

SPECTRAL BARRIERS IN CERTIFICATION PROBLEMS

by

Dmitriy Kunisky

A DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT

OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

DEPARTMENT OF MATHEMATICS

NEW YORK UNIVERSITY

MAY, 2021

Professor Afonso S. Bandeira

Professor Gérard Ben Arous

© DMITRIY KUNISKY

ALL RIGHTS RESERVED, 2021

DEDICATION

To my grandparents.

ACKNOWLEDGEMENTS

First and foremost, I am immensely grateful to Afonso Bandeira for his generosity with his time, his broad mathematical insight, his resourceful creativity, and of course his famous reservoir of open problems. Afonso's wonderful knack for keeping up the momentum—for always finding some new angle or somebody new to talk to when we were stuck, for mentioning “you know, just yesterday I heard a very nice problem...” right when I could use something else to work on—made five years fly by. I have Afonso to thank for my sense of taste in research, for learning to collaborate, for remembering to enjoy myself, and above all for coming to be ever on the lookout for something new to think about.

I am also indebted to Gérard Ben Arous, who, though I wound up working on topics far from his main interests, was always willing to hear out what I'd been doing, to encourage and give suggestions, and to share his truly encyclopedic knowledge of mathematics and its culture and history. On many occasions I've learned more tracking down what Gérard meant with the smallest aside in the classroom or in our meetings than I could manage to get out of a week on my own in the library. I couldn't have asked for a better role model of deep expertise, precise thinking, and genuine scholarship.

Thanks also to the other faculty mentors I was fortunate to have at Courant: Carlos Fernandez-Granda at the beginning of my PhD, whose patience with my endless requests for advice on an early project helped me find my footing and grow to be confident in my ideas, and Jon Niles-Weed at the end, who brought his great energy in full force to the department

and helped me stay on track at a difficult time. In the same breath I must also mention Cris Moore, who most generously arranged a visit at the beginning of my degree when I knew less than nothing about anything, and gently guided me over the years into his terrific interdisciplinary community. Thanks to Cris for administering another dose of spectacular taste in research, imparting a major appreciation of the gastronomy of Santa Fe, sharing his physicist's enthusiasm for a pretty calculation, and impressing upon me his conviction that a scientist is, always and first of all, a citizen.

Thanks to Oded Regev and Michael Overton, two of the most thoughtful teachers I've had the honor of learning from over the years, for being on my thesis committee. It's a pleasure to share my work with you.

My cohort at Courant and other collaborators over the years have taught me so many little bits of wisdom, tricks, tools, ideas big and small, and I am grateful for having had the opportunity to work or banter with all of them. A separate thanks to Alex Wein, Yunzi Ding, and Jess Banks, with whom I worked on the material that has ended up in several chapters of this thesis. All that work would have been impossible without their many insights and suggestions.

Thanks to my friends in New York, for their chorus of "so...how's the math?" and their knowing smiles when I would describe once again how it's "kind of related" to something "that's like something else that's somewhat useful," for their gently steering me away from work from time to time, and for many blissfully idle evenings. And to my family, their cheerful kindness, their love and support, their effortless confidence that I was making the right decisions, their sayings and jokes and stories around the kitchen table. I'm made of those stories, and that kitchen table is home for me.

Thanks finally to Eleanor: for all these happy, bright, and weightless years together, the adventures and travels and the quiet sturdy times in between, and for knowing just when to take me by the hand and lead me far away from all this nonsense.

ABSTRACT

The related tasks of *certifying* bounds on optimization problems and *refuting* unsatisfiable systems of constraints have a long history in computer science and optimization, and deep mathematical roots in the proof complexity of algebraic systems. From the algorithmic perspective, these problems differ fundamentally from ordinary optimization in that they ask not merely for a single high-quality solution, but rather for a simultaneous bound on the quality of all possible solutions. The computational complexity of certification and its connections with and distinctions from the complexity of optimization, especially in the average case, remain poorly understood for many problems.

The purpose of this thesis is to study the average-case complexity of certification for *constrained principal component analysis (PCA)* problems, eigenvalue-like problems optimizing quadratic forms over sets of structured vectors or low-rank matrices. As one notable example, several of our results concern the much-studied Sherrington-Kirkpatrick (SK) Hamiltonian, a random function over the hypercube originating in the statistical physics of spin glasses. In problems of this kind, a natural certification strategy is a bound by the spectral norm of the relevant coefficient matrix. We provide evidence for *spectral barriers* in constrained PCA problems, supporting the conjecture that, for reasonable instance distributions, no efficient certification algorithm can improve on the spectral bound.

In the first part of the thesis, we develop tools for proving reductions from certification over instances drawn from the Gaussian orthogonal ensemble (GOE) to hypothesis testing

in spiked matrix models. These models describe tiltings of the eigenspaces of the GOE towards *planted* constrained PCA solutions. We then provide evidence that these testing problems are hard via lower bounds against algorithms computing low-degree polynomials of the observations. In doing so, we develop new techniques for working with the low-degree likelihood ratio. Namely, we show that its L^2 norm, which governs the efficacy of low-degree tests, is in many models an expectation of a function of a scalar *overlap* or inner product of two draws of the random signal observed in the model. We use this to give a simple, unified account of old and new low-degree lower bounds, with tight results for testing in models including tensor PCA, Wigner and Wishart matrix PCA of arbitrary rank, sparse PCA, and the stochastic block model. Using our results for Wishart models, we deduce the conditional hardness of better-than-spectral certification for the SK Hamiltonian, the Potts spin glass Hamiltonian, and non-negative PCA.

In the second part, we investigate lower bounds against the *sum-of-squares* hierarchy of semidefinite programming relaxations in the hypercube setting of the SK Hamiltonian. We first study the structure of the degree 4 relaxation, giving constraints on the spectra of pseudomoments and deterministic examples from finite frame theory. We then propose a general means of extending degree 2 lower bounds to higher degrees, an alternative to the pseudocalibration framework of Barak et al. (2019). We show that this technique is exact for the deterministic lower bound of Grigoriev (2001) and Laurent (2003), and prove a conjecture of Laurent’s on the spectrum of the associated pseudomoments. Finally, we apply our extension to the SK Hamiltonian and give tight lower bounds for the degree 4 and 6 relaxations.

Contents

Dedication	iii
Acknowledgments	iv
Abstract	vi
List of Figures	xiv
List of Tables	xv
List of Notations	xvi
1 Introduction	1
1.1 Search and Certification	1
1.2 Average-Case Computational Complexity	6
1.3 A Motivating Example	10
1.4 Spectral Barriers	15
1.5 Summary of Contributions	17
1.5.1 Reductions, Planting, and Low-Degree Polynomials	18
1.5.2 The Sum-of-Squares Hierarchy	19
I Computationally-Quiet Planting and Low-Degree Polynomials	21
2 Spectral Planting and Reductions to Hypothesis Testing	22

2.1	Hypothesis Testing	24
2.1.1	Statistical-Computational Gaps	26
2.2	Pushout and Computationally-Quiet Planting	27
2.3	Tools for Reductions Under the GOE	34
2.4	Hard Regime in the Wishart Spiked Matrix Model	38
2.5	An Example of Spectral Planting in Graphs	40
3	Preliminaries on Low-Degree Polynomial Algorithms	44
3.1	Le Cam’s Contiguity and the Second Moment Method	45
3.2	Two Roads to the Low-Degree Likelihood Ratio	51
3.2.1	Computationally-Bounded Second Moment Method	51
3.2.2	Appearance in Sum-of-Squares Lower Bounds	54
3.3	Consequences of Low-Degree Lower Bounds	57
4	Low-Degree Likelihood Ratio Analysis: Overlap Formulae	60
4.1	Gaussian Wigner Models	62
4.1.1	Hermite Polynomials	66
4.1.2	Computing the Low-Degree Likelihood Ratio	68
4.2	Exponential Families and Variance Functions	71
4.3	NEF-QVF Models	76
4.3.1	Orthogonal Polynomials in NEF-QVFs	77
4.3.2	Computing the Low-Degree Likelihood Ratio	79
4.3.3	Channel Monotonicity	86
4.4	Gaussian Wishart Models	87
4.4.1	Hermite Polynomials, Revisited	91
4.4.2	Computing the Low-Degree Likelihood Ratio	94
5	Low-Degree Likelihood Ratio Analysis: Lower Bounds	101

5.1	Tools for Scalar Overlap Analysis	103
5.1.1	Well-Behaved Spike Priors	105
5.2	Gaussian Wigner and Wishart Models	112
5.2.1	Bounding Wishart by Wigner	112
5.2.2	Truncated Exponential Polynomials	114
5.2.3	Tensor PCA	115
5.2.4	Dense Matrix PCA	120
5.2.5	Sparse Matrix PCA	125
5.3	Consequences for Certification	128
5.4	NEF-QVF Models	133
5.4.1	Stochastic Block Model	133
5.4.2	A Non-Gaussian Matrix PCA Model	136
II	A Geometric Perspective on Sum-of-Squares Lower Bounds	148
6	Basic Notions and the Gramian Viewpoint	149
6.1	Sum-of-Squares over the Hypercube	149
6.2	The Geometry of Pseudomoments	156
7	Degree 4 Deterministic Structure and Examples	158
7.1	Gramian Description of Degree 4 Extensions	159
7.2	Constraints on Pseudomoment Extensions	171
7.3	Integrality and Separability	184
7.4	Examples from Equiangular Tight Frames	192
7.5	New Inequalities Certifiable at Degree 4	200
8	Spectral Pseudomoment Extensions	206
8.1	Notations and Assumptions	207

8.2	Gaussian Conditioning Interpretation at Degree 4	.208
8.3	Generalizing to Higher Degree	.215
8.3.1	Symmetric Tensors and Homogeneous Polynomials	.215
8.3.2	Surrogate Random Tensor Construction	.219
8.3.3	Homogeneous Ideals and Multiharmonic Polynomials	.221
8.3.4	Conditioning by Translating to Homogeneous Polynomials	.222
9	Grigoriev-Laurent Lower Bound	231
9.1	Pseudomoment Construction	.233
9.2	Representation Theory of the Symmetric Group	.236
9.2.1	Combinatorial Interpretations of Characters	.238
9.2.2	Combinatorial Class Functions	.239
9.2.3	The Simplex-Harmonic Representation	.246
9.3	Proof of Positivity	.249
9.4	Pseudomoment Spectrum and Laurent’s Conjecture	.251
10	Sum-of-Forests Pseudomoments and Lifting Theorems	261
10.1	Concrete Pseudomoment Extension	.262
10.1.1	Generalized Maxwell-Sylvester Representations	.263
10.1.2	Closing the Prediction	.265
10.1.3	Sum-of-Forests Pseudomoments	.270
10.2	Lifting 1: High Rank to High Degree	.275
10.3	Partial Ordering and Möbius Function of $\mathcal{F}(m)$.278
10.3.1	Möbius Functions of Partially Ordered Sets	.279
10.3.2	The Compositional Ordering	.282
10.4	Pseudomoment and Contractive Graphical Matrices	.288
10.5	Main and Error Terms	.291

10.6	Spectral Analysis in the Harmonic Basis: Outline of Theorem 10.2.3	299
10.7	Approximate Block Diagonalization: Towards Lemma 10.6.2	301
10.7.1	Stretched Forest Ribbon Diagrams	302
10.7.2	Tying Bound and Bowtie Forest Ribbon Diagrams	306
10.7.3	Simplification of Coefficients	310
10.8	Approximate Gram Factorization: Proof of Lemma 10.6.2	318
10.8.1	Partition Transport Plans and Ribbon Diagrams	319
10.8.2	Tying Bound	323
10.8.3	Simplification of Coefficients	328
10.9	Bound on the Error Term: Proof of Lemma 10.6.3	334
10.10	Obstacles to Low-Rank Extensions	339
10.10.1	Order-2 Tensor Orthonormality: The ϵ_{pow} Obstacle	340
10.10.2	Sided Error Terms: The ϵ_{err} Obstacle	342
10.11	Lifting 2: Low Rank to Low Degree	343
10.12	Miscellaneous Bounds	355
10.12.1	Matrix Inequalities	355
10.12.2	Combinatorial Bounds	357
10.13	Calculus of Contractive Graphical Matrices	361
10.13.1	Connected Components and Tensorization	362
10.13.2	Splitting	364
10.13.3	Pinning, Cutting, and Direct Sum Decomposition	365
10.13.4	Factorization	366
10.13.5	General-Purpose Norm Bounds	368
11	Applications of Lifting Theorems	373
11.1	Pattern Diagrams	374

11.2	Warmup 1: Grigoriev-Laurent Pseudomoments376
11.3	Warmup 2: Lifting Random High-Rank Projectors380
11.4	Degree 6 Lower Bound for Sherrington-Kirkpatrick Hamiltonian392
11.5	Tensorial Sample Covariance Matrices397
11.5.1	Gaussian Tensors: Proof of Lemma 11.4.2398
11.5.2	Haar Tensors401
A	Open Problems	410
A.1	Low-Degree Method Beyond Integrable Models410
A.2	Overlaps, Low-Degree Polynomials, and Statistical Physics411
A.3	Channel Monotonicity411
A.4	Fisher Information–Orthogonal Polynomial Identities412
A.5	Can Low-Degree Polynomials Explore?412
A.6	Sum-of-Squares and Entanglement413
A.7	Generalized Maxwell-Sylvester Representations414
A.8	Pseudocalibration Reconciliation414
A.9	Replicated Sum-of-Squares415
	Bibliography	416

List of Figures

2.1	Phase diagram of the Wishart spiked matrix model	39
4.1	Nonlinearities in low-degree likelihood ratio norm for NEF-QVFs	81
5.1	Low-degree lower bounds: dense matrix and tensor PCA	121
5.2	Hyperbolic secant and Laplace distribution densities	137
6.1	Cut polytope and generalized elliptopes in low dimension	154
10.1	Forests, polynomials, and coefficients in spectral pseudomoment extension	273
10.2	Ordering in the compositional partially ordered set $\mathcal{F}(m)$	283
10.3	Odd merge of partitions	304
10.4	Partition transport plan and ribbon diagram	321
10.5	Tying stretched forest and partition transport ribbon diagrams	323
10.6	Extra term from size 4 partition transport ribbon diagram	341
10.7	Diagrammatic manipulation of degree 4 pseudoexpectation adjustment	350

List of Tables

4.1	Summary of overlap formulae	61
4.2	Exponential families and variance functions	74

List of Notations

BBP	Baik–Ben Arous–Péché (transition)	29
CSP	Constraint satisfaction problem	6
ETF	Equiangular tight frame	192
GOE	Gaussian orthogonal ensemble	11
LDLR	Low-degree likelihood ratio	52
LP	Linear program(ming)	1
LR	Likelihood ratio	46
NEF	Natural exponential family	72
NEF-QVF	Natural exponential family with quadratic variance function	73
PCA	Principal component analysis	15
psd	Positive semidefinite	150
SDP	Semidefinite program(ming)	5
SK	Sherrington-Kirkpatrick	12
SOS	Sum-of-squares	5
UNTF	Unit-norm tight frame	192
$\mathbb{1}\{A\}, \mathbb{1}_A$	Indicator function	29
\mathbb{N}	Natural numbers including zero, $\{0, 1, 2, \dots\}$	65
\mathbb{N}_+	Natural numbers excluding zero, $\{1, 2, 3, \dots\}$	24
$[n]$	$\{1, 2, \dots, n\}$	27
$\binom{S}{d}$	Subsets of S of size d	137

$\binom{S}{\leq d}$	Subsets of S of size at most d	152
$\mathcal{M}(S)$	Multisets of elements of S	216
$\mathcal{M}_d(S)$	Multisets of elements of S of size d	216
$\mathcal{M}_{\leq d}(S)$	Multisets of elements of S of size at most d	216
$\text{Part}(S)$	Partitions of S	90
$\text{Plans}(\sigma, \tau)$	Partition transport plans (generalized permutations)	319
$\mathbb{R}_{\text{sym}}^{n \times n}$	Symmetric matrices	10
$\mathbb{R}[\mathbf{x}_1, \dots, \mathbf{x}_n]_{\leq d}$	Polynomials of degree at most d	52
$\mathbb{R}[\mathbf{x}_1, \dots, \mathbf{x}_n]_d^{\text{hom}}$	Homogeneous polynomials of degree d	215
\mathbb{S}^{n-1}	Unit sphere in \mathbb{R}^n	29
$\mathcal{P}(\Omega)$	Probability measures on Ω (with Borel σ -algebra)	27
$\text{Ber}(p)$	Bernoulli distribution	110
$\text{ER}(n, p)$	Erdős-Rényi random graph	6
$\mathcal{G}_d^{\text{poly}}(n, \sigma^2)$	Standard Gaussian degree d homogeneous polynomial	218
$\mathcal{G}_d^{\text{tens}}(n, \sigma^2)$	Standard Gaussian symmetric d -tensor	218
$\text{GOE}(n)$	Gaussian orthogonal ensemble, $n \times n$	11
$\text{Haar}(\Omega)$	Haar measure	401
$\text{Reg}(n, d)$	Random d -regular graph, n vertices	11
$\text{Unif}(\Omega)$	Uniform distribution over Ω	29
\mathbf{I}_n	Identity matrix, $n \times n$	14
$\ \mathbf{A}\ _F$	Matrix Frobenius norm	16
$\langle \mathbf{A}, \mathbf{B} \rangle$	Matrix Frobenius / Hilbert-Schmidt inner product	35
$\mathbf{A} \circ \mathbf{B}$	Matrix Hadamard (entrywise) product	181
$\mathbf{A}^{\circ k}$	Matrix Hadamard (entrywise) power	181
\mathbf{A}^Γ	Block matrix partial transpose	172
$\sigma(\mathbf{A})$	Vector of singular values	37

$\text{diag}(\mathbf{A})$	Diagonal vector	150
$\text{isovec}(\mathbf{A})$	Isometric vectorization of symmetric matrix	212
$\text{offdiag}(\mathbf{A})$	Off-diagonal vectorization of symmetric matrix	212
$\text{vec}(\mathbf{A})$	Vectorization (column-major order)	161
$\mathcal{O}(n)$	Orthogonal matrix group, $n \times n$	186
\mathbf{P}_V	Orthogonal projection matrix to V	194
$\mathbf{1}_n$	All-ones vector, length n	150
$\ \mathbf{a}\ $	Vector ℓ^2 norm	13
$\ \mathbf{a}\ _p$	Vector ℓ^p norm	37
$\text{diag}(\mathbf{a})$	Diagonal matrix	185
$\text{set}(\mathbf{a})$	Set of distinct entries	276
$\langle p, q \rangle_\circ$	Polynomial apolar inner product	217
$\mathbf{M}(\mathbf{W})$	Constrained PCA problem over $\{\pm 1\}^n$	149
$\mathbf{M}_\mathcal{X}(\mathbf{W})$	Constrained PCA problem over \mathcal{X}	15
n	Ambient dimension of constrained PCA problem	10
N	Number of samples in Wishart spiked matrix model	31
\mathbf{P}_*	Parisi number	11
\mathbf{W}	Gaussian orthogonal ensemble matrix	10
\mathbf{x}, \mathbf{X}	Constrained PCA feasible point; spike(s) in spiked model	10
\mathcal{X}	PCA constraint set	15
L_n	Likelihood ratio (of hypothesis testing problem)	48
$L_n^{\leq D}$	Low-degree likelihood ratio	52
\mathcal{P}_n	Spiked model spike prior	31
$\tilde{\mathcal{P}}_n$	Spiked model signal prior	87
\mathbb{P}_n	Hypothesis test planted / alternative distribution	24
\mathbb{Q}_n	Hypothesis test null distribution	24

$(\mathbb{P}_n) \triangleleft (\mathbb{Q}_n)$	Contiguity	48
p, q	Hypothesis test outputs	24
$\tilde{\mathbf{X}}$	Spiked model signal	87
\mathbf{Y}	Hypothesis test observation	24
β	Signal strength in spiked Wishart model	31
γ	Dimension-to-samples ratio in spiked Wishart model	31
λ	Signal-to-noise ratio in spiked models	115
$\phi_D^{\text{Mor}}, \phi_D^{\text{Wig}}, \phi_{N,D}^{\text{Wish}}$	Link functions	65
$\mathcal{B}^{n,r}$	Degree 4 Gram witness matrices	160
$\mathcal{B}_{\text{sep}}^{n,r}$	Separable degree 4 Gram witness matrices	187
\mathcal{C}^n	Cut polytope	150
\mathcal{E}_{2d}^n	Generalized elliptope of degree $2d$	151
$\tilde{\mathbb{E}}$	Pseudoexpectation	54
$\mathcal{F}(m), \mathcal{T}(m)$	Good forests and trees	272
\mathbf{M}	Degree 2 pseudomoment matrix	150
$\text{SOS}_{2d}(\mathbf{W})$	Hypercube sum-of-squares relaxation of degree $2d$	151
$\mathbf{v}_1, \dots, \mathbf{v}_n$	Degree 2 pseudomoment matrix Gram vectors	156
V^\bullet, V^\square	Diagram vertices	289
\mathbf{V}	Degree 2 pseudomoment matrix Gram factor	161
Z^F	Contractive graphical scalar	270
\mathbf{Z}^F	Contractive graphical matrix	290
$\kappa, \kappa_L, \kappa_R$	Diagram leaf labelling	270
$\mu_{\mathcal{P}}(x, y)$	Möbius function of partially ordered set \mathcal{P}	279

1 | INTRODUCTION

To describe it in detail would be a pleasure.

—Samuel Beckett, *Molloy*

1.1 SEARCH AND CERTIFICATION

Mathematical optimization, and combinatorial optimization in particular, have been intertwined with computational complexity since the very origins of both fields, Kantorovich studying the economic applications of linear programming (LP) [Kan39] just as Turing developed his machines and so the theory of computability [Tur37]. The two developed in tandem, with many of the classical hard problems of the theory of algorithms identified in the NP complexity class [Coo71, Kar72, Lev73] already having a long history in the theory and practice of optimization [Kar86, Sch05]. These problems, in their practical manifestations, were *search* problems. Given a mathematical object, find a near-optimal structure, they asked: given a graph, find a large cut; given distances between locations, find a short tour; given a set of points, find the furthest point in some direction.

The theory of computational complexity, however, often preferred *decision* problems: given a graph, determine whether or not (answering just “yes” or “no”) there exists a large cut, and so forth. This was perhaps a matter of convenience, as decision problems have

simpler outputs less specific to the problem at hand.¹ But, in fact, such problems already had a rich mathematical history, whose connections to computation slowly came to light and eventually gave rise to new practical algorithms. A prototypical example of a decision problem hiding in classical mathematics is Hilbert’s Nullstellensatz [Hil93]: either $f_1, \dots, f_k \in \mathbb{C}[z_1, \dots, z_n]$ have a common zero, or there are $g_1, \dots, g_k \in \mathbb{C}[z_1, \dots, z_n]$ such that $\sum_{i=1}^k f_i g_i = 1$, a polynomial sentence showing that the f_i *cannot* have a common zero. Precursors abound: Bézout’s identity over the integers, for instance, states that either $a, b \in \mathbb{Z}$ have a non-trivial common divisor, or there exist $x, y \in \mathbb{Z}$ such that $ax + by = 1$; such ideas are even apparent in the ancient so-called Chinese remainder theorem. The philosophy of these results is, towards deciding if an object exists (a common zero, a common divisor), to identify precisely what obstructions can prevent it from existing.

Gradually, the idea that such results could form the foundation for *algorithms* solving the underlying decision problems—do the f_i have a common zero or not—gained currency. For the Nullstellensatz, for instance, one may bound the degree of the g_i involved, thereby reducing the question of their existence to a (potentially very large) linear system. The first result on an “effective Nullstellensatz” giving bounds on the degree of the g_i and thereby indirectly studying the efficacy of this strategy is due to Hermann in 1926 ([Her98] is a modern translation).² Improvements followed [Bro87, Kol88], as well as degree lower bounds [MM82, BIK⁺96]. The question of polynomial-time algorithms—perhaps more efficient ones than bounding the degree and solving a linear system—was proposed as a fundamental one and related to $P \stackrel{?}{=} NP$ by Shub and Smale [SS94, Sma98]. These ideas joined a rich and active literature on “proof complexity” of various statements in different restricted, often computationally-tractable, proof systems; important results adjacent to those above include [CR79, GV01, Raz01, BSW01, GHP02, Raz02, Raz03, Ber18]. Often the decision problems one

¹For example, Levin’s version of this theory in [Lev73] was developed in terms of search problems.

²As the translation discusses, this remarkable paper makes use of computational language a decade before Turing, already recognizing the algorithmic implications of an effective Nullstellensatz.

solves in this way are called *refutation* problems, as above we prove that some structure cannot exist by manipulating its algebraic description into a contradiction.

However, this approach does not seem entirely adapted to the quantitative optimization problems that we began by considering. The techniques do not fit on two counts. First, while one may, for example, look for a large cut in a graph by repeatedly solving algebraic formulations of “there exists a cut of size exactly k ” and performing a search over k , this technique seems eminently wasteful and subject to subtleties of parity or integrality as k varies.³ Instead, we would like to somehow replay the idea of the effective Nullstellensatz while actually being able to *optimize* a quantity rather than only check if it may have a specific value. Second, and what is more material, often in applications one wishes to include inequality constraints, such as the simple ones appearing in LPs, perhaps together with a more complicated polynomial objective or some non-convex polynomial constraints.

The first steps hinting at a middle ground between proof systems and quantitative optimization came from practitioners of optimization. Knowing that LPs could be solved efficiently, the idea arose to formulate combinatorial problems as LPs by ignoring their constraints that variables be integer-valued—so-called *relaxation* of the constraints—and then gradually add constraints that integer solutions must satisfy. This iterative form of relaxation was first proposed in a general form by Gomory in 1958 [Gom10] following prior results for special cases. Gomory crystallizes the key idea:

“If the solution is not in integers, ingenuity is used to formulate a new constraint that can be shown to be satisfied by the still unknown integer solution but not by the noninteger solution already attained...What has been needed to transform this procedure into an algorithm is a systematic method for generating the new constraints.”

³For other problems, such as finding a clique of size at least k in a graph, there is a monotonicity (whenever there exists a clique of size k then there also exists one of any size $k' \leq k$) that makes this approach more plausible.

Different ways of systematizing this were then proposed by Chvátal [Chv73], and, in what became the dominant approaches, by Sherali and Adams [SA90] and Lovász and Schrijver [LS91], both groups proposing hierarchies of algorithms *lifting* problems to larger sets of variables, relaxing some resulting integrality constraints, and then *projecting* the solution of this relaxation by discarding the extra variables. These hierarchies blur the boundary between proof systems and optimization techniques. On the one hand, they yield bounds on the objective function, so they refute the existence of a feasible point with objective value superior to that bound; also, they may be viewed as manipulating proofs with certain inequalities derived from integrality of their variables. On the other hand, they are also merely certain large LPs, and so may be studied as convex optimizations. For instance, the latter viewpoint leads to the notion of *rounding* relaxed solutions—the variables from the lift-and-project procedure, which is the dual formulation to that where proofs are assembled by combining available inequalities—to genuine feasible solutions to the original problem. In this way these algorithms are directly usable for search as well as refutation. All of this remains in the toolbox of modern algorithm design; see the surveys [Tul10, CT12] for a theoretical perspective. We will refer to algorithms that yield bounds (upper bounds, in all cases we consider) on optimization problems as *certification* algorithms, and it is these algorithms and their limitations that are the subject of this thesis.

The specific certification algorithms we study, however, go one step further and come from a final line of work that again dates back to Hilbert, in this case to his seventeenth problem: can every non-negative rational function on \mathbb{R}^n be expressed as a sum of squares of rational functions? Hilbert asked about rational functions because he had shown that, in general, there are non-negative polynomials that cannot be expressed as sums of squares of polynomials; Hilbert’s proof was non-constructive, but Motzkin later gave a concrete example [Mot67]. While Artin eventually answered Hilbert’s question in the affirmative [Art27], what we will be more concerned with are refinements of Hilbert’s initial failed attempt at

a Nullstellensatz-like statement for non-negative polynomials. These so-called *Positivstellensätzen* gave conditions on when systems of equality and inequality constraints in polynomials admit Nullstellensatz-like refutations: they state that, under various conditions, if $f_1, \dots, f_k, g_1, \dots, g_\ell \in \mathbb{R}[x_1, \dots, x_n]$, then there exists no $\mathbf{x} \in \mathbb{R}^n$ with $f_i(\mathbf{x}) \geq 0$ and $g_j(\mathbf{x}) = 0$ if and only if there exist $s_{i,a}, r_j \in \mathbb{R}[x_1, \dots, x_n]$ such that

$$\sum_{i=1}^k \left(\sum_{a=1}^m s_{i,a}(\mathbf{x})^2 \right) f_i(\mathbf{x}) + \sum_{j=1}^{\ell} r_j(\mathbf{x}) g_j(\mathbf{x}) = -1. \quad (1.1)$$

There is a deep literature on such results; see [Ste74, Sch91, Put93, Rez95, Rez00, BPT12, Sch17] and numerous references therein.

The relevance of this to the optimization problems we began with is that, in a minor adjustment of the above, if we seek to maximize some $F \in \mathbb{R}[x_1, \dots, x_n]$ over the given constraint set (rather than showing that the constraint set is empty), we may consider

$$\begin{aligned} & \text{minimize} && c \\ & \text{subject to} && c - F(\mathbf{x}) = \sum_{i=1}^k \left(\sum_{a=1}^m s_{i,a}(\mathbf{x})^2 \right) f_i(\mathbf{x}) + \sum_{j=1}^{\ell} r_j(\mathbf{x}) g_j(\mathbf{x}). \end{aligned} \quad (1.2)$$

The given condition is a so-called *sum-of-squares (SOS) proof* that $F(\mathbf{x}) \leq c$ on the given constraint set. In a crucial insight, Lasserre [Las01] and Parrilo [Par00, Par03] concurrently observed that such problems, so long as the degrees involved are bounded (as in the premise of the effective Nullstellensätzen), can be solved efficiently with *semidefinite programming (SDP)*.⁴ This yields the *SOS hierarchy* (sometimes called the *Lasserre hierarchy*) of certification or refutation algorithms, graded by the degree allowed in the above equation.

⁴This approach can also be used to treat real-valued refutation problems instead of Nullstellensatz-based systems, and is strictly more powerful. Also, Lovász and Schrijver in [LS91] gave an SDP-based precursor in addition to the LP hierarchy mentioned above, and nascent versions of this idea were also present in [Sho87, Nes00]. It is Lasserre's variant of the SOS hierarchy that follows the lift-and-project formalism, while Parrilo's is the dual we give in (1.2).

SOS relaxations are, at least when measuring computational cost coarsely, as effective as both the LP hierarchies mentioned above and other SDP hierarchies; see [Lau03a] for specific comparisons for integer programming as well as the remarkable result on optimality of SOS among a general class of SDP relaxations of [LRS15]. Thus we will take lower bounds against the SOS hierarchy as the “gold standard” of difficulty of certification, which is at this point a widely-held consensus. See also the surveys [BS14, Moi20] for more specific discussion of various other lines of work that support this position.

1.2 AVERAGE-CASE COMPUTATIONAL COMPLEXITY

The other key aspect of the setting we will work in is that we will consider *random* instances of optimization problems. Such questions were proposed at least concurrently with the foundational results on worst-case computational complexity [Kar76, Kar86]. The most immediate reason to wonder if algorithms perform well on random instances of a problem is to avoid a false impression of intractability from contrived worst-case instances. One prominent example is that of the simplex algorithm for LP: first introduced by Dantzig in 1947 (see [Dan65]), the simplex method remains an extraordinarily successful practical algorithm. However, Klee and Minty in 1972 produced a sequence of instances on which it requires exponential time to terminate [KM72]. To reconcile the practical picture with the theoretical one, a series of works showed that the simplex method terminates quickly on various random models; e.g., the line of work of [Bor82, Sma83, Bor88] treated rotationally-invariant distributions of the data. Finally, [ST04] introduced the weaker notion of *smoothed analysis* where only a small random perturbation is made to input data, noting that i.i.d. distributions and the like seldom have much relevance to practical applications, and again proved that the simplex method converges in polynomial time under such a random model.

While such situations are the practical motivation for average-case analysis, even for

unrealistically uniform distributions of inputs, theorists were quick to observe that average-case behavior exhibits intriguing new phenomena. Perhaps the main theoretical convenience of average-case analysis for optimization problems specifically is that they allow a very concrete assessment of performance that behaves quite differently from the “approximation ratios” considered in the worst case. Namely, when a probability distribution over problem instances is introduced, the true value of an optimization problem—in expectation, we may usually safely say, since strong concentration often holds for large random problems by generic arguments—is a single number. For example, Erdős-Rényi graphs on n vertices with edge probability $\frac{1}{2}$ (which we denote $ER(n, p = \frac{1}{2})$) typically have largest cliques of size approximately $2 \log_2 n$. How close can a search algorithm reach to this number? Karp’s analysis in [Kar76] led him to conjecture that *no* efficient algorithm could typically find a clique of size $(1 + \epsilon) \log_2 n$; the same was confirmed for other search techniques by [Jer92]. Similar *barriers* to search have been observed in various random optimization problems, especially many of the classical discrete constraint satisfaction problems (CSPs) including graph coloring [GM75, COKV07b], largest independent set in graphs [COE15, GS14], and various abstract CSPs such as satisfiability (SAT) and its variants [MMZ05, ART06, KMRT⁺07, COKV07a, ACO08].

While conjectures such as Karp’s above were made on the basis of analysis of some straightforward algorithms and perhaps their most immediate improvements, various more systematic frameworks have since emerged for analyzing the average-case complexity of search problems. Some of these heuristics for hardness are based on the putative optimality of various classes of algorithms, while others consider other properties of the optimization landscape.⁵ They include:

- failure of Markov chain Monte Carlo methods [Jer92, DFJ02];

⁵Some heuristic approach or restricted model of computation appears necessary to make progress on such matters, as, even assuming $P \neq NP$ or similar complexity-theoretic conjectures, actual proofs of average-case hardness seem far out of reach.

- failure of local algorithms [GS14, DM15a, BGJ20, CGPR19];
- failure of low-degree polynomial algorithms [GJW20, Wei20];
- failure of approximate message passing variants and related analysis with methods of statistical physics [Mon18, AMS20, AM20];
- lower bounds against circuit models of computation [Ros10, Ros14];
- structural properties of the solution space and “shattering” of level sets [ACO08, KMRT⁺07, GS14, GZ19];
- geometric analysis of critical points and dynamics on non-convex optimization landscapes [ABAČ13, MKUZ19].

Two more classes of random computational problems often considered—no longer optimization problems—are motivated instead by statistics. Here, we either draw a problem instance from a distribution with a “planted” structure that we wish to recover (say, a large clique in a graph, or a large principal direction in a matrix), or we observe an instance from either such a planted distribution or a “null” distribution with no planted structure and must decide which distribution the observation came from. The former is called *estimation* in statistics and *recovery* in some more recent machine learning literature, and likewise the latter is called either *hypothesis testing* (or simply testing) or *detection*. For these problems there are further tools, some related to the above and some entirely different:

- failure of low-degree polynomial algorithms [HKP⁺17, HS17, BKW20b, DKWB19, SW20, BBK⁺20, DHS20] (see also Chapter 3 for an overview based on the notes [KWB19], as we will use these results later);
- failure of the local statistics hierarchy of semidefinite programs [BMR21, BBK⁺20];

- methods from statistical physics which suggest failure of belief propagation or approximate message passing algorithms [DKMZ11b, DKMZ11a, LKZ15a, LKZ15b] (see [ZK16] for a survey or [BPW18] for expository notes);
- geometric analysis of the optimization landscape of the maximum likelihood estimator (see “geometric analysis” entries above);
- reductions from the planted clique model (which has become a “canonical” problem believed to be hard in the average case) [BR13, HWX15, WBS16, BB19b, BB19a, BB20];
- lower bounds in the statistical query model [Kea98, KS07, FGR⁺17, FPV18, KV16, DKS17].

That all concerns search, testing, and estimation. Here we will instead ask: what are the barriers to *certification* for random optimization problems? In contrast to the wealth of resources above, for certification there is essentially one plan: to prove lower bounds against various hierarchies of convex relaxations, the SOS hierarchy if possible or the Sherali-Adams or Lovász-Schrijver hierarchies if not.⁶

Perhaps the most prominent line of work in that direction concerns refutation of random unsatisfiable CSPs. A notable paper of Feige [Fei02] used the difficulty of doing this for 3-SAT as a “hypothesis” to derive hardness of approximation results, drawing attention to this question, and [Gri01b, Sch08, KMOW17] showed that the SOS hierarchy cannot yield efficient refutations until the number of clauses is far larger than the threshold at which a random formula becomes unsatisfiable. Another prominent line of work concerns certifying bounds on the size of the largest clique in a graph drawn from $ER(n, \frac{1}{2})$, the certification side of the search problem discussed above. As we mentioned, the typical size of the largest clique is $2 \log_2 n$; however, a series of results [FK00, MPW15, DM15b, HKP⁺18, BHK⁺19] showed that

⁶One notable exception is [WBP16], who use a reduction strategy similar to what we will pursue, albeit for just one specific problem. See our discussion at the beginning of Chapter 2.

SOS cannot efficiently certify a bound better than $O(\sqrt{n})$. The clique problem was harder to address, being different from k -CSPs with $k \geq 3$ in two salient ways: first, it is a *quadratic* problem, concerning the quadratic form of a certain structured kind of vector with the adjacency matrix. As we will see, this relates the problem to the *spectrum* of the adjacency matrix, an additional structural feature that may be exploited by algorithms and complicates the task of proving lower bounds. Second, it has *globally-distributed* information, meaning that, unlike the specific and local information we learn from a CSP constraint about the few variables it involves, each vertex of a dense graph only carries a weak signal as to its propensity to belong to a large clique. All of this weak information must somehow be synthesized to search for a large clique, certify bounds on the size of the largest clique, or prove that either is impossible (see [Moi20] for lucid discussion of this point).

It is with these challenges that we pick up the thread. In this thesis, we will consider barriers to certification both in a general class of quadratic problems formally resembling the largest clique problem, and in a specific representative problem of this class that we discuss below. We will develop a general notion of a “spectral barrier” to certification, give a new kind of evidence for such a barrier based on reductions to hypothesis testing—expanding the limited arsenal of tools available for demonstrating average-case complexity of certification problems—and propose related new tools for proving lower bounds against the SOS hierarchy.

1.3 A MOTIVATING EXAMPLE

While we have presented quite a broad setting above, the topics we will consider all stem more or less directly from attempts to understand the hardness of certification for the

following specific random optimization problem: for some random $\mathbf{W} \in \mathbb{R}_{\text{sym}}^{n \times n}$,

$$\mathbf{M}_{\{\pm 1/\sqrt{n}\}^n}(\mathbf{W}) := \left\{ \begin{array}{l} \text{maximize } \mathbf{x}^\top \mathbf{W} \mathbf{x} \\ \text{subject to } \mathbf{x} \in \{\pm 1/\sqrt{n}\}^n \end{array} \right\}. \quad (1.3)$$

From the perspective of combinatorial optimization, a natural choice of \mathbf{W} is a graph Laplacian, in which case (1.3) computes (up to rescaling) the size of the largest cut. We are then led to consider \mathbf{W} the Laplacian of a random graph, and the most interesting regime turns out to be the case of *sparse* random graphs, either Erdős-Rényi graphs $\text{ER}(n, \frac{c}{n})$ or uniformly-random regular graphs with fixed integer degree $c \geq 3$, which we denote $\text{Reg}(c, n)$.

We recall that, in the earlier average-case examples we considered, our first step was to pin down the true typical value of the random optimization problem in question (e.g., that the size of the largest clique in our earlier example was of order $2 \log_2 n$). Here, for fixed c and $n \rightarrow \infty$, this is an endeavor unto itself: it is known that the size of the maximum cut is asymptotically $n(\frac{c}{4} + f(c)\sqrt{c})$ for some $f(c) \in [a, A]$ for some $0 < a < A$ for all c , and various quantitative bounds are available, but $f(c)$ has not been exactly determined for any specific c [DSW07, BGT10, GL18]. On the other hand, a remarkable result of [DMS17], first conjectured by [ZB10], established that

$$\lim_{c \rightarrow \infty} f(c) = \frac{1}{2} \mathbf{P}_* := \frac{1}{4} \lim_{n \rightarrow \infty} \mathbb{E}_{\mathbf{W} \sim \text{GOE}(n)} [\mathbf{M}_{\{\pm 1/\sqrt{n}\}^n}(\mathbf{W})] \approx 0.382, \quad (1.4)$$

where $\text{GOE}(n)$ denotes the *Gaussian orthogonal ensemble*, the probability distribution over symmetric matrices \mathbf{W} where $W_{ii} \sim \mathcal{N}(0, 2/n)$ and $W_{ij} = W_{ji} \sim \mathcal{N}(0, 1/n)$ when $i \neq j$, independently. That is, at least for the purposes of this optimization problem, $\text{GOE}(n)$ may be viewed as a kind of limit of the sparse random graph distributions as $n \rightarrow \infty$ followed by $c \rightarrow \infty$.⁷ Thus, to ease the technical difficulty of working with random graph distributions,

⁷One simpler hint that such a behavior might occur is the convergence of the Kesten-McKay law of the eigenvalues of the random c -regular graph to the semicircle law of the eigenvalues of $\text{GOE}(n)$ (suitably

it is reasonable to consider $\mathbf{W} \sim \text{GOE}(n)$ as a surrogate.

There are other reasons to be interested in this optimization problem. Most immediately, the function of \mathbf{x} in (1.3) when $\mathbf{W} \sim \text{GOE}(n)$ is the random Hamiltonian of the celebrated *Sherrington-Kirkpatrick (SK) model*, a mean-field model of spin glasses in statistical physics [SK75]. Optimizing it, up to a change of sign, computes the *ground state energy* of the model. The asymptotics of this value motivated a large body of fascinating mathematical and physical work; the quantity we cite in (1.4) comes from the predictions of Parisi [Par79, Par80, CR02] using the non-rigorous *replica method*, which were later in part validated mathematically by [Gue03, Tal06, Pan13]. The remarkable fact is that the SK model was the first model of its kind understood to exhibit a rich structure in its optimization landscape known as *continuous replica symmetry breaking*. This entails, roughly speaking, a hierarchical clustering of sets of \mathbf{x} having high objective value, so that given such a high-quality point there are points of comparable quality at many scales of distances to \mathbf{x} .

Alternatively, from a point of view more familiar in optimization and convex analysis, one may view this problem as choosing a uniformly random direction in the space of symmetric matrices—as the law of $\mathbf{W} \sim \text{GOE}(n)$ is isotropic with respect to the Frobenius inner product—and measuring along that direction the width (up to rescaling) of the *cut polytope*, the convex hull of the matrices $\mathbf{x}\mathbf{x}^\top$ over $\mathbf{x} \in \{\pm 1\}^n$. In this way we measure the *Gaussian width* or the closely-related *mean width*, basic statistics for measuring the size of a convex body and the first of the much-studied *intrinsic volumes* (see, e.g., [AS17]), for the cut polytope, one of the central objects of combinatorial optimization [DL09].

Whichever of these three interpretations the reader might find most appealing, the problem (1.3) with $\mathbf{W} \sim \text{GOE}(n)$ is evidently a central example of random combinatorial optimization. We have also seen that we know the typical value of this problem, which we rescaled) as $c \rightarrow \infty$ [McK81].

rewrite in its more usual normalization,

$$\lim_{n \rightarrow \infty} \mathbb{E}_{\mathbf{W} \sim \text{GOE}(n)} [\mathbb{M}_{\{\pm 1/\sqrt{n}\}^n}(\mathbf{W})] = 2P_* \approx \boxed{1.526}. \quad (1.5)$$

We are now ready to ask: is it possible to *efficiently* optimize this random function? As we have seen above, it is fruitful to think of “optimizing” both from above and from below—either searching for a \mathbf{x} achieving a large objective value, or using an algorithm that certifies upper bounds guaranteed to be valid for all \mathbf{x} . We may already evaluate some simple algorithmic approaches.

Example 1.3.1 (Spectral search and certification). *Two related spectral algorithms give simple examples of algorithms for both search and certification. For certification, writing λ_{\max} for the largest eigenvalue of \mathbf{W} , we may use the bound*

$$\mathbf{x}^\top \mathbf{W} \mathbf{x} \leq \lambda_{\max} \cdot \|\mathbf{x}\|^2 = \lambda_{\max} \approx \boxed{2} \quad (1.6)$$

for all $\mathbf{x} \in \{\pm 1/\sqrt{n}\}^n$, whereby λ_{\max} is a certifiable upper bound on (1.3). From classical random matrix theory [Gem80, AGZ10], it is known that $\lambda_{\max} \rightarrow 2$ almost surely as $n \rightarrow \infty$.

For search, for \mathbf{v}_{\max} the eigenvector of λ_{\max} , we may take $\mathbf{x} = \mathbf{x}(\mathbf{W}) := \text{sgn}(\mathbf{v}_{\max})/\sqrt{n}$ where sgn denotes the $\{\pm 1\}$ -valued sign function, applied entrywise. The vector \mathbf{v}_{\max} is distributed as an uniform random unit vector in \mathbb{R}^n , so the quality of this solution may be computed as

$$\mathbf{x}^\top \mathbf{W} \mathbf{x} = \lambda_{\max} \cdot \langle \mathbf{x}, \mathbf{v}_{\max} \rangle^2 + O\left(\frac{1}{\sqrt{n}}\right) = \lambda_{\max} \cdot \frac{\|\mathbf{v}_{\max}\|_1^2}{n} + O\left(\frac{1}{\sqrt{n}}\right) \approx \frac{4}{\pi} \approx \boxed{1.273} \quad (1.7)$$

with high probability as $n \rightarrow \infty$.⁸ (The error in the first equation is obtained as $\sum_i \lambda_i \langle \mathbf{v}_i, \mathbf{x} \rangle^2 \approx$

⁸We say that a sequence of events $(A_n)_{n \in \mathbb{N}}$ with $A_n \in \mathcal{F}_n$ the σ -algebra of an associated measurable space occurs with high probability (in n) if the probability of A_n tends to 1 as $n \rightarrow \infty$.

$\frac{1}{n} \text{tr}(\mathbf{W})(1 - \langle \mathbf{v}_{\max}, \mathbf{x} \rangle^2)$, where the sum is over all eigenvectors \mathbf{v}_i except \mathbf{v}_{\max} . This analysis appeared in [ALR87], an early rigorous mathematical work on the SK model.)

Comparing the three boxed numbers above, we see that neither spectral search nor spectral certification gives a tight approximation of the SK Hamiltonian.

Montanari [Mon18] recently showed that, using a variant of approximate message passing, for any $\epsilon > 0$ there in fact exists a polynomial-time search algorithm that with high probability achieves a value of $2P_* - \epsilon$ (assuming a technical conjecture on the SK model, and inspired by earlier results of [ABM20, Sub18]). This closed the question of efficient “perfect” search. For certification, for some time the only algorithm beyond the spectral bound that had been studied was the degree 2 SOS relaxation, an SDP also occurring in the seminal results of Goemans, Williamson and Nesterov [GW95, Nes98]. For this algorithm, [MS16] showed (in the course of applications to discrete problems in the vein of [DMS17]) that the value achieved is again 2, asymptotically—no better than the spectral algorithm. Finally, several works of the author [BKW20b, KB20, Kun20b] as well as the concurrent works [MRX20, GJJ⁺20] showed, in the case of the first reference, that conditional on a conjecture we will discuss in Chapter 3 no efficient certification algorithm can improve on the spectral bound, and, for the remaining references, that no polynomial-time SOS relaxation improves on the spectral bound.

Remark 1.3.2 (Proof complexity of the SK Hamiltonian). *The above hardness of certification is despite the fact that a relatively simple argument due to Guerra [Gue01], which may be viewed in our context even more simply as an application of the Fernique-Sudakov Gaussian comparison inequality (see Chapter 3 of [LT13]), gives a bound of $2\sqrt{2/\pi} < 2$ (called the “replica-symmetric prediction” in the spin glass literature) on the objective. The proof only involves an interpolation argument between the optimization of $\mathbf{x}^\top \mathbf{W} \mathbf{x}$ and a linear optimization of $\mathbf{x}^\top \mathbf{g}$ over $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ (see, in addition to Guerra’s argument, Chatterjee’s proof*

along these lines of the Fernique-Sudakov inequality in [Cha05]). Thus these results suggest that no simple algebraic “pointwise” argument can reproduce even this preliminary bound.

1.4 SPECTRAL BARRIERS

In fact, we will see that this kind of failure of certification algorithms is unique neither to optimization over the hypercube nor to the particular random function of the SK Hamiltonian. We will focus our initial efforts on the following broader class of problems, which conveniently is broad enough to encompass many interesting problems of applied and theoretical interest, while enforcing enough structure that a unified explanation can be given for the high computational costs of certification.

Definition 1.4.1 (Constrained PCA problem). *Let $\mathcal{X} \subseteq \mathbb{R}^{n \times k}$ and $\mathbf{W} \in \mathbb{R}_{\text{sym}}^{n \times n}$. We call an optimization problem of the following form a constrained principal component analysis (PCA) problem:*

$$M_{\mathcal{X}}(\mathbf{W}) := \left\{ \begin{array}{l} \text{maximize } \text{tr}(\mathbf{X}^{\top} \mathbf{W} \mathbf{X}) \\ \text{subject to } \mathbf{X} \in \mathcal{X} \end{array} \right\}. \quad (1.8)$$

In essence, constrained PCA asks us to search for a structured vector or small collection of vectors \mathbf{X} (we will exclusively think of k , the “rank parameter,” as constant while $n \rightarrow \infty$) aligned with those eigenvectors of \mathbf{W} that lie close to the top of its spectrum. These problems include finding large cuts and k -cuts in graphs [GW95, Tre12a], coloring graphs [Hof70, AG84, AK97], variants of the satisfiability constraint satisfaction problem [Zwi98, DMO⁺19], synchronization problems [Sin11], PCA with variously structured signals (non-negative [MR15], sparse [ZHT06], and others [DMR14]), and certain cases of tensor-valued PCA [RM14]. Despite this diversity of applications, for \mathbf{W} random, many constrained PCA problems appear to present one and the same difficulty to efficient certification algorithms. Let us first define this notion that we have discussed informally above.

Definition 1.4.2 (Certification algorithm). Let f be an algorithm, possibly randomized,⁹ that takes a square matrix \mathbf{W} as input and outputs a number $f(\mathbf{W}) \in \mathbb{R}$. We say that f certifies a value K on $\mathcal{M}_{\mathcal{X}}$ when $\mathbf{W} \sim \mathcal{Q}_n$ for some sequence of distributions $\mathcal{Q}_n \in \mathcal{P}(\mathbb{R}_{\text{sym}}^{n \times n})$ if

1. for any $\mathbf{W} \in \mathbb{R}_{\text{sym}}^{n \times n}$, $M_{\mathcal{X}}(\mathbf{W}) \leq f(\mathbf{W})$, and
2. if $\mathbf{W} \sim \mathcal{Q}_n$, then $f(\mathbf{W}) \leq K + o(1)$ with high probability as $n \rightarrow \infty$.

The obstacle we will study is that such algorithms often cannot certify a better bound than a spectral bound that follows from decoupling the dependences of constrained PCA on \mathbf{W} and \mathcal{X} :

$$\max_{\mathbf{X} \in \mathcal{X}} \text{tr}(\mathbf{X}^\top \mathbf{W} \mathbf{X}) \leq \lambda_{\max}(\mathbf{W}) \cdot \max_{\mathbf{X} \in \mathcal{X}} \|\mathbf{X}\|_F^2. \quad (1.9)$$

(Here $\|\mathbf{X}\|_F^2 := \text{tr}(\mathbf{X}^\top \mathbf{X})$.) When a spectral bound is optimal among some broad class of efficient certification algorithms, but does not yield a tight result, we say that a constrained PCA problem exhibits a *spectral barrier*. The main specific project of this thesis is to provide evidence for spectral barriers in numerous constrained PCA problems, and to explore the mathematical phenomena that appear to be responsible for these barriers.

We believe that spectral barriers are an exceptionally robust phenomenon. There will be many indications in this work that spectral barriers are sensitive neither to the constraint set \mathcal{X} nor to the instance \mathbf{W} (or, in the average case analysis we will mostly be concerned with, to its distribution). So long as both are “generic” or sufficiently “incoherent” with respect to the standard coordinate basis, spectral barriers seem to arise reliably.

As discussed above for the case of the SK Hamiltonian, spectral barriers sometimes arise when optimal search can be performed efficiently; other times, both search and certification appear unable to attain optimal performance. Beyond merely numerical gaps in the performance various algorithms are able to achieve, though, we emphasize that the mechanisms

⁹We allow f to be randomized; i.e., it may use randomness in its computations, but the output K must be an upper bound almost surely. We do not expect certification algorithms to require randomness, but it may be convenient, e.g., to randomly initialize an iterative optimization procedure.

giving rise to gaps in search and certification appear to be completely different. In the case of the SK Hamiltonian, the spectral barrier to certification is in essence a matter of the difficulty of refuting that a hypercube point might lie in the span of the top eigenvectors of \mathbf{W} , a uniformly-distributed low-dimensional subspace in the case of $\mathbf{W} \sim \text{GOE}(n)$. The probabilistic setting of the spectral barrier is thus relatively simple; indeed, this is related to the robustness mentioned above, since a sufficiently incoherent deterministic low-dimensional subspace shares all the features of a random subspace that are relevant to the analysis of the spectral barrier. On the other hand, the search algorithm of [Mon18] is intimately related to the “depths” of the landscape of the SK Hamiltonian and the complex structure of the low-temperature Gibbs measures related to continuous replica symmetry breaking that allow the landscape to remain navigable by an optimization algorithm to near-optimality. Because of this, the algorithm itself is also likely brittle, being calibrated by quantities computed from the Parisi formalism and other aspects of our theoretical understanding of the particular SK Hamiltonian.

1.5 SUMMARY OF CONTRIBUTIONS

We now outline the organization of the thesis and the content and main contributions of each chapter. At the beginning of each chapter we give a more detailed summary and its context at that point in the thesis, as well as specific references and lists of main results. Overall, all of the content is based on the publications [BKW20b, KB20], as well as the articles in submission [DKWB19, Kun20a, Kun20b, BBK⁺20, BKW20a], the article in revision [BK18] at the time of writing, and the expository notes [KWB19]. Chapter 9 is based on a short article in preparation.

1.5.1 REDUCTIONS, PLANTING, AND LOW-DEGREE POLYNOMIALS

We begin by developing *reductions* from certification in constrained PCA problems to certain associated hypothesis testing problems, and using a method based on analyzing algorithms computing low-degree polynomials to provide evidence that these hypothesis testing problems are hard.

In Chapter 2, we give a unified treatment of these reductions, which have been proved on an ad hoc basis in several publications. We also describe the *spectral planting* strategy underlying the reductions, which gives an analytically-tractable way to skew the top eigenspace of a GOE matrix towards a particular constrained PCA solution.

In Chapter 3, we introduce the method based on low-degree polynomials that we will use to treat the resulting hypothesis testing problems, giving two justifications (one based on the actual development of these ideas in prior literature on SOS optimization and another streamlined one devised post hoc in [KWB19]) and describing consequences of the lower bounds we will prove.

In Chapter 4, we prove bounds and formulae for the norm of the low-degree likelihood ratio, the main quantity governing lower bounds against algorithms computing low-degree polynomials. We explore the *overlap form* that these expressions often take, which show that, though they describe high-dimensional sums of moments of random variables, the norms can often be condensed into a single scalar expectation. We treat several observation models in this way using various properties of orthogonal polynomials.

In Chapter 5, we apply these results to prove lower bounds against algorithms computing low-degree polynomials. Using our tools we are able to treat a wide range of models beyond those needed for our applications to certification: we give tight results with a unified and simplified presentation for matrix and tensor PCA, sparse PCA, and the stochastic block model. We use suitable matrix PCA results to then deduce conditional hardness results for

certification for the SK Hamiltonian, the Potts spin glass Hamiltonian (related to a Gaussian approximation of graph coloring), and non-negative PCA over a GOE input matrix.

1.5.2 THE SUM-OF-SQUARES HIERARCHY

We then specialize to the hypercube constraints of the SK Hamiltonian, and study SOS relaxations over this constraint set and build up to proving SOS lower bounds.

In Chapter 6, we introduce basic properties of SOS relaxations over the hypercube and present some background results. We also motivate our program of building pseudomoment matrices—the objects underlying lower bounds against the SOS hierarchy—as Gram matrices, and argue that this is more broadly a productive approach to understanding SOS.

In Chapter 7, we study the structure of degree 4 SOS pseudomoment matrices viewed as Gram matrices. This yields constraints on the spectra of these matrices and intriguing connections to the notions of entanglement and separability from quantum information theory. We also use our characterization to demonstrate classes of examples of degree 2 Gram matrices that can and cannot be extended to degree 4 pseudomoments that are built from equiangular tight frames, highly-structured combinatorial packings of vectors.

In Chapter 8, we reinterpret our degree 4 construction for Gram matrices of equiangular tight frames as a Gaussian conditioning computation involving *surrogate random matrices* from which we propose building degree 4 pseudomoments. Generalizing these to surrogate random *tensors*, we propose a general scheme for extending degree 2 pseudomoments to higher degrees, and derive a description of the result in terms of the decomposition of polynomials into parts belonging to and orthogonal to an ideal under the apolar inner product. We note that few general techniques for building pseudomoments are known; the main other idea in this direction is the *pseudocalibration* technique of [BHK⁺19], to which our *spectral extension* method gives a plausible alternative.

In Chapter 9, we show that the spectral pseudomoment extension applies exactly to

a deterministic construction due to Grigoriev [Gri01a] and Laurent [Lau03b]. We give a streamlined representation-theoretic proof of this result, and prove a conjecture of Laurent’s concerning the eigenvalues of the associated pseudomoment matrix.

In Chapter 10, we propose a general closed-form pseudomoment extension derived from the spectral extension, by heuristically generalizing the Maxwell-Sylvester representation of harmonic polynomials to certain multiharmonic polynomials. This yields a combinatorial construction of *sum-of-forests pseudomoments*. We prove *lifting theorems* giving general (though rather technical) conditions under which degree 2 pseudomoments may be extended in this way to higher degrees. For “sufficiently incoherent” high-rank degree 2 pseudomoments our lifting succeeds to any constant degree of SOS, while for low-rank degree 2 pseudomoments we adapt it to succeed to degree 6.

In Chapter 11, we apply our lifting theorems. We first show that the sum-of-forests pseudomoments approximately recover, to leading order, the Grigoriev-Laurent pseudomoments, giving an enumerative combinatorial interpretation of the construction’s behavior. We also show that the lifting to high degree applies to uniformly-random high-rank projection matrices. Lastly, we show that the lifting to degree 6 applies to rescaled uniformly-random low-rank projection matrices, which results in a tight degree 6 lower bound for the SK Hamiltonian.

Finally, in Appendix A we collect several open problems concerning various topics discussed in the thesis.

Part I

Computationally-Quiet Planting and Low-Degree Polynomials

2 | SPECTRAL PLANTING AND REDUCTIONS TO HYPOTHESIS TESTING

To address the question of whether efficiently certifying better-than-spectral bounds on constrained PCA problems is possible, our first order of business will be to develop tools that allow us to make a prediction and provide some indirect evidence of hardness. This is helpful because analyzing specific convex relaxations, as we will do in Part II, both is technically complicated and entails entirely different types of arguments—looking for SOS proofs versus constructing pseudomoments, in our case—depending on whether we are trying to show that it is possible or impossible to cross the spectral barrier. Instead, we will see that we can consider a different class of lower bounds, which are governed by more straightforward and unified computations that we expect can identify either situation. It will also turn out that the tools we use to prove these lower bounds are quite broadly applicable, and allow us to prove similar lower bounds in various other interesting settings.

Because of these further applications, we will eventually develop the tools for our lower bounds in greater generality than just those relevant to constrained PCA. In this initial chapter where we establish the connection between these lower bounds and constrained PCA, we will also be quite general in allowing the constraint set \mathcal{X} to vary, but for the most part we will consider the same distribution of \mathbf{W} as in the SK Hamiltonian, $\mathbf{W} \sim \text{GOE}(n)$.

