



Universiteit
Leiden
The Netherlands

On the renormalization of random network models

Catanzaro, A.

Citation

Catanzaro, A. (2026, September 10). *On the renormalization of random network models*. Retrieved from <https://hdl.handle.net/1887/4309543>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4309543>

Note: To cite this publication please use the final published version (if applicable).

PART I

MULTI-SCALE NETWORK MODEL IN THE INFINITE RENORMALIZABILITY REGIME: CLUSTERING AND SPECTRUM

Spectra of random graphs with discrete scale invariance

This chapter is based on:

Alessio Catanzaro, Rajat Subhra Hazra, and Diego Garlaschelli. *"Spectra of random graphs with discrete scale invariance"*. arXiv preprint arXiv:2509.12407 (2025).

Abstract

Random graphs defined by an occurrence probability that is invariant under node aggregation have been identified recently in the context of network renormalization. The invariance property requires that edges are drawn with a specific probability that, in the annealed case, depends on a necessarily infinite-mean node fitness. The diverging mean determines many properties that are uncommon in models with independent edges, but at the same time widespread in real-world networks. Here we focus on the leading eigenvalues and eigenvectors of the adjacency matrix of the model, where the n nodes are assigned a Pareto (α)-distributed fitness with $0 < \alpha < 1$. We find that the leading eigenvalues are all of order $\sqrt{n}$, alternate in sign and are located at the intersection between the real axis and a logarithmic spiral in the complex plane, which we characterize analytically in terms of the Gamma function. We also calculate the associated eigenvectors, finding that they display complex-valued scaling exponents and log-periodicity, which are signatures of discrete scale invariance. In contrast with the typical finite-rank behaviour of random graphs with finite-mean variables, we find that a growing number of the leading eigenvalues emerges from the bulk, whose edge extends up to order $\sqrt{n}$ and therefore reaches the same scale as that of the structural eigenvalues.

§2.1 Introduction

Many models have been proposed to reproduce the heterogeneous structure ubiquitously observed in real-world networks. Among these, a wide class of random graph models assigns a weight (or fitness, or hidden variable) x_i to each vertex i ($i = 1, n$ where n is the number of nodes) from a usually fat-tailed distribution $\rho(x)$, and connects pairs of vertices i and j with a probability p_{ij} that is a certain positive function $f(x_i, x_j)$ [11, 12, 64, 65].

Recently, in the context of network renormalization [54], a specific such model that complements the heterogeneity of nodes with an invariance principle under graph coarse-graining has been identified [1, 2, 3, 66]. In this so-called (annealed) Multi-Scale Model (MSM), one may group nodes into equally sized ‘supernodes’, each inheriting the sum of the fitness of its members, and impose that both the connection probability $p_{ij} = f(x_i, x_j)$ and the fitness distribution $\rho(x)$ remain unchanged in their functional form, while their parameters follow exact renormalization rules. In other words, aggregating nodes into supernodes according to any homogeneous partition – and connecting two supernodes whenever there is a connection between any of their constituent nodes – defines a coarse-grained version of the graph, governed by an occurrence probability that has the same functional form as that of the original graph, up to a rescaling of the parameters [1, 3].

This renormalization scheme does not rely on any underlying geometry or distance notion between nodes. This stands in contrast to geometric renormalization approaches [20, 21], where the nodes to be merged need be spatially proximate, and a preliminary definition of metric distance is needed. Rather, when nodes are merged into supernodes of equal cardinality, the invariance requirement in the MSM is closer in spirit to the concept of discrete scale invariance [35], whose consequences for the graph properties have not been investigated so far and are one of our main focuses here.

In particular, we are interested in the spectral properties [25] of the MSM, which turn out to be quite unique due to the peculiar structure of the fitness. Indeed, the requirement that the node fitness is a positive random (‘annealed’) variable that is stable under addition implies that it has an α -stable distribution with infinite mean ($0 < \alpha < 1$) [1, 3], where classical random matrix results [25, 29, 30, 67, 68, 69, 31] break down. When a symmetric random matrix is decomposed into a low-rank informative part plus noise, strong signals can produce eigenvalue outliers outside the bulk spectrum, as first shown in [32]. In networks, such outliers reveal structures like communities or hubs [25, 33]. While infinite-mean regimes have been studied for Lévy matrices [70, 71, 72], random graphs add further randomness from adjacency realizations, not captured by those approaches. Our goal is to shed light on the surprising behaviour of the eigenmodes of the adjacency matrix in this infinite-mean regime, which follows from discrete node aggregation invariance.

§2.2 Preliminaries

The annealed MSM is an inhomogeneous random graph on n nodes, each node i being assigned a positive weight x_i , sampled independently from an α -stable distribution $\rho_\alpha(x)$, hence with power-law tails decaying as $x^{-1-\alpha}$ for $0 < \alpha < 1$ [1]. In this regime of α , every moment of $\rho_\alpha(x)$ diverges. Given x_i and x_j , the probability that an edge exists between i and j is

$$p_{ij} = f_n(x_i, x_j) = 1 - \exp(-\varepsilon_n x_i x_j), \quad i, j = 1, n \quad (2.1)$$

if $i \neq j$ and $p_{ii} = 0$ if $i = j$, where the parameter ε_n may be chosen to scale with n in a desired way. Under an arbitrary homogeneous partition whereby the n nodes are aggregated into supernodes (each with the same number of participating nodes), the form of both $p_{ij} = f_n(x_i, x_j)$ and $\rho_\alpha(x)$ is unchanged, up to a parameter shift in the latter – reflecting the fact that the fitness of a supernode is the sum of the fitnesses of its constituent nodes [1, 3]. Since $\rho(x)$ for positive α -stable variables is not known in closed-form (except for $\alpha = 1/2$), for analytical tractability we follow [3] and consider a pure Pareto(α) distribution with the same tail as an α -stable distribution:

$$\rho_\alpha(x) = \alpha x^{-1-\alpha} \quad 0 < \alpha < 1 \quad (2.2)$$

for $x \geq 1$, and $\rho_\alpha(x) = 0$ otherwise. We also consider $\varepsilon_n = n^{-1/\alpha}$ as a scaling that allows the graph to be arbitrarily large while remaining sparse, allowing precise asymptotic calculations for large n [3]. In this regime, the tail of the expected degree distribution $P(k)$ decays as k^{-2} irrespective of the value of α , the expected link density is of order $\ln n/n$, while the average local clustering coefficient remains finite [1, 3] – which are all widespread properties found in real-world networks.

§2.2.1 Realized versus expected matrix.

