

Random matrix theory

Manjunath Krishnapur
Indian Institute of Science, Bangalore

October 27, 2017

Contents

1	The simplest non-trivial matrices	8
	Jacobi matrices	8
	The 1-dimensional discrete Laplacian matrix	9
	Empirical spectral distribution and related notions	11
	The oscillator matrix	14
	A class of Jacobi matrices	16
	Exercises	18
2	The simplest non-trivial random matrix	20
	Parameterizing a Jacobi matrix by its spectral measure at e_1	20
	The Jacobian determinant	21
	A class of random Jacobi matrices	22
	Computation of the Jacobian determinant	24
	Remarks on the moment problem and Jacobi matrices	28
	Laguerre Beta ensembles: Another random tridiagonal matrix	30
	Exercises	32
	Notes	33
3	The beta log-gas	34
	The quadratic beta log-gas	35
	Mehta integral	36
	The range of the log-gas	37
	Exercises	39
4	The method of moments applied to Jacobi matrices, deterministic and random	40
	A class of deterministic Jacobi matrices	40

	The method of moments	40
	A more general Jacobi matrix	43
	The limiting distribution of beta log-gases	45
	Exercises	47
5	Stieltjes' transform method for Jacobi matrices	49
	Spectral measure of $T_n(f)$ at e_1	49
	Limiting spectral distribution of $T_n(f)$	50
	One dimensional Anderson model and the method of spectral averaging . .	52
6	Gaussian random matrices	54
	GOE and GUE	54
	Reduction to a Jacobi matrix	55
	Eigenvalue distribution of GOE and GUE	56
	A direct proof by change of variable? Some remarks	58
	Generalizations of GOE and GUE in two directions	60
	Exercises	61
7	Wigner matrices: The semi-circle law	63
	The invariance principle	63
	An illustration: The Lindeberg-Feller central limit theorem	64
	Proof of the invariance principle	65
	Semicircle law for Wigner matrices	66
8	Moment method applied to GOE and connections to enumeration problems	70
	Expected ESD of the GOE matrix	70
	Semi-circle law for a more general class of random matrices	76
9	Free probability and random matrices	78
	Cumulants and moments in classical probability	78
	Noncommutative probability spaces, free independence	83
	Moment-cumulant calculus	88
	Free convolution	91
	Integral transforms	92
	Free central limit theorem	94
	Random matrices and freeness	95

Spectrum of the sum of two matrices and free convolution	96
Exercises	97
10 Non-asymptotic questions	98
Gaussian matrices	99
Rectangular matrices with independent entries	104
Appendix 1: Weak convergence and techniques to show it	105
Probability measures on the real line	105
Stieltjes' transform of a probability measure	107
Examples	109
Bounding Lévy distance in terms of Stieltjes transform	111
Method of moments	112
Exercises	113
Appendix 2: Some linear algebra facts	114
Bounds on eigenvalues	114
Perturbations of eigenvalues	114
Block matrix inversion formula	117
Shooting description of eigenvectors and eigenvalues of a Jacobi matrix	118
Appendix 3: Gaussian random variables	120
Basics of Gaussians, moments, cumulants	120
Comparison inequalities	124
Gaussian isoperimetric inequality	129
Appendix 4: Some combinatorics facts	131
The Mobius function of a lattice	131

Acknowledgments: Lecture notes from a course on random matrix theory in the fall of 2017 at IISc, Bangalore. An earlier version was made during a course in the spring of 2011 at IISc, Bangalore. Thanks to those who attended the course (Rajesh Sundaresan, Tulasi Ram Reddy, Kartick Adhikari, Indrajit Jana and Subhamay Saha first time). Thanks to Anirban Dasgupta for pointing out some errors in the notes.

Chapter 1

The simplest non-trivial matrices

Random matrix theory is largely the study of eigenvalues and eigenvectors of matrices whose entries are random variables. In this chapter, we shall motivate the kinds of questions studied in random matrix theory, but using deterministic matrices. That will also help us to set up the language in which to phrase the questions and answers.

Spectral quantities have more geometric meaning for symmetric (and Hermitian and normal) matrices than for general matrices. Since there is not much to say about diagonal matrices, we are led to real symmetric tridiagonal matrices as the simplest non-trivial matrices.

Jacobi matrices

Given real numbers a_1, \dots, a_n and strictly positive numbers b_1, \dots, b_{n-1} , let

$$T = T_n(a, b) = \begin{bmatrix} a_1 & b_1 & 0 & 0 & 0 & 0 \\ b_1 & a_2 & b_2 & \ddots & 0 & 0 \\ 0 & b_2 & \ddots & \ddots & \ddots & 0 \\ 0 & \ddots & \ddots & \ddots & b_{n-2} & 0 \\ 0 & 0 & \ddots & b_{n-2} & a_{n-1} & b_{n-1} \\ 0 & 0 & 0 & 0 & b_{n-1} & a_n \end{bmatrix}. \quad (1)$$

This is the real symmetric $n \times n$ tridiagonal matrix with diagonal entries $T_{k,k} = a_k$ for $1 \leq k \leq n$ and $T_{k,k+1} = T_{k+1,k} = b_k$ for $1 \leq k \leq n-1$.

Why did we assume strict positivity of b_j s? If $b_k = 0$, the matrix breaks into a direct sum of two matrices, hence we impose the condition $b_k \neq 0$ for all k . By conjugating



Figure 1.1: Line graph on 8 vertices. For $f : V \mapsto \mathbb{R}$, define $\Delta f(k) = 2f(k) - f(k-1) - f(k+1)$ where $f(0)$ and $f(n+1)$ are interpreted as 0. The matrix of Δ in standard basis is $2I_n - T_n$ where T_n has $a_k = 0$ and $b_k = 1$.

with an appropriate diagonal matrix having ± 1 entries on the diagonal, any real symmetric tridiagonal matrix can be transformed into a (unique) Jacobi matrix. Indeed, if $D = \text{diag}(\varepsilon_1, \dots, \varepsilon_n)$ where $\varepsilon_i = \pm 1$, then $DT_n(a, b)D^{-1} = T_n(a, c)$ where $c_k = \varepsilon_k \varepsilon_{k+1} b_k$. From this it is clear how to choose ε_i s so that $c_k > 0$ for all k . As we are interested in eigenvalues and they don't change under conjugation (and eigenvectors transform in a simple way), we may as well start with the assumption that b_k are positive.

The 1-dimensional discrete Laplacian matrix

The single most important example is the case $a_k = 0$ and $b_k = 1$. In this case,

$$T_n(a, b) = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & \ddots & 0 & 0 \\ 0 & 1 & \ddots & \ddots & \ddots & 0 \\ 0 & \ddots & \ddots & \ddots & 1 & 0 \\ 0 & 0 & \ddots & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}_{n \times n}.$$

This matrix arises in innumerable contexts. For example, $2I_n - T_n$ is the discrete second derivative operator. It is also the Laplacian on the to the line graph with n vertices,¹ provided we add a loop at 1 and at n . It can also be considered as the generator of the simple random walk on this graph.

The eigenvalue equation $Tv = \lambda v$ can be written out explicitly as

$$v_{k-1} + v_{k+1} = \lambda v_k \text{ for } 1 \leq k \leq n,$$

¹In general, for a graph $G = (V, E)$, the laplacian is the linear operator $\Delta : \mathbb{R}^V \mapsto \mathbb{R}^V$ defined by $\Delta f(u) = \text{deg}(u) - \sum_{v: v \sim u} f(v)$. In standard basis its matrix is $D - A$ where A is the adjacency matrix and D is the $V \times V$ diagonal matrix whose (u, u) entry is $\text{deg}(u)$.

where $v^t = (v_1, \dots, v_n)$ and we adopt the convention that $v_0 = v_{n+1} = 0$. We try $v_k = \sin(k\theta)$ since

$$\sin((k-1)\theta) + \sin((k+1)\theta) = 2\cos\theta \sin(k\theta).$$

To satisfy the boundary conditions $v_0 = v_{n+1} = 0$, we take $\theta = \theta_\ell = \frac{\pi\ell}{n+1}$ for some $1 \leq \ell \leq n$. This gives us the eigenvalues $\lambda_\ell = 2\cos(\theta_\ell)$ with the corresponding eigenvectors (caution: they are not normalized) $v_\ell^t = (\sin(\theta_\ell), \sin(2\theta_\ell), \dots, \sin(n\theta_\ell))$. Then by the spectral theorem

$$T_n = \lambda_1 v_1 v_1^t + \dots + \lambda_n v_n v_n^t.$$

Strictly speaking, we should put the subscript n in λ_ℓ, v_ℓ etc., but for simplicity of notation we don't.

Histogram of the eigenvalues: The eigenvalues are all between -2 and $+2$. What proportion are inside an interval $[a, b] \subseteq [-2, 2]$? A calculation free way to find this is to note that the points $2e^{i\theta_\ell}$, $1 \leq \ell \leq n$, are equispaced points on the top half of the circle of radius 2 centered at the origin in the complex plane. Our eigenvalues are just the real parts of these points. Thus the proportion of eigenvalues in $[a, b]$ must converge to the normalized length of the circular arc between the lines $x = a$ and $x = b$. This we calculate as $\frac{1}{\pi}(\arcsin(b/2) - \arcsin(a/2))$. In other words, as $n \rightarrow \infty$, the histogram of eigenvalues approximates the arcsine measure whose distribution function is $\frac{1}{\pi}(\arcsin(x/2) + \frac{\pi}{2})$ with \arcsin taking values in $(-\pi/2, \pi/2)$. By differentiating, we get the density $\rho(x) = \frac{1}{\pi\sqrt{4-x^2}}$.

Spacing between eigenvalues in the bulk: Observe that

$$\begin{aligned} \lambda_\ell - \lambda_{\ell+1} &= 2\cos\left(\frac{2\pi\ell}{n+1}\right) - 2\cos\left(\frac{2\pi(\ell+1)}{n+1}\right) \\ &= \frac{4\pi}{n+1} \sin(\theta_\ell^*) \end{aligned}$$

for some $\theta_\ell^* \in (\theta_\ell, \theta_{\ell+1})$. This is of order at most $1/n$, which makes sense because if n eigenvalues are packed into an interval of length 4, many of the successive differences must be below $1/n$.

More precisely, suppose ℓ is such that λ_ℓ is close to a point $x \in (-2, 2)$. This means that $\cos(\theta_\ell^*)$ is close to $x/2$ and hence $\sin(\theta_\ell^*)$ is close to $\frac{1}{2}\sqrt{4-x^2}$. Hence,

$$\lambda_{\ell+1} - \lambda_\ell \approx \frac{2\pi}{(n+1)} \sqrt{4-x^2}.$$

In words, the eigenvalues near x look like $\frac{2}{n\rho(x)}\mathbb{Z}$, the integer lattice with spacing $\frac{2}{n\rho(x)}$.

The factor $1/n$ makes sense as n eigenvalues are packed into an interval of length 4. The factor $1/\rho(x)$ makes sense because the lower the density, farther apart are the eigenvalues. But interestingly, other than the scaling which depends on $\rho(x)$, the structure of the eigenvalues is the same everywhere, namely like \mathbb{Z} .

Spacing of eigenvalues at the edge: The previous calculations are not right if $x = \pm 2$ (for example, $\rho(2)$ is not defined). To see how eigenvalues are spaced at the right edge, let us consider a fixed ℓ , say $\ell = 1$. Using the Taylor expansion $\cos t = 1 - \frac{1}{2}t^2 + o(t^2)$ as $t \rightarrow 0$, we get

$$\lambda_1 - \lambda_2 = 2 \cos\left(\frac{2\pi}{n+1}\right) - 2 \cos\left(\frac{4\pi}{n+1}\right) \approx \frac{12\pi^2}{(n+1)^2}.$$

Thus, consecutive eigenvalues are only $1/n^2$ distance apart. This too makes sense because the arcsine density blows up near 2, hence the eigenvalues must be crowded more than in the interior $(-2, 2)$. More generally, if ℓ is fixed, then

$$\lambda_\ell - \lambda_{\ell+1} \approx \frac{4\pi^2(2\ell+1)}{(n+1)^2}.$$

This holds for $\ell = 0$ too, if we set $\lambda_0 = 2$. These conclusions may also be written as

$$n^2(2 - \lambda_1, 2 - \lambda_2, 2 - \lambda_3, \dots) \rightarrow 4\pi^2(1, 4, 9, 16, \dots).$$

Empirical spectral distribution and related notions

Much of this course will be asking the same three questions for eigenvalues of various random matrices. For this we introduce some terms.

Empirical distribution of eigenvalues: The *empirical spectral distribution* of an $n \times n$ matrix A which is the probability measure (on \mathbb{R} or \mathbb{C}) defined as

$$L_A = \frac{1}{n} \sum_{k=1}^n \delta_{\lambda_k},$$

where $\lambda_1, \dots, \lambda_n$ are the eigenvalues of A counted with multiplicity. If A_n is a sequence of $n \times n$ matrices and it so happens that $L_{A_n} \xrightarrow{d} \mu$ for some probability measure μ , then we say that μ is the *limiting spectral distribution* of the given sequence.




Figure 1.2: . Histogram of the eigenvalues of the 400×400 matrix. We almost see the arcsine density. The density is higher at the edges.




Figure 1.3: . The top picture shows the 12 largest eigenvalues. The bottom picture shows 12 eigenvalues close to 0. The top picture spans a length of 0.01 while the bottom picture spans a length of 0.18. The crowding of eigenvalues at the edge can also be seen in the varying length of spacings in the top picture (in fact the spacings are in arithmetic progression). In the bottom picture, the eigenvalues appear equispaced like a lattice.

These notions are applicable to random matrices too. If A is random, L_A will be a random probability measure. Define $\bar{L}_n(A) = \mathbf{E}[L_n(A)]$ for any Borel set A . Then \bar{L}_n is a (non-random) probability measure. We call it the *expected empirical distribution* of A .

More discussion of the notion of convergence in distribution will come later. For now, we simply note that the empirical distribution is a very reasonable way to talk about the eigenvalue of a matrix. Other possible ways to talk about eigenvalues are

1. As a vector $(\lambda_1, \dots, \lambda_n) \in \mathbb{R}^n$. This has the drawback that we must label the eigenvalues to say which is first, which is second etc., while the eigenvalues themselves come as a set (or multi-set) without any ordering.

2. As a set or a multi-set. Equivalently, as an element of \mathbb{R}^n/S_n , the quotient of \mathbb{R}^n under the action of permutations of co-ordinates. This avoids the question of labeling eigenvalues. However, when we want to consider a sequence of matrices A_n , the spaces in which we view the eigenvalues change with n . This is not convenient if we want to talk of limits.

By putting together all the information about the eigenvalues in the measure L_A , we avoid labeling individual eigenvalues and further, irrespective of the size of the matrix, the empirical measure is in the same space $\mathcal{P}(\mathbb{R})$ or $\mathcal{P}(\mathbb{C})$ (space of probability measures on \mathbb{R} or \mathbb{C}). These spaces are complete, separable metric spaces, which is the standard setting in probability theory when one wants to study weak convergence of probability measures.

Point process limits: The empirical distribution puts mass $1/n$ at each eigenvalue, and hence it is not the right object to consider when we are studying the largest eigenvalues or the spacings. In such cases, we go back to writing the eigenvalues as a vector in $\mathbb{R}^{\mathbb{N}}$ or $\mathbb{R}^{\mathbb{Z}}$ and talk about convergence of any finite number of co-ordinates. For example, when studying the edge, we considered the vector

$$n^2(2 - \lambda_1, 2 - \lambda_2, \dots).$$

In the limit, we get only the largest eigenvalues remain and the rest fly off to infinity. As we saw, we got $4\pi^2(1, 2, 3, \dots)$ in the special case above.

When studying in the bulk, we consider

$$n(\lambda_1 - x, \lambda_2 - x, \dots, \lambda_n - x)$$

but shift this vector so that the 0th co-ordinate is $n(\lambda_\ell - x)$ with ℓ such that λ_ℓ is the closest eigenvalue to x . In the limit we get the doubly infinite sequence \mathbb{Z} (with a scale factor and a shift).

An equivalent way to state these results are to consider the counting measure $\tilde{L}_A = nL_A = \sum_{k=1}^n \delta_{\lambda_k}$. Then the above statement can be written as $\tilde{L}(x + \frac{1}{n}A) \rightarrow \delta_{\sigma\mathbb{Z}+b}(A)$ for any compact $A \subseteq \mathbb{R}$ and $\delta_{\sigma\mathbb{Z}+b}$ is the counting measure on the set $\sigma\mathbb{Z} + b$. When we deal with random matrices, we shall be considering counting measure of a random discrete set of points, which is what is called a *point process* in probability.




Figure 1.4: Histogram of the eigenvalues of the 400×400 matrix. We see that it is approximately a semi-circle density. In contrast with arc-sine, it vanishes at the edge.

The oscillator matrix

Let $a_k = 0$ and $b_k = \sqrt{k}$. Thus,

$$T_n = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & \sqrt{2} & \ddots & 0 & 0 \\ 0 & \sqrt{2} & \ddots & \ddots & \ddots & 0 \\ 0 & \ddots & \ddots & \ddots & \sqrt{n-2} & 0 \\ 0 & 0 & \ddots & \sqrt{n-2} & 0 & \sqrt{n-1} \\ 0 & 0 & 0 & 0 & \sqrt{n-1} & 0 \end{bmatrix}_{n \times n} \quad (2)$$

In this case, the eigenvalues do not have a simple formula like for the discrete Laplacian matrix, but they may be described as the zeros of the n th Hermite polynomial. Recall that Hermite polynomials are got by orthogonalizing monomials $1, x, x^2, \dots$ in the L^2 space with respect to the Gaussian measure on the line. In general, Jacobi matrices are closely related to orthogonal polynomials. See Exercise 2.

Returning to the oscillator matrix (2), figures 1.4 and 1.5 show that there are interesting things to prove.

In contrast with the earlier example, observe that the range of eigenvalues is about $[-40, 40]$. In fact, drawing the histograms for other sizes of the matrix shows that the




Figure 1.5: The top picture shows the 12 largest eigenvalues. The bottom picture shows 12 eigenvalues close to 0. The top picture spans a length of 5 while the bottom picture spans a length of 2. Eigenvalues are more crowded in the center than at the edge. This can also be seen in the varying length of spacings in the top picture. In the bottom picture, they appear equispaced.

range is $[-2\sqrt{n}, 2\sqrt{n}]$. Hence, the more appropriate thing to consider is the matrix

$$\frac{1}{\sqrt{n}}T_n = \begin{bmatrix} 0 & \frac{1}{\sqrt{n}} & 0 & 0 & 0 & 0 \\ \frac{1}{\sqrt{n}} & 0 & \frac{\sqrt{2}}{\sqrt{n}} & \ddots & 0 & 0 \\ 0 & \frac{\sqrt{2}}{\sqrt{n}} & \ddots & \ddots & \ddots & 0 \\ 0 & \ddots & \ddots & \ddots & \frac{\sqrt{n-2}}{\sqrt{n}} & 0 \\ 0 & 0 & \ddots & \frac{\sqrt{n-2}}{\sqrt{n}} & 0 & \frac{\sqrt{n-1}}{\sqrt{n}} \\ 0 & 0 & 0 & 0 & \frac{\sqrt{n-1}}{\sqrt{n}} & 0 \end{bmatrix}_{n \times n}. \quad (3)$$

Then the histogram of eigenvalues of T_n/\sqrt{n} appears to converge to a density on $[-2, 2]$. We shall show that this is indeed the case, and that the limiting spectral density of T_n/\sqrt{n} is the *semi-circle density* $\rho(x) = \frac{1}{2\pi}\sqrt{4-x^2}$.

The spacing of eigenvalues in the bulk is again like the integer lattice. Indeed, we should expect that near a point $x \in (-2, 2)$, the eigenvalues of T_n/\sqrt{n} look like the lattice $\frac{1}{n\rho(x)}\mathbb{Z}$.

The edge is quite interesting. A heuristic calculation may be made as follows. The number of eigenvalues in $[b, 2]$ is approximately $n \int_b^2 \rho(x) dx$, for fixed $b < 2$, by the definition of the limiting spectral distribution. Without justification, we hope that the first eigenvalue will be at b_n , where $n \int_{b_n}^2 \rho(x) dx = 1$. Writing $\rho(x) \approx \frac{1}{\pi}\sqrt{2-x}$ for x close to 2, we see that $2 - b_n \approx \frac{3\pi}{2}n^{-2/3}$. Similarly for $2 - \lambda_2, 2 - \lambda_3$, etc. and we expect to get

$$\frac{2}{3\pi}n^{2/3}(2 - \lambda_1, 2 - \lambda_2, \dots) \rightarrow (1, 2^{2/3}, 3^{2/3}, \dots).$$

Let us emphasize again that we have not actually proved anything about the eigenvalues




Figure 1.6: Histogram of the eigenvalues of $T_n(f)$ for (a) $f(x) = x$, (b) $f(x) = \sqrt{x(1-x)}$ (c) $f(x) = (1-x)\log(1-x)$.

of the oscillator matrix. We just have gathered reasonable evidence, from the pictures and heuristics.

A class of Jacobi matrices

Motivated by the previous two examples (particularly the scaling that was required in the second), we define a class of Jacobi matrices as follows. Let $f : [0, 1] \mapsto \mathbb{R}$ be a function. Define $T_n(f)$ as the $n \times n$ Jacobi matrix with $a_k = 0$ and $b_k = f(k/n)$. The discrete Laplacian matrix is the case $f = 1$ and the matrix (3) is the case $f(x) = \sqrt{x}$. We shall later show that all such matrices have a limiting spectral distribution.

Limiting spectral distributions: In Figure 1.6 are a few pictures to show that the his-

Top: $f(x)=x^{1/4}$. Middle: $f(x)=x^{1/2}$. Bottom: $f(x)=x$.




Figure 1.7: Eigenvalues at the edge for several different f with 400×400 matrices. In all three cases, the limiting spectral density has square-root vanishing at the edge. In the picture it appears plausible that the spacings might be exactly the same in the limit (a scaling has been done to match the first spacing length though).

tograms do have limits for many different f . We shall prove later that this is indeed the case, and in fact give an explicit formula to calculate the limiting density in terms of f . For example, when $f(x) = \sqrt{x}$, the limiting spectral density does turn out to be the semi-circle density $\frac{1}{2\pi}\sqrt{4-x^2}$ (on $[-2, 2]$) and when $f(x) = \sqrt{x(1-x)}$, the limiting spectral density is uniform on $[-1, 1]$.

A new idea is required for the proof, since unlike for the discrete Laplacian, in most cases we do not have explicit exact formulas for the eigenvalues of T_n .

Spacings in the bulk: If we zoom in at any point x where the limiting density $\rho(x)$ is strictly positive, it appears that we just see equispaced points with spacings of $1/n\rho(x)$.

Spacings at the edge: Look at Figure 1.7. It appears that when we zoom in at the right end of the support, then the the points are not equispaced, and further, the inter-point spacings are not always the same. However, the inter-point spacing appear to be depend only on the behaviour of the density ρ near the edge of the support. This is partly justified by the heuristic we gave for the $n^{-2/3}$ in case of the oscillator matrix eigenvalues. If ρ looks like $C(E-x)^\alpha$ at the edge E , then $E - \lambda_1 \asymp n^{-1/(1+\alpha)}$.

Universality: The observations about the spacings in the edge and bulk look reasonable. But at the moment I don't know if they are entirely correct or if they are proved in the (very likely a full answer is in the papers of Bálint Virág). But the lesson here motivates a

discussion of *universality*, a notion driving a good part of recent research in probability.

As an analogy, let us return to the best understood situation in probability, sums of random variables, $S_n = X_1 + \dots + X_n$. If X_i are i.i.d. with mean μ and variance σ^2 , then S_n is - (A) constant on the scale of n , i.e., $n^{-1}S_n$ converges to a constant μ , (B) a $N(\mu, \sigma^2)$ random variable on the scale of \sqrt{n} and (C) dependent of the full distribution on unit scale. The most interesting behaviour is the intermediate one, where there is randomness in the limit (in this case Gaussian), but it is universal, remembering nothing of the details of the distribution of X_1 except for the two parameters μ and σ^2 .

In a similar way, the limiting spectral distribution for $T_n(f)$ is like the law of large numbers, the observation that the spacings in the bulk and at the edge do not depend on f (or only on a small number of parameters such as the behaviour of ρ near the edge) are like the central limit theorem. When we go to random matrices, the limiting spectral distribution will be non-random usually, while the spacings in the bulk and edge have *universal* distributions described by only a few parameters of the original model. The exact eigenvalue distribution of course depends on the precise model (like the unit scale behaviour for sums).

Exercises

Exercise 1. Let $T_n = T_n(a, b)$ and denote by T_k the top-left $k \times k$ principal submatrix of T_n . Let $\phi_k(z) = \det(zI_k - T_k)$ denote the characteristic polynomial of T_k . Define $\phi_0(z) = 1$ and $\phi_{-1} = 0$. Show that

$$\phi_k(z) = (z - a_k)\phi_{k-1}(z) - b_{k-1}^2\phi_{k-2}(z), \quad \text{for } 1 \leq k \leq n.$$

Define $\psi_k(z) = \frac{1}{b_1 \dots b_{k-1}}\phi_k(z)$ for $0 \leq k \leq n-1$. If λ is any eigenvalue of T_n , show that

$$(\psi_0(\lambda), \psi_1(\lambda), \dots, \psi_{n-1}(\lambda))^t$$

is a corresponding eigenvector. [Note: We shall show later that all eigenvalues of a Jacobi matrix are simple, hence this is the unique eigenvector, up to scaling].

Exercise 2. In the following exercises, use the recursion obtained in the previous exercise.

1. If T_n is the modification of the discrete Laplacian matrix by setting $a_1 = 1$ (and $a_k = 0$ for $k \geq 2$, $b_k = 1$ for all k), then show that $\phi_k(\cos \theta) = 2^{-k+1} \cos(k\theta)$.
2. If T_n is the oscillator matrix from (2), show that $\phi_k(x) = (-1)^k e^{x^2/2} \frac{d^k}{dx^k} e^{-x^2/2}$.

[*Remark:* The polynomial defined by $p_n(\cos \theta) = \cos(n\theta)$ is called the n th *Chebyshev polynomial of the first kind*. The polynomial $(-1)^k e^{x^2/2} \frac{d^k}{dx^k} e^{-x^2/2}$ is called the k th *Hermite polynomial*.]

Exercise 3. Using the notation of Exercise 1, show that the eigenvalues of T_k are k distinct real numbers and strictly interlace with those of T_{k-1} . That is, if $\lambda_1^{(k)} > \lambda_2^{(k)} > \dots > \lambda_k^{(k)}$, are the eigenvalues of T_k , then show that

$$\lambda_k^{(1)} > \lambda_1^{(k-1)} > \lambda_k^{(2)} > \dots > \lambda_{k-1}^{(k-1)} > \lambda_k^{(k)}.$$

[*Hint:* Put z to be a zero of φ_{k-1} in the three term recurrence for φ_k in terms of φ_{k-1} and φ_{k-2} .]

Exercise 4. Let G_n be the subgraph of \mathbb{Z}^2 with vertex set $[m] \times [n]$. Let $\Delta_{m,n}$ be its discrete Laplacian. Show that the limiting spectral distribution of $\Delta_{[na], [nb]}$ as $n \rightarrow \infty$ is $\mu_a \star \mu_b$ where μ_a is the arcsine measure on $[-2a, 2a]$.

Chapter 2

The simplest non-trivial random matrix

We asserted that Jacobi matrices are the simplest non-trivial matrices. Now we inject randomness into the picture.

Question 5. If a_k s and b_k s are random variables with some joint distribution, what will be the distribution of eigenvalues of $T_n(a, b)$?

If we find a matrix for which the joint distribution of eigenvalues is explicit (and tractable!), that would be the closest analogue to the deterministic situation of the discrete Laplacian matrix where we could compute the eigenvalues exactly. Just as the exact formulas for eigenvalues allowed us to find the limiting spectral distribution, spacings in the edge and the bulk, etc., one hopes that having explicit density for the eigenvalues would help to answer the same questions for a random matrix. In this chapter we shall find one such random Jacobi matrix. However, answering the three questions is far from trivial even after we have the exact density and it will be done gradually over the course.

Given a distribution of a_k s and b_k s, how to find the joint density of eigenvalues? In principle, it is just a matter of change of variables, but the execution requires work.

Parameterizing a Jacobi matrix by its spectral measure at e_1

Let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ denote the eigenvalues of T_n . We shall see shortly that they are distinct. Since T_n has $2n - 1$ variables, we need $n - 1$ additional auxiliary variables to complete the change of variables exercise. These are chosen as follows. Let $T_n = \lambda_1 v_1 v_1^t + \dots + \lambda_n v_n v_n^t$ be the spectral decomposition of T_n with some choice of orthonormal eigenvectors

v_1, \dots, v_n . Then for any integer $m \geq 0$

$$\langle T_n^m e_1, e_1 \rangle = \sum_{k=1}^n \lambda_k^m p_k, \quad \text{where } p_k = |\langle v_k, e_1 \rangle|^2. \quad (1)$$

As the eigenvectors form an orthonormal basis, and $\|e_1\| = 1$, we get that $p_1 + \dots + p_n = 1$. Thus, we may take $p = (p_1, \dots, p_{n-1})$ is as the auxiliary variables.

The concept behind this choice: If we write $T_n = WDW^t$ where W is orthogonal and $D = \text{diag}(\lambda_1, \dots, \lambda_n)$, then p_k s are just the squares of the entries in the first row of W . More conceptually, the identity (1) can be written as

$$\langle f(T)e_1, e_1 \rangle = \int f(x) d\nu_T(x)$$

where $\nu_T = \sum_{k=1}^n p_k \delta_{\lambda_k}$ is called the *spectral measure* of T_n at the vector e_1 . And f is any polynomial, but by the usual symbolic calculus (wherein $f(T)$ is defined to be $Wf(D)W^t$), the identity continues to hold for arbitrary functions f . In particular, for any $z \in \mathbb{H}$, we have

$$(zI_n - T_n)^{1,1} = \sum_{k=1}^n \frac{p_k}{z - \lambda_k}$$

where $B^{1,1} = \langle B^{-1}e_1, e_1 \rangle$ denotes the $(1, 1)$ entry of B^{-1} . The condition $z \in \mathbb{H}$ ensures that the inverse on the left exists.

The Jacobian determinant

The key point is the computation of the Jacobian determinant of the transformation that maps (a, b) to (λ, p) . Let us be more precise about the spaces involved.

The set of all $n \times n$ Jacobi matrices is naturally identified, via the parameters (a, b) , with $\mathcal{J}_n := \mathbb{R}^n \times \mathbb{R}_+^{n-1}$ where $\mathbb{R}_+ = (0, \infty)$. Next, define

$$\Delta_n = \{p \in \mathbb{R}^{n-1} : p_1 + \dots + p_{n-1} < 1, p_i > 0\}, \quad \text{and} \quad \mathbb{R}_+^n = \{x \in \mathbb{R}^n : x_1 > x_2 > \dots > x_n\}.$$

The set of all probability measures on \mathbb{R} whose support has exactly n points is naturally identified with $\mathcal{M}_n := \mathbb{R}_+^n \times \Delta_n$ by identifying $\sum_{k=1}^n p_k \delta_{\lambda_k}$ with (λ, p) . Here and everywhere below, wherever p_n appears, it is to be regarded as a short form for $1 - (p_1 + \dots + p_{n-1})$.

Lemma 6. Fix $n \geq 1$ and let $G : \mathcal{J}_n \mapsto \mathcal{M}_n$ be defined as $G(T) = \nu_T$.

(a) G is a bijection from \mathcal{J}_n onto \mathcal{M}_n .

(b) If $T = T_n(a, b)$ and $\mathbf{v}_T = \sum_{k=1}^n p_k \delta_{\lambda_k}$, then

$$\prod_{k=1}^{n-1} b_k^{2(n-k)} = \prod_{k=1}^n p_k \cdot \prod_{i < j} |\lambda_i - \lambda_j|^2. \quad (2)$$

(c) The Jacobian determinant of G^{-1} is equal to (up to a sign that depends on the ordering of variables)

$$J_{G^{-1}}(\lambda, p) = \frac{\prod_{k=1}^n p_k \prod_{i < j} |\lambda_i - \lambda_j|^4}{2^{n-1} \prod_{k=1}^{n-1} b_k^{4(n-k)-1}} = \frac{\prod_{k=1}^{n-1} b_k}{2^{n-1} \prod_{k=1}^n p_k}. \quad (3)$$

We postpone the proof of this Lemma and first talk about some consequences.

A class of random Jacobi matrices

If (a, b) has joint density $f(a, b)$ with respect to Lebesgue measure on $\mathbb{R}^n \times \mathbb{R}_+^n$, then by the change of variable formula, the density of (λ, p) with respect to Lebesgue measure on $\mathbb{R}_+^n \times \Delta_n$ is given by

$$g(\lambda, p) = f(a, b) \frac{\prod_{k=1}^{n-1} b_k}{2^{n-1} \prod_{k=1}^n p_k}.$$

On the right, (a, b) is written as a short form for the image of (λ, p) under the bijection from \mathcal{M}_n to \mathcal{T}_n . We would like to choose a density f that has a nice form and for which $g(\lambda, p)$ also has a nice form. Or at least so that the marginal density of λ given by $\int_{\Delta_n} g(\lambda, p) dp$ is nice. Here is such a choice:

$$f(a, b) = \frac{1}{Z_{\beta, n}} \exp \left\{ -\frac{1}{4} \left[\sum_{k=1}^n a_k^2 + 2 \sum_{k=1}^{n-1} b_k^2 \right] \right\} \prod_{k=1}^{n-1} b_k^{\beta(n-k)-1} \quad (4)$$

where $\beta > 0$ and $Z_{\beta, n}$ is the normalizing constant. The logic behind this choice is as follows: The factors $\prod_{k=1}^{n-1} b_k^{\beta(n-k)-1}$ partly cancels the similar factor in the Jacobian determinant, and the remaining expression can be written in terms of (λ, p) by (2). The choice of the

exponent (something is required to make it integrable) is because it can be written nicely in terms of the eigenvalues:

$$\sum_{k=1}^n a_k^2 + 2 \sum_{k=1}^{n-1} b_k^2 = \text{tr}(T^2) = \sum_{k=1}^n \lambda_k^2.$$

Thus, we arrive at

$$g(\lambda, p) = \frac{1}{Z'_{n,\beta}} \exp \left\{ -\frac{1}{4} \sum_{k=1}^n \lambda_k^2 \right\} \prod_{i < j} |\lambda_i - \lambda_j|^\beta \cdot \prod_{k=1}^n p_k^{\frac{\beta}{2}-1}. \quad (5)$$

Here $Z'_{n,\beta}$ is the normalization constant so that g is a probability density. Both the densities $f(a, b)$ and $g(\lambda, p)$ are remarkably nice.

- Writing

$$f(a, b) = \frac{1}{Z_{\beta,n}} \prod_{k=1}^n e^{-\frac{1}{4}a_k^2} \prod_{k=1}^{n-1} e^{-\frac{1}{2}b_k^2} b_k^{\beta(n-k)-1},$$

we see that a_k s and b_k s are independent, $a_k \sim N(0, 2)$ and $b_k^2 \sim \chi_{\beta(n-k)}^2$. Recall that χ_m^2 is the distribution of the sum of squares of m independent standard normal variables. It is the same as Gamma($m/2, 1/2$) distribution and the density is

$$\frac{1}{\Gamma(m/2)2^m} e^{-x/2} x^{m-1} \text{ for } x > 0.$$

- Under $g(\lambda, p)$, the vector λ is independent of the vector p . The density of λ is proportional to

$$\exp \left\{ -\frac{1}{4} \sum_{k=1}^n \lambda_k^2 \right\} \prod_{i < j} |\lambda_i - \lambda_j|^\beta$$

and the density of p is proportional to $\prod_{k=1}^n p_k^{\frac{\beta}{2}-1}$ for $p \in \Delta_n$. The latter is the well-known Dirichlet distribution¹ with parameters n and $(\beta/2, \dots, \beta/2)$.

We summarize the main conclusion in the following theorem.

¹The Dirichlet distribution with parameters n and $\beta_1, \dots, \beta_{n-1}$ is the density on Δ_n equal to $C_n p_1^{\beta_1-1} \dots p_n^{\beta_n-1}$. When $n = 2$, this is just the Beta distribution with parameters β_1, β_2 . Similar to that case, the normalizing constant for the Dirichlet density is given by $C_n = \frac{\Gamma(\beta_1 + \dots + \beta_n)}{\Gamma(\beta_1) \dots \Gamma(\beta_n)}$.

Theorem 7. Let a_{ks} be independent $N(0, 2)$ random variables and let b_k^2 s be independent $\chi_{\beta(n-k)}^2$ variables also independent of the a_{ks} . Then the eigenvalues of the Jacobi matrix $T_n(a, b)$ have density

$$\frac{1}{Z''_{\beta, n}} \exp \left\{ -\frac{1}{4} \sum_{k=1}^n \lambda_k^2 \right\} \prod_{i < j} |\lambda_i - \lambda_j|^\beta. \quad (6)$$

with respect to Lebesgue measure on \mathbb{R}_+^n .

Owing to the appearance of the Gaussian factor, the density (6) is called the *beta Hermite ensemble*.

Remark 8. In summary, we have found a one-parameter family of random Jacobi matrices whose eigenvalue densities are that of the beta-log gas on the real line. Can we answer the questions about the limiting spectral distribution or the spacings of eigenvalues in the bulk and at the edge? While the joint density does, in principle, contain the answers to these questions, it is not a trivial task to pry that kind of information from it (but it has been done). Pleasantly, the Jacobi matrix actually helps in finding answers to these questions about the eigenvalues. In the next chapter, we shall show two such applications of the Jacobi matrix.

Computation of the Jacobian determinant

It only remains to prove Lemma 6. We will present the proof with minimal digressions. But what is embedded here is a deep connection between Jacobi matrices, probability measures and orthogonal polynomials. This remark is explained briefly after the proof.

Proof of the second part (2). Let $T_n = T_n(a, b)$ be a Jacobi matrix. Let T_k (respectively \tilde{T}_k) denote the top-left (respectively bottom-right) $k \times k$ principal submatrix of T . Let φ_k (respectively $\tilde{\varphi}_k$) denote the characteristic polynomial of T_k , i.e., $\varphi_k(z) = \det(zI_k - T_k)$. Let $\lambda_j^{(k)}$, $1 \leq j \leq k$ denote the zeros of φ_k , or in other words, the eigenvalues of T_k .

Expanding w.r.t. the last row, we get the recursion

$$\varphi_k(z) = (z - a_k)\varphi_{k-1}(z) - b_{k-1}^2\varphi_{k-2}(z). \quad (7)$$

which is valid also for $k = 1$ and $k = 0$ provided we set $\varphi_0 = 1$ and $\varphi_{-1} = 0$. From Exercise 3 we know that the eigenvalues of T_k are distinct and strictly interlace with those of T_{k-1} . If you did not do that exercise, you may see it by observing that φ_k and φ_{k-2} have opposing

signs at the zeros of φ_{k-1} . Inductively, if the interlacing is assumed for φ_{k-1} and φ_{k-2} , then the interlacing follows for φ_k and φ_{k-1} .

Now put $z = \lambda_j^{(k-1)}$ in (7) and multiply over $j \leq k-1$ to get

$$\prod_{j=1}^{k-1} \varphi_k(\lambda_j^{(k-1)}) = (-1)^{k-1} b_{k-1}^{2(k-1)} \prod_{j=1}^{k-1} \varphi_{k-2}(\lambda_j^{(k-1)}).$$

Now, for any two monic polynomials $P(z) = \prod_{j=1}^p (z - \alpha_j)$ and $Q(z) = \prod_{j=1}^q (z - \beta_j)$,

$$\prod_{j=1}^q P(\beta_j) = \pm \prod_{j=1}^p Q(\alpha_j)$$

since both are equal (up to sign) to $\prod_i \prod_j (\alpha_i - \beta_j)$. Use this for φ_k and φ_{k-1} to get

$$\prod_{j=1}^k \varphi_{k-1}(\lambda_j^{(k)}) = \pm b_{k-1}^{2(k-1)} \prod_{j=1}^{k-1} \varphi_{k-2}(\lambda_j^{(k-1)}).$$

Take product over k and telescope to get (we write λ_j for $\lambda_j^{(n)}$)

$$\prod_{j=1}^n \varphi_{n-1}(\lambda_j) = \pm \prod_{j=1}^{n-1} b_j^{2j}.$$

Clearly this can be done in reverse for the $\tilde{\varphi}_k$ s to get

$$\prod_{j=1}^n \tilde{\varphi}_{n-1}(\lambda_j) = \pm \prod_{j=1}^{n-1} b_j^{2(n-j)}. \quad (8)$$

The spectral measure is related to $\tilde{\varphi}_{n-1}$ as follows.

$$\sum_{k=1}^n \frac{p_k}{z - \lambda_k} = (zI - T)^{1,1} = \frac{\tilde{\varphi}_{n-1}(z)}{\varphi_n(z)}.$$

Multiply by $z - \lambda_j$ and let $z \rightarrow \lambda_j$ to get $p_j = \tilde{\varphi}_{n-1}(\lambda_j) / \varphi_n'(\lambda_j)$. Multiply and use (8) to get

$$\begin{aligned} \prod_{j=1}^{n-1} b_j^{2(n-j)} &= \pm \prod_{j=1}^n p_j \prod_{j=1}^n \varphi_n'(\lambda_j) \\ &= \prod_{j=1}^n p_j \prod_{i < j} |\lambda_i - \lambda_j|^2 \end{aligned}$$

since $\varphi_n'(\lambda_j) = \prod_{i \neq j} (\lambda_j - \lambda_i)$. In the end, both sides are positive, so we did not have to follow the sign. This proves (2). ■

Proof of the first part. On the way, we have proved one side of the first part of Lemma 6 too. Indeed, if $T \in \mathcal{J}_n$, then we have noted the distinctness of eigenvalues. Further, $p_j = \tilde{\Phi}_{n-1}(\lambda_j)/\Phi'_n(\lambda_j)$ which cannot be zero because of the strict interlacing of eigenvalues of T_n and \tilde{T}_{n-1} . Thus, v_T belongs to \mathcal{M}_n . This shows that G maps \mathcal{J}_n into \mathcal{M}_n .