Our goal in this chapter will be to show that, if it is possible to cross the spectral barrier

for certification in a constrained PCA problem, then it is possible to solve certain *hypothesis testing* problems of distinguishing between pairs of probability distributions. The toolbox for identifying computational hardness in testing problems is richer than that for average-case certification problems, so this will give us a new avenue to argue the hardness of certification. More specifically, we will show that certifying bounds necessarily involves excluding the possibility of the problem instance being tampered with by *planting* high-quality solutions: if we can certify a better-than-spectral bound, then we must have verified that \mathbf{W} has not been changed to align its top eigenspaces especially well with some $\mathbf{X} \in \mathcal{X}$. Thus we will consider hypothesis testing between whether \mathbf{W} was drawn from $\text{GOE}(n)$ or some tampered-with variant thereof. Moreover, any fixed collection of eigenspaces of $\text{GOE}(n)$ span a uniformly random subspace (see Proposition 2.2.2), we may “factorize” this testing problem and test instead between a uniformly random subspace and one that has been skewed to align well with $\mathbf{X} \in \mathcal{X}$.

SUMMARY AND REFERENCES This chapter gives a unified and generalized treatment of a reduction argument that has been repeated for several problems, *mutatis mutandis*, in the following references: [BKW20b] for rank-one constrained PCA, [BBK⁺20] for higher-rank constrained PCA, and [BKW20a] for non-negative PCA, which is a type of rank-one constrained PCA but which we will see presents minor additional technical obstacles. In Section 2.5 we also discuss one example of this reduction when \mathbf{W} is *not* drawn from $\text{GOE}(n)$, which is also taken from [BBK⁺20]. The main results of this chapter that we will use in the sequel are Corollaries 2.3.3 and 2.3.4, which give general reductions from certification in constrained PCA problems under the GOE to hypothesis testing in Wishart spiked matrix models.

PRIOR WORK The idea of reducing hypothesis testing between null and planted models to certifying bounds on some quantity under the null model is implicit in many works on

the planted clique model, starting with its introduction by [Jer92, Kuč95] and becoming especially apparent in the results on SOS of [DM15a, RS15, HKP⁺18, BHK⁺19], and even more explicit in subsequent works using the pseudocalibration framework for other problems [HKP⁺17, MRX20, GJJ⁺20]. However, none of these results draw a formal connection between the two problems; rather, they only argue that if we believe it should be hard to distinguish a given null and planted model, then that should constrain the possible SOS pseudomoments in various ways. The only prior work we are aware of where a genuine reduction is conducted is [WBP16], who do this for the specific problem of certifying the *restricted isometry property* (see also [DKWB20], in which the author participated but which we will not discuss here, where the framework we describe is applied to the same problem). More broadly, that result belongs to a line of work on reductions among numerous average-case problems [BR13, HWX15, WBS16, BB19b, BB19a, BB20], but [WBP16] is the only result in this direction we are aware of that reduces to *certification* problems in particular.

2.1 HYPOTHESIS TESTING

We begin with some generalities about hypothesis testing. The basic setup that all of the testing problems we study will share is as follows. Suppose $(\mathbb{P}_n)_{n \in \mathbb{N}_+}$ and $(\mathbb{Q}_n)_{n \in \mathbb{N}_+}$ are two sequences of probability distributions over a common sequence of measurable spaces $((\Omega_n, \mathcal{F}_n))_{n \in \mathbb{N}_+}$. In statistical parlance, we will think throughout of \mathbb{P}_n as the model of the *alternative hypothesis* and \mathbb{Q}_n as the model of the *null hypothesis*. In our previous language, the distributions (\mathbb{P}_n) include a “**p**lanted” structure, making the notation a helpful mnemonic. Suppose we observe $\mathbf{Y} \in \Omega_n$ which is drawn from one of \mathbb{P}_n or \mathbb{Q}_n . We hope to recover this choice of distribution in the following sense.

Definition 2.1.1. *A test is a measurable function $f_n : \Omega_n \rightarrow \{p, q\}$.*

Definition 2.1.2. A sequence of tests f_n is said to (strongly) distinguish (\mathbb{P}_n) and (\mathbb{Q}_n) if $f_n(\mathbf{Y}) = p$ with high probability when $\mathbf{Y} \sim \mathbb{P}_n$, and $f_n(\mathbf{Y}) = q$ with high probability when $\mathbf{Y} \sim \mathbb{Q}_n$. If such f_n exist, we say that (\mathbb{P}_n) and (\mathbb{Q}_n) are statistically distinguishable.¹

In statistics one often draws the distinction between the probabilities of a test returning an incorrect result when $\mathbf{Y} \sim \mathbb{Q}_n$ and $\mathbf{Y} \sim \mathbb{P}_n$, called Type I and Type II errors respectively. We will briefly discuss this point and some related facts later in Section 3.1, but generally we will only consider tests successful when they drive both error probabilities to zero. We will be especially interested in differences in the computational cost of hypothesis testing, for which we introduce the following terminology.

Definition 2.1.3. If a sequence of tests f_n distinguishes (\mathbb{P}_n) and (\mathbb{Q}_n) and $f_n(\mathbf{Y})$ may be computed in time $T(n)$, then we say that (\mathbb{P}_n) and (\mathbb{Q}_n) are distinguishable in time $T(n)$. If $T(n)$ grows polynomially in n , then we say that (\mathbb{P}_n) and (\mathbb{Q}_n) are computationally distinguishable.

Remark 2.1.4 (Model of computation). For the sake of convenience, we suppose we may perform real arithmetic exactly, and view $T(n)$ above as counting the number of exact real operations. We also will later assume we may compute eigendecompositions exactly in polynomial time. However, all computations we perform should be numerically stable, and all matrices should be well-conditioned, so our reductions and other claims should also hold in weaker models of, say, floating point arithmetic. The so-called “real RAM” model common in computational geometry is close to the informal model we adopt.

¹We will only consider this so-called *strong* version of distinguishability, where the probability of success must tend to 1 as $n \rightarrow \infty$, as opposed to the *weak* version where this probability need only be bounded above $\frac{1}{2}$. For high-dimensional problems, the strong version typically coincides with important notions of estimating the planted signal. See, e.g., [BM17, EAKJ20, CL19] for some discussion of “weak detection” in literature on some of the models we will study.

2.1.1 STATISTICAL-COMPUTATIONAL GAPS

We emphasize that these computational distinctions are far from trivial. In fact, one of the central phenomena of the asymptotics of high-dimensional testing problems is that statistical and computational distinguishability do not always coincide. Such a situation is called a *statistical-computational* or *information-computation gap*. (See, e.g., [MM09, ZK16, BPW18] for general discussion of this phenomenon and common analytical tools for studying it drawn from statistical physics.) For this reason, while classical statistical theory provides tools for identifying statistically-indistinguishable distributions, for which the hypothesis testing problem is impossible, these tools do not always suffice to accurately identify computationally-indistinguishable distributions, for which the testing problem may be possible but is prohibitively difficult.

Indeed, as a concrete example, if we reduce the problem of certifying a bound on the SK Hamiltonian to a hard hypothesis testing problem as we have outlined above, then such a gap *must* occur. After all, in exponential time it is possible to certify a tight bound on $M_{\{\pm 1/\sqrt{n}\}^n}(\mathbf{W})$ by simply searching over all $\mathbf{x} \in \{\pm 1/\sqrt{n}\}^n$, so perfect certification is possible but (we believe) hard. Accordingly, any hard hypothesis testing problem we reduce to certification will be possible to solve, too. In this sense, our plan entails demonstrating that **a spectral barrier is a manifestation of a statistical-computational gap.**

As we will soon meet a problem exhibiting a somewhat involved statistical-computational gap, in the negatively-spiked Wishart matrix model, let us give a simpler preliminary example for the time being. Typically, such a gap arises in the following more specific way. Suppose the sequence (\mathbb{P}_n) has a further dependence on a *signal-to-noise* parameter $\lambda > 0$, forming a parametrized sequence $(\mathbb{P}_{\lambda,n})_{\lambda>0,n\geq 1}$. This parameter should describe, in some sense, the strength of the structure present under the planted distribution (or, in some cases including the Wishart model, the number of i.i.d. samples received from a fixed distri-

bution). The following is one of the best-studied examples.

Example 2.1.5 (Planted clique problem [Jer92, Kuč95]). *Under the null model \mathbb{Q}_n , we observe an Erdős-Rényi random graph on n vertices and with edge probability $\frac{1}{2}$ (that is, each pair of vertices is connected independently with probability $\frac{1}{2}$). The signal-to-noise parameter λ is an integer $1 \leq \lambda \leq n$. Under the planted model $\mathbb{P}_{\lambda,n}$, we first choose a random subset of vertices $S \subseteq [n]$ of size $|S| = \lambda$, uniformly at random. We then take the union of a graph sampled from \mathbb{Q}_n with a planted clique or complete subgraph on S .*

The typical size of the largest clique under \mathbb{Q}_n is $2 \log_2 n$ (with fluctuations of lower order), so \mathbb{P}_n and \mathbb{Q}_n are statistically distinguishable whenever $\lambda \geq (2 + \epsilon) \log_2 n$ for some $\epsilon > 0$. On the other hand, the best known algorithms for distinguishing \mathbb{Q}_n and \mathbb{P}_n in polynomial time [AKS98, FK00, DM15a] only succeed when $\lambda = \Omega(\sqrt{n})$, and there is plentiful further evidence that this performance is optimal [Jer92, Ros10, BHK⁺19, GZ19]. Thus there is a large regime of $(2 + \epsilon) \log_2 n \leq \lambda \ll \sqrt{n}$ where it is known that it is possible to distinguish \mathbb{Q}_n and \mathbb{P}_n , but conjectured that it is hard to do so.

2.2 PUSHOUT AND COMPUTATIONALLY-QUIET PLANTING

We now move towards developing a means of reducing testing problems to certification in constrained PCA problems, in such a way that we have reason to believe the testing problems involved are hard. To begin, the following is a broad and easy claim that shows how, in general, a certification algorithm can be used to test between a natural and a tampered-with distribution of instances \mathbf{W} .

Theorem 2.2.1 (Abstract reduction). *Let $C > 0$. Let $\mathbb{P}_n \in \mathcal{P}(\mathbb{R}_{\text{sym}}^{n \times n})$ be such that, under $\mathbf{W} \sim \mathbb{P}_n$, for all $\epsilon > 0$, with high probability $\mathbb{M}_X(\mathbf{W}) \geq C - \epsilon$ (for instance, if $\mathbb{M}_X(\mathbf{W}) \rightarrow C$ in probability). Let $\mathbb{Q}_n \in \mathcal{P}(\mathbb{R}_{\text{sym}}^{n \times n})$ be such that there exists $\epsilon > 0$ and an algorithm running in*

time $T(n)$ that can certify with high probability a bound of at most $C - \epsilon$ on $M_X(\mathbf{W})$ when $\mathbf{W} \sim \mathbb{Q}_n$. Then, there exists an algorithm running in time $T(n) + O(1)$ that can distinguish (\mathbb{P}_n) from (\mathbb{Q}_n) .

Proof. Let $g(\mathbf{W})$ be the certificate the algorithm computes. Then, we define tests $f_n : \mathbb{R}_{\text{sym}}^{n \times n} \rightarrow \{p, q\}$ by

$$f_n(\mathbf{W}) := \begin{cases} p & \text{if } g(\mathbf{W}) \geq C - \epsilon/2, \\ q & \text{if } g(\mathbf{W}) < C - \epsilon/2, \end{cases} \quad (2.1)$$

and clearly under the assumptions the f_n will distinguish (\mathbb{P}_n) from (\mathbb{Q}_n) . \square

Recall that we will be given \mathbb{Q}_n , usually $\mathbb{Q}_n = \text{GOE}(n)$, and the assumption that a better-than-spectral certification algorithm exists will be for the sake of contradiction. Therefore, the main matter remaining will be to design \mathbb{P}_n making $M_X(\mathbf{W})$ typically large in a way that is *computationally-quiet*, or difficult to distinguish from \mathbb{Q}_n . That is, we want to *plant* an $\mathbf{X} \in \mathcal{X}$ in $\mathbf{W} \sim \text{GOE}(n)$ such that $\text{tr}(\mathbf{X}^\top \mathbf{W} \mathbf{X})$ is nearly as large as possible (so that \mathbf{X} nearly saturates the spectral bound) in a way that is difficult to detect.

The notion of quiet planting has appeared previously in the literature, but mostly in the context of *statistically-quiet planting*, seeking a planting such that \mathbb{P}_n and \mathbb{Q}_n are close in some measurement of distance between probability measures, or indistinguishable by computationally-unbounded tests. This has been used in hopes of intentionally designing random hard instances of constraint satisfaction problems [JP00, KZ09, ZK11], and also as a mathematical technique for analyzing the unplanted distribution via the planted distribution from which it is indistinguishable [COKV07b, COKV07a, ACO08], in particular for establishing the “shattering” phenomenon in the solution spaces of random constrained satisfaction problems. In these latter situations, the underlying random problem usually has one or many solutions—satisfying assignments in a constraint satisfaction problem, say—with high probability already, and we wish to plant *another* such solution that we have

explicit control over. In our setting, in contrast, we are interested in planting something (the matrix \mathbf{X} , in a suitable sense) in \mathbb{P}_n that with high probability *does not exist* under \mathbb{Q}_n . Typically such a pair of distributions will be distinguishable by a brute-force search for the planted object, so the restriction to computationally-bounded quietness is essential.

Before proceeding to consider quiet planting in the GOE, let us recall some basic facts about it and its spectrum.

Proposition 2.2.2 (GOE spectrum [Wig93, Gem80, AGZ10]). *For $\lambda_1 \geq \dots \geq \lambda_n$ eigenvalues of $\mathbf{W} \sim \text{GOE}(n)$ and $\mathbf{w}_1, \dots, \mathbf{w}_n \in \mathbb{S}^{n-1}$ the corresponding eigenvectors, the following hold.*

- *The λ_i are almost surely distinct.*
- *The empirical spectral measure $\frac{1}{n} \sum_{i=1}^n \delta_{\lambda_i}$ converges weakly to the semicircle law on $[-2, 2]$ having density $\mathbb{1}\{x \in [-2, 2]\} \frac{1}{2\pi} \sqrt{4 - x^2} dx$.*
- *$\lambda_1 \rightarrow 2$ and $\lambda_n \rightarrow -2$ in probability.*
- *For any $S \subseteq [n]$ independent of \mathbf{W} (including deterministic sets), the law of the subspace $\text{span}(\{\mathbf{w}_i : i \in S\})$ is that of a uniformly-distributed $|S|$ -dimensional subspace of \mathbb{R}^n .*

Let us consider a naive strategy for planting in the GOE, for the moment working under the hypercube constraints $\mathcal{X} = \{\pm 1/\sqrt{n}\}^n$. A natural idea is *additive planting*, where we draw $\mathbf{W} \sim \mathbb{P}_n$ by drawing $\mathbf{W}^{(0)} \sim \text{GOE}(n)$ and $\mathbf{x} \sim \text{Unif}(\{\pm 1/\sqrt{n}\}^n)$, and then set $\mathbf{W} = \mathbf{W}^{(0)} + \lambda \mathbf{x} \mathbf{x}^\top$ for some $\lambda > 0$. (This forms the so-called *Wigner spiked matrix model*, to which we will return in greater detail in Example 4.1.2 and again in Section 5.2.4.) To achieve $\mathbf{x}^\top \mathbf{W} \mathbf{x} \approx 2$ to saturate the spectral bound, we want to take $\lambda \approx 2$, since $\mathbf{x}^\top \mathbf{W}^{(0)} \mathbf{x}$ is with high probability of subconstant order. Will this planting be computationally-quiet? Perhaps the most natural algorithm for distinguishing \mathbb{P}_n from \mathbb{Q}_n in this setting is to compute the largest eigenvalue of \mathbf{W} , since adding a sufficiently large rank-one component under \mathbb{P}_n will eventually make the largest eigenvalue larger than that under \mathbb{Q}_n . We refer to this as

the *PCA test* (more properly, we compute the largest eigenvalue and then threshold it to perform the test). Since typically $\lambda_{\max}(\mathbf{W}) \approx 2$ under $\mathbf{W} \sim \mathbb{Q}_n = \text{GOE}(n)$ and we are adding a matrix with largest eigenvalue approximately 2, we might expect that the PCA test would not notice the deformation of \mathbf{W} . However, results of random matrix theory tell us that this is not the case. Instead, the added rank-one component $\lambda \mathbf{x} \mathbf{x}^\top$ and the unstructured noise of $\mathbf{W}^{(0)} \sim \text{GOE}(n)$ “cooperate” to produce an eigenvalue larger than 2 even for $\lambda < 2$. The following result characterizes this so-called “pushout effect” in our setting, a variant of the celebrated *Baik-Ben Arous-Péché (BBP) transition* [BBAP05].

Proposition 2.2.3 ([FP07, BGN11]). *Write $\mathbf{v}_{\max}(\mathbf{W})$ for the eigenvector corresponding to the largest eigenvalue of a matrix \mathbf{W} . For $\lambda > 0$ and any $\mathbf{x} \in \mathbb{S}^{n-1}$, when $\mathbf{W} = \lambda \mathbf{x} \mathbf{x}^\top + \mathbf{W}^{(0)}$ for $\mathbf{W}^{(0)} \sim \text{GOE}(n)$, the following hold.*

- *If $\lambda \leq 1$, then $\lambda_{\max}(\mathbf{W}) \rightarrow 2$ and $\langle \mathbf{v}_{\max}(\mathbf{W}), \mathbf{x} \rangle^2 \rightarrow 0$ almost surely.*
- *If $\lambda > 1$, then $\lambda_{\max}(\mathbf{W}) \rightarrow \lambda + \lambda^{-1} > 2$ and $\langle \mathbf{v}_{\max}(\mathbf{W}), \mathbf{x} \rangle^2 \rightarrow 1 - \lambda^{-2} \in (0, 1)$ almost surely.*

Thus additive planting can be computationally-quiet only for $\lambda \leq 1$, but will actually plant a sufficiently high-quality solution to saturate the spectral bound only for $\lambda \geq 2$. Reducing certification to testing in this kind of additively-planted model then cannot satisfy all of our criteria simultaneously.

The crux of the issue is in the second parts of the two results of Proposition 2.2.3: when $\lambda_{\max}(\mathbf{W}) > 2$ under $\mathbf{W} \sim \mathbb{P}_n$, then \mathbf{x} is not actually the top eigenvector of \mathbf{W} . The additive planting therefore “wastes some of its effort,” in that it unintentionally plants a vector $\mathbf{x}' \neq \mathbf{x}$ with $\mathbf{x}'^\top \mathbf{W} \mathbf{x}' > \mathbf{x}^\top \mathbf{W} \mathbf{x}$, and this \mathbf{x}' will depend in part on \mathbf{x} and in part on the GOE matrix $\mathbf{W}^{(0)}$, so there is no guarantee that \mathbf{x}' is close to \mathcal{X} . To correct this misbehavior, we construct a *spectral planting*, where \mathbf{x} is instead planted in a more direct way to lie very close to the top few eigenspaces of $\mathbf{W} \sim \mathbb{Q}_n$, without changing the eigenvalue

distribution of \mathbf{W} and therefore without creating a pushout effect that can be detected by the PCA test.

We now describe how to achieve this in reasonable generality when $\mathbb{Q}_n = \text{GOE}(n)$, using a modification of a well-known spiked matrix model where we apply a negative spike instead of the more usual positive spike. (This model itself is quite similar to the Wigner spiked matrix model mentioned above, as we will also discuss at length later in Chapters 4 and 5, but the way we use it for spectral planting will be different from how we used the Wigner model for the additive planting.)

Definition 2.2.4 (Wishart spiked matrix model [Joh01]). *Let $k \in \mathbb{N}_+$ (not depending on n), and let $\mathcal{P}_n \in \mathcal{P}(\mathbb{R}^{n \times k})$. The Wishart spiked matrix model with signal strength $\beta > -1$, sampling ratio γ , and spike prior (\mathcal{P}_n) is specified by the following distributions over $(\mathbf{y}_1, \dots, \mathbf{y}_N) \in (\mathbb{R}^n)^N$ with $N = N(n) = \lfloor n/\gamma \rfloor$:*

- Under \mathbb{Q}_n , draw $\mathbf{y}_1, \dots, \mathbf{y}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ independently.
- Under \mathbb{P}_n , first draw $\mathbf{X}^{(0)} \sim \mathcal{P}_n$, and define

$$\mathbf{X} := \begin{cases} \mathbf{X}^{(0)} & \text{if } \beta \|\mathbf{X}\|^2 > -1, \\ \mathbf{0} & \text{otherwise.} \end{cases} \quad (2.2)$$

Then, draw $\mathbf{y}_1, \dots, \mathbf{y}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n + \beta \mathbf{X} \mathbf{X}^\top)$ independently.

More briefly, we say that $(\mathbb{Q}_n, \mathbb{P}_n)_{n \geq 1}$ form a Wishart spiked matrix model with parameters $(\beta, \gamma, \mathcal{P}_n)$. We call k the rank of such a model or of the spike prior, and we call the model negatively-spiked if $\beta < 0$.

We note that often the \mathbf{y}_i are viewed as being organized into a matrix $\mathbf{Y} \in \mathbb{R}^{n \times N}$, and in particular the sample covariance matrix $\frac{1}{N} \mathbf{Y} \mathbf{Y}^\top$ plays an important role in algorithms for estimating the spike \mathbf{X} , which is why this model is viewed as belonging to the family of

spiked matrix models. While the positively-spiked version of this model has been studied at great length, the negatively-spiked version is less common and apparently was first attended to only by [PWBM18]. We discuss their results on this case in Section 2.4 below.

We also introduce an assumption that we will always assume to hold when we are working with such a model in this section, saying that we are almost surely in the first case of (2.2) above. When we later arrive at our applications of these results in Section 5.3, we will present a simple way to adjust a spike prior to satisfy this.

Definition 2.2.5. *A spike prior (\mathcal{P}_n) is β -good for the Wishart spiked matrix model if $\beta\|\mathbf{X}\|^2 > -1$ almost surely when $\mathbf{X} \sim \mathcal{P}_n$.*

Using this model we describe how to achieve a spectral planting of \mathbf{X} in the top few eigenspaces of the GOE.

Definition 2.2.6 (Spectrally-planted GOE model). *Let $(\mathbb{Q}_n^{\text{Wish}}, \mathbb{P}_n^{\text{Wish}})_{n \geq 1}$ be the null and planted distributions for a Wishart spiked matrix model with parameters $(\beta, \gamma, \mathcal{P}_n)$ with $\gamma > 1$. Define $\mathbb{Q}_n, \mathbb{P}_n \in \mathcal{P}(\mathbb{R}_{\text{sym}}^{n \times n})$ in the following way. Suppose $\mathbb{D}_n \in \{\mathbb{Q}_n, \mathbb{P}_n\}$, and let $\mathbb{D}_n^{\text{Wish}}$ equal $\mathbb{Q}_n^{\text{Wish}}$ if $\mathbb{D}_n = \mathbb{Q}_n$ and equal $\mathbb{P}_n^{\text{Wish}}$ if $\mathbb{D}_n = \mathbb{P}_n$. To sample from \mathbb{D}_n , first draw $\tilde{\mathbf{W}} \sim \text{GOE}(n)$ and $\mathbf{y}_1, \dots, \mathbf{y}_N \sim \mathbb{D}_n^{\text{Wish}}$. Let $\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_N$ be an orthonormal basis for the span of the $\mathbf{y}_1, \dots, \mathbf{y}_N$, and $\hat{\mathbf{y}}_{N+1}, \dots, \hat{\mathbf{y}}_n$ be an orthonormal basis for the orthogonal complement of this span. Let $\lambda_1 > \dots > \lambda_n$ be the eigenvalues of $\tilde{\mathbf{W}}$. Then, draw $\mathbf{W} := \sum_{i=1}^n \lambda_i \hat{\mathbf{y}}_{n-i+1} \hat{\mathbf{y}}_{n-i+1}^\top$. We call the sequence of pairs of distributions $(\mathbb{Q}_n, \mathbb{P}_n)$ defined in this way the spectrally-planted GOE model with parameters $(\beta, \gamma, \mathcal{P}_n)$.*

Note that, when $\beta < 0$, then $\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_N$ span an N -dimensional subspace that is biased away from the directions of the columns of \mathbf{X} . Since we build \mathbf{W} above to have as its top eigenvectors a basis for the orthogonal complement of the span of these, in this definition we are, somewhat indirectly, biasing the top eigenspace of \mathbf{W} towards the directions of \mathbf{X} , which is precisely what we claimed the spectral planting would achieve. The closer β is to

-1 , the stronger this bias is; the closer γ is to 1, the smaller the fraction $1 - \gamma^{-1}$ of directions occupied by this top eigenspace.

The following simple result shows that this model is indeed essentially just a “cloaked” version of the Wishart spiked matrix model.

Proposition 2.2.7. *If it is possible to distinguish between \mathbb{P}_n and \mathbb{Q}_n in the spectrally-planted GOE model with parameters $(\beta, \gamma, \mathcal{P}_n)$ in time $T(n)$, then it is possible to distinguish between \mathbb{P}_n and \mathbb{Q}_n in the Wishart spiked matrix model with parameters $(\beta, \gamma, \mathcal{P}_n)$ in time $T(n) + O(\text{poly}(n))$ using a randomized test, and likewise with the roles of the models reversed.*

Proof. From the definition, we may sample from \mathbb{P}_n or \mathbb{Q}_n in the spectrally-planted GOE model by sampling (y_1, \dots, y_N) from \mathbb{P}_n or \mathbb{Q}_n (respectively) in the Wishart spiked matrix model, sampling $\mathbf{W} \sim \text{GOE}(n)$ independently, and outputting a function $g_n(\mathbf{y}_1, \dots, \mathbf{y}_N, \mathbf{W})$, where g_n is computable in polynomial time in n . Thus given tests f_n that distinguish with high probability in the spectrally-planted GOE model, randomized tests f'_n outputting $f'_n(\mathbf{Y}) = f_n(g_n(\mathbf{Y}, \mathbf{W}))$ for $\mathbf{W} \sim \text{GOE}(n)$ will distinguish with high probability in the Wishart spiked matrix model. For the second statement, the result follows immediately from the laws of the top eigenvectors under \mathbb{P}_n and \mathbb{Q}_n in the spectrally-planted GOE model being, by construction, precisely the laws of observations under \mathbb{P}_n and \mathbb{Q}_n respectively in the Wishart spiked matrix model with the same parameters. \square

Remark 2.2.8 (Reductions). *To simplify our language and work slightly informally without a precisely-specified computational model, we state results of the above form in terms of the existence of algorithms with prescribed runtimes. However, stronger formal statements in terms of reductions also hold; for example, above, there exists a randomized polynomial-time reduction from testing under the Wishart spiked matrix model to testing under the spectrally-planted GOE model and a deterministic polynomial-time reduction from testing under the spectrally-planted GOE model to testing under the Wishart spiked matrix model.*

2.3 TOOLS FOR REDUCTIONS UNDER THE GOE

We now proceed to more technical results justifying that spectral planting actually achieves what we have outlined above. The following is the key result showing that, under \mathbb{P}_n of the spectrally-planted GOE model, the sampled $\mathbf{X} \sim \mathcal{P}_n$ indeed lies near the top eigenspaces of the sampled matrix. In fact it will be useful to be slightly more general and consider \mathbf{X}' that is itself close to \mathbf{X} , for the purposes of our applications, but the key intuitions are captured by thinking of $\mathbf{X}' = \mathbf{X}$ below.

Theorem 2.3.1. *If $(\mathbb{Q}_n, \mathbb{P}_n)$ are the null and planted distributions of the spectrally-planted GOE model with parameters $(\beta, \gamma, \mathcal{P}_n)$, then $\mathbb{Q}_n = \text{GOE}(n)$. Moreover, for any $\epsilon > 0$, there exist $C > 0$, $\beta > -1$, and $\gamma > 1$ such that the following holds. Let $k \geq 1$ and $\mathcal{P}_n \in \mathcal{P}(\mathbb{R}^{n \times k})$ be a β -good spike prior. Suppose further that there exists constants $K, L > 0$ and a random variable \mathbf{X}' coupled to $\mathbf{X} \sim \mathcal{P}_n$ such that $\|\mathbf{X}'\|_F^2 \rightarrow K$ in probability and $\|\mathbf{X}^\top \mathbf{X}'\|_F^2 \geq L$ with high probability. Then, letting \mathbf{X}' be the variable coupled to the spike matrix \mathbf{X} drawn in the Wishart spiked matrix model, under $\mathbf{W} \sim \mathbb{P}_n$ we have $\text{tr}(\mathbf{X}'^\top \mathbf{W} \mathbf{X}') \geq 2(K - C(K - L) - \epsilon)$ with high probability.*

Our proof will use the following technical result, which, in the above setting, controls how well $\sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^\top$ approximates $\sum_{i=1}^N \hat{\mathbf{y}}_i \hat{\mathbf{y}}_i^\top$.

Proposition 2.3.2. *Suppose $\mathbf{y}_1, \dots, \mathbf{y}_N \sim \mathbb{P}_n$ in a Wishart spiked matrix model with $\gamma > 1$ and $\beta \leq 0$. Then, for any $\epsilon > 0$, with high probability*

$$(1 - \epsilon)(\sqrt{\gamma} - 1)^2 \sum_{i=1}^N \hat{\mathbf{y}}_i \hat{\mathbf{y}}_i^\top \preceq \frac{1}{N} \sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^\top \preceq (1 + \epsilon)(\sqrt{\gamma} + 1)^2 \sum_{i=1}^N \hat{\mathbf{y}}_i \hat{\mathbf{y}}_i^\top. \quad (2.3)$$

Proof. Since $\sum_{i=1}^N \hat{\mathbf{y}}_i \hat{\mathbf{y}}_i^\top$ is the orthogonal projector to the row space of $\frac{1}{N} \sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^\top$, and this

matrix has exactly N non-zero eigenvalues almost surely, it suffices to show that

$$(1 - \epsilon)(\sqrt{y} - 1)^2 \leq \lambda_N \left(\frac{1}{N} \sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^\top \right) \leq \lambda_1 \left(\frac{1}{N} \sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^\top \right) \leq (1 + \epsilon)(\sqrt{y} + 1)^2. \quad (2.4)$$

This follows immediately from Theorem 1.2 of [BS06], noting that the \mathbf{y}_i are drawn from a spiked Wishart model with a covariance matrix with a constant number k of eigenvalues different than 1, all of which are smaller than 1. \square

Proof of Theorem 2.3.1. Note that, since (\mathcal{P}_n) is β -good, almost surely we are in the first case of (2.2) in the definition of the Wishart spiked matrix model, so that $\mathbf{y}_1, \dots, \mathbf{y}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n + \beta \mathbf{X} \mathbf{X}^\top)$. We begin with a direct computation. We note that $\langle \mathbf{A}, \mathbf{B} \rangle := \text{tr}(\mathbf{A}^\top \mathbf{B})$ denotes the Frobenius inner product for any \mathbf{A}, \mathbf{B} of equal size.

$$\begin{aligned} \text{tr}(\mathbf{X}'^\top \mathbf{W} \mathbf{X}') &= \langle \mathbf{X}' \mathbf{X}'^\top, \mathbf{W} \rangle \\ &= \left\langle \mathbf{X}' \mathbf{X}'^\top, \sum_{i=1}^n \lambda_i \hat{\mathbf{y}}_{n-i+1} \hat{\mathbf{y}}_{n-i+1}^\top \right\rangle \\ &\geq \left\langle \mathbf{X}' \mathbf{X}'^\top, \lambda_{n-N} \sum_{i=N+1}^n \hat{\mathbf{y}}_i \hat{\mathbf{y}}_i^\top + \lambda_n \sum_{i=1}^N \hat{\mathbf{y}}_i \hat{\mathbf{y}}_i^\top \right\rangle \\ &= \left\langle \mathbf{X}' \mathbf{X}'^\top, \lambda_{n-N} \left(\mathbf{I}_n - \sum_{i=1}^N \hat{\mathbf{y}}_i \hat{\mathbf{y}}_i^\top \right) + \lambda_n \sum_{i=1}^N \hat{\mathbf{y}}_i \hat{\mathbf{y}}_i^\top \right\rangle \\ &= \lambda_{n-N} \|\mathbf{X}'\|_F^2 - (\lambda_{n-N} - \lambda_n) \left\langle \mathbf{X}' \mathbf{X}'^\top, \sum_{i=1}^N \hat{\mathbf{y}}_i \hat{\mathbf{y}}_i^\top \right\rangle. \end{aligned} \quad (2.5)$$

We consider the two terms here separately. In the first term, we have $\|\mathbf{X}'\|_F^2 \rightarrow K$ in probability by assumption. For the second term, by Proposition 2.3.2 we have

$$\left\langle \mathbf{X}' \mathbf{X}'^\top, \sum_{i=1}^N \hat{\mathbf{y}}_i \hat{\mathbf{y}}_i^\top \right\rangle \leq \frac{2}{(\sqrt{y} - 1)^2 N} \left\langle \mathbf{X}' \mathbf{X}'^\top, \sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^\top \right\rangle$$

and, viewing $\mathbf{y}_i = (\mathbf{I} + \beta \mathbf{X} \mathbf{X}^\top)^{1/2} \mathbf{g}_i$ for \mathbf{g}_i independent from \mathbf{X} and \mathbf{X}' , we have

$$\begin{aligned}
&= \frac{2}{(\sqrt{\gamma} - 1)^2} \left\langle (\mathbf{I} + \beta \mathbf{X} \mathbf{X}^\top)^{1/2} \mathbf{X}' \mathbf{X}'^\top (\mathbf{I} + \beta \mathbf{X} \mathbf{X}^\top)^{1/2}, \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i \mathbf{g}_i^\top \right\rangle \\
&\leq \frac{2}{(\sqrt{\gamma} - 1)^2} \left\| \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i \mathbf{g}_i^\top \right\| \operatorname{tr} \left((\mathbf{I} + \beta \mathbf{X} \mathbf{X}^\top)^{1/2} \mathbf{X}' \mathbf{X}'^\top (\mathbf{I} + \beta \mathbf{X} \mathbf{X}^\top)^{1/2} \right) \\
&= \frac{2}{(\sqrt{\gamma} - 1)^2} \left\| \frac{1}{N} \sum_{i=1}^N \mathbf{g}_i \mathbf{g}_i^\top \right\| (\|\mathbf{X}'\|_F^2 + \beta \|\mathbf{X}^\top \mathbf{X}'\|_F^2) \tag{2.6}
\end{aligned}$$

Here, the norm factor is bounded by $2(\sqrt{\gamma} + 1)^2$ with high probability again by Proposition 2.3.2, so for any given γ there will be a constant $C > 0$ such that, with high probability,

$$\leq \frac{C}{5} (\|\mathbf{X}'\|_F^2 + \beta \|\mathbf{X}^\top \mathbf{X}'\|_F^2). \tag{2.7}$$

Also for any γ , we will have $\lambda_{n-N} - \lambda_n \leq \lambda_1 - \lambda_n \leq 5$ with high probability. Thus, choosing γ sufficiently close to 1 depending on ϵ and then β sufficiently close to -1 , we will have that

$$\operatorname{tr}(\mathbf{X}'^\top \mathbf{W} \mathbf{X}') \geq \left(2 - \frac{\epsilon}{K}\right) K - C \left(K - L + \frac{\epsilon}{C}\right) = 2(K - C(K - L) - \epsilon), \tag{2.8}$$

as claimed. □

Finally, we may use this to derive a reduction from hypothesis testing in the Wishart spiked matrix model to certification, for suitable spike priors. We give two versions: first, using the flexibility introduced in Theorem 2.3.1, we allow the planted direction \mathbf{X} in the spike prior of the Wishart model to not be exactly in \mathcal{X} , but only nearby. This is helpful, for example, if \mathcal{X} is a cone that cannot support a non-trivial centered distribution, which will interfere with our arguments for arguing hardness of testing for Wishart models. We will encounter this type of situation in Section 5.3 when we discuss non-negative PCA, where \mathcal{X} is the positive orthant $\mathbb{R}_+^n \subset \mathbb{R}^n$.

Corollary 2.3.3 (Reduction: planting near \mathcal{X}). *Let $K > 0$. Suppose that, for any $\delta > 0$ and $\beta > -1$, there exist $\mathcal{P}_{\beta,\delta,n} \in \mathcal{P}(\mathbb{R}^{n \times k})$ β -good spike priors such that there is a random variable \mathbf{X}' coupled to $\mathbf{X} \sim \mathcal{P}_{\beta,\delta,n}$ with $\mathbf{X}' \in \mathcal{X}$ almost surely, $\|\mathbf{X}'\|_F^2 \rightarrow K$ in probability, and $\|\mathbf{X}^\top \mathbf{X}'\|_F^2 \geq K - \delta$ with high probability. Suppose also that there exists an algorithm that runs in time $T(n)$ and certifies a bound of at most $2(K - \epsilon)$ on $M_{\mathcal{X}}(\mathbf{W})$ when $\mathbf{W} \sim \text{GOE}(n)$ with high probability for some $\epsilon > 0$. Then, there exist $\beta \in (-1, 0)$, $\gamma > 1$, $\delta > 0$, and an algorithm that can distinguish (\mathbb{P}_n) and (\mathbb{Q}_n) in the Wishart spiked matrix model with parameters $(\beta, \gamma, \mathcal{P}_{\beta,\delta,n})$ in time $T(n) + O(\text{poly}(n))$.*

Proof. By Theorem 2.3.1, there exist $C > 0$, $\beta > -1$ and $\gamma > 1$ such that, when $\mathbf{W} \sim \mathbb{P}_n$ under the spectrally-planted GOE model with parameters $(\beta, \gamma, \mathcal{P}_{\beta,\delta,n})$, then $M_{\mathcal{X}}(\mathbf{W}) \geq \text{tr}(\mathbf{X}'^\top \mathbf{W} \mathbf{X}') \geq 2(K - C\delta - \frac{1}{3}\epsilon)$ with high probability. Choosing δ small enough, we may further ensure that $M_{\mathcal{X}}(\mathbf{W}) \geq 2(K - \frac{2}{3}\epsilon)$ with high probability. Then, by Theorem 2.2.1, it is possible to distinguish \mathbb{P}_n from \mathbb{Q}_n in this spectrally-planted GOE model in time $T(n) + O(1)$. Finally, by Proposition 2.2.7 it is possible to do the same in the Wishart spiked matrix model in time $T(n) + O(\text{poly}(n))$. \square

Sometimes, it will be possible to just take $\mathbf{X}' = \mathbf{X}$ except on an event of low probability, in which case we can dispense with the dependence on δ and simplify the statement in the following way. By $\boldsymbol{\sigma}(\mathbf{X})$ we denote the vector of singular values of \mathbf{X} .

Corollary 2.3.4 (Reduction: planting in \mathcal{X}). *Let $K > 0$. Suppose that, for any $\beta > -1$, there exist $\mathcal{P}_n \in \mathcal{P}(\mathbb{R}^{n \times k})$ β -good spike priors, when $\mathbf{X} \sim \mathcal{P}_n$ then $\mathbf{X} \in \mathcal{X}$ with high probability, and both $\|\boldsymbol{\sigma}(\mathbf{X})\|_2^2 \rightarrow K$ and $\|\boldsymbol{\sigma}(\mathbf{X})\|_4^4 \rightarrow K$ in probability. Suppose also that there exists an algorithm that runs in time $T(n)$ and certifies a bound of at most $2(K - \epsilon)$ on $M_{\mathcal{X}}(\mathbf{W})$ when $\mathbf{W} \sim \text{GOE}(n)$ with high probability for some $\epsilon > 0$. Then, there exist $\beta \in (-1, 0)$, $\gamma > 1$, and an algorithm that can distinguish (\mathbb{P}_n) and (\mathbb{Q}_n) in the Wishart spiked matrix model with parameters $(\beta, \gamma, \mathcal{P}_n)$ in time $T(n) + O(\text{poly}(n))$.*

The proof is immediate from Corollary 2.3.3 taking $\mathbf{X}' = \mathbf{X}$.

Remark 2.3.5. *The conditions $\|\sigma(\mathbf{X})\|_2^2 \rightarrow K$ and $\|\sigma(\mathbf{X})\|_4^4 \rightarrow K$ in probability imply, loosely speaking, that as $n \rightarrow \infty$, \mathbf{X} has K non-zero singular values that are all approximately equal to 1. Thus K should be seen as the “effective rank” of the spike \mathbf{X} , in contrast to the dimensionality k . We will not find this essential in any of our examples, though one case where it allows a more explicit representation of \mathcal{X} is in the constraint set of the Potts spin glass Hamiltonian that we treat in Section 5.3.*

2.4 HARD REGIME IN THE WISHART SPIKED MATRIX MODEL

As discussed earlier, our reductions only stand a chance of being useful if the Wishart spiked matrix model has a statistical-computational gap and corresponding *hard regime*, where it is possible but difficult to perform hypothesis testing. The kind of spike prior with respect to which we would like this to be true depends on the constraint set \mathcal{X} of the certification problem in question. For example, for the case $\mathcal{X} = \{\pm 1/\sqrt{n}\}^n$ of the SK Hamiltonian, the natural choice is the *Rademacher* spike prior, $\mathcal{P}_n = \text{Unif}(\{\pm 1/\sqrt{n}\}^n)$.

On the basis of much of the literature on spiked matrix models, it would be tempting to think that this hard regime is unlikely to occur. That is because by far the best-studied hard regimes are for *sparse* priors. We briefly recount these results below.

Example 2.4.1 (Constant-sparsity PCA). *Consider the Wigner spiked matrix model under the sparse Rademacher prior where $\mathbf{x} \sim \mathcal{P}_n$ has i.i.d. entries drawn as*

$$x_i = \begin{cases} -\frac{1}{\sqrt{\rho n}} & \text{with probability } \rho/2, \\ 0 & \text{with probability } 1 - \rho, \\ +\frac{1}{\sqrt{\rho n}} & \text{with probability } \rho/2. \end{cases} \quad (2.9)$$

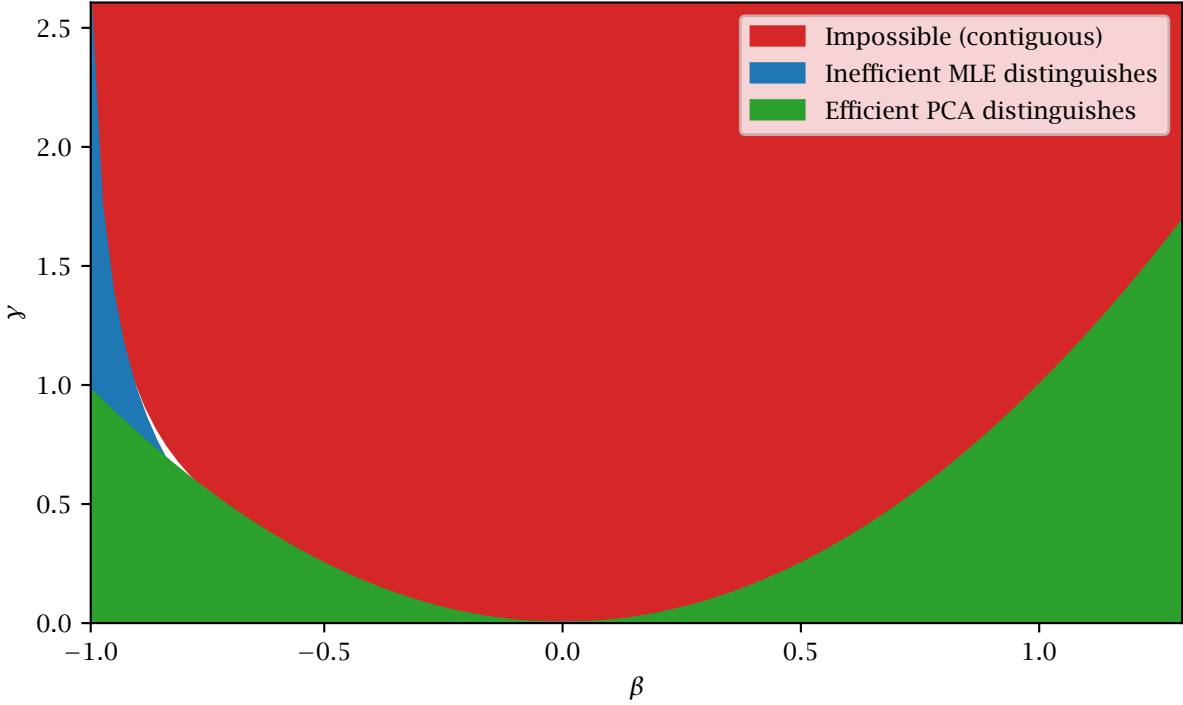


Figure 2.1: Phase diagram of the Wishart spiked matrix model. We illustrate the regimes in the (β, γ) plane where the Wishart spiked matrix model with Rademacher prior $\mathcal{P}_n = \text{Unif}(\{\pm 1/\sqrt{n}\}^n)$ has various computational complexities. This is identical up to cosmetics to Figure 3 of [PWBM18]; the small white region is not covered by their results and likely should be split between the red and blue regions under more precise analysis.

In these models a hard regime was first conjectured based on computations with the replica method for all $\rho \in (0, \rho^)$ with $\rho^* \approx 0.092$ [LKZ15b]. It is now rigorously established that, indeed, when $\rho < \rho^*$ then an inefficient test succeeds in a parameter range where the PCA test does not succeed (for all $\lambda \in (\lambda^*, 1)$ for a specific threshold $\lambda^* < 1$, in the notation of our earlier discussion in Section 2.2), while outside of this range no test succeeds [KXZ16, BDM⁺16, LM19, BMV⁺18, PWBM18, EAKJ20]. The same methods suggest that not only the PCA test but no polynomial-time test should succeed in these latter regimes (a conjecture that our results in Section 5.2.4 bolster with lower bounds against low-degree polynomial tests).*

This prominent line of work has led to a plausible intuition that sparsity is the key feature leading to hard regimes in spiked matrix models.

However, as observed by [PWBM18], hard regimes also arise for *dense* priors, like the Rademacher prior that we hope to work with, in the unusual case of the negatively-spiked Wishart model. They show, similar to the above discussion, that for various discrete spike priors including the Rademacher prior, there is a parameter regime where the efficient PCA test fails and an inefficient test succeeds. This inefficient test for the Rademacher prior searches by brute force over all $\mathbf{x} \in \{\pm 1/\sqrt{n}\}^n$ for the one minimizing $\sum_{i=1}^N \langle \mathbf{x}, \mathbf{y}_i \rangle^2$, which we note is compatible with our remark that, if we could reduce testing in the Wishart model to certification, then certification by the same brute force search over all feasible \mathbf{x} should yield a successful test. We illustrate the “phase diagram” of which values of (β, γ) give rise to what behavior in Figure 2.1; the blue region in the plot is the hard regime, which we note includes the region of γ slightly larger than 1 and β close to -1 that corresponds to the limiting cases of our reductions earlier.

2.5 AN EXAMPLE OF SPECTRAL PLANTING IN GRAPHS

It may seem from the preceding discussion that the strategy of spectral planting is quite specific to the GOE, or at least to the setting of continuous problem instances where the instance can be deformed freely. While it is certainly true that these kinds of problems make spectral planting more convenient and straightforward, we digress to present another surprising example where it may be performed, for the problem of *graph coloring*. These results are taken from [BBK⁺20]; the author was only modestly involved in formulating the arguments concerning graph models in this article, so we give only a brief overview.

A *k-coloring* of a graph $G = (V, E)$ is a map $\sigma : V \rightarrow [k]$ such that, whenever $i \sim j$ (meaning $\{i, j\} \in E$) then $\sigma(i) \neq \sigma(j)$. Coloring appears to be a combinatorial optimization problem, but it may also be formulated as constrained PCA in the following way. Suppose $|V| = n$. Let $\mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{S}^{k-1}$ be unit vectors pointing to the vertices of an equilateral

simplex, such that $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = -\frac{1}{k-1}$ whenever $i \neq j$. Let $\mathcal{X} \subset \mathbb{R}^{n \times (k-1)}$ be the set of matrices \mathbf{X} all of whose rows equal $\sqrt{\frac{k-1}{n}} \mathbf{v}_i$ for some $i \in [k]$. Note that this normalization makes the columns of \mathbf{X} have approximately unit norm for the rows chosen uniformly from $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$, following our convention. Then, letting \mathbf{A} be the adjacency matrix of G , we have

$$G \text{ is } k\text{-colorable} \Leftrightarrow M_{\mathcal{X}}(-\mathbf{A}) = \frac{2|E|}{n}, \quad (2.10)$$

where the right-hand side is the maximum possible value of $M_{\mathcal{X}}(-\mathbf{A})$.

Based on this formulation, there is a natural spectral bound for coloring (originally due to Hoffman [Hof70]). We note first that \mathbf{A} always has at least one non-positive eigenvalue since its trace is zero, so $\lambda_{\min}(\mathbf{A}) \leq 0$. Then, by the ordinary spectral bound on constrained PCA we have, since $\|\mathbf{X}\|_F^2 = k-1$ for any $\mathbf{X} \in \mathcal{X}$,

$$M_{\mathcal{X}}(\mathbf{W}) \leq |\lambda_{\min}(\mathbf{A})|(k-1), \quad (2.11)$$

and thus, letting $\chi(G)$ be the minimum k for which G is k -colorable, we have

$$\chi(G) \geq 1 + \frac{2|E|}{n|\lambda_{\min}(\mathbf{A})|}. \quad (2.12)$$

In particular, when G is d -regular, then $|E| = nd/2$, so the bound takes on the simpler form

$$\chi(G) \geq 1 + \frac{d}{|\lambda_{\min}(\mathbf{A})|}. \quad (2.13)$$

A natural random model to test the efficacy of this as a certifiable lower bound on $\chi(G)$ is the random graph distribution $\text{Reg}(n, d)$ of uniformly-random d -regular graphs on n vertices (see, e.g., [Wor99] for information on this distribution, including how to sample from it—which seems *a priori* unclear—using the “configuration model”). When $G \sim \text{Reg}(n, d)$,

then a series of combinatorial works have gradually shown that, for any given d , $\chi(G)$ concentrates on one or two consecutive numbers as $n \rightarrow \infty$ [AM04, AN04, SW07, KPGW10, COEH16], with asymptotic scaling $\chi(G) \sim \frac{1}{2} \frac{d}{\log d}$. On the other hand, by Friedman’s celebrated theorem [Fri03] we have $|\lambda_{\min}(\mathbf{A})| \approx 2\sqrt{d-1}$, and thus the spectral bound (2.13) can only certify the much smaller lower bound $\chi(G) \gtrsim \frac{1}{2}\sqrt{d}$. Moreover, a pair of works on the Lovász ϑ function SDP relaxation [CO03, BKM19], which is equivalent to the degree 2 SOS relaxation, showed that this relaxation achieves no better than the spectral bound, suggesting that the problem of certifying bounds on the chromatic number of random d -regular graphs might exhibit a spectral barrier.

As [BKM19] also observed, like the additive planting in Section 2.2, a natural strategy for planting a coloring is *not* quiet enough to plant one both quietly and saturating the spectral bound. This planting is an extremal variant of the *stochastic block model* (which we will discuss further in Section 5.4.1), which amounts to choosing σ uniformly at random and conditioning G on σ being a proper coloring. Such a planting with k colors, however, can be detected once $k < 1 + \sqrt{d}$, so this strategy cannot quietly plant a coloring that asymptotically meets the spectral bound but only one with roughly twice the number of colors of the spectral bound.

Thus it is natural to search for a quieter planting that places \mathbf{X} corresponding to σ (i.e., with the i th row of \mathbf{X} equal to $\mathbf{v}_{\sigma(i)}$) in the bottom eigenspaces of \mathbf{A} . While perhaps *a priori* this seems difficult, it is actually achieved by planting the following special type of coloring having an extra combinatorial structure.

Definition 2.5.1. *A coloring is equitable if each vertex $v \in V$ has an equal number of neighbors of each color; i.e., $|\{w \in V : w \sim v, \sigma(w) = i\}| = d/k$ for all $i \in [k] \setminus \{\sigma(v)\}$.*

The following pair of results bolster our hopes that this is the correct notion to work with.

Proposition 2.5.2. *If a coloring saturates Hoffman’s bound (2.13), then it is equitable.*

See, e.g., [Abi19] for a proof.

Proposition 2.5.3. *If a coloring σ of G is equitable, and $\mathbf{X} \in \mathcal{X}$ corresponds to σ , then every column of \mathbf{X} is an eigenvector of \mathbf{A} with eigenvalue $-\frac{d}{k-1}$.*

We note here that when $k \sim \frac{1}{2}\sqrt{d}$ then the corresponding eigenvalue is $\sim -2\sqrt{d}$, precisely the smallest eigenvalue a graph of $\text{Reg}(n, d)$ will have by Friedman's theorem.

Thus spectral planting of colorings in random regular graphs is achieved by choosing an equitable coloring σ uniformly at random (so long as the parameters n, d, k satisfy the necessary divisibility conditions) and then choosing a graph conditional on σ being a proper coloring. In [BBK⁺20], we go on to verify with various evidence that this planting is indeed quiet down to $k \approx \frac{1}{2}\sqrt{d}$, and to carry out a reduction strategy similar to that in this chapter to argue that graph coloring exhibits a spectral barrier. In Section 5.3, we will treat the Potts spin glass Hamiltonian, a variant of the SK Hamiltonian that is essentially the “Gaussian version” of graph coloring; this will entail working with the same constraint set \mathcal{X} , but the nuance of restricting our attention to equitable colorings will no longer play an important role. We leave it as an intriguing open problem to understand what other manifestations spectral planting can have in combinatorial problems.

3 | PRELIMINARIES ON LOW-DEGREE POLYNOMIAL ALGORITHMS

In the previous chapter we have seen how, when an efficient algorithm exists for better-than-spectral certification in constrained PCA, then we can often show that an efficient algorithm exists for hypothesis testing between a natural pair of associated models of the top eigenspaces of problem instances. To show that better-than-spectral certification is hard, it then suffices to show that this ancillary hypothesis testing problem is hard. Moreover, these eigenspace models may be put into the convenient form of *negatively-spiked Wishart models*, a Gaussian model that is relatively tractable to work with.

We have also seen in passing that to identify computational hardness in such problems it is not in general enough to identify the stronger property of statistical or information-theoretic impossibility. Instead, some testing problems exhibit *hard regimes*, where it is possible to test in exponential or subexponential-yet-superpolynomial time, but appears impossible to test in polynomial time. Therefore, we now take up the natural next question of how to diagnose this latter situation.

There are numerous approaches to producing evidence of hardness of hypothesis testing, which mostly boil down to showing (or sometimes making simplified computations suggesting) that a particular class of algorithms fails to distinguish between two distributions. Here we will describe how to analyze the testing problems arising from certification through

the lens of algorithms computing *low-degree polynomials*. The idea of considering these algorithms in the specific technical sense we will use here arose as an offshoot of a series of works on SOS lower bounds for the planted clique problem [MW13, MPW15, DM15b, RS15, HKP⁺18, BHK⁺19], subsequently elaborated in greater generality [HKP⁺17, HS17, Hop18]. It may also be viewed independently as a generalization of a related, older framework for identifying statistical indistinguishability, the second moment method for Le Cam’s notion of contiguity [LCY12]. We therefore first review this classical topic and its more recent applications as well as the motivation from the literature on SOS for low-degree polynomial algorithms. We then give some clarifying discussion of what, concretely, the proposed computations can say about the success or failure of low-degree polynomial tests, and outline the main technique for proving such lower bounds.

SUMMARY AND REFERENCES This expository chapter does not present any new results. Our first presentation of the low-degree likelihood ratio as an adjustment of the second moment method and our later results on consequent lower bounds for thresholded polynomials and spectral algorithms are drawn from the notes [KWB19]. Our second presentation motivated by SOS is a simplification of some of the discussion in [BHK⁺19], and is also informed by the exposition in [RSS18].

3.1 LE CAM’S CONTIGUITY AND THE SECOND MOMENT METHOD

We retain the same notations (\mathbb{P}_n) and (\mathbb{Q}_n) for sequences of probability measures from our earlier presentation of hypothesis testing in Section 2.1 in the later parts of our discussion, but at first we will only be concerned with a single pair of distributions \mathbb{P} and \mathbb{Q} , defined on a single mutual measurable space (Ω, \mathcal{F}) . For the sake of simplicity, we assume in either case that \mathbb{P}_n (or \mathbb{P}) is absolutely continuous with respect to \mathbb{Q}_n (or \mathbb{Q} , as appropriate).

We have mentioned before the following basic ways of measuring quantitatively how well a test distinguishes \mathbb{P} from \mathbb{Q} .

Definition 3.1.1. *Let f be a test. The Type I error of f is the event of falsely rejecting the null hypothesis, i.e., of having $f(\mathbf{Y}) = p$ when $\mathbf{Y} \sim \mathbb{Q}$. The Type II error of f is the event of falsely failing to reject the null hypothesis, i.e., of having $f(\mathbf{Y}) = q$ when $\mathbf{Y} \sim \mathbb{P}$. The probabilities of these errors are denoted*

$$\alpha(f) := \mathbb{Q}[f(\mathbf{Y}) = p],$$

$$\beta(f) := \mathbb{P}[f(\mathbf{Y}) = q].$$

The probability $1 - \beta(f)$ of correctly rejecting the null hypothesis is called the power of f .

There is a tradeoff between Type I and Type II errors. For instance, the trivial test that always outputs p will have maximal power, but will also have maximal probability of Type I error, and vice-versa for the trivial test that always outputs q . Thus, typically one fixes a tolerance for one type of error, and then attempts to design a test that minimizes the probability of the other type.

The following celebrated result shows that it is in fact possible to identify the test that is optimal in the sense of the above tradeoff.¹

Definition 3.1.2. *Let \mathbb{P} be absolutely continuous with respect to \mathbb{Q} . The likelihood ratio (LR) of \mathbb{P} and \mathbb{Q} is*

$$L(\mathbf{Y}) := \frac{d\mathbb{P}}{d\mathbb{Q}}(\mathbf{Y}). \tag{3.1}$$

¹Importantly, we are restricting our attention to deciding between two “simple” hypotheses, where each hypothesis consists of the dataset being drawn from a specific distribution. Optimal testing is more subtle for “composite” hypotheses in parametric families of probability distributions, a more typical setting in practice. The mathematical difficulties of this extended setting are discussed thoroughly in [LR06].

The thresholded likelihood ratio test with threshold η is the test

$$L_\eta(\mathbf{Y}) := \begin{cases} p & : L(\mathbf{Y}) > \eta \\ q & : L(\mathbf{Y}) \leq \eta \end{cases}. \quad (3.2)$$

Let us first present a heuristic argument for why thresholding the likelihood ratio might be a good idea. Specifically, we will show that the likelihood ratio is optimal in a particular sense measured in $L^2(\mathbb{Q})$, i.e., when its quality is measured in terms of first and second moments of a testing quantity. Below, and whenever we discuss hypothesis testing in the context of a model (\mathbb{Q}, \mathbb{P}) or sequence of models $(\mathbb{Q}_n, \mathbb{P}_n)$, the norm of a function of the observation variable \mathbf{Y} is the norm in $L^2(\mathbb{Q})$, $\|f(\mathbf{Y})\| := (\mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}} f(\mathbf{Y})^2)^{1/2}$, and likewise the inner product is the inner product in $L^2(\mathbb{Q})$, $\langle f(\mathbf{Y}), g(\mathbf{Y}) \rangle = \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}} f(\mathbf{Y})g(\mathbf{Y})$.

Proposition 3.1.3. *If \mathbb{P} is absolutely continuous with respect to \mathbb{Q} , then the unique solution f^* of the optimization problem*

$$\begin{aligned} & \text{maximize} && \mathbb{E}_{\mathbf{Y} \sim \mathbb{P}} [f(\mathbf{Y})] \\ & \text{subject to} && \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}} [f(\mathbf{Y})^2] = 1 \end{aligned} \quad (3.3)$$

is proportional to the likelihood ratio, $f^ = L/\|L\|$, and the value of the optimization problem is $\|L\|$.*

Proof. We may rewrite the objective as

$$\mathbb{E}_{\mathbf{Y} \sim \mathbb{P}} f(\mathbf{Y}) = \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}} [L(\mathbf{Y})f(\mathbf{Y})] = \langle L, f \rangle, \quad (3.4)$$

and rewrite the constraint as $\|f\| = 1$. The result now follows since $\langle L, f \rangle \leq \|L\| \cdot \|f\| = \|L\|$ by the Cauchy-Schwarz inequality, with equality if and only if f is a scalar multiple of L . \square

In words, this means that if we want a function to be as large as possible in expectation

under \mathbb{P} while remaining bounded (in the L^2 sense) under \mathbb{Q} , we can do no better than the likelihood ratio. We will soon return to this type of L^2 reasoning in order to devise computationally-bounded statistical tests.

