

MIT/WHOI 2007-01

**Massachusetts Institute of Technology
Woods Hole Oceanographic Institution**



**Joint Program
in Oceanography/
Applied Ocean Science
and Engineering**



DOCTORAL DISSERTATION

Applied Stochastic Eigen-Analysis

by

Rajesh Rao Nadakuditi

February 2007

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

MIT/WHOI
2007-01

Applied Stochastic Eigen-Analysis

by

Rajesh Rao Nadakuditi

Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

and

Woods Hole Oceanographic Institution
Woods Hole, Massachusetts 02543

February 2007

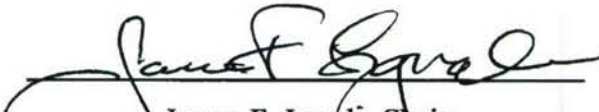
DOCTORAL DISSERTATION

Funding was provided by the National Science Foundation via grant DMS-0411962 and the Office of Naval Research Graduate Traineeship award.

Reproduction in whole or in part is permitted for any purpose of the United States Government. This thesis should be cited as: Rajesh Rao Nadakuditi, 2007. Applied Stochastic Eigen-Analysis. Ph.D. Thesis. MIT/WHOI, 2007-01.

Approved for publication; distribution unlimited.

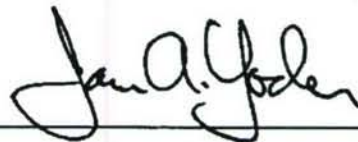
Approved for Distribution:



James F. Lynch, Chair
Department of Applied Ocean Physics and Engineering



Paola Malanotte-Rizzoli
MIT Director of Joint Program



James A. Yoder
WHOI Dean of Graduate Studies

Applied Stochastic Eigen-Analysis

by

Rajesh Rao Nadakuditi

B.S., Electrical Engineering, Lafayette College (1999)

S.M., Electrical Engineering and Computer Science, M.I.T. and W.H.O.I. (2002)

Submitted to the Department of Electrical Engineering and Computer Science
& the Joint Program in Applied Ocean Science and Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

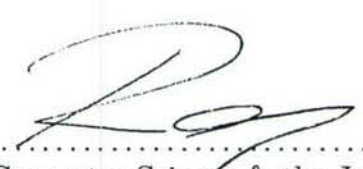
and the

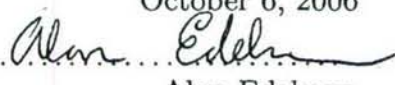
WOODS HOLE OCEANOGRAPHIC INSTITUTION

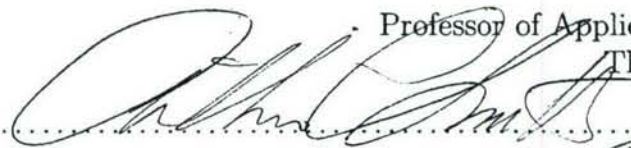
February 2007

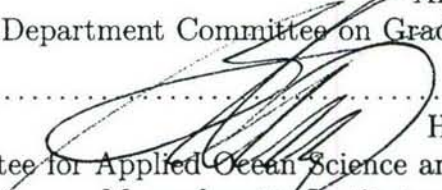
© Rajesh Rao Nadakuditi, MMVII. All rights reserved.

The author hereby grants to MIT and WHOI permission to reproduce and
distribute publicly paper and electronic copies of this thesis document in whole
or in part.

Author 
Department of Electrical Engineering and Computer Science & the Joint
Program in Applied Ocean Science and Engineering
October 6, 2006

Certified by 
Alan Edelman
Professor of Applied Mathematics
Thesis Supervisor

Accepted by 
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Accepted by 
Henrik Schmidt
Chair, Joint Committee for Applied Ocean Science and Engineering
Massachusetts Institute of Technology/
Woods Hole Oceanographic Institution

Applied Stochastic Eigen-Analysis

by

Rajesh Rao Nadakuditi

Submitted to the Department of Electrical Engineering and Computer Science & the
Joint Program in Applied Ocean Science and Engineering
on October 6, 2006, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

The first part of the dissertation investigates the application of the theory of large random matrices to high-dimensional inference problems when the samples are drawn from a multivariate normal distribution. A longstanding problem in sensor array processing is addressed by designing an estimator for the number of signals in white noise that dramatically outperforms that proposed by Wax and Kailath. This methodology is extended to develop new parametric techniques for testing and estimation. Unlike techniques found in the literature, these exhibit robustness to high-dimensionality, sample size constraints and eigenvector misspecification.

By interpreting the eigenvalues of the sample covariance matrix as an interacting particle system, the existence of a phase transition phenomenon in the largest (“signal”) eigenvalue is derived using heuristic arguments. This exposes a fundamental limit on the identifiability of low-level signals due to sample size constraints when using the sample eigenvalues alone.

The analysis is extended to address a problem in sensor array processing, posed by Baggeroer and Cox, on the distribution of the outputs of the Capon-MVDR beamformer when the sample covariance matrix is diagonally loaded.

The second part of the dissertation investigates the limiting distribution of the eigenvalues and eigenvectors of a broader class of random matrices. A powerful method is proposed that expands the reach of the theory beyond the special cases of matrices with Gaussian entries; this simultaneously establishes a framework for computational (non-commutative) “free probability” theory.

The class of “algebraic” random matrices is defined and the generators of this class are specified. Algebraicity of a random matrix sequence is shown to act as a certificate of the computability of the limiting eigenvalue distribution and, for a subclass, the limiting conditional “eigenvector distribution.” The limiting moments of algebraic random matrix sequences, when they exist, are shown to satisfy a finite depth linear recursion so that they may often be efficiently enumerated in closed form. The method is applied to predict the deterioration in the quality of the sample eigenvectors of large algebraic empirical covariance matrices due to sample size constraints.

Thesis Supervisor: Alan Edelman

Title: Professor of Applied Mathematics

Acknowledgments

It feels like it was yesterday when I first walked into Alan Edelman's office to inquire about random matrices. It took just one "dancing" eigenvalues histogram plot for me to get hooked. Working with him has been so stimulating and enjoyable that I look forward, every day, to coming into "work" and interacting with him. By infecting me with his enthusiasm for research he has given me something that I shall spend a lifetime being grateful for - the gift of happiness in doing what I have been doing since I met him. I now aspire to be an academic because of his inspiration.

I thank Alan, his wife Susan, and his two wonderful sons - the perennially, in my mind, "three year old" Jonnie and "five year old" Daniel for welcoming me into their home. I will cherish the memory of innumerable weekend mornings alternating between writing and playing with the kids. And, oh yeah, Susan, I really do like the pizza you make!

Arthur Baggeroer has played a vital role in my research - this thesis sprang to life during his array processing course when he posed a question that exposed the relevance of random matrix theory to practical array processing problems. Despite my acquaintance with him since when I first arrived at MIT, I have only really got to know him these past few years. He is a wonderful person, humble, and yet knows so much. He is open with questions and has taught me the value of looking at things from multiple viewpoints. I hope that I can internalize his example so that I am as generous with my time and advice to others during busier stages of my career.

I thank Jim Preisig for his mentoring from when I first arrived at WHOI and for continually challenging me to be a better researcher.

I owe multiple thanks to Jack Silverstein and Roland Speicher for numerous conversations and email exchanges. Though they were not formally a part of my dissertation process, they played a vital role. It was Jack Silverstein who, at a very early stage of my thesis work, instantly revived my flagging enthusiasm for random matrices by sending me a manuscript of unpublished work. Roland Speicher has been a mentor-by-email in many ways and it has been a special privilege to interact with him. He has been so gracious to me throughout the time I have known him. Just when I was struggling to emulate the clarity of his writing on free probability and random matrices, he let me off the hook by allowing me to reproduce his exposition verbatim - this forms the appendix of this dissertation.

I thank Greg Wornell and Lizhong Zheng for guiding my research as part of my thesis committee, and Sanjoy Mitter for his thought provoking comments on various subjects. Christ Richmond and Steven Smith advised me on research and other matters with a degree of diligence and thoughtfulness that made them committee members in spirit. I am grateful to Henrik Schmdit for graciously agreeing to chair my thesis defense and I wish to acknowledge the influence that Gil Strang's writing style has had on my striving to be a better expositor - this remains a work-in-progress! The work in this thesis has benefited from Richard Stanley pointing me towards the theory of D-finite series and from Pablo Parillo teaching me about resultants.

As I look back on this chapter of my life, I am filled with fond memories of friendships started and/or strengthened with Laura Dipietro, Manish Bhardwaj, Ola Ayaso, Julius Kusuma, Murtaza Zafer, Efe Arkayin, Vikramjit Sethi, Prem Kumar, Abhinav Shukla, Amit Hariyani, Darren Greninger, and Elise Ohayon. I value their friendship more than I know how to properly express. My fondest memory will be of the time spent with Laura.

I thank my academic “siblings,” Ioana Dumitriu, Brian Sutton, Plamen Koev, Per-Olof Persson, and Matthew Harding for many enjoyable conversations. Brett Coonley’s wizardry with L^AT_EX has made my tables and figures look so pretty (Figure 3-2, for example) that I feel pressure to make the written word worthy of it!

I thank Shirley Entzminger and Rhonda Culbert of MIT Mathematics, Ronni Schwartz of the MIT/WHOI Joint Program, Marilyn Pierce of MIT EECS, Julia Westwater and Marsha Gomes of WHOI, and Scott Perrigo of MIT Copytech for their assistance in helping me navigate administrative and other matters. I have greatly enjoyed my affiliation with WHOI and I attribute this in large part to the favorable first impression that Hanu Singh helped form during my time working with him as a summer student fellow.

I am grateful to the National Science Foundation for supporting this work via grant DMS-0411962 and the Office of Naval Research Graduate Traineeship award which supported my research at an earlier stage.

Last, but far from least, I wish to give thanks to my family. I would like to thank my father, my dearly departed mother, grandparents, uncles, aunts, cousins and external agents, seen and unseen, for their support and constant encouragement. In particular, I would like to single out my grandfather, a career scientist, for thanks. He piqued my interest in research at an early stage with fascinating anecdotes from his workplace. This thesis is dedicated to his memory.

Contents

1	Foreword	17
2	Introduction	19
2.1	Role of sample covariance matrices in signal processing	19
2.1.1	Statistical inference from random sample covariance matrices	20
2.1.2	Algorithmic performance measures	21
2.1.3	Impact of sample size constraints	22
2.1.4	Diagonal loading and subspace techniques	23
2.1.5	High-dimensional statistical inference	24
2.2	Random matrix theory	25
2.2.1	Beyond special cases: Free probability theory	26
2.2.2	Algebraic random matrices and a random matrix “calculator”	27
2.3	Contributions of this thesis	30
3	Statistical eigen-inference: Signals in white noise	33
3.1	Introduction	33
3.2	Problem formulation	34
3.3	Eigenvalues of the (null) Wishart matrix	37
3.3.1	An interacting particle system interpretation	37
3.3.2	Sample eigenvalues in the snapshot abundant regime	39
3.3.3	Limiting distribution of the sample eigenvalues	40
3.3.4	Gaussian fluctuations of sample eigenvalues	41
3.4	Signals in white noise	42
3.4.1	Interacting particle system interpretation	43
3.4.2	Repulsion and phase transition of largest (“signal”) eigenvalue	44
3.4.3	Gaussian fluctuations of largest (“signal”) eigenvalues	46
3.5	Estimating the number of signals	47
3.5.1	Akaike’s Information Criterion	47
3.5.2	Unknown noise variance	47
3.5.3	Known noise variance	48
3.6	Extensions to frequency domain and vector sensor arrays	49

3.7	Effective number of signals and consistency of the criteria	52
3.7.1	The asymptotic identifiability of two closely spaced signals . . .	53
3.8	Numerical simulations	54
3.8.1	Test of sphericity	54
3.8.2	Illustration of effective number of signals concept	54
3.9	Future work	55
4	Statistical eigen-inference: Large Wishart matrices	61
4.1	Problem formulation	61
4.1.1	Inferring the population eigenvalues from the sample eigenvalues	62
4.1.2	Testing for equality of population eigenvalues	64
4.1.3	Proposed statistical eigen-inference techniques	65
4.2	Preliminaries	67
4.3	Computational aspects	69
4.3.1	Computation of moments of limiting eigenvalue distribution . . .	69
4.3.2	Computation of covariance moments of second order limit distri- bution	71
4.4	Eigen-inference algorithms	73
4.4.1	Estimating θ for known model order	73
4.4.2	Testing $\theta = \theta_0$	75
4.4.3	Estimating θ and testing $\theta = \hat{\theta}$	76
4.4.4	Estimating θ for unknown model order	76
4.5	Numerical simulations	76
4.6	Inferential aspects of spiked covariance matrix models	81
4.6.1	Impact of the sample eigenvalue phase transition phenomenon .	82
4.7	Future work	83
5	The Capon beamformer: Approximating the output distribution	89
5.1	Introduction	89
5.2	Problem formulation	91
5.3	The Capon-Goodman result	94
5.3.1	Structure exploited in Capon and Goodman's analysis	95
5.4	Relevant result from random matrix theory	96
5.5	Distribution under diagonal loading: No sources	98
5.6	Distribution under diagonal loading: Single source	98
5.7	Distribution under diagonal loading: Two sources	101
5.8	Numerical examples	104
5.9	Future work	105
6	The polynomial method: Mathematical foundation	109
6.1	Transform representations	109
6.1.1	The Stieltjes transform and some minor variations	109
6.1.2	The moment transform	110

6.1.3	The R transform	111
6.1.4	The S transform	112
6.2	The Algebraic Framework	113
6.3	Algebraic manipulations of algebraic functions	120
6.4	Algebraic manipulations using the resultant	122
7	The polynomial method: Algebraic random matrices	127
7.1	Motivation	127
7.2	Definitions	128
7.3	Deterministic operations on algebraic random matrices	129
7.4	Gaussian-like matrix operations on algebraic random matrices	133
7.5	Sums and products of algebraic random matrices	137
8	The polynomial method: Computational aspects	139
8.1	Operational laws on bivariate polynomials	139
8.2	Interpreting the solution curves of polynomial equations	143
8.2.1	The atomic component	143
8.2.2	The non-atomic component	144
8.3	Enumerating the moments and free cumulants	147
8.4	Computational free probability	149
8.4.1	Implicitly encoding the free convolution computations	150
9	The polynomial method: Applications	153
9.1	The Jacobi random matrix	153
9.2	Random compression of a matrix	156
9.3	Free additive convolution of equilibrium measures	158
9.4	Other applications	159
10	The polynomial method: Eigenvectors of random matrices	161
10.1	The conditional “eigenvector distribution”	161
10.2	Algebraic conditional “eigenvector distributions”	163
10.3	Example	168
10.4	Algebraic empirical covariance matrices	175
10.4.1	Example	176
10.5	Future work	178
11	Afterword	181
A	Random Matrices and Free Probability Theory	183
A.1	Moments of random matrices and asymptotic freeness	183
A.2	Fluctuations of random matrices and asymptotic second order freeness	186
A.3	Wishart matrices and Proof of Proposition 4.27	190

List of Figures

2-1 Typical performance plot for the mean square estimation error of a signal processing algorithm. Three broad regimes that characterize the performance relative to the signal level are identified. Model mismatch refers to actual physical model differing from the assumed model. This behavior has been well documented by many authors including [115, 116]. 22

2-2 A random matrix calculator where a sequence of deterministic and stochastic operations performed on an “algebraically characterizable” random matrix sequence \mathbf{A}_N produces a “algebraically characterizable” random matrix sequence \mathbf{B}_N . The limiting eigenvalue density and moments of a “characterizable” matrix can be computed numerically, with the latter often in closed form. 28

2-3 A representative computation using the random matrix calculator. . . . 31

3-1 Blurring of sample eigenvalues due to finite number of snapshots. . . . 36

3-2 Interacting particle system interpretation of the eigenvalues of a (null) Wishart matrix. 38

3-3 Limiting distribution of the eigenvalues of the (null) Wishart matrix. . . 40

3-4 Interactions between the “signal” eigenvalue and the “noise” eigenvalues. 44

3-5 Sample realizations of the proposed criterion when $\lambda_1 = 10$, $\lambda_2 = 3$ and $\lambda_3 = \dots = \lambda_n = 1$ 49

3-6 “Signal” and “noise” eigenvalue fluctuations induced limits on signal identifiability. 55

3-7 Performance of proposed estimators when there are zero signals in white noise. Note that $10^{-0.01} \approx 0.9772$ while $10^{-0.04} \approx 0.9120$ 57

3-8 Performance of proposed estimators when there are 2 signals in white noise. 58

3-9 Comparison of the known vs. unknown estimators. 59

3-10 Comparison of the known vs. unknown estimators. 60

4-1 The challenge of estimating the population eigenvalues from the sample eigenvalues in high-dimensional settings. 66

4-2	Numerical simulations (when \mathbf{S} is complex) illustrating the robustness of test statistics formed with $\dim(\mathbf{v}) = 2$ to moderate dimensional settings.	74
4-3	Normal probability plots of the estimates of a and t (true values: $a = 2$, $t = 0.5$).	80
4-4	Normal probability plots of the spiked magnitude estimate (true value = 10).	85
4-5	Sphericity test.	86
5-1	Configuration of a uniform linear array with n sensors and inter-elements spacing of d units.	92
5-2	The Capon-MVDR spectral estimate when $n = 18$, $m = 36$.	93
5-3	Assessing the validity of the approximated distribution for $\hat{\psi}(\theta, \delta)$ when there are no sources in white noise.	99
5-4	Two equal power sources at 90 degrees and 70 degrees.	106
5-5	Two equal power sources at 90 degrees and 70 degrees.	107
5-6	Two Equal Power sources at 90 degrees and 95 degrees.	108
6-1	The six interconnected bivariate polynomials; transformations between the polynomials, indicated by the labelled arrows, are given in Table 6.3.	116
8-1	The real and imaginary components of the algebraic curve defined by the equation $L_{\text{mz}}(m, z) = 0$, where $L_{\text{mz}} \equiv czm^2 - (1 - c - z)m + 1$, which encodes the Marčenko-Pastur density. The curve is plotted for $c = 2$. The $-1/z$ behavior of the real part of the “correct solution” as $z \rightarrow \infty$ is the generic behavior exhibited by the real part of the Stieltjes transform of a valid probability density function.	146
9-1	The limiting density (solid line), $f_{A_6}(x)$, given by (9.7) with $c_1 = 0.1$ and $c_2 = 0.625$ is compared with the normalized histogram of the eigenvalues of a Jacobi matrix generated using the code in Table 9.1 over 4000 Monte-Carlo trials with $n = 100$, $N_1 = n/c_1 = 1000$ and $N_2 = n/c_2 = 160$.	155
9-2	The limiting eigenvalue density function (solid line) of the top $0.4N \times 0.4N$ block of a randomly rotated matrix is compared with the experimental histogram collected over 4000 trials with $N = 200$. Half of the eigenvalues of the original matrix were equal to one while the remainder were equal to zero.	157
9-3	Additive convolution of equilibrium measures corresponding to potentials $V_1(x)$ and $V_2(x)$.	158
10-1	Limiting eigenvalue density function of $\mathbf{A}_N + e\mathbf{W}_N$ for different values of e .	170
10-2	The composite limit eigenvalue density function is interpreted as the sum of the scaled individual kernel density functions.	172

10-3	Characterizing the relationship between the eigenvectors of $\mathbf{A} + 2\mathbf{W}$ with those of \mathbf{A}	173
10-4	The Markov transition kernel density $k_{A+B B}$ where $\mathbf{B} = 2\mathbf{W}$	174
10-5	A Wishart random matrix, \mathbf{S}_n , with spatio-temporal correlations. The spatial and the temporal covariance matrices have limiting eigenvalue distribution given by (10.28).	179
10-6	A Wishart random matrix, \mathbf{S}_n , with spatio-temporal correlations. The spatial and the temporal covariance matrices have limiting eigenvalue distribution given by (10.28).	180

List of Tables

3.1	Estimating the number of signals from the eigenvalues of a SCM formed from m snapshots.	51
4.1	Structure of proposed algorithms.	67
4.2	Relationship between the coefficients $\alpha_{i,j} = \alpha_{j,i}$ and α_i	72
4.3	Quality of estimation of $t = 0.5$ for different values of n (dimension of observation vector) and m (number of samples) – both real and complex case.	78
4.4	Quality of estimation of $a = 2$ for different values of n (dimension of observation vector) and m (number of samples) – both real and complex case.	79
4.5	Algorithm performance for different values of n (dimension of observation vector) and m (number of samples) – both real and complex case.	84
4.6	The null-hypothesis is accepted at the 95% significance level for χ_2^2 or whenever $h(\boldsymbol{\theta}) \leq 5.9914$	87
6.1	Making L_{uv} irreducible.	117
6.2	Bivariate polynomial representations of some algebraic distributions.	118
6.3	Transformations between the different bivariate polynomials. As a guide to MATLAB notation, the command <code>syms</code> declares a variable to be symbolic while the command <code>subs</code> symbolically substitutes every occurrence of the second argument in the first argument with the third argument. Thus, for example, the command <code>y=subs(x-a,a,10)</code> will yield the output <code>y=x-10</code> if we have previously declared x and a to be symbolic using the command <code>syms x a</code>	119
6.4	The companion matrix \mathbf{C}_{uv}^u , with respect to u , of the bivariate polynomial L_{uv} given by (6.32).	121
6.5	Formal and computational description of the \boxplus_u and \boxtimes_u operators acting on the bivariate polynomials $L_{uv}^1(u, v)$ and $L_{uv}^2(u, v)$ where \mathbf{C}_{uv}^{u1} and \mathbf{C}_{uv}^{u2} are their corresponding companion matrices constructed as in Table 6.4 and \otimes is the matrix Kronecker product.	122

6.6	Examples of \boxplus and \boxtimes operations on a pair of bivariate polynomials, L_{uv}^1 and L_{uv}^2	123
8.1	Operational laws on the bivariate polynomial encodings (and their computational realization in MATLAB) corresponding to a class of deterministic and stochastic transformations. The Gaussian-like random matrix \mathbf{G} is an $N \times L$, the Wishart-like matrix $\mathbf{W}(c) = \mathbf{G} \mathbf{G}'$ where $N/L \rightarrow c > 0$ as $N, L \rightarrow \infty$, and the matrix \mathbf{T} is a diagonal atomic random matrix.	140
8.2	Operational laws on the bivariate polynomial encodings for some deterministic random matrix transformations. The operations \boxplus_u and \boxtimes_u are defined in Table 6.5.	141
8.3	Operational laws on the bivariate polynomial encodings for some canonical random matrix transformations. The operations \boxplus_u and \boxtimes_u are defined in Table 6.5.	142
8.4	Implicit representation of the free convolution of two algebraic probability densities.	150
9.1	Sequence of MATLAB commands for sampling the Jacobi ensemble. The functions used to generate the corresponding bivariate polynomials symbolically are listed in Table 8.1	153
9.2	Parameters for determining the limiting eigenvalue density function using (9.7).	155
10.1	Symbolic code in MATLAB for computing the trivariate polynomial that encodes the Markov transition kernel that characterizes the conditional “eigenvector distribution” of sums and products of algebraic random matrices.	167
10.2	Symbolic code for computing the bivariate and trivariate polynomials which, respectively, encode the limiting conditional eigenvector and eigenvalue distribution of algebraic empirical covariance matrices.	176

This thesis applies random matrix theory to high-dimensional inference problems. For measurements modeled as

$$\mathbf{x}_i = \mathbf{A} \mathbf{s}_i + \mathbf{z}_i,$$

the term “high-dimensional” refers to the dimension n of the vector \mathbf{x}_i . In array processing applications such as radar and sonar, where the elements of \mathbf{x}_i are interpreted as spatial observations, n can range from ten up to a few thousand. The elements of \mathbf{x}_i are modeled as random variables and have different interpretations in different applications but the core problem can be succinctly summarized:

How does one use m samples (measurements), $\mathbf{x}_1, \dots, \mathbf{x}_m$ to estimate, as accurately as possible, the n -by- k matrix \mathbf{A} or the k -by-1 vectors $\mathbf{s}_1, \dots, \mathbf{s}_m$, or both, in the presence of random noise \mathbf{z}_i ?

In array processing applications such as radar and sonar, accurate estimation of the matrix \mathbf{A} leads to a commensurately accurate estimation of the location of an airplane. In an application, referred to as the “cocktail party problem,” [44] a sensor array is used to estimate \mathbf{A} and hence the positions of persons speaking in a room; this information is then used to isolate the voices of the various speakers.

Variations of this setup abound in applications such as time-series analysis, wireless communications, econometrics, geophysics, and many more. Consequently, this problem has been formulated, and “solved” by many research communities. Almost all the traditional solutions assume, however, that there are enough data samples available, relative to the number of sensors, so that an accurate statistical characterization can be performed on the measured data. When the number of sensors is relatively small (less than 8) this assumption is reasonable. However, as we keep adding sensors, this assumption is violated so that traditional algorithms perform considerably worse than expected.

This curse of high-dimensionality seemingly contradicts our expectation (hope, really) that adding more sensors translates into improved performance. Taking more samples is often not an option because of the time-varying nature of the problem (*e.g.*, tracking an airplane). Thus, devising techniques to counteract this effect will have a positive impact on many areas. This is where random matrices become relevant.

The n -by- m matrix obtained by stacking the m measurements made at n sensors alongside each other is a “random matrix” because its elements are random variables. As the dimensions of the matrix get large a remarkable phenomenon occurs; the behavior of a large class of random matrices becomes increasingly non-random in a manner than can be predicted analytically. In fact the larger the size of the matrix, the lesser the unpredictability, *i.e.*, the magnitude of the “fluctuation.”

This observation is the starting point for our research. The hypothesis we explore in this thesis is that high-dimensionality in such settings might just be a *blessing* – provided the underlying model is physically justifiable and the non-random part can be concretely predicted and taken advantage of. One of the main contributions of this thesis is the development of a framework based on this philosophy (see Chapters 6 - 10) to design implementable estimation and hypothesis testing algorithms (see Chapters 3 and 4) for physically motivated random matrix models.

In a setting of interest to many research communities, we are able to characterize the fundamental limit of these techniques – signals can be reliably detected using these techniques only if their power is above a threshold that is a simple function of the noise power, the dimensionality of the system and the number of samples available (see Chapter 3).

Along the way, we unlock the power of “free probability” – the mathematical theory that reveals the hidden structure lurking behind these high-dimensional objects. While in the past the non-random behavior for large random matrices could only be predicted for some special cases, the computational tools we develop ensure that concrete predictions can be made for a much broader class of matrices than thought possible. The tools reveal the full power of the theory in predicting the global behavior of the eigenvalues and eigenvectors of large random matrices (see Chapters 6 - 10).

The statistical techniques developed merely scratch the surface of this theory – our hope in presenting the software version, of the “free probability” or random matrix “calculator,” alongside the mathematics that facilitates the computational realization is that readers will take the code as a starting point for their own experimentation and develop additional applications of the theory on which our ideas are based. Readers interested in this latter framework may proceed directly to Chapter 6 and skip the preceding application oriented material.

We provide an overview of the sample covariance matrix based inference problems in signal processing in Chapter 2. Our point of departure will be inference problems in sensor array processing. Practitioners in other areas of science and engineering should easily be able to adapt the proposed techniques to their applications.

Introduction

The statistical theory of signal processing evolved in the 1930's and 1940's, spurred in large part by the successful consummation of mathematical theory and engineering practice [112]. Correlation techniques for time series analysis played a key role in the mathematics developed at the time by Wiener and colleagues. To quote Professor Henry J. Zimmerman, the Director of MIT's Research Lab of Electronics, (*italics added for emphasis*)

“... the potential significance of *correlation techniques* had fired the imagination ... the general enthusiasm was due to experimental evidence ... that *weak signals could be recovered in the presence of noise using correlation techniques*. From that point on the field evolved very rapidly [119].”

Covariance matrix based methods were the natural extension of correlation techniques to multi-channel signal processing algorithms and remain widely used to this day [102]. Array processing applications involving radar and sonar were amongst the first to use such techniques for tasks involving as detection, estimation, and classification. Representative applications include detecting airplanes, estimating environmental parameters using an array of sensors, and classifying objects based on surface reflections received at a sensor bank.

■ 2.1 Role of sample covariance matrices in signal processing

Typically, since the true covariance matrix is unknown, a sample covariance matrix is used. Hence, many modern multichannel signal processing algorithms used in practice can be labelled as sample covariance matrix based. The role of random matrices enters because of the statistical characterization of the sample covariance formed by summing over the outer products of the m observation (or “snapshot”) vectors $\mathbf{x}_1, \dots, \mathbf{x}_m$ when forming the $n \times n$ sample covariance matrix $\hat{\mathbf{R}}$ as

$$\hat{\mathbf{R}} = \frac{1}{m} \mathbf{X} \mathbf{X}', \quad (2.1)$$

where the $'$ denotes the Hermitian transpose and the data matrix $\mathbf{X} = [\mathbf{x}_1 | \dots | \mathbf{x}_m]$ is an $n \times m$ matrix whose rows represent measurements made at the sensors (spatial

and/or frequency samples) and columns represent snapshots (time samples). When the snapshots are modelled as a multivariate Gaussian with covariance \mathbf{R} , then the random sample covariance matrix in (2.1) is an instance of the Wishart matrix [114] extensively studied by statisticians.

■ 2.1.1 Statistical inference from random sample covariance matrices

Inference techniques posed on sample covariance matrices with the Wishart distribution include algorithms for testing hypothesis (*e.g.*, is there an airplane present?) and estimating values of parameters (*e.g.*, where is the airplane located?). In terms of the sample covariance matrix, these algorithms can be classified as either exploiting the *eigenvector* structure of the (assumed) true covariance matrix or the *eigenvalue* structure. When the physics of the operating environment are adequately modelled, a maximum likelihood technique can be used to estimate the unknown parameters. When algorithmic computational complexity is a factor, estimation in these settings is often reduced to the computation of a weight vector, which is given (up to a scale factor) by

$$\mathbf{w} \propto \hat{\mathbf{R}}^{-1} \mathbf{v}, \quad (2.2)$$

where \mathbf{v} is termed a replica, or matching signal vector or a spatial matched filter.

In recursive methods this weight is computed dynamically as data is accumulated with a “forgetting” factor which decreases the influence of older data; for example, using recursive least squares algorithms. Regardless of the method, the underlying problem in statistical signal processing [7] is that the non-stationarity and/or inhomogeneity of the data limits the number of samples m , which can be used to form the sample covariance matrix $\hat{\mathbf{R}}$. This non-stationarity/inhomogeneity can be caused by the motion of ships, aircraft, satellites, geophysical and/or oceanographic processes and regions; in other words, it is often inherent to the operating environment and cannot be “designed away.”

Examining the weight vector computation in (2.2) more carefully reveals why we label such techniques as exploiting eigenvector information. The weight vector can be written in terms of the projection of the replica vector \mathbf{v} onto the sample eigenspace as

$$\mathbf{w} \propto \sum_{i=1}^n \frac{1}{\hat{\lambda}_i} \hat{\mathbf{u}}_i \langle \hat{\mathbf{u}}_i, \mathbf{v} \rangle. \quad (2.3)$$

Clearly, as the expression in (2.3) indicates, the computation of the weight vector is directly affected by the projection $\langle \hat{\mathbf{u}}_i, \mathbf{v} \rangle$ of the signal bearing vector \mathbf{v} onto the sample eigenvectors $\hat{\mathbf{u}}_i$.

There is a whole class of statistical inference problems involving detection and estimation that do not rely on eigenvector information. Here, inference is performed directly on the eigenvalues of the sample covariance matrix. Examples include the ubiquitous (in signal processing applications) Wax-Kailath estimator [111] for the number of signals in white noise and Anderson’s tests and estimators developed in his landmark paper on principal components analysis [6].

These “classical” techniques exploit the large sample size asymptotics of the eigenvalues of Wishart distributed sample covariance matrices. However, the asymptotics developed are not adequate for high-dimensional settings so that they only work well when $m \gg n$.

An early motivation for such eigenvalue based inference techniques (“eigen-inference”) was the computational savings relative to maximum likelihood methods that incorporate information about the eigenvectors of the underlying covariance matrix. The inference methodologies developed in this thesis fall in the category of posing the estimation and hypothesis testing problem on the sample eigenvalues alone. Since we discard the information in the eigenvectors, this necessarily compromises their performance relative to algorithms that use high-quality parametric models for the eigenvectors. Conversely, this provides the justification for “robustifying” eigenvector dependent inferential algorithms when the models for the eigenvectors are of uncertain quality which is often the case in high-dimensional settings.

■ 2.1.2 Algorithmic performance measures

Regardless of whether the estimators exploit eigenvector information or not, for signal processing applications involving parameter estimation, the mean square estimation error is a commonly used performance metric. A common practice is to compare the simulated performance of an algorithm with the Cramer-Rao lower bound [25, 72] since the latter is the theoretically optimal performance achievable by the best possible (unbiased) estimator. Figure 2-1 shows a typical mean square error performance plot for an estimator as a function of the level of the signal (or parameter) that is being estimated. For asymptotically large signal levels, the performance of most algorithms matches the Cramer-Rao lower bound unless there is a saturation in performance because of a model misspecification [115, 116]. The latter issue motivates the bulk of the algorithms developed in this thesis - we shall return to it shortly.

There are three regimes in plots such as Figure 2-1 that need to be distinguished. The performance loss of an algorithm, as shown in Figure 2-1, is measured with respect to the difference between the achieved mean square estimation error and the Cramer-Rao lower bound.

In Figure 2-1, Regime III is referred to as the *asymptotic regime* and is characterized by an approximately linear (on appropriate rescaling) behavior. Regime II is characterized by a rapid, highly nonlinear deterioration in the performance of the algorithm - a phase transition, as it were. This breakdown has been observed and studied by several authors [43, 52, 60, 71, 100, 101, 116] and is referred to as a *threshold effect* [103]. The exact signal level where it occurs depends on the algorithm in question and can be complicated to compute [74, 118]. Different signal processing algorithms behave differently in this regime - some suffer more gradual deterioration in performance than others. The onset of this regime is characterized by ambiguous sidelobes in the parameter space accompanied by a deterioration in the reliability of the sample eigenvectors [43, 101, 103]. Regime I, sometimes referred to as the *no information regime*, occurs when the sidelobes




Figure 2-1. Typical performance plot for the mean square estimation error of a signal processing algorithm. Three broad regimes that characterize the performance relative to the signal level are identified. Model mismatch refers to actual physical model differing from the assumed model. This behavior has been well documented by many authors including [115,116].

in the parameter space lead to highly biased estimates leading to unacceptably high estimation error. Regimes II and III are thus of interest when designing algorithms.

The utility of such plots in practice comes from their being able to indicate how well the chosen algorithm is able to detect and estimate low level signals. When there are sample size constraints, there is a deterioration in performance; this is referred to in array processing literature as the “snapshot problem” which we discuss next.

■ 2.1.3 Impact of sample size constraints

For multi-channel signal processing algorithms the realized performance and the Cramer-Rao lower bound are, roughly speaking, a function of the true covariance matrix (which encodes the signal level and the number of sensors n) and the number of snapshots m [90]. It is quite natural to characterize the performance as a function of the ratio n/m of the number of sensors to the number of snapshots. For a fixed number of sensors n the performance of the algorithm for a chosen signal level improves as the number of snapshots increases. In other words, for a fixed n as $m \rightarrow \infty$, the ratio $n/m \rightarrow 0$ and the performance improves.

For array processing applications, there are two well-known results that capture this analytically - however, only in the scenario when the sample support exceeds the number of sensors, *i.e.*, $m > n$. The first, known as the Capon-Goodman result [20], states that the energy, or mean square value, of the projection of the weight vector \mathbf{w} for the so-

called Capon algorithm [19] onto a data vector has complex chi-squared distribution with $m - n + 1$ degrees of freedom where, as usual, m is the number of snapshots and n is number of sensors. The second, under the same conditions characterizes the signal to noise ratio, where the signal is the replica \mathbf{v} scaled by its amplitude and the noise level is the power obtained using the Capon-Goodman result. The results states that the power has beta distribution leading to often quoted Reed-Brennan-Mallett result [18] that one needs the ratio n/m between 1/2 to 1/3 for obtaining “adequate” signal to noise ratio.

There are many current applications where meeting this sample support requirement is just not possible. In large array processing applications in radar and sonar, the ratio of the number of sensors to snapshots is around between 2 – 100 though often it is much more. An immediate consequence of this for estimation problems posed in terms of the eigenvectors is that the weight vector cannot be computed using (2.2) since the sample covariance matrix $\hat{\mathbf{R}}$ given by (2.1) is singular. Practical strategies have been in place since the mid 1960’s to overcome this.

■ 2.1.4 Diagonal loading and subspace techniques

Several fields have developed a number of fixes to the problems that arise when $\hat{\mathbf{R}}$ is singular. Two methods dominate the approaches. In the first, the sample covariance matrix is “diagonally loaded” [24]. The sample covariance matrix $\hat{\mathbf{R}}$ in (2.2) is replaced with a diagonally loaded $\hat{\mathbf{R}}_\delta$ given by

$$\hat{\mathbf{R}}_\delta = \hat{\mathbf{R}} + \delta \mathbf{I}, \quad (2.4)$$

so that the weight vector computation in (2.3) becomes

$$\mathbf{w}_{\text{DL}}(\delta) \propto \sum_{i=1}^n \frac{1}{\hat{\lambda}_i + \delta} \hat{\mathbf{u}}_i \langle \hat{\mathbf{u}}_i, \mathbf{v} \rangle. \quad (2.5)$$

Sometimes this is termed “ridge regression” [47], “regularization,” “shrinkage parameter” [35] or “white noise gain control” [102]. It appears that every application has its own vocabulary. This approach has the impact of putting a “floor” on the low eigenvalues so when the inverse is taken, they do not dominate the solution. The choice of the diagonal loading or regularization parameter δ is an important factor that affects the statistical robustness and the sensitivity of the underlying algorithm.

The second approach is based on subspaces. Most often, the sample eigenvalues, $\hat{\lambda}_i$ for $i = 1, 2, \dots, n$ are ordered and only those above a threshold are used. This is termed dominant mode rejection (DMR) [69]. The processing is then done on the remaining subspace, either with each sample eigenvector/eigenvalue or with an appropriate transformation to reduce the signal dimensionality. The issue here is to establish the signal subspace; once done, most of the existing algorithms can be used. Some of these approaches are discussed in Van Trees [102]. In other words, the resulting weight vector

computation can be expressed as:

$$\mathbf{w}_{\text{SS}}(y) \propto \sum_{i=1}^n \left(g(\hat{\lambda}_i) \chi_y(\hat{\lambda}_i) \right) \hat{\mathbf{u}}_i \langle \hat{\mathbf{u}}_i, \mathbf{v} \rangle, \quad (2.6)$$

where $g(\cdot)$ is an appropriate function that depends on the algorithm used, and $\chi_y(\hat{\lambda}_i)$ is the indicator function that is equal to 1 when $\hat{\lambda}_i \geq y > 0$ and 0 otherwise.

Other variants of these processing techniques are also found in practice and in the literature. Van Trees [102] and Scharf [78] are good references on this and other related subjects. A common thread in all of these techniques, that can be discerned from the expressions in (2.5) and (2.6), is that they essentially represent different schemes for weighting the contribution of the individual sample eigenvectors in the computation of the weight vector \mathbf{w} . Such fixes affect the performance of eigenvector based parameter estimation algorithms.

Analyzing the distribution of the outputs when a diagonally loaded SCM is used for the computing the weight vector is the first step in analyzing its impact on performance. Lack on analytical results in this direction has been an outstanding problem for a while in the community [7]. In Chapter 5 we provide the first such analytical results for the Capon beamformer under diagonal loading. For inference problems posed using the eigenvalues alone, diagonal loading and other schemes are of little use since they modify the eigenvalues in a predictable manner. Hence, practitioners continue to use the Wax-Kailath estimator [111] and the algorithms proposed by Anderson [6] for eigenvalue based inference even though are clearly inadequate in high-dimensional, sample size constrained settings found in an increasing number of applications.

■ 2.1.5 High-dimensional statistical inference

There are many existing applications that already operate in a severely sample size constrained regime. Currently, engineers and scientists deploy array processing systems with a very large number of sensors. Arrays with up to 6000 geophones are now used in geophysical exploration for oil; US Navy towed arrays now have 100 to 1000 sensors, phased array radars have 100's of dipoles or turned helices. The current state of arrays now stretches the snapshot support and future ones certainly will only exacerbate the situation further.

In adaptive array processing applications, we are already in a situation where sample covariance matrix based estimators that rely on Anderson's eigen-analysis for $m \gg n$ perform inadequately. In situations with sample size constraints where the model is misspecified to the extent that eigenvector information is no longer reliable, we often witness performance saturation as depicted in Figure 2-1. It is important to develop sample eigenvalue based inference algorithms that supplant the methodologies proposed in [6] and [111]. One of the contributions of this thesis (see Chapters 3-4) is the development of such algorithms that are robust to high-dimensionality, sample size constraints and eigenvector misspecification.

We note that many practically important estimation problems can only be formulated in terms of the sample eigenvectors; in such cases the parameter of interest (*e.g.*, location of the airplane) resides in the “noisy” sample eigenvector. In such applications, parameter estimation is posed as a joint detection-estimation problem – the number of signals is first estimated so that the parameters associated with the signals can then be extracted. The former problem can be posed in the eigenvalue domain only and the techniques we have developed (see Chapter 3) will outperform other techniques in the literature. They will not, however, help improve the subsequent parameter estimation problem which will still be impacted by the sample eigenvector degradation due to sample size constraints. However, we note that in this thesis we have made noteworthy progress in this direction. We have developed the first computational framework for analyzing eigenvectors of a large class of random matrices (see Chapter 10) including those with Wishart distribution. We believe that this should pave the way for the development of new eigenvector based inference methodologies that are similar robust to high-dimensionality and eigenvector misspecification.

■ 2.2 Random matrix theory

It is worth emphasizing the nature of the stochastic eigen-analysis results being exploited. Finite random matrix theory, of the sort found in references such as Muirhead [66] is concerned with obtaining *exact* characterizations, for every n and m , of the distribution of the eigenvalues of random sample covariance matrices. Consequently, the finite random matrix theory results often focus on the Wishart distribution.

Infinite random matrix theory, on the other hand, is concerned with the characterization of the *limiting* distribution of the eigenvalues of random matrices. By posing the question in the asymptotic (with respect to n) regime concrete answers can be obtained for a much larger class of random matrices than can be handled by finite random matrix theory.

The theory of large random matrices arises naturally because the inference problems we are interested in are inherently high-dimensional. In that regard, a central object in the study of large random matrices is the empirical distribution function which is defined, for an $N \times N$ matrix \mathbf{A}_N with real eigenvalues, as

$$F^{\mathbf{A}_N}(x) = \frac{\text{Number of eigenvalues of } \mathbf{A}_N \leq x}{N}. \quad (2.7)$$

For a large class of random matrices, the empirical distribution function $F^{\mathbf{A}_N}(x)$ converges, for every x , almost surely (or in probability) as $N \rightarrow \infty$ to a non-random distribution function $F^A(x)$. In practice, $N \approx 8$ is “good enough” in the sense that the empirical histogram of the eigenvalues will very well approximate the distributional derivative of the limiting distribution function. The early literature on this subject used matrix theoretic arguments to determine the class of random matrices for which the limiting eigenvalue distribution could be determined. The techniques first used by

Marčenko-Pastur [59] and later perfected by Silverstein [83] and Girko [40] formed the foundation of these investigations (see [8] for a comprehensive review). Despite being able to characterize a rather broad class of practically applicable random matrices, the derivations had to be done on a case-by-case basis so that it was not clear if there was deeper structure in the random matrices that could permit “universal” computation.

■ 2.2.1 Beyond special cases: Free probability theory

The development of “free probability” by Voiculescu in the mid-1980’s changed all that by pinpointing the structure behind these high-dimensional objects that permits computation. Free probability has since emerged as a counterpart to “classical” probability. Some good references are [16, 46, 68, 109]. These references and even the name “free probability” are worthy of some introduction.

We begin with a viewpoint on classical probability. If we are given probability densities f and g for random variables X and Y respectively, and if we know that X and Y are independent, we can compute the moments of $X + Y$, and XY , for example, from the moments of X and Y .

Our viewpoint on free probability is similar. Given two random matrices, \mathbf{A} and \mathbf{B} with eigenvalue density functions f and g , we would like to compute the eigenvalue density functions for $\mathbf{A} + \mathbf{B}$ and \mathbf{AB} in terms of the moments of \mathbf{A} and \mathbf{B} .

Of course, \mathbf{A} and \mathbf{B} do not commute so we are in the realm of non-commutative algebra. Since all possible products of \mathbf{A} and \mathbf{B} are allowed we have the “free” product, *i.e.*, all words in \mathbf{A} and \mathbf{B} are allowed. (We recall that this is precisely the definition of the free product in algebra.) The theory of free probability allows us to compute the moments of these products in the large matrix limit, *i.e.*, $N \rightarrow \infty$ so long as \mathbf{A} and \mathbf{B} are (asymptotically) free. In that sense (asymptotic) freeness, for large random matrices, is considered the analogue of independence for scalar valued random variables. Remarkably, asymptotic freeness results whenever \mathbf{A} (or \mathbf{B}) has isotropically random eigenvectors so that they bear no relationship to the eigenvectors of \mathbf{B} (or \mathbf{A} , resp.). In other words, a sufficient condition is for asymptotic freeness of \mathbf{A} and \mathbf{B} is that that the eigenvectors of \mathbf{A} (or \mathbf{B}) are uniformly distributed with Haar measure.

When \mathbf{A} and \mathbf{B} are asymptotically free, the limiting eigenvalue density function of $\mathbf{A} + \mathbf{B}$ (or \mathbf{AB}) is the free additive (or multiplicative) convolution of the limiting eigenvalue density function of \mathbf{A} and \mathbf{B} , thereby mirroring the structure for the sums and products of independent scalar valued random variables. In this sense, the development of free probability theory constitutes a breakthrough in our understanding of the behavior of large random matrices. Despite this elegant formulation, researchers were only able to use the underlying free convolution machinery for concrete computations for some simple cases. In this thesis, we solve this problem by establishing a computational free probability framework (see Section 8.4).

■ 2.2.2 Algebraic random matrices and a random matrix “calculator”

The development of this framework accompanied our characterization of the class of *algebraic random matrices* for which the limiting eigenvalue distribution and the associated moments can be concretely computed. The class of algebraic random matrices is defined next.

Definition 1 (Algebraic random matrices). Let $F^A(x)$ denote the limiting eigenvalue distribution function of a sequence of random matrices \mathbf{A}_N . If a bivariate polynomial $L_{\text{mz}}(m, z)$ exists such that

$$m_A(z) = \int \frac{1}{x-z} dF^A(x) \quad z \in \mathbb{C}^+ \setminus \mathbb{R}$$

is a solution of $L_{\text{mz}}(m_A(z), z) = 0$ then \mathbf{A}_N is said to be an algebraic random matrix. The density function $f_A = dF^A$ is referred to as an algebraic density and we say that $\mathbf{A}_N \in \mathcal{M}_{\text{alg}}$, the class of algebraic random matrices.

The utility of this, admittedly technical, definition comes from the fact that we are able to concretely specify the generators of this class. We illustrate this with a simple example. Let \mathbf{G} be an $n \times m$ random matrix with i.i.d. standard normal entries with variance $1/m$. The matrix $\mathbf{W}(c) = \mathbf{G}\mathbf{G}'$ is the Wishart matrix parameterized by $c = n/m$. Let \mathbf{A} be an arbitrary algebraic random matrix independent of $\mathbf{W}(c)$.

Figure 2-2 identifies deterministic and stochastic operations that can be performed on \mathbf{A} so that the resulting matrix is algebraic as well. The calculator analogy is apt because once we start with an algebraic random matrix, if we keep pushing away at the buttons we still get an algebraic random matrix whose limiting eigenvalue distribution is concretely computable using the algorithms developed in Section 8.

The algebraicity definition is important because everything we want to know about the limiting eigenvalue distribution of \mathbf{A} is encoded in the bivariate polynomial $L_{\text{mz}}^A(m, z)$. Thus, in establishing the algebraicity of any of the transformations in Figure 2-2, we have in effect determined the operational law for the polynomial transformation $L_{\text{mz}}^A(m, z) \mapsto L_{\text{mz}}^B(m, z)$ corresponding to the random matrix transformation $\mathbf{A} \mapsto \mathbf{B}$.

The catalogue of admissible transformations and their software realization is found in Section 8. This then allows us to calculate the eigenvalue distribution functions of a large class of algebraic random matrices that are generated from other algebraic random matrices.

We illustrate the underlying technique of mapping canonical operations of random matrices into operations on the bivariate polynomials with a simple example. Suppose we take the Wigner matrix, sampled in MATLAB as:

```
G = sign(randn(N))/sqrt(N); A = (G+G')/sqrt(2);
```




Figure 2-2. A random matrix calculator where a sequence of deterministic and stochastic operations performed on an “algebraically characterizable” random matrix sequence \mathbf{A}_N produces a “algebraically characterizable” random matrix sequence \mathbf{B}_N . The limiting eigenvalue density and moments of a “characterizable” matrix can be computed numerically, with the latter often in closed form.

whose eigenvalues in the $N \rightarrow \infty$ limit follow the semicircle law, and the Wishart matrix which may be sampled in MATLAB as:

$$G = \text{randn}(N, 2*N) / \text{sqrt}(2*N); \quad B = G*G';$$

whose eigenvalues in the limit follow the Marčenko-Pastur law. The associated limiting eigenvalue distribution functions have Stieltjes transforms $m_A(z)$ and $m_B(z)$ that are solutions of the equations $L_{mz}^A(m, z) = 0$ and $L_{mz}^B(m, z) = 0$, respectively, where

$$L_{mz}^A(m, z) = m^2 + zm + 1, \quad L_{mz}^B(m, z) = m^2z - (-2z + 1)m + 2.$$

The sum and product of these random matrices have limiting eigenvalue distribution whose Stieltjes transform is a solution of the bivariate polynomial equations $L_{mz}^{A+B}(m, z) = 0$ and $L_{mz}^{AB}(m, z) = 0$, respectively, which can be calculated from L_{mz}^A and L_{mz}^B alone as shown below.

To obtain $L_{mz}^{A+B}(m, z)$ we apply the transformation labelled as “Add Atomic Wishart”

in Table 8.1 with $c = 2$, $p_1 = 1$ and $\lambda_1 = 0.5$ to obtain the operational law

$$L_{mz}^{A+B}(m, z) = L_{mz}^A \left(m, z - \frac{1}{1 + 0.5m} \right). \quad (2.8)$$

Substituting $L_{mz}^A = m^2 + zm + 1$ in (2.8) and clearing the denominator, yields the bivariate polynomial

$$L_{mz}^{A+B}(m, z) = m^3 + (z + 2)m^2 - (-2z + 1)m + 2. \quad (2.9)$$

Similarly, to obtain L_{mz}^{AB} , we apply the transformation labelled as ‘‘Multiply Wishart’’ in Table 8.1 with $c = 0.5$ to obtain the operational law

$$L_{mz}^{AB}(m, z) = L_{mz}^A \left((0.5 - 0.5zm)m, \frac{z}{0.5 - 0.5zm} \right). \quad (2.10)$$

Substituting $L_{mz}^A = m^2 + zm + 1$ in (2.10) and clearing the denominator, yields the bivariate polynomial

$$L_{mz}^{AB}(m, z) = m^4 z^2 - 2m^3 z + m^2 + 4mz + 4. \quad (2.11)$$

Figure 2-3 plots the density function associated with the limiting eigenvalue distribution for the Wigner and Wishart matrices as well as their sum and product extracted directly from $L_{mz}^{A+B}(m, z)$ and $L_{mz}^{AB}(m, z)$.

In this simple case, the polynomials were obtained by hand calculation. Along with the theory of algebraic random matrices we also develop a software realization that maps the entire catalog of transformations (see Tables 8.1 -8.3) into symbolic MATLAB code. Thus, for the example considered, the sequence of commands:

```
>> syms m z
>> LmZA = m^2+z*m+1;
>> LmzB = m^2-(-2*z+1)*m+2;
>> LmzApB = AplusB(LmZA,LmzB);
>> LmzAtB = AtimesB(LmZA,LmzB);
```

could also have been used to obtain L_{mz}^{A+B} and L_{mz}^{AB} . The commands **AplusB** and **AtimesB** implicitly use the free convolution machinery to perform the said computation.

To summarize, by defining the class of algebraic random matrices, we are able to extend the reach of infinite random matrix theory well beyond the special cases of matrices with Gaussian entries. The key idea is that by encoding probability densities as solutions of bivariate polynomial equations, and deriving the correct operational laws on this encoding, we can take advantage of powerful symbolic and numerical techniques to compute these densities and their associated moments. In particular, for the examples considered, algebraically extracting the roots of these polynomials using the cubic or quartic formulas would be of little use. Consequently, looking for special cases where

the limiting density function can be written in closed form is needlessly restrictive unless one is attempting to classify these special random matrix ensembles.

The statistical techniques developed in this thesis, do not, however, exploit the full scope of this method which is developed in Chapters 6-10. The possibility of being able to characterize matrices more complicated than those formed from entries with Gaussian elements makes it possible to start thinking about formulating inference procedures for these more complicated random matrix models. In this thesis, we leave the theory at that. While we illustrate the power of the method with some examples, we leave it to practitioners to motivate additional applications that exploit the full power of the stochastic eigen-analysis techniques developed.

■ 2.3 Contributions of this thesis


As an addition to the table of contents, we now itemize the results we consider most important and where they can be found. We remark that all statements labelled as theorems represent, we believe, new results, while important results from the literature are labelled as propositions or lemmas. Chapters 3-5 are self-contained and can be read independently. Chapters 6 - 10 describe the “polynomial method” for characterizing a broad class of random matrices and may be read separately from the preceding material. Where appropriate, every chapter contains a section on future work and other directions of research. The thesis contributions are:

- New algorithm for detecting number of signals in white noise from the sample eigenvalues alone that dramatically outperforms the Kailath-Wax estimator. (see Chapter 3 for details and Table 3.1 for the algorithm). This solves a long-standing open problem in sensor array processing.
- Heuristic explanation of the phase transition phenomenon for largest (“signal”) eigenvalues. This establishes a fundamental limit in detection using the signal eigenvalues alone. Roughly speaking, for large number of sensors and snapshots, the signals can be reliably detected using the method developed if


$$\text{Signal Power} > \text{Noise Power} \sqrt{\frac{\# \text{ Sensors}}{\# \text{ Snapshots}}}$$

Consistency of the estimators in Table 3.1 with respect to to the concept of *effective number of signals* is discussed in Section 3.7.

- New eigen-inference techniques for testing equality of population eigenvalues and parametrically estimating population eigenvalues are presented in Chapter 4. These techniques supplant Anderson’s techniques for high-dimensional, sample size constrained settings.



(a) The limiting eigenvalue density function for the GOE and Wishart matrices.



(b) The limiting eigenvalue density function for the sum and product of independent GOE and Wishart matrices.

Figure 2-3. A representative computation using the random matrix calculator.

- We provide an approximation for the distribution of the outputs of the diagonally loaded Capon-MVDR beamformer in Chapter 5 solving an open problem posed by Baggeroer and Cox in [7].
- We describe a class of “algebraic” random matrices (see Chapter 7). These are random matrices for which the Stieltjes transform of the limiting eigenvalue distribution function is an algebraic function. Wishart matrices with identity covariance are a special case. The practical utility of this definition is that if a random matrix is shown to be algebraic then its limiting eigenvalue density function can be computed using a simple root-finding algorithm. Furthermore, if the moments exist, then we will often be able to enumerate them efficiently in closed form. By specifying the class of such random matrices by its generators we solve an open problem in computational random matrix theory by extending the reach of the theory to concretely predict the limiting distribution of a much broader class of random matrices than thought possible.
- We describe the computation of the Markov transition kernel for certain classes of algebraic random matrices. The Markov transition kernel encodes the conditional “eigenvector distribution” (see Chapter 10 for a precise description) of algebraic random matrices. The computation facilitates analysis, for the first time, of the eigenvectors of a broad subclass of algebraic random matrices including those with Wishart distribution

Statistical eigen-inference: Signals in white noise

■ 3.1 Introduction

The observation vector, in many signal processing applications, can be modelled as a superposition of a finite number of signals embedded in additive noise. Detecting the number of signals present becomes a key issue and is often the starting point for the signal parameter estimation problem. When the signals and the noise are assumed to be samples of a stationary, ergodic Gaussian vector process, the sample covariance matrix formed from m observations has the Wishart distribution [114]. The proposed algorithm uses an information theoretic approach for determining the number of signals in white noise by examining the eigenvalues of the resulting sample covariance matrix. An essential component of the proposed estimator is its explicit dependence on the dimensionality of the observation vector and the number of samples used to form the sample covariance matrix. This makes the proposed estimator robust to high-dimensionality and sample size constraints.

The form of the estimator is motivated by results on the eigenvalues of large dimensional sample covariance matrices [10–12, 32, 49, 51, 70]. We are able to re-derive a portion of these results [11, 12, 70], reported by other authors in the literature, using an interacting particle system interpretation, thereby providing insight into the structure of the proposed solution and its shortcomings. The concept of *effective number of signals* is introduced (see Section 3.7), which depends in a simple manner on the noise variance, sample size and dimensionality of the system. This concept captures the fundamental limits of sample eigenvalue based detection by explaining why, asymptotically, if the signal level is below a threshold that depends on the noise variance, sample size and the dimensionality of the system, then reliable detection is not possible. More importantly, the proposed estimators dramatically outperforms the standard estimators found in the literature, particularly so in sample starved settings. While such a behavior is to be expected when the dimensionality of the system is large because of the nature of the random matrix results being exploited, this trend is observed in smaller dimensional settings as well.

This chapter is organized as follows. The problem formulation in Section 3.2 is

followed by an exposition in Section 3.3, using an interacting particle system interpretation, of the properties of the eigenvalues of large dimensional sample covariance matrices when there are no signals present. The analysis is extended in Section 3.4 to the case when there are signals present. The occurrence of a phase transition phenomenon in the identifiability of the largest (“signal”) eigenvalue is heuristically described and re-derived using an interacting particle system interpretation. An estimator for the number of signals present that exploits these results is derived in Section 3.5. An extension of these results to the frequency domain is discussed in Section 3.6. Consistency of the proposed estimators and the concept of *effective number of signals* is discussed in Section 3.7. Simulation results that illustrate the superior performance of the new method in high dimensional, sample size starved settings are presented in Section 3.8; some concluding remarks are presented in Section 3.9.

