

Random Matrix Statistics, Loop Equations, and Random Regular Graphs

University of Michigan Random Matrix Theory Summer School

Jiaoyang Huang

Abstract

These notes introduce global and local eigenvalue statistics, explain a microscopic loop-equation convergence framework, and describe its application to random regular graphs. They follow the content and order of the supplied nineteen handwritten pages. Calculations and conventions have been made explicit, and the sketches have been reconstructed as numerical experiments or mathematical diagrams. Deep results used as inputs are identified as such; the final local-resampling argument retains its role as a proof outline.

Contents

1	Introduction: random matrices and universality	2
2	Bulk and edge statistics	5
3	Examples beyond random matrices	8
4	A convergence framework for local random-matrix statistics	11
5	Bulk loop equations for real Wigner matrices	14
6	Moment bounds and tightness of microscopic transforms	19
7	Random regular graphs and their local geometry	24
8	Tree Green functions and the global spectral law	27
9	Expansion, spectral gaps, and Ramanujan graphs	30
10	Local resampling and the cavity self-consistent equation	34
11	The first edge loop equation	37
A	Numerical experiments and reproduction	39

1 Introduction: random matrices and universality

The lectures have three parts. We first introduce random matrix statistics and some of their appearances in probability, number theory, and quantum chaos. We then develop a convergence framework based on microscopic loop equations. Finally, we discuss its application to edge universality for sparse random regular graphs.

1.1 Wigner matrices and the Gaussian ensembles

Wigner introduced random matrices in the 1950s as models for the energy levels of heavy nuclei. The striking feature is that local spectral statistics depend much more on symmetry than on the precise distribution of the matrix entries.

A *Wigner matrix* is an $N \times N$ real symmetric or complex Hermitian random matrix $H = (h_{ij})$. The entries on and above the diagonal are independent, $h_{ji} = \overline{h_{ij}}$, and

$$\mathbb{E}h_{ij} = 0, \quad \mathbb{E}|h_{ij}|^2 = \frac{1}{N} \quad (i \neq j), \quad \mathbb{E}|h_{ii}|^2 \leq \frac{C}{N}. \quad (1.1)$$

Moment or tail assumptions will be specified when needed. For the elementary semicircle-law proof below, assume $\mathbb{E}|h_{ij}|^4 \leq C/N^2$ for the off-diagonal entries. When discussing complex bulk universality, we use the standard genuinely complex symmetry class, for example $\mathbb{E}h_{ij}^2 = 0$; Hermitian symmetry alone does not exclude a real ensemble.

The two basic examples are:

- The *Gaussian orthogonal ensemble* (GOE): $h_{ij} \sim \mathcal{N}(0, N^{-1})$ for $i < j$, and $h_{ii} \sim \mathcal{N}(0, 2N^{-1})$. For every deterministic orthogonal matrix O , $OH O^\top \stackrel{d}{=} H$.
- The *Gaussian unitary ensemble* (GUE): $h_{ij} = (X_{ij} + iY_{ij})/\sqrt{2N}$ for $i < j$, where X_{ij}, Y_{ij} are independent standard real Gaussians, and $h_{ii} \sim \mathcal{N}(0, N^{-1})$. For every deterministic unitary matrix U , $UHU^* \stackrel{d}{=} H$.

We order the eigenvalues throughout as $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N$ and write $H = U\Lambda U^*$, with normalized eigenvectors $u_1, \dots, u_N$.

1.2 The semicircle law and resolvents

The empirical spectral measure and semicircle density are

$$\mu_N = \frac{1}{N} \sum_{\alpha=1}^N \delta_{\lambda_\alpha}, \quad \rho_{\text{sc}}(x) = \frac{\sqrt{(4-x^2)_+}}{2\pi}.$$

The global semicircle law asserts $\mu_N \Rightarrow \rho_{\text{sc}}(x) dx$ in probability. Here and below $\Rightarrow$ denotes weak convergence of measures or convergence in distribution, according to context. For $z = E + i\eta \in \mathbb{C}_+$, define

$$G(z) = (H - zI)^{-1}, \quad m_N(z) = \frac{1}{N} \text{Tr } G(z) = \int \frac{\mu_N(dx)}{x - z}. \quad (1.2)$$

The spectral decomposition gives

$$G(z) = \sum_{\alpha=1}^N \frac{u_\alpha u_\alpha^*}{\lambda_\alpha - z}, \quad \|G(z)\| \leq \eta^{-1}.$$




Figure 1: New simulations of GOE and GUE spectra. Left: empirical eigenvalue densities pooled from five independent samples of size $N = 1600$ for each ensemble, compared with the semicircle density. Center: the first realization from each ensemble in the bulk coordinate $x = N\pi\rho_{\text{sc}}(0)\lambda = N\lambda$, whose limiting intensity is $1/\pi$. Right: the same realizations near the upper edge in the coordinate $x = N^{2/3}(\lambda - 2)$. Spectra were generated by the exact Gaussian β -ensemble tridiagonal model.

Since $\text{Im } G = (G - G^*)/(2i)$,

$$GG^* = \frac{\text{Im } G}{\eta}, \quad \sum_j |G_{ij}|^2 = \frac{\text{Im } G_{ii}}{\eta}. \quad (1.3)$$

These are the *Ward identities*. The same row and column identities hold because G is a normal matrix.

1.3 A resolvent proof of the global semicircle law

Let $H^{(i)}$ denote the minor obtained by deleting row and column i , let $G^{(i)} = (H^{(i)} - zI)^{-1}$, and let $h_i = (h_{ji})_{j \neq i}$. Importantly, h_i is independent of $H^{(i)}$. The normalization $m_N^{(i)} = N^{-1} \text{Tr } G^{(i)}$ retains the denominator N .

Step 1: Schur complement. For an invertible block D , the upper-left block of $\begin{pmatrix} A & B \\ C & D \end{pmatrix}^{-1}$ is $(A - BD^{-1}C)^{-1}$. Therefore

$$G_{ii} = \frac{1}{h_{ii} - z - h_i^* G^{(i)} h_i}. \quad (1.4)$$

Step 2: concentration of a quadratic form. If h has independent centered entries with variance $1/N$ and fourth moment at most C/N^2 , then for any deterministic matrix A ,

$$\mathbb{E}(h^* A h) = \frac{1}{N} \text{Tr } A, \quad \mathbb{E} \left| h^* A h - \frac{1}{N} \text{Tr } A \right|^2 \leq \frac{C'}{N^2} \text{Tr}(A A^*). \quad (1.5)$$

To check the second bound, expand the centered quadratic form as $\sum_{j,k} (\bar{h}_j h_k - \delta_{jk}/N) A_{jk}$. Independence and centering eliminate all terms except paired indices and fourth-moment diagonal terms. Each surviving coefficient is bounded by C'/N^2 ; the possible transpose pairing is bounded using $\sum_{j,k} |A_{jk} A_{kj}| \leq \sum_{j,k} |A_{jk}|^2$. Conditioning on $H^{(i)}$ and using $\|G^{(i)}\| \leq \eta^{-1}$ gives

$$Z_i := h_i^* G^{(i)} h_i - m_N^{(i)}, \quad \mathbb{E} |Z_i|^2 \leq \frac{C'}{N\eta^2}. \quad (1.6)$$

Step 3: removing a minor. The resolvent minor identity is

$$G_{jk}^{(i)} = G_{jk} - \frac{G_{ji}G_{ik}}{G_{ii}} \quad (j, k \neq i).$$

Consequently, by Cauchy–Schwarz and the row and column Ward identities,

$$\begin{aligned} |m_N - m_N^{(i)}| &= \left| \frac{1}{NG_{ii}} \sum_j G_{ji}G_{ij} \right| \\ &\leq \frac{\operatorname{Im} G_{ii}}{N\eta|G_{ii}|} \leq \frac{1}{N\eta}. \end{aligned} \quad (1.7)$$

In the Hermitian case one must keep the product $G_{ji}G_{ij}$ before taking absolute values; it is not generally G_{ij}^2 .

Step 4: the self-consistent equation. Set $\delta_i = h_{ii} + m_N - m_N^{(i)} - Z_i$. Then

$$G_{ii} = \frac{1}{-z - m_N + \delta_i}, \quad \mathbb{E} \frac{1}{N} \sum_i |\delta_i| \leq \frac{C}{\sqrt{N}} + \frac{1}{N\eta} + \frac{C}{\sqrt{N}\eta}.$$

The imaginary parts of both denominators $-z - m_N$ and $-z - m_N + \delta_i = h_{ii} - z - h_i^* G^{(i)} h_i$ are at most $-\eta$. Thus

$$\left| G_{ii} - \frac{1}{-z - m_N} \right| \leq \frac{|\delta_i|}{\eta^2}.$$

For each fixed $z \in \mathbb{C}_+$, averaging yields

$$m_N = -\frac{1}{z + m_N} + o_{\mathbb{P}}(1), \quad 1 + zm_N + m_N^2 = o_{\mathbb{P}}(1). \quad (1.8)$$

The correct denominator has a minus sign in front of both z and m_N . The equation $1 + zm + m^2 = 0$ has exactly one solution in $\mathbb{C}_+$:

$$m_{\text{sc}}(z) = \frac{-z + \sqrt{z^2 - 4}}{2}, \quad (1.9)$$

where the square root is analytic off $[-2, 2]$ and asymptotic to z at infinity. To justify the limiting step, use $|m_N| \leq \eta^{-1}$ and $\operatorname{Im} m_N > 0$. Every possible limit lies in the closed upper half-plane and solves the quadratic. No root is real when $z \in \mathbb{C}_+$, and the other root lies in the lower half-plane. Hence $m_N(z) \rightarrow m_{\text{sc}}(z)$ in probability. The continuity theorem for Stieltjes transforms proves the global semicircle law. This fixed- η argument by itself is not a local law.

1.4 Gaussian beta ensembles and the log-gas picture

For GOE and GUE the matrix density, with respect to the independent real coordinates, is proportional to $\exp\{-\beta N \operatorname{Tr} H^2/4\}$, where $\beta = 1$ and 2 , respectively. Changing variables to eigenvalues and eigenvectors and integrating out the eigenvectors produces the Vandermonde Jacobian:

$$p_{N,\beta}(\lambda) = \frac{1}{Z_{N,\beta}} \prod_{i < j} |\lambda_i - \lambda_j|^\beta \exp \left\{ -\frac{\beta N}{4} \sum_i \lambda_i^2 \right\}. \quad (1.10)$$

This defines a probability density for every $\beta > 0$, either on $\mathbb{R}^N$ or on an ordered chamber with the appropriate normalizing constant. It has the Gibbs form $Z_{N,\beta}^{-1} e^{-\beta \mathcal{H}_N}$, where

$$\mathcal{H}_N(\lambda) = - \sum_{i < j} \log |\lambda_i - \lambda_j| + \frac{N}{4} \sum_i \lambda_i^2.$$

The particles repel logarithmically and are confined by a quadratic potential. The parameter β is inverse temperature. At fixed N , as $\beta \rightarrow \infty$, the ordered particles concentrate around the unique minimizer in the ordered chamber. More general beta ensembles replace the quadratic confinement by a potential V .

Dumitriu and Edelman [1] gave a tridiagonal realization for all $\beta > 0$:

$$T_{N,\beta} = \frac{1}{\sqrt{N\beta}} \begin{pmatrix} g_1 & \chi_{(N-1)\beta} & & & \\ \chi_{(N-1)\beta} & g_2 & \chi_{(N-2)\beta} & & \\ & \chi_{(N-2)\beta} & g_3 & \ddots & \\ & & \ddots & \ddots & \chi_\beta \\ & & & \chi_\beta & g_N \end{pmatrix}. \quad (1.11)$$

Here all displayed random variables are independent, $g_i \sim \mathcal{N}(0, 2)$, and χ_ν is the positive square root of a chi-square variable with ν degrees of freedom. For noninteger $\nu > 0$, this means $\chi_\nu^2 \sim \text{Gamma}(\nu/2, 2)$. The eigenvalue density is exactly (1.10). This representation also makes numerical simulation efficient.

2 Bulk and edge statistics

2.1 Bulk scaling and the sine process

Fix $E \in (-2, 2)$ and put $\rho_E = \rho_{\text{sc}}(E)$. The typical original eigenvalue spacing near E is $1/(N\rho_E)$. Throughout these notes the rescaled bulk process is

$$\Xi_{N,E} = \sum_{i=1}^N \delta_{N\pi\rho_E(\lambda_i - E)}. \quad (2.1)$$

Its limiting intensity is $1/\pi$, so the mean rescaled spacing is π . For the Gaussian beta ensemble, $\Xi_{N,E}$ converges to the Sine $_\beta$ point process in this convention. Different references use intensity 1 or $1/(2\pi)$; the coordinates must be rescaled when comparing formulas.

For a simple point process $\Xi = \sum_i \delta_{x_i}$, its k th factorial moment measure is defined by

$$\int_{\mathbb{R}^k} f d\alpha^{(k)} = \mathbb{E} \sum_{i_1, \dots, i_k \text{ distinct}} f(x_{i_1}, \dots, x_{i_k}).$$

When this measure has a density, the density is the k -point correlation function ρ_k . For Sine $_2$,

$$\rho_k(x_1, \dots, x_k) = \det[K_{\text{sin}}(x_i, x_j)]_{i,j=1}^k, \quad K_{\text{sin}}(x, y) = \frac{\sin(x - y)}{\pi(x - y)}, \quad (2.2)$$

with diagonal value $1/\pi$. Thus Sine $_2$ is determinantal. The $\beta = 1, 4$ processes are Pfaffian. General $\beta > 0$ limits can be constructed from the tridiagonal models [2].