We denote the realized adjacency matrix of the graph as $\mathbf{A}$, where A_{ij} is a Bernoulli variable: $A_{ij} = 1$ with probability p_{ij} and 0 otherwise. Note that $\mathbf{A}$ is symmetric and $A_{ii} = 0$. By construction, the expectation of $\mathbf{A}$ (conditioning on the realized weights) is the matrix $\mathbf{P} = \mathbb{E}[\mathbf{A}|\{x_i\}]$ with entries $P_{ij} = p_{ij} = f_n(x_i, x_j)$ and captures the deterministic, structural signal imposed by the weights, while the actual $\mathbf{A}$ can be seen as a noisy realization of that signal.

Indeed, a central idea in our analysis is to decompose the adjacency matrix into the sum of an expected component and a fluctuation component, and then use Random Matrix Theory (RMT) intuition to separate the spectrum into bulk and outlier components. Let us write:

$$\mathbf{A} = \langle \mathbf{A} \rangle + (\mathbf{A} - \langle \mathbf{A} \rangle) = \mathbf{P} + \mathbf{H}, \quad (2.3)$$

where $\mathbf{H} = \mathbf{A} - \mathbf{P}$ is the zero-mean noise matrix. By construction, $\mathbf{P}$ is a dense matrix (encoding for the structure of the network) with a highly skewed pattern due to the heavy-tailed x_i , whereas $\mathbf{H}$ captures the randomness of edge realizations. If the matrix $\mathbf{P}$ were of small, finite rank and $\mathbf{H}$ were a ‘typical’ random matrix with




Figure 2.1: The full spectrum of $\mathbf{A}$, a single realization of the network adjacency matrix from $\mathbf{P}$, with $n = 10^4$ and $\alpha = 0.5$. In the top panel the histogram (blue) of $\mathbf{A}$. We see separation into atomic delta densities and a central, continuous bulk. In the bottom panel we show the separation of the spectrum into two regimes, the outliers being those of $\mathbf{A}$ in blue, and of $\mathbf{P}$ in orange. We also superimpose in green the eigenvalue histogram of $\mathbf{H}$, that replicates the bulk behaviour of $\mathbf{A}$.

a bounded spectral distribution, then standard results show that each eigenvalue of $\mathbf{P}$ above a certain magnitude would detach from the continuous spectrum of $\mathbf{H}$. This phenomenon, rigorously characterized in spiked random matrix theory, underlies applications like spectral clustering [26, 27, 28]. In our case, $\mathbf{P}$ is not finite-rank: as we argue in the following sections though, we find evidence of an ‘effective’ rank of order $\ln n$. In this sense, we are close in spirit to recent results such as Ref. [73], that are hinting at the fact that generalized BBP transitions could be happening even when the rank of the perturbation matrix is growing as $o(n)$. An illustration of the typical spectra of $\mathbf{A}$ and $\mathbf{P}$ is provided in Fig. 2.1. Our main goal is the characterization of these spectra and the associated eigenvectors, with a specific attention to the outlier modes carrying structural information.

§2.3 Leading eigenmodes from the expected matrix.

We are interested in the solution of the eigenmode equation [25]

$$\mathbf{P}v_k = \lambda_k v_k \quad k = 1, n \quad (2.4)$$

with $|\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_n|$. Denoting the i -th entry of each eigenvector v_k as $v_k^{(i)}$, we can rewrite Eq.(2.4) as

$$\sum_{j=1}^n p_{ij} v_k^{(j)} = \lambda_k v_k^{(i)} \quad i = 1, n \quad k = 1, n \quad (2.5)$$

Note that, since $\mathbf{P}$ is real and symmetric, its eigenvalues are real. Also note that any two vertices i, j with the same realized weight $x_i = x_j$ are statistically equivalent, as their (expected) properties in the graph depend only on the weight x . Similarly, if $x_i = x_j$ then the i -th and j -th entries of each eigenvector v_k will be equal: $v_k^{(i)} = v_k^{(j)}$. This allows us to write $v_k^{(i)} = \nu_k(x_i)$ where $\nu_k(x)$ must be a continuous function of the weight x . For large enough n , Eq.(2.5) can be recast as the integral expression

$$\Pi_{\alpha, n}[\nu_k(x)] = \frac{\lambda_k}{n} \nu_k(x). \quad (2.6)$$

where $\nu_k(x)$ is an eigenfunction (with eigenvalue λ_k/n) of the integral operator

$$\Pi_{\alpha, n}[g(x)] \equiv \int_0^{+\infty} dy \, \rho_{\alpha}(y) (1 - e^{-\varepsilon_n xy}) g(y), \quad (2.7)$$

corresponding to the eigenvector v_k (with eigenvalue λ_k) of the original matrix $\mathbf{P}$. A natural normalization for $\nu_k(x)$ will emerge from the calculations that follow. If an eigenfunction $\nu_k(x)$ is found, a convenient way to go back to the actual eigenvector entry $v_k^{(j)}$ is the following. We first order all vertices so that $x_1 \leq x_2 \leq \dots \leq x_n$. Then, since a pure Pareto(α) variable can be represented as a uniform random variable, raised to the power $-1/\alpha$, in the interval $[0, 1]$, we arrive at the relation $x_j = (n/j)^{1/\alpha}$ (we will use this relation as an alternative, deterministic assignment of the fitness of nodes to produce a ‘noise-free’ power law distribution even for finite n). Finally,

$$v_k^{(j)} = \nu_k(x_j) = \nu_k\left(\frac{n^{1/\alpha}}{j^{1/\alpha}}\right) \quad j = 1, n \quad k = 1, n. \quad (2.8)$$

§2.3.1 Largest eigenvalue and principal eigenvector

We start with the principal component $k = 1$. From Perron-Frobenius theorem, λ_1 is be positive and v_1 has all entries of the same sign. We will set the normalization of v_1 in such a way that all its entries are positive, so that $\nu_1(x)$ is a positive-valued function. When evaluating the operator 2.7 on the eigenfunction, the multiplication of $\rho_{\alpha}(x)$ by $\nu_1(x)$ formally generates a modified, effective distribution of the weights.

We then look for an eigenfunction such that $\nu_1(x)\rho_\alpha(x) = \rho_\beta(x)$ is again a Pareto pdf with a new index β . This is obtained for the eigenfunction

$$\nu_1(x) = \frac{\beta}{\alpha} x^{\alpha-\beta}, \quad \beta \in \mathbb{R}^+. \quad (2.9)$$

Setting $g(x) = \nu_1(x)$ into Eq. (2.7), we get

$$\Pi_{\alpha,n}[\nu_1(x)] = 1 - \varphi_\beta(\varepsilon_n x) \quad (2.10)$$

where we have introduced the Laplace Transform (LT) φ_α of the Pareto- α density, defined as

$$\varphi_\alpha(t) \equiv \int_0^{+\infty} dx \rho_\alpha(x) e^{-tx} = \alpha \int_1^{+\infty} dx x^{-1-\alpha} e^{-tx}, \quad (2.11)$$

and we have exploited the normalization property

$$\int_0^{+\infty} dy \rho_\alpha(y) \nu_1(y) = \int_0^{+\infty} dy \rho_\beta(y) = 1. \quad (2.12)$$

The above expression sets our normalization for $\nu_1(x)$, which is the analogous of the ℓ_1 norm $\sum_{i=1}^n |v_1^{(i)}| = 1$ for the original eigenvector v_1 .

