each sequence of r_js. To do this we note that given that x is fixed, there are $\binom{N-1}{C_2}$ ways of picking $\mathbf{y}$ and for the $z_i^{(j)}$ on each path there are $\binom{N-k-\sum_{i=1}^{j-1} r_i}{r_j}$ ways of picking them. Thus

$$\begin{aligned}&\mathbb{P}\left(\mathcal{I}_x\right) \\ &\leq \binom{N-1}{C_2} \sum_{r_1=0}^{s} \cdots \sum_{r_{C_2}=0}^{s} \binom{N - C_2 - 1}{r_1} \cdots \binom{N - C_2 - \sum_{i=1}^{C_2 - 1 r_i}}{r_{C_2}} \frac{d^{\sum_{j=1}^{C_2} r_j + 1}}{N^{\left(\sum_{j=1}^{C_2} r_j + 1\right) + \eta C_2(1+o(1))}}\end{aligned}$$

$$\leq C \frac{d^{C_2(s+1)}}{N^{\eta C_2(1+o(1))}}$$

where C is a constant that depends on the constants s and C_2.

By taking $C_2 > \frac{2}{\eta}$, and then taking a union bound over all x, this implies that all $B_s(x)$ only contain a constant number of disjoint paths ending at other vertices from $\mathcal{U}$ with high probability.

□

Proof of Lemma 3.8. To construct the pruned graph $\hat{G}$ we delete edges in the same manner as in Lemma 7.2 in [ADK21b]: For every vertex $x \in \mathcal{U}$, and its neighbor y, consider the set of vertices T_y that are connected to y by a path of length at most 3, without traversing the edge (x, y). If x is in this set, or the graph induced by T_y on G is not a tree, then we prune the edge (x, y). Denote the set of edges that are pruned in this way P_x. According to Lemma 3.7, with high probability, each vertex $x \in \mathcal{U}$ has less than C_1 "excess" edges that create cycles in $B_3(x)$. Thus by the above procedure we prune at most $C_1 - 1$ edges that are adjacent to x.

In the second step, we work with the graph on $[N]$ with edges $E(G)\backslash P_x$, in which $B_3(x)$ is a tree. In that graph we consider for each neighbor y of x, the vertices V_y in $B_3(x)$ that are connected to y by a path that does not use the edge (x, y). If any of the vertices in V_y is in $\mathcal{U}$, we prune the edge (x, y) and add it to P_x. By Lemma 3.7 we prune at most $C_2 - 1$ edges adjacent to x by doing this procedure.

We then apply these steps, by choosing an arbitrary order of vertices in $\mathcal{U}$, then pruning edges surrounding these vertices sequentially. Let H be the graph on $[N]$ that only consists of the edges $\cup_{x \in cU} P_x$ that we pruned. We then define our pruned graph $\hat{G}$ to be the graph G with edges $E(G) \backslash \cup_{x \in cU} P_x$. By construction $\hat{G}$ satisfies 1. and 2.

Note that only vertices $x \in \mathcal{U}$ and vertices $y \in \cup_{x \in \mathcal{U}} S_1(x)$ are not isolated in H. It is clear that at each step of this procedure we prune at most $C_2 + C_1 - 2$ edges adjacent to some x. Moreover, note that that any subsequent step cannot affect the degree of x in H: otherwise, if we have already pruned for $x \in \mathcal{U}$, if in a subsequent pruning for $x' \in \mathcal{U}$ we were to delete an edge adjacent to x, this would mean that (x', x) is an edge in G, in which case we would already have pruned it when doing the pruning for x.

Now let $y \in \cup_{x \in \mathcal{U}} S_1(x) \backslash \mathcal{U}$, i.e. let y be a vertex that is not in $\mathcal{U}$ and is a neighbor of some vertex $x \in \mathcal{U}$. By Lemma 3.7, y can be adjacent to at most $C_2 - 1$ additional vertices from $\mathcal{U}$, since otherwise $B_2(x)$ contains more than $C_2 - 1$ vertices from $\mathcal{U}$. Thus we prune at most $C_2 - 1$ edges adjacent to y. Hence the maximal degree of the graph H is $C_1 + C_2 - 2$, implying 3.