The following classical result shows that the above heuristic is accurate, in that the thresholded likelihood ratio tests achieve the optimal tradeoff between Type I and Type II errors.

Proposition 3.1.4 (Neyman-Pearson Lemma [NP33]). *Fix an arbitrary threshold $\eta \geq 0$. Among all tests f with $\alpha(f) \leq \alpha(L_\eta) = \mathbb{Q}(L(\mathbf{Y}) > \eta)$, L_η is the test that maximizes the power $1 - \beta(f)$.*

Since the likelihood ratio is, in the sense of the Neyman-Pearson lemma, an optimal statistical test, it also stands to reason that it should be possible to argue about statistical distinguishability solely by computing with the likelihood ratio. We present one simple method by which such arguments may be made, based on an asymptotic theory introduced by Le Cam. See, e.g., [LCY12] for a textbook treatment.

We return to working with sequences of probability measures (\mathbb{P}_n) and (\mathbb{Q}_n) over measurable spaces $(\Omega_n, \mathcal{F}_n)$, and we denote by L_n the likelihood ratio $d\mathbb{P}_n/d\mathbb{Q}_n$. Norms and inner products of functions are those of $L^2(\mathbb{Q}_n)$. The following is the crucial definition underlying the arguments to come.

Definition 3.1.5. *A sequence (\mathbb{P}_n) of probability measures is contiguous to a sequence (\mathbb{Q}_n) , written $(\mathbb{P}_n) \triangleleft (\mathbb{Q}_n)$, if whenever $A_n \in \mathcal{F}_n$ with $\mathbb{Q}_n[A_n] \rightarrow 0$, then $\mathbb{P}_n[A_n] \rightarrow 0$ as well.*

The next result gives the relevance of contiguity to asymptotically-successful testing.

Proposition 3.1.6. *If $(\mathbb{P}_n) \triangleleft (\mathbb{Q}_n)$ or $(\mathbb{Q}_n) \triangleleft (\mathbb{P}_n)$, then (\mathbb{Q}_n) and (\mathbb{P}_n) are statistically indistinguishable (in the sense of Definition 2.1.2, i.e., no test can have both Type I and Type II error probabilities tending to 0).*

Proof. We give the proof for the case $(\mathbb{P}_n) \triangleleft (\mathbb{Q}_n)$; the other case may be shown by a symmetric argument. For the sake of contradiction, let $(f_n)_{n \geq 1}$ be a sequence of tests distinguishing

(\mathbb{P}_n) and (\mathbb{Q}_n) , and let $A_n = \{f_n(\mathbf{Y}) = p\}$. Then, $\mathbb{P}_n[A_n^c] \rightarrow 0$ and $\mathbb{Q}_n[A_n] \rightarrow 0$. But, by contiguity, $\mathbb{Q}_n[A_n] \rightarrow 0$ implies $\mathbb{P}_n[A_n] \rightarrow 0$ as well, so $\mathbb{P}_n[A_n^c] \rightarrow 1$, a contradiction. \square

It therefore suffices to establish contiguity in order to prove negative results about statistical distinguishability. The following *second moment method* gives a means of establishing contiguity through a computation with the likelihood ratio.

Lemma 3.1.7 (Second moment method for contiguity). *If $\|L_n\|^2 = \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n}[L_n(\mathbf{Y})^2]$ remains bounded as $n \rightarrow \infty$ (i.e., $\limsup_{n \rightarrow \infty} \|L_n\|^2 < \infty$), then $(\mathbb{P}_n) \triangleleft (\mathbb{Q}_n)$.*

Proof. Let $A_n \in \mathcal{F}_n$. Then, using the Cauchy-Schwarz inequality,

$$\mathbb{P}_n[A_n] = \mathbb{E}_{\mathbf{Y} \sim \mathbb{P}_n}[\mathbb{1}_{A_n}(\mathbf{Y})] = \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n}[L_n(\mathbf{Y})\mathbb{1}_{A_n}(\mathbf{Y})] \leq \left(\mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n}[L_n(\mathbf{Y})^2] \right)^{1/2} (\mathbb{Q}_n[A_n])^{1/2}, \quad (3.5)$$

and so $\mathbb{Q}_n[A_n] \rightarrow 0$ implies $\mathbb{P}_n[A_n] \rightarrow 0$. \square

This second moment method has been used in recent literature to establish contiguity for various high-dimensional statistical problems (e.g., [MRZ15, PWB16, BMV⁺18, PWBM18]). Typically the null hypothesis \mathbb{Q}_n is a “simpler” distribution than \mathbb{P}_n and, as a result, $L_n = d\mathbb{P}_n/d\mathbb{Q}_n$ is easier to compute than $d\mathbb{Q}_n/d\mathbb{P}_n$. In general, and essentially for this reason, establishing $(\mathbb{Q}_n) \triangleleft (\mathbb{P}_n)$ is often more difficult than $(\mathbb{P}_n) \triangleleft (\mathbb{Q}_n)$, requiring tools such as the *small subgraph conditioning method* introduced in [RW92, RW94] and used in, e.g., [MNS15, BMNN16].² Fortunately, one-sided contiguity $(\mathbb{P}_n) \triangleleft (\mathbb{Q}_n)$ is sufficient for our purposes.

Let us remark on one limitation of the simple picture presented thus far. Note that $\|L_n\|$, the quantity that controls contiguity per the second moment method, is the same as the

²In these combinatorial results one sees the similarity between the ordinary second moment method of the probabilistic method and the second moment method for contiguity. For example, [RW92, RW94] show that random regular graphs with high probability have a Hamiltonian cycle; they may be viewed as either doing so by counting Hamiltonian cycles with the ordinary second moment method, or by showing that the random regular graph distribution with a planted Hamiltonian cycle is indistinguishable from the null distribution with the second moment method for contiguity, and at a mathematical level these approaches are equivalent.

optimal value of the L^2 optimization problem in Proposition 3.1.3:

$$\left\{ \begin{array}{l} \text{maximize } \mathbb{E}_{\mathbf{Y} \sim \mathbb{P}_n}[f(\mathbf{Y})] \\ \text{subject to } \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n}[f(\mathbf{Y})^2] = 1 \end{array} \right\} = \|L_n\|. \quad (3.6)$$

We might then be tempted to conjecture that (\mathbb{P}_n) and (\mathbb{Q}_n) are statistically distinguishable *if and only if* $\|L_n\| \rightarrow \infty$ as $n \rightarrow \infty$. However, this is incorrect: there are cases when (\mathbb{P}_n) and (\mathbb{Q}_n) are not distinguishable, yet a rare “bad” event under \mathbb{P}_n causes $\|L_n\|$ to diverge. To overcome this failure of the ordinary second moment method, some previous works (e.g., [BMNN16, BMV⁺18, PWB16, PWBM18]) have used *conditional* second moment methods to show indistinguishability, where the second moment method is applied to a modified (\mathbb{P}_n) that conditions on these bad events not occurring.

Fortunately, as we will see, the low-degree variant of this second moment method seems unafflicted by this challenge, because low-degree polynomials of \mathbf{Y} fluctuate less wildly on these bad events than do analytic functions like $L_n(\mathbf{Y})$. More generally, for settings such as sparse PCA as we have discussed in Example 2.4.1, the dependence of information-theoretic thresholds on the model parameters can be quite intricate, involving formulae coming from analysis with the replica method of statistical physics. When such results apply, second moment methods often give partial results weaker than the physics conjectures, but do not offer a sharp analysis (as demonstrated in the case of sparse PCA by the partial results of [BMV⁺18, PWBM18] before the conjectured behavior was fully established by [EAKJ20] with a method more directly informed by the physics computations). On the other hand, the low-degree predictions have proved effective in correctly identifying computational thresholds, which can be in this sense simpler to pin down.

3.2 TWO ROADS TO THE LOW-DEGREE LIKELIHOOD RATIO

We now develop an analog of the above discussion for computationally-bounded hypothesis testing. Our goal will be to justify a conjecture that an object called the *low-degree likelihood ratio*—a close relative of the classical likelihood ratio—and in particular its norm in $L^2(\mathbb{Q}_n)$ can be used to predict whether (\mathbb{P}_n) and (\mathbb{Q}_n) are computationally distinguishable. We give two ways to arrive at this idea. The first, arguably simpler approach, taken from [KWB19], develops this object by analogy with the second moment method discussed above. The second approach, more technical but historically prior to the first and stemming from the series of works [BHK⁺19, HKP⁺17, HS17, Hop18] on SOS optimization, shows that the low-degree likelihood ratio also appears in the context of the natural *pseudocalibration* strategy for proving SOS lower bounds.

3.2.1 COMPUTATIONALLY-BOUNDED SECOND MOMENT METHOD

The first approach proceeds by formulating *low-degree* analogs of the notions described in the previous section, which together constitute a method for restricting the classical decision-theoretic second moment analysis to computationally-bounded tests. The premise of this *low-degree method* is to take low-degree multivariate polynomials in the entries of the observation \mathbf{Y} as a proxy for efficiently-computable functions. (We note that the polynomials involved may have degree greater than two; the “second moment” here refers to our working with polynomials in $L^2(\mathbb{Q}_n)$, and in particular to our thinking of a sequence of polynomials being bounded under \mathbb{Q}_n as $n \rightarrow \infty$ if their second moments are bounded.)

In the computationally-unbounded case, Proposition 3.1.3 showed that the likelihood ratio optimally distinguishes (\mathbb{P}_n) from (\mathbb{Q}_n) in the L^2 sense. Following the same heuristic, we look for the low-degree polynomial that best distinguishes (\mathbb{P}_n) from (\mathbb{Q}_n) in the L^2

sense. In order for polynomials to be defined, we assume here that $\Omega_n \subseteq \mathbb{R}^N$ for some $N = N(n)$, i.e., that our data (drawn from \mathbb{P}_n or \mathbb{Q}_n) are real-valued, and we assume that \mathbb{Q}_n has finite moments so that $\mathbb{R}[\mathbf{Y}] \subset L^2(\mathbb{Q}_n)$.

Definition 3.2.1. Let $\mathbb{R}[\mathbf{Y}]_{\leq D} \subset L^2(\mathbb{Q}_n)$ denote the linear subspace of polynomials $\Omega_n \rightarrow \mathbb{R}$ of degree at most D . Let $P^{\leq D} : L^2(\mathbb{Q}_n) \rightarrow \mathbb{R}[\mathbf{Y}]_{\leq D}$ denote the orthogonal (with respect to the inner product of $L^2(\mathbb{Q}_n)$) projection operator to this subspace. Finally, define the low-degree likelihood ratio (LDLR) of degree D as $L_n^{\leq D} := P^{\leq D} L_n$.

We now have a low-degree analog of Proposition 3.1.3, a simple but conceptually important statement which first appeared in [HS17, HKP⁺17].

Proposition 3.2.2. The unique solution f^* of the optimization problem

$$\begin{aligned} & \text{maximize} && \mathbb{E}_{\mathbf{Y} \sim \mathbb{P}_n} [f(\mathbf{Y})] \\ & \text{subject to} && \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} [f(\mathbf{Y})^2] = 1, \\ & && f \in \mathbb{R}[\mathbf{Y}]_{\leq D}, \end{aligned} \tag{3.7}$$

is the (normalized) LDLR $f^* = L_n^{\leq D} / \|L_n^{\leq D}\|$, and the value of the optimization problem is $\|L_n^{\leq D}\|$.

Proof. As in the proof of Proposition 3.1.3, we can restate the optimization problem as maximizing $\langle L_n, f \rangle$ subject to $\|f\| = 1$ and $f \in \mathbb{R}[\mathbf{Y}]_{\leq D}$. Since $\mathbb{R}[\mathbf{Y}]_{\leq D}$ is a linear subspace of $L^2(\mathbb{Q}_n)$, the result is then simply a restatement of the variational description and uniqueness of the orthogonal projection in $L^2(\mathbb{Q}_n)$. \square

The following informal conjecture is at the heart of the low-degree method. It states that a computationally-bounded analog of the second moment method for contiguity holds, with $L_n^{\leq D}$ playing the role of the likelihood ratio. Furthermore, it postulates that polynomials of

degree roughly $\log(n)$ are a proxy for polynomial-time algorithms. This conjecture is based on [HS17, HKP⁺17, Hop18], and particularly Conjecture 2.2.4 of [Hop18].

Conjecture 3.2.3 (Informal). *For “sufficiently nice” sequences of probability measures (\mathbb{P}_n) and (\mathbb{Q}_n) , if there exists $\epsilon > 0$ and $D = D(n) \geq (\log n)^{1+\epsilon}$ for which $\|L_n^{\leq D}\|$ remains bounded as $n \rightarrow \infty$, then (\mathbb{P}_n) and (\mathbb{Q}_n) cannot be distinguished in polynomial time (Definition 2.1.3).*

Hopkins further proposed a more speculative “hypothesis” (Hypothesis 2.1.5 in [Hop18]) that captures a broader range of applications to subexponential-time algorithms. One conservative prediction we may extract from this is that, if $\|L_n^{\leq (\log n)^K D(n)}\| = O(1)$ for any $K > 0$, then it is impossible to distinguish (\mathbb{P}_n) from (\mathbb{Q}_n) in time $\exp(D(n))$. We will refer to this in the sequel as “the extended Conjecture 3.2.3.”

We give some clarification about the details left vague above. Hopkins proposed that “sufficiently nice” above should mean that (1) \mathbb{Q}_n is a product distribution, (2) \mathbb{P}_n is invariant under permutations,³ and (3) starting from such a distribution, \mathbb{P}_n then undergoes some further “noising” operation. We will see in our results that we obtain sensible predictions even when (1) and (2) are relaxed, and these are likely too conservative. However, the condition (3) is actually quite important; it excludes, for example, the well-known example of planted solutions in k -XORSAT systems being detectable with Gaussian elimination—as these are linear systems in a finite field—but invisible to most analytical methods oblivious to this algebraic structure. Adding a small amount of noise disturbs this structure and makes identifying it (conjecturally) hard again. See, e.g., [IKKM12] for some discussion of this specific example. The problems we consider, however, will not have special structure of this kind, so we ignore this constraint as well. (The graph coloring problem discussed in Section 2.5 does have a version of this difficulty, which is handled in [BBK⁺20], but we will not return to that problem here.)

³It is important to allow for an “arity” in the data, i.e., for \mathbb{P}_n to be interpreted as a measure over matrices, tensors, graphs, multigraphs, etc., in which case the permutations in question permute the coordinates simultaneously and are not arbitrary permutations of the entire collection of observations.

Hopkins also suggested that “degree” should be interpreted as “coordinate degree,” meaning the number of coordinates that a function depends on, rather than polynomial degree. This, too, seems to be important sometimes, especially in models with distributions other than Gaussian or Bernoulli, as we will discuss with an example of a non-Gaussian spiked matrix model in Section 5.4.2. However, this example is rather peripheral for us and the other models we consider will be Gaussian or Bernoulli, in which case ignoring this distinction appears not to change the results obtained in any instances we are aware of.

All of this said, we suggest to the reader that, compared to the vast range of the statistical literature, only a vanishingly small number of quite similarly-structured testing problems have been studied to date with the low-degree method. In light of this, more valuable than a conjecture formulated safely might be a larger and better-understood collection of examples. In that spirit, the approach we take is to perform the low-degree computations for interesting problems even when they do not quite fit the stated conjectures, compare these results to other evidence of algorithmic hardness when such is available, and use these results to make predictions about hardness in new problems when such is possible.

3.2.2 APPEARANCE IN SUM-OF-SQUARES LOWER BOUNDS

As an alternative motivation for studying $\|L_n^{\leq D}\|$, and one more faithful to the original development of these ideas, we now briefly explain how this quantity naturally arises in the *pseudocalibration* framework for SOS lower bounds, as developed by [BHK⁺19]. We will be informal in our presentation, giving just the key ideas. In the setting we are interested in, the hypothetical difficulty of distinguishing \mathbb{P}_n from \mathbb{Q}_n is meant to guide the design of a *pseudoexpectation*, a linear operator $\tilde{\mathbb{E}} = \tilde{\mathbb{E}}_Y : \mathbb{R}[\mathbf{X}]_{\leq D} \rightarrow \mathbb{R}$, which when $\mathbf{Y} \sim \mathbb{Q}_n$ makes it appear as if \mathbf{Y} was drawn from \mathbb{P}_n and has the associated planted structure. We will give formal definitions in the specific setting of optimization over the hypercube in Chapter 6; this informal understanding will suffice for our discussion here.

Let us be explicit about the planted object involved in \mathbb{P}_n , so that we speak of sampling pairs $(\mathbf{X}^*, \mathbf{Y}) \sim \mathbb{P}_n$. For example, when \mathbb{P}_n is the Wigner spiked matrix model, this pair is $(\mathbf{x}, \lambda \mathbf{x} \mathbf{x}^\top + \mathbf{G})$ for \mathbf{G} having i.i.d. Gaussian entries and \mathbf{x} drawn from the spike prior. In this model, $\tilde{\mathbb{E}}_Y[\mathbf{x}^k]$ should give a plausible moment of a spike \mathbf{x} in \mathbf{Y} , if such a spike were there. Or, in the planted clique model that [BHK⁺19] studied and that we presented in Example 2.1.5, $\mathbf{Y} = (V, E)$ is a graph, and $\tilde{\mathbb{E}}_Y[\mathbf{x}^k]$ is a plausible moment of the indicator variables $\mathbf{x} \in \{0, 1\}^V$ of membership in a large clique in the graph, if such a clique were there.

To devise a way to compute $\tilde{\mathbb{E}}_Y$, we combine two ideas. First, it should be impossible to distinguish \mathbb{P}_n from \mathbb{Q}_n using the output of our construction. In particular, $\tilde{\mathbb{E}}_Y[p(\mathbf{X})]$ should be close to the same (at least in expectation) regardless of whether \mathbf{Y} is drawn from \mathbb{P}_n or \mathbb{Q}_n . And second, when we draw $(\mathbf{X}^*, \mathbf{Y}) \sim \mathbb{P}_n$, then $\tilde{\mathbb{E}}_Y[p(\mathbf{X})]$ should behave like the expectation of $p(\mathbf{X}^*)$ conditional on \mathbf{Y} . That is, when there actually is a planted object in \mathbf{Y} , then $\tilde{\mathbb{E}}_Y$ should just compute moments over what that object could be. Combining these observations yields the *pseudocalibration relations*:

$$\mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} \tilde{\mathbb{E}}_Y[p(\mathbf{X})] \stackrel{(1)}{\approx} \mathbb{E}_{(\mathbf{X}^*, \mathbf{Y}) \sim \mathbb{P}_n} \tilde{\mathbb{E}}_Y[p(\mathbf{X})] \stackrel{(2)}{\approx} \mathbb{E}_{(\mathbf{X}^*, \mathbf{Y}) \sim \mathbb{P}_n} p(\mathbf{X}^*), \quad (3.8)$$

where (1) denotes the application of the first idea and (2) of the second.

Noting that the right-hand side does not actually involve $\tilde{\mathbb{E}}_Y$, we see that this fixes the means of all of the values of $\tilde{\mathbb{E}}_Y$ under $\mathbf{Y} \sim \mathbb{Q}_n$. In fact, we can say much more, noting that the same argument holds verbatim if $p(\mathbf{X})$ also depends on \mathbf{Y} . Evaluating this enhanced version of the relation with a factorized polynomial $p(\mathbf{X})q(\mathbf{Y})$ gives the following:

$$\mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} q(\mathbf{Y}) \tilde{\mathbb{E}}_Y[p(\mathbf{X})] \approx \mathbb{E}_{(\mathbf{X}^*, \mathbf{Y}) \sim \mathbb{P}_n} p(\mathbf{X}^*)q(\mathbf{Y}). \quad (3.9)$$

Here we see the true power of the pseudocalibration relations: taking, for instance, $q(\mathbf{Y})$ to

be a system of orthogonal polynomials under \mathbb{Q}_n , we see that this gives the entire orthogonal polynomial decomposition of $\tilde{\mathbb{E}}_Y[p(\mathbf{X})]$, as a function of \mathbf{Y} for any $p(\mathbf{X})$. (To be even clearer on this point, we could write the left-hand side above as $\langle q(\bullet), \tilde{\mathbb{E}}_\bullet[p(\mathbf{X})] \rangle$ with the inner product in $L^2(\mathbb{Q}_n)$.)

One crucial difficulty appears, to see which it suffices to consider the simplest $p(\mathbf{X})$ to insert above, namely $p(\mathbf{X}) = 1$. Reasonably, any pseudoexpectation must satisfy $\tilde{\mathbb{E}}_Y 1 = 1$. Let us write $f(\mathbf{Y}) := \tilde{\mathbb{E}}_Y 1$. Also, suppose \hat{q}_k are a system of orthonormal polynomials for \mathbb{Q}_n . Then, the above says

$$\langle f, \hat{q}_k \rangle = \mathbb{E}_{Y \sim \mathbb{Q}_n} \hat{q}_k(\mathbf{Y}) \tilde{\mathbb{E}}_Y 1 \approx \mathbb{E}_{(\mathbf{X}^*, \mathbf{Y}) \sim \mathbb{P}_n} \hat{q}_k(\mathbf{Y}) = \langle L_n, \hat{q}_k \rangle, \quad (3.10)$$

where we have introduced the likelihood ratio $L_n = d\mathbb{P}_n/d\mathbb{Q}_n$. The above, taken literally, would mean $\tilde{\mathbb{E}}_Y 1 = f(\mathbf{Y}) = L_n(\mathbf{Y})$. This is also supposed to equal 1, but, when \mathbb{P}_n and \mathbb{Q}_n are (statistically) distinguishable, then, while indeed in expectation $\mathbb{E}_{Y \sim \mathbb{Q}_n} L_n(\mathbf{Y}) = 1$, L_n becomes increasingly poorly concentrated as $n \rightarrow \infty$. For example, in this situation $\text{Var}_{Y \sim \mathbb{Q}_n}[L_n(\mathbf{Y})] = \|L_n\|^2 - 1 \rightarrow \infty$, since otherwise by the second moment method (see the previous section) we would have contiguity $(\mathbb{P}_n) \triangleleft (\mathbb{Q}_n)$ and \mathbb{P}_n and \mathbb{Q}_n would not be distinguishable.

To cross this last impasse in the construction, [BHK⁺19] introduce a final adjustment of assuming $\tilde{\mathbb{E}}_Y$ is a low-degree function of \mathbf{Y} , and only asking that (3.9) (and therefore (3.10)) hold for $q(\mathbf{Y})$ that are low-degree. One plausible reason they propose for this is that if $\tilde{\mathbb{E}}_Y$ is the output of the SOS relaxation then it must, after all, be efficiently computable by a solver of the SOS relaxation; therefore, it should not have high-degree components making it prohibitively expensive to compute. That justification aside, this constraint is perhaps best seen as a technical device; subsequent works using pseudocalibration have found it necessary to adjust the ways that this truncation is applied to their convenience [HKP⁺17].

If we admit this heuristic, however, then the connection with the low-degree likelihood ratio is clear: $f(\mathbf{Y})$ above is the low-degree polynomial that agrees with L_n on inner products with low-degree orthogonal polynomials, and thus is precisely $L_n^{\leq D}$ if D is the degree threshold. Moreover, whether $\|L_n^{\leq D}\|$ is bounded or not as $n \rightarrow \infty$ is then precisely the question of whether our putative value of $\tilde{\mathbb{E}}_{\mathbf{Y}} 1$ has bounded variance under $\mathbf{Y} \sim \mathbb{Q}_n$, making it a plausible proxy for whether we expect pseudocalibration to “work.” The technical details of the analysis in [BHK⁺19] and subsequent works have since confirmed that this is an accurate heuristic.

3.3 CONSEQUENCES OF LOW-DEGREE LOWER BOUNDS

Finally, since Conjecture 3.2.3 remains novel and somewhat provisional, to shore up our confidence in these methods it is helpful to have some more concrete consequences of low-degree lower bounds. In particular, low-degree lower bounds are supposed to be lower bounds against algorithms computing low-degree polynomials, but they only ensure certain properties relating to means and variances of those polynomials. Can we use such bounds to deduce something about whether more concrete algorithms actually succeed or fail in distinguishing \mathbb{P}_n from \mathbb{Q}_n ?

We present two results in this direction, both due to Wein and appearing in [KWB19]. The first concerns thresholded polynomials, in the style of the thresholded likelihood ratio tests featuring in the Neyman-Pearson Lemma, our Proposition 3.1.4.

Proposition 3.3.1. *Suppose \mathbb{Q} is a product distribution of N copies of either $\mathcal{N}(0, 1)$ or $\text{Unif}(\{\pm 1\}^n)$, and \mathbb{P} is absolutely continuous with respect to \mathbb{Q} . Suppose there exists $f \in \mathbb{R}[x_1, \dots, x_N]_{\leq d}$ and $0 < a < A$ such that $\mathbb{E}_{\mathbf{Y} \sim \mathbb{P}} f(\mathbf{Y}) \geq A$ and $\mathbb{Q}[|f(\mathbf{Y})| \leq a] \geq 1 - \delta$. Let*

$0 \leq k \leq \frac{1}{4d} \log_3(2/\delta)$ be an integer. Then,

$$\|L^{\leq 2kd}\| \geq \frac{1}{2} \left(\frac{A}{a}\right)^{2k}. \quad (3.11)$$

The substance of this result in our framework is that, if when this setup is parametrized by n it is possible to take $k = k(n) \rightarrow \infty$, then $\|L^{\leq 2k(n)d(n)}\| \rightarrow \infty$. What that requires is for $\delta = \delta(n)$ to decay superexponentially in $d = d(n)$. If we take $d = (\log n)^{1+\epsilon}$ as for analysis of polynomial-time algorithms, and $k = \omega(1)$, then we will need $\delta \leq n^{-\omega(1)}$; i.e., super-polynomial decay of δ . Thus, conversely, if $\|L_n^{\leq D}\|$ is bounded, then there is no sequence of polynomials $f(\mathbf{Y})$ achieving this quantitative notion of distinguishing \mathbb{P}_n from \mathbb{Q}_n . The decay condition on δ is a plausible one to impose, if an unusual one for the literature. Unfortunately, it remains unclear how to prove such results for softer notions of polynomial thresholding satisfying bounds only with high probability.

The second result, quite similar to the first, concerns *spectral algorithms* that compute the norm of a matrix built from the observations.

Proposition 3.3.2. *Suppose that the hypotheses of Proposition 3.3.2 are fulfilled for $f(\mathbf{Y})$ not a polynomial, but $f(\mathbf{Y}) = \|\mathbf{F}(\mathbf{y})\|$, where $\mathbf{F} : \mathbb{R}^N \rightarrow \mathbb{R}^{M \times M}$ such that each entry of \mathbf{F} is a polynomial of degree at most d . Then,*

$$\|L^{\leq 2kd}\| \geq \frac{1}{2M} \left(\frac{A}{a}\right)^{2k}. \quad (3.12)$$

Both results are proved using *hypercontractive* concentration inequalities, which ensure concentration of polynomials of inputs having favorable distributions. We will use these kinds of inequalities in some of our other results, and present a related one in Proposition 11.3.2. It is the details of these inequalities that give rise to the particular dependence on δ in the above results; it is an interesting open problem to investigate whether other

tools can give more general results.

Remark 3.3.3 (Towards average-case equivalences). *A recent result also suggests an equivalence (suitably interpreted) between low-degree polynomial tests and algorithms in the statistical query framework [BBH⁺ 20]; further equivalences of this kind are an interesting future direction. Moreover, the historical motivation of low-degree algorithms described above in Section 3.2.2 relates them to SOS in general, while [HKP⁺ 17] have related SOS to spectral algorithms in general. There is therefore an incipient hope that some equivalences may be drawn among the powers of all of these classes of algorithms; unfortunately, all results so far have various limitations and technical subtleties that, taken together, amount to such equivalences holding “morally” but not strongly enough to show that any one class of algorithms dominates the performance of another for many problems of interest. Nonetheless, we pose another open problem to this effect, suggested by the results of the next chapter, in Section A.2.*

4 | LOW-DEGREE LIKELIHOOD RATIO

ANALYSIS: OVERLAP FORMULAE

We are now in principle equipped to prove low-degree lower bounds for the situations we are interested in by following the strategy we have encountered tangentially in Section 3.2.2. That is, we wish to bound the quantity $\|L_n^{\leq D}\|$, and may do so directly by expanding L_n in orthogonal polynomials and bounding the sums of contributions to the norm made by orthogonal polynomials with degree at most D . This approach was applied successfully to several problems in early works proving low-degree lower bounds (see, e.g., Hopkins' thesis [Hop18] for several examples), and continues to be useful; we will carry out a calculation like this in Section 5.4.2.

However, in this chapter we will digress to observe, in greater generality than we will need for our immediate goals, a useful and independently intriguing phenomenon arising in these calculations which is that, often, it is possible to prove a formula or bound of the following form:

$$\|L_n^{\leq D}\|^2 \leq \mathbb{E}[p_{n,D}(R_n)]. \quad (4.1)$$

Here $p_{n,D}$ are suitable polynomials and R_n is a *scalar* random variable. We contrast this with the immediate output of the aforementioned orthogonal polynomial approach, which typically involves moments of high-dimensional random variables (as we will see, in the Wigner and Wishart spiked matrix models and related settings, these will be moments of

Model	Parameters	Overlap	Link Function	Reference
Wigner Spiked Matrix	λ	$\langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2$	$\exp(\lambda^2 t)$	[KWB19]
General Gaussian Wigner	(None)	$\langle \tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \rangle$	$\exp(t)$	[KWB19]
Morris Class / NEF-QVF	ν_2	$\langle z(\tilde{\mathbf{x}}^1), z(\tilde{\mathbf{x}}^2) \rangle$	$(1 - \nu_2 t)^{-1/\nu_2}$	[Kun20a]
Wishart Spiked Matrix	β, N	$\langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2$	$(1 - \beta^2 t)^{-N/2}$	[BKW20b]
General Gaussian Wishart	N	$\tilde{\mathbf{X}}^1 \tilde{\mathbf{X}}^2$	$\det(\mathbf{I}_n - \mathbf{T})^{-N/2}$	[BBK ⁺ 20]

Table 4.1: Summary of overlap formulae. We summarize the results of Chapter 4 on overlap formulae for the norm of the LDLR in different models. For NEF-QVF models ν_2 is a coefficient in the variance function of the exponential family, and $z(\tilde{\mathbf{x}})$ is a function computing z -scores with respect to the null model. For Wishart models N is the number of samples and β is the signal strength for the spiked matrix model. For the Wigner spiked matrix model λ is the signal-to-noise ratio. In overlap expressions, superscripts indicate independent copies, $\tilde{\mathbf{X}}^i$ or $\tilde{\mathbf{x}}^i$ indicate “signals” that can be arbitrary vectors or matrices, while \mathbf{x}^i indicate “spikes” that appear as factors of positive semidefinite matrices $\mathbf{x}^i \mathbf{x}^{i\top}$. In all cases, the norm of the LDLR is given by the expectation of a truncated Taylor series of the link function evaluated on the overlap.

the spike matrix \mathbf{X} for $\mathbf{X} \sim \mathcal{P}_n$). We will see that R_n gives an *overlap* between two independent draws of a signal, an inner product after some normalizations that vary from setting to setting; in the rank-one Wishart spiked matrix model (Definition 2.2.4), it is just $R_n = \langle \beta \mathbf{x}^1 \mathbf{x}^{1\top}, \beta \mathbf{x}^2 \mathbf{x}^{2\top} \rangle = \beta^2 \langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2$ for $\mathbf{x}^i \sim \mathcal{P}_n$ independently. Moreover, $p_{n,D}$ gives the Taylor series truncation to order D of a *link function*, whose full untruncated form is the function appearing in an analogous expression for $\|L_n\|^2$, the norm of the untruncated likelihood ratio. Recognizing such identities and bounds often gives a much-simplified and more conceptual analysis of the norm of the LDLR.

SUMMARY AND REFERENCES The overlap formula for simple Gaussian models (including the Wigner spiked matrix model) was first observed in [KWB19]. The discussion of analogous results in other scalar-valued exponential families is taken from [Kun20a]. The treatment of Wishart models given here is a streamlined derivation of the relevant results of both [BKW20b] (which treated the rank-one spike case) and [BBK⁺20] (which treated the higher-rank case); we give a somewhat cleaner derivation and make some new observations about

the Taylor series involved that substantially simplifies the analysis we take up in the next chapter. We summarize these results in advance in Table 4.1.

4.1 GAUSSIAN WIGNER MODELS

We will build up gradually from the simplest models to more complicated ones. Thus we begin with the following model, assuming many convenient properties: Gaussianity, noise applied additively, and entrywise independence of observations conditional on the signal.

Definition 4.1.1 (General Gaussian Wigner model). *Let $N = N(n) \in \mathbb{N}_+$ and let $\tilde{\mathbf{X}}$ (the “signal”) be drawn from some distribution $\tilde{\mathcal{P}}_n$ (the “signal prior”) over \mathbb{R}^N . Let $\mathbf{G} \in \mathbb{R}^N$ (the “noise”) have i.i.d. entries distributed as $\mathcal{N}(0, 1)$. Then, we define \mathbb{P}_n and \mathbb{Q}_n as follows.*

- Under \mathbb{Q}_n , observe $\mathbf{Y} = \mathbf{G}$.
- Under \mathbb{P}_n , observe $\mathbf{Y} = \tilde{\mathbf{X}} + \mathbf{G}$.

One typical case takes $\tilde{\mathbf{X}}$ to be a low-rank matrix or tensor. (This case is so typical that we reserve the notations \mathbf{X} or \mathbf{x} for the factors of $\tilde{\mathbf{X}}$, as below.) The following is a particularly important and well-studied example, analogous to though often simpler to work with than the Wishart spiked matrix model (Definition 2.2.4). In fact, in Section 5.2.1 we will see that the LDLR norm in the Wishart spiked matrix model can be bounded by that in this model, making the analyses nearly identical.

Example 4.1.2 (Wigner spiked matrix model). *Consider the Gaussian Wigner model with $N = n^2$, \mathbb{R}^N identified with $\mathbb{R}^{n \times n}$, and $\tilde{\mathcal{P}}_n$ defined by sampling $\tilde{\mathbf{X}} = \frac{\lambda}{\sqrt{2}} \mathbf{x} \mathbf{x}^\top \in \mathbb{R}^{n \times n}$, where $\lambda = \lambda(n) > 0$ is a signal-to-noise parameter and $\mathbf{x} \sim \mathcal{P}_n$ for some $\mathcal{P}_n \in \mathcal{P}(\mathbb{R}^n)$. Then, the task of distinguishing \mathbb{P}_n from \mathbb{Q}_n amounts to distinguishing the laws of $\lambda \mathbf{x} \mathbf{x}^\top + \mathbf{G}$ and \mathbf{G} where $\mathbf{G} \in \mathbb{R}^{n \times n}$ has i.i.d. entries distributed as $\mathcal{N}(0, 1)$. (This variant is equivalent for all relevant*

purposes to the more standard model in which the noise matrix is symmetric and distributed as $\sqrt{n} \cdot \text{GOE}(n)$; the reason for including the factor of $\sqrt{2}$ is to match the scaling in the typical definition of $\text{GOE}(n)$, as we will see in Section 5.2.4.)

We first show how to compute the likelihood ratio and, as a warmup, its L^2 norm without truncation to low-degree polynomials, under a Gaussian Wigner model. This is a standard calculation; see, e.g., [MRZ15, BMV⁺18, PWBM18, PWB16], where it is used together with the second moment method for contiguity to establish statistical impossibility for various testing problems.

Proposition 4.1.3 (Likelihood ratio for Gaussian Wigner model). *Suppose (\mathbb{P}_n) and (\mathbb{Q}_n) are a Gaussian Wigner model as in Definition 4.1.1, with signal prior $(\tilde{\mathcal{P}}_n)$. Then, the likelihood ratio of \mathbb{P}_n and \mathbb{Q}_n is*

$$L_n(\mathbf{Y}) = \frac{d\mathbb{P}_n}{d\mathbb{Q}_n}(\mathbf{Y}) = \mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \left[\exp \left(-\frac{1}{2} \|\tilde{\mathbf{X}}\|^2 + \langle \tilde{\mathbf{X}}, \mathbf{Y} \rangle \right) \right].$$

Proof. Write \mathcal{L} for the Lebesgue measure on \mathbb{R}^N . Then, expanding the gaussian densities,

$$\frac{d\mathbb{Q}_n}{d\mathcal{L}}(\mathbf{Y}) = (2\pi)^{-N/2} \cdot \exp \left(-\frac{1}{2} \|\mathbf{Y}\|^2 \right) \quad (4.2)$$

$$\begin{aligned} \frac{d\mathbb{P}_n}{d\mathcal{L}}(\mathbf{Y}) &= (2\pi)^{-N/2} \cdot \mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \left[\exp \left(-\frac{1}{2} \|\mathbf{Y} - \tilde{\mathbf{X}}\|^2 \right) \right] \\ &= (2\pi)^{-N/2} \cdot \exp \left(-\frac{1}{2} \|\mathbf{Y}\|^2 \right) \cdot \mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \left[\exp \left(-\frac{1}{2} \|\tilde{\mathbf{X}}\|^2 + \langle \tilde{\mathbf{X}}, \mathbf{Y} \rangle \right) \right], \end{aligned} \quad (4.3)$$

and L_n is given by the quotient of (4.3) and (4.2). \square

Before proceeding to compute the norm, we isolate the following simple but crucial formal idea, that we will use repeatedly in our calculations in this chapter.

Proposition 4.1.4 (Replica manipulation). *For any random variable X with finite k th moment,*

if X^1, \dots, X^k are independent copies of X , then

$$(\mathbb{E}X)^k = \mathbb{E}X^1 \dots X^k. \quad (4.4)$$

The proof is immediate from the definition of independence. Yet, this idea is very useful as a means of “linearizing” powers of expectations into single expectations over independent copies. Perhaps most notably, it is an important ingredient (though not the entirety, and not the notoriously unrigorous part) of the “replica trick” of statistical physics; the interested reader may consult the references [MPV87, MM09, BPW18].

Proposition 4.1.5 (LR norm for Gaussian Wigner model). *Suppose (\mathbb{P}_n) and (\mathbb{Q}_n) are a Gaussian Wigner model as in Definition 4.1.1, with signal prior $(\tilde{\mathcal{P}}_n)$. Then,*

$$\|L_n\|^2 = \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \sim \tilde{\mathcal{P}}_n} \exp(\langle \tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \rangle), \quad (4.5)$$

where $\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2$ are drawn independently from $\tilde{\mathcal{P}}_n$.

Proof. We apply Proposition 4.1.4:

$$\begin{aligned} \|L_n\|^2 &= \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} \left[\left(\mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \exp \left(\langle \mathbf{Y}, \tilde{\mathbf{X}} \rangle - \frac{1}{2} \|\tilde{\mathbf{X}}\|^2 \right) \right)^2 \right] \\ &= \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} \left[\mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \sim \tilde{\mathcal{P}}_n} \exp \left(\langle \mathbf{Y}, \tilde{\mathbf{X}}^1 + \tilde{\mathbf{X}}^2 \rangle - \frac{1}{2} \|\tilde{\mathbf{X}}^1\|^2 - \frac{1}{2} \|\tilde{\mathbf{X}}^2\|^2 \right) \right], \end{aligned}$$

where $\tilde{\mathbf{X}}^1$ and $\tilde{\mathbf{X}}^2$ are drawn independently from $\tilde{\mathcal{P}}_n$. We now swap the order of the expectations,

$$= \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \sim \tilde{\mathcal{P}}_n} \left[\exp \left(-\frac{1}{2} \|\tilde{\mathbf{X}}^1\|^2 - \frac{1}{2} \|\tilde{\mathbf{X}}^2\|^2 \right) \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} \exp \left(\langle \mathbf{Y}, \tilde{\mathbf{X}}^1 + \tilde{\mathbf{X}}^2 \rangle \right) \right],$$

and the inner expectation may be evaluated explicitly using the moment-generating function of a Gaussian distribution (if $y \sim \mathcal{N}(0, 1)$, then for any fixed $t \in \mathbb{R}$, $\mathbb{E}[\exp(ty)] = \exp(t^2/2)$),

$$= \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2} \exp\left(-\frac{1}{2}\|\tilde{\mathbf{X}}^1\|^2 - \frac{1}{2}\|\tilde{\mathbf{X}}^2\|^2 + \frac{1}{2}\|\tilde{\mathbf{X}}^1 + \tilde{\mathbf{X}}^2\|^2\right),$$

from which the result follows by expanding the term inside the exponential. \square

Next, we will show that the norm of the LDLR takes the following remarkable related form under a Gaussian Wigner model. The truncation to low-degree polynomials of L_n , a truncation in a high-dimensional space of polynomials, in fact has the simple effect on the norm of truncating the univariate Taylor series of the exponential function.

Definition 4.1.6 (Link functions for Gaussian Wigner model). *For $D \in \mathbb{N}$, we define the functions*

$$\phi_D^{\text{Wig}}(t) := \sum_{d=0}^D \frac{1}{d!} t^d, \quad (4.6)$$

$$\phi^{\text{Wig}}(t) = \phi_\infty^{\text{Wig}}(t) := \lim_{D \rightarrow \infty} \phi_D^{\text{Wig}}(t) = \exp(t). \quad (4.7)$$

Theorem 4.1.7 (LDLR norm for Gaussian Wigner model). *Suppose (\mathbb{P}_n) and (\mathbb{Q}_n) are a Gaussian Wigner model as in Definition 4.1.1, with signal prior $(\tilde{\mathcal{P}}_n)$. Let $L_n^{\leq D}$ be as in Definition 3.2.1. Then, for $D \in \mathbb{N} \cup \{+\infty\}$,*

$$\|L_n^{\leq D}\|^2 = \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \sim \tilde{\mathcal{P}}_n} \left[\phi_D^{\text{Wig}}(\langle \tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \rangle) \right], \quad (4.8)$$

where $\mathbf{X}^1, \mathbf{X}^2$ are drawn independently from \mathcal{P}_n .

In such an overlap formula, we will call ϕ^{Wig} the *link function* and ϕ_D^{Wig} the *link polynomials*, as these functions deform the overlap before the expectation is computed in a manner formally similar to the link function in a generalized linear model. Our proof of Theorem 4.1.7

will follow the strategy of [HS17, HKP⁺17, Hop18] of expanding L_n in a basis of orthogonal polynomials with respect to \mathbb{Q}_n , which in this case are the *Hermite polynomials*.

4.1.1 HERMITE POLYNOMIALS

We first review the essential features of Hermite polynomials that we will use (the reader may consult the standard reference [Sze39] for further information and omitted proofs).

Definition 4.1.8. *The univariate Hermite polynomials are the polynomials $h_k(x) \in \mathbb{R}[x]_k$ for $k \geq 0$ defined by the recursion*

$$h_0(x) = 1, \tag{4.9}$$

$$h_{k+1}(x) = xh_k(x) - h'_k(x). \tag{4.10}$$

The normalized univariate Hermite polynomials are $\hat{h}_k(x) = h_k(x)/\sqrt{k!}$.

The following is the key property of the Hermite polynomials, which allows functions in $L^2(\mathcal{N}(0, 1))$ to be expanded in terms of them.

Proposition 4.1.9. *The normalized univariate Hermite polynomials are a complete orthonormal system of polynomials for $L^2(\mathcal{N}(0, 1))$. In particular,*

$$\mathbb{E}_{\mathcal{Y} \sim \mathcal{N}(0,1)} \hat{h}_k(\mathcal{Y})\hat{h}_\ell(\mathcal{Y}) = \delta_{k\ell}. \tag{4.11}$$

The following are the multivariate product basis formed from the Hermite polynomials.

Definition 4.1.10. *The N -variate Hermite polynomials are $H_{\mathbf{k}}(\mathbf{X}) := \prod_{i=1}^N h_{k_i}(X_i)$ indexed by tuples $\mathbf{k} \in \mathbb{N}^N$. The normalized N -variate Hermite polynomials are $\hat{H}_{\mathbf{k}}(\mathbf{X}) := \prod_{i=1}^N \hat{h}_{k_i}(X_i) = (\prod_{i=1}^N k_i!)^{-1/2} \prod_{i=1}^N h_{k_i}(X_i)$ for $\mathbf{k} \in \mathbb{N}^N$.*

Again, the following is the key property justifying expansions in terms of these polynomials.

Proposition 4.1.11. *The normalized N -variate Hermite polynomials are a complete orthonormal system of (multivariate) polynomials for $L^2(\mathcal{N}(\mathbf{0}, \mathbf{I}_N))$. In particular,*

$$\mathbb{E}_{\mathbf{Y} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_N)} \hat{H}_{\mathbf{k}}(\mathbf{Y}) \hat{H}_{\ell}(\mathbf{Y}) = \delta_{\mathbf{k}\ell}. \quad (4.12)$$

As a straightforward corollary, the collection of those $\hat{H}_{\mathbf{k}}$ for which $|\mathbf{k}| := \sum_{i=1}^N k_i \leq D$ form an orthonormal basis for $\mathbb{R}[Y_1, \dots, Y_N]_{\leq D}$.

We also introduce a few algebraic identities satisfied by the Hermite polynomials, which will give various proofs of the key identity we show below.

Proposition 4.1.12 (Translation). *For any $k \geq 0$ and $\mu \in \mathbb{R}$,*

$$\mathbb{E}_{\mathcal{Y} \sim \mathcal{N}(\mu, 1)} [h_k(\mathcal{Y})] = \mu^k. \quad (4.13)$$

Proposition 4.1.13 (Gaussian integration by parts). *If $f : \mathbb{R} \rightarrow \mathbb{R}$ is k times continuously differentiable and $f(\mathcal{Y})$ and its first k derivatives are bounded by $O(\exp(|\mathcal{Y}|^\alpha))$ for some $\alpha \in (0, 2)$, then*

$$\mathbb{E}_{\mathcal{Y} \sim \mathcal{N}(0, 1)} [h_k(\mathcal{Y}) f(\mathcal{Y})] = \mathbb{E}_{\mathcal{Y} \sim \mathcal{N}(0, 1)} \left[\frac{d^k f}{d\mathcal{Y}^k}(\mathcal{Y}) \right]. \quad (4.14)$$

(The better-known special case that often goes by the name of “Gaussian integration by parts” is $k = 1$, where one may substitute $h_1(x) = x$.)

Proposition 4.1.14 (Generating function). *For any $x, y \in \mathbb{R}$,*

$$\exp\left(xy - \frac{1}{2}x^2\right) = \sum_{k=0}^{\infty} \frac{1}{k!} x^k h_k(y). \quad (4.15)$$

4.1.2 COMPUTING THE LOW-DEGREE LIKELIHOOD RATIO

We now proceed to a proof of Theorem 4.1.7. First, we may expand

$$L_n^{\leq D}(\mathbf{Y}) = \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \langle L_n, \hat{H}_{\mathbf{k}} \rangle \hat{H}_{\mathbf{k}}(\mathbf{Y}) = \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \frac{1}{\prod_{i=1}^N k_i!} \langle L_n, H_{\mathbf{k}} \rangle H_{\mathbf{k}}(\mathbf{Y}), \quad (4.16)$$

and in particular we have

$$\|L_n^{\leq D}\|^2 = \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \frac{1}{\prod_{i=1}^N k_i!} \langle L_n, H_{\mathbf{k}} \rangle^2. \quad (4.17)$$

Our main task is then to compute quantities of the form $\langle L_n, H_{\mathbf{k}} \rangle$. Note that, using the likelihood ratio for a change of measure, these can be expressed equivalently as

$$\langle L_n, H_{\mathbf{k}} \rangle = \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} [L_n(\mathbf{Y}) H_{\mathbf{k}}(\mathbf{Y})] = \mathbb{E}_{\mathbf{Y} \sim \mathbb{P}_n} [H_{\mathbf{k}}(\mathbf{Y})]. \quad (4.18)$$

We will give three techniques for carrying out this calculation, each depending on a different algebraic identity satisfied by the Hermite polynomials. Each will give a proof of the following remarkable formula, which shows that the quantities $\langle L_n, H_{\mathbf{k}} \rangle$ are simply the moments of $\tilde{\mathcal{P}}_n$.

Proposition 4.1.15 (LDLR components for Gaussian Wigner model). *For any $\mathbf{k} \in \mathbb{N}^N$,*

$$\langle L_n, H_{\mathbf{k}} \rangle = \mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \left[\prod_{i=1}^N \tilde{X}_i^{k_i} \right].$$

Before continuing with the various proofs of Proposition 4.1.15, let us show how to use the Proposition to complete the proof of Theorem 4.1.7.

Proof of Theorem 4.1.7. By Proposition 4.1.15 substituted into (4.17), we have

$$\|L_n^{\leq D}\|^2 = \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \frac{1}{\prod_{i=1}^N k_i!} \left(\mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \left[\prod_{i=1}^N \tilde{X}_i^{k_i} \right] \right)^2,$$

and applying Proposition 4.1.4, this may be written

$$\begin{aligned} &= \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \sim \tilde{\mathcal{P}}_n} \left[\sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \frac{1}{\prod_{i=1}^N k_i!} \prod_{i=1}^N (\tilde{X}_i^1 \tilde{X}_i^2)^{k_i} \right] \\ &= \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \sim \tilde{\mathcal{P}}_n} \left[\sum_{d=0}^D \frac{1}{d!} \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}|=d}} \binom{d}{k_1, \dots, k_N} \prod_{i=1}^N (\tilde{X}_i^1 \tilde{X}_i^2)^{k_i} \right] \\ &= \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \sim \tilde{\mathcal{P}}_n} \left[\sum_{d=0}^D \frac{1}{d!} \langle \tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \rangle^d \right], \end{aligned} \tag{4.19}$$

where the last step uses the multinomial theorem. \square

We now proceed to the three proofs of Proposition 4.1.15. The three proofs respectively use Propositions 4.1.12, 4.1.13, and 4.1.14, the three identities concerning Hermite polynomials we introduced above.

Proof 1 of Proposition 4.1.15. We rewrite $\langle L_n, H_k \rangle$ as an expectation with respect to \mathbb{P}_n :

$$\begin{aligned} \langle L_n, H_k \rangle &= \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} [L_n(\mathbf{Y}) H_k(\mathbf{Y})] \\ &= \mathbb{E}_{\mathbf{Y} \sim \mathbb{P}_n} [H_k(\mathbf{Y})] \\ &= \mathbb{E}_{\mathbf{Y} \sim \mathbb{P}_n} \left[\prod_{i=1}^N h_{k_i}(Y_i) \right] \end{aligned}$$

and recall $\mathbf{Y} = \mathbf{X} + \mathbf{G}$ for $\mathbf{X} \sim \mathcal{P}_n$ and $\mathbf{G} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_N)$ under \mathbb{P}_n ,

$$\begin{aligned}
&= \mathbb{E}_{\mathbf{X} \sim \mathcal{P}_n} \left[\mathbb{E}_{\mathbf{G} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_N)} \prod_{i=1}^N h_{k_i}(X_i + G_i) \right] \\
&= \mathbb{E}_{\mathbf{X} \sim \mathcal{P}_n} \left[\prod_{i=1}^N \mathbb{E}_{z \sim \mathcal{N}(X_i, 1)} h_{k_i}(z) \right] \\
&= \mathbb{E}_{\mathbf{X} \sim \mathcal{P}_n} \left[\prod_{i=1}^N X_i^{k_i} \right],
\end{aligned}$$

where we used Proposition 4.1.12 in the last step. □

Proof 2 of Proposition 4.1.15. We simplify using Proposition 4.1.13:

$$\langle L_n, H_{\mathbf{k}} \rangle = \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} \left[L_n(\mathbf{Y}) \prod_{i=1}^N h_{k_i}(Y_i) \right] = \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} \left[\frac{\partial^{|\mathbf{k}|} L_n}{\partial Y_1^{k_1} \dots \partial Y_N^{k_N}}(\mathbf{Y}) \right].$$

Differentiating L_n under the expectation, we have

$$\frac{\partial^{|\mathbf{k}|} L}{\partial Y_1^{k_1} \dots \partial Y_N^{k_N}}(\mathbf{Y}) = \mathbb{E}_{\mathbf{X} \sim \mathcal{P}_n} \left[\prod_{i=1}^N X_i^{k_i} \exp\left(-\frac{1}{2}\|\mathbf{X}\|^2 + \langle \mathbf{X}, \mathbf{Y} \rangle\right) \right].$$

Taking the expectation over \mathbf{Y} , we have $\mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} \exp(\langle \mathbf{X}, \mathbf{Y} \rangle) = \exp(\frac{1}{2}\|\mathbf{X}\|^2)$, so the entire second term cancels and the result follows. □

Proof 3 of Proposition 4.1.15. We may use Proposition 4.1.14 to expand L_n in the Hermite polynomials directly:

$$\begin{aligned}
L_n(\mathbf{Y}) &= \mathbb{E}_{\mathbf{X} \sim \mathcal{P}_n} \left[\exp\left(\langle \mathbf{X}, \mathbf{Y} \rangle - \frac{1}{2}\|\mathbf{X}\|^2\right) \right] \\
&= \mathbb{E}_{\mathbf{X} \sim \mathcal{P}_n} \left[\prod_{i=1}^N \left(\sum_{k=0}^{\infty} \frac{1}{k!} X_i^k h_k(Y_i) \right) \right] \\
&= \sum_{\mathbf{k} \in \mathbb{N}^N} \frac{1}{\prod_{i=1}^N k_i!} \mathbb{E}_{\mathbf{X} \sim \mathcal{P}_n} \left[\prod_{i=1}^N X_i^{k_i} \right] H_{\mathbf{k}}(\mathbf{Y}).
\end{aligned}$$

Comparing with the expansion (4.16) then gives the result. □

4.2 EXPONENTIAL FAMILIES AND VARIANCE FUNCTIONS

We have seen above the first instance of an overlap formula, an expression for $\|L_n^{\leq D}\|^2$ as an expectation of a function of $\langle \tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \rangle$ for $\tilde{\mathbf{X}}^i$ independent draws from $\tilde{\mathcal{P}}_n$, in the Gaussian Wigner model. As mentioned earlier, besides their aesthetic appeal, these kinds of identities will yield substantial simplifications in our later probabilistic analysis by allowing us to appeal to concentration inequalities for the overlaps. We now ask: does this simplifying phenomenon occur in other, more complicated observation models?

We recall that, ultimately, we will want to address this matter for the Wishart negatively-spiked matrix model, where the signal is applied by deforming the covariance rather than the mean of a Gaussian distribution.

Remark 4.2.1. *Interestingly, the Wishart positively-spiked matrix model may be viewed as a Gaussian Wigner model with a suitable Gaussian spike prior, since if $\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n + \tilde{\mathbf{X}})$ with $\tilde{\mathbf{X}} \succeq \mathbf{0}$, then we may view sampling \mathbf{y} as drawing $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \tilde{\mathbf{X}})$ and then $\mathbf{y} \sim \mathcal{N}(\mathbf{x}, \mathbf{I}_n)$.*

Perhaps the key difference between the Wigner and Wishart models is that, conditional on the signal, the observations in the Wigner model are entrywise independent, while those in the Wishart model are not. Nonetheless, there is another analogy suggesting a similarity between the models: conditional on the signal, both make observations in an *exponential family* (though the Wishart model must be adjusted to observe the sample covariance $\sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^\top$ instead of the samples $\mathbf{y}_1, \dots, \mathbf{y}_N$ to make this the case). The Gaussian exponential family with fixed variance and varying means—the “Wigner family” arising in the Gaussian Wigner model, to use our terminology—is, in a sense we will make precise below, perhaps the simplest of all exponential families. The matrix-valued Wishart exponential family is surely

more complicated, but still, in the same sense, is “not too complicated” and is formally similar to certain other scalar-valued non-Gaussian exponential families.

Below we develop this analogy, measuring the complexity of exponential families by their *variance functions*. We then study the low-degree likelihood ratio in the simplest exponential families, those whose variance function is a constant, linear, or quadratic polynomial. First we consider non-Gaussian scalar-valued exponential families—we will see later that these calculations have interesting applications of their own—and then proceed to the Wishart family. We will show that, in all of these situations, a suitable analog of the overlap formula from the Gaussian Wigner model holds.

Throughout this section and the next we follow Morris’ presentation in the seminal papers [Mor82, Mor83], which first recognized the many shared statistical properties of the scalar-valued exponential families we consider (though Meixner in [Mei34] already saw their similarity from the perspective of orthogonal polynomials). We start by recalling the basic notions of exponential families.

Definition 4.2.2. *Let $\nu_0 \in \mathcal{P}(\mathbb{R})$ not be a single atom. Let $\psi(\theta) := \log \mathbb{E}_{x \sim \nu_0}[\exp(\theta x)]$ and $\Theta := \{\theta \in \mathbb{R} : \psi(\theta) < \infty\}$. Then, the natural exponential family (NEF) generated by ν_0 is the family of probability measures ν_θ , for $\theta \in \Theta$, given by*

$$d\nu_\theta(x) := \exp(\theta x - \psi(\theta))d\nu_0(x). \tag{4.20}$$

Sometimes, the “natural parameter” θ is the mean of ν_θ or a translation thereof; however, as the next example shows, the mapping $\theta \mapsto \mathbb{E}_{x \sim \nu_\theta}[x]$ can be more complicated.

Example 4.2.3. *Taking $d\nu_0(x) = e^{-x}\mathbb{1}\{x \geq 0\}dx$, we have $\Theta = (-\infty, 1)$, and this generates the NEF of exponential distributions, $d\nu_\theta(x) = (1 - \theta)e^{-(1-\theta)x}\mathbb{1}\{x \geq 0\}dx$. The mean of ν_θ is $\mathbb{E}_{x \sim \nu_\theta}[x] = \frac{1}{1-\theta}$.*

Nonetheless, it is always possible to reparametrize any NEF in terms of the mean in the

following way. The cumulant generating functions of the ν_θ are merely translations of ψ , $\psi_\theta(\eta) := \log \mathbb{E}_{x \sim \nu_\theta}[\exp(\eta x)] = \psi(\theta + \eta) - \psi(\theta)$. Therefore, the means and variances of ν_θ are

$$\mu_\theta := \mathbb{E}_{x \sim \nu_\theta}[x] = \psi'_\theta(0) = \psi'(\theta), \quad (4.21)$$

$$\sigma_\theta^2 := \mathbb{E}_{x \sim \nu_\theta}[x^2] - (\mathbb{E}_{x \sim \nu_\theta}[x])^2 = \psi''_\theta(0) = \psi''(\theta). \quad (4.22)$$

Since ν_0 is not an atom, neither is any ν_θ , and thus $\psi''(\theta) = \sigma_\theta^2 > 0$ for all $\theta \in \mathbb{R}$. Therefore, ψ' is strictly increasing, and thus one-to-one. Letting $\Omega \subseteq \mathbb{R}$ equal the image of \mathbb{R} under ψ' (some open interval, possibly infinite on either side, of \mathbb{R}), we see that ν_θ admits an alternative mean parametrization, as follows.

Definition 4.2.4. *If ν_0 generates the NEF ν_θ , then we let $\rho_\mu = \nu_{(\psi')^{-1}(\mu)}$ over $\mu \in \Omega$. The mean-parametrized NEF generated by ν_0 is the family of probability measures ρ_μ , for $\mu \in \Omega$.*

By the same token, within an NEF, the variance is a function of the mean. In the above setting, we denote this function as follows.

Definition 4.2.5. *For $\mu \in \Omega$, define the variance function $V(\mu) := \sigma_{(\psi')^{-1}(\mu)}^2 = \psi''((\psi')^{-1}(\mu))$.*

The function $V(\mu)$ is simple for many NEFs that are theoretically important, and its simplicity appears to be a better measure of the “canonicity” of an NEF than, e.g., the simplicity of the probability density or mass function. Specifically, the most important NEFs have $V(\mu)$ a low-degree polynomial: $V(\mu)$ is constant only for the Gaussian NEF with some fixed variance, and linear only for the Poisson NEF and affine transformations thereof.

The situation becomes more interesting for $V(\mu)$ quadratic, which NEFs Morris gave the following name.