■ 3.2 Problem formulation

We observe m samples (“snapshots”) of possibly signal bearing n -dimensional snapshot vectors $\mathbf{y}_1, \dots, \mathbf{y}_m$ where for each i , $\mathbf{x}_i \sim \mathcal{N}_n(\mathbf{0}, \mathbf{R})$ and \mathbf{x}_i are mutually independent. The snapshot vectors are modelled as

$$\mathbf{x}_i = \begin{cases} \mathbf{z}_i & \text{No Signal} \\ \mathbf{A} \mathbf{s}_i + \mathbf{z}_i & \text{Signal Present} \end{cases} \quad \text{for } i = 1, \dots, m, \quad (3.1)$$

where $\mathbf{z}_i \sim \mathcal{N}_n(0, \sigma^2 \mathbf{I})$, denotes an n -dimensional (real or complex) Gaussian noise vector where σ^2 is generically unknown, $\mathbf{s}_i \sim \mathcal{N}_k(\mathbf{0}, \mathbf{I})$, $\mathbf{s}_i \sim \mathcal{N}_k(\mathbf{0}, \mathbf{R}_s)$ denotes a k -dimensional (real or complex) Gaussian signal vector with covariance \mathbf{R}_s , and \mathbf{A} is a $n \times k$ unknown non-random matrix. In array processing applications, the j -th column of the matrix \mathbf{A} encodes the parameter vector associated with the j -th signal whose magnitude is described by the j -th element of \mathbf{s}_i .

Since the signal and noise vectors are independent of each other, the covariance matrix of \mathbf{x}_i can hence be decomposed as

$$\mathbf{R} = \mathbf{\Psi} + \sigma^2 \mathbf{I} \quad (3.2)$$

where

$$\mathbf{\Psi} = \mathbf{A} \mathbf{R}_s \mathbf{A}', \quad (3.3)$$

with $'$ denoting the conjugate transpose. Assuming that the matrix \mathbf{A} is of full column rank, *i.e.*, the columns of \mathbf{A} are linearly independent, and that the covariance matrix of the signals \mathbf{R}_s is nonsingular, it follows that the rank of $\mathbf{\Psi}$ is k . Equivalently, the $n - k$ smallest eigenvalues of $\mathbf{\Psi}$ are equal to zero.

If we denote the eigenvalues of \mathbf{R} by $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ then it follows that the smallest $n - k$ eigenvalues of \mathbf{R} are all equal to σ^2 so that

$$\lambda_{k+1} = \lambda_{k+2} = \dots = \lambda_n = \lambda = \sigma^2. \quad (3.4)$$

Thus, if the true covariance matrix \mathbf{R} were known apriori, the dimension of the signal vector k can be determined from the multiplicity of the smallest eigenvalue of \mathbf{R} . When there is no signal present, all the eigenvalues of \mathbf{R} will be identical. The problem in practice is that the covariance matrix \mathbf{R} is unknown so that such a straight-forward algorithm cannot be used. The signal detection and estimation problem is hence posed in terms of an inference problem on m samples of n -dimensional multivariate real or complex Gaussian snapshot vectors.

Inferring the number of signals from these m samples reduces the signal detection problem to a model selection problem for which there are many approaches. A classical approach to this problem, developed by Bartlett [13] and Lawley [54], uses a sequence of hypothesis tests. Though this approach is sophisticated, the main problem is the subjective judgement needed by the practitioner in selecting the threshold levels for the different tests.

Information theoretic criteria for model selection such as those developed by Akaike [1,2], Schwartz [80] and Rissanen [76] address this problem by proposing the selection of the model which gives the minimum *information criteria*. The criteria for the various approaches is generically a function of the log-likelihood of the maximum likelihood estimator of the parameters of the model and a term which depends on the number of parameters of the model that penalizes overfitting of the model order.

For the problem formulated above, Wax and Kailath propose an estimator [111] for the number of signals (assuming $m > n$) based on the eigenvalues $l_1 \geq l_2 \geq \dots \geq l_n$ of the sample covariance matrix (SCM) defined by

$$\hat{\mathbf{R}} = \frac{1}{m} \sum_{i=1}^m \mathbf{x}_i \mathbf{x}_i' = \frac{1}{m} \mathbf{X} \mathbf{X}' \quad (3.5)$$

where $\mathbf{X} = [\mathbf{x}_1 | \dots | \mathbf{x}_m]$ is the matrix of observations (samples). The Akaike Information Criteria (AIC) form of the estimator is given by

$$\hat{k}_{\text{AIC}} = \arg \min_{k \in \mathbb{N}: 0 \leq k < n} -2(n-k)m \log \frac{g(k)}{a(k)} + 2k(2n-k) \quad (3.6)$$

while the Minimum Descriptive Length (MDL) criterion is given by

$$\hat{k}_{\text{MDL}} = \arg \min_{k \in \mathbb{N}: 0 \leq k < n} -(n-k)m \log \frac{g(k)}{a(k)} + \frac{1}{2}k(2n-k) \log m \quad (3.7)$$

where $g(k) = \prod_{j=k+1}^n l_j^{1/(n-k)}$ is the geometric mean of the $n-k$ smallest sample eigenvalues and $a(k) = \frac{1}{n-k} \sum_{j=k+1}^n l_j$ is their arithmetic mean.

These estimators perform adequately only when the sample size greatly exceeds the dimension of the system by a factor of 15 – 100. While their large sample consistency have been analytically established, these results do not lend any insight into the shortcomings in situations where the dimensionality of the system is large or the number of

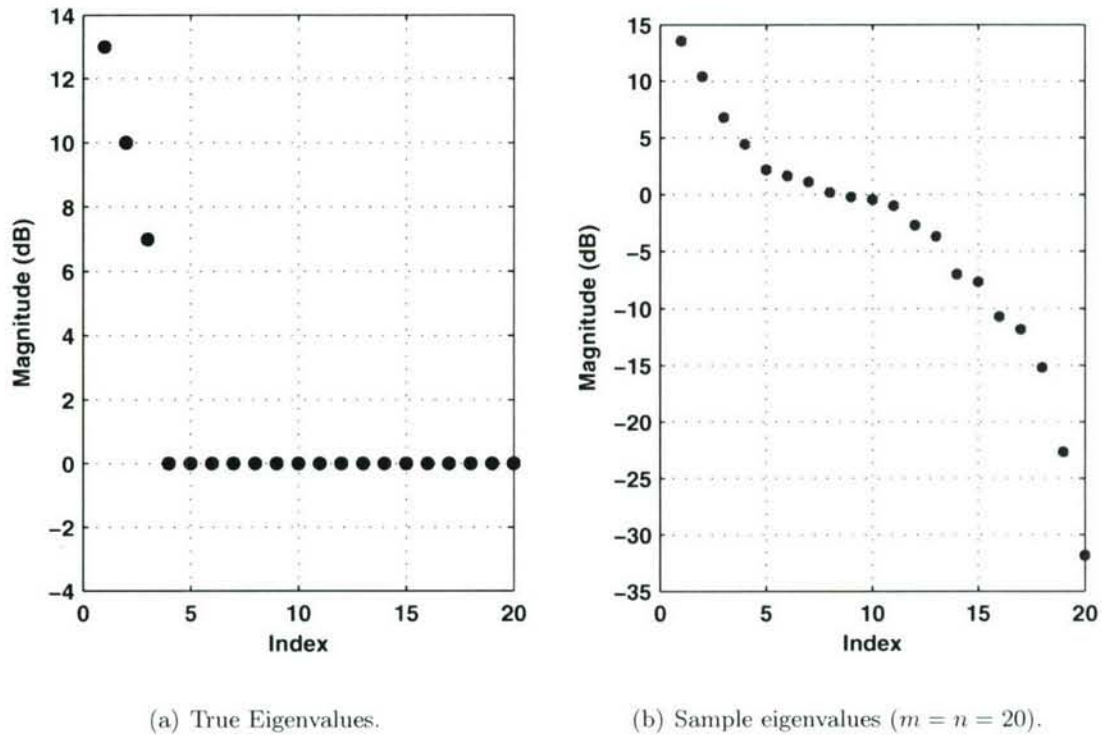


Figure 3-1. Blurring of sample eigenvalues due to finite number of snapshots.

samples size is on the order of the dimensionality or both, as is increasingly the case in many signal processing and scientific applications. The reason why these estimators perform so poorly while the proposed estimators, summarized in Table 3.1, perform so well in these settings is best illustrated by an example.

Figure 3-1 compares the 20 eigenvalues of the true covariance matrix (with noise variance equal to 1) with the eigenvalues of a single sample covariance matrix formed from 20 snapshots. The three “signal” eigenvalues can be readily distinguished in the true covariance eigen-spectrum; the distinction is less clear in the sample eigen-spectrum because of the significant *blurring* of the signal and noise eigenvalues.

Traditional estimators, including the Wax-Kailath algorithm, perform poorly in high-dimensional, sample size constrained settings because they do not account for this blurring; the proposed estimators are able to overcome this limitation by *explicitly* exploiting analytical results that capture the dependence of the blurring on the noise variance, the sample size and the dimensionality of the system. The applicability of the algorithms in scenarios where the sample size is less than the dimensionality of the system is a feature that makes it suitable for sensor array processing and other emerging applications in science and finance where such situations are routinely encountered.

Furthermore, the analytical results provide insight into the fundamental limit, due

to sample size constraints in high-dimensional settings, of reliable signal detection by eigen-inference, *i.e.*, by using the sample eigenvalues alone (see Section 3.7). This helps identify scenarios where algorithms that exploit any structure in the eigenvectors of the signals, such as the MUSIC [79] and the Capon-MVDR [19] algorithms in sensor array processing, might be better able to tease out lower level signals from the background noise. It is worth noting that the proposed approach remains relevant in situations where the eigenvector structure has been identified. This is because eigen-inference methodologies are inherently robust to eigenvector modelling errors that occur in high-dimensional settings. Thus the practitioner may use the proposed methodologies to complement and “robustify” the inference provided by algorithms that exploit the eigenvector structure.

■ 3.3 Eigenvalues of the (null) Wishart matrix

When there are no signals present, $\mathbf{R} = \lambda \mathbf{I}$ so that the SCM $\widehat{\mathbf{R}}$ is sampled from the (null) Wishart distribution [114]. The joint density function of the eigenvalues l_1, \dots, l_n of $\widehat{\mathbf{R}}$ when $m > n + 1$ is given by [67]

$$Z_{n,m}^\beta \exp\left(-\frac{\beta m}{2\lambda} \sum_{i=1}^n l_i\right) \prod_{i=1}^n l_i^{\beta(m-n+1)/2-1} \prod_{i<j}^n |l_i - l_j|^\beta \quad (3.8)$$

where $l_1 > \dots > l_n > 0$, $Z_{n,m}^\beta$ is the normalization constant and $\beta = 1$ (or 2) when $\widehat{\mathbf{R}}$ is real (or complex). Taking the negative logarithm of the joint density function in (3.8) and defining $\ell = (l_1, \dots, l_n)$ gives us the negative log-likelihood function

$$L(\ell) := -\log Z_{n,m} - \left(\frac{\beta(m-n+1)}{2} - 1\right) \sum_{i=1}^n \log l_i + \frac{\beta m}{2\lambda} \sum_{i=1}^n l_i - \beta \sum_{i<j}^n \log |l_i - l_j|. \quad (3.9)$$

■ 3.3.1 An interacting particle system interpretation

Let the sample eigenvalues l_1, \dots, l_n represent locations of particles, then (3.9) can be interpreted in statistical physics terms, as the *logarithmic energy* of this system of particles. Note that we constrain the particles to lie along the positive real axis so that $l_i > 0$.

The configuration of the particles that *minimizes* the logarithmic energy (assuming a unique minimum exists) is simply the maximum likelihood estimate of the sample eigenvalues. For the system represented in (3.9), it turns out that a unique minimum exists so we can proceed with trying to qualitatively predict the equilibrium configura-




Figure 3-2. Interacting particle system interpretation of the eigenvalues of a (null) Wishart matrix.

tion of the particles. Consider the rescaled (by $1/n^2$) logarithmic energy given by

$$\begin{aligned}
 V(\ell) := & \underbrace{-\frac{1}{n^2} \log Z_{n,m}^\beta}_{\text{constant}} - \underbrace{\left(\frac{\beta}{2} \left[\frac{1}{c_m} - 1 \right] + \frac{\beta - 2}{2n} \right) \frac{1}{n} \sum_i \log l_i}_I \\
 & - \underbrace{\beta \frac{1}{n^2} \sum_{i < j} \log |l_i - l_j|}_{\text{II}} + \underbrace{\frac{\beta}{2\lambda c_m} \frac{1}{n} \sum_i^n l_i}_{\text{III}}. \quad (3.10)
 \end{aligned}$$

with $c_m = n/m < 1$. The equilibrium position of the particles is the configuration that minimizes the logarithmic energy of the system given by (3.10) subject to the forces identified by the roman numerals. This involves balancing the three competing “forces” depicted in Figure 3-2. If the particles are placed in some arbitrary position as in Figure 3-2(a), they will be subjected to the competing forces described below, interact with each other and eventually reach an equilibrium configuration as in Figure 3-2(b). The term

$$T_1 := - \left(\frac{\beta}{2} \left[\frac{1}{c_m} - 1 \right] + \frac{\beta - 2}{2n} \right) \frac{1}{n} \sum_i \log l_i \approx -\frac{\beta}{2} \left(\frac{1}{c_m} - 1 \right) \frac{1}{n} \sum_i \log l_i \quad (3.11)$$

represents a *repulsion from the origin* that is minimized when the particles are further

away from the origin, *i.e.*, for larger values of l_i . The term

$$T_2 := -\beta \frac{1}{n^2} \sum_{i < j} \log |l_i - l_j| \quad (3.12)$$

represents an *inter-particle repulsion* that is minimized when the particles (sample eigenvalues) are spaced out as far apart as possible so that the difference $|l_i - l_j|$ is large. Finally, the term

$$T_3 := \frac{\beta}{2\lambda c_m} \frac{1}{n} \sum_i l_i \quad (3.13)$$

represents an *attraction to the origin* that is minimized when the particles are closer to the origin, *i.e.*, for small l_i . Generically speaking, for arbitrary $c_m < 1$ (and large m, n), since $\log l_i < l_i$, comparing (3.11) and (3.13), the particles experience an attraction to the origin that is greater than the repulsion away from the origin. Thus we can expect the sample eigenvalues to be distributed about $x = \lambda$ with a greater concentration closer towards the origin as depicted in Figure 3-2(b).

■ 3.3.2 Sample eigenvalues in the snapshot abundant regime

Continuing this physical analogy further, observe that in (3.10) the ratio $c_m = n/m < 1$ does not affect the (internal) repulsion between the particles (sample eigenvalues). Thus, for a fixed choice of λ , the value of c_m affects the equilibrium position by governing the manner in which the repulsion between the particles (T_2) is balanced by the repulsion/attraction of the origin (T_1 and T_3 respectively).

In the snapshot abundant regime, where the number of snapshots is significantly greater than the dimensionality of the system, we obtain values of $c_m = n/m$ very close to zero. Thus, since $m \gg n$, $1/c_m \gg 1$ so that the interaction between the particles can be neglected. In other words, the equilibrium configuration minimizes

$$V(l_1, \dots, l_n) \approx \frac{\beta}{2c_m} \frac{1}{n} \sum_i \log l_i - \frac{\beta}{2\lambda c_m} \frac{1}{n} \sum_i l_i$$

so that the equilibrium configuration of the i -th particle is determined by the condition

$$\frac{\partial}{\partial l_i} V(l_1, \dots, l_n) = 0.$$

This is equivalent to the condition

$$\frac{\beta}{l_i} - \frac{\beta}{\lambda} = 0,$$

resulting in $l_i = \lambda$ for both $\beta = 1$ (real) or $\beta = 2$ (complex), as expected. In other




Figure 3-3. Limiting distribution of the eigenvalues of the (null) Wishart matrix.

words, as $c_m \rightarrow 0$, $l_i \rightarrow \lambda$ for $i = 1, \dots, n$ so that the eigenvalues of the SCM converge to the (single) eigenvalue of the true covariance matrix.

■ 3.3.3 Limiting distribution of the sample eigenvalues

Generically speaking, for arbitrary values of $c_m = n/m$, the limiting distribution of the sample eigenvalues is influenced by all of the forces depicted in Figure 3-2. The limiting distribution exists and can be analytically computed. Define the empirical distribution function (e.d.f.) of the eigenvalues of an $n \times n$ self-adjoint matrix \mathbf{A}_n with n real eigenvalues (counted with multiplicity) as

$$F^{\mathbf{A}_n}(x) = \frac{\text{Number of eigenvalues of } \mathbf{A}_n \leq x}{n}. \quad (3.14)$$

Proposition 3.31. *Let $\hat{\mathbf{R}}$ denote a sample covariance matrix formed from an $n \times m$ matrix of observations with i.i.d. Gaussian samples of mean zero and variance λ . Then the e.d.f. $F^{\hat{\mathbf{R}}} \rightarrow F^W$ almost surely, as $m, n \rightarrow \infty$ and $c_m = n/m \rightarrow c$ where F^W is a non-random distribution function with density*

$$f_W(x) := dF^W(x) = \min\left(0, \left(1 - \frac{1}{c}\right)\right) \delta(x) + \frac{\sqrt{(x - a_-)(a_+ - x)}}{2\pi\lambda xc} \mathbb{I}_{[a_-, a_+]}(x) \quad (3.15)$$

with $a_{\pm} = \lambda(1 \pm \sqrt{c})^2$, $\mathbb{I}_{[a, b]}(x) = 1$ when $a \leq x \leq b$ and zero otherwise, and $\delta(x)$ is the Dirac delta function.

PROOF. This result was proved in [59, 110] in very general settings. Other proofs include [51, 86, 87]. The distribution may also be obtained by first determining the (non-random) equilibrium positions of the n particles and then determining the e.d.f. in the large n, m limit. These positions will precisely coincide with the appropriately normalized zeros of the n -th degree Laguerre polynomial. A proof of this fact follows readily from Szego's exposition on orthogonal polynomials [99] once the correspondence between the (null) Wishart distribution and the Laguerre orthogonal polynomial is recognized as in [34].

The approach taken in [45] explicitly relies on the interacting particle system interpretation by showing that the density $\mu \equiv \mu_{\mathbf{W}}$ is the unique minimizer of the functional obtained from the limit of (3.10)

$$V(\mu) := \text{Constant} - \frac{\beta}{2} \int Q(x)\mu(x)dx - \beta \iint \log|x-y|\mu(x)\mu(y)dxdy, \quad (3.16)$$

where $Q(x) = (1/c - 1) \log x - x/\lambda c$, $m, n \rightarrow \infty$ with $m/n \rightarrow c \in (0, 1]$. \square

The density $f_W(x)$, with $\lambda = 1$, is shown in Figure 3-3 for different values of $c \in (0, 1]$ confirming our qualitative prediction about its relative skewing towards the origin for moderate values of c and a localization about $\lambda = 1$ for values of c close to zero.

■ 3.3.4 Gaussian fluctuations of sample eigenvalues

The almost sure convergence of the e.d.f. of the (null) Wishart matrix implies that for any "well-behaved" function h ,

$$\frac{1}{n} \sum_{i=1}^n h(l_i) \rightarrow \int h(x)dF^W(x). \quad (3.17)$$

where the convergence in the above is almost surely. In particular, when h is a monomial, we obtain the moments associated with the density function

$$M_k^W := \int x^k dF^W(x) = \lambda^k \sum_{j=0}^{k-1} c^j \frac{1}{j+1} \binom{k}{j} \binom{k-1}{j}. \quad (3.18)$$

We can take this a step further by examining the fluctuations about these limiting results. Precisely speaking, for h a monomial as above (or more generally), once we subtract the expected average over the limiting eigenvalue density, *i.e.*, the right hand side of (3.17), the rescaled resulting quantity tends asymptotically to a normal distribution with mean and variance depending on h .

Proposition 3.32. *If $\widehat{\mathbf{R}}$ satisfies the hypotheses of Proposition 3.31 with $\lambda = 1$ then*

as $m, n \rightarrow \infty$ and $c_m = n/m \rightarrow c \in (0, \infty)$, then

$$\mathbf{t} = \begin{bmatrix} \sum_{i=1}^n l_i - n \\ \sum_{i=1}^n l_i^2 - n(1+c) - (\frac{2}{\beta} - 1)c \end{bmatrix} \rightarrow \mathbf{g} \sim \mathcal{N}\left(\mathbf{0}, \frac{2}{\beta} \mathbf{Q}\right) \quad (3.19)$$

where the convergence in distribution is almost surely, $\beta = 1$ (or 2) when \mathbf{x}_i is real (or complex) valued, respectively, and

$$\mathbf{Q} = \begin{bmatrix} c & 2c(c+1) \\ 2c(c+1) & 2c(2c^2 + 5c + 2) \end{bmatrix} \quad (3.20)$$

PROOF. This result appears in [49, 51] for the real case and in [10] for the real and complex cases. The result for general β appears in Dumitriu and Edelman [32]. \square It is worth remarking that while the limiting density of the SCM does not depend on whether the elements of the observation (snapshot) vectors are real or complex, the mean and variance of the fluctuations do. The Gaussianity of the eigenvalue fluctuations is consistent with our association of the limiting density with the maximum likelihood equilibrium configuration of the interacting particle system. The asymmetric interaction between the largest eigenvalue and the ‘‘bulk’’ of the eigen-spectrum accounts for the non-Gaussianity of the fluctuations of the largest eigenvalue which follow the Tracy-Widom distribution [50].

■ 3.4 Signals in white noise

For arbitrary covariance \mathbf{R} the joint density function of the eigenvalues l_1, \dots, l_n of the SCM $\widehat{\mathbf{R}}$ when $m > n + 1$ is given by

$$\tilde{Z}_{n,m}^\beta \sum_{i=1}^n l_i^{\beta(m-n+1)/2-1} \prod_{i<j}^n |l_i - l_j|^\beta \int_{\mathbf{Q}} \exp\left(-\frac{m\beta}{2} \text{Tr}(\mathbf{\Lambda}^{-1} \mathbf{Q} \mathbf{L} \mathbf{Q}')\right) d\mathbf{Q} \quad (3.21)$$

where $l_1 > \dots > l_n > 0$, $\tilde{Z}_{n,m}^\beta$ is a normalization constant, and $\beta = 1$ (or 2) when $\widehat{\mathbf{R}}$ is real (resp. complex). In (3.21), $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$, $\mathbf{L} = \text{diag}(l_1, \dots, l_n)$, $\mathbf{Q} \in \mathbf{O}(n)$ when $\beta = 1$ while $\mathbf{Q} \in \mathbf{U}(n)$ when $\beta = 2$ where $\mathbf{O}(n)$ and $\mathbf{U}(n)$ are, respectively, the set of $n \times n$ orthogonal and unitary matrices with Haar measure. The Haar measure is the unique uniform measure on orthogonal/unitary matrices; see Chapter 1 of Milman and Schechtman for a derivation [62].

It can be readily seen that when $\mathbf{R} = \mathbf{\Lambda} = \lambda \mathbf{I}$ so that

$$\int_{\mathbf{Q}} \exp\left(-\frac{m\beta}{2} \text{Tr}(\mathbf{\Lambda}^{-1} \mathbf{Q} \mathbf{L} \mathbf{Q}')\right) d\mathbf{Q} = \exp\left(-\frac{m\beta}{2\lambda} \sum_{i=1}^n l_i\right),$$

the joint density in (3.21) reduces to (3.8) when $Z_{n,m}^\beta$ and $\tilde{Z}_{n,m}^\beta$ are appropriately

defined.

In the general case, for arbitrary \mathbf{R} , the expression for the joint density in (3.21) is difficult to analyze because of the integral over the orthogonal (or unitary) group. For the problem we are interested in, when there are k signals in white noise and $\mathbf{R} = \text{diag}(\lambda_1, \dots, \lambda_k, \lambda, \dots, \lambda)$, an examination of the large m approximation of this integral can give us additional sights.

For this purpose, it suffices to examine the scenario when there is a single real-valued signal in white noise, *i.e.*, $k = 1$ and $\beta = 1$, for which we may employ the approximation stated in [67]

$$\int_{\mathbf{Q}} \exp\left(-\frac{m\beta}{2} \text{Tr}(\mathbf{A}^{-1} \mathbf{Q} \mathbf{L} \mathbf{Q}')\right) d\mathbf{Q} \approx C_{n,m} \exp\left(-\frac{ml_1}{2\lambda_1}\right) \exp\left(-\frac{m}{2\lambda} \sum_{i=2}^n l_i\right) \prod_{j=2}^n |l_1 - l_j|^{-1/2} \quad (3.22)$$

with $C_{n,m}$ being a normalization constant so that (3.21) may be approximated by

$$\underbrace{\tilde{C}_{n,m}^B \prod_{i=2}^n l_i^{(m-n-1)/2} \prod_{1 < i < j} |l_i - l_j| \exp\left(-\frac{m}{2\lambda} \sum_{i=2}^n l_i\right)}_{L_{\text{Bulk}}(l_2, \dots, l_n)} \times \underbrace{\tilde{C}_{n,m}^S l_1^{(m-n-1)/2} \exp\left(-\frac{ml_1}{2\lambda_1}\right) \prod_{j=2}^n |l_1 - l_j|^{1/2}}_{L_{\text{Spk}}^1(l_1 | l_2, \dots, l_n)}. \quad (3.23)$$

Note that the approximated joint density in (3.23) has been decomposed as shown, into the product of the joint density, $L_{\text{Bulk}}(l_2, \dots, l_n)$, of the “noise” eigenvalues and the conditional density, $L_{\text{Spk}}^1(l_1 | l_2, \dots, l_n)$ of the largest (signal) eigenvalue where $\tilde{C}_{n,m}^B$ and $\tilde{C}_{n,m}^S$ are normalization constants.

■ 3.4.1 Interacting particle system interpretation

As before, let the sample eigenvalues l_1, \dots, l_n represent locations of particles. Thus, the rescaled (by $1/n^2$) negative log-likelihood of the joint density function is interpreted as the logarithmic energy of the particle system whose n particles are located at positions l_1, \dots, l_n . From (3.23), the logarithmic energy may be approximated as

$$V(\ell) = \underbrace{-\frac{1}{n^2} \log C_{n,m}}_{\text{Constant}} + V_N(l_2, \dots, l_n) + \frac{1}{n} V_S(l_1 | l_2, \dots, l_n) \quad (3.24)$$



Figure 3-4. Interactions between the “signal” eigenvalue and the “noise” eigenvalues.

where the contributions due to the particles that represent the noise and signal eigenvalues are respectively given by

$$V_N(l_2, \dots, l_n) := \underbrace{-\frac{1}{2} \left(\frac{1}{c_m} - 1 - \frac{1}{n} \right) \frac{1}{n} \sum_{i=2}^n \log l_i}_{\text{I}} \underbrace{- \frac{1}{n^2} \sum_{1 < i < j} \log |l_i - l_j|}_{\text{II}} + \underbrace{\frac{1}{2\lambda c_m} \frac{1}{n} \sum_{i=2}^n l_i}_{\text{III}} \quad (3.25a)$$

and

$$V_S(l_1 | l_2, \dots, l_n) := \underbrace{-\frac{1}{2} \left(\frac{1}{c_m} - 1 - \frac{1}{n} \right) \log l_1}_{\text{I-Signal}} \underbrace{- \frac{1}{2} \frac{1}{n} \sum_{j=2}^n \log |l_1 - l_j|}_{\text{II-Signal}} + \underbrace{\frac{1}{2 c_m \lambda_1} l_1}_{\text{III-Signal}} \quad (3.25b)$$

with $c_m = n/m < 1$. The equilibrium configuration of the particles minimizes the logarithmic energy $V(\ell)$. The decomposition of the logarithmic energy as in (3.24) hints at the possibility of predicting the resulting configuration by using a two step approach. Specifically, for large enough n so that $V_N \gg (1/n)V_S$, the configuration of the $n - 1$ (“noise”) particles, l_2, \dots, l_n , that minimizes $V_N(l_2, \dots, l_n)$ should be a very good approximation of the configuration that minimizes $V(\ell)$.

■ 3.4.2 Repulsion and phase transition of largest (“signal”) eigenvalue

Conditioned on the resulting configuration of the $n - 1$ noise particles, the configuration of the n -th particle minimizes $V_S(l_1 | l_2, \dots, l_n)$. The underbraced terms in (3.25b) represent the forces that the n -th particle is subjected to. They denote, respectively, a force of repulsion away from the origin, a force of repulsion away from the $n - 1$ noise particles, and a force of attraction towards the origin as depicted in Figure 3-4. The equilibrium configuration of the particle l_1 is determined by the condition

$$\frac{\partial}{\partial l_1} V_S(l_1 | l_2, \dots, l_n) = 0 \quad (3.26)$$

which for large n , from (3.25b), reduces to

$$-\frac{1}{2} \left(\frac{1}{c_m} - 1 \right) \frac{1}{l_1} + \frac{1}{2 c_m \lambda_1} - \frac{1}{2} \frac{1}{n} \sum_{j=2}^n \frac{1}{l_1 - l_j} = 0. \quad (3.27)$$

In the large n, m limit, comparing (3.10) and (3.25a), it is clear that the equilibrium configuration of the $n - 1$ “noise” particles will be very well approximated by the configuration that results in the no signal case. Asymptotically, it is reasonable to replace the discrete condition for the minimization configuration in (3.25b) with its continuous limit so that equilibrium location of l_1 satisfies the equation

$$-\frac{1}{2} \left(\frac{1}{c} - 1 \right) \frac{1}{l_1} + \frac{1}{2c\lambda_1} - \frac{1}{2} \underbrace{\int \frac{1}{l_1 - x} dF^W}_{g_W(l_1)} = 0 \quad (3.28)$$

where $F^W(x)$ is the Marčenko-Pastur distribution in (3.15) and $c_m = n/m \rightarrow c < 1$ as $n, m \rightarrow \infty$. The Cauchy transform of the distribution function, F^A , is defined as

$$g_A(z) = \int \frac{1}{z - x} dF^A(x) \quad \text{for } z \in \mathbb{C}^+ \setminus \mathbb{R}. \quad (3.29)$$

Thus the underbraced term in (3.28) is the Cauchy transform of the distribution function $F^W(x)$ evaluated at $z = l_1$ so that $g_W(l_1)$ represents the *effective* repulsive force acting on the “signal” particle due to the “noise” particles. It can be seen from the definition of the Cauchy transform itself that $g(z) \sim 1/z$ for large $z \rightarrow \infty$ so that the effective repulsion felt by the n -th particle decreases the further away it is from the remaining $n - 1$ particles. The equilibrium configuration of the n -th particle is thus given by the force balancing condition

$$-\frac{1}{2} \left(\frac{1}{c} - 1 \right) \frac{1}{l_1} + \frac{1}{2c\lambda_1} - \frac{1}{2} g_W(l_1) = 0 \quad (3.30)$$

where, using the result stated in Proposition 3.31, it can be shown that

$$g_W(z) = \frac{-\lambda + \lambda c + z - \sqrt{\lambda^2 - 2\lambda^2 c - 2z\lambda + \lambda^2 c^2 - 2cz\lambda + z^2}}{2cz\lambda} \quad (3.31)$$

defined for all $z \notin [\lambda(1 - \sqrt{c})^2, \lambda(1 + \sqrt{c})^2]$ for all $c \in (0, \infty)$ and $z \neq 0$ for $c \geq 1$. Solving (3.30) gives us

$$l_1 = \lambda_1 \left(1 + \frac{\lambda c}{\lambda_1 - \lambda} \right). \quad (3.32)$$

To determine if this value of l_1 does indeed correspond to the minimum, we need to evaluate the derivative of the left hand side of (3.30) with respect to l_1 at the value given by (3.32). Symbolic manipulation using Maple yields the expression

$$\left. \frac{\partial^2}{\partial l_1^2} V_S(l_1 | l_2, \dots, l_n) \right|_{l_1 = \lambda_1 \left(1 + \frac{\lambda c}{\lambda_1 - \lambda} \right)} = \frac{(\lambda - \lambda_1)^2}{(\lambda^2 - 2\lambda_1 \lambda - \lambda^2 c + \lambda_1^2) \lambda_1^2 c} \quad (3.33)$$

which is positive iff $\lambda_1 > \lambda(1 + \sqrt{c})$ or $\lambda_1 < \lambda(1 - \sqrt{c})$. Since, by definition, $\lambda_1 > \lambda$, this implies that the equilibrium (equivalently, the logarithmic energy minimizing) configuration is described by (3.32) only when $\lambda_1 > \lambda(1 + \sqrt{c})$.

When $\lambda < \lambda_1 < \lambda(1 + \sqrt{c})$, (3.33) is negative so that l_1 given by (3.32) is a (local) *energy maximizing* configuration. The n -th particle, which starts out at $\lambda_1 > \lambda$ is unable to escape to infinity and hence minimizes its energy by “sliding down” towards the origin. However, it cannot get arbitrarily close to the origin because the equilibrium configuration of the 2-nd particle, as implied by Proposition 3.31 will, with high probability, be in a small neighborhood about $\lambda(1 + \sqrt{c})^2$. Hence, for large n , when $\lambda < \lambda_1 < \lambda(1 + \sqrt{c})$, the equilibrium configuration of the n -th particle will also be in a small neighborhood of $\lambda(1 + \sqrt{c})^2$. Thus a *phase transition phenomenon* occurs so that (asymptotically) the largest eigenvalue of the SCM is distinct from $\lambda(1 + \sqrt{c})^2$ only when the signal eigenvalues are greater than a certain threshold. This result is stated next in a more general setting, including the case when there are multiple signals, thereby lending credibility to the heuristic approximations and arguments we employed in our derivations.

Proposition 3.41. *Let $\widehat{\mathbf{R}}$ denote a sample covariance matrix formed from an $n \times m$ matrix of Gaussian observations whose columns are independent of each other and identically distributed with mean $\mathbf{0}$ and covariance \mathbf{R} . Denote the eigenvalues of \mathbf{R} by $\lambda_1 \geq \lambda_2 > \dots \geq \lambda_k > \lambda_{k+1} = \dots = \lambda_n = \lambda$. Let l_j denote the j -th largest eigenvalue of $\widehat{\mathbf{R}}$. Then as $n, m \rightarrow \infty$ with $c_m = n/m \rightarrow c \in (0, \infty)$,*

$$l_j \rightarrow \begin{cases} \lambda_j \left(1 + \frac{\lambda c}{\lambda_j - \lambda}\right) & \text{if } \lambda_j > \lambda(1 + \sqrt{c}) \\ \lambda(1 + \sqrt{c})^2 & \text{if } \lambda_j \leq \lambda(1 + \sqrt{c}) \end{cases} \quad (3.34)$$

where the convergence is almost surely.

PROOF. This result appears in [12] for very general settings. A matrix theoretic proof for when $c < 1$ for the real case may be found in [70] while a determinantal proof for the complex case may be found in [11]. \square

■ 3.4.3 Gaussian fluctuations of largest (“signal”) eigenvalues

Proposition 3.42. *Assume that $\widehat{\mathbf{R}}$ and \mathbf{R} satisfy the hypotheses of Proposition 3.41. If $\lambda_j > \lambda(1 + \sqrt{c})$ has multiplicity 1 and if $c_m = n/m \rightarrow c$ as $n, m \rightarrow \infty$ then*

$$s = \sqrt{n} \left(l_j - \lambda_j \left(1 + \frac{\lambda c}{\lambda_j - \lambda}\right) \right) \rightarrow g \sim \mathcal{N} \left(0, \frac{2}{\beta} \sigma_j^2 \right) \quad (3.35)$$

where the convergence in distribution is almost surely and

$$\sigma_j^2 = \lambda_j^2 \left(1 - \frac{c}{(\lambda_j - \lambda)^2} \right) \quad (3.36)$$

PROOF. A matrix theoretic proof for when $c < 1$ for the real case may be found in [70] while a determinantal proof for the complex case may be found in [11]. The result has been strengthened by Baik and Silverstein [81] for general $c \in (0, \infty)$. \square

■ 3.5 Estimating the number of signals

The key idea behind the proposed estimators can be succinctly summarized: When k signals are present and assuming $k \ll n$, then the fluctuations in the “noise” eigenvalues are not affected by the “signal” eigenvalues. Hence, “deviations” (on the $1/n^2$ scale) of the sample moments of subsets of sample eigenvalues subject to a criterion that penalizes overfitting of the model order should provide a good estimate of the number of signals. The Akaike Information Criterion (AIC) is applied to the noise eigenvalue fluctuations to obtain the relevant estimator.

■ 3.5.1 Akaike's Information Criterion

Given N observations $\mathbf{Y} = [\mathbf{y}(1), \dots, \mathbf{y}(N)]$ and a family of models, or equivalently a parameterized family of probability densities $f(\mathbf{Y}|\boldsymbol{\theta})$ indexed by the parameter vector $\boldsymbol{\theta}$, we select the model which gives the minimum AIC [2] defined by

$$\text{AIC}_k = -2 \log f(\mathbf{Y}|\hat{\boldsymbol{\theta}}) + 2k \quad (3.37)$$

where $\hat{\boldsymbol{\theta}}$ is the maximum likelihood estimate of $\boldsymbol{\theta}$, and k is the number of free parameters in $\boldsymbol{\theta}$. The idea behind this is that the AIC, given by (3.37), is an unbiased estimate of the mean Kullback-Liebler distance between the modelled density $f(\mathbf{Y}|\boldsymbol{\theta})$ and the estimated density $f(\mathbf{Y}|\hat{\boldsymbol{\theta}})$. We apply Akaike's information criteria on the fluctuations of the “noise” eigenvalues to detect the number of signals. The estimators presented, in effect, treat large departures (on the $1/n^2$ scale) of the sample moments of subsets of sample eigenvalues as reflecting the presence of a signal.

■ 3.5.2 Unknown noise variance

When the noise variance is unknown, the parameter vector of the model, denoted by $\boldsymbol{\theta}_k$, is given by

$$\boldsymbol{\theta}_k = [\lambda_1, \dots, \lambda_k, \sigma^2]'. \quad (3.38)$$

The number of free parameters in $\boldsymbol{\theta}_k$ is then $k + 1$. The maximum likelihood estimate of the noise variance is given by

$$\hat{\sigma}_{(k)}^2 = \frac{1}{n-k} \sum_{i=k+1}^n l_i \quad (3.39)$$

where $l_1 \geq \dots \geq l_n$ are the eigenvalues of $\hat{\mathbf{R}}$. We regard the observation y as

$$y = \sum_{i=k+1}^n \left(\frac{l_i}{\hat{\sigma}_{(k)}^2} \right)^2 - \left(1 + \frac{n}{m} \right) n + \left(\frac{2}{\beta} - 1 \right) \frac{n}{m} \quad (3.40)$$

where $\beta = 1$ (or 2) when the snapshots are real (or complex). The fluctuations of the $n - k$ smallest (“noise”) eigenvalues do not depend on the “signal” eigenvalues. The log-likelihood function is given by

$$-\log f(y|\hat{\boldsymbol{\theta}}) = \frac{y^2}{2q^2} + \frac{1}{2} \log 2\pi q^2 \quad (3.41)$$

where

$$q^2 = \frac{4}{\beta} \frac{n}{m} \left(2 \frac{n^2}{m^2} + 5 \frac{n}{m} + 2 \right) \quad (3.42)$$

Substituting this into (3.37) followed by some straightforward algebraic manipulations yields the criterion listed in the lower panel of Table 3.1(a).

■ 3.5.3 Known noise variance

When the noise variance, σ^2 , is known then the parameter vector of the model is given by


$$\boldsymbol{\theta}_k = [\lambda_1, \dots, \lambda_k]'. \quad (3.43)$$

The number of free parameters in $\boldsymbol{\theta}_k$ is then k . We regard the observation vector \mathbf{y} as

$$\mathbf{t}_k = \begin{bmatrix} \sum_{i=k+1}^n (l_i/\sigma^2) - n \\ \sum_{i=k+1}^n (l_i/\sigma^2)^2 - n \left(1 + \frac{n}{m} \right) - \left(\frac{2}{\beta} - 1 \right) \frac{n}{m} \end{bmatrix} \quad (3.44)$$

where $\beta = 1$ (or 2) when the snapshots are real (or complex). The fluctuations of the $n - k$ smallest (“noise”) eigenvalues do not depend on the “signal” eigenvalues. Following the procedure described earlier to obtain $\log f(\mathbf{t}_k|\boldsymbol{\theta})$ and substituting that expression into (3.37) yields the criterion listed in the upper panel of Table 3.1(a) with the overfitting penalty of $2k$ instead of $2(k + 1)$ as listed. Our usage of $2(k + 1)$ instead is motivated by aesthetic considerations.

Figure 3-5 shows sample realizations of the score function illustrating how large

(a) Complex signals: $n = 16, m = 32$.(b) Complex signals: $n = 32, m = 64$.**Figure 3-5.** Sample realizations of the proposed criterion when $\lambda_1 = 10, \lambda_2 = 3$ and $\lambda_3 = \dots = \lambda_n = 1$.

departures (on the $1/n^2$ scale) of the sample moments of subsets of sample eigenvalues when appropriately penalized can yield an accurate estimate of the number of signals present.

■ 3.6 Extensions to frequency domain and vector sensor arrays

When the m snapshot vectors $\mathbf{x}_i(w_j)$ for $j = 1, \dots, m$ represent Fourier coefficients vectors at frequency w_j then the sample covariance matrix

$$\hat{\mathbf{R}}(w_j) = \frac{1}{m} \sum_{i=1}^m \mathbf{x}_i(w_j) \mathbf{x}_i(w_j)' \quad (3.45)$$

is the periodogram estimate of the spectral density matrix at frequency w_j . The time-domain approach carries over to the frequency domain so that the estimators in Table 3.1 remain applicable with $l_i \equiv l_i(w_j)$ where $l_1(w_j) \geq l_2(w_j) \geq \dots \geq l_n(w_j)$ are the eigenvalues of $\hat{\mathbf{R}}(w_j)$.

When the signals are wideband and occupy M frequency bins, denoted by w_1, \dots, w_M ,

then the information on the number of signals present is contained in all the bins. The assumption that the observation time is much larger than the correlation times of the signals (sometimes referred to as the SPLOT assumption - stationary process, long observation time) ensures that the Fourier coefficients corresponding to the different frequencies are statistically independent.

Thus the AIC based criterion for detecting the number of wideband signals that occupy the frequency bands w_1, \dots, w_M , as given in Table 3.1(b), is obtained by summing the corresponding criterion in Table 3.1(a) over the frequency range of interest. Generically, we expect to use $\beta = 2$, representing the usual complex frequency domain representation, for the wideband frequency domain signal estimators. When the number of snapshots is severely constrained, the SPLOT assumption is likely to be violated so that the Fourier coefficients corresponding to different frequencies will not be statistically independent. This will likely degrade the performance of the proposed estimators.

When the measurement vectors represent quaternion valued narrowband signals, then $\beta = 4$ so that the estimators in Table 3.1(a) can be used. Quaternion valued vectors arise when the data collected from vector sensors is represented using quaternions as in [65].

Known Noise Variance	$\mathbf{t}_k = \begin{bmatrix} \sum_{i=k+1}^n (l_i/\sigma^2) - n \\ \sum_{i=k+1}^n (l_i/\sigma^2)^2 - n \left(1 + \frac{n}{m}\right) - \left(\frac{2}{\beta} - 1\right) \frac{n}{m} \end{bmatrix}$ $\ell_k = \frac{\beta}{2} \mathbf{t}'_k \begin{bmatrix} \frac{n}{m} & 2\frac{n}{m} \left(\frac{n}{m} + 1\right) \\ 2\frac{n}{m} \left(\frac{n}{m} + 1\right) & 2\frac{n}{m} \left(2\frac{n^2}{m^2} + 5\frac{n}{m} + 2\right) \end{bmatrix}^{-1} \mathbf{t}_k$
Unknown Noise Variance	$t_k = \left[(n-k) \frac{\sum_{i=k+1}^n l_i^2}{\left(\sum_{i=k+1}^n l_i\right)^2} - \left(1 + \frac{n}{m}\right) \right] n - \left(\frac{2}{\beta} - 1\right) \frac{n}{m}$ $\ell_k = \frac{\beta}{2} \frac{t_k^2}{2\frac{n}{m} \left(2\frac{n^2}{m^2} + 5\frac{n}{m} + 2\right)}$
	Number of Signals: $\hat{k} = \arg \min_{k \in \mathbb{N}: 0 \leq k < \min(m, n)} \ell_k + 2(k+1)$

(a) Time domain or narrow band frequency domain signals: $\beta = 1$ (or 2) when signals are real (or complex).

Known Noise Variance	$\mathbf{t}_{j,k} = \begin{bmatrix} \sum_{i=k+1}^n (l_i(w_j)/\sigma^2(w_j)) - n \\ \sum_{i=k+1}^n (l_i(w_j)/\sigma^2(w_j))^2 - n \left(1 + \frac{n}{m}\right) - \left(\frac{2}{\beta} - 1\right) \frac{n}{m} \end{bmatrix}$ $\ell_k = \sum_{j=1}^M \frac{\beta}{2} \mathbf{t}'_{j,k} \begin{bmatrix} \frac{n}{m} & 2\frac{n}{m} \left(\frac{n}{m} + 1\right) \\ 2\frac{n}{m} \left(\frac{n}{m} + 1\right) & 2\frac{n}{m} \left(2\frac{n^2}{m^2} + 5\frac{n}{m} + 2\right) \end{bmatrix}^{-1} \mathbf{t}_{j,k}$
Unknown Noise Variance	$t_{j,k} = \left[(n-k) \frac{\sum_{i=k+1}^n l_i(w_j)^2}{\left(\sum_{i=k+1}^n l_i(w_j)\right)^2} - \left(1 + \frac{n}{m}\right) \right] n - \left(\frac{2}{\beta} - 1\right) \frac{n}{m}$ $\ell_k = \sum_{j=1}^M \frac{\beta}{2} \frac{t_{j,k}^2}{2\frac{n}{m} \left(2\frac{n^2}{m^2} + 5\frac{n}{m} + 2\right)}$
	Number of Signals: $\hat{k} = \arg \min_{k \in \mathbb{N}: 0 \leq k < \min(m, n)} \ell_k + 2M(k+1)$

(b) Wideband signals occupying M frequency bins: $\beta = 1$ (or 2) when signals are real (or complex).

Table 3.1. Estimating the number of signals from the eigenvalues of a SCM formed from m snapshots.

■ 3.7 Effective number of signals and consistency of the criteria

Theorem 3.71. *Let \mathbf{R} and $\tilde{\mathbf{R}}$ be two $n \times n$ sized covariance matrices whose eigenvalues are related as*

$$\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_p, \lambda_{p+1}, \dots, \lambda_k, \lambda, \dots, \lambda) \quad (3.46a)$$

$$\tilde{\mathbf{\Lambda}} = \text{diag}(\lambda_1, \dots, \lambda_p, \lambda, \dots, \lambda) \quad (3.46b)$$

where for some $c \in (0, \infty)$, and all $i = p + 1, \dots, k$, $\lambda_i \leq \lambda(1 + \sqrt{c})$. Let $\hat{\mathbf{R}}$ and $\tilde{\hat{\mathbf{R}}}$ be the associated sample covariance matrices formed from m snapshots. Then for every $n, m(n) \rightarrow \infty$ such that $c_m = n/m \rightarrow c$,

$$\text{Prob}(\hat{k} = j | \mathbf{R}) \rightarrow \text{Prob}(\hat{k} = j | \tilde{\mathbf{R}}) \quad \text{for } j = 1, \dots, p \quad (3.47a)$$

and

$$\text{Prob}(\hat{k} > p | \mathbf{R}) \rightarrow \text{Prob}(\hat{k} > p | \tilde{\mathbf{R}}) \quad (3.47b)$$

where the convergence is almost surely and \hat{k} is the estimate of the number of signals using obtained the algorithms summarized in Table 3.1(a).

PROOF. The theorem follows from Proposition 3.41. The almost sure convergence of the sample eigenvalues $l_j \rightarrow \lambda(1 + \sqrt{c})^2$ for $j = p + 1, \dots, k$ implies that i -th largest eigenvalues of $\hat{\mathbf{R}}$ and $\tilde{\hat{\mathbf{R}}}$ converge to the same limit almost surely. The fluctuations about this limit will hence be identical so that (3.47) follows in the asymptotic limit. \square

Note that the rate of convergence to the asymptotic limit for $\text{Prob}(\hat{k} > p | \mathbf{R})$ and $\text{Prob}(\hat{k} > p | \tilde{\mathbf{R}})$ will, in general, depend on the eigenvalue structure and may be arbitrarily slow. Thus, Theorem 3.71 yields no insight into rate of convergence type issues which are important in practice. Rather, the theorem is a statement on the asymptotic equivalence, from an identifiability point of view, of sequences of sample covariance covariances which are related in the manner described. At this point, we are unable to prove the consistency of the proposed estimators as this would require more a refined analysis that characterizes the fluctuations of subsets of the (ordered) “noise” eigenvalues. The statement regarding consistency of the proposed estimator is presented as a conjecture with numerical simulations used as (non-definitive) evidence.

Conjecture 3.72. *Let \mathbf{R} be a $n \times n$ covariance matrix that satisfies the hypothesis of Proposition 3.41. Let $\hat{\mathbf{R}}$ be a sample covariance matrix formed from m snapshots. Define*

$$k_{\text{eff}}(c | \mathbf{R}) = \text{Number of eigenvalues of } \mathbf{R} > \lambda(1 + \sqrt{c}). \quad (3.48)$$

Then in $m, n \rightarrow \infty$ limit with $c_m = n/m \rightarrow c$, \hat{k} is a consistent estimator of $k_{\text{eff}}(c | \mathbf{R})$ where \hat{k} is the estimate of the number of signals obtained using the algorithms summarized in Table 3.1(a).

■ 3.7.1 The asymptotic identifiability of two closely spaced signals

Suppose there are two uncorrelated Gaussian (hence, independent) signals whose covariance matrix is the diagonal matrix $\mathbf{R}_s = \text{diag}(\sigma_{S1}^2, \sigma_{S2}^2)$. In (3.1) and (3.3), let $\mathbf{A} = [\mathbf{v}_1 \mathbf{v}_2]$. In a sensor array processing application, we think of $\mathbf{v}_1 \equiv \mathbf{v}(\theta_1)$ and $\mathbf{v}_2 \equiv \mathbf{v}(\theta_2)$ as encoding the array manifold vectors for a source and an interferer with powers σ_{S1}^2 and σ_{S2}^2 , located at θ_1 and θ_2 , respectively. The covariance matrix given by

$$\mathbf{R} = \sigma_{S1}^2 \mathbf{v}_1 \mathbf{v}_1' + \sigma_{S2}^2 \mathbf{v}_2 \mathbf{v}_2' + \sigma^2 \mathbf{I} \quad (3.49)$$

has the $n - 2$ smallest eigenvalues $\lambda_3 = \dots = \lambda_n = \sigma^2$ and the two largest eigenvalues

$$\lambda_1 = \sigma^2 + \frac{(\sigma_{S1}^2 \|\mathbf{v}_1\|^2 + \sigma_{S2}^2 \|\mathbf{v}_2\|^2)}{2} + \frac{\sqrt{(\sigma_{S1}^2 \|\mathbf{v}_1\|^2 - \sigma_{S2}^2 \|\mathbf{v}_2\|^2)^2 + 4\sigma_{S1}^2 \sigma_{S2}^2 |\langle \mathbf{v}_1, \mathbf{v}_2 \rangle|^2}}{2} \quad (3.50a)$$

$$\lambda_2 = \sigma^2 + \frac{(\sigma_{S1}^2 \|\mathbf{v}_1\|^2 + \sigma_{S2}^2 \|\mathbf{v}_2\|^2)}{2} - \frac{\sqrt{(\sigma_{S1}^2 \|\mathbf{v}_1\|^2 - \sigma_{S2}^2 \|\mathbf{v}_2\|^2)^2 + 4\sigma_{S1}^2 \sigma_{S2}^2 |\langle \mathbf{v}_1, \mathbf{v}_2 \rangle|^2}}{2} \quad (3.50b)$$

respectively. Applying the result in Proposition 3.41 allows us to express the effective number of signals as

$$k_{\text{eff}} = \begin{cases} 2 & \text{if } \sigma^2 \left(1 + \sqrt{\frac{n}{m}}\right) < \lambda_2 \\ 1 & \text{if } \lambda_2 < \sigma^2 \left(1 + \sqrt{\frac{n}{m}}\right) \leq \lambda_1 \\ 0 & \text{if } \lambda_1 \leq \sigma^2 \left(1 + \sqrt{\frac{n}{m}}\right) \end{cases} \quad (3.51)$$

In the special situation when $\|\mathbf{v}_1\| = \|\mathbf{v}_2\| = \|\mathbf{v}\|$ and $\sigma_{S1}^2 = \sigma_{S2}^2 = \sigma_S^2$, we can (in an asymptotic sense) reliably detect the presence of *both signals* from the sample eigenvalues alone whenever

Asymptotic identifiability condition : $\sigma_S^2 \|\mathbf{v}\|^2 \left(1 - \frac{|\langle \mathbf{v}_1, \mathbf{v}_2 \rangle|}{\|\mathbf{v}\|}\right) > \sigma^2 \sqrt{\frac{n}{m}}$

(3.52)

Equation (3.52) captures the tradeoff between the identifiability of two closely spaced signals, the dimensionality of the system, the number of available snapshots and the cosine of the angle between the vectors \mathbf{v}_1 and \mathbf{v}_2 . It may prove to be a useful heuristic for experimental design.

■ 3.8 Numerical simulations

■ 3.8.1 Test of sphericity

When $\mathbf{R} = \sigma^2 \mathbf{I}$, so that $k = 0$, we evaluate the performance of the proposed estimators by examining the probability that $\hat{k} = 0$ over 20,000 Monte-Carlo trials. Figure 3-7 compares the empirical results as a function of the number of snapshots, for different values of n . The Wax-Kailath estimators always over-estimate the number of signals for the sample sizes considered. In contrast, the proposed algorithms correctly predict the number of signals more than 90% of the time even when the dimensionality of the system is as small as 8. The simulations suggest large sample consistency of the estimators, with the complex signal case exhibiting a faster rate of convergence than the real signal case, as expected from Proposition 3.32. The characteristic U -shape of the performance curve appears because the noise eigenvalue fluctuations of the SCM with $n = 32$, $m = 16$ are identical to that of an SCM with $n = 16$, $m = 32$.

The superior detection performance of the estimator in the unknown noise variance scenario when $\mathbf{R} = \sigma^2 \mathbf{I}$ comes as no surprise since its criterion involves comparing the fluctuation of just a single moment of the SCM. Given the inherent symmetries of the null hypothesis, the degradation in the performance of the estimator when the noise variance is unknown and must be estimated will only be revealed in tests involving the detection of $k > 0$ signals.

■ 3.8.2 Illustration of effective number of signals concept

Consider the detection problem on a covariance matrix \mathbf{R} with $n - 2$ eigenvalues of magnitude $\sigma^2 = 1$, $\lambda_1 = 10$ and $\lambda_2 = 3$. Figure 3-8 compares the empirical probability (over 20,000 Monte-Carlo trials) of detecting 2 signals for the proposed estimators for a range of values of n and m . *The empirical probability of Wax-Kailath estimators detecting 2 signals over these trials is identically zero.* Note how the complex valued case performs better than the real valued case for the same (n, m) pair. This rate-of-convergence type effect is expected given the behavior of the associated fluctuations in Proposition 3.32. The simulation suggest that the estimators exhibit large sample consistency with a faster rate of convergence for complex signals than real signals. Figure 3-9 compares the performance of the two estimators; the case where the noise variance is known performs better, as expected.

Parsing the empirical data differently allows us to illustrate the relevance of effective number of signals concept in light of the discussion in Section 3.7. For the covariance matrix considered, when $n = 4m$, $c_m = 4$ so that the effective number of signals, from (3.48), equals 1. Figure 3-10 compares the empirical probability of detecting zero and one signals, when the signals are real and complex, for different values of n with $m = n/4$. The simulations illustrate the consistency of the proposed estimators in the $n, m(n) \rightarrow \infty$ limit, with respect to the effective number of signals, as conjectured.

The rate of convergence to the asymptotic result is faster for the complex signals case, as before. The probability of detecting zero signals decays to zero as the probabil-

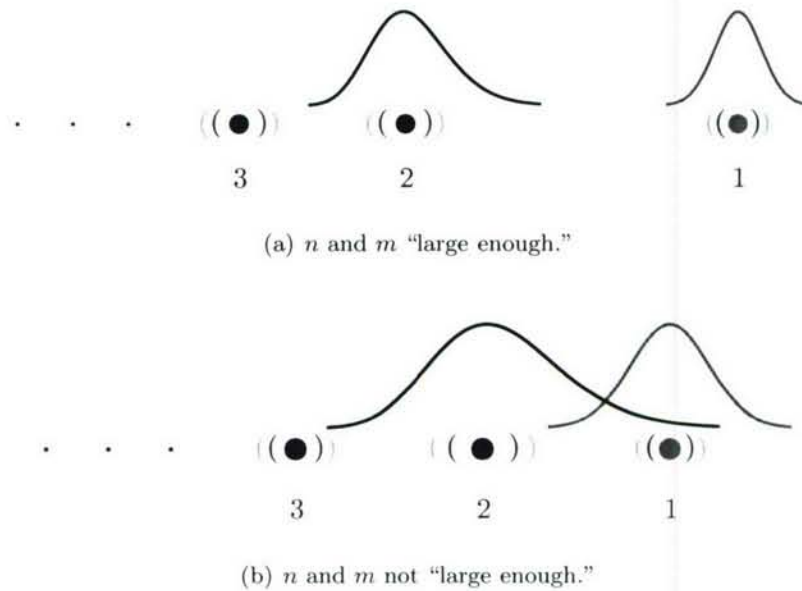


Figure 3-6. "Signal" and "noise" eigenvalue fluctuations induced limits on signal identifiability.

ity of detecting one signal approaches zero. This highlights the relevance of the effective number of signals concept in high-dimensional settings. In moderate dimensional settings, the fluctuations of the signal eigenvalues, when combined with the concept effective rank, best capture the inherent difficulties in the formulation of the detection problem.

For the example considered, when the signal is complex, the largest (and only) signal eigenvalue fluctuates about $10(1 - 4/9) \approx 14.4$ with a variance, given by Proposition 3.42, approximately equal to $10^2(1 - 4/9^2)/n \approx 95/n$. The largest noise eigenvalue fluctuates about $(1 + \sqrt{(4)})^2 = 9$. Reliable detection of the effective number of signals occurs in Figure 3-10 for values of n large enough that the separation between the signal eigenvalue and the largest noise eigenvalue is roughly 6–7 times the variance of the signal eigenvalue fluctuation (as in Figure 3-6(a)). For values of n smaller than that, the signal eigenvalue is insufficiently separated from the noise eigenvalue to be identified as such (as in Figure 3-6(b)). In this context, moderate dimensionality is a greater curse than high-dimensionality because the fluctuations of the signal and noise eigenvalues make the signal versus noise decidability issue even more challenging.


■ 3.9 Future work

We have developed an approach for detecting the number of signals in white noise from the sample eigenvalues alone. The proposed estimators explicitly take into account the blurring of the sample eigenvalues due to the finite size and are hence able to outperform


traditional estimators that do not exploit this information.