Recall the assumption that the maximum degree is at most $\mathfrak{u}$. Thus for each edge (x, y) that we prune, β_x is reduced by at most $\mathfrak{u}$. Additionally for each vertex $y \in S_1(x)$, we delete at most $C_2 - 1$ edges by doing the pruning procedure for other $x' \in \mathcal{U}$. This implies that $0 \leq \beta_x - \hat{\beta}_x \leq \alpha_x(C_2 - 1) + (C_1 + C_2 - 2)\mathfrak{u} = O(\mathfrak{u})$, which implies 4.

Lemma 3.6 gives a bound on the growth of the spheres in the original graph and since α_x and $\hat{\alpha}_x$ are of the same order, 5. follows immediately.

For the last statement we rewrite

$$\sum_{y\in\hat{S}_1(x)}\left(\hat{N}_y-\frac{\hat{\beta}}{\hat{\alpha}}\right)^2\leq 3\left[\sum_{y\in\hat{S}_1(x)}\left(\hat{N}_y-N_y\right)^2+\sum_{y\in\hat{S}_1(x)}\left(N_y-d\right)^2+\sum_{y\in\hat{S}_1(x)}\left(d-\frac{\hat{\beta}}{\hat{\alpha}}\right)^2\right].$$

The first term can be bounded by $O(\hat{\alpha})$ since $G-\hat{G}$ has a bounded degree by Lemma 3.8, 3, for the second term we use Lemma 3.6 and the last term can be bounded by $\hat{\alpha}\left(d-\frac{\hat{\beta}}{\hat{\alpha}}\right)^2$ and then be bounded using Lemma 3.8, 3.

□

A.2 Proofs of distributional identities

Binomial estimates

Proof of Lemma 4.3. Using the inequality $\log(1+x)\geq\frac{x}{x+1}$ for $x\geq 0$ [1],

$$I_p\left(\frac{p}{2}\right)=-\frac{p}{2}\log 2+\frac{2-p}{2}\log\left(1+\frac{p}{2(1-p)}\right)\geq-\frac{p}{2}\log 2+\frac{2-p}{2}\frac{\frac{p}{2(1-p)}}{1+\frac{p}{2(1-p)}}=\frac{p}{2}(1-\log 2).$$

The second inequality is a direct consequence of $I_p\left(\frac{p}{2}\right)=I_{1-p}\left(1-\frac{p}{2}\right)$ and the last statement uses the fact that $I_p(x)$ is decreasing on $x\in(0,p)$ [2]. □

Proof of Lemma 4.6. By Lemma 4.5,

$$\mathbb{P}\left(X\geq\tau\right)\leq e^{-mI_p(\frac{\tau}{m})}. \tag{A.4}$$

Using that $\log(1+x)\geq\frac{x}{1+x}$ for $x>-1$, we get that

$$\begin{aligned}I_p\left(\frac{\tau}{m}\right)&=\frac{\tau}{m}\log\left(\frac{\tau}{d}\frac{n}{m}\right)+\left(1-\frac{\tau}{m}\right)\log\left(\frac{m-\tau}{m}\frac{n}{n-d}\right)\\&\geq\frac{\tau}{m}\left[\log(\tau)-\log(d)+\frac{n-m}{m}\right]+\left(1-\frac{\tau}{m}\right)\left[-\frac{\tau}{m-\tau}+\frac{d}{n}\right]\end{aligned}$$

After multiplication with m this can be lower bounded by

$$\tau\left[\log(\tau)-\log(d)+\frac{n-m}{m}-1\right],$$

where all but the first term are $O\left(\tau\max\{\log d,1\}\right)$, which implies the bound. □

[1] $\frac{d}{dx}(\log(1+x)-\frac{x}{x+1})=\frac{x}{(x+1)^2}\geq 0$ for $x\geq 0$.
[2] $I_p'(x)=\log(\frac{x}{p}\cdot\frac{1-p}{1-x})<0$ for $x\in(0,p)$.