We can easily derive the asymptotics properties of the Laplace transform of a Pareto distribution. Indeed

$$\begin{aligned} \Pi_{\alpha,n}[\nu_1(x)] &= 1 - \varphi_\beta(\varepsilon_n x) \\ &= 1 - \int_1^{+\infty} dy \rho_\beta(y) e^{-\varepsilon_n xy} \\ &= 1 - \beta \int_1^{+\infty} dy y^{-1-\beta} e^{-\varepsilon_n xy} \\ &= 1 - \beta \int_{\varepsilon_n x}^{+\infty} dt \left(\frac{t}{\varepsilon_n x} \right)^{-1-\beta} \frac{e^{-t}}{\varepsilon_n x} \\ &= 1 - \frac{\beta}{(\varepsilon_n x)^{-\beta}} \int_{\varepsilon_n x}^{+\infty} dt t^{-1-\beta} e^{-t} \\ &= 1 - \frac{\beta}{(\varepsilon_n x)^{-\beta}} \Gamma(-\beta, \varepsilon_n x), \end{aligned} \quad (2.13)$$

where we have made the change of variables $\varepsilon_n xy \rightarrow t$ and introduced the incomplete Gamma function $\Gamma(z, s)$, which has the following definition and asymptotic relation with the (complete) Gamma function $\Gamma(z) \equiv \Gamma(z, 0)$ for $\Re[z] < 0$:

$$\Gamma(z, s) \equiv \int_s^{+\infty} dt t^{z-1} e^{-t} \sim \Gamma(z) + \frac{s^z}{z}, \quad s \rightarrow 0^+. \quad (2.14)$$

Since in our case $\Re[z] = -\alpha/2 < 0$, the use of the above relation is justified. Plugging

§2.3. Leading eigenmodes from the expected matrix.

Eq. (2.14) into Eq. (2.13), we obtain for $\varepsilon_n x \rightarrow 0^+$:

$$\begin{aligned}\Pi_{\alpha,n}[\nu_1(x)] &\sim 1 - \frac{\beta}{(\varepsilon_n x)^{-\beta}} \left[\Gamma(-\beta) - \frac{(\varepsilon_n x)^{-\beta}}{\beta} \right] \\ &= -\beta (\varepsilon_n x)^\beta \Gamma(-\beta) \\ &= (\varepsilon_n x)^\beta \Gamma(1 - \beta)\end{aligned}\tag{2.15}$$

This allows us to confirm that $\nu_1(x)$ is asymptotically an eigenfunction, for a suitable choice of β . Indeed, applying Eq. (2.15) into Eq. (2.10), we can recast Eq. (2.6) as

$$\Pi_{\alpha,n}[\nu_1(x)] \sim (\varepsilon_n x)^\beta \Gamma(1 - \beta) = \frac{\lambda_1}{n} \frac{\beta}{\alpha} x^{\alpha-\beta}\tag{2.16}$$

(provided we identify $\beta \equiv \alpha - \beta$, i.e. $\beta = \alpha/2$) to find

$$\lambda_1 = \frac{\alpha}{\beta} \Gamma(1 - \beta) n \varepsilon_n^\beta = -\alpha \Gamma\left(-\frac{\alpha}{2}\right) n^{1/2}.\tag{2.17}$$

This square-root growth of λ_1 can be confirmed via actual computations on the matrix $\mathbf{P}$ itself (see Fig. 2.2). Using Eq. (2.8), we also obtain all the eigenvector entries as

$$v_1^{(j)} = \nu_1(x_j) = \frac{1}{2} x_j^{\alpha/2} = \frac{1}{2} \left(\frac{n}{j}\right)^{1/2} \quad j = 1, n.\tag{2.18}$$

This is also confirmed numerically later on.




Figure 2.2: Largest eigenvalue λ_1 of $\mathbf{P}$ versus the number n of nodes, with weights $\{x_i\}_{i=1}^n$ either sampled i.i.d. from a Pareto distribution (in red, averaged over 5 samples) or assigned deterministically as $x_i = (n/i)^{1/\alpha}$ (in blue), compared with the analytic result in Eq. (2.17) (in black). Here $\alpha = 0.5$.

§2.3.2 Next eigenvalues

We now turn $k > 1$. Let us preliminarily introduce auxiliary complex-valued test functions $\nu_k^\pm(x)$ of the form still given by Eq. (2.9), with $\beta \in \mathbb{C}$:

$$\nu_k^\pm(x) = \frac{\beta_k^\pm}{\alpha} x^{\alpha-\beta_k^\pm}, \quad \beta_k^\pm \equiv \gamma_k \pm i\omega_k \in \mathbb{C} \quad (2.19)$$

where $\gamma_k = \Re[\beta_k^\pm] \in \mathbb{R}$ and $\pm\omega_k = \Im[\beta_k^\pm] \in \mathbb{R}$ are the real and imaginary parts of $\beta_k^\pm$, respectively. Note that β_k^+ and β_k^- are complex conjugate. Just like $\rho_\beta(x) \equiv \nu_1(x)\rho_\alpha(x)$ is a Pareto pdf with modified (real) index β , which has the same tail behaviour of a stable distribution with index β , $\rho_{\beta_k^\pm}(x) \equiv \nu_k^\pm(x)\rho_\alpha(x)$ is formally a generalized Pareto distribution with the same ‘tail behaviour’ as the recently defined stable distributions with complex-valued index $\beta_k^\pm$ [74, 75]. As we now show, these complex-valued scaling exponents generate log-periodicity, as typically found in systems with discrete scale invariance [35].