In principle, we could have formulated our algorithms to consider the fluctuations of the “signal” eigenvalues as well instead of focussing on the fluctuations of the “noise” eigenvalues alone. Such an algorithm would be computationally more complex because we would have to first obtain the maximum likelihood estimate of the signal eigenvalue using the results in Proposition 3.42. Since the distribution of the signal eigenvalue depends on its multiplicity, formulating the problem in terms of a joint signal and noise eigenvalue estimation-detection framework implies that the practitioner would be forced to make subjective assumptions on the multiplicity of the individual signal eigenvalues. This makes such a formulation less desirable than the noise eigenvalue only solution proposed.


Future work would be to analytically prove the conjecture stated regarding the consistency of the algorithms. It would be of value to compare the present AIC based formulation to the MDL/BIC based formulation for the proposed algorithms. It remains an open question to analyze and design such eigenvalue based signal detection algorithms in the Neyman-Pearson sense, *i.e.*, finding the most powerful test that does not exceed a threshold probability of false detection. Finer properties, perhaps buried in the rate of convergence to the asymptotic results used, might be useful in this context. Such estimators will require practitioners to set thresholds. Though this is something we instinctively shy away from, if the performance can be significantly improved then this might be a price we might be ready to pay, especially for the detection of low-level signals right around the threshold where the phase transition phenomenon kicks in.



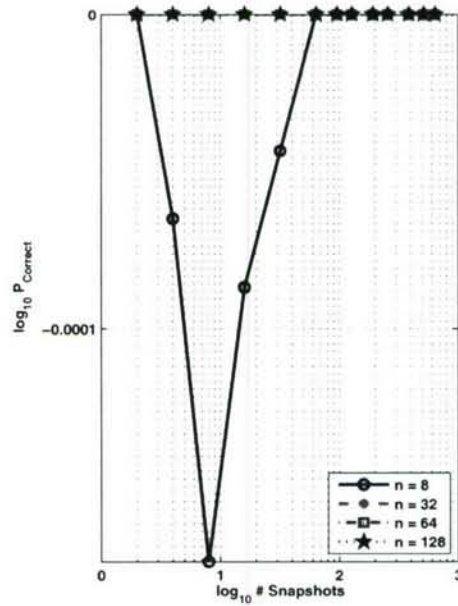
(a) Real signals: known noise variance.



(b) Complex signals: known noise variance.




(c) Real signals: unknown noise variance.




(d) Complex signals: unknown noise variance.


Figure 3-7. Performance of proposed estimators when there are zero signals in white noise. Note that $10^{-0.01} \approx 0.9772$ while $10^{-0.04} \approx 0.9120$.




(a) Real signals: known noise variance.



(b) Complex signals: known noise variance.

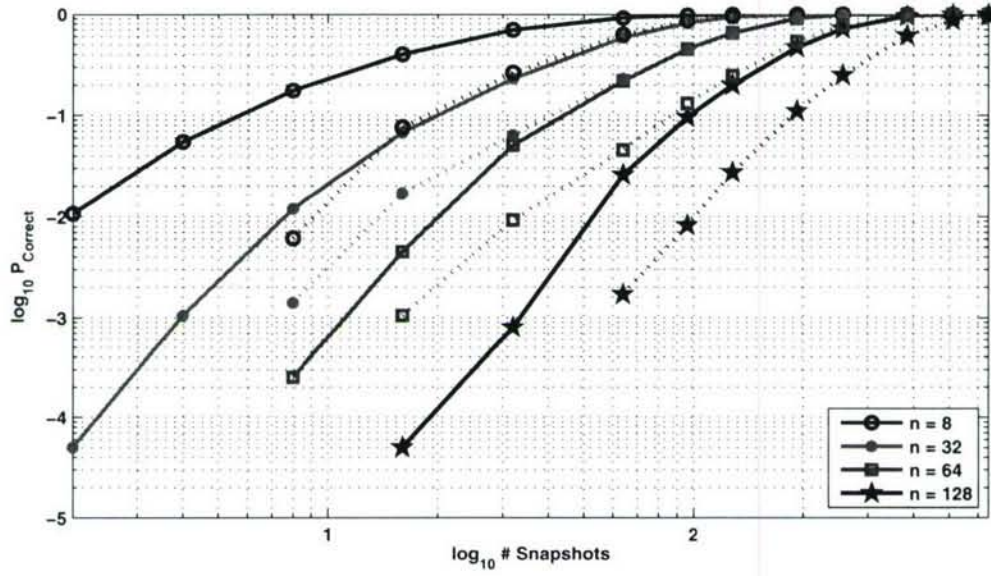


(c) Real signals: unknown noise variance.

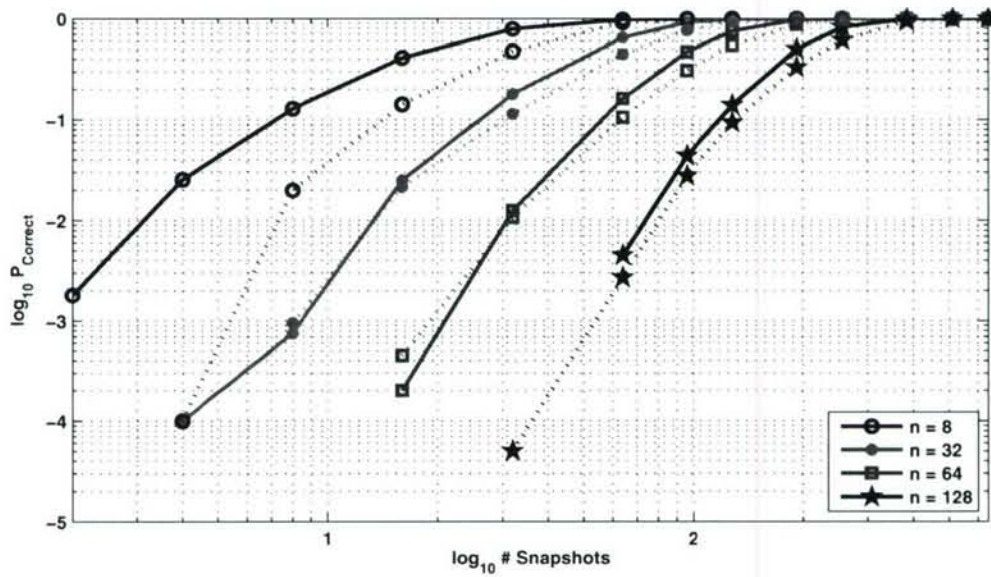


(d) Complex signals: unknown noise variance.

Figure 3-8. Performance of proposed estimators when there are 2 signals in white noise.




(a) Real signals.




(b) Complex signals


Figure 3-9. Comparison of the known vs. unknown estimators.




(a) Real signals: 1 signal hypothesis.



(b) Real signals: 0 signals hypothesis.



(c) Complex signals: 1 signal hypothesis.



(d) Complex signals: 0 signals hypothesis.

Figure 3-10. Comparison of the known vs. unknown estimators.

Statistical eigen-inference: Large Wishart matrices

In this chapter we expand the inference methodologies developed in Chapter 3 to a broader class of covariance matrices. This chapter is organized as follows. We motivate the problem in Section 4.1 and preview the structure of the proposed algorithms summarized in Table 4.1. In Section 4.2 we introduce the necessary definitions and summarize the relevant random matrix theorems that we exploit. Concrete algorithms for computing the analytic expectations that appear in the algorithms (summarized in Table 4.1) are presented in Section 4.3. The eigen-inference techniques are developed in Section 4.4. The performance of the algorithms is illustrated using Monte-Carlo simulations in Section 4.5. Some concluding remarks are presented in Section 4.7.

■ 4.1 Problem formulation

Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_m]$ be a $n \times m$ data matrix where $\mathbf{x}_1, \dots, \mathbf{x}_m$, denote m independent measurements, where for each i , \mathbf{x}_i has an n -dimensional (real or complex) Gaussian distribution with mean zero, and positive definite covariance matrix Σ . The sample covariance matrix (SCM) when formed from these m samples as

$$\mathbf{S} := \frac{1}{m} \sum_{i=1}^m \mathbf{x}_i \mathbf{x}_i' = \frac{1}{m} \mathbf{X} \mathbf{X}', \quad (4.1)$$

has the (central) Wishart distribution [114]. We focus on inference problems for parameterized covariance matrices modelled as $\Sigma_{\theta} = \mathbf{U} \Lambda_{\theta} \mathbf{U}'$ where

$$\Lambda_{\theta} = \begin{bmatrix} a_1 \mathbf{I}_{n_1} & & & \\ & a_2 \mathbf{I}_{n_2} & & \\ & & \ddots & \\ & & & a_k \mathbf{I}_{n_k} \end{bmatrix}, \quad (4.2)$$

where $a_1 > \dots > a_k$ and $\sum_{j=1}^k n_j = n$. Defining $t_i = n_i/n$, allows us to conveniently express the $2k - 1$ dimensional parameter vector as $\theta = (t_1, \dots, t_{k-1}, a_1, \dots, a_k)$ with

the obvious non-negativity constraints on the elements.

Models of the form in (4.2) arise whenever the measurements are of the form

$$\mathbf{x}_i = \mathbf{A} \mathbf{s}_i + \mathbf{z}_i \quad \text{for } i = 1, \dots, m \quad (4.3)$$

where $\mathbf{z}_i \sim \mathcal{N}_n(0, \sigma^2 \mathbf{I})$, denotes an n -dimensional (real or complex) Gaussian noise vector where σ^2 is generically unknown, $\mathbf{s}_i \sim \mathcal{N}_k(\mathbf{0}, \mathbf{I})$, $\mathbf{s}_i \sim \mathcal{N}_k(\mathbf{0}, \mathbf{R}_s)$ denotes a k -dimensional (real or complex) Gaussian signal vector with covariance \mathbf{R}_s , and \mathbf{A} is a $n \times k$ unknown non-random matrix. In array processing applications, the j -th column of the matrix \mathbf{A} encodes the parameter vector associated with the j -th signal whose magnitude is described by the j -th element of \mathbf{s}_i .

Since the signal and noise vectors are independent of each other, the covariance matrix of \mathbf{x}_i can hence be decomposed as

$$\mathbf{\Sigma} = \mathbf{\Psi} + \mathbf{\Sigma}_z \quad (4.4)$$

where $\mathbf{\Sigma}_z$ is the covariance of \mathbf{z} and $\mathbf{\Psi} = \mathbf{A} \mathbf{\Sigma}_s \mathbf{A}'$ with $'$ denoting the conjugate transpose. One way of obtaining $\mathbf{\Sigma}$ with eigenvalues of the form in (4.2) was described in Chapter 3. When $\mathbf{\Sigma}_z = \sigma^2 \mathbf{I}$ so that the $n - k$ smallest eigenvalues of $\mathbf{\Sigma}$ are equal to σ^2 . Then, if the matrix \mathbf{A} is of full column rank so and the covariance matrix of the signals $\mathbf{\Sigma}_s$ is nonsingular, the $n - k$ smallest eigenvalues of $\mathbf{\Psi}$ are equal to zero so that the eigenvalues of $\mathbf{\Sigma}$ will be of the form in (4.2). Alternately, if the eigenvalues of $\mathbf{\Psi}$ and $\mathbf{\Sigma}_z$ have the identical subspace structure, *i.e.*, in (4.2), $t_i^\Psi = t_i^{\Sigma_z}$ for all i , then whenever the eigenvectors associated with each of the subspaces of $\mathbf{\Psi}$ and $\mathbf{\Sigma}_z$ align, the eigenvalues of $\mathbf{\Sigma}$ will have the subspace structure in (4.2).

■ 4.1.1 Inferring the population eigenvalues from the sample eigenvalues

While inference problems for these models have been documented in texts such as [67], the inadequacies of classical algorithms in high-dimensional, (relatively) small sample size settings have not been adequately addressed. We highlight some of the prevalent issues in the context of statistical inference and hypothesis testing.

Anderson's landmark paper [6] develops the theory that describes the (large sample) asymptotics of the sample eigenvalues (in the real valued case) for such models when the true covariance matrix has eigenvalues of arbitrary multiplicity. Indeed, for arbitrary covariance \mathbf{R} , the joint density function of the eigenvalues l_1, \dots, l_n of the SCM \mathbf{S} when $m > n + 1$ is shown to be given by

$$\tilde{Z}_{n,m}^\beta \sum_{i=1}^n l_i^{\beta(m-n+1)/2-1} \prod_{i < j}^n |l_i - l_j|^\beta \int_{\mathbf{Q}} \exp\left(-\frac{m\beta}{2} \text{Tr}(\mathbf{\Sigma}^{-1} \mathbf{Q} \mathbf{S} \mathbf{Q}')\right) d\mathbf{Q} \quad (4.5)$$

where $l_1 > \dots > l_n > 0$, $\tilde{Z}_{n,m}^\beta$ is a normalization constant, and $\beta = 1$ (or 2) when \mathbf{S} is real (resp. complex). In (4.5), $\mathbf{Q} \in \mathbf{O}(n)$ when $\beta = 1$ while $\mathbf{Q} \in \mathbf{U}(n)$ when $\beta = 2$

where $\mathbf{O}(n)$ and $\mathbf{U}(n)$ are, respectively, the set of $n \times n$ orthogonal and unitary matrices with Haar measure. Anderson notes that

If the characteristic roots of Σ are different, the deviations ... from the corresponding population quantities are asymptotically normally distributed. When some of the roots of Σ are equal, the asymptotic distribution cannot be described so simply.

Indeed, the difficulty alluded to, arises due to the presence of the integral over orthogonal (or unitary) group on the right hand side of (4.5). This problem is compounded in situations when some of the eigenvalues of Σ are equal as is the case for the model considered in (4.2). Nonetheless, Anderson is able to use the (large sample) asymptotics to derive the maximum likelihood estimate of the population eigenvalues, a_l , as

$$\hat{a}_l \approx \frac{1}{n_l} \sum_{j \in N_l} \hat{\lambda}_j \quad \text{for } l = 1, \dots, k, \quad (4.6)$$

where $\hat{\lambda}_j$ are the sample eigenvalues (arranged in descending order) and N_l is the set of integers $n_1 + \dots + n_{l-1} + 1, \dots, n_1 + \dots + n_l$. This is a reasonable estimator that works well in practice when $m \gg n$. The large sample size asymptotics are, however, of limited utility because they ignore the (significant) effect of the dimensionality of the system on the behavior of the sample eigenvalues.

Consequently, (large sample size) asymptotic predictions, derived under the n fixed, $m \rightarrow \infty$ regime do not account for the additional complexities that arise in situations where the sample size m is large but the dimensionality n is of comparable order. Furthermore, the estimators developed using the classical large sample asymptotics invariably become degenerate whenever $n < m$, so that $n - m$ of the sample eigenvalues will identically equal to zero. For example, when $m = n/2$, and there are two distinct population eigenvalues each with multiplicity $n/2$ then the estimate of the smallest eigenvalue using (4.6) will be zero. Other such scenarios where the population eigenvalue estimates obtained using (4.6) are meaningless are easy to construct and are practically relevant in many applications such as radar and sonar signal processing [90,102], and many more.

There are, of course, other strategies one may employ for inferring the population eigenvalues. One might consider a maximum-likelihood technique based on maximizing the log-likelihood function of the observed data X which is given by (ignoring constants)

$$l(\mathbf{X}|\Sigma) := -m(\text{tr} \mathbf{S}\Sigma^{-1} + \log \det \Sigma),$$

or, equivalently, when $\Sigma = \mathbf{U}\Lambda\mathbf{U}'$, by minimizing the objective function

$$h(\mathbf{X}|\mathbf{U}, \Lambda) = (\text{tr} \mathbf{S}\mathbf{U}\Lambda^{-1}\mathbf{U}' + \log \det \Lambda). \quad (4.7)$$

What should be apparent on inspecting (4.7) is that the maximum-likelihood esti-

mation of the parameters of $\mathbf{\Lambda}$ of the form in (4.2) requires us to model the population eigenvectors \mathbf{U} as well (except when $k = 1$). If \mathbf{U} were known apriori, then an estimate of a_l obtained as

$$\hat{a}_l \approx \frac{1}{n_l} \sum_{j \in N_l} (\mathbf{U}' \mathbf{S} \mathbf{U})_{j,j} \quad \text{for } l = 1, \dots, k. \quad (4.8)$$

N_l is the set of integers $n_1 + \dots + n_{l-1} + 1, \dots, n_1 + \dots + n_l$ will provide a good estimate. In practical applications, the population eigenvectors might either be unknown or be misspecified leading to faulty inference. Hence it is important to have the ability to perform statistically sound, computationally feasible *eigen-inference* of the population eigenvalues, *i.e.*, from the sample eigenvalues alone, in a manner that is robust to high-dimensionality and sample size constraints.

We illustrate the difficulties encountered in high-dimensional settings with an example (summarized in Figure 4-1) of a SCM constructed from a covariance matrix modelled as $\mathbf{\Sigma} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}'$ with $n = 100$ and sample size $m = 300$. Half of the eigenvalues of $\mathbf{\Lambda}$ are of magnitude 3 while the remainder are of magnitude 1. The sample eigenvalues are significantly *blurred*, relative to the true eigenvalues as shown in Figure 4-1(a). Figures 4-1(b), and 4-1(d) plot the sample eigenvectors for the case when the true eigenvectors $\mathbf{U} = \mathbf{I}$, and an arbitrary \mathbf{U} , respectively. Figures 4-1(c) and 4-1(e) plot the diagonal elements $(\mathbf{S})_{j,j}$. Thus, if the true eigenvector was indeed $\mathbf{U} = \mathbf{I}$ then an estimate of the population eigenvalues formed as in (4.8) yields a good estimate; when $\mathbf{U} \neq \mathbf{I}$, however, the estimate is very poor.

■ 4.1.2 Testing for equality of population eigenvalues

Similar difficulties are encountered in problems of testing as well. In such situations, Anderson proposes the likelihood ratio criterion for testing the hypothesis

$$\lambda_{n_1 + \dots + n_{l-1} + 1} = \lambda_{n_1 + \dots + n_{l-1} + 1, \dots, n_1 + \dots + n_l}$$

given by

$$\left[\prod_{j \in N_l} \hat{\lambda}_j / (n_k^{-1} \sum_{j \in N_l} \hat{\lambda}_j)^{n_k} \right]^{\frac{1}{2}m} \quad \text{for } l = 1, \dots, k, \quad (4.9)$$

where $\hat{\lambda}_j$ are the sample eigenvalues (arranged in descending order) and N_l is the set of integers $n_1 + \dots + n_{l-1} + 1, \dots, n_1 + \dots + n_l$. The test in (4.9) suffers from the same deficiency as the population eigenvalue estimator in (4.6) - it becomes degenerate when $n > m$. When the population eigenvectors \mathbf{U} are known, (4.9) may be modified by forming the criterion

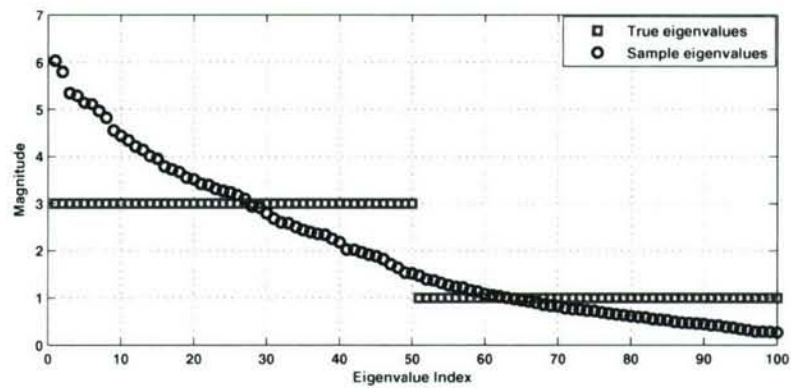
$$\left[\prod_{j \in N_l} (\mathbf{U}' \mathbf{S} \mathbf{U})_{j,j} / (n_k^{-1} \sum_{j \in N_l} (\mathbf{U}' \mathbf{S} \mathbf{U})_{j,j})^{n_k} \right]^{\frac{1}{2}m} \quad \text{for } l = 1, \dots, k, \quad (4.10)$$

where N_l is the set of integers $n_1 + \dots + n_{l-1} + 1, \dots, n_1 + \dots + n_l$. When the eigenvectors are misspecified the inference provided will be faulty. For the earlier example, Figure 4-1(e) illustrates this for the case when it is assumed that the population eigenvectors are \mathbf{I} when they are really $\mathbf{U} \neq \mathbf{I}$. Testing the hypothesis $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_0$, reduces to testing the null hypothesis $\boldsymbol{\Sigma} = \mathbf{I}$ when the transformation $\tilde{\mathbf{x}}_i = \boldsymbol{\Sigma}_0^{-1/2} \mathbf{x}_i$ is applied. The robustness of tests for sphericity in high dimensional settings has been extensively discussed in [55].

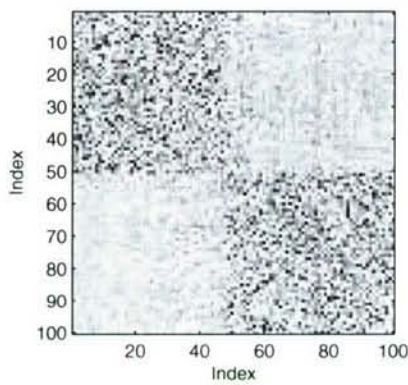
■ 4.1.3 Proposed statistical eigen-inference techniques

In this chapter our focus is on developing population eigenvalue estimation and testing algorithms for models of the form in (4.2) that are robust to high-dimensionality, sample size constraints and population eigenvector misspecification. We are able to develop such computationally feasible algorithms by exploiting the properties of the eigenvalues of large Wishart matrices. These results, analytically describe the non-random *blurring* of the sample eigenvalues, relative to the population eigenvalues, in the $n, m(n) \rightarrow \infty$ limit while compensating for the random *fluctuations* about the limiting behavior due to finite dimensionality effects. This allows us to handle the situation where the sample eigenvalues are blurred to the point that the block subspace structure of the population eigenvalues cannot be visually discerned, as in Figure 4-1(a), thereby extending the “signal” detection capability beyond the special cases tackled in [88]. The nature of the mathematics being exploited makes them robust to the high-dimensionality and sample size constraints while the reliance on the sample eigenvalues alone makes them insensitive to any assumptions on the population eigenvectors. In such situations where the eigenvectors are accurately modelled, the practitioner may use the proposed methodologies to complement and “robustify” the inference provided by estimation and testing methodologies that exploit the eigenvector structure.


We consider testing the hypothesis for the equality of the population eigenvalues and statistical inference about the population eigenvalues. In other words, for some unknown \mathbf{U} , if $\boldsymbol{\Sigma}_0 = \mathbf{U}\boldsymbol{\Lambda}_{\boldsymbol{\theta}_0}\mathbf{U}'$ where $\boldsymbol{\Lambda}_{\boldsymbol{\theta}}$ is modelled as in (4.2), techniques to 1) test if $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_0$, and 2) estimate $\boldsymbol{\theta}_0$ are summarized in Table 4.1. We note that inference on the population eigenvalues is performed using the *entire sample eigen-spectrum* unlike (4.6) and (4.9). This reflects the inherent non-linearities of the sample eigenvalue blurring induced by high-dimensionality and sample size constraints. An important implication of this in practice is that in high dimensional, sample size starved settings, inference performed on a subset of sample eigenvalues alone is likely to be inaccurate, or worse misleading. In such settings, practitioners are advised to consider tests (such as the ones proposed) for the equality of the entire population eigen-spectrum instead of testing for the equality of individual population eigenvalues.




(a) Sample eigenvalues versus true eigenvalues ($n = 100$, $m = 300$).




(b) Sample eigenvectors when $\mathbf{U} = \mathbf{I}$.



(c) Diagonal elements of \mathbf{S} when $\mathbf{U} = \mathbf{I}$.



(d) Sample eigenvectors for arbitrary \mathbf{U} .



(e) Diagonal elements of \mathbf{S} for arbitrary \mathbf{U} .

Figure 4-1. The challenge of estimating the population eigenvalues from the sample eigenvalues in high-dimensional settings.

Testing:	$H_{\theta_0} : h(\theta) := \mathbf{v}_\theta^T \mathbf{Q}_\theta^{-1} \mathbf{v}_\theta \sim \chi_2^2 \geq \gamma,$	$q = \dim(\mathbf{v}_\theta) = 2$
Estimation:	$\hat{\theta} = \arg \min_{\theta \in \Theta} \{ \mathbf{v}_\theta^T \mathbf{Q}_\theta^{-1} \mathbf{v}_\theta + \log \det \mathbf{Q}_\theta \},$	$q = \dim(\mathbf{v}_\theta) \geq \dim(\theta)$
Legend:	$(\mathbf{v}_\theta)_j = n \times \left(\frac{1}{n} \text{Tr} S^j - E \left[\frac{1}{n} \text{Tr} S^j \right] \right),$	$j = 1, \dots, q$
	$\mathbf{Q}_\theta = \text{cov} [\mathbf{v}_\theta \mathbf{v}_\theta']$	

Table 4.1. Structure of proposed algorithms.

■ 4.2 Preliminaries

Definition 4.21. Let $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ be an $N \times N$ matrix with real eigenvalues. The j -th sample moment is defined as

$$\text{tr}(\mathbf{A}^j) := \frac{1}{N} \text{Tr}(\mathbf{A}^j),$$

where Tr is the usual un-normalized trace.

Definition 4.22. Let $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ be a sequence of self-adjoint $N \times N$ -random matrices. If the limit of all moments defined as

$$\alpha_j^{\mathbf{A}} := \lim_{N \rightarrow \infty} E[\text{tr}(\mathbf{A}_N^j)] \quad (N \in \mathbb{N})$$

exists then we say that \mathbf{A} has a limit eigenvalue distribution.

Notation 4.23. For a random matrix \mathbf{A} with a limit eigenvalue distribution we denote by $M_{\mathbf{A}}(x)$ the moment power series, which we define by

$$M_{\mathbf{A}}(x) := 1 + \sum_{j \geq 1} \alpha_j^{\mathbf{A}} x^j.$$

Notation 4.24. For a random matrix ensemble \mathbf{A} with limit eigenvalue distribution we denote by $g_{\mathbf{A}}(x)$ the corresponding Cauchy-transform, which we define as formal power series by

$$g_{\mathbf{A}}(x) := \lim_{N \rightarrow \infty} E \left[\frac{1}{N} \text{Tr} (x \mathbf{I}_N - \mathbf{A}_N)^{-1} \right] = \frac{1}{x} M_{\mathbf{A}}(1/x).$$

Definition 4.25. Let $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ be a self-adjoint random matrix ensemble. We

say that it has a second order limit distribution if for all $i, j \in \mathbb{N}$ the limits

$$\alpha_j^A := \lim_{N \rightarrow \infty} c_1(\text{tr}(\mathbf{A}_N^j))$$

and

$$\alpha_{i,j}^A := \lim_{N \rightarrow \infty} c_2(\text{Tr}(\mathbf{A}_N^i), \text{Tr}(\mathbf{A}_N^j))$$

exist and if

$$\lim_{N \rightarrow \infty} c_r(\text{Tr}(\mathbf{A}_N^{j(1)}), \dots, \text{Tr}(\mathbf{A}_N^{j(r)})) = 0$$

for all $j \geq 3$ and all $j(1), \dots, j(r) \in \mathbb{N}$. In this definition, we denote the (classical) cumulants by c_m . Note that c_1 is just the expectation, and c_2 the covariance.

Notation 4.26. When $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ has a limit eigenvalue distribution, then the limits $\alpha_j^A := \lim_{N \rightarrow \infty} \mathbb{E}[\text{tr}(\mathbf{A}_N^j)]$ exist. When \mathbf{A}_N has a second order limit distribution, the fluctuation

$$\text{tr}(\mathbf{A}_N^j) - \alpha_j^A$$

is asymptotically Gaussian of order $1/N$. We consider the second order covariances defined as

$$\alpha_{i,j}^A := \lim_{N \rightarrow \infty} \text{cov}(\text{Tr}(\mathbf{A}_N^i), \text{Tr}(\mathbf{A}_N^j)),$$

and denote by $\mathcal{M}_A(x, y)$ the second order moment power series, which we define by:

$$\mathcal{M}_A(x, y) := \sum_{i,j \geq 1} \alpha_{i,j}^A x^i y^j.$$

Theorem 4.27. Assume that the $n \times n$ (non-random) covariance matrix $\Sigma = (\Sigma_n)_{n \in \mathbb{N}}$ has a limit eigenvalue distribution. Let \mathbf{S} be the (real or complex) sample covariance matrix formed from the m samples as in (4.1). Then for $n, m \rightarrow \infty$ with $n/m \rightarrow c \in (0, \infty)$, \mathbf{S} has both a limit eigenvalue distribution and a second order limit distribution. The Cauchy transform of the limit eigenvalue distribution, $g(x) \equiv g_\Sigma(x)$, satisfies the equation:

$$g(x) = \frac{1}{1 - c + c x g(x)} g_\Sigma \left(\frac{x}{1 - c + c x g(x)} \right), \quad (4.11)$$

with the corresponding power series $M_S(x) = 1/x g_\Sigma(1/x)$. Define $\tilde{\mathbf{S}} = \frac{1}{m} \mathbf{X}' \mathbf{X}$ so that its moment power series is given by

$$M_{\tilde{\mathbf{S}}}(y) = c(M_S(z) - 1) + 1. \quad (4.12)$$

The second order moment generating series is given by

$$\mathcal{M}_S(x, y) = \mathcal{M}_{\tilde{\mathbf{S}}}(x, y) = \frac{2}{\beta} \mathcal{M}_{\tilde{\mathbf{S}}}^\infty(x, y) \quad (4.13a)$$

where

$$\mathcal{M}_S^\infty(x, y) = xy \left(\frac{\frac{d}{dx}(xM_{\tilde{S}}(x)) \cdot \frac{d}{dy}(yM_{\tilde{S}}(y))}{(xM_{\tilde{S}}(x) - yM_{\tilde{S}}(y))^2} - \frac{1}{(x-y)^2} \right) \quad (4.13b)$$

where β equals 1 (or 2) when the elements of S are real (or complex).

PROOF. See Appendix A. \square

■ 4.3 Computational aspects

Proposition 4.31. For $\Sigma_\theta = \mathbf{U}\Lambda_\theta\mathbf{U}'$ as in (4.2), let $\theta = (t_1, \dots, t_{k-1}, a_1, \dots, a_k)$ where $t_i = n_i/n$. Then \mathbf{S} has a limit eigenvalue distribution as well as a second order limit distribution. The moments α_j^S , and hence $\alpha_{i,j}^S$, depend on θ and c . Let \mathbf{v}_θ be a q -by-1 vector whose j -th element is given by

$$(\mathbf{v}_\theta)_j = \text{Tr } \mathbf{S}^j - n \alpha_j^S.$$

Then for large n and m ,

$$\mathbf{v}_\theta \sim \mathcal{N}(\boldsymbol{\mu}_\theta, \mathbf{Q}_\theta) \quad (4.14)$$

where $\boldsymbol{\mu}_\theta = 0$ if \mathbf{S} is complex and $(\mathbf{Q}_\theta)_{i,j} = \alpha_{i,j}^S$.

PROOF. This follows directly from Theorem 4.27. From (4.15) and (4.17), the moments α_k^S depend on α^Σ and $c = n/m$ and hence on the unknown parameter vector θ . The existence of the non-zero mean when \mathbf{S} is real follows from the statement in [10]. \square

■ 4.3.1 Computation of moments of limiting eigenvalue distribution

Equation (4.11) expresses the relationship between the moment power series of Σ and that of \mathbf{S} via the limit of the ratio n/m . We can hence express the expected moments of \mathbf{S} in terms of the moments of Σ . The general form of the moments of $\tilde{\mathbf{S}}$, given by Corollary 9.12 in [68, pp.143], is

$$\alpha_{i_j}^{\tilde{S}} = \sum_{\substack{i_j \geq 0 \\ 1i_1 + 2i_2 + 3i_3 + \dots + ji_j = j}} c^{i_1 + i_2 + \dots + i_j} (\alpha_1^\Sigma)^{i_1} (\alpha_2^\Sigma)^{i_2} \dots (\alpha_j^\Sigma)^{i_j} \cdot \gamma_{i_1, i_2, \dots, i_j}^{(j)}, \quad (4.15)$$

where $\gamma_{i_1, \dots, i_j}^{(j)}$ is the multinomial coefficient given by

$$\gamma_{i_1, i_2, \dots, i_j}^{(j)} = \frac{j!}{i_1! i_2! \dots i_j! (j+1 - (i_1 + i_2 + \dots + i_j))!}. \quad (4.16)$$

The multinomial coefficient in (4.16) has an interesting combinatorial interpretation. Let j a positive integer, and let $i_1, \dots, i_j \in \mathbb{N} \cup \{0\}$ be such that $i_1 + 2i_2 + \dots + ji_j = j$. The number of partitions $\pi \in NC(j)$ which have i_1 blocks with 1 element, i_2 blocks with 2 elements, \dots , i_j blocks with j elements is given by the multinomial coefficient $\gamma_{i_1, \dots, i_j}^{(j)}$.

The moments of $\tilde{\mathbf{S}}$ are related to the moments of \mathbf{S} as

$$\alpha_j^{\tilde{\mathbf{S}}} = c \alpha_j^{\mathbf{S}} \quad \text{for } j = 1, 2, \dots \quad (4.17)$$

We can use (4.15) to compute the first few moments of \mathbf{S} in terms of the moments of Σ . This involves enumerating the partitions that appear in the computation of the multinomial coefficient in (4.16) For $j = 1$ only $i_1 = 1$ contributes with $\gamma_1^{(1)} = 1$, thus.

$$\alpha_1^{\tilde{\mathbf{S}}} = c \alpha_1^{\Sigma} \quad (4.18)$$

For $m = 2$ only $i_1 = 2, i_2 = 0$ and $i_1 = 0, i_2 = 1$ contribute with

$$\gamma_{2,0}^{(2)} = 1, \quad \gamma_{0,1}^{(2)} = 1,$$

and thus

$$\alpha_2^{\tilde{\mathbf{S}}} = c \alpha_2^{\Sigma} + c^2 (\alpha_1^{\Sigma})^2 \quad (4.19)$$

For $m = 3$ we have three possibilities for the indices, contributing with

$$\gamma_{3,0,0}^{(3)} = 1, \quad \gamma_{1,1,0}^{(3)} = 3, \quad \gamma_{0,0,1}^{(3)} = 1,$$

thus

$$\alpha_3^{\tilde{\mathbf{S}}} = c \alpha_3^{\Sigma} + 3c^2 \alpha_1^{\Sigma} \alpha_2^{\Sigma} + c^3 (\alpha_1^{\Sigma})^3 \quad (4.20)$$

For $m = 4$ we have five possibilities for the indices, contributing with

$$\gamma_{4,0,0,0}^{(4)} = 1, \quad \gamma_{2,1,0,0}^{(4)} = 6, \quad \gamma_{0,2,0,0}^{(4)} = 2, \quad \gamma_{1,0,1,0}^{(4)} = 4, \quad \gamma_{0,0,0,1}^{(4)} = 1$$

thus

$$\alpha_4^{\tilde{\mathbf{S}}} = c \alpha_4^{\Sigma} + 4c^2 \alpha_1^{\Sigma} \alpha_3^{\Sigma} + 2c^2 (\alpha_2^{\Sigma})^2 + 6c^3 (\alpha_1^{\Sigma})^2 \alpha_2^{\Sigma} + c^4 (\alpha_1^{\Sigma})^4. \quad (4.21)$$

For specific instances of Σ , we simply plug in the moments α_i^{Σ} into the above expressions to get the corresponding moments of \mathbf{S} . The general formula in (4.15) can be used to generate the expressions for higher order moments as well though such an explicit enumeration will be quite tedious even if symbolic software is used.

An alternate method is to use the software package RMTTool [73] based on the “polynomial method” developed in the second part of this dissertation. The software enables the moments of \mathbf{S} to be enumerate rapidly whenever the moment power series of Σ is an algebraic power series, *i.e.*, it is the solution of an algebraic equation. This is always the case when Σ is of the form in (4.2). For example, if $\theta = (t_1, t_2, a_1, a_2, a_3)$ then we can obtain the moments of \mathbf{S} by typing in the following sequence of commands in MATLAB once RMTTool has been installed. This eliminates the need to obtain manually obtain the expressions for the moments apriori.

```
>> startRMTTool
>> syms c t1 t2 a1 a2 a3
```

```

>> number_of_moments = 5;
>> LmzSigma = atomLmz([a1 a2 a3],[t1 t2 1-(t1+t2)]);
>> LmzS = AtimesWish(LmzSigma,c);
>> alpha_S = Lmz2MomF(LmzS,number_of_moments);
>> alpha_Stilde = c*alpha_S;

```

■ 4.3.2 Computation of covariance moments of second order limit distribution

Equations (4.13) and (4.13b) express the relationship between the covariance of the second order limit distribution and the moments of \mathbf{S} . Let $M(x)$ denote a moment power series as in Notation 4.23 with coefficients α_j . Define the power series $H(x) = x M(x)$ and let

$$\mathcal{H}(x,y) := \left(\frac{\frac{d}{dx}(H(x)) \cdot \frac{d}{dy}(H(y))}{(H(x) - H(y))^2} - \frac{1}{(x-y)^2} \right) \quad (4.22)$$

so that $\mathcal{M}^\infty(x,y) := xy \mathcal{H}(x,y)$. The (i,j) -th coefficient of $\mathcal{M}^\infty(x,y)$ can then be extracted from a multivariate Taylor series expansion of $\mathcal{H}(x,y)$ about $x = 0, y = 0$. From (4.13), we then obtain the coefficients $\alpha_{i,j}^S = (2/\beta)\alpha_{i,j}^{\mathcal{M}^\infty}$. This is best done using the MAPLE symbolic package where the following sequence of commands enumerates the coefficients $\alpha_{i,j}^S$ for $\beta = 1, 2$ and indices i and j such that $i+j \leq 2 \text{ max_coeff}$.

```

> with(numapprox):
> max_coeff := 5:
> H := x -> x*(1+sum(alpha[j]*x^2,j=1..2*max_coeff)):
> dHx := diff(H(x),x): dHy := diff(H(y),y):
> H2 := simplify(dHx*dHy/(H(x)-H(y))^2-1/(x-y)^2):
> H2series := mtaylor(H2,[x,y],2*max_coeff):
> i:=5: j =2:
> M2_infty_coeff[i,j] := simplify(coeff(coeff(H2series,x,i-1),y,j-1)):
> alphaS_second[i,j] := (2/beta)*M2_infty_coeff[i,j]:

```

Table 4.2 lists some of the coefficients of \mathcal{M}^∞ obtained using this procedure. When $\alpha_j = 1$ for all $j \in \mathbb{N}$, then $\alpha_{i,j} = 0$ as expected, since $\alpha_j = 1$ denotes the identity matrix. Note that the moments $\alpha_1, \dots, \alpha_{i+j}$ are needed to compute the second order covariance moments $\alpha_{i,j} = \alpha_{j,i}$.

The covariance matrix \mathbf{Q} with elements $\mathbf{Q}_{i,j} = \alpha_{i,j}$ gets increasingly ill-conditioned as $\dim(\mathbf{Q})$ increases; the growth in the magnitude of the diagonal entries $\alpha_{j,j}$ in Table 4.2 attests to this. This implies that the eigenvectors of \mathbf{Q} encode the information about the covariance of the second order limit distribution more efficiently than the matrix \mathbf{Q} itself. When $\Sigma = \mathbf{I}$ so that the SCM \mathbf{S} has the (null) Wishart distribution, the eigenvectors of \mathbf{Q} are the (appropriately normalized) Chebychev polynomials of the second kind [64]. The structure of the eigenvectors for arbitrary Σ is, as yet, unknown though research in that direction might yield additional insights.

Coefficient	Expression
$\alpha_{1,1}$	$\alpha_2 - \alpha_1^2$
$\alpha_{2,1}$	$-4\alpha_1\alpha_2 + 2\alpha_1^3 + 2\alpha_3$
$\alpha_{2,2}$	$16\alpha_1^2\alpha_2 - 6\alpha_2^2 - 6\alpha_1^4 - 8\alpha_1\alpha_3 + 4\alpha_4$
$\alpha_{3,1}$	$9\alpha_1^2\alpha_2 - 6\alpha_1\alpha_3 - 3\alpha_2^2 + 3\alpha_4 - 3\alpha_1^4$
$\alpha_{3,2}$	$6\alpha_5 + 30\alpha_1\alpha_2^2 - 42\alpha_1^3\alpha_2 - 18\alpha_2\alpha_3 + 12\alpha_1^5 + 24\alpha_1^2\alpha_3 - 12\alpha_1\alpha_4$
$\alpha_{3,3}$	$-18\alpha_3^2 - 27\alpha_2\alpha_4 + 9\alpha_6 - 30\alpha_1^6 + 21\alpha_2^3 + 36\alpha_1^2\alpha_4 - 72\alpha_1^3\alpha_3 + 126\alpha_1^4\alpha_2 - 135\alpha_1^2\alpha_2^2 + 108\alpha_1\alpha_2\alpha_3 - 18\alpha_1\alpha_5$
$\alpha_{4,1}$	$12\alpha_1\alpha_2^2 - 16\alpha_1^3\alpha_2 - 8\alpha_2\alpha_3 + 12\alpha_1^2\alpha_3 - 8\alpha_1\alpha_4 + 4\alpha_1^5 + 4\alpha_5$
$\alpha_{4,2}$	$-12\alpha_3^2 - 24\alpha_2\alpha_4 + 8\alpha_6 - 20\alpha_1^6 + 16\alpha_2^3 + 32\alpha_1^2\alpha_4 - 56\alpha_1^3\alpha_3 + 88\alpha_1^4\alpha_2 - 96\alpha_1^2\alpha_2^2 + 80\alpha_1\alpha_2\alpha_3 - 16\alpha_1\alpha_5$
$\alpha_{4,3}$	$96\alpha_2^2\alpha_3 + 60\alpha_1^7 + 84\alpha_1\alpha_3^2 + 432\alpha_1^3\alpha_2^2 + 180\alpha_1^4\alpha_3 - 48\alpha_3\alpha_4 + 12\alpha_7 - 36\alpha_2\alpha_5 - 24\alpha_1\alpha_6 + 144\alpha_1\alpha_2\alpha_4 + 48\alpha_1^2\alpha_5 - 96\alpha_1^3\alpha_4 - 156\alpha_1\alpha_2^3 - 300\alpha_1^5\alpha_2 - 396\alpha_1^2\alpha_2\alpha_3$
$\alpha_{4,4}$	$-140\alpha_1^8 - 76\alpha_2^4 - 48\alpha_6\alpha_2 + 256\alpha_3\alpha_4\alpha_1 - 40\alpha_4^2 + 16\alpha_8 - 64\alpha_3\alpha_5 - 32\alpha_1\alpha_7 + 1408\alpha_1^3\alpha_2\alpha_3 - 336\alpha_1^2\alpha_3^2 + 256\alpha_1^4\alpha_4 + 144\alpha_2^2\alpha_4 - 480\alpha_1^5\alpha_3 + 160\alpha_2\alpha_3^2 + 64\alpha_1^2\alpha_6 - 128\alpha_1^3\alpha_5 - 1440\alpha_1^4\alpha_2^2 + 832\alpha_1^2\alpha_2^3 + 800\alpha_1^6\alpha_2 - 768\alpha_1\alpha_2^2\alpha_3 - 576\alpha_1^2\alpha_2\alpha_4 + 192\alpha_1\alpha_2\alpha_5$
$\alpha_{5,1}$	$-5\alpha_3^2 - 10\alpha_2\alpha_4 + 5\alpha_6 - 5\alpha_1^6 + 5\alpha_2^3 + 15\alpha_1^2\alpha_4 - 20\alpha_1^3\alpha_3 + 25\alpha_1^4\alpha_2 - 30\alpha_1^2\alpha_2^2 + 30\alpha_1\alpha_2\alpha_3 - 10\alpha_1\alpha_5$
$\alpha_{5,2}$	$60\alpha_2^2\alpha_3 + 30\alpha_1^7 + 50\alpha_1\alpha_3^2 + 240\alpha_1^3\alpha_2^2 + 110\alpha_1^4\alpha_3 - 30\alpha_3\alpha_4 + 10\alpha_7 - 30\alpha_2\alpha_5 - 20\alpha_1\alpha_6 + 100\alpha_1\alpha_2\alpha_4 + 40\alpha_1^2\alpha_5 - 70\alpha_1^3\alpha_4 - 90\alpha_1\alpha_2^3 - 160\alpha_1^5\alpha_2 - 240\alpha_1^2\alpha_2\alpha_3$
$\alpha_{5,3}$	$-105\alpha_1^8 - 60\alpha_2^4 - 45\alpha_6\alpha_2 + 210\alpha_3\alpha_4\alpha_1 - 30\alpha_4^2 + 15\alpha_8 - 60\alpha_3\alpha_5 - 30\alpha_1\alpha_7 + 1140\alpha_1^3\alpha_2\alpha_3 - 270\alpha_1^2\alpha_3^2 + 225\alpha_1^4\alpha_4 + 120\alpha_2^2\alpha_4 - 390\alpha_1^5\alpha_3 + 135\alpha_2\alpha_3^2 + 60\alpha_1^2\alpha_6 - 120\alpha_1^3\alpha_5 - 1125\alpha_1^4\alpha_2^2 + 660\alpha_1^2\alpha_2^3 + 615\alpha_1^6\alpha_2 - 630\alpha_1\alpha_2^2\alpha_3 - 495\alpha_1^2\alpha_2\alpha_4 + 180\alpha_1\alpha_2\alpha_5$
$\alpha_{5,4}$	$-900\alpha_1^2\alpha_4\alpha_3 + 80\alpha_1^2\alpha_7 - 160\alpha_1^3\alpha_6 - 620\alpha_1^5\alpha_4 - 3200\alpha_1^3\alpha_2^3 + 700\alpha_1\alpha_2^4 + 3960\alpha_1^5\alpha_2^2 - 720\alpha_1^2\alpha_5\alpha_2 + 1840\alpha_1^3\alpha_4\alpha_2 - 4100\alpha_1^4\alpha_3\alpha_2 + 3600\alpha_1^2\alpha_2^2\alpha_3 - 1140\alpha_1\alpha_3^2\alpha_2 + 1040\alpha_1^3\alpha_3^2 - 440\alpha_2^3\alpha_3 + 440\alpha_3\alpha_4\alpha_2 + 240\alpha_1\alpha_6\alpha_2 + 320\alpha_1\alpha_5\alpha_3 - 1020\alpha_1\alpha_2^2\alpha_4 + 20\alpha_9 - 1820\alpha_1^7\alpha_2 + 180\alpha_2^2\alpha_5 + 320\alpha_1^4\alpha_5 + 180\alpha_1\alpha_4^2 + 1120\alpha_1^6\alpha_3 + 80\alpha_3^3 + 280\alpha_1^9 - 40\alpha_1\alpha_8 - 60\alpha_7\alpha_2 - 80\alpha_3\alpha_6 - 100\alpha_4\alpha_5$
$\alpha_{5,5}$	$2400\alpha_2\alpha_5\alpha_1^3 - 1350\alpha_2^2\alpha_5\alpha_1 + 600\alpha_3\alpha_5\alpha_2 + 300\alpha_1\alpha_7\alpha_2 - 900\alpha_6\alpha_2\alpha_1^2 - 1200\alpha_3\alpha_5\alpha_1^2 + 400\alpha_1\alpha_6\alpha_3 + 3000\alpha_3\alpha_4\alpha_1^3 + 5100\alpha_1^2\alpha_2^2\alpha_4 + 12300\alpha_1^5\alpha_2\alpha_3 + 5700\alpha_1^2\alpha_2\alpha_3^2 + 4400\alpha_1\alpha_2^3\alpha_3 + 400\alpha_1^4\alpha_6 - 15000\alpha_1^3\alpha_2^2\alpha_3 - 5750\alpha_1^4\alpha_2\alpha_4 - 200\alpha_1^3\alpha_7 + 500\alpha_1\alpha_4\alpha_5 + 225\alpha_6\alpha_2^2 - 675\alpha_4^2\alpha_1^2 - 3250\alpha_1^4\alpha_3^2 - 625\alpha_2^3\alpha_4 + 350\alpha_3^2\alpha_4 - 600\alpha_1\alpha_3^3 - 1050\alpha_2^2\alpha_3^2 - 2800\alpha_3\alpha_1^7 - 11550\alpha_1^6\alpha_2^2 - 3300\alpha_3\alpha_4\alpha_1\alpha_2 - 800\alpha_5\alpha_1^5 + 325\alpha_4^2\alpha_2 - 4375\alpha_1^2\alpha_2^4 - 630\alpha_1^{10} + 100\alpha_8\alpha_1^2 - 75\alpha_5^2 + 255\alpha_2^5 + 12000\alpha_1^4\alpha_2^3 + 4550\alpha_1^8\alpha_2 + 1550\alpha_1^6\alpha_4 + 25\alpha_{10} - 50\alpha_1\alpha_9 - 75\alpha_2\alpha_8 - 100\alpha_3\alpha_7 - 125\alpha_4\alpha_6$

Table 4.2. Relationship between the coefficients $\alpha_{i,j} = \alpha_{j,i}$ and α_i .

■ 4.4 Eigen-inference algorithms

■ 4.4.1 Estimating θ for known model order

Estimating the unknown parameter vector θ follows from the asymptotic result in Proposition 4.31. For large n, m , since \mathbf{v}_θ is (approximately) normally distributed we can obtain the estimate $\hat{\theta}$ by the principle of maximum-likelihood. When \mathbf{S} is real, Bai and Silverstein provide a formula, expressed as a difficult to compute contour integral, for the correction term μ_θ in (4.14). The log-likelihood of \mathbf{v}_θ is (ignoring constants and the correction term for the mean when \mathbf{S} is real) given by

$$\ell(\mathbf{v}_\theta|\theta) \approx -\mathbf{v}_\theta^T \mathbf{Q}_\theta^{-1} \mathbf{v}_\theta - \log \det \mathbf{Q}_\theta, \quad (4.23)$$

which allows us to obtain the maximum-likelihood estimate of θ as

$$\hat{\theta}_{(q)} = \arg \min_{\theta \in \Theta} \mathbf{v}_\theta^T \mathbf{Q}_\theta^{-1} \mathbf{v}_\theta + \log \det \mathbf{Q}_\theta \quad \text{for } q = \dim(\mathbf{v}_\theta) \geq \dim(\theta) \quad (4.24)$$

where Θ represents the parameter space for the elements of θ and \mathbf{v}_θ and \mathbf{Q}_θ are constructed as in Proposition 4.31.

Canonically, the parameter vector θ of models such as (4.2) is of length $2k - 1$ so that $q = \dim(\mathbf{v}_\theta) \geq 2k - 1$. In principle, estimation accuracy should increase with q since the covariance of \mathbf{v}_θ is explicitly accounted for via the weighting matrix \mathbf{Q}_θ .

Figure 4-2 compares the quantiles of the test statistic $\mathbf{v}_\theta^T \mathbf{Q}_\theta \mathbf{v}_\theta$ for $\dim(\mathbf{v}_\theta) = q$ with the quantiles of the chi-square distribution with q degrees of freedom when $q = 2, 3$ for the model in (4.2) with $\theta = (0.5, 2, 1)$, $m = n$ for $m = 40$ and $m = 320$. While there is good agreement with the theoretical distribution for large m, n , the deviation from the limiting result is not insignificant for moderate m, n . This justifies setting $q = 2$ for the testing procedures developed herein.

Hence, we suggest that for the estimation in (4.24), $q = \dim(\mathbf{v}_\theta) = \dim(\theta)$. This choice provide robustness in low to moderate dimensional settings where the deviations from the asymptotic result in Theorem 4.27 are not insignificant. Numerical simulations suggest that the resulting degradation in estimation accuracy in high dimensional settings, from such a choice, is relatively small. This loss in performance is offset by an increase in the speed of the underlying numerical optimization routine. This is the case because, though the dimensionality of θ is the same, the matrix \mathbf{Q} gets increasingly ill-conditioned for higher values of q thereby reducing the efficiency of optimization methods .





(a) $n = m = 40$.(b) $n = m = 320$.(c) $n = m = 40$.(d) $n = m = 320$.

Figure 4-2. Numerical simulations (when \mathbf{S} is complex) illustrating the robustness of test statistics formed with $\dim(\mathbf{v}) = 2$ to moderate dimensional settings.

■ 4.4.2 Testing $\theta = \theta_0$

Proposition 4.41. Define the vector \mathbf{v}_θ and the covariance matrix \mathbf{Q}_θ as

$$\mathbf{v}_\theta = \begin{bmatrix} \text{Tr } \mathbf{S} - n \alpha_1^\Sigma \\ \text{Tr } \mathbf{S}^2 - n \left(\alpha_2^\Sigma + \frac{n}{m} (\alpha_1^\Sigma)^2 \right) - \left(\frac{2}{\beta} - 1 \right) \alpha_2^\Sigma \frac{n}{m} \end{bmatrix} \quad (4.25a)$$

$$\mathbf{Q}_\theta = \frac{2}{\beta} \begin{bmatrix} \tilde{\alpha}_2 - \tilde{\alpha}_1^2 & 2\tilde{\alpha}_1^3 + 2\tilde{\alpha}_3 - 4\tilde{\alpha}_1\tilde{\alpha}_2 \\ 2\tilde{\alpha}_1^3 + 2\tilde{\alpha}_3 - 4\tilde{\alpha}_1\tilde{\alpha}_2 & 4\tilde{\alpha}_4 - 8\tilde{\alpha}_1\tilde{\alpha}_3 - 6\tilde{\alpha}_2^2 + 16\tilde{\alpha}_2\tilde{\alpha}_1^2 - 6\tilde{\alpha}_1^4 \end{bmatrix} \quad (4.25b)$$

with $\beta = 1$ (or 2) when S is real (or complex) and $\tilde{\alpha}_i \equiv \alpha_i^{\tilde{S}}$ given by

$$\tilde{\alpha}_1 = \frac{n}{m} \alpha_1^\Sigma \quad (4.26a)$$

$$\tilde{\alpha}_2 = \frac{n}{m} \alpha_2^\Sigma + \frac{n^2}{m^2} (\alpha_1^\Sigma)^2 \quad (4.26b)$$

$$\tilde{\alpha}_3 = \frac{n}{m} \alpha_3^\Sigma + 3 \frac{n^2}{m^2} \alpha_1^\Sigma \alpha_2^\Sigma + \frac{n^3}{m^3} (\alpha_1^\Sigma)^3 \quad (4.26c)$$

$$\tilde{\alpha}_4 = \frac{n}{m} \alpha_4^\Sigma + 4 \frac{n^2}{m^2} \alpha_1^\Sigma \alpha_3^\Sigma + 2 \frac{n^2}{m^2} (\alpha_2^\Sigma)^2 + 6 \frac{n^3}{m^3} (\alpha_1^\Sigma)^2 \alpha_2^\Sigma + \frac{n^4}{m^4} (\alpha_1^\Sigma)^4. \quad (4.26d)$$

and $\alpha_i^\Sigma = (1/n) \text{Tr } \Sigma^i$. Thus, for large n and m , $\mathbf{v}_\theta \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_\theta)$ so that

$$h(\theta) := \mathbf{v}_\theta^T \mathbf{Q}_\theta^{-1} \mathbf{v}_\theta \sim \chi_2^2$$

PROOF. This follows from Proposition 4.31. The correction term for the real case is discussed in a different context in [31]. \square

We test for $\theta = \theta_0$ by obtaining the test statistic

$$\boxed{H_{\theta_0} : h(\theta_0) = \mathbf{v}_{\theta_0}^T \mathbf{Q}_{\theta_0}^{-1} \mathbf{v}_{\theta_0}} \quad (4.27)$$

where the \mathbf{v}_{θ_0} and \mathbf{Q}_{θ_0} are constructed as in (4.25a) and (4.25b), respectively. We reject the hypothesis for large values of H_{θ_0} . For a choice of threshold γ , the asymptotic convergence of the test statistic to the χ_2^2 distribution, implies that

$$\text{Prob.}(H_{\theta_0} = 1 | \theta = \theta_0) \approx F^{\chi_2^2}(\gamma). \quad (4.28)$$

Thus, for large n and m , when $\gamma = 5.9914$, $\text{Prob.}(H_{\theta_0} = 1 | \theta = \theta_0) \approx 0.95$.

■ 4.4.3 Estimating θ and testing $\theta = \hat{\theta}$

When an $\hat{\theta}$ is obtained using (4.24) then we may test for $\theta = \hat{\theta}$ by forming the testing statistic

$$H_{\hat{\theta}} : h(\hat{\theta}) = \mathbf{u}_{\hat{\theta}}^T \mathbf{W}_{\hat{\theta}}^{-1} \mathbf{u}_{\hat{\theta}} \quad (4.29)$$

where the $\mathbf{u}_{\hat{\theta}}$, and $\mathbf{W}_{\hat{\theta}}$ are constructed as in (4.25a) and (4.25b), respectively. However, the sample covariance matrix \mathbf{S} can no longer be used since the estimate $\hat{\theta}$ was obtained from it. Instead, we form a test sample covariance matrix constructed from $\lceil (m/2) \rceil$ randomly chosen samples. Equivalently, since the samples are assumed to be mutually independent and identically distributed, we can form the test matrix from the first $\lceil (m/2) \rceil$ samples as

$$\bar{\mathbf{S}} = \frac{1}{\lceil \frac{m}{2} \rceil} \sum_{i=1}^{\lceil \frac{m}{2} \rceil} \mathbf{x}_i \mathbf{x}_i' \quad (4.30)$$

Note that $\alpha_k^{\bar{\mathbf{S}}}$ will have to be recomputed using $\Sigma_{\hat{\theta}}$ and $\bar{c} = n / \lceil (m/2) \rceil$. The hypothesis $\theta = \hat{\theta}$ is tested by rejecting values of the test statistic greater than a threshold γ . The threshold is selected using the approximation in (4.28).

■ 4.4.4 Estimating θ for unknown model order

Suppose we have a family of models parameterized by the vector $\theta^{(\bar{k})}$. The elements of $\theta^{(\bar{k})}$ are the free parameters of the model. For the model in (4.2), in the canonical case $\theta = (t_1, \dots, t_{k-1}, a_1, \dots, a_k)$ since $t_1 + \dots + t_{k-1} + t_k = 1$ so that $\dim(\theta^{(\bar{k})}) = 2k - 1$. If some of the parameters in (4.2) are known, then the parameter vector is modified accordingly.

When the model order is unknown, we select the model which has the minimum Akaike Information Criterion . For the situation at hand we propose that

$$\hat{\theta} = \hat{\theta}^{(\hat{k})} \quad \text{where} \quad \hat{k} = \arg \min_{k \in \mathbb{N}} \left\{ \mathbf{u}_{\hat{\theta}^{(k)}}^T \mathbf{W}_{\hat{\theta}^{(k)}}^{-1} \mathbf{u}_{\hat{\theta}^{(k)}} + \log \det \mathbf{W}_{\hat{\theta}^{(k)}} \right\} + 2 \dim(\theta^{(k)}) \quad (4.31)$$

where $\mathbf{u}_{\hat{\theta}^{(k)}}$ and $\mathbf{W}_{\hat{\theta}^{(k)}}$ are constructed as described in Section 4.4.3 using the test sample covariance matrix in (4.30). The Bayesian Information Criterion (BIC) may also be used for model order selection. It would be useful to compare the performance of these two criterion in situations of practical interest.