Poisson Approximation

Proof of Lemma 4.1. We simplify using Stirling's approximation and the fact that $\frac{e^c}{(1+\frac{c}{n})^n} = 1 + O(\frac{c^2}{n})$ for $|c| < n$:

$$\begin{aligned}
\mathbb{P}(X = k) &= \binom{n}{k} p^k (1-p)^{n-k} \\
&= \left(1 + O\left(\frac{1}{n}\right)\right) \frac{1}{k!} \frac{n^n}{(n-k)^{n-k} e^k} \sqrt{\frac{1}{1-\frac{k}{n}}} p^k (1-p)^{n-k} \\
&= \left(1 + O\left(\frac{k^2 + (np)^2 + 1}{n}\right)\right) \frac{e^{-np}(np)^k}{k!} \\
&= \left(1 + O\left(\frac{k^2 + (np)^2 + 1}{n}\right)\right) \mathbb{P}(Y = k)
\end{aligned}$$

□

Proof of Lemma 4.1. We have

$$\mathbb{P}\Big(X \geq k\Big) = \left(1 + O\left(\frac{k^2 + (np)^2 + 1}{n}\right)\right) \mathbb{P}\Big(Y \geq k\Big) + O\left(\mathbb{P}\left(X \geq \sqrt{n}\right) + \mathbb{P}\left(Y \geq \sqrt{n}\right)\right)$$

and the latter term satisfies

$$O\left(\mathbb{P}\left(X \geq \sqrt{n}\right) + \mathbb{P}\left(Y \geq \sqrt{n}\right)\right) = O\left(\binom{n}{\sqrt{n}} p^{\sqrt{n}} + \frac{(np)^{\sqrt{n}}}{\sqrt{n}!}\right) = O\left(\left(ep\sqrt{n}\right)^{\sqrt{n}}\right).$$

□

Proof of Lemma 4.2. **Upper bound:** We start with

$$\begin{aligned}
\mathbb{P}\Big(X \geq \lambda(1+\delta)\Big) &= \sum_{k=\lambda(1+\delta)}^{\infty} \mathbb{P}(X = k) \\
&= e^{-\lambda} \sum_{k=\lambda(1+\delta)}^{\infty} \frac{\lambda^k}{k!}.
\end{aligned}$$

Note that for k in our range, $\lambda/k \leq 1/(1+\delta)$. Therefore, we can upper bound this with a geometric series.

$$e^{-\lambda} \sum_{k=0}^{\infty} \frac{\lambda^k}{k!} \quad \leq \quad e^{-\lambda} \frac{\lambda^{\lambda(1+\delta)}}{(\lambda(1+\delta))!} \sum_{k=0}^{\infty} \frac{1}{(1+\delta)^k}$$

$$\leq \quad e^{-\lambda}\frac{\lambda^{\lambda(1+\delta)}}{(\lambda(1+\delta))!}(\frac{1+\delta}{\delta}).$$

Using a Stirling approximation, we obtain

$$\begin{aligned}
e^{-\lambda}\frac{\lambda^{\lambda(1+\delta)}}{(\lambda(1+\delta))!}\left(\frac{1+\delta}{\delta}\right) &= (1+o_\lambda(1))e^{-\lambda}\frac{e^{\lambda(1+\delta)}}{\sqrt{2\pi\lambda(1+\delta)}}\frac{\lambda^{\lambda(1+\delta)}}{(\lambda(1+\delta))^{\lambda(1+\delta)}}\left(\frac{1+\delta}{\delta}\right) \\
&= (1+o_\lambda(1))\frac{\exp(-\lambda(1+\delta)\log(1+\delta)+\lambda(1+\delta)-\lambda)}{\sqrt{2\pi\lambda(1+\delta)}\frac{\delta}{1+\delta}} \\
&\leq \frac{\exp(-\lambda(1+\delta)\log(1+\delta)+\lambda\delta)}{\sqrt{\lambda}\min\{\sqrt{\delta},\delta\}}
\end{aligned}$$

for sufficiently large λ.

Lower bound: We once again have that

$$\begin{aligned}
\mathbb{P}\Big(X \geq \lambda(1+\delta)\Big) &= \sum_{k=\lambda(1+\delta)}^{\infty}\mathbb{P}(X=k) \\
&= e^{-\lambda}\sum_{k=\lambda(1+\delta)}^{\infty}\frac{\lambda^k}{k!} \\
&= \frac{(1+o_\lambda(1))e^{-\lambda}}{\sqrt{2\pi}}\sum_{k\geq\lambda(1+\delta)}\frac{\lambda^k e^k}{k^{k+1/2}}
\end{aligned}$$

by a Stirling approximation. We write

$$c := \lambda(1+\delta), \qquad f(x) := -(x+1/2)\log x + x\log\lambda + x.$$

Therefore, we have

$$f'(x) = -\log x + \log\lambda + \frac{1}{2x}, \qquad f''(x) = -\frac{1}{x}+\frac{1}{2x^2}$$