Definition 4.2.6. *If $V(\mu) = v_0 + v_1\mu + v_2\mu^2$ for some $v_i \in \mathbb{R}$, then we say that ν_0 generates a natural exponential family with quadratic variance function (NEF-QVF).*

Name	$d\nu_0(x)$	Support	$V(\mu)$
Gaussian (variance $\sigma^2 > 0$)	$\frac{1}{\sqrt{2\pi\sigma^2}} \exp(-\frac{x^2}{2\sigma^2}) dx$	\mathbb{R}	σ^2
Poisson	$\frac{1}{e} \frac{1}{x!}$	$\mathbb{Z}_{\geq 0}$	μ
Gamma (shape $\alpha > 0$)	$\frac{1}{\Gamma(\alpha)} x^{\alpha-1} e^{-x} dx$	$(0, +\infty)$	$\frac{1}{\alpha} \mu^2$
Binomial (m trials)	$\frac{1}{2^m} \binom{m}{x}$	$\{0, \dots, m\}$	$-\frac{1}{m} \mu^2 + \mu$
Negative Binomial (m successes)	$\frac{1}{2^{m+x}} \binom{x+m-1}{x}$	$\mathbb{Z}_{\geq 0}$	$\frac{1}{m} \mu^2 + \mu$
Hyperbolic Secant (shape $r > 0$)	([Mor82], Section 5)	\mathbb{R}	$\frac{1}{r} \mu^2 + r$
Wishart ($n \times n$, $N \geq n$ samples)	$\frac{\det(x)^{\frac{N-n-1}{2}} \exp(-\frac{\text{tr}(x)}{2})}{2^{nN/2} \Gamma_n(N/2)} dx$	$\mathbb{R}_{\geq 0}^{n \times n}$	$\mathbf{A} \mapsto \frac{2}{N} \text{tr}(\mu \mathbf{A} \mu \mathbf{A})$

Table 4.2: Exponential families and variance functions. We describe the six scalar natural exponential families with quadratic variance function from which, per the results of [Mor82], any such family can be generated by an affine transformation. We also show the analogous quantities for the Wishart matrix-valued natural exponential family; see Proposition 4.2.10 for further explanation.

NEF-QVFs are also sometimes called the *Morris class* of exponential families. One of the main results of [Mor82] is a complete classification of the NEF-QVFs, as follows.

Proposition 4.2.7. *Any NEF-QVF can be obtained by an affine transformation ($X \mapsto aX + b$ applied to the underlying random variables) of one of the six families listed in Table 4.2. Conversely, any affine transformation of an NEF-QVF yields another NEF-QVF.*

Other common distributions occur as special cases: Bernoulli is a special case of binomial, geometric is a special case of negative binomial, and exponential and chi-squared are both special cases of gamma. The sixth “generalized hyperbolic secant” family is more complicated to describe, but one representative distribution generating the $r = 1$ family has density $\frac{1}{2} \text{sech}(\frac{\pi}{2}x) dx$, and may be thought of as a smoothed Laplace distribution (see Figure 5.2 later in the text for an illustration). Indeed, its appearance is one of the major surprises of Morris’ results; since his work these distributions have found some applications (see, e.g., [Fis13]) but remain rather obscure.

The value of the quadratic coefficient ν_2 will play an important role in our results on the

low-degree likelihood ratio, so we make the following definition of the possible values this coefficient can take in advance.

Definition 4.2.8 (Possible values of v_2). *Let $\mathcal{V} := [0, +\infty) \cup \{-\frac{1}{m} : m \in \mathbb{N}_+\} \subset \mathbb{R}$.*

Proposition 4.2.9. *For any NEF-QVF, $v_2 \in \mathcal{V}$. Conversely, for any $v \in \mathcal{V}$, there exists an NEF-QVF with $v_2 = v$. The only NEF-QVFs with $v_2 < 0$ are the binomial families (including Bernoulli), and the only NEF-QVFs with $v_2 = 0$ are the Gaussian and Poisson families.*

Finally, let us show the analogy between the NEF-QVFs and the Wishart family, which also has a “quadratic variance function” when properly interpreted as a function taking the mean as input and outputting a linear functional on matrices. The reader may consult [LM08] for a deep discussion of this exponential family and its generalizations.

Proposition 4.2.10 (Variance function of Wishart family). *Suppose $N \geq n$. Let $\nu_0^{(n,N)} \in \mathcal{P}(\mathbb{R}_{>0}^{n \times n})$ be the law of $\sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^\top$ when $\mathbf{y}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ independently. Then, the NEF generated by $\nu_0^{(n,N)}$ (that is, if $p(\mathbf{Y})$ is the density of $\nu_0^{(n,N)}$, the collection of tilted measures with density proportional to $p(\mathbf{Y}) \exp(\langle \mathbf{A}, \mathbf{Y} \rangle)$ for $\mathbf{A} \in \mathbb{R}_{\text{sym}}^{n \times n}$, when this is integrable) consists of the laws of $\sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^\top$ when $\mathbf{y}_i \sim \mathcal{N}(\mathbf{0}, \Sigma)$ independently, for $\Sigma \in \mathbb{R}_{>0}^{n \times n}$. Calling the above law $\rho_{N\Sigma}^{(n,N)}$, we have that the mean of this law is $\boldsymbol{\mu} := N\Sigma$, so this gives the mean parametrization of the NEF. The associated variance function is then*

$$\text{Var}_{\mathbf{Y} \sim \rho_{\boldsymbol{\mu}}^{(n,N)}} [\langle \mathbf{Y}, \mathbf{A} \rangle] = \frac{2}{N} \text{tr}(\boldsymbol{\mu} \mathbf{A} \boldsymbol{\mu} \mathbf{A}). \quad (4.23)$$

We note that, rather remarkably, not only is this a quadratic polynomial of the entries of the matrix $\boldsymbol{\mu}$, but it is even the trace of a quadratic matrix polynomial of $\boldsymbol{\mu}$ and the input \mathbf{A} .

Proof. The density of $\nu_0^{(n,N)}$ is proportional to $\det(\mathbf{Y})^{(N-n-1)/2} \exp(-\frac{1}{2} \text{tr}(\mathbf{Y}))$. Thus, the density of the tilting by $\exp(\langle \mathbf{A}, \mathbf{Y} \rangle)$ is proportional to $\det(\mathbf{Y})^{(N-n-1)/2} \exp(-\frac{1}{2} \text{tr}((\mathbf{I}_n - 2\mathbf{A})\mathbf{Y}))$,

which gives a normalizable density if and only if $\mathbf{A} \prec \frac{1}{2}\mathbf{I}_n$. In this case, the above is proportional to $\det((\mathbf{I}_n - 2\mathbf{A})^{1/2}\mathbf{Y}(\mathbf{I}_n - 2\mathbf{A})^{1/2})^{(N-n-1)/2} \exp(-\frac{1}{2}\text{tr}((\mathbf{I}_n - 2\mathbf{A})^{1/2}\mathbf{Y}(\mathbf{I}_n - 2\mathbf{A})^{1/2}))$, which is the law of $\sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^\top$ for $\mathbf{y}_i \sim \mathcal{N}(\mathbf{0}, (\mathbf{I}_n - 2\mathbf{A})^{-1})$, giving the characterization of the NEF generated by $\nu_0^{(n,N)}$.

For the variance function, we compute that, if $\mathbf{y}_i \sim \mathcal{N}(\mathbf{0}, \Sigma)$, so that the mean is $\mathbb{E}\mathbf{Y} = N\Sigma =: \boldsymbol{\mu}$, then

$$\begin{aligned}
\mathbb{E}\langle \mathbf{Y}, \mathbf{A} \rangle^2 &= \mathbb{E} \left(\sum_{i=1}^N \mathbf{y}_i^\top \mathbf{A} \mathbf{y}_i \right)^2 \\
&= N(N-1) \langle \mathbf{A}, \Sigma \rangle^2 + N \mathbb{E}(\mathbf{y}_1^\top \mathbf{A} \mathbf{y}_1)^2 \\
&= N(N-1) \langle \mathbf{A}, \Sigma \rangle^2 + N \sum_{i,j,k,\ell=1}^N A_{ij} A_{k\ell} \mathbb{E} \mathcal{Y}_i \mathcal{Y}_j \mathcal{Y}_k \mathcal{Y}_\ell \\
&= N(N-1) \langle \mathbf{A}, \Sigma \rangle^2 + N \sum_{i,j,k,\ell=1}^N A_{ij} A_{k\ell} (\Sigma_{ij} \Sigma_{k\ell} + \Sigma_{ik} \Sigma_{j\ell} + \Sigma_{i\ell} \Sigma_{jk}) \quad (\text{Wick's formula}) \\
&= N \langle \mathbf{A}, \Sigma \rangle^2 + 2N \text{tr}(\Sigma \mathbf{A} \Sigma \mathbf{A}) \\
&= \langle \mathbf{A}, \boldsymbol{\mu} \rangle^2 + \frac{2}{N} \text{tr}(\boldsymbol{\mu} \mathbf{A} \boldsymbol{\mu} \mathbf{A}) \tag{4.24}
\end{aligned}$$

and thus

$$\text{Var}\langle \mathbf{Y}, \mathbf{A} \rangle = \mathbb{E}\langle \mathbf{Y}, \mathbf{A} \rangle^2 - \langle \mathbf{A}, \boldsymbol{\mu} \rangle^2 = \frac{2}{N} \text{tr}(\boldsymbol{\mu} \mathbf{A} \boldsymbol{\mu} \mathbf{A}), \tag{4.25}$$

as claimed. □

4.3 NEF-QVF MODELS

In this section, we give an overlap *bound* for the norm of the LDLR in NEF-QVFs, similar to the exact overlap *formula* we obtained for the Gaussian Wigner model. We consider here a model we call *kin-spiked*, where \mathbb{P}_n belongs to the same NEF as \mathbb{Q}_n but has a different mean. Later in Section 5.4.2 we will also work briefly with an *additively-spiked* model, where \mathbb{P}_n is a

translation of \mathbb{Q}_n , possibly not belonging to the same NEF. While these two models coincide for the Gaussian Wigner NEF, they do not for the other NEF-QVFs, and we will see that the kin-spiked model is far more natural mathematically.

Definition 4.3.1 (Kin-spiked NEF-QVF model). *Let ρ_μ be a mean-parametrized NEF-QVF over $\mu \in \Omega \subseteq \mathbb{R}$.¹ Let $N = N(n) \in \mathbb{N}$ and $\mu_{n,i} \in \Omega$ for each $n \in \mathbb{N}$ and $i \in [N(n)]$. Let $\tilde{\mathcal{P}}_n \in \mathcal{P}(\Omega^{N(n)})$. Then, define sequences of probability measures $\mathbb{P}_n, \mathbb{Q}_n$ as follows:*

- Under \mathbb{Q}_n , draw $y_i \sim \rho_{\mu_{n,i}}$ independently for $i \in [N(n)]$.
- Under \mathbb{P}_n , draw $\tilde{x} \sim \tilde{\mathcal{P}}_n$, and then draw $y_i \sim \rho_{\tilde{x}_i}$ independently for $i \in [N(n)]$.

4.3.1 ORTHOGONAL POLYNOMIALS IN NEF-QVFS

Our main tool will be that the orthogonal polynomials of NEF-QVFs satisfy variants of some of the identities of Hermite polynomials that we used in Section 4.1.1 to treat the Gaussian Wigner model, in particular the “translation formula” (Proposition 4.1.12) and the generating function (Proposition 4.1.14). Indeed, as far as the generating function is concerned, as for the Gaussian Wigner family, in an NEF-QVF there is a remarkable connection between the likelihood ratio and the orthogonal polynomials of ν_θ . The likelihood ratio in any NEF is simple:

$$\frac{d\nu_\theta}{d\nu_0}(y) = \exp(y\theta - \psi(\theta)), \quad (4.26)$$

where $\psi(\theta) = \mathbb{E}_{x \sim \nu_0}[\exp(\theta x)]$. We may also reparametrize in terms of the mean:

$$L(y; \mu) := L(y; (\psi')^{-1}(\mu)) = \exp(y(\psi')^{-1}(\mu) - \psi((\psi')^{-1}(\mu))). \quad (4.27)$$

As the following result of Morris shows, in an NEF-QVF, $L(y; \mu)$ is a kind of generating function of the orthogonal polynomials of ρ_μ .

¹It will not matter for our purposes what the base measure ν_0 is.

Definition 4.3.2. For $v \in \mathbb{R}$, define the sequences of constants

$$\hat{a}_k(v) := \prod_{j=0}^{k-1} (1 + vj), \quad (4.28)$$

$$a_k(v) := k! \cdot \hat{a}_k(v). \quad (4.29)$$

Proposition 4.3.3 (NEF-QVF Rodrigues Formula; Theorem 4 of [Mor82]). Let $\mu_0 = \psi'(0) = \mathbb{E}_{x \sim \nu_0}[x]$. Define the polynomials

$$p_k(\mathcal{Y}; \mu_0) := \frac{V(\mu_0)^k}{L(\mathcal{Y}, \mu_0)} \cdot \frac{d^k L}{d\mu^k}(\mathcal{Y}, \mu_0). \quad (4.30)$$

Then, $p_k(\mathcal{Y}; \mu_0)$ is a degree k monic polynomial in \mathcal{Y} , and this family satisfies the orthogonality relation

$$\mathbb{E}_{\mathcal{Y} \sim \rho_{\mu_0}} p_k(\mathcal{Y}; \mu_0) p_\ell(\mathcal{Y}; \mu_0) = \delta_{k\ell} \cdot a_k(v_2) V(\mu_0)^k. \quad (4.31)$$

In particular, defining the normalized polynomials

$$\hat{p}_k(\mathcal{Y}; \mu_0) := \frac{1}{V(\mu_0)^{k/2} \sqrt{a_k(v_2)}} p_k(\mathcal{Y}; \mu_0), \quad (4.32)$$

the $\hat{p}_k(\mathcal{Y}; \mu_0)$ are orthonormal polynomials for ρ_{μ_0} .

The main property of these polynomials that will be useful for us is the following identity, also obtained by Morris, giving the expectation of a given orthogonal polynomial under the kin spiking operation, i.e., under a different distribution from the same NEF-QVF.

Proposition 4.3.4 (Corollary 1 of [Mor82]). For all $k \in \mathbb{N}$ and $x, \mu \in \Omega$,

$$\mathbb{E}_{\mathcal{Y} \sim \rho_x} p_k(\mathcal{Y}; \mu) = \hat{a}_k(v_2) (x - \mu)^k. \quad (4.33)$$

This should be viewed as the analog of Proposition 4.1.12 for NEF-QVFs. We may obtain a

straightforward further corollary by including the normalization, which allows us to incorporate the variance factor into a *z-score*, as follows.

Definition 4.3.5 (NEF-QVF *z-score*). *For $\mu, x \in \Omega$, define the *z-score* as*

$$z_\mu(x) := \frac{x - \mu}{\sqrt{V(\mu)}}. \quad (4.34)$$

Corollary 4.3.6 (Kin-spiked expectation). *For all $k \in \mathbb{N}$ and $x, \mu \in \Omega$,*

$$\mathbb{E}_{y' \sim \rho_x} \hat{p}_k(y; \mu) = \sqrt{\frac{\hat{a}_k(\nu_2)}{k!}} z_\mu(x)^k. \quad (4.35)$$

4.3.2 COMPUTING THE LOW-DEGREE LIKELIHOOD RATIO

Returning to the multivariate setting, let \mathbb{Q}_n and \mathbb{P}_n be as in Definition 4.3.1 of the kin-spiked NEF-QVF model. Then, the likelihood ratio is

$$L_n(\mathbf{y}) := \frac{d\mathbb{P}_n}{d\mathbb{Q}_n}(\mathbf{y}) = \mathbb{E}_{\tilde{x} \sim \mathcal{P}_n} \left[\prod_{i=1}^N \frac{d\rho_{\tilde{x}_i}}{d\rho_{\mu_{n,i}}}(\mathcal{Y}_i) \right]. \quad (4.36)$$

An orthonormal system of polynomials for \mathbb{Q}_n is given by the product basis formed from the $\hat{p}_k(y; \mu_{n,i})$ that we defined in Proposition 4.3.3:

$$\hat{P}_k(\mathbf{y}; \boldsymbol{\mu}_n) := \prod_{i=1}^N \hat{p}_{k_i}(y_i; \mu_{n,i}) \quad (4.37)$$

for $\mathbf{k} \in \mathbb{N}^N$, where $\boldsymbol{\mu}_n := (\mu_{n,1}, \dots, \mu_{n,N(n)})$.

We then show that the projection of L_n onto any component $\hat{P}_k(\cdot; \boldsymbol{\mu}_n)$ admits the following convenient expression in terms of the *z-score*.

Lemma 4.3.7 (LDLR components for NEF-QVF model). *For L_n the likelihood ratio of the kin-*

spiked NEF-QVF model, for all $\mathbf{k} \in \mathbb{N}^N$,

$$\langle L_n, \hat{P}_{\mathbf{k}}(\cdot; \boldsymbol{\mu}_n) \rangle = \sqrt{\frac{\prod_{i=1}^N \hat{\mathbf{a}}_{k_i}(\nu_2)}{\prod_{i=1}^N k_i!}} \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_n} \left[\prod_{i=1}^N z_{\mu_{n,i}}(\mathbf{x}_i)^{k_i} \right]. \quad (4.38)$$

This is the direct analog of Proposition 4.1.15 for the Gaussian Wigner model and its Hermite polynomials.

Proof. Performing a change of measure using the likelihood ratio and factorizing the inner product using independence of coordinates under \mathbb{Q}_n , we find

$$\begin{aligned} \langle L_n, \hat{P}_{\mathbf{k}}(\cdot; \boldsymbol{\mu}_n) \rangle &= \mathbb{E}_{\mathbf{y} \sim \mathbb{Q}_n} \left[L_n(\mathbf{y}) \hat{P}_{\mathbf{k}}(\mathbf{y}; \boldsymbol{\mu}_n) \right] \\ &= \mathbb{E}_{\mathbf{y} \sim \mathbb{P}_n} \left[\hat{P}_{\mathbf{k}}(\mathbf{y}; \boldsymbol{\mu}_n) \right] \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_n} \left[\prod_{i=1}^N \mathbb{E}_{\mathcal{Y}_i \sim \tilde{\rho}_{\mathbf{x}_i}} \left[\hat{p}_{k_i}(\mathcal{Y}_i; \mu_{n,i}) \right] \right] \end{aligned}$$

and using Corollary 4.3.6,

$$= \sqrt{\frac{\prod_{i=1}^N \hat{\mathbf{a}}_{k_i}(\nu_2)}{\prod_{i=1}^N k_i!}} \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_n} \left[\prod_{i=1}^N z_{\mu_{n,i}}(\mathbf{x}_i)^{k_i} \right], \quad (4.39)$$

completing the proof. \square

First, we give an exact formula for the norm of the untruncated likelihood ratio in a kin-spiked NEF-QVF model. These involve a collection of link functions playing the role of $\phi^{\text{Wig}}(t) = \exp(t)$ from the Gaussian Wigner model.

Definition 4.3.8 (Link functions for NEF-QVF model). *For $t \in \mathbb{R}$ and $\nu \in \mathcal{V}$, define*

$$\phi^{\text{Mor}}(t; \nu) := \begin{cases} e^t & \text{if } \nu = 0, \\ (1 - \nu t)^{-1/\nu} & \text{if } \nu \neq 0 \text{ and } t < 1/|\nu|, \\ +\infty & \text{if } \nu > 0 \text{ and } t \geq 1/|\nu|. \end{cases} \quad (4.40)$$




Figure 4.1: NEF-QVF LDLR nonlinearities. We plot the functions $\phi^{\text{Mor}}(t; \nu)$ associated to the norm of the low-degree likelihood ratio of an NEF-QVF with variance function having quadratic coefficient ν . We emphasize the monotonicity in ν and the appearance of the exponential function for $\nu = 0$.

Moreover, for $D \in \mathbb{N}$, let $\phi_D^{\text{Mor}}(t; \nu)$ denote the order- D Taylor expansion of $\phi^{\text{Mor}}(t; \nu)$ about $t = 0$ for fixed ν , and let $\phi_{+\infty}^{\text{Mor}}(t; \nu) := \phi^{\text{Mor}}(t; \nu)$.

See Figure 4.1 for an illustration of these functions “sandwiching” the exponential.

Theorem 4.3.9 (LR norm for NEF-QVF model). *In the kin-spiked NEF-QVF model, for all $n \in \mathbb{N}$,*

$$\|L_n\|^2 = \mathbb{E}_{x^1, x^2 \sim \mathcal{P}_n} \left[\prod_{i=1}^N \phi^{\text{Mor}}(z_{\mu_{n,i}}(x_i^1) z_{\mu_{n,i}}(x_i^2); \nu_2) \right]. \quad (4.41)$$

The key technical step is to recognize that the function $\phi^{\text{Mor}}(t; \nu)$ is in fact the exponential generating function of the $\hat{a}_k(\nu)$, as follows.

Proposition 4.3.10. For all $t \in \mathbb{R}$ and $v \in \mathcal{V}$,

$$\phi^{\text{Mor}}(t; v) = \sum_{k=0}^{\infty} \frac{\hat{a}_k(v)}{k!} t^k. \quad (4.42)$$

Proof. Call the right-hand side $f(t; v)$. Differentiating the power series termwise and using the formula from Definition 4.3.2 gives the differential equation

$$\frac{\partial f}{\partial t}(t; v) = f(t; v) + vt \frac{\partial f}{\partial t}(t; v) \quad (4.43)$$

with initial condition $f(0; v) = 1$, and the result follows upon solving the equation. \square

Proof of Theorem 4.3.9. We have by Lemma 4.3.7

$$\begin{aligned} \|L_n\|^2 &= \sum_{\mathbf{k} \in \mathbb{N}^N} \langle L_n, \hat{P}_{\mathbf{k}}(\cdot; \boldsymbol{\mu}_n) \rangle^2 \\ &= \sum_{\mathbf{k} \in \mathbb{N}^N} \frac{\prod_{i=1}^N \hat{a}_{k_i}(v_2)}{\prod_{i=1}^N k_i!} \left(\mathbb{E}_{x \sim \mathcal{P}_n} \left[\prod_{i=1}^N z_{\mu_{n,i}}(x_i)^{k_i} \right] \right)^2 \\ &= \mathbb{E}_{x^1, x^2 \sim \mathcal{P}_n} \left[\sum_{\mathbf{k} \in \mathbb{N}^N} \prod_{i=1}^N \left\{ \frac{\hat{a}_{k_i}(v_2)}{k_i!} (z_{\mu_{n,i}}(x_i^1) z_{\mu_{n,i}}(x_i^2))^{k_i} \right\} \right] \\ &= \mathbb{E}_{x^1, x^2 \sim \mathcal{P}_n} \left[\prod_{i=1}^N \left\{ \sum_{k=0}^{\infty} \frac{\hat{a}_k(v_2)}{k!} (z_{\mu_{n,i}}(x_i^1) z_{\mu_{n,i}}(x_i^2))^k \right\} \right], \end{aligned} \quad (4.44)$$

and the result follows from Proposition 4.3.10. \square

We note that this is not yet an ‘‘overlap bound’’ except in the case $v_2 = 0$, since we do not have $\phi^{\text{Mor}}(s; v_2) \phi^{\text{Mor}}(t; v_2) = \phi^{\text{Mor}}(s+t; v_2)$ when $v_2 \neq 0$. However, we show in the next result that, both for the norm of the full and low-degree likelihood ratios, we may obtain such a bound from either above or below, depending on the sign of v_2 .

Theorem 4.3.11 (LDLR norm for NEF-QVF model). *Let $\rho_{\boldsymbol{\mu}}$ be a mean-parametrized NEF-QVF over $\boldsymbol{\mu} \in \Omega \subseteq \mathbb{R}$, with variance function $V(\boldsymbol{\mu}) = v_0 + v_1 \boldsymbol{\mu} + v_2 \boldsymbol{\mu}^2$. Let $\mu_{n,i} \in \Omega$ and $\tilde{\mathcal{P}}_n$ be as in*

Definition 4.3.1 of the kin-spiked NEF-QVF model. Define the z-score overlap,

$$R_n := \sum_{i=1}^{N(n)} z_{\mu_{n,i}}(\tilde{\mathbf{x}}_i^1) z_{\mu_{n,i}}(\tilde{\mathbf{x}}_i^2), \quad (4.45)$$

where $\tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2 \sim \tilde{\mathcal{P}}_n$ independently. Let $L_n^{\leq D}$ denote the low-degree likelihood ratio.

- If $v_2 \geq 0$, then for any $n \in \mathbb{N}$ and $D \in \mathbb{N} \cup \{+\infty\}$,

$$\|L_n^{\leq D}\|^2 \leq \mathbb{E} \left[\phi_D^{\text{Mor}}(R_n; v_2) \right], \quad (4.46)$$

and equality holds if $v_2 = 0$ (i.e., in the Gaussian and Poisson NEFs).

- If $v_2 < 0$, then for any $n \in \mathbb{N}$ and $D \in \mathbb{N} \cup \{+\infty\}$,

$$\mathbb{E} \left[\phi_D^{\text{Mor}}(R_n; v_2) \right] \leq \|L_n^{\leq D}\|^2 \leq \mathbb{E} \left[\phi_D^{\text{Mor}}(R_n; 0) \right]. \quad (4.47)$$

To prove this result, we first establish two more ancillary facts about the power series coefficients $\hat{a}_k(v)$.

Proposition 4.3.12 (Monotonicity). *For $k \in \mathbb{N}$, $\hat{a}_k(v)$ is non-negative and monotonically non-decreasing in v over $v \in \mathcal{V}$.*

Proof. Recall from (4.28) that, by definition,

$$\hat{a}_k(v) = \prod_{j=0}^{k-1} (1 + vj). \quad (4.48)$$

Thus clearly $\hat{a}_k(v)$ is monotonically non-decreasing over $v \geq 0$, since each factor is monotonically non-decreasing.

If $v \in \mathcal{V}$ with $v < 0$, then $v = -\frac{1}{m}$ for some $m \in \mathbb{Z}_{\geq 1}$. Thus for $k \geq m + 1$, $\hat{a}_k(v) = 0$. So,

in this case we may rewrite

$$\hat{a}_k(\nu) = \mathbb{1}\{k \leq m\} \prod_{j=0}^{\min\{k-1, m-1\}} (1 + \nu j). \quad (4.49)$$

Now, each factor belongs to $[0, 1)$, and again each factor is monotonically non-decreasing with ν , so the result follows. \square

Proposition 4.3.13 (Multiplicativity relations). *For all $\mathbf{k} \in \mathbb{N}^N$ and $\nu \in \mathcal{V}$,*

$$\prod_{i=1}^N \hat{a}_{k_i}(\nu) \leq \hat{a}_{\sum_{i=1}^N k_i}(\nu) \text{ if } \nu > 0, \quad (4.50)$$

$$\prod_{i=1}^N \hat{a}_{k_i}(\nu) = \hat{a}_{\sum_{i=1}^N k_i}(\nu) \text{ if } \nu = 0, \quad (4.51)$$

$$\prod_{i=1}^N \hat{a}_{k_i}(\nu) \geq \hat{a}_{\sum_{i=1}^N k_i}(\nu) \text{ if } \nu < 0. \quad (4.52)$$

Proof. When $\nu = 0$, then $\hat{a}_k(\nu) = 1$ for all k , so the result follows immediately. When $\nu > 0$, we have

$$\begin{aligned} \prod_{i=1}^N \hat{a}_{k_i}(\nu) &= \prod_{i=1}^N \prod_{j=0}^{k_i-1} (1 + \nu j) \\ &\leq \prod_{i=1}^N \prod_{j=\sum_{a=1}^{i-1} k_a}^{\sum_{a=1}^i k_a} (1 + \nu j) \\ &= \prod_{j=1}^{\sum_{i=1}^N k_i} (1 + \nu j) \\ &= \hat{a}_{\sum_{i=1}^N k_i}(\nu). \end{aligned} \quad (4.53)$$

When $\nu < 0$, a symmetric argument together with the observations from Proposition 4.3.12 gives the result. \square

Proof of Theorem 4.3.11. Suppose first that $v_2 \geq 0$. We have by Lemma 4.3.7

$$\begin{aligned}
\|L_n^{\leq D}\|^2 &= \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \langle L_n, \hat{P}_{\mathbf{k}}(\cdot; \boldsymbol{\mu}_n) \rangle^2 \\
&= \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \frac{\prod_{i=1}^N \hat{a}_{k_i}(v_2)}{\prod_{i=1}^N k_i!} \left(\mathbb{E}_{\tilde{\mathbf{x}} \sim \mathcal{P}_n} \left[\prod_{i=1}^N z_{\mu_{n,i}}(\tilde{x}_i)^{k_i} \right] \right)^2 \\
&= \mathbb{E}_{\tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2 \sim \tilde{\mathcal{P}}_n} \left[\sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \frac{\prod_{i=1}^N \hat{a}_{k_i}(v_2)}{\prod_{i=1}^N k_i!} \prod_{i=1}^N (z_{\mu_{n,i}}(\tilde{x}_i^1) z_{\mu_{n,i}}(\tilde{x}_i^2))^{k_i} \right],
\end{aligned}$$

Using Proposition 4.3.13,

$$\leq \mathbb{E}_{\tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2 \sim \mathcal{P}_n} \left[\sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \frac{\hat{a}_{|\mathbf{k}|}(v_2)}{\prod_{i=1}^N k_i!} \prod_{i=1}^N (z_{\mu_{n,i}}(\tilde{x}_i^1) z_{\mu_{n,i}}(\tilde{x}_i^2))^{k_i} \right]$$

and following the same manipulations from the proof of Theorem 4.1.7,

$$\begin{aligned}
&= \mathbb{E}_{\tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2 \sim \mathcal{P}_n} \left[\sum_{d=0}^D \frac{\hat{a}_d(v_2)}{d!} \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}|=d}} \binom{d}{k_1, \dots, k_N} \prod_{i=1}^N (z_{\mu_{n,i}}(\tilde{x}_i^1) z_{\mu_{n,i}}(\tilde{x}_i^2))^{k_i} \right] \\
&= \mathbb{E}_{\tilde{\mathbf{x}}^1, \tilde{\mathbf{x}}^2 \sim \mathcal{P}_n} \left[\sum_{d=0}^D \frac{\hat{a}_d(v_2)}{d!} \left(\sum_{i=1}^N z_{\mu_{n,i}}(\tilde{x}_i^1) z_{\mu_{n,i}}(\tilde{x}_i^2) \right)^d \right], \tag{4.54}
\end{aligned}$$

giving the upper bound from (4.46) for $v_2 > 0$. When $v_2 = 0$, then equality holds above, so we obtain equality in (4.46). Also, when $v_2 < 0$, then the above argument holds with the inequality reversed, giving the lower bound of (4.47).

Finally, for the upper bound of (4.47), note that when $v_2 < 0$, we may bound $\|L_n^{\leq D}\|^2$

using Proposition 4.3.12 and the result for $v_2 = 0$ by

$$\begin{aligned}
\|L_n^{\leq D}\|^2 &= \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \frac{\prod_{i=1}^N \hat{a}_{k_i}(v_2)}{\prod_{i=1}^N k_i!} \left(\mathbb{E}_{\tilde{\mathcal{P}}_n} \left[\prod_{i=1}^N z_{\mu_{n,i}}(\tilde{\mathcal{X}}_i)^{k_i} \right] \right)^2 \\
&\leq \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \frac{\prod_{i=1}^N \hat{a}_{k_i}(0)}{\prod_{i=1}^N k_i!} \left(\mathbb{E}_{\tilde{\mathcal{P}}_n} \left[\prod_{i=1}^N z_{\mu_{n,i}}(\tilde{\mathcal{X}}_i)^{k_i} \right] \right)^2 \\
&= \mathbb{E}_{\tilde{\mathcal{X}}^1, \tilde{\mathcal{X}}^2 \sim \tilde{\mathcal{P}}_n} \phi_D^{\text{Mor}}(R_n; 0), \tag{4.55}
\end{aligned}$$

giving the result. □

4.3.3 CHANNEL MONOTONICITY

The monotonicity of the functions $\phi^{\text{Mor}}(t; v)$ in v evident in Figure 4.1 suggests that we might expect $\|L_n^{\leq D}\|$ to be monotone across different kin-spiked NEF-QVF models with the same signal prior $\tilde{\mathcal{P}}_n$. While this does not follow directly from the above result, a slightly more careful argument shows that it is indeed the case.

Theorem 4.3.14. *Suppose $L_n^{(i)}$ for $i \in \{1, 2\}$ are the likelihood ratios for the hypothesis testing problems in two kin-spiked NEF-QVF models, with mean domains $\Omega^{(i)}$ and variance functions $V^{(i)}(\mu) = v_0^{(i)} + v_1^{(i)}\mu + v_2^{(i)}\mu^2$. Suppose that the null means $\mu_{n,j}$ and the signal prior $\tilde{\mathcal{P}}_n$ are the same in both problems (in particular, $\Omega^{(1)} \cap \Omega^{(2)}$ must contain the support of $\tilde{\mathcal{P}}_n$). If $v_2^{(1)} \leq v_2^{(2)}$, then, for any $D \in \mathbb{N} \cup \{+\infty\}$, $\|(L_n^{(1)})^{\leq D}\|^2 \leq \|(L_n^{(2)})^{\leq D}\|^2$.*

Informally, this says that if $v_2^{(1)} \leq v_2^{(2)}$, then “Problem 1 is at least as hard as Problem 2,” for any given computational budget. For example, for a fixed collection of null means $\mu_{n,i}$ and a fixed signal prior $\tilde{\mathcal{P}}_n$, we would predict the following relationships among output “channels” or observation distributions, with “ \geq ” denoting greater computational difficulty:

$$\text{Bernoulli} \geq \text{Binomial} \geq \text{Gaussian} = \text{Poisson} \geq \text{Exponential}. \tag{4.56}$$

The result is a simple consequence of the arguments we have made already to prove Theorem 4.3.11.

Proof of Theorem 4.3.14. We have by Lemma 4.3.7

$$\begin{aligned} \|(L_n^{(i)})^{\leq D}\|^2 &= \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \langle L_n^{(i)}, \hat{P}_{\mathbf{k}}(\cdot; \boldsymbol{\mu}_n) \rangle^2 \\ &= \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}| \leq D}} \frac{\prod_{j=1}^N \hat{a}_{k_j}(\mathbf{v}_2^{(i)})}{\prod_{j=1}^N k_j!} \left(\mathbb{E}_{\tilde{\mathbf{x}} \sim \tilde{P}_n} \left[\prod_{j=1}^N z_{\mu_{n,j}} (\tilde{x}_j)^{k_j} \right] \right)^2. \end{aligned} \quad (4.57)$$

In each term on the right-hand side, the only factor that depends on i is $\prod_{j=1}^N \hat{a}_{k_j}(\mathbf{v}_2^{(i)})$, so the result follows from the monotonicity described by Proposition 4.3.12. (Indeed, this shows slightly more, that the monotonicity holds even for the norm of the projection of $L_n^{(i)}$ onto the orthogonal polynomial of any given index \mathbf{k} .) \square

4.4 GAUSSIAN WISHART MODELS

Finally, we proceed to considering the prospect of an overlap formula in the more unusual setting of the Wishart spiked matrix model of Definition 2.2.4. In fact, we will be able to treat a more general setting, which allows arbitrary, possibly sign-indefinite signal matrices.

Definition 4.4.1 (Gaussian Wishart model). *Let $\tilde{P}_n \in \mathcal{P}(\mathbb{R}_{\text{sym}}^{n \times n})$ be such that $\tilde{\mathbf{X}} \succ -\mathbf{I}_n$ almost surely. The Gaussian Wishart model with sampling ratio γ and spike prior \tilde{P}_n is specified by the following distributions over $(\mathbf{y}_1, \dots, \mathbf{y}_N) \in (\mathbb{R}^n)^N$ with $N = N(n) = \lfloor n/\gamma \rfloor$:*

- Under \mathbb{Q}_n , draw $\mathbf{y}_1, \dots, \mathbf{y}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ independently.
- Under \mathbb{P}_n , first draw $\tilde{\mathbf{X}} \sim \tilde{P}_n$, and then draw $\mathbf{y}_1, \dots, \mathbf{y}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n + \tilde{\mathbf{X}})$ independently.

We might hope that $\|L_n^{\leq D}\|$ in this model would be governed by an expectation of a function of $\langle \tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \rangle$, as for the Gaussian Wigner model. Unfortunately, all is not so simple:

for Gaussian Wishart models, we will see that we must instead work with an *overlap matrix*, given by $\tilde{\mathbf{X}}^{-1}\tilde{\mathbf{X}}^2$. This may appear to be at odds with our goal of reducing the dimensionality of the expectations arising in the evaluation of $\|L_n^{\leq D}\|$. However, a further nuance comes to the rescue: the link function applied to the overlap matrix turns out to be $\det(\mathbf{I}_n - \tilde{\mathbf{X}}^{-1}\tilde{\mathbf{X}}^2)^{-N/2}$, whereby, by Sylvester's determinant identity, when $\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n$ is low-rank, say rank k , we would be able to reduce our formula to an expectation of a function of a random $k \times k$, rather than $n \times n$, matrix. Actually, in Section 5.2.1 we will also deduce from this a bound comparing Wishart spiked matrix models to corresponding Wigner models, which will, up to a usually-negligible correction, realize our initial hope to work with expectations over $\langle \tilde{\mathbf{X}}^{-1}, \tilde{\mathbf{X}}^2 \rangle$, at least in our applications where the signal is sign-definite.

Before proceeding, let us give the derivation of the untruncated likelihood ratio and its norm under such a model. This will confirm the link function mentioned above, and will show what we expect from the computations to come. (A similar computation, though constrained to the case of $\tilde{\mathbf{X}}$ having rank one, appears in [PWBM18].) Let us view the observations as a collection $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_N) \in (\mathbb{R}^n)^N$.

Proposition 4.4.2 (Likelihood ratio in Gaussian Wishart model). *Suppose (\mathbb{P}_n) and (\mathbb{Q}_n) are a Gaussian Wishart model as in Definition 4.4.1, with signal prior $(\tilde{\mathcal{P}}_n)$. Then,*

$$\begin{aligned} L_n(\mathbf{Y}) &= \frac{d\mathbb{P}_n}{d\mathbb{Q}_n}(\mathbf{Y}) \\ &= \mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \left[\det(\mathbf{I}_n + \tilde{\mathbf{X}})^{-N/2} \exp \left(-\frac{1}{2} \sum_{i=1}^N \mathbf{y}_i^\top ((\mathbf{I}_n + \tilde{\mathbf{X}})^{-1} - \mathbf{I}_n) \mathbf{y}_i \right) \right]. \end{aligned} \quad (4.58)$$

Proof. Expanding the Gaussian densities as for the Gaussian Wigner model,

$$\frac{d\mathbb{Q}_n}{d\mathcal{L}}(\mathbf{y}_1, \dots, \mathbf{y}_N) = (2\pi)^{-nN/2} \cdot \exp\left(-\frac{1}{2} \sum_{i=1}^N \|\mathbf{y}_i\|^2\right) \quad (4.59)$$

$$\begin{aligned} \frac{d\mathbb{P}_n}{d\mathcal{L}}(\mathbf{y}_1, \dots, \mathbf{y}_N) &= (2\pi)^{-nN/2} \cdot \mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \left[\det(\mathbf{I}_n + \tilde{\mathbf{X}})^{-N/2} \exp\left(-\frac{1}{2} \sum_{i=1}^N \mathbf{y}_i^\top (\mathbf{I}_n + \tilde{\mathbf{X}})^{-1} \mathbf{y}_i\right) \right] \\ &= (2\pi)^{-N/2} \cdot \exp\left(-\frac{1}{2} \|\mathbf{Y}\|^2\right) \cdot \mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \left[\exp\left(-\frac{1}{2} \|\tilde{\mathbf{X}}\|^2 + \langle \tilde{\mathbf{X}}, \mathbf{Y} \rangle\right) \right], \end{aligned} \quad (4.60)$$

and taking the quotient gives the result. \square

Proposition 4.4.3 (LR norm for Gaussian Wishart model). *Suppose (\mathbb{P}_n) and (\mathbb{Q}_n) are a Gaussian Wishart model as in Definition 4.4.1, with signal prior $(\tilde{\mathcal{P}}_n)$. Let $L_n = d\mathbb{P}_n/d\mathbb{Q}_n$. Then,*

$$\|L_n\|^2 = \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \sim \tilde{\mathcal{P}}_n} \det(\mathbf{I}_n - \tilde{\mathbf{X}}^1 \tilde{\mathbf{X}}^2)^{-N/2}. \quad (4.61)$$

This is essentially Lemma 7 of [CMW15], and the rank-one case appeared later as Proposition 5.11 of [PWBM18].

Proof. We compute directly, following the proof of Proposition 4.1.5 from the Gaussian Wigner model,

$$\|L_n\|^2 = \mathbb{E}_{\mathbf{y}_1, \dots, \mathbf{y}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)} \left(\mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \left[\det(\mathbf{I}_n + \tilde{\mathbf{X}})^{-N/2} \exp\left(-\frac{1}{2} \sum_{i=1}^N \mathbf{y}_i^\top ((\mathbf{I}_n + \tilde{\mathbf{X}})^{-1} - \mathbf{I}_n) \mathbf{y}_i\right) \right] \right)^2$$

Using the replica manipulation (Proposition 4.1.4),

$$\begin{aligned} &= \mathbb{E}_{\mathbf{y}_1, \dots, \mathbf{y}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)} \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2} \det(\mathbf{I}_n + \tilde{\mathbf{X}}^1)^{-N/2} \det(\mathbf{I}_n + \tilde{\mathbf{X}}^2)^{-N/2} \\ &\quad \exp\left(-\frac{1}{2} \sum_{i=1}^N \mathbf{y}_i^\top ((\mathbf{I}_n + \tilde{\mathbf{X}}^1)^{-1} + (\mathbf{I}_n + \tilde{\mathbf{X}}^2)^{-1} - 2\mathbf{I}_n) \mathbf{y}_i\right) \end{aligned}$$

and swapping the order of expectations,

$$\begin{aligned}
&= \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2} \det(\mathbf{I}_n + \tilde{\mathbf{X}}^1)^{-N/2} \det(\mathbf{I}_n + \tilde{\mathbf{X}}^2)^{-N/2} \\
&\quad \mathbb{E}_{\mathbf{y}_1, \dots, \mathbf{y}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)} \exp \left(-\frac{1}{2} \sum_{i=1}^N \mathbf{y}_i^\top \left((\mathbf{I}_n + \tilde{\mathbf{X}}^1)^{-1} + (\mathbf{I}_n + \tilde{\mathbf{X}}^2)^{-1} - 2\mathbf{I}_n \right) \mathbf{y}_i \right)
\end{aligned}$$

Applying an orthogonal change of basis to diagonalize this matrix and evaluating the resulting χ^2 moment generating function then gives

$$\begin{aligned}
&= \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2} \det(\mathbf{I}_n + \tilde{\mathbf{X}}^1)^{-N/2} \det(\mathbf{I}_n + \tilde{\mathbf{X}}^2)^{-N/2} \det \left((\mathbf{I}_n + \tilde{\mathbf{X}}^1)^{-1} + (\mathbf{I}_n + \tilde{\mathbf{X}}^2)^{-1} - \mathbf{I}_n \right)^{-N/2} \\
&= \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2} \det \left((\mathbf{I}_n + \tilde{\mathbf{X}}^1) \left((\mathbf{I}_n + \tilde{\mathbf{X}}^1)^{-1} + (\mathbf{I}_n + \tilde{\mathbf{X}}^2)^{-1} - \mathbf{I}_n \right) (\mathbf{I}_n + \tilde{\mathbf{X}}^2) \right)^{-N/2} \\
&= \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2} \det(\mathbf{I}_n - \tilde{\mathbf{X}}^1 \tilde{\mathbf{X}}^2)^{-N/2}, \tag{4.62}
\end{aligned}$$

as claimed. □

Therefore, by analogy with our previous results, we expect that the LDLR norm in the Gaussian Wishart model should involve the truncation of a Taylor series (in some suitable sense) of the function $\mathbf{T} \mapsto \det(\mathbf{I}_n - \mathbf{T})^{-N/2}$. We derive this series representation now. We write $\text{Part}(S)$ for the set of partitions of a set S .

Definition 4.4.4. For $d \in \mathbb{N}$ and $\mathbf{T} \in \mathbb{R}_{\text{sym}}^{n \times n}$, define the polynomials

$$r_{N,d}(\mathbf{T}) := \frac{1}{d!} \sum_{\pi \in \text{Part}([d])} \left(\frac{N}{2} \right)^{|\pi|} \prod_{S \in \pi} (|S| - 1)! \text{tr}(\mathbf{T}^{|S|}). \tag{4.63}$$

Note that $r_{N,d}(\mathbf{T})$ is homogeneous of degree d in the entries of \mathbf{T} .

Proposition 4.4.5. If $\mathbf{T} \in \mathbb{R}^{n \times n}$ satisfies $\|\mathbf{T}\| < 1$, then

$$\det(\mathbf{I}_n - \mathbf{T})^{-N/2} = \sum_{d=0}^{\infty} r_{N,d}(\mathbf{T}). \tag{4.64}$$

Proof. We manipulate, using that the matrix logarithm $\log(\mathbf{I}_n - \mathbf{T})$ is well-defined thanks to the norm constraint on \mathbf{T} ,

$$\begin{aligned}
\det(\mathbf{I}_n - \mathbf{T})^{-N/2} &= \exp(\log \det(\mathbf{I}_n - \mathbf{T}))^{-N/2} \\
&= \exp\left(-\frac{N}{2} \operatorname{tr} \log(\mathbf{I}_n - \mathbf{T})\right) \\
&= \exp\left(\frac{N}{2} \operatorname{tr} \left(\sum_{k=1}^{\infty} \frac{1}{k} \mathbf{T}^k\right)\right) \\
&= \exp\left(\frac{N}{2} \sum_{k=1}^{\infty} \frac{1}{k!} (k-1)! \operatorname{tr}(\mathbf{T}^k)\right), \tag{4.65}
\end{aligned}$$

and composing the exponential power series with the inner one gives the result (see, e.g., [FS09] for the analogous result for generating functions, from which this follows). \square

We also register the following cyclic property that the $r_{N,d}(\mathbf{A})$ inherit from their being polynomials of traces of powers, which will be crucial in our computations.

Proposition 4.4.6. *For any matrices \mathbf{A}, \mathbf{B} , not necessarily square but of compatible dimension to form a square product \mathbf{AB} , $r_{N,d}$ has the same cyclic property as the trace,*

$$r_{N,d}(\mathbf{AB}) = r_{N,d}(\mathbf{BA}). \tag{4.66}$$

Remark 4.4.7. *Incidentally, this reasoning also gives a simple account of Sylvester's determinant identity $\det(\mathbf{I} + \mathbf{AB}) = \det(\mathbf{I} + \mathbf{BA})$: by the same introduction of $\exp(\log(\cdot))$, these functions (for sufficiently small \mathbf{A}, \mathbf{B} in norm) expand in polynomials of $\operatorname{tr}((\mathbf{AB})^k)$, so the cyclic property is inherited from that of the trace.*

4.4.1 HERMITE POLYNOMIALS, REVISITED

We now work from the other direction, towards deriving the components of the LDLR norm. As we have seen earlier, the most robust approach to evaluating $\|L_n^{\leq D}\|$ with orthogonal

polynomials appears to be evaluating expectations of orthogonal polynomials for a given parametrized family of distributions with the “wrong” parameter: this is what we did for translations of Hermite polynomials in Proposition 4.1.12, and for distributions in a given NEF-QVF with different means in Corollary 4.3.6. In the Wishart model, the relevant orthogonal polynomials are again Hermite polynomials, only now we must consider their expectations under Gaussians with differing *variance* rather than differing mean. The key tool for our analysis will therefore generalize the following beautiful fact from the “umbral calculus” of Hermite polynomials (a proof will be subsumed in our result below).

Proposition 4.4.8 (Hermite mismatched variance identity). *Let $a > -1$. Then,*

$$\mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, 1+a)} h_k(\mathbf{y}) = \begin{cases} (k-1)!! \cdot a^{k/2} & \text{if } k \text{ is even,} \\ 0 & \text{if } k \text{ is odd.} \end{cases} \quad (4.67)$$

Note that the formula on the right-hand side is that for the moments of a Gaussian random variable with variance a , but we extend it to apply even for negative a , which is the “umbral” case of the result, admitting an interpretation in terms of a fictitious Gaussian of negative variance—even if $a \in (-1, 0)$, the right-hand side may be viewed formally as the value of “ $\mathbb{E}_{g \sim \mathcal{N}(0, a)} g^k$.” A thorough exposition of such analogies arising in combinatorics and associated to various polynomial sequences is given in [Rom05].

In fact, the same holds even for multivariate Gaussians. The correct result in this case is given by imitating the formula for the moments of a multivariate Gaussian, via Wick’s formula. While Proposition 4.4.8 is well-known in the literature on Hermite polynomials and the umbral calculus, we are not aware of previous appearances of the formula below (though it is likely folklore).

Proposition 4.4.9 (Multivariate Hermite mismatched variance identity). *Let $\mathbf{A} \in \mathbb{R}_{\text{sym}}^{n \times n}$ with $\mathbf{A} \succ -\mathbf{I}_n$. For $\mathbf{k} \in \mathbb{N}^n$ viewed as a multiset of elements of $[n]$, let $\text{Part}(\mathbf{k}; 2)$ be the set of*

matchings of elements of \mathbf{k} (the empty set for n odd), and for each $M \in \text{Part}(\mathbf{k}; 2)$ write \mathbf{A}^M for the product of the entries of \mathbf{A} located at positions indexed by pairs in M ; i.e., if $M = \{\{i_1, j_1\}, \dots, \{i_m, j_m\}\}$, then $\mathbf{A}^M = \prod_{a=1}^m A_{i_a j_a}$. Then,

$$\mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, \mathbf{I}_n + \mathbf{A})} H_{\mathbf{k}}(\mathbf{y}) = \sum_{M \in \text{Part}(\mathbf{k}; 2)} \mathbf{A}^M. \quad (4.68)$$

Similarly to before, if $\mathbf{A} \geq 0$ then the right-hand side equals $\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{A})} \mathbf{g}^{\mathbf{k}}$ by Wick's formula, but again we have an umbral extension to indefinite matrices \mathbf{A} .

We will use the following standard Gaussian integration by parts result, generalized in a different direction from Proposition 4.1.13 presented before.

Proposition 4.4.10 (Gaussian integration by parts: general covariance). *Let $\Sigma \in \mathbb{R}_{\geq 0}^{n \times n}$ and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be continuously differentiable with $f(\mathbf{y})$ and $\partial_i f(\mathbf{y})$ bounded by $O(\exp(|\mathbf{y}|^\alpha))$ for some $\alpha \in (0, 2)$. Then,*

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Sigma)} x_i f(\mathbf{x}) = \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Sigma)} \sum_{j=1}^n \Sigma_{ij} \frac{\partial f}{\partial x_j}(\mathbf{x}). \quad (4.69)$$

In matrix notation,

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Sigma)} \mathbf{x} f(\mathbf{x}) = \Sigma \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Sigma)} \nabla f(\mathbf{x}). \quad (4.70)$$

Proof of Proposition 4.4.9. Define

$$\ell_{\mathbf{k}} := \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, \mathbf{I}_n + \mathbf{X})} H_{\mathbf{k}}(\mathbf{y}). \quad (4.71)$$

Let $\mathbf{e}_i \in \mathbb{N}^n$ have i th coordinate equal to 1 and all other coordinates equal to zero, and write $\mathbf{0} \in \mathbb{N}^n$ for the vector with all coordinates equal to zero. Clearly $\ell_{\mathbf{0}} = 1$ and $\ell_{\mathbf{e}_i} = 0$ for any

$i \in [n]$. We then proceed by induction on $|\mathbf{k}|$:

$$\begin{aligned}
\ell_{\mathbf{k}+e_i} &= \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n + \mathbf{X})} h_{k_i+1}(\mathbf{y}_i) \prod_{j \in [n] \setminus \{i\}} h_{k_j}(\mathbf{y}_j) \\
&= \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n + \mathbf{X})} \left(\mathbf{y}_i h_{k_i}(\mathbf{y}_i) - h'_{k_i}(\mathbf{y}_i) \right) \prod_{j \in [n] \setminus \{i\}} h_{k_j}(\mathbf{y}_j) && \text{(Definition 4.1.8)} \\
&= \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n + \mathbf{X})} \left[\sum_{a=1}^n (\mathbf{I}_n + \mathbf{X})_{ia} \prod_{j \in [n]} h_{k_j}^{(\delta_{a,j})}(\mathbf{y}_j) - \prod_{j \in [n]} h_{k_j}^{(\delta_{i,j})}(\mathbf{y}_j) \right] && \text{(Proposition 4.4.10)} \\
&= \sum_{\substack{a \in [n] \\ k_a > 0}} X_{ia} \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n + \mathbf{X})} \prod_{j \in [n]} h_{k_j}^{(\delta_{j,a})}(\mathbf{y}_j) \\
&= \sum_{\substack{a \in [n] \\ k_a > 0}} k_a X_{ia} \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n + \mathbf{X})} \prod_{j \in [n]} h_{k_j - \delta_{j,a}}(\mathbf{y}_j) \\
&= \sum_{\substack{a \in [n] \\ k_a > 0}} k_a X_{ia} \ell_{\mathbf{k} - e_a}. && \text{(inductive hypothesis)}
\end{aligned}$$

Thus $\ell_{\mathbf{k}}$ satisfy the same recursion and initial condition as the sum-of-products formula on the right-hand side of (4.68), completing the proof. \square

4.4.2 COMPUTING THE LOW-DEGREE LIKELIHOOD RATIO

This tool in hand, we proceed towards evaluating the LDLR norm. Let (\mathbb{Q}_n) and (\mathbb{P}_n) be as in the Gaussian Wishart model (Definition 2.2.4), and let L_n be the associated likelihood ratio. Recall that the ambient parameter N is the number of samples \mathbf{y}_i we observe. Recall also that we assume the signal prior to satisfy that $\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n$ satisfies $\tilde{\mathbf{X}} \succ -\mathbf{I}_n$ almost surely.

Definition 4.4.11 (Link functions for Gaussian Wishart model). *For $N, D \in \mathbb{N}$ and $\mathbf{T} \in \mathbb{R}^{n \times n}$, we define the functions*

$$\phi_{N,D}^{\text{Wish}}(\mathbf{T}) := \sum_{d=0}^{\lfloor D/2 \rfloor} r_{N,d}(\mathbf{T}), \tag{4.72}$$

$$\phi_N^{\text{Wish}}(\mathbf{T}) = \phi_{N,\infty}^{\text{Wish}}(\mathbf{T}) := \det(\mathbf{I} - \mathbf{T})^{-N/2}. \tag{4.73}$$

Theorem 4.4.12 (LDLR norm for Gaussian Wishart model). *Suppose (\mathbb{P}_n) and (\mathbb{Q}_n) are a Gaussian Wishart model as in Definition 4.4.1, with signal prior $(\tilde{\mathcal{P}}_n)$ and sampling ratio $\gamma > 0$, and let $L_n^{\leq D}$ be the low-degree likelihood ratio (Definition 3.2.1). Then, for $D \in \mathbb{N} \cup \{+\infty\}$,*

$$\|L_n^{\leq D}\|^2 = \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \sim \tilde{\mathcal{P}}_n} \phi_{N,D}^{\text{Wish}}(\tilde{\mathbf{X}}^1 \tilde{\mathbf{X}}^2). \quad (4.74)$$

Proof. For $\mathbf{k} \in (\mathbb{N}^n)^N$, we denote

$$\mathbf{m}_{\mathbf{k}} := \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} H_{\mathbf{k}}(\mathbf{Y})L(\mathbf{Y}), \quad (4.75)$$

$$\widehat{\mathbf{m}}_{\mathbf{k}} := \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} \widehat{H}_{\mathbf{k}}(\mathbf{Y})L(\mathbf{Y}) = \frac{1}{\sqrt{\mathbf{k}!}} \mathbf{m}_{\mathbf{k}}. \quad (4.76)$$

We may compute these numbers as follows. For any $\mathbf{k} \in (\mathbb{N}^n)^N$, we have, passing to an expectation under \mathbb{P}_n rather than \mathbb{Q}_n and then using Proposition 4.4.9,

$$\begin{aligned} \mathbf{m}_{\mathbf{k}} &= \mathbb{E}_{\mathbf{Y} \sim \mathbb{P}} H_{\mathbf{k}}(\mathbf{Y}) \\ &= \mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \prod_{i=1}^N \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, \mathbf{I} + \tilde{\mathbf{X}})} H_{\mathbf{k}_i}(\mathbf{y}) \\ &= \mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \prod_{i=1}^N \left(\sum_{M \in \text{Part}(\mathbf{k}_i; 2)} \tilde{\mathbf{X}}^M \right). \end{aligned} \quad (4.77)$$

We then have, applying the usual replica manipulation,

$$\begin{aligned} \|L_n^{\leq D}\|^2 &= \sum_{\substack{\mathbf{k} \in (\mathbb{N}^n)^N \\ |\mathbf{k}| \leq D}} \widehat{\mathbf{m}}_{\mathbf{k}}^2 \\ &= \sum_{\substack{\mathbf{k} \in 2(\mathbb{N}^n)^N \\ |\mathbf{k}| \leq D}} \prod_{i=1}^N \frac{1}{\mathbf{k}_i!} \left(\mathbb{E}_{\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n} \sum_{M \in \mathcal{M}(\mathbf{k}_i)} \tilde{\mathbf{X}}^M \right)^2 \quad (\text{by (4.77)}) \\ &= \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \sim \tilde{\mathcal{P}}_n} \sum_{\substack{\mathbf{k} \in (\mathbb{N}^n)^N \\ |\mathbf{k}| \leq D}} \prod_{i=1}^N \frac{1}{\mathbf{k}_i!} \left(\sum_{M \in \mathcal{M}(\mathbf{k}_i)} (\tilde{\mathbf{X}}^1)^M \right) \left(\sum_{M \in \mathcal{M}(\mathbf{k}_i)} (\tilde{\mathbf{X}}^2)^M \right) \quad (\text{Proposition 4.1.4}) \end{aligned}$$

And, since if any $|\mathbf{k}_i|$ is odd then $\text{Part}(\mathbf{k}_i; 2)$ is empty and the corresponding term is zero, we may restrict

$$= \mathbb{E}_{\tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \sim \tilde{\mathcal{P}}_n} \sum_{d=0}^{\lfloor D/2 \rfloor} \sum_{\substack{\mathbf{k} \in 2(\mathbb{N}^n)^N \\ |\mathbf{k}|=2d}} \prod_{i=1}^N \frac{1}{\mathbf{k}_i!} \left(\sum_{M \in \mathcal{M}(\mathbf{k}_i)} (\tilde{\mathbf{X}}^1)^M \right) \left(\sum_{M \in \mathcal{M}(\mathbf{k}_i)} (\tilde{\mathbf{X}}^2)^M \right). \quad (4.78)$$

Thus we will be finished if we can show the identity

$$r_{N,d}(\mathbf{A}\mathbf{B}) = \sum_{\substack{\mathbf{k} \in (\mathbb{N}^n)^N \\ |\mathbf{k}|=2d}} \prod_{i=1}^N \frac{1}{\mathbf{k}_i!} \left(\sum_{M \in \text{Part}(\mathbf{k}_i; 2)} \mathbf{A}^M \right) \left(\sum_{M \in \text{Part}(\mathbf{k}_i; 2)} \mathbf{B}^M \right) =: s_{N,d}(\mathbf{A}, \mathbf{B}), \quad (4.79)$$

which will occupy the rest of the proof. We have $s_{N,d}(\mathbf{A}, \mathbf{B}) = 0$ for all odd d since in this case some $|\mathbf{k}_i|$ must be odd, whereby $\text{Part}(\mathbf{k}_i; 2) = \emptyset$. It remains to show that, for even d , $r_{N,d/2}(\mathbf{A}\mathbf{B}) = s_{N,d}(\mathbf{A}, \mathbf{B})$. Since either side is a polynomial, it suffices to show this for (\mathbf{A}, \mathbf{B}) belonging to a subset of $(\mathbb{R}_{\text{sym}}^{n \times n})^2$ having positive measure. We will show that the equality holds when $\mathbf{A}, \mathbf{B} \geq \mathbf{0}$ and $\|\mathbf{A}\|, \|\mathbf{B}\| < 1$.