■ 4.5 Numerical simulations

Let $\Sigma_{\bar{\theta}}$ be as in (4.2) with $\bar{\theta} = (t_1, a_1, a_2)$. When $t_1 = 0.5$, $a_1 = 2$ and $a_2 = 1$ then half of the population eigenvalues are of magnitude two while the remainder are of magnitude one. Let the unknown parameter vector $\theta = (t, a)$ where $t \equiv t_1$ and $a \equiv a_1$.

Using the procedure described in Section 4.3.1, the first four moments can be obtained as (here $c = n/m$)

$$\alpha_1^S = 1 + t(a - 1) \quad (4.32a)$$

$$\alpha_2^S = (-2ac + a^2c + c)t^2 + (-1 + 2ac - 2c + a^2)t + 1 + c \quad (4.32b)$$

$$\begin{aligned} \alpha_3^S = & (-3c^2a^2 + a^3c^2 - c^2 + 3ac^2)t^3 + (3c^2 + 3c^2a^2 - 3ac - 6ac^2 - 3a^2c + 3a^3c + 3c)t^2 \\ & + (-3c^2 + a^3 - 1 - 6c + 3ac + 3a^2c + 3ac^2)t + 1 + c^2 + 3c \end{aligned} \quad (4.32c)$$

$$\begin{aligned} \alpha_4^S = & (6a^2c^3 + a^4c^3 - 4ac^3 - 4a^3c^3 + c^3)t^4 \\ & + (-6c^2 - 12a^3c^2 + 12ac^3 - 12a^2c^3 + 4a^3c^3 + 12ac^2 + 6a^4c^2 - 4c^3)t^3 + \\ & (-4a^2c - 4ac - 12ac^3 - 24ac^2 + 6a^4c + 6a^2c^3 + 12a^3c^2 + 6c - 6c^2a^2 + 6c^3 + 18c^2 - 4a^3c)t^2 \\ & + (-4c^3 + 4ac + 6c^2a^2 + 4ac^3 - 1 + 12ac^2 - 18c^2 + 4a^2c - 12c + 4a^3c + a^4)t \\ & + 1 + c^3 + 6c + 6c^2 \end{aligned} \quad (4.32d)$$

From the discussion in Section 4.3.2, we obtain the covariance of the second order limit distribution

$$\mathbf{Q}_\theta = \frac{2}{\beta} \begin{bmatrix} c^2(\alpha_2^S - \alpha_1^2) & c^3(2(\alpha_1^S)^3 + 2\alpha_3^S - 4\alpha_1^S\alpha_2^S) \\ c^3((2(\alpha_1^S)^3 + 2\alpha_3^S - 4\alpha_1^S\alpha_2^S) & c^4(4\alpha_4^S - 8\alpha_1^S\alpha_3^S - 6(\alpha_2^S)^2 + 16\alpha_2^S(\alpha_1^S)^2 - 6(\alpha_1^S)^4) \end{bmatrix}. \quad (4.33)$$

where $\beta = 1$ when S is real valued and $\beta = 2$ when \mathbf{S} is complex valued.

We then use (4.24) to estimate θ and hence the unknown parameters t and a . Table 4.3 and 4.4 compares the bias and mean squared error of the estimates for a and t respectively. Note the $1/n^2$ type decay in the mean squared error and how the real case has twice the variance as the complex case. As expected by the theory of maximum likelihood estimation, the estimates become increasingly normal for large n and m . This is evident from Figure 4-3. As expected, the performance improves as the dimensionality of the system increases.

n	m	Bias	Complex Case MSE	MSE x $n^2/100$	Bias	Real Case MSE	MSE x $n^2/100$
10	5						
20	10	0.0455	0.3658	1.4632	0.4862	1.2479	4.9915
40	20	-0.0046	0.1167	1.8671	0.2430	0.3205	5.1272
80	40	-0.0122	0.0337	2.1595	0.1137	0.08495	5.437
160	80	-0.0024	0.0083	2.1250	0.0598	0.02084	5.335
320	160	0.0008	0.0021	2.1790	0.0300	0.00528	5.406

(a) $m = 0.5n$.

n	m	Bias	Complex Case MSE	MSE x $n^2/100$	Bias	Real Case MSE	MSE x $n^2/100$
10	10						
20	20	-0.0137	0.1299	0.5196	0.2243	0.3483	1.3932
40	40	-0.0052	0.0390	0.6233	0.1083	0.0901	1.4412
80	80	-0.0019	0.0093	0.5941	0.0605	0.0231	1.4787
160	160	-0.0005	0.0024	0.6127	0.0303	0.0055	1.4106
320	320	-0.0001	0.0006	0.6113	0.0162	0.0015	1.5155

(b) $m = n$.

n	m	Bias	Complex Case MSE	MSE x $n^2/100$	Bias	Real Case MSE	MSE x $n^2/100$
10	20	-					
20	40	-0.0119	0.0420	0.1679	0.1085	0.1020	0.4081
40	80	-0.0017	0.0109	0.1740	0.0563	0.0255	0.4079
80	160	-0.0005	0.0028	0.1765	0.0290	0.0063	0.4056
160	320	-0.0004	0.0007	0.1828	0.0151	0.0016	0.4139
320	640	0.0001	0.0002	0.1752	0.0080	0.0004	0.4024

(c) $m = 2n$.

Table 4.3. Quality of estimation of $t = 0.5$ for different values of n (dimension of observation vector) and m (number of samples) – both real and complex case.

n	m	Bias	Complex Case MSE	MSE x $n^2/100$	Bias	Real Case MSE	MSE x $n^2/100$
10	5						
20	10	0.1278	0.1046	0.4185	0.00748	0.1024	0.4097
40	20	0.0674	0.0478	0.7647	-0.01835	0.04993	0.7989
80	40	0.0238	0.0111	0.7116	-0.02240	0.01800	1.1545
160	80	0.0055	0.0022	0.5639	-0.02146	0.00414	1.0563
320	160	0.0007	0.0005	0.5418	-0.01263	0.00112	1.1692

(a) $m = 0.5n$.





n	m	Bias	Complex Case MSE	MSE x $n^2/100$	Bias	Real Case MSE	MSE x $n^2/100$
10	10						
20	20	0.0750	0.0525	0.2099	-0.0019	0.0577	0.2307
40	40	0.0227	0.0127	0.2028	-0.0206	0.0187	0.2992
80	80	0.0052	0.0024	0.1544	-0.0206	0.0047	0.3007
160	160	0.0014	0.0006	0.1499	-0.0126	0.0012	0.3065
320	320	0.0003	0.0001	0.1447	-0.0074	0.0003	0.3407

(b) $m = n$.

n	m	Bias	Complex Case MSE	MSE x $n^2/100$	Bias	Real Case MSE	MSE x $n^2/100$
10	20						
20	40	0.0251	0.0134	0.0534	-0.0182	0.0205	0.0821
40	80	0.0049	0.0028	0.0447	-0.0175	0.0052	0.0834
80	160	0.0015	0.0007	0.0428	-0.0115	0.0014	0.0865
160	320	0.0004	0.0002	0.0434	-0.0067	0.0004	0.0920
320	640	0.0000	0.0000	0.0412	-0.0038	0.0001	0.0932

(c) $m = 2n$.

Table 4.4. Quality of estimation of $a = 2$ for different values of n (dimension of observation vector) and m (number of samples) – both real and complex case.

(a) $\hat{a}: n = 320, m = 640$.(b) $\hat{t}: n = 320, n = 640$.(c) $\hat{a}: n = 320, m = 640$. (Real valued)(d) $\hat{t}: n = 320, m = 640$. (Real valued)Figure 4-3. Normal probability plots of the estimates of a and t (true values: $a = 2, t = 0.5$).

■ 4.6 Inferential aspects of spiked covariance matrix models

Consider covariance matrix models whose eigenvalues are of the form $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k > \lambda_{k+1} = \dots = \lambda_n = \lambda$. Such models arise when the signal occupies a k -dimensional subspace and the noise has covariance $\lambda \mathbf{I}$. Such models are referred to as *spiked covariance matrix models*. When $k \ll n$, then for large n , for \mathbf{v}_θ defined as in Proposition 4.31, the matrix \mathbf{Q}_θ may be constructed from the moments of the (null) Wishart distribution [33] instead, which are given by

$$\alpha_k^W = \lambda^k \sum_{j=0}^{k-1} c^j \frac{1}{j+1} \binom{k}{j} \binom{k-1}{j} \quad (4.34)$$

where $c = n/m$. Thus, for $q = 2$, \mathbf{Q}_θ is given by

$$\mathbf{Q}_\theta \equiv \mathbf{Q}_\lambda = \frac{2}{\beta} \begin{bmatrix} \lambda^2 c & 2\lambda^3 (c+1)c \\ 2\lambda^3 (c+1)c & 2\lambda^4 (2c^2 + 5c + 2)c \end{bmatrix}. \quad (4.35)$$

This substitution is motivated by Bai and Silverstein's analysis [10] where it is shown that when k is small relative to n , then the second order fluctuation distribution is asymptotically independent of the "spikes." When the multiplicities of the spike is known (say 1), then we let $t_i = 1/n$ and compute the moments α_j^S accordingly. The estimation problem thus reduces to

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \mathbf{v}_\theta^T \mathbf{Q}_\lambda^{-1} \mathbf{v}_\theta \quad \text{with } q = \dim(\mathbf{v}_\theta) = \dim(\theta) + 1 \quad (4.36)$$

where λ is an element of θ when it is unknown.

Consider the problem of estimating the magnitude of the spike for the model in (4.2) with $t_1 = 1/n$, and $a_2 = 1$ known and $a_1 = 10$ unknown so that $\theta = a \equiv a_1$. We obtain the estimate $\hat{\theta}$ from (4.36) with $\lambda = 1$ wherein the moments α_k^S given by

$$\alpha_1^S = \frac{-1 + a + n}{n} \quad (4.37a)$$

$$\alpha_2^S = \frac{a^2 n - 2pc + c - 2ac + cn^2 + n^2 - n + 2pac + a^2 c}{n^2} \quad (4.37b)$$

are obtained by plugging in $t = 1/n$ into (4.32).

Table 4.5 summarizes the estimation performance for this example. Note the $1/n$ scaling of the mean squared error and how the complex case has half the mean squared error. The estimates produced are asymptotically normal as seen in Figure 4-4.

■ 4.6.1 Impact of the sample eigenvalue phase transition phenomenon

Consider testing for the hypothesis that $\Sigma = \mathbf{I}$. For the model in (4.2), which is equivalent to testing $\theta = (1, 1)$, from the discussion in Section 4.4.2, we form the test statistic

$$H_{\text{Sph.}} : h(\theta) = \mathbf{v}_\theta^T \mathbf{Q}_\theta^{-1} \mathbf{v}_\theta \quad (4.38)$$

where \mathbf{Q}_θ is given by (4.35) with $\lambda = 1$ and

$$\mathbf{v}_\theta = \begin{bmatrix} \text{Tr } \mathbf{S} - n \\ \text{Tr } \mathbf{S}^2 - n \left(1 + \frac{n}{m}\right) - \left(\frac{2}{\beta} - 1\right) \frac{n}{m} \end{bmatrix}$$

where $c = n/m$, as usual. Figure 4-5 compares quantiles of the test statistic, collected over 4000 Monte-Carlo simulations, with the theoretical quantiles of the χ_2^2 distribution. The agreement validates distributional approximation for modest values of n and m .

We set a threshold $\gamma = 5.9914$ so that we accept the sphericity hypothesis whenever $h(\theta) \leq \gamma$. This corresponds to the 95-th percentile of the χ_2^2 distribution. Table 4.6(a) demonstrates how the test is able to accept the hypothesis when $\Sigma = \mathbf{I}$ close to the 0.95 significance level it was designed for.

Table 4.6(b) shows the acceptance of the sphericity hypothesis when $\Sigma = \Sigma = \text{diag}(10, 1, \dots, 1)$ instead. Note how when n/m is large, the test erroneously accepts the null hypothesis an inordinate number of times. The faulty inference provided by the test based on the methodologies developed is not surprising given the phase transition phenomenon for the sample eigenvalues described by the following result due to Baik-Silverstein [12], Paul [70] and others [11].

Proposition 4.61. *Let \mathbf{S} denote a sample covariance matrix formed from an $n \times m$ matrix of Gaussian observations whose columns are independent of each other and identically distributed with mean $\mathbf{0}$ and covariance Σ . Denote the eigenvalues of Σ by $\lambda_1 \geq \lambda_2 > \dots \geq \lambda_k > \lambda_{k+1} = \dots = \lambda_n = \lambda$. Let l_j denote the j -th largest eigenvalue of $\widehat{\mathbf{R}}$. Then as $n, m \rightarrow \infty$ with $c_m = n/m \rightarrow c \in (0, \infty)$,*

$$l_j \rightarrow \begin{cases} \lambda_j \left(1 + \frac{\lambda c}{\lambda_j - \lambda}\right) & \text{if } \lambda_j > \lambda(1 + \sqrt{c}) \\ \lambda(1 + \sqrt{c})^2 & \text{if } \lambda_j \leq \lambda(1 + \sqrt{c}) \end{cases} \quad (4.39)$$

where the convergence is almost surely.

Since the inference methodologies we propose in this paper exploit the distributional properties of traces of powers of the sample covariance matrix, Proposition 4.61 pinpoints the fundamental inability of the sphericity test proposed to reject the hypothesis

$\Sigma = \mathbf{I}$ whenever (for large n, m),

$$\lambda_i \leq 1 + \sqrt{\frac{n}{m}}$$

For the example considered, $\lambda_1 = 10$, so that the above condition is met whenever $n/m > c_t = 81$. For n/m on the order of c_t , the resulting inability to correctly reject the null hypothesis can be attributed to this phenomenon and the fluctuations of the largest eigenvalue.

Canonically speaking, eigen-inference methodologies which rely on traces of powers of the sample covariance matrix will be unable to differentiate between closely spaced population eigenvalues in high-dimensional, sample sized starved settings. This impacts the quality of the inference in a fundamental manner that is difficult to overcome. At the same time, however, the results in [12] suggest that if the practitioner has reason to believe that the population eigenvalues can be split into several clusters about $a_i \pm \sqrt{n/m}$, then the use of the model in (4.2) with a block subspace structure, where the individual blocks of sizes n_1, \dots, n_k are comparable to n , is justified. In such situations, the benefit of the proposed eigen-methodologies will be most apparent and might motivate experimental design that ensures that this condition is met.

■ 4.7 Future work

In the development of the estimation procedures in this chapter, we ignored the correction term for the mean that appears in the real covariance matrix case (see Proposition 4.31). This was because Bai and Silverstein expressed it as a contour integral which appeared challenging to compute (see Eq. (1.6) in [10]). It is desirable to include this extra term in the estimation procedure if it can be computed efficiently using symbolic techniques. The recent work of Anderson and Zeitouni [5], despite its ambiguous title, represents a breakthrough on this and other fronts.

Anderson and Zeitouni encode the correction term in the coefficients of a power series that can be directly computed from the limiting moment series of the sample covariance matrix (see Theorem 3.4 [5]). Furthermore, they have expanded the range of the theory for the fluctuations of traces of powers of large Wishart-like sample covariance matrices, in the real sample covariance matrix case, to the situation when the entries are composed from a broad class of admissible non-Gaussian distributions. In such a scenario, the correction term takes into account the fourth moment of the distribution (see Eq. (5) and Theorems 3.3-3.4 in [5]). This latter development might be of use in some practical settings where the non-Gaussianity is well characterized. We have yet to translate their results into a computational recipe for determining the correction term though we intend to do so at a later date. The numerical results presented show the consistency of the proposed estimators; it would be of interest to establish this analytically and identify conditions in the real covariance matrix case, where ignoring the correction term in the mean can severely degrade the quality of estimation.

n	m	Bias	Complex Case MSE	MSE x n	Bias	Real Case MSE	MSE x n
10	10	-0.5528	9.3312	93.3120	-0.5612	18.4181	184.1808
20	20	-0.2407	4.8444	96.8871	-0.2005	9.6207	192.4143
40	40	-0.1168	2.5352	101.4074	-0.0427	4.9949	199.7965
80	80	-0.0833	1.2419	99.3510	-0.03662	2.4994	199.9565
160	160	-0.0371	0.6318	101.0949	0.03751	1.2268	196.3018
320	320	-0.0125	0.3186	101.9388	0.04927	0.6420	204.4711

(a) $m = n$.


n	m	Bias	Complex Case MSE	MSE x n	Bias	Real Case MSE	MSE x n
10	15	-0.3343	6.6954	66.9537	-0.3168	12.7099	127.0991
20	30	-0.1781	3.2473	64.9454	-0.1454	6.4439	128.8798
40	60	-0.1126	1.6655	66.6186	-0.08347	3.2470	129.88188
80	120	-0.0565	0.8358	66.8600	-0.02661	1.6381	131.04739
160	240	-0.0287	0.4101	65.6120	0.02318	0.8534	136.5475
320	480	-0.0135	0.2083	66.6571	0.02168	0.4352	139.2527

(b) $m = 1.5n$.


n	m	Bias	Complex Case MSE	MSE x n	Bias	Real Case MSE	MSE x n
10	20	-0.2319	4.9049	49.0494	-0.2764	9.6992	96.9922
20	40	-0.1500	2.5033	50.0666	-0.1657	4.6752	93.5043
40	80	-0.0687	1.2094	48.3761	-0.03922	2.5300	101.2007
80	160	-0.0482	0.6214	49.7090	-0.02426	1.2252	98.0234
160	320	-0.0111	0.3160	50.5613	0.01892	0.6273	100.3799
320	640	-0.0139	0.1580	50.5636	0.02748	0.3267	104.5465

(c) $m = 2n$.

Table 4.5. Algorithm performance for different values of n (dimension of observation vector) and m (number of samples) – both real and complex case.

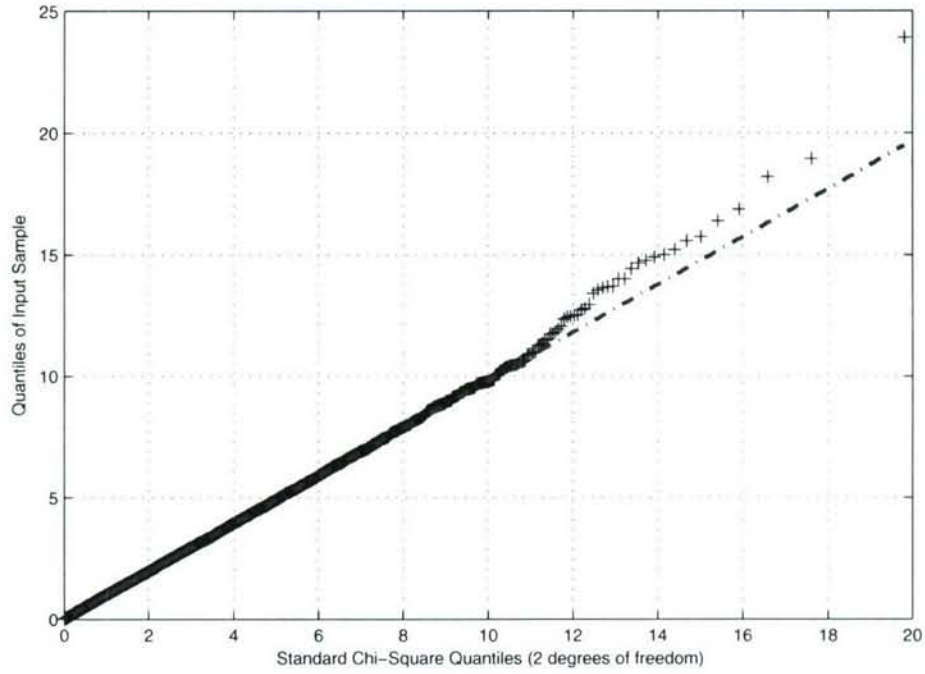



(a) $n = 320, m = 640$ (Complex S).



(b) $n = 320, m = 640$ (Real S).

Figure 4-4. Normal probability plots of the spiked magnitude estimate (true value = 10).

(a) $n = 40, m = 20$ (Complex S).(b) $n = 320, m = 160$ (Complex S).**Figure 4-5.** Sphericity test.

	m = 10	m = 20	m = 40	m = 80	m = 160	m = 320	m = 640
n = 10	0.9491	0.9529	0.9490	0.9465	0.9489	0.9508	0.9498
n = 20	0.9510	0.9478	0.9495	0.9514	0.9493	0.9511	0.9465
n = 40	0.9534	0.9521	0.9480	0.9497	0.9514	0.9473	0.9483
n = 80	0.9491	0.9457	0.9514	0.9547	0.9507	0.9512	0.9489
n = 160	0.9507	0.9472	0.9490	0.9484	0.9464	0.9546	0.9482
n = 320	0.9528	0.9458	0.9448	0.9509	0.9479	0.9486	0.9510

(a) Empirical probability of accepting the null hypothesis when $\Sigma = \mathbf{I}$.

	m = 10	m = 20	m = 40	m = 80	m = 160	m = 320	m = 640
n = 10	0.0009	-	-	-	-	-	-
n = 20	-	-	-	-	-	-	-
n = 40	0.0189	-	-	-	-	-	-
n = 80	0.0829	0.0011	-	-	-	-	-
n = 160	0.2349	0.0258	0.0002	-	-	-	-
n = 320	0.4793	0.1568	0.0062	-	-	-	-

(b) Empirical probability of accepting the null hypothesis when $\Sigma = \Sigma_T$.

Table 4.6. The null-hypothesis is accepted at the 95% significance level for χ_2^2 or whenever $h(\theta) \leq 5.9914$.

The Capon beamformer: Approximating the output distribution

■ 5.1 Introduction

Given a set of independent identically distributed (i.i.d.) signal bearing observations $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_m]$ where each vector is $n \times 1$ zero mean complex circular Gaussian, *i.e.* $\mathbf{x}_i \sim \mathcal{CN}_n(\mathbf{0}, \mathbf{R})$, $i = 1, 2, \dots, m$, Capon proposed a filter-bank approach to power spectral estimation in which he suggested the optimal design of linear filters that pass the desired signal undistorted, while minimizing the power from all other sources of interference [19].

Formally, when the $n \times n$ data covariance matrix is given by \mathbf{R} and the assumed $n \times 1$ array response for a desired signal originating from angle θ is $\mathbf{v}(\theta)$, the solution to the following constrained optimization problem

$$\min_{\mathbf{w}} \mathbf{w}^H \mathbf{R} \mathbf{w} \quad \text{such that} \quad \mathbf{w}^H \mathbf{v}(\theta) = 1 \quad (5.1)$$

satisfies the minimum variance distortionless response (MVDR) criterion leading to the Capon-MVDR filter

$$\mathbf{w}_{MVDR} = \mathbf{R}^{-1} \mathbf{v}(\theta) / \mathbf{v}^H(\theta) \mathbf{R}^{-1} \mathbf{v}(\theta). \quad (5.2)$$

The average output power of this optimal filter is given by

$$P_{Capon}(\theta) \triangleq E \left[|\mathbf{w}_{MVDR}^H \mathbf{x}|^2 \right] = \frac{1}{\mathbf{v}^H(\theta) \mathbf{R}^{-1} \mathbf{v}(\theta)} \quad (5.3)$$

leading to the power spectral estimator

$$\boxed{\hat{P}_{Capon}(\theta) = \frac{1}{\mathbf{v}^H(\theta) \hat{\mathbf{R}}^{-1} \mathbf{v}(\theta)}} \quad (5.4)$$

where $\hat{\mathbf{R}} = (1/m) \mathbf{X} \mathbf{X}^H$ and $m \geq n$ is assumed. This estimator results when $\hat{\mathbf{R}}$ replaces \mathbf{R} in the expression for \mathbf{w}_{MVDR} in (5.2), and this filter is subsequently applied to

the same data used to obtain the covariance estimate $\widehat{\mathbf{R}}$.

When the number of snapshots is greater than but on the order of the number of sensors, *i.e.*, $m \approx n$ the sample covariance matrix is ill-conditioned. Moreover, when the number of snapshots is less than the number of sensors, *i.e.*, $m < n$ the sample covariance matrix is rank deficient (singular). In both of these scenarios the sample covariance matrix is diagonally loaded with a loading value δ (it is necessary in the latter case) to yield the estimate

$$\widehat{\mathbf{R}}_\delta \triangleq \frac{1}{m} \mathbf{X} \mathbf{X}^H + \delta \mathbf{I}. \quad (5.5)$$

The justification for using $\widehat{\mathbf{R}}_\delta$ in place of $\widehat{\mathbf{R}}$ even when $m > n$ is the observation that doing so “robustifies” the signal processing [21, 24, 36]. The covariance matrix estimate $\widehat{\mathbf{R}}_\delta$ thus formed may be interpreted as a structured covariance estimator; in statistics literature, such structured estimators are encountered in the context of shrinkage based approaches to covariance matrix estimation (*e.g.*, [26, 58]).

Two power spectral estimators naturally follow from this modified covariance estimate. The simplest power spectral estimate is obtained by replacing $\widehat{\mathbf{R}}$ with $\widehat{\mathbf{R}}_\delta$ in (5.4) yielding the expression

$$\widehat{P}_{\text{Capon}}(\theta, \delta) \equiv \widehat{P}_{\text{Capon}}^I(\theta, \delta) = \frac{1}{\mathbf{v}^H(\theta) \widehat{\mathbf{R}}_\delta^{-1} \mathbf{v}(\theta)}. \quad (5.6)$$

This estimator was demonstrated to possess inherent robustness properties and yield performance commensurate with the Multiple Signal Classification (MUSIC) algorithm [36].

The alternate estimator is obtained by reformulating Capon’s approach to obtain the constrained optimization problem

$$\min_{\mathbf{w}} \mathbf{w}^H \widehat{\mathbf{R}}_\delta \mathbf{w} \quad \text{such that} \quad \mathbf{w}^H \mathbf{v}(\theta) = 1. \quad (5.7)$$

This leads to the filter

$$\mathbf{w}_\delta = \widehat{\mathbf{R}}_\delta^{-1} \mathbf{v}(\theta) / \mathbf{v}^H(\theta) \widehat{\mathbf{R}}_\delta^{-1} \mathbf{v}(\theta). \quad (5.8)$$

The average output power of this filter conditioned on $\widehat{\mathbf{R}}_\delta$ is given by

$$E \left[|\mathbf{w}_\delta^H \mathbf{x}|^2 \middle| \widehat{\mathbf{R}}_\delta \right] = \frac{\mathbf{v}^H(\theta) \widehat{\mathbf{R}}_\delta^{-1} \mathbf{R} \widehat{\mathbf{R}}_\delta^{-1} \mathbf{v}(\theta)}{\left[\mathbf{v}^H(\theta) \widehat{\mathbf{R}}_\delta^{-1} \mathbf{v}(\theta) \right]^2}. \quad (5.9)$$

Replacing \mathbf{R} with $\widehat{\mathbf{R}}$ in (5.9) yields the second form of a diagonally loaded Capon

spectral estimator

$$\hat{P}_{Capon}^{II}(\theta, \delta) = \frac{\mathbf{v}^H(\theta) \hat{\mathbf{R}}_\delta^{-1} \hat{\mathbf{R}} \hat{\mathbf{R}}_\delta^{-1} \mathbf{v}(\theta)}{\left[\mathbf{v}^H(\theta) \hat{\mathbf{R}}_\delta^{-1} \mathbf{v}(\theta) \right]^2}. \quad (5.10)$$

Note that

$$\lim_{\delta \rightarrow \infty} P_{Capon}^{II}(\theta, \delta) \propto \mathbf{v}^H(\theta) \hat{\mathbf{R}} \mathbf{v}(\theta) = P_{Bartlett}(\theta), \quad (5.11)$$

i.e., as the diagonal loading value increases, this adaptive spectral estimate approaches its conventional beamforming counterpart, known as the Bartlett spectral estimator [102]. While both power spectral estimators are of interest to the array processing community, we focus on the estimator of the form in (5.4) in this chapter.

■ 5.2 Problem formulation

Consider the situation where the i.i.d. observation vectors \mathbf{x}_i for $i = 1, \dots, m$, distributed as $\mathcal{CN}(\mathbf{0}, \mathbf{R})$, have covariance matrix of the form

$$\mathbf{R} = \mathbf{V}(\boldsymbol{\theta}) \mathbf{R}_s \mathbf{V}(\boldsymbol{\theta})^H + \sigma^2 \mathbf{I}, \quad (5.12)$$

where the $n \times k$ matrix $\mathbf{V}(\boldsymbol{\theta}) = [\mathbf{v}(\theta_1), \dots, \mathbf{v}(\theta_k)]$, \mathbf{R}_s is the $k \times k$ covariance matrix of the amplitudes of the k signals, and σ^2 is the variance of the noise process. In array processing applications this models a situation where there are k Gaussian random sources at $\theta_1, \dots, \theta_k$ with array manifold vectors $\mathbf{v}(\theta_1), \dots, \mathbf{v}(\theta_k)$ and we can treat the observation vector \mathbf{x}_i as the superposition of these k Gaussian signals embedded in white noise.

The manifold vector $\mathbf{v}(\theta_i)$ associated with the i -th source is parametrized by the angular location of the source with respect to a chosen coordinate system. The elements of the manifold vector encode how the waves (*e.g.*, electromagnetic or acoustic) impinge on the elements of the sensor array. The manifold (or replica) vector captures the degree of correlation between wavefronts arriving from different directions at the various elements of a sensor array.

The geometry, the relative placement of the sensors on the array, and the propagation characteristics of the operating medium thus play an important role when determining the dependence of the manifold vector on the direction of arrival. This dependence on the direction of arrival θ can be explicitly represented for many array configurations [102, Chapters 2-4]. The simplest array configuration is the uniform linear array depicted in Figure 5-1. Here, as the name suggests, the n sensors are placed uniformly along a line. The manifold vector for this configuration is the $n \times 1$ vector

$$\mathbf{v}(\theta) \equiv \mathbf{v}_\ell(\theta) = \left[1 \quad e^{j2\pi \frac{d}{\ell} \cos \theta} \quad \dots \quad e^{j2\pi (i-1) \frac{d}{\ell} \cos \theta} \quad \dots \quad e^{j2\pi (n-1) \frac{d}{\ell} \cos \theta} \right]^T, \quad (5.13)$$




Figure 5-1. Configuration of a uniform linear array with n sensors and inter-elements spacing of d units.

where d is the inter-element spacing of the sensors and ℓ is the wavelength of the propagating wave (same units as k) so that determining ratio d/ℓ is the inter-element spacing in wavelengths.

When the manifold vector is known, the Capon power spectral estimator can be used to detect the number of signals in white noise. The Capon estimator $\hat{P}_{Capon}(\theta)$ is an estimate of the spatial power spectrum as a function of the scan angle θ . The number of signals present can be estimated by scanning the angle space and determining the number of peaks obtained.

This is illustrated in Figure 5-2 where the theoretical power spectral estimate $P_{Capon}(\theta)$ in (5.3) is compared with the estimates $\hat{P}_{Capon}(\theta, \delta)$ formed using (5.4) for $\delta = 0$ and $\delta = 10$ when $n = m/2 = 18$. Here the observation vectors were sampled in the scenarios where $k = 2$ and the two (independent) sources, *i.e.*, $\mathbf{R}_s = \text{diag}(\sigma_1^2, \sigma_2^2)$, where $\sigma_1^2 = \sigma_2^2 = 100$ with $\sigma^2 = 1$, $\theta_1 = 90^\circ$, $\theta_2 = 70^\circ$, and $d/\ell = 0.45$. As in Figure 5-2, the sources, will (generally) manifest as peaks in the spatial power spectrum estimate.

Note that underlying setup is identical to that considered in Chapter 3; however, unlike the eigen-inference solution proposed in Chapter 3, the Capon-MVDR beamformer exploits information about the eigenvectors of \mathbf{R} encoded by means of the manifold vector $\mathbf{v}(\theta)$. Consequently, provided there is no mismatch between the assumed manifold vector and the true manifold vector, the Capon-MVDR beamformer should be able to identify signals with power levels below the identifiability threshold in Section 3.7.

Detecting the number of sources from the spatial power estimate is challenging because the estimate $\hat{P}_{Capon}(\theta, \delta)$ is a random variable that is a function of the random sample covariance matrix $\hat{\mathbf{R}}$, which has the (complex) Wishart distribution. Thus it becomes important to characterize the distribution of the output $\hat{P}_{Capon}(\theta, \delta)$ and its dependence on $P_{Capon}(\theta)$, the loading level δ , the number of sensors n , and the sample size m . This can facilitate the judicious selection of thresholds for testing the hypothesis that a signal is present while controlling the false discovery rate.




Figure 5-2. The Capon-MVDR spectral estimate when $n = 18$, $m = 36$.

In the situation when $m \geq n$ and $\delta = 0$, the distribution of the outputs is given by the famous Capon-Goodman result [20] which we shall revisit in Section 5.3. The result captures the bias and the variance in $\hat{P}_{Capon}(\theta)$ relative to $P_{Capon}(\theta)$ (the bias can be seen in Figure 5-2) due to finite sample size. The corresponding question for the situation when the Capon-MVDR beamformer is diagonally loaded, *i.e.*, when $\delta > 0$ and general n and m , has remained outstanding in the array processing literature for over four decades. Baggeroer and Cox emphasize its importance and the analytical void in [7, pp. 105]. In their words:

“When using a limited number of snapshots and diagonal loading, significant biases are introduced which can be misleading *vis a vis* the level where a weak signal can be detected. The Capon-Goodman formula is valid only for the case of no loading with $m \geq n$ which is typically not the case for sonars ... Except for the very special case of a single snapshot, we are not aware of any analytic results for the bias when $m < n$ and/or loading is applied even when $m \geq n$.”

In this chapter, we solve this problem by providing analytical expressions for the bias for when $m < n$ and $m \geq n$ and a loading value of δ is applied. We provide stochastic approximations for the distribution of $\hat{\psi}_{Capon}(\theta, \delta) = 1/\hat{P}_{Capon}(\theta, \delta)$ as a function of n , m , and δ for the situation where there are no sources, a single source, and two sources (or a source and an interferer). The results apply for arbitrary array configuration and include the manifold vector mismatch case.

In principle, the results can be extended to approximate the distribution of the outputs when there are an arbitrary number of sources in white noise. However, as we shall shortly see, there will be an accompanying combinatorial explosion in the number of terms needed to obtain an “accurate” approximation. When the expressions start becoming that cumbersome to write down, it is perhaps reasonable to question what, if any, analytical insight they yield that can help the user in practical matters such as determining the “optimal” diagonal loading value or compensating for the induced biases. We suggest that extensions of this work focus on using the relatively simpler approximations in the canonical two or less signals in white noise scenario to piece together a usable approximation when there are more than two signals. This is a matter we shall defer to a later, more thorough investigation.

We note that in contrast, the distribution of $\widehat{P}_{\text{Capon}}(\theta, \delta)$ in the sidelobes, *i.e.*, for values of θ that are not in the proximity of the true signal directions can be approximated by a Normal distribution with mean and variance that are related, in closed form, to the loading value δ , the number of sensors n and the number of snapshots m . The sidelobe level distribution statistic is hence likely to be of greatest utility to the practitioner since it facilitates the setting of sidelobe level thresholds that avoid false signal discovery.

The remainder of this chapter is organized as follows. We review the Capon-Goodman result for the case when $\delta = 0$ and $m \geq n$ in Section 5.3 with the objective of identifying what makes it difficult to extend their analysis to the situation when $\delta \neq 0$. The relevant result from random matrix theory, due to Jack Silverstein [81], is isolated in Section 5.4 and applied in Sections 5.5, 5.6, and 5.7 to characterize the distribution of the Capon-MVDR beamformer outputs where there is no source, a single source and two sources in white noise, respectively. The results are validated using numerical simulations in Section 5.8; extensions and directions for future research are briefly discussed in Section 5.9.

■ 5.3 The Capon-Goodman result

When $\delta = 0$ and $m \geq n$ the distribution of $\widehat{P}_{\text{Capon}}(\theta)$ (when appropriately normalized) is equal to a chi-squared distribution with $2(m - n + 1)$ degrees of freedom (this is equivalent to a so-called *complex* chi-squared distribution with $m - n + 1$ degrees of freedom). This is implicitly stated in the result that follows.

Proposition 5.31. *Let $\widehat{\mathbf{R}}$ be the sample covariance matrix formed from m independent, identically distribution complex valued observation vectors $\mathbf{x}_1, \dots, \mathbf{x}_m$ where, for each $i = 1, \dots, m$, $\mathbf{x} \sim \mathcal{CN}(\mathbf{0}, \mathbf{R})$. When $m \geq n$, and if $\text{Prob.}(\mathbf{a} = \mathbf{0}) = 0$ then,*

$$2m \frac{\mathbf{a}^H \mathbf{R}^{-1} \mathbf{a}}{\mathbf{a}^H \widehat{\mathbf{R}}^{-1} \mathbf{a}} \sim \chi_{2(m-n+1)}^2$$

PROOF. The statement, in the real valued case, follows from Theorem 3.2.12 in [67, pp. 96]. The complex case was derived by Capon and Goodman in [20]. \square

From Proposition 5.31, an application of the well-known formulas for the mean and variance of the chi-squared distribution leads to the famous Capon and Goodman expressions, in (5.14a), for the mean and variance of the Capon-MVDR spectral estimator. The mean and variance of $\psi(\theta, \delta)$ is computed as well using the properties of the inverse chi-squared distribution.

Corollary 5.32. *Assume $\widehat{\mathbf{R}}$, as in Proposition 5.31, is an estimate of the true covariance matrix \mathbf{R} . Let $P_{\text{Capon}}(\theta)$ and $\widehat{P}_{\text{Capon}}(\theta) \equiv \widehat{P}_{\text{Capon}}(\theta, 0)$ be defined as in (5.3) and (5.4) respectively. Then, assuming $m \geq n$,*

$$E \left[\widehat{P}_{\text{Capon}}(\theta) \right] = \frac{m - n + 1}{m} P_{\text{Capon}}(\theta) \quad (5.14a)$$

$$\text{var} \left[\widehat{P}_{\text{Capon}}(\theta) \right] = \frac{m - n + 1}{m^2} P_{\text{Capon}}(\theta)^2. \quad (5.14b)$$

Define $\psi_{\text{Capon}}(\theta) = 1/P_{\text{Capon}}(\theta)$ and $\widehat{\psi}_{\text{Capon}}(\theta) = 1/\widehat{P}_{\text{Capon}}(\theta)$. Then, assuming $m \geq n + 2$,

$$E \left[\widehat{\psi}_{\text{Capon}}(\theta) \right] = \frac{m}{m - n} \psi_{\text{Capon}}(\theta) \quad (5.15a)$$

$$\text{var} \left[\widehat{\psi}_{\text{Capon}}(\theta) \right] = \frac{m^2}{(m - n)^2(m - n - 1)} \psi_{\text{Capon}}(\theta)^2. \quad (5.15b)$$

Equation (5.14a) captures the degradation in the quality of the power spectral estimate due to sample size constraints. Our objective is to mimic the Capon-Goodman result by characterizing the distribution of the Capon-MVDR outputs for the case when $\delta > 0$ and $m < n$. Before we do so, we revisit the special case when $\delta = 0$ with the goal of identifying the key property that allows us to characterize the output distribution as simply as in Proposition 5.31. This will provide insight into why, when $\delta \neq 0$, the characterization has eluded researchers for over four decades.

■ 5.3.1 Structure exploited in Capon and Goodman's analysis

We first consider the scenario with no diagonal loading, *i.e.*, $\delta = 0$. The inverse of the MVDR beamformer output is then given by

$$\widehat{\psi}_{\text{Capon}}(\theta) \equiv \widehat{\psi}_{\text{Capon}}(\theta, 0) = \mathbf{v}^H(\theta) \widehat{\mathbf{R}}^{-1} \mathbf{v}(\theta). \quad (5.16)$$

Assume that $m > n$ so that the sample covariance matrix is not singular and can be decomposed as

$$\widehat{\mathbf{R}} = \mathbf{R}^{1/2} \mathbf{W}(c) \mathbf{R}^{1/2}, \quad (5.17)$$

where \mathbf{R} is the true covariance matrix and $\mathbf{W}(c)$ (here and henceforth) is the complex Wishart random matrix with identity covariance. We parameterize the Wishart matrix by $c = n/m$ which is the ratio of the number of sensors to snapshots.

When there is no diagonal loading, $\hat{\psi}(\theta)$ in (5.16), can be decomposed as

$$\hat{\psi}_{\text{Capon}}(\theta) = \underbrace{\mathbf{v}^H(\theta)\mathbf{R}^{-1}\mathbf{v}(\theta)}_{\text{Deterministic term}} \underbrace{\mathbf{u}^H(\theta)\mathbf{W}(c)^{-1}\mathbf{u}(\theta)}_{\text{Stochastic term}} \quad (5.18)$$

where $\mathbf{u}(\theta)$ is a unit vector such that $\mathbf{R}^{-1/2}\mathbf{v}(\theta) = \alpha(\theta)\mathbf{u}(\theta)$ and $|\alpha(\theta)|^2 = \psi_{\text{Capon}}(\theta) = \mathbf{v}^H(\theta)\mathbf{R}^{-1}\mathbf{v}(\theta)$.

Recall that $\psi_{\text{Capon}}(\theta) = 1/P_{\text{Capon}}(\theta)$ where $P_{\text{Capon}}(\theta)$ is the spectral estimate when the true covariance matrix \mathbf{R} is known. The stochastic term is a quadratic form involving the Wishart matrix. Thus, as a function of θ , when there is no diagonal loading, the probability distribution of the MVDR beamformer is completely characterized by the single stochastic term in (5.18) which has an inverse chi-squared distributions from Proposition 5.31.

In essence the decomposability of the quadratic form into the stochastic and the deterministic components is exploited in the derivation of the chi-squared distribution for $\hat{P}_{\text{Capon}}(\theta)$ in the famous Capon-Goodman paper [20]. The ability to do so implies the true covariance matrix \mathbf{R} appears in the solution only in the form of a deterministic scale factor as in (5.18). This means that the relative bias and variance of the outputs will be identical across the entire scan angle space as demonstrated as can be seen in (5.14a). More importantly the distribution thus computed applies for arbitrary \mathbf{R} so that the model in (5.12) is merely a special case.

When the Capon-MVDR processor is diagonally loaded, it is no longer possible to decouple the stochastic part from the deterministic part. In particular, the distribution of $\hat{\psi}_{\text{Capon}}(\theta, \delta)$ will explicitly depend on the structure of the true covariance matrix and the approximations we develop will only apply for the model in (5.12).

■ 5.4 Relevant result from random matrix theory

The distributional approximations for $\hat{\psi}(\theta, \delta)$ that we shall develop rely on the following asymptotic characterization of quadratic forms of functions of complex Wishart matrices with identity covariance.

Proposition 5.41. *Let \mathbf{u} and \mathbf{u}_\perp be two fixed mutually orthogonal $n \times 1$ unit vectors. Let $\mathbf{W}_\delta(c) = (\mathbf{W}(c) + \delta\mathbf{I}_n)$ where $\mathbf{W}(c)$ is a complex Wishart matrix with covariance identity. Then, as $n, m \rightarrow \infty$, with $n/m \rightarrow c > 0$,*

$$\sqrt{n}(\mathbf{u}^H\mathbf{W}_\delta^{-1}(c)\mathbf{u} - \mu_\delta) \xrightarrow{\mathcal{D}} q_1 \sim \mathcal{N}(0, \sigma_\delta^2) \quad (5.19a)$$

$$\sqrt{n}(\mathbf{u}^H\mathbf{W}_\delta^{-1}(c)\mathbf{u}_\perp) \xrightarrow{\mathcal{D}} q_2 \sim \mathcal{CN}(0, \sigma_\delta^2/2) \quad (5.19b)$$

$$\sqrt{n}(\mathbf{u}_\perp^H\mathbf{W}_\delta^{-1}(c)\mathbf{u}_\perp - \mu_\delta) \xrightarrow{\mathcal{D}} q_3 \sim \mathcal{N}(0, \sigma_\delta^2) \quad (5.19c)$$

where the convergence in distribution is almost surely and

$$\mu_\delta = \frac{-1 + c - \delta + \sqrt{1 - 2c + 2\delta + c^2 + 2c\delta + \delta^2}}{2c\delta} \quad (5.20)$$

$$\sigma_\delta^2 = -\frac{\partial \mu_\delta}{\partial \delta} - \mu_\delta^2. \quad (5.21)$$

PROOF. The result follows from an extension [81] of the techniques developed by Silverstein in [84, 85]. The mean and variance are obtained by evaluating the integrals

$$\mu_\delta = \int \frac{1}{x + \delta} dF^{W(c)}(x) \quad (5.22a)$$

$$\sigma_\delta^2 = \int \frac{1}{x + \delta} dF^{W(c)}(x) - \mu_\delta^2 \quad (5.22b)$$

where $dF^{W(c)}$ is the Marčenko-Pastur density in (3.15). \square

Thus, for large enough n and m , Proposition 5.41 suggests that we can approximate the quadratic forms $r_1 = \mathbf{u}^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}$, $r_2 = \mathbf{u}^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_\perp$ and $r_3 = \mathbf{u}_\perp^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_\perp$ by independent Gaussian random variables where r_1 and r_3 are identically distributed real-valued Gaussian random variables with mean μ_δ and variance σ_δ^2/n and r_2 is a complex-valued Gaussian random variable whose real and imaginary parts are independent and identically distributed Gaussian random variables with mean μ_δ and variance $\sigma_\delta^2/(4n)$. In deriving the distributional approximation, whenever we encounter such quadratic forms formed from orthogonal unit vectors, we shall replace them with independent normally distributed variables as in Proposition 5.41.

The accuracy of this asymptotic approximation even when n and m are of moderate size can be discerned by comparing (5.15) with the result obtained using Proposition 5.41. From (5.18) and (5.19a) we have

$$\widehat{\psi}_{Capon}(\theta) = \mathbf{v}^H(\theta) \mathbf{R}^{-1} \mathbf{v}(\theta) \mathbf{u}^H \mathbf{W}_0^{-1}(c) \mathbf{u} = \psi_{Capon}(\theta) r_1 \quad (5.23)$$

where

$$r_1 \stackrel{d}{\simeq} \mathcal{N} \left(\frac{1}{1-c}, \frac{1}{n} \frac{c^2}{(1-c)^3} \right), \quad (5.24)$$

so that we obtain the approximations

$$E \left[\widehat{\psi}_{Capon}(\theta) \right] \approx \frac{n}{m-n} \psi_{Capon}(\theta) \quad (5.25a)$$

$$\text{var} \left[\widehat{\psi}_{Capon}(\theta) \right] \approx \frac{m^2}{(m-n)^3} \psi_{Capon}(\theta)^2 \quad (5.25b)$$

where we have substituted $c = n/m$ in (5.24). Comparing (5.15) with (5.25) shows the accuracy of the approximation used.

■ 5.5 Distribution under diagonal loading: No sources

Consider the null hypothesis - where there are no sources in white noise of variance $\sigma^2 = 1$ so that $\mathbf{R} = \mathbf{I}$. We assume that diagonal loading is applied, *i.e.*, $\delta > 0$. The statistic

$$\widehat{\psi}_{\text{Capon}}(\theta, \delta) = \mathbf{v}^H(\theta) (\mathbf{W}(c) + \delta \mathbf{I}_n)^{-1} \mathbf{v}(\theta) \quad (5.26)$$

can be approximated using Proposition 5.41 as follows. Figure 5-3 validates the approximation for moderate m and n . This approximation can also be used in the sidelobe region when there are many sources in white noise.

Approximation 5.51 (No sources/sidelobe region).

$$\widehat{\psi}(\theta, \delta) = r_1 \stackrel{D}{\simeq} \mathcal{N}(\|\mathbf{v}(\theta)\|^2 \mu_\delta, \|\mathbf{v}(\theta)\|^4 \sigma_\delta^2/n) \quad (5.27)$$

■ 5.6 Distribution under diagonal loading: Single source

Consider the scenario where there is a single source. Assume, without loss of generality, that $\sigma^2 = 1$ so that the covariance matrix $\mathbf{R} = \sigma_S^2 \mathbf{v}(\theta_S) \mathbf{v}(\theta_S)^H + \mathbf{I}$ where θ_S is the direction of the source and σ_S^2 is the corresponding source power. Given a vector $\mathbf{v}(\theta)$, we construct the $\mathbf{u}(\theta)$ as

$$\mathbf{u}(\theta) = \frac{\mathbf{v}(\theta)}{\|\mathbf{v}(\theta)\|}. \quad (5.28)$$

Note that, $\mathbf{u}(\theta_S)$ is an eigenvector of \mathbf{R} . The covariance matrix \mathbf{R} can hence be decomposed as

$$\mathbf{R} = (\alpha + 1) \mathbf{u}(\theta_S) \mathbf{u}^H(\theta_S) + \mathbf{U}_\perp(\theta_S) \mathbf{U}_\perp^H(\theta_S), \quad (5.29)$$

where $\alpha = \sigma_S^2 \|\mathbf{v}(\theta)\|^2$ and $\mathbf{U}_\perp(\theta_S)$ is an $n \times (n - 1)$ matrix orthogonal to $\mathbf{u}(\theta_S)$ such that $\mathbf{U}_\perp^H(\theta_S) \mathbf{U}_\perp(\theta_S) = \mathbf{I}_{n-1}$. Hence, we have

$$\mathbf{R}^{-1} = \frac{1}{\alpha + 1} \mathbf{u}(\theta_S) \mathbf{u}^H(\theta_S) + \mathbf{U}_\perp(\theta_S) \mathbf{U}_\perp^H(\theta_S), \quad (5.30)$$

so that

$$\widehat{\psi}(\theta, \delta) = \mathbf{v}^H(\theta) \widehat{\mathbf{R}}_\delta^{-1} \mathbf{v}(\theta). \quad (5.31)$$

We can rewrite (5.31) as

$$\widehat{\psi}(\theta, \delta) = \mathbf{v}^H(\theta) \mathbf{R}^{-\frac{1}{2}} (\mathbf{W}(c) + \delta \mathbf{R}^{-1})^{-1} \mathbf{R}^{-\frac{1}{2}} \mathbf{v}(\theta). \quad (5.32)$$



(a) $n = 36, m = 18, \delta = 10$ (b) $n = 108, m = 54, \delta = 10$

Figure 5-3. Assessing the validity of the approximated distribution for $\hat{\psi}(\theta, \delta)$ when there are no sources in white noise.

Let us first consider the matricial term $(\mathbf{W}(c) + \delta \mathbf{R}^{-1})^{-1}$, which may be rewritten as

$$\begin{aligned} (\mathbf{W}(c) + \delta \mathbf{R}^{-1})^{-1} &= (\mathbf{W}(c) + \delta \mathbf{I} + \delta \mathbf{R}^{-1} - \mathbf{I})^{-1} \\ &= \left(\mathbf{W}_\delta(c) + \underbrace{\delta \mathbf{R}^{-1} - \delta \mathbf{I}}_{\mathbf{D}} \right)^{-1}. \end{aligned}$$

From (5.30), we note that the matrix \mathbf{D} is a rank-one matrix of the form

$$\mathbf{D} = \delta \underbrace{\left(\frac{1}{\alpha + 1} - 1 \right)}_{1/d} \mathbf{u}(\theta_S) \mathbf{u}^H(\theta_S). \quad (5.33)$$

Using the Sherman-Morrison-Woodbury matrix inversion lemma [41], we have

$$(\mathbf{W}(c) + \delta \mathbf{R}^{-1})^{-1} = \mathbf{W}_\delta^{-1}(c) - \frac{\mathbf{W}_\delta^{-1}(c) \mathbf{u}(\theta_S) \mathbf{u}^H(\theta_S) \mathbf{W}_\delta^{-1}(c)}{d + \mathbf{u}^H(\theta_S) \mathbf{W}_\delta^{-1}(c) \mathbf{u}(\theta_S)}. \quad (5.34)$$

In (5.32), the term $\mathbf{v}^H(\theta) \mathbf{R}^{-\frac{1}{2}}$ can also be written in terms of $\mathbf{u}(\theta_S)$ as

$$\mathbf{v}^H(\theta) \mathbf{R}^{-\frac{1}{2}} = \|\mathbf{v}(\theta)\| \frac{\langle \mathbf{u}(\theta), \mathbf{u}(\theta_S) \rangle}{\sqrt{\alpha + 1}} \mathbf{u}^H(\theta_S) + \|\mathbf{v}(\theta)\| \langle \mathbf{u}(\theta), \mathbf{U}_\perp(\theta_S) \rangle \mathbf{U}_\perp^H(\theta_S) \quad (5.35)$$

or, equivalently as

$$\mathbf{v}^H(\theta) \mathbf{R}^{-\frac{1}{2}} = \beta \mathbf{u}^H(\theta_S) + \gamma \mathbf{u}_\perp^H(\theta_S), \quad (5.36)$$

where $\beta = \|\mathbf{v}(\theta)\| \langle \mathbf{u}(\theta), \mathbf{u}(\theta_S) \rangle / \sqrt{\alpha + 1}$, and $\mathbf{u}_\perp(\theta_S)$ is an $n \times 1$ unit vector such that $\gamma \mathbf{u}_\perp(\theta_S)$ is equal to the second term on the right hand side of (5.35).

On substituting (5.34) and (5.36) into (5.32), and performing some algebraic manipulations we obtain the expression in (5.37) below. Applying Proposition 5.41 gives us the stochastic approximation for the distribution of $\hat{\psi}(\theta, \delta)$ composed using independent normally distributed random variables as described below.

Approximation 5.61 (Single source in white noise).

$$\widehat{\psi}(\theta, \delta) = |\beta|^2 \left(\frac{d r_1}{d + r_1} \right) + 2 \operatorname{Re} \left(\beta \gamma^* \left[\frac{d r_2}{d + r_1} \right] \right) + |\gamma|^2 \left(r_3 - \frac{|r_2|^2}{d + r_1} \right). \quad (5.37)$$

where

$$r_1 = \mathbf{u}^H(\theta_S) \mathbf{W}_\delta^{-1}(c) \mathbf{u}(\theta_S) \stackrel{D}{\simeq} \mathcal{N}(\mu_\delta, \sigma_\delta^2/n) \quad (5.38a)$$

$$r_2 = \mathbf{u}^H(\theta_S) \mathbf{W}_\delta^{-1}(c) \mathbf{u}_\perp(\theta_S) \stackrel{D}{\simeq} \mathcal{CN}(0, \sigma_\delta^2/2n) \quad (5.38b)$$

$$r_3 = \mathbf{u}_\perp^H(\theta_S) \mathbf{W}_\delta^{-1}(c) \mathbf{u}_\perp(\theta_S) \stackrel{D}{\simeq} \mathcal{N}(\mu_\delta, \sigma_\delta^2/n) \quad (5.38c)$$

It is worth noting how much more complicated the structure of (5.37) is compared to (5.27). This is evidence of the fact that when diagonal loading is applied, the probability distribution of the outputs depends in a more complicated manner on the underlying structure of the covariance matrix \mathbf{R} . Nonetheless, (5.37) is an exact expression for $\widehat{\psi}_{\text{Capon}}(\theta, \delta)$ and the *approximated stochastic representation* relies on treating the variables r_1 , r_2 , and r_3 as independent Gaussian random variables. Though the distribution does not have a nicely expressible density function, we can efficiently sample from it using this approximate stochastic representation.

■ 5.7 Distribution under diagonal loading: Two sources

When there are sources in white noise of variance $\sigma^2 = 1$ we have

$$\mathbf{R} = [\mathbf{v}(\theta_1) \ \mathbf{v}(\theta_2)] \mathbf{R}_s [\mathbf{v}(\theta_1) \ \mathbf{v}(\theta_2)]^H + \mathbf{I} \quad (5.39)$$

$$= [\mathbf{u}(\theta_1) \ \mathbf{u}(\theta_2)] \begin{bmatrix} \alpha_1 & 0 \\ 0 & \alpha_2 \end{bmatrix} [\mathbf{u}(\theta_1) \ \mathbf{u}(\theta_2)]^H + \mathbf{I} \quad (5.40)$$

where \mathbf{u}_1 and \mathbf{u}_2 are the eigenvectors corresponding to the two largest “signal” eigenvalues of \mathbf{R} . The inverse covariance matrix is given by

$$\mathbf{R}^{-1} = \frac{1}{1 + \alpha_1} \mathbf{u}_1 \mathbf{u}_1^H + \frac{1}{1 + \alpha_2} \mathbf{u}_2 \mathbf{u}_2^H + \mathbf{I}, \quad (5.41)$$

so that the matrix $\mathbf{D} \triangleq \mathbf{R}^{-1} - \delta \mathbf{I}$ given by

$$\mathbf{D} = \underbrace{\left(\frac{\delta}{\alpha_1 + 1} - 1\right)}_{1/d_1} \mathbf{u}_1 \mathbf{u}_1^H + \underbrace{\left(\frac{\delta}{\alpha_2 + 1} - 1\right)}_{1/d_2} \mathbf{u}_2 \mathbf{u}_2^H + \mathbf{I} \quad (5.42)$$

$$= [\mathbf{u}_1 \ \mathbf{u}_2] \begin{bmatrix} 1/d_1 & 0 \\ 0 & 1/d_2 \end{bmatrix} [\mathbf{u}_1 \ \mathbf{u}_2]^H, \quad (5.43)$$

is a rank-two matrix. Applying the Sherman-Morrison-Woodbury matrix inversion lemma to $(\mathbf{W}_\delta(c) + \mathbf{D})^{-1}$ we have

$$(\mathbf{W}_\delta(c) + \mathbf{D})^{-1} = \mathbf{W}_\delta(c)^{-1} - \mathbf{W}_\delta^{-1}(c) [\mathbf{u}_1 \ \mathbf{u}_2] \mathbf{T}_\delta^{-1} [\mathbf{u}_1 \ \mathbf{u}_2]^H, \quad (5.44)$$

where

$$\mathbf{T}_\delta^{-1} = \left(\begin{bmatrix} d_1 & 0 \\ 0 & d_2 \end{bmatrix} + \begin{bmatrix} \mathbf{u}_1^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_1 & \mathbf{u}_1^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_2 \\ \mathbf{u}_2^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_1 & \mathbf{u}_2^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_2 \end{bmatrix} \right)^{-1}. \quad (5.45)$$

To simplify the analysis we assume that δ is large enough so that $\mathbf{u}_1^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_1 \gg \mathbf{u}_2^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_1$ so that the approximation

$$\mathbf{T}_\delta^{-1} \approx \begin{bmatrix} \frac{1}{d_1 + \mathbf{u}_1^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_1} & 0 \\ 0 & \frac{1}{d_2 + \mathbf{u}_2^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_2} \end{bmatrix} \quad (5.46)$$

holds. Substituting (5.46) into (5.44) we have

$$(\mathbf{W}_\delta(c) + \mathbf{D})^{-1} \approx \mathbf{W}_\delta^{-1}(c) - \frac{\mathbf{W}_\delta^{-1}(c) \mathbf{u}_1 \mathbf{u}_1^H \mathbf{W}_\delta^{-1}(c)}{d_1 + \mathbf{u}_1^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_1} - \frac{\mathbf{W}_\delta^{-1}(c) \mathbf{u}_2 \mathbf{u}_2^H \mathbf{W}_\delta^{-1}(c)}{d_2 + \mathbf{u}_2^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_2}. \quad (5.47)$$

Assume that we can decompose $\mathbf{v}^H(\theta) \mathbf{R}^{-\frac{1}{2}}$ as

$$\mathbf{v}^H(\theta) \mathbf{R}^{-\frac{1}{2}} = \|\mathbf{v}(\theta)\| \frac{\langle \mathbf{u}(\theta), \mathbf{u}_1 \rangle}{\sqrt{\alpha_1 + 1}} \mathbf{u}_1^H + \|\mathbf{v}(\theta)\| \frac{\langle \mathbf{u}(\theta), \mathbf{u}_2 \rangle}{\sqrt{\alpha_2 + 1}} \mathbf{u}_2^H + \|\mathbf{v}(\theta)\| \langle \mathbf{u}(\theta), \mathbf{U}_\perp \rangle \mathbf{U}_\perp^H \quad (5.48)$$

so that we have

$$\mathbf{v}^H(\theta) \mathbf{R}^{-\frac{1}{2}} = \beta_1 \mathbf{u}_1^H + \beta_2 \mathbf{u}_2^H + \gamma \mathbf{u}_\perp^H, \quad (5.49)$$

where β_1 , β_2 and γ are non-random parameters obtained by comparing the representations in (5.48) and (5.49), term by term. Once these parameters are computed for the non-random \mathbf{u}_1 and \mathbf{u}_2 given by (5.39) we can apply Proposition 5.41 to obtain the

approximated stochastic representation for $\widehat{\psi}(\theta, \delta)$ given below.