The Wigner–Dyson–Mehta universality principle says that the same local statistics occur for non-Gaussian Wigner matrices in the same symmetry class. Under standard independence, variance, nondegeneracy, and sufficiently strong moment assumptions, this is a theorem: the bulk limit is Sine₁ in the real case and Sine₂ in the complex case [15]. Two major approaches are moment comparison and local relaxation through Dyson Brownian motion. Later lectures discuss a loop-equation approach.

2.2 The spectral edge and Tracy–Widom laws

At the right edge,

$$\rho_{\text{sc}}(2-s) = \frac{1}{\pi} \sqrt{s} (1 + O(s)), \quad \int_{2-\Delta}^2 \rho_{\text{sc}}(x) dx = \frac{2}{3\pi} \Delta^{3/2} (1 + O(\Delta)).$$

An interval containing order one eigenvalues therefore has width of order $N^{-2/3}$. For the Gaussian beta ensemble,

$$\sum_{i=1}^N \delta_{N^{2/3}(\lambda_i - 2)} \Rightarrow \text{Airy}_\beta, \quad N^{2/3}(\lambda_1 - 2) \Rightarrow \text{TW}_\beta. \quad (2.3)$$

Here TW_β is the distribution of the largest point. The name Airy₁ refers here to the beta-one *point process*, not to the differently named temporal Airy₁ process of growth models. For $\beta = 2$, the point process is determinantal with kernel

$$K_{\text{Ai}}(x, y) = \frac{\text{Ai}(x) \text{Ai}'(y) - \text{Ai}'(x) \text{Ai}(y)}{x - y} = \int_0^\infty \text{Ai}(x+t) \text{Ai}(y+t) dt. \quad (2.4)$$

The Airy function satisfies $\text{Ai}''(x) = x \text{Ai}(x)$ and is normalized by $\text{Ai}(x) \sim (2\sqrt{\pi})^{-1} x^{-1/4} e^{-2x^{3/2}/3}$ as $x \rightarrow +\infty$. On the diagonal, $K_{\text{Ai}}(x, x) = \text{Ai}'(x)^2 - x \text{Ai}(x)^2$. Far to the left this intensity behaves as $\sqrt{-x}/\pi$.

For general $\beta > 0$ [3], let $\Lambda_0 < \Lambda_1 < \dots$ be the eigenvalues of the stochastic Airy operator on $[0, \infty)$ with Dirichlet boundary at zero,

$$\mathcal{A}_\beta = -\frac{d^2}{dx^2} + x + \frac{2}{\sqrt{\beta}} b'(x),$$

where b is standard Brownian motion and the white-noise potential is defined through the appropriate quadratic-form construction. Then $\text{Airy}_\beta = \sum_{j \geq 0} \delta_{-\Lambda_j}$. The minus sign is essential: the Airy points decrease to $-\infty$. At the level of leading logarithmic tails,

$$\log \mathbb{P}(\text{TW}_\beta > s) \sim -\frac{2\beta}{3} s^{3/2}, \quad \log \mathbb{P}(\text{TW}_\beta < -s) \sim -\frac{\beta}{24} s^3 \quad (s \rightarrow \infty).$$

Edge universality also holds for Wigner matrices under suitable tail conditions. For an iid off-diagonal entry distribution with $x_{ij} = \sqrt{N} h_{ij}$ independent of N , the sharp off-diagonal condition in the standard Wigner setting [6] is $\lim_{s \rightarrow \infty} s^4 \mathbb{P}(|x_{12}| \geq s) = 0$, with the usual diagonal assumptions. A finite fourth moment is sufficient and is stronger than this tail condition.




Figure 2: Finite-dimensional edge simulations in the normalization $x = N^{2/3}(\lambda - 2)$. Left: kernel density estimates for the largest eigenvalue from 3000 independent Gaussian β -ensemble samples of size $N = 800$, for each $\beta \in \{1, 2, 4\}$. These are finite- N empirical approximations, not evaluations of the exact Tracy–Widom densities. Right: the six largest eigenvalues in seven GUE realizations. The $\beta = 4$ curve uses N particles and the density proportional to $\exp[-N\beta \sum_i \lambda_i^2/4] \prod_{i < j} |\lambda_i - \lambda_j|^\beta$.




Figure 3: Deterministic numerical evaluations of the Airy function (left) and the Airy₂ one-point intensity $K_{\text{Ai}}(x, x) = \text{Ai}'(x)^2 - x \text{Ai}(x)^2$ (right). The dashed curve is $\pi^{-1} \sqrt{(-x)_+}$, the leading intensity asymptotic as $x \rightarrow -\infty$.

3 Examples beyond random matrices

3.1 Longest increasing subsequences and last-passage percolation

For a uniformly random permutation σ of $\{1, \dots, N\}$, define

$$L_N = \max\{k : i_1 < \dots < i_k, \sigma(i_1) < \dots < \sigma(i_k)\}.$$

The Baik–Deift–Johansson theorem [7] states

$$\frac{L_N - 2\sqrt{N}}{N^{1/6}} \Rightarrow \text{TW}_2. \quad (3.1)$$

This is one of the first striking appearances of a random matrix edge law outside matrix theory.

For independent nonnegative weights w_{ij} , directed last-passage percolation is defined by

$$G(m, n) = \max_{\pi: (1,1) \nearrow (m,n)} \sum_{(i,j) \in \pi} w_{ij},$$

where each path step increases exactly one coordinate by one and both endpoints are included. The dynamic-programming recursion is $G(i, j) = w_{ij} + \max\{G(i-1, j), G(i, j-1)\}$, with the boundary initialized by cumulative sums. For independent mean-one exponential weights [8],

$$\frac{G(n, n) - 4n}{2^{4/3}n^{1/3}} \Rightarrow \text{TW}_2. \quad (3.2)$$

Geometric weights have an analogous theorem with distribution-dependent constants. These models belong to the KPZ universality class. A general universality statement needs assumptions on the weight distribution: the claim cannot hold for arbitrary iid weights, for example deterministic or sufficiently heavy-tailed weights.

3.2 The Riemann zeta function

The Dirichlet series $\zeta(s) = \sum_{n \geq 1} n^{-s}$ initially converges for $\text{Re } s > 1$; the zeta function elsewhere is defined by analytic continuation. The Riemann hypothesis says that every nontrivial zero has real part $1/2$. Write the positive ordinates as t_j , so that, under this hypothesis, the zeros are $1/2 + it_j$. Near a large height T , their average spacing is asymptotic to $2\pi/\log(T/(2\pi))$. Thus coordinates compatible with our intensity $1/\pi$ convention are

$$x_j = \frac{\log(T/(2\pi))}{2}(t_j - T).$$

The GUE conjecture predicts Sine_2 statistics after this rescaling. A probabilistic formulation averages the center T over a long high window before taking the limit. Montgomery's pair-correlation conjecture [9] is the two-point part of this picture; full convergence of the point process is a stronger conjecture. Random matrix theory provides related predictions for other L -functions as well.




Figure 4: A uniform random permutation of 90 elements. The red polygonal line joins one longest increasing subsequence, which has length 16 in this realization. Connecting segments only emphasize the ordering of the selected permutation points.




Figure 5: Exponential last-passage percolation on a 24×24 grid. Left: independent mean-one exponential weights and a maximizing up-right path from $(1, 1)$ to $(24, 24)$. Right: the dynamic-programming values $G(i, j)$, with the same path superimposed. Both endpoint weights are included. The last-passage value is $G(24, 24) \approx 87.3$ in this realization.




Figure 6: The first forty computed nontrivial zeta zeros and their locally unfolded consecutive gaps. These are deterministic low-height evaluations, not a simulation or verification of the high-height GUE conjecture.

3.3 Billiards and quantum chaos

A classical billiard particle moves in straight lines and reflects specularly at the boundary: the normal component of velocity changes sign and the tangential component is preserved. In a circular table, successive collision points differ by a fixed angular increment. If that increment divided by 2π is rational, the orbit is periodic. If it is irrational, the collision points are dense on the boundary and the chords are dense in the annular region outside a fixed circular caustic. This is an integrable system.

A standard dispersing example is a Sinai billiard: a square table with a circular obstacle, with specular reflection at the walls and obstacle. Typical trajectories explore the available table. Ergodicity is a statement about almost every initial condition with respect to invariant phase-space measure, not about every orbit. For an observable depending only on position, its consequence is

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t f(x(s)) ds = \frac{1}{|\Omega|} \int_{\Omega} f(x) dx$$

for almost every initial condition, under the standard billiard hypotheses.

The associated quantum problem is, for example, the Dirichlet Laplacian $-\Delta\psi_k = E_k\psi_k$ on Ω , with $0 < E_1 \leq E_2 \leq \dots$. The Bohigas–Giannoni–Schmit conjecture [10] predicts random matrix local statistics for suitable classically chaotic systems after unfolding by the mean spectral density. For a time-reversal-invariant scalar billiard one expects GOE statistics, after separating spatial symmetry sectors. Without time reversal, the expected symmetry class is typically GUE. Classical chaos alone is not an unconditional theorem of random matrix statistics; additional symmetries and arithmetic structure matter.




Figure 7: Event-driven billiard trajectories computed with the specular reflection rule $v \mapsto v - 2(v \cdot n)n$. In the unit circle, the first trajectory has rotation number $2/7$ and closes after seven reflections; the second has rotation number $(\sqrt{5} - 1)/4$ and is shown for 160 reflections. The dashed circles are their circular caustics. The last panel shows 160 reflections in the square $[-1, 1]^2$ with the central disk of radius 0.30 removed. Its initial position is $(-0.78, -0.62)$ and its initial velocity is proportional to $(1, 0.37)$. The computation illustrates trajectories; it does not establish spectral universality.

4 A convergence framework for local random-matrix statistics

The next goal is to study the eigenvalues on the scale of their typical spacing. At a bulk energy $E \in (-2, 2)$, write

$$\rho_E = \rho_{\text{sc}}(E) = \frac{\sqrt{4 - E^2}}{2\pi}, \quad \sigma_E = \pi \rho_E = \frac{\sqrt{4 - E^2}}{2}.$$

If $\lambda_1 \geq \dots \geq \lambda_N$, the microscopic point measure is

$$\Xi_{N,E} = \sum_{i=1}^N \delta_{N\sigma_E(\lambda_i - E)}.$$

This convention gives limiting intensity $1/\pi$, rather than intensity one. Near a regular spectral edge the appropriate scale is $N^{-2/3}$; the corresponding limits are the Airy point processes.

The organizing idea is analogous to Stein's characterization of a Gaussian random variable:

$$Z \sim \mathcal{N}(0, 1) \iff \mathbb{E}[Zf(Z)] = \mathbb{E}[f'(Z)]$$

for every suitable test function f . An approximate identity leads to a convergence theorem only after tightness and the passage of expectations to the limit have been justified. In the present setting, the test functions are products of resolvents, and integration by parts produces a hierarchy of *loop equations*, also called Schwinger–Dyson equations.

4.1 The finite-dimensional loop equations

For the Gaussian β -ensemble with density

$$p_{N,\beta}(\lambda) = \frac{1}{Z_{N,\beta}} \prod_{i < j} |\lambda_i - \lambda_j|^\beta \exp\left(-\frac{\beta N}{4} \sum_{i=1}^N \lambda_i^2\right), \quad \beta > 0,$$

set

$$m_N(z) = \frac{1}{N} \sum_{i=1}^N \frac{1}{\lambda_i - z}, \quad z \in \mathbb{C} \setminus \mathbb{R}.$$

Integration by parts gives

$$\mathbb{E} \left[m_N(z)^2 + z m_N(z) + 1 + \frac{1}{N} \left(\frac{2}{\beta} - 1 \right) m'_N(z) \right] = 0. \quad (4.1)$$

More generally, for $p \geq 1$, put $z_1 = z$ and interpret empty products as one. Then

$$\begin{aligned} 0 = \mathbb{E} & \left[\left(m_N(z)^2 + z m_N(z) + 1 + \frac{1}{N} \left(\frac{2}{\beta} - 1 \right) m'_N(z) \right) \prod_{j=2}^p m_N(z_j) \right] \\ & + \frac{2}{\beta N^2} \sum_{j=2}^p \mathbb{E} \left[\partial_{z_j} \frac{m_N(z) - m_N(z_j)}{z - z_j} \prod_{\substack{2 \leq k \leq p \\ k \neq j}} m_N(z_k) \right]. \end{aligned} \quad (4.2)$$

Divided differences at coincident arguments are interpreted by analytic continuation. The first equation controls one resolvent; the entire hierarchy encodes joint resolvent correlations. Rescaling these identities in the bulk and at the edge leads to the two microscopic hierarchies below.