We proceed by computing the (ordinary) generating functions of either collection of polynomials. We already know from Proposition 4.4.5 that

$$R(t) := \sum_{d=0}^{\infty} r_{N,d}(\mathbf{A}\mathbf{B}) t^{2d} = \det(\mathbf{I}_n - t^2 \mathbf{A}\mathbf{B})^{-N/2} \quad (4.80)$$

for $|t| < 1$. Thus it suffices to show the same equality for

$$S(t) := \sum_{d=0}^{\infty} s_{N,d}(\mathbf{A}, \mathbf{B}) t^d. \quad (4.81)$$

We note that, under our assumptions, by Wick's formula

$$\begin{aligned}
s_{N,d}(\mathbf{A}, \mathbf{B}) &= \sum_{\substack{\mathbf{k} \in (\mathbb{N}^n)^N \\ |\mathbf{k}|=d}} \prod_{i=1}^N \frac{1}{\mathbf{k}_i!} \left(\mathbb{E}_{g \sim \mathcal{N}(\mathbf{0}, \mathbf{A})} \mathbf{g}^{\mathbf{k}_i} \right) \left(\mathbb{E}_{h \sim \mathcal{N}(\mathbf{0}, \mathbf{B})} \mathbf{h}^{\mathbf{k}_i} \right) \\
&= \mathbb{E}_{\substack{\mathbf{g}_1, \dots, \mathbf{g}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{A}) \\ \mathbf{h}_1, \dots, \mathbf{h}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{B})}} \sum_{\substack{\mathbf{k} \in (\mathbb{N}^n)^N \\ |\mathbf{k}|=d \\ |\mathbf{k}_i| \text{ even for all } i \in [N]}} \prod_{i=1}^N \frac{1}{\mathbf{k}_i!} (\mathbf{g}_i)^{\mathbf{k}_i} (\mathbf{h}_i)^{\mathbf{k}_i}
\end{aligned}$$

and, grouping by the values of $|\mathbf{k}_i|$,

$$= \mathbb{E}_{\substack{\mathbf{g}_1, \dots, \mathbf{g}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{A}) \\ \mathbf{h}_1, \dots, \mathbf{h}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{B})}} \sum_{\substack{d_1, \dots, d_N \in 2\mathbb{N} \\ \sum_{i=1}^N d_i = d}} \frac{1}{\prod_{i=1}^N d_i!} \sum_{\substack{\mathbf{k}_1, \dots, \mathbf{k}_N \in \mathbb{N}^n \\ |\mathbf{k}_i| = d_i}} \prod_{i=1}^N \binom{d_i}{\mathbf{k}_i} (\mathbf{g}_i)^{\mathbf{k}_i} (\mathbf{h}_i)^{\mathbf{k}_i}$$

so that by the multinomial theorem,

$$\begin{aligned}
&= \mathbb{E}_{\substack{\mathbf{g}_1, \dots, \mathbf{g}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{A}) \\ \mathbf{h}_1, \dots, \mathbf{h}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{B})}} \sum_{\substack{d_1, \dots, d_N \in 2\mathbb{N} \\ \sum_{i=1}^N d_i = d}} \prod_{i=1}^N \frac{\langle \mathbf{g}_i, \mathbf{h}_i \rangle^{d_i}}{d_i!} \\
&= \sum_{\substack{d_1, \dots, d_N \in 2\mathbb{N} \\ \sum_{i=1}^N d_i = d}} \prod_{i=1}^N \frac{1}{d_i!} \mathbb{E}_{\substack{g \sim \mathcal{N}(\mathbf{0}, \mathbf{A}) \\ h \sim \mathcal{N}(\mathbf{0}, \mathbf{B})}} \langle \mathbf{g}, \mathbf{h} \rangle^{d_i} \tag{4.82}
\end{aligned}$$

We introduce the moment-generating function of the remaining expectations:

$$f(t) := \sum_{d=0}^{\infty} \frac{t^d}{d!} \mathbb{E}_{\substack{g \sim \mathcal{N}(\mathbf{0}, \mathbf{A}) \\ h \sim \mathcal{N}(\mathbf{0}, \mathbf{B})}} \langle \mathbf{g}, \mathbf{h} \rangle^d = \mathbb{E}_{\substack{g \sim \mathcal{N}(\mathbf{0}, \mathbf{A}) \\ h \sim \mathcal{N}(\mathbf{0}, \mathbf{B})}} \exp(t \langle \mathbf{g}, \mathbf{h} \rangle). \tag{4.83}$$

Then, $s_{N,d}(\mathbf{A}, \mathbf{B})$ above is simply the coefficient of t^d in the Taylor series of $f(t)^N$, and thus $S(t) = f(t)^N$. So, it will suffice for us to establish that

$$f(t) \stackrel{?}{=} \det(\mathbf{I}_n - t^2 \mathbf{A} \mathbf{B})^{-1/2}. \tag{4.84}$$

However, $f(t)$ may be evaluated with a direct computation:

$$\begin{aligned} f(t) &= \mathbb{E}_{g, h \sim \mathcal{N}(0, I_n)} \exp \left(t \mathbf{g}^\top \sqrt{\mathbf{A}} \sqrt{\mathbf{B}} \mathbf{h} \right) \\ &= \mathbb{E}_{g \sim \mathcal{N}(0, I_{2n})} \exp \left(\mathbf{g}^\top \begin{bmatrix} \mathbf{0} & \frac{t}{2} \sqrt{\mathbf{A}} \sqrt{\mathbf{B}} \\ \frac{t}{2} \sqrt{\mathbf{B}} \sqrt{\mathbf{A}} & \mathbf{0} \end{bmatrix} \mathbf{g} \right) \end{aligned}$$

and applying an orthogonal change of basis to diagonalize this matrix and evaluating the χ^2 moment generating function then gives

$$\begin{aligned} &= \det \left(\begin{bmatrix} \mathbf{I}_n & -t \sqrt{\mathbf{A}} \sqrt{\mathbf{B}} \\ -t \sqrt{\mathbf{B}} \sqrt{\mathbf{A}} & \mathbf{I}_n \end{bmatrix} \right)^{-1/2} \\ &= \det \left(\mathbf{I}_n - t^2 \sqrt{\mathbf{A}} \mathbf{B} \sqrt{\mathbf{A}} \right)^{-1/2} \end{aligned}$$

and by Sylvester's identity,

$$= \det \left(\mathbf{I}_n - t^2 \mathbf{A} \mathbf{B} \right)^{-1/2}, \quad (4.85)$$

which completes the proof. \square

Lastly, we make two remarks that are specific to the Wishart spiked matrix model, where the signal is sign-definite. First, we observe that, for the purposes of bounding the LDLR norm, we may ignore the ‘‘truncation clause’’ of the spiked matrix model, as this only makes the norm smaller.

Proposition 4.4.13. *In the Wishart spiked matrix model with parameters $(\beta, \gamma, \mathcal{P}_n)$ for \mathcal{P}_n not necessarily β -good (Definition 2.2.5), for any (finite) $D \in \mathbb{N}$,*

$$\|L_n^{\leq D}\|^2 \leq \mathbb{E}_{\mathbf{X}^1, \mathbf{X}^2 \sim \mathcal{P}_n} \phi_{N, D}^{\text{Wish}}(\beta^2 \mathbf{X}^1 \mathbf{X}^{1\top} \mathbf{X}^2 \mathbf{X}^{2\top}). \quad (4.86)$$

Proof. For $\mathbf{X}^i \sim \mathcal{P}_n$ independently for $i \in \{1, 2\}$, let $\beta \mathbf{A}^i$ be the “effective spikes” in the Wishart spiked matrix model,

$$\mathbf{A}^i = \begin{cases} \mathbf{X}^i \mathbf{X}^{i\top} & \text{if } \beta \|\mathbf{X}^i\|^2 > -1, \\ \mathbf{0} & \text{otherwise.} \end{cases} \quad (4.87)$$

In either case, $\mathbf{A}^i \geq \mathbf{0}$. Then, by Theorem 4.4.12 and using the cyclic property of the link functions, we have

$$\begin{aligned} \|L_n^{\leq D}\|^2 &= \mathbb{E} \phi_{N,D}^{\text{Wish}}(\beta^2 \mathbf{A}^1 \mathbf{A}^2) \\ &= \mathbb{E} \phi_{N,D}^{\text{Wish}}(\beta^2 (\sqrt{\mathbf{A}^1} \sqrt{\mathbf{A}^2}) (\sqrt{\mathbf{A}^1} \sqrt{\mathbf{A}^2})^\top) \end{aligned} \quad (\text{Proposition 4.4.6})$$

Here, the argument of the link function is zero if either \mathbf{A}^i is zero, and is always positive semidefinite. Since the $r_{N,d}$ are non-negative on positive semidefinite matrices and zero on the zero matrix, the link function has the same property. In particular, the link function does not decrease from replacing a $\mathbf{0}$ input with any positive semidefinite matrix, so

$$\begin{aligned} &\leq \mathbb{E} \phi_{N,D}^{\text{Wish}}(\beta^2 (\sqrt{\mathbf{X}^1 \mathbf{X}^{1\top}} \sqrt{\mathbf{X}^2 \mathbf{X}^{2\top}}) (\sqrt{\mathbf{X}^1 \mathbf{X}^{1\top}} \sqrt{\mathbf{X}^2 \mathbf{X}^{2\top}})^\top) \\ &= \mathbb{E} \phi_{N,D}^{\text{Wish}}(\beta^2 \mathbf{X}^1 \mathbf{X}^{1\top} \mathbf{X}^2 \mathbf{X}^{2\top}), \end{aligned} \quad (\text{Proposition 4.4.6})$$

completing the proof. □

Second, to illustrate the result above in a more familiar setting, we consider the special case when $\tilde{\mathcal{P}}_n$ is supported on rank-one matrices, for which this result simplifies substantially to a scalar formula similar to those we have seen in other models. These results were also derived in [BKW20b] more directly before the general framework presented here was developed in [BBK⁺20].

Corollary 4.4.14. *In the Wishart spiked matrix model with rank one and parameters $(\beta, \gamma, \mathcal{P}_n)$*

for $\mathcal{P}_n \in \mathcal{P}(\mathbb{R}^n)$ a β -good spike prior, for any $D \in \mathbb{N} \cup \{+\infty\}$,

$$\|L_n^{\leq D}\|^2 = \mathbb{E}_{\mathbf{x}^1, \mathbf{x}^2 \sim \mathcal{P}_n} \phi_{N,D}^{\text{Wish}}(\beta^2 \langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2) = \mathbb{E}_{\mathbf{x}^1, \mathbf{x}^2 \sim \mathcal{P}_n} \sum_{d=0}^{\lfloor D/2 \rfloor} a_{N,d} (\beta^2 \langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2)^d, \quad (4.88)$$

where the coefficients are

$$a_{N,d} = \binom{N/2 + d - 1}{d} = \frac{1}{d!} \prod_{k=0}^{d-1} \left(\frac{N}{2} + k \right), \quad (4.89)$$

with the binomial coefficient interpreted in the generalized manner of the final right-hand side even when $N/2$ is not an integer. These coefficients satisfy

$$\sum_{d=0}^{\infty} a_{N,d} t^d = (1 - t)^{-N/2}. \quad (4.90)$$

We note that the original argument in [BKW20b] did not take advantage of the simple expression (4.89) for the coefficients appearing here, which we will see in Section 5.2.1 dramatically simplifies the treatment of Wishart spiked matrix models, completely reducing them to Wigner models. We note also that by Proposition 4.4.13, even for a prior that is not necessarily β -good, (4.88) holds as an inequality for finite D .

Proof. Let $\tilde{\mathbf{X}}^i \sim \tilde{\mathcal{P}}_n$ independently be formed as $\tilde{\mathbf{X}}^i = \beta \mathbf{x}^i \mathbf{x}^{i\top}$ with $\mathbf{x}^i \sim \mathcal{P}_n$ independently for $i \in \{1, 2\}$. Then, by the cyclic property from Proposition 4.4.6, for any N, d we have

$$r_{N,d}(\tilde{\mathbf{X}}^1 \tilde{\mathbf{X}}^2) = r_{N,d}(\beta^2 \mathbf{x}^1 \mathbf{x}^{1\top} \mathbf{x}^2 \mathbf{x}^{2\top}) = r_{N,d}(\beta^2 \langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2). \quad (4.91)$$

Thus it suffices to show that $r_{N,d}(t^2) = a_{N,d} t^{2d}$ for $t \in \mathbb{R}$. The generating function (4.90) for these coefficients follows from the series expression of Proposition 4.4.5. The formula (4.89) then follows from the generalized binomial theorem, which gives $a_{N,d} = (-1)^d \binom{-N/2}{d}$, which, upon expanding and cancelling signs, gives the stated product. \square

5 | LOW-DEGREE LIKELIHOOD RATIO

ANALYSIS: LOWER BOUNDS

In this chapter, we will reap the benefits of the machinery developed in the rest of Part I to prove lower bounds for the hypothesis testing problems to which certification for various constrained PCA problems reduces. More broadly, we will also provide streamlined and sometimes improved results on low-degree lower bounds for related problems, including Wigner and Wishart matrix PCA (i.e., spiked matrix models), tensor PCA, and the stochastic block model. We will see how, using the tools we have developed and some further tools we introduce below for analyzing overlap formulae for the norm of the low-degree likelihood ratio, all of these problems may be treated within one framework. Our approach also mostly avoids the explicit moment computations, often having a combinatorial character, common to these analyses carried out in previous literature. The applications of the results for Wishart models to constrained PCA are given in Section 5.3.

SUMMARY AND REFERENCES The results in this chapter are taken mostly from the same publications as the previous chapter; we have only reorganized the material to group the overlap formulae together and their consequences for lower bounds together. However, this chapter also describes some proof techniques that have not yet appeared in print, including a low-degree comparison inequality between Wigner and Wishart models (Section 5.2.1) that

dramatically simplifies the analysis of Wishart models, and the generic approach to Wigner models through integrated tail bounds rather than moment computations (used throughout). Similar applications to Gaussian Wigner models appeared in [KWB19], to sparse PCA in [DKWB19], and the applications to Wishart models are streamlined versions of the arguments of [BKW20b, BBK⁺20], which also respectively treated the applications to certification for the SK and Potts spin glass Hamiltonians. The application to certification for non-negative PCA is based on [BKW20a]. The computations for the stochastic block model (which sharpen previous computations appearing in [HS17]) also appeared in [BBK⁺20], while the non-Gaussian spiked matrix model of Section 5.4.2 appeared in [Kun20a] and is inspired by the work of [PWBM18] on similar models.

The following is a summary of our main results in this chapter.

1. (Theorem 5.2.6) A tight low-degree lower bound for tensor PCA, agreeing with previous results on the performance of conjecturally-optimal algorithms [RM14, HSS15, ADGM17, HSS16], SOS lower bounds [HSS15, HKP⁺17], and results on subexponential-time algorithms [BGG⁺16, BGL16, WEM19].
2. (Theorem 5.2.10, Theorem 5.2.12) Tight low-degree lower bounds for the Wigner and Wishart spiked matrix models, giving strong evidence for the optimality of PCA among all subexponential-time algorithms in these Gaussian models (for example, giving further evidence for conjectures on constant-sparsity PCA discussed in Example 2.4.1).
3. (Theorem 5.2.14, Theorem 5.2.15) A tight low-degree lower bound for Wigner and Wishart sparse PCA (with sub-constant sparsity, unlike the above), matching the behavior of similar sub-exponential time algorithms proposed by [DKWB19, HSV20].
4. (Corollary 5.3.6) A conditional result that better-than-spectral certification is hard for the SK Hamiltonian, giving evidence towards an answer to a question of Montanari and

Jain, Koehler, and Risteski [JKR19].¹

5. (Corollary 5.3.7) A conditional result that better-than-spectral certification is hard for non-negative PCA on GOE matrices, giving evidence towards an answer to a question of [MR15].
6. (Corollary 5.3.8) A conditional result that better-than-spectral certification is hard for the Potts spin glass Hamiltonian.

5.1 TOOLS FOR SCALAR OVERLAP ANALYSIS

We have seen that, in many cases of interest, a bound (and indeed, often a formula with equality) of the following form holds for the norm of the LDLR:

$$\|L_n^{\leq D}\|^2 \leq \mathbb{E}[p_{n,D}(R_n)]. \quad (5.1)$$

We now take up the technical matter of how to bound expectations such as those on the right-hand side above in the situations we will be interested in. We decompose this task into three parts: first, an inequality that decouples the dependence of this quantity on the growth of $p_{n,D}(t)$ and that on the tail decay of R_n ; second, a notion of “modest growth” for $p_{n,D}(t)$; and third, a notion of “fast decay” for R_n . Moreover, these latter analytic notions turn out to have the following piecewise character. For R_n we will be concerned separately with “small deviations” and “large deviations,” following the style of computation of [PWB16, PWBM18]. For $p_{n,D}(t)$, as these will always be truncations of a Taylor series in our applications, we will be concerned separately with small inputs t where $p_{n,D}$ accurately approximates the entire series, which grows exponentially in t , and large inputs where its more modest polynomial $O(t^D)$ growth becomes evident.

¹“Hard” in this result and the below means requiring nearly-exponential time.

It is tempting to attempt to give general conditions on R_n and $p_{n,D}$ that ensure boundedness of the low-degree likelihood ratio. However, we will mostly be concerned with only two different observation models, the Gaussian Wigner and Wishart models, and even between these two models we will see shortly that the Wigner model controls the Wishart model (and reduces the higher-rank Wishart models where it may seem on the basis of Section 4.4 that we must deal with overlap matrices to simpler scalar problems). Thus we illustrate the basic themes of these computations on a case-by-case basis; developing abstract statements that work for many distributions and associated polynomial sequences “automatically” is an interesting project for future investigation.

Our decoupling inequality, the most generic of the above ingredients, is the following simple but general integrated tail bound.

Lemma 5.1.1. *Suppose that $p \in \mathbb{R}[t]$ has non-negative coefficients and $\phi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is strictly increasing. Suppose also that $X \geq 0$ is a random variable satisfying the tail bound that, for all $t \geq 0$, $\mathbb{P}[X \geq t] \leq f(t)$. Then,*

$$\mathbb{E}[p(\phi(X))] \leq \int_0^\infty f(\phi^{-1}(t))p'(t)dt. \quad (5.2)$$

Proof. Suppose that $p(t) = \sum_{d=0}^D a_d t^d$. We expand directly,

$$\begin{aligned} \mathbb{E}[g(\phi(X))] &= \sum_{d=0}^D a_d \mathbb{E}\phi(X)^d \\ &= \sum_{d=0}^D a_d \int_0^\infty \mathbb{P}[\phi(X)^d \geq t] dt \end{aligned}$$

and, applying the tail bound,

$$\leq \sum_{d=0}^D a_d \int_0^\infty f(\phi^{-1}(t^{1/d})) dt$$

and making the substitution $s = t^{1/d}$ we find

$$\begin{aligned}
&= \sum_{d=0}^D da_d \int_0^\infty f(\phi^{-1}(s)) s^{d-1} ds \\
&= \int_0^\infty f(\phi^{-1}(s)) p'(s) ds
\end{aligned} \tag{5.3}$$

as desired. □

5.1.1 WELL-BEHAVED SPIKE PRIORS

We also introduce some definitions that will be useful for specifying spike priors that yield well-behaved overlaps R_n . We always think of spike variables $\mathbf{x} \sim \mathcal{P}_n$ in rank-one spiked models (both the Wigner and Wishart spiked matrix models and the tensor models we will discuss in Section 4.4) as having norm approximately 1, so that the typical magnitude of the overlap is, for a dense prior, $|\langle \mathbf{x}^1, \mathbf{x}^2 \rangle| \sim n^{-1/2}$. Likewise, when the rank is $k \geq 1$, we think of $\mathbf{X} \sim \mathcal{P}_n$ as having k columns each of norm approximately 1. Below we work with the arbitrary-rank setting, so we use matrix notation, but the same notions apply to vector spikes.

We first recall the following standard tail bound property.

Definition 5.1.2 (Subgaussian). *A centered random variable R is σ^2 -subgaussian if, for all $t \in \mathbb{R}$, $\mathbb{E}[\exp(tR)] \leq \exp(\frac{\sigma^2}{2}t^2)$.*

Proposition 5.1.3. *Let R be a centered random variable. Then, the following conditions are equivalent.*

1. R is σ^2 -subgaussian for some $\sigma^2 > 0$.
2. There exist $C, \epsilon > 0$ such that $\mathbb{P}[|R| \geq t] \leq C \exp(-\epsilon t^2)$ for all $t \geq 0$.
3. There exists $\epsilon > 0$ such that $\mathbb{E} \exp(\epsilon R^2) < \infty$.

Proof. Clearly (2) implies (3) by integrating the tail bound. Conversely, (3) implies (2) by Chernoff bound, since

$$\mathbb{P}[|R| \geq t] = \mathbb{P}[\exp(sR^2) \geq \exp(st^2)] \leq (\mathbb{E} \exp(sR^2)) \exp(-st^2) \quad (5.4)$$

for any $s < \epsilon$. Thus (2) and (3) are equivalent.

Moreover, by the same argument as above but choosing s sufficiently small, (3) implies a stronger version of (2) with the constant C arbitrarily close to 1, since $\lim_{s \rightarrow 0} \mathbb{E} \exp(sR^2) = 1$ so long as this expectation is finite for some positive s . With $C = 2$, (2) implies (1) by a slightly involved argument of bounding moments; see, e.g., Lemma 1.5 of [RH17] for this computation. Finally, (1) implies (2) by another Chernoff bound, completing the proof. \square

Subgaussianity will be a useful assumption later, but most immediately in spike priors we will work with the following weaker notion.

Definition 5.1.4 (Locally subgaussian). *A sequence of random variables $R_n \geq 0$ is locally subgaussian with speed $\rho(n) \leq 1$ if, for any $\eta > 0$, there exists $\delta > 0$ such that, for all $t \in [0, \delta\rho(n)]$,*

$$\mathbb{P}[R_n \geq t] \leq \exp\left(-n \frac{1-\eta}{2} t^2\right). \quad (5.5)$$

When $\rho(n)$ is not specified, we take $\rho(n) = 1$. A spike prior (\mathcal{P}_n) is locally subgaussian if the sequence $\|\mathbf{X}^{1\top} \mathbf{X}^2\|_F$ is locally subgaussian for $\mathbf{X}^i \sim \mathcal{P}_n$ independently.

Often this definition has been given in other works with $2\sigma^2$ in the rate denominator, giving a notion of “locally subgaussian with variance proxy σ^2 ” by analogy with the usual “global” notion of subgaussianity (Definition 5.1.2), but such factors can always be absorbed into the spike prior itself in our setting, so we work with this simpler and stricter definition. We use the notation $\rho(n)$ here and below because this will turn out to correspond to the sparsity of a prior, so that $\rho(n) \cdot n$ is the typical number of non-zero entries.

The assumption on spike priors we will usually work with combines local subgaussianity controlling small deviations with a modest bound on large deviations.

Definition 5.1.5 (Tame prior). *A spike prior (\mathcal{P}_n) is tame with speed $\rho(n) \leq 1$ if there exist $\xi, T > 0$ such that, for any $\eta > 0$, there exists a further $\delta > 0$ such that*

$$\mathcal{P}_n^{\otimes 2} \left[\|\mathbf{X}^{\top} \mathbf{X}^2\|_F \geq t \right] \lesssim \exp \left(-n \cdot \left\{ \begin{array}{ll} \frac{1-\eta}{2} t^2 & \text{if } 0 \leq t \leq \delta\rho, \\ \rho^2 \frac{1-\eta}{2} \delta^2 & \text{if } \delta\rho \leq t \leq T\rho, \\ \rho^2 (\log(1+t))^{1+\xi} & \text{if } t \geq T\rho. \end{array} \right. \right). \quad (5.6)$$

In this case, we call T the upper threshold. As above, when $\rho(n)$ is not specified, we take $\rho(n) = 1$.

The middle condition for $t \in [\delta\rho, T\rho]$ follows from the bound at $t = \delta\rho$, so this is really only a combination of local subgaussianity and the bound on $t \geq T\rho$. We also note that, in the dense case $\rho(n) = 1$, one simple way for a spike prior to satisfy the large deviations condition of being tame is for it to be bounded in norm.

Lastly, we give some broad classes of natural examples of tame priors. All of our examples will be drawn from priors having at least the following independence structure.

Definition 5.1.6 (Priors with i.i.d. rows). *Let $\pi \in \mathcal{P}(\mathbb{R}^k)$. Then, we denote by $\mathcal{P}_n^\pi \in \mathcal{P}(\mathbb{R}^{n \times k})$ the prior formed by sampling $\mathbf{X} \sim \mathcal{P}_n^\pi$ as having i.i.d. rows $\hat{\mathbf{r}}_1, \dots, \hat{\mathbf{r}}_n$, with $\hat{\mathbf{r}}_k = \frac{1}{\sqrt{n}} \mathbf{r}_k$ for $\mathbf{r}_k \sim \pi$ independently.*

Below we will use the fact that a centered random variable X is subgaussian with some variance proxy if and only if $\mathbb{E} \exp(\epsilon X^2) < \infty$ for some $\epsilon > 0$; see Proposition 5.1.3.

Proposition 5.1.7. *Let $\pi \in \mathcal{P}(\mathbb{R}^k)$ have $\mathbb{E}_{v \sim \pi} \mathbf{v} = \mathbf{0}$, $\mathbb{E}_{v \sim \pi} \mathbf{v} \mathbf{v}^\top \leq \mathbf{I}_k$, and $\|\mathbf{v}\|$ subgaussian with any variance proxy. Then, (\mathcal{P}_n^π) is a tame spike prior (with speed $\rho(n) = 1$).*

Proof. To lighten the notation, let us write $\mathcal{P}_n = \mathcal{P}_n^\pi$ for this proof. Suppose $\mathbf{X}^1, \mathbf{X}^2 \sim \mathcal{P}_n$ independently, and let $\mathbf{r}_1^1, \dots, \mathbf{r}_n^1 \sim \pi$ and $\mathbf{r}_1^2, \dots, \mathbf{r}_n^2 \sim \pi$, all independently, be the rows of $\sqrt{n}\mathbf{X}^1$ and $\sqrt{n}\mathbf{X}^2$, respectively. Then, we note that

$$\|\mathbf{X}^{1\top} \mathbf{X}^2\|_F = \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{r}_i^1 \mathbf{r}_i^{2\top} \right\|_F = \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{r}_i^1 \otimes \mathbf{r}_i^2 \right\|_F. \quad (5.7)$$

Since $\|\mathbf{r}_i\|$ is subgaussian, $\mathbb{E} \exp(\epsilon \|\mathbf{r}_i\|^2) < \infty$ for some $\epsilon > 0$. Therefore, the moment generating function of the law of $\mathbf{r}_i^1 \otimes \mathbf{r}_i^2$, i.e., the function $M : \mathbb{R}^{k^2} \rightarrow \mathbb{R}$ defined by $M(\mathbf{u}) := \mathbb{E}_{\mathbf{v}^1, \mathbf{v}^2 \sim \pi} \exp(\langle \mathbf{u}, \mathbf{v}^1 \otimes \mathbf{v}^2 \rangle)$, exists on the ball of radius ϵ centered at the origin, since $\langle \mathbf{u}, \mathbf{v}^1 \otimes \mathbf{v}^2 \rangle \leq \|\mathbf{u}\| \cdot \|\mathbf{v}^1 \otimes \mathbf{v}^2\| = \|\mathbf{u}\| \cdot \|\mathbf{v}^1\| \cdot \|\mathbf{v}^2\| \leq \frac{1}{2} \|\mathbf{u}\| (\|\mathbf{v}^1\|^2 + \|\mathbf{v}^2\|^2)$. Since this law is centered and has covariance matrix $\leq \mathbf{I}_{k^2}$ (as this is the tensor square of the covariance matrix of π), we have $\nabla M(\mathbf{0}) = \mathbf{0}$ and $\nabla^2 M(\mathbf{0}) \leq \mathbf{I}_{k^2}$. Thus we have that, for any $\eta > 0$, there exists $\delta > 0$ such that, whenever $\|\mathbf{u}\| \leq \delta$, then

$$M(\mathbf{u}) \leq \exp\left(\frac{1}{2\sqrt{1-\eta}} \|\mathbf{u}\|^2\right), \quad (5.8)$$

by comparing the second-order Taylor expansions of either side.

Then, the moment generating function of the sum above is bounded by

$$M^{(n)}(\mathbf{u}) := \mathbb{E}_{\mathbf{r}_i^1, \mathbf{r}_i^2} \exp\left(\left\langle \mathbf{u}, \frac{1}{n} \sum_{i=1}^n \mathbf{r}_i^1 \otimes \mathbf{r}_i^2 \right\rangle\right) = M\left(\frac{1}{n} \mathbf{u}\right)^n \leq \exp\left(\frac{1}{n} \frac{1}{2\sqrt{1-\eta}} \|\mathbf{u}\|^2\right) \quad (5.9)$$

whenever $\|\mathbf{u}\| \leq \delta n$. Accordingly, taking a Chernoff bound, for any $\|\mathbf{u}\| \leq 1$ and $t \leq \delta$, we have

$$\begin{aligned} \mathcal{P}_n^{\otimes 2} \left[\left\langle \mathbf{u}, \frac{1}{n} \sum_{i=1}^n \mathbf{r}_i^1 \otimes \mathbf{r}_i^2 \right\rangle \geq t \right] &= \mathcal{P}_n^{\otimes 2} \left[\left\langle \sqrt{1-\eta} n t \mathbf{u}, \frac{1}{n} \sum_{i=1}^n \mathbf{r}_i^1 \otimes \mathbf{r}_i^2 \right\rangle \geq \sqrt{1-\eta} n t^2 \right] \\ &\leq \exp\left(-n \frac{\sqrt{1-\eta}}{2} t^2\right). \end{aligned} \quad (5.10)$$

Now, let \mathcal{N} be a net of \mathbb{S}^{k^2-1} , such that for any $\mathbf{v} \in \mathbb{S}^{k^2-1}$ there exists $\mathbf{u} \in \mathcal{N}$ with $\langle \mathbf{u}, \mathbf{v} \rangle \geq (1 - \eta)^{1/4}$. We then have, for any $t \leq \delta$, by union bound,

$$\begin{aligned} \mathcal{P}_n^{\otimes 2} \left[\|\mathbf{X}^{1\top} \mathbf{X}^2\|_F \geq t \right] &= \mathcal{P}_n^{\otimes 2} \left[\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{r}_i^1 \otimes \mathbf{r}_i^2 \right\| \geq t \right] \\ &\leq \mathcal{P}_n^{\otimes 2} \left[\max_{\mathbf{u} \in \mathcal{N}} \left\langle \mathbf{u}, \frac{1}{n} \sum_{i=1}^n \mathbf{r}_i^1 \otimes \mathbf{r}_i^2 \right\rangle \geq (1 - \eta)^{1/4} t \right] \\ &\leq |\mathcal{N}| \exp \left(-n \frac{1 - \eta}{2} t^2 \right), \end{aligned} \quad (5.11)$$

giving the local subgaussianity condition (since $|\mathcal{N}|$ may be taken to depend only on k and η and thus subsumed into the leading constant).

For the bound on large deviations, we note that, since $\|\mathbf{X}^{1\top} \mathbf{X}^2\|_F \leq \|\mathbf{X}^1\|_F \|\mathbf{X}^2\|_F$, if $\|\mathbf{X}^{1\top} \mathbf{X}^2\|_F \geq t$ then either $\|\mathbf{X}^1\|_F^2 \geq t$ or $\|\mathbf{X}^2\|_F^2 \geq t$. Therefore, by union bound,

$$\begin{aligned} \mathcal{P}_n^{\otimes 2} \left[\|\mathbf{X}^{1\top} \mathbf{X}^2\|_F \geq t \right] &\leq 2\mathcal{P}_n \left[\|\mathbf{X}\|^2 \geq t \right] \\ &= 2\mathcal{P}_n \left[\frac{1}{n} \sum_{i=1}^n \|\mathbf{r}_i\|^2 \geq t \right] \end{aligned}$$

and by Chernoff bound, for $C = \mathbb{E}_{\mathbf{v} \sim \pi} \exp(\epsilon \|\mathbf{v}\|^2)$,

$$\begin{aligned} &\leq 2 \left(\mathbb{E}_{\mathbf{r} \sim \pi} \exp(\epsilon \|\mathbf{r}\|^2) \right)^n \exp(-\epsilon n t) \\ &\lesssim \exp(-n(\epsilon t - \log C)), \end{aligned} \quad (5.12)$$

which gives the required bound for upper threshold T sufficiently large. \square

As a special case, this automatically treats the simpler situation of priors with i.i.d. entries.

Corollary 5.1.8. *Let $\pi \in \mathcal{P}(\mathbb{R})$ have mean zero, variance at most 1, and be subgaussian with any variance proxy. Then, for any $k \geq 1$, $(\mathcal{P}_n^{(\pi^{\otimes k})})$ is a tame spike prior (with speed $\rho(n) = 1$).*

Finally, we also consider the analogous situation for sparse priors, where we use the case

$\rho(n) \ll 1$ of the definitions of tameness and subgaussianity. We restrict our attention to the rank-one case, as it is unclear how to most naturally extend the notion below to “sparsifying” priors with i.i.d. rows.

Definition 5.1.9 (Sparse priors with i.i.d. entries). *Let $\pi \in \mathcal{P}(\mathbb{R})$ and $\rho(n) \leq 1$. Then, we denote by $\mathcal{P}_n^{\pi, \rho(n)}$ the spike prior formed by sampling $x \sim \mathcal{P}_n^{\pi, \rho(n)}$ as having $x_i = \frac{1}{\sqrt{\rho(n) \cdot n}} w_i s_i$, for $w_i \sim \pi$ and $s_i \sim \text{Ber}(\rho(n))$ all drawn independently for $i \in [n]$.*

Proposition 5.1.10. *Let $\pi \in \mathcal{P}(\mathbb{R})$ have mean zero, variance at most 1, and be subgaussian with any variance proxy. Then, for any $\rho(n) \leq 1$, $(\mathcal{P}_n^{\pi, \rho(n)})$ is a tame spike prior with speed $\rho(n)$.*

Proof. Again, let us write $\mathcal{P}_n = \mathcal{P}_n^{\pi, \rho(n)}$ to lighten the notation. Following the beginning of the proof of Proposition 5.1.7, we find that π itself satisfies that for any $\eta > 0$ there exists $\delta > 0$ such that, for all $|t| \leq \delta$,

$$\mathbb{E}_{w^1, w^2 \sim \pi} \exp(t w^1 w^2) \leq \exp\left(\frac{1}{2\sqrt{1-\eta}} t^2\right). \quad (5.13)$$

Thus we have

$$\begin{aligned} \mathbb{E}_{x^1, x^2 \sim \mathcal{P}_n} \exp(t \langle x^1, x^2 \rangle) &= \mathbb{E}_{\substack{w_i^1, w_i^2 \sim \pi \\ s_i^1, s_i^2 \sim \text{Ber}(\rho)}} \exp\left(\frac{t}{\rho n} \sum_{i=1}^n w_i^1 w_i^2 s_i^1 s_i^2\right) \\ &= \left(\mathbb{E}_{\substack{w^1, w^2 \sim \pi \\ s^1, s^2 \sim \text{Ber}(\rho)}} \exp\left(\frac{t}{\rho n} w^1 w^2 s^1 s^2\right) \right)^n \end{aligned}$$

and since $s^1 s^2 = 1$ with probability ρ^2 and otherwise equals zero, taking the expectation over these variables first gives

$$= \left(1 - \rho^2 + \rho^2 \mathbb{E}_{w^1, w^2 \sim \pi} \exp\left(\frac{t}{\rho n} w^1 w^2\right) \right)^n$$

so if $t \leq \delta \rho n$, then

$$\leq \left(1 + \rho^2 \left(\exp \left(\frac{t^2}{2\sqrt{1-\eta}\rho^2 n^2} \right) - 1 \right) \right)^n$$

Choosing δ sufficiently small that for all $|x| \leq \delta$ we furthermore have $\exp\left(\frac{x^2}{2\sqrt{1-\eta}}\right) - 1 \leq \frac{1}{2(1-\eta)}x^2$, which is possible by comparing Taylor expansions, we then have

$$\begin{aligned} &\leq \left(1 + \frac{t^2}{2(1-\eta)n^2} \right)^n \\ &\leq \exp \left(\frac{t^2}{2(1-\eta)n} \right). \end{aligned} \tag{5.14}$$

Finally, applying a Chernoff bound we have, for $t \leq \delta \rho$,

$$\begin{aligned} \mathcal{P}_n^{\otimes 2}[\langle \mathbf{x}^1, \mathbf{x}^2 \rangle \geq t] &= \mathcal{P}_n^{\otimes 2}[(1-\eta)nt \langle \mathbf{x}^1, \mathbf{x}^2 \rangle \geq (1-\eta)nt^2] \\ &\leq \exp \left(-n \frac{1-\eta}{2} t^2 \right), \end{aligned} \tag{5.15}$$

which gives the local subgaussianity with speed $\rho(n)$, with the other tail bound following by a symmetric argument.

For the large deviations bound, we follow the proof of Proposition 5.1.7. If $|\langle \mathbf{x}^1, \mathbf{x}^2 \rangle| \geq t$, then either $\|\mathbf{x}^1\|^2 \geq t$ or $\|\mathbf{x}^2\|^2 \geq t$, so, by union bound,

$$\begin{aligned} \mathcal{P}_n^{\otimes 2}[|\langle \mathbf{x}^1, \mathbf{x}^2 \rangle| \geq t] &\leq 2\mathcal{P}_n[\|\mathbf{x}\|^2 \geq t] \\ &= 2\mathcal{P}_n \left[\frac{1}{\rho n} \sum_{i=1}^n w_i^2 s_i \geq t \right] \end{aligned}$$

and by Chernoff bound, if $C = \mathbb{E} \exp(\epsilon w^2)$,

$$\begin{aligned}
&\leq 2 \left(\mathbb{E} \exp(\epsilon w^2 s) \right)^n \exp(-\epsilon \rho n t) \\
&= 2 \left(1 + \rho^2 (\mathbb{E} \exp(\epsilon w^2) - 1) \right)^n \exp(-\epsilon \rho n t) \\
&\leq 2 \exp \left(C \rho^2 n - \epsilon \rho n t \right), \tag{5.16}
\end{aligned}$$

giving the upper bound for $t \geq T\rho$ for sufficiently large T . \square

We note that a fairly similar analysis has appeared in [LWB20] using Bernstein’s inequality to control the overlaps of sparse priors. Our application of Proposition 5.1.10 serves the same purpose—Bernstein’s inequality similarly interpolates smoothly between different tail regimes for small and large deviations—but is slightly more flexible in treating subgaussian underlying distributions rather than only bounded ones.

5.2 GAUSSIAN WIGNER AND WISHART MODELS

5.2.1 BOUNDING WISHART BY WIGNER

We first show a very useful relationship between Gaussian Wigner and Wishart models, which essentially states that, up to a rescaling and small shift of the magnitude of the signal, **a sign-definite Wishart model is at least as hard as a Wigner model with the same signal prior**. Moreover, this simple relationship even holds pointwise for the link polynomials $\phi_{N,D}^{\text{Wish}}$ and ϕ_D^{Wig} , without taking expectations over signal priors.

Lemma 5.2.1. *For $N, D \in \mathbb{N}_+$ and $\mathbf{T} \succeq \mathbf{0}$,*

$$\phi_{N,D}^{\text{Wish}}(\mathbf{T}) \leq \phi_D^{\text{Wig}} \left(\left(\frac{N}{2} + D \right) \text{tr}(\mathbf{T}) \right). \tag{5.17}$$

Proof. Since $\mathbf{T} \succeq \mathbf{0}$, we have $\text{tr}(\mathbf{T}^k) \leq \text{tr}(\mathbf{T})^k$. Therefore, working from Definition 4.4.4 of the polynomials $r_{N,d}$ appearing in $\phi_{N,D}^{\text{Wish}}$, we have

$$\begin{aligned} r_{N,d}(\mathbf{T}) &= \frac{1}{d!} \sum_{\pi \in \text{Part}([d])} \left(\frac{N}{2}\right)^{|\pi|} \prod_{S \in \pi} (|S| - 1)! \text{tr}(\mathbf{T}^{|S|}) \\ &\leq \frac{\text{tr}(\mathbf{T})^d}{d!} \sum_{\pi \in \text{Part}([d])} \left(\frac{N}{2}\right)^{|\pi|} \prod_{S \in \pi} (|S| - 1)! \end{aligned}$$

and now, noting that the remaining sum by Definition 4.4.4 and Corollary 4.4.14 is just equal to $d! a_{N,d}$ from the statement of the Corollary (alternatively, one may compute this directly as an evaluation of one definition of Stirling numbers of the first kind), we have

$$= \frac{\text{tr}(\mathbf{T})^d}{d!} \prod_{k=0}^{d-1} \left(\frac{N}{2} + k\right)$$

and if $d \leq D$, then

$$\leq \frac{1}{d!} \left(\left(\frac{N}{2} + D\right) \text{tr}(\mathbf{T}) \right)^d. \quad (5.18)$$

Thus, we have

$$\phi_{N,D}^{\text{Wish}}(\mathbf{T}) = \sum_{d=0}^D r_{N,d}(\mathbf{T}) \leq \sum_{d=0}^D \frac{1}{d!} \left(\left(\frac{N}{2} + D\right) \text{tr}(\mathbf{T}) \right)^d = \phi_D^{\text{Wig}} \left(\left(\frac{N}{2} + D\right) \text{tr}(\mathbf{T}) \right), \quad (5.19)$$

as claimed. □

To see briefly why this relation will be so useful, recall that we will consider Wishart models with signals $\tilde{\mathbf{X}} = \beta \mathbf{X} \mathbf{X}^\top$, whereby the overlaps input into $\phi_{N,D}^{\text{Wish}}$ will be of the form $\beta^2 \mathbf{X}^1 \mathbf{X}^{1^\top} \mathbf{X}^2 \mathbf{X}^{2^\top}$. By the cyclic property of $\phi_{N,D}^{\text{Wish}}$ (following from the cyclic property of the $r_{N,d}$ from Proposition 4.4.6), we will then have

$$\phi_{N,D}^{\text{Wish}}(\beta^2 \mathbf{X}^1 \mathbf{X}^{1^\top} \mathbf{X}^2 \mathbf{X}^{2^\top}) = \phi_{N,D}^{\text{Wish}}(\beta^2 (\mathbf{X}^{1^\top} \mathbf{X}^2) (\mathbf{X}^{1^\top} \mathbf{X}^2)^\top), \quad (5.20)$$

where the input is now positive semidefinite and Lemma 5.2.1 applies.

5.2.2 TRUNCATED EXPONENTIAL POLYNOMIALS

In light of the previous result, the only link polynomials whose analytic properties we will need to consider carefully are the $\phi_D^{\text{Wig}}(t) = \sum_{d=0}^D \frac{1}{d!} t^d$. One useful device is the following integral form of these “truncated exponential polynomials;” see, e.g., [DCS03] for this identity and further information. Recall that, at least for large t , we expect $\phi_D^{\text{Wig}}(t) \approx \frac{t^D}{D!}$; this result gives an exact integral expression for the excess in this approximation.

Proposition 5.2.2. *For any $t > 0$,*

$$\phi_D^{\text{Wig}}(t) = \frac{t^D}{D!} \int_0^\infty e^{-s} \left(1 + \frac{s}{t}\right)^D ds, \quad (5.21)$$

and $\phi_D^{\text{Wig}}(0) = 1$.

Proof. Expanding the integral with the binomial theorem,

$$\frac{1}{D!} \int_0^\infty e^{-s} (s+t)^D ds = \sum_{d=0}^D \frac{t^d}{d!(D-d)!} \int_0^\infty e^{-s} s^{D-d} ds = \sum_{d=0}^D \frac{t^d}{d!} = \phi_D^{\text{Wig}}(t) \quad (5.22)$$

with the integrals evaluated as gamma functions. □

Corollary 5.2.3. *For any $t \geq 0$, we have the bounds*

$$\frac{t^D}{D!} \leq \phi_D^{\text{Wig}}(t) \leq \min \left\{ \exp(t), 2 \frac{(2D \vee t)^D}{D!} \right\}. \quad (5.23)$$

Proof. The lower bound is immediate since all terms in ϕ_D^{Wig} are non-negative. For the upper bound, $\phi_D^{\text{Wig}}(v) \leq \exp(v)$ is again immediate for the same reason. For the remaining inequality, since $\phi_D^{\text{Wig}}(v)$ is increasing in v , it suffices to consider the case $v > 2D$. In this

case, we bound

$$\begin{aligned}
\phi_D^{\text{Wig}}(v) &= \frac{v^D}{D!} \int_0^\infty e^{-u} \left(1 + \frac{u}{v}\right)^D du \\
&\leq \frac{v^D}{D!} \int_0^\infty \exp\left(-u + \frac{u}{v}D\right) du \\
&\leq \frac{v^D}{D!} \cdot \frac{1}{1 - \frac{D}{v}} \\
&\leq 2 \frac{v^D}{D!},
\end{aligned} \tag{5.24}$$

and the result follows. \square

We also make the following simple observation, which we will use repeatedly to handle the differentiation of the link function appearing in Lemma 5.1.1.

Proposition 5.2.4. *For any $D \in \mathbb{N}$ and $t \geq 0$, $\frac{d}{dt} \phi_D^{\text{Wig}}(t) = \phi_{D-1}^{\text{Wig}}(t) \leq \phi_D^{\text{Wig}}(t)$, where we interpret $\phi_{-1}^{\text{Wig}} = 0$.*

5.2.3 TENSOR PCA

The following model, first introduced by [RM14], has become quite popular as a tensor-valued analog of spiked matrix models. We present results on this model before those for spiked matrix models since, as we will see, the latter are just a special case of the former where a more precise analysis is possible.

Definition 5.2.5 (Spiked tensor model). *Let $p \geq 2$ and $\lambda = \lambda(n) > 0$. Given a sequence $\mathcal{P}_n \in \mathcal{P}(\mathbb{R}^n)$, the spiked tensor model of order p is the Gaussian Wigner model with $N(n) = n^p$ and signal prior $\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n$ given by sampling $\tilde{\mathbf{X}} = \lambda \mathbf{x}^{\otimes p}$ for $\mathbf{x} \sim \mathcal{P}_n$. That is, the null and planted distributions are given respectively by:*

- Under \mathbb{Q}_n , draw $\mathbf{Y} \in (\mathbb{R}^n)^{\otimes p}$ with independent entries distributed as $\mathcal{N}(0, 1)$.

- Under \mathbb{P}_n , first draw $\mathbf{x} \sim \mathcal{P}_n$, draw $\mathbf{G} \in (\mathbb{R}^n)^{\otimes p}$ with independent entries distributed as $\mathcal{N}(0, 1)$, and observe $\mathbf{Y} = \lambda \mathbf{x}^{\otimes p} + \mathbf{G}$.

More briefly, we say $(\mathbb{Q}_n, \mathbb{P}_n)_{n \geq 1}$ form a spiked tensor model with parameters (λ, \mathcal{P}_n) .

Theorem 5.2.6 (Spiked tensor lower bound: dense, rank one). *Suppose (\mathcal{P}_n) is a tame rank-one spike prior. Then, there exists $c > 0$ depending only on p and the tail bounds in the tameness of the prior such that, whenever $D \leq cn$ and $\lambda \leq (9p)^{-p/4} n^{p/4} D^{(2-p)/4}$, then $\|L_n^{\leq D}\| = O(1)$ in the spiked tensor model with parameters (λ, \mathcal{P}_n) .*

This suggests that, per the extended Conjecture 3.2.3, in time $\exp(O(n^\delta))$ we may distinguish in the spiked tensor model so long as $\lambda \gg n^{-p/4 - \delta(p-2)/4}$ (neglecting logarithmic factors). For $\delta = 0$ this concerns polynomial-time algorithms, and the threshold of $\lambda \gg n^{-p/4}$ is the same as that achieved by various algorithms in the literature [RM14, HSS15, ADGM17, HSS16]. For $\delta > 0$ this concerns subexponential-time algorithms, and the corresponding relationship between computational cost and the threshold λ coincides to other subexponential-time algorithms [BGG⁺16, BGL16, RRS17, WEM19]. Finally, for $\delta = 1$ this concerns exponential-time algorithms, and thus at least informally we expect to recover the threshold of statistical distinguishability, since we do not expect super-exponential time algorithms to be of use when exponential-time algorithms cannot distinguish \mathbb{P}_n from \mathbb{Q}_n . Indeed, this gives the threshold $\lambda \ll n^{(1-p)/2}$, which is the threshold below which distinguishing is impossible [RM14, PWB16, LML⁺17, JLM20].

Proof. Let $\tilde{\mathbf{X}}^1 = \lambda(\mathbf{x}^1)^{\otimes p}$, $\tilde{\mathbf{X}}^2 = \lambda(\mathbf{x}^2)^{\otimes p}$ be two independent draws from $\tilde{\mathcal{P}}_n$, for two independent draws $\mathbf{x}^i \sim \mathcal{P}_n$. Then, the overlap is $\langle \tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \rangle = \lambda^2 \langle \mathbf{x}^1, \mathbf{x}^2 \rangle^p$. Let us write $\psi(t)$ for the “rate function” of the tail bound on the overlap that tameness provides: for a given η, δ

parameters of local subgaussianity,

$$\psi(t) = \begin{cases} \frac{1-\eta}{2}t^2 & \text{if } 0 \leq t \leq \delta, \\ \frac{1-\eta}{2}\delta^2 & \text{if } \delta \leq t \leq T, \\ (\log(1+t))^{1+\xi} & \text{if } t \geq T. \end{cases} \quad (5.25)$$

Then, we recall that $\mathcal{P}_n^{\otimes 2}[|\langle \mathbf{x}^1, \mathbf{x}^2 \rangle| \geq t] \leq \exp(-n\psi(t))$. Therefore, we have

$$\|L_n^{\leq D}\|^2 = \mathbb{E}\phi_D^{\text{Wig}}(\lambda^2 \langle \mathbf{x}^1, \mathbf{x}^2 \rangle^p) \quad (\text{Theorem 4.1.7})$$

and, before proceeding, we note that this is monotone in λ , whereby we may suppose without loss of generality that $\lambda = (9p)^{-p/4} n^{p/4} D^{(2-p)/4} = (n/9pD)^{p/4} \sqrt{D}$. Continuing,

$$\leq \int_0^\infty \exp(-n\psi(\lambda^{-2/p} t^{1/p})) \phi_D^{\text{Wig}}(t) dt \quad (\text{Lemma 5.1.1})$$

Let local subgaussianity hold over $[0, \delta]$ with $\eta = \frac{1}{3}$. Let us choose $c \leq \frac{\delta^2}{18p}$, so that $D \leq \frac{\delta^2}{18p} n$, whereby we have $\delta^p \lambda^2 = (\delta^2 n/9pD)^{p/2} D \geq 2^{p/2} D \geq 2D$. Therefore, letting T be the upper threshold in the tameness assumption on the spike prior, assuming without loss of generality that $T > \delta$, we may divide the integral into four regions and substitute in the corresponding behavior of ψ and ϕ_D^{Wig} , as follows:

$$\begin{aligned} &\leq \int_0^{2D} \exp\left(t - \frac{1}{3}n\lambda^{-4/p}t^{2/p}\right) dt \\ &\quad + \int_{2D}^{\delta^p \lambda^2} \exp\left(-\frac{1}{3}n\lambda^{-4/p}t^{2/p}\right) \frac{t^D}{D!} dt \\ &\quad + \int_{\delta^p \lambda^2}^{T^p \lambda^2} \exp\left(-\frac{1}{3}n\delta^2\right) \frac{t^D}{D!} dt \\ &\quad + \int_{T^p \lambda^2}^\infty \exp\left(-n(\log(1+t))^{1+\xi}\right) \frac{t^D}{D!} dt \\ &=: A_1 + A_2 + A_3 + A_4. \end{aligned} \quad (5.26)$$

We now treat the four terms individually, showing that each is $O(1)$ as $n \rightarrow \infty$. We also note before proceeding that $n\lambda^{-4/p} = 9pD^{\frac{p-2}{p}}$.

For A_1 , since $t \mapsto t^{2/p}$ is a concave function, on the interval $t \in [0, 2D]$ we have $t^{2/p} \geq (2D)^{-\frac{p-2}{p}} t$ as this line is the secant of the function on this interval. Thus we find

$$\begin{aligned} A_1 &\leq \int_0^\infty \exp\left(t \left[1 - \frac{2^{-\frac{p-2}{p}}}{3} n\lambda^{-4/p} D^{2/p-1}\right]\right) dt \\ &= \int_0^\infty \exp\left(t \left[1 - 2^{-\frac{p-2}{p}} \cdot 3p\right]\right) dt \end{aligned} \quad (5.27)$$

The remaining rate is strictly negative, so this integral is a finite constant and $A_1 = O(1)$.

For A_2 , we extend the domain of integration and compute the full integral over $\mathbb{R}_{\geq 0}$,

$$A_2 \leq \frac{1}{D!} \int_0^\infty \exp\left(-3pD^{\frac{p-2}{p}} t^{2/p}\right) t^D dt$$

If $p = 2$, then the integral is at most $D!$ by evaluating it as a gamma function, whereby we find $A_2 \leq 1$. Thus let us suppose $p \geq 3$. Performing the change of variables $s = 3pD^{\frac{p-2}{p}} t^{2/p}$, we find

$$= \frac{1}{6} \frac{1}{D!} (3p)^{-\frac{pD+p-2}{2}} D^{-\frac{p-2}{2}(D+1)} \int_0^\infty e^{-s} s^{\frac{pD+p-2}{2}} ds$$

Evaluating again as a gamma function and bounding the result, we find

$$\begin{aligned} &\leq \frac{1}{D!} D^{-\frac{p-2}{2}(D+1)} \left(\frac{pD+p-2}{6p}\right)^{\frac{pD+p-2}{2}} \\ &\leq \frac{1}{D!} D^{-\frac{p-2}{2}(D+1)} \left(\frac{D}{3}\right)^{\frac{pD+p-2}{2}} \\ &\leq \frac{1}{D!} \left(\frac{D}{3}\right)^D, \end{aligned} \quad (5.28)$$

which is bounded by Stirling's approximation, so $A_2 = O(1)$.

For A_3 , we note that the rate term is constant, so we simply have

$$\begin{aligned}
A_3 &\leq \exp\left(-\frac{1}{3}n\delta^2\right) \int_{\delta^p \lambda^2}^{T^p \lambda^2} \frac{t^D}{D!} dt \\
&\leq \exp\left(-\frac{1}{3}n\delta^2\right) \cdot \frac{T^{p(D+1)} \lambda^{2(D+1)}}{(D+1)!} \\
&\leq \exp\left(-\frac{1}{3}n\delta^2\right) \cdot \frac{(T^2 \frac{n}{D})^{p(D+1)/2} D^{D+1}}{(D+1)!}
\end{aligned}$$

and using Stirling's approximation and suppressing a constant,

$$\leq \exp\left(-\frac{1}{3}n\delta^2\right) \cdot \left(eT^2 \frac{n}{D}\right)^{p(D+1)/2}$$

Assuming without loss of generality that $T \geq e$, we see upon taking a derivative that the second term is increasing with D over all $0 \leq D \leq n$. Therefore, if $c < 1$ and $D \leq cn$, over these permissible D the second factor will be maximized by $D = cn$, with which we find

$$\leq \exp\left(-n \left[\frac{1}{3}\delta^2 - pc \log\left(\frac{eT^2}{c}\right) \right]\right), \quad (5.29)$$

whereby choosing c small enough will make the exponent negative, and thus $A_3 = O(1)$.

Finally, for A_4 we may again compare to the integral extended to all of $\mathbb{R}_{\geq 0}$, and bound coarsely using $D \leq n$:

$$A_4 \leq \int_0^\infty \exp\left(-n \left[(\log(1+t))^{1+\xi} - \log(1+t) \right]\right) dt, \quad (5.30)$$

whereby $A_4 = O(1)$, completing the proof. □

We illustrate this proof in Figure 5.1, comparing the rate function in the tail bound on the overlap and the truncated exponential polynomials on the exponential scale and showing how our decomposition corresponds to the different regimes of these functions.

Remark 5.2.7 (Leading-order behavior and prefactor). *The decomposition we have introduced in the proof above into integrals A_i conveniently corresponds to determining increasingly fine-grained behavior of the threshold of computational hardness: the analysis of A_1 gives the dependence of the critical λ upon n and D , the analysis of A_2 gives the $\exp(\Theta(-p \log p))$ dependence of the prefactor on p , and the analyses of A_3 and A_4 are essentially generic and do not use any special properties of the model.*

Remark 5.2.8 (Wishart spiked tensor models). *It is reasonable to wonder whether there is a model whose LDLR is given by the Wishart link function ϕ_D^{Wish} applied to the tensor overlap $\langle \mathbf{x}^1, \mathbf{x}^2 \rangle^p$. In fact, a similar model appeared recently in [CHK⁺20], under the name of “single-spike block mixtures,” where in the planted model one draws $\mathbf{s} \sim \text{Unif}(\{\pm 1\}^d)$, $\mathbf{x} \sim \text{Unif}(\{\pm 1/\sqrt{n}\}^n)$, and then $\mathbf{y}_{i1}, \dots, \mathbf{y}_{im} \sim \mathcal{N}(\mathbf{0}, \mathbf{I} + \lambda s_i \mathbf{x} \mathbf{x}^\top)$ for $i \in [d]$. This may be written as a Wishart spiked matrix model with a block-structured covariance matrix, but one may check that, after applying the comparison bound of Lemma 5.2.1, its LDLR is bounded by $\|L_n^{\leq D}\|^2 \leq \mathbb{E} \phi_D^{\text{Wig}}(\lambda^2 \langle \mathbf{s}^1, \mathbf{s}^2 \rangle \langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2)$, like a Wigner tensor model with spike $\lambda \mathbf{s} \otimes \mathbf{x} \otimes \mathbf{x}$. Our tools should yield a streamlined proof of the results on this model used there; it would be interesting to understand if such models arise in other settings.*

5.2.4 DENSE MATRIX PCA

We now consider spiked matrix models. We recall that we have seen two variants: the Wishart spiked matrix model of Definition 2.2.4, and the Wigner spiked matrix model of Example 4.1.2. We restate more precisely and generally the definition of the latter model below, in particular allowing spikes of arbitrary fixed rank.

Definition 5.2.9 (Wigner spiked matrix model). *Let $\lambda = \lambda(n) > 0$ and $k \in \mathbb{N}_+$ not depending on n . Given a sequence $\mathcal{P}_n \in \mathcal{P}(\mathbb{R}^{n \times k})$, the Wigner spiked matrix model is the Gaussian Wigner model with $N(n) = n^2$ and signal prior $\tilde{\mathbf{X}} \sim \tilde{\mathcal{P}}_n$ given by sampling $\tilde{\mathbf{X}} = \frac{\lambda}{\sqrt{2}} \mathbf{X} \mathbf{X}^\top$ for*




Figure 5.1: Low-degree lower bounds: dense matrix and tensor PCA. We illustrate the proofs of Theorems 5.2.6 and 5.2.10, showing the relationship on an exponential scale between the tail bound on a tame prior and the truncated exponential polynomial. “Knots” where these functions change behavior are marked with solid circles, and bounds from the proofs are marked with dotted lines.

$\mathbf{X} \sim \mathcal{P}_n$. Equivalently, this is the spiked tensor model with $p = 2$ and λ rescaled by $\sqrt{2}$, or the model with null and planted distributions given respectively by:

- Under \mathbb{Q}_n , draw $\mathbf{Y} \in \mathbb{R}^{n \times n}$ with independent entries distributed as $\mathcal{N}(0, 1)$.
- Under \mathbb{P}_n , draw $\mathbf{G} \in \mathbb{R}^{n \times n}$ with independent entries distributed as $\mathcal{N}(0, 1)$, and observe
$$\mathbf{Y} = \frac{\lambda}{\sqrt{2}} \mathbf{X} \mathbf{X}^\top + \mathbf{G}.$$

More briefly, we say $(\mathbb{Q}_n, \mathbb{P}_n)_{n \geq 1}$ form a Wigner spiked matrix model with parameters (λ, \mathcal{P}_n) .