To prove the converse, start with a measure $\nu = \sum_{j=1}^n p_j \delta_{\lambda_j} \in \mathcal{M}_n$. Observe that $L^2(\nu)$ has dimension exactly equal to n and that $1, x, \dots, x^{n-1}$ are linear independent in $L^2(\nu)$. Therefore, we may apply Gram-Schmidt procedure to get $\psi_0, \dots, \psi_{n-1}$, where ψ_j is a polynomial with degree j . Fix some k and expand $x\psi_k(x)$ in this orthonormal basis to write (note that there is no ψ_n)

$$\begin{aligned} x\psi_k(x) &= c_{k,k+1}\psi_{k+1}(x) + \dots + c_{k,0}\psi_0(x) \quad \text{for } k \leq n-2, \\ x\psi_{n-1}(x) &= c_{n,n-1}\psi_{n-1}(x) + \dots + c_{n,0}\psi_0(x). \end{aligned}$$

For $k \leq n-2$, observe that $c_{k,k+1}$ is strictly positive, since ψ_k and ψ_{k+1} both have strictly positive leading coefficients. Further, observe that $\langle x\psi_k(x), \psi_j(x) \rangle = \langle \psi_k(x), x\psi_j(x) \rangle$ which is zero if $j < k-1$ as ψ_k is orthogonal to all polynomials of degree lower than k . That leave

$$c_{k,k+1} = \int x\psi_k(x)\psi_{k+1}(x)d\nu(x), \quad c_{k,k} = \int x\psi_k^2(x)d\nu(x).$$

From this it is clear that $c_{k,k+1} = c_{k+1,k}$ for $k \leq n-1$. Set $a_k = c_{k-1,k-1}$, $1 \leq k \leq n$ and $b_k = c_{k-1,k}$, $1 \leq k \leq n-1$. We have already shown that $b_k > 0$ for all $k \leq n-1$. Thus, if we define $H(\nu)$ to be the Jacobi matrix $T = T_n(a, b)$, then H maps \mathcal{M}_n into \mathcal{J}_n .

With all this, the recursions are now written as

$$\begin{aligned} x\psi_k(x) &= b_{k-1}\psi_{k-1}(x) + a_k\psi_k(x) + b_k\psi_{k+1}(x), \quad \text{for } k \leq n-2, \\ x\psi_{n-1}(x) &= b_{n-1}\psi_{n-2}(x) + a_n\psi_{n-1}(x). \end{aligned}$$

The equalities are in $L^2(\nu)$, meaning that it holds for $x \in \{\lambda_1, \dots, \lambda_n\}$ (the first line of identities then extends to all x by considering the degree, but the last one cannot possibly!). In short, the above equations are saying that T_n has eigenvalues λ_j with eigenvector

$$v_j = \sqrt{p_j}(\psi_0(\lambda_j), \dots, \psi_{n-1}(\lambda_j))^t.$$

We have introduced the factor $\sqrt{p_j}$ because then the rows of the matrix $[v_1 \ v_2 \ \dots \ v_n]$ become orthonormal. As $\psi_0 = 1$, we get $|v_j(1)|^2 = p_j$ and hence the spectral measure at e_1 is $\sum_{j=1}^n p_j \delta_{\lambda_j} = \nu$. Thus, $G \circ H$ is the identity map from \mathcal{M}_n into itself. In particular, G maps \mathcal{J}_n onto \mathcal{M}_n .

The proof will be complete if we show that G is one-one, which can be done in many ways. We simply refer to the equations (9) from which it is clear that if we know (λ, p) , then we can successively recover $a_1, b_1, a_2, b_2, \dots$

This completes the proof that the map G from \mathcal{J}_n to \mathcal{M}_n is a bijection. \blacksquare

Proof of the third part (3). It remains to prove the formula for the Jacobian determinant. Let $T_n = T_n(a, b) \in \mathcal{J}_n$ correspond to $v = \sum_{j=1}^n p_j \delta_{\lambda_j} \in \mathcal{M}_n$. We write the identities $(T_n^m)_{1,1} = \sum_{j=1}^n p_j \lambda_j^m$ for $m = 1, 2, \dots, 2n-1$.

$$\begin{aligned} \sum p_j \lambda_j &= T_{1,1} = a_1 & \sum p_j \lambda_j^2 &= (T^2)_{1,1} = b_1^2 + [\dots] \\ \sum p_j \lambda_j^3 &= (T^3)_{1,1} = a_2 b_1^2 + [\dots] & \sum p_j \lambda_j^4 &= (T^4)_{1,1} = b_2^2 b_1^2 + [\dots] \\ \sum p_j \lambda_j^5 &= (T^5)_{1,1} = a_3 b_2^2 b_1^2 + [\dots] & \sum p_j \lambda_j^6 &= (T^6)_{1,1} = b_3^2 b_2^2 b_1^2 + [\dots] \\ & \dots & & \dots \end{aligned} \tag{9}$$

Here the $[\dots]$ include many terms, but all the a_k, b_k that appear there have appeared in previous equations. For example, $(T^2)_{1,1} = b_1^2 + a_1^2$ and as a_1 appeared in the first equation, we have brushed it under $[\dots]$ as they will not matter.

Let $u = (u_1, \dots, u_{2n-1})$ where $u_j = (T^j)_{1,1}$. The right hand sides of the above equations express u as $F(a, b)$ while the left hand sides as $u = H(\lambda, p)$. We find the Jacobian determinants of F and H as follows.

Jacobian determinant of F : Note that u_{2k} is a function of $a_i, i \leq k$ and $b_j, j \leq k$ while u_{2k-1} is a function of $a_i, i \leq k$ and $b_j, j \leq k-1$. Thus, ordering (a, b) as $(a_1, b_1, a_2, b_2, \dots, b_{n-1}, a_n)$ and u as u_1, \dots, u_{2n-1} , the derivative matrix of u with respect to a, b becomes upper triangular with determinant

$$J_F(a, b) = 2^{n-1} \prod_{k=1}^{n-1} b_k^{4(n-k)-1}. \tag{10}$$

Jacobian determinant of H : The equations above give the derivative of H to be

$$D_H(\lambda, p) = \begin{bmatrix} p_1 & \dots & p_n & \lambda_1 - \lambda_n & \dots & \lambda_{n-1} - \lambda_n \\ 2p_1 \lambda_1 & \dots & 2p_n \lambda_n & \lambda_1^2 - \lambda_n^2 & \dots & \lambda_{n-1}^2 - \lambda_n^2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ (2n-1)p_1 \lambda_1^{2n-2} & \dots & (2n-1)p_n \lambda_n^{2n-2} & \lambda_1^{2n-1} - \lambda_n^{2n-1} & \dots & \lambda_{n-1}^{2n-1} - \lambda_n^{2n-1} \end{bmatrix}.$$

Let C_i denote the i th column of this matrix. Factor out p_i from the C_i . The resulting matrix is of the same form (as if $p_i = 1$ for all i) and its determinant is clearly a polynomial in

$\lambda_1, \dots, \lambda_n$. It must also symmetric in $\lambda_{k,s}$, because the original problem we started with was symmetric in $\lambda_{k,s}$ (although in the matrix λ_n appears superficially to have a different role).

If $h := \lambda_1 - \lambda_n \rightarrow 0$, then $C_{n+1} = O(h)$, $C_1 - C_n = O(h)$. Further, it is easy to check that $C_{n+1} - h(C_1 + C_2)/2 = O(h^2)$. Thus for fixed $\lambda_k, k \geq 2$, the polynomial in λ_1 has (at least) a four fold zero at λ_n . By symmetry, the determinant has a factor $\Delta(\lambda)^4$. However, the determinant above and $\Delta(\lambda)^4 = \prod_{i < j} (\lambda_i - \lambda_j)^4$ are both polynomials of degree $4(n-1)$. Further, the coefficient of λ_1^{4n-4} in both is the same. Therefore we get the Jacobian determinant

$$J_H(\lambda, p) = \pm |\Delta(\lambda)|^4 \prod_{i=1}^n p_i. \quad (11)$$

From (10) and (11) we deduce that

$$|J_{G^{-1}}(\lambda, p)| = \frac{J_H(\lambda, p)}{J_F(a, b)} = \frac{\prod_{i=1}^n p_i \prod_{i < j} |\lambda_i - \lambda_j|^4}{2^{n-1} \prod_{k=1}^{n-1} b_k^{4(n-k)-1}}.$$

This proves the first equality in (3). The second equality follows by using (2). ■

Remarks on the moment problem and Jacobi matrices

Consider the following four objects.

1. \mathcal{M}_n , the set of probability measures on \mathbb{R} with support of cardinality n .
2. \mathcal{P}_n , the space of positive definite sequences $\alpha = (\alpha_0, \alpha_1, \dots)$ of rank n . This just means that the infinite Hankel matrix $H = (\alpha_{i+j})_{i,j \geq 0}$ is positive semi-definite and has rank n . To be even more explicit, this just means that every principal submatrix of H is positive semi-definite and the maximal rank of a principal submatrix is n .
3. OP_n , the set of sequence of polynomials $(\psi_0, \psi_1, \dots, \psi_{n-1})$ such that ψ_j has degree j and has positive leading coefficient and such that if an inner product is defined on the space of polynomials of degree at most $n-1$ by declaring ψ_j to be orthogonal, then we have $\langle x\psi_j, \psi_k \rangle = \langle \psi_j, x\psi_k \rangle$ for $0 \leq j, k \leq n-2$ (in short, multiplication operator by x is self-adjoint. However, $x\psi_{n-1}$ has degree n , that is why the condition $j, k \leq n-2$).

4. \mathcal{J}_n , the set of $n \times n$ Jacobi matrices.

The fact is that these four sets are in natural bijections with each other. We briefly explain how.

- Given $\mu \in \mathcal{M}_n$, let $\alpha_k = \int x^k d\mu(x)$ be its k th moment. Then $\alpha \in \mathcal{P}_n$.
- Given $\alpha \in \mathcal{P}_n$, use H to define an inner product on the space of polynomials of degree at most $n - 1$ polynomials. Then apply Gram-Schmidt procedure to $1, x, \dots, x^{n-1}$ to get $(\psi_0, \dots, \psi_{n-1}) \in OP_n$.
- Given $(\psi_0, \dots, \psi_{n-1}) \in OP_n$, prove that they satisfy a three term recurrence

$$x\psi_k(x) = b_{k-1}\psi_{k-1}(x) + a_k\psi_k(x) + b_k\psi_{k+1}(x),$$

as in the proof of Lemma 6 (the self-adjointness of multiplication by x is required to see that $c_{k,k+1} = c_{k+1,k}$ which we then define to be a_k). Thus we get a Jacobi matrix $T(a, b) \in \mathcal{J}_n$.

- Given $T \in \mathcal{J}_n$, define μ to be its spectral measure at e_1 . Then $\mu \in \mathcal{M}_n$. This completes the cycle.

The classical moment problem is analogous, except that the probability measure need not have support of finite cardinality, moment sequences need not have finite rank, polynomial sequences need not end, Jacobi matrices need not be finite. A cycle similar to the above exists, except for the last link, from infinite Jacobi matrices to measures. In fact, an infinite Jacobi matrix defines a symmetric (unbounded) operator on $\ell^2(\mathbb{N})$ (say defined on the dense subspace of sequences that vanish eventually). It will always have a self-adjoint extension, but it may have several such extensions. Each self-adjoint extension has a possibly different spectral measure at e_1 . When the extension is unique, the measure is uniquely defined.

Thus the solution to the classical moment problem is this. Given a positive semi-definite sequence α , the question is if it is the moment sequence of a unique measure on \mathbb{R} . Construct the corresponding orthogonal polynomial sequence and then the (infinite) Jacobi matrix. At this point, there is either uniqueness (of the spectral measure at e_1) or not, depending on the uniqueness of the self-adjoint extension of the Jacobi matrix. Existence is assured because there is always a self-adjoint extension (but the existence of a measure is easy to prove, by standard Riesz-like representation theorems or Helly's selection principle. It is uniqueness that is subtle).

Laguerre Beta ensembles: Another random tridiagonal matrix

Now consider a bidiagonal matrix

$$S_n = \begin{bmatrix} A_1 & 0 & 0 & 0 & 0 & 0 \\ B_1 & A_2 & 0 & \ddots & 0 & 0 \\ 0 & B_2 & A_3 & \ddots & \ddots & 0 \\ 0 & \ddots & \ddots & \ddots & 0 & 0 \\ 0 & 0 & \ddots & B_{n-2} & A_{n-1} & 0 \\ 0 & 0 & 0 & 0 & B_{n-1} & A_n \end{bmatrix}. \quad (12)$$

Then $T_n = T_n(a, b) = S_n S_n^t$ is a tridiagonal matrix with

$$a_k = A_k^2 + B_{k-1}^2, \quad 1 \leq k \leq n, \quad \text{and} \quad b_k = A_k B_k, \quad 1 \leq k \leq n-1, \quad (13)$$

with the convention that $B_0 = 0$.

The goal is to find a nice distribution on (A, B) so that the eigenvalues of T_n have a nice density (these are also the same as the squared singular values of S_n). Let (λ, p) be the variables associated to $T_n(a, b)$ as before. We already know how to change variables from (a, b) to (λ, p) by Lemma 6. To change from (A, B) to (a, b) , we have from (13),

$$\begin{aligned} da_1 &= 2A_1 dA_1, \quad db_1 = A_1 dB_1, \quad da_2 = 2A_2 dA_2 + [\dots], \dots \\ \dots, db_{n-1} &= A_{n-1} dB_{n-1} + [\dots], \quad da_n = 2A_n dA_n + [\dots]. \end{aligned}$$

Therefore, we get the Jacobian determinant

$$\det \left[\frac{\partial(a_1, b_1, \dots, b_{n-1}, a_n)}{\partial(A_1, B_1, \dots, B_{n-1}, A_n)} \right] = 2^n A_n \prod_{j=1}^{n-1} A_j^2.$$

Thus, if (A, B) has joint density $f(A, B)$, then (λ, p) has joint density

$$\begin{aligned} g(\lambda, p) &= f(A, B) \frac{\prod_{k=1}^{n-1} b_k}{2^{n-1} \prod_{k=1}^n p_k} \frac{1}{2^n A_n \prod_{k=1}^{n-1} A_k^2} \\ &= \frac{1}{2^{2n-1}} f(A, B) \prod_{k=1}^{n-1} B_k \prod_{k=1}^n \frac{1}{A_k} \prod_{k=1}^n \frac{1}{p_k} \end{aligned}$$

by using the relations $b_k = A_k B_k$.

Now let $A_k^2 \sim \chi_{p_k}^2$ and $B_k^2 \sim \chi_{q_k}^2$. Then,

$$f(A, B) = \frac{1}{Z} e^{-\frac{1}{2}[\sum_{k=1}^n A_k^2 + \sum_{k=1}^{n-1} B_k^2]} \prod_{k=1}^n A_k^{p_k-1} \prod_{k=1}^{n-1} B_k^{q_k-1}.$$

The normalizing constant Z is easy to compute explicitly. Now we note the relations,

$$\begin{aligned} \sum_{k=1}^n A_k^2 + \sum_{k=1}^{n-1} B_k^2 &= \text{tr}(S_n S_n^t) = \text{tr}(T_n) = \sum_{k=1}^n \lambda_k, \\ \prod_{k=1}^n A_k^2 &= \det(S_n)^2 = \det(T_n) = \prod_{k=1}^n \lambda_k. \end{aligned}$$

Thus, we get

$$\begin{aligned} g(\lambda, p) &= \frac{1}{Z'} e^{-\frac{1}{2} \sum_{k=1}^n \lambda_k} \prod_{k=1}^n A_k^{p_k-2} \prod_{k=1}^{n-1} B_k^{q_k} \prod_{k=1}^n \frac{1}{p_k} \\ &= \frac{1}{Z'} e^{-\frac{1}{2} \sum_{k=1}^n \lambda_k} \prod_{k=1}^n A_k^{p_k - q_k - 2} \prod_{k=1}^{n-1} (A_k B_k)^{q_k} \prod_{k=1}^n \frac{1}{p_k} \end{aligned}$$

if we adopt the convention that $q_n = 0$. Now we make the choice $q_k = \beta(n-k)$ so that

$$\prod_{k=1}^{n-1} (A_k B_k)^{q_k} = \prod_{k=1}^{n-1} b_k^{2(n-k)} = \prod_{k=1}^n p_k^{\frac{1}{2}\beta} \cdot \prod_{i<j} |\lambda_i - \lambda_j|^\beta$$

where the last equality follows from (2). In fact this identity motivates the choice of q_k s. We are left with $\prod A_k^{p_k - q_k - 2}$ and this can be written in terms of eigenvalues if the exponents are equal. Hence we take $p_k = 2\alpha + \beta(n-k)$ so that

$$\prod_{k=1}^n A_k^{p_k - q_k - 2} = \prod_{k=1}^n A_k^{2(\alpha-1)} = \prod_{k=1}^n \lambda_k^{\alpha-1}.$$

Putting everything together, we have

$$g(\lambda, p) = \frac{1}{Z''} e^{-\frac{1}{2} \sum_{k=1}^n \lambda_k} \prod_{k=1}^n \lambda_k^{\alpha-1} \prod_{i<j} |\lambda_i - \lambda_j|^\beta \cdot \prod_{k=1}^n p_k^{\frac{1}{2}\beta-1}$$

Again, we see that λ and p are independent, p has Dirichlet distribution, and λ has joint density

$$\frac{1}{Z_{n,\alpha,\beta}} \prod_{i<j} |\lambda_i - \lambda_j|^\beta \prod_{k=1}^n h(\lambda_k), \quad \text{for } \lambda_k \in \mathbb{R}_+, \quad (14)$$

where $h(x) = x^{\alpha-1} e^{-\frac{1}{2}x}$ is the Gamma density. This is very similar to the earlier case, except that eigenvalues are now restricted to \mathbb{R}_+ and the Gaussian factor $e^{-x^2/2}$ is replaced by the Gamma factor $h(x)$. For this reason, this joint density is called the *Laguerre beta ensemble* in contrast to (6) which is called *Hermite beta ensemble*. Of course, now we have an extra parameter α , but that features in the product factor h and not in the interaction $|\lambda_i - \lambda_j|^\beta$ (hence, β is the “more important parameter”).

Exercises

Exercise 9. Let A be an $n \times n$ real symmetric (or Hermitian or normal) matrix. Let e_1, \dots, e_j be an orthonormal basis. If ν_j is the spectral measure of A at e_j , show that $L_A = \frac{1}{n}(\nu_1 + \dots + \nu_n)$.

Exercise 10. For the discrete Laplacian matrix T_n with $a_k = 0$ and $b_k = 1$, and $e = e_1$ (the first coordinate vector), find the spectral measure explicitly and draw its histogram. What shape do you see? What about the spectral measure at e_m where $m = \lfloor n/2 \rfloor$?

Exercise 11. Let ξ_1, \dots, ξ_n be independent random variables with $\xi_j \sim \text{Gamma}(\alpha_j, 1)$. Let $S = \xi_1 + \dots + \xi_n$ and let $p_i = \xi_i/S$.

1. Make a change of variables to show that the density of (S, p_1, \dots, p_{n-1}) on $(0, \infty) \times \Delta_n$ is

$$\frac{1}{\prod_{j=1}^n \Gamma(\alpha_j)} e^{-S} S^{\alpha_1 + \dots + \alpha_n - 1} \prod_{j=1}^n p_j^{\alpha_j - 1}.$$

2. Deduce that the normalizing constant in the Dirichlet density $\frac{1}{D_n(\alpha_1, \dots, \alpha_n)} p_1^{\alpha_1 - 1} \dots p_n^{\alpha_n - 1}$ is

$$D_n(\alpha_1, \dots, \alpha_n) = \frac{\Gamma(\alpha_1 + \dots + \alpha_n)}{\Gamma(\alpha_1) \dots \Gamma(\alpha_n)}.$$

Exercise 12. With $a_k \sim N(0, 2)$, $b_k^2 \sim \chi_{\beta(n-k)}^2$, all independent, follow the constants to deduce the following normalization constants.

1. For (a, b) the density is given by (4) with $Z_{\beta, n} = 2^{1 + \frac{1}{2}\beta n(n-1)} \pi^{\frac{1}{2}n} \prod_{k=1}^{n-1} \Gamma(\frac{1}{2}\beta k)$.

2. For (λ, p) , the density is given by (5) with $Z'_{\beta, n} = 2^{n + \frac{1}{2}\beta n(n-1)} \pi^{\frac{1}{2}n} \prod_{k=1}^{n-1} \Gamma(\frac{1}{2}\beta k)$.

3. For λ , the density is given by (6) with

$$Z''_{\beta, n} = 2^{n + \frac{1}{2}\beta n(n-1)} \pi^{\frac{1}{2}n} \frac{\prod_{k=1}^n \Gamma(\frac{1}{2}\beta k)}{\Gamma(\frac{1}{2}\beta)^n}.$$

Exercise 13. Similarly to the previous exercise, work out the normalization constant in (14)

$$Z_{n,\alpha,\beta} = 2^{\beta(n-1)n+n\alpha} \frac{\Gamma(\alpha)\Gamma(\frac{1}{2}\beta n)}{\Gamma(\frac{1}{2}\beta)^n} \prod_{k=1}^{n-1} \Gamma\left(\frac{1}{2}\beta k\right) \Gamma\left(\alpha + \frac{1}{2}\beta k\right).$$

possibly wrong, check!

Notes

Trotter was the first to consider random tridiagonal matrices. He derived the limiting spectral distribution of $T_n(f)$, and deduces the semi-circle law for the GUE matrix (to be introduced later) by reducing it to a tridiagonal matrix (we shall see this in Chapter ??). The beta development is due to Dumitriu and Edelman. This led to many developments, including an approach to the study of spacing in the bulk and edge, even for the classical eigenvalue ensembles. The proof of Lemma 6 is taken from Forrester's book. I was not able to find a shorter or more conceptual argument for the identity (2).

Chapter 3

The beta log-gas

The beta log-gas is the joint density on \mathbb{R}^n given by

$$p_{n,\beta}^V(x) = \frac{1}{Z_{\beta,n}^V} \prod_{i<j} |x_i - x_j|^\beta \prod_{k=1}^n e^{-\beta n V(x_k)}.$$

This can also be written as

$$p_{n,\beta}^V(x) = \frac{1}{Z_{\beta,n}^V} \exp\{-\beta H_n^V(x)\}, \quad H_n^V(x) = n \sum_{k=1}^n V(x_k) - \sum_{i<j} \log |x_i - x_j|. \quad (1)$$

This is in accordance with the general prescription in statistical mechanics wherein any system (here \mathbb{R}^n or the space of configurations of n particles) is defined by an energy function H (here H_n^V) and the probability or probability density of a configuration is proportional to $\exp\{-\beta H(x)\}$, where β is a tunable parameter. At $\beta = 0$ all configurations are equally likely, while as $\beta \rightarrow +\infty$, the probability concentrates on configurations with the lowest possible energy. Its reciprocal $1/\beta$ is what is called temperature in physics.

A simpler example is the probability density $\exp\{-\sum_{k=1}^n V(x_k)\}$, which corresponds to independent random variables with density $e^{-V(x)}$. Physically it describes n non-interacting charges in an electric potential well given by V (i.e., a particle at x has energy $V(x)$) so that the total energy is just $V(x_1) + \dots + V(x_n)$.

In the same way, the log-gas has the physical interpretation of n unit charges in an electric potential V and with interaction energy $\log(1/|x - y|)$. Then the total energy is given by H_n^V . Actual charges in space have the same form of the energy, except that the interaction energy is $1/|x - y|$, which blows up to $+\infty$ when x and y get close. Hence, such configurations are highly unlikely, indicating that under this probability distribution the

points tend to stay away from each other. This *repulsion* is one of the key features of the log-gas.

Mathematically and physically, if we study charges in d -dimensional space, the right choice of the interaction energy is

$$G_d(x, y) = \begin{cases} \frac{1}{|x-y|^{d-2}} & \text{if } d \geq 3, \\ \log \frac{1}{|x-y|} & \text{if } d = 2, \\ |x-y| & \text{if } d = 1. \end{cases}$$

The key point which dictates this choice is that $\Delta_x G_d(x, y) = c_d \delta_y(x)$, where $\Delta_x = \sum_{i=1}^d \frac{\partial^2}{\partial x_i^2}$ is the Laplacian and c_d is a constant (easy to find but irrelevant now). The equation should be interpreted in weak sense¹.

Any probability density can be written as $p(x) = \exp\{\log p(x)\}$. What makes these situations “physical” is that $\log p$ is made of individual and pairwise contributions.

The three questions discussed in the first chapter (limit of the histogram, spacings in the bulk, spacings at the edge) are of great interest for the beta log-gas. A good fraction of this course will be devoted to these questions. The answers have been found recently by several methods, none too easy. The quadratic potential case $V(x) = x^2$ turns out to be relatively easier. One reason is the representation of the log-gas as the eigenvalues of the Jacobi matrix as in Theorem 7.

The quadratic beta log-gas

A particularly important case is when $V(x) = x^2/4$, which we refer to as the quadratic beta log-gas. The density is

$$\begin{aligned} p_{n,\beta}(x) &= \frac{1}{Z_{\beta,n}^V} \prod_{i < j} |x_i - x_j|^\beta \prod_{k=1}^n e^{-\beta n x_k^2/4} \\ &= \frac{1}{Z_{\beta,n}^V} \exp \left\{ -\beta \left[\frac{n}{4} \sum_{k=1}^n x_k^2 - \sum_{j < k} \log |x_j - x_k| \right] \right\}. \end{aligned}$$

By Theorem 7, this is exactly the same as the density of eigenvalues of $T_n/\sqrt{\beta n}$, where $T_n = T_n(a, b)$ is the Jacobi matrix with independent entries, $a_k \sim N(0, 2)$ and $b_k^2 \sim \chi_{\beta(n-k)}^2$.

¹Meaning: For any smooth, compactly supported function $\varphi : \mathbb{R}^d \mapsto \mathbb{R}$, we have $\int_{\mathbb{R}^d} G_d(x, y) \Delta \varphi(x) dx = c_d \varphi(y)$. If G_d was smooth, the left hand side could be written after integrating by parts twice as $\int_{\mathbb{R}^d} \varphi(x) \Delta_x G_d(x, y) dx$. But the right side is $c_d \int \varphi(x) \delta_y(dx)$, hence the identity “ $\Delta_x G_d(x, y) = c_d \delta_y(x)$ ”.

The density of the log-gas itself has complicated dependent variables, which make it hard to analyse. But since we understand independent random variables better, the tridiagonal matrix helps. In particular, here are three features of the quadratic beta log-gas.

1. An explicit formula for the normalization constant.
2. The empirical distribution of the points converges to the semicircle density on $[-2, 2]$.
3. With high probability, the largest point is close to 2.

We shall prove the first and the third statement in this chapter, and the second statement in the next chapter.

Mehta integral

We did not explicitly give the normalization constant for the quadratic beta log-gas. It can be found explicitly, as given in the following identity (conjecture by Mehta, proved by Selberg 20 years earlier!).

$$\int_{\mathbb{R}^n} \prod_{i < j} |x_i - x_j|^\beta \prod_{k=1}^n e^{-\frac{1}{2}x_k^2} dx = (2\pi)^{n/2} \prod_{k=1}^n \frac{\Gamma(1 + \frac{\beta k}{2})}{\Gamma(1 + \frac{\beta}{2})}. \quad (2)$$

Working out the integral on the left is not a trivial task. Mehta and Dyson were unable to do it in the 1960s, although they conjectured an exact answer. It later turned out to be a consequence of a more general integral identity proved by Selberg in the 1940s:

Theorem 14 (Selberg). *If $\alpha, \beta, \gamma \in \mathbb{C}$ with $\operatorname{Re}(\alpha) > 0$, $\operatorname{Re}(\beta) > 0$, $\operatorname{Re}(\gamma) > -\min\{\frac{1}{n}, \frac{\alpha}{n-1}, \frac{\beta}{n-1}\}$.*

$$\int_{[0,1]^n} \prod_{j < k} |x_j - x_k|^{2\gamma} \prod_{k=1}^n x_k^{\alpha-1} (1-x_k)^{\beta-1} dx_1 \dots dx_n = \prod_{k=0}^{n-1} \frac{\Gamma(\alpha + j\gamma)\Gamma(\beta + j\gamma)\Gamma(1 + (j+1)\gamma)}{\Gamma(1 + \gamma)\Gamma(\alpha + \beta + (n+j-1)\gamma)}.$$

Selberg's integral is a famous identity now, with developments in many directions. Mehta's integral can be obtained as a limiting case of Selberg's integral for special values of parameters, see Exercise 21.

The Jacobi matrix approach has already given us a direct proof of (2). Indeed, in (6), put $\lambda_k = \sqrt{2}x_k$ and use Exercise 12 to get

$$\int_{\mathbb{R}_+^n} \prod_{i < j} |x_i - x_j|^\beta \prod_{k=1}^n e^{-\frac{1}{2}x_k^2} dx = \frac{1}{(2\pi)^{n/2}} \frac{\prod_{k=1}^n \Gamma(\frac{1}{2}k\beta)}{\Gamma(\frac{1}{2}\beta)^n}.$$

If the integral is extended to \mathbb{R}^n , the result will be $n!$ times the expression on the right, which is equal to the right hand side of (2) by writing $\Gamma(1 + \frac{1}{2}\beta k) = \frac{1}{2}\beta k \Gamma(\frac{1}{2}\beta k)$.

The range of the log-gas

Let $x = (x_1, \dots, x_n)$ be the quadratic beta log-gas. We now show that all the points are essentially inside the interval $[-2, 2]$. The precise statement is as follows.

Theorem 15. *There exist δ_n and ε_n converging to 0 such that ,*

$$-2 - \delta_n \leq \min_i x_i \leq \max_i x_i \leq 2 + \delta_n \text{ with probability more than } 1 - \varepsilon_n.$$

I do not know any simple deduction of it from the log-gas density itself. But the tridiagonal matrix makes it almost trivial! We start with a fact in linear algebra².

Fact 16. Let $A_n = (a_{i,j})_{i,j \leq n}$ be any real (or complex) matrix. Let $s_k = \sum_{j=1}^n |a_{j,k}|$ for $1 \leq k \leq n$ and let $s = \max\{s_1, \dots, s_n\}$. Then all eigenvalues of A_n are bounded in absolute value by s .

To see this, suppose $Av = \lambda v$, $v \neq 0$. Pick a k such that $|v_k| = \max\{|v_1|, \dots, |v_n|\}$. Then

$$|\lambda v_k| = \left| \sum_{j=1}^n a_{k,j} v_j \right| \leq \sum_{j=1}^n |a_{j,k}| |v_j| \leq |v_k| s_k$$

showing that $|\lambda| \leq s_k$ for some k . Thus all eigenvalues are of absolute value at most s . ■

Let us return to the log-gas. We know that the quadratic beta log-gas is the joint density of eigenvalues of $\frac{1}{\sqrt{\beta n}} T_n(a, b)$ where T_n is as in Theorem 7. If we apply the fact proved above to this random matrix, we see that the points of the quadratic beta log-gas all lie in the interval $[-B_n, B_n]$ where

$$B_n = \frac{1}{\sqrt{\beta n}} \left\{ \max_{k \leq n} |a_k| + 2 \max_{k \leq n-1} b_k \right\}.$$

Theorem 15 is proved if we show that $\limsup B_n \leq 2$ a.s. We shall be more quantitative and get some explicit (but not optimal!) δ_n and ε_n . Clearly, what we need are tail bounds on Gaussian and chi-squared random variables.

²It is a good place to recall a more powerful and beautiful theorem in linear algebra. The *Gershgorin circles theorem* states that all the eigenvalues of A_n are contained in the union of the closed disks $\mathbb{D}(a_{k,k}, r_k)$ where $r_k = s_k - |a_{k,k}|$. This of course implies the fact we stated.

Fact 17. Let $Z \sim N(0, 1)$ and $W \sim \chi_m^2$.

1. $\mathbf{P}\{Z > t\} \leq e^{-t^2/2}$ for any $t > 1$.
2. $\mathbf{P}\{W > mt\} \leq \exp\{-mc(t)\}$ for $t > 1$ where $c(t) = \frac{1}{2}(t - 1 - \log t)$.

Proof of Fact 17. The key idea is (as often in probability) Chebyshev's inequality.

1. $\mathbf{P}(Z > t) = \frac{1}{\sqrt{2\pi}} \int_t^\infty e^{-x^2/2} dx \leq \frac{1}{\sqrt{2\pi}} \int_t^\infty \frac{x}{t} e^{-x^2/2} dx = \frac{1}{\sqrt{2\pi}} \frac{1}{t} e^{-t^2/2}$. For $t > 1$ just drop the denominator to get the desired inequality.
2. We may write $W = Z_1^2 + \dots + Z_m^2$ where Z_i are i.i.d. standard Gaussian random variables. For $0 < \theta < \frac{1}{2}$ (so that $\mathbf{E}[e^{\theta Z_1^2}]$ is finite), we write

$$\mathbf{P}\{W > mt\} \leq e^{-\theta mt} \mathbf{E} \left[e^{\theta Z_1^2} \right]^m = e^{-\theta mt} \left(\frac{1}{\sqrt{1-2\theta}} \right)^m = \exp \left\{ -\frac{1}{2} m [2\theta t + \log(1-2\theta)] \right\}.$$

Using the optimal value $\theta = \frac{1}{2} - \frac{1}{2t}$, the exponent on the right hand side becomes $\frac{1}{2} m [t - 1 - \log t]$. ■

Using these tail bounds, we now find the asymptotics of B_n . For $h > 0$ observe that

$$\mathbf{P} \left\{ \max_{k \leq n} |a_k| \geq \sqrt{2(1+h) \log n} \right\} \leq n \mathbf{P}\{|a_1| \geq \sqrt{2(1+h) \log n}\} \leq n e^{-(1+h) \log n} = \frac{1}{n^h}.$$

Next, recall that $b_k^2 \sim \chi_{\beta(n-k)}^2$ and that χ_m^2 is stochastically smaller than χ_ℓ^2 if $m < \ell$ (by writing as sums of squares of Gaussians for example). Hence,

$$\mathbf{P}\{b_k^2 \geq \beta n t\} \leq \mathbf{P}\{\chi_{\beta n}^2 \geq \beta n t\} \leq \exp\{-c(t)\beta n\}.$$

Therefore, by the union bound

$$\mathbf{P} \left\{ \max_{k \leq n-1} b_k \geq \sqrt{\beta n} \sqrt{t} \right\} \leq n e^{-c(t)\beta n}.$$

Put $\sqrt{t} = 1 + h$ and take h small enough so that $c(t) \geq \frac{1}{4} h^2$. Then, outside an event of probability $\frac{1}{n^h} + n e^{-h^2 n/4}$, we have

$$B_n \leq \frac{\sqrt{2 \log n (1+h)}}{\sqrt{\beta n}} + 2\sqrt{1+h} \leq 2 + O(h) + O\left(\frac{\sqrt{\log n}}{\sqrt{n}}\right).$$

If we take $h = h_n \rightarrow 0$ so that $n^{h_n} \rightarrow \infty$ (eg., $h = n^{-1/4}$), we get the conclusion of the theorem.

Exercises

Exercise 18. Carry out the computations suggested in the text and prove the identity (2).

Exercise 19. If Z_1, \dots, Z_n are i.i.d $N(0, 1)$ and $Z_n^* = \max\{Z_1, \dots, Z_n\}$, show that $\frac{Z_n^*}{\sqrt{2 \log n}} \xrightarrow{P} 1$ as $n \rightarrow \infty$.

Exercise 20. If $W_n \sim \chi_n$, show that $W_n - \sqrt{n}$ converges (without normalization) to a mean zero normal distribution and find the limiting variance.

Exercise 21. Make the substitution $x_k = \frac{1}{2} - \frac{y_k}{2L}$ in the Selberg integral formula, choose α, β appropriately and deduce the Mehta integral formula.

Exercise 22. Similarly, make an appropriate substitution in the Selberg integral and choose the parameters to deduce the value of

$$\int_{\mathbb{R}_+^n} \prod_{j < k} |x_j - x_k|^\beta \prod_{k=1}^n x_k^{\alpha-1} e^{-\frac{1}{2}x_k} dx_1 \dots dx_n.$$

as given in Exercise 13.

Chapter 4

The method of moments applied to Jacobi matrices, deterministic and random

A class of deterministic Jacobi matrices

Fix a continuous function $f : [0, 1] \mapsto \mathbb{R}$ and recall the Jacobi matrix with $a_k = 0$ and $b_k = f(k/n)$.

$$T_n(f) = \begin{bmatrix} 0 & f(\frac{1}{n}) & 0 & 0 & 0 & 0 \\ f(\frac{1}{n}) & 0 & f(\frac{2}{n}) & \ddots & 0 & 0 \\ 0 & f(\frac{2}{n}) & \ddots & \ddots & \ddots & 0 \\ 0 & \ddots & \ddots & \ddots & f(\frac{n-2}{n}) & 0 \\ 0 & 0 & \ddots & f(\frac{n-2}{n}) & 0 & f(\frac{n-1}{n}) \\ 0 & 0 & 0 & 0 & f(\frac{n-1}{n}) & 0 \end{bmatrix}. \quad (1)$$

We shall find the limiting spectral distribution of T_n in terms of f . We do this to illustrate the method of moments and the method of Stieltjes' transform.

The method of moments

The key point is that the eigenvalues (unknown) are related to the entries (known) by the remarkable identity

$$\sum_{k=1}^n a_{k,k} = \text{tr}(A) = \sum_{k=1}^n \lambda_k.$$

valid for any matrix $A_n = (a_{i,j})_{i,j \leq n}$ with eigenvalues $\lambda_1, \dots, \lambda_n$. Applying the identity to A^p whose eigenvalues are λ_k^p , we get

$$\int x^p dL_{A_n}(x) = \frac{1}{n} \sum_{k=1}^n \lambda_k^n = \frac{1}{n} \text{tr}(A_n^p) = \frac{1}{n} \sum_{i_1, \dots, i_p \leq n} a_{i_1, i_2} a_{i_2, i_3} \dots a_{i_p, i_1}.$$

This identity allows us to apply the method of moments to the empirical spectral distributions L_{A_n} . For the particular case of the Jacobi matrix $T_n = T_n(f)$, we get

$$\int x^p dL_{T_n}(x) = \frac{1}{n} \sum_{1 \leq i_1, \dots, i_p \leq n} T_n(i_1, i_2) \dots T_n(i_p, i_1).$$

Since $T_n(i, j)$ is zero unless $|i - j| = 1$, the sum is over lattice paths (or “simple random walk paths” if that makes it more clear) in \mathbb{Z} , constrained to be inside $\{1, \dots, n\}$. From this or otherwise, one may see that $\text{tr}(T_n^p) = 0$ whenever p is odd. If $p = 2q$, then we split the sum based on the starting point as follows.

$$\frac{1}{n} \sum_{\ell=1}^n \sum_{i_2, \dots, i_{2q}} T_n(\ell, i_2) T_n(i_2, i_3) \dots T_n(i_{2q-1}, i_{2q}) T_n(i_{2q}, \ell).$$

For a given starting point $q < \ell < n - q$, the full set of $\binom{2q}{q}$ paths appear in the sum (the constraint to stay inside $\{1, \dots, n\}$ is irrelevant for them). And for a given ℓ , all the indices i_2, \dots, i_{2q} are within $\ell - q$ and $\ell + q$ (ignoring the zero terms, of course). Therefore, by the continuity of f , we may write¹

$$|T_n(i_1, i_2) \dots T_n(i_{2q}, i_1) - f(\ell/n)^{2q}| \leq \delta_n$$

where $\delta_n \rightarrow 0$ as $n \rightarrow \infty$. The entire contribution of paths starting within q distance of ℓ/n is bounded from above by $\frac{1}{n} \binom{2q}{q} 2q \|f\|_{\text{sup}}^{2q}$ which also goes to zero as $n \rightarrow \infty$. Thus

$$\left| \frac{1}{n} \text{tr}(T_n^{2q}) - \frac{1}{n} \sum_{\ell=1}^n \binom{2q}{q} f(\ell/n)^{2q} \right| \rightarrow 0.$$

The Riemann sum approaches the corresponding integral and hence we get

$$\int x^{2q} dL_{T_n}(x) \rightarrow \binom{2q}{q} \int_0^1 f(x)^{2q} dx.$$

¹If $|x_i - y_i| \leq \varepsilon$ and $|x_i|, |y_i| \leq M$, then $|x_1 \dots x_k - y_1 \dots y_k| \leq k\varepsilon M^{k-1}$. In our case, $M = \|f\|_{\text{sup}}$ and $\varepsilon = \omega_f(q/n)$ where $\omega_f(h) = \sup\{|f(x) - f(y)| : |x - y| \leq h\}$.