4.2 Nevanlinna functions and point measures

A Nevanlinna (or Herglotz) function is an analytic function $S : \mathbb{C}_+ \rightarrow \overline{\mathbb{C}_+}$. Its representation is

$$S(z) = c + dz + \int_{\mathbb{R}} \left(\frac{1}{x - z} - \frac{x}{1 + x^2} \right) \mu(dx), \quad (4.3)$$

where $c \in \mathbb{R}$, $d \geq 0$, and μ is a positive Borel measure satisfying

$$\int_{\mathbb{R}} \frac{\mu(dx)}{1 + x^2} < \infty.$$

We extend S to the lower half-plane by $S(\bar{z}) = \overline{S(z)}$. For a finite measure, the ordinary Stieltjes transform $\int (x - z)^{-1} \mu(dx)$ corresponds to $d = 0$ and $c = \int x(1 + x^2)^{-1} \mu(dx)$. Microscopic point configurations have *counting measures* $\sum_i \delta_{x_i}$, which are generally infinite measures; a subtraction or principal-value prescription is therefore necessary. They should not be confused with the normalized probability measure $N^{-1} \sum_i \delta_{\lambda_i}$.

The imaginary part determines the measure:

$$\operatorname{Im} S(u + i\eta) = d\eta + \int_{\mathbb{R}} \frac{\eta}{(x - u)^2 + \eta^2} \mu(dx), \quad \eta > 0,$$

and

$$\frac{1}{\pi} \operatorname{Im} S(u + i\eta) du \xrightarrow[\eta \downarrow 0]{\text{vague}} \mu(du). \quad (4.4)$$

Here vague convergence means convergence against continuous functions with compact support.

4.3 Bulk and edge loop equations

For a bulk configuration of asymptotic density $1/\pi$, the natural transform is

$$S(z) = \text{pv} \sum_i \frac{1}{x_i - z}.$$

The principal value is taken with a symmetric spatial cutoff, whenever this limit exists. Its normalization is suggested by

$$\lim_{R \rightarrow \infty} \frac{1}{\pi} \int_{-R}^R \frac{dx}{x - z} = i, \quad z \in \mathbb{C}_+.$$

The Sine_β process has intensity $1/\pi$ in this convention. For $\beta = 2$ its determinantal kernel is

$$K_{\text{sine}}(x, y) = \frac{\sin(x - y)}{\pi(x - y)}, \quad K_{\text{sine}}(x, x) = \frac{1}{\pi}.$$

Its loop hierarchy takes the form

$$\begin{aligned} 0 = \mathbb{E} & \left[\left(S(z)^2 + 1 + \left(\frac{2}{\beta} - 1 \right) S'(z) \right) \prod_{j=2}^p S(z_j) \right] \\ & + \frac{2}{\beta} \sum_{j=2}^p \mathbb{E} \left[\partial_{z_j} \frac{S(z) - S(z_j)}{z - z_j} \prod_{\substack{2 \leq k \leq p \\ k \neq j}} S(z_k) \right]. \end{aligned} \quad (4.5)$$

For an edge configuration, order the particles as $x_1 > x_2 > \dots \rightarrow -\infty$. The Airy density behaves as $\sqrt{-x}/\pi$ far to the left, so the expected location scale is

$$x_i \sim - \left(\frac{3\pi i}{2} \right)^{2/3}.$$

Let $a_1 > a_2 > \dots$ be the negative zeros of the Airy function Ai , which satisfies $\text{Ai}''(z) = z \text{Ai}(z)$. The regularized transform used in the notes is

$$S(z) = \sum_{i=1}^{\infty} \left(\frac{1}{x_i - z} - \frac{1}{a_i} \right) - \frac{\text{Ai}'(0)}{\text{Ai}(0)}. \quad (4.6)$$

This formula requires enough control of $x_i - a_i$ to ensure local uniform convergence; for example, $|x_i - a_i| = O(i^{-1/3+\varepsilon})$ for some $\varepsilon < 2/3$ is sufficient. The common leading-order asymptotic $x_i \sim a_i$ alone is not sufficient. In its intended Airy setting the transform is a Nevanlinna function whose representing measure is $\sum_i \delta_{x_i}$. The normalization can be checked on the deterministic configuration of Airy zeros, for which

$$S(z) = - \frac{\text{Ai}'(z)}{\text{Ai}(z)}.$$

The corresponding leading asymptotic is $S(z) \sim \sqrt{z}$ in appropriate sectors away from the negative real axis; it is an asymptotic statement, not an identity. The branch is chosen positive on $(0, \infty)$.

The Airy hierarchy is

$$0 = \mathbb{E} \left[\left(S(z)^2 - z + \left(\frac{2}{\beta} - 1 \right) S'(z) \right) \prod_{j=2}^p S(z_j) \right] \quad (4.7)$$

$$+ \frac{2}{\beta} \sum_{j=2}^p \mathbb{E} \left[\partial_{z_j} \frac{S(z) - S(z_j)}{z - z_j} \prod_{\substack{2 \leq k \leq p \\ k \neq j}} S(z_k) \right].$$

Characterization theorem (Bourgade–Huang). Call S *particle-generated* if its representation (4.3) has measure $\mu = \sum_{x \in \mathcal{P}} \delta_x$, where $\mathcal{P}$ is a locally finite multiset and $\sum_{x \in \mathcal{P}} (1 + x^2)^{-1} < \infty$. Thus its poles on the real axis encode the particles, with the multiplicity of a point equal to the negative of the corresponding residue in the variable z . For rational $\beta > 0$, if a random particle-generated Nevanlinna function satisfies the entire bulk hierarchy (4.5), for every $p \geq 1$ and every choice of arguments in both half-planes, with the stated expectations well defined, then its particle configuration is distributed as Sine_β . Replacing the bulk hierarchy by (4.7) gives Airy_β . These are external characterization theorems [16, Theorems 1.5 and 1.11]; the lecture concentrates on $\beta = 1, 2$. No additional rigidity assumption is needed in these exact characterizations. Both the complete hierarchy and the availability of mixed upper- and lower-half-plane arguments are essential to the formulation used here.

5 Bulk loop equations for real Wigner matrices

Let $H = (h_{ij})_{i,j=1}^N$ be real symmetric. Assume that the upper triangular entries are independent, centered, and satisfy

$$\mathbb{E} h_{ij}^2 = \frac{1}{N}, \quad \mathbb{E} |\sqrt{N} h_{ij}|^q \leq C_q \quad (q \geq 2),$$

uniformly in i, j, N . The choice of diagonal variance is inessential for the limiting argument; here it is $1/N$ as in the handwritten notes. The Gaussian orthogonal ensemble usually has diagonal variance $2/N$. Set

$$G(z) = (H - zI)^{-1}, \quad m_N(z) = \frac{1}{N} \text{Tr } G(z).$$

Fix $E \in (-2, 2)$, put $\sigma = \sigma_E$, and define

$$S_N(w) = \sum_{i=1}^N \frac{1}{N\sigma(\lambda_i - E) - w} + \frac{E}{2\sigma} = \frac{1}{\sigma} \left(m_N \left(E + \frac{w}{N\sigma} \right) + \frac{E}{2} \right). \quad (5.1)$$

The real constant $E/(2\sigma)$ removes the real part of the macroscopic semicircle transform. It does not change the associated point measure.

The desired approximate hierarchy is the following: for every fixed $p \geq 1$ and compact $K \subset \mathbb{C} \setminus \mathbb{R}$, uniformly for $w = w_1, w_2, \dots, w_p \in K$,

$$\mathbb{E} \left[\left(S'_N(w) + S_N(w)^2 + 1 \right) \prod_{j=2}^p S_N(w_j) \right] \quad (5.2)$$

$$+ 2 \sum_{j=2}^p \mathbb{E} \left[\partial_{w_j} \frac{S_N(w) - S_N(w_j)}{w - w_j} \prod_{\substack{2 \leq k \leq p \\ k \neq j}} S_N(w_k) \right] = o(1).$$

One useful intermediate version has right-hand side bounded by $\varepsilon_N(1 + \mathbb{E}R_N^{p+1})$, where $\varepsilon_N \rightarrow 0$ and $R_N = \max_j |S_N(w_j)|$. In the Wigner setting, the estimates below also give a direct $o(1)$ bound for each fixed p , by choosing the local-law exponent sufficiently small. Subsequential limits, moment bounds, and the characterization theorem then identify the bulk limit.

5.1 Local law, rigidity, and delocalization

For nonnegative random variables $X = X_N$ and $Y = Y_N$, the notation $X \prec Y$ means that for every $\varepsilon, D > 0$,

$$\mathbb{P}(X > N^\varepsilon Y) \leq N^{-D}$$

for all sufficiently large N . Uniform stochastic domination means that the same estimate holds uniformly over the stated deterministic parameters.

We use the standard Wigner local law as an input [14]. For fixed $\delta > 0$, let

$$\mathcal{D}_\delta = \{u + i\eta : |u| \leq 3, N^{-1+\delta} \leq \eta \leq 1\}.$$

Uniformly on this domain,

$$\begin{aligned} \max_{i,j} |G_{ij}(z) - \delta_{ij} m_{\text{sc}}(z)| &\prec \sqrt{\frac{\text{Im } m_{\text{sc}}(z)}{N\eta}} + \frac{1}{N\eta}, \\ |m_N(z) - m_{\text{sc}}(z)| &\prec \frac{1}{N\eta}. \end{aligned} \tag{5.3}$$

The distinction from the global semicircle law is that η is allowed to shrink almost to the eigenvalue-spacing scale. Equivalently, one can count eigenvalues in intervals whose lengths depend on N .

Define the descending classical locations by

$$\int_{\gamma_i}^2 \rho_{\text{sc}}(x) dx = \frac{i - \frac{1}{2}}{N}, \quad i = 1, \dots, N.$$

The standard rigidity consequence [13, 14] is

$$|\lambda_i - \gamma_i| \prec N^{-2/3} \min\{i, N + 1 - i\}^{-1/3}. \tag{5.4}$$

This is the size of a local classical gap, up to the arbitrarily small powers of N permitted by stochastic domination.

The local law also implies complete eigenvector delocalization. Let u_α be a normalized eigenvector with eigenvalue λ_α . On an event of overwhelming probability, $\max_{i,j} |G_{ij}(z)| \leq 2$ throughout $\mathcal{D}_\delta$. The spectral theorem gives

$$\text{Im } G_{ii}(u + i\eta) = \sum_{\alpha=1}^N \frac{\eta |u_\alpha(i)|^2}{(\lambda_\alpha - u)^2 + \eta^2}.$$

Taking $u = \lambda_\alpha$ and $\eta = N^{-1+\delta}$ yields

$$|u_\alpha(i)|^2 \leq \eta \text{Im } G_{ii}(\lambda_\alpha + i\eta) \leq 2N^{-1+\delta}.$$

Here rigidity ensures that all eigenvalues lie in $[-3, 3]$ with overwhelming probability. Since $\delta > 0$ is arbitrary, $\max_{\alpha,i} |u_\alpha(i)|^2 \prec N^{-1}$. This scale agrees with $\sum_i |u_\alpha(i)|^2 = 1$.




Figure 8: Eigenvalue rigidity illustrated by one simulated GUE spectrum with $N = 1600$. Left: descending eigenvalues and semicircle quantiles. Right: deviations divided by the deterministic scale in (5.4). This finite sample illustrates the scale; the stochastic bound is a theorem, not a numerical conclusion.




Figure 9: Two schematic diagrams from the convergence argument. Left: the local-law domain $\mathcal{D}_\delta$ and a microscopic spectral window near E ; vertical scales are not proportional. Right: a finite δ -net of a compact set K separated from the real axis by $\eta_0 > 0$, as used in the tightness argument. The disk is only an illustration of a compact set; K need not be a disk.

5.2 Extending entrywise bounds to the microscopic scale

The spectral parameter in (5.1) has imaginary part of order N^{-1} , slightly below $\mathcal{D}_\delta$. The following deterministic argument bridges that gap. On the local-law event, for $|u| \leq 3$ and $\eta \neq 0$,

$$\max_{a,b} |G_{ab}(u + i\eta)| \leq 2 \max \left\{ 1, \frac{N^\delta}{N|\eta|} \right\}. \quad (5.5)$$

To prove this for $\eta > 0$, write $M(s) = \max_{a,b} |G_{ab}(u + is)|$. For $N^{-1+\delta} \leq s \leq 1$ the local law gives $M(s) \leq 2$, while for $s \geq 1$ the operator norm bound gives $M(s) \leq 1/s \leq 1$. Furthermore,

$$\partial_s G(u + is) = iG(u + is)^2.$$

Using Cauchy–Schwarz and the Ward identity,

$$\begin{aligned} |\partial_s G_{ab}(u + is)| &\leq \left(\sum_k |G_{ak}|^2 \sum_k |G_{kb}|^2 \right)^{1/2} \\ &= \frac{\sqrt{\operatorname{Im} G_{aa} \operatorname{Im} G_{bb}}}{s} \leq \frac{M(s)}{s}. \end{aligned}$$

The function M is locally Lipschitz and satisfies $|M'(s)| \leq M(s)/s$ almost everywhere. Integrating between s and $s_0 = N^{-1+\delta}$ gives $M(s) \leq M(s_0)s_0/s \leq 2N^\delta/(Ns)$ for $0 < s \leq s_0$. The lower half-plane follows by conjugation. Consequently, for fixed $K \subset \mathbb{C} \setminus \mathbb{R}$,

$$\max_{a,b} \left| G_{ab} \left(E + \frac{w}{N\sigma} \right) \right| \leq C_K N^\delta, \quad w \in K, \quad (5.6)$$

with overwhelming probability.