We then approximate our probability by an integral, which serves as a lower bound as our function is decreasing. We then perform a Laplace method type bound.

$$\begin{aligned}
\sum_{k=\lambda(1+\delta)}^{c+c^{1/3}}\frac{\lambda^k e^k}{k^{k+1/2}} &\geq \int_c^{c+c^{1/3}}\frac{\lambda^x e^x}{(x)^{x+1/2}}dx \\
&= \int_c^{c+c^{1/3}}\exp\left(f(x)\right)dx \\
&= e^{f(c)}\int_c^{c+c^{1/3}}\exp\left(f(x)-f(c)\right)dx
\end{aligned}$$

$$= e^{f(c)} \int_c^{c+c^{1/3}} \exp\left(f'(c)(x-c) + O_\lambda\left(f''(c)(x-c)^2\right)\right) dx$$

where the last statement follows from Taylor expanding and the formula of $f''(x)$. By the choice of our window, $f''(c)(x-c)^2 = O(\lambda^{-1/3})$. Therefore, this is

$$\begin{aligned}
&e^{f(c)} \int_c^{c+c^{1/3}} \exp\left(f'(c)(x-c) + O_\lambda\left(f''(c)(x-c)^2\right)\right) dx \\
=&(1+o_\lambda(1))e^{f(c)-cf'(c)} \int_c^{c+c^{1/3}} \exp(f'(c)x)dx \\
=&-(1+o_\lambda(1))\frac{e^{f(c)-cf'(c)+cf'(c)}}{f'(c)}(1-e^{c^{1/3}f'(c)}) \\
=&-(1+o_\lambda(1))\frac{e^{f(c)}}{f'(c)}.
\end{aligned}$$

By our choice of c, and under the assumption that $\delta > \frac{1}{\sqrt{\lambda}}$, we have

$$-\frac{1}{f'(c)} = (1+o_\lambda(1))\frac{1}{\log(1+\delta)}.$$

Putting this together with the fact that $\mathbb{P}\big(X \geq \lambda(1+\delta)\big) \geq \mathbb{P}\big(X = \lambda(1+\delta)\big)$, we have that

$$\begin{aligned}
\mathbb{P}\big(X \geq \lambda(1+\delta)\big) &\geq (1+o_\lambda(1))e^{-\lambda}\frac{\lambda^{\lambda(1+\delta)}e^{\lambda(1+\delta)}}{(\lambda(1+\delta))^{\lambda(1+\delta)+1/2}} \cdot \max\left\{\frac{1}{\log(1+\delta)}, 1\right\} \\
&= (1+o_\lambda(1))c_{4.2}\frac{\exp(-\lambda h(\delta))}{\sqrt{\lambda}\min\{\sqrt{\delta}, \delta\}}
\end{aligned}$$

for some constant $c_{4.2}$. □

Proof of Corollary 4.2. For the upper bound,

$$\mathbb{P}\big(X \geq \lambda(1+\delta)\big) \leq \frac{e^{-\lambda h(\delta)}}{\sqrt{\lambda}\min\{\delta, \delta^2\}} \leq \big(1+o_\lambda(1)\big)\frac{e^{-\frac{\lambda\delta^2}{2}}}{\delta\sqrt{\lambda}}$$

by the Taylor expansion of $h(\delta)$.

For the lower bound, define $\delta' = \delta + \frac{1}{\lambda}$. Then $\lambda(1+\delta)+1 = \lambda(1+\delta')$. We have

$$\mathbb{P}\big(X \geq \lambda(1+\delta)\big) = \mathbb{P}\big(X \geq \lceil\lambda(1+\delta)\rceil\big) \geq c_{4.2}\frac{e^{-\lambda h(\delta')}}{\sqrt{\lambda}\min\{\delta', \delta'^2\}} \geq \big(1-o_\lambda(1)\big)c_{4.2}\frac{e^{-\frac{\lambda\delta^2}{2}}}{\delta\sqrt{\lambda}}.$$

□

Tails of sums

Proof of 4.10. We will only work with the upper tail as the lower tail is similar. The bound we get from Lemma 4.7 is

$$\Pr\left(X > t + \mathbb{E}[X]\right) \leq \exp\left(-\frac{t^2}{2d + \frac{2}{3}t}\right).$$