The reason for rescaling λ is that, more often, this model is taken with symmetric matrix observations and $\mathbf{G} \sim \sqrt{n} \cdot \text{GOE}(n)$ above (or, what is equivalent, $\mathbf{G} \sim \text{GOE}(n)$ and λ a constant independent of n). The above model is *a priori* more “favorable” for testing, since this symmetric model may be recovered by forming $\mathbf{Y}^{\text{sym}} = \frac{1}{\sqrt{2}}(\mathbf{Y} + \mathbf{Y}^\top)$, with the normalizing factor chosen so that \mathbf{Y}^{sym} has the law $\sqrt{n} \cdot \text{GOE}(n)$ when $\mathbf{Y} \sim \mathbb{Q}_n$. Actually, it is straightforward to show that the two models are equivalent for the purposes of hypothesis testing. Under this symmetrizing transformation the $\sqrt{2}$ in the definition above cancels, so we expect the model defined above to have the same critical λ as the GOE model, namely $\lambda(n) = \sqrt{n}$.

Theorem 5.2.10 (Wigner spiked matrix lower bound: dense, rank k). *Suppose (\mathcal{P}_n) is a tame spike prior. Then, for any $\epsilon > 0$, there exists $c > 0$ depending only on ϵ and the tail bounds in the tameness of the prior such that, whenever $\lambda \leq (1 - \epsilon)\sqrt{n}$ and $D \leq cn$, then $\|L_n^{\leq D}\| = O(1)$ in the Wigner spiked matrix model with parameters (λ, \mathcal{P}_n) .*

We note that this matches the behavior of the PCA test mentioned earlier, for which the threshold of distinguishing is $\lambda \sim \sqrt{n}$ per Proposition 2.2.3.

Remark 5.2.11 (The “Gaussian heuristic”). *We mention a useful heuristic argument that can be used to quickly predict the outcome of such a computation quite easily. Suppose we are in*

the rank-one case of the Wigner spiked matrix model. Then, at the critical scaling, we have $\|L_n^{\leq D}\|^2 = \mathbb{E}\phi_D^{\text{Wig}}(\frac{1}{2}\hat{\lambda}n\langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2)$ for $\hat{\lambda}$ of order constant. Since $\|\mathbf{x}\| \approx 1$ for $\mathbf{x} \sim \mathcal{P}_n$, for “nice” spike priors we may suppose by a central limit theorem heuristic that $\sqrt{n}\langle \mathbf{x}^1, \mathbf{x}^2 \rangle \rightarrow \mathcal{N}(0, 1)$ in distribution. In particular, for large n , $n\langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2$ is approximately distributed as a χ^2 random variable (with one degree of freedom). Supposing that $D \rightarrow \infty$ slowly enough that $\phi_D^{\text{Wig}} \rightarrow \phi^{\text{Wig}}$ “after” this convergence in distribution, we might expect $\limsup_{n \rightarrow \infty} \|L_n^{\leq D}\|^2 \leq \mathbb{E}\exp(\frac{1}{2}\hat{\lambda}g^2)$ for $g \sim \mathcal{N}(0, 1)$. This is finite if and only if $\hat{\lambda} < 1$, recovering our result above. It is an intriguing open problem to find a direct way to make this reasoning rigorous with a quantitative version of the central limit theorem for priors of interest.

Proof of Theorem 5.2.10. We essentially repeat the proof of Theorem 5.2.6 with a few adjustments. First, because of our rescaling and the more general rank- k spike prior, we will have $\langle \tilde{\mathbf{X}}^1, \tilde{\mathbf{X}}^2 \rangle = \frac{1}{2}\lambda^2\|\mathbf{X}^{1\top}\mathbf{X}^2\|_F^2$. Recall, though, that by the definition of tameness of priors, the overlap will satisfy the same tail bound as in the rank-one case used in the proof of Theorem 5.2.6,

$$\mathcal{P}_n^{\otimes 2}[\|\mathbf{X}^{1\top}\mathbf{X}^2\|_F \geq t] \leq \exp(-n\psi(t)), \quad (5.31)$$

for $\psi(t)$ as in (5.25).

Second, we will make full use of local subgaussianity: let $\eta > 0$ be a small constant to be fixed later, and δ corresponding width of the interval of local subgaussianity. Then,

$$\begin{aligned} \|L_n^{\leq D}\|^2 &= \mathbb{E}\phi_D^{\text{Wig}}\left(\frac{\lambda^2}{2}\|\mathbf{X}^{1\top}\mathbf{X}^2\|_F^2\right) && \text{(Theorem 4.1.7)} \\ &\leq \int_0^\infty \exp\left(-n\psi\left(\frac{\sqrt{2t}}{\lambda}\right)\right)\phi_D^{\text{Wig}}(t)dt && \text{(Lemma 5.1.1)} \end{aligned}$$

and decomposing as before,

$$\begin{aligned}
&\leq \int_0^{2D} \exp\left(t - (1 - \eta) \frac{n}{\lambda^2} t\right) dt \\
&\quad + \int_{2D}^{\delta^2 \lambda^2} \exp\left(- (1 - \eta) \frac{n}{\lambda^2} t\right) \frac{t^D}{D!} dt \\
&\quad + \int_{\delta^2 \lambda^2}^{c^2 \lambda^2} \exp\left(- (1 - \eta) n \delta^2\right) \frac{t^D}{D!} dt \\
&\quad + \int_{c^2 \lambda^2}^{\infty} \exp\left(-n(\log(1 + t))^{1+\xi}\right) \frac{t^D}{D!} dt \\
&=: A_1 + A_2 + A_3 + A_4.
\end{aligned} \tag{5.32}$$

We have $A_3 = O(1)$ and $A_4 = O(1)$ by the same coarse bounds as in the proof of Theorem 5.2.6. We choose η sufficiently small that $\frac{1-\eta}{(1-\epsilon)^2} > 1$. Thus we will have $A_1 = O(1)$ and $A_2 \leq 1$ by extending both to full integrals over $\mathbb{R}_{\geq 0}$, and integrating directly for A_1 and evaluating A_2 as a gamma function. \square

We also give a graphical illustration of this argument in Figure 5.1 as for the tensor case, showing how the sharp threshold arises from the small deviation tails exhibiting exponential decay (with linear exponent) independent of D .

Next, displaying the power of Lemma 5.2.1, we deduce with almost no further work the corresponding lower bound for the Wishart model.

Theorem 5.2.12 (Wishart spiked matrix lower bound: dense, rank k). *Suppose (\mathcal{P}_n) is a tame spike prior. Then, for any $\beta > -1$ and $\gamma > 0$ such that $\beta^2/\gamma < 1$, there exists $c > 0$ depending only on β, γ , and the tail bounds in the tameness of the prior such that, whenever $D \leq cn$, then $\|L_n^{\leq D}\| = O(1)$ in the Wishart spiked matrix model with parameters $(\beta, \gamma, \mathcal{P}_n)$.*

Proof. Recall by Proposition 4.4.13 that we may ignore the truncation in the Wishart spiked matrix model for the purposes of these bounds. Using the cyclic property of Proposition 4.4.6 to transform the Wishart link polynomials, and then bounding with Lemma 5.2.1,

we find

$$\begin{aligned}
\|L_n^{\leq D}\|^2 &= \mathbb{E} \phi_{n/\gamma, D}^{\text{Wish}} \left(\beta^2 \mathbf{X}^1 \mathbf{X}^{1\top} \mathbf{X}^2 \mathbf{X}^{2\top} \right) \\
&= \mathbb{E} \phi_{n/\gamma, D}^{\text{Wish}} \left(\beta^2 (\mathbf{X}^{1\top} \mathbf{X}^2) (\mathbf{X}^{1\top} \mathbf{X}^2)^\top \right) \\
&\leq \mathbb{E} \phi_D^{\text{Wig}} \left(\frac{\beta^2}{2\gamma} \left(n + \frac{\gamma}{2} D \right) \|\mathbf{X}^{1\top} \mathbf{X}^2\|_F^2 \right) \tag{5.33}
\end{aligned}$$

$$\leq \mathbb{E} \phi_D^{\text{Wig}} \left(\frac{\beta^2}{2\gamma} \left(1 + \frac{\gamma}{2} c \right) n \|\mathbf{X}^{1\top} \mathbf{X}^2\|_F^2 \right) \tag{5.34}$$

whereby the result holds by taking c sufficiently small that $\frac{\beta^2}{\gamma} (1 + \frac{\gamma}{2} c) < 1$ and applying the argument of Theorem 5.2.10. \square

Remark 5.2.13. *As is well-understood in the random matrix literature, for the Wishart model β^2/γ behaves as the “effective signal-to-noise ratio,” playing the role of λ/\sqrt{n} from the Wigner model. Lemma 5.2.1 gives us a pleasantly direct way of applying this equivalence to derive computational lower bounds, as we have done above. We risk belaboring this point because, while in all relevant statistical and computational behaviors the analogy between the Wigner and Wishart spiked matrix models is quite strong, we are not aware of any earlier works that have found tools to directly compare one to the other.*

5.2.5 SPARSE MATRIX PCA

Finally, we also consider the case of the setting of the previous section where the spike prior is *sparse* (here we mean having sparsity $\rho(n) = o(1)$, in contrast to the constant sparsity discussed in Example 2.4.1).

Theorem 5.2.14 (Wigner spiked matrix lower bound: sparse, rank one). *Suppose (\mathcal{P}_n) is a tame rank-one spike prior with speed $\rho = \rho(n)$ satisfying that $\liminf_{n \rightarrow \infty} \rho(n)^2 n > 0$. Then, for any $\epsilon > 0$, there exists $c > 0$ depending only on ϵ and the tail bounds in the tameness of*

the prior such that, whenever $\lambda \leq (1 - \epsilon)\sqrt{n}$ and $D \leq c\rho^2 n$, then $\|L_n^{\leq D}\| = O(1)$ in the Wigner spiked matrix model with parameters (λ, \mathcal{P}_n) .

The result suggests that it is possible to distinguish in the spiked matrix model (Wigner in this case, or Wishart below) in time $\exp(O(\rho^2 n))$ for a prior of sparsity ρ . This matches the threshold achieved by algorithms that enumerate principal submatrices of size $\rho^2 n \times \rho^2 n$ and search by brute force for such a submatrix with a large eigenvalue, as proposed concurrently by [DKWB19, HSV20]. For polynomial-time algorithms this suggests that $\rho \lesssim 1/\sqrt{n}$ is required when $\lambda \leq (1 - \epsilon)\sqrt{n}$ (below the PCA or BBP threshold), which is compatible with the algorithms of [JL09, DM14], the SOS lower bounds of [MW15, HKP⁺17], and the reduction arguments of [BR13, WBS16, BBH18, BB19b].

Proof. We recapitulate the proof of Theorem 5.2.10 with a few adjustments. As there, let $\eta > 0$ be a small constant to be fixed later, and $\delta\rho$ the corresponding width of the interval of local subgaussianity. Then,

$$\|L_n^{\leq D}\|^2 = \mathbb{E} \phi_D^{\text{Wig}} \left(\frac{\lambda^2}{2} \langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2 \right) \quad (\text{Theorem 4.1.7})$$

where we note that this expression is increasing in λ , so we may suppose without loss of generality that we have $\lambda = (1 - \epsilon)\sqrt{n}$. Then, choosing c sufficiently small depending on δ , we will have $2D \leq 2c\rho^2 n \leq \delta^2 \rho^2 \lambda^2 = (1 - \epsilon)^2 \delta^2 \rho^2 n$, whereby we may proceed with our usual bound

$$\leq \int_0^\infty \exp \left(-n\psi \left(\frac{\sqrt{2t}}{\lambda} \right) \right) \phi_D^{\text{Wig}}(t) dt \quad (\text{Lemma 5.1.1})$$

and our usual decomposition of the integral:

$$\begin{aligned}
&\leq \int_0^{2D} \exp\left(t - (1 - \eta)\frac{n}{\lambda^2}t\right) dt \\
&\quad + \int_{2D}^{\delta^2\rho^2\lambda^2} \exp\left(- (1 - \eta)\frac{n}{\lambda^2}t\right) \frac{t^D}{D!} dt \\
&\quad + \int_{\delta^2\rho^2\lambda^2}^{T^2\rho^2\lambda^2} \exp\left(- (1 - \eta)\delta^2\rho^2n\right) \frac{t^D}{D!} dt \\
&\quad + \int_{T^2\rho^2\lambda^2}^{\infty} \exp\left(-\rho^2n(\log(1+t))^{1+\xi}\right) \frac{t^D}{D!} dt \\
&=: A_1 + A_2 + A_3 + A_4.
\end{aligned} \tag{5.35}$$

The proof is now mostly identical to that of Theorem 5.2.10: $A_1, A_2, A_3 = O(1)$ by the same arguments, with n replaced by ρ^2n throughout and noting that the assumption $D \leq cn$ has been replaced with $D \leq c\rho^2n$ accordingly. The only small difference is in A_4 , where we have

$$A_4 \leq \int_0^{\infty} \exp\left(-\rho^2n\left[(\log(1+t))^{1+\xi} - \log(1+t)\right]\right), \tag{5.36}$$

and here we use the assumption that ρ^2n is bounded above 0 for sufficiently large n so that this is $O(1)$ as well. □

Theorem 5.2.15 (Wishart spiked matrix model: sparse, rank one). *Suppose (\mathcal{P}_n) is a tame rank-one spike prior with speed $\rho = \rho(n)$ satisfying that $\liminf_{n \rightarrow \infty} \rho(n)^2n > 0$. Then, for any $\beta > -1$ and $\gamma > 0$ such that $\beta^2/\gamma < 1$, there exists $c > 0$ depending only on β, γ , and the tail bounds in the tameness of the prior such that, whenever $D \leq c\rho^2n$, then $\|L_n^{\leq D}\| = O(1)$ in the Wishart spiked matrix model with parameters $(\beta, \gamma, \mathcal{P}_n)$.*

The proof is identical to that of Theorem 5.2.12, essentially following immediately from Lemma 5.2.1.

5.3 CONSEQUENCES FOR CERTIFICATION

We finally arrive at our applications to certification for constrained PCA problems. Here, of the above results, we will draw only on the Wishart spiked matrix model, Theorem 5.2.12 above. This will be combined with the reduction results from Chapter 2.

It will be convenient to have a tool to truncate priors to force them to be β -good, as we have seen is important for our reduction arguments with the Wishart spiked matrix model.

Definition 5.3.1 (β -truncation). *For $\beta > -1$ and $\mathcal{P} \in \mathcal{P}(\mathbb{R}^{n \times k})$, define $\text{trunc}_\beta(\mathcal{P}) \in \mathcal{P}(\mathbb{R}^{n \times k})$ by sampling $\mathbf{X} \sim \text{trunc}_\beta(\mathcal{P})$ by first drawing $\mathbf{X}^{(0)} \sim \mathcal{P}$, and setting $\mathbf{X} = \mathbf{X}^{(0)}$ if $\beta \|\mathbf{X}^{(0)}\|^2 > -1$, and $\mathbf{X} = \mathbf{0}$ otherwise.*

The following are the important properties of this definition: truncation makes a spike prior β -good and does not change the Wishart spiked matrix model.

Proposition 5.3.2. *$\text{trunc}_\beta(\mathcal{P})$ is β -good for any $\mathcal{P} \in \mathcal{P}(\mathbb{R}^{n \times k})$.*

Proposition 5.3.3. *The Wishart spiked matrix models with parameters $(\beta, \gamma, \mathcal{P}_n)$ and with parameters $(\beta, \gamma, \text{trunc}_\beta(\mathcal{P}_n))$ are identical (in the sense of having the same sequences of probability measures $(\mathbb{Q}_n, \mathbb{P}_n)$).*

Thus any hardness result for the Wishart spiked matrix model with spike prior (\mathcal{P}_n) will also hold with spike prior $(\text{trunc}_\beta(\mathcal{P}_n))$.

We also introduce the following notion of “ β -goodness in probability,” which all of the spike priors we consider will satisfy.

Definition 5.3.4 (Weakly β -good). *For $\beta > -1$, call a spike prior (\mathcal{P}_n) weakly β -good if, under $\mathbf{X} \sim \mathcal{P}_n$, $\beta \|\mathbf{X}\|^2 > -1$ with high probability.*

Indeed, it suffices to have $\|\mathbf{X}\| \rightarrow 1$ in probability to have (\mathcal{P}_n) weakly β -good for all $\beta > -1$, and this is usually just a matter of normalization as convergence to *some* limiting norm should hold for most spike priors of interest. Under this condition, we have the following.

Proposition 5.3.5. *If (\mathcal{P}_n) is weakly β -good, then \mathcal{P}_n and $\text{trunc}_\beta(\mathcal{P}_n)$ may be coupled such that draws from either are equal with high probability.*

We will use these simple facts to argue as follows: suppose \mathcal{P}_n is one of the tame priors with independence structure as in Definitions 5.1.6 and 5.1.9. If (\mathcal{P}_n) satisfies all of the conditions of the reductions of Chapter 2 (Corollaries 2.3.3 and 2.3.4) except for being β -good, but (\mathcal{P}_n) is weakly β -good, then $\text{trunc}_\beta(\mathcal{P}_n)$ will satisfy all of the conditions, and, assuming that better-than-spectral certification is possible for an associated constrained PCA problem, we can infer that it is possible to distinguish in the Wishart spiked matrix model with parameters $(\beta, \gamma, \text{trunc}_\beta(\mathcal{P}_n))$. But this model is identical to the Wishart spiked matrix model with parameters $(\beta, \gamma, \mathcal{P}_n)$, so our hardness results concerning Wishart models with spike prior (\mathcal{P}_n) apply. We elide this essentially trivial but somewhat convoluted reasoning in the proofs below.

Our first result concerns certification in the SK Hamiltonian. This gives rigorous evidence addressing the question of whether better-than-spectral certification is possible, that was first raised by Montanari following the publication of [MS16] (for example it is mentioned more explicitly in [Mon18]), and later repeated in [JKR19].

Corollary 5.3.6 (SK Hamiltonian [BKW20b]). *Let $\mathcal{X} = \{\pm 1/\sqrt{n}\}^n$. If the extended Conjecture 3.2.3 holds,² then there exists no algorithm that runs in time $O(\exp(n^{1-\delta}))$ for $\delta > 0$ and certifies a bound on $M_{\mathcal{X}}(\mathbf{W})$ that is with high probability at most $2 - \epsilon$ for $\epsilon > 0$ when $\mathbf{W} \sim \text{GOE}(n)$. On the other hand, $M_{\mathcal{X}}(\mathbf{W}) \rightarrow 2P_* \approx 1.526$ in probability.*

Proof. The final statement is the deep result of [Gue03, Tal06]; see also [Pan13] for a textbook treatment and [CR02] for numerics justifying the number we give. A simpler upper bound of $2\sqrt{2/\pi}$ follows from the Fernique-Sudakov inequality (Corollary 3.12 of [LT13]).

²Really, we only need the conjecture to apply to the particular Wishart spiked matrix model described in the proof. The same also applies to the remaining results in this section.

Let $\pi = \text{Unif}(\{1, -1\}) \in \mathcal{P}(\mathbb{R})$ be the Rademacher distribution. The mean of π is zero, the variance is 1, and the distribution is bounded, so, by Proposition 5.1.7, $(\mathcal{P}_n^\pi) = (\text{Unif}(\{\pm 1/\sqrt{n}\}^n))$ forms a tame spike prior. This spike prior is also β -good for any $\beta > -1$, as $\mathbf{x} \sim \mathcal{P}_n^\pi$ has $\|\mathbf{x}\| = 1$ almost surely. Lastly, we have $\mathbf{x} \in \mathcal{X} = \{\pm 1/\sqrt{n}\}^n$ almost surely when $\mathbf{x} \sim \mathcal{P}_n^\pi$.

Suppose there exists an algorithm certifying a bound of at most $2 - \epsilon$ on $M_{\mathcal{X}}(\mathbf{W})$ with high probability when $\mathbf{W} \sim \text{GOE}(n)$ and running in time $T(n)$. Then, by Corollary 2.3.4, there exist $\beta \in (-1, 0)$, $\gamma > 1$, and an algorithm that can distinguish (\mathbb{P}_n) from (\mathbb{Q}_n) in the Wishart spiked matrix model with parameters $(\beta, \gamma, \mathcal{P}_n^\pi)$ in time $T(n) + O(1)$. On the other hand, since $\beta^2/\gamma < 1$ for these parameters, by Theorem 5.2.12 we have $\|L_n^{\leq D}\| = O(1)$ for any $D = D(n) = o(n)$ in this model. Therefore, assuming Conjecture 3.2.3 holds, we must have $T(n) \gtrsim \exp(n^{1-\delta})$ for any $\delta > 0$. \square

Clearly, the argument also goes through unchanged for any \mathcal{X} such that there exists some π satisfying the hypotheses of Proposition 5.1.7 (in the rank-one case) such that when $\mathbf{x} \sim \pi^{\otimes n}$ then $\mathbf{x}/\sqrt{n} \in \mathcal{X}$ with high probability. For example, this treats the constraint set $\mathcal{X} = \{\mathbf{x} \in \mathbb{S}^{n-1} : \|\mathbf{x}\|_0 \leq \rho n\}$ for any $\rho > 0$, showing the same result.

Next, we present an example where the additional flexibility of the “planting near \mathcal{X} ” reduction in Corollary 2.3.3 is useful. This addresses a question posed by Montanari and Richard in [MR15] concerning the performance of semidefinite programming relaxations for non-negative PCA.

Corollary 5.3.7 (Non-negative PCA [BKW20a]). *Let $\mathcal{X} = \mathbb{R}_+^n \cap \mathbb{S}^{n-1}$. If the extended Conjecture 3.2.3 holds, then there exists no algorithm that runs in time $\exp(n^{1-\delta})$ for $\delta > 0$ and certifies a bound on $M_{\mathcal{X}}(\mathbf{W})$ that is with high probability at most $2 - \epsilon$ for $\epsilon > 0$ when $\mathbf{W} \sim \text{GOE}(n)$. On the other hand, $M_{\mathcal{X}}(\mathbf{W}) \rightarrow \sqrt{2}$ in probability.*

Proof. The final statement is the result of Theorem 2 of [RM14]; the upper bound alone

follows from a straightforward application of the Fernique-Sudakov inequality.

Fix some $\rho \in (0, 1)$. Let $\pi \in \mathcal{P}(\mathbb{R})$ be the centered Bernoulli distribution, that which samples $x \sim \pi$ as

$$x = \begin{cases} \sqrt{\frac{1-\rho}{\rho}} & \text{with probability } \rho \\ -\sqrt{\frac{\rho}{1-\rho}} & \text{with probability } 1 - \rho \end{cases} \quad (5.37)$$

The mean of π is zero, the variance is 1, and the distribution is bounded, so, by Proposition 5.1.7, (\mathcal{P}_n^π) forms a tame spike prior.

\mathcal{P}_n is not itself β -good, for any $\beta > -1$, since under $x \sim \mathcal{P}_n$ we may have $\|x\|^2 = \frac{1-\rho}{\rho}$ at the largest. However, $\|x\|^2$ is an average of n i.i.d. bounded random variables, and its mean is 1, so by the law of large numbers $\|x\|^2 \rightarrow 1$ in probability and thus \mathcal{P}_n is weakly β -good.

For $x \sim \mathcal{P}_n$, let x' have entries $x'_i = 0 \vee x_i \geq 0$. Then, $x' \in \mathcal{X} = \mathbb{R}_+^n$ almost surely. We also have $\langle x, x' \rangle = \frac{1-\rho}{\rho} \cdot \#\{i \in [n] : x_i > 0\}/n \rightarrow 1 - \rho$ in probability. Thus, for any $\beta > -1$ and $\delta > 0$, taking $\rho = \delta/2$ above we will have that $\text{trunc}_\beta(\mathcal{P}_n^\pi)$ satisfies the hypothesis of Corollary 2.3.3, with $K = 1$.

We may then conclude as before: suppose there exists an algorithm certifying a bound of at most $2 - \epsilon$ on $M_X(\mathbf{W})$ with high probability when $\mathbf{W} \sim \text{GOE}(n)$ and running in time $T(n)$. Then, by Corollary 2.3.4, there exist $\beta \in (-1, 0)$, $\gamma > 1$, and an algorithm that can distinguish (\mathbb{P}_n) from (\mathbb{Q}_n) in the Wishart spiked matrix model with parameters $(\beta, \gamma, \mathcal{P}_n^\pi)$ in time $T(n) + O(1)$. On the other hand, since $\beta^2/\gamma < 1$ for these parameters, by Theorem 5.2.12 we have $\|L_n^{\leq D}\| = O(1)$ for any $D = D(n) = o(n)$ in this model (noting that applying trunc_β to the spike prior does not change the spiked Wishart model). Therefore, assuming Conjecture 3.2.3 holds, we must have $T(n) \gtrsim \exp(n^{1-\delta})$ for any $\delta > 0$. \square

Lastly, we treat an example where both the constrained PCA problem and the associated Wishart spiked matrix model have rank greater than one. This is the Gaussian analog of the coloring problem discussed in Section 2.5 (or, more broadly, of the problem of finding the largest “ k -cut” or “ k -multisection” in a graph, which is just an improper coloring with the

number of monochromatic edges minimized), and also is the analog of the SK Hamiltonian with “Potts spins.” Under the latter interpretation this model is prominent in the statistical physics literature [ES83, GKS85] (see also [Pan18] for a modern mathematically-rigorous analysis), and through that connection it has been related to the problem of finding maximum multisections in sparse random graphs in the same way that the SK Hamiltonian is related to maximum cuts [Sen18].

Corollary 5.3.8 (Potts spin glass Hamiltonian [BBK⁺20]). *Let $k \geq 2$. Let $\mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{S}^{k-1}$ be unit vectors pointing to the vertices of an equilateral simplex, such that $\|\mathbf{v}_i\| = 1$ and $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = -\frac{1}{k-1}$ whenever $i \neq j$. Let $\mathcal{X} \subset \mathbb{R}^{n \times (k-1)}$ be the set of matrices \mathbf{X} all of whose rows equal $\sqrt{\frac{k-1}{n}} \mathbf{v}_i$ for some $i \in [k]$. If the extended Conjecture 3.2.3 holds, then there exists no algorithm that runs in time $O(\exp(n^{1-\delta}))$ for $\delta > 0$ and certifies a bound on $M_{\mathcal{X}}(\mathbf{W})$ that is with high probability at most $2(k-1) - \epsilon$ for $\epsilon > 0$ when $\mathbf{W} \sim \text{GOE}(n)$.*

We note that $2(k-1)$ is indeed the spectral bound, since $\|\mathbf{X}\|_F^2 = k-1$ for any $\mathbf{X} \in \mathcal{X}$. Also, the case $k = 2$ recovers our above result for the SK Hamiltonian. What is more typically called the “Potts spin glass Hamiltonian” is the function $H(\sigma) = \sum_{i,j=1}^n \mathbb{1}\{\sigma(i) = \sigma(j)\} W_{ij}$ over $\sigma \in [k]^n$, but this is merely a rescaling and negligible shift depending only on \mathbf{W} of the quantity optimized by $M_{\mathcal{X}}(\mathbf{W})$. A Parisi formula analogous to the SK Hamiltonian has been established rigorously in this model by [Pan18], but the numerical value of the ground state energy does not appear to have been studied extensively.

Proof. The proof is essentially identical to that of Corollary 5.3.6, with suitable higher-rank notions substituted in as needed. Let $\pi = \text{Unif}(\{\sqrt{k-1}\mathbf{v}_1, \dots, \sqrt{k-1}\mathbf{v}_k\}) \in \mathcal{P}(\mathbb{R}^{k-1})$. The mean of π is zero, the covariance is $\frac{k-1}{k} \sum_{i=1}^k \mathbf{v}_i \mathbf{v}_i^\top = \mathbf{I}_{k-1}$, and the distribution is bounded, so, by Proposition 5.1.7, (\mathcal{P}_n^π) forms a tame spike prior. This spike prior is also weakly β -good for any $\beta > -1$, as $\mathbf{X} \sim \mathcal{P}_n^\pi$ has $\mathbf{X}^\top \mathbf{X} = \frac{k-1}{n} \sum_{i=1}^n \mathbf{v}_{\sigma(i)} \mathbf{v}_{\sigma(i)}^\top$ for $\sigma(i) \sim \text{Unif}([k])$ independently, so $\mathbf{X}^\top \mathbf{X} \rightarrow (k-1) \mathbb{E} \mathbf{v}_{\sigma(i)} \mathbf{v}_{\sigma(i)}^\top = \mathbf{I}_{k-1}$ in operator norm in probability, and

thus $\|\mathbf{X}\| \rightarrow 1$ in probability. We also have $\|\mathbf{X}\|_F^2 = k - 1$ and $\mathbf{X} \in \mathcal{X}$ almost surely when $\mathbf{X} \sim \mathcal{P}_n^\pi$.

Suppose there exists an algorithm certifying a bound of at most $2(k - 1) - \epsilon$ on $M_{\mathcal{X}}(\mathbf{W})$ with high probability when $\mathbf{W} \sim \text{GOE}(n)$ and running in time $T(n)$. Then, by Corollary 2.3.4, there exist $\beta \in (-1, 0)$, $\gamma > 1$, and an algorithm that can distinguish (\mathbb{P}_n) from (\mathbb{Q}_n) in the Wishart spiked matrix model with parameters $(\beta, \gamma, \text{trunc}_\beta(\mathcal{P}_n^\pi))$ in time $T(n) + O(1)$, and so the same holds with parameters $(\beta, \gamma, \mathcal{P}_n^\pi)$ by our remarks following Definition 5.3.4. On the other hand, since $\beta^2/\gamma < 1$ for these parameters, by Theorem 5.2.12 we have $\|L_n^{\leq D}\| = O(1)$ for any $D = D(n) = o(n)$ in this model. Therefore, assuming Conjecture 3.2.3 holds, we must have $T(n) \gtrsim \exp(n^{1-\delta})$ for any $\delta > 0$. \square

5.4 NEF-QVF MODELS

Finally, we give some ancillary applications of the tools we have developed for working with low-degree polynomial algorithms in NEF-QVF models. These do not concern certification, but give new results for related models of general interest.

5.4.1 STOCHASTIC BLOCK MODEL

First, we show how to use our results for comparing NEF-QVF models (Section 4.3.3) to recover the well-known *Kesten-Stigum* computational threshold in the symmetric stochastic block model with k communities (see, e.g., [Abb17, Moo17] for surveys of this model). We also sharpen previous low-degree lower bounds of [HS17] for this model, showing that polynomials of degree $\Omega(n)$ are required to distinguish \mathbb{P}_n from \mathbb{Q}_n in the conjectural hard regime, and give a simpler argument based on the general principles from Part I.

Definition 5.4.1 (Stochastic block model). *The symmetric stochastic block model with k communities and parameters $a, b > 0$ is specified by the following null and planted distributions*

over graphs on vertex set $[n]$.

- Under \mathbb{Q}_n , each edge occurs independently with probability $\frac{a+b(k-1)}{kn}$.
- Under \mathbb{P}_n , assign each vertex $i \in [n]$ a random community $\sigma(i) \sim \text{Unif}([k])$ independently. Then, an edge between vertices $i, j \in [n]$ occurs with probability $\frac{a}{n}$ if $\sigma(i) = \sigma(j)$ and with probability $\frac{b}{n}$ if $\sigma(i) \neq \sigma(j)$.

Note that the model is calibrated so that the average degree of any vertex is $\frac{a+b(k-1)}{k}$ under either the null or planted distribution.

Corollary 5.4.2. *If $(a-b)^2 < k(a+(k-1)b)$, then there exists a constant $c > 0$ depending only on a, b , and k such that, whenever $D \leq cn$, then $\|L_n^{\leq D}\| = O(1)$ in the symmetric stochastic block model.*

Proof. We may rewrite the null model as making independent observations from $\text{Ber}(\mu)$ with $\mu = \frac{a+(k-1)b}{kn}$. For the planted model, as we did for the Potts spin glass, let us write $\mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{S}^{k-1}$ for the unit vectors pointing to the vertices of an equilateral simplex, so that $\|\mathbf{v}_i\| = 1$ and $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = -\frac{1}{k-1}$ whenever $i \neq j$. Then, we may rewrite the planted model as making independent observations from $\text{Ber}(x_{\{i,j\}})$, where

$$x_{\{i,j\}} = \underbrace{\frac{a+b(k-1)}{kn}}_{\mu} + \frac{(k-1)(a-b)}{kn} \langle \mathbf{v}_{\sigma(i)}, \mathbf{v}_{\sigma(j)} \rangle. \quad (5.38)$$

Let us write $\mathbf{X} \in \mathbb{R}^{n \times k}$ for the matrix whose i th row is $\sqrt{\frac{k-1}{n}} \mathbf{v}_{\sigma(i)}$. Then, the z -scores are

$$z_{\{i,j\}} = \frac{x_{\{i,j\}} - \mu}{\sqrt{\mu(1-\mu)}} = \sqrt{n} \frac{a-b}{\sqrt{k(a+(k-1)b)}} \frac{1}{\sqrt{1-\mu}} (\mathbf{X} \mathbf{X}^\top)_{ij}. \quad (5.39)$$

Let us suppose that we also make n further observations, corresponding to the formal “diagonal” case $i = j$ above, with the same distribution $\text{Ber}(\mu)$ under \mathbb{Q}_n , and with $x_{\{i,i\}} =$

$2(\mu + \frac{(k-1)(a-b)}{kn})$. Clearly including these further observations will only increase $\|L_n^{\leq D}\|$ (in the orthogonal polynomial decomposition of the squared norm, this change will only add new non-negative terms). With this adjustment, the overlap between the z -scores of two independent \mathbf{x} drawn as above is

$$R_n := \langle \mathbf{z}^1, \mathbf{z}^2 \rangle = \frac{n}{2} \frac{(a-b)^2}{k(a+(k-1)b)} \frac{1}{1-\mu} \text{tr}(\mathbf{X}^1 \mathbf{X}^{1\top} \mathbf{X}^2 \mathbf{X}^{2\top}), \quad (5.40)$$

for \mathbf{X}^i drawn independently as functions of two draws σ^i of the community labels.

We note that the Bernoulli NEF-QVF has $v_2 = -1$, as its variance function is $V(\mu) = \mu(1-\mu) = -\mu^2 + \mu$ (recall that v_2 is the quadratic coefficient in this polynomial). Then, by Theorem 4.3.11, we have

$$\|L_n^{\leq D}\|^2 = \mathbb{E} \phi_D^{\text{Mor}}(R_n; -1) \quad (5.41)$$

and by the ‘‘channel monotonicity’’ comparison of Theorem 4.3.14 we may bound

$$\leq \mathbb{E} \phi_D^{\text{Mor}}(R_n; 0) \quad (5.42)$$

and we recall that $\phi_D^{\text{Mor}}(\cdot; 0) = \phi_D^{\text{Wig}}$ since the NEF-QVF with $v_2 = 0$ is the Gaussian Wigner family, so

$$= \mathbb{E} \phi_D^{\text{Wig}} \left(\frac{n}{2} \frac{(a-b)^2}{k(a+(k-1)b)} \frac{1}{1-\mu} \text{tr}(\mathbf{X}^1 \mathbf{X}^{1\top} \mathbf{X}^2 \mathbf{X}^{2\top}) \right) \quad (5.43)$$

Finally, we note that $\mu \rightarrow 0$ as $n \rightarrow \infty$, so under our assumptions for sufficiently large n there will be some $\hat{\lambda} \in (0, 1)$ such that

$$\leq \mathbb{E} \phi_D^{\text{Wig}} \left(\frac{\hat{\lambda}n}{2} \cdot \text{tr}(\mathbf{X}^1 \mathbf{X}^{1\top} \mathbf{X}^2 \mathbf{X}^{2\top}) \right). \quad (5.44)$$

This is precisely the norm of the LDLR in the Wigner spiked matrix model (Definition 5.2.9) with a rank- $(k - 1)$ spike prior (\mathcal{P}_n) given by the law of \mathbf{X}^i and with signal-to-noise ratio $\lambda = \sqrt{\hat{\lambda}n}$. By our discussion in the proof of Corollary 5.3.8 the spike prior (\mathcal{P}_n) is tame, so by Theorem 5.2.10 we have $\|L_n^{\leq D}\| = O(1)$ under the stated conditions, as $\hat{\lambda} < 1$. \square

5.4.2 A NON-GAUSSIAN MATRIX PCA MODEL

Finally, we show how some of the ideas developed for NEF-QVF models can be adapted to treat a Wigner spiked matrix model with non-Gaussian noise. We treat this model as a “stress test” of the low-degree method and Conjecture 3.2.3, since, as we will detail below, it is already known from [PWBM18] that the model we consider has no hard regime. However, we will see that the low-degree analysis still sheds light on what kinds of tests can or cannot successfully distinguish these models, and perhaps more broadly on the power and limitations of low-degree polynomials.

Though we work with a distribution borrowed from an NEF-QVF, we will still be interested in noise applied additively like in the Gaussian Wigner model, rather than within an NEF-QVF as for the kin-spiked NEF-QVF model of Definition 4.3.1, so we make the following alternative definition for this application.

Definition 5.4.3 (Additively-spiked NEF-QVF model). *In the same setting as Definition 4.3.1 (the kin-spiked model) but with \mathcal{P}_n now a probability measure over $\mathbb{R}^{N(n)}$, define:*

- Under \mathbb{Q}_n , draw $y_i \sim \tilde{\rho}_{\mu_{n,i}}$ independently for $i \in [N(n)]$.
- Under \mathbb{P}_n , first draw $\mathbf{x} \sim \mathcal{P}_n$ and $z_i \sim \tilde{\rho}_{\mu_{n,i}}$ independently for $i \in [N(n)]$, and observe $y_i = x_i + z_i$.

In particular, we will study a model with noise distributed according to ρ^{sech} the probabil-




Figure 5.2: Hyperbolic secant distribution. We plot the density $w(x)$ of the hyperbolic secant distribution used in Theorem 5.4.7, showing that it is a smoothed variant of the better-known Laplace density $\frac{1}{2} \exp(-|x|)$.

ity measure on \mathbb{R} which has the following density $w(x)$ with respect to Lebesgue measure:

$$w(x) := \frac{1}{2 \cosh(\pi x/2)} = \frac{1}{2} \operatorname{sech}(\pi x/2). \quad (5.45)$$

This density belongs to the rather obscure class of “generalized hyperbolic secant” NEFs mentioned in Table 4.2. It may be viewed as a smoothing of the Laplace distribution; see Figure 5.2.³

We next specify the spiked matrix model we will study. For the sake of convenience here and in some further discussion of open problems inspired by this model that we give in Section A.5, we give a definition rather orthogonal to our previous treatment, fixing the spike distribution and considering varying noise rather than vice-versa.

Definition 5.4.4 (Rademacher-spiked Wigner matrix models). *Given a probability measure ρ over \mathbb{R} and $\lambda > 0$, the Wigner Rademacher-spiked matrix model with parameters (ρ, λ) is*

³This density has some other remarkable mathematical properties: (1) like the Gaussian density, up to dilation $w(x)$ is its own Fourier transform, and (2) $w(x)$ is the Poisson kernel over the strip $\{z : \operatorname{Im}(z) \in [-1, 1]\} \subset \mathbb{C}$.

specified by the following probability distributions over $\mathbb{R}^{\binom{n}{2}}$:

- Under \mathbb{Q}_n , we draw $Y_{\{i,j\}} \sim \rho$ independently for all $i < j$.
- Under \mathbb{P}_n , we first draw $\mathbf{x} \sim \text{Unif}(\{\pm 1\}^n)$ and $Y_{\{i,j\}}^{(0)} \sim \rho$ independently for all $i < j$, and then set $Y_{\{i,j\}} = \frac{\lambda}{\sqrt{n}} \mathbf{x}_i \mathbf{x}_j + Y_{\{i,j\}}^{(0)}$.

We omit the diagonal observations with $i = j$ for the sake of convenience; it is straightforward but tedious to adapt the argument to include these, but with this change we would have $\mathbf{Y} = \frac{\lambda}{\sqrt{n}} \mathbf{x} \mathbf{x}^\top + \mathbf{Y}^{(0)}$ as symmetric matrices under \mathbb{P}_n .

In Section 2.2 we presented results concerning testing in such a model, particularly using the *PCA test* of thresholding the largest eigenvalue of \mathbf{Y} , for the same model with ρ the standard Gaussian measure. How, if at all, does that picture change for a different noise distribution? The following characterizes two testing algorithms related to computing the largest eigenvalue. For \mathbf{Y} itself, for sufficiently large λ , the largest eigenvalue undergoes the same pushout effect as in the Gaussian model under \mathbb{P}_n and becomes larger than the typical largest eigenvalue under \mathbb{Q}_n . It turns out, however, that it is suboptimal to merely compute and threshold the largest eigenvalue of \mathbf{Y} ; instead, the optimal algorithm is to first apply an entrywise transformation and only then compute and threshold the largest eigenvalue.

Proposition 5.4.5 (Better-than-BBP testing [CDMF09, PWBM18]). *Define $\lambda_* := 2\sqrt{2}/\pi \approx 0.9$. In all of the statements below we refer to \mathbb{P}_n and \mathbb{Q}_n in the Wigner Rademacher-spiked matrix model with parameters $(\rho^{\text{sech}}, \lambda)$.*

- If $\lambda > 1$, then \mathbb{P}_n may be distinguished from \mathbb{Q}_n in polynomial time by the PCA test,

$$f^{\text{PCA}}(\mathbf{Y}) := \begin{cases} p & \text{if } \frac{1}{\sqrt{n}} \lambda_{\max}(\mathbf{Y}) \geq \frac{1}{2}(2 + \lambda + \lambda^{-1}), \\ q & \text{otherwise.} \end{cases} \quad (5.46)$$

- If $\lambda < 1$, then f^{PCA} fails to distinguish \mathbb{P}_n from \mathbb{Q}_n .

- If $\lambda > \lambda_*$, then \mathbb{P}_n may be distinguished from \mathbb{Q}_n in polynomial time by the pre-transformed PCA test:

$$f^{\text{tPCA}}(\mathbf{Y}) := \begin{cases} P & \text{if } \frac{1}{\sqrt{n}} \lambda_{\max} \left(\frac{\pi}{2} \tanh \left(\frac{\pi}{2} \mathbf{Y} \right) \right) \geq \frac{1}{2} (2\lambda_* + \lambda_*^2 \cdot \lambda + \lambda^{-1}), \\ Q & \text{otherwise.} \end{cases} \quad (5.47)$$

Here, $\tanh(\cdot)$ is applied entrywise to the matrix argument.

- If $\lambda < \lambda_*$, then \mathbb{P}_n and \mathbb{Q}_n are statistically indistinguishable.

The threshold λ_* is related to the *Fisher information* in the family of translates of ρ^{sech} as $\lambda_* = \left(\int_{-\infty}^{\infty} w'(x)^2 / w(x) dx \right)^{-1/2}$, and the optimal entrywise transformation is the logarithmic derivative $\frac{\pi}{2} \tanh(\frac{\pi}{2} x) = -w'(x)/w(x)$; the results of [PWBM18] show that both relationships hold for optimal tests in non-Gaussian spiked matrix models for a broad class of noise measures.

Let us consider how these facts interact with low-degree predictions. Heuristically speaking, while low-degree polynomials can approximate the test f^{PCA} via the power method, the transcendental entrywise $\tanh(\cdot)$ transformation used by f^{tPCA} seems rather ill-suited to low-degree polynomials. We show below that, indeed, if we attempt to carry out the low-degree prediction for this problem while bounding the *entrywise degree* of the polynomials involved—the greatest power with which any given entry of \mathbf{Y} can appear—then we obtain an incorrect threshold. Loosely speaking, this suggests that some analytic computation like the transcendental $\tanh(\cdot)$ operation is in fact *necessary* to obtain an optimal test.

Definition 5.4.6 (Entrywise degree). *For a polynomial $p \in \mathbb{R}[y_1, \dots, y_N]$, write $\deg_i(p)$ for the greatest power with which y_i occurs in a monomial having non-zero coefficient in p .*

Theorem 5.4.7. *Suppose $D \in \mathbb{N}$ and $0 < \lambda < \lambda_* + \frac{1}{20D}$. Then,*

$$\limsup_{n \rightarrow \infty} \left\{ \begin{array}{l} \text{maximize } \mathbb{E}_{\mathbf{Y} \sim \mathbb{P}_n} f_n(\mathbf{Y}) \\ \text{subject to } f_n \in \mathbb{R}[\mathbf{Y}], \\ \text{deg}_{\{i,j\}}(f_n) \leq D \text{ for all } \{i,j\} \in \binom{[n]}{2}, \\ \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}_n} f_n(\mathbf{Y})^2 = 1 \end{array} \right\} < +\infty. \quad (5.48)$$

That is, when we restrict our attention to polynomials of entrywise degree at most D a constant not growing with n , the apparent computational threshold suggested by the corresponding low-degree calculation shifts by $\Omega(1/D)$ from the true value.

This limitation applies, for example, to an approach suggested by [DHS20]. The authors propose to build tests and estimators for spiked matrix models that remain effective under heavy-tailed noise distributions by using polynomials that sum over monomials indexed by *self-avoiding walks* on the matrix \mathbf{Y} . In particular, they show that, for $\lambda > 1$ —the optimal “BBP threshold” for Gaussian noise—such polynomials can successfully distinguish \mathbb{P}_n from \mathbb{Q}_n in Wigner Rademacher-spiked matrix models with parameters (ρ, λ) for a wide variety of measures ρ , ranging from Gaussian ρ to very heavy-tailed ρ for which f^{PCA} fails severely. However, our result implies that, since these polynomials have entrywise degree 1 (that is, they are multilinear), such polynomials (and many generalizations thereof to higher but bounded entrywise degree) *cannot* distinguish for all $\lambda > \lambda_*$, and thus are suboptimal for this model.

To prove Theorem 5.4.7, we will need to develop some analogs for the additively-spiked model to the tools we developed earlier in Section 4.3.1 for working with the orthogonal polynomials of NEF-QVFs for the kin-spiked model. We return here to the setting and terminology of Section 4.3. First, we write the precise generating function relation between the likelihood ratio and the orthogonal polynomials (as given in [Mor82]).

Proposition 5.4.8 (Generating function). *Let $\mu \in \Omega$ and write $\psi(\eta) = \mathbb{E}_{x \sim \tilde{\rho}_\mu}[\exp(\eta x)]$. Then,*

$$\sum_{k \geq 0} \frac{z_\mu(t)^k}{k!} p(y; \mu) = \exp(y(\psi')^{-1}(t) - \psi((\psi')^{-1}(t))). \quad (5.49)$$

Note that here we are “rebasings” the NEF-QVF to have $\tilde{\rho}_\mu$ as the base measure by our definition of $\psi(\cdot)$. One may view this result as generalizing to NEF-QVFs the generating function $\exp(ty - \frac{1}{2}t^2)$ for Hermite polynomials (Proposition 4.1.14). The key property of such generating functions is that y appears *linearly* in the exponential. (Indeed, as early as 1934, Meixner had essentially discovered the NEF-QVFs, albeit only recognizing their significance in terms of this distinctive property of their orthogonal polynomials [Mei34, Lan75].)

This linearity allows us to prove an “addition formula,” expanding the translation operator in orthogonal polynomials.

Definition 5.4.9 (Translation polynomials). *Let $\tau_k(y; \mu) \in \mathbb{R}[y]$ be defined by the generating function*

$$\sum_{k \geq 0} \frac{z_\mu(t)^k}{k!} \tau_k(y; \mu) := \exp(y(\psi')^{-1}(t)). \quad (5.50)$$

Also, define the normalized versions

$$\hat{\tau}(y; \mu) := \frac{1}{V(\mu)^{k/2} \sqrt{a_k(v_2)}} \tau_k(x; \mu). \quad (5.51)$$

Proposition 5.4.10 (Addition formula). *For all $x, y \in \mathbb{R}$ and $\mu \in \Omega$,*

$$p_k(x + y; \mu) = \sum_{\ell=0}^k \binom{k}{\ell} \tau_{k-\ell}(x; \mu) p_\ell(y; \mu). \quad (5.52)$$

Proof. This follows from expanding the generating function (5.49) at $x + y$ as a product of two exponential generating functions. □

Finally, we obtain the additively-spiked version of Corollary 4.3.6 by taking expectations

and using the orthogonality of the p_k .

Proposition 5.4.11 (Additively-spiked expectation). *For all $k \in \mathbb{N}$, $\mu \in \Omega$, and $x \in \mathbb{R}$,*

$$\mathbb{E}_{y \sim \tilde{\rho}_\mu} \hat{p}_k(x + y; \mu) = \hat{\tau}_k(x; \mu). \quad (5.53)$$

Proof. This follows from taking expectations on either side of (5.52), observing that the only non-zero term is for $\ell = 0$ by the orthogonality of the p_ℓ , and noting that $p_0(y; \mu) = 1$. \square

Following the argument of Lemma 4.3.7 for the additively-spiked model and using Proposition 5.4.11 instead of Corollary 4.3.6 gives the following result for the orthogonal polynomial components of the likelihood ratio.

Lemma 5.4.12 (Components under additive spiking). *In the additively-spiked NEF-QVF model, for all $\mathbf{k} \in \mathbb{N}^N$,*

$$\langle L_n, \hat{P}_{\mathbf{k}}(\cdot; \boldsymbol{\mu}_n) \rangle = \mathbb{E}_{x \sim P_n} \left[\prod_{i=1}^N \hat{\tau}_{k_i}(x_i; \mu_{n,i}) \right]. \quad (5.54)$$

We now analyze the translation polynomials τ_k from Definition 5.4.9 for the NEF generated by ρ^{sech} . First, note that the mean and variance of ρ^{sech} are $\mu = 0$ and $V(0) = 1$, and more generally the variance function in the generated NEF is $V(\mu) = \mu^2 + 1$ (per Table 4.2), where in particular the quadratic coefficient is $v_2 = 1$. Thus the associated normalizing constants are $a_k(v_2) = (k!)^2$ and $\hat{a}_k(v_2) = k!$.

Recall that the translation polynomials admit a generating function expressed in terms of the cumulant generating function of ρ^{sech} . We therefore compute

$$\psi(\theta) := \mathbb{E}_{y \sim \rho^{\text{sech}}} \exp(\theta y) = \frac{1}{2} \int_{-\infty}^{\infty} \text{sech}\left(\frac{\pi y}{2}\right) \exp(\theta y) dy = -\log(\cos \theta), \quad (5.55)$$

$$\psi'(\theta) = \tan(\theta), \quad (5.56)$$

whereby the translation polynomials for $\mu = 0$ (the mean of ρ^{sech}) have the generating func-

tion

$$\sum_{k \geq 0} \frac{t^k}{k!} \tau_k(y; 0) = \sum_{k \geq 0} t^k \hat{\tau}_k(y; 0) = \exp(y \tan^{-1}(t)). \quad (5.57)$$

Before proceeding, we also establish some preliminary bounds on the coefficients and values of these polynomials. We denote by $[x^\ell](p(x))$ the coefficient of x^ℓ in a polynomial or formal power series $p(x)$.

Proposition 5.4.13. *For all $k \geq 1$ and $\ell \geq 0$,*

$$|[x^\ell](\hat{\tau}_k(x))| \leq \mathbb{1}\{k \equiv \ell \pmod{2}, \ell > 0\} \frac{(2 \log(ek))^{\ell-1}}{k \ell!}. \quad (5.58)$$

Proof. Expanding the generating function, we have

$$[x^\ell](\hat{\tau}_k(x)) = [t^k x^\ell](\exp(x \tan^{-1}(t))) = \frac{1}{\ell!} [t^k]((\tan^{-1}(t))^\ell). \quad (5.59)$$

If $k \geq 1$ and $\ell = 0$, then this is zero. Since the coefficients in the Taylor series of $\tanh^{-1}(t)$ are $[t^k](\tanh^{-1}(t)) = \mathbb{1}\{k \equiv 1 \pmod{2}\} (-1)^{(k-1)/2} / k$, we may bound

$$|[x^\ell](\hat{\tau}_k(x))| \leq \mathbb{1}\{k \equiv \ell \pmod{2}, \ell > 0\} \frac{1}{\ell!} \sum_{\substack{a_1, \dots, a_\ell \geq 1 \\ a_1 + \dots + a_\ell = k}} \frac{1}{\prod_{i=1}^{\ell} a_i}. \quad (5.60)$$

$\underbrace{\hspace{10em}}_{c(k, \ell)}$

We now show that $c(k, \ell) \leq (2 \log(ek))^{\ell-1} / k$ by induction on ℓ . Since $c(k, 1) = 1/k$, the base case holds. We note the bound on harmonic numbers

$$\sum_{a=1}^k \frac{1}{a} \leq \log(ek) \text{ for all } k \geq 1. \quad (5.61)$$

Supposing the result holds for $c(k, \ell - 1)$, we expand $c(k, \ell)$ according to the value that a_ℓ

takes:

$$\begin{aligned}
c(k, \ell) &\leq \sum_{a=1}^{k-1} \frac{1}{a} c(k-a, \ell-1) \\
&\leq (2 \log(ek))^{\ell-2} \sum_{a=1}^{k-1} \frac{1}{a} \cdot \frac{1}{k-a} && \text{(inductive hypothesis)} \\
&\leq \frac{(2 \log(ek))^{\ell-2}}{k} \sum_{a=1}^{k-1} \left(\frac{1}{a} + \frac{1}{k-a} \right) \\
&\leq \frac{(2 \log(ek))^{\ell-2}}{k} \cdot 2 \log(ek), && \text{(by (5.61))}
\end{aligned}$$

completing the argument. □

This yields the following pointwise bound. As we will ultimately be evaluating this on quantities of order $O(n^{-1/2})$, what is most important to us is the precision for very small arguments.

Corollary 5.4.14. *For all $k \geq 1$ and $x > 0$,*

$$|\hat{\tau}_k(x)| \leq \begin{cases} x \cdot \frac{1}{k} \cdot (ek)^{2x} & \text{if } k \text{ odd,} \\ x^2 \cdot \frac{2 \log(ek)}{k} \cdot (ek)^{2x} & \text{if } k \text{ even.} \end{cases} \quad (5.62)$$

Proof. Write $\ell_0 = 1$ if k is odd and $\ell_0 = 2$ if k is even. We bound by Proposition 5.4.13,

$$\begin{aligned}
|\hat{\tau}_k(x)| &\leq \frac{1}{k} \sum_{\ell=\ell_0}^k \frac{(2 \log(ek))^{\ell-1}}{\ell!} x^\ell \\
&\leq \frac{x^{\ell_0} (2 \log(ek))^{\ell_0-1}}{k} \sum_{\ell=\ell_0}^k \frac{(2 \log(ek)x)^{\ell-\ell_0}}{(\ell-\ell_0)!} \\
&\leq \frac{x^{\ell_0} (2 \log(ek))^{\ell_0-1}}{k} \exp((2 \log(ek)x), \quad (5.63)
\end{aligned}$$

and the result follows upon rearranging. □

We now proceed with the proof of our main result.

Proof of Theorem 5.4.7. First, applying Lemma 5.4.12 to the hyperbolic secant spiked matrix model, the coefficients of the likelihood ratio are given by, for any $\mathbf{k} \in \mathbb{N}^{\binom{[n]}{2}}$,

$$\langle L_n(\mathbf{Y}), \hat{P}_{\mathbf{k}} \rangle = \mathbb{E}_{\mathbf{X} \sim \mathcal{P}_n} \left[\prod_{1 \leq i < j \leq n} \hat{\tau}_{k_{\{i,j\}}}(X_{\{i,j\}}) \right] = \mathbb{E}_{\mathbf{x} \sim \text{Unif}(\{\pm 1\}^n)} \left[\prod_{1 \leq i < j \leq n} \hat{\tau}_{k_{\{i,j\}}}\left(\frac{\lambda}{\sqrt{n}} \mathbf{x}_i \mathbf{x}_j\right) \right]. \quad (5.64)$$

First, we observe that our specific choice of $\mathbf{x} \in \{\pm 1\}^n$ allows an interesting further simplification: thanks to this choice, we can decouple the dependence of the components of L_n on λ from the dependence on \mathbf{x} . Note that, by the generating function identity (5.57), for all $k \geq 0$ we have that $\tau_k(\mathbf{x})$ contains only monomials of the same parity as k . Therefore, for all $\mathbf{k} \in \mathbb{N}^{\binom{[n]}{2}}$, we have

$$\langle L_n, \hat{P}_{\mathbf{k}} \rangle = \prod_{i < j} \hat{\tau}_{k_{ij}}\left(\frac{\lambda}{\sqrt{n}}\right) \cdot \mathbb{E}_{\mathbf{x}} \left[\prod_{i < j} (\mathbf{x}_i \mathbf{x}_j)^{k_{ij}} \right]. \quad (5.65)$$

Here and in the remainder of the proof, we write $k_{ij} = k_{\{i,j\}}$ and $i < j$ for $1 \leq i < j \leq n$ to lighten the notation. Let us also write $|\mathbf{k}|_{\infty} := \max_{i < j} k_{ij}$.