Approximation 5.71 (Two sources in white noise).

$$\begin{aligned} \widehat{\psi}(\theta, \delta) \approx & |\beta_1|^2 \left(\frac{d_1 r_1}{d_1 + r_1} - \frac{|r_1|^2}{d_2 + r_4} \right) + |\beta_2|^2 \left(\frac{d_2 r_4}{d_2 + r_4} - \frac{|r_2|^2}{d_1 + r_1} \right) \\ & + |\gamma|^2 \left(r_6 - \frac{|r_3|^2}{d_1 + r_1} - \frac{|r_5|^2}{d_2 + r_4} \right) + 2\text{Re} \left(\beta_1 \beta_2^* \left[\frac{d_1 r_2}{d_1 + r_1} - \frac{r_2 r_4}{d_2 + r_4} \right] \right) \\ & + 2\text{Re} \left(\beta_1 \gamma^* \left[\frac{d_1 r_3}{d_1 + r_1} - \frac{r_2 r_5}{d_2 + r_4} \right] \right) + 2\text{Re} \left(\beta_2 \gamma^* \left[\frac{d_2 r_5}{d_2 + r_4} - \frac{r_2^* r_3}{d_1 + r_1} \right] \right) \end{aligned} \quad (5.50)$$

$$r_1 = \mathbf{u}_1^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_1 \stackrel{D}{\simeq} \mathcal{N}(\mu_\delta, \sigma_\delta^2/n) \quad (5.51a)$$

$$r_2 = \mathbf{u}_1^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_2 \stackrel{D}{\simeq} \mathcal{CN}(0, \sigma_\delta^2/2n) \quad (5.51b)$$

$$r_3 = \mathbf{u}_1^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_\perp \stackrel{D}{\simeq} \mathcal{CN}(0, \sigma_\delta^2/2n) \quad (5.51c)$$

$$r_4 = \mathbf{u}_2^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_2 \stackrel{D}{\simeq} \mathcal{N}(\mu_\delta, \sigma_\delta^2/n) \quad (5.51d)$$

$$r_5 = \mathbf{u}_2^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_\perp \stackrel{D}{\simeq} \mathcal{CN}(0, \sigma_\delta^2/2n) \quad (5.51e)$$

$$r_6 = \mathbf{u}_\perp^H \mathbf{W}_\delta^{-1}(c) \mathbf{u}_\perp \stackrel{D}{\simeq} \mathcal{N}(\mu_\delta, \sigma_\delta^2/n) \quad (5.51f)$$

Note that there are two levels of approximation in the stochastic representation derived. Firstly that \mathbf{T}_δ^{-1} could be written in the simpler form as in (5.46) and secondly that the six quadratic forms r_1, \dots, r_6 can be treated as independent random variables. The conditions where the first assumption holds needs to be rigorously investigated.

We can also use the approximation (5.50) to derive an approximation for the covariance matrix $E \left[(\widehat{\mathbf{R}} + \delta \mathbf{I})^{-1} \right]$. Note that in (5.50), when $\mathbf{v}(\theta) = \mathbf{u}_1$ then $\beta_1 \neq 0$ but $\beta_2 = \gamma = 0$ in (5.49) so that we have

$$E[\mathbf{u}_1^H (\widehat{\mathbf{R}} + \delta \mathbf{I})^{-1} \mathbf{u}_1] = E \left[|\beta_1|^2 \left(\frac{d_1 r_1}{d_1 + r_1} - \frac{|r_1|^2}{d_2 + r_4} \right) \right].$$

The quadratic forms $\mathbf{u}_2^H (\widehat{\mathbf{R}} + \delta \mathbf{I})^{-1} \mathbf{u}_2$ and $\mathbf{u}_\perp^H (\widehat{\mathbf{R}} + \delta \mathbf{I})^{-1} \mathbf{u}_\perp$ exhibit a similar simple structure. Correspondingly, the quadratic form $\mathbf{u}_1^H (\widehat{\mathbf{R}} + \delta \mathbf{I})^{-1} \mathbf{u}_\perp$ can be expressed as

$$\mathbf{u}_1^H (\widehat{\mathbf{R}} + \delta \mathbf{I})^{-1} \mathbf{u}_\perp = \beta_1 \beta_2^* \left[r_2 \left(\frac{d_1}{d_1 + r_1} - \frac{r_4}{d_2 + r_4} \right) \right] \quad (5.52)$$

so that since $E[r_2] \approx 0$, we have

$$E[\mathbf{u}_1^H (\widehat{\mathbf{R}} + \delta \mathbf{I})^{-1} \mathbf{u}_2] \approx 0. \quad (5.53)$$

This suggests that the covariance matrix $E[(\widehat{\mathbf{R}} + \delta \mathbf{I})^{-1}]$ is (approximately) diagonalized by the eigenvectors of \mathbf{R} so that it can be approximated as in (5.55). We can use Taylor series expansions to obtain the diagonal elements of Ψ_δ that captures the effect of the diagonal loading value δ , and the ratio $c = n/m$. Once we compute this covariance matrix, it is easy to compute to $E[\widehat{\psi}(\theta, \delta)]$ since we have

$$E[\widehat{\psi}(\theta, \delta)] = \mathbf{v}(\theta)^H E[(\widehat{\mathbf{R}} + \delta \mathbf{I})^{-1}] \mathbf{v}(\theta). \quad (5.54)$$

Approximation 5.72 (Two sources in white noise).

$$E[(\widehat{\mathbf{R}} + \delta \mathbf{I})^{-1}] \approx \mathbf{U} \Psi_\delta \mathbf{U} \quad (5.55)$$

where \mathbf{U} diagonalizes \mathbf{R} and Ψ_δ is a diagonal matrix with

$$(\Psi_\delta)_{i,i} = \begin{cases} E \left[|\beta_1|^2 \left(\frac{d_1 r_1}{d_1 + r_1} - \frac{|r_1|^2}{d_2 + r_4} \right) \right] & \text{for } i = 1, \\ E \left[|\beta_2|^2 \left(\frac{d_2 r_4}{d_2 + r_4} - \frac{|r_2|^2}{d_1 + r_1} \right) \right] & \text{for } i = 2, \\ E \left[|\gamma|^2 \left(r_6 - \frac{|r_3|^2}{d_1 + r_1} - \frac{|r_5|^2}{d_2 + r_4} \right) \right] & \text{for } i = 3, \dots, n. \end{cases} \quad (5.56)$$

■ 5.8 Numerical examples

Consider a scenario involving a single source plus interferer and a set of signal bearing snapshots $\mathbf{x}_i \sim \mathcal{CN}[0, \mathbf{I}_n + \sigma_S^2 \mathbf{v}(\theta_T) \mathbf{v}(\theta_T)^H + \sigma_I^2 \mathbf{v}(\theta_I) \mathbf{v}(\theta_I)^H]$, for $n = 1, 2, \dots, m$ for an $n = 18$ element uniform linear array (ULA) with slightly less than $\lambda/2$ element spacing. The array has a 3 dB beam width of 7.2 degrees and the desired target signal is arbitrarily placed at $\theta_T = 90$ degrees (array broadside) while the interferer is arbitrarily placed at $\theta_I = 70$ degrees.

Figures 5-4 and 5-5 illustrate the success of the predicted (inverse) of the diagonally Capon-MVDR spectral estimator, *i.e.*, the denominator of (5.4) and the results obtained for the same from 4000 Monte Carlo simulations (red circles) for snapshot deficient cases. Note that the Capon-Goodman result cannot be used here since they require $m \geq n$ and no diagonal loading. The denominator of the Capon estimator in (5.3) constructed

using the true covariance matrix is plotted for reference.

Consider yet another such scenario involving a single source plus interferer and a set of signal bearing snapshots $\mathbf{x}_i \sim \mathcal{CN}[0, I + \sigma_S^2 \mathbf{v}(\theta_T) \mathbf{v}(\theta_T) + \sigma_I^2 \mathbf{v}(\theta_I) \mathbf{v}(\theta_I)]$, for $n = 1, 2, \dots, m$ for an $n = 18$ element uniform linear array (ULA) with slightly less than $\lambda/2$ element spacing. The array has a 3 dB beamwidth of 7.2 degrees and the desired target signal is arbitrarily placed at $\theta_T = 90$ degrees (array broadside) while the interferer is arbitrarily placed at $\theta_I = 95$ degrees so that the source and interferer are closely spaced.

Figure 5-6 illustrates the success of the predicted (inverse) of the diagonally Capon-MVDR spectral estimator, *i.e.*, the denominator of (5.4) and the results obtained for the same from 4000 Monte Carlo simulations (red circles) for snapshot deficient cases. Note how as the source and interferer power reduces from 20 dB to 0 dB, the resolution of the Capon-MVDR beamformer is adversely affected as observed in practice. The denominator of the Capon estimator in (5.3) constructed using the true covariance matrix is plotted for reference. The availability and great accuracy of these analytical predictions in the snapshot deficient case promises to have a major impact on the analysis of the Capon-MVDR algorithm beyond the threshold SNR where its performance is known to degrade dramatically.

■ 5.9 Future work

In the spirit of the original Capon-Goodman result for the case with no diagonal loading, we were able to use the knowledge of the distribution to analytically predict the beampattern induced by diagonal loading. In other words, we approximated the distribution of the random variable $P_{Capon}^I(\theta, \delta)$ as a function of θ for a given value of δ . The predictions were shown to be accurate *vis a vis* the numerical simulations. The most important implication of this for practice is that this understanding can help facilitate the analysis, for the first time, of the performance of applications such as direction-of-arrival (DOA) estimation or direction finding (DF) that use a diagonally loaded Capon-MVDR processor in the snapshot deficient case. Initial results in this direction may be found in [75].



(a) $n = 18, m = 4, \delta = 0.1, \sigma_S^2 = \sigma_I^2 = 100$ (b) $n = 18, m = 4, \delta = 10, \sigma_S^2 = \sigma_I^2 = 100$

Figure 5-4. Two equal power sources at 90 degrees and 70 degrees.

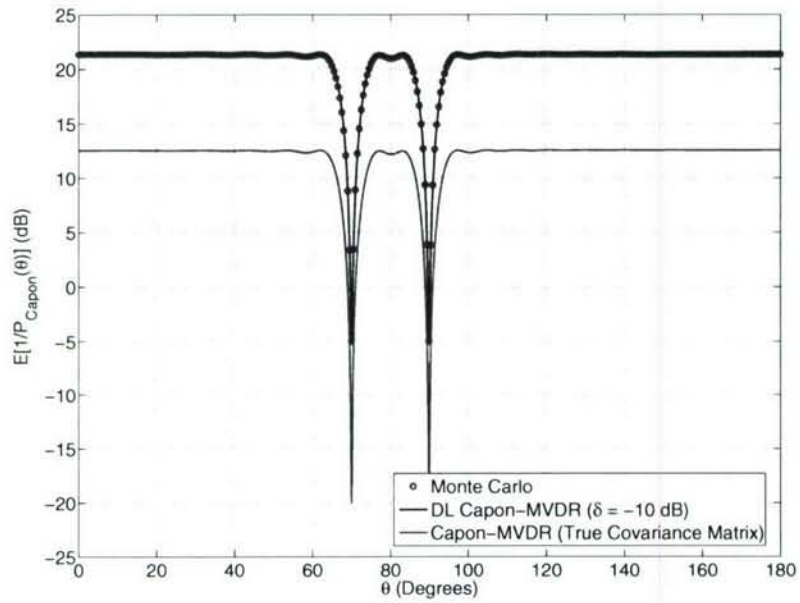

(a) $n = 18, m = 6, \delta = 0.1, \sigma_S^2 = \sigma_I^2 = 100$ (b) $n = 18, m = 6, \delta = 10, \sigma_S^2 = \sigma_I^2 = 100$

Figure 5-5. Two equal power sources at 90 degrees and 70 degrees.



(a) $n = 18, m = 4, \delta = 10, \sigma_S^2 = \sigma_I^2 = 100$ (b) $n = 18, m = 4, \delta = 10, \sigma_S^2 = \sigma_I^2 = 1$

Figure 5-6. Two Equal Power sources at 90 degrees and 95 degrees.

The polynomial method: Mathematical foundation

This chapter marks the start of the second part of the dissertation. Here, we lay the foundation for a powerful method that allows us to calculate the limiting eigenvalue distribution of a large class of random matrices. We see this method as allowing us to expand our reach beyond the well known special random matrices whose limiting distributions have the semi-circle density [113], the Marčenko-Pastur density [59], the McKay density [61] or their close cousins [22, 82].

In particular, we encode transforms of the limiting eigenvalue distribution function as solutions of a bivariate polynomial equation. Then canonical operations on the random matrices become operations on the bivariate polynomials. Before delving into a description of a class of random matrices for which this characterization applies, we describe the various ways in which transforms of the underlying probability distribution function can be encoded and manipulated.

■ 6.1 Transform representations

■ 6.1.1 The Stieltjes transform and some minor variations

The Stieltjes transform of the distribution function $F^A(x)$ is given by

$$m_A(z) = \int \frac{1}{x-z} dF^A(x) \quad \text{for } z \in \mathbb{C}^+ \setminus \mathbb{R}. \quad (6.1)$$

The Stieltjes transform may be interpreted as the expectation

$$m_A(z) = E_x \left[\frac{1}{x-z} \right],$$

with respect to the random variable x with distribution function $F^A(x)$. Consequently, for any invertible function $h(x)$ continuous over the support of $dF^A(x)$, the Stieltjes transform $m_A(z)$ can also be written in terms of the distribution of the random variable

$y = h(x)$ as

$$m_A(z) = E_x \left[\frac{1}{x-z} \right] = E_y \left[\frac{1}{h^{(-1)}(y)-z} \right], \quad (6.2)$$

where $h^{(-1)}(\cdot)$ is the inverse of $h(\cdot)$ with respect to composition i.e. $h(h^{(-1)}(x)) = x$. Equivalently, for $y = h(x)$, we obtain the relationship

$$E_y \left[\frac{1}{y-z} \right] = E_x \left[\frac{1}{h(x)-z} \right]. \quad (6.3)$$

The well-known Stieltjes-Perron inversion formula [3]

$$f_A(x) \equiv dF^A(x) = \frac{1}{\pi} \lim_{\xi \rightarrow 0^+} \text{Im } m_A(x + i\xi). \quad (6.4)$$

can be used to recover the probability density function $f_A(x)$ from the Stieltjes transform. Here and for the remainder of this thesis, the density function is assumed to be distributional derivative of the distribution function. In a portion of the literature on random matrices, the Cauchy transform is defined as

$$g_A(z) = \int \frac{1}{z-x} dF^A(x) \quad \text{for } z \in \mathbb{C}^{-1} \setminus \mathbb{R}.$$

The Cauchy transform is related to the Stieltjes transform, as defined in (6.1), by

$$g_A(z) = -m_A(z). \quad (6.5)$$

■ 6.1.2 The moment transform

When the probability distribution is compactly supported, the Stieltjes transform can also be expressed as the series expansion

$$m_A(z) = -\frac{1}{z} - \sum_{j=1}^{\infty} \frac{M_j^A}{z^{j+1}}, \quad (6.6)$$

about $z = \infty$, where $M_j^A := \int x^j dF^A(x)$ is the j -th moment. The ordinary moment generating function, $\mu_A(z)$, is the power series

$$\mu_A(z) = \sum_{j=0}^{\infty} M_j^A z^j, \quad (6.7)$$

with $M_0^A = 1$. The moment generating function, referred to as the moment transform, is related to the Stieltjes transform by

$$\mu_A(z) = -\frac{1}{z}m_A\left(\frac{1}{z}\right). \quad (6.8)$$

The Stieltjes transform can be expressed in terms of the moment transform as

$$m_A(z) = -\frac{1}{z}\mu_A\left(\frac{1}{z}\right). \quad (6.9)$$

The eta transform, introduced by Tulino and Verdù in [104], is a minor variation of the moment transform. It can be expressed in terms of the Stieltjes transform as

$$\eta_A(z) = \frac{1}{z}m_A\left(-\frac{1}{z}\right), \quad (6.10)$$

while the Stieltjes transform can be expressed in terms of the eta transform as

$$m_A(z) = -\frac{1}{z}\eta_A\left(-\frac{1}{z}\right). \quad (6.11)$$

■ 6.1.3 The R transform

The R transform is defined in terms of the Cauchy transform as

$$r_A(z) = g_A^{\langle -1 \rangle}(z) - \frac{1}{z}, \quad (6.12)$$

where $g_A^{\langle -1 \rangle}(z)$ is the functional inverse of $g_A(z)$ with respect to composition. It will often be more convenient to use the expression for the R transform in terms of the Cauchy transform given by

$$r_A(g) = z(g) - \frac{1}{g}. \quad (6.13)$$

The R transform can be written as a power series whose coefficients K_j^A are known as the “free cumulants.” For a combinatorial interpretation of free cumulants, see [92]. Thus the R transform is the (ordinary) free cumulant generating function

$$r_A(g) = \sum_{j=0}^{\infty} K_{j+1}^A g^j. \quad (6.14)$$

■ 6.1.4 The S transform

The S transform is relatively more complicated. It is defined as

$$s_A(z) = \frac{1+z}{z} \Upsilon_A^{(-1)}(z) \quad (6.15)$$

where $\Upsilon_A(z)$ can be written in terms of the Stieltjes transform $m_A(z)$ as

$$\Upsilon_A(z) = -\frac{1}{z} m_A(1/z) - 1. \quad (6.16)$$

This definition is quite cumbersome to work with because of the functional inverse in (6.15). It also places a technical restriction (to enable series inversion) that $M_1^A \neq 0$. We can, however, avoid this by expressing the S transform algebraically in terms of the Stieltjes transform as shown next. We first plug in $\Upsilon_A(z)$ into the left-hand side of (6.15) to obtain

$$s_A(\Upsilon_A(z)) = \frac{1 + \Upsilon_A(z)}{\Upsilon_A(z)} z.$$

This can be rewritten in terms of $m_A(z)$ using the relationship in (6.16) to obtain

$$s_A\left(-\frac{1}{z} m(1/z) - 1\right) = \frac{z m(1/z)}{m(1/z) + z}$$

or, equivalently:

$$s_A(-z m(z) - 1) = \frac{m(z)}{z m(z) + 1}. \quad (6.17)$$

We now define $y(z)$ in terms of the Stieltjes transform as $y(z) = -z m(z) - 1$. It is clear that $y(z)$ is an invertible function of $m(z)$. The right hand side of (6.17) can be rewritten in terms of $y(z)$ as

$$s_A(y(z)) = -\frac{m(z)}{y(z)} = \frac{m(z)}{z m(z) + 1}. \quad (6.18)$$

Equation (6.18) can be rewritten to obtain a simple relationship between the Stieltjes transform and the S transform

$$m_A(z) = -y s_A(y). \quad (6.19)$$

Noting that $y = -z m(z) - 1$ and $m(z) = -y s_A(y)$ we obtain the relationship

$$y = z y s_A(y) - 1$$

or, equivalently

$$z = \frac{y+1}{y s_A(y)}. \quad (6.20)$$

■ 6.2 The Algebraic Framework

Notation 6.21 (Bivariate polynomial). Let L_{uv} denote a bivariate polynomial of degree D_u in u and D_v in v defined as

$$L_{uv} \equiv L_{uv}(\cdot, \cdot) = \sum_{j=0}^{D_u} \sum_{k=0}^{D_v} c_{jk} u^j v^k = \sum_{j=0}^{D_u} l_j(v) u^j. \quad (6.21)$$

The scalar coefficients c_{jk} are real valued.

The two letter subscripts for the bivariate polynomial L_{uv} provide us with a convention of which dummy variables we will use. We will generically use the first letter in the subscript to represent a transform of the density with the second letter acting as a mnemonic for the dummy variable associated with the transform. By consistently using the same pair of letters to denote the bivariate polynomial that encodes the transform and the associated dummy variable, this abuse of notation allows us to readily identify the encoding of the distribution that is being manipulated.

Remark 6.22 (Irreducibility). Unless otherwise stated it will be understood that $L_{uv}(u, v)$ is “irreducible” in the sense that the conditions:

- $l_0(v), \dots, l_{D_u}(v)$ have no common factor involving v ,
- $l_{D_u}(v) \neq 0$,
- $\text{disc}_L(v) \neq 0$,

are satisfied, where $\text{disc}_L(v)$ is the discriminant of $L_{uv}(u, v)$ thought of as a polynomial in v .

We are particularly focused on the solution “curves,” $u_1(v), \dots, u_{D_u}(v)$, i.e.,

$$L_{uv}(u, v) = l_{D_u}(v) \prod_{i=1}^{D_u} (u - u_i(v)).$$

Informally speaking, when we refer to the bivariate polynomial equation $L_{uv}(u, v) = 0$ with solutions $u_i(v)$ we are actually considering the equivalence class of rational functions with this set of solution curves.

Remark 6.23 (Equivalence class). The equivalence class of $L_{uv}(u, v)$ may be characterized as functions of the form $L_{uv}(u, v)g(v)/h(u, v)$ where h is relatively prime to $L_{uv}(u, v)$ and $g(v)$ is not identically 0.

A few technicalities (such as poles and singular points) that will be catalogued later in Chapter 8 remain, but this is sufficient for allowing us to introduce rational transformations of the arguments and continue to use the language of polynomials.

Definition 6.24 (Algebraic distributions). Let $F(x)$ be a probability distribution function and $f(x)$ be its distributional derivative (here and henceforth). Consider the Stieltjes transform $m(z)$ of the distribution function, defined as

$$m(z) = \int \frac{1}{x-z} dF(x) \quad \text{for } z \in \mathbb{C}^+ \setminus \mathbb{R}. \quad (6.22)$$

If there exists a bivariate polynomial L_{mz} such that $L_{mz}(m(z), z) = 0$ then we refer to $F(x)$ as algebraic (probability) distribution function, $f(x)$ as an algebraic (probability) density function and say the $f \in \mathcal{P}_{alg}$. Here \mathcal{P}_{alg} denotes the class of algebraic (probability) distributions.

Definition 6.25 (Atomic distribution). Let $F(x)$ be a probability distribution function of the form

$$F(x) = \sum_{i=1}^K p_i \mathbb{I}_{[\lambda_i, \infty)},$$

where the K atoms at $\lambda_i \in \mathbb{R}$ have (non-negative) weights p_i subject to $\sum_i p_i = 1$ and $\mathbb{I}_{[x, \infty)}$ is the indicator (or characteristic) function of the set $[x, \infty)$. We refer to $F(x)$ as an atomic (probability) distribution function. Denoting its distributional derivative by $f(x)$, we say that $f(x) \in \mathcal{P}_{atom}$. Here \mathcal{P}_{atom} denotes the class of atomic distributions.

Example 6.26. An atomic probability distribution, as in Definition 6.25, has a Stieltjes transform

$$m(z) = \sum_{i=1}^K \frac{p_i}{\lambda_i - z}$$

which is the solution of the equation $L_{mz}(m, z) = 0$ where

$$L_{mz}(m, z) \equiv \prod_{i=1}^K (\lambda_i - z) m - \sum_{i=1}^K \prod_{\substack{j \neq i \\ j=1}}^K p_j (\lambda_j - z).$$

Hence it is an algebraic distribution; consequently $\mathcal{P}_{atom} \subset \mathcal{P}_{alg}$.

Example 6.27. The Cauchy distribution whose density

$$f(x) = \frac{1}{\pi(x^2 + 1)},$$

has a Stieltjes transform $m(z)$ which is the solution of the equation $L_{mz}(m, z) = 0$ where

$$L_{mz}(m, z) \equiv (z^2 + 1) m^2 + 2z m + 1.$$

Hence it is an algebraic distribution.

It is often the case that the probability density functions of algebraic distributions, according to our definition, will also be algebraic functions themselves. We conjecture that this is a necessary but not sufficient condition. We show that it is not sufficient by providing the counter-example below.

Counter-example 6.28. Consider the quarter-circle distribution with density function

$$f(x) = \frac{\sqrt{4-x^2}}{\pi} \quad \text{for } x \in [0, 2].$$

Its Stieltjes transform :

$$m(z) = -\frac{4 - 2\sqrt{-z^2+4} \ln\left(-\frac{2+\sqrt{-z^2+4}}{z}\right) + z\pi}{2\pi},$$

is clearly not an algebraic function. Thus $f(x) \notin \mathcal{P}_{alg}$.

We now define six interconnected bivariate polynomials denoted by L_{mz} , L_{gz} , L_{rg} , L_{sy} , $L_{\mu z}$, and $L_{\eta z}$. We assume that $L_{uv}(u, v)$ is an irreducible bivariate polynomial of the form in (6.21). The main protagonist of the transformations we consider is the bivariate polynomial L_{mz} which implicitly defines the Stieltjes transform $m(z)$ via the equation $L_{mz}(m, z) = 0$. Starting off with this polynomial we can obtain the polynomial L_{gz} using the relationship in (6.5) as

$$L_{gz}(g, z) = L_{mz}(-g, z). \quad (6.23)$$

Perhaps we should explain our abuse of notation once again, for the sake of clarity. Given any one polynomial, all the other polynomials can be obtained. The two letter subscripts not only tell us which of the six polynomials we are focusing on, it provides a convention of which dummy variables we will use. The first letter in the subscript represents the transform; the second letter is a mnemonic for the variable associated with the transform that we use consistently in the software based on this framework. With this notation in mind, we can obtain the polynomial L_{rg} from L_{gz} using (6.13) as

$$L_{rg}(r, g) = L_{gz}\left(g, r + \frac{1}{g}\right). \quad (6.24)$$

Similarly, we can obtain the bivariate polynomial L_{sy} from L_{mz} using the expressions in (6.19) and (6.20) to obtain the relationship

$$L_{sy} = L_{mz}\left(-ys, \frac{y+1}{sy}\right). \quad (6.25)$$




Figure 6-1. The six interconnected bivariate polynomials; transformations between the polynomials, indicated by the labelled arrows, are given in Table 6.3.

Based on the transforms discussed in Section 6.1, we can derive transformations between additional pairs of bivariate polynomials represented by the bidirectional arrows in Figure 6-1 and listed in the third column of Table 6.3. Specifically, the expressions in (6.8) and (6.11) can be used to derive the transformations between L_{mz} and $L_{\mu z}$ and L_{mz} and $L_{\eta z}$ respectively. The fourth column of Table 6.3 lists the MATLAB function, implemented using its MAPLE based Symbolic Toolbox, corresponding to the bivariate polynomial transformations represented in Figure 6-1. In the MATLAB functions, the function `irreducLuv(u,v)` listed in Table 6.2 ensures that the resulting bivariate polynomial is irreducible by clearing the denominator and making the resulting polynomial square free.

Example: Consider an atomic probability distribution with

$$F(x) = 0.5 \mathbb{I}_{[0,\infty)} + 0.5 \mathbb{I}_{[1,\infty)}, \quad (6.26)$$

whose Stieltjes transform

$$m(z) = \frac{0.5}{0-z} + \frac{0.5}{1-z},$$

is the solution of the equation

$$m(0-z)(1-z) - 0.5(1-2z) = 0,$$

or equivalently, the solution of the equation $L_{mz}(m, z) = 0$ where

$$L_{mz}(m, z) \equiv m(2z^2 - 2z) - (1 - 2z). \quad (6.27)$$

We can obtain the bivariate polynomial $L_{gz}(g, z)$ by applying the transformation in

Procedure	MATLAB Code
Simplify and clear the denominator	<code>function Luv = irreducLuv(Luv,u,v)</code>
Make square free	<code>L = numden(simplify(expand(Luv)));</code> <code>L = Luv / maple('gcd',L,diff(L,u));</code>
Simplify	<code>L = simplify(expand(L));</code> <code>L = Luv / maple('gcd',L,diff(L,v));</code> <code>Luv = simplify(expand(L));</code>

Table 6.1. Making L_{uv} irreducible.

(6.23) to the bivariate polynomial L_{mz} given by (6.27) so that

$$L_{gz}(g, z) = -g(2z^2 - 2z) - (1 - 2z). \quad (6.28)$$

Similarly, by applying the transformation in (6.24) we obtain

$$L_{rg}(r, g) = -g \left(2 \left(r + \frac{1}{g} \right) - 2 \left(r + \frac{1}{g} \right)^2 \right) - \left(1 - 2 \left(r + \frac{1}{g} \right) \right). \quad (6.29)$$

which, on clearing the denominator and invoking the equivalence class representation of our polynomials (see Remark 6.23), gives us the irreducible bivariate polynomial

$$L_{rg}(r, g) = -1 + 2gr^2 + (2 - 2g)r. \quad (6.30)$$

By applying the transformation in (6.25) to the bivariate polynomial L_{mz} , we obtain

$$L_{sy} \equiv (-sy) \left(2 \frac{y+1}{sy} - 2 \left(\frac{y+1}{sy} \right)^2 \right) - \left(1 - 2 \frac{y+1}{sy} \right)$$

which on clearing the denominator gives us the irreducible bivariate polynomial

$$L_{sy}^A(s, y) = (1 + 2y)s - 2 - 2y. \quad (6.31)$$

Table 6.2 tabulates the six bivariate polynomial encodings in Figure 6-1 for the distribution in (6.26), the semi-circle distribution for Wigner matrices and the Marčenko-Pastur distribution for Wishart matrices.

L	Bivariate Polynomials
L_{mz}	$m(2z^2 - 2z) - (1 - 2z)$
L_{gz}	$-g(2z^2 - 2z) - (1 - 2z)$
L_{rg}	$-1 + 2gr^2 + (2 - 2g)r$
L_{sy}	$(1 + 2y)s - 2 - 2y$
$L_{\mu z}$	$(-2 + 2z)\mu + 2 - z$
$L_{\eta z}$	$(2z + 2)\eta - 2 - z$

(a) The atomic distribution in (6.26).

L	Bivariate Polynomials
L_{mz}	$czm^2 - (1 - c - z)m + 1$
L_{gz}	$czg^2 + (1 - c - z)g + 1$
L_{rg}	$(cg - 1)r + 1$
L_{sy}	$(cy + 1)s - 1$
$L_{\mu z}$	$\mu^2zc - (zc + 1 - z)\mu + 1$
$L_{\eta z}$	$\eta^2zc + (-zc + 1 - z)\eta - 1$

(b) The Marčenko-Pastur distribution.

L	Bivariate polynomials
L_{mz}	$m^2 + mz + 1$
L_{gz}	$g^2 - gz + 1$
L_{rg}	$r - g$
L_{sy}	$s^2y - 1$
$L_{\mu z}$	$\mu^2z^2 - \mu + 1$
$L_{\eta z}$	$z^2\eta^2 - \eta + 1$

(c) The semi-circle distribution.

Table 6.2. Bivariate polynomial representations of some algebraic distributions.

<i>Label</i>	<i>Conversion</i>	<i>Transformation</i>	<i>MATLAB Code</i>
I	$L_{mz} \Leftrightarrow L_{gz}$	$L_{mz} = L_{gz}(-m, z)$	function l.mz = l.gz2l.mz(l.gz) syms m g z l.mz = subs(l.gz,g,-m);
		$L_{gz} = L_{mz}(-g, z)$	function l.gz = l.mz2l.gz(l.mz) syms m g z l.gz = subs(l.mz,m,-g);
II	$L_{gz} \Leftrightarrow L_{rg}$	$L_{gz} = L_{rg}(z - \frac{1}{g}, z)$	function l.gz = l.rg2l.gz(l.rg) syms r g z l.gz = subs(l.rg,r,z-1/g); l.gz = irreducl.uv(l.gz,g,z);
		$L_{rg} = L_{gz}(g, r + \frac{1}{g})$	function l.rg = l.gz2l.rg(l.gz) syms r g z l.rg = subs(l.gz,g,r+1/g); l.rg = irreducl.uv(l.rg,r,g);
III	$L_{mz} \Leftrightarrow L_{rg}$	$L_{mz} \Leftrightarrow L_{gz} \Leftrightarrow L_{rg}$	function l.mz = l.rg2l.mz(l.rg) syms m z r g l.gz = l.rg2l.gz(l.rg); l.mz = l.gz2l.mz(l.gz); ----- function l.rg = l.mz2l.rg(l.mz) syms m z r g l.gz = l.mz2l.gz(l.mz); l.rg = l.gz2l.rg(l.gz);
IV	$L_{mz} \Leftrightarrow L_{sy}$	$L_{mz} = L_{sy}(\frac{m}{z m + 1}, -z m - 1)$	function l.mz = l.sy2l.mz(l.sy) syms m z s y l.mz = subs(l.sy,s,m/(z*m+1)); l.mz = subs(l.mz,y,-z*m-1); l.mz = irreducl.uv(l.mz,m,z);
		$L_{sy} = L_{mz}(-y s, \frac{y + 1}{s y})$	function l.sy = l.mz2l.sy(l.mz) syms m z s y l.sy = subs(l.mz,m,-y*s); l.sy = subs(l.sy,z,(y+1)/y/s); l.sy = irreducl.uv(l.sy,s,y);
V	$L_{mz} \Leftrightarrow L_{\mu z}$	$L_{mz} = L_{\mu z}(-m z, \frac{1}{z})$	function l.mz = l.muzy2l.mz(l.muzy) syms m myu z l.mz = subs(l.muzy,z,1/z); l.mz = subs(l.mz,myu,-m*z); l.mz = irreducl.uv(l.mz,m,z);
		$L_{\mu z} = L_{mz}(-\mu z, \frac{1}{z})$	function l.muzy = l.mz2l.muzy(l.mz) syms m myu z l.muzy = subs(l.mz,z,1/z); l.muzy = subs(l.muzy,m,-myu*z); l.muzy = irreducl.uv(l.muzy,myu,z);
VI	$L_{mz} \Leftrightarrow L_{\eta z}$	$L_{mz} = L_{\eta z}(-z m, -\frac{1}{z})$	function l.mz = l.etaz2l.mz(l.etaz) syms m eta z l.mz = subs(l.etaz,z,-1/z); l.mz = subs(l.mz,eta,-z*m); l.mz = irreducl.uv(l.mz,m,z);
		$L_{\eta z} = L_{mz}(z \eta, -\frac{1}{z})$	function l.etaz = l.mz2l.etaz(l.mz) syms m eta z l.etaz = subs(l.mz,z,-1/z); l.etaz = subs(l.etaz,m,z*eta); l.etaz = irreducl.uv(l.etaz,eta,z);

Table 6.3. Transformations between the different bivariate polynomials. As a guide to MATLAB notation, the command `syms` declares a variable to be symbolic while the command `subs` symbolically substitutes every occurrence of the second argument in the first argument with the third argument. Thus, for example, the command `y=subs(x-a,a,10)` will yield the output `y=x-10` if we have previously declared `x` and `a` to be symbolic using the command `syms x a`.

■ 6.3 Algebraic manipulations of algebraic functions

Algebraic functions are closed under addition and multiplication. Hence we can add (or multiply) two algebraic functions and obtain another algebraic function. We show, using purely matrix theoretic arguments, how to obtain the polynomial equation whose solution is the sum (or product) of two algebraic functions without ever actually computing the individual functions. In Section 6.4, we interpret this computation using the concept of resultants [98] from elimination theory. These tools will feature prominently in Chapter 7 when we encode the transformations of the random matrices as algebraic operations on the appropriate form of the bivariate polynomial that encodes their limiting eigenvalue distributions.

Definition 6.31 (Companion Matrix). *The companion matrix $\mathbf{C}_{a(x)}$ to a monic polynomial*

$$a(x) \equiv a_0 + a_1 x + \dots + a_{n-1} x^{n-1} + x^n$$

is the $n \times n$ square matrix

$$\mathbf{C}_{a(x)} = \begin{bmatrix} 0 & \dots & \dots & \dots & -a_0 \\ 1 & \dots & \dots & \dots & -a_1 \\ 0 & \ddots & & & -a_2 \\ \vdots & & \ddots & & \vdots \\ 0 & \dots & \dots & 1 & -a_{n-1} \end{bmatrix}$$

with ones on the sub-diagonal and the last column given by the negative coefficients of $a(x)$.

Remark 6.32. *The eigenvalues of the companion matrix are the solutions of the equation $a(x) = 0$. This is intimately related to the observation that the characteristic polynomial of the companion matrix equals $a(x)$, i.e.,*

$$a(x) = \det(x \mathbf{I}_n - \mathbf{C}_{a(x)}).$$

Consider the bivariate polynomial L_{uv} as in (6.21). By treating it as a polynomial in u whose coefficients are polynomials in v , i.e., by rewriting it as

$$L_{uv}(u, v) \equiv \sum_{j=0}^{D_u} l_j(v) u^j, \quad (6.32)$$

we can create a companion matrix \mathbf{C}_{uv}^u whose characteristic polynomial as a function of u is the bivariate polynomial L_{uv} . The companion matrix \mathbf{C}_{uv}^u is the $D_u \times D_u$ matrix in Table 6.4.

Remark 6.33. *Analogous to the univariate case, the characteristic polynomial of \mathbf{C}_{uv}^u*

\mathbf{C}_{uv}^u	MATLAB code
$\begin{bmatrix} 0 & \dots & \dots & \dots & -l_0(v)/l_{D_u}(v) \\ 1 & \dots & \dots & \dots & -l_1(v)/l_{D_u}(v) \\ 0 & \ddots & & & -l_2(v)/l_{D_u}(v) \\ \vdots & & \ddots & & \vdots \\ 0 & \dots & \dots & 1 & -l_{D_u-1}(v)/l_{D_u}(v) \end{bmatrix}$	<pre>function Cu = Luv2Cu(Luv,u) Du = double(maple('degree',Luv,u)); LDu = maple('coeff',Luv,u,Du); Cu = sym(zeros(Du))+ .. +diag(ones(Du-1,1),-1)); for Di = 0:Du-1 LtuDi = maple('coeff',Lt,u,Di); Cu(Di+1,Du) = -LtuDi/LDu; end</pre>

Table 6.4. The companion matrix \mathbf{C}_{uv}^u , with respect to u , of the bivariate polynomial L_{uv} given by (6.32).

is $\det(u\mathbf{I} - \mathbf{C}_{uv}^u) = L_{uv}(u, v)/l_{D_u}(v)^{D_u}$. Since $l_{D_u}(v)$ is not identically zero, we say that $\det(u\mathbf{I} - \mathbf{C}_{uv}^u) = L_{uv}(u, v)$ where the equality is understood to be with respect to the equivalence class of L_{uv} as in Remark 6.23. The eigenvalues of \mathbf{C}_{uv}^u are the solutions of the algebraic equation $L_{uv}(u, v) = 0$; specifically, we obtain the algebraic function $u(v)$.

Definition 6.34 (Kronecker product). If \mathbf{A}_m (with entries a_{ij}) is an $m \times m$ matrix and \mathbf{B}_n is an $n \times n$ matrix then the Kronecker (or tensor) product of \mathbf{A}_m and \mathbf{B}_n , denoted by $\mathbf{A}_m \otimes \mathbf{B}_n$, is the $mn \times mn$ matrix defined as:

$$\mathbf{A}_m \otimes \mathbf{B}_n = \begin{bmatrix} a_{11}\mathbf{B}_n & \dots & a_{1n}\mathbf{B}_n \\ \vdots & \ddots & \vdots \\ a_{m1}\mathbf{B}_n & \dots & a_{mn}\mathbf{B}_n \end{bmatrix}$$

Lemma 6.35. If α_i and β_j are the eigenvalues of \mathbf{A}_m and \mathbf{B}_n respectively, then

1. $\alpha_i + \beta_j$ is an eigenvalue of $(\mathbf{A}_m \otimes \mathbf{I}_n) + (\mathbf{I}_m \otimes \mathbf{B}_n)$,
2. $\alpha_i \beta_j$ is an eigenvalue of $\mathbf{A}_m \otimes \mathbf{B}_n$,

for $i = 1, \dots, m$, $j = 1, \dots, n$.

PROOF. This is a standard result in linear algebra that may be found in several standard texts including [48]. \square

Proposition 6.36. Let $u_1(v)$ be a solution of the algebraic equation $L_{uv}^1(u, v) = 0$, or equivalently an eigenvalue of the $D_u^1 \times D_u^1$ companion matrix $\mathbf{C}_{uv}^{u_1}$. Let $u_2(v)$ be a solution of the algebraic equation $L_{uv}^2(u, v) = 0$, or equivalently an eigenvalue of the $D_u^2 \times D_u^2$ companion matrix $\mathbf{C}_{uv}^{u_2}$. Then

1. $u_3(v) = u_1(v) + u_2(v)$ is an eigenvalue of the matrix $\mathbf{C}_{uv}^{u_3} = (\mathbf{C}_{uv}^{u_1} \otimes \mathbf{I}_{D_u^2}) + (\mathbf{I}_{D_u^1} \otimes \mathbf{C}_{uv}^{u_2})$,
2. $u_3(v) = u_1(v)u_2(v)$ is an eigenvalue of the matrix $\mathbf{C}_{uv}^{u_3} = \mathbf{C}_{uv}^{u_1} \otimes \mathbf{C}_{uv}^{u_2}$.

Equivalently $u_3(v)$ is a solution of the algebraic equation $L_{uv}^3 = 0$ where $L_{uv}^3 = \det(u\mathbf{I} - \mathbf{C}_{uv}^{u_3})$.

PROOF. This follows directly from Lemma 6.35. \square

We represent the binary addition and multiplication operators on the space of algebraic functions by the symbols \boxplus_u and \boxtimes_u respectively. We define addition and multiplication as in Table 6.5 by applying Proposition 6.36. Note that the subscript ‘u’ in \boxplus_u and \boxtimes_u provides us with an indispensable convention of which dummy variable we are using. Table 6.6 illustrates the \boxplus and \boxtimes operations on a pair of bivariate polynomials and underscores the importance of the symbolic software developed. The $(D_u + 1) \times (D_v + 1)$ matrix \mathbf{T}_{uv} lists only the coefficients c_{ij} for the term $u^i v^j$ in the polynomial $L_{uv}(u, v)$. Note that the indexing for i and j starts with zero.

Operation: $L_{uv}^1, L_{uv}^2 \mapsto L_{uv}^3$	MATLAB Code
$L_{uv}^3 = L_{uv}^1 \boxplus_u L_{uv}^2 \equiv \det(u\mathbf{I} - \mathbf{C}_{uv}^{u_3}), \text{ where}$ $\mathbf{C}_{uv}^{u_3} = \begin{cases} 2 \mathbf{C}_{uv}^{u_1} & \text{if } L_{uv}^1 = L_{uv}^2, \\ (\mathbf{C}_{uv}^{u_1} \otimes \mathbf{I}_{D_u^2}) + (\mathbf{I}_{D_u^1} \otimes \mathbf{C}_{uv}^{u_2}) & \text{otherwise.} \end{cases}$	<pre>function Luv3 = L1plusL2(Luv1,Luv2,u) Cu1 = Luv2Cu(Luv1,u); if (Luv1 == Luv2) Cu3 = 2*Cu1; else Cu2 = Luv2Cu(Luv2,u); Cu3 = kron(Cu1,eye(length(Cu2))) + .. +kron(eye(length(Cu1)),Cu2); end Luv3 = det(u*eye(length(Cu3))-Cu3);</pre>
$L_{uv}^3 = L_{uv}^1 \boxtimes_u L_{uv}^2 \equiv \det(u\mathbf{I} - \mathbf{C}_{uv}^{u_3}), \text{ where}$ $\mathbf{C}_{uv}^{u_3} = \begin{cases} \mathbf{C}_{uv}^{u_3} = (\mathbf{C}_{uv}^{u_1})^2 & \text{if } L_{uv}^1 = L_{uv}^2, \\ \mathbf{C}_{uv}^{u_3} = \mathbf{C}_{uv}^{u_1} \otimes \mathbf{C}_{uv}^{u_2} & \text{otherwise.} \end{cases}$	<pre>function Luv3 = L1timesL2(Luv1,Luv2,u) Cu1 = Luv2Cu(Luv1,u); if (Luv1 == Luv2) Cu3 = Cu1^2; else Cu2 = Luv2Cu(Luv2,u); Cu3 = kron(Cu1,Cu2); end Luv3 = det(u*eye(length(Cu3))-Cu3);</pre>

Table 6.5. Formal and computational description of the \boxplus_u and \boxtimes_u operators acting on the bivariate polynomials $L_{uv}^1(u, v)$ and $L_{uv}^2(u, v)$ where $\mathbf{C}_{uv}^{u_1}$ and $\mathbf{C}_{uv}^{u_2}$ are their corresponding companion matrices constructed as in Table 6.4 and \otimes is the matrix Kronecker product.

■ 6.4 Algebraic manipulations using the resultant

Addition (and multiplication) of algebraic functions produces another algebraic function. We now demonstrate how the concept of resultants from elimination theory can be used to obtain the polynomial whose zero set is the required algebraic function.

Definition 6.41 (Resultant). *Given a polynomial*

$$a(x) \equiv a_0 + a_1 x + \dots + a_{n-1} x^{n-1} + a_n x^n$$

L_{uv}	T_{uv}	C_{uv}^u	C_{uv}^v
$L_{uv}^1 \equiv u^2v + u(1-v) + v^2$	$1 \begin{bmatrix} 1 & v & v^2 \\ \cdot & \cdot & 1 \end{bmatrix}$ $u \begin{bmatrix} 1 & -1 & \cdot \\ \cdot & \cdot & \cdot \end{bmatrix}$ $u^2 \begin{bmatrix} \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot \end{bmatrix}$	$\begin{bmatrix} 0 & -v \\ 1 & \frac{-1+v}{v} \end{bmatrix}$	$\begin{bmatrix} 0 & -u \\ 1 & -u^2 + u \end{bmatrix}$
$L_{uv}^2 \equiv u^2(v^2 - 3v + 1) + u(1+v) + v^2$	$1 \begin{bmatrix} 1 & v & v^2 \\ \cdot & \cdot & 1 \end{bmatrix}$ $u \begin{bmatrix} 1 & 1 & \cdot \\ \cdot & \cdot & \cdot \end{bmatrix}$ $u^2 \begin{bmatrix} 1 & -3 & 1 \\ \cdot & \cdot & \cdot \end{bmatrix}$	$\begin{bmatrix} 0 & \frac{-v^2}{v^2 - 3v + 1} \\ 1 & \frac{-1 - v}{v^2 - 3v + 1} \end{bmatrix}$	$\begin{bmatrix} 0 & \frac{-u^2 - u}{u^2 + 1} \\ 1 & \frac{3u^2 - u}{u^2 + 1} \end{bmatrix}$

$L_{uv}^1 \boxplus_u L_{uv}^2$	$1 \begin{bmatrix} 1 & v & v^2 & v^3 & v^4 & v^5 & v^6 & v^7 & v^8 \\ \cdot & \cdot & 2 & -6 & 11 & -10 & 18 & -8 & 1 \end{bmatrix}$ $u \begin{bmatrix} 2 & \cdot & 2 & -8 & 4 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^2 \begin{bmatrix} 5 & \cdot & 1 & -4 & 2 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^3 \begin{bmatrix} 4 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^4 \begin{bmatrix} 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$
$L_{uv}^1 \boxtimes_u L_{uv}^2$	$1 \begin{bmatrix} 1 & v & v^2 & v^3 & v^4 & v^5 & v^6 & v^7 & v^8 & v^9 & v^{10} & v^{11} & v^{12} & v^{13} & v^{14} \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & -6 & 11 & -6 & 1 \end{bmatrix}$ $u \begin{bmatrix} \cdot & \cdot & \cdot & \cdot & -1 & 3 & \cdot & -3 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^2 \begin{bmatrix} \cdot & \cdot & 1 & -4 & 10 & -6 & 7 & -2 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^3 \begin{bmatrix} -1 & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^4 \begin{bmatrix} 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$

$L_{uv}^1 \boxplus_v L_{uv}^2$	$L_{uv}^2 \boxtimes_v L_{uv}^2$
$1 \begin{bmatrix} 1 & v & v^2 & v^3 & v^4 \\ \cdot & \cdot & \cdot & \cdot & 1 \end{bmatrix}$ $u \begin{bmatrix} \cdot & \cdot & 4 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^2 \begin{bmatrix} \cdot & \cdot & 1 & -4 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^3 \begin{bmatrix} \cdot & -8 & 6 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^4 \begin{bmatrix} 1 & -2 & 3 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^5 \begin{bmatrix} 8 & -12 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^6 \begin{bmatrix} 3 & 2 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^7 \begin{bmatrix} 2 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^8 \begin{bmatrix} -1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$	$1 \begin{bmatrix} 1 & v & v^2 & v^3 & v^4 \\ \cdot & \cdot & \cdot & \cdot & 1 \end{bmatrix}$ $u \begin{bmatrix} \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^2 \begin{bmatrix} \cdot & \cdot & -2 & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^3 \begin{bmatrix} \cdot & \cdot & \cdot & -4 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^4 \begin{bmatrix} 1 & 1 & -9 & 3 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^5 \begin{bmatrix} 2 & -3 & 7 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^6 \begin{bmatrix} 3 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^7 \begin{bmatrix} 4 & \cdot & -1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^8 \begin{bmatrix} 3 & -1 & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^9 \begin{bmatrix} 2 & 3 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$ $u^{10} \begin{bmatrix} 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{bmatrix}$

Table 6.6. Examples of \boxplus and \boxtimes operations on a pair of bivariate polynomials, L_{uv}^1 and L_{uv}^2 .

of degree n with roots α_i , for $i = 1, \dots, n$ and a polynomial

$$b(x) \equiv b_0 + b_1 x + \dots + b_{m-1} x^{m-1} + b_m x^m$$

of degree m with roots β_j , for $j = 1, \dots, m$, the resultant is defined as

$$\text{Res}_x(a(x), b(x)) = a_n^m b_m^n \prod_{i=1}^n \prod_{j=1}^m (\beta_j - \alpha_i).$$

From a computational standpoint, the resultant can be directly computed from the coefficients of the polynomials itself. The computation involves the formation of the Sylvester matrix and exploiting an identity that relates the determinant of the Sylvester matrix to the resultant.

Definition 6.42 (Sylvester matrix). Given polynomials $a(x)$ and $b(x)$ with degree n and m respectively and coefficients as in Definition 6.41, the Sylvester matrix is the $(n+m) \times (n+m)$ matrix

$$\mathbf{S}(a, b) = \begin{bmatrix} a_n & 0 & \cdots & 0 & 0 & b_m & 0 & \cdots & 0 & 0 \\ a_{n-1} & a_n & \cdots & 0 & 0 & b_{m-1} & b_m & \cdots & 0 & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & a_0 & a_1 & 0 & 0 & \cdots & b_0 & b_1 \\ 0 & 0 & \cdots & 0 & a_0 & 0 & 0 & \cdots & 0 & b_0 \end{bmatrix}$$

Proposition 6.43. The resultant of two polynomials $a(x)$ and $b(x)$ is related to the determinant of the Sylvester matrix by

$$\det(\mathbf{S}(a, b)) = \text{Res}_x(a(x), b(x))$$

PROOF. This identity can be proved using standard linear algebra arguments. A proof may be found in [4]. \square

For our purpose, the utility of this definition is that the \boxplus_u and \boxtimes_u operations can be expressed in terms of resultants. Suppose we are given two bivariate polynomials L_{uv}^1 and L_{uv}^2 . By using the definition of the resultant and treating the bivariate polynomials as polynomials in u whose coefficients are polynomials in v , we obtain the identities

$$\boxed{L_{uv}^3(t, v) = L_{uv}^1 \boxplus_u L_{uv}^2 \equiv \text{Res}_u(L_{uv}^1(t-u, v), L_{uv}^2(u, v)),} \quad (6.33)$$

and

$$\boxed{L_{uv}^3(t, v) = L_{uv}^1 \boxtimes_u L_{uv}^2 \equiv \text{Res}_u(u^{D_u^1} L_{uv}^1(t/u, v), L_{uv}^2(u, v)),} \quad (6.34)$$

where D_u^1 is the degree of L_{uv}^1 with respect to u . By Proposition 6.43, evaluating the \boxplus_u

and \boxtimes_u operations via the resultant formulation involves computing the determinant of the $(D_u^1 + D_u^2) \times (D_u^1 + D_u^2)$ Sylvester matrix. When $L_{uv}^1 \neq L_{uv}^2$, this results in a steep computational saving relative to the companion matrix based formulation in Table 6.5 which involves computing the determinant of a $(D_u^1 D_u^2) \times (D_u^1 D_u^2)$ matrix. Fast algorithms for computing the resultant exploit this and other properties of the Sylvester matrix formulation. In MAPLE, the computation $L_{uv}^3 = L_{uv}^1 \boxplus_u L_{uv}^2$ may be performed using the command:

```
Luv3 = subs(t=u,resultant(subs(u=t-u,Luv1),Luv2,u));
```

The computation $L_{uv}^3 = L_{uv}^1 \boxtimes_u L_{uv}^2$ can be performed via the sequence of commands:

```
Du1 = degree(Luv1,u);
Luv3 = subs(t=u,resultant(simplify(u^Du1*subs(u=t/u,Luv1)),Luv2,u));
```

When $L_{uv}^1 = L_{uv}^2$, however, the \boxplus_u and \boxtimes_u operations are best performed using the companion matrix formulation in Table 6.5. The software implementation of the operations in Table 6.5 in [73] uses the companion matrix formulation when $L_{uv}^1 = L_{uv}^2$ and the resultant formulation otherwise.

In this chapter we established our ability to encode algebraic distribution as solutions of bivariate polynomial equations and to manipulate the solutions. This sets the stage for defining the class of “algebraic” random matrices next.

The polynomial method: Algebraic random matrices

■ 7.1 Motivation

A random matrix is a matrix whose elements are random variables. Let \mathbf{A}_N be an $N \times N$ symmetric/Hermitian random matrix. Its empirical distribution function (e.d.f.) is given by

$$F^{\mathbf{A}_N}(x) = \frac{1}{N} \sum_{i=1}^N \mathbb{I}_{[\lambda_i, \infty)}. \quad (7.1)$$

where $\lambda_1, \dots, \lambda_N$ are the eigenvalues of \mathbf{A}_N (counted with multiplicity) and $\mathbb{I}_{[\lambda_i, \infty)} = 1$ when $x \geq \lambda_i$ and zero otherwise. For a large class of random matrices, the empirical distribution function $F^{\mathbf{A}_N}(x)$ converges, for every x , almost surely (or in probability) as $N \rightarrow \infty$ to a non-random distribution function $F^A(x)$. The associated eigenvalue density function, denoted by $f_A(x)$, is its distributional derivative.

We are interested in identifying canonical random matrix operations for which the limiting eigenvalue distribution of the resulting matrix is an algebraic distribution. This is equivalent to identifying operations for which the transformations in the random matrices can be mapped into transformations of the bivariate polynomial that encodes the limiting eigenvalue distribution function. This motivates the construction of the class of “algebraic” random matrices which we shall define next.

The practical utility of this definition, which will become apparent in Chapters 8 and 9 can be succinctly summarized: if a random matrix is shown to be algebraic then its limiting eigenvalue density function can be computed using a simple root-finding algorithm. Furthermore, if the moments exist, they will satisfy a finite depth linear recursion (see Theorem 8.36) with polynomial coefficients so that we will often be able to enumerate them efficiently in closed form. Algebraicity of a random matrix thus acts as a certificate of the computability of its limiting eigenvalue density function and the associated moments. In this chapter our objective is to specify the class of algebraic random matrices by its generators.

■ 7.2 Definitions

Let \mathbf{A}_N , for $N = 1, 2, \dots$ be a sequence of $N \times N$ random matrices with real eigenvalues. Let $F^{\mathbf{A}_N}$ denote the e.d.f., as in (7.1). Suppose $F^{\mathbf{A}_N}(x)$ converges almost surely (or in probability), for every x , to $F^A(x)$ as $N \rightarrow \infty$, then we say that $\mathbf{A}_N \mapsto A$. We denote the associated (non-random) limiting probability density function by $f_A(x)$.

Notation 7.21 (Mode of convergence of the empirical distribution function).

When necessary we highlight the mode of convergence of the underlying distribution function thus: if $\mathbf{A}_N \xrightarrow{\text{a.s.}} A$ then it is shorthand for the statement that the empirical distribution function of \mathbf{A}_N converges almost surely to the distribution function F^A ; likewise $\mathbf{A}_N \xrightarrow{P} A$ is shorthand for the statement that the empirical distribution function of \mathbf{A}_N converges in probability to the distribution function F^A . When the distinction is not made then almost sure convergence is assumed.

Remark 7.22. The element A above is not to be interpreted as a matrix. There is no convergence in the sense of an $\infty \times \infty$ matrix. The notation $\mathbf{A}_N \xrightarrow{\text{a.s.}} A$ is shorthand for describing the convergence of the associated distribution functions and not of the matrix itself. We think of A as being an (abstract) element of a probability space with distribution function F^A and associated density function f_A .

Definition 7.23 (Atomic random matrix). If $f_A \in \mathcal{P}_{\text{atom}}$ then we say that \mathbf{A}_N is an atomic random matrix. We represent this as $\mathbf{A}_N \mapsto A \in \mathcal{M}_{\text{atom}}$ where $\mathcal{M}_{\text{atom}}$ denotes the class of atomic random matrices.