5.3 An approximate integration-by-parts formula

Let h be a centered real random variable with $\mathbb{E}h^2 = N^{-1}$. For a twice continuously differentiable function f ,

$$\mathbb{E}[hf(h)] = \frac{1}{N} \mathbb{E}[f'(h)] + R, \quad |R| \leq \mathbb{E} \left[\left(|h|^3 + \frac{|h|}{N} \right) \sup_{|x| \leq |h|} |f''(x)| \right], \quad (5.7)$$

provided the displayed expectations are finite. Indeed, Taylor's formula gives

$$f(h) = f(0) + hf'(0) + h^2 \int_0^1 (1-t) f''(th) dt, \quad f'(h) = f'(0) + h \int_0^1 f''(th) dt.$$

Taking expectations after multiplying the first equation by h and subtracting N^{-1} times the second gives (5.7). For a Wigner entry, the coefficient of the supremum has expectation $O(N^{-3/2})$.

There is one technical point when applying this formula to a resolvent: $f''(x)$ is evaluated after one entry has been changed from h to x . The bound (5.6) must therefore hold for these modified matrices as well. This follows from the resolvent identity. On the event $|h_{ij}| \leq N^{-1/2+\varepsilon}$, changing one independent entry by at most $2N^{-1/2+\varepsilon}$ is a perturbation supported on at most two coordinates. The corresponding two-dimensional inverse has the form $(I + VG_{\{i,j\},\{i,j\}})^{-1}$, which is bounded when $\varepsilon + \delta < 1/2$. The entrywise bound remains $O_K(N^\delta)$. Moment assumptions make the complement negligible; there one uses the deterministic resolvent norm bound $\|G(z)\| \leq |\operatorname{Im} z|^{-1}$. Thus the Taylor remainder is controlled uniformly along the entire entry interpolation.

5.4 Derivation of the first microscopic loop equation

Taking the trace of $I = G(H - zI)$ and dividing by N gives

$$1 + zm_N(z) = \frac{1}{N} \sum_{i,j} h_{ij} G_{ji}(z).$$

When $i \neq j$, differentiation with respect to the independent entry $h_{ij} = h_{ji}$ changes both symmetric positions, so

$$\partial_{h_{ij}} G_{ji} = -G_{jj} G_{ii} - G_{ji}^2, \quad \partial_{h_{ii}} G_{ii} = -G_{ii}^2.$$

Using (5.7), summing over the ordered pairs (i, j) , and using (5.6), we obtain

$$\begin{aligned} \mathbb{E}[1 + zm_N] &= \frac{1}{N^2} \sum_{i,j} \mathbb{E}[\partial_{h_{ij}} G_{ji}] + O_K(N^{-1/2+3\delta}) + o(1) \\ &= -\mathbb{E}\left[m_N^2 + \frac{1}{N^2} \text{Tr} G^2\right] + \frac{1}{N^2} \sum_i \mathbb{E} G_{ii}^2 + o(1) \\ &= -\mathbb{E}\left[m_N^2 + \frac{1}{N} m'_N(z)\right] + o(1). \end{aligned}$$

In the first line, the notation $\partial_{h_{ij}}$ for $i > j$ denotes differentiation with respect to the same independent symmetric entry. The second derivative of a resolvent entry is a sum of products of three resolvent entries. There are N^2 terms and an overall factor $1/N$, explaining the error $N^{-1/2+3\delta}$. The diagonal correction is $O_K(N^{-1+2\delta})$. Choosing $\delta < 1/6$ makes both errors vanish. Thus

$$\mathbb{E}\left[1 + zm_N(z) + m_N(z)^2 + \frac{1}{N} m'_N(z)\right] = o(1). \quad (5.8)$$

Now put $z = E + w/(N\sigma)$. From $m_N(z) = \sigma S_N(w) - E/2$ and $N^{-1} m'_N(z) = \sigma^2 S'_N(w)$, we find

$$1 + zm_N + m_N^2 + \frac{1}{N} m'_N = \sigma^2 (S_N^2 + 1 + S'_N) + \frac{wm_N}{N\sigma},$$

because $\sigma^2 = 1 - E^2/4$. The last term is $o(1)$ in expectation by (5.6). Dividing by σ^2 proves

$$\mathbb{E}[S'_N(w) + S_N(w)^2 + 1] = o(1). \quad (5.9)$$

5.5 How the higher equations arise

The handwritten notes indicate the extension by multiplying the trace identity by additional resolvents. We record the calculation, since it explains why the full hierarchy is necessary. Multiply by $P = \prod_{j=2}^p S_N(w_j)$ before applying integration by parts. The terms where the derivative falls on $G_{ba}(z)$ reproduce the first line of (5.2). For an off-diagonal entry,

$$\partial_{h_{ab}} m_N(z_j) = -\frac{2}{N} (G(z_j)^2)_{ab}, \quad a \neq b,$$

whereas the coefficient is $-1/N$ on the diagonal. The main sum is therefore expressed through

$$\sum_{a,b} G_{ba}(z) (G(z_j)^2)_{ab} = \text{Tr}(G(z)G(z_j)^2) = N \partial_{z_j} \frac{m_N(z) - m_N(z_j)}{z - z_j}.$$

Here we used $G(z)G(v) = (G(z) - G(v))/(z - v)$. Since $\partial_{h_{ab}} S_N(w_j) = \sigma^{-1} \partial_{h_{ab}} m_N(z_j)$, the extra terms before rescaling have coefficient $2/(N^2 \sigma)$. The substitution $z_j = E + w_j/(N\sigma)$ then gives exactly the coefficient 2 in (5.2) after division by σ^2 .

For completeness, all remainder estimates remain small for fixed p . Derivatives of a normalized trace obey

$$|\partial_{h_{ab}} m_N(z)| \leq C_K N^\delta, \quad |\partial_{h_{ab}}^2 m_N(z)| \leq C_K N^{2\delta}$$

on the event above. To see this, use

$$|(G^2)_{ab}| \leq \frac{\sqrt{|\operatorname{Im} G_{aa}| |\operatorname{Im} G_{bb}|}}{|\operatorname{Im} z|} \leq C_K N^{1+\delta}$$

and the resolvent derivative formula; normalization by N cancels the factor N . Consequently the second derivative of $G_{ba}P$ is bounded by $O_{p,K}(N^{(p+2)\delta})$. The summed Taylor remainder is

$$O_{p,K}(N^{-1/2+(p+2)\delta}).$$

Diagonal corrections and the term $(z - E)m_N P$ are bounded by negative powers of N after choosing $\delta > 0$ sufficiently small depending on p . The bad event is again negligible by the moment assumptions and deterministic resolvent bounds. This proves the fixed-order approximate hierarchy (5.2).

6 Moment bounds and tightness of microscopic transforms

We next explain the compactness step in the convergence framework. Two elementary facts are useful.

First, locally uniform convergence of Nevanlinna functions implies vague convergence of their representing measures. Thus, if $S_n \rightarrow S$ locally uniformly in $\mathbb{C}_+$ and μ_n, μ are the measures in (4.3), then $\mu_n \rightarrow \mu$ vaguely. One way to prove this is to use (4.4) against smooth compactly supported test functions, first at fixed height and then let the height tend to zero. The bound on $\operatorname{Im} S_n(i)$ supplies the uniform local mass and weighted tail bounds needed for this argument. If the μ_n are locally finite integer-valued point measures, so is μ (possibly with multiplicities). Simplicity is not automatic. Neither μ_n nor μ is required to be a probability measure.

Second, from the Nevanlinna representation, for $z = u + i\eta$ with $\eta > 0$,

$$|S'(z)| \leq d + \int \frac{\mu(dx)}{|x - z|^2} = \frac{\operatorname{Im} S(z)}{\eta} \leq \frac{|S(z)|}{\eta}. \quad (6.1)$$

The same bound with $|\eta|$ holds in the lower half-plane. Also, for $z, v \in \mathbb{C} \setminus \mathbb{R}$,

$$\partial_v \frac{S(z) - S(v)}{z - v} = \int_{\mathbb{R}} \frac{\mu(dx)}{(x - z)(x - v)^2}. \quad (6.2)$$

When $v = z$ or $v = \bar{z}$, its absolute value is at most $|S(z)|/|\operatorname{Im} z|^2$.

6.1 Pointwise moment estimates

Fix an integer $q \geq 2$ and $w \in \mathbb{C}_+$. In (5.2), take $p = 2q - 1$, with $q - 2$ of the arguments $w_2, \dots, w_p$ equal to w and q equal to $\bar{w}$. Then

$$S_N(w)^2 \prod_{j=2}^p S_N(w_j) = |S_N(w)|^{2q}.$$

The derivative estimates (6.1) and (6.2) give, uniformly when w stays in a fixed compact subset of the upper half-plane,

$$\mathbb{E}|S_N(w)|^{2q} \leq C_{q,K} \mathbb{E} \left[|S_N(w)|^{2q-1} + |S_N(w)|^{2q-2} \right] + o(1).$$

Young's inequality absorbs the lower powers into the left-hand side: for every fixed C ,

$$Ct^{2q-1} + Ct^{2q-2} \leq \frac{1}{2}t^{2q} + C_q, \quad t \geq 0.$$

Therefore

$$\sup_N \sup_{w \in K} \mathbb{E}|S_N(w)|^{2q} < \infty. \quad (6.3)$$

The finitely many small values of N can be absorbed into the constant. The case $q = 1$ follows from $q = 2$ by Cauchy-Schwarz, and conjugation handles the lower half-plane. If one uses instead the relative error $\varepsilon_N(1 + \mathbb{E}R_N^{p+1})$, then in this choice of arguments $R_N = |S_N(w)|$ and $p + 1 = 2q$, so that term is also absorbed for large N .

6.2 Uniform bounds on compact sets

Let $K \subset \mathbb{C} \setminus \mathbb{R}$ be compact and put $\eta_0 = \text{dist}(K, \mathbb{R}) > 0$. Choose a finite δ -net $\{w_1, \dots, w_M\} \subset K$ with $0 < \delta < \eta_0/2$. If $|z - w_j| \leq \delta$, the straight line segment joining them stays at distance at least $\eta_0/2$ from the real axis. Applying (6.1) along that segment and using Grönwall's inequality gives

$$|S_N(z)| \leq \exp\left(\frac{2\delta}{\eta_0}\right) |S_N(w_j)|.$$

Consequently,

$$\sup_{z \in K} |S_N(z)|^{2q} \leq e^{4q\delta/\eta_0} \sum_{j=1}^M |S_N(w_j)|^{2q},$$

and (6.3) implies

$$\sup_N \mathbb{E} \left[\sup_{z \in K} |S_N(z)|^{2q} \right] < \infty. \quad (6.4)$$

In particular, Markov's inequality gives compact-uniform bounds in probability. Montel's theorem, applied to an exhaustion by compact sets, now yields tightness in the topology of locally uniform convergence on $\mathbb{C} \setminus \mathbb{R}$. This corrects the handwritten net estimate: a difference bound must be integrated along a path, and cannot be replaced by a bound having the small prefactor δ/η_0 multiplying the function itself.

6.3 Passing to the exact bulk loop equations

We now complete the compactness argument. Continue with the centered bulk transform S_N and microscopic counting measure $\Xi_{N,E}$ defined above. In this subsection write $\mu_N = \Xi_{N,E}$. For example, at the center $E = 0$, these reduce to

$$\mu_N = \sum_{i=1}^N \delta_{N\lambda_i}, \quad S_N(w) = m_N(w/N),$$

since $\pi\rho_{\text{sc}}(0) = 1$. The normalization gives limiting intensity $1/\pi$, at every bulk energy.

Assume that the preceding tightness and moment estimates hold. From any subsequence we can extract a further subsequence along which the analytic functions S_N converge in distribution, locally uniformly on $\mathbb{C} \setminus \mathbb{R}$. On a Skorokhod coupling we may take this convergence to hold almost surely. The corresponding particle measures converge vaguely to the pole measure of a particle-generated limiting function S . This identification uses the compactness statement for particle-generated functions; local uniform convergence by itself should not be substituted for that statement.

For clarity, consider first $\beta = 1$. The limiting identity is

$$\begin{aligned} & \mathbb{E} \left[(S'(w) + S(w)^2 + 1) \prod_{j=2}^p S(w_j) \right] \\ & + 2 \sum_{j=2}^p \mathbb{E} \left[\partial_{w_j} \frac{S(w) - S(w_j)}{w - w_j} \prod_{\substack{k=2 \\ k \neq j}}^p S(w_k) \right] = 0. \end{aligned} \quad (6.5)$$

The variables may lie in either half-plane, and coincident variables are interpreted by analytic continuation. In particular, $(S(w) - S(v))/(w - v)$ extends to $S'(w)$ at $v = w$.

Here is the justification for passing to the limit. Cauchy's integral formula converts local uniform convergence into local uniform convergence of derivatives and divided differences. Uniform moment bounds on every compact set imply uniform integrability of the products appearing in (6.5). For example, for a finite Stieltjes sum and $\eta = |\operatorname{Im} w|$,

$$|S'_N(w)| \leq \frac{|\operatorname{Im} S_N(w)|}{\eta} \leq \frac{|S_N(w)|}{\eta}.$$

Also, writing $\eta_j = |\operatorname{Im} w_j|$,

$$\left| \partial_{w_j} \frac{S_N(w) - S_N(w_j)}{w - w_j} \right| \leq \frac{|\operatorname{Im} S_N(w_j)|}{\eta \eta_j}.$$

Alternatively, the needed bounds follow directly from Cauchy's formula on slightly larger compact sets. Thus almost sure convergence of the integrands, together with uniform integrability, gives convergence of expectations. For an approximate finite- N identity, its residual must also vanish in precisely the moment-normalized sense required by the characterization theorem.