We will use the fact that the Gamma function of the complex conjugate of its argument is the complex conjugate of the Gamma function. The discussion in the previous section extends to complex-valued parameters $\beta_k^\pm = \gamma_k \pm i\omega_k$, leading to the following generalization of Eq. (2.15):

$$\Pi_{\alpha,n}[\nu_k^\pm(x)] \sim (\varepsilon_n x)^{\beta_k^\pm} \Gamma(1 - \beta_k^\pm) = (\varepsilon_n x)^\gamma (\varepsilon_n x)^{\pm i\omega} \Gamma(1 - \gamma \mp i\omega_k^\pm) \quad (2.20)$$

We then introduce the candidate eigenfunction

$$\nu_k(x) = \Re[\Gamma(1 - \beta_k^-) \nu_k^+(x)] \quad (2.21)$$

$$= \frac{1}{2} [\Gamma(1 - \beta_k^-) \nu_k^+(x) + \Gamma(1 - \beta_k^+) \nu_k^-(x)] \quad (2.22)$$

then Eq. (2.20) leads to the following asymptotics:

$$\Pi_{\alpha,n}[\nu_k(x)] \sim \frac{1}{2} \Gamma(1 - \beta_k^+) \Gamma(1 - \beta_k^-) [(\varepsilon_n x)^{\beta_k^+} + (\varepsilon_n x)^{\beta_k^-}] \quad (2.23)$$

$$= |\Gamma(1 - \beta_k^+)|^2 \Re[(\varepsilon_n x)^{\beta_k^+}] \quad (2.24)$$

$$= |\Gamma(1 - \gamma_k \pm i\omega_k)|^2 (\varepsilon_n x)^{\gamma_k} [(\varepsilon_n x)^{+i\omega} + (\varepsilon_n x)^{-i\omega}]. \quad (2.25)$$

To enforce the eigenvector equation $\Pi_{\alpha,n}[\nu_k(x)] = \frac{\lambda_k}{n} \nu_k(x)$, the above expression has to equal

$$\frac{\lambda_k}{n} \nu_k(x) = \frac{\lambda_k}{2n\alpha} [\Gamma(1 - \beta_k^-) \beta_k^+ x^{\alpha-\beta_k^+} + \Gamma(1 - \beta_k^+) \beta_k^- x^{\alpha-\beta_k^-}] \quad (2.26)$$

$$= \frac{\lambda_k}{n\alpha} \Re[\Gamma(1 - \beta_k^-) \beta_k^+ x^{\alpha-\beta_k^+}]. \quad (2.27)$$

This requires either $\beta_k^+ \equiv \alpha - \beta_k^+$ and $\beta_k^- \equiv \alpha - \beta_k^-$, which would however retrieve the case $k = 1$ that we have already considered for the largest eigenvalue

§2.3. Leading eigenmodes from the expected matrix.

($\beta_k^\pm = \gamma_k = \alpha/2$ and $\omega_k = 0$, since $\alpha \in \mathbb{R}$), or $\beta_k^+ \equiv \alpha - \beta_k^-$ and $\beta_k^- \equiv \alpha - \beta_k^+$, which still leads to $\gamma_k = \alpha/2$ (for all k) but puts no restriction on ω_k , so that $\beta_k^\pm = \alpha/2 \pm i\omega_k$. The defining restriction on ω_k is obtained by matching the other terms in Eqs. (2.25) and (2.27), yielding the two simultaneous conditions

$$\lambda_k = -\alpha \Gamma\left(-\frac{\alpha}{2} \mp i\omega_k\right) n^{1/2 \pm i\omega_k/\alpha} = -\alpha \Gamma\left(-\frac{\alpha}{2} \mp i\omega_k\right) n^{1/2} e^{\pm i(\omega_k/\alpha) \ln n}. \quad (2.28)$$

Since the right-hand sides of the above expression are complex conjugates of each other, the two simultaneous conditions ensure that λ_k is real, as required for the eigenvalues of a symmetric matrix. This means that ω_k has to be such that $\arg[\lambda_k] = 0, \pi$ modulo 2π , i.e.

$$\arg\left[\Gamma\left(-\frac{\alpha}{2} + i\omega_k\right)\right] = \frac{\omega_k}{\alpha} \ln n - k\pi, \quad k \in \mathbb{N}, \quad -\pi \leq \arg\left[\Gamma\left(-\frac{\alpha}{2} + i\omega_k\right)\right] < +\pi, \quad (2.29)$$

Note that the (necessarily real-valued) eigenvalues λ_k ($k \geq 1$) can be found graphically as the intersections between the real axis and the two spirals

$$\sigma^\pm(\omega) \equiv -\alpha \Gamma\left(-\frac{\alpha}{2} \mp i\omega\right) n^{1/2} e^{\pm i(\omega/\alpha) \ln n} \quad (2.30)$$

defined parametrically in the complex plane by extending the ω_k in Eq.(2.28) to any real value $\omega \geq 0$. See Fig. 2.3 for an illustration of this result, for three representative values of the index $\alpha = 0.2, 0.5, 0.8$. When the solutions for ω_k that realize Eq. (2.29) (or equivalently the intersections) are inserted into Eq.(2.28), they produce all the eigenvalues as

$$\lambda_k = (-1)^k \alpha \left| \Gamma\left(-\frac{\alpha}{2} + i\omega_k\right) \right| n^{1/2}, \quad (2.31)$$

which nicely extends the expression for λ_1 to values $k > 1$.

Graphical inspection

Unfortunately, Eq. (2.29) cannot be solved explicitly for ω_k , except for $k = 0$ (for which there is no solution for ω_0) and $k = 1$, for which $\omega_1 = 0$ is a solution, leading to the value λ_1 when inserted into Eq.(2.31): indeed, $\Gamma(-\frac{\alpha}{2}) < 0$, so that $\arg \Gamma(-\frac{\alpha}{2}) = -\pi$ solves Eq. (2.29) for $k = 1$ and $\omega_1 = 0$. The solutions for $k > 1$ can be inspected graphically, either as the aforementioned intersections between the real axis and the complex spiral, or as the intersections between the two curves defined as the l.h.s. and r.h.s. of Eq. (2.29). Specifically, the l.h.s. defines the function

$$f(\omega) \equiv \arg \Gamma\left(-\frac{\alpha}{2} + i\omega\right), \quad (2.32)$$

while the r.h.s. defines the family of straight lines

$$g_k(\omega) \equiv \omega \frac{\ln n}{\alpha} - k\pi, \quad k \geq 0. \quad (2.33)$$

The solutions ω_k to Eq. (2.29) are the values at the intersections

$$f(\omega_k) = g_k(\omega_k), \quad k \geq 0. \quad (2.34)$$




Figure 2.3: Comparison between the actual real-valued eigenvalues (red) of the expected adjacency matrix $\mathbf{P}$, obtained by solving Eq. (2.29) numerically, and the predicted eigenvalues obtained as the intersection between the real axis and the logarithmic spiral σ_+ (blue) defined in Eq. (2.30). Here $n = 10^4$ and $\alpha = 0.2, 0.5, 0.8$ from top to bottom.