Definition 7.24 (Algebraic random matrix). If $f_A \in \mathcal{P}_{\text{alg}}$ then we say that \mathbf{A}_N is an algebraically characterizable random matrix (often suppressing the word characterizable for brevity). We represent this as $\mathbf{A}_N \mapsto A \in \mathcal{M}_{\text{alg}}$ where \mathcal{M}_{alg} denotes the class of algebraic random matrices. Note that, by definition, $\mathcal{M}_{\text{atom}} \subset \mathcal{M}_{\text{alg}}$.

The ability to describe the class of algebraic random matrices and the technique needed to compute the associated bivariate polynomial is at the crux our investigation. In the theorems that follow, we accomplish the former by cataloguing random matrix operations that preserve algebraicity of the limiting distribution. The following property of the convergence of distributions will prove useful.

Proposition 7.25 (Continuous mapping theorem). Let $\mathbf{A}_N \mapsto A$. Let f_A and S_A^δ denote the corresponding limiting density function and the atomic component of the support, respectively. Consider the mapping $y = h(x)$ continuous everywhere on the real line except on the set of its discontinuities denoted by \mathcal{D}_h . If $D_h \cap S_A^\delta = \emptyset$ then $\mathbf{B}_N = h(\mathbf{A}_N) \mapsto B$. The associated non-random distribution function, F^B is given by $F^B(y) = F^A(h^{(-1)}(y))$. The associated probability density function is its distributional derivative.

PROOF. This is a restatement of continuous mapping theorem which follows from well-known facts about the convergence of distributions [17]. \square

■ 7.3 Deterministic operations on algebraic random matrices

We first consider some simple deterministic transformations on an algebraic random matrix \mathbf{A}_N that produce an algebraic random matrix \mathbf{B}_N .

Theorem 7.31. *Let $\mathbf{A}_N \mapsto A \in \mathcal{M}_{alg}$ and $p, q, r,$ and s be real-valued scalars. Then,*

$$\mathbf{B}_N = (p \mathbf{A}_N + q \mathbf{I}_N)/(r \mathbf{A}_N + s \mathbf{I}_N) \mapsto B \in \mathcal{M}_{alg},$$

provided f_A does not contain an atom at $-s/r$ and r, s are not zero simultaneously.

PROOF. Here we have $h(x) = (px + r)/(qx + s)$ which is continuous everywhere except at $x = -s/r$ for s and r not simultaneously zero. From Proposition 7.25, unless $f_A(x)$ has an atomic component at $-s/r$, $\mathbf{B}_N \mapsto B$. The Stieltjes transform of F^B can be expressed as

$$m_B(z) = E_y \left[\frac{1}{y - z} \right] = E_x \left[\frac{rx + s}{px + q - z(rx + s)} \right]. \quad (7.2)$$

Equation (7.2) can be rewritten as

$$m_B(z) = \int \frac{rx + s}{(p - rz)x + (q - sz)} dF^A(x) = \frac{1}{p - rz} \int \frac{rx + s}{x + \frac{q - sz}{p - rz}} dF^A(x). \quad (7.3)$$

With some algebraic manipulations, we can rewrite (7.3) as

$$\begin{aligned} m_B(z) &= \beta_z \int \frac{rx + s}{x + \alpha_z} dF^A(x) = \beta_z \left(r \int \frac{x}{x + \alpha_z} dF^A(x) + s \int \frac{1}{x + \alpha_z} dF^A(x) \right) \\ &= \beta_z \left(r \int dF^A(x) - r \alpha_z \int \frac{1}{x + \alpha_z} dF^A(x) + s \int \frac{1}{x + \alpha_z} dF^A(x) \right). \end{aligned} \quad (7.4)$$

where $\beta_z = 1/(p - rz)$ and $\alpha_z = (q - sz)/(p - rz)$. Using the definition of the Stieltjes transform and the identity $\int dF^A(x) = 1$, we can express $m_B(z)$ in (7.4) in terms of $m_A(z)$ as

$$m_B(z) = \beta_z r + (\beta_z s - \beta_z r \alpha_z) m_A(-\alpha_z). \quad (7.5)$$

Equation (7.5) can, equivalently, be rewritten as

$$m_A(-\alpha_z) = \frac{m_B(z) - \beta_z r}{\beta_z s - \beta_z r \alpha_z}. \quad (7.6)$$

Equation (7.6) can be expressed as an operational law on L_{mz}^A as

$$\boxed{L_{mz}^B(m, z) = L_{mz}^A((m - \beta_z r)/(\beta_z s - \beta_z r \alpha_z), -\alpha_z)}. \quad (7.7)$$

Since L_{mz}^A exists, we can obtain L_{mz}^B by applying the transformation in (7.7), and clearing the denominator to obtain the irreducible bivariate polynomial consistent with Remark 6.23. Since L_{mz}^B exists, this proves that $f_B \in \mathcal{P}_{alg}$ and $\mathbf{B}_N \mapsto B \in \mathcal{M}_{alg}$. \square

Appropriate substitutions for the scalars p, q, r and s in Theorem 7.31 leads to the following Corollary.

Corollary 7.32. *Let $\mathbf{A}_N \mapsto A \in \mathcal{M}_{alg}$ and let α be a real-valued scalar. Then,*

1. $\mathbf{B}_N = \mathbf{A}_N^{-1} \mapsto B \in \mathcal{M}_{alg}$, provided f_A does not contain an atom at 0,
2. $\mathbf{B}_N = \alpha \mathbf{A}_N \mapsto B \in \mathcal{M}_{alg}$,
3. $\mathbf{B}_N = \mathbf{A}_N + \alpha \mathbf{I}_N \mapsto B \in \mathcal{M}_{alg}$.

Theorem 7.33. *Let $\mathbf{X}_{n,N}$ be an $n \times N$ matrix. If $\mathbf{A}_N = \mathbf{X}_{n,N} \mathbf{X}'_{n,N} \mapsto A \in \mathcal{M}_{alg}$ then*

$$\mathbf{B}_N = \mathbf{X}'_{n,N} \mathbf{X}_{n,N} \mapsto B \in \mathcal{M}_{alg}.$$

PROOF. Here $\mathbf{X}_{n,N}$ is an $n \times N$ matrix, so that \mathbf{A}_n and \mathbf{B}_N are $n \times n$ and $N \times N$ sized matrices respectively. Let $c_N = n/N$. When $c_N < 1$, \mathbf{B}_N will have $N - n$ eigenvalues of magnitude zero while the remaining n eigenvalues will be identically equal to the eigenvalues of \mathbf{A}_n . Thus, the e.d.f. of \mathbf{B}_N is related to the e.d.f. of \mathbf{A}_n as

$$\begin{aligned} F^{\mathbf{B}_N}(x) &= \frac{N-n}{N} \mathbb{I}_{[0,\infty)} + \frac{n}{N} F^{\mathbf{A}_n}(x) \\ &= (1 - c_N) \mathbb{I}_{[0,\infty)} + c_N F^{\mathbf{A}_n}(x). \end{aligned} \tag{7.8}$$

where $\mathbb{I}_{[0,\infty)}$ is the indicator function that is equal to 1 when $x \geq 0$ and is equal to zero otherwise.

Similarly, when $c_N > 1$, \mathbf{A}_n will have $n - N$ eigenvalues of magnitude zero while the remaining N eigenvalues will be identically equal to the eigenvalues of \mathbf{B}_N . Thus the e.d.f. of \mathbf{A}_n is related to the e.d.f. of \mathbf{B}_N as

$$\begin{aligned} F^{\mathbf{A}_n}(x) &= \frac{n-N}{n} \mathbb{I}_{[0,\infty)} + \frac{N}{n} F^{\mathbf{B}_N}(x) \\ &= \left(1 - \frac{1}{c_N}\right) \mathbb{I}_{[0,\infty)} + \frac{1}{c_N} F^{\mathbf{B}_N}(x). \end{aligned} \tag{7.9}$$

Equation (7.9) is (7.8) rearranged; so we do not need to differentiate between the case when $c_N < 1$ and $c_N > 1$.

Thus, as $n, N \rightarrow \infty$ with $c_N = n/N \rightarrow c$, if $F^{\mathbf{A}_n}$ converges to a non-random d.f.

F^A , then $F^{\mathbf{B}_N}$ will also converge to a non-random d.f. F^B related to F^A by

$$F^B(x) = (1-c)\mathbb{I}_{[0,\infty)} + cF^A(x). \quad (7.10)$$

From (7.10), it is evident that the Stieltjes transform of the limiting distribution functions F^A and F^B are related as

$$m_A(z) = -\left(1 - \frac{1}{c}\right)\frac{1}{z} + \frac{1}{c}m_B(z). \quad (7.11)$$

Rearranging the terms on either side of (7.11) allows us to express $m_B(z)$ in terms of $m_A(z)$ as

$$m_B(z) = -\frac{1-c}{z} + cm_A(z). \quad (7.12)$$

Equation (7.12) can be expressed as an operational law on L_{mz}^A as

$$\boxed{L_{\text{mz}}^B(m, z) = L_{\text{mz}}^A\left(-\left(1 - \frac{1}{c}\right)\frac{1}{z} + \frac{1}{c}m, z\right)}. \quad (7.13)$$

Given L_{mz}^A , we can obtain L_{mz}^B by using (7.13). Hence $\mathbf{B}_N \mapsto B \in \mathcal{M}_{\text{alg}}$. \square

Theorem 7.34. *Let $\mathbf{A}_N \mapsto A \in \mathcal{M}_{\text{alg}}$. Then*

$$\mathbf{B}_N = (\mathbf{A}_N)^2 \mapsto B \in \mathcal{M}_{\text{alg}}.$$

PROOF. Here we have $h(x) = x^2$ which is continuous everywhere. From Proposition 7.25, $\mathbf{B}_N \mapsto B$. The Stieltjes transform of F^B can be expressed as

$$m_B(z) = E_Y\left[\frac{1}{y-z}\right] = E_X\left[\frac{1}{x^2-z}\right]. \quad (7.14)$$

Equation (7.14) can be rewritten as

$$m_B(z) = \frac{1}{2\sqrt{z}} \int \frac{1}{x-\sqrt{z}} dF^A(x) - \frac{1}{2\sqrt{z}} \int \frac{1}{x+\sqrt{z}} dF^A(x) \quad (7.15)$$

$$= \frac{1}{2\sqrt{z}} m_A(\sqrt{z}) - \frac{1}{2\sqrt{z}} m_A(-\sqrt{z}). \quad (7.16)$$

Equation (7.15) leads to the operational law

$$\boxed{L_{\text{mz}}^B(m, z) = L_{\text{mz}}^A(2m\sqrt{z}, \sqrt{z}) \boxplus_m L_{\text{mz}}^A(-2m\sqrt{z}, \sqrt{z})}. \quad (7.17)$$

Given L_{mz}^A , we can obtain L_{mz}^B by using (7.17). This proves that $\mathbf{B}_N \mapsto B \in \mathcal{M}_{\text{alg}}$. \square

Theorem 7.35. *Let $\mathbf{A}_n \mapsto A \in \mathcal{M}_{alg}$ and $\mathbf{B}_N \mapsto B \in \mathcal{M}_{alg}$. Then,*

$$\mathbf{C}_M = \text{diag}(\mathbf{A}_n, \mathbf{B}_N) \mapsto C \in \mathcal{M}_{alg},$$

where $M = n + N$ and $n/N \rightarrow c > 0$ as $n, N \rightarrow \infty$.

PROOF. Let \mathbf{C}_N be an $N \times N$ block diagonal matrix formed from the $n \times n$ matrix \mathbf{A}_n and the $M \times M$ matrix \mathbf{B}_M . Let $c_N = n/N$. The e.d.f. of \mathbf{C}_N is given by

$$F^{\mathbf{C}_N} = c_N F^{\mathbf{A}_n} + (1 - c_N) F^{\mathbf{B}_M}.$$

Let $n, N \rightarrow \infty$ and $c_N = n/N \rightarrow c$. If $F^{\mathbf{A}_n}$ and $F^{\mathbf{B}_M}$ converge in distribution almost surely (or in probability) to non-random d.f.'s F^A and F^B respectively, then $F^{\mathbf{C}_N}$ will also converge in distribution almost surely (or in probability) to a non-random distribution function F^C given by

$$F^C(x) = c F^A(x) + (1 - c) F^B(x). \quad (7.18)$$

The Stieltjes transform of the distribution function F^C can hence be written in terms of the Stieltjes transforms of the distribution functions F^A and F^B as

$$m_C(z) = c m_A(z) + (1 - c) m_B(z) \quad (7.19)$$

Equation (7.19) can be expressed as an operational law on the bivariate polynomial $L_{mz}^A(m, z)$ as

$$L_{mz}^C = L_{mz}^A \left(\frac{m}{c}, z \right) \boxplus_m L_{mz}^B \left(\frac{m}{1-c}, z \right). \quad (7.20)$$

Given L_{mz}^A and L_{mz}^B , and the definition of the \boxplus_m operator in Section 6.3, L_{mz}^C is a polynomial which can be constructed explicitly. This proves that $\mathbf{C}_N \mapsto C \in \mathcal{M}_{alg}$. \square

Theorem 7.36. *If $\mathbf{A}_n = \text{diag}(\mathbf{B}_N, \alpha \mathbf{I}_{n-N})$ and α is a real valued scalar. Then,*

$$\mathbf{B}_N \mapsto B \in \mathcal{M}_{alg},$$

as $n, N \rightarrow \infty$ with $c_N = n/N \rightarrow c$,

PROOF. Assume that as $n, N \rightarrow \infty$, $c_N = n/N \rightarrow c$. As we did in the proof of Theorem 7.35, we can show that the Stieltjes transform $m_A(z)$ can be expressed in terms of $m_B(z)$ as

$$m_A(z) = \left(\frac{1}{c} - 1 \right) \frac{1}{\alpha - z} + \frac{1}{c} m_B(z). \quad (7.21)$$

This allows us to express $L_{mz}^B(m, z)$ in terms of $L_{mz}^A(m, z)$ using the relationship in (7.21) as

$$L_{mz}^B(m, z) = L_{mz}^A\left(-\left(\frac{1}{c}-1\right)\frac{1}{\alpha-z} + \frac{1}{c}m, z\right). \quad (7.22)$$

We can hence obtain L_{mz}^B from L_{mz}^A using (7.22). This proves that $\mathbf{B}_N \mapsto B \in \mathcal{M}_{\text{alg}}$. \square

Corollary 7.37. *Let $\mathbf{A}_N \mapsto A \in \mathcal{M}_{\text{alg}}$. Then*

$$\mathbf{B}_N = \text{diag}(\mathbf{A}_n, \alpha \mathbf{I}_{N-n}) \mapsto B \in \mathcal{M}_{\text{alg}},$$

for $n/N \rightarrow c > 0$ as $n, N \rightarrow \infty$.

PROOF. This follows directly from Theorem 7.35. \square

■ 7.4 Gaussian-like matrix operations on algebraic random matrices

We now consider some simple stochastic transformations that “blur” the eigenvalues of \mathbf{A}_N by injecting additional randomness. We show that canonical operations involving an algebraic random matrix \mathbf{A}_N and Gaussian-like and Wishart-like random matrices (defined next) produce an algebraic random matrix \mathbf{B}_N .

Definition 7.41 (Gaussian-like random matrix). *Let $\mathbf{Y}_{N,L}$ be an $N \times L$ matrix with independent, identically distributed (i.i.d.) elements having zero mean, unit variance and bounded higher order moments. We label the matrix $\mathbf{G}_{N,L} = \frac{1}{\sqrt{L}}\mathbf{Y}_{N,L}$ as a Gaussian-like random matrix.*

We can sample a Gaussian-like random matrix in MATLAB as

```
G = sign(randn(N,L))/sqrt(L);
```

Gaussian-like matrices are labelled thus because they exhibit the same limiting behavior in the $N \rightarrow \infty$ limit as “pure” Gaussian matrices which may be sampled in MATLAB as

```
G = randn(N,L)/sqrt(L);
```

Definition 7.42 (Wishart-like random matrix). *Let $\mathbf{G}_{N,L}$ be a Gaussian-like random matrix. We label the matrix $W_N = \mathbf{G}_{N,L} \times \mathbf{G}'_{N,L}$ as a Wishart-like random matrix. Let $c_N = N/L$. We denote a Wishart-like random matrix thus formed by $\mathbf{W}_N(c_N)$.*

Remark 7.43 (Algebraicity of Wishart-like random matrices). *The limiting eigenvalue distribution of the Wishart-like random matrix has the Marčenko-Pastur density which is an algebraic density since L_{mz}^W exists (see Table 6.2(b)).*

Proposition 7.44. *Assume that $\mathbf{G}_{N,L}$ is an $N \times L$ Gaussian-like random matrix. Let $\mathbf{A}_N \xrightarrow{\text{a.s.}} A$ be an $N \times N$ symmetric/Hermitian random matrix and $\mathbf{T}_L \xrightarrow{\text{a.s.}} T$ be an $L \times L$ diagonal atomic random matrix respectively. If $\mathbf{G}_{N,L}$, \mathbf{A}_N and \mathbf{T}_L are independent then $\mathbf{B}_N = \mathbf{A}_N + \mathbf{G}'_{N,L} \mathbf{T}_L \mathbf{G}_{N,L} \xrightarrow{\text{a.s.}} B$, as $c_L = N/L \rightarrow c$ for $N, L \rightarrow \infty$. The Stieltjes transform $m_B(z)$ of the unique distribution function F^B satisfies the equation*

$$m_B(z) = m_A \left(z - c \int \frac{x dF^T(x)}{1 + x m_B(z)} \right). \quad (7.23)$$

PROOF. This result may be found in Marčenko-Pastur [59] and Silverstein [86]. \square

We can reformulate Proposition 7.44 to obtain the following result on algebraic random matrices.

Theorem 7.45. *Let \mathbf{A}_N , $\mathbf{G}_{N,L}$ and \mathbf{T}_L be defined as in Proposition 7.44. Then*

$$\mathbf{B}_N = \mathbf{A}_N + \mathbf{G}'_{L,N} \mathbf{T}_L \mathbf{G}_{L,N} \xrightarrow{\text{a.s.}} B \in \mathcal{M}_{\text{alg}},$$

as $c_L = N/L \rightarrow c$ for $N, L \rightarrow \infty$.

PROOF. Let \mathbf{T}_L be an atomic matrix with d atomic masses of weight p_i and magnitude λ_i for $i = 1, 2, \dots, d$. From Proposition 7.44, $m_B(z)$ can be written in terms of $m_A(z)$ as

$$m_B(z) = m_A \left(z - c \sum_{i=1}^d \frac{p_i \lambda_i}{1 + \lambda_i m_B(z)} \right). \quad (7.24)$$

where we have substituted $F^T(x) = \sum_{i=1}^d p_i \mathbb{I}_{[\lambda_i, \infty)}$ into (7.23) with $\sum_i p_i = 1$.

Equation (7.24) can be expressed as an operational law on the bivariate polynomial L_{mz}^A as

$$L_{mz}^B(m, z) = L_{mz}^A(m, z - \alpha_m). \quad (7.25)$$

where $\alpha_m = c \sum_{i=1}^d p_i \lambda_i / (1 + \lambda_i m)$. This proves that $\mathbf{B}_N \xrightarrow{\text{a.s.}} B \in \mathcal{M}_{\text{alg}}$. \square

Proposition 7.46. *Assume that $\mathbf{W}_N(c_N)$ is an $N \times N$ Wishart-like random matrix. Let $\mathbf{A}_N \xrightarrow{\text{a.s.}} A$ be an $N \times N$ random Hermitian non-negative definite matrix. If $\mathbf{W}_N(c_N)$ and \mathbf{A}_N are independent, then $\mathbf{B}_N = \mathbf{A}_N \times \mathbf{W}_N(c_N) \xrightarrow{\text{a.s.}} B$ as $c_N \rightarrow c$. The Stieltjes*

transform $m_B(z)$ of the unique distribution function F^B satisfies

$$m_B(z) = \int \frac{dF^A(x)}{\{1 - c - cz m_B(z)\}x - z}. \quad (7.26)$$

PROOF. This result may be found in Bai and Silverstein [9, 86]. \square

We can reformulate Proposition 7.46 to obtain the following result on algebraic random matrices.

Theorem 7.47. *Let \mathbf{A}_N and $\mathbf{W}_N(c_N)$ satisfy the hypothesis of Proposition 7.46. Then,*

$$\mathbf{B}_N = \mathbf{A}_N \times \mathbf{W}_N(c_N) \xrightarrow{\text{a.s.}} B \in \mathcal{M}_{\text{alg}},$$

as $c_N \rightarrow c$.

PROOF. By rearranging the terms in the numerator and denominator, (7.26) can be rewritten as

$$m_B(z) = \frac{1}{1 - c - cz m_B(z)} \int \frac{dF^A(x)}{x - \frac{z}{1 - c - cz m_B(z)}}. \quad (7.27)$$

Let $\alpha_{m,z} = 1 - c - cz m_B(z)$ so that (7.27) can be rewritten as

$$m_B(z) = \frac{1}{\alpha_{m,z}} \int \frac{dF^A(x)}{x - (z/\alpha_{m,z})}. \quad (7.28)$$

We can express $m_B(z)$ in (7.28) in terms of $m_A(z)$ as

$$m_B(z) = \frac{1}{\alpha_{m,z}} m_A(z/\alpha_{m,z}). \quad (7.29)$$

Equation (7.29) can be rewritten as

$$m_A(z/\alpha_{m,z}) = \alpha_{m,z} m_B(z). \quad (7.30)$$

Equation (7.30) can be expressed as an operational law on the bivariate polynomial L_{mz}^A as

$$L_{mz}^B(m, z) = L_{mz}^A(\alpha_{m,z} m, z/\alpha_{m,z}). \quad (7.31)$$

This proves that $\mathbf{B}_N \xrightarrow{\text{a.s.}} B \in \mathcal{M}_{\text{alg}}$. \square

Proposition 7.48. *Assume that $\mathbf{G}_{N,L}$ is an $N \times L$ Gaussian-like random matrix. Let $\mathbf{A}_N \xrightarrow{\text{a.s.}} A$ be an $N \times N$ symmetric/Hermitian random matrix independent of $\mathbf{G}_{N,L}$, \mathbf{A}_N . Let $\mathbf{A}_N^{1/2}$ denote an $N \times L$ matrix. If s is a positive real-valued scalar then*

$\mathbf{B}_N = (\mathbf{A}_N^{1/2} + \sqrt{s} \mathbf{G}_{N,L})(\mathbf{A}_N^{1/2} + \sqrt{s} \mathbf{G}_{N,L})' \xrightarrow{\text{a.s.}} B$, as $c_L = N/L \rightarrow c$ for $N, L \rightarrow \infty$. The Stieltjes transform, $m_B(z)$ of the unique distribution function F^B satisfies the equation

$$m_B(z) = - \int \frac{dF^A(x)}{z \{1 + s c m_B(z)\} - \frac{x}{1 + s c m_B(z)} + s(c-1)}. \quad (7.32)$$

PROOF. This result is found in Dozier and Silverstein [30]. \square

We can reformulate Proposition 7.48 to obtain the following result on algebraic random matrices.

Theorem 7.49. *Assume \mathbf{A}_N , $\mathbf{G}_{N,L}$ and s satisfy the hypothesis of Proposition 7.48. Then*

$$\mathbf{B}_N = (\mathbf{A}_N^{1/2} + \sqrt{s} \mathbf{G}_{N,L})(\mathbf{A}_N^{1/2} + \sqrt{s} \mathbf{G}_{N,L})' \xrightarrow{\text{a.s.}} B \in \mathcal{M}_{\text{alg}},$$

as $c_L = N/L \rightarrow c$ for $N, L \rightarrow \infty$.

PROOF. By rearranging the terms in the numerator and denominator, (7.32) can be rewritten as

$$m_B(z) = \int \frac{\{1 + s c m_B(z)\} dF^A(x)}{x - \{1 + s c m_B(z)\}(z \{1 + s c m_B(z)\} + (c-1)s)}. \quad (7.33)$$

Let $\alpha_m = 1 + s c m_B(z)$ and $\beta_m = \{1 + s c m_B(z)\}(z \{1 + s c m_B(z)\} + (c-1)s)$, so that $\beta = \alpha_m^2 z + \alpha_m s(c-1)$. Equation (7.33) can hence be rewritten as

$$m_B(z) = \alpha_m \int \frac{dF^A(x)}{x - \beta_m}. \quad (7.34)$$

Using the definition of the Stieltjes transform in (6.1), we can express $m_B(z)$ in (7.34) in terms of $m_A(z)$ as

$$\begin{aligned} m_B(z) &= \alpha_m m_A(\beta_m) \\ &= \alpha_m m_A(\alpha_m^2 z + \alpha_m s(c-1)). \end{aligned} \quad (7.35)$$

Equation (7.35) can, equivalently, be rewritten as

$$m_A(\alpha_m^2 z + \alpha_m s(c-1)) = \frac{1}{\alpha_m} m_B(z). \quad (7.36)$$

Equation (7.36) can be expressed as an operational law on the bivariate polynomial L_{mz} as

$$\boxed{L_{mz}^B(m, z) = L_{mz}^A(m/\alpha_m, \alpha^2 z + \alpha_m s(c-1))}. \quad (7.37)$$

This proves that $\mathbf{B}_N \xrightarrow{\text{a.s.}} B \in \mathcal{M}_{\text{alg}}$. \square

■ 7.5 Sums and products of algebraic random matrices

Proposition 7.51. *Let $\mathbf{A}_N \xrightarrow{\mathbb{P}} A$ and $\mathbf{B}_N \xrightarrow{\mathbb{P}} B$ be $N \times N$ symmetric/Hermitian random matrices. Let \mathbf{Q}_N be a Haar distributed unitary/orthogonal matrix independent of \mathbf{A}_N and \mathbf{B}_N . Then $\mathbf{C}_N = \mathbf{A}_N + \mathbf{Q}_N \mathbf{B}_N \mathbf{Q}'_N \xrightarrow{\mathbb{P}} C$. The associated distribution function F^C is the unique distribution function whose R transform satisfies*

$$r_C(g) = r_A(g) + r_B(g). \quad (7.38)$$

PROOF. This result was obtained by Voiculescu in [106]. \square

We can reformulate Proposition 7.51 to obtain the following result on algebraic random matrices.

Theorem 7.52. *Assume that \mathbf{A}_N , \mathbf{B}_N and \mathbf{Q}_N satisfy the hypothesis of Proposition 7.51. Then,*

$$\mathbf{C}_N = \mathbf{A}_N + \mathbf{Q}_N \mathbf{B}_N \mathbf{Q}'_N \xrightarrow{\mathbb{P}} C \in \mathcal{M}_{alg}$$

PROOF. Equation (7.38) can be expressed as an operational law on the bivariate polynomials L_{rg}^A and L_{rg}^B as

$$\boxed{L_{rg}^C = L_{rg}^A \boxplus_r L_{rg}^B} \quad (7.39)$$

If L_{mz} exists then so does L_{rg} and vice-versa. This proves that $\mathbf{C}_N \xrightarrow{\mathbb{P}} C \in \mathcal{M}_{alg}$. \square

Proposition 7.53. *Let $\mathbf{A}_N \xrightarrow{\mathbb{P}} A$ and $\mathbf{B}_N \xrightarrow{\mathbb{P}} B$ be $N \times N$ symmetric/Hermitian random matrices. Let \mathbf{Q}_N be a Haar distributed unitary/orthogonal matrix independent of \mathbf{A}_N and \mathbf{B}_N . Then $\mathbf{C}_N = \mathbf{A}_N \times \mathbf{Q}_N \mathbf{B}_N \mathbf{Q}'_N \xrightarrow{\mathbb{P}} C$ where \mathbf{C}_N is defined only if \mathbf{C}_N has real eigenvalues for every sequence \mathbf{A}_N and \mathbf{B}_N . The associated distribution function F^C is the unique distribution function whose S transform satisfies*

$$s_C(y) = s_A(y) s_B(y). \quad (7.40)$$

PROOF. This result was obtained by Voiculescu in [107, 108]. \square

We can reformulate Proposition 7.53 to obtain the following result on algebraic random matrices.

Theorem 7.54. *Assume that \mathbf{A}_N , and \mathbf{B}_N satisfy the hypothesis of Proposition 7.53. Then*

$$\mathbf{C}_N = \mathbf{A}_N \times \mathbf{Q}_N \mathbf{B}_N \mathbf{Q}'_N \xrightarrow{\mathbb{P}} C \in \mathcal{M}_{alg}.$$

PROOF. Equation (7.40) can be expressed as an operational law on the bivariate poly-

nomials L_{sy}^A and L_{sy}^B as

$$\boxed{L_{\text{sy}}^C = L_{\text{sy}}^A \boxtimes_s L_{\text{sy}}^B} \quad (7.41)$$

If L_{mz} exists then so does L_{sy} and vice versa. This proves that $\mathbf{B}_N \xrightarrow{\text{p}} B \in \mathcal{M}_{\text{alg}}$. \square

Definition 7.55 (Orthogonally/Unitarily invariant random matrix). *If the joint distribution of the elements of a random matrix \mathbf{A}_N is invariant under orthogonal/unitary transformations, it is referred to as an orthogonally/unitarily invariant random matrix.*

If \mathbf{A}_N (or \mathbf{B}_N) or both are an orthogonally/unitarily invariant sequences of random matrices then Theorems 7.52 and 7.54 can be stated more simply.

Corollary 7.56. *Let $\mathbf{A}_N \xrightarrow{\text{p}} A \in \mathcal{M}_{\text{alg}}$ and $\mathbf{B}_N \rightarrow B \xrightarrow{\text{p}} \mathcal{M}_{\text{alg}}$ be a orthogonally/unitarily invariant random matrix independent of \mathbf{A}_N . Then,*

1. $\mathbf{C}_N = \mathbf{A}_N + \mathbf{B}_N \xrightarrow{\text{p}} C \in \mathcal{M}_{\text{alg}}$
2. $\mathbf{C}_N = \mathbf{A}_N \times \mathbf{B}_N \xrightarrow{\text{p}} C \in \mathcal{M}_{\text{alg}}$

Here multiplication is defined only if \mathbf{C}_N has real eigenvalues for every sequence \mathbf{A}_N and \mathbf{B}_N .

When both the limiting eigenvalue distributions of \mathbf{A}_N and \mathbf{B}_N have compact support, it is possible to strengthen the mode of convergence in Theorems 7.52 and 7.54 to almost surely [46]. We suspect that almost sure convergence must hold when the distributions are not compactly supported; this remains an open problem.

The polynomial method: Computational aspects

■ 8.1 Operational laws on bivariate polynomials

The key idea behind the definition of algebraic random matrices in Chapter 7 was that when the limiting eigenvalue distribution of a random matrix can be encoded by a bivariate polynomial, then for the broad class of random matrix operations identified in Chapter 7, algebraicity of the eigenvalue distribution is preserved under the transformation.

Our proofs relied on exploiting the fact that some random matrix transformations, say $\mathbf{A}_N \mapsto \mathbf{B}_N$, could be most naturally expressed as transformations of $L_{\text{mz}}^A \mapsto L_{\text{mz}}^B$; others as $L_{\text{rg}}^A \mapsto L_{\text{rg}}^B$ while some as $L_{\text{sy}}^A \mapsto L_{\text{sy}}^B$. Hence, we manipulate the bivariate polynomials to the form needed to apply the appropriate operational law, which we ended up deriving as part of the proof, and then reverse the transformations to obtain the bivariate polynomial L_{mz}^B . Once we have derived the operational law for computing L_{mz}^B from L_{mz}^A , we have established the algebraicity of the limiting eigenvalue distribution of \mathbf{B}_N and we are done.

These operational laws, the associated random matrix transformation and the symbolic MATLAB code for the operational law are summarized in Tables 8.1-8.3. The remainder of this chapter discusses techniques for extracting the density function from the polynomial and the special structure in the moments that allows them to be efficiently enumerated using symbolic methods.

B	Operation	$L_{mz}^B(m, z)$	MATLAB code
Deterministic Transformations			
$\begin{bmatrix} p\mathbf{A}+q\mathbf{I} \\ r\mathbf{A}+s\mathbf{I} \end{bmatrix}$	“Möbius”	$L_{mz}^A \left(\frac{m - \beta_2 r}{\beta_2 s - \beta_2 r \alpha_2}, -\alpha_2 \right)$, where $\alpha_2 = (q - s z) / (p - r z)$, and $\beta_2 = 1 / (p - r z)$.	function LmzB = mobiusA(LmzA,p,q,r,s) syms m z alpha = (q-s*z)/(p-r*z);beta=1/(p-r*z); temp_pol = subs(LmzA,z,-alpha); LmzB = irreducLuv(temp_pol,m,z);
\mathbf{A}^{-1}	“Invert”	$L_{mz}^A \left(-z - z^2 m, \frac{1}{z} \right)$	function LmzB = invA(LmzA) LmzB = mobiusA(LmzA,0,1,1,0);
$\mathbf{A} + \alpha \mathbf{I}$	“Translate”	$L_{mz}^A(m, z - \alpha)$	function LmzB = shiftA(LmzA,alpha) LmzB = mobiusA(LmzA,1,alpha,0,1);
$\alpha \mathbf{A}$	“Scale”	$L_{mz}^A \left(\alpha m, \frac{z}{\alpha} \right)$	function LmzB = scaleA(LmzA) LmzB = mobiusA(LmzA,alpha,0,0,1);
$\begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \alpha \mathbf{I} \end{bmatrix}$	“Projection/ Transpose”	Size of $\mathbf{A} \rightarrow c > 1$ Size of $\mathbf{B} \rightarrow c < 1$ $L_{mz}^A \left(\left(1 - \frac{1}{c}\right) \frac{1}{\alpha - z} + \frac{m}{c}, z \right)$	function LmzB = projectA(LmzA,c,alpha) syms m z mb = (1-(1/c))*(1/(alpha-z))+m/c; temp_pol = subs(LmzA,m,mb); LmzB = irreducLuv(temp_pol,m,z);
$\mathbf{A} = \begin{bmatrix} \mathbf{B} & \mathbf{0} \\ \mathbf{0} & \alpha \mathbf{I} \end{bmatrix}$	“Augmentation”	Size of $\mathbf{A} \rightarrow c < 1$ Size of $\mathbf{B} \rightarrow c > 1$	function LmzB = augmentA(LmzA,c,alpha) syms m z mb = (1-(1/c))*(1/(alpha-z))+m/c; temp_pol = subs(LmzA,m,mb); LmzB = irreducLuv(temp_pol,m,z);
Stochastic Transformations			
$\mathbf{A} + \mathbf{G}' \mathbf{T} \mathbf{G}$	“Add Atomic Wishart”	$L_{mz}^A(m, z - \alpha_m)$, where $\alpha_m = c \sum_{i=1}^d \frac{p_i \lambda_i}{1 + \lambda_i m}$, with $\sum_i p_i = 1$.	function LmzB = AplusWish(LmzA,c,p,lambdas) syms m z alpha = z-c*sum(p.*(lambda./(1+lambda*m))); temp_pol = subs(LmzA,z,alpha); LmzB = irreducLuv(temp_pol,m,z);
$\mathbf{A} \times \mathbf{W}(c)$	“Multiply Wishart”	$L_{mz}^A \left(\alpha_m, z \frac{m}{\alpha_m} \right)$, where $\alpha_{m,z} = (1 - c - c z m)$.	function LmzB = AtimesWish(LmzA,c) syms m z z1 alpha = (1-c*z1*m); temp_pol = subs(LmzA,m,m*alpha); temp_pol = subs(temp_pol,z1/alpha); temp_pol = subs(temp_pol,z1,z); % Replace dummy variable LmzB = irreducLuv(temp_pol,m,z);
$\begin{pmatrix} \mathbf{A}^{1/2} + \sqrt{s} \mathbf{G} \\ \times \\ \mathbf{A}^{1/2} + \sqrt{s} \mathbf{G}' \end{pmatrix}$	“Grammian”	$L_{mz}^A \left(\frac{m}{\alpha_m}, \alpha_m^2 z + \alpha_m s(c - 1) \right)$, where $\alpha_m = 1 + s c m$.	function LmzB = AgramG(LmzA,c,s) syms m z alpha = (1+s*c*m); beta = alpha*(z*alpha+s*(c-1)); temp_pol = subs(subs(LmzA,m,m/alpha),z,beta); LmzB = irreducLuv(temp_pol,m,z);

Table 8.1. Operational laws on the bivariate polynomial encodings (and their computational realization in MATLAB) corresponding to a class of deterministic and stochastic transformations. The Gaussian-like random matrix \mathbf{G} is an $N \times L$, the Wishart-like matrix $\mathbf{W}(c) = \mathbf{G} \mathbf{G}'$ where $N/L \rightarrow c > 0$ as $N, L \rightarrow \infty$, and the matrix \mathbf{T} is a diagonal atomic random matrix.

Operational Law	MATLAB Code
<p style="text-align: center;"> $L_{mz}^A(2m\sqrt{z}, \sqrt{z})$ and $L_{mz}^A(-2m\sqrt{z}, -\sqrt{z})$ \searrow and \swarrow \boxplus_m \downarrow L_{mz}^B </p>	<pre>function LmzB = squareA(LmzA) syms m z Lmz1 = subs(LmzA,z,sqrt(z)); Lmz1 = subs(Lmz1,m,2*m*sqrt(z)); Lmz2 = subs(LmzA,z,-sqrt(z)); Lmz2 = subs(Lmz2,m,-2*m*sqrt(z)); LmzB = L1plusL2(Lmz1,Lmz2,m); LmzB = irreducLuv(LmzB,m,z);</pre>

(a) $L_{mz}^A \mapsto L_{mz}^B$ for $\mathbf{A} \mapsto \mathbf{B} = \mathbf{A}^2$.

Operational Law	MATLAB Code
<p style="text-align: center;"> $L_{mz}^A(\frac{m}{c}, z)$ and $L_{mz}^B(\frac{m}{1-c}, z)$ \searrow and \swarrow \boxplus_m \downarrow L_{mz}^C </p>	<pre>function LmzC = AblockB(LmzA,LmzB,c) syms m z mu LmzA1 = subs(LmzA,m,m/c); LmzB1 = subs(LmzB,m,m/(1-c)); LmzC = L1plusL2(LmzA1,LmzB1,m); LmzC = irreducLuv(LmzC,m,z);</pre>

(b) $L_{mz}^A, L_{mz}^B \mapsto L_{mz}^C$ for $\mathbf{A}, \mathbf{B} \mapsto \mathbf{C} = \text{diag}(\mathbf{A}, \mathbf{B})$ where Size of \mathbf{A} / Size of $\mathbf{C} \rightarrow c$.

Table 8.2. Operational laws on the bivariate polynomial encodings for some deterministic random matrix transformations. The operations \boxplus_u and \boxtimes_u are defined in Table 6.5.

Operational Law	MATLAB Code
$ \begin{array}{ccc} L_{mz}^A & & L_{mz}^B \\ \downarrow & & \downarrow \\ L_{rg}^A & & L_{rg}^B \\ \swarrow & & \searrow \\ & \boxplus_r & \\ \downarrow & & \\ L_{rg}^C & & \\ \downarrow & & \\ L_{mz}^C & & \end{array} $	<pre> function LmzC = AplusB(LmzA,LmzB) syms m z r g LrgA = Lmz2Lrg(LmzA); LrgB = Lmz2Lrg(LmzB); LrgC = L1plusL2(LrgA,LrgB,r); LmzC = Lrg2Lmz(LrgC); </pre>

(a) $L_{mz}^A, L_{mz}^B \mapsto L_{mz}^C$ for $A, B \mapsto C = A + QBQ'$.

Operational Law	MATLAB Code
$ \begin{array}{ccc} L_{mz}^A & & L_{mz}^B \\ \downarrow & & \downarrow \\ L_{sy}^A & & L_{sy}^B \\ \swarrow & & \searrow \\ & \boxtimes_s & \\ \downarrow & & \\ L_{sy}^C & & \\ \downarrow & & \\ L_{mz}^C & & \end{array} $	<pre> function LmzC = AtimesB(LmzA,LmzB) syms m z s y LsyA = Lmz2Lsy(LmzA); LsyB = Lmz2Lsy(LmzB); LsyC = L1timesL2(LsyA,LsyB,s); LmzC = Lsy2Lmz(LsyC); </pre>

(b) $L_{mz}^A, L_{mz}^B \mapsto L_{mz}^C$ for $A, B \mapsto C = A \times QBQ'$.

Table 8.3. Operational laws on the bivariate polynomial encodings for some canonical random matrix transformations. The operations \boxplus_u and \boxtimes_u are defined in Table 6.5.

■ 8.2 Interpreting the solution curves of polynomial equations

Consider a bivariate polynomial L_{mz} . Let D_m be the degree of $L_{mz}(m, z)$ with respect to m and $l_k(z)$, for $k = 0, \dots, D_m$, be polynomials in z that are the coefficients of m^k . For every z along the real axis, there are at most D_m solutions to the polynomial equation $L_{mz}(m, z) = 0$. The solutions of the bivariate polynomial equation $L_{mz} = 0$ define a locus of points (m, z) in $\mathbb{C} \times \mathbb{C}$ referred to as a complex algebraic curve. Since the limiting density is over \mathbb{R} , we may focus on real values of z .

For almost every $z \in \mathbb{R}$, there will be D_m values of m . The exception consists of the singularities of $L_{mz}(m, z)$. A singularity occurs at $z = z_0$ if:

- There is a reduction in the degree of m at z_0 so that there are less than D_m roots for $z = z_0$. This occurs when $l_{D_m}(z_0) = 0$. Poles of $L_{mz}(m, z)$ occur if some of the m -solutions blow up to infinity.
- There are multiple roots of L_{mz} at z_0 so that some of the values of m coalesce.

The singularities constitute the so-called exceptional set of $L_{mz}(m, z)$. Singularity analysis, in the context of algebraic functions, is a well studied problem [37] from which we know that the singularities of $L_{mz}^A(m, z)$ are constrained to be *branch points*.

A *branch* of the algebraic curve $L_{mz}(m, z) = 0$ is the choice of a locally analytic function $m_j(z)$ defined outside the exceptional set of $L_{mz}^A(m, z)$ together with a connected region of the $\mathbb{C} \times \mathbb{R}$ plane throughout which this particular choice $m_j(z)$ is analytic. These properties of singularities and branches of algebraic curve are helpful in determining the atomic and non-atomic component of the encoded probability density from L_{mz} . We note that, as yet, we do not have a fully automated algorithm for extracting the limiting density function from the bivariate polynomial. Development of efficient computational algorithms that exploit the algebraic properties of the solution curve would be of great benefit to the community.

■ 8.2.1 The atomic component

If there are any atomic components in the limiting density function, they will necessarily manifest themselves as poles of $L_{mz}(m, z)$. This follows from the definition of the Stieltjes transform in (6.1). As mentioned in the discussion on the singularities of algebraic curves, the poles are located at the roots of $l_{D_m}(z)$. These may be computed in MAPLE using the sequence of commands:

```
> Dm := degree(LmzA, m);
> lDmz := coeff(LmzA, m, Dm);
> poles := solve(lDmz=0, z);
```

We can then compute the Puiseux expansion about each of the poles at $z = z_0$. This can be computed in MAPLE using the `algcures` package as:

```
> with(algcures):
> puiseux(Lmz, z=pole, m, 1);
```

For the pole at $z = z_0$, we inspect the Puiseux expansions for branches with leading term $1/(z_0 - z)$. An atomic component in the limiting spectrum occurs if and only if the coefficient of such a branch is non-negative and not greater than one. This constraint ensures that the branch is associated with the Stieltjes transform of a valid probability distribution function.

Of course, as is often the case with algebraic curves, pathological cases can be easily constructed. For example, more than one branch of the Puiseux expansion might correspond to a candidate atomic component, *i.e.*, the coefficients are non-negative and not greater than one. In our experimentation, whenever this has happened it has been possible to eliminate the spurious branch by matrix theoretic arguments. Demonstrating this rigorously using analytical arguments remains an open problem.

Sometimes it is possible to encounter a double pole at $z = z_0$ corresponding to two admissible weights. In such cases, empirical evidence suggests that the branch with the largest coefficient (less than one) is the “right” Puiseux expansion though we have no theoretical justification for this choice.

■ 8.2.2 The non-atomic component

The probability density function can be recovered from the Stieltjes transform by applying the inversion formula in (6.4). Since the Stieltjes transform is encoded in the bivariate polynomial L_{mz} , we accomplish this by first computing all D_m roots along $z \in \mathbb{R}$ (except at poles or singularities). There will be D_m roots of which one solution curve will be the “correct” solution, *i.e.*, the non-atomic component of the desired density function is the imaginary part of the correct solution normalized by π . In MATLAB, the D_m roots can be computed using the sequence of commands:

```
Lmz_roots = [];
x_range = [x_start:x_step:x_end];
for x = x_range
    Lmz_roots_unnorm = roots(sym2poly(subs(Lmz,z,x)));
    Lmz_roots = [Lmz_roots;
                 real(Lmz_roots_unnorm) + i*imag(Lmz_roots_unnorm)/pi];
end
```

The density of the limiting eigenvalue distribution function can be, generically, be expressed in closed form when $D_m = 2$. When using root-finding algorithms, for $D_m = 2, 3$, the correct solution can often be easily identified; the imaginary branch will always appear with its complex conjugate. The density is just the scaled (by $1/\pi$) positive imaginary component.

When $D_m \geq 4$, except when L_{mz} is bi-quadratic for $D_m = 4$, there is no choice but to manually isolate the correct solution among the numerically computed D_m roots of the polynomial $L_{mz}^{(m)}(z)$ at each $z = z_0$. The class of algebraic random matrices whose eigenvalue density function can be expressed in closed form is thus a much smaller subset of the class of algebraic random matrices. When the underlying density function is compactly supported, the boundary points will be singularities of the algebraic curve.

In particular, when the probability density function is compactly supported and the boundary points are not poles, they occur at points where some values of m coalesce. These points are the roots of the discriminant of L_{mz} , computed in MAPLE as:

```
> PossibleBoundaryPoints = solve(discrim(Lmz,m),z);
```

We suspect that “nearly all” algebraic random matrices with compactly supported eigenvalue distribution will exhibit a square root type behavior near boundary points at which there are no poles. In the generic case, this will occur whenever the boundary points correspond to locations where two branches of the algebraic curve coalesce.


For a class of random matrices that includes a subclass of algebraic random matrices, this has been established in [87]. This endpoint behavior has also been observed orthogonally/unitarily invariant random matrices whose distribution has the element-wise joint density function of the form

$$f(\mathbf{A}) = C_N \exp(-N \text{Tr} V(\mathbf{A})) d\mathbf{A}$$


where V is an even degree polynomial with positive leading coefficient and $d\mathbf{A}$ is the Lebesgue measure on $N \times N$ symmetric/Hermitian matrices. In [27], it is shown that these random matrices have a limiting mean eigenvalue density in the $N \rightarrow \infty$ limit that is algebraic and compactly supported. The behavior at the endpoint typically vanishes like a square root, though higher order vanishing at endpoints is possible and a full classification is made in [28]. In [53] it is shown that square root vanishing is generic. A similar classification for the general class of algebraic random matrices remains an open problem.

Whether the encoded distribution is compactly supported or not, the $-1/z$ behavior of the real part of Stieltjes transform (the principal value) as $z \rightarrow \pm\infty$ helps isolate the correct solution. In our experience, while multiple solution curves might exhibit this behavior, invariably only one solution will have an imaginary branch that, when normalized, will correspond to a valid probability density. Why this always appears to be the case for the operational laws described is a bit of a mystery to us.

Example: Consider the Marčenko-Pastur density encoded by L_{mz} given in Table 6.2(b). The Puiseux expansion about the pole at $z = 0$ (the only pole!), has coefficient $(1 - 1/c)$ which corresponds to an atom only when $c > 1$ (as expected using a matrix theoretic argument). Finally, the branch points at $(1 \pm \sqrt{c})^2$ correspond to boundary points of the compactly supported probability density. Figure 8-1 plots the real and imaginary parts of the algebraic curve for $c = 2$.



(a) Real component. The singularity at zero corresponds to an atom of weight $1/2$. The branch points at $(1 \pm \sqrt{2})^2$ correspond to the boundary points of the region of support.



(b) Imaginary component normalized by π . The positive component corresponds to the encoded probability density function.

Figure 8-1. The real and imaginary components of the algebraic curve defined by the equation $L_{mz}(m, z) = 0$, where $L_{mz} \equiv czm^2 - (1 - c - z)m + 1$, which encodes the Marčenko-Pastur density. The curve is plotted for $c = 2$. The $-1/z$ behavior of the real part of the “correct solution” as $z \rightarrow \infty$ is the generic behavior exhibited by the real part of the Stieltjes transform of a valid probability density function.

■ 8.3 Enumerating the moments and free cumulants

In principle, the moments generating function can be extracted from $L_{\mu z}$ by a Puiseux expansion of the algebraic function $\mu(z)$ about $z = 0$. When the moments of an algebraic probability distribution exist, there is additional structure in the moments and free cumulants that allows us to enumerate them efficiently. For an algebraic probability distribution, we conjecture that the moments of all order exist if and only if the distribution is compactly supported.

Definition 8.31 (Rational generating function). Let $\mathbb{R}[[x]]$ denote the ring of formal power series (or generating functions) in x with real coefficients. A formal power series (or generating function) $v \in \mathbb{R}[[u]]$ is said to be rational if there exist polynomials in u , $P(u)$ and $Q(u)$, $Q(0) \neq 0$ such that

$$v(u) = \frac{P(u)}{Q(u)}.$$

Definition 8.32 (Algebraic generating function). Let $\mathbb{R}[[x]]$ denote the ring of formal power series (or generating functions) in x with real coefficients. A formal power series (or generating function) $v \in \mathbb{R}[[u]]$ is said to be algebraic if there exist polynomials in u , $P_0(u), \dots, P_{D_v}(u)$, not all identically zero, such that

$$P_0(u) + P_1(u)v + \dots + P_{D_v}(u)v^{D_v} = 0.$$

The degree of v is said to be D_v .

Definition 8.33 (D-finite generating function). Let $v \in \mathbb{R}[[u]]$. If there exist polynomials $p_0(u), \dots, p_d(u)$, such that

$$p_d(u)v^{(d)} + p_{d-1}(u)v^{(d-1)} + \dots + p_1(u)v^{(1)} + p_0(u) = 0, \quad (8.1)$$

where $v^{(j)} = d^j v / du^j$. Then we say that v is a D -finite (short for differentiably finite) generating function (or power series). The generating function, $v(u)$, is also referred to as a holonomic function.

Definition 8.34 (P-recursive coefficients). Let a_n for $n \geq 0$ denote the coefficients of a D -finite series v . If there exist polynomials $P_0, \dots, P_e \in \mathbb{R}[n]$ with $P_e \neq 0$, such that

$$P_e(n)a_{n+e} + P_{e-1}(n)a_{n+e-1} + \dots + P_0(n)a_n = 0,$$

for all $n \in \mathbb{N}$, then the coefficients a_n are said to be P -recursive (short for polynomially recursive).

Proposition 8.35. Let $v \in \mathbb{R}[[u]]$ be an algebraic power series of degree D_v . Then v is D -finite and satisfies an equation (8.1) of order D_v .

PROOF. A proof appears in Stanley [94, pp.187]. \square

The structure of the limiting moments and free cumulants associated with algebraic densities is described next.

Theorem 8.36. *If $f_A \in \mathcal{P}_{alg}$, and the moments exist, then the moment and free cumulant generating functions are algebraic power series. Moreover, both generating functions are D-finite and the coefficients are P-recursive.*

PROOF. If $f_A \in \mathcal{P}_{alg}$, then L_{mz}^A exists. Hence $L_{\mu z}^A$ and L_{rg}^A exist, so that $\mu_A(z)$ and $r_A(g)$ are algebraic power series. By Theorem 8.35 they are D-finite; the moments and free cumulants are hence P-recursive. \square

There are powerful symbolic tools available for enumerating the coefficients of algebraic power series. The MAPLE based package `gfun` is one such example [77]. From the bivariate polynomial $L_{\mu z}$, we can obtain the series expansion up to degree `expansion_degree` by using the commands:

```
> with(gfun):
> MomentSeries = algeqtoseries(Lmyuz,z,myu,expansion_degree,'pos_slopes');
```

The option `pos_slopes` computes only those branches tending to zero. Similarly, the free cumulants can be enumerated from L_{rg} using the commands:

```
> with(gfun):
> FreeCumulantSeries = algeqtoseries(Lrg,g,r,expansion_degree,'pos_slopes');
```

For computing expansions to a large order, it is best to work with the recurrence relation. For an algebraic power series $v(u)$, the first `number_of_terms` coefficients can be computed from L_{uv} using the sequence of commands:

```
> with(gfun):
> deq := algeqtodiffeq(Luv,v(u));
> rec := diffeqtorec(deq,v(u),a(n));
> p_generator := rectoproc(rec,a(n),list):
> p_generator(number_of_terms);
```

Example: Consider the Marčenko-Pastur density encoded by the bivariate polynomials listed in Table 6.2(b). Using the above sequence of commands, we can enumerate the first five terms of its moment generating function as

$$\mu(z) = 1 + z + (c + 1)z^2 + (3c + c^2 + 1)z^3 + (6c^2 + c^3 + 6c + 1)z^4 + O(z^5).$$

The moment generating function is a D-Finite power series and satisfies the second order differential equation

$$-z + zc - 1 + (-z - zc + 1)\mu(z) + (z^3c^2 - 2z^2c - 2z^3c + z - 2z^2 + z^3) \frac{d}{dz}\mu(z) = 0,$$

with initial condition $\mu(0) = 1$. The moments $M_n = a(n)$ themselves are P-recursive satisfying the finite depth recursion

$$(-2c + c^2 + 1)na(n) + ((-2 - 2c)n - 3c - 3)a(n+1) + (3+n)a(n+2) = 0$$

with the initial conditions $a(0) = 1$ and $a(1) = 1$. The free cumulants can be analogously computed.

What we find rather remarkable is that for algebraic random matrices, it is often possible to enumerate the moments in closed form even when the limiting density function cannot. The linear recurrence satisfied by the moments may be used to analyze their asymptotic growth.

When using the sequence of commands described, sometimes more than one solution might emerge. In such cases, we have often found that one can identify the correct solution by checking for the positivity of even moments or the condition $\mu(0) = 1$. More sophisticated arguments might be needed for pathological cases. It might involve verifying, using techniques such as those in [3], that the coefficients enumerated correspond to the moments a valid distribution function.

■ 8.4 Computational free probability

There is a deep connection between eigenvalue distributions of random matrices and “free probability” (See Appendix A for a brief discussion). We now clarify the connection between the operational law of a subclass of algebraic random matrices and the convolution operations of free probability. This will bring into sharp focus how the polynomial method constitutes a framework for computational free probability theory.

Proposition 8.41. *Let $\mathbf{A}_N \xrightarrow{P} A$ and $\mathbf{B}_N \xrightarrow{P} B$ be two asymptotically free random matrix sequences as in Definition A.1. Then $\mathbf{A}_N + \mathbf{B}_N \xrightarrow{P} A + B$ and $\mathbf{A}_N \times \mathbf{B}_N \xrightarrow{P} AB$ (where the product is defined whenever $\mathbf{A}_N \times \mathbf{B}_N$ has real eigenvalues for every \mathbf{A}_N and \mathbf{B}_N) with the corresponding limit eigenvalue density functions, f_{A+B} and f_{AB} given by*

$$f_{A+B} = f_A \boxplus f_B \tag{8.2a}$$

$$f_{AB} = f_A \boxtimes f_B \tag{8.2b}$$

where \boxplus denotes free additive convolution and \boxtimes denotes free multiplicative convolution. These convolution operations can be expressed in terms of the R and S transforms as described in Propositions 7.51 and 7.53 respectively.

Free additive convolution	$f_{A+B} = f_A \boxplus f_B$	$L_{\text{rg}}^{A+B} = L_{\text{rg}}^A \boxplus_r L_{\text{rg}}^B$
Free multiplicative convolution	$f_{A \times B} = f_A \boxtimes f_B$	$L_{\text{sy}}^{A \times B} = L_{\text{sy}}^A \boxtimes_s L_{\text{sy}}^B$

Table 8.4. Implicit representation of the free convolution of two algebraic probability densities.

PROOF. This result appears for density functions with compact support in [106, 107]. It was later strengthened to the case of density functions with unbounded support. See [46] for additional details and references. \square

In Theorems 7.52 and 7.54 we, in effect, showed that the free convolution of algebraic densities produces an algebraic density. This stated succinctly next.

Corollary 8.42. *Algebraic probability distributions form a semi-group under free additive convolution.*

Corollary 8.43. *Algebraic distributions with positive semi-definite support form a semi-group under free multiplicative convolution.*

This establishes a framework for computational free probability theory by identifying the class of distributions for which the free convolution operations produce a “computable” distribution.

■ 8.4.1 Implicitly encoding the free convolution computations

The computational framework established relies on being able to implicitly encode free convolution computations as a resultant computation on appropriate bivariate polynomials as in Table 8.4. This leads to the obvious question: Are there other more *effective* ways to implicitly encode free convolution computations? The answer to this rhetorical question will bring into sharp focus the reason why the bivariate polynomial encoding at the heart of the polynomial method is indispensable for any symbolic computational implementation of free convolution. First, we answer the analogous question about the most effective encoding for classical convolution computations.

Recall that classical convolution can be expressed in terms of the Laplace transform of the distribution function. In what follows, we assume that the distributions have finite moments¹. Hence the Laplace transform can be written as a formal exponential moment generating function. Classical additive and multiplicative convolution of two distributions produces a distribution whose exponential moment generating function equals the series (or Cauchy) product and the coefficient-wise (or Hadamard) product of

¹In the general case, tools from complex analysis can be used to extend the argument.

the individual exponential moment generating functions, respectively. Often, however, the Laplace transform of either or both the individual distributions being convolved cannot be written in closed form. The next best thing to do then is to find an implicit way to encode the Laplace transform and to do the convolution computations via this representation.

When this point of view is adopted, the task of identifying candidate encodings is reduced to finding the class of representations of the exponential generating function that remains closed under the Cauchy and Hadamard product. Clearly, rational generating functions (see Definition 8.31) satisfy this requirement. It is shown in Theorem 6.4.12 [94, pp.194], that D-finite generating functions (see Definition 8.33) satisfy this requirement as well.

Proposition 8.35 establishes that all algebraic generating functions (see Definition 8.32) and by extension, rational generating functions, are also D-finite. However, not all D-finite generating functions are algebraic (see Exercise 6.1 [94, pp. 217] for a counter-example) so that algebraic generating functions do not satisfy the closure requirement. Furthermore, from Proposition 6.4.3 and Theorem 6.4.12 in [94], if the *ordinary* generating function is D-finite then so is the *exponential* generating function and vice versa. Thus D-finite generating functions are the largest class of generating functions for which classical convolution computations can be performed via an implicit representation.