In other words, the moments of the empirical spectral distribution of T_n converge to α_p , where $\alpha_p = 0$ if p is odd and $\alpha_{2q} := \binom{2q}{q} \int_0^1 f(x)^{2q} dx$. Since f is bounded, we see that

$$|\alpha_{2q}| \leq 2^{2q} \|f\|_{\text{sup}}^{2q},$$

which shows that $(\alpha_q)_{q \geq 1}$ are the moments of a unique compactly supported probability measure μ_f and that $L_{T_n} \xrightarrow{d} \mu_f$.

Alternately, from the fact that $\|f\|_{\text{sup}} < \infty$, we see that the eigenvalues of T_n must be bounded by $2\|f\|_{\text{sup}}$ (maximum of the row sums of absolute values of the entries is an upper bound for the largest eigenvalue) and hence L_{T_n} and the limiting measure must all be supported in $[-\|f\|_{\text{sup}}, \|f\|_{\text{sup}}]$.

As it happens, it is easy to write down the measure μ_f . We say it in terms of random variables.

Exercise 23. Let $V \sim \text{unif}[0, 1]$ and $X \sim \text{arcsine}[-2, 2]$ (the density is $\frac{1}{\pi\sqrt{4-x^2}}$) be independent random variables and let $Y = Xf(V)$. Then

$$\mathbf{E}[Y^p] = \begin{cases} \binom{2q}{q} \int_0^1 f(x)^{2q} dx & \text{if } p = 2q \text{ is even,} \\ 0 & \text{if } p \text{ is odd.} \end{cases}$$

As a special case, if f is the constant function 1, we get the discrete Laplacian matrix that we considered earlier. The limiting distribution is arcsine measure on $[-2, 2]$, which agrees with our derivation in Chapter 1.

Another special case is of the (scaled) oscillator matrix for which $f(x) = \sqrt{x}$ (the function $\sqrt{1-x}$ will give the same results). The even moments are the Catalan numbers $\frac{1}{q+1} \binom{2q}{q}$. We know that the semi-circle on $[-2, 2]$ is the only such measure. Thus we have justified the theorem suggested in Figure 1.4 and Figure 1.6.

Remark 24. A shortcoming of the method of moments is seen in this class of examples. From the moments (even when it determines the measure), it is hard to see whether the measure is absolutely continuous, whether the density is smooth or bounded, etc. The cases where the density can be guessed are essentially the only cases where one is able to answer these questions. We shall see some open questions of this type later.

In the case at hand we get the absolute continuity of the limiting spectral distribution from the representation $Z = Xf(V)$. If we condition on V , the distribution is arcsine on

$[-2f(V), 2f(V)]$ which has density

$$\frac{\mathbf{1}_{-2f(V) < x < 2f(V)}}{\pi \sqrt{4f(V)^2 - x^2}}.$$

Integrate over V to get the density of the limiting spectral distribution to be

$$\rho(x) = \frac{1}{\pi} \int_0^1 \frac{\mathbf{1}_{-2f(v) < x < 2f(v)}}{\sqrt{4f(v)^2 - x^2}} dv$$

To cross check the answer, take the case $f(v) = \sqrt{v}$ to get

$$\begin{aligned} \rho(x) &= \frac{1}{\pi} \int_0^1 \frac{\mathbf{1}_{-2\sqrt{v} < x < 2\sqrt{v}}}{\sqrt{4v - x^2}} dv = \frac{1}{2\pi} \int_{x^2/4}^1 \frac{1}{\sqrt{v - \frac{1}{4}x^2}} dv \\ &= \frac{1}{2\pi} \frac{1}{2} \sqrt{v - \frac{1}{4}x^2} \Big|_{v=\frac{1}{4}x^2}^{v=1} = \frac{1}{2\pi} \sqrt{4 - x^2}. \end{aligned}$$

which agrees with our earlier derivation.

Remark 25. The method of moments is quite effective and robust. In Exercise 27, you are asked to follow the same idea as above to find the limit of the spectral distribution of $T_n(f)$ at e_1 . It is worth noting that the limit is always a semi-circle distribution provided $f(0)$ is strictly positive.

Another example where the method of moments is effective is the one-dimensional Anderson model, see Exercise 28. However, that exercise shows one limitation of the method of moments. It is difficult to determine absolute continuity of the limiting spectral distribution or the extent to which the density is smooth from a knowledge of the moments.

A more general Jacobi matrix

Let $f : [0, 1] \mapsto \mathbb{R}$ and $g : [0, 1] \mapsto \mathbb{R}_+$ be continuous functions. Let $T_n(f, g) = T_n(a, b)$ where $a_k = f(k/n)$ and $b_k = g(k/n)$. Then we want to find the limiting spectral distribution of T_n . Instead of going over the moments as before, let us put forth another way to think of the situation. When we are focusing on indices k such that k/n is close to $x \in (0, 1)$, we have

before us a matrix of the form

$$\begin{bmatrix} \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \\ \ddots & 0 & g(x) & f(x) & g(x) & 0 & \ddots & \ddots & \ddots \\ \ddots & \ddots & 0 & g(x) & f(x) & g(x) & 0 & \ddots & \ddots \\ \ddots & \ddots & \ddots & 0 & g(x) & f(x) & g(x) & 0 & \ddots \\ \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \end{bmatrix}$$

which is really $f(x)I + g(x)L$ where L is the discrete Laplacian matrix. Hence its eigenvalue distribution is like the arcsine law on $[f(x) - 2g(x), f(x) + 2g(x)]$ (since the eigenvalues of L are like arcsine on $[-2, 2]$).

Overall, we have a superposition of these arcsine densities as x varies over $[0, 1]$. A linear change of variables gives us the arcsine density on an interval $[a - 2b, a + 2b]$ to be

$$\frac{\mathbf{1}_{|y-a|\leq 2b}}{\pi\sqrt{4b^2 - (y-a)^2}}.$$

Superposing such densities, we conclude that the limiting spectral distribution of $T_n(f, g)$ must have density

$$\rho(t) = \frac{1}{\pi} \int_0^1 \frac{\mathbf{1}_{|t-f(x)|\leq 2g(x)}}{\sqrt{4g(x)^2 - (t-f(x))^2}} dx$$

When $f = 0$, we recover the result obtained earlier.

If this derivation does not satisfy your demands of rigour, then complete the following steps to get a full proof:

1. Show that for $k \approx nx$ with $0 < x < 1$ and fixed $p \geq 1$,

$$(T^p)_{k,k} \approx \sum_{\ell=0}^{p/2} \binom{p}{2\ell} \binom{2\ell}{\ell} f(x)^{p-2\ell} g(x)^{2\ell}.$$

2. Show that

$$\int x^p dL_{T_n}(x) = \frac{1}{n} \text{tr}(T_n^p) \rightarrow \sum_{\ell=0}^{p/2} \binom{p}{2\ell} \binom{2\ell}{\ell} \int_0^1 f(x)^{p-2\ell} g(x)^{2\ell} dx.$$

3. Use the Fourier identity $\int_0^1 e^{2\pi i m \theta} d\theta = \delta_{m,0}$ (for integer m), to prove the identity

$$\int_0^1 \left(a + be^{2\pi i \theta} + be^{-2\pi i \theta} \right)^p d\theta = \sum_{\ell=0}^{p/2} \binom{p}{2\ell} \binom{2\ell}{\ell} a^{p-2\ell} b^{2\ell}.$$

4. Put the previous three steps together to show that

$$\lim_{n \rightarrow \infty} \int x^p dL_{T_n}(x) = \mathbf{E}[Z^p]$$

where $Z = f(U) + 2g(U) \cos(2\pi V)$ where U, V are independent uniform $[0, 1]$ random variables.

5. As Z is a bounded random variable, convergence of moments of L_{T_n} to the moments of Z implies the convergence in distribution of L_{T_n} to the distribution of Z .

Observe that $\cos(2\pi V)$ has arcsine distribution on $[-1, 1]$. Therefore, by conditioning on U , we see that the distribution of Z is a mixture of arcsine distributions (not necessarily symmetric about zero). This is the density formula we wrote earlier as $\rho(t)$.

The limiting distribution of beta log-gases

If $a_k \sim N(0, 1)$ and $b_k^2 \sim \chi_{\beta(n-k)}^2$ are independent, then we know that the eigenvalues of $T_n(a, b)$ form the beta log-gas. Finding the limiting distribution of the latter is thus equivalent to finding the limiting spectral distribution of T_n . We shall do this by the method of moments, for which there is hope since the entries of T_n are independent random variables with well-understood distributions while the log-gas is a joint density of highly dependent random variables.

First and foremost, we must figure out the scaling. The second moment of the empirical spectral distribution is given by $\frac{1}{n} \sum_{k=1}^n \lambda_k^2 = \frac{1}{n} \text{tr}(T_n^2)$. If we scale the eigenvalues (or equivalently the matrix T_n) down by a factor of s_n , then the second moment becomes $\frac{1}{ns_n^2} \sum_{k=1}^n \lambda_k^2 = \frac{1}{ns_n^2} \text{tr}(T_n^2)$. We would like to choose s_n so that this quantity stays bounded away from 0 or ∞ (why?). Taking expectations, we see that

$$\mathbf{E}[\text{tr}(T_n^2)] = \mathbf{E} \left[\sum_{k=1}^n a_k^2 + 2 \sum_{k=1}^{n-1} b_k^2 \right] = n + 2\beta \sum_{k=1}^{n-1} (n-k) \sim \beta n^2.$$

Thus we take $s_n = \sqrt{n}$ so that $\mathbf{E}[\frac{1}{ns_n^2} \text{tr}(T_n^2)] \sim \beta$. With this scaling, we are working with the matrix $\frac{1}{\sqrt{n}} T_n$ and its empirical distribution $L_n = L_{T_n/\sqrt{n}}$.

A back of the envelope calculation: The diagonal entries of T_n/\sqrt{n} are a_k/\sqrt{n} which has $N(0, 1/n)$ distribution. The super-diagonal entries are b_k/\sqrt{n} . Exercise 31 tells us the behaviour of the square-root of a chi-squared random variable as the number of degrees of

freedom grows large:

$$\sqrt{Y_n} - \sqrt{n} \xrightarrow{d} N(0, \frac{1}{2}).$$

Thus, $\frac{1}{\sqrt{n}}b_k \approx \sqrt{\beta}\sqrt{1 - \frac{k}{n}} + \frac{c_k}{\sqrt{2n}}$ where $c_k \sim N(0, 1)$ (except when $n - k$ is small).

Now we ignore the terms a_k/\sqrt{n} and c_k/\sqrt{n} , as they look small. Then T_n/\sqrt{n} becomes precisely a matrix of the form (1) with $f(x) = \sqrt{\beta}\sqrt{1-x}$. Now it is clear that the parameter β only affects an overall scaling (it would have been wiser to take $s_n = \sqrt{\beta n}$). Taking $\beta = 1$ for simplicity, we have the scaled oscillator matrix whose limiting spectral distribution is the the semi-circle distribution on $[-2, 2]$.

Fixing the loose parts in the argument: In the heuristic above, we must justify ignoring the terms a_k/\sqrt{n} and c_k/\sqrt{n} as well as the normal approximation for b_k . All this can be done in one step. Let A_n be the matrix in (1) with $f(x) = \sqrt{1-x}$ and let T_n be the matrix as before but with $\beta = 1$. We compare $\hat{T}_n := T_n/\sqrt{n}$ and A_n .

$$\text{tr}(\hat{T}_n - A_n)^2 = \frac{1}{n} \sum_{k=1}^n a_k^2 + \frac{2}{n} \sum_{k=1}^{n-1} (b_k - \sqrt{n-k})^2.$$

We take expectations next. To avoid long calculations, we bound the second term as follows. Let $Y_m^2 \sim \chi_m^2$. Then²,

$$\mathbf{E}[(Y_m - \sqrt{m})^2] = \mathbf{E}\left[\frac{(Y_m^2 - m)^2}{(Y_m + \sqrt{m})^2}\right] \leq \frac{1}{m} \mathbf{E}[(Y_m^2 - m)^2] = \frac{\text{Var}(Y_m^2)}{m} = 2.$$

This is true for every m . Hence we get $\mathbf{E}[\text{tr}(\hat{T}_n - A_n)^2] \leq 5$. By the Hoffman-Wielandt inequality, if λ_i and μ_i are the eigenvalues of \hat{T}_n and A_n (in increasing order), then

$$\mathbf{E}\left[\sum_{k=1}^n (\lambda_k - \mu_k)^2\right] \leq 5.$$

From Lemma 26, it follows that

$$\mathbf{E}\left[\mathcal{D}(L_{A_n}, L_{\hat{T}_n})^3\right] \leq \frac{5}{n}$$

which converges to 0 as $n \rightarrow \infty$. Hence we also get $\mathcal{D}(L_{A_n}, L_{\hat{T}_n}) \xrightarrow{P} 0$. Since $L_{A_n} \rightarrow \mu_{s.c.}$, the semi-circle measure on $[-2, 2]$, it follows that $\mathcal{D}(L_{T_n}, \mu_{sc}) \xrightarrow{P} 0$.

²Note that \sqrt{m} is not equal to $\mathbf{E}[\chi_m]$. The latter is in fact $\sqrt{2}\Gamma((m+1)/2)/\Gamma(m/2)$, which is asymptotically the same as \sqrt{m} . By the trick of writing $(Y_m - \sqrt{m})^2 \leq (Y_m^2 - m)/(Y_m + \sqrt{m})$, we avoided having to use any facts about the Gamma function.

Lemma 26. Suppose $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ and $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$. Let $L_\lambda = \frac{1}{n} \sum_{k=1}^n \delta_{\lambda_k}$ and $L_\mu = \frac{1}{n} \sum_{k=1}^n \delta_{\mu_k}$ be the corresponding empirical measures. Then,

$$\mathcal{D}(L_\lambda, L_\mu) \leq \left(\frac{1}{n} \sum_{k=1}^n (\lambda_k - \mu_k)^2 \right)^{\frac{1}{3}}.$$

Proof. Now, for any $x \in \mathbb{R}$, it is clear that

$$\begin{aligned} \#\{k : \lambda_k \leq x, \mu_k \geq x + \delta\} &\leq \frac{1}{\delta^2} \sum_{k=1}^n (\lambda_k - \mu_k)^2, \\ \#\{k : \mu_k \leq x, \lambda_k \geq x + \delta\} &\leq \frac{1}{\delta^2} \sum_{k=1}^n (\lambda_k - \mu_k)^2. \end{aligned}$$

If $\mathcal{D}(L_\lambda, L_\mu) \geq \delta$, then, by the definition of the Lévy metric, there is an $x \in \mathbb{R}$ such that the left hand sides of one of these inequalities exceeds $n\delta$. Therefore,

$$n\delta^3 \leq \sum_{k=1}^n (\lambda_k - \mu_k)^2,$$

implying the statement of the lemma. ■

Exercises

Exercise 27. Let $f : [0, 1] \mapsto \mathbb{R}$ be as before. Assume that $f(0) > 0$. Let ν_{T_n} be the spectral distribution of $T_n(f)$ at e_1 . Then, show that ν_{T_n} converges to the semi-circle distribution on the interval $[-2f(0), 2f(0)]$.

Exercise 28. (One dimensional Anderson model). Let X_i be i.i.d. random variables from a distribution determined by its moments. Let $T_n = T_n(a, b)$ where $a_i = X_i$ and $b_i = 1$. Show that T_n has a limiting spectral distribution.

Exercise 29. If $Z = Xf(V)$, where $X \sim \text{arcsine}[-2, 2]$ and $V \sim \text{uniform}[0, 1]$ are independent, then show that Z has density

$$\rho_Z(t) = \frac{1}{\pi} \int_0^1 \frac{\mathbf{1}_{|t| \leq 2f(v)}}{\sqrt{4f(v)^2 - t^2}} dv.$$

Exercise 30. Let X, V be as in the previous exercise. If $f, g : [0, 1] \mapsto \mathbb{R}$ are two continuous functions such that $f(V)X \stackrel{d}{=} g(V)X$, then show that there is a measure-preserving transformation $\phi : [0, 1] \mapsto$

$[0, 1]$ such that $f = g \circ \phi$. Conclude that for any f , there is a g that is decreasing on $[0, 1]$ and such that $T_n(g)$ has the same limiting spectral distribution as $T_n(f)$.

[Note: It is obvious that if $f = g \circ \phi$, then $f(V)X$ and $g(V)X$ have the same distribution. What we are saying here shows precisely the equivalence classes of $f \in C[0, 1]$ for which the limiting spectral distributions of $T_n(f)$ are the same.]

Exercise 31. The following probability facts relate to the tail estimates we used in bounding the range of the log-gas.

1. If Z_1, Z_2, \dots are i.i.d. standard Normal random variables and $Z_n^* = \max\{Z_1, \dots, Z_n\}$, then show that $\frac{1}{\sqrt{2 \log n}} Z_n^* \xrightarrow{P} 1$ as $n \rightarrow \infty$.
2. If $Y_n \sim \chi_n$, then $Y_n - \sqrt{n} \xrightarrow{d} N(0, \frac{1}{2})$ as $n \rightarrow \infty$.

Chapter 5

Stieltjes' transform method for Jacobi matrices

We now show how to get limiting spectral distributions for deterministic and random Jacobi matrices using the method of Stieltjes' transform. First we start with the Jacobi matrix $T_n(f)$ where $f : [0, 1] \mapsto \mathbb{R}_+$ is continuous. Before going to the empirical spectral distribution, we find the limit of the spectral measure at e_1 . In other words, we solve Exercise 27 by this method.

Spectral measure of $T_n(f)$ at e_1

Without loss of generality (since $f(0) > 0$) scale and assume that $f(0) = 1$. Let $\nu_n = \nu_{T_n(f)}$ be the spectral measure of T_n at e_1 . Then,

$$G_{\nu_n}(z) := \int \frac{1}{z-x} d\nu_n(x) = (zI - T_n)^{1,1}.$$

Let \tilde{T} be the matrix got from T by deleting the first row and first column. We know that¹

$$(zI - T)^{1,1} = \frac{1}{z - a_1 - b_1^2(zI - \tilde{T})^{1,1}}. \quad (1)$$

Let us write $G_n(z) = (zI - T)^{1,1}$ and $\tilde{G}_n(z) = (zI - \tilde{T})^{1,1}$. From Lemma 108, we get **Arguments shaky here. Need to fix**

¹Let $M = \begin{bmatrix} a & u^t \\ v & B \end{bmatrix}$ where $a \in \mathbb{C}$, $u, v \in \mathbb{C}^{n-1}$ and B is an $(n-1) \times (n-1)$ complex matrix. If M is non-singular, then the $M^{1,1} = \frac{1}{a - u^t B^{-1} v}$. This can be seen in many ways, for example, first writing the entry in the inverse as $\det(B) / \det(A)$ and then expanding the determinant of A with respect to the first row.

The eigenvalues of T and eigenvalues of \tilde{T} interlace and hence, (see Lemma 105) we have $d_{KS}(L_T, L_{\tilde{T}}) \leq \frac{1}{n}$. From this and the properties of Stieltjes' transform, we get $|G_n(z) - \tilde{G}_n(z)| \leq \frac{1}{nv}$ where $z = u + iv$. From (1) we get

$$G_n(z) = \frac{1}{z - (1 + \delta_n)(G_n(z) + \varepsilon_n)}$$

where $\delta_n = f(1/n) - 1$ and $|\varepsilon_n| \leq \frac{1}{nv}$. Thus, any subsequential limit w of $G_n(z)$ satisfies $w(z - w) = 1$. As $G_n(z)$, $n \geq 1$, are contained in the closed disk $\mathbb{D}(0, 1/v)$, a compact set, it follows that $G(z) := \lim_{n \rightarrow \infty} G_n(z)$ exists and satisfies

$$G(z)(z - G(z)) = 1.$$

Solving the quadratic equation for $G(z)$ we get

$$G(z) = \frac{z - \sqrt{z^2 - 4}}{2}.$$

where $\sqrt{\cdot}$ is defined on $\mathbb{C} \setminus [0, \infty)$ by $re^{i\theta} \mapsto \sqrt{r}e^{i\theta/2}$ with $0 < \theta < 2\pi$. The other root of the quadratic equation is discarded because it does not satisfy $G(z) \sim \frac{1}{z}$ as $z \rightarrow \infty$ (why must this condition be satisfied?). We have already seen that this is the Stieltjes' transform of the semi-circle distribution on $[-2, 2]$. Thus, ν_{T_n} converge to this measure.

Limiting spectral distribution of $T_n(f)$

The Stieltjes' transform of the empirical spectral distribution of $T_n(f)$ is given by

$$H_n(z) = \frac{1}{n} \text{tr}(zI - T_n)^{-1} = \frac{1}{n} \sum_{k=1}^n (zI - T)^{k,k}.$$

Fix k , let u be the k th row of T_n after deleting the k th term in it, let S be the matrix got from T_k by deleting the k th row and the k th column. Then, by the same formula as previously,

$$(zI - T)^{k,k} = \frac{1}{z - u^t S^{-1} u}.$$

But $u = (0 \dots 0 b_{k-1} 0 b_k 0 \dots 0)^t$ and

$$S = \begin{bmatrix} zI - T_{[k]} & 0 \\ 0 & zI - \tilde{T}_{[n-k]} \end{bmatrix}$$

where $T_{[k]}$ is the top $(k-1) \times (k-1)$ principal submatrix of T_n and T_{n-k} is the bottom $(n-k) \times (n-k)$ principal submatrix of T_n . Hence

$$u^t S^{-1} u = (zI - T_{[k]})^{k-1, k-1} + (zI - \tilde{T}_{[n-k]})^{1, 1}.$$

Observe that the two terms are the Stieltjes' transforms of the spectral measures of $T_{[k]}$ and $\tilde{T}_{[n-k]}$ at e_{k-1} and e_1 , respectively. Denoting them by G_k and \tilde{G}_k , we have,

$$H_n(z) = \frac{1}{n} \sum_{k=1}^n \frac{1}{z - b_{k-1}^2 G_k(z) - b_{k-1}^2 \tilde{G}_k(z)} \quad (2)$$

From Exercise 27, we see that if k and $n-k$ are both large, then

$$\begin{aligned} G_k(z) &\approx \frac{1}{b_{k-1}} G\left(\frac{z}{b_{k-1}}\right), \\ \tilde{G}_k(z) &\approx \frac{1}{b_{k-1}} G\left(\frac{z}{b_{k-1}}\right). \end{aligned}$$

where G is the Stieltjes' transform on the semi-circle distribution on $[-2, 2]$. Recall that $G(z) = (z - \sqrt{z^2 - 4})/2$ to see that

$$z - b_{k-1}^2 G_k(z) \tilde{G}_k(z) \approx z - 2b_{k-1} G(z/b_{k-1}) = \sqrt{z^2 - 4b^2}.$$

Plugging these approximations back into (2) and then using the integral approximation to the Riemann sum, we conclude that

$$H_n(z) \rightarrow H(z) := \int_0^1 \frac{1}{\sqrt{z^2 - 4f(t)^2}} dt.$$

We had computed that the arcsine measure on $[-2, 2]$ has Stieltjes' transform $1/\sqrt{z^2 - 4}$. Therefore, $H(z)$ is a superposition of arcsine measures on $[-2f(t), 2f(t)]$. This directly tells us that H is the Stieltjes transform of the distribution of $f(V)X$ where $V \sim \text{Uniform}[-1, 1]$ and $X \sim \text{arcsine}[-2, 2]$ (when we fix V , we get arcsine on $[-2f(V), 2f(V)]$).

We have left out a couple of steps in the above proof. One is approximating G_k and \tilde{G}_k by the scaled semi-circle Stieltjes' transforms. The second is approximating the Riemann sum by the integral. Both are standard and we omit the justification.

Relative advantages of the method of moments and the method of Stieltjes' transforms:

Clearly, method of moments requires the underlying measures to have all moments. Not

entirely true, because it is possible to truncate a measure and work with it. It works, but is often tedious. The Stieltjes' transform does not require moments to exist. More interestingly, if you consider f such that the limiting moments do not determine a measure, then the method of moments does not have a conclusion. But the Stieltjes' transform finds the limiting measure (as the superposition of arcsine laws on $[-2f(V), 2f(V)]$) even in such cases! The virtue of the method of moments is that it is quite flexible.

Let us give another example where Stieltjes' transform gives more information about a measure than the moments do.

One dimensional Anderson model and the method of spectral averaging

Let $T_n = T_n(a, b)$ where $a_k = X_k$ are i.i.d. random variables with distribution θ and $b_k = 1$ for all k . We introduced this model in Exercise 28. if you did that exercise, you would have shown that if θ has moments $\alpha_p, p \geq 0$, then the limiting expected spectral distribution of T_n exists (call it μ) and has n th moment

$$\beta_n = \sum_{P \in \Pi_n} \prod_{k \in \mathbb{Z}} \alpha_{n_k(P)}$$

where the sum is over lazy-paths in \mathbb{Z} of length n starting and ending at 0 (i.e., $P = (x_0, x_1, \dots, x_{n-1}, x_n)$ with $x_0 = x_n = 0$ and $x_i - x_{i-1} \in \{0, 1, -1\}$) and where $n_k(P)$ denotes the number of steps in P from k to itself (i.e., $\#\{i : x_i = x_{i+1} = k\}$).

As we have said before, it is hard to say anything about the continuity properties of μ from these moments. We shall use Stieltjes' transforms to prove the following theorem.

The method of spectral averaging: Fix a real symmetric (or Hermitian) matrix A and a vector v . For $\lambda \in \mathbb{R}$ (or \mathbb{C}), define $A_\lambda = A + \lambda v v^*$. Then, for any $z \in \mathbb{H}$, we have

$$\begin{aligned} (zI - A_\lambda)^{-1} - (zI - A)^{-1} &= (zI - A)^{-1} \{(zI - A) - (zI - A_\lambda)\} (zI - A_\lambda)^{-1} \\ &= \lambda (zI - A)^{-1} v v^* (zI - A_\lambda)^{-1}. \end{aligned}$$

Multiply on the left by v^* and on the right by v and write $q(\lambda) := v^*(zI - A_\lambda)^{-1}v$ to get $q(\lambda) - q(0) = \lambda q(0)q(\lambda)$, which gives

$$q(\lambda) = \frac{-1}{\lambda - \frac{1}{q(0)}}.$$

Now suppose A is fixed and λ is a random variable with bounded density g . Then, the Stieltjes' transform of λ is bounded by $\pi\|g\|_{\text{sup}}$ on \mathbb{H} . Since $\frac{1}{q(0)} \in \mathbb{H}$, we get

$$|\mathbf{E}[q(\lambda)]| = |G_\lambda(1/q(0))| \leq \pi\|g\|_{\text{sup}}.$$

But $q(\lambda)$ is the Stieltjes transform of the spectral measure of A_λ at z . Invoking Lemma 97, we conclude that the expected spectral measure of A_λ at v is absolutely continuous and its density is bounded by $\|g\|_{\text{sup}}$. This is what is called *spectral averaging*.

Theorem 32. *In the one dimensional Anderson model, assume that θ has bounded density g . Then μ also has bounded density (with the same bound).*

Proof. Let L_n be the discrete Laplacian matrix. Then we write T_n as $L + \sum_{k=1}^n X_k v_k v_k^t$. Condition on all random variables except X_k and use the spectral averaging as above to conclude that

$$\mathbf{E}[e_k^t(zI - T_n)e_k \mid X_i, i \neq k] \leq \pi\|g\|_{\text{sup}}.$$

Take another expectation over the conditioned variables and sum over k to get

$$\frac{1}{n} \sum_{k=1}^n \mathbf{E}[e_k^t(zI - T_n)e_k] \leq \pi\|g\|_{\text{sup}}.$$

But the left hand side is equal to $\frac{1}{n}\mathbf{E}[\text{tr}(zI - T_n)^{-1}]$, the Stieltjes' transform of the expected empirical distribution of eigenvalues. As the bound is uniform over n , and convergence in distribution implies pointwise convergence of Stieltjes' transforms, it follows that $|G_\mu(z)| \leq \pi\|g\|_{\text{sup}}$. Invoking Lemma 97 we get the absolute continuity of μ and that its density is bounded by $\|g\|_{\text{sup}}$. ■

Remark 33. [TO CHECK!] It is not that the absolute continuity of θ is necessary. Even if it is a Bernoulli distribution, the limiting distribution is absolutely continuous, but this is harder to prove.

Chapter 6

Gaussian random matrices

GOE and GUE

We now introduce the most well-studied of all random matrices (most courses in the subject start with these matrices!). We shall say that a complex valued random variable Z has $\mathbb{CN}(0, 1)$ distribution if its real and imaginary parts have independent $N(0, 1/2)$ distribution. Equivalently, the density in the complex plane is $\frac{1}{\pi}e^{-|z|^2}$. Then σZ is said to have $\mathbb{CN}(0, \sigma^2)$ distribution for $\sigma > 0$.

Definition 34. Let $A = (a_{i,j})_{i,j \leq n}$ and $B = (b_{i,j})_{i,j \leq n}$ where $a_{i,j}$ are i.i.d $\mathbb{CN}(0, 1)$ random variables and $b_{i,j}$ are i.i.d $N(0, 1)$ random variables. The matrix $X := \frac{A+A^*}{\sqrt{2}}$ is called a *GUE matrix* and $Y := \frac{H+H^t}{\sqrt{2}}$ is called a *GOE matrix*.

Equivalently, we could have said that X is a random Hermitian matrix whose entries on and above the diagonal are independent, the diagonal entries have $N(0, 2)$ distribution and the off-diagonal entries have $\mathbb{CN}(0, 1)$ distribution. Similarly Y is a real symmetric random matrix whose entries on and above the diagonal are independent, the diagonal entries have $N(0, 2)$ distribution and the off-diagonal entries have $N(0, 1)$ distribution.

The names stand for *Gaussian unitary ensemble* and *Gaussian orthogonal ensemble*, stemming from the following invariance property under unitary or orthogonal conjugation.

Lemma 35. Let X be a GOE (or GUE) matrix. Let P be a non-random orthogonal (respectively, unitary) $n \times n$ matrix. Then $P^*XP \stackrel{d}{=} X$.

Proof. Let X be GOE. Then the joint density of $X_{i,j}$, $i, j \leq n$ can be written as

$$\prod_{k=1}^n \frac{1}{\sqrt{2\pi}\sqrt{2}} e^{-\frac{1}{4}X_{k,k}^2} \prod_{i < j} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}X_{i,j}^2} = \frac{1}{2^{n/2}(2\pi)^{n(n+1)/2}} e^{-\frac{1}{4}\text{tr}(X^2)}.$$

This density is with respect to Lebesgue measure on $\mathbb{R}^{n(n+1)/2}$. Equivalently, we may say that X has standard Gaussian distribution on the Hilbert space \mathcal{H}_n of $n \times n$ real symmetric matrices endowed with the inner product $\langle A, B \rangle = \text{tr}(AB)$.

Now if P is an $n \times n$ orthogonal matrix and $T_P : \mathcal{H}_n \mapsto \mathcal{H}_n$ is defined by $T_P(A) = P^t A P$, then

$$\langle T_P(A), T_P(B) \rangle = \text{tr}(P^t A P P^t B P) = \text{tr}(AB) = \langle A, B \rangle$$

showing that T_P is an orthogonal transformation on \mathcal{H}_n . Since standard Gaussian measure is invariant under orthogonal transformations, we get the orthogonal invariance of GOE. The proof is almost identical for GUE. ■

Reduction to a Jacobi matrix

Given a symmetric matrix, there is a standard way to reduce it to a Jacobi matrix by a sequence of similarity transformations. Then, the resulting Jacobi matrix will have the same eigenvalues as the original matrix. As Jacobi matrices are easier to deal with, this is apparently quite useful in numerical algorithms for finding eigenvalues of a real symmetric matrix. For us it will be useful in a different way.

Given a real symmetric matrix $A_{n \times n}$, write it in block form as

$$A = \begin{bmatrix} a & v^t \\ v & B \end{bmatrix}.$$

Find an orthogonal matrix $P \in O(n-1)$ such that $Pv = re_1$ with $r = \|v\|$ and where e_1 is the first co-ordinate vector in \mathbb{R}^{n-1} . There are many such orthogonal matrices. An explicit one is reflection on the affine hyperplane that bisects the line joining re_1 and v given by

$$H_u := I - 2uu^t, \text{ where } u = \frac{v - re_1}{\|v - re_1\|}.$$

In general, H_u reflects about the hyperplane u^\perp . It is called a *Householder reflection*. Extend P to an $n \times n$ orthogonal matrix as

$$\hat{P} = \begin{bmatrix} 1 & 0^t \\ 0 & P \end{bmatrix}$$

and set $A_1 = \hat{P} A \hat{P}^t$ which is equal to

$$\begin{bmatrix} a & v^t P^t \\ P v & P B P^t \end{bmatrix} = \begin{bmatrix} a & re_1^t \\ re_1 & C \end{bmatrix}$$

where $C = PBP^t$ is an $(n-1) \times (n-1)$ matrix.

Repeat the whole procedure for C and find an $(n-1) \times (n-1)$ orthogonal matrix Q such that

$$QCQ^t = \begin{bmatrix} b & se_1^t \\ re_1 & D \end{bmatrix}$$

where $b = C_{1,1}$, $s^2 = C_{1,2}^2 + \dots + C_{1,n-1}^2$, and D is conjugate by an orthogonal matrix to the $(n-2) \times (n-2)$ matrix got from C by removing the first row and first column. Extending Q to an $n \times n$ orthogonal matrix \hat{Q} just as we did P to \hat{P} , we get $A_1 = \hat{Q}A_1\hat{Q}$ which is equal to

$$\left[\begin{array}{cc|cccc} a & r & 0 & 0 & \dots & 0 \\ r & b & s & 0 & \dots & 0 \\ \hline 0 & s & & & & \\ 0 & 0 & & & & \\ \vdots & \vdots & & & D & \\ 0 & 0 & & & & \end{array} \right].$$

Now the procedure is clear. Continuing it, one ends with a Jacobi matrix $T_n(a,b)$ (eg., $a_1 = a$ and $b_1 = r$). It is got by conjugating with orthogonal matrices at each step, hence must have the same eigenvalues as A .

The exact same procedure can be carried out for a Hermitian matrix, except that the conjugation is by unitary matrices. In the end we get a Jacobi matrix (the entries are real, even though the original matrix may be complex).

Eigenvalue distribution of GOE and GUE

Let A be an $n \times n$ GOE matrix, and apply the procedure outlined above to reduce it to a Jacobi matrix. First of all, note that a, v, B are independent. At the first step, A is reduced to

$$\begin{bmatrix} a & re_1^t \\ re_1 & C \end{bmatrix}$$

where $r = \|v\|$ and $C = PBP^t$. Observe that P is obtained as a function of v , and v is independent of B and a . Further, B is just a GOE matrix of order $n-1$. Hence, it follows that a, r, C are independent, $a \sim N(0, 2)$, $r^2 \sim \chi_{n-1}^2$, and C is an $(n-1) \times (n-1)$ GOE matrix.

Writing $a_1 = a$ and $b_1 = r$, we have got

$$\begin{bmatrix} a_1 & b_1 e_1^t \\ b_1 e_1 & C \end{bmatrix}.$$

Now apply the same procedure to C . The only difference is that the length of the vector v is reduced by 1, hence, after two steps we end up with

$$\left[\begin{array}{cc|cccc} a_1 & b_1 & 0 & 0 & \dots & 0 \\ b_1 & a_2 & b_2 & 0 & \dots & 0 \\ \hline 0 & b_2 & & & & \\ 0 & 0 & & & & \\ \vdots & \vdots & & & D & \\ 0 & 0 & & & & \end{array} \right].$$

Here a_1, a_2, b_1, b_2, D are independent, a_1, a_2 are i.i.d. $N(0, 2)$ random variables, $b_1^2 \chi_{n-1}^2$ and $b_2^2 \sim \chi_{n-2}^2$ and D is an $(n-2) \times (n-2)$ GOE matrix.

Thus, the end result of the procedure is a random Jacobi matrix $T_n(a, b)$ where all the entries are independent, $a_k \sim N(0, 2)$ for $k \leq n$, and $b_k^2 \sim \chi_{n-k}^2$ for $k \leq n-1$.

But this is precisely the matrix corresponding to $\beta = 1$ in Theorem 7, whose eigenvalue density is the $\beta = 1$ log-gas. Hence, the eigenvalue density of GOE is precisely the same. In a similar manner, show that the GUE matrix eigenvalues form the $\beta = 2$ log-gas. We summarize the results.

Theorem 36. *The GOE and GUE random matrices have eigenvalue densities proportional to*

$$\exp \left\{ -\frac{1}{4} \sum_{k=1}^n \lambda_k^2 \right\} \prod_{i < j} |\lambda_i - \lambda_j|^\beta$$

for $\beta = 1$ and $\beta = 2$, respectively.

There is another one, called the GSE (Gaussian symplectic ensemble) whose eigenvalue density corresponds to $\beta = 4$. However, there are no corresponding matrices for general β . The random Jacobi matrices are better in this respect - there is nothing special about $\beta = 1$ or $\beta = 2$ in Theorem 7.

It is then clear that whatever we have proved for the log-gases in general, apply also to the GOE and GUE eigenvalue densities. In particular, we have

Theorem 37. *Let X_n be the GOE or GUE random matrix. Then, the limiting spectral distribution of $\frac{1}{\sqrt{\beta n}} X_n$ (with $\beta = 1$ for GOE and $\beta = 2$ for GUE) is the semicircle law on $[-2, 2]$.*

A direct proof by change of variable? Some remarks

It is natural to ask if we could have obtained the eigenvalue density of GOE and GUE directly, by a change of variables from the matrix entries to eigenvalues (and suitably chosen auxiliary variables). Historically that was how it was done, and the proof is presented in many books on random matrices (Mehta ?, ?, ? etc.). Hence we give only a brief account of it and refer you to these books.

As expected, the idea is to make a change of variables from the entries to eigenvalues and some auxiliary variables. These auxiliary variables will come from the eigenvectors. For definiteness, we consider the case of GOE, i.e., real symmetric matrices. We recall,

1. $\mathbb{R}_\uparrow^n = \{\lambda \in \mathbb{R}^n : \lambda_1 > \lambda_2 > \dots > \lambda_n\}$. This is an open subset of \mathbb{R}^n .
2. $O(n)$, the group of $n \times n$ orthogonal matrices. This is a subset of $M_n(\mathbb{R}) = \mathbb{R}^{n^2}$, defined by $n(n+1)/2$ equations, namely $P^t P = I$ (as $P^t P$ is symmetric, we only consider $(P^t P)_{i,j} = \delta_{i,j}$ for $i \leq j$). Thus, $O(n)$ is a manifold of dimension $n^2 - \frac{1}{2}n(n+1) = \frac{1}{2}n(n-1)$.
3. \mathcal{H}_n , the space of $n \times n$ real symmetric matrices. This is identified with $\mathbb{R}^{n(n+1)/2}$ as we have already seen.
4. Define the map $T : \mathbb{R}_\uparrow^n \times O(n) \mapsto \mathcal{H}_n$ by $T(\lambda, P) = P\Lambda P^t$, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$. This map is neither one-one nor onto, but nearly both.

Injectivity fails because if $Q = PD$ where D is a diagonal matrix with ± 1 on the diagonal, then $Q\Lambda Q^t = P\Lambda P^t$ (as D and Λ commute, both being diagonal matrices). But this just means that every point in the image of T has exactly 2^n pre-images. This is just as good as injectivity when making change of variables.

Surjectivity fails because the image consists precisely of symmetric matrices having distinct eigenvalues. But the complement inside \mathcal{H}_n has zero Lebesgue measure (in fact a lower dimensional manifold). Hence we may ignore that part of the space.

We now want the Jacobian determinant of T . We write $X = T(\lambda, P)$. Then $X_{i,j} : \mathbb{R}_\uparrow^n \times O(n) \mapsto \mathbb{R}$ are smooth functions. We write their differentials as

$$dX = Pd\Lambda P^t + (dP)\Lambda P^t + P\Lambda(dP^t).$$

Here $dX = (dX_{i,j})$ is the matrix of differentials. Since $P \mapsto P_{i,j}$ are smooth functions on $O(n)$, the elements $dP_{i,j}$ are also differentials (one-forms) on $O(n)$. The matrix equation

is just a short form of writing n^2 different equations (actually $n(n+1)/2$ equations as the matrices are symmetric). From the identity $PP^t = I$, we have $dP^t = -P^t(dP)P^t$. Define $\Omega = P^t dP$. This is a skew symmetric matrix of differentials. That is, if $\omega_{i,j}$ is the (i,j) entry of Ω , then $\omega_{i,j} = -\omega_{j,i}$ (in particular $\omega_{i,i} = 0$ for all i). Then, the earlier equation becomes

$$P^t(dX)P = d\Lambda + \Omega\Lambda - \Lambda\Omega.$$

Now we take the wedge product of all the differentials on the upper-triangular part. We get

$$\bigwedge_{i \leq j} (P^t dXP)_{i,j} = \bigwedge_i d\lambda_i \bigwedge_{i < j} (\lambda_i - \lambda_j) \omega_{i,j}.$$

In wedge products, the order is important. In the above equation, we are unclear about it, as a change in order only make a difference of change in sign, which we can recover in the end. Now, the forms $(P^t dXP)_{i,j} = \sum_{k,\ell} P_{k,i} P_{\ell,j} dX_{k,\ell}$ is a linear combination of $dX_{i,j}$. Whenever we have a set of equations like $du_i = \sum_{j=1}^n B_{i,j} dx_j$ for $1 \leq i \leq n$, it is clear that

$$du_1 \wedge \dots \wedge du_n = \det(B) dx_1 \wedge \dots \wedge dx_n.$$

This is simply a consequence of the alternating property $dx \wedge dy = -dy \wedge dx$ which is also the key property of determinant (changes sign upon exchange of rows). In our situation, we have $n(n+1)/2$ equations,

$$(P^t dXP)_{i,j} = \sum_k P_{k,i} P_{k,j} dX_{k,k} + \sum_{k < \ell} (P_{k,i} P_{\ell,j} + P_{\ell,i} P_{k,j}) dX_{k,\ell}.$$

Show that the determinant of the linear transformation here is ± 1 . Hence

$$\bigwedge_{i \leq j} (P^t dXP)_{i,j} = \pm \bigwedge_{i \leq j} dX_{i,j}.$$

Putting everything together, we have arrived at the Jacobian determinant formula

$$\bigwedge_{i \leq j} dX_{i,j} = \pm \prod_{i < j} |\lambda_i - \lambda_j| \bigwedge_i d\lambda_i \bigwedge_{i < j} \omega_{i,j}. \quad (1)$$

Locally this is fine, but when done globally we must include a factor of 2^n to account for the multiplicity of the map T .