§2.3. *Leading eigenmodes from the expected matrix.*

In Fig.2.4 we plot $f(\omega)$ and $g_k(\omega)$ for a few small values of k . The plot confirms that there is no intersection point ω_0 for $k = 0$, while the intersection for $k = 1$ is found at $\omega_1 = 0$, for which $f(\omega_1) = g_1(\omega_1) = -\pi$. What is more useful from the graphical inspection is that the next few intersections for $1 < k \leq k^*$ occur at a range of values ω_k close to some value ω_α at which $f(\omega)$ has a stationary point and varies very slowly around the corresponding value $\phi_\alpha \equiv f(\omega_\alpha)$, which is approximately a local ‘plateau’ for the function (the value k^* is defined such that for $k > k^*$ the value $f(\omega_k)$ is away from the plateau). We can therefore approximate the solutions ω_k for the first few values $k = 2, 3, \dots, k^*$ as the points at which $g_k(\omega)$ intersects the constant line $f(\omega) \approx \phi_\alpha$, and replace Eq. (2.34) with

$$\phi_\alpha \approx g_k(\omega_k), \quad 1 < k \leq k^*. \quad (2.35)$$

which, using Eq. (2.33), is solved by the values

$$\omega_k \approx \alpha \frac{k\pi + \phi_\alpha}{\ln n}, \quad \phi_\alpha \equiv f(\omega_\alpha), \quad 1 < k \leq k^*. \quad (2.36)$$

To complete our characterization of these approximate solutions, we need to estimate both k^* and ω_α .

Estimation of ω_α

We start from the determination of ω_α . Since the latter is defined as a stationary point of $f(\omega)$, we look for its value by calculating the derivative $f'(\omega)$ and requiring $f'(\omega_\alpha) = 0$, or more precisely:

$$\frac{\partial}{\partial \omega} \arg \Gamma \left(-\frac{\alpha}{2} + \mathbf{i}\omega \right) \Big|_{\omega=\omega_\alpha} = 0. \quad (2.37)$$

We now rearrange the identity $z = |z|e^{\mathbf{i} \arg z} = \sqrt{z \bar{z}} e^{\mathbf{i} \arg z}$ (valid for every complex number z , where $\bar{z}$ is the complex conjugate of z) to express

$$\arg z = -\mathbf{i} \ln \frac{z}{|z|} = \frac{\mathbf{i}}{2} \ln \frac{\bar{z}}{z}. \quad (2.38)$$

If $z = \arg \Gamma \left(-\frac{\alpha}{2} + \mathbf{i}\omega \right)$, we get

$$\arg \Gamma \left(-\frac{\alpha}{2} + \mathbf{i}\omega \right) = \frac{\mathbf{i}}{2} \ln \frac{\Gamma \left(-\frac{\alpha}{2} - \mathbf{i}\omega \right)}{\Gamma \left(-\frac{\alpha}{2} + \mathbf{i}\omega \right)} = \frac{\mathbf{i}}{2} \ln \Gamma \left(-\frac{\alpha}{2} - \mathbf{i}\omega \right) - \frac{\mathbf{i}}{2} \ln \Gamma \left(-\frac{\alpha}{2} + \mathbf{i}\omega \right). \quad (2.39)$$

In order to be able to differentiate the above expression with respect to ω , we need a closed-form expression for $\ln \Gamma(z)$ (also called the ‘log-Gamma function’), which is available as follows:

$$\ln \Gamma(z) = -\gamma z - \ln z + \sum_{m=1}^{+\infty} \left[\frac{z}{m} - \ln \left(1 + \frac{z}{m} \right) \right], \quad (2.40)$$

where

$$\gamma \equiv \lim_{m \rightarrow +\infty} \left[-\ln m + \sum_{l=1}^m \frac{1}{l} \right] = 0.5772 \dots \quad (2.41)$$




Figure 2.4: Graphical Solution of Eq. (2.29). In orange, the function $f(\omega)$ defined in Eq. (2.32), while the straight lines $g_k(\omega)$ defined by Eq. (2.33) are in red, with k from 0 to 10. The dashed red line corresponds to $g_0(\omega)$, which has no intersection ω_0 with $f(\omega)$. The solid red line at its immediate right is $g_1(\omega)$, which intersects $f(\omega)$ at the exact value $\omega_1 = 0$ where $f(\omega_1) = g_1(\omega_1) = -\pi$. Further on the right, the next straight lines $g_k(\omega)$ (with $k > 1$) are illustrated in light red. The solution to Eq. (2.29) for these values of k can be obtained graphically where the orange and light red lines meet. In this regime, $f(\omega)$ can be approximated as the horizontal line $f(\omega) \approx \phi_\alpha \equiv f(\omega_\alpha)$ (blue line), where the value ω_α is the stationary point such that $f'(\omega_\alpha) = 0$. The horizontal line is no longer a good approximation for $f(\omega_k)$ for larger k , but that is already the regime for which the value of $|\Gamma(-\frac{\alpha}{2} + i\omega_k)|$ (plotted in green) is exponentially suppressed, as one can verify from Eqs. (2.53).

is the Euler-Mascheroni constant. Using Eq. (2.40) with $z = -\alpha/2 \pm i\omega$ into Eq. (2.39), we arrive at

$$\arg \Gamma\left(-\frac{\alpha}{2} + i\omega\right) = -\gamma\omega - \arg\left(-\frac{\alpha}{2} + i\omega\right) + \sum_{m=1}^{+\infty} \left[\frac{\omega}{m} - \arg\left(m - \frac{\alpha}{2} + i\omega\right) \right]. \quad (2.42)$$

Now, in order to differentiate with respect to ω as prescribed by Eq. (2.37), we preliminarily calculate

$$\frac{\partial}{\partial \omega} \arg\left(m - \frac{\alpha}{2} + i\omega\right) = \frac{i}{2} \frac{\partial}{\partial \omega} \ln\left(\frac{m - \frac{\alpha}{2} - i\omega}{m - \frac{\alpha}{2} + i\omega}\right) = \frac{m - \frac{\alpha}{2}}{(m - \frac{\alpha}{2})^2 + \omega^2}, \quad (2.43)$$

where we have used Eq. (2.38) again. Using the above expression into Eq. (2.42), we obtain

$$\frac{\partial}{\partial \omega} \arg \Gamma\left(-\frac{\alpha}{2} + i\omega\right) = -\gamma + \frac{\frac{\alpha}{2}}{(\frac{\alpha}{2})^2 + \omega^2} + \sum_{m=1}^{+\infty} \left[\frac{1}{m} - \frac{m - \frac{\alpha}{2}}{(m - \frac{\alpha}{2})^2 + \omega^2} \right]. \quad (2.44)$$

Now, imposing Eq. (2.37) translates into the following:

$$-\gamma + \frac{\frac{\alpha}{2}}{\left(\frac{\alpha}{2}\right)^2 + \omega_\alpha^2} + \sum_{m=1}^{+\infty} \left[\frac{1}{m} - \frac{m - \frac{\alpha}{2}}{\left(m - \frac{\alpha}{2}\right)^2 + \omega_\alpha^2} \right] = 0, \quad (2.45)$$

which, if the sum is truncated to any finite order for m , can be solved algebraically for ω_α . In general, the number of solutions of the above equation (i.e. stationary points of the function f) may depend on α . How many (and how closely) such solutions are recovered after truncating the sum will depend on the order of truncation. For our purposes here (i.e. identifying the first stationary point as shown in Fig. 2.4), we simply truncate to zeroth order and neglect the sum altogether, to get the condition

$$-\gamma + \frac{\frac{\alpha}{2}}{\left(\frac{\alpha}{2}\right)^2 + \omega_\alpha^2} = 0, \quad (2.46)$$

which is solved by the single (positive) value

$$\omega_\alpha \equiv \sqrt{\frac{\alpha}{2} \left(\frac{1}{\gamma} - \frac{\alpha}{2} \right)}. \quad (2.47)$$