Next, we note that, since the second factor above is either 0 or 1, we may further bound

$$|\langle L_n, \hat{P}_{\mathbf{k}} \rangle| \leq \left| \prod_{i < j} \hat{\tau}_{k_{ij}}\left(\frac{\lambda}{\sqrt{n}}\right) \right| \cdot \mathbb{E}_{\mathbf{x}} \left[\prod_{i < j} (\mathbf{x}_i \mathbf{x}_j)^{k_{ij}} \right]$$

When $|\mathbf{k}|_{\infty} \leq D$, then, by Corollary 5.4.14, we may continue

$$\leq \prod_{\substack{i < j \\ k_{ij} > 0}} \frac{(eD)^{\frac{\lambda}{\sqrt{n}}}}{k_{ij}} (2 \log(e k_{ij}))^{\mathbb{1}_{\{k_{ij} \text{ even}\}}} \left(\frac{\lambda}{\sqrt{n}}\right)^{1 + \mathbb{1}_{\{k_{ij} \text{ even}\}}} \mathbb{E}_{\mathbf{x}} \left[\prod_{i < j} (\mathbf{x}_i \mathbf{x}_j)^{k_{ij}} \right]. \quad (5.66)$$

We now follow our usual strategy from previous examples. Squaring and rewriting this

as an expectation over two independent $\mathbf{x}^1, \mathbf{x}^2 \sim \text{Unif}(\{\pm 1\}^n)$ (Proposition 4.1.4), we find

$$|\langle L_n, \hat{P}_k \rangle|^2 \leq \mathbb{E} \prod_{\substack{i < j \\ k_{ij} > 0}} \frac{(eD)^{2\frac{\lambda}{\sqrt{n}}}}{k_{ij}^2} (2 \log(ek_{ij}))^{2\mathbb{1}\{k_{ij} \text{ even}\}} \left(\frac{\lambda}{\sqrt{n}}\right)^{2(1+\mathbb{1}\{k_{ij} \text{ even}\})} (x_i^1 x_i^2 x_j^1 x_j^2)^{k_{ij}}. \quad (5.67)$$

Summing over $|\mathbf{k}|_\infty \leq D$, we then find

$$\begin{aligned} & \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}|_\infty \leq D}} |\langle L_n, \hat{P}_k \rangle|^2 \\ & \leq \mathbb{E}_{\mathbf{x}^1, \mathbf{x}^2} \prod_{i < j} \left(1 + (eD)^{2\frac{\lambda}{\sqrt{n}}} \frac{\lambda^2}{n} \sum_{\substack{k=1 \\ k \text{ odd}}}^D \frac{1}{k^2} (x_i^1 x_i^2 x_j^1 x_j^2)^k + (eD)^{2\frac{\lambda}{\sqrt{n}}} \frac{\lambda^4}{n^2} \sum_{\substack{k=1 \\ k \text{ even}}}^D \frac{4 \log(ek)^2}{k^2} (x_i^1 x_i^2 x_j^1 x_j^2)^k \right) \end{aligned}$$

and, using that the x_i^a are Rademacher-valued,

$$= \mathbb{E}_{\mathbf{x}^1, \mathbf{x}^2} \prod_{1 \leq i < j \leq n} \left(1 + (eD)^{2\frac{\lambda}{\sqrt{n}}} \frac{\lambda^2}{n} x_i^1 x_i^2 x_j^1 x_j^2 \sum_{\substack{k=1 \\ k \text{ odd}}}^D \frac{1}{k^2} + (eD)^{2\frac{\lambda}{\sqrt{n}}} \frac{\lambda^4}{n^2} \sum_{\substack{k=1 \\ k \text{ even}}}^D \frac{4 \log(ek)^2}{k^2} \right)$$

Here, using that $\sum_{\ell \geq 0} \frac{1}{(2\ell+1)^2} = \frac{\pi^2}{8} = \lambda_*^{-2}$ and $\sum_{\ell \geq D/2} \frac{1}{(2\ell+1)^2} \geq \int_{D/2}^{\infty} \frac{dx}{(2x+1)^2} = \frac{1}{2D+2} \geq \frac{1}{3D}$, we may write

$$\begin{aligned} & = \mathbb{E}_{\mathbf{x}^1, \mathbf{x}^2} \prod_{1 \leq i < j \leq n} \left(1 + (eD)^{2\frac{\lambda}{\sqrt{n}}} \frac{\lambda^2}{n} x_i^1 x_i^2 x_j^1 x_j^2 \left(\frac{1}{\lambda_*^2} - \frac{1}{3D} \right) + O\left(\frac{1}{n^2}\right) \right) \\ & \leq \mathbb{E}_{\mathbf{x}^1, \mathbf{x}^2} \exp \left((eD)^{2\frac{\lambda}{\sqrt{n}}} \frac{\lambda^2}{n} \left(\frac{1}{\lambda_*^2} - \frac{1}{3D} \right) \sum_{1 \leq i < j \leq n} x_i^1 x_i^2 x_j^1 x_j^2 + O(1) \right) \\ & = \mathbb{E}_{\mathbf{x}^1, \mathbf{x}^2} \exp \left((eD)^{2\frac{\lambda}{\sqrt{n}}} \frac{\lambda^2}{2} \left(\frac{1}{\lambda_*^2} - \frac{1}{3D} \right) \frac{\langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2}{n} + O(1) \right), \end{aligned}$$

where we absorb the diagonal terms from $\langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2/n$ into the $O(1)$ term. Finally, by our assumption we have $\lambda < \lambda_* + \frac{1}{20D}$. Therefore, $\lambda^2 < \lambda_*^2 + \frac{41}{400D}$. So, $\lambda^2(\lambda_*^{-2} - \frac{1}{3D}) < 1 - \frac{1}{6D} + \frac{41}{400D} < 1 - \frac{1}{20D}$, and thus, for sufficiently large n , we will have

$$\leq \mathbb{E}_{\mathbf{x}^1, \mathbf{x}^2} \exp \left(\frac{1}{2} \left(1 - \frac{1}{20D} \right) \frac{\langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2}{n} + O(1) \right). \quad (5.68)$$

The remaining expectation is precisely the quantity arising in the application of the second moment method for contiguity to the Wigner Rademacher-spiked matrix model with Gaussian noise in [PWBM18], where it is shown that this quantity is bounded as $n \rightarrow \infty$, since the factor multiplying $\langle \mathbf{x}^1, \mathbf{x}^2 \rangle^2/n$ is strictly smaller than $\frac{1}{2}$. Thus we find

$$\limsup_{n \rightarrow \infty} \sum_{\substack{\mathbf{k} \in \mathbb{N}^N \\ |\mathbf{k}|_\infty \leq D}} |\langle L_n, \hat{P}_{\mathbf{k}} \rangle|^2 < +\infty, \quad (5.69)$$

as claimed. □

Remark 5.4.15 (A general “Rademacher trick”). *The first step in the proof above, where we take advantage of the Rademacher prior to decouple the dependence of the likelihood ratio’s components on the signal-to-noise ratio λ from that on the actual spike vector \mathbf{x} , should apply in much greater generality. Indeed, we expect a similar property to hold in any additive model where (1) the spike distribution $\mathbf{x} \sim \mathcal{P}_n$ has the property that $|x_i| = \lambda(n)$ for some constant $\lambda(n)$ for all $i \in [N(n)]$, and (2) the noise distribution is symmetric, whereby the polynomials playing the role of $\hat{\tau}_{\mathbf{k}}$ will be even polynomials for even k and odd polynomials for odd k . Thus a similar analysis is likely possible in a wide range of models with “flat” signals \mathbf{x} , reducing the low-degree analysis to analytic questions about $\hat{\tau}_{\mathbf{k}}$.*

Part II

A Geometric Perspective on Sum-of-Squares

Lower Bounds

6 | BASIC NOTIONS AND THE GRAMIAN

VIEWPOINT

We will focus in the second part of the thesis on SOS relaxations of the problem of optimizing a quadratic form over the hypercube $\{\pm 1\}^n$. In this initial chapter, we give the basic definitions that will be common to the following chapters, and also motivate our particular interest in the “Gramian structure” of pseudomoment matrices (we have also given a broad overview previously in Chapter 1). This discussion is taken from the introductions of [BK18, KB20, Kun20b] and informed by numerous books and reviews including [Lau09, BPT12, BS14, RSS18].

6.1 SUM-OF-SQUARES OVER THE HYPERCUBE

For the sake of notational convenience, here and in the chapters to come we will consider a rescaled version of the hypercube optimization that we encountered earlier,

$$M(\mathbf{W}) = M_{\{\pm 1\}^n}(\mathbf{W}) := \max_{\mathbf{x} \in \{\pm 1\}^n} \mathbf{x}^\top \mathbf{W} \mathbf{x} = n \cdot M_{\{\pm 1/\sqrt{n}\}^n}(\mathbf{W}). \quad (6.1)$$

We will study *convex relaxations* of this problem, which are best viewed from a slightly different perspective. Namely, this optimization may equivalently be viewed as optimizing

a linear objective over the following convex set.

Definition 6.1.1 (Cut polytope). *The cut polytope is the set*

$$\mathcal{C}^n := \text{conv}(\{\mathbf{x}\mathbf{x}^\top : \mathbf{x} \in \{\pm 1\}^n\}) \subset \mathbb{R}_{\text{sym}}^{n \times n}. \quad (6.2)$$

Then, we have

$$M(\mathbf{W}) = \max_{M \in \mathcal{C}^n} \langle M, \mathbf{W} \rangle. \quad (6.3)$$

Though it is convex, this problem is nonetheless difficult to solve exactly (e.g., NP-hard for \mathbf{W} a graph Laplacian, which computes the maximum cut [Kar72]) due to the intricate discrete geometry of the cut polytope [DL09].

A popular algorithmic choice for approximating $M(\mathbf{W})$ and estimating its optimizer is to form *relaxations* of \mathcal{C}^n , larger convex sets admitting simpler descriptions. Often, the relaxed sets may be described concisely in terms of positive semidefiniteness (psd) conditions, which leads to SDP relaxations of $M(\mathbf{W})$. The most common way to execute this strategy is to optimize over the *elliptope*,

$$\mathcal{E}^n = \mathcal{E}_2^n := \{\mathbf{M} \in \mathbb{R}_{\text{sym}}^{n \times n} : \mathbf{M} \succeq 0, \text{diag}(\mathbf{M}) = \mathbf{1}_n\} \supseteq \mathcal{C}^n. \quad (6.4)$$

For example, the well-known approximation algorithms of Goemans-Williamson [GW95] and Nesterov [Nes98] optimize over \mathcal{E}_2^n and then perform a *rounding* procedure to recover an approximately optimal $\mathbf{x} \in \{\pm 1\}^n$ from $\mathbf{M} \in \mathcal{E}_2^n$.

As our notation suggests, \mathcal{E}_2^n is only the first of a sequence of increasingly tighter relaxations of \mathcal{C}^n , corresponding to *sum-of-squares (SOS)* relaxations of $M(\mathbf{W})$. We describe these relaxations now.

Definition 6.1.2 (Hypercube pseudoexpectation). *Let $\tilde{\mathbb{E}} : \mathbb{R}[x_1, \dots, x_n]_{\leq 2d} \rightarrow \mathbb{R}$ be a linear operator. We say $\tilde{\mathbb{E}}$ is a degree $2d$ pseudoexpectation over $\mathbf{x} \in \{\pm 1\}^n$, or, more precisely,*

with respect to the constraint polynomials $\{x_i^2 - 1\}_{i=1}^n$, if the following conditions hold:

1. $\tilde{\mathbb{E}}[1] = 1$ (normalization),
2. $\tilde{\mathbb{E}}[(x_i^2 - 1)p(\mathbf{x})] = 0$ for all $i \in [n]$, $p \in \mathbb{R}[x_1, \dots, x_n]_{\leq 2d-2}$ (ideal annihilation),
3. $\tilde{\mathbb{E}}[p(\mathbf{x})^2] \geq 0$ for all $p \in \mathbb{R}[x_1, \dots, x_n]_{\leq d}$ (positivity),
4. $\tilde{\mathbb{E}}[p(-\mathbf{x})] = \tilde{\mathbb{E}}[p(\mathbf{x})]$ for all $p \in \mathbb{R}[x_1, \dots, x_n]_{\leq d}$ (symmetry).

Since we always work over the hypercube constraints, we abbreviate and simply call such $\tilde{\mathbb{E}}$ a degree $2d$ pseudoexpectation.

Definition 6.1.3 (Generalized elliptopes). *The degree $2d$ generalized elliptope is the set*

$$\mathcal{E}_{2d}^n := \left\{ \mathbf{M} \in \mathbb{R}_{\text{sym}}^{n \times n} : \mathbf{M} = \tilde{\mathbb{E}}[\mathbf{x}\mathbf{x}^\top] \text{ for some degree } 2d \text{ pseudoexpectation } \tilde{\mathbb{E}} \right\}. \quad (6.5)$$

Definition 6.1.4 (Hypercube SOS relaxation). *The degree $2d$ SOS relaxation of the optimization problem $\mathbf{M}(\mathbf{W})$ is the problem*

$$\text{SOS}_{2d}(\mathbf{W}) := \max_{\substack{\tilde{\mathbb{E}} \\ \text{degree } 2d \\ \text{pseudoexpectation}}} \tilde{\mathbb{E}}[\mathbf{x}^\top \mathbf{W} \mathbf{x}] = \max_{\mathbf{M} \in \mathcal{E}_{2d}^n} \langle \mathbf{W}, \mathbf{M} \rangle. \quad (6.6)$$

We note that, for our purposes, the assumption of symmetry of pseudoexpectations is without loss of generality, since given a $\tilde{\mathbb{E}}$ satisfying all constraints but symmetry we may define $\tilde{\mathbb{E}}'[p(\mathbf{x})] := \tilde{\mathbb{E}}[p(-\mathbf{x})]$, which will also satisfy all constraints but symmetry, and take $\frac{1}{2}(\tilde{\mathbb{E}} + \tilde{\mathbb{E}}')$, which will have the same objective value and satisfy symmetry. It is not standard to include this in the definition of a pseudoexpectation, but we will only be interested in pseudoexpectations satisfying this extra condition in all of our applications.

The reason for the term ‘‘pseudoexpectation’’ is that these objects are a relaxed version of a genuine expectation with respect to a probability measure over $\{\pm 1\}^n$.

Definition 6.1.5 (Integrality). A pseudoexpectation is integral if there exists $\mu \in \mathcal{P}(\{\pm 1\}^n)$ such that $\tilde{\mathbb{E}}[p(\mathbf{x})] = \mathbb{E}_{\mathbf{x} \sim \mu}[p(\mathbf{x})]$.

Clearly, for any such μ that is symmetric about the origin (to satisfy our extra symmetry assumption) the associated $\tilde{\mathbb{E}}$ will indeed be a pseudoexpectation.

The positivity condition on a pseudoexpectation is often easiest to work with in linear-algebraic terms, as follows.

Definition 6.1.6 (Hypercube pseudomoment matrix). A matrix $\mathbf{Y} \in \mathbb{R}_{\text{sym}}^{\binom{[n]}{\leq d} \times \binom{[n]}{\leq d}}$ is a degree $2d$ pseudomoment matrix if the following conditions hold:

1. $Y_{\emptyset, \emptyset} = 1$.
2. $Y_{S, T}$ depends only on $S \Delta T$.
3. $\mathbf{Y} \succeq \mathbf{0}$.
4. $Y_{S, T} = 0$ if $|S \Delta T|$ is odd.

We have arranged the definitions such that the following equivalence between pseudoexpectations and pseudomoment matrices holds.

Proposition 6.1.7. A linear operator $\tilde{\mathbb{E}} : \mathbb{R}[x_1, \dots, x_n]_{\leq 2d} \rightarrow \mathbb{R}$ is a degree $2d$ pseudoexpectation if and only if the matrix $(\tilde{\mathbb{E}}[\mathbf{x}^S \mathbf{x}^T])_{S, T \in \binom{[n]}{\leq d}}$ is a degree $2d$ pseudomoment matrix. We call the latter the pseudomoment matrix of $\tilde{\mathbb{E}}$.

Indeed, the constraints in Definition 6.1.6 and Definition 6.1.2 are equivalent to one another in the order they are listed.

The above gives an explicit way to write $\text{SOS}_{2d}(\mathbf{W})$ as an SDP over a matrix variable of size $n^{d/2} \times n^{d/2}$, which may therefore be solved in time $n^{O(d)}$. (This is not entirely immediate and is subject to various nuances, but it is indeed the case in all situations we will

consider; see, e.g., [O'D17] where this point was first raised and [RW17] where our setting is addressed.)

It is a central and highly non-trivial result that the generalized elliptopes give a sequence of strictly tightening relaxations of \mathcal{C}^n , and that at degree n (up to parity considerations) they achieve integrality:

$$\mathcal{E}_2^n \supseteq \mathcal{E}_4^n \supseteq \cdots \supseteq \mathcal{E}_{n+1\{n \text{ odd}\}}^n = \mathcal{C}^n. \quad (6.7)$$

See Figure 6.1 for a depiction of these containments in low dimension. The strictness of the inequalities here is due to Laurent [Lau03b], though really her result rediscovered in a different guise a slightly earlier result of Grigoriev [Gri01a], while the final equality was conjectured Laurent but proved only a decade later by [FSP16]. We will discuss the Grigoriev-Laurent result underlying the strictness of the inequalities at length in Chapter 9, giving a new proof and clarifying some of the structure of the associated pseudomoments.

In light of these results, optimizing over generalized elliptopes of higher degree may yield better approximations of $M(\mathbf{W})$; however, it is also costlier, since the associated SDP is over ever larger matrices as d grows. It is therefore important to know whether optimizing over generalized elliptopes of constant (or, from the theoretical point of view, growing at various rates with n) degree $d > 2$ actually improves the bounds on $M(\mathbf{W})$ achieved by optimizing over \mathcal{E}_2^n on specific classes of optimization problems as $n \rightarrow \infty$. There is an extensive literature relating this question to the Unique Games Conjecture [KV05, Tre12b], which implies for several problems, most notably maximum cut, that optimizing over generalized elliptopes of constant degree cannot improve the worst-case approximation ratio achieved by optimizing over \mathcal{E}_2^n and then rounding (see, e.g., [KKMO07, Rag08]). However, we will study the average-case rather than worst-case versions of such questions, where there is not yet such a general theory.

The perspective on SOS optimization given above as a bounded-degree variant of a mo-

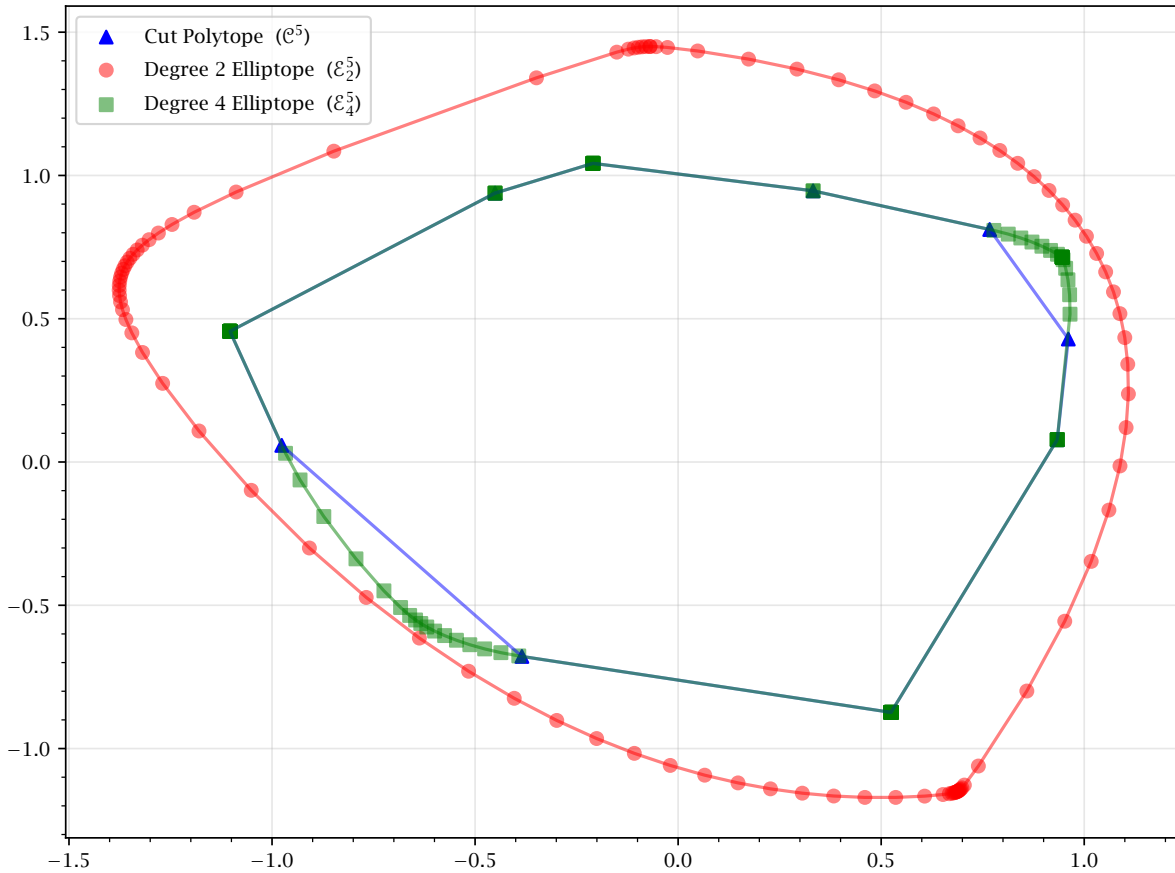


Figure 6.1: Cut polytope and elliptopes in low dimension. We plot the cross-section of \mathcal{C}^5 , \mathcal{E}_2^5 , and \mathcal{E}_4^5 by an isotropic random subspace (in the off-diagonal matrix entries) by numerically solving suitable linear and semidefinite programs. This behaves differently from the projection of these sets onto such a subspace, for which we observe that \mathcal{E}_4^5 is very often indistinguishable from \mathcal{C}^5 . (Note that $n = 5$ is the lowest dimension where $\mathcal{E}_4^n \neq \mathcal{C}^n$.)

ment problem was first developed by Lasserre [Las01, Lau09]. There is also another picture, dual to this one, that explains the term “sum-of-squares” and was developed around the same time by Parillo [Par03, BPT12]. The following is the fundamental duality statement relating these formulations in our setting.

Proposition 6.1.8 (SOS duality). *For some $m = m(n) = O(n^d)$, for any $\mathbf{W} \in \mathbb{R}_{\text{sym}}^{n \times n}$,*

$$\text{SOS}_{2d}(\mathbf{W}) = \left\{ \begin{array}{l} \text{minimize } c \\ \text{subject to } c = \mathbf{x}^\top \mathbf{W} \mathbf{x} + \sum_{i=1}^n (x_i^2 - 1)r_i(\mathbf{x}) + \sum_{j=1}^m s_j(\mathbf{x})^2 \\ \deg(r_i) \leq 2d - 2 \\ \deg(s_j) \leq d \end{array} \right\}. \quad (6.8)$$

The basic idea is that the polynomial equation on the right-hand side is a simple *proof* that $\mathbf{x}^\top \mathbf{W} \mathbf{x} \leq c$ whenever $\mathbf{x} \in \{\pm 1\}^n$, and $\text{SOS}_{2d}(\mathbf{W})$ may be viewed as optimizing this upper bound over all such *sum-of-squares proofs*. This may again be written as an SDP, which is dual to the pseudomoment SDP discussed above. This gives rise to a pleasantly ergonomic way of reasoning about SOS algorithms, sometimes called the *proofs-to-algorithms* paradigm [BS14, BS16]: if we, the theorist, can bound a quantity using only simple algebraic manipulations that admit SOS proofs, then the SOS algorithm itself will also be able to implement our reasoning and achieving a bound at least as strong. We will only occasionally mention such reasoning, but it has given rise to a rich literature of algorithmic applications of SOS (see the above surveys for references).

There is also a related literature on the *proof complexity* of various bounds or refutations in various restricted proof systems, even those that cannot be automatized efficiently; see our discussion in Chapter 1. As some of our results will concern certifying bounds on objective functions with elaborate mathematical treatments (like the SK Hamiltonian), it can be instructive to view our results as showing that there is no “very elementary” argument

giving non-trivial bounds on these functions; see Remark 1.3.2.

6.2 THE GEOMETRY OF PSEUDOMOMENTS

We now propose a somewhat idiosyncratic perspective on the first, primal interpretation of SOS discussed above, which we refer to as a “Gramian” perspective. This may be more broadly applicable, but is especially germane to our hypercube setting thanks to its symmetries. The ellipsope admits the following elegant geometric or Gramian description:

$$\mathcal{E}_2^n = \left\{ \mathbf{M} \in \mathbb{R}_{\text{sym}}^{n \times n} : \mathbf{M} = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n) \text{ for some } \mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{S}^{r-1}, r \leq n \right\}. \quad (6.9)$$

This description is central to the rounding procedures of [GW95, Nes98] as well as the efficient rank-constrained approximations of [BM03]. It is also a useful conceptual tool, allowing us to think of such a relaxation of, say, the maximum cut problem as optimizing over a softer notion of a “vector coloring” of the vertices.

We propose two questions inspired by this observation. First, motivated if not by the mathematical intrigue then at least by the importance of this Gramian description for practical optimization over \mathcal{E}_2^n , we ask the direct analogous question over the generalized ellipsope:

1. Under what conditions does $\text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathcal{E}_{2d}^n$ for $2d > 2$?

We will take up this question in Chapter 7 for the case $2d = 4$, and find that this perspective reveals new general structure in \mathcal{E}_4^n as well as interesting classes of examples of Gram matrices belonging or not belonging to \mathcal{E}_4^n .

We also ask a less specific, perhaps more philosophical question. In proving lower bounds against the SOS hierarchy, one typically constructs a pseudoexpectation $\tilde{\mathbb{E}}$, or, equivalently, a pseudomoment matrix \mathbf{Y} . Per Definition 6.1.6, this is a large positive semidefi-

nite matrix with intricate entrywise symmetries. We propose to think of any such positive semidefinite matrix as a Gram matrix, or, more precisely, as encoding some relative geometry up to orthogonal transformations. For example, \mathcal{E}_2^n may be viewed as the set of all relative configurations of collections of n unit vectors. Thus we ask:

2. What objects' relative geometry is encoded in the large pseudomoment matrices produced in SOS lower bounds? More whimsically, if the v_i above are "pretending" to be the x_i , then who is pretending to be the higher-degree monomials x^S ?

We study this question in Chapters 8 through 11, pursuing a program suggesting that these objects are (at least approximately) certain *multiharmonic polynomials* generalizing the familiar spherical harmonics. While this perspective may not be the most expedient for proving SOS lower bounds, we hope to convince the reader that it begins to reveal new mathematical structure in the pseudomoments involved in those lower bounds, especially structure connecting their entrywise symmetry to their positivity.

7 | DEGREE 4 DETERMINISTIC STRUCTURE AND EXAMPLES

We begin our investigation of SOS relaxations over the hypercube and the associated generalized ellipsoids by studying the Gramian structure of the first non-trivial case, \mathcal{E}_4^n . That is, we know $\text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathcal{E}_2^n$ if and only if the $\mathbf{v}_i \in \mathbb{R}^r$ all have unit norm, and now ask what further constraints on the \mathbf{v}_i must be satisfied to have $\text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathcal{E}_4^n$. In fact, we will be able to give an equivalent condition for membership in \mathcal{E}_4^n in terms of an ancillary semidefinite program parametrized by the \mathbf{v}_i , from which we obtain constraints on possible pseudomoment extensions of a low-rank Gram matrix. We will then use this to derive examples of Gram matrices both extensible and inextensible to degree 4. Those examples, while they concern only highly structured combinatorial Gram matrices, will later guide our constructions for more generic Gram matrices in the chapters to come.

SUMMARY AND REFERENCES This chapter is based on the reference [BK18]. Further applications of the results on equiangular tight frames to the properties of those frames as well as to new constructions of few-distance tight frames were developed in [BK19a, BK19b], but we will not discuss these here to avoid straying too far from our main themes. The following is a summary of our main results in this chapter.

1. (Theorem 7.1.4) An SDP parametrized by $\mathbf{v}_1, \dots, \mathbf{v}_n$ that determines whether the Gram

matrix $\text{Gram}(v_1, \dots, v_n)$ belongs to \mathcal{E}_4^n and which is smaller than the usual feasibility SDP when $r < n$.

2. (Theorem 7.2.15) Constraints on the spectrum and rank of a degree 4 pseudomoment matrix extending $\text{Gram}(v_1, \dots, v_n)$ when $r < n$.
3. (Theorem 7.3.3) An equivalence between integrality of degree 4 pseudomoments and separability of certain associated bipartite quantum states.
4. (Theorem 7.4.5) A complete characterization of when Gram matrices of equiangular tight frames can be extended to degree 4 pseudomoments, and a closed form for an extension when such exists.
5. (Theorem 7.5.1) A new class of inequalities certifiable by degree 4 SOS over the hypercube that are not implied by the triangle inequalities.

7.1 GRAMIAN DESCRIPTION OF DEGREE 4 EXTENSIONS

We first search for a characterization of $\text{Gram}(v_1, \dots, v_n) \in \mathcal{E}_4^n$ that depends more directly on the v_i than the usual definition of a pseudomoment matrix. We begin with some preliminary definitions. Our approach will be based on observing some symmetries in the pseudomoment matrix that are clearer if we allow indexing not by sets of indices but by tuples, introducing some redundancy.

Definition 7.1.1 (Redundant pseudomoment matrix). *For $s, t \in [n]^d$, write $s \bullet t$ for the concatenation, and $\text{odd}(s) \subseteq [n]$ for the symbols occurring an odd number of times in s . A matrix $Y \in \mathbb{R}_{\text{sym}}^{[n]^d \times [n]^d}$ is a degree $2d$ redundant pseudomoment matrix if the following conditions hold:*

1. $Y_{(1 \dots 1)(1 \dots 1)} = 1$.

2. $Y_{s,t}$ depends only on $\text{odd}(s \bullet t)$.

3. $Y \succeq 0$.

4. $Y_{s,t} = 0$ if $|\text{odd}(s \bullet t)|$ is odd.

We say that Y extends its upper left $n \times n$ block.

It is straightforward to see that such Y is determined by its minor indexed by subsets in $\binom{[n]}{\leq d}$ in ascending order extended to length d by adding occurrences of 1, and that Y is a redundant pseudomoment matrix if and only if that minor is an ordinary pseudomoment matrix (Definition 6.1.6). We give a concrete version of this definition for degree 4, which is what we will be concerned with here.

Proposition 7.1.2. *Let $Y \in \mathbb{R}^{n^2 \times n^2}$, with the row and column indices of Y identified with pairs $(ij) \in [n]^2$ ordered lexicographically. Then, Y is a degree 4 pseudomoment matrix if and only if the following conditions hold:*

1. $Y \succeq 0$.

2. $Y_{(ij)(kk)}$ does not depend on the index k .

3. $Y_{(ii)(ii)} = 1$ for every $i \in [n]$.

4. $Y_{(ij)(k\ell)}$ is invariant under permutations of the indices i, j, k, ℓ .

To state our criterion for $\text{Gram}(v_1, \dots, v_n) \in \mathcal{E}_4^n$, we will use the following ancillary set of matrices.

Definition 7.1.3. *For $A \in \mathbb{R}^{rn \times rn}$, we write $A_{[ij]}$ with $i, j \in [n]$ for the $r \times r$ block in position (i, j) when A is viewed as a block matrix. With this notation, let $\mathcal{B}^{n,r} \subset \mathbb{R}_{\text{sym}}^{rn \times rn}$ consist of matrices A satisfying the following properties:*

1. $A \succeq 0$.

2. $\mathbf{A}_{[ii]} = \mathbf{I}_r$ for all $i \in [n]$.
3. $\mathbf{A}_{[ij]} = \mathbf{A}_{[ij]}^\top$ for all $i, j \in [n]$.

In terms of these matrices, \mathcal{E}_4^n admits the following description.

Theorem 7.1.4. *Let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$, let $\mathbf{M} := \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathbb{R}_{\text{sym}}^{n \times n}$, let $\mathbf{V} \in \mathbb{R}^{r \times n}$ have $\mathbf{v}_1, \dots, \mathbf{v}_n$ as its columns, and let $\mathbf{v} := \text{vec}(\mathbf{V}) \in \mathbb{R}^{rn}$ be the concatenation of the \mathbf{v}_i . Then, $\mathbf{M} \in \mathcal{E}_4^n$ if and only if $\|\mathbf{v}\|_2^2 = \sum_{i=1}^n \|\mathbf{v}_i\|_2^2 = n$ and there exists $\mathbf{A} \in \mathcal{B}^{n,r}$ such that $\mathbf{v}^\top \mathbf{A} \mathbf{v} = n^2$. Moreover, if $\mathbf{M} \in \mathcal{E}_4^n$ and a redundant degree 4 pseudomoment matrix $\mathbf{Y} \in \mathbb{R}_{\text{sym}}^{n^2 \times n^2}$ extends \mathbf{M} , then there exists $\mathbf{A} \in \mathcal{B}^{n,r}$ with $\mathbf{v}^\top \mathbf{A} \mathbf{v} = n^2$ and*

$$\mathbf{Y} = (\mathbf{I}_n \otimes \mathbf{V})^\top \mathbf{A} (\mathbf{I}_n \otimes \mathbf{V}), \text{ i.e.} \quad (7.1)$$

$$Y_{(ij)(k\ell)} = \mathbf{v}_i^\top \mathbf{A}_{[jk]} \mathbf{v}_\ell \text{ for all } i, j, k, \ell \in [n]. \quad (7.2)$$

Conversely, if $\sum_{i=1}^n \|\mathbf{v}_i\|_2^2 = n$ and $\mathbf{A} \in \mathcal{B}^{n,r}$ with $\mathbf{v}^\top \mathbf{A} \mathbf{v} = n^2$, then \mathbf{Y} as defined by (7.1) is a degree 4 pseudomoment matrix extending \mathbf{M} .

We will show below that $\|\mathbf{A}\| \leq n$ for all $\mathbf{A} \in \mathcal{B}^{n,r}$, so the condition $\mathbf{v}^\top \mathbf{A} \mathbf{v} = n^2$ is equivalent to \mathbf{v} being a top eigenvector, with eigenvalue n , of \mathbf{A} . We think of \mathbf{A} as a *witness* of the fact that $\mathbf{M} \in \mathcal{E}_4^n$, a Gramian alternative to the conventional pseudoexpectation or pseudomoment witness. The second, more detailed part of Theorem 7.1.4 gives one direction of the equivalence between these two types of witness; the other direction will be described in the course of the proof below.

Before proceeding to the proof, we establish some preliminary facts about the matrices of $\mathcal{B}^{n,r}$.

Proposition 7.1.5. *Let $\mathbf{A} \in \mathcal{B}^{n,r}$. Then,*

1. $\|\mathbf{A}_{[ij]}\| \leq 1$ for all $i, j \in [n]$;

2. $\|\mathbf{A}\| \leq n$;

3. if $\mathbf{A}\mathbf{v} = n\mathbf{v}$, and $\mathbf{0} \neq \mathbf{v} \in \mathbb{R}^{rn}$ is the concatenation of $\mathbf{v}_i \in \mathbb{R}^r$, then the norms $\|\mathbf{v}_i\|_2$ are all equal, and $\mathbf{A}_{[ij]}\mathbf{v}_j = \mathbf{v}_i$ for all $i, j \in [n]$.

Proof. Let $\mathbf{A} \in \mathcal{B}^{n,r}$. To obtain the spectral bound on the blocks $\|\mathbf{A}_{[ij]}\| \leq 1$, note that the claim is trivial for $i = j$, so let us fix $i, j \in [n]$ with $i \neq j$ and denote $\mathbf{S} := \mathbf{A}_{[ij]} \in \mathbb{R}_{\text{sym}}^{r \times r}$.

Taking a suitable submatrix of \mathbf{A} , we find

$$\begin{bmatrix} \mathbf{I}_r & \mathbf{S} \\ \mathbf{S} & \mathbf{I}_r \end{bmatrix} \succeq \mathbf{0}. \quad (7.3)$$

Taking a quadratic form with this matrix, we find that for any $\mathbf{v} \in \mathbb{R}^r$ with $\|\mathbf{v}\|_2 = 1$,

$$0 \leq \begin{bmatrix} \pm\mathbf{v} \\ \mathbf{v} \end{bmatrix}^\top \begin{bmatrix} \mathbf{I}_r & \mathbf{S} \\ \mathbf{S} & \mathbf{I}_r \end{bmatrix} \begin{bmatrix} \pm\mathbf{v} \\ \mathbf{v} \end{bmatrix} = 2 \pm 2\mathbf{v}^\top \mathbf{S}\mathbf{v}, \quad (7.4)$$

thus $|\mathbf{v}^\top \mathbf{S}\mathbf{v}| \leq 1$, and the result follows.

From this, the bound $\|\mathbf{A}\| \leq n$ follows from a simple case of the ‘‘block Gershgorin circle theorem’’ [FV62], which may be deduced directly in this case as follows: suppose $\mathbf{v} \in \mathbb{R}^{rn}$ is the concatenation of $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$, then

$$\mathbf{v}^\top \mathbf{A}\mathbf{v} \leq \sum_{i=1}^n \sum_{j=1}^n |\mathbf{v}_i^\top \mathbf{A}_{[ij]}\mathbf{v}_j| \leq \sum_{i=1}^n \sum_{j=1}^n \|\mathbf{v}_i\|_2 \|\mathbf{v}_j\|_2 = \left(\sum_{i=1}^n \|\mathbf{v}_i\|_2 \right)^2 \leq n \sum_{i=1}^n \|\mathbf{v}_i\|_2^2 = n\|\mathbf{v}\|_2^2, \quad (7.5)$$

giving the result.

For the final statement of the Proposition, if $\mathbf{A}\mathbf{v} = n\mathbf{v}$, then all of the inequalities in (7.5) must be equalities. For the third inequality to be an equality requires all of the $\|\mathbf{v}_i\|_2$ to be equal for $i \in [n]$. For the first inequality to be an equality requires $\mathbf{v}_i^\top \mathbf{A}_{[ij]}\mathbf{v}_j \geq 0$ for all $i, j \in [n]$. For the second inequality to be an equality requires $\mathbf{A}_{[ij]}\mathbf{v}_j = \mathbf{v}_i$ for all $i, j \in [n]$,

completing the proof. □

Proposition 7.1.6. *Let $\mathbf{A} \in \mathcal{B}^{n,r}$. Then, there exists $\mathbf{U} \in \mathbb{R}^{r' \times rn}$ for some $r \leq r' \leq rn$ such that $\mathbf{A} = \mathbf{U}^\top \mathbf{U}$, where*

$$\mathbf{U} = \begin{bmatrix} \mathbf{S}_1 & \mathbf{S}_2 & \cdots & \mathbf{S}_n \\ \mathbf{R}_1 & \mathbf{R}_2 & \cdots & \mathbf{R}_n \end{bmatrix} \quad (7.6)$$

for some $\mathbf{S}_i \in \mathbb{R}_{\text{sym}}^{r' \times r}$, $\mathbf{S}_1 = \mathbf{I}_r$, $\mathbf{R}_i \in \mathbb{R}^{(r'-r) \times r}$, $\mathbf{R}_1 = \mathbf{0}$, which satisfy the relations

$$\mathbf{S}_i^2 + \mathbf{R}_i^\top \mathbf{R}_i = \mathbf{I}_r, \quad (7.7)$$

$$\mathbf{S}_i \mathbf{S}_j - \mathbf{S}_j \mathbf{S}_i + \mathbf{R}_i^\top \mathbf{R}_j - \mathbf{R}_j^\top \mathbf{R}_i = \mathbf{0}. \quad (7.8)$$

(The latter relations encode the conditions $\mathbf{A}_{[ii]} = \mathbf{I}_r$ and $\mathbf{A}_{[ij]}^\top = \mathbf{A}_{[ij]}$, respectively.)

Proof. Let $\mathbf{A} \in \mathcal{B}^{n,r}$ and let $r' := \text{rank}(\mathbf{A})$. Since \mathbf{A} contains \mathbf{I}_r as a principal submatrix, $r' \geq r$, and since rn is the dimension of \mathbf{A} , $r' \leq rn$. Then, there exists $\mathbf{U} \in \mathbb{R}^{r' \times rn}$ such that $\mathbf{A} = \mathbf{U}^\top \mathbf{U}$. Let us expand in blocks

$$\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 & \mathbf{U}_2 & \cdots & \mathbf{U}_n \end{bmatrix}, \quad (7.9)$$

for $\mathbf{U}_i \in \mathbb{R}^{r' \times r}$. Then, $\mathbf{U}_i^\top \mathbf{U}_i = \mathbf{A}_{[ii]} = \mathbf{I}_r$.

This factorization is unchanged by multiplying \mathbf{U} on the left by any matrix of $\mathcal{O}(r')$. Since \mathbf{U}_1 has orthogonal columns, by choosing a suitable such multiplication we may assume without loss of generality that the columns of \mathbf{U}_1 are the first r standard basis vectors $\mathbf{e}_1, \dots, \mathbf{e}_r \in \mathbb{R}^{r'}$. Equivalently,

$$\mathbf{U}_1 = \begin{bmatrix} \mathbf{I}_r \\ \mathbf{0} \end{bmatrix} \begin{matrix} \} r \\ \} r' - r \end{matrix}. \quad (7.10)$$

Let us expand each U_i in blocks of the same dimensions,

$$U_i =: \begin{bmatrix} \mathbf{S}_i \\ \mathbf{R}_i \end{bmatrix} \begin{matrix} \} r \\ \} r' - r \end{matrix}, \quad (7.11)$$

then $\mathbf{S}_1 = \mathbf{I}_r$ and $\mathbf{R}_1 = \mathbf{0}$. We first show that the \mathbf{S}_i are all symmetric. Expanding the block $\mathbf{A}_{[1i]}$, we have

$$\mathbf{A}_{[1i]} = \mathbf{U}_1^\top \mathbf{U}_i = \mathbf{S}_1^\top \mathbf{S}_i + \mathbf{R}_1^\top \mathbf{R}_i = \mathbf{S}_i, \quad (7.12)$$

and since $\mathbf{A}_{[1i]}$ is symmetric, \mathbf{S}_i is symmetric as well.

It remains to show the relations (7.7) and (7.8). For the former, we expand $\mathbf{A}_{[ii]}$:

$$\mathbf{I}_r = \mathbf{A}_{[ii]} = \mathbf{U}_i^\top \mathbf{U}_i = \mathbf{S}_i^2 + \mathbf{R}_i^\top \mathbf{R}_i. \quad (7.13)$$

For the latter, we expand $\mathbf{A}_{[ij]}$ and $\mathbf{A}_{[ji]}$:

$$\mathbf{0} = \mathbf{A}_{[ij]} - \mathbf{A}_{[ji]} = \mathbf{U}_i^\top \mathbf{U}_j - \mathbf{U}_j^\top \mathbf{U}_i = \mathbf{S}_i \mathbf{S}_j - \mathbf{S}_j \mathbf{S}_i + \mathbf{R}_i^\top \mathbf{R}_j - \mathbf{R}_j^\top \mathbf{R}_i, \quad (7.14)$$

completing the proof. □

We now proceed to the main proof of this section. The basic idea of the proof is that the U_i that appear above as the factors of $\mathcal{B}^{n,r}$ encode the Gram vectors of the larger degree 4 pseudomoment matrix \mathbf{Y} as various isometric embeddings of the vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$ of \mathbf{M} . The Gram witness $\mathbf{A} \in \mathcal{B}^{n,r}$ is one step further removed, describing the *relative* orthogonal transformations relating these isometric embeddings. The proof is long but straightforward, and amounts simply to checking that the stated conditions on the Gram witness enforce sufficient “rigidity” in these Gram vectors that the various symmetries required of a degree 4 pseudomoment matrix are satisfied.

Proof of Theorem 7.1.4. We give the proof in two parts, first showing how to construct the Gram witness \mathbf{A} from a pseudomoment witness \mathbf{Y} , and then vice-versa.

PART 1: GRAM WITNESS TO PSEUDOMOMENT WITNESS Let $\mathbf{Y} \in \mathbb{R}^{n^2 \times n^2}$ be a degree 4 redundant pseudomoment matrix extending $\mathbf{M} \in \mathbb{R}^{n \times n}$, where for some $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$, $\mathbf{M} = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n)$. Let $\mathbf{V} \in \mathbb{R}^{r \times n}$ have the \mathbf{v}_i as its columns, and let $\mathbf{v} = \text{vec}(\mathbf{V}) \in \mathbb{R}^{rn}$ be the concatenation of $\mathbf{v}_1, \dots, \mathbf{v}_n$. We will then show that there exists $\mathbf{A} \in \mathcal{B}^{n, r}$ with $\mathbf{v}^\top \mathbf{A} \mathbf{v} = n^2$ and

$$\mathbf{Y} = (\mathbf{I}_n \otimes \mathbf{V})^\top \mathbf{A} (\mathbf{I}_n \otimes \mathbf{V}). \quad (7.15)$$

We first analyze the special case $r = \text{rank}(\mathbf{M})$, then extend to the general case.

CASE 1: $r = \text{rank}(\mathbf{M})$. We build \mathbf{A} based on a suitable factorization of \mathbf{Y} . Let $r' := \text{rank}(\mathbf{Y}) \geq r$, then there exists $\mathbf{A} \in \mathbb{R}^{r' \times n^2}$ such that $\mathbf{Y} = \mathbf{A}^\top \mathbf{A}$. Let us expand in blocks

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 & \cdots & \mathbf{A}_n \end{bmatrix}, \quad (7.16)$$

for $\mathbf{A}_i \in \mathbb{R}^{r' \times n}$. Since $\mathbf{A}_1^\top \mathbf{A}_1 = \mathbf{Y}_{[11]} = \mathbf{M} = \mathbf{V}^\top \mathbf{V}$, there exists $\mathbf{Z} \in \mathbb{R}^{r' \times r}$ such that $\mathbf{A}_1 = \mathbf{Z} \mathbf{V}$ and $\mathbf{Z}^\top \mathbf{Z} = \mathbf{I}_r$. By adding extra columns, we may extend \mathbf{Z} to an orthogonal matrix $\tilde{\mathbf{Z}} \in \mathcal{O}(r')$. The factorization $\mathbf{Y} = \mathbf{A}^\top \mathbf{A}$ is unchanged by multiplying \mathbf{A} on the left by any element of $\mathcal{O}(r')$. By performing this transformation with $\tilde{\mathbf{Z}}$, we may assume without loss of generality that \mathbf{A} is chosen such that

$$\mathbf{A}_1 = \begin{bmatrix} \mathbf{V} \\ \mathbf{0} \end{bmatrix} \begin{matrix} \} r \\ \} r' - r \end{matrix} \quad (7.17)$$

where the numbers following the braces show the dimensionality of the matrix blocks.

Now, since $\mathbf{A}_i^\top \mathbf{A}_i = \mathbf{Y}_{[ii]} = \mathbf{M} = \mathbf{A}_1^\top \mathbf{A}_1$ for every $i \in [n]$ (since, by the degree 4 pseudo-

moment conditions, $Y_{(ik)(i\ell)} = Y_{(ii)(k\ell)} = Y_{(11)(k\ell)} = M_{k\ell}$, there must exist $I_{r'} = Q_1, \dots, Q_n \in \mathcal{O}(r')$ such that $A_i = Q_i A_1$. Let us expand Q_i in blocks,

$$Q_i = \begin{bmatrix} \underbrace{U_i}_r & \underbrace{U'_i}_{r'-r} \end{bmatrix}. \quad (7.18)$$

We then have

$$A_i = Q_i A_1 = U_i V. \quad (7.19)$$

(The extra variable U'_i will not be used in the argument.) Therefore, the blocks of Y are given by

$$Y_{[ij]} = A_i^\top A_j = V^\top U_i^\top U_j V. \quad (7.20)$$

By the permutation symmetry of Y , every such block is symmetric. Since V has full rank, VV^\top is invertible, and therefore the matrix $(VV^\top)^{-1} V Y_{[ij]} V^\top (VV^\top)^{-1} = U_i^\top U_j$ is also symmetric.

We now define A blockwise by

$$A_{[ij]} := U_i^\top U_j. \quad (7.21)$$

Then $A \geq 0$ by construction, $A_{[ii]} = I_r$ since this is the upper left block of $Q_i^\top Q_i = I_{r'}$, and $A_{[ij]}$ is symmetric by the preceding derivation. Thus, $A \in \mathcal{B}^{n,r}$. By (7.20), we also have

$$Y = (I_n \otimes V)^\top A (I_n \otimes V). \quad (7.22)$$

It remains only to check that $v^\top A v = n^2$:

$$v^\top A v = \sum_{i=1}^n \sum_{j=1}^n v_i^\top A_{[ij]} v_j = \sum_{i=1}^n \sum_{j=1}^n (V^\top U_i^\top U_j V)_{ij} = \sum_{i=1}^n \sum_{j=1}^n Y_{(ii)(jj)} = n^2, \quad (7.23)$$

completing the proof of the first case.

CASE 2: $r > \text{rank}(M)$. We will reduce this case to the previous case. Let $r_0 = \text{rank}(M) < r$. Fix Gram vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^{r_0}$ such that $M = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n)$, and, by the previous argument, choose $A \in \mathcal{B}^{n, r_0}$ having $\mathbf{v}^\top A \mathbf{v} = n^2$.

Suppose that $\mathbf{v}'_1, \dots, \mathbf{v}'_n \in \mathbb{R}^r$ such that $M = \text{Gram}(\mathbf{v}'_1, \dots, \mathbf{v}'_n)$. Let \mathbf{v}' be the concatenation of $\mathbf{v}'_1, \dots, \mathbf{v}'_n$. Since the Gram matrices of $\mathbf{v}_1, \dots, \mathbf{v}_n$ and $\mathbf{v}'_1, \dots, \mathbf{v}'_n$ are equal, there must exist $Z \in \mathbb{R}^{r \times r_0}$ with $Z \mathbf{v}_i = \mathbf{v}'_i$ for each $i \in [n]$ and $Z^\top Z = I_{r_0}$. Define $A' \in \mathbb{R}^{rn \times rn}$ to have blocks

$$A'_{[ij]} := \begin{cases} Z A_{[ij]} Z^\top & : i \neq j, \\ I_r & : i = j. \end{cases} \quad (7.24)$$

Equivalently,

$$A' = (I_n \otimes Z) A (I_n \otimes Z)^\top + I_n \otimes (I_r - Z Z^\top). \quad (7.25)$$

Since $Z Z^\top \preceq I_r$ (the left-hand side is a projection matrix), $A' \succeq 0$, and by construction $A'_{[ii]} = I_r$ and $A'_{[ij]}$ is symmetric. Thus, $A' \in \mathcal{B}^{n, r}$.

We also have

$$\mathbf{v}'^\top A' \mathbf{v}' = \sum_{i=1}^n \|\mathbf{v}'_i\|_2^2 + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \mathbf{v}'_i^\top A'_{[ij]} \mathbf{v}'_j = n + \sum_{\substack{1 \leq i, j \leq n \\ i \neq j}} \mathbf{v}'_i^\top A_{[ij]} \mathbf{v}'_j = n^2. \quad (7.26)$$

Lastly, we check the formula for the entries of Y , distinguishing the cases $i = j$ and $i \neq j$:

$$Y_{(ii)(k\ell)} = M_{k\ell} = \langle \mathbf{v}'_k, \mathbf{v}'_\ell \rangle = \mathbf{v}'_k{}^\top A'_{[ii]} \mathbf{v}'_\ell, \quad (7.27)$$

$$\begin{aligned} Y_{(ij)(k\ell)} &= \mathbf{v}'_k{}^\top A_{[ij]} \mathbf{v}'_\ell \\ &= \mathbf{v}'_k{}^\top Z A_{[ij]} Z^\top \mathbf{v}'_\ell \\ &= \mathbf{v}'_k{}^\top A'_{[ij]} \mathbf{v}'_\ell \quad (\text{for } i \neq j), \end{aligned} \quad (7.28)$$

completing the proof.

PART 2: PSEUDOMOMENT WITNESS TO GRAM WITNESS. We now show how to construct the Gram witness \mathbf{A} from the pseudomoment witness \mathbf{Y} . Suppose that we have $\mathbf{M} = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathbb{R}^{n \times n}$ for some $\mathbf{v}_i \in \mathbb{R}^r$ having $\sum_{i=1}^n \|\mathbf{v}_i\|_2^2 = n$. Let \mathbf{v} be the concatenation of $\mathbf{v}_1, \dots, \mathbf{v}_n$. Suppose also that $\mathbf{A} \in \mathcal{B}^{n,r}$ with $\mathbf{v}^\top \mathbf{A} \mathbf{v} = n^2$. We will show that $\mathbf{Y} \in \mathbb{R}^{n^2 \times n^2}$ defined by

$$Y_{(ij)(k\ell)} = \mathbf{v}_i^\top \mathbf{A}_{[jk]} \mathbf{v}_\ell \quad (7.29)$$

is a degree 4 redundant pseudomoment matrix. Recall that this requires the following properties to hold:

1. $\mathbf{Y} \succeq \mathbf{0}$.
2. $Y_{(ij)(kk)}$ does not depend on the index k .
3. $Y_{(ii)(ii)} = 1$ for every $i \in [n]$.
4. $Y_{(ij)(k\ell)}$ is invariant under permutations of the indices i, j, k, ℓ .

(That the upper left $n \times n$ block of \mathbf{Y} is \mathbf{M} follows from Property 4 and that $\mathbf{A}_{[ii]} = \mathbf{I}_r$.) We will obtain these one by one below. This essentially just entails reversing the derivation of the previous part; however, verifying some of the properties of \mathbf{Y} will require a more detailed understanding of the factorization of \mathbf{A} that we used.

The simplest is Property 1: since $\mathbf{A} \succeq \mathbf{0}$, there exist some $\mathbf{U}_1, \dots, \mathbf{U}_n \in \mathbb{R}^{r' \times r}$ for some $r' \geq 1$ such that $\mathbf{A}_{[jk]} = \mathbf{U}_j^\top \mathbf{U}_k$. Thus,

$$Y_{(ij)(k\ell)} = \mathbf{v}_i^\top \mathbf{U}_j^\top \mathbf{U}_k \mathbf{v}_\ell = \langle \mathbf{U}_j \mathbf{v}_i, \mathbf{U}_k \mathbf{v}_\ell \rangle, \quad (7.30)$$

so $\mathbf{Y} = \text{Gram}(\mathbf{U}_1 \mathbf{v}_1, \dots, \mathbf{U}_n \mathbf{v}_n) \succeq \mathbf{0}$.

For Properties 2 and 3, we will use Proposition 7.1.5. From Claim 2 in the Proposition, since $\|\mathbf{v}\|_2^2 = \text{tr}(\mathbf{M}) = n$, then if $\mathbf{v}^\top \mathbf{A} \mathbf{v} = n^2$ we must have $\mathbf{A} \mathbf{v} = n \mathbf{v}$. Therefore, by Claim 3,

$\|\mathbf{v}_i\|_2 = 1$ for each $i \in [n]$. Also by Claim 3, we have

$$Y_{(ij)(kk)} = \mathbf{v}_i^\top \mathbf{A}_{[jk]} \mathbf{v}_k = \langle \mathbf{v}_i, \mathbf{v}_j \rangle. \quad (7.31)$$

This gives Property 2, and taking $i = j = k$ gives Property 3 since $\|\mathbf{v}_i\|_2 = 1$.

Property 4 is more subtle to establish. First, for a moment treating i, j, k, ℓ as merely four distinct symbols, note that the symmetric group on $\{i, j, k, \ell\}$ is generated by the three transpositions (ij) , (jk) , and $(k\ell)$. Therefore, to establish Property 4 it suffices to show the three equalities

$$Y_{(ij)(k\ell)} = Y_{(ji)(k\ell)} = Y_{(ij)(\ell k)} = Y_{(ik)(j\ell)} \quad (7.32)$$

for all $i, j, k, \ell \in [n]$. One equality follows directly from both $\mathbf{A}_{[jk]}$ and \mathbf{A} being symmetric, whereby $\mathbf{A}_{[jk]} = \mathbf{A}_{[kj]}$:

$$Y_{(ij)(k\ell)} = \mathbf{v}_i^\top \mathbf{A}_{[jk]} \mathbf{v}_\ell = \mathbf{v}_i^\top \mathbf{A}_{[kj]} \mathbf{v}_\ell = Y_{(ik)(j\ell)}. \quad (7.33)$$

For the others, combining Proposition 7.1.5's Claim 3 and Proposition 7.1.6, we find that, following the notation for the factorization of Proposition 7.1.6 in matrices \mathbf{S}_i and \mathbf{R}_i ,

$$\mathbf{v}_i = \mathbf{A}_{[i1]} \mathbf{v}_1 = \mathbf{S}_i \mathbf{v}_1, \quad (7.34)$$

$$\mathbf{v}_1 = \mathbf{A}_{[1i]} \mathbf{v}_i = \mathbf{S}_i \mathbf{v}_i. \quad (7.35)$$

We expand the entries of \mathbf{Y} in terms of the matrices \mathbf{S}_i and \mathbf{R}_i and the single vector \mathbf{v}_1 :

$$\begin{aligned} Y_{(ij)(k\ell)} &= \mathbf{v}_i^\top \mathbf{A}_{[jk]} \mathbf{v}_\ell \\ &= \mathbf{v}_1^\top \mathbf{S}_i (\mathbf{S}_j \mathbf{S}_k + \mathbf{R}_j^\top \mathbf{R}_k) \mathbf{S}_\ell \mathbf{v}_1 \\ &= \mathbf{v}_1^\top \mathbf{S}_i \mathbf{S}_j \mathbf{S}_k \mathbf{S}_\ell \mathbf{v}_1 + \mathbf{v}_1^\top \mathbf{S}_i \mathbf{R}_j^\top \mathbf{R}_k \mathbf{S}_\ell \mathbf{v}_1. \end{aligned} \quad (7.36)$$

To show the first two equalities of (7.32), it then suffices to show that for any $i, j \in [n]$, we have

$$\mathbf{S}_i \mathbf{S}_j \mathbf{v}_1 \stackrel{?}{=} \mathbf{S}_j \mathbf{S}_i \mathbf{v}_1, \quad (7.37)$$

$$\mathbf{R}_i \mathbf{S}_j \mathbf{v}_1 \stackrel{?}{=} \mathbf{R}_j \mathbf{S}_i \mathbf{v}_1. \quad (7.38)$$

Observe first that, by (7.34) and (7.35), we have

$$\mathbf{S}_i^2 \mathbf{v}_1 = \mathbf{v}_1. \quad (7.39)$$

Taking (7.7) as a quadratic form with \mathbf{v}_1 , we find

$$1 = \|\mathbf{v}_1\|_2^2 = \mathbf{v}_1^\top \mathbf{S}_i^2 \mathbf{v}_1 + \|\mathbf{R}_i \mathbf{v}_1\|_2^2 = 1 + \|\mathbf{R}_i \mathbf{v}_1\|_2^2, \quad (7.40)$$

hence $\mathbf{R}_i \mathbf{v}_1 = \mathbf{0}$ for all $i \in [n]$. Then, multiplying (7.8) on the right by \mathbf{v}_1 establishes (7.37).

Next, taking (7.7) as a quadratic form with $\mathbf{v}_i = \mathbf{S}_i \mathbf{v}_1$, we find

$$1 = \|\mathbf{v}_i\|_2^2 = \|\mathbf{S}_i \mathbf{v}_i\|_2^2 + \|\mathbf{R}_i \mathbf{v}_i\|_2^2 = 1 + \|\mathbf{R}_i \mathbf{v}_i\|_2^2, \quad (7.41)$$

so $\mathbf{R}_i \mathbf{S}_i \mathbf{v}_1 = \mathbf{R}_i \mathbf{v}_i = \mathbf{0}$ for each $i \in [n]$ as well. Also, evaluating (7.7) as a quadratic form with $\mathbf{v}_j = \mathbf{S}_j \mathbf{v}_1$, we have

$$1 = \|\mathbf{v}_j\|_2^2 = \|\mathbf{S}_i \mathbf{S}_j \mathbf{v}_1\|_2^2 + \|\mathbf{R}_i \mathbf{S}_j \mathbf{v}_1\|_2^2. \quad (7.42)$$

Taking (7.7) as a bilinear form with $\mathbf{S}_i \mathbf{v}_1$ and $\mathbf{S}_j \mathbf{v}_1$ and using the preceding observations

gives

$$\begin{aligned}
0 &= \mathbf{v}_1^\top \mathbf{S}_j (\mathbf{S}_i \mathbf{S}_j - \mathbf{S}_j \mathbf{S}_i + \mathbf{R}_i^\top \mathbf{R}_j - \mathbf{R}_j^\top \mathbf{R}_i) \mathbf{S}_i \mathbf{v}_1 \\
&= \|\mathbf{S}_i \mathbf{S}_j \mathbf{v}_1\|_2^2 - 1 + \langle \mathbf{R}_i \mathbf{S}_j \mathbf{v}_1, \mathbf{R}_j \mathbf{S}_i \mathbf{v}_1 \rangle \\
&= -\|\mathbf{R}_i \mathbf{S}_j \mathbf{v}_1\|_2^2 + \langle \mathbf{R}_i \mathbf{S}_j \mathbf{v}_1, \mathbf{R}_j \mathbf{S}_i \mathbf{v}_1 \rangle.
\end{aligned} \tag{7.43}$$

The same holds with indices i and j exchanged, so we find

$$\langle \mathbf{R}_i \mathbf{S}_j \mathbf{v}_1, \mathbf{R}_j \mathbf{S}_i \mathbf{v}_1 \rangle = \|\mathbf{R}_i \mathbf{S}_j \mathbf{v}_1\|_2^2 = \|\mathbf{R}_j \mathbf{S}_i \mathbf{v}_1\|_2^2 = \|\mathbf{R}_i \mathbf{S}_j \mathbf{v}_1\|_2 \|\mathbf{R}_j \mathbf{S}_i \mathbf{v}_1\|_2. \tag{7.44}$$

Thus the Cauchy-Schwarz inequality holds tightly between the vectors $\mathbf{R}_i \mathbf{S}_j \mathbf{v}_1$ and $\mathbf{R}_j \mathbf{S}_i \mathbf{v}_1$, so $\mathbf{R}_i \mathbf{S}_j \mathbf{v}_1 = \mathbf{R}_j \mathbf{S}_i \mathbf{v}_1$, establishing (7.38) and completing the proof. \square

7.2 CONSTRAINTS ON PSEUDOMOMENT EXTENSIONS

Through Theorem 7.1.4, we will next connect the structure of degree 4 pseudomoment extensions of $M \in \mathcal{E}_2^n$ and the local geometry of \mathcal{E}_2^n near M . Theorem 7.1.4 describes the membership of $\text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n)$ in \mathcal{E}_4^n in terms of a semidefinite program, whose variable is $\mathbf{A} \in \mathcal{B}^{n,r}$, as follows.

Definition 7.2.1. *Given $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$ and $\mathbf{v} \in \mathbb{R}^{rn}$ their concatenation, define the following two semidefinite programs parametrized by the \mathbf{v}_i :*

$$\text{GramSDP}(\mathbf{v}_1, \dots, \mathbf{v}_n) := \left\{ \begin{array}{l} \text{maximize } \langle \mathbf{v} \mathbf{v}^\top, \mathbf{A} \rangle \\ \text{subject to } \mathbf{A} \succeq 0, \\ \mathbf{A}_{[ii]} = \mathbf{I}_r, \\ \mathbf{A}_{[ij]} = \mathbf{A}_{[ij]}^\top \text{ for } i \neq j. \end{array} \right\}, \tag{7.45}$$

$$\text{GramSDP}^*(v_1, \dots, v_n) := \left\{ \begin{array}{l} \text{minimize } \text{tr}(D) \\ \text{subject to } D \succeq vv^\top, \\ D_{[ij]} = -D_{[ij]}^\top \text{ for } i \neq j. \end{array} \right\}. \quad (7.46)$$

It is easy to verify that these two SDPs are each other's duals, which we will exploit in the argument below.