Lastly, writing $\text{tr}(X^2) = \sum_{k=1}^n \lambda_k^2$, we arrive at

$$e^{-\frac{1}{4} \text{tr}(X^2)} dX = e^{-\frac{1}{2} \sum_{k=1}^n \lambda_k^2} \prod_{i < j} |\lambda_i - \lambda_j| d\lambda d\mu(P)$$

where dX is Lebesgue measure on \mathcal{H}_n and μ is a measure on $O(n)$. Integrating out with respect to μ , we get the density of eigenvalues of the GOE matrix. In a similar manner, one may derive the density of eigenvalues of the GUE matrix.

Just as the Jacobi matrix was useful in studying the quadratic log-gas, the GOE and GUE matrices can also be used to study the same (but only for $\beta = 1$ and $\beta = 2$). For example, the method of moments can be applied directly to the GOE or GUE matrix to prove Theorem 37. The combinatorics is a bit more involved, but very interesting. We shall do it later, to show the interesting connection between random matrices and certain enumeration problems in combinatorics.

Generalizations of GOE and GUE in two directions

What is most natural for probabilists is to relax the assumption of Gaussian distribution but keep independence as much as possible. This leads to the first generalization.

Definition 38. A random matrix $A = (a_{j,k})_{j,k \leq n}$ is called a *Wigner matrix* if A is real symmetric (or complex Hermitian), the entries $a_{j,k}$, $1 \leq j \leq k \leq n$, are all independent, the diagonal entries are identically distributed, the off-diagonal entries are identically distributed, both $a_{1,1}$ and $a_{1,2}$ have finite second moment, $a_{1,2}$ has zero mean and unit variance (i.e., $\mathbf{E}[|a_{1,2}|^2] = 1$).

It is natural to ask whether the asymptotic properties of eigenvalues of GOE (or GUE) also remain valid for the correspond real Wigner (or complex Wigner) matrix. So far we have only seen the semi-circle law for GOE, hence that is one question. But even more interesting are the questions of bulk and edge spacing of eigenvalues (which we have not yet seen even in the GOE case).

GOE is the only real Wigner matrix that has orthogonal invariance (see Exercise ??). As that invariance was crucial in getting the exact eigenvalue density, in general we cannot hope to find the exact eigenvalue distribution for any other Wigner matrix. In fact, I am not aware of the exact eigenvalue distribution of any Wigner matrices other than GOE and GUE. This makes their study more difficult and interesting.

The second kind of generalization gives us exact eigenvalue densities but the distribution of the entries will be entirely unclear.

Definition 39. Let $V : \mathbb{R} \mapsto \mathbb{R}$ be a piecewise continuous function such that $\frac{V(x)}{(\log|x|)^{1+\varepsilon}} \rightarrow \infty$ for some $\varepsilon > 0$. Then, we may consider the random $n \times n$ matrix A with density on

\mathcal{H}_n proportional to $e^{-\text{tr}V(X)}$. Such random matrices are referred to as *unitarily invariant ensembles*.

It shares the following properties with the GOE/GUE. For any orthogonal (or unitary) matrix P , we have $P^*XP \stackrel{d}{=} X$. Therefore, the eigenvalues and eigenvectors are independent of each other. The exact density of eigenvalues is proportional

$$e^{-\sum_{k=1}^n V(\lambda_k)} \prod_{j < k} |\lambda_j - \lambda_k|^\beta$$

where β is 1 or 2, depending on whether we are considering real symmetric or complex Hermitian matrices. When studying asymptotics as the size of the matrix goes to infinity, there are two possible ways: either keep a fixed V (in which case, the right scaling must be found, say to get a limiting spectral distribution) or to take $\beta n V$ in place of V when considering $n \times n$ matrices. In the latter case, the eigenvalue density agrees exactly with the log-gas as we defined it in (1) (but only for $\beta = 1, 2$). I have never seen any use of the matrix itself, the entries. People study the eigenvalue density directly, or using the Jacobi matrices (for quadratic V).

In this course, we shall see a bit of both - the Wigner matrix eigenvalues and the unitarily invariant ensembles. But not the deepest results available. In fact, we shall discuss the limiting spectral distribution completely (for Wigner, may be not for general V in the unitarily invariant case). But when it comes to spacings between eigenvalues, we shall do it for GUE and then for the quadratic β log-gas. Corresponding results for Wigner matrices or for general V are some of the big achievements of random matrix theory (the few breakthroughs available in the question of *universality*), but we may not have time to discuss them.

Exercises

Exercise 40. Let $m \leq n$ and let $A_{m,n}$ be a random matrix whose entries are i.i.d. $N(0, 1)$ random variables. Show that there exist orthogonal matrices $P_{m \times m}$ and $Q_{n \times n}$ such that

$$(PAQ)_{i,j} = \begin{cases} a_i & \text{if } j = i, \\ b_i & \text{if } j = i + 1, \\ 0 & \text{otherwise,} \end{cases}$$

where $a_1, \dots, a_m, b_1, \dots, b_m$ are independent, $a_k^2 \sim \chi_{m-k+1}^2$ and $b_k^2 \sim \chi_{n-k}^2$.

Exercise 41. Show that the density of singular values of the matrix of the previous exercise is of the form $\prod_{j < k} |s_j - s_k| e^{-\sum_j s_j}$.

Chapter 7

Wigner matrices: The semi-circle law

In this chapter we shall prove that Wigner matrices, upon scaling down by \sqrt{n} , have semi-circle distribution as the limiting spectral distribution. We do this by a powerful probabilistic idea of replacement that was first introduced by Lindeberg in his proof of the central limit theorem. Chatterjee generalized this idea and applied it to many questions in probability. One of these applications was to derive semi-circle law for Wigner matrices under optimal conditions.

The invariance principle

Invariance here is the general idea is that for many function $f : \mathbb{R}^n \mapsto \mathbb{R}$, then the distribution of the random variable $f(X_1, \dots, X_n)$ is approximately the same for any i.i.d. X_i s having zero mean and unit variance. The important requirement on f is that it should not depend too much on any single variable.

Theorem 42. *Let $X_k, Y_k, k \leq n$ be independent real-valued random variables. Assume that $\mathbf{E}[X_k] = \mathbf{E}[Y_k]$ and $\mathbf{E}[X_k^2] = \mathbf{E}[Y_k^2]$ for all $k \leq n$. Let A be a positive constant. Let $h : \mathbb{R}^n \mapsto \mathbb{R}$ be a function with bounded derivatives up to order 3. Then, with $U = h(X_1, \dots, X_n)$ and $V = h(Y_1, \dots, Y_n)$, we have*

$$\begin{aligned} |\mathbf{E}[h(U)] - \mathbf{E}[g(V)]| &\leq B_2(h) \sum_{k=1}^n \mathbf{E}[X_k^2 \mathbf{1}_{|X_k| > A}] + \mathbf{E}[Y_k^2 \mathbf{1}_{|Y_k| > A}] \\ &\quad + \frac{1}{6} B_3(h) \sum_{k=1}^n \mathbf{E}[|X_k|^3 \mathbf{1}_{|X_k| \leq A}] + \mathbf{E}[|Y_k|^3 \mathbf{1}_{|Y_k| \leq A}] \end{aligned}$$

where $B_p(h) = \max\{\|\partial_k^p h\|_{\text{sup}} : k \leq n\}$.

The bound looks a bit complicated, hence the following corollary may help to parse it.

Corollary 43. *Along with the assumptions in the theorem, make an additional assumption that $\mathbf{E}[|X_k|^3], \mathbf{E}[|Y_k|^3]$ all exist and are bounded by γ . Then,*

$$|\mathbf{E}[h(U)] - \mathbf{E}[g(V)]| \leq \frac{1}{3}\gamma B_3(h)n.$$

As written, the invariance principle shows the closeness of expectations of the same function applied to two sets of random variables. This also allows us to show closeness of distributions, since the distance between distributions can be measured by the differences in expectations over sufficiently rich classes of functions. To be concrete, suppose $f : \mathbb{R}^n \mapsto \mathbb{R}$ and we want to show the closeness in distribution of $W = f(X)$ and $Z = f(Y)$. We take $g : \mathbb{R} \mapsto \mathbb{R}$ and apply the invariance principle to $h = g \circ f$ to get

$$|\mathbf{E}[g(W)] - \mathbf{E}[g(Z)]| \leq B_2(h)[\dots] + B_3(h)[\dots].$$

If in a given situation, we can show that the right hand side is small for any g , then it follows that W and Z are close in distribution. In this regard, it is convenient to note that

$$\begin{aligned} \partial_k^2 h(x) &= g''(f(x))(\partial_k f(x))^2 + g'(f(x))\partial_k^2 f(x), \\ \partial_k^3 h(x) &= g'''(f(x))(\partial_k f(x))^3 + 3g''(f(x))\partial_k^2 f(x)\partial_k f(x) + g'(f(x))\partial_k^3 f(x), \end{aligned}$$

whence

$$B_2(h) \leq C_2(g)\lambda_2(f), \quad B_3(h) \leq C_3(g)\lambda_3(f),$$

where $C_2(g) = \|g'\|_{\text{sup}} + \|g''\|_{\text{sup}}$ and $C_3(g) = \|g'\|_{\text{sup}} + \|g''\|_{\text{sup}} + \|g'''\|_{\text{sup}}$ while

$$\lambda_p(f) = \max\{\|\partial_k^r f\|_{\text{sup}}^{p/r} : 1 \leq r \leq p, k \leq n\}.$$

An illustration: The Lindeberg-Feller central limit theorem

Fix $g \in C_b^3(\mathbb{R})$ and let $h(x_1, \dots, x_n) = g(x_1 + \dots + x_n)$. Then, writing $s_n = x_1 + \dots + x_n$, we have

$$\partial_k h(x) = g'(s_n), \quad \partial_k^2 h(x) = g''(s_n), \quad \partial_k^3 h(x) = g'''(s_n)$$

from which it follows that $B_2(h) = \|g''\|_{\text{sup}}$ and $B_3(h) = \|g'''\|_{\text{sup}}$. Now we apply Theorem 42 with $X_k = X_{n,k}$ having zero mean and variance $\sigma_{n,k}^2$. We assume that the total variance $\sum_{k=1}^n \mathbf{E}[X_{n,k}^2] = \sigma^2$ stays constant. Further, we choose $Y_k = Y_{n,k}$ to be Normal with the same

mean and variance as X_k , so that $S_n^Y \sim N(0, 1)$. Then, setting $A = \varepsilon$, and writing $|x|^3 \mathbf{1}_{|x| \leq \varepsilon} \leq \varepsilon |x|^2$, we get

$$|\mathbf{E}[g(S_n^X)] - \mathbf{E}[g(S_n^Y)]| \leq \|g''\|_{\sup} \sum_{k=1}^n \mathbf{E}[X_{n,k}^2 \mathbf{1}_{|X_{n,k}| > \varepsilon}] + \mathbf{E}[Y_{n,k}^2 \mathbf{1}_{|Y_{n,k}| > \varepsilon}] + 2\|g'''\|_{\sup} \sigma^2 \varepsilon.$$

Now impose the *Lindeberg condition*

$$\lim_{n \rightarrow \infty} \sum_{k=1}^n \mathbf{E}[X_{n,k}^2 \mathbf{1}_{|X_{n,k}| > \varepsilon}] = 0 \quad \text{for any } \varepsilon > 0.$$

From this, it follows that $\max_{k \leq n} \sigma_{n,k}^2 \rightarrow 0$ as $n \rightarrow \infty$ (why?) and from that it follows that the Lindeberg condition holds for $Y_{n,k}$ s in place of $X_{n,k}$ s. Putting all this together, letting $n \rightarrow \infty$ first and then letting $\varepsilon \rightarrow 0$ we arrive at

$$\mathbf{E}[g(S_n^X)] - \mathbf{E}[g(S_n^Y)] \rightarrow 0$$

as $n \rightarrow \infty$. As $S_n^Y \sim N(0, \sigma^2)$ and this is true for all $g \in C_b^3$, we get $S_n^X \xrightarrow{d} N(0, \sigma^2)$. This is the Lindeberg-Feller central limit theorem.

Remark 44. It is easy to find examples to show that the central limit theorem may fail if the Lindeberg condition is not satisfied. The Lindeberg condition expresses the idea that no single $X_{n,k}$ is too large. More generally, in Theorem 42, the quantities $\lambda_2(h)$ and $\lambda_3(h)$ measure the maximum “influence” of any variable on the value of h . This is seen from the definition of the partial derivative as the change in h proportional to the change in x_k .

Proof of the invariance principle

Define the vectors

$$W_k = (X_1, \dots, X_{k-1}, Y_k, \dots, Y_n), \quad W_k^0 = (X_1, \dots, X_{k-1}, 0, Y_{k+1}, \dots, Y_n).$$

Then,

$$U - V = \sum_{k=0}^n (h(W_{k+1}) - h(W_k^0)) - \sum_{k=0}^n (h(W_k) - h(W_k^0)).$$

By Taylor expansion we write the k^{th} summands as

$$h(W_{k+1}) - h(W_k^0) = \begin{cases} \partial_k h(W_k^0)X_k + \partial_k^2 h(W_k^0)\frac{X_k^2}{2} + \partial_k^3 h(W_k^*)\frac{X_k^3}{6} & \text{if } |X_k| \leq A, \\ \partial_k h(W_k^0)X_k + \partial_k^2 h(W_k^{**})\frac{X_k^2}{2} & \text{if } |X_k| > A, \end{cases}$$

$$h(W_k) - h(W_k^0) = \begin{cases} \partial_k h(W_k^0)Y_k + \partial_k^2 h(W_k^0)\frac{Y_k^2}{2} + \partial_k^3 h(W_k^\#)\frac{Y_k^3}{6} & \text{if } |Y_k| \leq A, \\ \partial_k h(W_k^0)Y_k + \partial_k^2 h(W_k^{\#\#})\frac{Y_k^2}{2} & \text{if } |Y_k| > A. \end{cases}$$

where W_k^*, W_k^{**} are in $[0, X_k]$ and $W_k^\#, W_k^{\#\#}$ are in $[0, Y_k]$.

Observe that X_k and Y_k are independent of W_k^0 , hence upon taking expectations, certain terms factor (this is the purpose of introducing W_k^0 instead of using Taylor expansion around W_k). Take expectations in the above equations and subtract the second set from the first set. As the first two moments of X_k match with those of Y_k , the first terms cancel and we also have

$$\mathbf{E}[X_k^2 \mathbf{1}_{|X_k| \leq A}] - \mathbf{E}[Y_k^2 \mathbf{1}_{|Y_k| \leq A}] = -\mathbf{E}[X_k^2 \mathbf{1}_{|X_k| > A}] + \mathbf{E}[Y_k^2 \mathbf{1}_{|Y_k| > A}].$$

Thus, after a little manipulation (note that when terms do not factor, we put absolute values inside and then we don't get a difference but a sum as the bound), we get

$$\begin{aligned} |\mathbf{E}[h(W_{k+1})] - \mathbf{E}[h(W_k)]| &\leq \|\partial_k^2 h\|_{\text{sup}} (\mathbf{E}[X_k^2 \mathbf{1}_{|X_k| > A}] + \mathbf{E}[Y_k^2 \mathbf{1}_{|Y_k| > A}]) \\ &\quad + \frac{1}{6} \|\partial_k^3 h\|_{\text{sup}} (\mathbf{E}[|X_k|^3 \mathbf{1}_{|X_k| \leq A}] + \mathbf{E}[|Y_k|^3 \mathbf{1}_{|Y_k| \leq A}]). \end{aligned}$$

Summing over k , we get the statement in the theorem. ■

Semicircle law for Wigner matrices

We already know the semi-circle law for GOE. We show that the Stieltjes' transform of any Wigner matrix is close to that of the GOE matrix, using the invariance principle. To this end, fix $z = u + iv \in \mathbb{H}$ and define $h : \mathbb{R}^{n(n+1)/2} \mapsto \mathbb{R}$ by

$$h(x) = \frac{1}{n} \text{tr}(zI - M(x))^{-1}$$

where $M(x)$ is the symmetric matrix whose (i, j) entry (for $i \leq j$) is $x_{i,j}$ (here we think of coordinates in $\mathbb{R}^{n(n+1)/2}$ as indexed by (i, j) , $i \leq j$ and correspondingly write $\partial_{(i,j)}$ for the partial derivative w.r.t. $x_{i,j}$).

To calculate the derivatives of h , introduce the matrix $H_{(i,j)}$ which has 1 at the (i, j) and (j, i) entries and zeros elsewhere. Then, $\partial_{(i,j)}M(x) = H_{(i,j)}$. Hence,

$$\begin{aligned}\partial_{(i,j)}h(x) &= \frac{1}{n}\text{tr}\{(zI - M(x))^{-2}H_{(i,j)}\}, \\ \partial_{(i,j)}^2h(x) &= \frac{1}{n}\text{tr}\{(zI - M(x))^{-3}H_{(i,j)}^2\}, \\ \partial_{(i,j)}^3h(x) &= \frac{1}{n}\text{tr}\{(zI - M(x))^{-4}H_{(i,j)}^3\}.\end{aligned}$$

We need bounds for these. If u, w are vectors, observe that

$$\text{tr}\{(zI - M(x))^{-p}uw^t\} = w^t(zI - M(x))^{-1}u = \sum_{k=1}^n \frac{w^t y_k y_k^t u}{(z - \lambda_k)^p}$$

where $M(x) = \sum_{k=1}^n \lambda_k y_k y_k^t$ is the spectral decomposition of M . As $|z - \lambda_k| \geq v$, we get

$$|\text{tr}\{(zI - M(x))^{-p}uw^t\}| \leq \frac{1}{v^p} \sum_{k=1}^n |\langle y_k, v \rangle| \cdot |\langle y_k, u \rangle| \leq \frac{1}{v^p} \|u\| \|w\|$$

since the orthonormality of y_k s implies that $\|u\|^2 = \sum_{k=1}^n |\langle y_k, u \rangle|^2$ and similarly for w .

To use this, note that $H_{(i,j)} = e_i e_j^t + e_j e_i^t$ where e_i are the standard co-ordinate vectors. Hence also $H_{(i,j)}^2 = e_i e_i^t + e_j e_j^t$ and $H_{(i,j)}^3 = H_{(i,j)}$. Thus, we arrive at the bounds

$$|\partial_{(i,j)}h(x)| \leq \frac{2}{v^2 n}, \quad |\partial_{(i,j)}^2h(x)| \leq \frac{2}{v^3 n}, \quad |\partial_{(i,j)}^3h(x)| \leq \frac{2}{v^4 n}.$$

Consequently, $B_2(h) \leq 2v^{-3}n^{-1}$ and $B_3(h) \leq 2v^{-4}n^{-1}$.

Now we are ready to apply the invariance principle. Let $X = (\frac{1}{\sqrt{n}}X_{i,j})_{i \leq j}$ and $Y = (\frac{1}{\sqrt{n}}Y_{i,j})_{i \leq j}$, where $X_{i,j}$ are independent and $Y_{i,j}$ are independent and $\mathbf{E}[X_{i,j}] = \mathbf{E}[Y_{i,j}]$ and $\mathbf{E}[X_{i,j}^2] = \mathbf{E}[Y_{i,j}^2]$. For simplicity, first assume that all the variables have third moments bounded by γ . By Corollary 43, we get

$$|\mathbf{E}[h(X)] - \mathbf{E}[h(Y)]| \leq \frac{1}{3} \frac{\gamma}{n^{5/2}} B_3(h) \frac{n(n+1)}{2} \leq \frac{\gamma B_3(h)}{3} \frac{1}{\sqrt{n}}.$$

In particular, if $\mathbf{E}[X_{i,j}] = 0$ for all $i \leq j$ and $\mathbf{E}[X_{i,j}^2] = 1$ for $i < j$ and $\mathbf{E}[X_{i,i}^2] = 2$, then taking Y to be the GOE matrix, we see that

$$\left| \mathbf{E} \left[\text{tr} \left(zI - \frac{1}{\sqrt{n}} X_n \right)^{-1} \right] - \mathbf{E} \left[\text{tr} \left(zI - \frac{1}{\sqrt{n}} Y_n \right)^{-1} \right] \right| \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Since the ESD of Y_n/\sqrt{n} converges in probability to $\mu_{s.c.}$, its Stieltjes' transform $\frac{1}{n}\text{tr}(zI - \frac{1}{\sqrt{n}}Y_n)^{-1}$ converges in probability to $G_{\mu_{s.c.}}$, the Stieltjes' transform of the semi-circle distribution. Therefore, from the above comparison of expectations, we see that $\frac{1}{n}\text{tr}(zI - \frac{1}{\sqrt{n}}Y_n)^{-1}$ also converges in probability to $G_{\mu_{s.c.}}$. Hence the ESD of X_n/\sqrt{n} converges in probability to $\mu_{s.c.}$. In particular, this proves semi-circle law for Wigner matrices, under third moment assumption.

But we do not need the third moment assumption. If we apply the invariance principle in its original form, we get

$$\begin{aligned} |\mathbf{E}[h(X)] - \mathbf{E}[h(Y)]| &\leq \frac{2}{v^3 n^2} \sum_{i \leq j} \mathbf{E}[X_{i,j}^2 \mathbf{1}_{|X_{i,j}| > A\sqrt{n}}] + \mathbf{E}[Y_{i,j}^2 \mathbf{1}_{|Y_{i,j}| > A\sqrt{n}}] \\ &\quad + \frac{1}{2v^4 n^{5/2}} \sum_{i \leq j} \mathbf{E}[|X_{i,j}|^3 \mathbf{1}_{|X_{i,j}| \leq A\sqrt{n}}] + \mathbf{E}[|Y_{i,j}|^3 \mathbf{1}_{|Y_{i,j}| \leq A\sqrt{n}}] \end{aligned}$$

We make the following assumptions on moments.

1. $\mathbf{E}[X_{i,j}] = \mathbf{E}[Y_{i,j}]$ for all $i \leq j$.
2. $\mathbf{E}[X_{i,j}^2] = \mathbf{E}[Y_{i,j}^2] = 1$ for all $i < j$ and $\mathbf{E}[X_{i,i}^2] = \mathbf{E}[Y_{i,i}^2] = \sigma^2$ for all i .
3. Pastur's condition: $\frac{1}{n^2} \sum_{i \leq j} \mathbf{E}[X_{i,j}^2 \mathbf{1}_{|X_{i,j}| \geq \varepsilon\sqrt{n}}] \rightarrow 0$ as $n \rightarrow \infty$ and similarly for Ys .

Then, take $A = \varepsilon$ to be small and use the bound $|x|^3 \mathbf{1}_{|x| \leq \varepsilon\sqrt{n}}$ to write

$$\begin{aligned} |\mathbf{E}[h(X)] - \mathbf{E}[h(Y)]| &\leq \frac{2}{v^3 n^2} \sum_{i \leq j} \mathbf{E}[X_{i,j}^2 \mathbf{1}_{|X_{i,j}| > \varepsilon\sqrt{n}}] + \mathbf{E}[Y_{i,j}^2 \mathbf{1}_{|Y_{i,j}| > \varepsilon\sqrt{n}}] \\ &\quad + \frac{\varepsilon}{2v^4 n^2} \sum_{i \leq j} \mathbf{E}[|X_{i,j}|^2] + \mathbf{E}[|Y_{i,j}|^2]. \end{aligned}$$

As $n \rightarrow \infty$, the first summand converges to zero (by Pastur's condition) and the second summand converges to $C\varepsilon/v^4$. Then let $\varepsilon \rightarrow 0$. The conclusion is that semi-circle law holds for Wigner matrices where $X_{i,j}$, $i < j$, are allowed to have different distributions (but zero mean, unit variance) provided Pastur's condition is satisfied.

Remark 45. We skipped a couple of points in the end. What we must do is assume the conditions on X , and take Y to be the GOE matrix (in which case the conditions are satisfied). Then from the convergence in probability of ESD of Y_n to $\mu_{s.c.}$ we go to the convergence in probability of the Stieltjes' transform of Y_n to $G_{\mu_{s.c.}}$ and then using the comparison, to the convergence in probability of the Stieltjes' transform of X_n to $G_{\mu_{s.c.}}$, and thence to the convergence in probability of the ESD of X_n to $\mu_{s.c.}$

One may question applying the invariance principle to X_n and GOE, unless it is assumed that $\mathbf{E}[X_{i,i}^2] = 2$. In fact, the diagonal terms are irrelevant, as will be seen in Exercise ??.

Chapter 8

Moment method applied to GOE and connections to enumeration problems

Expected ESD of the GOE matrix

Let X_n be the $n \times n$ GOE matrix. Let L_n be the empirical spectral distribution of X_n/\sqrt{n} and let \bar{L}_n be the expected empirical spectral distribution. We already know that L_n converges in probability to the semi-circle distribution. That also implies that \bar{L}_n converges to $\mu_{s.c.}$ (why?). But we shall prove the latter directly now, by the method of moments. The goal is to bring out certain interesting combinatorics that comes out of Gaussian matrix integrals.

The starting point for the method of moments is the relationship

$$\int x^p \bar{L}_n(dx) = \frac{1}{n^{1+\frac{p}{2}}} \sum_{i_1, \dots, i_p=1}^n \mathbf{E}[X_{i_1, i_2} \dots X_{i_p, i_1}] \quad (1)$$

To evaluate the right hand side, we need the following important fact about expectations of products of Gaussians. To state it, we need the notion of a *matching* of the set $\{1, 2, \dots, n\}$ which is any partitioning of the set into pairs (subsets of size 2). The collection of all matchings of this set will be denoted by \mathcal{M}_n . Clearly this is empty if n is odd. For $n = 2q$, it may be checked that \mathcal{M}_n has cardinality equal to $(2q-1) \times (2q-3) \times \dots \times 3 \times 1$ (why?).

Example 46. \mathcal{M}_4 consists of 3 elements,

$$\{\{1, 2\}, \{3, 4\}\}, \{\{1, 3\}, \{2, 4\}\}, \{\{1, 4\}, \{2, 3\}\}.$$

Lemma 47 (Wick formula). *Let (X_1, \dots, X_n) be jointly complex Gaussian with zero means and covariance matrix $\Sigma = (\sigma_{i,j})_{i,j \leq n}$. Then*

$$\mathbf{E}[X_1 \dots X_n] = \sum_{M \in \mathcal{M}_n} \prod_{\{i,j\} \in M} \sigma_{i,j}.$$

As an example,

$$\mathbf{E}[X_1 X_2 X_3 X_4] = \sigma_{1,2} \sigma_{3,4} + \sigma_{1,3} \sigma_{2,4} + \sigma_{1,4} \sigma_{2,3}.$$

Proof of Lemma 47. When n is odd, the right side is an empty sum, zero by definition. The left side is zero because of the symmetry $(X_1, \dots, X_n) \stackrel{d}{=} (-X_1, \dots, -X_n)$. Henceforth, $n = 2q$.

We start with the joint characteristic function

$$\mathbf{E}[e^{i(t_1 X_1 + \dots + t_n X_n)}] = \exp \left\{ -\frac{1}{2} \sum_{i,j=1}^n t_i t_j \sigma_{i,j} \right\}$$

which follows from noticing that $t_1 X_1 + \dots + t_n X_n \sim N(0, \sum_{i,j=1}^n t_i t_j \sigma_{i,j})$. Differentiate with respect to t_i s and to get

$$i^n \mathbf{E} \left[\prod_{i=1}^n X_i \right] = \frac{\partial^n}{\partial t_1 \dots \partial t_n} \exp \left\{ -\frac{1}{2} \sum_{i,j=1}^n t_i t_j \sigma_{i,j} \right\} \Big|_{t=0}.$$

The left side is $(-1)^q \mathbf{E}[X_1 \dots X_{2q}]$. On the right, the partial derivative with respect to t_1 brings down a factor of $\sum_{j=1}^n t_j \sigma_{1,j}$. One of the other derivatives $\partial/\partial t_j$ must operate on this factor, otherwise the final expression will vanish when we set $t = 0$. This gives a match $\{1, j\}$. Continuing to argue similarly with the remaining variables, we get the expression given in the statement of the lemma. ■

The idea is to use the formula (1) and evaluate the expectation on the right hand side with the help of the Wick. The rest of the work is in keeping track of the combinatorics to see how the semicircle moments emerge. To get the idea, we first do it by hand for a few small values of q in (1). Remember that $X_{i,i} \sim N(0, 2)$ and $X_{i,j} \sim N(0, 1)$ for $i < j$. Also recall that the even moments of the semi-circle distribution are given by the Catalan numbers.

- (i) Case, $q = 1$. $\mathbf{E}[X_{i,j} X_{j,i}] = 1$ for $j \neq i$ and 2 for $j = i$. Hence $\mathbf{E}[\text{tr}(X^2)] = 2n + 2 \binom{n}{2} = n^2 + n$ and

$$\int x^2 \bar{L}_n(dx) = \frac{1}{n^2} \mathbf{E}[\text{tr} X^2] = 1.$$

(ii) Case $q = 2$. From the Wick formula for real Gaussians, $\mathbf{E}[X_{i,j}X_{j,k}X_{k,\ell}X_{\ell,i}]$ becomes

$$\begin{aligned} &= \mathbf{E}[X_{i,j}X_{j,k}]\mathbf{E}[X_{k,\ell}X_{\ell,i}] + \mathbf{E}[X_{i,j}X_{k,\ell}]\mathbf{E}[X_{j,k}X_{\ell,i}] + \mathbf{E}[X_{i,j}X_{\ell,i}]\mathbf{E}[X_{j,k}X_{k,\ell}] \\ &= (\delta_{i,k} + \delta_{i,j,k}) + (\delta_{i,k}\delta_{j,\ell} + \delta_{i,\ell}\delta_{j,k})(\delta_{i,k}\delta_{j,\ell} + \delta_{i,j}\delta_{k,\ell}) + (\delta_{j,\ell} + \delta_{i,j,\ell})(\delta_{j,\ell} + \delta_{j,k,\ell}) \end{aligned}$$

corresponding to the three matchings $\{\{1,2\},\{3,4\}\}$, $\{\{1,3\},\{2,4\}\}$, $\{\{1,4\},\{2,3\}\}$ respectively. Observe that the diagonal entries are also taken care of, since their variance is 2. This looks messy, but look at the first few terms. When we sum over all i, j, k, ℓ , we get

$$\sum_{i,j,k,\ell} \delta_{i,k} = n^3, \quad \sum_{i,j,k,\ell} \delta_{i,j,k} = n^2, \quad \sum_{i,j,k,\ell} (\delta_{i,k}\delta_{j,\ell})^2 = n^2.$$

It is clear that what matters is how many of the indices i, j, k, ℓ are forced to be equal by the delta functions. The more the constraints, the smaller the contribution upon summing. Going back, we can see that only two terms ($\delta_{i,k}$ in the first summand and $\delta_{j,\ell}$ term in the third summand) contribute n^3 , while the other give n^2 or n only.

$$\int x^4 \bar{L}_n(dx) = \frac{1}{n^3} \mathbf{E}[\text{tr}X^4] = \frac{1}{n^3} \sum_{i,j,k,\ell} (\delta_{i,k} + \delta_{j,\ell}) + \frac{1}{n^3} O(n^2) = 2 + O(n^{-1}).$$

Observe that the two non-crossing matchings $\{\{1,2\},\{3,4\}\}$ and $\{\{1,4\},\{2,3\}\}$ contributed 1 each, while the crossing-matching $\{\{1,3\},\{2,4\}\}$ contributed zero in the limit. Thus, we find that $\int x^4 \bar{L}_n(dx) \rightarrow \int x^4 \mu_{s.c.}(dx)$

(iii) Case $q = 3$. We need to evaluate $\mathbf{E}[X_{i_1,i_2}X_{i_2,i_3} \dots X_{i_6,i_1}]$. By the wick formula, we get a sum over matching of [6]. Consider two of these matchings.

(a) $\{1,4\},\{2,3\},\{5,6\}$: This is a non-crossing matching. We get

$$\begin{aligned} &\mathbf{E}[X_{i_1,i_2}X_{i_4,i_5}]\mathbf{E}[X_{i_2,i_3}X_{i_3,i_4}]\mathbf{E}[X_{i_5,i_6}X_{i_6,i_1}] \\ &= (\delta_{i_1,i_4}\delta_{i_2,i_5} + \delta_{i_1,i_5}\delta_{i_2,i_4})(\delta_{i_2,i_4} + \delta_{i_2,i_3,i_4})(\delta_{i_5,i_1} + \delta_{i_5,i_1,i_6}) \\ &= \delta_{i_1,i_5}\delta_{i_2,i_4} + [\dots]. \end{aligned}$$

When we sum over i_1, \dots, i_6 , the first summand gives n^4 while all the other terms (pushed under $[\dots]$) give $O(n^3)$. Thus the contribution from this matching is $n^4 + O(n^3)$.

(b) $\{1, 5\}, \{2, 6\}, \{3, 4\}$: A crossing matching. We get which is equal to

$$\begin{aligned} & \mathbf{E}[X_{i_1, i_2} X_{i_5, i_6}] \mathbf{E}[X_{i_2, i_3} X_{i_6, i_1}] \mathbf{E}[X_{i_3, i_4} X_{i_4, i_5}] \\ &= (\delta_{i_1, i_5} \delta_{i_2, i_6} + \delta_{i_1, i_6} \delta_{i_2, i_5}) (\delta_{i_2, i_6} \delta_{i_3, i_1} + \delta_{i_2, i_1} \delta_{i_3, i_6}) (\delta_{i_3, i_5} + \delta_{i_3, i_4, i_5}) \end{aligned}$$

It is easy to see that all terms are $O(n^3)$. Thus the total contribution from this matching is $O(n^3)$.

We leave it as an exercise to check that all crossing matchings of [6] give $O(n^3)$ contribution while the non-crossing ones give $n^4 + O(n^3)$. Thus,

$$\int x^6 \bar{L}_n(dx) = \frac{1}{n^4} \mathbf{E}[\text{tr} X^6] = \frac{1}{n^4} (C_6 n^4 + O(n^3)) \rightarrow C_6 = \int x^6 \mu_{s.c.}(dx).$$

The general case: We need some preparation in combinatorics.

Definition 48. Let P be a polygon with $2q$ vertices labeled $1, 2, \dots, 2q$. A *gluing* of P is a matching of the edges into pairs along with an assignment of sign $\{+, -\}$ to each matched pair of edges. Let \mathcal{M}_{2q}^\dagger denote the set of all gluings of P . Thus, there are $2^q(2q-1)!!$ gluings of a polygon with $2q$ sides.

Further, let us call a gluing $M \in \mathcal{M}_{2q}^\dagger$ to be *good* if the underlying matching of edges is non-crossing and the orientations are such that matched edges are oriented in opposite directions. That is, $[r, r+1]$ can be matched by $[s+1, s]$ but not with $[s, s+1]$. The number of good matchings is the Catalan number C_q .

Example 49. Let P be a quadrilateral with vertices $1, 2, 3, 4$. Consider the gluing $M = \{\{[1, 2], [4, 3]\}, \{[2, 3], [1, 4]\}\}$. It means that the edge $[1, 2]$ is identified with $[4, 3]$ and the edge $[2, 3]$ is identified with $[1, 4]$. If we actually glue the edges of the polygon according to these rules, we get a torus¹. The gluing $M' = \{\{[1, 2], [3, 4]\}, \{[2, 3], [1, 4]\}\}$ is different from M . What does the gluing give us? We identify the edges $[2, 3]$ and $[1, 4]$ as before, getting a cylinder. Then we glue the two circular ends in *reverse* orientation. Hence the resulting surface is Klein's bottle.

¹Informally, gluing means just that. Formally, gluing means that we fix homeomorphism $f: [1, 2] \rightarrow [3, 4]$ such that $f(1) = 3$ and $f(2) = 4$ and a homeomorphism $g: [2, 3] \rightarrow [1, 4]$ such that $g(2) = 1$ and $g(3) = 4$. Then define the equivalences $x \sim f(x)$, $y \sim g(y)$. The resulting quotient space is what we refer to as the glued surface. It is locally homeomorphic to \mathbb{R}^2 which justifies the word "surface". The quotient space does not depend on the choice of homeomorphisms f and g . In particular, if we reverse the orientations of all the edges, we get the same quotient space.

For a polygon P and a gluing M , let V_M denote the number of distinct vertices in P after gluing by M . In other words, the gluing M gives an equivalence relationship on the vertices of P , and V_M is the number of equivalence classes.

Lemma 50. *Let P be a polygon with $2q$ edges and let $M \in \mathcal{M}_{2q}^\dagger$. Then $V_M \leq q + 1$ with equality if and only if M is good.*

Assuming the lemma we prove the convergence of \bar{L}_n to semicircle.

$$\begin{aligned} \mathbf{E}[X_{i_1, i_2} \cdots X_{i_{2q}, i_1}] &= \sum_{M \in \mathcal{M}_{2q}} \prod_{\{r, s\} \in M} \mathbf{E}[X_{i_r, i_{r+1}} X_{i_s, i_{s+1}}] \\ &= \sum_{M \in \mathcal{M}_{2q}} \prod_{\{r, s\} \in M} (\delta_{i_r, i_s} \delta_{i_{r+1}, i_{s+1}} + \delta_{r, s+1} \delta_{r+1, s}) \\ &= \sum_{M \in \mathcal{M}_{2q}^\dagger} \prod_{\{e, f\} \in M} \delta_{i_e, i_f}. \end{aligned}$$

Here for two edges e, f , if $e = [r, r + 1]$ and $s = [s, s + 1]$ (or $f = [s + 1, s]$), then δ_{i_e, i_f} is just $\delta_{i_r, i_s} \delta_{i_{r+1}, i_{s+1}}$ (respectively $\delta_{i_r, i_{s+1}} \delta_{i_{r+1}, i_s}$). Also observe that diagonal entries are automatically taken care of since they have variance 2 (as opposed to variance 1 for off-diagonal entries).

Sum (8) over i_1, \dots, i_{2q} and compare with Recall (1) to get

$$\int x^{2q} \bar{L}_n(dx) = \frac{1}{n^{1+q}} \sum_{M \in \mathcal{M}_{2q}^\dagger} \sum_{i_1, \dots, i_{2q}} \prod_{\{e, f\} \in M} \delta_{i_e, i_f} = \frac{1}{n^{1+q}} \sum_{M \in \mathcal{M}_{2q}^\dagger} n^{V_M}. \quad (2)$$

We explain the last equality. Fix M , and suppose some two vertices r, s are identified by M . If we choose indices i_1, \dots, i_{2q} so that some $i_r \neq i_s$, then the δ -functions force the term to vanish. Thus, we can only choose one index for each equivalence class of vertices. This can be done in n^{V_M} ways.

Invoke Lemma 50, and let $n \rightarrow \infty$ in (2). Good matchings contribute 1 and others contribute zero in the limit. Hence, $\lim_{n \rightarrow \infty} \int x^{2q} \bar{L}_n(dx) = C_q$. The odd moments of \bar{L}_n as well as $\mu_{s.c}$ are obviously zero. This completes the proof that $\bar{L}_n \rightarrow \mu_{s.c}$. ■

It remains to prove Lemma 50. If one knows a little algebraic topology, this is clear. First we describe this “high level picture”. For the benefit of those not unfamiliar with Euler characteristic and genus of a surface, we give a self-contained proof later².

²However, the connection given here is at the edge of something deep. Note the exact formula for GOE

A detour into algebraic topology: Recall that a *surface* is a topological space in which each point has a neighbourhood that is homeomorphic to the open disk in the plane. For example, a polygon (where we mean the interior of the polygon as well as its boundary) is not a surface, since points on the boundary do not have disk-like neighbourhoods. A sphere, torus, Klein bottle, projective plane are all surfaces. In fact, these can be obtained from the square P_4 by the gluing edges appropriately.