This finally allows us to compute

$$\phi_\alpha \equiv f(\omega_\alpha) = f \left(\sqrt{\frac{\alpha}{2} \left(\frac{1}{\gamma} - \frac{\alpha}{2} \right)} \right), \quad (2.48)$$

which, when inserted into Eq. (2.36), produces our approximate solutions for ω_k to Eq. (2.29) with $1 < k \leq k^*$. A comparison between the actual solutions ω_k of Eq. (2.29) and the approximate solutions given in (2.36) is shown in Fig. 2.5, confirming a very good agreement, especially for smaller values of α . Note that our only use of these approximate solutions is the approximate proportionality between ω_k and k , which is used below to prove the exponential decay of the prefactor of the leading eigenvalues with k . This linearity only requires that $f(\omega)$ is sufficiently flat around some value ϕ_α before $|\Gamma(-\frac{\alpha}{2} + \mathbf{i}\omega)|$ is exponentially suppressed (see Fig. 2.4). A more accurate estimate of ϕ_α would improve the estimate of ϕ_α used in Eq. (2.36) (hence improving the slope of the predicted lines in Fig. 2.5), but not modify the approximate proportionality between ω_k and k , which is all we need in what follows.

Exponential decay of the prefactor of the leading eigenvalues

We now turn to the estimation of the order of k^* . Let us consider the Stirling approximation [159, (5.11.9)] for the absolute value of the Gamma function:

$$|\Gamma(a + \mathbf{i}b)| \sim \sqrt{2\pi} |b|^{a-\frac{1}{2}} e^{-\frac{\pi|b|}{2}}, \quad (2.49)$$

where $a, b \in \mathbb{R}$. Choosing $a = -\alpha/2$ and $b = \omega_k \geq 0$, this implies

$$\left| \Gamma \left(-\frac{\alpha}{2} + \mathbf{i}\omega_k \right) \right| \sim \sqrt{2\pi} \omega_k^{-\frac{\alpha+1}{2}} e^{-\frac{\pi\omega_k}{2}}, \quad (2.50)$$




Figure 2.5: Values of ω_k as a function of k . The isolated points are the actual, numerical solutions of Eq. (2.29), for $\alpha = 0.2, 0.5, 0.8$ (in blue, orange and green respectively). The solid lines are the approximate linear solutions defined by Eq. (2.36) with $k > 1$ (plus the exact solution $\omega_1 = 0$) for the same values of $\alpha = 0.2, 0.5, 0.8$ (and same colors).

showing an exponential decay with increasing ω_k . This exponential decay is visible in Fig. 2.4. If we further use the expression we derived in Eq. (2.36) for the approximate solutions ω_k for $1 < k \leq k^*$, we see that the above behaviour translates into an exponential decay as a function of $k/\ln n$:

$$\left| \Gamma \left(-\frac{\alpha}{2} + \mathbf{i} \omega_k \right) \right| \sim \sqrt{2\pi} \left(\alpha \frac{k\pi + \phi_\alpha}{\ln n} \right)^{-\frac{\alpha+1}{2}} e^{-\frac{\alpha\pi}{2} \frac{k\pi + \phi_\alpha}{\ln n}}, \quad 1 < k \leq k^*. \quad (2.51)$$

Looking at Fig. 2.4, we see that when $k \gtrsim k^*$ (i.e. when the intersections between $f(\omega)$ and $g_k(\omega)$ depart from the approximate plateau of constant height ϕ_α), the value of $\left| \Gamma \left(-\frac{\alpha}{2} + \mathbf{i} \omega_k \right) \right|$ is already exponentially suppressed. Therefore Eq. (2.51) indicates that k^* is of order $\ln n$. Obviously, the magnitude of the corresponding eigenvalue λ_k of $\mathbf{P}$, which is given by Eq. (2.31), is also exponentially (in $k/\ln n$) suppressed:

$$|\lambda_k| = \alpha \left| \Gamma \left(-\frac{\alpha}{2} + \mathbf{i} \omega_k \right) \right| \sqrt{n} \quad (2.52)$$

$$\sim \alpha \sqrt{2\pi n} \left(\alpha \frac{k\pi + \phi_\alpha}{\ln n} \right)^{-\frac{\alpha+1}{2}} e^{-\frac{\alpha\pi}{2} \frac{k\pi + \phi_\alpha}{\ln n}}. \quad (2.53)$$

This means that, for $k = O(\ln n)$, the approximation in Eq. (2.36) breaks down, but on the other hand the corresponding values of λ_k become exponentially small. This is also where, moving from the expected matrix $\mathbf{P}$ to a realized random matrix $\mathbf{A}$, the randomness will create noisy eigenvalues ‘eating up’ the eigenvalues of $\mathbf{P}$.

Indeed, generalized BBP-type results typically assume that the expected matrix $\mathbf{P}$ has a fixed, finite rank in order to guarantee a clean separation between informative (structural) eigenvalues and the random bulk. In our model this assumption fails literally, since $\mathbf{P}$ is not finite-rank. Nevertheless, both our numerics and the above arguments indicate that the k -th largest (in absolute value) eigenvalue λ_k of $\mathbf{P}$ decays rapidly with $k/\ln n$, so that the rank of the matrix is *effectively* of order $\ln n$, and

hence $o(n)$. This aligns with recent literature suggesting that a BBP-like transition can persist to order $o(n)$, hence also beyond strict finite rank [73, 76]. Numerics also indicate that the edge of the random spectrum contributed by $\mathbf{H}$ scales as $C(\alpha)\sqrt{n}$ with $C(\alpha)$ independent of n . Consequently, only those signal eigenvalues that are not exponentially suppressed can rise above the bulk, i.e., indices up to $k^* \sim \ln n$. In this sense, the signal behaves as if it had *effective rank* $\Theta(\ln n)$: the spectrum of $\mathbf{A}$ consists of order $\ln n$ outliers (delta spikes from $\mathbf{P}$) on top of the bulk determined by $\mathbf{H}$.