The bulk characterization theorem [16] then identifies the pole measure of S as the Sine₁ point process. Every subsequence has a further subsequence with this same limit, and consequently the whole sequence μ_N converges in distribution to Sine₁.

6.4 The bulk hierarchy for general β

The scaling calculation is especially transparent at $E = 0$. Since $S_N(w) = m_N(w/N)$, we have

$$\partial_z m_N(z) = N S'_N(w), \quad z = w/N,$$

and

$$\partial_{z_j} \frac{m_N(z) - m_N(z_j)}{z - z_j} = N^2 \partial_{w_j} \frac{S_N(w) - S_N(w_j)}{w - w_j}.$$

The Gaussian β -ensemble loop equation therefore becomes

$$0 = \mathbb{E} \left[\left(S_N(w)^2 + 1 + \left(\frac{2}{\beta} - 1 \right) S'_N(w) + \frac{w}{N} S_N(w) \right) \prod_{j=2}^p S_N(w_j) \right] \quad (6.6)$$

$$+ \frac{2}{\beta} \sum_{j=2}^p \mathbb{E} \left[\partial_{w_j} \frac{S_N(w) - S_N(w_j)}{w - w_j} \prod_{\substack{k=2 \\ k \neq j}}^p S_N(w_k) \right].$$

Under the compactness and moment assumptions, the $w S_N(w)/N$ term vanishes. The limiting hierarchy is

$$0 = \mathbb{E} \left[\left(S(w)^2 + 1 + \left(\frac{2}{\beta} - 1 \right) S'(w) \right) \prod_{j=2}^p S(w_j) \right] \quad (6.7)$$

$$+ \frac{2}{\beta} \sum_{j=2}^p \mathbb{E} \left[\partial_{w_j} \frac{S(w) - S(w_j)}{w - w_j} \prod_{\substack{k=2 \\ k \neq j}}^p S(w_k) \right].$$

The Sine $_{\beta}$ process satisfies these identities. In the characterization result used in these notes, uniqueness is available for positive rational β , within the stated particle-generated class and with the full hierarchy, including mixed-half-plane spectral parameters. The first equation alone is not a characterization.

6.5 Recovering correlation functions from the resolvent

Let $\mu = \sum_i \delta_{x_i}$ be a simple point process. Its k -point correlation function is defined, when it exists, by

$$\mathbb{E} \sum_{i_1, \dots, i_k \text{ distinct}} f(x_{i_1}, \dots, x_{i_k}) = \int_{\mathbb{R}^k} f(t_1, \dots, t_k) \rho_k(t_1, \dots, t_k) dt_1 \cdots dt_k.$$

Thus ρ_k describes the joint intensity of distinct particles, rather than the joint density of arbitrarily chosen particle labels. For the Poisson kernel

$$P_{\eta}(t) = \frac{1}{\pi} \frac{\eta}{t^2 + \eta^2},$$

the imaginary part of the resolvent gives

$$\frac{1}{\pi} \operatorname{Im} S(x + i\eta) = \sum_i P_{\eta}(x - x_i).$$

A real subtraction used to define a renormalized Stieltjes transform does not affect this formula. In particular,

$$\mathbb{E} \left[\frac{1}{\pi} \operatorname{Im} S(x + i\eta) \right] = (P_{\eta} * \rho_1)(x).$$

As $\eta \downarrow 0$, this converges to $\rho_1(x)$ at continuity points, under the usual integrability assumption for the convolution.

For higher correlations the exact smoothed identity is

$$\mathbb{E} \sum_{i_1, \dots, i_k} \prod_{\text{distinct } j=1}^k P_{\eta_j}(t_j - x_{i_j}) = \int_{\mathbb{R}^k} \rho_k(y_1, \dots, y_k) \prod_{j=1}^k P_{\eta_j}(t_j - y_j) d\mathbf{y}.$$

The commonly used notation

$$\rho_k(t_1, \dots, t_k) = \mathbb{E} \prod_{j=1}^k \frac{1}{\pi} \operatorname{Im} S(t_j + i0) \quad (t_1, \dots, t_k \text{ distinct})$$

is a boundary-value shorthand. It is understood after taking the appropriate limit, for example distributionally away from the diagonals $t_i = t_j$, with sufficient local moment control. An ordinary product at positive η also contains repeated-particle terms, so these terms must not be silently identified with a factorial correlation function.

6.6 Why edge loop equations require a smaller error

At the right semicircle edge, use the $N^{-2/3}$ scale:

$$z = 2 + \frac{w}{N^{2/3}}, \quad S_N^{\text{edge}}(w) = N^{1/3}(m_N(z) + 1) = \sum_{i=1}^N \frac{1}{N^{2/3}(\lambda_i - 2) - w} + N^{1/3}.$$

The additive term compensates for $m_{\text{sc}}(2) = -1$. It is part of the normalization, not an additional particle. Write $\gamma_\beta = 2/\beta - 1$ and suppose the first unscaled loop equation has residual

$$r_N(z) = \mathbb{E} \left[1 + z m_N(z) + m_N(z)^2 + \frac{\gamma_\beta}{N} m'_N(z) \right].$$

A direct substitution, using

$$m_N(z) = -1 + N^{-1/3} S_N^{\text{edge}}(w), \quad m'_N(z) = N^{1/3} (S_N^{\text{edge}})'(w),$$

gives the exact identity

$$N^{2/3} r_N(z) = \mathbb{E} \left[(S_N^{\text{edge}}(w))^2 - w + \gamma_\beta (S_N^{\text{edge}})'(w) + N^{-1/3} w S_N^{\text{edge}}(w) \right]. \quad (6.8)$$

Thus, assuming the required moment control, $r_N(z) = o(N^{-2/3})$ yields the first edge loop equation

$$\mathbb{E} \left[S(w)^2 - w + \gamma_\beta S'(w) \right] = 0.$$

In particular, for $\beta = 1$ this is $\mathbb{E}[S(w)^2 - w + S'(w)] = 0$. An $O(N^{-2/3})$ residual would only give a bounded rescaled residual and would not suffice for this conclusion. The Gaussian ensemble has an exact loop equation, so its r_N is zero.

The full edge hierarchy is obtained by including products of resolvents, just as in the bulk. Subject to the precise compactness, moment, and approximate-equation hypotheses, the edge characterization theorem for positive rational β identifies the limiting pole process as Airy_β . This paragraph explains the scaling of the required error; it does not supply the additional model-specific estimates needed to verify the full hierarchy.

7 Random regular graphs and their local geometry

7.1 Two random graph models

The Erdős–Rényi graph $G(N, p_N)$ has vertex set $\{1, \dots, N\}$, and each unordered pair of distinct vertices is an edge independently with probability p_N . Its adjacency matrix A satisfies

$$A_{ij} = A_{ji} = \mathbf{1}_{\{i,j\} \text{ is an edge}}, \quad A_{ii} = 0.$$

The upper-triangular entries are independent Bernoulli random variables. After centering and normalizing they form a Wigner-type model, although sparsity creates new issues as p_N tends to zero.

A uniform simple d -regular graph is chosen uniformly from the simple graphs on the same labeled vertex set in which every vertex has degree d . Its adjacency matrix has exactly d ones in each row. These constraints destroy entrywise independence and introduce correlations throughout the matrix. We take $d \geq 3$ fixed and Nd even unless otherwise indicated.

Spectral information is useful for both models. Outlying eigenvalues and their eigenvectors can reveal large-scale structure, such as communities in variants of the Erdős–Rényi model. Random regular graphs provide a rich family of sparse expanders and also serve as models for spectral questions motivated by quantum chaos.

An expander is a sparse graph with strong connectivity. For a d -regular graph this can be expressed by a lower bound on an edge isoperimetric constant, or equivalently up to constants by a spectral gap. If A has eigenvalues $d = \theta_1 \geq \theta_2 \geq \dots$, a one-sided spectral expansion condition is

$$d - \theta_2 \geq \varepsilon > 0$$

uniformly along a sequence of graphs. A bound on the absolute values of all nontrivial eigenvalues is a stronger, two-sided condition and also excludes bipartite behavior.

7.2 The configuration model and switchings

To generate a regular multigraph, attach d labeled half-edges to each vertex and choose a uniform perfect matching of the Nd half-edges. Each matched pair is an edge. The number of matchings is

$$(Nd - 1)!! = (Nd - 1)(Nd - 3) \cdots 1.$$

Loops and multiple edges are possible. For fixed d , the standard configuration-model asymptotic [17] is

$$\mathbb{P}(\text{the resulting graph is simple}) = \exp\left(-\frac{d^2 - 1}{4}\right) (1 + o(1)).$$

Conditioned on simplicity, the graph is uniform among simple regular graphs: every such graph is represented by exactly $(d!)^N$ matchings. The displayed probability is a standard input here; it is not proved by the construction alone.

Another useful operation is a simple switching. Choose two oriented edges (u, v) and (x, y) and replace them by (u, y) and (x, v) whenever this creates a simple graph. Otherwise leave the graph unchanged. Every vertex keeps the same degree. A symmetric choice of proposals gives detailed balance with respect to the uniform measure. Adding a holding probability gives a lazy Markov chain; together with irreducibility, this implies convergence to the uniform regular graph law. Quantitative polynomial mixing bounds require an additional theorem and do not follow merely from stationarity.




Figure 10: Mathematical schematics: a 4-regular graph and the configuration model with three half-edges per vertex. Pairing may create loops or multiple edges; conditioning on simplicity gives the uniform simple model.

7.3 The Kesten–McKay law

Normalize the adjacency matrix by

$$H = \frac{A}{\sqrt{d-1}}, \quad \lambda_1 = \frac{d}{\sqrt{d-1}} \geq \lambda_2 \geq \dots \geq \lambda_N.$$

The deterministic eigenvalue λ_1 corresponds to the constant vector. The empirical spectral measure converges to the Kesten–McKay law:

$$\frac{1}{N} \sum_{i=1}^N \delta_{\lambda_i} \implies \rho_d(x) dx,$$

where

$$\rho_d(x) = \frac{d(d-1)}{d^2 - (d-1)x^2} \frac{\sqrt{4-x^2}}{2\pi} \mathbf{1}_{\{|x| \leq 2\}}. \quad (7.1)$$

The factor in front of the semicircle density tends to one as $d \rightarrow \infty$. The top eigenvalue lies outside $[-2, 2]$ for $d \geq 3$, but its mass in the empirical measure is only $1/N$. The global law alone does not locate every nontrivial eigenvalue: controlling the spectral gap or the extreme eigenvalues requires further estimates.

7.4 Tree neighborhoods

The local geometry explains the appearance of (7.1). Let $B_r(v)$ denote the induced ball of graph radius r around v . Fix $0 < c < 1$ and put

$$r = \left\lfloor \frac{c}{4} \log_{d-1} N \right\rfloor.$$

With probability tending to one, the following hold:

1. All but at most N^c vertices have a tree neighborhood $B_r(v)$.




Figure 11: Twelve independent uniform simple cubic graphs on $N = 1000$ vertices, normalized by $\sqrt{2}$. The histogram of nontrivial eigenvalues is compared with the Kesten–McKay density; the trivial eigenvalue is marked separately.

2. Every radius- r neighborhood has at most one cycle.

Here a connected graph has at most one cycle if its excess $|E| - |V| + 1$ is at most one.

These claims follow by exploring the configuration model. A radius- r exploration reveals $O((d-1)^r)$ half-edges. A collision occurs when the next matching partner belongs to the part already explored. The probability of at least one collision at a fixed root is $O((d-1)^{2r}/N) = O(N^{c/2-1})$. Hence the expected number of roots with a non-tree ball is $O(N^{c/2})$, and Markov's inequality gives the first claim. Two collisions at a fixed root have probability $O((d-1)^{4r}/N^2) = O(N^{c-2})$. A union bound over roots gives the second claim, since $c < 1$. Conditioning on simplicity preserves these high-probability statements because its probability stays bounded away from zero.




Figure 12: Actual radius-three neighborhoods in the first simulated cubic graph. One is a tree and one is unicyclic; the extra edge is highlighted. These are sampled neighborhoods, not abstract tree drawings.

8 Tree Green functions and the global spectral law

8.1 The regular tree and the cavity tree

Let $\mathcal{X}_d$ be the infinite d -regular tree. Let $\mathcal{Y}_d$ be the rooted tree obtained by cutting an edge of $\mathcal{X}_d$ and taking one component: its root has $d - 1$ children, and every other vertex also has $d - 1$ children in addition to its parent. On both trees each edge has weight $1/\sqrt{d - 1}$. Their adjacency operators are bounded and self-adjoint, so the resolvent and Schur-complement identities apply for $z \in \mathbb{C} \setminus \mathbb{R}$.

Write $h(z)$ for the root diagonal resolvent on $\mathcal{Y}_d$. Removing the root leaves $d - 1$ disjoint copies of the same rooted tree, and the Schur complement gives

$$h(z) = \frac{1}{-z - (d - 1) \frac{1}{d-1} h(z)} = \frac{1}{-z - h(z)}.$$

The solution with positive imaginary part for $z \in \mathbb{C}_+$ and $h(z) \sim -1/z$ at infinity is

$$h(z) = m_{\text{sc}}(z) = \frac{-z + \sqrt{z^2 - 4}}{2}.$$

The square-root branch is chosen by $\sqrt{z^2 - 4} \sim z$ at infinity.