In the context of developing a computational framework based on the chosen implicit representation, it is important to consider computability and algorithmic efficiency issues. The class of D-finite functions is well-suited in that regard as well [77] so that we regard it as the most *effective* class of representations in which the classical convolution computations may be performed implicitly.

However, this class is inadequate for performing free convolution computations implicitly. This is a consequence of the prominent role occupied in this theory by ordinary generating functions. Specifically, the ordinary formal R and S power series, are obtained from the ordinary moment generating function by functional inversion (or reversion), and are the key ingredients of free additive and multiplicative convolution (see Propositions 8.41, 7.51 and 7.53). The task of identifying candidate encodings is thus reduced to finding the class of representations of the ordinary moment generating function that remains closed under addition, the Cauchy product, *and* reversion. D-finite functions only satisfy the first two conditions and are hence unsuitable representations.

Algebraic functions do, however, satisfy all three conditions. The algorithmic efficiency of computing the resultant (see Section 6.4) justifies our labelling of the bivariate polynomial encoding as the most *effective* way of implicitly encoding free convolution computations. The candidacy of constructibly D-finite generating functions [14], which do not contain the class of D-finite functions but do contain the class of algebraic functions, merits further investigation since they are closed under reversion, addition and multiplication. Identifying classes of representations of generating functions for which *both* the classical and free convolution computations can be performed implicitly and effectively remains an important open problem.

The polynomial method: Applications

We illustrate the use of the computational techniques developed in Chapter 8 with some examples.

■ 9.1 The Jacobi random matrix

The Jacobi matrix ensemble is defined in terms of two independent Wishart matrices $\mathbf{W}_1(c_1)$ and $\mathbf{W}_2(c_2)$ as $\mathbf{J} = (\mathbf{I} + \mathbf{W}_2(c_2) \mathbf{W}_1^{-1}(c_1))^{-1}$. The subscripts are not to be confused for the size of the matrices. Listing the computational steps needed to generate a realization of this ensemble, as in Table 9.1, is the easiest way to identify the sequence of random matrix operations needed to obtain L_{mz}^J .

Transformation	Numerical MATLAB code	Symbolic MATLAB code
Initialization	% Pick n, c1, c2 N1=n/c1; N2=n/c2;	% Define symbolic variables syms m c z;
$\mathbf{A}_1 = \mathbf{I}$	A1 = eye(n,n);	Lmz1 = m*(1-z)-1;
$\mathbf{A}_2 = \mathbf{W}_1(c_1) \times \mathbf{A}_1$	G1 = randn(n,N1)/sqrt(N1); W1 = G1*G1'; A2 = W1*A1;	Lmz2 = AtimesWish(Lmz1,c1);
$\mathbf{A}_3 = \mathbf{A}_2^{-1}$	A3 = inv(A2);	Lmz3 = invA(Lmz2);
$\mathbf{A}_4 = \mathbf{W}_2(c_2) \times \mathbf{A}_3$	G2 = randn(n,N2)/sqrt(N2); W2 = G2*G2'; A4 = W2*A3;	Lmz4 = AtimesWish(Lmz3,c2);
$\mathbf{A}_5 = \mathbf{A}_4 + \mathbf{I}$	A5 = A4+I;	Lmz5 = shiftA(Lmz4,1);
$\mathbf{A}_6 = \mathbf{A}_5^{-1}$	A6 = inv(A5);	Lmz6 = invA(Lmz5);

Table 9.1. Sequence of MATLAB commands for sampling the Jacobi ensemble. The functions used to generate the corresponding bivariate polynomials symbolically are listed in Table 8.1

We first start off with $\mathbf{A}_1 = \mathbf{I}$. The bivariate polynomial that encodes the Stieltjes transform of its eigenvalue distribution function is given by

$$L_{mz}^1(m, z) = (1 - z)m - 1. \quad (9.1)$$

For $\mathbf{A}_2 = \mathbf{W}_1(c_1) \times \mathbf{A}_1$, we can use (7.31) to obtain the bivariate polynomial

$$L_{mz}^2(m, z) = z c_1 m^2 - (-c_1 - z + 1)m + 1. \quad (9.2)$$

For $\mathbf{A}_3 = \mathbf{A}_2^{-1}$, from (7.7), we obtain the bivariate polynomial

$$L_{mz}^3(m, z) = z^2 c_1 m^2 + (c_1 z + z - 1)m + 1. \quad (9.3)$$

For $\mathbf{A}_4 = \mathbf{W}_2(c_2) \times \mathbf{A}_3$. We can use (7.31) to obtain the bivariate polynomial

$$L_{mz}^4(m, z) = (c_1 z^2 + c_2 z) m^2 + (c_1 z + z - 1 + c_2)m + 1. \quad (9.4)$$

For $\mathbf{A}_5 = \mathbf{A}_4 + \mathbf{I}$, from (7.7), we obtain the bivariate polynomial

$$L_{mz}^5(m, z) = \left((z - 1)^2 c_1 + c_2 (z - 1) \right) m^2 + (c_1 (z - 1) + z - 2 + c_2)m + 1. \quad (9.5)$$

Finally, for $\mathbf{J} = \mathbf{A}_6 = \mathbf{A}_5^{-1}$, from (7.7), we obtain the required bivariate polynomial

$$\begin{aligned} L_{mz}^J(m, z) \equiv L_{mz}^6(m, z) &= (c_1 z + z^3 c_1 - 2 c_1 z^2 - c_2 z^3 + c_2 z^2) m^2 \\ &+ (-1 + 2 z + c_1 - 3 c_1 z + 2 c_1 z^2 + c_2 z - 2 c_2 z^2) m - c_2 z - c_1 + 2 + c_1 z. \end{aligned} \quad (9.6)$$

Using matrix theoretic arguments, it is clear that the random matrix ensembles $\mathbf{A}_3, \dots, \mathbf{A}_6$ are defined only when $c_1 < 1$. There will be an atomic mass of weight $(1 - 1/c_2)$ at 1 whenever $c_2 > 1$. The non-atomic component of the distribution will have a region of support $\mathcal{S}^\cap = (a_-, a_+)$. The limiting density function for each of these ensembles can be expressed as

$$\boxed{f_{A_i}(x) = \frac{\sqrt{(x - a_-)(a_+ - x)}}{2\pi l_2(x)} \quad \text{for } a_- < x < a_+,} \quad (9.7)$$

for $i = 2, \dots, 6$, where a_-, a_+ , where the polynomials $l_2(x)$ are listed in Table 9.2. The moments for the general case when $c_1 \neq c_2$ can be enumerated using the techniques described; they will be quite messy. Instead, consider the special case when $c_1 = c_2 = c$. Using the tools described, the first four terms of the moment series, $\mu(z) = \mu_J(z)$, can

	$l_2(x)$	a_{\pm}
A_2	$x c_1$	$(1 \pm \sqrt{c_1})^2$
A_3	$x^2 c_1$	$\frac{1}{(1 \mp \sqrt{c_1})^2}$
A_4	$c_1 x^2 + c_2 x$	$\frac{1 + c_1 + c_2 - c_1 c_2 \pm 2\sqrt{c_1 + c_2 - c_1 c_2}}{(1 - c_1)^2}$
A_5	$c_1(x-1)^2 + c_2(x-1)$	$\frac{c_1^2 - c_1 + 2 + c_2 - c_1 c_2 \pm 2\sqrt{c_1 + c_2 - c_1 c_2}}{(1 - c_1)^2}$
A_6	$(c_1 x + x^3 c_1 - 2 c_1 x^2 - c_2 x^3 + c_2 x^2)$	$\frac{c_1^2 - c_1 + 2 + c_2 - c_1 c_2 \mp 2\sqrt{c_1 + c_2 - c_1 c_2}}{(1 - c_1)^2}$

Table 9.2. Parameters for determining the limiting eigenvalue density function using (9.7).




Figure 9-1. The limiting density (solid line), $f_{A_6}(x)$, given by (9.7) with $c_1 = 0.1$ and $c_2 = 0.625$ is compared with the normalized histogram of the eigenvalues of a Jacobi matrix generated using the code in Table 9.1 over 4000 Monte-Carlo trials with $n = 100$, $N_1 = n/c_1 = 1000$ and $N_2 = n/c_2 = 160$.

be computed directly from $L_{\mu z}^J$ as

$$\begin{aligned} \mu(z) = & \frac{1}{2} + \left(\frac{1}{8}c + \frac{1}{4}\right)z + \left(\frac{3}{16}c + \frac{1}{8}\right)z^2 + \left(\frac{1}{32}c^2 + \frac{3}{16}c - \frac{1}{128}c^3 + \frac{1}{16}\right)z^3 \\ & + \left(-\frac{5}{256}c^3 + \frac{5}{64}c^2 + \frac{5}{32}c + \frac{1}{32}\right)z^4 + O(z^5). \end{aligned}$$

The moment generating function satisfies the differential equation

$$\begin{aligned} -3z + 2 + zc + (-6z^2 + z^3 + 10z + z^3c^2 - 2z^3c - 4)\mu(z) \\ + (z^4 - 5z^3 - 2z^4c + 8z^2 + z^4c^2 + 2z^3c - 4z - z^3c^2) \frac{d}{dz}\mu(z) = 0, \end{aligned}$$

with the initial condition $\mu(0) = 1$. The moments $a(n) = M_n$ themselves are P-recursive and obtained by the recursion

$$\begin{aligned} &(-2c + c^2 + 1 + (-2c + c^2 + 1)n)a(n) + ((-5 + 2c - c^2)n - 11 + 2c - c^2)a(n+1) \\ &\quad + (26 + 8n)a(n+2) + (-16 - 4n)a(n+3) = 0, \end{aligned}$$

with the initial conditions $a(0) = 1/2$, $a(1) = 1/8c + 1/4$, and $a(2) = 3/16c + 1/8$. We can similarly compute the recursion for the free cumulants, $a(n) = K_{n+1}$, as

$$nc^2a(n) + (12 + 4n)a(n+2) = 0,$$

with the initial conditions $a(0) = 1/2$, and $a(1) = 1/8c$.

■ 9.2 Random compression of a matrix

Theorem 9.21. *Let $\mathbf{A}_N \mapsto A \in \mathcal{P}_{\text{alg}}$. Let \mathbf{Q}_N be an $N \times N$ Haar unitary/orthogonal random matrix independent of \mathbf{A}_N . Let \mathbf{B}_n be the upper $n \times n$ block of $\mathbf{Q}_N \mathbf{A}_N \mathbf{Q}'_N$. Then*

$$\mathbf{B}_n \mapsto B \in \mathcal{P}_{\text{alg}}$$

as $n/N \rightarrow c$ for $n, N \rightarrow \infty$.

PROOF. Let \mathbf{P}_N be an $N \times N$ projection matrix

$$\mathbf{P}_N \equiv \mathbf{Q}_N \begin{bmatrix} \mathbf{I}_n & \\ & \mathbf{0}_{N-n} \end{bmatrix} \mathbf{Q}'_N.$$

By definition, \mathbf{P}_N is an atomic matrix so that $\mathbf{P}_N \rightarrow P \in \mathcal{M}_{\text{alg}}$ as $n/N \rightarrow c$ for $n, N \rightarrow \infty$. Let $\tilde{\mathbf{B}}_N = \mathbf{P}_N \times \mathbf{A}_N$. By Corollary 7.56, $\tilde{\mathbf{B}}_N \rightarrow \tilde{B} \in \mathcal{M}_{\text{alg}}$. Finally, from Theorem 7.36, we have that $\mathbf{B}_n \rightarrow B \in \mathcal{M}_{\text{alg}}$. \square

The proof above provides a recipe for computing the bivariate polynomial L_{mz}^B explicitly as a function of L_{mz}^A and the compression factor c . For this particular application, however, one can use first principles [93] to directly obtain the relationship

$$r_B(g) = r_A(cg),$$

expressed in terms of the R transform. This translates into the operational law

$$\boxed{L_{\text{rg}}^B(r, g) = L_{\text{rg}}^A(r, cg)}. \quad (9.8)$$

Example: Consider the atomic matrix \mathbf{A}_N half of whose eigenvalues are equal to one while the remainder are equal to zero. Its eigenvalue distribution function is given by




Figure 9-2. The limiting eigenvalue density function (solid line) of the top $0.4N \times 0.4N$ block of a randomly rotated matrix is compared with the experimental histogram collected over 4000 trials with $N = 200$. Half of the eigenvalues of the original matrix were equal to one while the remainder were equal to zero.

(6.26). From the bivariate polynomial, L_{rg}^A in Table 6.2(a) and (9.8) it can be shown that the limiting eigenvalue distribution function of \mathbf{B}_n , constructed from \mathbf{A}_N as in Theorem 9.21, is encoded by the polynomial

$$L_{\text{mz}}^B = (-2cz^2 + 2cz)m^2 - (-2c + 4cz + 1 - 2z)m - 2c + 2,$$

where c is the limiting compression factor. Poles occur at $z = 0$ and $z = 1$. The leading terms of the Puiseux expansion of the two branches about the poles at $z = z_0$ are

$$\left\{ \left(\frac{z - z_0}{-2c + 4c^2} + \frac{1 - 2c}{2c} \right) \frac{1}{z - z_0}, \frac{2c - 2}{-1 + 2c} \right\}.$$

It can be easily seen that when $c > 1/2$, the Puiseux expansion about the poles $z = z_0$ will correspond to an atom of weight $w_0 = (2c - 1)/2c$. Thus the limiting eigenvalue distribution function has density

$$f_B(x) = \max\left(\frac{2c-1}{2c}, 0\right) \delta(x) + \frac{1}{\pi} \frac{\sqrt{(x-a_-)(a_+ - x)}}{2xc - 2cx^2} I_{[a_-, a_+]} + \max\left(\frac{2c-1}{2c}, 0\right) \delta(x-1),$$

(9.9)

where $a_{\pm} = 1/2 \pm \sqrt{-c^2 + c}$. Figure 9.2 compares the theoretical prediction in (9.9) with a Monte-Carlo experiment for $c = 0.4$.




Figure 9-3. Additive convolution of equilibrium measures corresponding to potentials $V_1(x)$ and $V_2(x)$.

From the associated bivariate polynomial

$$L_{\mu z}^B \equiv (-2c + 2cz)\mu^2 + (z - 2 - 2cz + 4c)\mu - 2c + 2,$$

we obtain two series expansions whose branches tend to zero. The first four terms of the series are given by

$$1 + \frac{1}{2}z + \frac{1+c}{4}z^2 + \frac{3+c}{8}z^3 + O(z^4), \quad (9.10)$$

and,

$$\frac{c-1}{c} + \frac{c-1}{2c}z - \frac{(c-1)(-2+c)}{4c}z^2 - \frac{(c-1)(3c-4)}{8c}z^3 + O(z^4), \quad (9.11)$$

respectively. Since $c \leq 1$, the series expansion in (9.11) can be eliminated since $\mu(0) := \int dF^B(x) = 1$. Thus the coefficients of the series in (9.10) are the correct moments of the limiting eigenvalue distribution. A recursion for the moments can be readily derived using the techniques developed earlier.

■ 9.3 Free additive convolution of equilibrium measures

Equilibrium measures are a fascinating topic within random matrix theory. They arise in the context of research that examines why very general random models for random matrices exhibit universal behavior in the large matrix limit. Suppose we are given a potential $V(x)$ then we consider a sequence of Hermitian, unitarily invariant random matrices \mathbf{A}_N , the joint distribution of whose elements is of the form

$$P(\mathbf{A}_N) \propto \exp(-N \operatorname{Tr} V(\mathbf{A}_N)) d\mathbf{A}_N,$$

where $d\mathbf{A}_N = \prod_{i \leq j} (d\mathbf{A}_N)_{ij}$. The equilibrium measure, when it exists, is the unique probability distribution function that minimizes the logarithmic energy (see [29] for additional details). The resulting equilibrium measure depends explicitly on the potential $V(x)$ and can be explicitly computed for some potentials. In particular, for potentials of the form $V(x) = t x^{2m}$, the Stieltjes transform of the resulting equilibrium measure is an algebraic function [29, Chp. 6.7, pp. 174-175] so that the equilibrium measure is an algebraic distribution. Hence we can formally investigate the additive convolution of equilibrium measures corresponding to two different potentials. For $V_1(x) = x^2$, the equilibrium measure is the (scaled) semi-circle distribution encoded by the bivariate polynomial

$$L_{mz}^A \equiv m^2 + 2 m z + 2.$$

For $V_2(x) = x^4$, the equilibrium measure is encoded by the bivariate polynomial

$$L_{mz}^B \equiv 1/4 m^2 + m z^3 + z^2 + 2/9 \sqrt{3}.$$

Since \mathbf{A}_N and \mathbf{B}_N are unitarily invariant random matrices, if \mathbf{A}_N and \mathbf{B}_N are independent, then the limiting eigenvalue distribution function of $\mathbf{C}_N = \mathbf{A}_N + \mathbf{B}_N$ can be computed from L_{mz}^A and L_{mz}^B . The limiting eigenvalue density function $f_C(x)$ is the free additive convolution of f_A and f_B . The MATLAB command `LmzC = AplusB(LmzA, LmzB)`; will produce the bivariate polynomial

$$L_{mz}^C = -9 m^4 - 54 m^3 z + (-108 z^2 - 36) m^2 - (72 z^3 + 72 z) m - 72 z^2 - 16 \sqrt{3}.$$

Figure 9.3 plots the probability density function for the equilibrium measure for the potentials $V_1(x) = x^2$ and $V_2(x) = x^4$ as well as the free additive convolution of these measures. The interpretation of the resulting measuring in the context of potential theory is not clear. The matrix \mathbf{C}_N will no longer be unitarily invariant so it might not sense to look for a potential $V_3(x)$ for which F^C is an equilibrium measure. The tools and techniques developed in this article might prove useful in further explorations.

■ 9.4 Other applications

There is often a connection between well-known combinatorial numbers and random matrices. For example, the even moments of the Wigner matrix are the famous Catalan numbers. Similarly, if $\mathbf{W}_N(c)$ denotes the Wishart matrix with parameter c , other combinatorial correspondences can be easily established using the techniques developed. For instance, the limiting moments of $\mathbf{W}_N(1) - \mathbf{I}_N$ are the Riordan numbers, the large Schröder numbers correspond to the limiting moments of $2\mathbf{W}_N(0.5)$ while the small Schröder numbers are the limiting moments of $4\mathbf{W}_N(0.125)$. Combinatorial identities along the lines of those developed in [33] might result from these correspondences.

The polynomial method: Eigenvectors of random matrices

Proposition 8.41 succinctly captures an important connection between free probability and random matrices. Specifically, free probability provides the analytic machinery for computing the limiting eigenvalue distribution of $\mathbf{A}_N + \mathbf{B}_N$ and $\mathbf{A}_N \mathbf{B}_N$ from the limiting eigenvalue distribution of \mathbf{A}_N and \mathbf{B}_N when they are asymptotically free. A less well-known fact is that it also provides us with a machinery for computing the limiting conditional “eigenvector distribution” of the eigenvectors of $\mathbf{A}_N + \mathbf{B}_N$.

Note that if \mathbf{B}_N is small (in some appropriate norm sense) compared to \mathbf{A}_N , then the eigenvectors of $\mathbf{A}_N + \mathbf{B}_N$ should be close to those of \mathbf{A}_N so that standard perturbation theory as in [96] should be able to adequately describe the transformation in the eigenvectors. The power of the free probabilistic framework is that it makes no assumptions on the relative norms of \mathbf{A}_N and \mathbf{B}_N except that their limiting eigenvalue distributions exist. The machinery for analytically characterizing the eigenvectors was developed by Biane in [15] in the context of his investigation of processes with free increments. The applicability of these results for describing the conditional “eigenvector distribution” is mentioned in [16, pp. 70].

In this chapter, we summarize Biane’s relevant results from [15], and define the subclass of algebraic random matrices for which the conditional “eigenvector distribution” is algebraic as well. As before, algebraicity of this subclass acts as a certificate of the computability of the limiting conditional “eigenvector distribution.”

■ 10.1 The conditional “eigenvector distribution”

Consider the random matrices $\mathbf{A} \equiv \mathbf{A}_N$ and $\mathbf{B} \equiv \mathbf{B}_N$ with limiting eigenvalue distribution functions given by F^A and F^B , respectively. Let $\mathbf{u}_1, \dots, \mathbf{u}_N$ and $\mathbf{v}_1, \dots, \mathbf{v}_N$ be the eigenvectors of \mathbf{A}_N and $\mathbf{A}_N + \mathbf{B}_N$, associated with the eigenvalues $\lambda_1^A, \dots, \lambda_N^A$ and $\lambda_1^{A+B}, \dots, \lambda_N^{A+B}$, respectively.

The passage from the old basis to the new basis is given by the $N \times N$ sized transition matrix whose (i, j) -th entry is the projection $\langle \mathbf{v}_i, \mathbf{u}_j \rangle$. Since the eigenvectors are only defined up to some complex number with modulus one, Biane considers the numbers $|\langle \mathbf{v}_i, \mathbf{u}_j \rangle|^2$, which form a bistochastic matrix.

While it is not meaningful to speak of the limit of the entries of the bistochastic matrix itself, it does make sense to ask if the entries have some definite asymptotic behavior as $N \rightarrow \infty$. Let g and h be smooth function on \mathbb{R} and consider the asymptotic behavior of the expression

$$\frac{1}{N} \sum \text{Tr} (g(\mathbf{A}_N)h(\mathbf{A}_N + \mathbf{B}_N)) = \sum_{1 \leq i, j \leq N} h(\lambda_i^{\mathbf{A}+\mathbf{B}})g(\lambda_j^{\mathbf{A}})|\langle \mathbf{v}_i, \mathbf{u}_j \rangle|^2. \quad (10.1)$$

In what follows it is established how the conditional ‘‘eigenvector distribution’’ is encoded by a Markov transition kernel density function.

Proposition 10.11. *Let \mathbf{A}_N and \mathbf{B}_N be asymptotically free sequences of random matrices that satisfy the hypotheses in Proposition 8.41. Let g and h be smooth functions on \mathbb{R} . If \mathbf{A}_N and \mathbf{B}_N are chosen at random, then*

$$\frac{1}{N} \text{Tr} (h(\mathbf{A}_N + \mathbf{B}_N)g(\mathbf{A}_N)) \rightarrow \int g(x)h(y)\rho_{A+B}(x, y)dx dy, \quad (10.2a)$$

$$\frac{1}{N} \text{Tr} \left(h(\mathbf{A}_N^{1/2}\mathbf{B}_N\mathbf{A}_N^{1/2})g(\mathbf{A}_N) \right) \rightarrow \int g(x)h(y)\rho_{AB}(x, y)dx dy, \quad (10.2b)$$

where the convergence is in probability as $N \rightarrow \infty$ and ρ_{A+B} and ρ_{AB} are bivariate probability density functions on \mathbb{R}^2 that can be decomposed as

$$\rho_{A+B}(x, y) = k_{A+B|A}(x, y)f_A(x) \quad (10.3a)$$

$$\rho_{AB}(x, y) = k_{AB|A}(x, y)f_A(x), \quad (10.3b)$$

where $f_A := dF^A(x)$ is the limiting eigenvalue density function of \mathbf{A}_N and $k_{A+B}(x, y)$ and $k_{AB}(x, y)$ are Markov transition kernel density functions.

PROOF. This result appears in Biane [15] in the context of processes with free increments. The connection with eigenvectors is mentioned in [16, pp. 70]. \square

The Markov transition kernels obtained may be intuitively thought of as the limit of the bistochastic matrix $|\langle \mathbf{v}_i, \mathbf{u}_j \rangle|^2$ that appears on the right hand side of (10.1). The propositions that follow describe the procedure for computing these Markov transition kernels.

Proposition 10.12. *Let $k_{A+B|A}(x, y)$ be the Markov transition kernel density function as defined in Proposition 10.11. Then $k_{A+B|A}(x, y)$ is the probability density function on $\mathbb{R} \times \mathbb{R}$ with support $\mathcal{S}_A \times \mathcal{S}_{AB}$ associated with the analytic function q defined on $\mathbb{C}^- \setminus \mathbb{R}$. Both are uniquely determined by the relations*

$$G_{A+B|A}(x, y) = \int \frac{1}{y-z} k_{A+B|A}(x, z) dz \quad (10.4a)$$

$$G_{A+B|A}(x, y) = \frac{1}{q(y) - x} \quad (10.4b)$$

$$g_A(q(y)) = g_{A+B}(y) \quad (10.4c)$$

for all $z \in \mathbb{C}^- \setminus \mathbb{R}$.

PROOF. These relations are derived in [15, pp. 151–153]. \square

Proposition 10.13. *Let $k_{AB|A}(x, y)$ be the Markov transition kernel as defined in Proposition 10.11. Then $k_{AB|A}(x, y)$ is the probability density function on $\mathbb{R} \times \mathbb{R}$ with support $\mathcal{S}_A \times \mathcal{S}_{AB}$ associated with the analytic function q defined on $\mathbb{C}^- \setminus \mathbb{R}$. Both are uniquely determined by the relations*

$$G_{AB|A}(x, y) = \int \frac{1}{y - z} k_{AB|A}(x, z) dz \quad (10.5a)$$

$$G_{AB|A}(x, y) = \frac{1}{y - q(1/y)xy} \quad (10.5b)$$

$$\frac{1}{q(1/y)} g_A \left(\frac{1}{q(1/y)} \right) = y g_{AB}(y) \quad (10.5c)$$

for all $z \in \mathbb{C}^- \setminus \mathbb{R}$.

PROOF. These relations are derived in [15, pp. 158]. \square

■ 10.2 Algebraic conditional "eigenvector distributions"

A closer inspection of the analytical procedures, described in Propositions 10.12 and 10.13, for computing the Markov transition kernels $k_{A+B|A}$ and $k_{AB|A}$ reveals the difficulty of concretely computing these kernels. Specifically, the conditions in (10.4c) and (10.5c) will be satisfied by a function q that can be expressed in closed form in only some special cases. However, when the probability density functions f_A and f_{A+B} (or f_{AB}) are algebraic so that we can encode their Stieltjes (or Cauchy) transform as a solution of a bivariate polynomial equation, the Markov transition kernels can be readily computed.

Remark 10.21 (Terminology). *We shall often informally use the phrase conditional "eigenvector distribution" when referring to the Markov transition kernel that emerges from Proposition 10.11. The phrase "eigenvector distribution" is enclosed in quotes because the kernel characterization is not a distribution in the usual sense of the word, i.e., it does not describe the probability distribution of the eigenvectors of $\mathbf{A}_N + \mathbf{B}_N$. It is qualified by affixing the label conditional because, in the sense of Proposition 10.11, it encodes how the eigenvectors of $\mathbf{A}_N + \mathbf{B}_N$ (or $\mathbf{A}_N^{1/2} \mathbf{B}_N \mathbf{A}_N^{1/2}$) are related to the eigenvectors of \mathbf{A}_N .*

Notation 10.22 (Trivariate polynomial). Let L_{uvw} denote a trivariate polynomial of degree D_u in u , D_v in v and D_w in w given by

$$L_{uvw} \equiv L_{uvw}(\cdot, \cdot, \cdot) = \sum_{i=0}^{D_u} \sum_{j=0}^{D_v} \sum_{k=0}^{D_w} c_{ijk} u^i v^j w^k. \quad (10.6)$$

The scalar coefficients c_{ijk} are real valued.

Definition 10.23 (Algebraic Markov transition kernels). Let $k(x, y)$ be a Markov transition kernel density function. Consider the analytic function $G(x, y)$ defined on $\mathbb{C}^- \setminus \mathbb{R} \times \mathbb{C}^- \setminus \mathbb{R}$ as

$$G(x, y) = \int \frac{1}{y-z} k(x, z) dz. \quad (10.7)$$

If there exists a trivariate polynomial L_{Gxy} such that $L_{Gxy}(G(x, y), x, y) = 0$ then we refer to as $k(x, y)$ as an algebraic Markov transition kernel and say that $k(x, y) \in \mathcal{K}_{\text{alg}}$. Here \mathcal{K}_{alg} denotes the class of algebraic Markov transition kernels.

Remark 10.24 (Equivalent representation). Let $k(x, y)$ be a Markov transition kernel density function. Consider the analytic function $M(x, y)$ defined on $\mathbb{C}^+ \setminus \mathbb{R} \times \mathbb{C}^+ \setminus \mathbb{R}$ as

$$M(x, y) = \int \frac{1}{z-y} k(x, z) dz. \quad (10.8)$$

The function $M(x, y)$ is related to $G(x, y)$, defined in (10.7), as $M(x, y) = -G(x, y)$. If $k(x, y) \in \mathcal{K}_{\text{alg}}$ then L_{Gxy} exists so that L_{Mxy} exists and is given by

$$L_{Mxy}(M, x, y) = L_{Gxy}(-M, x, y). \quad (10.9)$$

Remark 10.25 (Property of Markov transition kernels). Let $k(x, y)$ be a Markov transition kernel with support $\mathcal{S}_x \times \mathcal{S}_y$. Then, by definition, for every $x_0 \in \mathcal{S}_x$, $k(x_0, y)$ is a positive probability density function on \mathcal{S}_y and $M(x_0, y)$ is its Stieltjes transform. Similarly, for every $y_0 \in \mathcal{S}_y$, $k(x, y_0)$ is a positive probability density function with support on \mathcal{S}_x and $M(x, y_0)$ is its Stieltjes transform.

Remark 10.26 (Property of algebraic Markov transition kernels). If $k(x, y) \in \mathcal{K}_{\text{alg}}$ with support $\mathcal{S}_x \times \mathcal{S}_y$ then for every $x_0 \in \mathcal{S}_x$ and $y_0 \in \mathcal{S}_y$, it follows that $k(x_0, y) \in \mathcal{P}_{\text{alg}}$ and $k(x, y_0) \in \mathcal{P}_{\text{alg}}$.

The main result, stated below, is that the Markov transition kernel that emerges when characterizing the eigenvectors of the sum and product of asymptotically free algebraic random matrices is algebraic as well. The value of this statement is that when combined with Remark 10.26 it allows us to concretely compute the Markov transition kernel numerically using the techniques discussed in Chapter 8.

Theorem 10.27. *Let \mathbf{A}_N and \mathbf{B}_N be asymptotically free random matrices that satisfy the hypothesis in Proposition 8.41. If $f_A \in \mathcal{P}_{\text{alg}}$ and $f_B \in \mathcal{P}_{\text{alg}}$ and $k_{A+B|A}$ is a Markov transition kernel as defined in Proposition 10.11 then $k_{A+B|A} \in \mathcal{K}_{\text{alg}}$.*

PROOF. Since $f_A \in \mathcal{P}_{\text{alg}}$ and $f_B \in \mathcal{P}_{\text{alg}}$, from Corollary 8.42, $f_{A+B} \in \mathcal{P}_{\text{alg}}$. Hence L_{gz}^A and L_{gz}^{A+B} exist. Equation (10.4c) implies that the set of polynomial equations $L_{\text{gz}}^A(g, q) = 0$ and $L_{\text{gz}}^{A+B}(g, y) = 0$ share a common solution. Hence the resultant, given by Definition 6.41, of the polynomials will equal zero, *i.e.*, $L_{\text{qy}}^{A+B|A}(q, y) = 0$ where

$$L_{\text{qy}}^{A+B|A}(q, y) = \text{Res}_g (L_{\text{gz}}^A(g, q), L_{\text{gz}}^{A+B}(g, y)). \quad (10.10)$$

Equation (10.4b) yields the relationship

$$q(y) = x + \frac{1}{G(x, y)}. \quad (10.11)$$

Thus $G(x, y)$ is a solution of the trivariate polynomial equation $L_{\text{Gxy}}^{A+B|A}(G, x, y) = 0$ where

$$L_{\text{Gxy}}^{A+B|A}(G, x, y) = L_{\text{qy}}^{A+B|A} \left(x + \frac{1}{G}, y \right) \quad (10.12)$$

is the polynomial obtained by clearing the denominator or, equivalently, multiplying the right hand side by G^{D_q} where D_q is the degree of q in $L_{\text{qy}}^{A+B|A}$. The trivariate polynomial $L_{\text{Gxy}}^{A+B|A}$ thus obtained proves that $k_{A+B|A}(x, y) \in \mathcal{K}_{\text{alg}}$. \square

Theorem 10.28. *Let \mathbf{A}_N and \mathbf{B}_N be asymptotically free random matrices that satisfy the hypothesis in Proposition 8.41. If $f_A \in \mathcal{P}_{\text{alg}}$ and $f_B \in \mathcal{P}_{\text{alg}}$ and $k_{AB|A}$ is a Markov transition kernel as defined in Proposition 10.11 then $k_{AB|A} \in \mathcal{K}_{\text{alg}}$.*

PROOF. Since $f_A \in \mathcal{P}_{\text{alg}}$ and $f_B \in \mathcal{P}_{\text{alg}}$, from Corollary 8.43, $f_{AB} \in \mathcal{P}_{\text{alg}}$. Hence L_{gz}^A and L_{gz}^{AB} exist. The function $\tilde{q}(y) := q(1/y)$ given by the relation (10.5c) is an algebraic function, *i.e.*, it satisfies the algebraic equation $L_{\text{qy}}^{AB|A}(\tilde{q}, y) = 0$. The bivariate polynomial $L_{\text{qy}}^{AB|A}$ is obtained as follows. First we obtain the bivariate polynomial $L_{\text{gq}}^A(g, \tilde{q})$ given by

$$L_{\text{gq}}^A(g, \tilde{q}) = \tilde{q}^{D_z^A} L_{\text{gz}}^A(g\tilde{q}, 1/\tilde{q}). \quad (10.13)$$

where D_z^A is the degree of z in the polynomial L_{gz}^A . Equation (10.5c) implies that the set of polynomial equations $L_{\text{gq}}^A(g, \tilde{q}) = 0$ and $L_{\text{gz}}^{AB}(gy, y) = 0$ share a common solution. Hence the resultant, given by Definition 6.41, of the polynomials will equal zero, *i.e.*,

$L_{\tilde{q}y}^{AB|A}(\tilde{q}, y) = 0$ where

$$L_{\tilde{q}y}^{AB|A}(\tilde{q}, y) = \text{Res}_g \left(L_{g\tilde{q}}^A(g, \tilde{q}), y^{D_g^{AB}} L_{gz}^{AB}(gy, y) \right) \quad (10.14)$$

and D_g^{AB} is the degree g in the polynomial L_{gz}^{AB} . Equation (10.5b) yields the relationship

$$\tilde{q}(y) = h(1/y) = \frac{1}{x} - \frac{1}{xyG(x, y)}. \quad (10.15)$$

Hence, $G(x, y)$ satisfies the trivariate polynomial equation $L_{Gxy}^{AB|A}(G, x, y) = 0$ where

$$L_{Gxy}^{AB|A}(G, x, y) = L_{\tilde{q}y}^{AB|A} \left(\frac{1}{x} - \frac{1}{Gxy}, y \right) \quad (10.16)$$

is the polynomial obtained by clearing the denominator of the rational function obtained, or equivalently multiplying the right hand side by $(Gxy)^{D_{\tilde{q}}}$ where $D_{\tilde{q}}$ is the degree of \tilde{q} in $L_{\tilde{q}y}^{AB|A}$. The trivariate polynomial $L_{Gxy}^{AB|A}$ thus be obtained proves that $k_{A+B|A}(x, y) \in \mathcal{K}_{\text{alg}}$. \square

Corollary 10.29. *Let \mathbf{A}_N and \mathbf{B}_N be asymptotically free algebraic random matrices. Then*

- $k_{A+B|B} \in \mathcal{K}_{\text{alg}}$ and $k_{A+B|A} \in \mathcal{K}_{\text{alg}}$,
- $k_{AB|B} \in \mathcal{K}_{\text{alg}}$ and $k_{AB|A} \in \mathcal{K}_{\text{alg}}$.

PROOF. This first part of the statement follows directly from Theorem 10.27. The kernels $k_{A+B|B}$ and $k_{A+B|A}$ will generically be different unless $f_A = f_B$ almost everywhere. The second part of the statement directly from Theorem 10.28. The kernels $k_{AB|B}$ and $k_{AB|A}$ will generically be different unless $f_A = f_B$ almost everywhere. \square

The proofs of Theorems 10.27 and 10.28 reveal the symbolic code need to compute the trivariate polynomial that encodes the Markov transition kernel. These are listed in Table 10.1. This allow us to identify the subclass of algebraic random matrices for which the conditional “eigenevector distribution” is algebraic, in the sense of having algebraic Markov transition kernels.

Theorem 10.210. *Sums and (admissible) products of asymptotically free algebraic random matrices have conditional “eigenevector distributions” that are encoded by algebraic Markov transition kernels.*

MATLAB Code
<pre>function [LmxyApB,LmzApB] = AplusBkernel(LmzA,LmzB) syms m g q x y z LgzA = Lmz2Lgz(LmzA); LmzApB = AplusB(LmzA,LmzB); LgzApB = Lmz2Lgz(LmzApB); LgqA = subs(LgzA,z,q); LgyApB = subs(Lgz,z,y); LqyApB = maple('resultant',LgqA,LgyApB,g); LgxyApB = subs(LqyApB,q,x+1/g); LgxyApB = irreducLuv(LgxyApB,g,x); LmxyApB = subs(LgxyApB,g,-m);</pre>

(a) $L_{mz}^A, L_{mz}^B \mapsto L_{mxy}^{A+B|A}$ for $A, B \mapsto A + QBQ'$.

MATLAB Code
<pre>function [LmxyAtb,LmzAtb] = AtimesBkernel(LmzA,LmzB) syms m g q x y z LgzA = Lmz2Lgz(LmzA); LmzAtb = AtimesB(LmzA,LmzB); LgzAtB = Lmz2Lgz(LmzAtb); LgqA = irreducLuv(subs(LgzA,{g,z},{g*q,1/q}),g,q); LgyAtB = irreducLuv(subs(LgzAtB,{g,z},{g/y,y}),g,y); LqyAtB = maple('resultant',LgqA,LgyAtB,g); LgxyAtB = subs(LqyAtB,q,1/x-1/(y*g*x)); LgxyAtb = irreducLuv(LgxyAtB,g,y); LmxyAtB = subs(LgxyAtB,g,-m);</pre>

(b) $L_{mz}^A, L_{mz}^B \mapsto L_{mxy}^{AB|A}$ for $A, B \mapsto A \times QBQ'$.

Table 10.1. Symbolic code in MATLAB for computing the trivariate polynomial that encodes the Markov transition kernel that characterizes the conditional "eigenvector distribution" of sums and products of algebraic random matrices.

■ 10.3 Example

Suppose \mathbf{A}_N is the $N \times N$ matrix

$$\mathbf{A}_N = \begin{bmatrix} 5\mathbf{I}_{N/2} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{N/2} \end{bmatrix} \quad (10.17)$$

where we have assumed without loss of generality, for the eigenvector discussion to follow, that \mathbf{A}_N is a diagonal matrix. There are two eigenspaces associated with this matrix. We think of the eigenspace associated with the eigenvalue equal to 5 as being the “signal” subspace. Consider the Wishart matrix constructed from an $N \times L$ random matrix \mathbf{G}_N with standard normal entries as

$$\mathbf{W}_N = \frac{1}{L} \mathbf{G}_N \mathbf{G}_N^H. \quad (10.18)$$

We employ the machinery developed to describe the eigenvectors of the matrix

$$\mathbf{C}_N = \mathbf{A}_N + e \mathbf{W}_N. \quad (10.19)$$

as a function of e (a mnemonic for ϵ) in the $N \rightarrow \infty$ limit when $N = 2L$. Note that this choice of N and L makes \mathbf{W}_N singular with rank L (with high probability). Both \mathbf{A}_N and \mathbf{W}_N are algebraic random matrices. The limiting eigenvalue distribution function of \mathbf{A}_N has Stieltjes transform

$$m_A(z) = \frac{0.5}{5-z} + \frac{0.5}{1-z},$$

which is the solution of the equation $L_{mz}^A(m, z) = 0$ where

$$L_{mz}^A(m, z) = m(5-z)(1-z) - (3-z). \quad (10.20)$$

The limiting eigenvalue distribution function of \mathbf{W}_N has Stieltjes transform that is the solution of the equation $L_{mz}^W(m, z) = 0$ where

$$L_{mz}^W(m, z) = 2zm^2 + (1+z)m + 1, \quad (10.21)$$

is obtained by plugging $c = N/L = 2$ into the appropriate polynomial in Table 6.2(b). Let $\mathbf{B}_N = e \mathbf{W}_N$. By Corollary 7.32, \mathbf{B}_N is algebraic. Since \mathbf{W}_N has Haar distributed eigenvectors, it is orthogonally invariant. This makes \mathbf{B}_N asymptotically free with respect to \mathbf{A}_N so that by Corollary 7.56 and Theorem 10.210, $\mathbf{C}_N = \mathbf{A}_N + \mathbf{B}_N$ has an algebraic conditional “eigenvector distribution.”

To predict the distortion in the “signal” subspace of \mathbf{A}_N when additively perturbed by \mathbf{B}_N we compute $k_{A+B|A}(x, y)$ and evaluate the Markov transition kernel density function at $x = 1$ and $x = 5$.

Using the MATLAB tools developed, the trivariate polynomial that encodes the Markov transition kernel is computed using the sequence of commands:

```
>> syms m x y z e
>> LmzA = m - 0.5/(5-z) - 0.5/(1-z);
>> LmzA = irreducLuv(LmzA,m,z);
>> LmzW = 2*z*m^2+(1+z)*m+1;
>> LmzB = scaleA(LmzW,e);
>> [LmxC,LmzC] = AplusBkernel(LmzA,LmzB);
```

from which we obtain the bivariate polynomial

$$L_{mz}^C(m, z) = (20e^2 + 4e^2z^2 - 24e^2z)m^3 + (-24ze + 20e + 4ez^2)m^2 + (5 + z^2 - e^2 - 6z + 2ze - 6e)m - 3 - e + z \quad (10.22)$$

and the trivariate polynomial

$$L_{Mxy}^{C|A}(m, x, y) = (-5x - x^3 - 5e + 6x^2 + 6ye - 6zx + ex^2 + yx^2 + 5y - 2yex)m^3 + (6y + 2ye - 2yx + 3x^2 + 5 - 12x - 2ex)m^2 + (6 + e + y - 3x)m + 1 \quad (10.23)$$

Figure 10-1 compares the density function associated with limiting eigenvalue distribution of \mathbf{C}_N for different values of e . These curves were computed using the techniques described in Section 8.2 from the bivariate polynomial L_{mz}^C . The curves reveal the extent of the distortion in the eigen-spectrum of \mathbf{A}_N induced by the low rank perturbation $\mathbf{B}_N = e\mathbf{W}_N$. As $e \rightarrow 0$, the distortion lessens and the limiting eigenvalue distribution of \mathbf{C}_N will resemble that of \mathbf{A}_N .

The Stieltjes transform of the density function $k(1, y)$ is the solution of the algebraic equation $L_{Mxy}^{C|A}(m, 1, y) = 0$, *i.e.*,

$$(-4e + 4ye)m^3 + (4y + 2ye - 4 - 2e)m^2 + (3 + e + y)m + 1 = 0 \quad (10.24)$$

We can compute the density and the moments of $k(1, y)$ from (10.24) by using the techniques described in Chapter 8. Its first 4 moments are, respectively

$$\begin{bmatrix} 1 + e \\ 1 + 2e + 3e^2 \\ 1 + 3e + 13e^2 + 11e^3 \\ 1 + 4e + 76e^3 + 50e^2 + 45e^4 \end{bmatrix}.$$

Similarly, the Stieltjes transform of the density function $k(5, y)$ is a solution of the algebraic equation $L_{Mxy}^{C|A}(m, 5, y) = 0$, *i.e.*,

$$(20e - 4ye)m^3 + (-4y + 2ye + 20 - 10e)m^2 + (-9 + e + y)m + 1 = 0. \quad (10.25)$$




Figure 10-1. Limiting eigenvalue density function of $\mathbf{A}_N + \epsilon \mathbf{W}_N$ for different values of ϵ .

We can similarly recover the density function $k(5, y)$ and its moments from (10.25). From Proposition 10.11, the relationship between the limiting eigenvalue distribution of $\mathbf{C} = \mathbf{A} + \mathbf{B}$ and the kernel density functions $k(1, y)$ and $k(5, y)$ can be deduced. In general, for $k(x, y) \equiv k_{\mathbf{A}+\mathbf{B}|\mathbf{A}}(x, y)$ the relationship

$$dF^{\mathbf{A}+\mathbf{B}}(y) = \int k(x, y) dF^{\mathbf{A}}(x), \quad (10.26)$$

reduces, in our case to

$$dF^{\mathbf{A}+\mathbf{B}} = 0.5 k(1, y) + 0.5 k(5, y). \quad (10.27)$$

Figure 10-2(a) illustrates the relationship in (10.27). Figure 10-2(b) compares the function $0.5 k(5, x)$ with the *weighted* empirical histogram of the eigenvalues of \mathbf{C}_N , collected over 4000 trials with $N = 100 = 2L$. The weight used to compute the histogram is the norm square of the projection of the each eigenvector of \mathbf{C}_N onto the “signal” subspace, *i.e.*, for \mathbf{A}_N given by (10.17), the $N \times N/2$ projection matrix with ones along the diagonal and zeros elsewhere.

Figure 10-2(a) provides insight into the how the eigenvectors of the $\mathbf{A} + 2\mathbf{W}$ are related to the eigenvectors of \mathbf{A} . For a large enough \mathbf{A} , suppose we obtain an eigenvalue of magnitude approximately 4.25 (where the curve representing $0.5 k(5, \cdot)$ intersects with the curve representing $0.5 k(1, \cdot)$). What Figure 10-2(a) conveys is that *the corre-*

sponding eigenvector of $\mathbf{A} + 2\mathbf{W}$ will have a projection of equal norm onto each of the eigenspaces of \mathbf{A} .

In other words, conditioned on an eigenvalue of $\mathbf{A} + 2\mathbf{W}$ having magnitude z , the projection of the corresponding eigenvector onto the eigenspace of \mathbf{A} spanned by the eigenvalue of magnitude 5 will have square norm that will be very well approximated (for large N) by the expression

$$\frac{0.5 k(5, z)}{0.5 k(1, z) + 0.5 k(5, z)}.$$


Figure 10-3(a) compares this expression for values of z in the support of the limiting distribution (shown in Figure 10-3(b)) with the norm square of the projection of the sample eigenvectors of a *single realization* of \mathbf{C}_N formed with $N = 100 = 2L$ onto the subspace spanned by the eigenvalues of magnitude 5 in \mathbf{A} in (10.17). It is clear that despite the predictions being asymptotic in nature, they accurately predict the behavior for finite sized matrices as well.

The experiments and the theory capture many interesting features about the behavior of sample eigenvectors:


1. If we consider the square norm of the projection onto each of the subspaces to be a “reliability metric,” then it is immediately apparent that all the sample eigenvectors are not equally reliable.
2. The behavior of the eigenvectors corresponding to the smallest and the largest eigenvalues is very different. In fact, the middle eigenvectors have the greatest projection on the eigenspace of \mathbf{A} spanned by the eigenvalue of magnitude 5.

Taken together, the computational tools developed allow us to use the machinery of free probability to *analytically describe the deterioration in the “reliability” of the sample eigenvectors* induced by the additive random moderate rank subspace perturbation. Tools from “classical” perturbation theory [96] would have been inadequate in the scenario considered because the matrix norm of \mathbf{B} is comparable to the matrix norm of \mathbf{A} and the perturbation in question is not, by any stretch of the imagination, of low rank.

Figure 10.3 plots the kernel density function $k_{A+B|B}(x, y)$. The fact that the kernel density function is a true bivariate probability density function over \mathbb{R}^2 follows from Remark 10.25.




(a) The density function dF^{A+B} , and the scaled kernel density functions $0.5k_{A+B|A}(1, \cdot)$ and $0.5k_{A+B|A}(5, \cdot)$.




(b) Empirical validation of the theoretical scaled kernel density function $0.5k_{A+B|A}(5, \cdot)$.

Figure 10-2. The composite limit eigenvalue density function is interpreted as the sum of the scaled individual kernel density functions.



(a) Norm square of the projection of the eigenvectors of C_N onto the signal subspace of A_N .



(b) Limiting eigenvalue density function of C_N .

Figure 10-3. Characterizing the relationship between the eigenvectors of $A + 2W$ with those of A .




Figure 10-4. The Markov transition kernel density $k_{A+B|B}$ where $B = 2W$.

■ 10.4 Algebraic empirical covariance matrices

We conclude by applying this machinery to predict the deterioration of the eigenvectors of empirical covariance matrices due to sample size constraints. The (broader) class of algebraic Wishart sample covariance matrices for which this framework applies is described next.

Theorem 10.41. *Let $\mathbf{A}_n \xrightarrow{p} A \in \mathcal{M}_{\text{alg}}$, and $\mathbf{B}_N \xrightarrow{p} B \in \mathcal{M}_{\text{alg}}$ be algebraic covariance matrices with $\mathbf{G}_{n,N}$ denoting an $n \times N$ (pure) Gaussian random matrix (see Definition 7.41). Let $\mathbf{X}_{n,N} = \mathbf{A}_n^{1/2} \mathbf{G}_{n,N} \mathbf{B}_N^{1/2}$. Then*

$$\mathbf{S}_n = \mathbf{X}_{n,N} \mathbf{X}'_{n,N} \xrightarrow{p} S \in \mathcal{M}_{\text{alg}} \quad \text{and} \quad k_{S|A}(x, y) \in \mathcal{K}_{\text{alg}},$$

as $n, N \rightarrow \infty$ and $c_N = n/N \rightarrow c$.

PROOF. Let $\mathbf{Y}_{n,N} \equiv \mathbf{G}_{n,N} \mathbf{B}_N^{1/2}$, $\mathbf{T}_n \equiv \mathbf{Y}_{n,N} \mathbf{Y}'_{n,N}$ and $\tilde{\mathbf{T}}_N = \mathbf{Y}'_{n,N} \mathbf{Y}_{n,N}$. Thus $\mathbf{S}_n = \mathbf{A}_n \times \mathbf{T}_n \equiv \mathbf{A}_n^{1/2} \mathbf{T}_n \mathbf{A}_n^{1/2}$. The matrix \mathbf{T}_n , as defined, is invariant under orthogonal/unitary transformations, though the matrix $\tilde{\mathbf{T}}_N$ is not. Hence, by Corollary 7.56, and since $\mathbf{A}_n \mapsto A \in \mathcal{M}_{\text{alg}}$, $\mathbf{S}_n \mapsto S \in \mathcal{M}_{\text{alg}}$ whenever $\mathbf{T}_n \mapsto T \in \mathcal{M}_{\text{alg}}$. From Theorem 7.36, $\mathbf{T}_n \mapsto T \in \mathcal{M}_{\text{alg}}$ if $\tilde{\mathbf{T}}_N \mapsto \tilde{T} \in \mathcal{M}_{\text{alg}}$. The matrix $\tilde{\mathbf{T}}_N = \mathbf{B}_N^{1/2} \mathbf{G}'_{n,N} \mathbf{G}_{n,N} \mathbf{B}_N^{1/2}$ is clearly algebraic by application of Corollary 7.56 and Theorem 7.31 since \mathbf{B}_N is algebraic and $\mathbf{G}'_{n,N} \mathbf{G}_{n,N}$ is algebraic and unitarily invariant.

From Theorem 10.210, since \mathbf{T}_n and \mathbf{A}_n are algebraic, the conditional “eigenvector distribution” of $\mathbf{S}_n = \mathbf{A}_n \times \mathbf{T}_n$ is algebraic. This proves that $k_{S|A}(x, y) \in \mathcal{K}_{\text{alg}}$. \square

The theorem can be restated more succinctly.

Corollary 10.42. *Algebraic sample covariance matrices with Wishart distribution have limiting eigenvalue and conditional “eigenvector distributions” that are algebraic.*

In high-dimensional inference applications, n is often interpreted as the number of variables (spatial dimension) while N is the number of measurements (temporal dimension). The matrices \mathbf{A}_n and \mathbf{B}_N then model the spatial and temporal covariance structure of the collected data. The parameter $c_N = n/N$ is the ratio of the number of variables to the number of measurements. In a sample size constrained setting, we expect c_N to be significantly greater than zero.

The proof of Theorem 10.41 provides us with a recipe for computing the polynomials that encode the limiting eigenvalue and conditional eigenvector distributions of \mathbf{S} in the far more general situation where the observation vectors are modelled as samples of a

MATLAB Code
<pre>function [LmxyS,LmzS] = AtimesWishtimesB(LmzA,LmzB,c) syms m x y z LmzW = c*z*m^2-(1-c-z)*m+1; LmzWt = transposeA(LmzW,c); LmzT = AtimesB(LmzWt,LmzB); LmzTt = transposeA(LmzT,1/c); [LmxyS,LmzS] = AtimesBkernel(LmzA,LmzTt);</pre>

Table 10.2. Symbolic code for computing the bivariate and trivariate polynomials which, respectively, encode the limiting conditional eigenvector and eigenvalue distribution of algebraic empirical covariance matrices.

multivariate Gaussian with spatio-temporal correlations. The limiting eigenvalue and eigenvector distribution of \mathbf{S} depends on the limiting (algebraic) eigenvalue distributions of \mathbf{A} and \mathbf{B} . The symbolic code for computing these polynomials is listed in table 10.2. When there are no temporal correlations, *i.e.*, $\mathbf{B} = \mathbf{I}$, then we set $L_{mz}^B = m(1-z) - 1$ in the computations and proceed to extract the density and moments from L_{mz}^S as usual. Note the dependence on the limiting value of the ratio $c := \lim c_N$.

Thus the methods developed allow the practitioner to analytically predict the quality of the eigenvectors of \mathbf{S} relative to the eigenvectors of the (spatial) covariance matrix \mathbf{A} for $c \in (0, \infty)$. This provides a window into how sample size constraints affects the estimation of the eigenvectors in high-dimensional settings.

■ 10.4.1 Example

Consider the sample covariance matrix \mathbf{S} formed as in Theorem 10.41. Assume that \mathbf{A}_n and \mathbf{B}_N have the same the limiting eigenvalue distribution function given by

$$F^A(x) = F^B(x) = 0.5 \mathbb{1}_{[1,\infty)} + 0.5 \mathbb{1}_{[2,\infty)}. \quad (10.28)$$

The Stieltjes transform of the limiting eigenvalue distribution function is

$$m_A(z) = m_B(z) = \frac{0.5}{2-z} + \frac{0.5}{1-z},$$

which satisfies the polynomial equation $L_{mz}^A(m, z) = L_{mz}^B(m, z) = 0$ where

$$L_{mz}^A = L_{mz}^B = (-6z + 2z^2 + 4)m + 2z - 3.$$

Using the symbolic tools developed, we can obtain the polynomials that encode the limiting eigenvalue and eigenvector distribution from the sequence of commands:

```
>> syms m x y z c
>> LmzA = m - 0.5/(2-z) - 0.5/(1-z);
>> LmzA = irreducLuv(LmzA,m,z);
>> LmzB = LmzA;
>> [LmxyS,LmzS] = AtimesWishtimesB(LmzA,LmzB,c);
```

from which we obtain the bivariate polynomial

$$L_{mz}^S(m, z) = \sum_{j=1}^6 \sum_{k=1}^4 [\mathbf{T}_{mz}^S]_{jk} m^{j-1} z^{k-1},$$

where

$$\mathbf{T}_{mz}^S \equiv \begin{bmatrix} -18c + 18c^2 & 18c - 9 & 4 & 0 \\ -108c^2 + 36c + 72c^3 & -112c + 18 + 130c^2 & -18 + 54c & 4 \\ 64c^2 + 64c^4 - 128c^3 & 72c - 324c^2 + 288c^3 & 224c^2 - 112c & 36c \\ 0 & 64c^2 - 256c^3 + 192c^4 & 360c^3 - 216c^2 & 112c^2 \\ 0 & 0 & 192c^4 - 128c^3 & 144c^3 \\ 0 & 0 & 0 & 64c^4 \end{bmatrix}.$$

Using the sequence of commands described in Section 8.3, we obtain the first four terms, parameterized by c , of the moment generating function:

$$\begin{aligned} \mu_C(z) = 1 + \frac{9}{4}z + \left(\frac{45}{8}c + \frac{45}{8}\right)z^2 + \left(\frac{675}{16}c + \frac{243}{16}c^2 + \frac{243}{16}\right)z^3 \\ + \left(\frac{3555}{16}c^2 + \frac{1377}{32}c^3 + \frac{3555}{16}c + \frac{1377}{32}\right)z^4 + O(z^5). \end{aligned}$$

Note how the moments explicitly capture the impact of the c on the limiting distribution. This is remarkable, since, for this particular example, it is just not possible to express the density function in closed form.

Figure 10-5(a) plots the limiting eigenvalue density function of \mathbf{S}_n for different values of c . Note the convergence of the distribution, as $c \rightarrow 0$, to an atomic distribution with two equally weighted atoms at 1.5 and 3. Figure 10-5(b) compares theory with experiment for $c = 0.25$.

The trivariate polynomial L_{Mxy}^S is too messy to print. The kernel density functions $k_{S|A}(2, y)$ and $k_{S|A}(1, y)$ have Stieltjes transforms which are solutions of the equation $L_{mz}^{S|i}(m, z) = 0$ where

$$L_{mz}^{S|i}(m, z) = \sum_{j=1}^5 \sum_{k=1}^4 [\mathbf{T}_{mz}^{S|i}]_{jk} m^{j-1} z^{k-1},$$

for $i = 1, 2$ with

$$\mathbf{T}_{mz}^{S|2} = \begin{bmatrix} 8 & 0 & 0 & 0 \\ -24 + 48c & -8 & 0 & 0 \\ 64c^2 - 64c & 48 - 24c & -8 & 0 \\ 0 & 96c & -24 - 48c & 8 \\ 0 & 0 & -32c - 48c^2 & 24c \\ 0 & 0 & 0 & 16c^2 \end{bmatrix}$$


and

$$\mathbf{T}_{mz}^{S|1} = \begin{bmatrix} 8 & 0 & 0 & 0 \\ -12 + 24c & 16 & 0 & 0 \\ -16c + 16c^2 & 42c - 12 & 10 & 0 \\ 0 & 24c^2 - 12c & 21c - 3 & 2 \\ 0 & 0 & 9c^2 - 2c & 3c \\ 0 & 0 & 0 & c^2 \end{bmatrix}.$$


Figure 10-6(a) compares experiment with theory, over values of z in the support of the limiting distribution (shown in Figure 10-6(b)) for the norm square of the projection of the sample eigenvectors of a *single realization* of \mathbf{S}_n formed with $N = 400 = 4n$ onto the subspace spanned by the eigenvalues of value 2 in \mathbf{A} . It is clear that despite the predictions being asymptotic in nature, they accurately predict the behavior for finite sized matrices as well.