§2.3.3 Next Eigenvectors

We now turn to the eigenvectors for $k > 1$. Plugging ω_k into Eq. (2.22), we get

$$\nu_k(x) = \frac{x^{\alpha/2}}{\alpha} \Re \left[\left(\frac{\alpha}{2} - \mathbf{i}\omega_k \right) \Gamma \left(1 - \frac{\alpha}{2} - \mathbf{i}\omega_k \right) x^{\mathbf{i}\omega_k} \right]. \quad (2.54)$$

Using Eq. (2.8), we re-express the eigenvector entries as

$$v_k^{(j)} = \sqrt{\frac{n}{j\alpha^2}} \Re \left[\left(\frac{\alpha}{2} - \mathbf{i}\omega_k \right) \Gamma \left(1 - \frac{\alpha}{2} - \mathbf{i}\omega_k \right) \left(\frac{n}{j} \right)^{\frac{\mathbf{i}\omega_k}{\alpha}} \right] \quad (2.55)$$

$$= -\sqrt{\frac{n}{j}} \left(\frac{\alpha}{4} + \frac{\omega_k^2}{\alpha} \right) \Re \left[\Gamma \left(-\frac{\alpha}{2} - \mathbf{i}\omega_k \right) \left(\frac{n}{j} \right)^{\frac{\mathbf{i}\omega_k}{\alpha}} \right] \quad (2.56)$$

Since $x^{\mathbf{i}\omega_k} = e^{\mathbf{i}\omega_k \ln x}$ and $j^{\mathbf{i}\omega_k/\alpha} = e^{\mathbf{i}(\omega_k/\alpha) \ln j}$, the structure of the above eigenfunctions and eigenvectors is oscillating, log-periodic in their arguments (x and j respectively), and modulated by a power-law profile ($x^{\alpha/2}$ and $j^{-1/2}$ respectively). In Figs. 2.6 2.7 2.8 we show that our analytically derived log-periodic formula tracks the actual eigenvectors of both $\mathbf{P}$ and $\mathbf{A}$ remarkably well for small k .

These eigenvectors do not display traditional delta-like localization, as they are broadly supported on a wide range of weight percentiles with largest amplitude around the hub node ($j = 1$), followed by power-law decay $j^{-1/2}$. As k grows large enough (below we provide evidence that this happens for k growing with n but of order at most $\ln n$), we find a departure of the eigenvectors of $\mathbf{A}$ from those of $\mathbf{P}$, due to the entrance of the associated eigenvalues into the random bulk where eigenvectors are expected to be more uniformly spread and less structured.

We also note that, when combined with Eq. (2.31), the above expression allows to rewrite the eigenvector entries in terms of the corresponding eigenvalue λ_k as

$$v_k^{(j)} = \frac{(-1)^{k+1} \lambda_k}{\sqrt{j}} \left(\frac{\alpha}{4} + \frac{\omega_k^2}{\alpha} \right) \Re \left[j^{\frac{\mathbf{i}\omega_k}{\alpha}} \right] = \frac{(-1)^{k+1} \lambda_k}{\sqrt{j}} \left(\frac{\alpha}{4} + \frac{\omega_k^2}{\alpha} \right) \cos \left(\frac{\mathbf{i}\omega_k}{\alpha} \ln j \right). \quad (2.57)$$

In particular, the first entry $v_k^{(1)}$ of each eigenvector v_k (which is the entry that localizes around the hub node $j = 1$) is

$$v_k^{(1)} = (-1)^{k+1} \lambda_k \left(\frac{\alpha}{4} + \frac{\omega_k^2}{\alpha} \right), \quad (2.58)$$

so that, for $j = 2, n$,

$$v_k^{(j)} = v_k^{(1)} \frac{\cos\left(\frac{i\omega_k}{\alpha} \ln j\right)}{\sqrt{j}}. \quad (2.59)$$

The above expressions illustrate how each (leading) eigenvector v_k has a structure where the first entry $v_k^{(1)}$ (located at the hub node $j = 1$) has an amplitude governed by the corresponding eigenvalue λ_k , and the next entries $j > 1$ follow a power-law decay $\sim j^{-1/2}$ with log-periodic oscillations.

§2.4 Bulk of the Spectrum

As we have seen in Fig. 2.1 for the eigenvalues and from Figs. 2.6 2.7 2.8 for the eigenvectors, there is a value of k beyond which the spectral properties of $\mathbf{P}$ are no longer able to describe those of $\mathbf{A}$. We refer to this as the ‘entrance’ into the random bulk, i.e. the noisy part of the spectrum generated primarily by the fluctuation of $\mathbf{H} = \mathbf{A} - \mathbf{P}$. The analysis of this part of the spectrum is highly non-trivial, and eludes the purposes of this chapter. What we are mainly interested in here is the estimation of where the signal ‘meets’ the bulk. In particular, we look for an upper bound for the spectral norm of $\mathbf{H}$, whose elements $H_{ij} = A_{ij} - P_{ij}$ are independent but not identically distributed random variables in the range $[-1, 1]$, the variance of H_{ij} being $p_{ij}(1 - p_{ij})$. The study of such so-called random matrices with a variance profile, or structured random matrices, is of great interest in the mathematical ecology community, and has fostered a lively discussion in the mathematical literature [77, 78, 79, 80, 81, 82, 83, 84]. In particular, we are interested in the bound derived in [83], that can be extended to matrices with entries having a different subgaussian distributions, like the centered Bernoulli of our case.

Indeed, let us condition on the latent variables so that the edge probabilities (p_{ij}) are deterministic and $\{A_{ij}\}_{i < j}$ are independent. Let

$$v_{ij} := p_{ij}(1 - p_{ij}), \quad (2.60)$$

$$\sigma := \max_{1 \leq i \leq n} \left(\sum_{j \neq i} v_{ij} \right)^{1/2}, \quad (2.61)$$

$$\sigma_* := \max_{i \neq j} \sqrt{v_{ij}} \leq \frac{1}{2}. \quad (2.62)$$

Then there exist absolute constants $C, c > 0$ such that

$$\mathbb{E} \|\mathbf{H}\| \leq C \left(\sigma + \sigma_* \sqrt{\log n} \right). \quad (2.63)$$

Moreover, for any $0 < \varepsilon \leq \frac{1}{2}$ and any $t \geq 0$,

$$\mathbb{P} \left(\|\mathbf{H}\| \geq (1 + \varepsilon) 2\sqrt{2} \sigma + t \right) \leq n \exp \left(- \frac{t^2}{c_\varepsilon \sigma_*^2} \right), \quad (2.64)$$

where $c_\varepsilon > 0$ is universal.




Figure 2.6: Plot of eigenvector entries $v_k^{(j)}$ versus j , for the first 10 eigenvectors $k = 1, \dots, 10$ (from left to right and top to bottom) and for all entries $j = 1, \dots, n$. Blue dots correspond to the actual eigenvectors entries of $\mathbf{P}$ obtained by solving Eq. (2.29) explicitly, while red solid lines represent the corresponding predicted entries obtained from Eq. (2.56). Orange dots are the entries of the eigenvectors of $\mathbf{A}$, and we can see that in the regime outside the bulk they are in quite a good accordance with those of $\mathbf{P}$ and hence the ones we obtained analytically. Here $\alpha = 0.2$ and $n = 10^4$.