The root of $\mathcal{X}_d$ has d branches. Consequently its diagonal Green function is

$$m_d(z) = \frac{1}{-z - \frac{d}{d-1} m_{\text{sc}}(z)}. \quad (8.1)$$

Taking $\pi^{-1} \text{Im } m_d(x + i0)$ recovers (7.1).




Figure 13: Deterministic tree schematics. A full root has d branches, a cavity root has $d - 1$, and a finite ball acquires boundary self-energies when its exterior is eliminated by a Schur complement.

The off-diagonal entries are also explicit:

$$G_{uv}^{\mathcal{X}_d}(z) = m_d(z) \left(-\frac{m_{\text{sc}}(z)}{\sqrt{d-1}} \right)^{\text{dist}(u,v)}, \quad (8.2)$$

$$G_{ou}^{\mathcal{Y}_d}(z) = m_{\text{sc}}(z) \left(-\frac{m_{\text{sc}}(z)}{\sqrt{d-1}} \right)^{\text{dist}(o,u)}. \quad (8.3)$$

Indeed, at each step along the unique path, the off-diagonal Schur-complement formula contributes a factor $-m_{\text{sc}}(z)/\sqrt{d-1}$. Iterating gives the displayed formulas. They exhibit exponential decay in graph distance.

8.2 Why one vertex is too small a neighborhood

For a finite regular graph, remove a vertex o and let $G^{(o)}(z)$ be the resolvent of the resulting minor. The exact identity is

$$G_{oo}(z) = \frac{1}{-z - \frac{1}{d-1} \sum_{u \sim o} \sum_{v \sim o} G_{uv}^{(o)}(z)}.$$

For fixed d , the denominator contains only finitely many neighbors, so there is no averaging over a growing number of terms at this stage. This is why the usual large-degree concentration heuristic does not immediately yield a local law. We instead take the Schur complement of a larger, tree-like neighborhood.

8.3 The Schur complement of a ball

Suppose $T = B_r(o)$ is a tree. Partition the vertices into T and its complement, and write

$$H = \begin{pmatrix} H_T & B^* \\ B & D \end{pmatrix}, \quad G^{(T)}(z) = (D - z)^{-1}.$$

Here B contains the normalized adjacency entries crossing the boundary. The compression of the full resolvent to T is

$$G_T(z) = (H_T - z - B^* G^{(T)}(z) B)^{-1}. \quad (8.4)$$

The compression G_T is not the resolvent $(H_T - z)^{-1}$ of the induced subgraph.

Let ∂T be the vertices at distance r from o , and let $I_{\partial T}$ be the diagonal projection onto these vertices. If the exterior consisted of disjoint infinite cavity trees, each of the $d - 1$ external edges at a boundary vertex would contribute $m_{\text{sc}}(z)/(d - 1)$ to the diagonal self-energy. The total would be $m_{\text{sc}}(z)$, with no off-diagonal terms. Thus the comparison matrix is

$$P(z) = (H_T - z - m_{\text{sc}}(z)I_{\partial T})^{-1}.$$

It is exactly the compression to T of the infinite-tree resolvent, so

$$P_{uv}(z) = m_d(z) \left(-\frac{m_{\text{sc}}(z)}{\sqrt{d-1}} \right)^{\text{dist}(u,v)}.$$

For the actual finite graph define the boundary error

$$\Delta = B^* G^{(T)} B - m_{\text{sc}} I_{\partial T}, \quad G_T = (P^{-1} - \Delta)^{-1}.$$

Its entries are

$$\Delta_{uv} = \frac{1}{d-1} \sum_{\substack{a \notin T, a \sim u \\ b \notin T, b \sim v}} G_{ab}^{(T)} - m_{\text{sc}} \mathbf{1}_{\{u=v \in \partial T\}}.$$

In particular, Δ is supported on $\partial T \times \partial T$.

8.4 A rigorous macroscopic comparison

The formal expansion $G_T = P + P\Delta P + \dots$ is useful intuition, but its Neumann series need not converge. An exact two-step resolvent identity avoids that issue:

$$G_T - P = P\Delta P + P\Delta G_T \Delta P. \quad (8.5)$$

For $z = E + i\eta$ with $\eta > 0$, we have $\|G_T\| \leq \eta^{-1}$ and

$$\|\Delta\| \leq \frac{\|B\|^2}{\eta} + |m_{\text{sc}}(z)| \leq \frac{d^2}{(d-1)\eta} + 1.$$

The last bound uses $\|B\| \leq \|H\| \leq d/\sqrt{d-1}$. Consequently,

$$|G_{oo} - m_d| \leq \left(\|\Delta\| + \frac{\|\Delta\|^2}{\eta} \right) \sum_{u \in \partial T} |P_{ou}|^2. \quad (8.6)$$

Since $|\partial T| = d(d-1)^{r-1}$ for $r \geq 1$,

$$\sum_{u \in \partial T} |P_{ou}|^2 = \frac{d}{d-1} |m_d(z)|^2 |m_{\text{sc}}(z)|^{2r}.$$

For every fixed $z \in \mathbb{C}_+$, $|m_{\text{sc}}(z)| < 1$, so

$$|G_{oo}(z) - m_d(z)| \leq C_{d,z} |m_{\text{sc}}(z)|^{2r} = o(1) \quad \text{as } r \rightarrow \infty.$$

The same argument handles entries G_{uv} when u, v stay in a fixed smaller ball around o : their distance from the boundary still tends to infinity.

Average the diagonal estimate over roots with tree neighborhoods. On the high-probability event above there are at most N^c remaining roots, and $|G_{oo}(z)| \leq \eta^{-1}$ for all of them. Therefore

$$|m_N(z) - m_d(z)| \leq C_{d,z} |m_{sc}(z)|^{2r} + N^{c-1}(\eta^{-1} + |m_d(z)|) \rightarrow 0$$

in probability, for each fixed $z \in \mathbb{C}_+$. The Stieltjes continuity theorem proves the global Kesten–McKay law. Because r is a positive multiple of $\log N$, this argument even gives a polynomial error at a fixed spectral parameter, with an exponent that depends on that parameter. It is not a proof of a local law down to a microscopic imaginary part.

8.5 The role of local resampling

For a sharper result one must use the randomness of the exterior self-energy. Heuristically, a typical exterior endpoint has $G_{aa}^{(T)} \approx m_{sc}$, while resolvent entries between well-separated exterior endpoints should be small. Neither assertion follows just by removing the neighborhood; they are targets of the local resampling argument.

Enumerate the edges leaving T as

$$(\ell_\alpha, a_\alpha), \quad 1 \leq \alpha \leq \mu, \quad \mu = d(d-1)^r,$$

where $\ell_\alpha \in \partial T$ and $a_\alpha \notin T$. For each such edge, choose a distant oriented edge (b_α, c_α) in the exterior, and perform the switching

$$(\ell_\alpha, a_\alpha), (b_\alpha, c_\alpha) \mapsto (\ell_\alpha, c_\alpha), (a_\alpha, b_\alpha).$$

The internal tree T is preserved and the exterior endpoints attached to its boundary are randomized. Admissibility conditions are needed to avoid loops, multiple edges, and interference between selected switchings. The next part of the notes develops how this operation can be used to estimate the boundary error and derive a local law.

9 Expansion, spectral gaps, and Ramanujan graphs

In this section we first work with the *unnormalized* adjacency matrix A of a connected d -regular graph $\mathcal{G} = (V, E)$, where $|V| = N$. Write its eigenvalues as

$$d = \mu_1 > \mu_2 \geq \cdots \geq \mu_N.$$

The constant vector is an eigenvector with eigenvalue d . The gap $d - \mu_2$ measures how difficult it is to separate the graph into two large pieces.

9.1 The Cheeger constant

For $S \subset V$, let $E(S, S^c)$ be the set of edges with one endpoint in S and one endpoint outside S . Define

$$h(\mathcal{G}) := \min_{\emptyset \neq S \subsetneq V} \frac{|E(S, S^c)|}{\min\{|S|, |S^c|\}} = \min_{0 < |S| \leq N/2} \frac{|E(S, S^c)|}{|S|}.$$

A connected finite graph has $h(\mathcal{G}) > 0$. A small value signals a bottleneck: comparatively few edges join a set of vertices to the rest of the graph. A family of graphs of a fixed degree d is an *expander family* if its Cheeger constants are bounded below by a positive constant independent of the number of vertices.

Theorem 9.1 (Discrete Cheeger inequalities). *For a connected d -regular graph,*

$$\frac{d - \mu_2}{2} \leq h(\mathcal{G}) \leq \sqrt{2d(d - \mu_2)}. \quad (9.1)$$

Consequently, for fixed d , a family is an expander family if and only if its spectral gaps $d - \mu_2$ are bounded below by a positive constant.

Proof. Put $L = dI - A$. Every sum over $\{i, j\} \in E$ below counts an undirected edge once. The variational characterization is

$$d - \mu_2 = \min_{v \perp \mathbf{1}, v \neq 0} \frac{v^\top L v}{\|v\|_2^2} = \min_{v \perp \mathbf{1}, v \neq 0} \frac{\sum_{\{i,j\} \in E} (v_i - v_j)^2}{\sum_i v_i^2}. \quad (9.2)$$

Indeed, $v^\top A v = 2 \sum_{\{i,j\} \in E} v_i v_j$ and each vertex has degree d . Choose S realizing $h(\mathcal{G})$, with $s = |S| \leq N/2$, and set

$$v_i = \begin{cases} s^{-1}, & i \in S, \\ -(N - s)^{-1}, & i \notin S. \end{cases}$$

Then $\sum_i v_i = 0$, and its Rayleigh quotient is

$$\frac{|E(S, S^c)|(s^{-1} + (N - s)^{-1})^2}{s^{-1} + (N - s)^{-1}} = \frac{|E(S, S^c)|}{s} \frac{N}{N - s} \leq 2h(\mathcal{G}).$$

This proves the first inequality.

For the second inequality, let $v \perp \mathbf{1}$ satisfy $Lv = \gamma v$, where $\gamma = d - \mu_2$. After replacing v by $-v$ if necessary, its positive part $u_i = \max\{v_i, 0\}$ is nonzero and is supported on at most $N/2$ vertices. Truncation gives

$$\sum_{\{i,j\} \in E} (u_i - u_j)^2 \leq \sum_{\{i,j\} \in E} (u_i - u_j)(v_i - v_j) = \sum_i u_i (Lv)_i = \gamma \sum_i u_i^2.$$

For $S_t = \{i : u_i^2 > t\}$, integration over level sets gives

$$\sum_{\{i,j\} \in E} |u_i^2 - u_j^2| = \int_0^\infty |E(S_t, S_t^c)| dt \geq h(\mathcal{G}) \int_0^\infty |S_t| dt = h(\mathcal{G}) \sum_i u_i^2.$$

On the other hand, Cauchy–Schwarz and regularity imply

$$\sum_{\{i,j\} \in E} |u_i^2 - u_j^2| \leq \left(\sum_{\{i,j\} \in E} (u_i - u_j)^2 \right)^{1/2} \left(\sum_{\{i,j\} \in E} (u_i + u_j)^2 \right)^{1/2} \leq \sqrt{2d\gamma} \sum_i u_i^2.$$

Canceling $\sum_i u_i^2$ proves the claim. This proof also explains the picture in the handwritten notes: level sets of an eigenvector associated with μ_2 supply candidate cuts. A threshold other than zero may be needed. $\square$

There is a useful L^1 form of expansion. For every real function f on V , the same level-set argument gives

$$N \sum_{\{i,j\} \in E} |f_i - f_j| \geq h(\mathcal{G}) \sum_{i < j} |f_i - f_j| = \frac{h(\mathcal{G})}{2} \sum_{i,j} |f_i - f_j|. \quad (9.3)$$

For example, at a level set of size s , the two sums count the boundary edges and $s(N - s)$ separated vertex pairs, respectively; the inequality follows from $s(N - s) \leq N \min\{s, N - s\}$. The factor $1/2$ on the right is necessary when vertex pairs are ordered.




Figure 14: Two cuts in cubic graphs: a bottleneck with two crossing edges and a cut with eight crossing edges. These are illustrative cuts, not a computation of the global Cheeger minimum. Expansion compares every cut with the smaller number of vertices on its two sides.

9.2 The largest possible spectral gap

For fixed d , the Alon–Boppana theorem [18] states that a growing family of d -regular graphs cannot have μ_2 uniformly below $2\sqrt{d-1}$. One diameter version, sufficient for the discussion here, is

$$\mu_2 \geq 2\sqrt{d-1} - \frac{C_d}{\text{diam}(\mathcal{G})}, \quad \text{diam}(\mathcal{G}) = \max_{i,j \in V} \text{dist}(i,j).$$

The variational idea is to construct test vectors supported near vertices far apart and compare their Rayleigh quotients with those of a large truncated regular tree. This is a theorem being quoted, rather than a consequence of the short test-vector calculation above.

A connected d -regular graph is called *Ramanujan* if every nontrivial adjacency eigenvalue satisfies

$$|\mu_j| \leq 2\sqrt{d-1}.$$

Here d is always trivial, and $-d$ is also trivial when the graph is bipartite. Thus one excludes the eigenvalues of absolute value d in this definition. Ramanujan graphs attain the asymptotically optimal spectral bound at both ends of the nontrivial spectrum.