Therefore,

$$\Pr\left(X^2 \geq t^2 + 2t\mathbb{E}[X] + \mathbb{E}[X]^2\right) \leq \exp\left(-\frac{t^2}{2d + \frac{2}{3}t}\right).$$

which we rewrite as

$$\Pr\left(X^2 - (d^2 + d) \geq t^2 + (2t - 1)\mathbb{E}[X]\right) \leq \exp\left(-\frac{t^2}{2d + \frac{2}{3}t}\right).$$

By using the substitution $t = \sqrt{y + d^2 + d} - d$, we can rewrite this as

$$\Pr\left(X^2 - (d^2 + d) \geq y\right) \leq \exp\left(-\frac{y + 2d^2 + d - 2d\sqrt{y + d^2 + d}}{\frac{2}{3}\sqrt{y + d^2 + d} + \frac{4}{3}d}\right).$$

As mentioned, this will follow from [BMdlP23], Theorem 1. First, we wish to show that, using the notation in the paper, $v(L, \beta) \to \mathrm{Var}(X^2)$ for large L. We first show that the strategy of proof of Lemma 4 extends to our scenario. The probability that $X^2 - (d^2 + d) \geq y$ is at most $e^{-\frac{\sqrt{y}}{12}}$ for $y \geq 4(d^2 + d)$. Therefore, we choose our Y to be

$$Y = \left(X^2 - \mathbb{E}\left[X^2\right]\right)^2 \exp(X/12)\mathbf{1}_{X^2 > E(X^2)}.$$

Y is integrable as the moment generating function $\mathbb{E}(e^{cX})$ is finite for all c. Therefore, for sufficiently large mt, $v(L, \beta) \approx \mathrm{Var}(X^2)$.

This makes the formulation of Theorem 1 in [BMdlP23] much simpler, and $t_{\max}$ is such that

$$t_{\max} = \mathrm{Var}(X^2)\frac{mt_{\max} + 2d^2 + d - 2d\sqrt{mt_{\max} + d^2 + d}}{\frac{2}{3}mt_{\max}\sqrt{mt_{\max} + d^2 + d} + \frac{4}{3}d}.$$

Therefore $t_{\max} = (1 + o_m(1))(\frac{3}{2}\mathrm{Var}(X^2))^{2/3}m^{-1/3}$. Using this on the four terms in Theorem 1 from [BMdlP23], as well as the bound that $I(t) \geq \sqrt{t}/12$ if $t \geq 4(d^2 + d)$, gives the result. □

To verify Lemma 4.11 we need the following easy lemma regarding the discretization of a tail.

Lemma A.2. Suppose that $k \in \mathbb{N}$ and $t > k$. Then,

$$\Big\{x_1 + \cdots + x_k \geq t, x_1, \cdots, x_k \geq 0\Big\} \subseteq \bigcup_{\substack{(t_1,\ldots,t_k)\in\mathbb{N}_{\geq 0}^k \\ t_1+\cdots+t_k=\lfloor t\rfloor-k+1}} \Big\{x_1 \geq t_1, \ldots, x_k \geq t_k\Big\}. \tag{A.5}$$

Here, $\mathbb{N}_{\geq 0}$ denotes the set of non-negative integers.

Proof. Since $0 \leq x - \lfloor x \rfloor < 1$, we have $x_1 - \lfloor x_1 \rfloor + \cdots + x_k - \lfloor x_k \rfloor < k$. Using this fact, for any non-negative $x_1, \cdots, x_k$,

$$\begin{aligned}\Big\{x_1 + \cdots + x_k \geq t\Big\} &\subseteq \Big\{\lfloor x_1 \rfloor + \cdots + \lfloor x_k \rfloor > \lfloor t \rfloor - k\Big\} \\ &= \Big\{\lfloor x_1 \rfloor + \cdots + \lfloor x_k \rfloor \geq \lfloor t \rfloor - k + 1\Big\} \\ &= \bigcup_{\substack{(t_1,\ldots,t_k)\in\mathbb{N}_{\geq 0}^k \\ t_1+\cdots+t_k=\lfloor t\rfloor-k+1}} \Big\{\lfloor x_1 \rfloor \geq t_1, \ldots, \lfloor x_k \rfloor \geq t_k\Big\}.\end{aligned}$$