We can show this by writing $H_{ij} = \sqrt{v_{ij}} \xi_{ij}$ with $\xi_{ij} = (A_{ij} - p_{ij})/\sqrt{v_{ij}}$. Conditionally on (p_{ij}) , the ξ_{ij} are independent, centered, and uniformly sub-Gaussian (bounded support), so the sub-Gaussian extension of the main bound applies with




Figure 2.7: Plot of eigenvector entries $v_k^{(j)}$ versus j , for the first 10 eigenvectors $k = 1, \dots, 10$ (from left to right and top to bottom) and for all entries $j = 1, \dots, n$. Blue dots correspond to the actual eigenvectors entries of $\mathbf{P}$ obtained by solving Eq. (2.29) explicitly, while red solid lines represent the corresponding predicted entries obtained from Eq. (2.56). Orange dots are the entries of the eigenvectors of $\mathbf{A}$, and we can see that in the regime outside the bulk they are in quite a good accordance with those of $\mathbf{P}$ and hence the ones we obtained analytically. Here $\alpha = 0.5$ and $n = 10^4$.

scale matrix $b_{ij} = \sqrt{v_{ij}}$. This yields (2.63) from the *Gaussian* Theorem 1.1 together with its sub-Gaussian transfer [83](Cor. 3.3):

$$\mathbb{E}\|\mathbf{H}\| \lesssim \sigma + \sigma_* \sqrt{\log n}. \quad (2.65)$$




Figure 2.8: Plot of eigenvector entries $v_k^{(j)}$ versus j , for the first 10 eigenvectors $k = 1, \dots, 10$ (from left to right and top to bottom) and for all entries $j = 1, \dots, n$. Blue dots correspond to the actual eigenvectors entries of $\mathbf{P}$ obtained by solving Eq. (2.29) explicitly, while red solid lines represent the corresponding predicted entries obtained from Eq. (2.56). Orange dots are the entries of the eigenvectors of $\mathbf{A}$, and we can see that in the regime outside the bulk they are in quite a good accordance with those of $\mathbf{P}$ and hence the ones we obtained analytically. Here $\alpha = 0.8$ and $n = 10^4$.

For tails, note that $|H_{ij}| \leq 1$ and $\mathbb{E}H_{ij} = 0$. Now, [83] (Corollary 3.12) gives, for *symmetric* bounded entries, a variance-sensitive inequality of the form

$$\mathbb{P}(\|\mathbf{H}\| \geq (1 + \varepsilon)2\sigma + t) \leq n \exp(-t^2/\tilde{c}_\varepsilon \sigma_*^2). \quad (2.66)$$




Figure 2.9: Edge of the bulk of the spectrum versus the number n of nodes. Points are the actual edge eigenvalues (largest eigenvalue of $\mathbf{H}$, averaged over 10 realizations), for various values of $\alpha = 0.3$ (red), $\alpha = 0.4$ (blue), $\alpha = 0.5$ (green), $\alpha = 0.6$ (yellow) and $\alpha = 0.7$ (purple). The dashed line is the upper bound $\sqrt{n}/2$.

Our entries are not symmetric, but [83](Remark 3.13) shows that symmetrization inflates only the leading constant by a factor $\sqrt{2}$, which yields (2.64). As $v_{ij} \leq p_{ij}$, we have $\sigma^2 \leq d_{\max}$ where $d_{\max} := \max_i \sum_{j \neq i} p_{ij}$ is the maximal expected degree. Hence

$$\mathbb{E}\|\mathbf{H}\| \leq C(\sqrt{d_{\max}} + \sqrt{\log n}), \quad (2.67)$$

and

$$\|\mathbf{H}\| \leq C'(\sqrt{d_{\max}} + \sqrt{\log n}) \quad \text{w.h.p.} \quad (2.68)$$

As a crude but universal bound, since $v_{ij} \leq \frac{1}{4}$, $\sigma \leq \sqrt{n}/2$, we get

$$\|\mathbf{H}\| \lesssim \frac{\sqrt{n}}{2} + O(\sqrt{\log n}) \quad (2.69)$$

In Fig. 2.9 we show that the square root growth (that is commonplace in Wigner matrices) seems to be the right one, and that the coefficient coming from the maximization of the row/column sum of the variance profile matrix is already a quite good envelope for all values of α . Nonetheless this crude bound needs refinement, at least to include the dependence on α . This calls for the full characterization of the bulk spectrum, which we leave to future work. What is important for our purposes here is noticing that the edge of the bulk lives on the same scale $\sqrt{n}$ as the leading eigenvalues given by Eq. (2.31). Since the prefactor of $\sqrt{n}$ for those outliers decays exponentially with ω_k (and, via Eq. (2.36), with $k/\ln n$), we see that the bulk will necessarily ‘eat out’ the leading eigenvalues for some value of k of order at most $\ln n$.

§2.5 Concluding remarks

In this chapter, we provided a characterization of the spectral properties of the Multi-Scale Model (MSM) in the infinite-mean regime ($0 < \alpha < 1$), which has been recently identified as an annealed random graph model that is exactly invariant under node aggregation [1, 3]. This invariance is a key concept in the field of network renormalization [54], but its connection with discrete scale invariance and the associated idea of complex scaling exponents [35, 85] had not been investigated so far. We showed that the leading eigenvalues of the expected adjacency matrix $\mathbf{P}$ scale as $\sqrt{n}$ and accurately proxy the outlier eigenvalues (whose number is growing and of order at most $\ln n$) of the realized adjacency matrix $\mathbf{A}$. Remarkably, these outliers have an approximately constant spacing on a logarithmic scale of their index k and can be identified as the intersection between the real axis and a *spira mirabilis* in the complex plane, a self-similar object naturally emerging in presence of complex scaling exponents and stable laws with complex stability index [74, 75]. We also derived a closed-form expression for the outlier eigenvectors, showing they are delocalized but biased toward high-weight nodes. The leading eigenvectors beyond the principal (Perron–Frobenius) one exhibit log-periodic oscillations, a feature naturally produced by discrete scale invariance [35, 85, 86] which is here present in the property that the same random graph model describes exactly (up to a parameter rescaling) any coarse-grained version of the original graph, where sets of nodes of the same cardinality are aggregated into supernodes [1, 3]. Our results qualitatively confirm the findings of [33, 34], while partially extending them to the regime of infinite-rank perturbations of random matrices as in more recent attempts [73, 76]. In summary, the MSM illustrates how infinite-mean heterogeneity generates spectral phenomena far from classical paradigms. Our analysis suggests parallels with other models such as Chung–Lu, where broad degree distributions induce hub-driven outliers. Open directions include testing the universality of the $\sqrt{n}$ scaling and log-periodicity across other networks and models. More broadly, our results provide a framework for understanding spectra in multi-scale and heavy-tailed systems, and lay the ground for future work in these fields.