The notes mention two construction methods. The number-theoretic constructions of Lubotzky–Phillips–Sarnak [19] yield explicit families, including degrees $d = p + 1$ for suitable primes p (in their classical presentation, $p \equiv 1 \pmod{4}$). The interlacing-polynomial method of Marcus–Spielman–Srivastava [20] gives infinite families of bipartite Ramanujan graphs for every integer $d \geq 3$.

The second method can be pictured as a tower of *two-lifts*. Replace each vertex v by $(v, 0), (v, 1)$. For each edge $\{u, v\}$, choose either the two parallel edges $\{(u, 0), (v, 0)\}, \{(u, 1), (v, 1)\}$ or the two crossed edges $\{(u, 0), (v, 1)\}, \{(u, 1), (v, 0)\}$. Degrees remain unchanged. The spectrum splits into the old adjacency spectrum and the spectrum of the signed adjacency matrix specifying these choices. For a bipartite base, their theorem supplies a choice for which the new eigenvalues obey the two-sided Ramanujan bound. Starting with $K_{d,d}$ therefore gives

$$\underbrace{K_{d,d}}_{2d \text{ vertices}} \longleftarrow \underbrace{\mathcal{G}_1}_{4d \text{ vertices}} \longleftarrow \underbrace{\mathcal{G}_2}_{8d \text{ vertices}} \longleftarrow \cdots,$$

where each arrow is a two-to-one covering map.

9.3 Random graphs and the two spectral edges

Return to the normalization

$$H = \frac{A}{\sqrt{d-1}}, \quad \lambda_j = \frac{\mu_j}{\sqrt{d-1}}, \quad \mathcal{A}_d = \frac{d(d-1)}{(d-2)^2}.$$

The nontrivial Ramanujan interval is now $[-2, 2]$. For a uniformly random simple d -regular graph, with fixed $d \geq 3$ and Nd even, the edge universality theorem of Huang–McKenzie–Yau gives

$$(\mathcal{A}_d N)^{2/3}(\lambda_2 - 2) \implies \text{TW}_1. \quad (9.4)$$

The lower edge has the same limiting law after reflection. The two limiting edge fluctuations are independent [11]. Consequently, if F_1 is the distribution function of TW_1 ,

$$\mathbb{P}(\lambda_2 \leq 2) \longrightarrow F_1(0) \approx 0.83, \quad \mathbb{P}(\lambda_2 \leq 2, \lambda_N \geq -2) \longrightarrow F_1(0)^2 \approx 0.69.$$

The second event, rather than the one-sided event alone, gives the Ramanujan property for the typical nonbipartite random graph. These are limits as $N \rightarrow \infty$, not exact percentages at finite N .




Figure 15: A finite-size simulation of the two edges. The graph sample consists of 200 independent uniform cubic graphs on $N = 1000$ vertices; the horizontal variable is $(6N)^{2/3}(\lambda_2 - 2)$ at the upper edge and $(6N)^{2/3}(-\lambda_N - 2)$ at the lower edge. The comparison curve uses 3000 finite GOE samples of dimension 800. The observed two-edge Ramanujan fraction is 0.795 (Monte Carlo standard error approximately 0.029). Its difference from the asymptotic value near 0.69 illustrates the finite-size effect and must not be mistaken for a new limiting value.

10 Local resampling and the cavity self-consistent equation

10.1 Two resolvent averages

Write $G(z) = (H - z)^{-1}$ and use $G^{(i)}$ for the resolvent after deleting vertex i . Besides the usual Stieltjes transform, introduce the average cavity Green function

$$m_N(z) = \frac{1}{N} \sum_i G_{ii}(z), \quad Q(z) = \frac{1}{Nd} \sum_{(i,j) \in \vec{E}} G_{jj}^{(i)}(z), \quad (10.1)$$

where $\vec{E}$ consists of the Nd directed edges. Deleting the parent of a vertex changes a regular tree into a rooted $(d-1)$ -ary tree. Accordingly, the deterministic approximations are different:

$$m_N(z) \approx m_d(z), \quad Q(z) \approx m_{\text{sc}}(z), \quad m_d(z) = \frac{1}{-z - \frac{d}{d-1} m_{\text{sc}}(z)}.$$

Here $m_{\text{sc}} = (-z + \sqrt{z^2 - 4})/2$, with the branch for which $m_{\text{sc}}(z) \sim -1/z$ at infinity. In particular, m_d must not be replaced by m_{sc} when the degree is fixed.

For context, a convenient weak form of the local law is as follows. Fix $d \geq 3$, $\delta > 0$, and a spectral window $|\operatorname{Re} z| \leq 2 + \kappa_d$, where $0 < \kappa_d < d/\sqrt{d-1} - 2$. There are small positive constants c, a such that, for $r = \lfloor c \log_{d-1} N \rfloor$, with probability tending to one,

$$\begin{aligned} |m_N(z) - m_d(z)| + |Q(z) - m_{\text{sc}}(z)| &\leq N^{-a}, \\ \max_{i,j} |G_{ij}(z) - P_{ij}(\mathcal{B}_r(\{i, j\}), z)| &\leq N^{-a}, \end{aligned} \quad (10.2)$$

uniformly in that window for $N^{-1+\delta} \leq \operatorname{Im} z \leq 1$. In this statement P is the resolvent of the local graph with infinite regular-tree branches attached to restore the missing degrees. Unlike a single universal tree, it retains any cycles in the local neighborhood. This local law is an input to the discussion below [12].

The handwritten estimate $(\log N)^C((d-1)^{-r/2} + (N\eta)^{-1/4})$, with $\eta = \operatorname{Im} z$, describes the two expected sources of a weak error: a finite-depth error and a sampling error. The precise uniform local-law theorem includes stability and graph-geometry conditions; the short resampling argument below is an explanation of its mechanism, not a proof of that theorem.

10.2 A finite-tree fixed-point map

Let $\mathcal{Y}_r$ be the rooted $(d-1)$ -ary tree of depth r , with root o , and all edge weights equal to $1/\sqrt{d-1}$. At each leaf place a self-energy Δ . Equivalently, for $r \geq 1$ define

$$f_r(\Delta) = (H_{\mathcal{Y}_r} - z - \Delta I_{\partial \mathcal{Y}_r})_{oo}^{-1}.$$

For $r = 0$ the same convention means $f_0(\Delta) = (-z - \Delta)^{-1}$. Eliminating the leaves by the Schur complement gives the particularly simple recursion

$$F_z(w) := \frac{-1}{z + w}, \quad f_r(\Delta) = F_z^{o(r+1)}(\Delta). \quad (10.3)$$

Since $m = m_{\text{sc}}(z)$ satisfies $m = F_z(m)$, this proves the exact identity

$$f_r(m_{\text{sc}}) = m_{\text{sc}}.$$

For a full d -regular root with a depth- r neighborhood, the corresponding map is

$$\chi_r(\Delta) := \frac{-1}{z + \frac{d}{d-1} F_z^{\text{or}}(\Delta)}, \quad \chi_r(m_{\text{sc}}) = m_d. \quad (10.4)$$

Thus the natural random-graph self-consistent relations are

$$Q \approx f_r(Q), \quad m_N \approx \chi_r(Q).$$

Selecting the solution near m_{sc} also requires the resolvent sign, large- z behavior, and stability; an approximate fixed-point identity alone does not identify the desired branch.

To connect this recursion with the block-matrix picture in the notes, partition a graph into a neighborhood T and its complement:

$$H = \begin{pmatrix} H_T & B^\top \\ B & D \end{pmatrix}, \quad G|_T = (H_T - z - B^\top(D - z)^{-1}B)^{-1}.$$

On an infinite tree the outside branches are disjoint. Therefore the boundary self-energy is diagonal; each leaf receives $(d-1)$ equal contributions $m_{\text{sc}}/(d-1)$, totaling m_{sc} . This explains exactly why the boundary value in (10.3) is the semicircle transform even though the full-root value is m_d .

10.3 Switching the boundary edges

Let T be a tree neighborhood and enumerate its outgoing boundary edges as

$$(\ell_\alpha, a_\alpha), \quad 1 \leq \alpha \leq M, \quad \ell_\alpha \in T, \quad a_\alpha \notin T.$$

The number M is of order $(d-1)^{r+1}$ for a depth- r neighborhood. Choose directed edges (b_α, c_α) in the rest of the graph and perform the switch

$$\{(\ell_\alpha, a_\alpha), (b_\alpha, c_\alpha)\} \longleftrightarrow \{(\ell_\alpha, c_\alpha), (b_\alpha, a_\alpha)\}. \quad (10.5)$$

Every participating vertex loses and gains one edge, so degrees are preserved. If endpoints collide or one of the new edges already exists, the proposal must be rejected. Geometrically, the neighborhood is replanted at random distant vertices $c_1, \dots, c_M$.

For this interpretation to be useful, the chosen locations should have tree neighborhoods and should be separated from each other and from the original neighborhood. If all but N^ε vertices have a tree neighborhood of a prescribed radius R , elementary exposure bounds have the schematic form

$$\mathbb{P}(\text{a specified sampled location is bad} \mid \mathcal{G}) \lesssim \frac{N^\varepsilon}{N} + \frac{|\mathcal{F}|(d-1)^{2R}}{N},$$

where $\mathcal{F}$ is the previously exposed set. A union bound is also needed over the M proposals. Taking the radii to be sufficiently small multiples of $\log_{d-1} N$ makes the total error small. The precise exponent depends on the separation convention; it is not automatically $N^{-1+o(1)}$ for every logarithmic radius.

The resampling procedure is arranged so that the original and resampled graphs $\mathcal{G}, \tilde{\mathcal{G}}$ form an *exchangeable pair*:

$$(\mathcal{G}, \tilde{\mathcal{G}}) \stackrel{\text{d}}{=} (\tilde{\mathcal{G}}, \mathcal{G}). \quad (10.6)$$




Figure 16: The elementary switching in (10.5). The internal neighborhood stays fixed while its attachment point moves from a to c . Both graphs have the same vertex degrees.

The mechanism is a measure-preserving involution on the augmented space consisting of a graph together with its sampled switching data. For one switch, the reverse data use the edges (ℓ_α, c_α) and (b_α, a_α) , and switching again returns the original graph. Uniform sampling gives equal weights to the forward and reverse marked configurations. For several switches, the eligibility rules must themselves be invariant under this reversal. Involutivity of the edge replacement alone is not sufficient if arbitrary, nonsymmetric rejection rules are imposed. The fully specified local-resampling construction has this reversibility property [11].

10.4 What the first-order resolvent expansion explains

Consider a cavity neighborhood T , with the parent i of o removed. All block decompositions in this subsection are taken inside the parent-deleted graph. After resampling write its boundary self-energy as

$$\tilde{B}^\top \tilde{G}^{(T \cup \{i\})} \tilde{B} = Q I_{\partial T} + \Delta, \quad R_Q = (H_T - z - Q I_{\partial T})^{-1}.$$

For leaves u, v , the entries of the error are

$$\Delta_{uv} = \frac{1}{d-1} \sum_{\alpha: \ell_\alpha = u} \sum_{\beta: \ell_\beta = v} \tilde{G}_{c_\alpha c_\beta}^{(T \cup \{i\})} - Q \mathbf{1}_{u=v}.$$

The exact two-step resolvent identity is

$$\tilde{G}^{(i)}|_T = R_Q + R_Q \Delta R_Q + R_Q \Delta R_Q \tilde{G}^{(i)}|_T. \quad (10.7)$$

It follows that

$$\tilde{G}_{oo}^{(i)} = f_r(Q) + \sum_{u,v \in \partial T} (R_Q)_{ou} \Delta_{uv} (R_Q)_{vo} + \text{an exact quadratic remainder.}$$

A convergent infinite Neumann series requires an additional norm bound; (10.7) does not.

At the deterministic value $Q = m_{\text{sc}}$, the cavity tree kernel is

$$(R_m)_{ou} = m \left(-\frac{m}{\sqrt{d-1}} \right)^{\text{dist}(o,u)}. \quad (10.8)$$

The minus sign comes from the convention $G = (H - z)^{-1}$. For leaves at depth r , its magnitude is at most $(d-1)^{-r/2}$, since $|m| \leq 1$. Thus the linear correction is a weighted average of roughly $(d-1)^{2r}$ boundary resolvent entries.

The diagonal and off-diagonal terms have different cancellation mechanisms. For a uniformly sampled directed edge (B, C) of a fixed graph,

$$\mathbb{E}_{(B,C)} G_{CC}^{(B)} = Q$$

is an exact identity. An average over independent samples would therefore have a centered diagonal contribution and a square-root sampling gain. For the off-diagonal part, if U, V are independent uniform vertices of a fixed graph, then

$$\begin{aligned} \mathbb{E}_{U,V} G_{UV} &= \frac{1}{N^2} \mathbf{1}^\top G \mathbf{1} = \frac{1}{N(d/\sqrt{d-1} - z)}, \\ \mathbb{E}_{U,V} |G_{UV}|^2 &= \frac{1}{N^2} \sum_{u,v} |G_{uv}|^2 = \frac{\text{Im } m_N(z)}{N\eta}. \end{aligned} \quad (10.9)$$

The first identity uses the constant eigenvector, and the second is the Ward identity. Near the nontrivial edge 2, the denominator in the first line is bounded away from zero for fixed d , so this average is $O(N^{-1})$.