1. Let $P = P_{2q}$ and $M \in \mathcal{M}_{2q}^\dagger$. After gluing P according to M , we get a surface (means a topological space that is locally homeomorphic to an open disk in the plane) which we denote P/M . See examples 49.
2. If we project the edges of P via the quotient map to P/M , we get a graph G_M drawn (or “embedded”) on the surface P/M . A graph is a combinatorial object, defined by a set of vertices V and a set of edges E . An *embedding of a graph* on a surface is a collection of function $f : V \rightarrow S$ and $f_e : [0, 1] \rightarrow S$ for each $e \in E$ such that f is one-one, for $e = (u, v)$ the function f_e is a homeomorphism such that $f_e(0) = f(u)$ and $f_e(1) = f(v)$, and such that $f_e((0, 1))$ are pairwise disjoint. For an embedding, each connected component of $S \setminus \cup_{e \in E} f_e[0, 1]$ is called a *face*. A *map* is an embedding of the graph such that each face is homeomorphic to a disk.
3. For any surface, there is a number χ called the *Euler characteristic* of the surface, such that for any map drawn on the surface, $V - E + F = \chi$, where V is the number of vertices, E is the number of edges and F is the number of faces of the graph. For example, the sphere has $\chi = 2$ and the torus has $\chi = 0$. The Klein bottle also has $\chi = 0$. The *genus* of the surface is related to the Euler characteristic by $\chi = 2 - 2g$.
4. A general fact is that $\chi \leq 2$ for any surface, with equality if and only if the surface is simply connected (in which case it is homeomorphic to the sphere).
5. The graph G_M has $F = 1$ face (the interior of the polygon is the one face, as it is home-

$\int t^{2q} d\bar{L}_n(t) = \sum_{g=0}^q n^{-g} A_{q,g}$, where $A_{q,g}$ is the number of gluings of P_{2q} that lead to a surface with Euler characteristic $2 - 2g$. The number g is called the genus. The right hand side can be thought of as a generating function for the number $A_{q,g}$ in the variable n^{-1} . This, and other related formulas express generating functions for maps drawn on surfaces of varying genus in terms of Gaussian integrals over hermitian matrices, which is what the left side is. In particular, such formulas have been used to study “random quadrangulations of the sphere”, and other similar objects, using random matrix theory. Random planar maps are a fascinating and active research area in probability, motivated by the notion of “quantum gravity” in physics.

omorphically mapped under the quotient map), $E = q$ edges (since we have merged $2q$ edges in pairs) and $V = V_M$ vertices. Thus, $V_M = \chi(G_M) - 1 + q$. By the previous remark, we get $V_M \leq q + 1$ with equality if and only if P/M is simply connected.

6. Only good gluings lead to simply connected P/M .

From these statements, it is clear that Lemma 50 follows. However, for someone unfamiliar with algebraic topology, it may seem that we have restated the problem without solving it. Therefore we give a self-contained proof of the lemma now.

Proof of Lemma 50. After gluing by M , certain vertices of P are identified. If $V_M > q$, there must be at least one vertex, say r , of P that was not identified with any other vertex. Clearly, then M must glue $[r - 1, r]$ with $[r, r + 1]$. Glue these two edges, and we are left with a polygon Q with $2q - 2$ sides with an edge sticking out. For r to remain isolated, it must not enter the gluing at any future stage. This means, the gluing will continue within the polygon Q . Inductively, we conclude that Q must be glued by a good gluing. Retracing this to P , we see that M must be a good gluing of P . Conversely, if M is a good gluing, it is easy to see that $V_M = q + 1$ ³. ■

Remark 51. We showed the method of moments for GOE to emphasize the exact combinatorial formulas for finite n . However, if we care only about the limits of moments, it is possible to carry out the same proof for general Wigner matrices. Some points to note however. The Wick formula is not strictly necessary - here the random variables are either equal or uncorrelated. By throwing away terms which are not pair matchings (they can be shown to be negligible), we can recover the limit. Secondly, for general Wigner matrices, the entries may not have all moments, hence an initial truncation argument is employed. This is all done in the more general setting of the next section.

Semi-circle law for a more general class of random matrices

Let $G_n = (V_n, E_n)$ be a sequence of graphs with vertex set V_n and edge set E_n such that $|V_n| \rightarrow \infty$. Without loss of generality, we take $V_n = [n]$ and make the assumption that G_n is

³Thanks to R. Deepak for this neat proof. Another way to state it is as follows. Consider the polygon P (now a topological space homeomorphic to the closed disk). Glue it by M to get a quotient space P/M . Consider the graph G formed by the edges of P (so G is a cycle). Project to G to P/M . The resulting graph G_M is connected (since G was), and has q edges. Hence it can have at most $q + 1$ vertices, and it has $q + 1$ vertices if and only if the G_M is a tree. Work backwards to see that M must be good. The induction step is implicit in proving that a graph has $V \leq E + 1$ with equality for and only for trees.

regular, i.e., all vertices of G_n have d_n neighbours (then $|E_n| = nd_n/2$) and that $d_n \rightarrow \infty$.

Let $X_n = (X_{i,j})_{i,j \leq n}$ be a real symmetric random matrix satisfying the following assumptions. 1. $X_{i,i}$, $i \leq n$, are i.i.d. with zero mean and finite variance σ^2 , 2. $X_{i,j}$, $i \sim j$, are i.i.d. with mean zero and variance 1, 3. $X_{i,j} = 0$ if $i \not\sim j$ in G_n .

Theorem 52. *With the above setting, the expected empirical spectral distributions of $\frac{1}{\sqrt{d_n}}A_n$ converge to the semi-circle distribution on $[-2, 2]$.*

In proving the theorem, first assume that $X_{i,i}$ and $X_{i,j}$ are bounded random variables. Then, all moments exist, and we only need to show that $\frac{1}{nd_n^{p/2}}\mathbf{E}[\text{tr}(A_n^p)]$ converges to $C_{p/2}$ for even p and to 0 for odd p . As always, we start with

$$\frac{1}{nd_n^{p/2}} \sum_{i_1 \sim i_2 \sim \dots \sim i_p \sim i_1} \mathbf{E}[X_{i_1, i_2} X_{i_2, i_3} \dots X_{i_p, i_1}].$$

Consider a summand. If any variable $X_{j,k}$ occurs only once, then the expectation is zero. If any term occurs

Chapter 9

Free probability and random matrices

The appearance of the semicircle law as the limiting spectral distribution of Wigner matrices will be made more natural and put in a larger context by introducing the calculus of free probability. It is a form of non-commutative probability, with a specific definition of independence. The combinatorial approach here will require the notion of Mobius function on a lattice, particularly as applied to (a) \mathcal{P}_n , the lattice of partitions of $[n]$ and (b) NC_n , the lattice of non-crossing partitions of $[n]$. These are explained in Appendix 10.

Cumulants and moments in classical probability

Let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space. For random variables X_i on this probability space having moments of all orders, define $m_n[X_1, \dots, X_n] = \mathbf{E}[X_1 X_2 \dots X_n]$, whenever the expectation exists. We will also write $m_0 = 1$. The function $m[\bullet]$ is called the *moment* function.

The lattice of partitions will play an important role. A partition $\Pi \in \mathcal{P}_n$ of $[n]$ will be written as $\{\Pi_1, \dots, \Pi_\ell\}$, where Π_j are the blocks of Π , arranged (just for definiteness) in increasing order of the smallest elements in the blocks. We write $|\Pi_j|$ for the cardinality of Π_j and $\ell(\Pi) = \ell$ for the number of blocks in Π .

Definition 53. Define the *cumulant function* $\kappa_n[X_1, \dots, X_n]$ by the equations

$$m_n[X_1, \dots, X_n] = \sum_{\Pi \in \mathcal{P}_n} \prod_{j=1}^{\ell(\Pi)} \kappa_{|\Pi_j|}[X_{\Pi_j}]. \quad (1)$$

Here $[X_{\Pi_j}]$ is the short form for $[X_{k_1}, \dots, X_{k_r}]$.

To see that this defines the cumulant function unambiguously, consider the first three equations

$$\kappa_1[X] = m_1[X]$$

$$\kappa_2[X, Y] = m_2[X, Y] - \kappa_1[X]\kappa_1[Y]$$

$$\kappa_3[X, Y, Z] = m_3[X, Y, Z] - \kappa_2[X, Y]\kappa_1[Z] - \kappa_2[X, Z]\kappa_1[Y] - \kappa_2[Y, Z]\kappa_1[X] + \kappa_1[X]\kappa_1[Y]\kappa_1[Z]$$

It is clear that we can define $\kappa_1[\bullet]$ from the first equation, $\kappa_2[\bullet]$ from the second and so on, inductively. To write various formulas more succinctly, for any $\Pi \in \mathcal{P}_n$ define

$$m_\Pi[X_1, \dots, X_n] = \prod_{j=1}^{\ell(\Pi)} m_{|\Pi_j|}[X_{\Pi_j}], \quad \kappa_\Pi[X_1, \dots, X_n] = \prod_{j=1}^{\ell(\Pi)} \kappa_{|\Pi_j|}[X_{\Pi_j}].$$

In this notation, the equations defining cumulants may be written as $m_n[X] = \sum_{\Pi \in \mathcal{P}_n} \kappa_\Pi[X]$ where $X = (X_1, \dots, X_n)$. Further, it follows that (below, let $\mathcal{P}(S)$ denote the collection of partitions of the set S)

$$m_\Pi[X] = \prod_{j=1}^{\ell(\Pi)} \sum_{\Gamma_j \in \mathcal{P}(\Pi_j)} \kappa_{\Gamma_j}[X_{\Pi_j}] = \sum_{\Gamma \leq \Pi} \kappa_\Gamma[X].$$

In the last equality, we just used that a choice of $\Gamma_1, \dots, \Gamma_{\ell(\Pi)}$, when put together, gives a $\Gamma \in \mathcal{P}_n$ that is a refinement of Π . If we fix the random variables, then $m_\bullet[X]$ and $\kappa_\bullet[X]$ become functions on the lattice \mathcal{P}_n related by the the above relation. By the Mobius inversion formula,

$$\kappa_\Pi[X] = \sum_{\Gamma \leq \Pi} \mu(\Gamma, \Pi) m_\Gamma[X].$$

From the explicit form of the Mobius function for this lattice, we get

$$\kappa_n[X] = \sum_{\Gamma \in \mathcal{P}_n} (-1)^{\ell(\Gamma)-1} (\ell(\Gamma) - 1)! m_\Gamma[X].$$

Univariate situation: A special case to note is when all X_i s are equal to one random variable ξ . In that case write $m_n(\xi) = m_n[X]$ (where $[X] = [\xi, \dots, \xi]$) and $\kappa_n(\xi) = \kappa_n[X]$. Then $m_n(\xi) = \mathbf{E}[\xi^n]$ and

$$\kappa_n(\xi) = \sum_{\Pi \in \mathcal{P}_n} (-1)^{\ell(\Pi)-1} (\ell(\Pi) - 1)! \prod_{j=1}^{\ell(\Pi)} \mathbf{E}[\xi^{|\Pi_j|}].$$

In particular, $\kappa_1(\xi) = \mathbf{E}[\xi]$ and $\kappa_2(\xi) = \text{Var}(\xi)$.

Example 54. Let $\xi \sim N(0, 1)$. Then, we claim that $\kappa_n(\xi) = 1$ if $n = 2$ and $\kappa_n(\xi) = 0$ for all other n . Indeed, with this definition of κ , we have

$$\sum_{\Pi \in \mathcal{P}_n} \prod_{j=1}^{\ell(\Pi)} \kappa_{|\Pi_j|}[X_{\Pi_j}] = \#\{\Pi \in \mathcal{P}_n : \text{each block of } \Pi \text{ has two elements}\}.$$

We know that the right hand side is $m_n(\xi)$, hence the moment-cumulant relations are satisfied.

Properties of classical cumulants: The following lemma collects some basic properties of the cumulant function. Of particular importance is the relationship between independence and the cumulant function.

Lemma 55. 1. *Multilinearity:* $\kappa_{\Pi}[\bullet]$ is multilinear in its arguments.

2. *Symmetry:* $\kappa_{\Pi}[\bullet]$ is invariant if variables within a single block of Π are permuted. In particular, κ_n is symmetric in X_1, \dots, X_n .

3. Assume that $X = (X_1, \dots, X_d)$ is such that $\mathbf{E}[e^{\langle t, X \rangle}] < \infty$ for $t = (t_1, \dots, t_d)$ in a neighbourhood of 0 in \mathbb{R}^d . Let $\varphi_X(t) = \mathbf{E}[e^{\langle t, X \rangle}]$ and $\psi_X(t) = \log \mathbf{E}[e^{\langle t, X \rangle}]$. Then,

$$\begin{aligned} \varphi_X(t) &= \sum_{n=0}^{\infty} \sum_{i_1, \dots, i_n=1}^d \frac{t_{i_1} \dots t_{i_n}}{n!} m_n[X_{i_1}, \dots, X_{i_n}], \\ \psi_X(t) &= \sum_{n=1}^{\infty} \sum_{i_1, \dots, i_n=1}^d \frac{t_{i_1} \dots t_{i_n}}{n!} \kappa_n[X_{i_1}, \dots, X_{i_n}]. \end{aligned}$$

4. Let $U = (X_1, \dots, X_k)$ and $V = (X_{k+1}, \dots, X_d)$. Then, the following are equivalent.

(i) U and V are independent.

(ii) $\kappa_n[X_{i_1}, \dots, X_{i_n}] = 0$ for any $n \geq 1$ and any $i_1, \dots, i_n \in [d]$ whenever there is least one p such that $i_p \leq k$ and at least one q such that $i_q > k$.

Proof. 1. The moment function $m_n[\bullet]$ is clearly multilinear in each of its co-ordinates. Hence, m_{Π} is also multilinear in its co-ordinates (although it is a product of $m_{|\Pi_j|}$ s, observe that only one of the blocks contains a particular index. Then the multilinearity of κ_{Π} follows from the expression for it in terms of the moments.

2. Follows because m_n is symmetric in all its arguments and m_Π is symmetric under permutations of its arguments within blocks.
3. Expand $e^{\langle t, X \rangle} = \sum_n \langle t, X \rangle^n / n!$ and $\langle t, X \rangle^n = \sum_{i_1, \dots, i_n=1}^d t_{i_1} \dots t_{i_n} X_{i_1} \dots X_{i_n}$. Taking expectations gives the expansion for $\varphi_X(t)$. To get the expansion for $\Psi_X(t)$, let $\Psi(t) = \sum_{n=1}^{\infty} \sum_{i_1, \dots, i_n=1}^d \frac{t_{i_1} \dots t_{i_n}}{n!} \kappa_n[X_{i_1}, \dots, X_{i_n}]$ and consider

$$e^{\Psi(t)} = \sum_{n=1}^{\infty} \frac{1}{n!} \sum_{k_1, \dots, k_n=1}^d \kappa_{k_1}$$

4. $U = (X_1, \dots, X_m)$ is independent of $V = (X_{m+1}, \dots, X_n)$ if and only if $\Psi_{(U,V)}(t, s) = \Psi_U(t) + \Psi_V(s)$ for all $t \in \mathbb{R}^m, s \in \mathbb{R}^{n-m}$. By part (b), Ψ_U (respectively, Ψ_V) has an expansion involving $\kappa_k[X_{i_1}, \dots, X_{i_k}]$ where $i_1, \dots, i_k \leq m$ (respectively, $i_1, \dots, i_k > m$). However, $\Psi_{(U,V)}$ has coefficients $\kappa_k[X_{i_1}, \dots, X_{i_k}]$ where i_r range over all of $[n]$. Thus, U and V are independent if and only if $\kappa_k[X_{i_1}, \dots, X_{i_k}] = 0$ whenever there are p, q such that $i_p \leq m$ and $i_q > m$. This proves the equivalence of the two statements. ■

The consequences for the univariate situation is worth summarizing.

Corollary 56. *Let ξ, η be real-valued random variables having finite moment generating functions.*

1. $\mathbf{E}[e^{t\xi}] = \sum_{n=0}^{\infty} \frac{m_n(\xi)}{n!} t^n$ and $\log \mathbf{E}[e^{t\xi}] = \sum_{n=1}^{\infty} \frac{\kappa_n(\xi)}{n!} t^n$.
2. The variables ξ and η are independent if and only if $\kappa_n(\xi + \eta) = \kappa_n(\xi) + \kappa_n(\eta)$ for all $n \geq 1$.
3. If c is a constant, $\kappa_n(\xi + c) = \kappa_n(\xi) + c\delta_{n,1}$.

The proofs are obvious. In the second part, observe that $\kappa_n(\xi + \eta) = \kappa_n[\xi + \eta, \dots, \xi + \eta]$ has 2^n terms when expanded by multilinearity, and all but two terms vanish by independence. The third follows from the second since a constant is independent of any random variable. Since κ_n (except for $n = 1$) remain unchanged under translations, cumulants are also called *semi-invariants*.

Example 57. Let $X \sim \exp(1)$. Then $\varphi_X(t) = (1 - t)^{-1} = \sum_{n \geq 0} t^n$ for $t < 1$. Hence $m_n = n!$. $\Psi_X(t) = -\log(1 - t) = \sum_{n \geq 1} n^{-1} t^n$ which shows that $\kappa_n = (n - 1)!$. If $Y \sim \text{Gamma}(v, 1)$ then for integer values of v it is a sum of i.i.d exponentials, hence $\kappa_n(Y) = v(n - 1)!$. It may be verified directly that this is also true for any $v > 0$.

Example 58. Let $X \sim \text{Pois}(1)$. Then $\mathbf{E}[e^{tX}] = e^{-1+e^t}$. Expanding this, one can check that $m_n = e^{-1} \sum_{k=0}^{\infty} \frac{k^n}{k!}$. It is even easier to see that $\psi_X(t) = -1 + e^t$ and hence $\kappa_n = 1$ for all $n \geq 1$ and hence also $\kappa_{\Pi} = 1$. But then, the defining equation for cumulants in terms of moments shows that $m_n = \sum_{\Pi \in \mathcal{P}_n} \kappa_{\Pi} = |\mathcal{P}_n|$. Thus as a corollary, we have the non-trivial relation $|\mathcal{P}_n| = e^{-1} \sum_{k=0}^{\infty} \frac{k^n}{k!}$, known as Dobinsky's formula.

Remark 59 (Logarithm of exponential generating functions). The relationship between m_n and κ_n just comes from the connection that $\log \varphi = \psi$ where $m_n/n!$ are the coefficient of φ and $\kappa_n/n!$ are coefficients of ψ . The same is true for coefficients of any two power series related this way. A closer look at the expressions for m_n in terms of κ_n or the reverse one shows that if m_n counts some combinatorial objects, then κ_n counts the connected pieces of the same combinatorial object.

For example, in Example 57, $m_n = n!$ counts the number of permutations on n letters while $\kappa_n = (n-1)!$ counts the number of cyclic permutations. As any permutation may be written as a product of disjoint cycles, it makes sense to say that cycles are the only connected permutations.

In Example 58, $m_n = |\mathcal{P}_n|$ while $\kappa_n = 1$. Indeed, the only "connected partition" is the one having only one block $\{1, 2, \dots, n\}$.

In case of $N(0, 1)$, we know that m_n counts the number of matching of $[n]$. What are connected matchings? If $n > 2$, there are no connected matchings! Hence, $\kappa_n = 0$ for $n \geq 3$.

Now we show how cumulants may be used to write a neat proof of (a restricted version of) the central limit theorem.

Proof of central limit theorem assuming mgf exists. Suppose X_1, X_2, \dots are i.i.d with zero mean and unit variance and such that the mgf of X_1 exists in a neighbourhood of zero, then for any fixed $p \geq 1$,

$$\kappa_p(S_n/\sqrt{n}) = n^{-\frac{p}{2}} \kappa[S_n, \dots, S_n] = n^{-\frac{p}{2}} \sum_{1 \leq i_1, \dots, i_p \leq n} \kappa_p[X_{i_1}, \dots, X_{i_p}]$$

by multilinearity of cumulants. If $X_{i_r} \neq X_{i_s}$, the corresponding summand will vanish by the independence of X_j s. Therefore,

$$\kappa_p(S_n/\sqrt{n}) = n^{-\frac{p}{2}} \sum_{j=1}^n \kappa_p[X_j, X_j, \dots, X_j] = n^{-\frac{p}{2}+1} \kappa_p(X_1)$$

which goes to zero for $p \geq 3$. As the first two cumulants are 0 and 1 respectively, we see that the cumulants of S_n/\sqrt{n} converge to cumulants of $N(0, 1)$ and hence the moments converge also. Thus, S_n/\sqrt{n} converges in distribution to $N(0, 1)$. ■

As we have said before, the moment method is very flexible, and in many cases cumulants are a better book-keeping device than moments. See Exercise ???. In this context, we mention a fact that is often useful in that it allows us to consider only cumulants of high enough order.

Fact 60 (Marcinkiewicz). Let X be a random variable with $\kappa_p(X) = 0$ for $p \geq p_0$ for some p_0 . Then $X \sim N(\kappa_1, \kappa_2)$. Hence, if X_n is a sequence of random variables such that $\kappa_p(X_n) \rightarrow 0$ for all large enough p , as $n \rightarrow \infty$. Then X_n converge in distribution to a Gaussian.

Noncommutative probability spaces, free independence

A *unital algebra* \mathcal{A} is a vector space over \mathbb{C} endowed with a multiplication operation $(a, b) \rightarrow ab$ which is assumed to be associative and also distributive over addition and scalar multiplication. In addition we assume that there is a *unit*, denoted 1 , such that $a1 = a = 1a$ for all $a \in \mathcal{A}$. If in addition, there is an *involution operation* $\star : \mathcal{A} \mapsto \mathcal{A}$ that is conjugate linear ($(\alpha a + \beta b)^* = \bar{\alpha}a^* + \bar{\beta}b^*$) and idempotent ($((a^*)^* = a)$), then we say that \mathcal{A} is a \star -algebra.

Definition 61. A *non-commutative probability space* is a pair (\mathcal{A}, φ) where \mathcal{A} is a unital algebra over complex numbers and φ is a linear functional on \mathcal{A} such that $\varphi(\mathbf{1}) = 1$.

If \mathcal{A} is a \star -algebra and φ is positive in the sense that $\varphi(aa^*) \geq 0$ for all $a \in \mathcal{A}$, then we shall say that (\mathcal{A}, φ) is a \star -NCPS.

Elements of \mathcal{A} will take the place of complex-valued random variables and φ will take the place of expectation. The involution allows us to define real-valued variables, i.e., self-adjoint variables ($a^* = a$) and positive variables (those of the form aa^*). This way, the basic properties of expectation: linearity, positivity ($\mathbf{E}[X] \geq 0$ for $X \geq 0$) and unicity ($\mathbf{E}[\mathbf{1}] = 1$) - have the right analogues in the non-commutative setting.

Example 62. Let $\mathcal{A} = \mathbb{C}[x_1, \dots, x_n]$ and $\varphi(P) = \int P(x) d\mu(x)$ where μ is a Borel measure on \mathbb{C}^n with total mass 1. Complex conjugation is an involution on \mathcal{A} . In that setting we shall require μ to be a probability measure for positivity.

Example 63. Let x_1, \dots, x_n be non-commuting variables and let \mathcal{A} be the space of complex polynomials in these variables. This just means that elements of \mathcal{A} are formal linear combinations (with complex coefficients) of the monomials $x_{i_1}^{p_1} x_{i_2}^{p_2} \dots x_{i_k}^{p_k}$ where $1 \leq i_1, \dots, i_k \leq n$ and $p_1, \dots, p_k \geq 1$. When $k = 0$, the empty product is 1. How to add and multiply is clear.

Define $\varphi(P)$ to be the coefficient of the constant term in P . Then (\mathcal{A}, φ) is an NCPS. Again complex conjugation is an involution, and φ is positive.

Example 64. Natural examples of \star -algebras are C^* -algebras¹ If that is the case, we may also call it a C^* -probability space.

Example 65. Let $\mathcal{A} = M_n(\mathbb{C})$ be the space of $n \times n$ complex matrices with the usual operations. For any $u \in \mathbb{C}^n$, we can define $\varphi_u(A) = \langle Au, u \rangle$, which makes this into a \star -probability space $\langle A^*Au, u \rangle = \|Au\|^2 \geq 0$ for any $A \in M_n(\mathbb{C})$.

Another example of expectation is $\varphi(A) = \frac{1}{n}\text{tr}(A)$. Indeed this is nothing but $\frac{1}{n}(\varphi_{u_1} + \dots + \varphi_{u_n})$ where $\{u_1, \dots, u_n\}$ is any orthonormal basis of \mathbb{C}^n .

The previous example generalizes to infinite dimensions.

Example 66. Let $\mathcal{A} := \mathcal{B}(H)$ be the algebra of bounded linear operators on a Hilbert space H . This is a C^* -algebra where the identity I is the unit and taking adjoints is the involution. Let $u \in H$ be a unit vector and define $\varphi(T) = \langle Tu, u \rangle$. Then, φ is a linear functional and $\varphi(I) = 1$. Further, $\varphi(T^*T) = \|Tu\|^2 \geq 0$. Thus, (\mathcal{A}, φ) is a C^* -probability space. Here multiplication is truly non-commutative.

If $\psi(T) = \langle Tv, v \rangle$ for a different unit vector v , then for $0 < s < 1$, the pair $(\mathcal{A}, s\varphi + (1-s)\psi)$ is also a C^* -probability space. φ is called a pure state while $s\varphi + (1-s)\psi$ is called a mixed state. Trace is not well-defined, in general.

The following examples are commutative.

Example 67. Let K be a compact metric space and let $\mathcal{A} = C(K)$ (continuous complex-valued functions). Let μ be any Borel probability measure on K and define $\varphi(f) = \int_K f d\mu$. Then (\mathcal{A}, φ) is a commutative C^* -probability space.

Example 68. Let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space and let $\mathcal{A} = L^\infty(\mathbf{P})$. Let $\mathbf{1}$ be the constant random variable 1. Then \mathcal{A} is a unital algebra. Let $\varphi(X) := \mathbf{E}[X]$ for $X \in \mathcal{A}$. This is also a \star -probability space.

¹By definition, this means that \mathcal{A} has three structures. (a) That of a complex Banach space, (b) that of an algebra and finally, (c) an *involution* $*$: $\mathcal{A} \rightarrow \mathcal{A}$. These operations respect each other as follows. The algebra operations are continuous and respect the norm in the sense that $\|ab\| \leq \|a\|\|b\|$. The involution is also continuous, norm-preserving, and is conjugate linear. Further $(ab)^* = b^*a^*$. Lastly, we have the identity $\|aa^*\| = \|a\|^2$ for all $a \in \mathcal{A}$.

Distributions of random variables: In classical probability, under some conditions, the moment sequence determines the distribution of a random variable, which is a measure on \mathbb{R} . In the non-commutative setting, there is nothing other than the moments. This is summarized below.

- For a single element a in a NCPS, we can compute its moment sequence $m_n(a) = \varphi(a^n)$, $n \geq 0$. This sequence will be called the distribution of a . For several variables a_1, \dots, a_k we compute their joint moments, i.e., the value of φ on all monomials generated by these (non-commuting, in general) variables. That collection is called their joint distribution. If you prefer a more precise definition, the linear functional $L : \mathbb{C}[X_1, \dots, X_n] \mapsto \mathbb{C}$ (where X_i are non-commuting variables) define by $L(P(X_1, \dots, X_k)) := \varphi(P(a_1, \dots, a_k))$, is called the joint distribution of (a_1, \dots, a_k) . The notion extends to infinitely many variables naturally.
- In a \star -NCPS, for a single element, we compute the joint moments of a and a^* , which we call the \star -distribution of a . In the special case when $aa^* = a^*a$ (we then say that a is *normal*), this reduces to a two-dimensional array of number $\varphi(a^k a^{*\ell})$. In particular, for a self-adjoint variable (i.e., if $a = a^*$), then all the information is in the sequence of moments $\varphi(a^n)$. For several variables in a \star -NCPS, we compute joint moments of the elements and their involutions. This whole data comprises their joint distribution.

We said that there is nothing more to the distribution than the moments. There is an important exception. Suppose a is a self-adjoint variable in a \star -NCPS. Then, there exists² a probability measure whose classical moments are equal to the (non-commutative) moments of a . In case this measure is unique (for that issue see Appendix 10), we denote it as μ_a and refer to it as the distribution of a .

If a is not self-adjoint but normal, again we can associate a probability measure μ_a in the complex plane such that $\int z^k \bar{z}^\ell d\mu_a(z) = \varphi(a^k a^{*\ell})$. That measure (if unique) is called the distribution of a . It is worth remarking that in the special case when $\mathcal{A} = \mathcal{B}(H)$ and $\varphi(T) = \langle Te, e \rangle$ for a unit vector e in the Hilbert space H , the distribution of a self-adjoint or normal $T \in \mathcal{A}$ is nothing but its spectral measure at the vector e .

However, for an element that is not self-adjoint or normal, or even for an n -tuple of

²Recall that the necessary condition for existence of a probability measure on \mathbb{R} with given moments is the positive semi-definiteness of the Hankel matrix $(m_{i+j})_{i,j \geq 0}$. In our case, for any scalars c_i (only finitely many non-zero), we have $\sum_{i,j} c_i c_j m_{i+j}(a) = \varphi(bb^*)$ with $b = \sum_i c_i a^i$. Positivity of φ gives positive semi-definiteness of the moment sequence.

self-adjoint variables, there is no meaningful way to associate a distribution on \mathbb{C} or a joint distribution on \mathbb{R}^n . All we have are the moments.

Example 69. Let $H = \ell^2(\mathbb{Z})$ and $\mathcal{A} = \mathcal{B}(H)$ and let $\varphi(T) = \langle Te_0, e_0 \rangle$. Define the left-shift operator L by $(Lx)(n) = x(n+1)$ for $n \in \mathbb{Z}$. Its adjoint is the right-shift $(L^*x)(n) = x(n-1)$. Since $L^n e_0 = e_{-n}$ and $L^{*n} e_0 = e_n$, we see that $\varphi(L^n) = \varphi(L^{*n}) = 0$ for $n \geq 1$. More interestingly,

$$\varphi((L+L^*)^n) = \sum_{\varepsilon \in \{+, -\}^n} \varphi(L^{\varepsilon_1} \dots L^{\varepsilon_n})$$

where L^+, L^- denote L and L^* , respectively. But $L^{\varepsilon_1} \dots L^{\varepsilon_n} e_0 = e_k$ if $\sum_i \varepsilon_i = k$. Thus, $m_n(L+L^*) = 0$ if n is odd and $m_{2n}(L+L^*) = \binom{2n}{n}$. Therefore, $L+L^*$ has arcsine($[-2, 2]$) distribution.

Example 70. Let $H = \ell^2(\mathbb{N})$ and $\mathcal{A} = \mathcal{B}(H)$ and $\varphi(T) = \langle Te_0, e_0 \rangle$. Let L be the left-shift operator as before. Then, L^* is the right-shift operator, i.e., $L^*(x_1, x_2, \dots) = (0, x_1, x_2, \dots)$. Again $\varphi(L^n) = 0$ and $\varphi((L^*)^n) = 0$ for $n \geq 1$. Also, $L+L^*$ has vanishing odd moments and its even moments are $m_{2n}(L+L^*) = C_n$, then n th Catalan number. To see this, observe that $\varphi(L^{\varepsilon_1} \dots L^{\varepsilon_n})$ is zero whenever there is some k such that there are more pluses than minuses in $\{\varepsilon_k, \dots, \varepsilon_n\}$. Any other ε contributes 1. The counting problem here is well known to give the Catalan number.

Example 71. Let μ be any compactly supported probability measure on \mathbb{R} . Then there is a bounded measurable function $f : \mathbb{R} \mapsto \mathbb{R}$ such that $\mu_{s.c} \circ f^{-1} = \mu$. Return to the previous example and take $T = f(L+L^*)$. By definition, this satisfies

$$m_n(T) = \langle T^n e_0, e_0 \rangle = \int f(x)^n d\mu_{s.c}(x) = \int u^n d\mu(u).$$

Thus, T has distribution μ .

Free independence: The all-important concept of independence must be now defined in the non-commutative setting. There are multiple options, of which the only one we need to consider is that of *free independence*.

Definition 72. Let (\mathcal{A}, φ) be an NCPS. Let $\mathcal{A}_i, i \in I$, be unital sub-algebras of \mathcal{A} . We say that these subalgebras are *freely independent* if for any $n \geq 1$ and any $a_1 \in \mathcal{A}_{i_1}, \dots, a_n \in \mathcal{A}_{i_n}$ with $\varphi(a_i) = 0$ for all i , we have

$$\varphi(a_1 a_2 \dots a_n) = 0 \quad \text{provided } i_1 \neq i_2 \neq i_3 \dots \neq i_n.$$

In a \star -NCPS, we shall usually require the subalgebras to be closed under involution too.

We say that a, b, c, \dots (elements of the algebra \mathcal{A}) are free if the algebras $\mathcal{A}_a, \mathcal{A}_b, \dots$ generated by a, b, \dots are free if the algebras $\mathcal{A}_{a, a^*}, \mathcal{A}_{b, b^*}, \dots$ are freely independent.

In classical probability, we could define independence by $\mathbf{E}[P(X)Q(Y)] = 0$ for any two polynomials P, Q with $\mathbf{E}[P(X)] = \mathbf{E}[Q(Y)] = 0$. However, even if we regard a classical probability space as an NCPS (as in the examples above), two classically independent random variables are not freely independent except in trivial situations. This is because if X, Y are independent zero mean variables, free independence requires $\mathbf{E}[XYXY] = 0$ which is almost never satisfied in commutative situation.

If we know the marginal distributions of two independent variables, then their joint distribution can be determined. The same holds for several variables. Here is an indication how. Assume that a, b are freely independent variables (for simplicity let us ignore involutions and work with a NCPS). Then

$$\varphi((a^{k_1} - \varphi(a^{k_1}))(b^{\ell_1} - \varphi(b^{\ell_1})) \dots (a^{k_m} - \varphi(a^{k_m}))(b^{\ell_m} - \varphi(b^{\ell_m}))) = 0,$$

for any k_j, ℓ_j s. Expand the product and observe that $\varphi(a^{k_1} b^{\ell_1} \dots a^{k_m} b^{\ell_m})$ in terms of φ applied to lower degree monomials. Inductively, it is clear that we can recover φ applied to any monomial in a and b , which is the joint distribution of these two variables.

Existence of freely independent variables: We have not yet given any example of freely independent variables. In fact, one would like to see a theorem (analogous to product measure construction in classical probability) that given any two distributions (i.e., moment sequences), there exists an algebra and variables that have these distributions and are freely independent. Let us state the result in full generality.

Result 73. Given non-commutative probability spaces $(\mathcal{A}_i, \varphi_i)$, $i \in I$, does there exist a NCPS (\mathcal{A}, φ) and subalgebras $\mathcal{B}_i \subseteq \mathcal{A}$, $i \in I$, such that (a) $(\mathcal{B}_i, \varphi|_{\mathcal{B}_i})$ is isomorphic to $(\mathcal{A}_i, \varphi_i)$ and (b) \mathcal{B}_i , $i \in I$, are freely independent. The same statement holds with \star -NCPS in place of NCPS.

We give proof in a special setting, for simplicity. We shall simply construct two freely independent variables with given distributions.

Let x, y be two symbols and let G be the collection of all expressions of the form $x^{k_1} y^{\ell_1} \dots x^{k_m} y^{\ell_m}$ for any m and any k_j, ℓ_j s (where k_1 or ℓ_m can be zero, but others are non-zero

integers). The empty expression is denoted 1. Then G is nothing but the free group³ generated by two elements x and y . Since $G_x := \{x^k : k \in \mathbb{Z}\}$ is isomorphic to \mathbb{Z} (and similarly G_y), we also say that G is the free product of \mathbb{Z} with itself and write $G = G_x \star G_y$.

Define H, H_x, H_y , as complex Hilbert spaces with orthonormal bases G, G_x and G_y , respectively. In short, $H_x = H_y = \ell^2(\mathbb{Z})$ and $H = \ell^2(\mathbb{Z} \star \mathbb{Z})$. We have the non-commutative expectations $\varphi, \varphi_x, \varphi_y$, all given by $T \mapsto \langle T1, 1 \rangle$, where 1 is the identity.

On H_x and H_y , we have the natural left-shifts L_x and L_y . For example, $L_x(x^k) = x^{k-1}$. We extend these to H by setting

$$\tilde{L}_x(x^{k_1}y^{\ell_1} \dots x^{k_m}y^{\ell_m}) = \begin{cases} x^{k_1-1}y^{\ell_1} \dots x^{k_m}y^{\ell_m} & \text{if } k_1 \neq 0, \\ y^{\ell_1} \dots y^{\ell_m} & \text{if } k_1 = 0. \end{cases}$$

Similarly, \tilde{L}_y acts by “dividing by y on the left”, if the expression starts with y , and keeps it intact if the expression starts with x . The claim is that \tilde{L}_x and \tilde{L}_y are \star -free in $(\mathcal{B}(H), \varphi)$.

More generally, to $T \in \mathcal{B}(H_x)$ associate $\tilde{T} \in \mathcal{B}(H)$ as follows: If a basis element starts with y , then \tilde{T} keeps it intact. If it starts with x , we define $\tilde{T}(x^{k_1}y^{\ell_1} \dots y^{\ell_m}) = T(x^{k_1})y^{\ell_1} \dots y^{\ell_m}$. Let \mathcal{A}_x be the image of $\mathcal{B}(H_x)$ under this map. Similarly define \mathcal{A}_y . It is clear that $(\mathcal{A}_x, \varphi|_{\mathcal{A}_x})$ is isomorphic to $(\mathcal{B}(H_x), \varphi_x)$ and similarly for y .

We claim that \mathcal{A}_x and \mathcal{A}_y are freely independent in $\mathcal{B}(H)$ with respect to φ . Since the algebra generated by $\{\tilde{L}_x, \tilde{L}_x^*\}$ is contained in \mathcal{A}_x (and similarly for y), this implies the earlier claim that \tilde{L}_x and \tilde{L}_y are \star -free.

Proof. Let $\tilde{T}_i \in \mathcal{A}_x$ and $\tilde{S}_i \in \mathcal{B}(H_y)$ with $\varphi(\tilde{T}_i) = 0$ and $\varphi(\tilde{S}_i) = 0$. Consider $\varphi(T_1 S_1 \dots T_m S_m)$.
complete this ■

Moment-cumulant calculus

Let (\mathcal{A}, φ) be an NCPS. For $a_1, \dots, a_n \in \mathcal{A}$, define

$$m_n[a_1, \dots, a_n] = \varphi(a_1, \dots, a_n).$$

We also set $m_0 = 1$. This is called the moment function. In particular, for a single variable a , we have its moments $m_n(a) := \varphi(a^n)$ and joint moments of a and a^* such as $m_5[a, a^*, a^*, a]$

³In general, given two groups G and H , we can form a group $G \star H$ by considering all finite words of the form $g_1 h_1 \dots g_k h_k$ with $g_i \in G$ and $h_i \in H$. Multiplication is just juxtaposition of words, with the obvious simplification when two elements of G (or both of H) are adjacent to each other. This free product construction is the reason behind the choice of the name free probability.

etc. This collection of numbers is what substitutes the distribution of a . There are no measures associated, except in an important special case.

Further development is analogous to the classical case, with the lattice of non-crossing partitions replacing the lattice of all partitions. For $\Pi \in NC_n$ define m_Π as before (as $NC_n \subseteq \mathcal{P}_n$ anyway) and define *free cumulants* by the expressions,

$$m_n[a_1, \dots, a_n] = \sum_{\Pi \in NC_n} \prod_{j=1}^{\ell(\Pi)} \kappa_{|\Pi_j|}[a_{\Pi_j}], \quad \kappa_\Pi[a] = \prod_{j=1}^{\ell(\Pi)} \kappa_{|\Pi_j|}[a_{\Pi_j}] \text{ for } \Pi \in NC_n.$$

Thus, the relationship between the moments and cumulants (we say free when there is need to disambiguate) is given by

$$m_\Pi[a] = \sum_{\Gamma \leq \Pi} \kappa_\Gamma[a], \quad \kappa_\Pi[a] = \sum_{\Gamma \leq \Pi} (-1)^{\ell(\Pi)-1} C_{\ell(\Pi)-1} m_\Gamma[a]$$

where the sums are over non-crossing partitions finer than Π . In the second expression, we used the formula for the Mobius function of NC_n .

Properties of free cumulants: As in the classical case, it is easy to see that $\kappa_\Pi[\bullet]$ is multilinear in its arguments. There is no symmetry, either in κ_n or m_n , since the variables are not commutative. In some of our examples where φ was a “trace”, there is circular symmetry coming from the identity $\text{tr}(AB) = \text{tr}(BA)$, but that is about it. The key analogy that we wish to carry through is the characterization of free independence in terms of free cumulants.

Lemma 74. *Let (\mathcal{A}, φ) be an NCPS. Let \mathcal{B} and \mathcal{C} be two unital subalgebras of \mathcal{A} . Then \mathcal{B} and \mathcal{C} are freely independent if and only if $\kappa[x_1, \dots, x_n] = 0$ for any $n \geq 1$ and any x_i s coming from $\mathcal{B} \cup \mathcal{C}$ and there is at least one $x_i \in \mathcal{B}$ and one $x_j \in \mathcal{C}$.*

This proof will be different from the one we gave in the classical case in that we shall do it entirely combinatorially. Further, we have not yet introduced the analogue of the moment generating function and its logarithm in the noncommutative setting. We first prove a special case, when one of the variables is a constant.

Lemma 75. *In an NCPS (\mathcal{A}, φ) , for any $n \geq 2$ and any $a_1, \dots, a_n \in \mathcal{A}$, if one of the a_i s is constant, then $\kappa_n[a_1, \dots, a_n] = 0$.*

Proof. Let $n = 2$ and consider $\kappa_2[1, a]$ (without loss of generality). By the moment cumulant relations, we have

$$m_1[a] = \kappa_1[a], \quad m_2[1, a] = \kappa_2[1, a] + \kappa_1[1]\kappa_1[a].$$

Since $m_2[1, a] = m_1[a]$ and $\kappa_1[1] = 1$, we get $\kappa_2[1, a] = 0$ as desired. Now inductively suppose that the lemma has been proved for $n \leq k - 1$.