Since t_is are integers, $\lfloor x_i \rfloor \geq t_i$ is equivalent to $x_i \geq t_i$. Thus, we are done. □

Proof of Lemma 4.11. For the lower bound, note that

$$\mathbb{P}\left(Y_1^2 + \cdots + Y_k^2 \geq t\right) \geq \mathbb{P}\Big(Y_1^2 \geq \frac{t}{k}, \ldots, Y_k^2 \geq \frac{t}{k}\Big) = \mathbb{P}\left(Y_1^2 \geq \frac{t}{k}\right)^k \geq C_1^k e^{-t^{\frac{\alpha}{2}} k^{1-\frac{\alpha}{2}}}.$$

The upper bound is obtained as an application of the discretization from Lemma A.2. In fact,

$$\mathbb{P}\left(Y_1^2 + \cdots + Y_k^2 \geq t\right) \leq \sum \mathbb{P}\left(Y_1^2 \geq t_1, \ldots, Y_k^2 \geq t_k\right) \leq C_2^k \sum e^{-\sum_{i=1}^k t_i^{\frac{\alpha}{2}}},$$

where the summation is taken over $(t_1, \ldots, t_k) \in \mathbb{N}_{\geq 0}^k$ with $t_1 + \cdots + t_k = \lfloor t \rfloor - k + 1$. Since the function $f(x) = x^{\frac{\alpha}{2}}$ with $\alpha > 2$ is convex, by Jensen's inequality,

$$\frac{1}{k}\sum_{i=1}^k t_i^{\frac{\alpha}{2}} \geq \left(\frac{1}{k}\sum_{i=1}^k t_i\right)^{\frac{\alpha}{2}}.$$

Hence,

$$\mathbb{P}\left(Y_1^2 + \cdots + Y_k^2 \geq t\right) \leq C_2^k \sum e^{-(\lfloor t \rfloor - k + 1)^{\frac{\alpha}{2}} k^{1-\frac{\alpha}{2}}} < C_2^k \sum e^{-(t-k)^{\frac{\alpha}{2}} k^{1-\frac{\alpha}{2}}}. \tag{A.6}$$

We now bound the number of summands, i.e. the number of tuples $(t_1, \ldots, t_k) \in \mathbb{N}_{\geq 0}^k$ such that $t_1 + \cdots + t_k = \lfloor t \rfloor - k + 1$. This is known as the number of weak compositions of $\lfloor t \rfloor - k + 1$

into k terms, and is given by $\binom{\lfloor t\rfloor-k+1+k-1}{k-1} = \binom{\lfloor t\rfloor}{k-1}$. Using the fact that $\binom{n}{k} \leq \left(\frac{en}{k}\right)^k$ and the condition $t > k$, the number of summands is thus bounded by

$$\binom{\lfloor t\rfloor}{k-1} \leq \left(\frac{e\lfloor t\rfloor}{k-1}\right)^{k-1} \leq \left(\frac{2et}{k}\right)^k.$$

The upper bound is established by applying this to (A.6).

To obtain (4.14), we simply plug the given values into (4.13) and then use that $(\frac{2et}{k})^k = N^{o(1)}$ and

$$(t-k)^{\frac{\alpha}{2}}k^{1-\frac{\alpha}{2}} = (1+o(1))d^\alpha \frac{2}{\alpha-2}\Big(1-\frac{2}{\alpha}\Big)^{\frac{\alpha}{2}} b^{1-\frac{\alpha}{2}} \log N.$$

Noting in addition that

$$\begin{aligned}\mathbb{P}\left(\tilde{Y}_1^2+\cdots+\tilde{Y}_k^2 \geq t\right) &= \frac{\mathbb{P}\left(Y_1^2+\cdots+Y_k^2 \geq t \text{ and } Y_i \geq (\varepsilon \log\log N)^{\frac{1}{\alpha}} \text{ for all } i=1,\cdots,k\right)}{\mathbb{P}\left(|Y_1| \geq (\varepsilon\log\log N)^{\frac{1}{\alpha}}\right)^k} \\ &\leq \frac{\mathbb{P}\left(Y_1^2+\cdots+Y_k^2 \geq t\right)}{\mathbb{P}\left(|Y_1| \geq (\varepsilon\log\log N)^{\frac{1}{\alpha}}\right)^k},\end{aligned}$$

and recalling the tail probabilities of Y_i in (4.12), we establish (4.15).