■ 10.5 Future work

The ability of predict the deterioration in the quality of the eigenvectors of algebraic empirical covariance matrices due to sample size constraints raises the possibility of whether the results can be used to formulate new high-dimensional covariance matrix estimation algorithms. Covariance estimation from frequentist and Bayesian perspectives is an established topic, *e.g.* [42, 95, 117]; recently some authors have begun to address large N questions, for example by in effect using linear shrinkage on eigenvalues (e.g [26, 58, 89]), with applications for example in empirical finance [56, 57]. There has been interest in structured covariance matrix estimation for signal processing applications as in [38, 39]. Smith treats covariance matrix estimation from a geometric point of view in [90]. Combining the insights of these various authors with the analytical results that capture the degradation in the estimated eigenvectors offers a possibility of attacking this problem from a fresh perspective. A wide open problem is the understanding the nature of the fluctuations, for both the eigenvalues and eigenvectors, for a broader class of random matrix. Related questions include characterizing the rate of convergence.




(a) The density of the limiting eigenvalue distribution function of \mathbf{S}_n for different values of c . When $c = 0.001$ it means that there are roughly 1000 times as many temporal measurements as there are spatial observations and so on.




(b) The theoretical limiting density function (solid line) for $c = 0.25$ is compared with the normalized histogram of the eigenvalues of S_n collected over 4000 Monte-Carlo trials with $n = 100$ and $N = 400$.

Figure 10-5. A Wishart random matrix, \mathbf{S}_n , with spatio-temporal correlations. The spatial and the temporal covariance matrices have limiting eigenvalue distribution given by (10.28).



(a) Norm square of the projection of the eigenvectors of a single random matrix realization onto the eigenspace of \mathbf{A} spanned by the eigenvalue of value 5.



(b) The theoretical limiting density function (solid line) for $c = 0.25$ is compared with the normalized histogram of the eigenvalues of \mathbf{S}_n collected over 4000 Monte-Carlo trials with $n = 100$ and $N = 400$.

Figure 10-6. A Wishart random matrix, \mathbf{S}_n , with spatio-temporal correlations. The spatial and the temporal covariance matrices have limiting eigenvalue distribution given by (10.28).

Afterword

In the first part of this dissertation, we applied random matrix theory to inference problems where the measurements were drawn from a multivariate normal distribution. By exploiting the properties of the eigenvalues of large dimensional Wishart distributed random matrices we were able to design algorithms that turned the underlying high-dimensionality into an advantage.

In the second part of this dissertation, we developed a powerful method that allowed us to characterize the eigenvalues of a broad class of random matrices well beyond the special case of matrices with Wishart distribution. A natural question then arises: Can these more complicated matrix models be physically justified so that the results can be applied? The conclusion of this dissertation makes this the opportune moment for us to pose this question and share some of our thoughts on this matter.

We feel that an important extension of the work in this thesis is the development of random matrix models that adequately capture the essential complexities of the real-world high-dimensional inference problem without being so complicated that we cannot get answers for them. We anticipate that initially this will have to be done on an application-by-application basis in close collaboration with experts in the field who can ensure that aspects of the problem that could affect the solution are not missed.

Progress on this front is likely to be deliberate because there is an art to model building which makes it difficult to rush, although like other artistic endeavors, as Gil Strang puts it [97, pp. 9], “people who do it well will agree when it is done well.”

Researchers used to producing differential equations to model their problem might have to work a bit longer and squint a bit harder to discern the random matrix that is buried in their problem, if at all. The incentive for their effort is that if their random matrix model fits into the general framework developed, then the full power of the methods developed in this dissertation can be brought to bear on their problems. The blessings of high-dimensionality, in an inferential context, will then fully manifest.

Of course, we acknowledge the possibility that practitioners experimenting with our framework while model building might discover that the current theory does not suffice and that additional theory and methods are needed. In concluding this dissertation, we adopt the position that we would greet such a contrary discovery with a healthy dose of gratification. After all, it would be a testimony to our belief that this dissertation can be the starting point for this and other explorations.

Random Matrices and Free Probability Theory

This material in this appendix is based on an exposition by Roland Speicher. It has been included (almost) verbatim with his permission.

■ A.1 Moments of random matrices and asymptotic freeness

What can we say about the eigenvalue distribution of the sum $\mathbf{A} + \mathbf{B}$ of the matrices? Of course, the latter is not just determined by the eigenvalues of \mathbf{A} and the eigenvalues of \mathbf{B} , but also by the relation between the eigenspaces of \mathbf{A} and of \mathbf{B} . Actually, it is a quite hard problem (Horn's conjecture) — which was only solved recently — to characterize all possible eigenvalue distributions of $\mathbf{A} + \mathbf{B}$. However, if one is asking this question in the context of $N \times N$ -random matrices, then in many situations the answer becomes deterministic in the limit $N \rightarrow \infty$.

Definition A.11. *Let $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ be a sequence of $N \times N$ -random matrices. We say that \mathbf{A} has a limit eigenvalue distribution if the limit of all moments*

$$\alpha_n := \lim_{N \rightarrow \infty} E[\text{tr}(\mathbf{A}_N^n)] \quad (n \in \mathbb{N})$$

exists, where E denotes the expectation and tr the normalized trace.

Using the language of limit eigenvalue distribution as in Definition A.11, our question becomes: Given two random matrix ensembles of $N \times N$ -random matrices, $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ and $\mathbf{B} = (\mathbf{B}_N)_{N \in \mathbb{N}}$, with limit eigenvalue distribution, does their sum $\mathbf{C} = (\mathbf{C}_N)_{N \in \mathbb{N}}$, with $\mathbf{C}_N = \mathbf{A}_N + \mathbf{B}_N$, also have a limit eigenvalue distribution, and furthermore, can we calculate the limit moments $\alpha_n^{\mathbf{C}}$ of \mathbf{C} from the limiting moments $(\alpha_k^{\mathbf{A}})_{k \geq 1}$ of \mathbf{A} and the limiting moments $(\alpha_k^{\mathbf{B}})_{k \geq 1}$ of \mathbf{B} in a deterministic way. It turns out that this is the case if the two ensembles are in generic position, and then the rule for calculating the limit moments of \mathbf{C} are given by Voiculescu's concept of "freeness".

Lemma A.12 (Voiculescu [108]). *Let \mathbf{A} and \mathbf{B} be two random matrix ensembles of $N \times N$ -random matrices, $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ and $\mathbf{B} = (\mathbf{B}_N)_{N \in \mathbb{N}}$, each of them with a limit*

eigenvalue distribution. Assume that \mathbf{A} and \mathbf{B} are independent (i.e., for each $N \in \mathbb{N}$, all entries of \mathbf{A}_N are independent from all entries of \mathbf{B}_N), and that at least one of them is unitarily invariant (i.e., for each N , the joint distribution of the entries does not change if we conjugate the random matrix with an arbitrary unitary $N \times N$ matrix). Then \mathbf{A} and \mathbf{B} are asymptotically free in the sense of the following definition.

Definition A.13 (Voiculescu [105]). Two random matrix ensembles $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ and $\mathbf{B} = (\mathbf{B}_N)_{N \in \mathbb{N}}$ with limit eigenvalue distributions are asymptotically free if we have for all $p \geq 1$ and all integers $n(1), m(1), \dots, n(p), m(p) \geq 1$ that

$$\lim_{N \rightarrow \infty} E \left[\text{tr} \left\{ (\mathbf{A}_N^{n(1)} - \alpha_{n(1)}^A \mathbf{I}) \cdot (\mathbf{B}_N^{m(1)} - \alpha_{m(1)}^B \mathbf{I}) \cdots \right. \right. \\ \left. \left. \cdots (\mathbf{A}_N^{n(p)} - \alpha_{n(p)}^A \mathbf{I}) \cdot (\mathbf{B}_N^{m(p)} - \alpha_{m(p)}^B \mathbf{I}) \right\} \right] = 0$$

Thus, for example if \mathbf{A} and \mathbf{B} are asymptotically free then, we necessarily have

$$\lim_{N \rightarrow \infty} E \left[\text{tr} \left\{ (\mathbf{A}_N^1 - \alpha_1^A \mathbf{I}) \cdot (\mathbf{B}_N^2 - \alpha_2^B \mathbf{I}) \cdot (\mathbf{A}_N^3 - \alpha_3^A \mathbf{I}) \cdot (\mathbf{B}_N^4 - \alpha_4^B \mathbf{I}) \right\} \right] = 0$$

where we have inserted $n(1) = 1, n(2) = 3, m(1) = 2, m(2) = 4$ in Definition A.13. Embedded in the definition of asymptotic freeness is a rule which allows us to calculate all mixed moments in \mathbf{A} and \mathbf{B} , i.e., all expressions of the form

$$\lim_{N \rightarrow \infty} E[\text{tr}(\mathbf{A}^{n(1)} \mathbf{B}^{m(1)} \mathbf{A}^{n(2)} \mathbf{B}^{m(2)} \cdots \mathbf{A}^{n(p)} \mathbf{B}^{m(p)})]$$

out of the limit moments of \mathbf{A} and the limit moments of \mathbf{B} . In particular, this means that all limit moments of $\mathbf{A} + \mathbf{B}$ (which are sums of mixed moments) exist, thus $\mathbf{A} + \mathbf{B}$ has a limit distribution, and are actually determined in terms of the limit moments of \mathbf{A} and the limit moments of \mathbf{B} . The actual calculation rule is not directly clear from the above definition but a basic result of Voiculescu shows how this can be achieved by going over from the moments α_n to new quantities κ_n . In [91], the combinatorial structure behind these κ_n was revealed and the name “free cumulants” was coined for them.

Definition A.14 (Voiculescu [106], Speicher [91]). Given the moments $(\alpha_n)_{n \geq 1}$ of some distribution (or limit moments of some random matrix ensemble), we define the corresponding free cumulants $(\kappa_n)_{n \geq 1}$ by the following relation between their generating power series: If we put

$$M(x) := 1 + \sum_{n \geq 1} \alpha_n x^n \quad \text{and} \quad C(x) := 1 + \sum_{n \geq 1} \kappa_n x^n,$$

then we require as a relation between these formal power series that

$$C(xM(x)) = M(x).$$

Voiculescu actually formulated the relation above in a slightly different way using the so-called R -transform $\mathcal{R}(x)$, which is related to $C(x)$ by the relation

$$C(x) = 1 + z\mathcal{R}(x)$$

and in terms of the Cauchy transform $G(x)$ corresponding to a measure with moments α_n , which is related to $M(x)$ by

$$G(x) = \frac{M(\frac{1}{x})}{x}.$$

In these terms the equation $C(xM(x)) = M(x)$ says that

$$\frac{1}{G(x)} + \mathcal{R}(G(x)) = x, \tag{A.1}$$

i.e., that $G(x)$ and $K(x) := \frac{1}{x} + \mathcal{R}(x)$ are inverses of each other under composition.

One should also note that the relation $C(xM(x)) = M(x)$ determines the moments uniquely in terms of the cumulants and the other way around. The relevance of the κ_n and the R -transform for our problem comes from the following result of Voiculescu, which provides, together with (A.1), a very efficient way for calculating eigenvalue distributions of the sum of asymptotically free random matrices.

Lemma A.15 (Voiculescu [106]). *Let \mathbf{A} and \mathbf{B} be two random matrix ensembles which are asymptotically free. Denote by $\kappa_n^A, \kappa_n^B, \kappa_n^{A+B}$ the free cumulants of $\mathbf{A}, \mathbf{B}, \mathbf{A} + \mathbf{B}$, respectively. Then one has for all $n \geq 1$ that*

$$\kappa_n^{A+B} = \kappa_n^A + \kappa_n^B.$$

Alternatively,

$$\mathcal{R}^{A+B}(x) = \mathcal{R}^A(x) + \mathcal{R}^B(x).$$

This lemma is one reason for calling the κ_n cumulants (as they linearize the “free convolution” in the same way as the usual convolution is linearized by classical cumulants), but there is also another justification for this, namely they are also the limit of classical cumulants of the entries of our random matrix, in the case that this is unitarily invariant.

Proposition A.16. *Let $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ be a unitarily invariant random matrix ensemble of $N \times N$ random matrices \mathbf{A}_N whose limit eigenvalue distribution exists. Then the free cumulants of this matrix ensemble can also be expressed as the limit of special classical cumulants of the entries of the random matrices: If $\mathbf{A}_N = (a_{ij}^{(N)})_{i,j=1}^N$, then*

$$\kappa_n^A = \lim_{N \rightarrow \infty} N^{n-1} c_n(a_{i(1)i(2)}^{(N)}, a_{i(2)i(3)}^{(N)}, \dots, a_{i(n),i(1)}^{(N)})$$

for any choice of distinct $i(1), \dots, i(n)$.

PROOF. This appears in Collins, Mingo, Śniady, Speicher [23]. \square

■ A.2 Fluctuations of random matrices and asymptotic second order freeness

There are many more refined questions about the limiting eigenvalue distribution of random matrices. In particular, questions around fluctuations have received a lot of interest in the last decade or so. The main motivation for Speicher and colleagues introducing the concept of “second order freeness” was to understand the global fluctuations of the eigenvalues, which means that we look at the probabilistic behavior of traces of powers of our matrices. The limiting eigenvalue distribution, as considered in the last section, gives us the limit of the average of this traces. However, one can make more refined statements about their distributions. Consider a random matrix $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ and look on the normalized traces $\text{tr}(\mathbf{A}_N^k)$. Our assumption of a limit eigenvalue distribution means that the limits $\alpha_k := \lim_{N \rightarrow \infty} E[\text{tr}(\mathbf{A}_N^k)]$ exist. It turned out that in many cases the fluctuation around this limit,

$$\text{tr}(\mathbf{A}_N^k) - \alpha_k$$

is asymptotically Gaussian of order $1/N$; i.e., the random variable

$$N \cdot (\text{tr}(\mathbf{A}_N^k) - \alpha_k) = \text{Tr}(\mathbf{A}_N^k) - N\alpha_k = \text{Tr}(\mathbf{A}_N^k - \alpha_k \mathbf{1})$$

(where Tr denotes the unnormalized trace) converges for $N \rightarrow \infty$ to a normal variable. Actually, the whole family of centered unnormalized traces $(\text{Tr}(\mathbf{A}_N^k) - N\alpha_k)_{k \geq 1}$ converges to a centered Gaussian family.

Note that in Speicher and colleagues theory the formulation is in terms of complex random matrices; in the case of real random matrices there are additional complications which their theory does not currently account for but which are likely be resolved in future investigations.

Thus the main information about fluctuations of our considered ensemble is contained in the covariance matrix of the limiting Gaussian family, i.e., in the quantities

$$\alpha_{m,n} := \lim_{N \rightarrow \infty} \text{cov}(\text{Tr}(\mathbf{A}_N^m), \text{Tr}(\mathbf{A}_N^n)).$$

Let us emphasize that the α_n and the $\alpha_{m,n}$ are actually limits of classical cumulants of traces; namely of the expectation as first and the variance as second cumulant. Nevertheless, the α 's will behave and will also be treated like moments; accordingly we will call the $\alpha_{m,n}$ ‘fluctuation moments’. We will below define some other quantities $\kappa_{m,n}$, which take the role of cumulants in this context.

This kind of convergence to a Gaussian family was formalized in [64] by the notion of “second order limit distribution” (see Definition 4.25).

Definition A.21. Let $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ be an ensemble of $N \times N$ random matrices \mathbf{A}_N .

We say that it has a second order limit distribution if for all $m, n \geq 1$ the limits

$$\alpha_n := \lim_{N \rightarrow \infty} c_1(\text{tr}(\mathbf{A}_N^n))$$

and

$$\alpha_{m,n} := \lim_{N \rightarrow \infty} c_2(\text{Tr}(\mathbf{A}_N^m), \text{Tr}(\mathbf{A}_N^n))$$

exist and if

$$\lim_{N \rightarrow \infty} c_r(\text{Tr}(\mathbf{A}_N^{n(1)}), \dots, \text{Tr}(\mathbf{A}_N^{n(r)})) = 0$$

for all $r \geq 3$ and all $n(1), \dots, n(r) \geq 1$.

We can now ask the same kind of question for the limit fluctuations as for the limit moments; namely, if we have two random matrix ensembles \mathbf{A} and \mathbf{B} and we know the second order limit distribution of \mathbf{A} and the second order limit distribution of \mathbf{B} , does this imply that we have a second order limit distribution for $\mathbf{A} + \mathbf{B}$, and, if so, is there an effective way for calculating it. Again, we can only hope for a positive solution to this if \mathbf{A} and \mathbf{B} are in a kind of generic position. As it turned out, the same requirements as before are sufficient for this. The rule for calculating mixed fluctuations constitutes the essence of the definition of the concept of second order freeness.

Proposition A.22. *Let \mathbf{A} and \mathbf{B} be two random matrix ensembles of $N \times N$ -random matrices, $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ and $\mathbf{B} = (\mathbf{B}_N)_{N \in \mathbb{N}}$, each of them having a second order limit distribution. Assume that \mathbf{A} and \mathbf{B} are independent and that at least one of them is unitarily invariant. Then \mathbf{A} and \mathbf{B} are asymptotically free of second order in the sense of the following definition.*

PROOF. This appears in Mingo, Śniady, Speicher [63]. \square

Definition A.23 (Mingo, Speicher [64]). *Consider two random matrix ensembles $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ and $\mathbf{B} = (\mathbf{B}_N)_{N \in \mathbb{N}}$, each of them with a second order limit distribution. Denote by*

$$Y_N(n(1), m(1), \dots, n(p), m(p))$$

the random variable

$$\text{Tr}((\mathbf{A}_N^{n(1)} - \alpha_{n(1)}^A \mathbf{1})(\mathbf{B}_N^{m(1)} - \alpha_{m(1)}^B \mathbf{1}) \cdots (\mathbf{A}_N^{n(p)} - \alpha_{n(p)}^A \mathbf{1})(\mathbf{B}_N^{m(p)} - \alpha_{m(p)}^B \mathbf{1})).$$

The random matrices $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ and $\mathbf{B} = (\mathbf{B}_N)_{N \in \mathbb{N}}$ are asymptotically free of second order if for all $n, m \geq 1$

$$\lim_{N \rightarrow \infty} c_2(\text{Tr}(\mathbf{A}_N^n - \alpha_n^A \mathbf{1}), \text{Tr}(\mathbf{B}_N^m - \alpha_m^B \mathbf{1})) = 0$$

and for all $p, q \geq 1$ and $n(1), \dots, n(p), m(1), \dots, m(p), \tilde{n}(1), \dots, \tilde{n}(q), \tilde{m}(1), \dots, \tilde{m}(q) \geq 1$ we have

$$\lim_{N \rightarrow \infty} c_2\left(Y_N(n(1), m(1), \dots, n(p), m(p)), Y_N(\tilde{n}(1), \tilde{m}(2), \dots, \tilde{n}(q), \tilde{m}(q))\right) = 0$$

if $p \neq q$, and otherwise (where we count modulo p for the arguments of the indices, i.e., $n(i+p) = n(i)$)

$$\begin{aligned} & \lim_{N \rightarrow \infty} c_2 \left(Y_N(n(1), m(1), \dots, n(p), m(p)), Y_N(\tilde{n}(p), \tilde{m}(p), \dots, \tilde{n}(1), \tilde{m}(1)) \right) \\ &= \sum_{k=1}^p \prod_{i=1}^p (\alpha_{\tilde{n}(i+k)+\tilde{n}(i)}^A - \alpha_{\tilde{n}(i+k)}^A \alpha_{\tilde{n}(i)}^A) (\alpha_{m(i+k)+\tilde{m}(i+1)}^B - \alpha_{m(i+k)}^B \alpha_{\tilde{m}(i+1)}^B). \end{aligned}$$

Again, it is crucial to realize that this definition allows one (albeit in a complicated way) to express every second order mixed moment, i.e., a limit of the form

$$\lim_{N \rightarrow \infty} c_2(\text{Tr}(\mathbf{A}_N^{n(1)} \mathbf{B}_N^{m(1)} \cdots \mathbf{A}_N^{n(p)} \mathbf{B}_N^{m(p)}), \text{Tr}(\mathbf{A}_N^{\tilde{n}(1)} \mathbf{B}_N^{\tilde{m}(1)} \cdots \mathbf{A}_N^{\tilde{n}(q)} \mathbf{B}_N^{\tilde{m}(q)}))$$

in terms of the second order limits of \mathbf{A} and the second order limits of \mathbf{B} . In particular, asymptotic freeness of second order also implies that the sum $\mathbf{A} + \mathbf{B}$ of our random matrix ensembles has a second order limit distribution and allows one to express them in principle in terms of the second order limit distribution of \mathbf{A} and the second order limit distribution of \mathbf{B} . As in the case of first order freeness, it is not clear at all how this calculation of the fluctuations of $\mathbf{A} + \mathbf{B}$ out of the fluctuations of \mathbf{A} and the fluctuations of \mathbf{B} can be performed effectively. In [23] Speicher and colleagues were able to solve this problem by providing a second order cumulant machinery, similar to the first order case. Again, the idea is to go over to quantities which behave like cumulants in this setting. The actual description of those relies on combinatorial objects (annular non-crossing permutations), but as before this can be reformulated in terms of formal power series. The definition can be spelled out in this form below.

Definition A.24 (Collins, Mingo, Śniady, Speicher [23]). Let $(\alpha_n)_{n \geq 1}$ and $(\alpha_{m,n})_{m,n \geq 1}$ describe the first and second order limit moments of a random matrix ensemble. We define the corresponding first and second order free cumulants $(\kappa_n)_{n \geq 1}$ and $(\kappa_{m,n})_{m,n \geq 1}$ by the following requirement in terms of the corresponding generating power series. Put

$$C(x) := 1 + \sum_{n \geq 1} \kappa_n x^n, \quad C(x, y) := \sum_{m,n \geq 1} \kappa_{m,n} x^m y^n$$

and

$$M(x) := 1 + \sum_{n \geq 1} \alpha_n x^n, \quad M(x, y) := \sum_{m,n \geq 1} \alpha_{m,n} x^m y^n.$$

Then we require as relations between these formal power series that

$$C(xM(x)) = M(x) \tag{A.2}$$

and for the second order

$$M(x, y) = H(xM(x), yM(y)) \cdot \frac{\frac{d}{dx}(xM(x))}{M(x)} \cdot \frac{\frac{d}{dy}(yM(y))}{M(y)}, \quad (\text{A.3})$$

where

$$H(x, y) := C(x, y) - xy \frac{\partial^2}{\partial x \partial y} \log \left(\frac{x C(y) - y C(x)}{x - y} \right), \quad (\text{A.4})$$

or equivalently,

$$\begin{aligned} M(x, y) = C(xM(x), yM(y)) \cdot \frac{\frac{d}{dx}(xM(x))}{M(x)} \cdot \frac{\frac{d}{dy}(yM(y))}{M(y)} \\ + xy \left(\frac{\frac{d}{dx}(xM(x)) \cdot \frac{d}{dy}(yM(y))}{(xM(x) - yM(y))^2} - \frac{1}{(x - y)^2} \right). \end{aligned} \quad (\text{A.5})$$

As in the first order case, instead of the moment power series $M(x, y)$ one can consider a kind of second order Cauchy transform, defined by

$$G(x, y) := \frac{M(\frac{1}{x}, \frac{1}{y})}{xy}.$$

If we also define a kind of second order R transform $\mathcal{R}(x, y)$ by

$$\mathcal{R}(x, y) := \frac{1}{xy} C(x, y),$$

then the formula (A.5) takes on a particularly nice form:

$$G(x, y) = G'(x)G'(y) \left\{ \mathcal{R}(G(x), G(y)) + \frac{1}{(G(x) - G(y))^2} \right\} - \frac{1}{(x - y)^2}. \quad (\text{A.6})$$

$G(x)$ is here, as before, the first order Cauchy transform, $G(x) = \frac{1}{x} M(1/x)$.

The $\kappa_{m,n}$ defined above deserve the name ‘‘cumulants’’ as they linearize the problem of adding random matrices which are asymptotically free of second order. Namely, we have the following lemma, which provides, together with (A.6), an effective machinery for calculating the fluctuations of the sum of asymptotically free random matrices.

Lemma A.25. *Let \mathbf{A} and \mathbf{B} be two random matrix ensembles which are asymptotically free. Then one has for all $m, n \geq 1$ that*

$$\kappa_n^{A+B} = \kappa_n^A + \kappa_n^B \quad \text{and} \quad \kappa_{m,n}^{A+B} = \kappa_{m,n}^A + \kappa_{m,n}^B.$$

Alternatively,

$$\mathcal{R}^{A+B}(x) = \mathcal{R}^A(x) + \mathcal{R}^B(x)$$

and

$$\mathcal{R}^{A+B}(x, y) = \mathcal{R}^A(x, y) + \mathcal{R}^B(x, y).$$

PROOF. This was proved by Collins, Mingo, Śniady, Speicher in [23]. \square

Again, one can express the second order cumulants as limits of classical cumulants of entries of a unitarily invariant matrix. In contrast to the first order case, we have now to run over two disjoint cycles in the indices of the matrix entries.

Lemma A.26. *Let $\mathbf{A} = (\mathbf{A}_N)_{N \in \mathbb{N}}$ be a unitarily invariant random matrix ensemble which has a second order limit distribution. Then the second order free cumulants of this matrix ensemble can also be expressed as the limit of classical cumulants of the entries of the random matrices: If $\mathbf{A}_N = (a_{ij}^{(N)})_{i,j=1}^N$, then*

$$\kappa_{m,n}^A = \lim_{N \rightarrow \infty} N^{m+n} c_{m+n}(a_{i(1)i(2)}^{(N)}, a_{i(2)i(3)}^{(N)}, \dots, a_{i(m),i(1)}^{(N)}, \\ a_{j(1)j(2)}^{(N)}, a_{j(2)j(3)}^{(N)}, \dots, a_{j(n),j(1)}^{(N)})$$

for any choice of distinct $i(1), \dots, i(m), j(1), \dots, j(n)$.

PROOF. This was proved by Collins, Mingo, Śniady, Speicher in [23]. \square

■ A.3 Wishart matrices and Proof of Proposition 4.27

Wishart matrices, in the large size limit, fit quite well into the framework of first and second order free probability theory. In particular, their free cumulants of first and second order are quite easy to determine and are of a particularly nice form. We will use this to give a proof of Proposition 4.27. The statements in that proposition go back to the work of Bai and Silverstein, see *e.g.*, [10] who give a more direct proof via analytic calculations of the Cauchy transforms. We prefer here, however, to show how Wishart matrices fit conceptually into the frame of free probability theory.

Let us remark that whereas the results around first order freeness are valid for complex as well as real random matrices, this is not the case any more for the second order; there are some complications to be dealt with in this case and at the moment the theory of second order freeness for real random matrices has not yet been developed. Thus our proof of the fluctuation formula (4.13b) will only cover the complex case. The fact that the real case differs from the complex case by a factor 2 can be found in the work of Bai and Silverstein [10].

Instead of looking on the Wishart matrix $\mathbf{S} := \frac{1}{n} \mathbf{X} \mathbf{X}'$ from Equation (4.1) we will consider the closely related matrix

$$\mathbf{T} := \frac{1}{m} \mathbf{X}' \mathbf{X}.$$

Note that \mathbf{S} is a $m \times m$ -matrix, whereas \mathbf{T} is an $n \times n$ matrix. The relation between the spectral behavior of those two matrices is quite straightforward, namely they have

the same non-zero eigenvalues, which are filled up with additional zeros for the larger one. Thus the transition between these two matrices is very easy; their eigenvalue distributions are related by a rescaling (since the first order moments α_n go with the normalized trace) and their fluctuations are the same (since the second order moments $\alpha_{m,n}$ go with the unnormalized trace). The reason for considering \mathbf{T} instead of \mathbf{S} is the following nice description of its first and second order distribution. In this theorem we will realize the Wishart matrix $\mathbf{S} = \frac{1}{m}\mathbf{X}\mathbf{X}'$ with covariance matrix Σ in the form $\Sigma^{1/2}\mathbf{Y}\mathbf{Y}'\Sigma^{1/2}$ where \mathbf{Y} is a $n \times m$ Gaussian random matrix with independent entries of mean zero and variance $1/m$. The matrix \mathbf{T} takes then on the form

$$\mathbf{T} = \mathbf{Y}'\Sigma\mathbf{Y}.$$

Note that we allow Σ to be itself random in the following theorem.

Proposition A.31. *Let $\Sigma = (\Sigma_n)_{n \in \mathbb{N}}$ be a random matrix ensemble of selfadjoint $n \times n$ -matrices and consider in addition a Gaussian ensemble $\mathbf{Y} = (\mathbf{Y}_n)_{n \in \mathbb{N}}$ of non-selfadjoint rectangular Gaussian $n \times m$ -random matrices (with mean zero and variance $1/m$ for the entries) such that \mathbf{Y} and Σ are independent. Put*

$$\mathbf{T} := (\mathbf{Y}'_n \Sigma_n \mathbf{Y}_n)_{n \in \mathbb{N}}.$$

In the following we consider the limit

$$n, m \rightarrow \infty \quad \text{such that} \quad \lim \frac{n}{m} =: c$$

for some fixed $c \in (0, \infty)$.

(1) Assume that the limit eigenvalue distribution of $\Sigma = (\Sigma_n)_{n \in \mathbb{N}}$ exists for $n \rightarrow \infty$. Then \mathbf{T} , considered as an ensemble of $m \times m$ -random matrices $\mathbf{Y}'_n \Sigma_n \mathbf{Y}_n$, has a limit eigenvalue distribution. This limit eigenvalue distribution is determined by the fact that its free cumulants are given by the scaled corresponding limit moments of Σ , i.e., for all $j \geq 1$ we have

$$\kappa_j^{\mathbf{T}} = c\alpha_j^{\Sigma}.$$

(2) Assume that we are in the complex case and that $\Sigma = (\Sigma_n)_{n \in \mathbb{N}}$ has a second order limit distribution for $n \rightarrow \infty$. Then \mathbf{T} has a second order limit distribution, which is determined as follows: for all $i, j \geq 1$ we have

$$\kappa_i^{\mathbf{T}} = c\alpha_i^{\Sigma} \quad \text{and} \quad \kappa_{i,j}^{\mathbf{T}} = \alpha_{i,j}^{\Sigma}.$$

PROOF. The first order statement of this theorem is due to Nica and Speicher, see [68], the second order statement follows from the calculations in [64]. \square

We will now use this theorem to prove our Proposition 4.27 in the complex case.

PROOF. If

$$M^{\Sigma}(x) = 1 + \sum_{i \geq 1} \alpha_i^{\Sigma} x^i$$

is the generating power series for the limit moments of Σ , then the above proposition says that the generating power series $C^{\mathbf{T}}(x)$ for the free cumulants of \mathbf{T} is related with $M^{\Sigma}(x)$ by

$$\begin{aligned} C^{\mathbf{T}}(x) &= 1 + \sum_{i \geq 1} \kappa_i^{\mathbf{T}} x^i \\ &= 1 + c \sum_{i \geq 1} \alpha_i^{\Sigma} x^i \\ &= (1 - c) + cM^{\Sigma}(x). \end{aligned}$$

Thus, by the general relation $C^{\mathbf{T}}(xM^{\mathbf{T}}(x)) = M^{\mathbf{T}}(x)$, we get the generating power series $M^{\mathbf{T}}(x)$ for the limit moments of \mathbf{T} as a solution to the equation

$$1 - c + cM^{\Sigma}[xM^{\mathbf{T}}(x)] = M^{\mathbf{T}}(x). \quad (\text{A.7})$$

Let us now rewrite this for the Wishart matrix \mathbf{S} . Recall that the moments of \mathbf{S} and the moments of \mathbf{T} are related by a simple scaling factor, resulting in a relation of the form

$$M^{\mathbf{T}}(x) = c(M^{\mathbf{S}}(x) - 1) + 1.$$

This gives

$$M^{\mathbf{T}}(x) = M^{\Sigma}[x(cM^{\mathbf{T}}(x) - c + 1)].$$

Rewriting this in terms of

$$g(x) := \frac{1}{x}M^{\mathbf{S}}(1/x) \quad \text{and} \quad g^{\Sigma}(x) := \frac{1}{x}M^{\Sigma}(1/x)$$

yields formula (4.11).

In order to get the result for second order one only has to observe that the fluctuations of a non-random covariance matrix vanish identically, hence $C^{\mathbf{T}}(x, y) = C^{\mathbf{S}}(x, y) = 0$, and thus (A.5) reduces directly to (4.13).

□

Bibliography

- [1] H. AKAIKE, *Information theory and an extension of the maximum likelihood principle*, in Second International Symposium on Inform. Theory (Tsahkadsor, 1971), Akadémiai Kiadó, Budapest, 1973, pp. 267–281.
- [2] H. AKAIKE, *A new look at the statistical model identification*, IEEE Trans. Automatic Control, AC-19 (1974), pp. 716–723. System identification and time-series analysis.
- [3] N. I. AKHIEZER, *The classical moment problem and some related questions in analysis*, Hafner Publishing Co., New York, New York, 1965. Translated by N. Kemmer.
- [4] A. G. AKRITAS, *Sylvester’s forgotten form of the resultant*, Fibonacci Quart., 31 (1993), pp. 325–332.
- [5] G. W. ANDERSON AND O. ZEITOUNI, *A CLT for a band matrix model*, Probab. Theory Related Fields, 134 (2006), pp. 283–338.
- [6] T. W. ANDERSON, *Asymptotic theory of principal component analysis*, Annals of Math. Statistics, 34 (1963), pp. 122–248.
- [7] A. B. BAGGEROER AND H. COX, *Passive sonar limits upon nulling multiple moving ships with large aperture arrays*, in Conference Record of the Thirty-Third Asilomar Conference on Signals, Systems, and Computers, vol. 1, Pacific Grove, CA, 1999, pp. 103–108.
- [8] Z. D. BAI, *Methodologies in spectral analysis of large-dimensional random matrices, a review*, Statist. Sinica, 9 (1999), pp. 611–677. With comments by G. J. Rodgers and Jack W. Silverstein; and a rejoinder by the author.
- [9] Z. D. BAI AND J. W. SILVERSTEIN, *On the empirical distribution of eigenvalues of a class of large dimensional random matrices*, Journal of Multi. Analysis, 54 (1995), pp. 175–192.

- [10] —, *CLT for linear spectral statistics of large-dimensional sample covariance matrices*, *Ann. Probab.*, 32 (2004), pp. 553–605.
- [11] J. BAIK, G. BEN AROUS, AND S. PÉCHÉ, *Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices*, *Ann. Probab.*, 33 (2005), pp. 1643–1697.
- [12] J. BAIK AND J. W. SILVERSTEIN, *Eigenvalues of large sample covariance matrices of spiked population models*, <http://arxiv.org/math.ST/048165>, (2004).
- [13] M. S. BARTLETT, *A note on the multiplying factors for various χ^2 approximations*, *J. Roy. Stat. Soc., ser. B*, 16 (1954), pp. 296–298.
- [14] F. BERGERON AND C. REUTENAUER, *Combinatorial resolution of systems of differential equations. III. A special class of differentially algebraic series*, *European J. Combin.*, 11 (1990), pp. 501–512.
- [15] P. BIANE, *Processes with free increments*, *Math. Z.*, 227 (1998), pp. 143–174.
- [16] —, *Free probability for probabilists*, in *Quantum probability communications*, Vol. XI (Grenoble, 1998), QP-PQ, XI, World Sci. Publishing, River Edge, NJ, 2003, pp. 55–71.
- [17] P. BILLINGSLEY, *Convergence of probability measures*, *Wiley Series in Probability and Statistics: Probability and Statistics*, John Wiley & Sons Inc., New York, second ed., 1999. A Wiley-Interscience Publication.
- [18] L. BRENNAN, I. REED, AND J. MALLETT, *Rapid convergence rate in adaptive arrays*, *IEEE Trans. on Aero. and Electr. Sys.*, 10 (1974), pp. 850–863.
- [19] J. CAPON, *High-resolution frequency-wavenumber spectrum analysis*, *Proc. of the IEEE*, 57 (1969), pp. 1408–1418.
- [20] J. CAPON AND N. R. GOODMAN, *Probability distributions for estimator of the frequency wavenumber spectrum*, *Proc. of the IEEE*, 58 (1970), pp. 1785–1786.
- [21] B. D. CARLSON, *Covariance matrix estimation errors and diagonal loading in adaptive arrays*, *IEEE Trans. on Aero. and Elec. Systems*, 24 (1988), pp. 397–401.
- [22] B. COLLINS, *Product of random projections, Jacobi ensembles and universality problems arising from free probability*, *Probab. Theory Related Fields*, 133 (2005), pp. 315–344.
- [23] B. COLLINS, J. MINGO, P. ŚNIADY, AND R. SPEICHER, *Second order freeness and fluctuations of random matrices: III. higher order freeness and free cumulants*. <http://arxiv.org/math.OA/0606431>.

- [24] H. COX, R. ZESKIND, AND M. OWEN, *Robust adaptive beamforming*, IEEE Trans. on Acoustics, Speech, and Signal Proc., 35 (1987), pp. 1365–1376.
- [25] H. CRAMER, *Mathematical methods of Statistics*, Princeton University Press, Princeton, New Jersey, 1946.
- [26] M. J. DANIELS AND R. E. KASS, *Shrinkage estimators for covariance matrices*, Biometrics, 57 (2001), pp. 1173–1184.
- [27] P. DEIFT, T. KRIECHERBAUER, AND K. T.-R. MCCLAUGHLIN, *New results on the equilibrium measure for logarithmic potentials in the presence of an external field*, J. Approx. Theory, 95 (1998), pp. 388–475.
- [28] P. DEIFT, T. KRIECHERBAUER, K. T.-R. MCCLAUGHLIN, S. VENAKIDES, AND X. ZHOU, *Uniform asymptotics for polynomials orthogonal with respect to varying exponential weights and applications to universality questions in random matrix theory*, Comm. Pure Appl. Math., 52 (1999), pp. 1335–1425.
- [29] P. A. DEIFT, *Orthogonal polynomials and random matrices: a Riemann-Hilbert approach*, vol. 3 of Courant Lecture Notes in Mathematics, New York University Courant Institute of Mathematical Sciences, New York, 1999.
- [30] W. B. DOZIER AND J. W. SILVERSTEIN, *On the empirical distribution of eigenvalues of large dimensional information-plus-noise type matrices*. <http://www4.ncsu.edu/~jack/infnoise.pdf>, 2004.
- [31] I. DUMITRIU AND A. EDELMAN, *MOPS: Multivariate Orthogonal Polynomials (symbolically)*. <http://arxiv.org/math-ph/0409066>, 2004.
- [32] ———, *Global spectrum fluctuations for the β -Hermite and β -Laguerre ensembles via matrix models*, J. Math. Phys., 47 (2006), pp. 063302, 36.
- [33] I. DUMITRIU AND E. RASSART, *Path counting and random matrix theory*, Electronic Journal of Combinatorics, 7 (2003). R-43.
- [34] A. EDELMAN AND N. R. RAO, *Random matrix theory*, in Acta Numerica, vol. 14, Cambridge University Press, 2005, pp. 233–297.
- [35] B. EFRON AND C. MORRIS, *Limiting the risk of Bayes and Empirical Bayes Estimators - Part I: The Bayes case*, Journal of the American Statistical Association, 66 (1971), pp. 807–815.
- [36] W. FEATHERSTONE, H. J. STRANGWAYS, AND H. MEWES, *An improved Capon estimator for HF DF using limited sample data sets*, in IEE Colloquium on Propagation Characteristics and Related System Techniques for Beyond Line-of-Sight Radio, November 1997, pp. 2/1–2/5. Ref. No. 1997/390.

- [37] P. FLAJOLET AND R. SEDGEWICK, *Analytic combinatorics: Functional equations, rational and algebraic functions*, Research Report 4103, INRIA, 2001. <http://algo.inria.fr/flajolet/Publications/FlSe01.pdf>.
- [38] D. FUHRMANN, *Application of structure covariance estimation to adaptive detection*, tech. rep., Washington University, 1990.
- [39] D. R. FUHRMANN AND M. I. MILLER, *On the existence of positive-definite maximum-likelihood estimates of structured covariance matrices*, IEEE Transactions on Inform. Theory, IT-4 (1998), pp. 722–729.
- [40] V. L. GIRKO, *An introduction to statistical analysis of random arrays*, VSP, Utrecht, 1998. Translated from the Russian, Revised by the author.
- [41] G. H. GOLUB AND C. F. VAN LOAN, *Matrix computations*, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, third ed., 1996.
- [42] L. R. HAFF, *The variational form of certain Bayes estimators*, Ann. Statist., 19 (1991), pp. 1163–1190.
- [43] M. HAWKES, A. NEHORAI, AND P. STOICA, *Performance breakdown of subspace-based methods: prediction and cure*, in Proceedings of ICASSP, vol. 6, 2001, pp. 4005–4008.
- [44] S. HAYKIN AND Z. CHEN, *The cocktail party problem*, Neural Computation, 17 (2005), pp. 1875–1902.
- [45] F. HIAI AND D. PETZ, *Eigenvalue density of the Wishart matrix and large deviations*, Infin. Dimens. Anal. Quantum Probab. Relat. Top., 1 (1998), pp. 633–646.
- [46] ———, *The semicircle law, free random variables and entropy*, vol. 77, American Mathematical Society, 2000.
- [47] A. E. HOERL AND R. W. KENNARD, *Ridge regression: Biased estimation for non-orthogonal problems*, Technometrics, 12 (1970), pp. 55–67.
- [48] R. A. HORN AND C. R. JOHNSON, *Topics in matrix analysis*, Cambridge University Press, Cambridge, 1991.
- [49] K. JOHANSSON, *On fluctuations of random Hermitian matrices*, Duke Math. J., 91 (1998), pp. 151–203.
- [50] I. M. JOHNSTONE, *On the distribution of the largest eigenvalue in principal components analysis*, Annals of Statistics, 29(2) (2001), pp. 295–327.
- [51] D. JONSSON, *Some limit theorems for the eigenvalues of a sample covariance matrix*, J. of Multivar. Anal., 12 (1982), pp. 1–38.

- [52] L. KNOCKAERT, *The Barankin bound and threshold behavior in frequency estimation*, IEEE Trans. on Signal Process., 45 (1997), pp. 2398–2401.
- [53] A. KUIJLAARS AND K. T.-R. MCLAUGHLIN, *Generic behavior of the density of states in random matrix theory and equilibrium problems in the presence of real analytic external fields*, Commun. Pure Appl. Math., 53 (2000), pp. 736–785.
- [54] D. N. LAWLEY, *Tests of significance of the latent roots of the covariance and correlation matrices*, Biometrika, 43 (1956), pp. 128–136.
- [55] O. LEDOIT AND M. WOLF, *Some hypothesis tests for the covariance matrix when the dimension is large compared to the sample size*, Ann. Statist., 30 (2002), pp. 1081–1102.
- [56] ———, *Improved estimation of the covariance matrix of stock returns with an application to portfolio selection.*, Journal of Empirical Finance, 10 (2003), pp. 603–621.
- [57] ———, *Honey, I shrunk the sample covariance matrix.*, Journal of Portfolio Management, 30 (2004), pp. 110–119.
- [58] ———, *A well-conditioned estimator for large-dimensional covariance matrices*, J. Multivariate Anal., 88 (2004), pp. 365–411.
- [59] V. A. MARČENKO AND L. A. PASTUR, *Distribution of eigenvalues in certain sets of random matrices*, Mat. Sb. (N.S.), 72 (114) (1967), pp. 507–536.
- [60] R. MCAULAY AND E. HOFSTETTER, *Barankin bounds on parameter estimation*, IEEE Trans. on Inform. Theory, 17 (1971), pp. 669–676.
- [61] B. D. MCKAY, *The expected eigenvalue distribution of a large regular graph*, Linear Algebra And Its Applications, 40 (1981), pp. 203–216.
- [62] V. D. MILMAN AND G. SCHECHTMAN, *Asymptotic theory of finite-dimensional normed spaces*, vol. 1200 of Lecture Notes in Mathematics, Springer-Verlag, Berlin, 1986. With an appendix by M. Gromov.
- [63] J. A. MINGO, P. SNIADY, AND R. SPEICHER, *Second order freeness and fluctuations of random matrices: II. unitary random matrices*. To appear in Adv. in Math.
- [64] J. A. MINGO AND R. SPEICHER, *Second order freeness and fluctuations of random matrices: I. Gaussian and Wishart matrices and cyclic Fock spaces*, Journal of Functional Analysis, 235 (2006), pp. 226–270.
- [65] S. MIRON, N. LE BIHAN, AND J. MARS, *Quaternion-MUSIC for vector-sensor array processing*, IEEE Trans. on Signal Process., 54 (2006), pp. 1218–1229.

- [66] R. J. MUIRHEAD, *Aspects of Multivariate Statistical Theory*, Wiley, New York, 1982.
- [67] R. J. MUIRHEAD, *Aspects of multivariate statistical theory*, John Wiley & Sons Inc., New York, 1982. Wiley Series in Probability and Mathematical Statistics.
- [68] A. NICA AND R. SPEICHER, *Lectures on the Combinatorics of Free Probability*, London Mathematical Society Lecture Note Series, New York : Cambridge University Press, 2006.
- [69] N. L. OWSLEY, *(Dominant Mode) Minimum Variance (Rejection) Beamforming*, in *Underwater Acoustic Data Processing*, Y. T. Chan, ed., Kluwer Academic Publishers, 1989, pp. 285–291.
- [70] D. PAUL, *Asymptotics of sample eigenstructure for a large dimensional spiked covariance model*, technical report, Stanford University, 2005. <http://anson.ucdavis.edu/~debashis/techrep/eigenlimit.pdf>.
- [71] B. QUINN AND P. KOOTSOOKOS, *Threshold behavior of the maximum likelihood estimator of frequency*, *IEEE Trans. on Signal Process.*, 42 (1994), pp. 3291–3294.
- [72] C. R. RAO, *Information and the accuracy attainable in the estimation of statistical parameters*, *Bull. Calcutta Math. Soc.*, 37 (1945), pp. 81–89.
- [73] N. R. RAO, *RMTool: A random matrix and free probability calculator in MATLAB*. <http://www.mit.edu/~raj/rmtool/>.
- [74] C. D. RICHMOND, *Capon and Bartlett beamforming: Threshold effect in direction-of-arrival estimation error and on the probability of resolution*, tech. rep., Lincoln Laboratory, M.I.T., May 2005.
- [75] C. R. RICHMOND, R. R. NADAKUDITI, AND A. EDELMAN, *Asymptotic mean squared error performance of diagonally loaded Capon-MVDR processor*, in *Conference Record of the Thirty-Ninth Asilomar Conference on Signals, Systems and Computers*, 2005, 2005, pp. 1711–1716.
- [76] J. RISSANEN, *Modeling by shortest data description*, *Automatica*, 14 (1978), pp. 465–471.
- [77] B. SALVY AND P. ZIMMERMANN, *Gfun: a Maple package for the manipulation of generating and holonomic functions in one variable*, *ACM Trans. on Math. Software*, 20 (1994), pp. 163–177.
- [78] L. L. SCHARF, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*, Addison-Wesley, Reading, Massachusetts, 1991.
- [79] R. O. SCHMIDT, *Multiple emitter location and signal parameter estimation*, *IEEE Trans. on Antennas Prop.*, AP-34 (1986), pp. 276–280.

- [80] G. SCHWARTZ, *Estimating the dimension of a model*, Annals of Statistics, 6 (1978), pp. 461–464.
- [81] J. W. SILVERSTEIN, *Private communication*.
- [82] —, *The limiting eigenvalue distribution of a multivariate F matrix*, SIAM Journal on Math. Anal., 16 (1985), pp. 641–646.
- [83] —, *Random matrices and their applications*, Amer. Math. Soc., Providence, RI, 1986, ch. Eigenvalues and eigenvectors of large dimensional sample covariance matrices, pp. 153–159.
- [84] —, *On the eigenvectors of large-dimensional sample covariance matrices*, J. Multivariate Anal., 30 (1989), pp. 1–16.
- [85] —, *Weak convergence of random functions defined by the eigenvectors of sample covariance matrices*, Ann. Probab., 18 (1990), pp. 1174–1194.
- [86] —, *Strong convergence of the empirical distribution of eigenvalues of large dimensional random matrices*, J. of Multivariate Anal., 55(2) (1995), pp. 331–339.
- [87] J. W. SILVERSTEIN AND S.-I. CHOI, *Analysis of the limiting spectral distribution of large-dimensional random matrices*, J. Multivariate Anal., 54 (1995), pp. 295–309.
- [88] J. W. SILVERSTEIN AND J. W. COMBETTES, *Signal detection via spectral theory of large dimensional random matrices*, IEEE Trans. on Signal Process., 40 (1992), pp. 2100–2105.
- [89] M. SMITH AND R. KOHN, *Parsimonious covariance matrix estimation for longitudinal data*, J. Amer. Statist. Assoc., 97 (2002), pp. 1141–1153.
- [90] S. T. SMITH, *Covariance, subspace, and intrinsic Cramér-Rao bounds*, IEEE Trans. Signal Process., 53 (2005), pp. 1610–1630.
- [91] R. SPEICHER, *Multiplicative functions on the lattice of noncrossing partitions and free convolution*, Math. Ann., 298 (1994), pp. 611–628.
- [92] —, *Free probability theory and non-crossing partitions*, Sémin. Lothar. Combin., 39 (1997), pp. Art. B39c, 38 pp. (electronic).
- [93] —, *Free probability theory and random matrices*, in Asymptotic combinatorics with applications to mathematical physics (St. Petersburg, 2001), vol. 1815 of Lecture Notes in Math., Springer, Berlin, 2003, pp. 53–73.
- [94] R. P. STANLEY, *Enumerative combinatorics. Vol. 2*, vol. 62 of Cambridge Studies in Advanced Mathematics, Cambridge University Press, Cambridge, 1999. With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin.

- [95] C. M. STEIN, *Estimation of the parameters of a multivariate normal distribution, ii; estimation of a covariance matrix*. Rietz Lecture, unpublished ms., 1975.
- [96] G. W. STEWART AND J. G. SUN, *Matrix perturbation theory*, Computer Science and Scientific Computing, Academic Press Inc., Boston, MA, 1990.
- [97] G. STRANG, *The changing face of applied mathematics. An interview by Y. K. Leong*, Newsletter of the IMS, NUS, Issue 2, (2003), pp. 8–11. <http://www.ims.nus.edu.sg/imprints/interviews/GilbertStrang.pdf>.
- [98] B. STURMFELS, *Introduction to resultants*, in Applications of computational algebraic geometry (San Diego, CA, 1997), vol. 53 of Proc. Sympos. Appl. Math., Amer. Math. Soc., Providence, RI, 1998, pp. 25–39.
- [99] G. SZEGÖ, *Orthogonal Polynomials*, American Mathematical Society, Providence, 1975. 4th edition.
- [100] R. TAUSWORTHE, *A method for calculating phase-locked loop performance near threshold*, IEEE Trans. on Comms., 15 (1967), pp. 502–506.
- [101] J. THOMAS, L. L. SCHARF, AND D. TUFTS, *The probability of a subspace swap in the SVD*, IEEE Trans. on Signal Process., 43 (1995), pp. 730–736.
- [102] H. L. V. TREES, *Detection, Estimation, and Modulation Theory Part IV: Optimum Array Processing*, John Wiley and Sons, Inc., New York, 2002.
- [103] D. W. TUFTS, A. C. KOT, AND R. J. VACCARO, *The threshold effect in signal processing algorithms which use an estimated subspace*, in SVD and Signal Processing, II: Algorithms, Analysis, and Applications, R. J. Vaccaro, ed., Elsevier Science Publishers B.V. New York, 1991, ch. 19.
- [104] A. M. TULINO AND S. VERDÚ, *Random matrices and wireless communications*, Foundations and Trends in Communications and Information Theory, 1 (2004).
- [105] D. VOICULESCU, *Symmetries of some reduced free product C^* -algebras*, in Operator algebras and their connections with topology and ergodic theory (Buşteni, 1983), vol. 1132 of Lecture Notes in Math., Springer, Berlin, 1985, pp. 556–588.
- [106] —, *Addition of certain noncommuting random variables*, J. Funct. Anal., 66 (1986), pp. 323–346.
- [107] —, *Multiplication of certain noncommuting random variables*, J. Operator Theory, 18 (1987), pp. 223–235.
- [108] —, *Limit laws for random matrices and free products*, Invent. Math., 104 (1991), pp. 201–220.

- [109] D. V. VOICULESCU, K. J. DYKEMA, AND A. NICA, *Free random variables*, vol. 1 of CRM Monograph Series, American Mathematical Society, Providence, RI, 1992.
- [110] K. W. WACHTER, *The strong limits of random matrix spectra for sample matrices of independent elements*, *Annals of Probab.*, 6 (1978), pp. 1–18.
- [111] M. WAX AND T. KAILATH, *Detection of signals by information theoretic criteria*, *IEEE Trans. Acoust. Speech Signal Process.*, 33 (1985), pp. 387–392.
- [112] N. WIENER, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series. With Engineering Applications*, The Technology Press of the Massachusetts Institute of Technology, Cambridge, Mass, 1949.
- [113] E. P. WIGNER, *Characteristic vectors of bordered matrices with infinite dimensions*, *Annals of Math.*, 62 (1955), pp. 548–564.
- [114] J. WISHART, *The generalized product moment distribution in samples from a normal multivariate population*, *Biometrika*, 20 A (1928), pp. 32–52.
- [115] W. XU, A. B. BAGGEROER, AND K. L. BELL, *A bound on mean-square estimation error with background parameter mismatch*, *IEEE Trans. Inform. Theory*, 50 (2004), pp. 621–632.
- [116] W. XU, A. B. BAGGEROER, AND C. D. RICHMOND, *Bayesian bounds for matched-field parameter estimation*, *IEEE Trans. Signal Process.*, 52 (2004), pp. 3293–3305.
- [117] R. YANG AND J. O. BERGER, *Estimation of a covariance matrix using the reference prior*, *Ann. Statist.*, 22 (1994), pp. 1195–1211.
- [118] Q. T. ZHANG, *Probability of resolution of the MUSIC algorithm*, *IEEE Trans. on Signal Process.*, 43 (1995), pp. 978–987.
- [119] H. J. ZIMMERMAN, *Archives of the Research Lab of Electronics at M.I.T.* <http://www.rle.mit.edu/>.

Index

- Akaike Information Criterion, 47, 76
- algebraic curves, 143
 - singularities, 143
 - eigenvalue density, 143, 144
- algebraic power series
 - definition, 147
 - free cumulant power series, 148
 - moment power series, 148
 - polynomially recursive coefficients, 147
- Anderson, 63
- Bai-Silverstein
 - fluctuations, 190
- Biane, 161
- Capon-MVDR beamformer
 - diagonal loading, 90, 91
 - formulation, 89
- computational free probability, 149
- D-finite generating functions
 - definition, 147
 - free cumulant power series, 148
 - moment power series, 148
- D-finite power series
 - algebraic power series, 147
- eigen-inference, 33, 61
 - algorithms, 65
- eigenvalue density function
 - atomic component, 143
 - endpoint behavior, 145
 - conjecture, 145
 - non-atomic component, 144
- equilibrium measure
 - convolution, 158
- estimation
 - population eigenvalues, 65
 - known model order, 73, 76
 - unknown model order, 76
- free convolution
 - implicit encoding, 150
- free cumulant power series
 - enumeration, 147
- free probability
 - asymptotic freeness, 183
 - convolutions, 150
 - free cumulants, 184
 - polynomial method, 149
 - R transform, 185
 - random matrices, 183
 - second order freeness, 186
- freeness
 - first order, 183
 - second order, 186
- hypothesis testing
 - population eigenvalues, 65, 75, 76
- Jacobi random matrix, 153
- joint density of eigenvalues, 62
- limit eigenvalue distribution, 68
- MANOVA, 153
- Markov transition kernel, 161

- model order
 - estimation, 76
- moment power series, 67
 - enumeration, 147
- moments
 - existence
 - conjecture, 147
- Multivariate F matrix, 153

- phase transition
 - impact on inference, 82
 - population eigenvalue modelling, 83
- polynomial method
 - computational software, 139
 - eigenvectors, 161
 - free probability, 139, 149
 - MATLAB code, 139
 - operational law, 139
- polynomial recursive coefficients
 - free cumulants, 148
 - moments, 148
- polynomially recursive coefficients
 - algebraic power series, 147
 - definition, 147

- R transform, 185
- random matrices
 - compression, 156
 - fluctuations, 186
- rational power series
 - definition, 147
- resultants
 - free probability, 150

- sample covariance matrix, 61
- Speicher
 - free cumulants, 184
- spiked covariance matrix, 81
 - phase transition, 82

- Voiculescu
 - freeness, 183

- Wishart distribution, 62

- Wishart matrices
 - moments, 190
 - quadratic forms, 96

REPORT DOCUMENTATION PAGE	1. REPORT NO. MIT/WHOI 2007-01	2.	3. Recipient's Accession No.
4. Title and Subtitle Applied Stochastic Eigen-Analysis		5. Report Date February 2007	
7. Author(s) Rajesh Rao Nadakuditi		6.	
9. Performing Organization Name and Address MIT/WHOI Joint Program in Oceanography/Applied Ocean Science & Engineering		8. Performing Organization Rept. No.	
12. Sponsoring Organization Name and Address Massachusetts Institute of Technology National Science Foundation		10. Project/Task/Work Unit No. MIT/WHOI 2007-01	
		11. Contract(C) or Grant(G) No. (C) (G)	
		13. Type of Report & Period Covered Ph.D. Thesis	
15. Supplementary Notes This thesis should be cited as: Rajesh Rao Nadakuditi, 2007. Applied Stochastic Eigen-Analysis. Ph.D. Thesis. MIT/WHOI, 2007-01.		14.	
16. Abstract (Limit: 200 words) <p>The first part of the dissertation investigates the application of the theory of large random matrices to high-dimensional inference problems when the samples are drawn from a multivariate normal distribution. New parametric techniques for testing and estimation of signals in white noise are developed that exhibit robustness to high-dimensionality, sample size constraints and eigenvector misspecification.</p> <p>By interpreting the eigenvalues of the sample covariance matrix as an interacting particle system, the existence of a phase transition phenomenon in the largest ("signal") eigenvalue is derived using heuristic arguments. This exposes a fundamental limit on the identifiability of low-level signals due to sample size constraints when using the sample eigenvalues alone.</p> <p>The analysis is extended to address a problem in sensor array processing, posed by Baggeroer and Cox, on the distribution of the outputs of the Capon-MVDR beamformer when the sample covariance matrix is diagonally loaded.</p> <p>The second part of the dissertation investigates the limiting distribution of the eigenvalues and eigenvectors of a broader class of random matrices. A powerful method is proposed that expands the reach of the theory beyond the special cases of matrices with Gaussian entries; this simultaneously establishes a framework for computational (non-commutative) "free probability" theory.</p>			
17. Document Analysis a. Descriptors Random matrices high-dimensional inference signal processing			
b. Identifiers/Open-Ended Terms			
c. COSATI Field/Group			
18. Availability Statement Approved for publication; distribution unlimited.		19. Security Class (This Report) UNCLASSIFIED	21. No. of Pages 204
		20. Security Class (This Page)	22. Price