Consider a_1, \dots, a_k with $a_\ell = 1$. Let $[a] = [a_1, \dots, a_k]$ and $\hat{a} = [a_1, \dots, a_{\ell-1}, a_{\ell+1}, \dots, a_k]$. From the moment cumulant relations,

$$m_k[a] = \sum_{\Pi \in NC_k} \kappa_\Pi[a], \quad m_{k-1}[\hat{a}] = \sum_{\Gamma \in NC_{k-1}} \kappa_\Gamma[\hat{a}].$$

Clearly $m_k[a] = m_{k-1}[\hat{a}]$. Further, $\kappa_\Pi[a] = \kappa_\Gamma[\hat{a}]$ if $\Pi = \Gamma \sqcup \{\ell\}$ (i.e., a singleton block $\{\ell\}$ appended to Γ). Thus, subtracting the above identities, we see that the sum of $\kappa_\Pi[a]$ over all Π in which ℓ is not a singleton, must be zero. Of these, all terms in which ℓ is in a block of size less than k vanish, by the induction hypothesis. That leaves only $\Pi = [k]$, showing that $\kappa_k[a] = 0$. This completes the induction. \blacksquare

Lemma 76. *Let a_1, \dots, a_n , where $n \geq 2$, be elements in a NCPS (\mathcal{A}, φ) . Then for any $1 \leq k \leq n - 1$, we have*

$$\kappa_n[a_1, \dots, a_n] = \kappa_{n-1}[a_1, \dots, a_{k-1}, a_k a_{k+1}, a_{k+2}, \dots, a_n] - \sum_{\substack{\Pi \in NC(n), \ell(\Pi)=2 \\ k \in \Pi_1, k+1 \in \Pi_2}} \kappa_\Pi[a_1, \dots, a_n].$$

Proof. For $n = 2$, check that this is true. For general n , we write the moment-cumulant relations as

$$m_n[a_1, \dots, a_n] = \sum_{\Pi \in NC_n} \kappa_\Pi[a_1, \dots, a_n] = \kappa_n[a] + \sum_{\Pi: \ell(\Pi) \geq 2} \kappa_\Pi[a], \quad (2)$$

$$\begin{aligned} m_{n-1}[a_1, \dots, a_{k-1}, a_k a_{k+1}, a_{k+2}, \dots, a_n] &= \sum_{\Gamma \in NC_{n-1}} \kappa_\Gamma[a_1, \dots, a_{k-1}, a_k a_{k+1}, a_{k+2}, \dots, a_n] \quad (3) \\ &= \kappa_{n-1}[a_1, \dots, a_{k-1}, a_k a_{k+1}, a_{k+2}, \dots, a_n] + \sum_{\Gamma: \ell(\Gamma) \geq 2} \kappa_\Gamma[a_1, \dots, a_{k-1}, a_k a_{k+1}, a_{k+2}, \dots, a_n] \end{aligned}$$

The left hand sides are of course equal, both being $\varphi(a_1 \dots a_n)$. Thus, the right hand sides are equal. We cancel common terms. Indeed, if Γ has at least two blocks, then inductively we may write for

$$\kappa_\Gamma[a_1, \dots, a_{k-1}, a_k a_{k+1}, a_{k+2}, \dots, a_n] = \kappa_{\Gamma'}[a_1, \dots, a_n] + \sum_{\Gamma''} \kappa_{\Gamma''}[a_1, \dots, a_n]$$

where, if Γ_j is the block of Γ that contains⁴ $a_k a_{k+1}$, then

⁴Here it would be better to think of Γ as a partition of the set of variables $a_1, \dots, a_{k-1}, a_k a_{k+1}, a_{k+2}, \dots, a_n$ rather than $1, 2, \dots, n - 1$. The latter makes the phrasing more complicated, as when we separate the two terms then indices will have to be shifted.

- $\Gamma' \in NC_n$ is got from Γ by replacing the block Γ_j by a block of size $|\Gamma_j| + 1$ in which a_k and a_{k+1} are distinct elements (and the rest are untouched),
- Γ'' varies over all partitions in NC_{n-1} got by breaking the j th block of Γ' (constructed in the previous step) into exactly two blocks, one of which contains a_k and the other contains a_{k+1} .

What all partitions of NC_n occur in this way? We claim that every $\Pi \in NC_n$ occurs exactly once, except for those Π with exactly two blocks, one of which contains a_k and the other contains a_{k+1} . Indeed, if Π contains a_k, a_{k+1} in the same block, then it occurs as Γ' for the Γ got by amalgamating a_k and a_{k+1} as one element $a_k a_{k+1}$. On the other hand, if a_k and a_{k+1} occur in different blocks of Π , then it is got as a Γ'' where $\Gamma \in NC_{n-1}$ is got by merging those two blocks of Π and then amalgamating a_k and a_{k+1} . The only issue is if the resulting Γ has only one block.

Thus, subtracting the two equations in (2), we get the conclusion of the lemma. ■

Proof of Lemma 74. By Lemma 75 and multilinearity of the free cumulant, for any $n \geq 2$ and any x_i 's, $\kappa_n[x_1, \dots, x_n] = \kappa_n[y_1, \dots, y_n]$ where $y_i = x_i - \phi(x_i)\mathbf{1}$. Thus, without loss of generality, we may assume that x_i are centered.

Again inductively assume that the lemma is proved for κ_m , $m \leq n - 1$ (we leave the base case as an exercise). Now consider the case of κ_n and let S and T be the subsets of indices for which x_i belongs to \mathcal{B} and \mathcal{C} respectively. Then,

$$m_n[x] = \sum_{\Pi \in NC_n} \kappa_{\Pi}[x].$$

Since κ_{Π} factors over the blocks of Π , if there is any of these blocks has size less than n and intersects both S and T , then $\kappa_{\Pi}[x] = 0$. That leaves only the full partition $\{[n]\}$ and all those partitions in which S and T are unions of blocks. ■

Free convolution

Let a, b be self-adjoint elements of a \star -NCPS (\mathcal{A}, ϕ) having distributions μ and ν . Assume that these measures are compactly supported. If a and b are freely independent, then the moments of $a + b$ can be calculated in terms of the moments of a and the moments of b . Hence, the distribution θ of $a + b$ is determined by μ and ν . This gives a binary operation called *free convolution* and we write $\theta = \mu \boxplus \nu$.

In principle, the definition gives us an algorithm to calculate the free convolution of two probability measures. However, the formulas for the moments of $a + b$ have not been given explicitly in terms of the individual moments of a and b . Indeed, such formulas would be horribly complicated. It is to simplify this that free cumulants were introduced. From Lemma 74, we see that

$$\begin{aligned}\kappa_n(a + b) &= \kappa_n[a + b, \dots, a + b] \\ &= \kappa_n[a, \dots, a] + \kappa_n[b, \dots, b] = \kappa_n(a) + \kappa_n(b).\end{aligned}$$

Thus, the free cumulants of θ are expressed in a very simple way in terms of the free cumulants of μ and ν .

Example 77. Let a, b be freely independent semi-circular elements. Then $\kappa_n(a + b) = 2\delta_{n,2}$. Since $\kappa_n(cx) = c^n \kappa_n(x)$ (by multi-linearity), we see that $\kappa_n(\frac{a+b}{\sqrt{2}}) = \delta_{n,2}$. This means that $(a + b)/\sqrt{2}$ also has the semi-circle distribution!

This is the analogue of the situation in classical probability where $(X + Y)/\sqrt{2} \sim N(0, 1)$ if X, Y are i.i.d. $N(0, 1)$. Thus, the semi circle law takes the central place in free probability, like the Gaussian distribution in classical probability.

Integral transforms

We want to find the analogues in free probability of the moment generating function and its logarithm in classical probability. We shall restrict ourselves to the univariate setting.

For a probability measure $\mu \in \mathcal{P}(\mathbb{R})$, recall that its Stieltjes' transform $G_\mu(z) = \int \frac{1}{z-x} d\mu(x)$. If μ is compactly supported, say $\text{spt}(\mu) \subseteq [-R, R]$, then for $|z| > R$, we have the series expansion near infinity,

$$G_\mu(z) = \sum_{n=0}^{\infty} \frac{m_n(\mu)}{z^{n+1}}$$

where $m_n(\mu)$ is the n th moment of μ . We now define the R -transform of μ as

$$R_\mu(w) = \sum_{n=1}^{\infty} \kappa_n(\mu) w^{n-1}$$

where $\kappa_n(\mu)$ is the n th free cumulant of μ . Also define the K -transform of μ as $K_\mu(w) = \frac{1}{w} + R_\mu(w)$. Do these series converge? Clearly $|m_n(\mu)| \leq R^n$. Further, from the moment cumulant relations, $|\kappa_n(\mu)| \leq$ (Show exponential bound for κ_n)

Thus, G_μ encodes the moments and R_μ (or K_μ) encode the free cumulants. How does the moment-cumulant relation carry over to the level of these transforms?⁵

Theorem 78. G_μ and K_μ are inverses of each other.

More precisely, the series for G_μ converges in a neighbourhood of ∞ in $\mathbb{C} \cup \{\infty\}$ and the series for K_μ converges in a neighbourhood of 0. Further, G_μ maps a neighbourhood of ∞ to a neighbourhood of 0 and there the inverse relationship holds.

Proof. We write the moment-cumulant relationship in a less explicit but more convenient form. Start with

$$m_n = \sum_{\Pi \in NC_n} \prod_{j=1}^{\ell(\Pi)} \kappa_{|\Pi_j|}.$$

Fix the block containing 1 to be $V = \{1, i_1 + 2, i_2 + 3, \dots, i_\ell + \ell + 1\}$ where $i_j + j$ are in increasing order, so that $[n] \setminus V$ has segments of lengths i_1, \dots, i_ℓ . Then, the restriction of Π to each of the segments $\{2, \dots, i_1 + 1\}, \dots, \{i_\ell + 2, \dots, n\}$ is a non-crossing partition in itself. Thus, when we sum over all Π with the first block equal to V , we get

$$\kappa_\ell \sum_{\Pi: \Pi_1=V} \prod_{j=2}^{\ell(\Pi)} \kappa_{|\Pi_j|} = \kappa_\ell m_{i_1} \dots m_{i_\ell}.$$

Thus, we arrive at

$$m_n = \sum_{\substack{\ell \geq 1, i_1, \dots, i_\ell \geq 0, \\ i_1 + \dots + i_\ell = n - \ell}} \kappa_\ell m_{i_1} \dots m_{i_\ell}.$$

Now multiply by z^{-n-1} and sum over $n \geq 0$ to get

$$G(z) = z^{-1} + \sum_{\ell \geq 1} \kappa_\ell z^{-1} \prod_{j=1}^{\ell} \left(\sum_{i_j \geq 0} m_{i_j} z^{-i_j-1} \right) = \sum_{\ell \geq 1} \kappa_\ell G(z)^\ell = \frac{1}{z} + \frac{G(z)}{z} R(G(z)).$$

⁵Lagrange's inversion formula: Let $f(w) = w/\varphi(w)$ where φ is analytic near zero with $\varphi(0) = 1$. Let g be the functional inverse of f in a neighbourhood of the origin (exists because $f(w) \sim w$ near $w = 0$). Then, $[z^n]g(z) = \frac{1}{n+1} [w^n] \varphi(w)^n$. In our case, $f(w) = \frac{1}{K(w)} = \frac{w}{1+wR(w)}$ and $g(z) = G(1/z) = \sum_{n=0}^{\infty} m_n z^{n+1}$. The inversion formula gives the relationship between the coefficients of G and coefficients of R . This formula is $m_n = \frac{1}{n+1} \sum \kappa_{\ell_1} \dots \kappa_{\ell_{n+1}}$ where the sum is over $\ell_j \geq 0$ such that $\ell_1 + \dots + \ell_{n+1} = n$, with the convention that $\kappa_0 = 0$. **Two things to do:** (1) From here, go to the formula in terms of NC_n . (2) Rewrite the chapter just starting from the Lagrange inversion, which naturally leads to non-crossing matchings and the moment-freecumulant relations, in the same way that the relationship between coefficients of an analytic function and its logarithm can be made the starting point of the moment-cumulant relations in the classical setting. Is there a way to also make the link to random matrix addition directly from here?

This is the same as $z = \frac{1}{G(z)} + R(G(z)) = K(G(z))$. Thus, K is the inverse of G . ■

It may be noted that the R -transform (we have defined it only for compactly supported measures) also determines the measure, since it determines the Stieltjes' transform in a neighbourhood of infinity.

Corollary 79. $R_{\mu \boxplus \nu} = R_\mu + R_\nu$.

Proof. Immediate consequence of additivity of free cumulants under free convolution. ■

We illustrate the usefulness by computing the free convolution of a Bernoulli measure with itself.

Example 80. Let $\mu = \frac{1}{2}(\delta_1 + \delta_{-1})$. Then $G(z) = \frac{z}{z^2-1}$. Solve $G(z) = w$ to get $z = \frac{1 \pm \sqrt{1+4w^2}}{2w}$. Choosing the right branch, we get $R(w) = \frac{-1 + \sqrt{1+4w^2}}{2w}$. If $\theta = \mu \boxplus \mu$, then $R_\theta(w) = 2R_\mu(w) = \frac{-1 + \sqrt{1+4w^2}}{w}$, hence $K_\theta(w) = \frac{\sqrt{1+4w^2}}{w}$. Solving for the inverse, we get $G_\theta(z) = \frac{1}{\sqrt{z^2-4}}$. Recognizing this as the free convolution of the arcsine law, we see that θ is arcsine on $[-2, 2]$.

As another example, let us re-derive the free convolution of semi-circle measure with itself.

Example 81. If μ is the semi-circle measure on $[-2, 2]$, then we know that $G_\mu(z) = \frac{z - \sqrt{z^2-4}}{2}$. More usefully, $G_\mu(z)$ satisfies the quadratic equation $G_\mu(z)^2 - zG_\mu(z) + 1 = 0$. Hence, $z = G_\mu(z) + \frac{1}{G_\mu(z)}$ showing that $K_\mu(w) = w + \frac{1}{w}$ and $R_\mu(w) = w$. This is an alternate way to derive that the free cumulants are all zero except for κ_2 which is 1. Hence $R_{\mu \boxplus \mu}(w) = 2w$ which is also the R -transform of semi-circle on $[-2\sqrt{2}, 2\sqrt{2}]$ (work out the relationship between R -transforms of a probability measure ν and its scaling $\nu_t(A) := \nu(tA)$).

Free central limit theorem

We have said before that the semicircle plays a role in free probability very analogous to the Gaussian in classical probability. Now we prove a free version of the central limit theorem. Suppose a_k are freely independent and identically distributed elements in an algebra \mathcal{A} . Does $(a_1 + \dots + a_n)/\sqrt{n}$ converge in distribution to some variable? Firstly note that $\kappa_2[a_1 + \dots + a_n] = n\kappa_2[a_1]$ and hence \sqrt{n} is the right scaling factor. Secondly, if we assume that $(a_1 + \dots + a_n)/\sqrt{n}$ does converge in distribution to some variable a , then for two freely independent copies a, b of this variable $a + b$ must have the same distribution

as $\sqrt{2}a$. Just as we saw earlier for classical random variables, this forces the free cumulants to satisfy the relationship $2^{\frac{p}{2}}\kappa_p[a] = 2\kappa_p[a]$ which implies $\kappa_p[a] = 0$ for $p \neq 2$ which implies that a is a semicircular variable. Now we actually prove that the convergence does happen.

Theorem 82. *Let a, a_k be freely independent, identically distributed self-adjoint variables in a \star -NCPS (\mathcal{A}, φ) . Assume that the distribution of a is non-degenerate. Then,*

$$\frac{a_1 + \dots + a_n - n\kappa_1(a)}{\sqrt{n}\sqrt{\kappa_2(a)}} \xrightarrow{d} \mu_{s.c.},$$

the standard semicircle law supported on $[-2, 2]$.

Proof. Without loss of generality assume that $\kappa_1(a) = 0$ and $\kappa_2(a) = 1$. The proof is word for word the same as we gave for classical CLT using cumulants (wisely we did not even change the notation for cumulants!). We conclude that $\kappa_p(S_n/\sqrt{n}) \rightarrow \delta_{p,2}$. The only non-commutative variable whose free cumulants are $\delta_{p,2}$ is the standard semicircle law. Hence the conclusion. ■

Random matrices and freeness

We have now seen Voiculescu's world of free probability with objects and theorems analogous to those in classical probability theory (we saw only a tiny sample of this. There is a free version of nearly everything, free Poisson, free Brownian motion, free Lévy process, free entropy, ... even free graduate students).

Apart from analogy, there is connection between the classical and free worlds, and that is provided by random matrix theory. Indeed, one of our motivations for introducing free probability theory is to explain the occurrence of semicircle law and other limit laws in random matrices, from a more conceptual algebraic framework. The essential connection is in the following theorem (and other such statements asserting free independence of classically independent large random matrices).

Theorem 83. *Consider $M_n(\mathbb{C}) \otimes L^\infty(\mathbf{P})$, the algebra of $n \times n$ random complex matrices with the state $\varphi(A) = n^{-1}\mathbf{E}[\text{tr}(A)]$. Let $X_n = (X_{i,j})$ and $Y_n = (Y_{i,j})_{i,j \leq n}$ be random Hermitian matrices on a common probability space taking values in $M_n(\mathbb{C})$. We consider two scenarios.*

1. X_n and Y_n are Wigner matrices with $X_{1,1}$ and $X_{1,2}$ having exponential tails.

2. $X_n = A_n$ and $Y_n = U_n B_n U_n^*$ where A_n, B_n are real diagonal matrices and U_n is a Haar distributed unitary matrix. We assume that the ESD of A_n and B_n are tight????

In either of these two situations, X_n and Y_n are asymptotically freely independent.

In particular, if X_n and Y_n have limiting spectral distributions μ and ν respectively, then $X_n + Y_n$ has limiting spectral distribution $\mu \boxplus \nu$.

Now suppose X_n and Y_n are independent copies of GOE matrix. By properties of normals, $X_n + Y_n$ has the same distribution as $\sqrt{2}X_n$.

Spectrum of the sum of two matrices and free convolution

Let a, b be two self-adjoint, freely independent variables in a non-commutative probability space (\mathcal{A}, φ) . Then, $\kappa_n(a + b) = \kappa_n(a) + \kappa_n(b)$. Hence the distribution of a and b determine the distribution of $a + b$. The procedure to find the distribution of $a + b$ is as follows.

1. Let μ and ν be the distributions of a and b respectively. This means $\varphi(a^n) = \int x^n \mu(dx)$ and $\varphi(b^n) = \int x^n \nu(dx)$ for all n .
2. From the moments $m_n(a) := \varphi(a^n)$ and $m_n(b) = \varphi(b^n)$ find the free cumulants $\kappa_n[a]$ and $\kappa_n[b]$. This can be done using the relations (??).
3. Find $\kappa_n := \kappa_n[a] + \kappa_n[b]$ and insert into formulas (??) to find m_n .
4. Find the measure θ whose moments are m_n . Then θ is the distribution of $a + b$.

An analogous procedure can be described in classical probability, to find the sum of two independent random variables using their cumulants. But there are also other useful techniques for dealing with sums of random variables such as the characteristic function (which is multiplicative under independence) or the logarithm of the characteristic function (which is additive). There are also such analytic objects associated to non-commutative random variables, which we describe now.

Let μ be a compactly supported on \mathbb{R} with Stieltjes' transform $G_\mu(z) = \int (z - x)^{-1} \mu(dx)$ for the Stieltjes' transform of μ . From properties of Stieltjes transforms, we know that knowing G_μ in a neighbourhood of ∞ one can recover all the moments of μ and hence recover μ itself. Further, G_μ is one-one in a neighbourhood of ∞ and has an analytic inverse K_μ defined in a neighbourhood of 0. Since $G_\mu(z) = z^{-1} + m_1 z^{-2} + \dots$ (where m_k are the moments of μ) for z close to ∞ , we see that $K_\mu(w) = w^{-1} + R_\mu(w)$ for some analytic function R (defined in a neighbourhood of 0). R_μ is called the R -transform of μ .

Lemma 84. $R_\mu(w) = \sum_{n=1}^{\infty} \kappa_n^\mu w^{n-1}$, where κ_n^μ are the free cumulants of μ .

Proof. Let $S(w) = \sum_{n=1}^{\infty} \kappa_n^\mu w^{n-1}$. We show that $G(w^{-1} + S(w)) = w$ for w close to 0 and this clearly implies that $S = R_\mu$. ■

Exercises

Exercise 85. If P, Q are finite posets, then $P \times Q$ is a poset with the order $(a_1, a_1) \leq (b_1, b_2)$ if $a_1 \leq a_2$ and $b_1 \leq b_2$. Show that $\mu_{P \times Q}((a_1, a_2), (b_1, b_2)) = \mu_P(a_1, b_1)\mu_Q(a_2, b_2)$. Use this to deduce that for the Boolean poset $\mu(A, B) = (-1)^{|B \setminus A|}$. [Hint: Write the Boolean poset as a product of n posets].

Exercise 86. Give an alternate derivation of the Mobius function for \mathcal{P}_n following these steps.

1. Fix $x \in \mathbb{N}$.

(a) Let $f_x(\Pi)$ be the number of $\Sigma \in \mathcal{P}_n$ such that $\Pi \leq \Sigma$ and $\ell(\Sigma) \leq x$.

(b) Let $g_x(\Pi)$ be the number of $\Sigma \in \mathcal{P}_n$ such that $\Pi \leq \Sigma$ and $\ell(\Sigma) = x$.

Argue that $g_x(\Pi) = x(x-1)\dots(x-\ell(\Pi)+1)$ and $f_x(\Pi) = x^{\ell(\Pi)}$.

2. Prove that $x(x-1)\dots(x-n+1) = \sum_{\Sigma \in \mathcal{P}_n} \mu(0, \Sigma)x^{\ell(\Sigma)}$ where 0 is the partition of $[n]$ into singletons.

3. Equate powers of x to deduce that $\mu(0, 1) = (-1)^{n-1}(n-1)!$, where 1 is the single-block partition.

4. Deduce the general formula for $\mu(\Sigma, \Pi)$.

Exercise 87. Show that $|NC_n| = C_n$ as by arguing that

$$|NC_n| = \sum_{i=2}^n |NC_{i-1}| \times |NC_{n-i}|.$$

[Hint: Consider the largest element in the block containing 1.]

Exercise 88. Let X_1, X_2, \dots be i.i.d. random variables with a finite moment generating function. Use cumulants to show that $\frac{1}{\sqrt{n}}(X_1 X_2 + \dots + X_n X_{n+1})$ converges in distribution to a Gaussian.

Chapter 10

Non-asymptotic questions

So far we have only looked at asymptotic questions about random matrices, as the dimension goes to infinity. It is always an important question as to how well an asymptotic result can be used for finite size. For example, is semicircle distribution a reasonable for the empirical distribution of a 50×50 matrix? If not, 100×100 ? Rates of convergence are one kind of answer. Another kind is to look for bounds that are valid for finite dimensions, for instance questions of concentration of measure.

To illustrate with examples from basic probability, the strong law of large numbers and Cramer's theorem of large deviations are asymptotic statements. Chebyshev's inequality and Hoeffding's inequality are examples of non-asymptotic statements. In this course we have seen asymptotic theorems like Wigner's semi-circle law. A non-asymptotic statement that we have seen is Theorem 15, where we got explicit probability bounds for the event that the maximal eigenvalue of the quadratic beta gas is more than $2 + \varepsilon$.

In this chapter, we look at bounds for the extreme singular values of random matrices with independent entries. More precisely, we want upper bounds for the largest singular value and lower bounds for the smallest singular value. These bounds will be quite different for rectangular matrices as compared to square matrices. We discuss both, first in the context of Gaussian matrices and then for general distributions.

Throughout the chapter, A will be an $m \times n$ matrix with i.i.d. entries and $m \geq n$. The case of square matrices, when $m = n$, is usually harder. When m is much larger than n , the matrix will be called *tall*. We shall place restrictions on the distribution of entries where necessary, either for the validity of results or for simplicity of exposition. We shall write the singular values as $0 \leq s_1(A) \leq \dots \leq s_n(A)$. Recall that s_i^2 are the eigenvalues of $A^t A$.

We also write s_{\min} and s_{\max} for s_1 and s_n , respectively¹. We shall often use the variational formulas

$$s_n(A) = \max_{u \in \mathbb{R}^n, \|u\|=1} \|Au\| = \max_{\substack{u \in \mathbb{R}^n, v \in \mathbb{R}^m \\ \|u\|=1=\|v\|}} v^T Au, \quad (1)$$

$$s_1(A) = \min_{u \in \mathbb{R}^n, \|u\|=1} \|Au\| = \min_{u \in \mathbb{R}^n, \|u\|=1} \max_{v \in \mathbb{R}^m, \|v\|=1} v^T Au. \quad (2)$$

Gaussian matrices

As in many other contexts, it is easier to deal with matrices with Gaussian entries. Let $a_{i,j}$ be i.i.d. standard Gaussian random variables. The key property that simplifies proofs in this situation is the orthogonal invariance, $PAQ \stackrel{d}{=} A$ for any $P \in O(m)$, $Q \in O(n)$.

Many questions about singular values can be answered by techniques we have seen before:

1. It is possible to reduce A to a bidiagonal matrix with independent χ -entries with various degrees of freedom (see Exercise 40). This makes $A^T A$ a Jacobi matrix, and bounds for its eigenvalues can be got by using Gershgorin-type theorems, just as we showed the bounds for extreme eigenvalues in the quadratic β -ensembles.
2. It is possible to get the exact distribution of the singular values. From (14) which gives the density of s_i^2 , we can deduce the joint density of s_i s to be

$$\prod_{j < k} |s_j^2 - s_k^2|^2 \prod_k e^{-s_k^2} s_k^{m-n-1}.$$

3. One can use method of moments to get bounds for $s_n(A)$. Indeed, $\text{tr}[(A^T A)^p] \geq s_n^{2p}$ for any $p \geq 1$. When p is large. For a fixed p , the left hand side (its expectation, for example) can be handled by the method of moments, but the bound obtained on $s_n(A)$ or its expectation is loose. To get better bounds, we must use p growing appropriately with n , which makes the execution of the method of moments harder but not impossible. The reason we get better bounds is simply that $(\text{tr}[(A_n^T A_n)^p])^{1/2p} \rightarrow s_n(A)$, if n is fixed and $p \rightarrow \infty$.

¹Almost everything here is taken from various superb expositions by Rudelson and Vershynin, who are also the discoverers of many of the results in the subject.

The first two methods do not apply when we go beyond Gaussian matrices. Further, occasionally we shall consider matrices of the form $M + A$, where M is a deterministic matrix, and in that case all these techniques are harder. Hence we shall use other more geometric techniques. Some of these are fundamental techniques in studying Gaussian processes in general.

Rectangular Gaussian matrices

Theorem 89. *Let $A_{m \times n}$ have i.i.d. standard Gaussian entries.*

1. For every $m \geq n$, we have $\mathbf{E}[s_{\max}(A)] \leq \sqrt{m} + \sqrt{n}$ and $\mathbf{E}[s_{\min}(A)] \geq \sqrt{m} - \sqrt{n}$.
2. For any $t > 0$, we have

$$\mathbf{P}\{s_{\max} \geq \sqrt{m} + \sqrt{n} + t\} \leq 7e^{-t^2/8}, \quad \text{and} \quad \mathbf{P}\{s_{\min} \leq \sqrt{m} - \sqrt{n} - t\} \leq 7e^{-t^2/8}.$$

Observe that the lower bound for s_{\min} is vacuous when $m = n$. But if $m = n + 1$ (or $m = n + k$ for a fixed k), then the lower bound is of the order of $1/\sqrt{n}$, which is actually the right bound even for square matrices. We shall see that later. From the above theorem, we can deduce that very tall Gaussian matrices are nearly isometries.

Corollary 90. *Let $A_{m \times n}$ be a matrix with i.i.d. standard Gaussian entries and let $B = \frac{1}{\sqrt{m}}A$. Fix $\delta < 1$. If $n \leq m\delta^2/4$, then for sufficiently large m , with probability at least $1 - e^{-m\delta^2/32}$, we have*

$$1 - \delta \leq \|Bu\| \leq 1 + \delta \quad \text{for all unit vectors } u \in \mathbb{R}^n.$$

Why should a tall matrix be almost an isometry? By the law of large numbers, we see that

$$\frac{1}{m} \sum_{i=1}^m a_{i,1}^2 \approx 1, \quad \frac{1}{m} \sum_{i=1}^m a_{i,1}a_{i,2} \approx 0.$$

Thus, the columns of B are nearly orthonormal in \mathbb{R}^m . While this is true for every pair, the approximations are about $1/\sqrt{m}$, and when there are many columns, it is not clear how the errors mount. The corollary above asserts that for tall enough matrices, we do get approximate isometry property. Still, the heuristic here is the best “explanation” I know. The proof below does not really illuminate why the bounds should hold.

To prove the theorem, we need some important facts about Gaussian random variables, given in Appendix 10. The key theorems are Gaussian concentration inequality (Theorem 137) and two comparison inequalities (Sudakov-Fernique inequality of Theorem 129 and Gordon’s inequality of Theorem 132).

Second part of Theorem 89 assuming the first. The key point is that $s_1(A)$ and $s_n(A)$ are Lip(1) functions from \mathbb{R}^{mn} to \mathbb{R} . This follows from the variational formulas (1) and (2), since $A \mapsto v^t A u$ is a Lip(1) function for every u, v , and maxima and minima of Lip(1) functions are necessarily Lip(1). Hence the inequalities in the second part of Theorem 89 follow from the Gaussian isoperimetric inequality and the bounds for the expectations in the first part. ■

Next we prove the upper bound for the expectations of the largest singular value using Sudakov-Fernique inequality.

Proof of the upper bound for expectation of s_{\max} . Let $X(u, v) = v^t A u$ for $(u, v) \in I = S^{n-1} \times S^{m-1}$. Then X is a Gaussian process and

$$\begin{aligned} \mathbf{E}[|X(u, v) - X(u', v')|^2] &= \mathbf{E} \left[\left(\sum_{i,j} a_{i,j} (v_i u_j - v'_i u'_j) \right)^2 \right] = \sum_{i \leq m, j \leq n} (v_i u_j - v'_i u'_j)^2 \\ &= 2 - 2\langle u, u' \rangle \langle v, v' \rangle. \end{aligned}$$

From (1), we know that $s_n(A) = X^*$. We compare X to a simpler Gaussian process Y on I defined by $Y(u, v) = v^t \xi + u^t \eta$ where $\xi \sim N_m(0, I_m)$ and $\eta \sim N_n(0, I_n)$. Then,

$$\begin{aligned} \mathbf{E}[|Y(u, v) - Y(u', v')|^2] &= \mathbf{E}[(v - v')^t \xi + (u - u')^t \eta]^2 = \|u - u'\|^2 + \|v - v'\|^2 \\ &= 4 - 2\langle u, u' \rangle - 2\langle v, v' \rangle. \end{aligned}$$

Thus,

$$\mathbf{E}[|Y(u, v) - Y(u', v')|^2] - \mathbf{E}[|X(u, v) - X(u', v')|^2] = 2(1 - \langle u, u' \rangle)(1 - \langle v, v' \rangle)$$

which is non-negative. By the Sudakov-Fernique inequality (Theorem 129), and we get $\mathbf{E}[s_n(A)] \leq \mathbf{E}[Y^*]$. But $Y^* = \|\xi\| + \|\eta\|$ (attained when $v = \xi/\|\xi\|$ and $u = \eta/\|\eta\|$). Further, $\mathbf{E}[\|\xi\|] \leq \mathbf{E}[\|\xi\|^2]^{1/2} = \sqrt{m}$ and similarly $\mathbf{E}[\|\eta\|] \leq \sqrt{n}$. Thus, $\mathbf{E}[s_n(A)] \leq \sqrt{m} + \sqrt{n}$. ■

To get the lower bound for $\mathbf{E}[s_1(A)]$, we use Gordon's inequality (Exercise 132)

Proof of the lower bound for expectation of s_{\min} . Define X and Y as in the previous proof. We have already seen that

$$\mathbf{E}[|X(u, v) - X(u', v')|^2] \leq \mathbf{E}[|Y(u, v) - Y(u', v')|^2]$$

for all $(u, v) \in I$. Further, if $u = u'$, then equality holds as both sides are equal to $2 - 2\langle u, u' \rangle$. Hence, Gordon's inequality applies and we get $\mathbf{E}[\min_u \max_v X(u, v)] \geq \mathbf{E}[\min_u \max_v Y(u, v)]$. By (2), $\min_u \max_v X(u, v) = s_1(A)$ while $\min_u \max_v Y(u, v) = \|\xi\| - \|\eta\|$. Hence $\mathbf{E}[s_1(A)] \geq \mathbf{E}[\|\xi\|] - \mathbf{E}[\|\eta\|]$. Since $\|\xi\|^2 \sim \chi_m^2$,

$$\mathbf{E}[\|\xi\|] = \frac{1}{2^{m/2}\Gamma(m/2)} \int_0^\infty \sqrt{x} e^{-x} x^{m/2-1} dx = \frac{\sqrt{2}\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m}{2})}$$

and similarly $\mathbf{E}[\|\eta\|] = \frac{\sqrt{2}\Gamma(\frac{n+1}{2})}{\Gamma(\frac{n}{2})}$. Thus the theorem is proved if we show that $\mathbf{E}[\|\xi\|] - \mathbf{E}[\|\eta\|] \geq \sqrt{m} - \sqrt{n}$. This follows if we prove that $v \rightarrow \frac{\sqrt{2}\Gamma(\frac{v+1}{2})}{\Gamma(\frac{v}{2})} - \sqrt{v}$ is an increasing function. **Finish this** ■

Square Gaussian matrix

As already remarked, we don't have a lower bound for the smallest singular value for a square Gaussian matrix. The following theorem shows that $\sqrt{n}/s_1(A)$ is *tight*, and in fact gives an upper bound of $1/x$ for the tail probability $P(\sqrt{n}/s_1 > x)$.

Theorem 91. *Let $A_{n \times n}$ have i.i.d. standard Gaussian entries. Then $\mathbf{P}\{s_1(A) \leq \frac{\varepsilon}{\sqrt{n}}\} \leq \varepsilon$ for any $\varepsilon > 0$. In fact, for any deterministic matrix $M_{n \times n}$, we have the same bound for $s_1(A + M)$.*

The reason for adding M is to make the point that even if we start with a matrix with very small singular values (even zeros), a Gaussian perturbation of it will have singular values not less than $1/\sqrt{n}$ (in order). Perhaps it will be better to write down the statement for $M + \sigma A$, where σ is a small number (to make it feel more like a perturbation of M)².

²This is the theme of *smoothed analysis*, introduced by Spielman and Teng. In studying performance of algorithms two common criteria are to consider the worst case input or a completely random input. Smoothed analysis is between the two, and considers an arbitrary input and perturbs it slightly, and studies the worst performance over all choices of that input.

For example, when studying an algorithm for solving linear equations $Ax = b$, it is known that how badly it performs depends on the magnitude of *condition number* of the input matrix, $\kappa(A) := s_n(A)/s_1(A)$. There are ill-conditioned matrices, for example singular matrices, hence the worst case performance is infinitely bad! What the theorem here says is that for any M , the smallest singular value of $M + \sigma A$ is at least of order $\sigma n^{-1/2}$, which is surprisingly good. One also needs a bound on the largest singular value, but that is easier.

Perturbation here makes sense, because there are necessarily numerical approximations due to rounding off of the entries of a matrix. However, that perturbation should not be modeled by a Gaussian, but perhaps a discrete distribution. The corresponding theorem is harder and will come later in the chapter.

It is possible to prove this theorem by the exact density of singular values or the Jacobi random matrix with the same singular value distribution. But such proofs work only when $M = 0$. We shall instead give two proofs, first of a weaker statement, and then the actual theorem. The proof of the weaker statement is given to illustrate an issue that will recur later, the loss that accumulates due to naive union bounds.

Proof of a weaker version of Theorem 91. Let $A = [u_1 \dots u_n]$ and $M = [m_1 \dots m_n]$ where u_i and m_i are the columns. If $s_1(A) \leq t$, then there is a unit vector $w \in \mathbb{R}^n$ such that $\|(A + M)w\| \leq t$. There will be some co-ordinate, say k , such that $|w_k| \geq \frac{1}{\sqrt{n}}$. Then, writing $(A + M)w = \sum_{j=1}^n w_j(u_j + m_j)$, and dividing out by w_k , we get

$$\|u_k + m_k + \sum_{j \neq k} \frac{w_j}{w_k}(u_j + m_j)\| \leq t\sqrt{n}.$$

Condition on $u_j + w_j$, $j \neq k$, and let P be the projection onto the orthogonal complement of the span of these vectors (with probability 1 this is a rank one projection). Then $X = P(u_k + m_k)$ is a one-dimensional Gaussian with some mean and unit variance. Hence, the probability that $|X| < t\sqrt{n}$ is at most $\frac{\sqrt{2nt}}{\sqrt{\pi}}$.

Taking into account that this event must happen for some k , the probability that $s_1(A) \leq t\sqrt{n}$ is at most $\sqrt{\frac{2}{\pi}}tn^{\frac{3}{2}}$. Taking $t = \varepsilon/\sqrt{n}$, we get the bound $n\varepsilon$ which is weaker than the claim. ■

Even though not optimal, this already shows that the smallest singular value is of order at least $n^{-3/2}$, a non-trivial statement. We now strengthen this. We shall use the following important property of a multivariate standard Gaussian vector $X \sim N_n(0, I)$ coming from orthogonal invariance. For any k -dimensional subspace W of \mathbb{R}^n , the projection of X onto W is a standard k -dimensional Gaussian (inside W), and in particular the length of the projection is a χ_k random variable.

Proof of Theorem 91. Let $B = M + A$.

Claim: For $w \in S^{n-1}$ fixed, $\mathbf{P}\{\|B^{-1}w\| > t\} \leq 1/t$.

Proof of the claim: Replacing B with PBQ for some $P, Q \in O(n)$, we may assume that $w = e_1$. This is because PQM has the same distribution as $A + M$ for a different M . Now, $\langle B^{-1}e_1, B^t e_j \rangle = \delta_{j,1}$, hence 1. $B^{-1}e_1$ is orthogonal to u_2, \dots, u_n and 2. $\|B^{-1}e_1\|$ is the reciprocal of the length of the projection of the u_1 onto the orthogonal complement of

$\text{span}\{u_2, \dots, u_n\}$. Condition on u_2, \dots, u_n and choose a vector v normal to all these columns. Then, $\|B^{-1}e_1\| = |\langle u, v \rangle|$. Since $\langle u, v \rangle$ is $N(\mu, 1)$ for some μ (which is fixed by M and v), it follows that $\mathbf{P}\{|\langle u, v \rangle| < 1/t \mid u_2, \dots, u_n\} \leq 1/t$. Take expectation over u_2, \dots, u_n to get $\mathbf{P}\{\|B^{-1}e_1\| > t\} \leq 1/t$. This completes the proof of the claim.

Now let $w \sim \text{unif}(S^{n-1})$ be chosen independently of A . By conditioning on w and applying the claim, we see that $\mathbf{P}\{\|B^{-1}w\| \leq t\} \leq t/\sqrt{2\pi}$. Write the singular value decomposition of B as $B = s_1 v_1 y_1^t + \dots + s_n v_n y_n^t$, we see that $\|B^{-1}w\|^2 \geq s_1^{-2} |\langle w, v_1 \rangle|^2$. As w is uniform on the sphere, $|\langle w, v_1 \rangle|^2$ is at least $1/n$ with probability

■

A curious observation: Let $A_{m \times n}$ be the Gaussian matrix and write $a_{i,j} = \varepsilon_{i,j} |a_{i,j}|$ where $\varepsilon_{i,j} = \text{sgn}(a_{i,j})$. Observe that $\varepsilon_{i,j}$ and $|a_{i,j}|$ are independent. From the convexity of the norm, using Jensen's inequality for fixed $\varepsilon_{i,j}$ s, we get

$$\mathbf{E}[\|A\| \mid (\varepsilon_{i,j})] \geq \|(\varepsilon_{i,j} \mathbf{E}[|a_{i,j}|])\| = \sqrt{\frac{2}{\pi}} \|B\|$$

where $b_{i,j} = \varepsilon_{i,j}$. Thus, from the bound for the norm of a Gaussian matrix, we get the bound for the Bernoulli matrix B ,

$$\mathbf{E}[s_n(B)] = \mathbf{E}[\|B\|] \leq \sqrt{\frac{\pi}{2}} (\sqrt{m} + \sqrt{n}).$$

For what other distributions can we do this? If we can write a normal random variable as XY where X and Y are independent random variables, then by the same proof as above, we get $\mathbf{E}[\|B\|] \leq \frac{1}{\mathbf{E}[Y]} (\sqrt{m} + \sqrt{n})$, where the entries of B are i.i.d. with the same distribution as X . **I don't know, but would like to know, what random variables X can occur like this. Apart from Bernoulli, another example I see is that of uniform, by writing a $N(0, 1)$ variable as $\varphi^{-1}(U)V$ where $U \sim \text{unif}[0, \frac{1}{\sqrt{2\pi}}]$, $V \sim \text{unif}[-1, 1]$ and U, V are independent.**

Rectangular matrices with independent entries

Now suppose $a_{i,j}$ are i.i.d. We assume that they have a subgaussian distribution, meaning that $\mathbf{P}\{|a_{1,1}| \geq t\} \leq Ke^{-\kappa t^2}$ for all t , for some large constant K and some small constant κ . Various other constants that will appear in the statements will depend only on K and κ , unless otherwise stated.