These computations motivate the rough relation recorded in the notes,

$$Q - f_r(Q) \text{ of size roughly } (d-1)^{-r/2} + (N\eta)^{-1/2},$$

up to logarithms, stability losses, and geometry errors. They are not by themselves a proof: the entries of $\tilde{G}^{(T \cup \{i\})}$ depend on the same sampling data as the locations c_α , and the cavity minors differ from the full resolvent used in (10.9). One must control these dependencies, the quadratic remainder, exceptional switches, and high moments. The local law and the sharper loop equations require this additional resampling analysis.

11 The first edge loop equation

The preceding self-consistent equation gives the leading deterministic behavior. At an edge, the next-order correction has the same size as the fluctuations one is trying to study. The final handwritten page illustrates this using the first loop equation.

For the standard GOE normalization, Gaussian integration by parts gives

$$\mathbb{E} \left[1 + z m_N(z) + m_N(z)^2 + \frac{1}{N} \partial_z m_N(z) \right] = 0. \quad (11.1)$$

For a general Wigner matrix, additional cumulant errors may be present. The derivative term shown here is the real-symmetric ($\beta = 1$) correction; it is absent in the corresponding exact GUE ($\beta = 2$) identity with its standard normalization.

For random regular graphs, the corresponding first refined identity near the right edge has the structure

$$\mathbb{E} \left[\frac{\mathcal{A}_d^2}{r+1} (Q - f_r(Q)) + \frac{1}{N} \partial_z m_N \right] = \text{a controlled smaller error.} \quad (11.2)$$

This refinement is an input from the local-resampling analysis [11]; it does not follow merely by inserting the rough error bound from the previous subsection. We now explain the deterministic Taylor algebra that converts it into the familiar edge loop equation.

11.1 Taylor expansion of the finite-tree maps

Put $m = m_{\text{sc}}(z)$, $u = Q - m$, and $k = r + 1$. Since

$$F_z(m + u) = m + m^2 u + m^3 u^2 + O(u^3),$$

iteration of (10.3) yields

$$f_r(m + u) = m + m^{2k} u + m^{2k+1} \left(\sum_{j=0}^{k-1} m^{2j} \right) u^2 + O_r(u^3). \quad (11.3)$$

To check the quadratic coefficient, write an iterate as $m + a_k u + b_k u^2 + O(u^3)$. Then $a_{k+1} = m^2 a_k$ and $b_{k+1} = m^2 b_k + m^3 a_k^2$, with $a_0 = 1, b_0 = 0$. This gives exactly the displayed finite sum. Similarly,

$$\chi_r(m + u) - m_d = \frac{d}{d-1} m_d^2 m^{2r} u + O_r(u^2). \quad (11.4)$$

Set $s = \sqrt{z-2}$, using the branch inherited from the upper half-plane. As $z \rightarrow 2$,

$$\begin{aligned} m &= -1 + s + O(s^2), & m^2 &= 1 - 2s + O(s^2), \\ m_d &= -\frac{d-1}{d-2} + \mathcal{A}_d s + O(s^2). \end{aligned} \quad (11.5)$$

For fixed r , substitution into (11.3) gives

$$\frac{Q - f_r(Q)}{r+1} = 2su + u^2 + O_r(|s|^2|u| + |s||u|^2 + |u|^3). \quad (11.6)$$

The factor $r+1$ has a simple origin: at $m = -1$ each of the $r+1$ iterations contributes the same quadratic term. Also (11.4) becomes

$$\chi_r(Q) - m_d = \mathcal{A}_d u + O_r(|s||u| + |u|^2).$$

Together with the suitably accurate approximation $m_N \approx \chi_r(Q)$, this identifies $m_N - m_d$ with $\mathcal{A}_d(Q - m)$ to leading order. If r grows with N , the remainder constants must be tracked; the fixed- r notation O_r alone does not justify a growing-radius argument.

Substituting these relations into (11.2) gives the first edge loop equation in the form recorded in the lecture:

$$\mathbb{E} \left[(m_N - m_d)^2 + 2\mathcal{A}_d \sqrt{z-2} (m_N - m_d) + \frac{1}{N} \partial_z m_N \right] = o(N^{-2/3}). \quad (11.7)$$

Here $z = 2 + w/(\mathcal{A}_d N)^{2/3}$ with w in a fixed compact subset of $\mathbb{C}^+$. The $o(N^{-2/3})$ is the precision supplied by the full refined resampling argument, not an error bound proved by the preceding Taylor

calculation. At the microscopic scale $z - 2 \asymp N^{-2/3}$, all three principal terms are of order $N^{-2/3}$. Hence an error that is only $o(1)$ is inadequate: the error in this identity must be negligible relative to $N^{-2/3}$. The theorem establishing the requisite estimates, and the higher loop equations needed to identify the full limiting edge process, are separate parts of the edge-universality proof. The calculations above explain the normalization and the first equation without treating those missing estimates as established by a formal expansion.

The scale factor can also be checked directly. Set

$$z = 2 + \frac{w}{(\mathcal{A}_d N)^{2/3}}, \quad \mathcal{S}_N(w) = \frac{N^{1/3}}{\mathcal{A}_d^{2/3}} \left(m_N(z) + \frac{d-1}{d-2} \right).$$

The poles of $\mathcal{S}_N$ are the rescaled graph eigenvalues $(\mathcal{A}_d N)^{2/3}(\lambda_j - 2)$, each with residue -1 . Using (11.5) in the refined identity and dividing by $\mathcal{A}_d^{4/3} N^{-2/3}$ gives, under the requisite moment bounds,

$$\mathbb{E}[\mathcal{S}_N(w)^2 - w + \mathcal{S}'_N(w)] = o(1).$$

This is precisely the first Airy_1 loop equation. The trivial graph eigenvalue moves to $+\infty$ in these coordinates and has a vanishing resolvent contribution on compact sets; edge statistics concern the nontrivial eigenvalues.

A Numerical experiments and reproduction

Every numerical picture in these notes is a new computation. The original handwritten pictures are qualitative sketches and do not specify datasets that can be recovered exactly. Random samples, deterministic evaluations, and graph schematics are distinguished in the captions.

Experiment	Parameters and method
Semicircle and local scales	GOE/GUE: $N = 1600$, five samples per ensemble; exact Gaussian beta tridiagonal model. Seed 2026101201.
Edge fluctuations	$N = 800$, 3000 samples for each $\beta = 1, 2, 4$; six largest eigenvalues retained. Seed 2026101202.
Rigidity	First GUE realization above; descending midpoint semicircle quantiles.
Last-passage percolation	24×24 independent mean-one exponential weights; exact dynamic programming. Seed 2026101203.
Increasing subsequence	Uniform permutation of 90 elements; exact longest increasing subsequence. Seed 2026101204.
Billiards	Specular reflections; two circle rotation numbers $2/7$ and $(\sqrt{5} - 1)/4$; Sinai table $[-1, 1]^2$ minus a radius-0.30 disk.
Regular graphs	200 simple cubic graphs on 1000 labeled vertices; 12 complete spectra. Exact configuration-model rejection. Seed 260612.
Airy and zeta plots	Deterministic special-function evaluations; the zeta picture uses the first 40 computed nontrivial zeros.

The Gaussian ensembles are generated by (1.11), and the tridiagonal eigenvalue solver is used instead of forming dense matrices. This changes computational cost but not the eigenvalue distribution. For each graph sample, all Nd labeled half-edges are paired uniformly. If any loop or multiple edge occurs, the entire pairing is rejected and a new one is drawn. Each simple labeled graph has exactly $(d!)^N$ representations, so accepted samples have the desired uniform law. The experiment needed 1581 pairing attempts for its 200 accepted graphs.

The edge density curves are kernel density estimates with Scott bandwidth. The graph comparison uses empirical cumulative distribution functions. The observed two-sided event $\lambda_2 \leq 2$, $\lambda_N \geq -2$ occurred in 159 of 200 graph samples, giving 0.795 with estimated Monte Carlo standard error $\sqrt{0.795(1 - 0.795)/200} \approx 0.0285$. This estimates the finite- N probability; it does not remove finite-size bias or test a convergence rate toward $F_1(0)^2$.

The companion source archive contains the single main LaTeX file, all figure PDFs and PNGs, the Python scripts, and raw numerical arrays. Run the scripts from the archive directory in the following order:

```
python3 simulations.py
python3 graph_simulations.py
python3 framework_diagrams.py
python3 zeta_zeros.py
pdflatex rmt_lecture_notes.tex
pdflatex rmt_lecture_notes.tex
```

The main numerical dependencies are NumPy, SciPy, Matplotlib, and mpmath. The exact versions used for the main matrix computations and the seeds are recorded in `data/metadata.json`; graph parameters are in `data/graph_parameters.json`. Raw arrays are stored in the `data` directory. Figures are also supplied, so compiling the notes does not require rerunning the simulations.

References

- [1] I. Dumitriu and A. Edelman, *Matrix models for beta ensembles*, Journal of Mathematical Physics **43** (2002), 5830–5847. <https://arxiv.org/abs/math-ph/0206043>.
- [2] B. Valkó and B. Virág, *Continuum limits of random matrices and the Brownian carousel*, Inventiones Mathematicae **177** (2009), 463–508. <https://arxiv.org/abs/0712.2000>.
- [3] J. A. Ramírez, B. Rider, and B. Virág, *Beta ensembles, stochastic Airy spectrum, and a diffusion*, Journal of the American Mathematical Society **24** (2011), 919–944. <https://arxiv.org/abs/math/0607331>.
- [4] C. A. Tracy and H. Widom, *Level-spacing distributions and the Airy kernel*, Communications in Mathematical Physics **159** (1994), 151–174. <https://arxiv.org/abs/hep-th/9211141>.
- [5] C. A. Tracy and H. Widom, *On orthogonal and symplectic matrix ensembles*, Communications in Mathematical Physics **177** (1996), 727–754. <https://arxiv.org/abs/solv-int/9509007>.
- [6] J. O. Lee and J. Yin, *A necessary and sufficient condition for edge universality of Wigner matrices*, Duke Mathematical Journal **163** (2014), 117–173. <https://arxiv.org/abs/1206.2251>.
- [7] J. Baik, P. Deift, and K. Johansson, *On the distribution of the length of the longest increasing subsequence of random permutations*, Journal of the American Mathematical Society **12** (1999), 1119–1178. <https://arxiv.org/abs/math/9810105>.
- [8] K. Johansson, *Shape fluctuations and random matrices*, Communications in Mathematical Physics **209** (2000), 437–476. <https://arxiv.org/abs/math/9903134>.

- [9] H. L. Montgomery, *The pair correlation of zeros of the zeta function*, in *Analytic Number Theory* (St. Louis, 1972), Proceedings of Symposia in Pure Mathematics **24**, American Mathematical Society, 1973, 181–193. <https://websites.umich.edu/~hlm/paircor1.pdf>.
- [10] O. Bohigas, M. J. Giannoni, and C. Schmit, *Characterization of chaotic quantum spectra and universality of level fluctuation laws*, Physical Review Letters **52** (1984), 1–4. <https://doi.org/10.1103/PhysRevLett.52.1>.
- [11] J. Huang, T. McKenzie, and H.-T. Yau, *Ramanujan property and edge universality of random regular graphs*, preprint, 2024. <https://arxiv.org/abs/2412.20263>.
- [12] J. Huang and H.-T. Yau, *Spectrum of random d -regular graphs up to the edge*, Communications on Pure and Applied Mathematics **77** (2024), 1635–1723. <https://arxiv.org/abs/2102.00963>.
- [13] L. Erdős, H.-T. Yau, and J. Yin, *Rigidity of eigenvalues of generalized Wigner matrices*, Advances in Mathematics **229** (2012), 1435–1515. <https://arxiv.org/abs/1007.4652>.
- [14] L. Erdős, A. Knowles, H.-T. Yau, and J. Yin, *The local semicircle law for a general class of random matrices*, Electronic Journal of Probability **18** (2013), no. 59, 1–58. <https://arxiv.org/abs/1212.0164>.
- [15] P. Bourgade, L. Erdős, H.-T. Yau, and J. Yin, *Fixed energy universality for generalized Wigner matrices*, Communications on Pure and Applied Mathematics **69** (2016), 1815–1881. <https://arxiv.org/abs/1407.5606>.
- [16] P. Bourgade and J. Huang, *Loop equations characterize random matrix statistics*, preprint, 2026. <https://arxiv.org/abs/2607.07617>.
- [17] S. Janson, *The probability that a random multigraph is simple*, Combinatorics, Probability and Computing **18** (2009), 205–225. <https://arxiv.org/abs/math/0609802>.
- [18] A. Nilli, *On the second eigenvalue of a graph*, Discrete Mathematics **91** (1991), 207–210. [https://doi.org/10.1016/0012-365X\(91\)90112-F](https://doi.org/10.1016/0012-365X(91)90112-F).
- [19] A. Lubotzky, R. Phillips, and P. Sarnak, *Ramanujan graphs*, Combinatorica **8** (1988), 261–277. <https://doi.org/10.1007/BF02126799>.
- [20] A. W. Marcus, D. A. Spielman, and N. Srivastava, *Interlacing families I: Bipartite Ramanujan graphs of all degrees*, Annals of Mathematics **182** (2015), 307–325. <https://arxiv.org/abs/1304.4132>.