The final statement (4.16) follows by similar calculations. Note that formally sending $b \to 0$ in (4.14) and (4.15) gives (4.16) (recall that $\alpha > 2$). □

Remark A.1. The proof of Lemma 4.11 shows that the lower bound for the probability $\mathbb{P}(Y_1^2+\cdots+Y_k^2 \geq t)$, obtained by forcing $Y_i^2 \geq \frac{t}{k}$ for all $i = 1, \cdots, k$, is of the same order (at the exponential scale) as the upper bound.

Proof of Lemma 4.12. By the exponential Chebyshev bound, for any $s > 0$,

$$\mathbb{P}\left(|\tilde{Y}_1|^\alpha+\cdots+|\tilde{Y}_m|^\alpha \geq L\right) \leq e^{-sL}\mathbb{E}\left[e^{s|\tilde{Y}_1|^\alpha}\right]^m. \tag{A.7}$$

Recalling that $\tilde{Y}_1$ is a random variable Y_1 conditioned to have an absolute value greater than $(\varepsilon \log\log N)^{\frac{1}{\alpha}}$, using a tail bound (4.12), there exists a constant $C \geq 1$ such that for $x > (\varepsilon\log\log N)^{\frac{1}{\alpha}}$,

$$\mathbb{P}\left(|\tilde{Y}_1| > x\right) \leq Ce^{\varepsilon\log\log N}e^{-x^\alpha},$$

and for $0 \leq x \leq (\varepsilon\log\log N)^{\frac{1}{\alpha}}$, $\mathbb{P}(|\tilde{Y}_1| > x) = 1$. Hence, for any $0 < s < 1$, using the tail bounds (4.12)

$$\mathbb{E}\left[e^{s|\tilde{Y}_1|^\alpha}\right] = 1 + \int_0^\infty e^{sx^\alpha} s\alpha x^{\alpha-1}\mathbb{P}\left(|\tilde{Y}_1| > x\right)dx$$

$$
\begin{aligned}
&\leq 1 + \int_0^{(\varepsilon \log\log N)^{\frac{1}{\alpha}}} e^{sx^\alpha} s\alpha x^{\alpha-1} dx + C \int_{(\varepsilon \log\log N)^{\frac{1}{\alpha}}}^{\infty} e^{sx^\alpha} s\alpha x^{\alpha-1} e^{\varepsilon \log\log N} e^{-x^\alpha} dx \\
&\leq e^{\varepsilon s \log\log N} + C e^{\varepsilon \log\log N} \frac{s}{1-s} e^{-\varepsilon(1-s)\log\log N} \leq \Big(1 + \frac{Cs}{1-s}\Big) e^{\varepsilon s \log\log N}.
\end{aligned}
$$

Note that the tail bound (4.12) implies that the moment generating function is infinity for $s \geq 1$.

Applying this to (A.7), by Chernoff's bound, for any $0 < s < 1$,

$$
\mathbb{P}\left(|\tilde{Y}_1|^\alpha + \cdots + |\tilde{Y}_m|^\alpha \geq L\right) \leq e^{-sL}\Big(1 + \frac{Cs}{1-s}\Big)^m e^{\varepsilon s m \log\log N}.
$$

Now recall that $L > m$ and set $s := 1 - \frac{m}{L} \in (0, 1)$ in order to balance the two terms e^{-sL} and $(1 + \frac{Cs}{1-s})^m$. Then,

$$
\begin{aligned}
\mathbb{P}\left(|\tilde{Y}_1|^\alpha + \cdots + |\tilde{Y}_m|^\alpha \geq L\right) &\leq e^{-L} e^m \left(1 + \frac{C(L-m)}{m}\right)^m e^{\varepsilon m \log\log N} \\
&\leq C^m e^{-L} e^m \left(\frac{L}{m}\right)^m e^{\varepsilon m \log\log N},
\end{aligned}
$$

which concludes the proof of (4.17). Since the function $x \mapsto (\frac{L}{x})^x$ is increasing on $(0, \frac{L}{e})$, by taking $m = b\frac{\log N}{\log\log N} + c$ and $L = a \log N$ in the RHS of (4.17), we establish (4.18).

□