Appendix 1: Weak convergence and techniques to show it

Probability measures on the real line

Let $\mathcal{P}(\mathbb{R})$ denote the space of all Borel probability measures on \mathbb{R} . They can be parameterized by *distribution functions*, which are functions $F : \mathbb{R} \mapsto [0, 1]$, that are non-decreasing, right-continuous and satisfy $F(+\infty) = 1$ and $F(-\infty) = 0$. The parameterization (correspondence) is given by sending the measure μ to the distribution function F_μ defined by $F_\mu(t) := \mu(-\infty, t]$ for $t \in \mathbb{R}$. The fact that every distribution function is the distribution function of a probability measure is told and proved in any first course in measure theoretic probability.

The use of this correspondence is that probability measures are more complicated objects, being functions on the large and intangible set $\mathcal{B}(\mathbb{R})$ (Borel sigma-algebra of \mathbb{R}) while distribution functions are (very restricted) functions on a smaller, more structured, familiar set, namely the real line. It is possible to work without this correspondence, but it helps, for example in defining the Lévy-Prohorov metric on $\mathcal{P}(\mathbb{R})$ by

$$\mathcal{D}(\mu, \nu) = \inf\{r > 0 : F_\mu(t+r) + r > F_\nu(t) \text{ and } F_\nu(t+r) + r > F_\mu(t) \text{ for all } t\}.$$

The form of the metric is less important than that it exists, and that convergence in this metric, denoted $\mu_n \xrightarrow{d} \mu$, is equivalent to $F_{\mu_n}(t) \rightarrow F_\mu(t)$ for some dense subset of $t \in \mathbb{R}$. This is the notion of *weak convergence* or *convergence in distribution*. We recall the following basic result from basic probability class which we recall without proof.

Lemma 92. *Let $\mu_n, \mu \in \mathcal{P}(\mathbb{R})$. The following are equivalent.*

1. $\mu_n \xrightarrow{d} \mu$, i.e., $\mathcal{D}(\mu_n, \mu) \rightarrow 0$.
2. $F_{\mu_n}(t) \rightarrow F_\mu(t)$ for all t where F_μ is continuous.

3. $\int f d\mu_n \rightarrow \int f d\mu$ for all $f \in C_b(\mathbb{R})$, the space of bounded continuous functions on \mathbb{R} .

4. $\int f d\mu_n \rightarrow \int f d\mu$ for all $f \in C_c^\infty(\mathbb{R})$.

Sometimes the third statement is taken as the definition of weak convergence, but the point to note is that as shown in the fourth statement, it suffices to prove the convergence of integrals for a much smaller class of integrands. This is of much practical importance. Another famous theorem from probability class shows that an even smaller class of functions suffices. Let $e_t(x) = e^{itx}$. The function $\hat{\mu}(t) := \int e_t d\mu$ is called the characteristic function (or Fourier transform) of μ . If X is a random variable with distribution μ , then $\hat{\mu}(t) = \mathbf{E}[e^{itX}]$.

Lemma 93 (Lévy's continuity theorem). *Let μ_n, μ be probability measures. Then $\mu_n \xrightarrow{d} \mu$ if and only if $\int e_t d\mu_n \rightarrow \int e_t d\mu$ for all $t \in \mathbb{R}$.*

A more basic fact than this is that the characteristic function uniquely determines the distribution.

Lemma 94 (Fourier inversion). *Let $\mu, \nu \in \mathcal{P}(\mathbb{R})$. The $\hat{\mu} = \hat{\nu}$ if and only if $\mu = \nu$. In fact, one can recover μ from $\hat{\mu}$ by*

$$\mu(a, b) + \frac{1}{2}\mu\{a, b\} = \lim_{L \rightarrow \infty} \int_{-L}^L \frac{e^{-iat} - e^{-ibt}}{2\pi it} \hat{\mu}(t) dt.$$

The practical use of having a small class of functions is seen, for instance, in the proof of central limit theorem, where the last two lemmas play a crucial role. In random matrix theory however, the characteristic function itself is of limited use, but there are two other techniques of great use.

1. The Stieltjes' transform, which is another integral transform akin to the Fourier transform.
2. The method of moments. This is the idea that we show that $\mu_n \xrightarrow{d} \mu$ by showing that the moments of μ_n converge to the corresponding moments of μ . Conditions apply, as we shall see.

The rest of the chapter is devoted to explaining the generalities of these two techniques.

Stieltjes' transform of a probability measure

Definition 95. For $\mu \in \mathcal{P}(\mathbb{R})$, its *Stieltjes' transform* is defined as

$$G_\mu(z) = \int \frac{1}{z-x} d\mu(x) = \mathbf{E} \left[\frac{1}{z-X} \right]$$

where X is a random variable with distribution μ . The Stieltjes' transform is certainly well-defined for $z \in \mathbb{C} \setminus \text{support}(\mu)$ as the integrand is bounded in that case. In particular, $G_\mu(z)$ makes sense for $z \in \mathbb{H}$.

Some simple observations on Stieltjes' transforms.

1. For any $\mu \in \mathcal{P}(\mathbb{R})$, $|G_\mu(z)| \leq \frac{1}{\text{Im}z}$ for $z \in \mathbb{H}$. This is because $|z-x| \geq \text{Im}z$ for $z \in \mathbb{H}$, $x \in \mathbb{R}$.
2. G_μ is holomorphic on $\mathbb{C} \setminus \text{support}(\mu)$. Indeed, if γ is any closed contour that does not enclose any point of $\text{spt}(\mu)$, then using the uniform boundedness of $1/(z-x)$ over $z \in \gamma$ and $x \in \text{spt}(\mu)$ to interchange integrals, we get

$$\int_\gamma G_\mu(z) dz = \int_\gamma \int_{\mathbb{R}} \frac{1}{z-x} d\mu(x) dz = \int_{\mathbb{R}} \int_\gamma \frac{1}{z-x} dz d\mu(x) = 0.$$

By Morera's theorem, G_μ is holomorphic on $\mathbb{C} \setminus \text{spt}(\mu)$.

3. Suppose μ is supported on a compact interval $[-a, a]$. Then, its moments $m_k := \int x^k \mu(dx)$ satisfy $|m_k| \leq a^k$ and hence $\sum m_k z^{-k-1}$ converges for $|z| > a$ and uniformly for $|z| \geq a + \delta$ for any $\delta > 0$. Hence,

$$\sum_{k=0}^{\infty} \frac{m_k}{z^{k+1}} = \mathbf{E} \left[\sum_{k=0}^{\infty} \frac{X^k}{z^k} \right] = \mathbf{E} \left[\frac{1}{z-X} \right] = G_\mu(z) \quad (3)$$

where the first equality follows by DCT. One can legitimately define $G_\mu(\infty) = 0$ and then (3) just gives the power series expansion of $w \rightarrow G_\mu(1/w)$ around 0.

4. If μ is compactly supported, $G_\mu(z) \sim \frac{1}{z}$ as $z \rightarrow \infty$. If μ is not compactly supported, the same is true for $z = iy$ as $y \uparrow \infty$.
5. If μ, ν are compactly supported and $G_\mu(z) = G_\nu(z)$ for all z in some open subset of \mathbb{H} , then by (3) the two measures have the same moments and hence $\mu = \nu$. But the condition of compact support is not required, as we shall see in Lemma 96.

As remarked earlier, the role of Stieltjes' transform in random matrix theory is analogous to the role of characteristic functions in classical limit theorems. But in fact it is a bit nicer/easier than the Fourier transform, as it can be seen as the convolution of the given probability measure with symmetric Cauchy distributions. As this fact will be crucial in the proofs below, we state it now.

The symmetric Cauchy distribution with parameter $y > 0$ is the probability measure $C_y \in \mathcal{P}(\mathbb{R})$ having density $\frac{y}{\pi(y^2+x^2)}$. As y approaches 0, the measure $C_y \xrightarrow{d} \delta_0$ (check!). It may also be noted that these distributions form a scale family, i.e., if $X \sim C_1$, then $yX \sim C_y$. Now,

$$\frac{-1}{\pi} \operatorname{Im} G_\mu(x+iy) = \frac{-1}{\pi} \int_{\mathbb{R}} \operatorname{Im} \left\{ \frac{1}{x+iy-t} \right\} \mu(dt) = \int_{\mathbb{R}} \frac{1}{\pi} \frac{y}{(x-t)^2+y^2} \mu(dt). \quad (4)$$

The last quantity is the density of $\mu \star C_y$. In other words, for each y , the function $\operatorname{Im} G_\mu(\cdot + iy)$ is (up to a factor $-1/\pi$) the density of the measure $\mu \star C_y$. As $y \rightarrow 0$ this should approach $\mu \star \delta_0 = \mu$. This is the essential point in the proof of the following lemma which gives the Stieltjes' transform analogues of Fourier inversion and Lévy's continuity theorems.

Lemma 96. *Let μ, ν be probability measures on \mathbb{R} .*

1. *For any $a < b$*

$$\lim_{y \downarrow 0} \frac{-1}{\pi} \int_a^b \operatorname{Im} G_\mu(x+iy) dx = \mu(a, b) + \frac{1}{2} \mu\{a\} + \frac{1}{2} \mu\{b\}. \quad (5)$$

2. *If $G_\mu(z) = G_\nu(z)$ for all z in an open subset of \mathbb{H} , then $\mu = \nu$.*

3. *If $\mu_n \rightarrow \mu$, then $G_{\mu_n} \rightarrow G_\mu$ pointwise on \mathbb{H} .*

4. *If $G_{\mu_n} \rightarrow G$ pointwise on \mathbb{H} for some $G : \mathbb{H} \rightarrow \mathbb{C}$, then G is the Stieltjes' transform of a possibly defective measure. If further, $iyG(iy) \rightarrow 1$ as $y \uparrow \infty$, then $G = G_\mu$ for a probability measure μ and $\mu_n \rightarrow \mu$.*

Proof. 1. Let X and Z be independent random variables on some probability space such that $X \sim \mu$ and $Z \sim C_1$. From (4), we know that

$$\frac{-1}{\pi} \int_a^b \operatorname{Im} G_\mu(x+iy) dx = \mathbf{E} [\mathbf{1}_{a \leq X+Y Z \leq b}].$$

Now $\lim_{y \downarrow 0} \mathbf{1}_{X+Y Z \in [a, b]} = \mathbf{1}_{X \in (a, b)} + \mathbf{1}_{X=a, Z>0} + \mathbf{1}_{X=b, Z<0}$. Take expectations, apply DCT, use the previous identity the independence of X and Z to get (5).

2. Follows immediately from the first part.
3. For fixed $z \in \mathbb{H}$, the function $x \mapsto \frac{1}{z-x}$ is a bounded and continuous on \mathbb{R} . Hence by one of the equivalent forms of the definition of weak convergence, $G_{\mu_n}(z) \rightarrow G_\mu(z)$.
4. Suppose that $G_{\mu_n} \rightarrow G$ pointwise for some function G . By Helly's selection principle, some subsequence μ_{n_k} converges vaguely to a possibly defective measure μ . As $(z-x)^{-1}$ is continuous and *vanishes at infinity*, $G_{\mu_{n_k}}(z) \rightarrow G_\mu(z)$ for all $z \in \mathbb{H}$.

Hence $G_\mu = G$ which shows that all subsequential limits have the same Stieltjes transform G . Further $iyG(iy) \rightarrow 1$ which shows that μ is a probability measure (for a general positive measure we have $G_\mu(iy) \sim \mu(\mathbb{R})/iy$ as $y \uparrow \infty$). By uniqueness of Stieltjes transforms, all subsequential limits are the same and hence $\mu_n \rightarrow \mu$. ■

In Lemma 101, we shall see a sharper version of the uniqueness theorem, by getting a bound on the distance between two probability measures in terms of the difference between their Stieltjes transforms. Here is a corollary that is sometimes helpful in proving absolute continuity of a measure from its Stieltjes' transform. This is to be contrasted against the difficulty of obtaining such information from moments³.

Corollary 97. *If $|G_\mu(z)| \leq B$ for some B and all $z \in \mathbb{H}$, then μ is absolutely continuous and has a density bounded by B/π .*

Proof. From the inversion formula, letting $y \downarrow 0$ in (5) we see that $\mu(a, b) \leq \frac{B}{\pi}(b-a)$ for all $a < b$. This proves the statement in the corollary. ■

Examples

Example 98. If $\mu = p_1\delta_{\lambda_1} + \dots + p_n\delta_{\lambda_n}$ is a probability measure on \mathbb{R} , its Stieltjes' transform is given by

$$G_\mu(z) = \sum_{k=1}^n \frac{p_k}{z - \lambda_k}.$$

For $y > 0$, it is easy to see that $x \mapsto \frac{-1}{\pi} \operatorname{Im} G_\mu(x + iy)$ is a mixture of Cauchy densities with scale y centered at $\lambda_1, \dots, \lambda_n$, and with p_k s as the mixing weights.

³Difficult, but not impossible. Markov proved a necessary and sufficient condition for having a bounded density in terms of the moments! But yes, it is difficult to check and impossible unless the moments are known very explicitly.

Example 99. Let μ be the arc-sine measure having density $\frac{1}{\pi\sqrt{1-x^2}}$ on $[-1, 1]$. Its Stieltjes' transform is

$$G_\mu(z) = \frac{1}{\sqrt{z^2 - 4}}.$$

Before explaining the result, we explain the meaning of the function on the right. Observe that $z \mapsto z^2 - 4$ maps \mathbb{H} onto $\mathbb{C} \setminus [-4, \infty)$. On this region, a holomorphic square-root can be defined, for example by taking $re^{i\theta} \mapsto re^{i\theta/2}$ where $r > 0$ and $0 < \theta < 2\pi$. The composition of these two functions is what is meant by $\sqrt{z^2 - 4}$.

One way to arrive at the result is to use the series expansion (3) that is valid for $|z| > 2$. The odd moments are zero while the $2p$ moment is $\binom{2p}{p}$ which may also be written as $(-1)^p 2^{2p} \binom{-1/2}{p}$. Hence

$$G_\mu(z) = \sum_{p=0}^{\infty} \binom{2p}{p} \frac{1}{z^{2p+1}} = \frac{1}{z} \sum_{p=0}^{\infty} \binom{-1/2}{p} \frac{(-1)^p 2^{2p}}{z^{2p}} = \frac{1}{z} \left(1 - \frac{4}{z^2}\right)^{-\frac{1}{2}}.$$

This is the same as $1/\sqrt{z^2 - 4}$ (don't simply write $\sqrt{ab} = \sqrt{a}\sqrt{b}$ to conclude this, the choice of square roots is dictated by the $1/z$ behaviour at infinity).

Example 100. If μ is the semi-circle measure having density $\frac{1}{2\pi}\sqrt{4-x^2}dx$ on $[-2, 2]$, then

$$G_\mu(z) = \frac{1}{2\pi} \int_{-2}^2 \frac{\sqrt{4-x^2}}{z-x} dx = \frac{z - \sqrt{z^2 - 4}}{2}.$$

The last integral can be computed by computing the series (3) again. We leave this as an exercise. Just remember that the odd moments are zero and the $2p$ moment is the Catalan numbers $C_p = \frac{1}{p+1} \binom{2p}{p}$.

The inversion formula can be verified in these cases directly. For example, consider the Stieltjes' transform $G(z) = 1/\sqrt{z^2 - 4}$. Let $z = x + iv$ where we shall let $v \rightarrow 0$. Hence $z^2 - 4 = x^2 - 4 - 2ixv + O(v^2)$. Hence, for $|x| < 2$,

$$\sqrt{z^2 - 4} = \sqrt{4 - x^2}(-1 - O(v))^{-1/2}$$

from which it is easy to see that

$$\frac{-1}{\pi} \operatorname{Im} G(x + iv) \rightarrow \frac{1}{\pi\sqrt{4 - x^2}} \quad \text{if } |x| < 2.$$

When $|x| > 2$, $G(x + iv)$ approaches the real line as $v \downarrow 0$, hence the above limit is zero.

Exactly the same way, for $G(z) = \frac{1}{2}(z - \sqrt{z^2 - 4})$, we get

$$\frac{-1}{\pi} \operatorname{Im} G(x + iv) = \frac{-1}{\pi} \frac{v - \operatorname{Im} \sqrt{(x + iv)^2 - 4}}{2} \rightarrow \frac{1}{2\pi} \sqrt{4 - x^2}.$$

This verifies the inversion formula in both cases.

Bounding Lévy distance in terms of Stieltjes transform

The following lemma is a quantitative statement that implies parts (2) and (4) of Lemma 96 as easy corollaries (how do you get part (4) of Lemma 96?). The Fourier transform analogue of this is a well-known lemma that is used in the proof of Berry-Esseen theorem (see ?, chapter.??).

We introduce two stronger metrics on probability measures. Recall the *Kolmogorov-Smirnov* distance

$$d_{KS}(\mu, \nu) = \sup_{x \in \mathbb{R}} |F_\mu(x) - F_\nu(x)|.$$

and the *total variation* distance

$$d_{TV}(\mu, \nu) = \frac{1}{2} \int |f(x) - g(x)| dx$$

if μ and ν have densities f and g . It can be defined more generally, but we shall use total variation distance only when densities exist. It is an easy exercise to check that

$$\mathcal{D}(\mu, \nu) \leq d_{KS}(\mu, \nu) \leq d_{TV}(\mu, \nu). \quad (6)$$

Lemma 101. *Let $\mu, \nu \in \mathcal{P}(\mathbb{R})$. Then, for any $y > 0$, we have*

$$\mathcal{D}(\mu, \nu) \leq 3\sqrt{y} + \frac{1}{\pi} \int_{\mathbb{R}} |\operatorname{Im} G_\mu(x + iy) - \operatorname{Im} G_\nu(x + iy)| dx.$$

Proof. Let $\mu_y = \mu \star C_y$ and $\nu_y = \nu \star C_y$. We bound the Lévy distance between μ and ν in three stages.

$$\mathcal{D}(\mu, \nu) \leq \mathcal{D}(\mu_y, \mu) + \mathcal{D}(\nu_y, \nu) + \mathcal{D}(\mu_y, \nu_y).$$

By the proof of Lemma 96 we know that μ_y has density $-\pi^{-1} \operatorname{Im} G_\mu(x + iy)$ and similarly for ν_y . Hence, by the second inequality in (6),

$$\mathcal{D}(\mu_y, \nu_y) \leq \frac{1}{\pi} \int_{\mathbb{R}} |\operatorname{Im} G_\mu(x + iy) - \operatorname{Im} G_\nu(x + iy)| dx.$$

Next we control $\mathcal{D}(\mu_y, \mu)$. Let $X \sim \mu$ and $Z \sim C_1$ so that $V = X + yZ \sim \mu_y$. For $t > 0$ observe that $\mathbf{P}(Z > t) = \int_t^\infty \pi^{-1}(1+u^2)^{-1} du \leq \int_t^\infty \pi^{-1}u^{-2} du = \pi^{-1}t^{-1}$. Thus, for any $\delta > 0$, we get

$$\mathbf{P}(X \leq t, V > t + \delta) \leq \mathbf{P}(Z > y^{-1}\delta) \leq \pi^{-1}\delta^{-1}y$$

$$\mathbf{P}(V \leq t, X > t + \delta) \leq \mathbf{P}(Z < -y^{-1}\delta) \leq \pi^{-1}\delta^{-1}y.$$

These immediately gives a bound of $\delta + \pi^{-1}\delta^{-1}y$ for $\mathcal{D}(\mu, \mu_y)$. Choose $\delta = \sqrt{y/\pi}$ to get

$$\mathcal{D}(\mu, \mu_y) \leq \frac{2}{\sqrt{\pi}}\sqrt{y}.$$

The same bound holds for $\mathcal{D}(\nu, \nu_y)$. Combine with (10) to get the inequality in the statement. ■

Method of moments

If $\mu \in \mathcal{P}(\mathbb{R})$, its moments are given by $\alpha_p = \alpha_p(\mu) = \int x^p d\mu(x)$, provided the integral exists. Two natural questions are,

1. If $\alpha_p(\mu) = \alpha_p(\nu)$ for all $p \in \mathbb{N}$, then is it necessarily true that $\mu = \nu$?
2. If $\alpha_p(\mu_n) \rightarrow \alpha_p(\mu)$ as $n \rightarrow \infty$, for each $p \in \mathbb{N}$, is it necessarily true that $\mu_n \xrightarrow{d} \mu$?

The answer to the first question is “No, in general”. But what is practically useful for us is that there are easy-to-check sufficient conditions under which the answer is “Yes”. This is Theorem 102

The answer to the second question is surprisingly clean: Yes, if and only if μ is determined by its moments (i.e., there is no other measure with the same moments as μ). This is Theorem 103.

Theorem 102. *Let $\mu \in \mathcal{P}(\mathbb{R})$ with all moments and let $\alpha_p = \int x^p d\mu(x)$. Then, μ is determined by its moments if any of the following (progressively weaker) conditions is satisfied.*

1. μ is compactly supported. This is equivalent to the condition that $\limsup_{m \rightarrow \infty} \alpha_{2m}^{1/2m} < \infty$.
2. The moment generating function of μ exists in a neighbourhood of 0, i.e., $\int e^{tx} d\mu(x) < \infty$ for $|t| < \delta$ for some $\delta > 0$. This is equivalent to the condition that $\limsup_{m \rightarrow \infty} \frac{\alpha_{2m}^{1/2m}}{m} < \infty$

3. The moments satisfy Carleman's condition: $\sum_{m \geq 1} \frac{1}{\alpha_{2m}^{1/2m}} = \infty$.

The reason that all conditions are stated only in terms of the even moments is that odd moments can be bounded simply by Cauchy-Schwarz inequality: $\alpha_m^2 \leq \alpha_{2m}$. Hence, in the first two conditions, we may as well write $|\alpha_m|^{1/m}$ instead of $\alpha_{2m}^{1/2m}$. In the last condition, it is better to leave it as it is, since the odd moments could well be zero (which causes the sum to diverge for silly reasons).

One small point in addition to what we have already said - if moments of a sequence of probability measures converge, the resulting sequence of numbers is necessarily a moment sequence of a probability measure. This is an easy fact, because a necessary and sufficient condition for a sequence of numbers to be a moment sequence is that it must be positive semi-definite, i.e., $\sum_{i,j=0}^p c_i c_j \alpha_{i+j} \geq 0$ for all $c_i \in \mathbb{R}$.

Putting all this together, we have the following theorem. Nothing more than the statement of this theorem will be required in this course when using the method of moments.

Theorem 103. Let μ_n be a sequence in $\mathcal{P}(\mathbb{R})$. Assume that each μ_n has all moments which we denote as $\alpha_p^{(n)}$. Assume that $\alpha_p = \lim_{n \rightarrow \infty} \alpha_p^{(n)}$ exists for all p .

1. There exists a $\mu \in \mathcal{P}(\mathbb{R})$ having moments $(\alpha_p)_{p \geq 0}$.
2. If $(\alpha_p)_{p \geq 0}$ satisfy Carleman's condition, then such a measure μ is unique and $\mu_n \xrightarrow{d} \mu$.

Exercises

Exercise 104. If μ has a continuous density f , then show that $f(x) = -\frac{1}{\pi} \lim_{y \downarrow 0} \text{Im}\{G_\mu(x + iy)\}$.

Appendix 2: Some linear algebra facts

Bounds on eigenvalues

Let $\lambda_1, \dots, \lambda_n$ be eigenvalues of $A = (a_{i,j})_{i,j \leq n}$. Then,

$$\max_{k \leq n} |\lambda_k| \leq \max_{k \leq n} r_k \quad (7)$$

where $r_k = \sum_{j=1}^n |a_{k,j}|$. Indeed, if $Av = \lambda v$, then choosing an index k for which $|v_k| \geq |v_j|$ for all j , we get

$$|\lambda v_k| = |(Av)_k| \leq \sum_{j=1}^n |a_{k,j}| |v_j| \leq |v_k| r_k.$$

Thus $|\lambda| \leq r_k$ for at least one k , which proves (7).

A better theorem is the Gershgorin circles theorem which asserts that all the eigenvalues of A are contained in the union of the closed disks $\bar{\mathbb{D}}(a_{k,k}, s_k)$ where $s_k = r_k - |a_{k,k}|$. Since $|a_{k,k}| + s_k = r_k$, this theorem implies (7). In some lucky situations, the Gershgorin theorem allows one to give bounds on the locations of various other eigenvalues also.

Perturbations of eigenvalues

How do eigenvalues change when the matrix changes? It is easy to see that eigenvalues vary continuously as a function of the matrix entries. But to say more is usually difficult. The standard example is the matrix

$$A_{n,\varepsilon} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & \ddots & 0 & 0 \\ 0 & 0 & \ddots & \ddots & \ddots & 0 \\ 0 & \ddots & \ddots & \ddots & 1 & 0 \\ 0 & 0 & \ddots & 0 & 0 & 1 \\ \varepsilon & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

The characteristic polynomial is $z^n - \varepsilon$ whose roots are equispaced points on the circle $|z| = \varepsilon^{1/n}$. For fixed n , as $\varepsilon \rightarrow 0$, the eigenvalues converge to those of $A_{n,0}$. However, we see that the continuity gets worse as n increases. This is typically the problem in dealing with non-Hermitian matrices.

But for Hermitian matrices, the situation is much better. We prove two useful inequalities. The first is useful when only a few entries are perturbed, possibly by large amounts. The second is useful when many entries are changed, but by a small amount.

Lemma 105. *Let A and B be real symmetric (or Hermitian) matrices. If $A - B$ has rank r , then*

$$\mathcal{D}(L_A, L_B) \leq d_{KS}(L_A, L_B) \leq \frac{r}{n}.$$

Proof. Any matrix of rank r can be written as a sum of r rank 1 matrices. Hence it suffices to prove the lemma for $r = 1$.

Fix any x and let V be the span of eigenvectors of A having eigenvalue less than or equal to x and let W be the span of eigenvectors of B having eigenvalue greater x . If $v \in V \cap W$, then $\langle Av, v \rangle \leq x\langle v, v \rangle$ and $\langle Bv, v \rangle > x\langle v, v \rangle$. Therefore, v is not in the kernel of $A - B$. Since $A - B$ has rank 1, this shows that $\dim(V \cap W) \leq 1$. But $\dim(V) = nF_A(x)$ and $\dim(W) = n(1 - F_B(x))$. Hence,

$$nF_A(x) + n(1 - F_B(x)) = \dim(V) + \dim(W) \leq 1 + n,$$

which gives $n|F_A(x) - F_B(x)| \leq 1$. ■

Lemma 106 (Hoffman-Wielandt inequality). *Let A and B be real symmetric (or Hermitian) matrices. Let $\lambda_1 \geq \dots \geq \lambda_n$ be the eigenvalues of A and let $\mu_1 \geq \dots \geq \mu_n$ be the eigenvalues of B . Then,*

$$\sum_{k=1}^n (\lambda_k - \mu_k)^2 \leq \text{tr}(A - B)^2.$$

Proof. Square both sides and use $\text{tr}(A^2) = \sum_{k=1}^n \lambda_k^2$ and $\text{tr}(B^2) = \sum_{k=1}^n \mu_k^2$ to rewrite the desired inequality as $\text{tr}(AB) \leq \sum_{k=1}^n \lambda_k \mu_k$. Let $A = \sum_{j=1}^n \lambda_j v_j v_j^t$ and $B = \sum_{k=1}^n \mu_k w_k w_k^t$ be the spectral decompositions of A and B . Then,

$$\text{tr}(AB) = \sum_{j,k=1}^n \lambda_j \mu_k \langle v_j, w_k \rangle^2.$$

The matrix $(\langle v_j, w_k \rangle^2)_{j,k \leq n}$ is doubly stochastic, since $\{v_j\}$ and $\{w_k\}$ are orthonormal bases of \mathbb{R}^n . Not all doubly stochastic matrices are of this form, but if the maximum of $f(D) =$

$\sum_{j,k=1}^n \lambda_j \mu_k D_{j,k}$ over all doubly stochastic D is attained within the set of *orthostochastic matrices* (those of the form $(\langle v_j, w_k \rangle^2)$), then that is also the maximum over orthostochastic matrices.

The set of doubly stochastic matrices DS_n is a compact convex set whose extreme points are permutation matrices (this is known as the Birkoff-von Neumann theorem). And f is a convex function (in fact linear!), hence it attains its maximum on DS_n at a permutation matrix. It may be noted that permutation matrices are orthostochastic. Now, among all permutations π , the quantity $f(\pi) = \sum_{k=1}^n \lambda_k \mu_{\pi(k)}$ is maximized when π is the identity. To see this, observe that if there are $i < j$ such that $\pi(i) > \pi(j)$, then

$$\lambda_j \mu_{\pi(j)} + \lambda_k \mu_{\pi(k)} < \lambda_j \mu_j + \lambda_k \mu_k$$

since $\lambda_j > \mu_j$ and $\lambda_k > \mu_k$. Hence, any inversion only decreases the value of f , showing that identity maximizes $f(\pi)$. And the maximum value of f is $\sum_j \lambda_j \mu_j$. ■

Here is how to get a bound for the Lévy-Prohorov distance between empirical measures from the bound given in Hoffman-Wielandt inequality.

Corollary 107. *In the notation of Lemma 106, we have $\mathcal{D}(L_A, L_B) \leq (\frac{1}{n} \text{tr}(A - B)^2)^{1/3}$.*

Proof. If $\mathcal{D}(L_A, L_B) > \delta$, then there is some x such that $F_A(x) > F_B(x + \delta) + \delta$. This means that there must be at least $n\delta$ indices i for which $\lambda_i \leq x$ but $\mu_i > x + \delta$. But then, $\sum_i (\lambda_i - \mu_i)^2 \geq n\delta^3$. From the Hoffman-Wielandt inequality, this is at most $\text{tr}(A - B)^2$ which gives the bound $\delta \leq (\frac{1}{n} \text{tr}(A - B)^2)^{1/3}$. ■

Some times, we need to compare eigenvalues of a matrix to those of a submatrix, say by removing one row and one column. Then the perturbation inequalities can be modified as follows. Let

$$A = \begin{bmatrix} a & v^t \\ v & B \end{bmatrix}, \quad C = \begin{bmatrix} 0 & 0^t \\ 0 & B \end{bmatrix}.$$

As A and C are matrices of the same size, we have

1. $d_{KS}(L_A, L_C) \leq \frac{1}{n}$ from Lemma 105.
2. $\sum_{k=1}^n (\lambda_k^A - \lambda_k^C)^2 \leq a^2 + 2\|v\|^2$ from Lemma 106.

But the eigenvalues of C are the eigenvalues of B along with an extra zero eigenvalue. That is, for some $1 \leq \ell \leq n-1$, we have

$$\lambda_i^C = \lambda_i^B \text{ for } i < \ell, \quad \lambda_\ell^C = 0, \quad \lambda_i^C = \lambda_{i-1}^B \text{ for } i > \ell.$$

Putting these together, we also get a comparison of eigenvalues of A with those of B . We summarize this as a lemma.

Lemma 108. *If A is a real symmetric (or Hermitian) $n \times n$ matrix and B is the matrix got from A by deleting the first row and first column, then*

1. $d_{KS}(L_A, L_B) \leq \frac{2}{n}$.
2. $\sum_{k=1}^{\ell-1} (\lambda_k^A - \lambda_k^B)^2 + \sum_{k=\ell+1}^n (\lambda_k^A - \lambda_{k-1}^B)^2 \leq a^2 + 2\|v\|^2$.
3. $\mathcal{D}(L_A, L_B) \leq \frac{(a^2 + 2\|v\|^2)^{1/3}}{(n-1)^{1/3}}$.

Proof. 1. This follows from the fact that $d_{KS}(L_A, L_C) \leq \frac{1}{n}$ and the exact relationship between eigenvalues of C and B .

2. This follows simply by dropping the $(\lambda_\ell^A - \lambda_\ell^C)^2$ term in the bound for $\sum_{k=1}^n (\lambda_k^A - \lambda_k^C)^2$.
3. Repeat the argument in the proof of Corollary 107. Just note that L_B has atoms of size $1/(n-1)$. ■

The second statement followingThe first bound can be improved slightly because the eigenvalues of A and B interlace. Instead of setting the first row and column of C to zero, we could have set them to any value etc. But this bound is good enough for us.

Block matrix inversion formula

Consider an $(m+n) \times (m+n)$ matrix written in block form as

$$X = \begin{bmatrix} A_{m \times m} & B_{m \times n} \\ C_{n \times m} & D_{n \times n} \end{bmatrix}.$$

Then, assuming all relevant matrices are invertible,

$$X^{-1} = \begin{bmatrix} (A - BD^{-1}C)^{-1} & \star \\ \star & (D - CA^{-1}B)^{-1} \end{bmatrix}. \quad (8)$$

The off-diagonal blocks can also be written explicitly but we omit that here. Particularly useful is the case when $m = 1$.

$$\begin{bmatrix} a & b^t \\ c & D \end{bmatrix}^{-1} = \begin{bmatrix} \frac{1}{a-b^t D^{-1}c} & \star \\ \star & (D - \frac{1}{a}cb^t)^{-1} \end{bmatrix}. \quad (9)$$

In particular, $X^{1,1} = \frac{1}{a-b^t D^{-1}c}$. This can be seen from the fact that $X^{1,1} = |D|/|A|$ and $|A| = |D|(a - b^t D^{-1}c)$. We omit the proofs, which are easy and can be looked up in many introductory linear algebra books (at least those that do a bit of matrices).

Shooting description of eigenvectors and eigenvalues of a Jacobi matrix

Let $T_n = T_n(a, b)$ be a Jacobi matrix. Fix a number $x \in \mathbb{R}$. Suppose we want to count how many eigenvalues of T_n are above x . How do we do it? As we shall see, this does not require us to compute the eigenvalues at all!

For $1 \leq k \leq n$, let φ_k be the characteristic polynomial of the top-left $k \times k$ principal submatrix of T . We also set $\varphi_0 = 1$ and $\varphi_{-1} = 0$. As we saw in (7), they satisfy the three-term recurrence

$$\varphi_k(x) = (x - a_k)\varphi_{k-1}(x) - b_{k-1}^2\varphi_{k-2}(x).$$

Thus, to compute $\varphi_k(x)$ for a fixed x , we have a simple recursive formula. Now we claim that we can count the number of eigenvalues of T_n that are above x by just looking at this sequence of number $\varphi_0(x), \dots, \varphi_n(x)$ which sounds surprising!

Claim 109. *The number of eigenvalues of T_n in (x, ∞) is equal to the number of sign-changes of the sequence $\varphi_0(x), \dots, \varphi_n(x)$.*

Indeed, the three term recurrences imply that the roots of φ_k and the roots of φ_{k-1} strictly interlace. This argument was given soon after (7). The key point was that when x is a root of φ_k , the recurrence formula shows that φ_{k+1} and φ_{k-1} have opposing signs. Therefore, if inductively we assume that the roots of φ_{k-1} interlace with those of φ_k , then it follows that the roots of φ_{k+1} also interlace with those of φ_k .

Now the stated claim follows by staring at Figure 1. But if you prefer words, here is a verbal argument.

Proof. Suppose the sequence $\varphi_0(x), \dots, \varphi_n(x)$ has k strict sign-changes. Then, there exist indices $0 = i_0 < i_1 < \dots < i_k \leq n$, such that $\varphi_{i_{2r}}(x) > 0$ and $\varphi_{i_{2r-1}}(x) < 0$. Since all the φ_j s are




Figure 1: The top line represents $\text{sgn}(\varphi_0)$, the next one $\text{sgn}(\varphi_1)$ and so on till the bottom line showing $\text{sgn}(\varphi_n)$. Red indicates positive, Blue indicates negative. Note the interlacing of zeros. From the vertical line drawn at x we can read off the signs of $\varphi_0(x), \dots, \varphi_n(x)$. For each of the three vertical lines, note that the number of times the color changes along the vertical line is equal to the number of zeros on the bottom line that lie to the right of the vertical line. Convince yourself that this is always the case (at least if the vertical line does not pass through any of the zeros on any of the lines).

positive eventually on the right (they are monic polynomials), this shows that φ_{i_1} must have a root in (x, ∞) , say $\lambda_{i_1,1}$. But then, φ_{i_2} must have a root above $\lambda_{i_1,1}$ as well as one in $(x, \lambda_{i_1,1})$ (otherwise $\varphi_{i_2}(x)$ would have been negative). Continuing this way, inductively we see that φ_{i_k} has at least k roots in (x, ∞) . By the interlacing property, each of $\varphi_{i_k+1}, \dots, \varphi_n$ must also have at least k roots in (x, ∞) . ■

Appendix 3: Gaussian random variables

Basics of Gaussians, moments, cumulants

Standard normal: A standard normal or Gaussian random variable is one with density $\varphi(x) := \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}x^2}$. Its distribution function is $\Phi(x) = \int_{-\infty}^x \varphi(t)dt$ and its tail distribution function is denoted $\bar{\Phi}(x) := 1 - \Phi(x)$. If X_i are i.i.d. standard normals, then $X = (X_1, \dots, X_n)$ is called a standard normal vector in \mathbb{R}^n . It has density $\prod_{i=1}^n \varphi(x_i) = (2\pi)^{-n/2} \exp\{-|\mathbf{x}|^2/2\}$ and the distribution is denoted by γ_n , so that for every Borel set A in \mathbb{R}^n we have $\gamma_n(A) = (2\pi)^{-n/2} \int_A \exp\{-|\mathbf{x}|^2/2\}d\mathbf{x}$.

Exercise 110. [Rotation invariance] If $P_{n \times n}$ is an orthogonal matrix, then $\gamma_n P^{-1} = \gamma_n$ or equivalently, $PX \stackrel{d}{=} X$. Conversely, if a random vector with independent co-ordinates has a distribution invariant under orthogonal transformations, then it has the same distribution as cX for some (non-random) scalar c .

Multivariate normal: If $Y_{m \times 1} = \mu_{m \times 1} + B_{m \times n} X_{n \times 1}$ where X_1, \dots, X_n are i.i.d. standard normal, then we say that $Y \sim N_m(\mu, \Sigma)$ with $\Sigma = BB^t$. Implicit in this notation is the fact that the distribution of Y depends only on Σ and not on the way in which Y is expressed as a linear combination of standard normals (this follows from Exercise 110). It is a simple exercise that $\mu_i = \mathbf{E}[X_i]$ and $\sigma_{i,j} = \text{Cov}(X_i, X_j)$. Since matrices of the form BB^t are precisely positive semi-definite matrices (defined as those $\Sigma_{m \times m}$ for which $\mathbf{v}^t \Sigma \mathbf{v} \geq 0$ for all $\mathbf{v} \in \mathbb{R}^m$), it is clear that covariance matrices of normal random vectors are precisely p.s.d. matrices. Clearly, if $Y \sim N_m(\mu, \Sigma)$ and $Z_{p \times 1} = C_{p \times m} Y + \theta_{p \times 1}$, then $Z \sim N_p(\theta + C\mu, C\Sigma C^t)$. Thus, affine linear transformations of normal random vectors are again normal.

Exercise 111. The random vector Y has density if and only if Σ is non-singular, and in that

case the density is

$$\frac{1}{(2\pi)^{n/2} \sqrt{\det(\Sigma)}} \exp \left\{ -\frac{1}{2} \mathbf{y}' \Sigma^{-1} \mathbf{y} \right\}.$$

If Σ is singular, then X takes values in a lower dimensional subspace in \mathbb{R}^n and hence does not have density.

Exercise 112. Irrespective of whether Σ is non-singular or not, the characteristic function of Y is given by

$$\mathbf{E} \left[e^{i\langle \lambda, Y \rangle} \right] = e^{-\frac{1}{2} \lambda' \Sigma \lambda}, \quad \text{for } \lambda \in \mathbb{R}^m.$$

In particular, if $X \sim N(0, \sigma^2)$, then its characteristic function is $\mathbf{E}[e^{i\lambda X}] = e^{-\frac{1}{2} \sigma^2 \lambda^2}$ for $\lambda \in \mathbb{R}$.

Exercise 113. If $U_{k \times 1}$ and $V_{(m-k) \times 1}$ are such that $Y^t = (U^t, V^t)$, and we write $\mu = (\mu_1, \mu_2)$ and $\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}$ are partitioned accordingly, then

1. $U \sim N_k(\mu_1, \Sigma_{11})$.
2. $U \Big|_V \sim N_k \left(\mu_1 - \Sigma_{12} \Sigma_{22}^{-1/2} V, \Sigma_{11} - \Sigma_{12} \Sigma_{22}^{-1} \Sigma_{21} \right)$ (assume that Σ_{22} is invertible).

Moments: All questions about a centered Gaussian random vector must be answerable in terms of the covariance matrix. In some cases, there are explicit answers.

Exercise 114. Prove the *Wick formula* (also called *Feynman diagram formula*) for moments of centered Gaussians.

1. Let $X \sim N_n(0, \Sigma)$. Then, $\mathbf{E}[X_1 \dots X_n] = \sum_{M \in \mathcal{M}_n} \prod_{\{i,j\} \in M} \sigma_{i,j}$, where \mathcal{M}_n is the collection of all matchings of the set $[n]$ (thus \mathcal{M}_n is empty if n is odd) and the product is over all matched pairs. For example, $\mathbf{E}[X_1 X_2 X_3 X_4] = \sigma_{12} \sigma_{34} + \sigma_{13} \sigma_{24} + \sigma_{14} \sigma_{23}$.
2. If $\xi \sim N(0, 1)$, then $\mathbf{E}[\xi^{2n}] = (2n-1)(2n-3) \dots (3)(1)$.

Cumulants: Let X be a real-valued random variable with $\mathbf{E}[e^{tX}] < \infty$ for t in a neighbourhood of 0. Then, we can write the power series expansions

$$\mathbf{E}[e^{i\lambda X}] = \sum_{k=0}^{\infty} m_n(X) \frac{\lambda^n}{n!}, \quad \log \mathbf{E}[e^{i\lambda X}] = \sum_{k=1}^{\infty} \kappa_n[X] \frac{\lambda^n}{n!}.$$

Here $m_n[X] = \mathbf{E}[X^n]$ are the moments while $\kappa_n[X]$ is a linear combination of the first n moments ($\kappa_1 = m_1$, $\kappa_2 = m_2 - m_1^2$, etc). Then κ_n is called the n th cumulant of X . If X and Y are independent, then it is clear that $\kappa_n[X + Y] = \kappa_n[X] + \kappa_n[Y]$.

Exercise 115. (optional). Prove the following relationship between moments and cumulants. The sums below are over partitions Π of the set $[n]$ and $\Pi_1, \dots, \Pi_{\ell_\Pi}$ denote the blocks of Π .

$$m_n[X] = \sum_{\Pi} \prod_i \kappa_{|\Pi_i|}[X], \quad \kappa_n[X] = \sum_{\Pi} (-1)^{\ell_\Pi - 1} \prod_i m_{|\Pi_i|}[X].$$

Thus $\kappa_1 = m_1$, $\kappa_2 = m_2 - m_1^2$,

Exercise 116. If $\xi \sim N(0, 1)$, then $\kappa_1 = 0$, $\kappa_2 = 1$ and $\kappa_n = 0$ for all $n \geq 3$.

The converse of this result is also true and often useful in proving that a random variable is normal. For instance, the theorem below implies that to show that a sequence of random variables converges to normal, it suffices to show that cumulants $\kappa_m[X_n] \rightarrow 0$ for all $m \geq m_0$ for some m_0 .

Result 117 (Marcinkiewicz). If X is a random variable with finite moments of all orders and $\kappa_n[X] = 0$ for all $n \geq n_0$ for some n_0 , then X is Gaussian.

Convergence and Gaussians:

Exercise 118. The family of distributions $N(\mu, \sigma^2)$, where $\mu \in \mathbb{R}$ and $0 \leq \sigma^2 < \infty$, is closed under convergence in distribution (for this statement to be valid we include $N(\mu, 0)$ which means δ_μ). Indeed, $N(\mu_n, \sigma_n^2) \xrightarrow{d} N(\mu, \sigma^2)$ if and only if $\mu_n \rightarrow \mu$ and $\sigma_n^2 \rightarrow \sigma^2$.

A vector space of Gaussian random variables: Let $Y \sim N_m(0, \Sigma)$ be a random vector in some probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Then, for every vector $\mathbf{v} \in \mathbb{R}^m$, define the random variable $Y_{\mathbf{v}} := \mathbf{v}^t Y$. Then, for any $\mathbf{v}_1, \dots, \mathbf{v}_j$, the random variables $Y_{\mathbf{v}_1}, \dots, Y_{\mathbf{v}_j}$ are jointly normal. The joint distribution of $\{Y_{\mathbf{v}}\}$ is fully specified by noting that $Y_{\mathbf{v}}$ have zero mean and $\mathbf{E}[Y_{\mathbf{v}} Y_{\mathbf{u}}] = \mathbf{v}^t \Sigma \mathbf{u}$.

We may interpret this as follows. If Σ is p.d. (p.s.d. and non-singular), then $(\mathbf{v}, \mathbf{u})_{\Sigma} := \mathbf{v}^t \Sigma \mathbf{u}$ defines an inner product on \mathbb{R}^m . On the other hand, the set $L_0^2(\Omega, \mathcal{F}, \mathbf{P})$ of real-valued random variables on Ω with zero mean and finite variance, is also an inner product space

under the inner product $\langle U, V \rangle := \mathbf{E}[UV]$. The observation in the previous paragraph is that $\mathbf{v} \rightarrow Y_{\mathbf{v}}$ is an isomorphism of $(\mathbb{R}^m, \langle \cdot, \cdot \rangle_{\Sigma})$ into $L_0^2(\Omega, \mathcal{F}, \mathbf{P})$.

In other words, given any finite dimensional inner-product space $(V, \langle \cdot, \cdot \rangle)$, we can find a collection of Gaussian random variables on some probability space, such that this collection is isomorphic to the given inner-product space. Later we shall see the same for Hilbert spaces⁴.

Tails of the Gaussian distribution: Recall the standard Gaussian density $\varphi(x)$. The corresponding cumulative distribution function is denoted by Φ and the tail is denoted by $\bar{\Phi}(x) := \int_x^\infty \varphi(t)dt$. The following estimate will be used very often.

Exercise 119. For all $x > 0$, we have $\frac{1}{\sqrt{2\pi}} \frac{x}{1+x^2} e^{-\frac{1}{2}x^2} \leq \bar{\Phi}(x) \leq \frac{1}{\sqrt{2\pi}} \frac{1}{x} e^{-\frac{1}{2}x^2}$. In particular⁵, $\bar{\Phi}(x) \sim x^{-1}\varphi(x)$ as $x \rightarrow \infty$. Most often the following simpler bound, valid for $x \geq 1$, suffices.

$$\frac{1}{10x} e^{-\frac{1}{2}x^2} \leq \bar{\Phi}(x) \leq e^{-\frac{1}{2}x^2}.$$

For $t > 0$, let $p_t(x) := \frac{1}{\sqrt{t}}\varphi(x/\sqrt{t})$ be the $N(0, t)$ density. We interpret $p_0(x)dx$ as the degenerate measure at 0. These densities have the following interesting properties.

Exercise 120. Show that $p_t \star p_s = p_{t+s}$, i.e., $\int_{\mathbb{R}} p_t(x-y)p_s(y)dy = p_{t+s}(x)$.

Exercise 121. Show that $p_t(x)$ satisfies the heat equation: $\frac{\partial}{\partial t} p_t(x) = \frac{1}{2} \frac{\partial^2}{\partial x^2} p_t(x)$ for all $t > 0$ and $x \in \mathbb{R}$.

Remark 122. Put together, these facts say that $p_t(x)$ is the *fundamental solution* to the heat equation. This just means that the heat equation $\frac{\partial}{\partial t} u(t, x) = \frac{1}{2} \frac{\partial^2}{\partial x^2} u(t, x)$ with the initial condition $u(0, x) = f(x)$ can be solved simply as $u(t, x) = (f \star p_t)(x) := \int_{\mathbb{R}} f(y)p_t(x-y)dy$. This works for reasonable f (say $f \in L^1(\mathbb{R})$).

We shall have many occasions to use the following “integration by parts” formula.

Exercise 123. Let $X \sim N_n(0, \Sigma)$ and let $F : \mathbb{R}^n \rightarrow \mathbb{R}$. Under suitable conditions on F (state sufficient conditions), show that $\mathbf{E}[X_i F(X)] = \sum_{j=1}^n \sigma_{ij} \mathbf{E}[\partial_j F(X)]$. As a corollary, deduce the Wick formula of Exercise 114.

⁴This may seem fairly pointless, but here is one thought-provoking question. Given a vector space of Gaussian random variables, we can multiply any two of them and thus get a larger vector space spanned by the given normal random variables and all pair-wise products of them. What does this new vector space correspond to in terms of the original $(V, \langle \cdot, \cdot \rangle)$?

⁵The notation $f(x) \sim g(x)$ means that $\lim_{x \rightarrow \infty} \frac{f(x)}{g(x)} = 1$.

Stein's equation: Here we may revert to $t = 1$, thus $p_1 = \phi$. Then, $\phi'(x) = -x\phi(x)$. Hence, for any $f \in C_b^1(\mathbb{R})$, we integrate by parts to get $\int f'(x)\phi(x)dx = -\int f(x)\phi'(x)dx = \int f(x)x\phi(x)dx$. If $X \sim N(0, 1)$, then we may write this as

$$\mathbf{E}[(Tf)(X)] = 0 \quad \text{for all } f \in C_b^1(\mathbb{R}), \text{ where } (Tf)(x) = f'(x) - xf(x). \quad (10)$$

The converse is also true. Suppose (10) holds for all $f \in C_b^1(\mathbb{R})$. Apply it to $f(x) = e^{i\lambda x}$ for any fixed $\lambda \in \mathbb{R}$ to get $\mathbf{E}[Xe^{i\lambda X}] = i\lambda\mathbf{E}[e^{i\lambda X}]$. Thus, if $\psi(\lambda) := \mathbf{E}[e^{i\lambda X}]$ is the characteristic function of X , then $\psi'(\lambda) = -\lambda\psi(\lambda)$ which has only one solution, $e^{-\lambda^2/2}$. Hence X must have standard normal distribution.

Digression - central limit theorem: One reason for the importance of normal distribution is of course the central limit theorem. The basic central limit theorem is for $W_n := (X_1 + \dots + X_n)/\sqrt{n}$ where X_i are i.i.d. with zero mean and unit variance. Here is a sketch of how central limit theorem can be proved using Stein's method. Let $f \in C_b^1(\mathbb{R})$ and observe that $\mathbf{E}[W_n f(W_n)] = \sqrt{n}\mathbf{E}[X_1 f(W_n)]$. Next, write

$$f\left(\frac{X_1 + \dots + X_n}{\sqrt{n}}\right) \approx f\left(\frac{X_2 + \dots + X_n}{\sqrt{n}}\right) + \frac{X_1}{\sqrt{n}}f'\left(\frac{X_2 + \dots + X_n}{\sqrt{n}}\right)$$

where we do not make precise the meaning of the approximation. Let $\hat{W}_n = \frac{X_2 + \dots + X_n}{\sqrt{n}}$. Then,

$$\mathbf{E}[W_n f(W_n)] \approx \sqrt{n}\mathbf{E}[X_1]\mathbf{E}[f(\hat{W}_n)] + \mathbf{E}[X_1^2]\mathbf{E}[f'(\hat{W}_n)] = \mathbf{E}[f'(\hat{W}_n)].$$

Since $\hat{W}_n \approx W_n$, this shows that $\mathbf{E}[Tf(W_n)] \approx 0$. We conclude that $W_n \approx N(0, 1)$.

There are missing pieces here, most important being the last statement - that if a random variable satisfies Stein's equation approximately, then it must be approximately normal. When included, one does get a proof of the standard CLT.

Comparison inequalities

The study of the maximum (or supremum) of a collection of Gaussian random variables is of fundamental importance. In such cases, certain comparison inequalities are helpful in reducing the problem at hand to the same problem for a simpler correlation matrix. We start with a lemma of this kind and from which we derive two important results - Slepian's inequality, Gordon's inequality and Sudakov-Fernique inequality⁶.

⁶The presentation here is cooked up from Ledoux-Talagrand (the book titled *Probability on Banach spaces*) and from Sourav Chatterjee's paper on Sudakov-Fernique inequality. Chatterjee's proof can be used to

Lemma 124 (J.P. Kahane). *Let X and Y be $n \times 1$ multivariate Gaussian vectors with equal means, i.e., $\mathbf{E}[X_i] = \mathbf{E}[Y_i]$ for all i . Let $A = \{(i, j) : \sigma_{ij}^X < \sigma_{ij}^Y\}$ and let $B = \{(i, j) : \sigma_{ij}^X > \sigma_{ij}^Y\}$. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be any C^2 function all of whose partial derivatives up to second order have subgaussian growth and such that $\partial_i \partial_j f \geq 0$ for all $(i, j) \in A$ and $\partial_i \partial_j f \leq 0$ for all $(i, j) \in B$. Then, $\mathbf{E}[f(X)] \leq \mathbf{E}[f(Y)]$.*

Proof. First assume that both X and Y are centered. Without loss of generality we may assume that X and Y are defined on the same probability space and independent of each other.

Interpolate between them by setting $Z(\theta) = (\cos \theta)X + (\sin \theta)Y$ for $0 \leq \theta \leq \frac{\pi}{2}$ so that $Z(0) = X$ and $Z(\pi/2) = Y$. Then,

$$\mathbf{E}[f(Y)] - \mathbf{E}[f(X)] = \mathbf{E} \left[\int_0^{\pi/2} \frac{d}{d\theta} f(Z(\theta)) d\theta \right] = \int_0^{\pi/2} \frac{d}{d\theta} \mathbf{E}[f(Z_\theta)] d\theta.$$

The interchange of expectation and derivative etc., can be justified by the conditions on f but we shall omit these routine checks. Further,

$$\frac{d}{d\theta} \mathbf{E}[f(Z_\theta)] = \mathbf{E}[\nabla f(Z_\theta) \cdot \dot{Z}(\theta)] = \sum_{i=1}^n \{ -(\sin \theta) \mathbf{E}[X_i \partial_i f(Z_\theta)] + (\cos \theta) \mathbf{E}[Y_i \partial_i f(Z_\theta)] \}.$$

Now use Exercise 123 to deduce (apply the exercise after conditioning on X or Y and using the independence of X and Y) that

$$\begin{aligned} \mathbf{E}[X_i \partial_i f(Z_\theta)] &= (\cos \theta) \sum_{j=1}^n \sigma_{ij}^X \mathbf{E}[\partial_i \partial_j f(Z_\theta)] \\ \mathbf{E}[Y_i \partial_i f(Z_\theta)] &= (\sin \theta) \sum_{j=1}^n \sigma_{ij}^Y \mathbf{E}[\partial_i \partial_j f(Z_\theta)]. \end{aligned}$$

Consequently,

$$\frac{d}{d\theta} \mathbf{E}[f(Z_\theta)] = (\cos \theta)(\sin \theta) \sum_{i,j=1}^n \mathbf{E}[\partial_i \partial_j f(Z_\theta)] (\sigma_{ij}^Y - \sigma_{ij}^X). \quad (11)$$

The assumptions on $\partial_i \partial_j f$ ensure that each term is non-negative. Integrating, we get $\mathbf{E}[f(X)] \leq \mathbf{E}[f(Y)]$.

It remains to consider the case when the means are not zero. Let $\mu_i = \mathbf{E}[X_i] = \mathbf{E}[Y_i]$ and set $\hat{X}_i = X_i - \mu_i$ and $\hat{Y}_i = Y_i - \mu_i$ and let $g(x_1, \dots, x_n) = f(x_1 + \mu_1, \dots, x_n + \mu_n)$. Then $f(X) = g(\hat{X})$ and $f(Y) = g(\hat{Y})$ while $\partial_i \partial_j g(x) = \partial_i \partial_j f(x + \mu)$. Thus, the already proved statement for centered variables implies the one for non-centered variables. ■

prove Kahane's inequality too, and consequently Slepian's, and that is the way we present it here.

Special cases of this lemma are very useful. We write X^* for $\max_i X_i$.

Corollary 125 (Slepian's inequality). *Let X and Y be $n \times 1$ multivariate Gaussian vectors with equal means, i.e., $\mathbf{E}[X_i] = \mathbf{E}[Y_i]$ for all i . Assume that $\sigma_{ii}^X = \sigma_{ii}^Y$ for all i and that $\sigma_{ij}^X \geq \sigma_{ij}^Y$ for all i, j . Then,*

1. For any real t_1, \dots, t_n , we have $\mathbf{P}\{X_i < t_i \text{ for all } i\} \geq \mathbf{P}\{Y_i < t_i \text{ for all } i\}$.
2. $X^* \prec Y^*$, i.e., $\mathbf{P}\{X^* > t\} \leq \mathbf{P}\{Y^* > t\}$ for all t .

Proof. In the language of Lemma 124 by taking $B \subseteq \{(i, i) : 1 \leq i \leq n\}$ while $A = \emptyset$. We would like to say that the first conclusion follows by simply taking $f(x_1, \dots, x_n) = \prod_{i=1}^n \mathbf{1}_{x_i < t_i}$. The only wrinkle is that it is not smooth. by approximating the indicator with smooth increasing functions, we can get the conclusion.

To elaborate, let $\psi \in C^\infty(\mathbb{R})$ be an increasing function $\psi(t) = 0$ for $t < 0$ and $\psi(t) = 1$ for $t > 1$. Then $\psi_\varepsilon(t) = \psi(t/\varepsilon)$ increases to $\mathbf{1}_{t < 0}$ as $\varepsilon \downarrow 0$. If $f_\varepsilon(x_1, \dots, x_n) = \prod_{i=1}^n \psi_\varepsilon(x_i - t_i)$, then $\partial_{ij} f \geq 0$ and hence Lemma 124 applies to show that $\mathbf{E}[f_\varepsilon(X)] \leq \mathbf{E}[f_\varepsilon(Y)]$. Let $\varepsilon \downarrow 0$ and apply monotone convergence theorem to get the first conclusion.

Taking $t_i = t$, we immediately get the second conclusion from the first. ■

Here is a second corollary which generalizes Slepian's inequality (take $m = 1$).

Corollary 126 (Gordon's inequality). *Let $X_{i,j}$ and $Y_{i,j}$ be $m \times n$ arrays of joint Gaussians with equal means. Assume that*

1. $\text{Cov}(X_{i,j}, X_{i,\ell}) \geq \text{Cov}(Y_{i,j}, Y_{i,\ell})$,
2. $\text{Cov}(X_{i,j}, X_{k,\ell}) \leq \text{Cov}(Y_{i,j}, Y_{k,\ell})$ if $i \neq k$,
3. $\text{Var}(X_{i,j}) = \text{Var}(Y_{i,j})$.

Then

1. For any real $t_{i,j}$ we have $\mathbf{P}\left\{\bigcap_j \bigcup_i \{X_{i,j} < t_{i,j}\}\right\} \geq \mathbf{P}\left\{\bigcap_j \bigcup_i \{Y_{i,j} < t_{i,j}\}\right\}$,
2. $\min_i \max_j X_{i,j} \prec \min_i \max_j Y_{i,j}$.

Exercise 127. Deduce this from Lemma 124.

Remark 128. The often repeated trick that we referred to is of constructing the two random vectors independently on the same space and interpolating between them. Then the comparison inequality reduces to a differential inequality which is simpler to deal with. Quite often different parameterizations of the same interpolation are used, for example $Z_t = \sqrt{1-t^2}X + tY$ for $0 \leq t \leq 1$ or $Z_s = \sqrt{1-e^{-2s}}X + e^{-s}Y$ for $-\infty \leq s \leq \infty$.

Studying the maximum of a Gaussian process is a very important problem. Slepian's (or Gordon's) inequality helps to control the maximum of our process by that of a simpler process. For example, if X_1, \dots, X_n are standard normal variables with positive correlation between any pair of them, then $\max X_i$ is stochastically smaller than the maximum of n independent standard normals (which is easy). However, the conditions of Slepian's inequality are sometimes restrictive, and the conclusions are much stronger than often required. The following theorem is a more applicable substitute.

Theorem 129 (Sudakov-Fernique inequality). *Let X and Y be $n \times 1$ Gaussian vectors satisfying $\mathbf{E}[X_i] = \mathbf{E}[Y_i]$ for all i and $\mathbf{E}[(X_i - X_j)^2] \leq \mathbf{E}[(Y_i - Y_j)^2]$ for all $i \neq j$. Then, $\mathbf{E}[X^*] \leq \mathbf{E}[Y^*]$.*

Remark 130. Assume that the processes are centered. If the two processes had the same variances, then the condition $\mathbf{E}[(X_i - X_j)^2] \leq \mathbf{E}[(Y_i - Y_j)^2]$ would be the same as $\text{Cov}(X_i, X_j) \geq \text{Cov}(Y_i, Y_j)$. In that case, Slepian's inequality would apply and we would get the much stronger conclusion of $X^* \prec Y^*$. The point here is that we relax the assumption of equal variances and settle for the weaker conclusion which only compares expectations of the maxima.

For non-centered processes one may wonder whether it would not be more appropriate to compare $\text{Var}(X_i - X_j)$ with $\text{Var}(Y_i - Y_j)$ in the assumption. But since $\mathbf{E}[(X_i - X_j)^2] = \text{Var}(X_i - X_j) + (\mathbf{E}[X_i] - \mathbf{E}[X_j])^2$, and the means are assumed to be equal, that would be the same condition!

Proof. The proof of Lemma 124 can be copied exactly to get (11) for any smooth function f with appropriate growth conditions. Now we specialize to the function $f_\beta(x) = \frac{1}{\beta} \log \sum_{i=1}^n e^{\beta x_i}$ where $\beta > 0$ is fixed. Let $p_i(x) = \frac{e^{\beta x_i}}{\sum_{i=1}^n e^{\beta x_i}}$, so that $(p_1(x), \dots, p_n(x))$ is a probability vector for each $x \in \mathbb{R}^n$. Observe that

$$\begin{aligned} \partial_i f(x) &= p_i(x) \\ \partial_i \partial_j f(x) &= \beta p_i(x) \delta_{i,j} - \beta p_i(x) p_j(x). \end{aligned}$$

Thus, (11) gives

$$\begin{aligned} \frac{1}{\beta(\cos\theta)(\sin\theta)} \frac{d}{d\theta} \mathbf{E}[f_\beta(Z_\theta)] &= \sum_{i,j=1}^n (\sigma_{ij}^Y - \sigma_{ij}^X) \mathbf{E}[p_i(x)\delta_{i,j} - p_i(x)p_j(x)] \\ &= \sum_{i=1}^n (\sigma_{ii}^Y - \sigma_{ii}^X) \mathbf{E}[p_i(x)] - \sum_{i,j=1}^n (\sigma_{ij}^Y - \sigma_{ij}^X) \mathbf{E}[p_i(x)p_j(x)] \end{aligned}$$

Since $\sum_i p_i(x) = 1$, we can write $p_i(x) = \sum_j p_i(x)p_j(x)$ and hence

$$\begin{aligned} \frac{1}{\beta(\cos\theta)(\sin\theta)} \frac{d}{d\theta} \mathbf{E}[f_\beta(Z_\theta)] &= \sum_{i,j=1}^n (\sigma_{ii}^Y - \sigma_{ii}^X) \mathbf{E}[p_i(x)p_j(x)] - \sum_{i,j=1}^n (\sigma_{ij}^Y - \sigma_{ij}^X) \mathbf{E}[p_i(x)p_j(x)] \\ &= \sum_{i<j} \mathbf{E}[p_i(x)p_j(x)] (\sigma_{ii}^Y - \sigma_{ii}^X + \sigma_{jj}^Y - \sigma_{jj}^X - 2\sigma_{ij}^Y + 2\sigma_{ij}^X) \\ &= \sum_{i<j} \mathbf{E}[p_i(x)p_j(x)] (\gamma_{ij}^X - \gamma_{ij}^Y) \end{aligned}$$

where $\gamma_{ij}^X = \sigma_{ii}^X + \sigma_{jj}^X - 2\sigma_{ij}^X = \mathbf{E}[(X_i - \mu_i - X_j + \mu_j)^2]$. Of course, the latter is equal to $\mathbf{E}[(X_i - X_j)^2] - (\mu_i - \mu_j)^2$. Since the μ_i are the same for X as for Y we get $\gamma_{ij}^X \leq \gamma_{ij}^Y$. Clearly $p_i(x) \geq 0$ too. Therefore, $\frac{d}{d\theta} \mathbf{E}[f_\beta(Z_\theta)] \geq 0$ and we get $\mathbf{E}[f_\beta(X)] \leq \mathbf{E}[f_\beta(Y)]$. Letting $\beta \uparrow \infty$ we get $\mathbf{E}[X^*] \leq \mathbf{E}[Y^*]$. ■

Remark 131. This proof contains another useful idea - to express $\max_i x_i$ in terms of $f_\beta(x)$. The advantage is that f_β is smooth while the maximum is not. And for large β , the two are close because $\max_i x_i \leq f_\beta(x) \leq \max_i x_i + \frac{\log n}{\beta}$.

If Sudakov-Fernique inequality is considered a modification of Slepian's inequality, the analogous modification of Gordon's inequality is the following. We leave it as exercise as we may not use it in the course.

Exercise 132. Let $X_{i,j}$ and $Y_{i,j}$ be $n \times m$ arrays of joint Gaussians with equal means. Assume that

1. $\mathbf{E}[|X_{i,j} - X_{i,\ell}|^2] \geq \mathbf{E}[|Y_{i,j} - Y_{i,\ell}|^2]$,
2. $\mathbf{E}[|X_{i,j} - X_{k,\ell}|^2] \leq \mathbf{E}[|Y_{i,j} - Y_{k,\ell}|^2]$ if $i \neq k$.

Then $\mathbf{E}[\min_i \max_j X_{i,j}] \geq \mathbf{E}[\min_i \max_j Y_{i,j}]$.

Remark 133. All through this section, we have stated comparison inequalities for two Gaussian vectors of the same dimension. What about infinite, even uncountable, index sets? Indeed, suppose T is an index set and $X = (X_t)_{t \in T}$ is a Gaussian process on T . By this we just mean that all finite linear combinations $c_1 X_{t_1} + \dots + c_n X_{t_n}$ are Gaussian. We want to talk about $X_T^* = \sup_{t \in T} X_t$. One possible issue is that this is not measurable. For instance, if $T = [0, 1]$ and $X_t, t \in T$, are i.i.d. $N(0, 1)$. But this is hardly of significance.

If there is any reasonable sample-path regularity in $t \mapsto X_t$ (for example almost sure continuity if T has a topology), then X^* turns out to be measurable trivially. And then it follows that $\mathbf{E}[X_T^*] = \sup_F \mathbf{E}[X_F^*]$ where the supremum is over finite $F \subseteq T$. For example, in this book we used Gaussian processes such as $X(u, v) = v^t A u$ where A is an $m \times n$ matrix with i.i.d. Gaussian entries and $u \in S^{n-1}, v \in S^{m-1}$. Continuity of X is clear and hence X^* is well-defined and measurable.

The point of all this is that we may apply all the comparison inequalities we have obtained to Gaussian processes on arbitrary index sets.

Gaussian isoperimetric inequality

Let γ_m denotes the standard Gaussian measure on \mathbb{R}^m . For a set $A \subseteq \mathbb{R}^m$ and $\varepsilon > 0$, let A^ε denote the ε -neighbourhood of A . Let $\bar{\Phi}(t) = \gamma_1(t, \infty)$ be the tail cumulative distribution function of the standard Gaussian on the line.

Theorem 134 (Borell, Tsirelson-Ibragimov-Sudakov (1970s)). *Let A be any Borel subset of \mathbb{R}^m with $\gamma_m(A) > 0$ and let H be a half-space in \mathbb{R}^m with $\gamma_m(H) = \gamma_m(A)$. Then $\gamma_m(A^\varepsilon) \geq \gamma_m(H^\varepsilon)$ for all $\varepsilon > 0$. If A is a closed set with $\gamma_m(A) > 0$, then equality holds for some $\varepsilon > 0$ if and only if A is a half-space.*

The isoperimetric inequality implies concentration inequalities for various functions of Gaussian random variables.

Theorem 135. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a $\text{Lip}(\kappa)$ function. Let M_f be a median of f , defined by $\gamma_n\{f \geq M_f\} \geq \frac{1}{2}$ and $\gamma_n\{f \leq M_f\} \geq \frac{1}{2}$. Then, for every $t > 0$, we have*

$$\gamma_n\{f - M_f \geq t\} \leq \bar{\Phi}\left(\frac{t}{\kappa}\right) \leq e^{-\frac{t^2}{2\kappa^2}}, \quad (12)$$

$$\gamma_n\{|f - M_f| \geq t\} \leq 2\bar{\Phi}\left(\frac{t}{\kappa}\right) \leq 2e^{-\frac{t^2}{2\kappa^2}}. \quad (13)$$

Proof. If $A = \{f \leq M_f\}$ then $A^t \subseteq \{f \leq M_f + \kappa t\}$. But $\Phi^{-1}(\gamma_n(A)) \geq 0$ and hence by (??) we get $\Phi^{-1}(\gamma_n(A^t)) \geq t$. Hence $\gamma_n\{f \geq M_f + \kappa t\} \leq \bar{\Phi}(t)$ which shows the first claim. The second follows by adding the same estimate for $\gamma_n\{f \leq M_f - t\}$. ■

Remark 136. Since $\bar{\Phi}(t)$ is strictly smaller than $\frac{1}{2}$ for every $t > 0$, it follows that the median is unique! Some examples of Lipschitz functions of interest are $\max_i x_i$ and $d(\mathbf{x}, A)$ for a fixed closed set A . A smooth function is Lipschitz if and only if its gradient is bounded. It is also useful to observe that if $f = \sup_{i \in I} f_i$ and each f_i is $\text{Lip}(\kappa)$, then so is f .

In many situations, the mean is easier to compute than the median. Here is a simple way to get a (sub-optimal) concentration inequality around the mean for the same setting as above. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a $\text{Lip}(\kappa)$ function and let M_f be its median under γ_n and let $E_f = \int f(\mathbf{x}) d\gamma_n(\mathbf{x})$ be its expectation. Using the bound in Theorem 135 we get

$$\mathbf{E}[(f - M_f)_+] = \int_0^\infty \gamma_n\{f > M_f + t\} dt \leq \int_0^\infty \bar{\Phi}(t/\kappa) dt = \frac{\kappa}{\sqrt{2\pi}}.$$

The same bound holds for $\mathbf{E}[(f - M_f)_-]$ and we get $\mathbf{E}[|f - M_f|] \leq \sqrt{\frac{2}{\pi}} \kappa < \kappa$. In particular, $|E_f - M_f| < \kappa$. Therefore, for $t \geq 2$, we get

$$\gamma_n\{f - E_f > t\kappa\} \leq \gamma_n\left\{f - M_f > \frac{t}{2}\kappa\right\} \leq \bar{\Phi}(t/2),$$

by another application of Theorem 135. For $t \leq 2$, we use the trivial bound $\gamma_n\{f - E_f > t\kappa\} \leq 1$. Hence $C\bar{\Phi}(t/2)$ is a valid bound for all t if we set $C = 1/\bar{\Phi}(1)$. Putting all this together and using the same for deviations below E_f we arrive at the following result.

Theorem 137. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a $\text{Lip}(\kappa)$ function. Let $E_f = \int f d\gamma_n$. Then, for every $t > 0$, we have*

$$\gamma_n\{f - E_f \geq t\} \leq 7\bar{\Phi}\left(\frac{t}{2\kappa}\right) \leq 7e^{-\frac{t^2}{8\kappa^2}}, \quad (14)$$

$$\gamma_n\{|f - E_f| \geq t\} \leq 7\bar{\Phi}\left(\frac{t}{2\kappa}\right) \leq 7e^{-\frac{t^2}{8\kappa^2}}. \quad (15)$$

We simply replaced $\bar{\Phi}(1)$ by the smaller number $1/7$.

Appendix 4: Some combinatorics facts

The Mobius function of a lattice

Let (L, \leq) be a finite partially ordered set. If for every $x, y \in L$, there is a unique least upper bound and a unique greatest lower bound, then we say that L is a lattice. Some of the generalities below apply to any finite poset, but all our applications of interest are to lattices. We define the *Mobius function* of L as the function on $L \times L$ satisfying

$$\mu(a, b) := \begin{cases} 1 & \text{if } a = b \\ -\sum_{a \leq x < b} \mu(a, x) & \text{if } a < b, \\ 0 & \text{if } b < a. \end{cases}$$

Note that this is defined inductively. First we know $\mu(a, a)$. Then we compute $\mu(a, b)$ where a is immediately below b . Then, when there is at most one intermediate element in any chain connecting a to b , etc. The key property of the Mobius function is that

$$\sum_{x: a \leq x \leq b} \mu(a, x) = \begin{cases} 1 & \text{if } a = b, \\ 0 & \text{if } a < b. \end{cases}$$

If μ satisfies $\mu(a, a) = 1$ and the above identities, then it must be the Mobius function.

Lemma 138 (Mobius inversion formula). *Let $f, g : L \mapsto \mathbb{C}$. If $f(b) = \sum_{a \leq b} g(a)$ for all b , then,*

$$g(b) = \sum_{a \leq b} f(a) \mu(a, b) \text{ for all } b.$$

Proof. Let $h(b)$ denote the right hand side of the above equation. For any $b \in L$, we have

$$\sum_{a: a \leq b} h(a) = \sum_{a: a \leq b} \sum_{c: c \leq a} f(c) \mu(c, a) = \sum_{c: c \leq a} f(c) \sum_{a: c \leq a \leq b} \mu(c, a) = f(b).$$

Thus, $\sum_{a: a \leq b} [g(a) - h(a)] = 0$ for all b . Starting with the minimal elements and proceeding upwards inductively, we see that $g = h$. This proves the lemma. ■

Example 139. If $L = \{0, 1, \dots, n\}$ with the order from \mathbb{Z} , then it is easy to work out that $\mu(a, a) = 1$, $\mu(a-1, a) = -1$ and $\mu(a, b) = 0$ for all other (a, b) . Then the lemma just says the obvious thing that if $f(k) = g(0) + \dots + g(k)$, then $g(k) = f(k) - f(k-1)$.

Example 140. Let L be the Boolean lattice consisting of all subsets of $[n]$ with inclusion as the partial order. Then, one can work out that $\mu(A, B) = (-1)^{|B \setminus A|}$. To see this, suppose we have proved it whenever $A \subseteq B$ and $|B \setminus A| \leq k$. Then take $A \subseteq B$ with $|B \setminus A| = k+1$. By definition of the Mobius function,

$$\mu(A, B) = - \sum_{C: A \subseteq C \subsetneq B} (-1)^{|C \setminus A|} = - \sum_{D: D \subsetneq B \setminus A} (-1)^{|D|} = (-1)^{|B \setminus A|}$$

where the last equality follows from the fact that for any set S ,

$$\sum_{D: D \subseteq S} (-1)^{|D|} = \begin{cases} 1 & \text{if } S = \emptyset, \\ 0 & \text{if } S \neq \emptyset. \end{cases}$$

Example 141. Let $L = \mathbb{N}$ with $a \leq b$ if a divides b . Then $\mu(m, n)$ is equal to $(-1)^k$ if n/m is a product of k distinct primes, and equal to zero if n/m has a square factor (or if m does not divide n). Check this inductively as in the previous example.

Usually $\mu(1, n)$ is simply written as $\mu(n)$. The Mobius inversion formula is used frequently in number theory.

Two lattices will be relevant to us in what follows.

The lattice of partitions: Let \mathcal{P}_n denote the set of all set-partitions of $[n]$. The sets that make up a partition are referred to as *blocks*. Note that the order of the blocks, or of the elements in individual blocks are irrelevant.

Example 142. The set \mathcal{P}_3 consists of the five partitions $\{\{1, 2, 3\}\}$, $\{\{1, 2\}, \{3\}\}$, $\{\{1, 3\}, \{2\}\}$, $\{\{2, 3\}, \{1\}\}$ and $\{\{1\}, \{2\}, \{3\}\}$.

For a partition Π we denote the number of blocks by $\ell(\Pi)$ and the individual blocks by Π_j , $1 \leq j \leq \ell(\Pi)$. If we ever need to be more definite, we shall define Π_1 be the block containing 1, and Π_2 to be the block containing the least element not in Π_1 , etc. We shall write $|\Pi_j|$ for the cardinality of the block Π_j .

The most important structure on \mathcal{P}_n is the *partial order* defined by refinement. That is, $\Gamma \leq \Pi$ if every block of Γ is contained inside a block of Π . In this partial order, there is a unique maximal element $\{[n]\}$ and a unique minimal element $\{\{1\}, \dots, \{n\}\}$. Further, this order makes \mathcal{P}_n into a *lattice*, meaning that for any Π, Γ , there is a unique least upper bound (denoted $\Pi \vee \Gamma$) and a unique greatest lower bound (denoted $\Pi \wedge \Gamma$).

Lemma 143. *The Mobius function of \mathcal{P}_n is given by*

$$\mu(\Gamma, \Pi) = \prod_{j=1}^{\ell(\Pi)} (-1)^{\ell(\Gamma_j)-1} (\ell(\Gamma_j) - 1)! \quad \text{for } \Gamma \leq \Pi.$$

Here $\Gamma_j \in \mathcal{P}(\Pi_j)$ is the partition of Π_j induced by Γ . In particular,

$$\mu(\Gamma, \{[n]\}) = (-1)^{\ell(\Gamma)-1} (\ell(\Gamma) - 1)!$$

Proof of Lemma 143. We claim that for any $m \geq 1$,

$$\sum_{\Pi \in \mathcal{P}_m} (-1)^{\ell(\Pi)-1} (\ell(\Pi) - 1)! = 0. \quad (16)$$

Assume this, and fix a pair of partitions $\Gamma \leq \Pi$. Let Π_j be a union of k_j blocks of Γ . Then if $\Gamma \leq \Theta \leq \Pi$, then Θ is naturally identified with a tuple $(\theta^1, \dots, \theta^{\ell(\Pi)})$, where $\theta^i \in \mathcal{P}_{k_i}$, $i \leq \ell(\Pi)$ (by considering how the k_j blocks of Γ inside Π_j are combined in Θ). Therefore,

$$\sum_{\Theta: \Gamma \leq \Theta \leq \Pi} \prod_{j=1}^{\ell(\Pi)} (-1)^{\ell(\Theta^j)-1} (\ell(\Theta^j) - 1)! = \prod_{j=1}^{\ell(\Pi)} \sum_{\theta \in \mathcal{P}_{k_j}} (-1)^{\ell(\theta)-1} (\ell(\theta) - 1)! = 0.$$

This shows that the given formula for μ makes it into the Mobius function. It only remains to prove (16). Check validity for small m . Assuming the result for \mathcal{P}_m , we prove it for \mathcal{P}_{m+1} . For this, observe the following way to build \mathcal{P}_{m+1} from \mathcal{P}_m . Take any Π in \mathcal{P}_m with $\ell = \ell(\Pi)$ and consider the $\ell(\Pi) + 1$ partitions $\Pi^0, \Pi^1, \dots, \Pi^\ell$ of $[m+1]$ where Π^j is got by inserting the element $m+1$ into the j th block of Π and Π^0 by appending $\{m+1\}$ as a singleton block. As Π varies over \mathcal{P}_m , we get each element of \mathcal{P}_{m+1} exactly once.

Now, write $\mu_m(\Pi) = (-1)^{\ell(\Pi)-1} (\ell(\Pi) - 1)!$ for $\Pi \in \mathcal{P}_m$. Then

$$\mu_{m+1}(\Pi^0) = -\mu_m(\Pi)\ell, \quad \mu_m(\Pi^j) = \mu_m(\Pi), \quad 1 \leq j \leq m,$$

whence $\mu_{m+1}(\Pi^0) + \dots + \mu_{m+1}(\Pi^m) = 0$. Thus, $\sum_{\Gamma \in \mathcal{P}_{m+1}} \mu_{m+1}(\Gamma) = 0$. **Did we use the result for m ?** ■

The lattice of non-crossing partitions: Let NC_n denote the set of all non-crossing set-partitions of $[n]$. By this we mean $\Pi \in \mathcal{P}_n$ for which there does not exist and $1 \leq i < j < k < \ell \leq n$ such that i, k are in one block of Π and j, ℓ are in a different block. From \mathcal{P}_n , the refinement order is inherited by NC_n . It is also a lattice.

Note that by definition, $NC_n = \mathcal{P}_n$ for $n \leq 3$. The first non-trivial case is $n = 4$.

Example 144. In \mathcal{P}_4 there is only one crossing partition $\{1, 3\}, \{2, 4\}$. Thus, NC_n consists of all elements of \mathcal{P}_n except this one.

Now recall the Catalan numbers $C_n = \binom{2n}{n} \frac{1}{n+1}$.

Lemma 145. *The Mobius function of NC_n is given by*

$$\mu(\Gamma, \Pi) = \prod_{j=1}^{\ell(\Pi)} (-1)^{\ell(\Gamma_j)-1} C_{\ell(\Gamma_j)-1} \quad \text{for } \Gamma \leq \Pi.$$

In particular,

$$\mu(\Gamma, \{[n]\}) = (-1)^{\ell(\Gamma)-1} C_{\ell(\Gamma)-1}.$$

Proof. As in Lemma 143, it suffices to prove that

$$\sum_{\Gamma \in NC_m} \mu_m(\Gamma) = 0$$

where $\mu_m(\Gamma) = (-1)^{\ell(\Gamma)-1} C_{\ell(\Gamma)-1}$. Here NC_{m+1} is constructed from NC_m as follows. Take any $\Pi \in NC_m$ with $\ell(\Pi) = \ell$ and construct $\Pi^0, \Pi^1, \dots, \Pi^\ell$ by either keeping $\{m+1\}$ as a singleton block (that is Π^0) or by adding $m+1$ to one of the blocks of Π so that the resulting partition remains non-crossing. The latter may be possible for only $\ell' = \ell'(\Pi)$ blocks among the ℓ blocks. Then,

$$\mu_{m+1}(\Pi^0) = -\mu_m(\Pi) \frac{C_\ell}{C_{\ell-1}}, \quad \mu_{m+1}(\Pi^j) = \mu_m(\Pi), 1 \leq j \leq \ell'.$$

Thus

$$\sum_{\Gamma \in \mathcal{P}_{m+1}} \mu_{m+1}(\Gamma) = \sum_{\Pi \in \mathcal{P}_m} \mu_m(\Pi) (\ell'(\Pi) - \ell(\Pi))$$

■