Remark 7.2.2 (Orthogonal cut SDP). *The variant of our primal GramSDP without the latter constraint of blockwise symmetry, $A_{[ij]} = A_{[ij]}^\top$, was previously considered in [NRV13, BRS15, BKS16] as a natural semidefinite programming relaxation of “orthogonal cut” problems, where one seeks to maximize $\sum_{i,j=1}^n \text{tr}(Q_i^\top C_{[i,j]} Q_j)$ over $Q_i \in \mathcal{O}(r)$ for some matrix coefficients $C_{[i,j]} \in \mathbb{R}^{r \times r}$. This is a natural matrix-valued generalization of the ordinary maximum cut problem which can encode problems such as optimally aligning point clouds (a so-called “Procrustes problem”).*

We also observe that the following operation on block matrices is clearly intimately connected to constraints of these SDPs.

Definition 7.2.3 (Partial transpose). *For $A \in \mathbb{R}^{rn \times rn}$ divided into $n \times n$ many $r \times r$ blocks, let the partial transpose A^Γ denote¹ the matrix where each $r \times r$ block is transposed:*

$$A^\Gamma := \left[A_{[ij]}^\top \right]_{i,j=1}^n \in \mathbb{R}^{rn \times rn}. \quad (7.47)$$

Partial transposition is of great importance in quantum information theory, where it yields a basic technique for detecting *entanglement*. We will return to this connection in Section 7.3. For now, we recall some elementary but perhaps not widely known facts of linear algebra

¹The notation comes from Γ being “half” of the transpose symbol \top .

that originate in applications to quantum information theory. We include proofs for the sake of completeness. The first is the following, a rewriting of the singular value decomposition.

Proposition 7.2.4 (Schmidt Decomposition, Section 2.2.2 of [AS17]). *Let $r \leq n$, $\mathbf{V} \in \mathbb{R}^{r \times n}$ having singular value decomposition $\mathbf{V} = \sum_{i=1}^r \sigma_i \mathbf{y}_i \mathbf{z}_i^\top$, where the $\mathbf{y}_i \in \mathbb{R}^r$ and $\mathbf{z}_i \in \mathbb{R}^n$ each form orthonormal sets and $\sigma_i \geq 0$. Then,*

$$\text{vec}(\mathbf{V}) = \sum_{i=1}^r \sigma_i \mathbf{z}_i \otimes \mathbf{y}_i. \quad (7.48)$$

Proof. This result is simply a matter of applying the vectorization operation vec to the singular value decomposition: if $\mathbf{V} = \sum_{i=1}^r \sigma_i \mathbf{y}_i \mathbf{z}_i^\top$ for $\mathbf{y}_i \in \mathbb{R}^r$ and $\mathbf{z}_i \in \mathbb{R}^n$, then, noting that $\text{vec}(\mathbf{y}_i \mathbf{z}_i^\top) = \mathbf{z}_i \otimes \mathbf{y}_i$ and $\text{vec} : \mathbb{R}^{r \times n} \rightarrow \mathbb{R}^{rn}$ is linear, the result follows. \square

This representation makes it convenient to work with the partial transpose; in particular, using the Schmidt decomposition, it is possible to diagonalize the partial transpose of a rank one matrix explicitly, as follows. (This result appears to be folkloric in the quantum information literature; the references we give are unlikely to be the earliest.)

Proposition 7.2.5 (Lemma III.3 of [Hil07]; Lemma 1 of [JP18]). *Let $\mathbf{V} \in \mathbb{R}^{r \times n}$ with $r \leq n$ and $\mathbf{V} = \sum_{i=1}^r \sigma_i \mathbf{y}_i \mathbf{z}_i^\top$ where $\mathbf{y}_i \in \mathbb{R}^r$ and $\mathbf{z}_i \in \mathbb{R}^n$ form orthonormal sets and $\sigma_i \geq 0$. Let $\mathbf{v} = \text{vec}(\mathbf{V})$. Then,*

$$(\mathbf{v} \mathbf{v}^\top)^\Gamma = \sum_{i=1}^r \sigma_i^2 \mathbf{d}_i \mathbf{d}_i^\top + \sum_{1 \leq i < j \leq r} \sigma_i \sigma_j \mathbf{s}_{ij} \mathbf{s}_{ij}^\top - \sum_{1 \leq i < j \leq r} \sigma_i \sigma_j \mathbf{a}_{ij} \mathbf{a}_{ij}^\top \quad (7.49)$$

where

$$\mathbf{d}_i = \mathbf{z}_i \otimes \mathbf{y}_i, \quad (7.50)$$

$$\mathbf{s}_{ij} = (\mathbf{z}_i \otimes \mathbf{y}_j + \mathbf{z}_j \otimes \mathbf{y}_i) / \sqrt{2}, \quad (7.51)$$

$$\mathbf{a}_{ij} = (\mathbf{z}_i \otimes \mathbf{y}_j - \mathbf{z}_j \otimes \mathbf{y}_i) / \sqrt{2}. \quad (7.52)$$

The r^2 vectors $\mathbf{d}_i, \mathbf{s}_{ij}, \mathbf{a}_{ij}$ moreover have unit norm and are mutually orthogonal, so (7.49) is a spectral decomposition (ignoring terms whose coefficient is zero if \mathbf{V} is not full rank).

Proof. Note that $\mathbf{V} = \sum_{i=1}^r \sigma_i \mathbf{y}_i \mathbf{z}_i^\top$ is a singular value decomposition. By Proposition 7.2.4, we may write

$$\begin{aligned} \mathbf{v} \mathbf{v}^\top &= \left(\sum_{i=1}^r \sigma_i \mathbf{z}_i \otimes \mathbf{y}_i \right) \left(\sum_{i=1}^r \sigma_i \mathbf{z}_i \otimes \mathbf{y}_i \right)^\top \\ &= \sum_{i=1}^r \sum_{j=1}^r \sigma_i \sigma_j (\mathbf{z}_i \mathbf{z}_j^\top) \otimes (\mathbf{y}_i \mathbf{y}_j^\top). \end{aligned} \quad (7.53)$$

Therefore, the partial transpose is

$$\begin{aligned} (\mathbf{v} \mathbf{v}^\top)^\Gamma &= \sum_{i=1}^r \sum_{j=1}^r \sigma_i \sigma_j (\mathbf{z}_i \mathbf{z}_j^\top) \otimes (\mathbf{y}_j \mathbf{y}_i^\top) \\ &= \sum_{i=1}^r \sum_{j=1}^r \sigma_i \sigma_j (\mathbf{z}_i \otimes \mathbf{y}_j) \otimes (\mathbf{z}_j \otimes \mathbf{y}_i)^\top \\ &= \sum_{i=1}^r \sigma_i^2 (\mathbf{z}_i \otimes \mathbf{y}_i) (\mathbf{z}_i \otimes \mathbf{y}_i)^\top \\ &\quad + \sum_{1 \leq i < j \leq r} \sigma_i \sigma_j \left((\mathbf{z}_i \otimes \mathbf{y}_j) (\mathbf{z}_j \otimes \mathbf{y}_i)^\top + (\mathbf{z}_j \otimes \mathbf{y}_i) (\mathbf{z}_i \otimes \mathbf{y}_j)^\top \right), \end{aligned} \quad (7.54)$$

and the result follows by diagonalizing the rank-two matrices in the second sum. \square

Finally, we introduce the following related result characterizing the subspace on which a certain matrix inequality involving the partial transpose is tight, which appears to be original.

Proposition 7.2.6. *Let $\mathbf{V} \in \mathbb{R}^{r \times n}$ with $r \leq n$ have full rank, and let $\mathbf{v} = \text{vec}(\mathbf{V})$. Then,*

$$\mathbf{I}_n \otimes (\mathbf{V} \mathbf{V}^\top) \succeq (\mathbf{v} \mathbf{v}^\top)^\Gamma. \quad (7.55)$$

The subspace on which this inequality is tight is given by

$$\ker \left(\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\Gamma \right) = \left\{ \text{vec}(\mathbf{S}\mathbf{V}) : \mathbf{S} \in \mathbb{R}_{\text{sym}}^{r \times r} \right\} =: V_{\text{sym}}. \quad (7.56)$$

Letting $\mathbf{V} = \sum_{i=1}^r \sigma_i \mathbf{y}_i \mathbf{z}_i^\top$ for $\mathbf{y}_i \in \mathbb{R}^r$ an orthonormal basis, $\mathbf{z}_i \in \mathbb{R}^n$ an orthonormal set, and $\sigma_i > 0$ be the singular decomposition, an orthonormal basis for V_{sym} is given by the $\frac{r(r+1)}{2}$ vectors

$$\mathbf{z}_i \otimes \mathbf{y}_i \text{ for } 1 \leq i \leq n, \quad (7.57)$$

$$\frac{1}{\sqrt{\sigma_i^2 + \sigma_j^2}} (\sigma_i \mathbf{z}_i \otimes \mathbf{y}_j + \sigma_j \mathbf{z}_j \otimes \mathbf{y}_i) \text{ for } 1 \leq i < j \leq n. \quad (7.58)$$

Proof. Let us extend $\mathbf{z}_1, \dots, \mathbf{z}_r$ with $\mathbf{z}_{r+1}, \dots, \mathbf{z}_n$ to a full orthonormal basis. Since $\mathbf{V}\mathbf{V}^\top = \sum_{i=1}^r \sigma_i^2 \mathbf{y}_i \mathbf{y}_i^\top$, we may expand

$$\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) = \left(\sum_{i=1}^n \mathbf{z}_i \mathbf{z}_i^\top \right) \otimes \left(\sum_{j=1}^r \sigma_j^2 \mathbf{y}_j \mathbf{y}_j^\top \right) = \sum_{i=1}^n \sum_{j=1}^r \sigma_j^2 (\mathbf{z}_i \otimes \mathbf{y}_j) (\mathbf{z}_i \otimes \mathbf{y}_j)^\top. \quad (7.59)$$

Dividing this sum into those summands with $i \leq r$ and those with $i > r$ and subtracting (7.54), we may write

$$\begin{aligned} \mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\Gamma &= \sum_{1 \leq i < j \leq r} \left(\frac{1}{2} \sigma_i^2 (\mathbf{z}_j \otimes \mathbf{y}_i) (\mathbf{z}_j \otimes \mathbf{y}_i)^\top + \frac{1}{2} \sigma_j^2 (\mathbf{z}_i \otimes \mathbf{y}_j) (\mathbf{z}_i \otimes \mathbf{y}_j)^\top \right. \\ &\quad \left. - \sigma_i \sigma_j (\mathbf{z}_i \otimes \mathbf{y}_j) (\mathbf{z}_j \otimes \mathbf{y}_i)^\top - \sigma_i \sigma_j (\mathbf{z}_j \otimes \mathbf{y}_i) (\mathbf{z}_i \otimes \mathbf{y}_j)^\top \right) \\ &\quad + \sum_{i=r+1}^n \sum_{j=1}^r \sigma_j^2 (\mathbf{z}_i \otimes \mathbf{y}_j) (\mathbf{z}_i \otimes \mathbf{y}_j)^\top \\ &= \frac{1}{2} \sum_{1 \leq i < j \leq r} (\sigma_i \mathbf{z}_j \otimes \mathbf{y}_i - \sigma_j \mathbf{z}_i \otimes \mathbf{y}_j) (\sigma_i \mathbf{z}_j \otimes \mathbf{y}_i - \sigma_j \mathbf{z}_i \otimes \mathbf{y}_j)^\top \\ &\quad + \sum_{i=r+1}^n \sum_{j=1}^r \sigma_j^2 (\mathbf{z}_i \otimes \mathbf{y}_j) (\mathbf{z}_i \otimes \mathbf{y}_j)^\top. \end{aligned} \quad (7.60)$$

One may show more directly that $\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\Gamma \geq \mathbf{0}$ with a Cauchy-Schwarz argument, but the benefit of this approach is that it allows us to read off the subspace we are interested in directly: note that up to rescaling the expression (7.60) is a spectral decomposition, and thus

$$\begin{aligned} & \ker \left(\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\Gamma \right)^\perp \\ &= \text{span} \left(\left\{ \frac{1}{\sqrt{\sigma_i^2 + \sigma_j^2}} (\sigma_i \mathbf{z}_j \otimes \mathbf{y}_i - \sigma_j \mathbf{z}_i \otimes \mathbf{y}_j) \right\}_{1 \leq i < j \leq r} \cup \{ \mathbf{z}_i \otimes \mathbf{y}_j \}_{i \in [n] \setminus [r], j \in [r]} \right), \end{aligned} \quad (7.61)$$

$$\begin{aligned} & \ker \left(\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\Gamma \right) \\ &= \text{span} \left(\left\{ \frac{1}{\sqrt{\sigma_i^2 + \sigma_j^2}} (\sigma_i \mathbf{z}_j \otimes \mathbf{y}_i + \sigma_j \mathbf{z}_i \otimes \mathbf{y}_j) \right\}_{1 \leq i < j \leq r} \cup \{ \mathbf{z}_i \otimes \mathbf{y}_i \}_{i \in [r]} \right), \end{aligned} \quad (7.62)$$

where the first equality follows from (7.60) and the second may be checked by counting dimensions and verifying mutual orthogonalities. It is also straightforward to verify that the vectors enumerated in (7.62) are orthonormal, and thus give an orthonormal basis for $\ker(\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\Gamma)$.

The only remaining task is to check the alternate description

$$\ker \left(\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\Gamma \right) \stackrel{?}{=} \{ \text{vec}(\mathbf{S}\mathbf{V}) : \mathbf{S} \in \mathbb{R}_{\text{sym}}^{r \times r} \} =: V_{\text{sym}}. \quad (7.63)$$

We have $\dim(\ker(\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\Gamma)) = \frac{r(r+1)}{2}$ by (7.62). Since \mathbf{v}_i are a spanning set, if $\text{vec}(\mathbf{S}\mathbf{V}) = \mathbf{0}$ then $\mathbf{S} = \mathbf{0}$, so the map $\mathbf{S} \mapsto \text{vec}(\mathbf{S}\mathbf{V})$ is injective and thus $\dim(V_{\text{sym}}) = \dim(\mathbb{R}_{\text{sym}}^{r \times r}) = \frac{r(r+1)}{2}$ as well. Therefore, to show (7.63) it suffices to show one inclusion.

Suppose that $\mathbf{S} \in \mathbb{R}_{\text{sym}}^{r \times r}$, then

$$\begin{aligned}
((\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\top) \text{vec}(\mathbf{S}\mathbf{V}))_{[i]} &= (\mathbf{V}\mathbf{V}^\top) \mathbf{S}\mathbf{v}_i - \sum_{j=1}^n \mathbf{v}_j \mathbf{v}_i^\top \mathbf{S}\mathbf{v}_j \\
&= \sum_{j=1}^n \mathbf{v}_j \mathbf{v}_j^\top \mathbf{S}\mathbf{v}_i - \sum_{j=1}^n \mathbf{v}_j \mathbf{v}_i^\top \mathbf{S}\mathbf{v}_j \\
&= \mathbf{0},
\end{aligned} \tag{7.64}$$

where in the last step we use that \mathbf{S} is symmetric. Thus, $\text{vec}(\mathbf{S}\mathbf{V}) \in \ker(\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\top)$, so $V_{\text{sym}} \subseteq \ker(\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\top)$, which completes the proof by the previous dimension counting argument. \square

Using these results to study the duality and complementary slackness of the semidefinite programs from Definition 7.2.1, we find that any optimal \mathbf{A} in GramSDP is highly constrained, as follows.

Lemma 7.2.7. *Let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{S}^{r-1}$ be a spanning set, let $\mathbf{V} \in \mathbb{R}^{r \times n}$ have the \mathbf{v}_i as its columns, let $\mathbf{v} := \text{vec}(\mathbf{V}) \in \mathbb{R}^{rn}$ be the concatenation of $\mathbf{v}_1, \dots, \mathbf{v}_n$, let $\mathbf{M} := \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathcal{E}_2^n$, and let $\mathbf{A}^* \in \mathcal{B}^{n,r}$ be such that $\mathbf{v}^\top \mathbf{A}^* \mathbf{v} = n^2$. Then, all eigenvectors of \mathbf{A}^* with nonzero eigenvalue belong to the subspace*

$$V_{\text{sym}} := \{ \text{vec}(\mathbf{S}\mathbf{V}) : \mathbf{S} \in \mathbb{R}_{\text{sym}}^{r \times r} \} \subset \mathbb{R}^{rn}. \tag{7.65}$$

Additionally, \mathbf{v} is an eigenvector of \mathbf{A}^ with eigenvalue n , and all eigenvectors of \mathbf{A}^* with nonzero eigenvalue that are orthogonal to \mathbf{v} belong to the subspace*

$$V'_{\text{sym}} := \{ \text{vec}(\mathbf{S}\mathbf{V}) : \mathbf{S} \in \mathbb{R}_{\text{sym}}^{r \times r}, \mathbf{v}_i^\top \mathbf{S}\mathbf{v}_i = 0 \text{ for } i \in [n] \} \subset \mathbb{R}^{rn}. \tag{7.66}$$

Proof of Lemma 7.2.7. Suppose that $\mathbf{M} = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathcal{E}_4^n$ for some $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{S}^{r-1}$,

\mathbf{v} is the concatenation of the \mathbf{v}_i , and $\mathbf{v}^\top \mathbf{A}^* \mathbf{v} = n^2$ for some $\mathbf{A}^* \in \mathcal{B}^{n,r}$. Then, \mathbf{A}^* is an optimizer for $\text{GramSDP}(\mathbf{v}_1, \dots, \mathbf{v}_n)$, as defined in Definition 7.2.1. We next apply basic convex optimization results to this SDP. Background on these general facts may be found in [BTN01, BV04]. First, the dual SDP is GramSDP^* from Definition 7.2.1. Next, it is simple to verify that the Slater condition holds, implying strong duality between these SDPs, whereby $\text{GramSDP}(\mathbf{v}_1, \dots, \mathbf{v}_n) = \text{GramSDP}^*(\mathbf{v}_1, \dots, \mathbf{v}_n) = n^2$. Finally, if \mathbf{A}^* and \mathbf{D}^* are primal and dual variables achieving the optimal values of GramSDP and GramSDP^* respectively, then *complementary slackness* must hold between them, $\mathbf{A}^*(\mathbf{D}^* - \mathbf{v}\mathbf{v}^\top) = \mathbf{0}$.

The key to the proof is that, while constructing \mathbf{A}^* achieving a value of n^2 in GramSDP from $\mathbf{v}_1, \dots, \mathbf{v}_n$ (when their Gram matrix belongs to \mathcal{E}_4^n) is difficult (by the more detailed part of Theorem 7.1.4 it is equivalent to constructing the degree 4 pseudomoments themselves), constructing \mathbf{D}^* achieving a value of n^2 in GramSDP^* turns out to be straightforward.

The construction uses the partial transpose operation from Definition 7.2.3. Namely, we define \mathbf{D}^* as

$$\mathbf{D}^* := \mathbf{v}\mathbf{v}^\top - (\mathbf{v}\mathbf{v}^\top)^\Gamma + \mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top). \quad (7.67)$$

We have $\text{tr}(\mathbf{D}^*) = \text{tr}(\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top)) = n^2$, and for $i \neq j$, $\mathbf{D}^*_{[ij]} = \mathbf{v}_i \mathbf{v}_j^\top - \mathbf{v}_j \mathbf{v}_i^\top$, which is antisymmetric as required. The final feasibility condition $\mathbf{D}^* \succeq \mathbf{v}\mathbf{v}^\top$ is equivalent to $(\mathbf{v}\mathbf{v}^\top)^\Gamma \preceq \mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top)$, which follows from Proposition 7.2.6. Thus \mathbf{D}^* is indeed feasible and optimal for GramSDP^* . By complementary slackness, any \mathbf{A}^* optimal for GramSDP must have positive eigenvectors in $\ker(\mathbf{D}^* - \mathbf{v}\mathbf{v}^\top) = \ker(\mathbf{I}_n \otimes (\mathbf{V}\mathbf{V}^\top) - (\mathbf{v}\mathbf{v}^\top)^\Gamma) = V_{\text{sym}}$ by the other result of Proposition 7.2.6.

For the second part of the statement, first note that if $\mathbf{v}^\top \mathbf{A}^* \mathbf{v} = n^2$ then by Proposition 7.1.5 $\mathbf{A}^* \mathbf{v} = n\mathbf{v}$, so $\mathbf{A}^* = \mathbf{v}\mathbf{v}^\top + \mathbf{A}'$ for some $\mathbf{A}' \succeq \mathbf{0}$. Suppose that $\mathbf{w} \in \mathbb{R}^{rn}$ is an eigenvector of \mathbf{A}' with eigenvalue $\lambda > 0$. Then, $\mathbf{w} \in V_{\text{sym}}$ by the above reasoning, so

$\mathbf{w} = \text{vec}(\mathbf{S}\mathbf{V})$ for some $\mathbf{S} \in \mathbb{R}_{\text{sym}}^{r \times r}$. Also,

$$\mathbf{I}_r = \mathbf{A}_{[ii]}^* \succeq (\mathbf{v}\mathbf{v}^\top + \lambda\mathbf{w}\mathbf{w}^\top)_{[ii]} = \mathbf{v}_i\mathbf{v}_i^\top + \lambda\mathbf{S}\mathbf{v}_i\mathbf{v}_i^\top\mathbf{S}, \quad (7.68)$$

and taking this as a quadratic form with \mathbf{v}_i shows that $\mathbf{v}_i^\top\mathbf{S}\mathbf{v}_i = 0$. Since this holds for each $i \in [n]$, we obtain the conclusion, that

$$\mathbf{w} \in V'_{\text{sym}} := \left\{ \text{vec}(\mathbf{S}\mathbf{V}) : \mathbf{S} \in \mathbb{R}_{\text{sym}}^{r \times r}, \mathbf{v}_i^\top\mathbf{S}\mathbf{v}_i = 0 \text{ for } i \in [n] \right\}, \quad (7.69)$$

completing the proof. □

We next apply (7.1) from Theorem 7.1.4, which shows how a spectral decomposition of \mathbf{A}^* gives an expression for the associated pseudomoment matrix \mathbf{Y} as a sum of $\text{rank}(\mathbf{A}^*)$ (not necessarily orthogonal) rank one matrices, which are constrained by Lemma 7.2.7. It turns out that these latter constraints are similar to those appearing in results of [LT94, LP96] connecting the smallest face of \mathcal{E}_2^n containing \mathbf{M} to $\text{span}(\{\mathbf{v}_i\mathbf{v}_i^\top\}_{i=1}^n)$, which lets us describe the constraints on \mathbf{Y} concisely in terms of the local geometry of \mathcal{E}_2^n near \mathbf{M} .

Because of these connections, let us review the basic notions of convex geometry that will be involved before proceeding. In what follows, let $K \subseteq \mathbb{R}^d$ be a compact convex set.

Definition 7.2.8. *The dimension of K is the dimension of the affine hull of K , denoted $\dim(K)$.*

Definition 7.2.9. *A convex subset $F \subseteq K$ is a face of K if whenever $\theta\mathbf{X} + (1 - \theta)\mathbf{Y} \in F$ with $\theta \in (0, 1)$ and $\mathbf{X}, \mathbf{Y} \in K$, then $\mathbf{X}, \mathbf{Y} \in F$.*

Definition 7.2.10. *$\mathbf{X} \in K$ is an extreme point of K if $\{\mathbf{X}\}$ is a face of K (of dimension zero).*

Definition 7.2.11. *The intersection of all faces of K containing $\mathbf{X} \in K$ is the unique smallest face of K containing \mathbf{X} , denoted $\text{face}_K(\mathbf{X})$.*

Definition 7.2.12. *The perturbation of \mathbf{X} in K is the subspace*

$$\text{pert}_K(\mathbf{X}) := \left\{ \mathbf{A} \in \mathbb{R}^d : \mathbf{X} \pm t\mathbf{A} \in K \text{ for all } t > 0 \text{ sufficiently small} \right\}. \quad (7.70)$$

The perturbation will come up naturally in our results, so we present the following useful fact giving its connection to the more intuitive objects from facial geometry.

Proposition 7.2.13. *Let $\mathbf{X} \in K$. Then,*

$$\text{face}_K(\mathbf{X}) = K \cap (\mathbf{X} + \text{pert}_K(\mathbf{X})). \quad (7.71)$$

In particular, the affine hull of $\text{face}_K(\mathbf{X})$ is $\mathbf{X} + \text{pert}_K(\mathbf{X})$, and therefore

$$\dim(\text{face}_K(\mathbf{X})) = \dim(\text{pert}_K(\mathbf{X})) \quad (7.72)$$

(in which there is a harmless reuse of notation between the dimension of a convex set and the dimension of a subspace).

Proof. Let $G := K \cap (\mathbf{X} + \text{pert}_K(\mathbf{X}))$. It is simple to check that $\mathbf{Y} \in G$ if and only if there exists $\mathbf{Y}' \in K$ and $\theta \in (0, 1]$ with $\mathbf{X} = \theta\mathbf{Y} + (1 - \theta)\mathbf{Y}'$ (and if $\theta < 1$ then $\mathbf{Y}' \in G$ as well).

Then, if F is any face of K containing \mathbf{X} , and $\mathbf{Y} \in G$, there exists $\mathbf{Y}' \in K$ and $\theta \in (0, 1]$ such that $\mathbf{X} = \theta\mathbf{Y} + (1 - \theta)\mathbf{Y}'$. If $\theta = 1$, then $\mathbf{Y} = \mathbf{X} \in F$. Otherwise, $\mathbf{Y} \in F$ by the definition of a face. Thus, in any case $\mathbf{Y} \in F$, so $G \subseteq F$. Since this holds for any face F containing \mathbf{X} , in fact $G \subseteq \text{face}_K(\mathbf{X})$.

It then suffices to show that G is a face of K . Suppose $\mathbf{Y} \in G$, and $\mathbf{Y}_1, \mathbf{Y}_2 \in K$ and $\theta \in (0, 1)$ with $\mathbf{Y} = \theta\mathbf{Y}_1 + (1 - \theta)\mathbf{Y}_2$. Since $\mathbf{Y} \in G$, there exists $\mathbf{Z} \in K$ and $\phi \in (0, 1]$ such

that

$$\begin{aligned}
\mathbf{X} &= \phi \mathbf{Y} + (1 - \phi) \mathbf{Z} \\
&= \phi(\theta \mathbf{Y}_1 + (1 - \theta) \mathbf{Y}_2) + (1 - \phi) \mathbf{Z} \\
&= \phi \theta \mathbf{Y}_1 + \phi(1 - \theta) \mathbf{Y}_2 + (1 - \phi) \mathbf{Z}.
\end{aligned} \tag{7.73}$$

This is a convex combination of three points where the coefficients of \mathbf{Y}_1 and \mathbf{Y}_2 are strictly positive, so by the previous characterization we have $\mathbf{Y}_1, \mathbf{Y}_2 \in G$, completing the proof. \square

The following result is the particular application of these definitions to the ellipotope that relates to our result.

Proposition 7.2.14 (Theorem 1(a) of [LT94]). *Let $\mathbf{M} = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathcal{E}_2^n$ for $\mathbf{v}_1, \dots, \mathbf{v}_n \in S^{r-1}$ having $\text{rank}(\mathbf{M}) = r$, and let $\mathbf{V} \in \mathbb{R}^{r \times n}$ have the \mathbf{v}_i as its columns, so that $\mathbf{M} = \mathbf{V}^\top \mathbf{V}$. Then,*

$$\text{pert}_{\mathcal{E}_2^n}(\mathbf{M}) = \left\{ \mathbf{V}^\top \mathbf{S} \mathbf{V} : \mathbf{S} \in \mathbb{R}_{\text{sym}}^{r \times r} \right\} \cap \left\{ \mathbf{A} \in \mathbb{R}^{n \times n} : \text{diag}(\mathbf{A}) = \mathbf{0} \right\} \tag{7.74}$$

$$= \left\{ \mathbf{V}^\top \mathbf{S} \mathbf{V} : \mathbf{S} \in \mathbb{R}_{\text{sym}}^{r \times r}, \mathbf{v}_i^\top \mathbf{S} \mathbf{v}_i = 0 \text{ for } i \in [n] \right\}. \tag{7.75}$$

Theorem 7.2.15. *Suppose $\mathbf{M} \in \mathcal{E}_4^n$ and \mathbf{Y} is a degree 4 redundant pseudomoment matrix extending \mathbf{M} . Then, $\mathbf{Y} \succeq \text{vec}(\mathbf{M}) \text{vec}(\mathbf{M})^\top$, and all eigenvectors of $\mathbf{Y} - \text{vec}(\mathbf{M}) \text{vec}(\mathbf{M})^\top$ with nonzero eigenvalue belong to the subspace $\text{vec}(\text{pert}_{\mathcal{E}_2^n}(\mathbf{M}))$. Consequently,*

$$\text{rank}(\mathbf{Y}) \leq \dim \left(\text{pert}_{\mathcal{E}_2^n}(\mathbf{M}) \right) + 1 \tag{7.76}$$

$$= \frac{\text{rank}(\mathbf{M})(\text{rank}(\mathbf{M}) + 1)}{2} - \text{rank}(\mathbf{M}^{\circ 2}) + 1 \tag{7.77}$$

$$\leq \frac{\text{rank}(\mathbf{M})(\text{rank}(\mathbf{M}) + 1)}{2}, \tag{7.78}$$

where $\mathbf{M}^{\circ 2} = \mathbf{M} \circ \mathbf{M}$ is the entrywise square of \mathbf{M} . In particular, if \mathbf{M} is an extreme point

of \mathcal{E}_2^n and is extensible to a degree 4 pseudomoment matrix \mathbf{Y} , then $\text{rank}(\mathbf{Y}) = \text{rank}(\mathbf{M}) = 1$, and $\mathbf{M} = \mathbf{x}\mathbf{x}^\top$ and $\mathbf{Y} = (\mathbf{x} \otimes \mathbf{x})(\mathbf{x} \otimes \mathbf{x})^\top$ for some $\mathbf{x} \in \{\pm 1\}^n$.

The equality (7.77) is the result Proposition 7.2.14 of [LT94, LP96]. We recall also that $\text{pert}_{\mathcal{E}_2^n}(\mathbf{M})$ as a subspace has the same dimension as $\text{face}_{\mathcal{E}_2^n}(\mathbf{M})$ as a convex set. The final claim then gives a strong, albeit non-quantitative, suggestion that \mathcal{E}_4^n is a substantially tighter relaxation of \mathcal{C}^n than \mathcal{E}_2^n : it implies that no “spurious” extreme points of \mathcal{E}_2^n that are not already extreme points of \mathcal{C}^n persist after constraining to \mathcal{E}_4^n .

The bounds (7.77) and (7.78) are similar in form to the Pataki bound on the rank of extreme points of feasible sets of general SDPs [Pat98]. Because of the very large number of linear constraints in SDPs arising from SOS optimization, however, the Pataki bound is less effective in this setting; it also only applies to extreme points of the set of degree 4 pseudomoment matrices. It is simple to check, for example, that the Pataki bound is far inferior to ours when $\text{rank}(\mathbf{M}) \leq \delta n$ for n large and δ a small constant.

Remark 7.2.16 (Pseudocovariance matrix). *The matrix $\mathbf{Y} - \text{vec}(\mathbf{M})\text{vec}(\mathbf{M})^\top$ is quite natural in the pseudomoment framework: entry $(ij)(k\ell)$ of this matrix contains the difference $\tilde{\mathbb{E}}[x_i x_j x_k x_\ell] - \tilde{\mathbb{E}}[x_i x_j] \tilde{\mathbb{E}}[x_k x_\ell]$. It is natural to think of this quantity as the pseudocovariance of $x_i x_j$ and $x_k x_\ell$, and it is then not surprising that the SOS constraints imply that the pseudocovariance matrix is psd. We are not aware, however, of previous results on SOS optimization that make direct use of the pseudocovariance matrix. It would be interesting to understand what role higher “pseudocumulants” might play in SOS reasoning. The combinatorics of the “sum-of-forests pseudomoments” we construct later in Chapter 10 will suggest that these quantities might be useful to consider in that context; see Remark 10.1.14.*

Proof of Theorem 7.2.15. Suppose $\mathbf{M} \in \mathcal{E}_4^n$ with $\mathbf{M} = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n)$, and $\mathbf{v}_i \in \mathbb{R}^r$ with $r = \text{rank}(\mathbf{M})$. Then if $\mathbf{V} \in \mathbb{R}^{r \times n}$ has the \mathbf{v}_i as its columns, \mathbf{V} is full-rank. If $\mathbf{Y} \in \mathbb{R}^{n^2 \times n^2}$ is any degree 4 pseudomoment matrix extending \mathbf{M} , then there is $\mathbf{A} \in \mathcal{B}^{n,r}$ with

$\mathbf{v}^\top \mathbf{A} \mathbf{v} = n^2$. Suppose $r' = \text{rank}(\mathbf{A})$, then let us write the spectral decomposition $\mathbf{A} = \mathbf{v} \mathbf{v}^\top + \sum_{m=1}^{r'-1} \lambda_m \mathbf{w}_m \mathbf{w}_m^\top$ for some $\lambda_m > 0$.

By Lemma 7.2.7, $\mathbf{w}_m \in V'_{\text{sym}}$. Therefore, $\mathbf{w}_m = \text{vec}(\mathbf{S}_m \mathbf{V})$ for some $\mathbf{S}_m \in \mathbb{R}_{\text{sym}}^{r \times r}$ with $\mathbf{v}_i^\top \mathbf{S}_m \mathbf{v}_i = 0$ for all $i \in [n], m \in [r' - 1]$. By (7.1) from Theorem 7.1.4, we may therefore expand

$$\mathbf{Y} = \tilde{\mathbf{v}} \tilde{\mathbf{v}}^\top + \sum_{m=1}^{r'-1} \lambda_m \tilde{\mathbf{w}}_m \tilde{\mathbf{w}}_m^\top, \quad (7.79)$$

$$(\tilde{\mathbf{v}})_{(ij)} = \langle \mathbf{v}_i, \mathbf{v}_j \rangle, \quad (7.80)$$

$$(\tilde{\mathbf{w}}_m)_{(ij)} = \langle \mathbf{S}_m \mathbf{v}_i, \mathbf{v}_j \rangle. \quad (7.81)$$

Thus, we simply have $\tilde{\mathbf{v}} = \text{vec}(\mathbf{M})$ and $\tilde{\mathbf{w}}_m = \text{vec}(\mathbf{V}^\top \mathbf{S}_m \mathbf{V})$.

Using Proposition 7.2.14, we find that for each $m \in [r' - 1]$, $\tilde{\mathbf{w}}_m = \text{vec}(\mathbf{W}_m)$ for some $\mathbf{W}_m \in \text{pert}_{\mathcal{E}_2^n}(\mathbf{M})$. Hence, every eigenvector of $\mathbf{Y} - \text{vec}(\mathbf{M}) \text{vec}(\mathbf{M})^\top$ having nonzero eigenvalue must lie in $\text{vec}(\text{pert}_{\mathcal{E}_2^n}(\mathbf{M}))$, establishing the first part of the result.

The second part of the result controls $\text{rank}(\mathbf{Y}) \leq r'$. By the first part of the result,

$$r' \leq \dim(\text{pert}_{\mathcal{E}_2^n}(\mathbf{M})) + 1, \quad (7.82)$$

so it suffices to compute the right-hand side. Since \mathbf{V} is full-rank, the map $\mathbf{S} \mapsto \mathbf{V}^\top \mathbf{S} \mathbf{V}$ is injective, so this may be computed as

$$\dim(\text{pert}_{\mathcal{E}_2^n}(\mathbf{M})) = \dim(\text{span}(\{\mathbf{v}_i \mathbf{v}_i^\top\}_{i=1}^n)^\perp) = \frac{r(r+1)}{2} - \dim(\text{span}(\{\mathbf{v}_i \mathbf{v}_i^\top\}_{i=1}^n)). \quad (7.83)$$

Since $\text{Gram}(\mathbf{v}_1 \mathbf{v}_1^\top, \dots, \mathbf{v}_n \mathbf{v}_n^\top) = \mathbf{M}^{\circ 2}$, we equivalently have

$$\dim(\text{pert}_{\mathcal{E}_2^n}(\mathbf{M})) = \frac{r(r+1)}{2} - \text{rank}(\mathbf{M}^{\circ 2}), \quad (7.84)$$

a previously known corollary of Proposition 7.2.14 used in [LT94, LP96].

The final part of the result concerns the special case where $M \in \mathcal{E}_2^n$ is an extreme point, whereby $\dim(\text{pert}_{\mathcal{E}_2^n}(M)) = 0$. Then, if Y is a degree 4 pseudomoment matrix extending M we have $\text{rank}(Y) = r' = 1$, so $\text{rank}(M) = 1$ as well since M is a principal submatrix of Y . Since $M \in \mathcal{E}_2^n$, in fact $M = \mathbf{x}\mathbf{x}^\top$ for some $\mathbf{x} \in \{\pm 1\}^n$, and it is simple to check that the only possible degree 4 extension of rank one is then $Y = (\mathbf{x} \otimes \mathbf{x})(\mathbf{x} \otimes \mathbf{x})^\top$. \square

7.3 INTEGRALITY AND SEPARABILITY

We next make a small detour to investigate more deeply the role of ideas from quantum information theory in our description of \mathcal{E}_4^n , which we first glimpsed in the role of the partial transpose operation above. Since \mathcal{E}_4^n may be seen as a relaxation of the cut polytope \mathcal{C}^n , one expects that the description of \mathcal{E}_4^n in terms of an SDP over the matrices of $\mathcal{B}^{n,r}$, as stated in Theorem 7.1.4, should itself relax a description of \mathcal{C}^n in terms of a similar SDP with additional non-convex constraints. In this section, we show that the most naive such description one might expect is in fact incorrect, and give the correct description, which is related to *separability* and *entanglement* of quantum states.

Naively, by analogy with the fact that if $M \in \mathcal{E}_2^n$ with $\text{rank}(M) = 1$ then $M = \mathbf{x}\mathbf{x}^\top$ for $\mathbf{x} \in \{\pm 1\}^n$, one might expect that constraining the rank of $A \in \mathcal{B}^{n,r}$ in Theorem 7.1.4 to be as small as possible, namely to equal r , would give a description of \mathcal{C}^n . Unfortunately, as the following result shows, this only holds in one direction: if the Gram witness A has rank r then the associated $M \in \mathcal{C}^n$, but there exist $M \in \mathcal{C}^n$ with $\text{rank}(M) = r$ whose membership in \mathcal{E}_4^n does not admit a Gram witness A with $\text{rank}(A) = r$.

Proposition 7.3.1. *Let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$, let $M = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n)$, and let $\mathbf{v} \in \mathbb{R}^{rn}$ be the concatenation of $\mathbf{v}_1, \dots, \mathbf{v}_n$. Then, if $\sum_{i=1}^n \|\mathbf{v}_i\|_2^2 = n$ and there exists $A \in \mathcal{B}^{n,r}$ with $\text{rank}(A) = r$ and $\mathbf{v}^\top A \mathbf{v} = n^2$, then $M \in \mathcal{C}^n$. On the other hand, if $n \notin \{1, 2\}$ and n is not divisible by*

4, then $I_n \in \mathcal{C}^n$ with $I_n = \text{Gram}(e_1, \dots, e_n)$, but letting v be the concatenation of e_1, \dots, e_n , there does not exist $A \in \mathcal{B}^{n,n}$ with $v^\top A v = n^2$ and $\text{rank}(M) = n$.

The unusual arithmetic condition on n in the negative result is probably superfluous if one searches for counterexamples other than the identity; the question is related to the relationship between the rank of a matrix in \mathcal{C}^n and the minimum number of cut matrices to whose convex hull it belongs. The latter quantity is similar to the notions of *completely-positive rank* and *non-negative rank*, and appears to behave counterintuitively sometimes; see [FP16, Liu] for some discussion.

Proof. For the positive direction, suppose $v_1, \dots, v_n \in \mathbb{R}^r$, $M = \text{Gram}(v_1, \dots, v_n)$, $v \in \mathbb{R}^{rn}$ is the concatenation of v_1, \dots, v_n , $\sum_{i=1}^n \|v_i\|_2^2 = n$, and $A \in \mathcal{B}^{n,n}$ with $\text{rank}(A) = r$ and $v^\top A v = n^2$. By Proposition 7.1.5, $\|v_i\|_2 = 1$ for each $i \in [n]$ and $A_{[ij]} v_j = v_i$ for each $i, j \in [n]$.

Since $A \geq 0$ and $\text{rank}(A) = r$, there exist $Q_i \in \mathbb{R}^{r \times r}$ such that $A_{[ij]} = Q_i^\top Q_j$. Moreover, since $Q_i^\top Q_i = A_{[ii]} = I_r$, $Q_i \in \mathcal{O}(r)$ for each $i \in [n]$. The above factorization is unchanged by multiplying each Q_i on the left by an orthogonal matrix, so we may assume without loss of generality that $Q_1 = I_r$.

Thus, $A_{[1i]} = Q_1^\top Q_i = Q_i$, which must be symmetric, so Q_i is symmetric for each $i \in [n]$. And, $A_{[ij]} = Q_i Q_j$ is also symmetric, so Q_1, \dots, Q_n are a commuting family of symmetric orthogonal matrices. Therefore, there exists some $Q \in \mathcal{O}(r)$ and $\mathbf{1}_r = d_1, \dots, d_n \in \{\pm 1\}^r$ such that $Q_i = Q D_i Q^\top$ where $D_i = \text{diag}(d_i)$.

We have $v_i = A_{[i1]} v_1 = Q_i v_1 = Q D_i Q^\top v_1$ for each $i \in [n]$. Thus,

$$M_{ij} = \langle v_i, v_j \rangle = \langle D_i Q^\top v_1, D_j Q^\top v_1 \rangle = \langle D_i D_j, Q^\top v_1 v_1^\top Q \rangle. \quad (7.85)$$

Let $\rho = \text{diag}(Q^\top v_1 v_1^\top Q)$, then since $Q^\top v_1 v_1^\top Q \geq 0$, $\rho \geq 0$, and $\sum_{i=1}^r \rho_i = \text{tr}(Q^\top v_1 v_1^\top Q) = 1$.

Therefore, letting $\tilde{\mathbf{d}}_k := ((\mathbf{d}_i)_k)_{k=1}^n \in \{\pm 1\}^n$, (7.85) is

$$M_{ij} = \sum_{k=1}^r \rho_k (\mathbf{d}_i)_k (\mathbf{d}_j)_k, \quad (7.86)$$

$$\mathbf{M} = \sum_{k=1}^r \rho_k \tilde{\mathbf{d}}_k \tilde{\mathbf{d}}_k^\top \in \mathbb{C}^n, \quad (7.87)$$

completing the proof.

For the negative direction, take $\mathbf{M} = \mathbf{I}_n$. We have $\mathbf{I}_n \in \mathbb{C}^n$ since $\mathbf{I}_n = \frac{1}{2^n} \sum_{\mathbf{x} \in \{\pm 1\}^n} \mathbf{x} \mathbf{x}^\top$, as each off-diagonal entry occurs an equal number of times with a positive sign as with a negative sign in the summation. We will view $\mathbf{I}_n = \text{Gram}(\mathbf{e}_1, \dots, \mathbf{e}_n)$, let $\mathbf{v} = \sum_{i=1}^n \mathbf{e}_i \otimes \mathbf{e}_i$ be the concatenation of the \mathbf{e}_i , and will show that if $\mathbf{A} \in \mathcal{B}^{n,n}$ with $\mathbf{v}^\top \mathbf{A} \mathbf{v} = n^2$, then $\text{rank}(\mathbf{A}) > n$ when $n \notin \{1, 2\} \cup 4\mathbb{N}$.

Suppose otherwise. Then, as in the argument above, $\mathbf{A} \in \mathcal{B}^{n,n}$ has $\mathbf{A}_{[ij]} = \mathbf{Q}_i \mathbf{Q}_j$ for some $\mathbf{Q}_i \in \mathcal{O}(n) \cap \mathbb{R}_{\text{sym}}^{n \times n}$, with $\mathbf{Q}_1 = \mathbf{I}_n$, and where $\mathbf{Q}_1, \dots, \mathbf{Q}_n$ commute. We may then write $\mathbf{Q}_i = \mathbf{Q} \mathbf{D}_i \mathbf{Q}^\top$ for $\mathbf{Q} \in \mathcal{O}(n)$ and $\mathbf{D}_i = \text{diag}(\mathbf{d}_i)$ for $\mathbf{d}_i \in \{\pm 1\}^n$. Let us also write $\mathbf{q}_1, \dots, \mathbf{q}_n$ for the rows of \mathbf{Q} , which form an orthonormal basis of \mathbb{R}^n .

We have

$$n^2 = \mathbf{v}^\top \mathbf{A} \mathbf{v} = \sum_{i=1}^n \sum_{j=1}^n (\mathbf{e}_i \otimes \mathbf{e}_i)^\top \mathbf{A} (\mathbf{e}_j \otimes \mathbf{e}_j) = \sum_{i=1}^n \sum_{j=1}^n (\mathbf{A}_{[ij]})_{ij}. \quad (7.88)$$

Since $\mathbf{A} \geq 0$ and $\text{diag}(\mathbf{A}) = \mathbf{1}_{n^2}$, all entries of \mathbf{A} are at most 1, so each term in this sum must equal 1, i.e. $(\mathbf{A}_{[ij]})_{ij} = 1$ for all $i, j \in [n]$. We then have, for any i, j ,

$$1 = (\mathbf{A}_{[ij]})_{ij} = \mathbf{e}_i^\top \mathbf{Q} \mathbf{D}_i \mathbf{D}_j \mathbf{Q}^\top \mathbf{e}_j = \langle \mathbf{D}_i \mathbf{q}_i, \mathbf{D}_j \mathbf{q}_j \rangle, \quad (7.89)$$

whereby $\mathbf{D}_i \mathbf{q}_i = \mathbf{D}_j \mathbf{q}_j$ for all i, j . In other words, there exists some $\mathbf{q} \in \mathbb{R}^n$ with $\|\mathbf{q}\|_2 = 1$ such that $\mathbf{D}_i \mathbf{q}_i = \mathbf{q}$, or $\mathbf{q}_i = \mathbf{D}_i \mathbf{q}$. Thus, the \mathbf{q}_i are sign flips of a fixed vector.

On the other hand, the \mathbf{q}_i are the rows of $\mathbf{Q} \in \mathcal{O}(n)$, whose columns must also form an orthonormal basis. Therefore, every entry of \mathbf{q} must have the same norm, so each entry of

Q also has equal norm; in other words, Q is, up to a scaling depending on definitions, a Hadamard matrix with real entries [CD06]. A real-valued Hadamard matrix of order n can only exist when $n \in \{1, 2\} \cup 4\mathbb{N}$, so this is a contradiction. \square

The correct way to “repair” this first attempt is quite surprising: the key condition on the Gram witness $A \in \mathcal{B}^{n,r}$ that is equivalent to $M \in \mathcal{C}^n$ is not minimal rank, but *separability*, another notion from quantum information theory. The full extent of this connection remains unclear and is an intriguing subject for future work. We note that the language we use below for our real-valued objects is used in the physics literature almost exclusively to describe similar settings over complex numbers.

Definition 7.3.2. A matrix $A \in \mathbb{R}^{rn \times rn}$ with $\text{tr}(A) = 1$ is separable if there exist $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^n$ with $\|\mathbf{a}_i\|_2 = 1$, $\mathbf{b}_1, \dots, \mathbf{b}_m \in \mathbb{R}^r$ with $\|\mathbf{b}_i\|_2 = 1$, and $\rho_1, \dots, \rho_m \geq 0$ with $\sum_i \rho_i = 1$ such that

$$A = \sum_{i=1}^m \rho_i (\mathbf{a}_i \otimes \mathbf{b}_i)(\mathbf{a}_i \otimes \mathbf{b}_i)^\top. \quad (7.90)$$

If it is not possible to write A in this way, A is entangled. (More properly, A is the density matrix representing, with respect to a particular choice of basis, a bipartite quantum state, and it is the state that is entangled or separable.) We write $\mathcal{B}_{\text{sep}}^{n,r} \subseteq \mathcal{B}^{n,r}$ for the matrices $A \in \mathcal{B}^{n,r}$ such that $\frac{1}{rn}A$ is separable.

Theorem 7.3.3. Let $v_1, \dots, v_n \in \mathbb{R}^r$, let $M = \text{Gram}(v_1, \dots, v_n)$, and let $v \in \mathbb{R}^{rn}$ be the concatenation of v_1, \dots, v_n . Then, $M \in \mathcal{C}^n$ if and only if $\sum_{i=1}^n \|v_i\|_2^2 = n$ and there exists $A \in \mathcal{B}_{\text{sep}}^{n,r}$ such that $v^\top A v = n^2$.

By corollary, if $M \in \mathcal{E}_4^n \setminus \mathcal{C}^n$, then any Gram witness A (suitably scaled) must be the density matrix of an entangled state which, by the definition of $\mathcal{B}^{n,r}$, has the *positive partial transpose (PPT)* property that its partial transpose remains psd (indeed, $A = A^\top \geq 0$). If the partial transpose of a density matrix of a state fails to be psd, it follows that the state

is entangled, but the converse does not hold in general [Per96, HHH96]. The structure of states for which this test does not prove entanglement but which are nonetheless entangled has received considerable attention in the quantum information literature (see e.g. [LKCH00, SBŽ06, LMS10, CD12], as well as [Jae07, BŽ17, AS17] for more general discussion). It is therefore striking that these objects are, per our results, rather commonplace in SOS optimization—for every hypercube optimization problem for which degree 4 SOS is not tight (i.e. for which the optimizer $M^* \in \mathcal{E}_4^n \setminus \mathcal{C}^n$), there is an underlying entangled PPT state that may be recovered from M^* .

Proof of Theorem 7.3.3. Suppose first that $M = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n)$ for some $\mathbf{v}_i \in \mathbb{R}^r$ with $\sum_{i=1}^n \|\mathbf{v}_i\|_2^2 = n$, and $\mathbf{A} \in \mathcal{B}_{\text{sep}}^{n,r}$ such that $\mathbf{v}^\top \mathbf{A} \mathbf{v} = n^2$. By Proposition 7.1.5, $\|\mathbf{v}_i\|_2 = 1$ for each $i \in [n]$. By absorbing constants and rearranging tensor products, the condition $\mathbf{A} \in \mathcal{B}_{\text{sep}}^{n,r}$ may be rewritten as

$$\mathbf{A} = \sum_{i=1}^m \mathbf{A}_i \otimes (\mathbf{b}_i \mathbf{b}_i^\top) \quad (7.91)$$

for some $\mathbf{A}_i \in \mathbb{R}_{\text{sym}}^{n \times n}$ with $\mathbf{A}_i \geq \mathbf{0}$ and such that, letting $\mathbf{a}_i = \text{diag}(\mathbf{A}_i)$,

$$\sum_{i=1}^m (\mathbf{a}_i)_j \mathbf{b}_i \mathbf{b}_i^\top = \mathbf{I}_r \quad (7.92)$$

for each $j \in [n]$.

Let $\mathbf{V} \in \mathbb{R}^{r \times n}$ have the \mathbf{v}_i as its columns. Then,

$$\mathbf{v}^\top \mathbf{A} \mathbf{v} = \sum_{i=1}^m \sum_{j=1}^n \sum_{k=1}^n (\mathbf{A}_i)_{jk} \langle \mathbf{b}_i, \mathbf{v}_j \rangle \langle \mathbf{b}_i, \mathbf{v}_k \rangle = \sum_{i=1}^m \mathbf{b}_i^\top \mathbf{V} \mathbf{A}_i \mathbf{V}^\top \mathbf{b}_i. \quad (7.93)$$

We now bound $\mathbf{b}_i^\top \mathbf{V} \mathbf{A}_i \mathbf{V}^\top \mathbf{b}_i$ by applying a simple matrix inequality; the rather complicated formulation below is only to handle carefully the possibility of certain diagonal entries of \mathbf{A}_i equaling zero. Let $\tilde{\mathbf{A}}_i$ be the maximal strictly positive definite principal submatrix of \mathbf{A}_i , of dimension n_i , and let \mathbf{w}_i be the restriction of $\mathbf{V}^\top \mathbf{b}_i$ to the same indices. Then, $\text{diag}(\tilde{\mathbf{A}}_i) > 0$.

Let $\pi_i : [n_i] \rightarrow [n]$ map the indices of this submatrix to the original indices, and let us define a diagonal matrix $\mathbf{D}_i \in \mathbb{R}^{n_i \times n_i}$ by

$$(\mathbf{D}_i)_{jj} := \left(\sum_{j'=1}^{n_i} \sqrt{(\tilde{\mathbf{A}}_i)_{j'j'}} \cdot |\langle \mathbf{b}_i, \mathbf{v}_{\pi_i(j')} \rangle| \right) \frac{|\langle \mathbf{b}_i, \mathbf{v}_{\pi_i(j)} \rangle|}{\sqrt{(\tilde{\mathbf{A}}_i)_{jj}}}. \quad (7.94)$$

Then, we claim $\mathbf{D}_i \succeq \mathbf{w}_i \mathbf{w}_i^\top$. This is a matter of applying a weighted Cauchy-Schwarz inequality: for $\mathbf{x} \in \mathbb{R}^n$, we have

$$\begin{aligned} \mathbf{x}^\top \mathbf{w}_i \mathbf{w}_i^\top \mathbf{x} &= \left(\sum_{j=1}^{n_i} x_j \langle \mathbf{b}_i, \mathbf{v}_{\pi_i(j)} \rangle \right)^2 \\ &\leq \left(\sum_{j'=1}^{n_i} \sqrt{(\tilde{\mathbf{A}}_i)_{j'j'}} \cdot |\langle \mathbf{b}_i, \mathbf{v}_{\pi_i(j')} \rangle| \right) \left(\sum_{j=1}^{n_i} \frac{|\langle \mathbf{b}_i, \mathbf{v}_{\pi_i(j)} \rangle|}{\sqrt{(\tilde{\mathbf{A}}_i)_{jj}}} x_j^2 \right) \\ &= \sum_{j=1}^{n_i} (\mathbf{D}_i)_{jj} x_j^2. \end{aligned} \quad (7.95)$$

Therefore,

$$\begin{aligned} \mathbf{b}_i^\top \mathbf{V} \mathbf{A}_i \mathbf{V}^\top \mathbf{b}_i &= \mathbf{w}_i^\top \tilde{\mathbf{A}}_i \mathbf{w}_i \\ &\leq \langle \mathbf{D}_i, \tilde{\mathbf{A}}_i \rangle \\ &= \left(\sum_{j=1}^{n_i} \sqrt{(\tilde{\mathbf{A}}_i)_{jj}} \cdot |\langle \mathbf{b}_i, \mathbf{v}_{\pi_i(j)} \rangle| \right)^2 \\ &= \left(\sum_{j=1}^n \sqrt{(\mathbf{a}_i)_j} \cdot |\langle \mathbf{b}_i, \mathbf{v}_j \rangle| \right)^2. \end{aligned} \quad (7.96)$$

Now, combining (7.96) with (7.92) and (7.93) and using the Cauchy-Schwarz inequality, we find

$$\mathbf{v}^\top \mathbf{A} \mathbf{v} \leq n \sum_{i=1}^m \sum_{j=1}^n (\mathbf{a}_i)_j \langle \mathbf{b}_i, \mathbf{v}_j \rangle^2 = n \sum_{j=1}^n \|\mathbf{v}_j\|_2^2 = n^2. \quad (7.97)$$

Thus, the Cauchy-Schwarz inequality in (7.97) must be tight, whereby there exist $\kappa_i \geq 0$ with

$\sum_{i=1}^m \kappa_i = 1$ such that

$$(\mathbf{a}_i)_j \langle \mathbf{b}_i, \mathbf{v}_j \rangle^2 = \kappa_i \quad (7.98)$$

for every $i \in [m]$ and $j \in [n]$. Note in particular that if $\kappa_i > 0$ for some $i \in [m]$, then $\langle \mathbf{b}_i, \mathbf{v}_j \rangle \neq 0$ for all $j \in [n]$. We may then define vectors $\beta_{jk} \in \mathbb{R}^m$ by

$$(\beta_{jk})_i := \begin{cases} \sqrt{\kappa_i} \frac{\langle \mathbf{b}_i, \mathbf{v}_j \rangle}{\langle \mathbf{b}_i, \mathbf{v}_k \rangle} & : \kappa_i > 0, \\ 0 & : \kappa_i = 0. \end{cases} \quad (7.99)$$

Then,

$$\begin{aligned} \|\beta_{jk}\|_2^2 &= \sum_{i:\kappa_i>0} \kappa_i \frac{\langle \mathbf{b}_i, \mathbf{v}_j \rangle^2}{\langle \mathbf{b}_i, \mathbf{v}_k \rangle^2} \\ &= \sum_{i:\kappa_i>0} (\mathbf{a}_i)_k \langle \mathbf{b}_i, \mathbf{v}_j \rangle^2 \\ &\leq \sum_{i=1}^m (\mathbf{a}_i)_k \langle \mathbf{b}_i, \mathbf{v}_j \rangle^2 \\ &= 1, \end{aligned} \quad (\text{by (7.92)})$$

$$\langle \beta_{jk}, \beta_{kj} \rangle = \sum_{i:\kappa_i>0} \kappa_i = 1. \quad (7.100)$$

Thus, in fact $\|\beta_{jk}\|_2 = 1$ and $\beta_{jk} = \beta_{kj}$ for all $j, k \in [n]$. This implies first that whenever $\kappa_i > 0$ then $\langle \mathbf{b}_i, \mathbf{v}_j \rangle^2$ does not depend on j , and second that whenever $\kappa_i = 0$ then $(\mathbf{a}_i)_k \langle \mathbf{b}_i, \mathbf{v}_j \rangle^2 = 0$ for all $j, k \in [n]$. We may assume without loss of generality that $\mathbf{A}_i \neq \mathbf{0}$, so $\mathbf{a}_i \neq \mathbf{0}$, and thus the latter implies that whenever $\kappa_i = 0$, then $\langle \mathbf{b}_i, \mathbf{v}_j \rangle^2 = 0$ for all $j \in [n]$. Therefore, in all cases, $\langle \mathbf{b}_i, \mathbf{v}_j \rangle^2$ does not depend on j .

Let us write $\eta_i := \langle \mathbf{b}_i, \mathbf{v}_j \rangle^2$. For i where $\eta_i \neq 0$, by (7.98) $(\mathbf{a}_i)_j$ does not depend on j either. For these i , let us write $\phi_i := (\mathbf{a}_i)_j$. Evaluating (7.92) as a bilinear form on \mathbf{v}_j and \mathbf{v}_k ,

we then find

$$M_{jk} = \langle \mathbf{v}_j, \mathbf{v}_k \rangle = \sum_{i:\eta_i \neq 0} \phi_i \langle \mathbf{b}_i, \mathbf{v}_j \rangle \langle \mathbf{b}_i, \mathbf{v}_k \rangle = \sum_{i:\eta_i \neq 0} \phi_i \eta_i \operatorname{sgn}(\langle \mathbf{b}_i, \mathbf{v}_j \rangle) \operatorname{sgn}(\langle \mathbf{b}_i, \mathbf{v}_k \rangle). \quad (7.101)$$

When $\eta_i \neq 0$, then $\phi_i \eta_i = \kappa_i$, and when $\eta_i = 0$ then $\kappa_i = 0$. Therefore, we have in fact

$$M_{jk} = \sum_{i=1}^m \kappa_i \operatorname{sgn}(\langle \mathbf{b}_i, \mathbf{v}_j \rangle) \operatorname{sgn}(\langle \mathbf{b}_i, \mathbf{v}_k \rangle), \quad (7.102)$$

showing $M \in \mathcal{C}^n$.

The converse is simpler: suppose that $M \in \mathcal{C}^n$ and $M = \operatorname{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathcal{C}^n$ for $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$. Let $\mathbf{v} \in \mathbb{R}^{rn}$ be the concatenation of the $\mathbf{v}_1, \dots, \mathbf{v}_n$. We will build $\mathbf{A} \in \mathcal{B}_{\text{sep}}^{n,r}$ by essentially reversing the process described in the proof of Proposition 7.3.1. Let $\rho_1, \dots, \rho_m \geq 0$ with $\sum_{i=1}^m \rho_i = 1$ and $\tilde{\mathbf{d}}_1, \dots, \tilde{\mathbf{d}}_m \in \{\pm 1\}^n$ be such that

$$M = \sum_{k=1}^m \rho_k \tilde{\mathbf{d}}_k \tilde{\mathbf{d}}_k^\top. \quad (7.103)$$

We may assume without loss of generality that $m \geq r$, by adding extra terms with zero coefficient to this expression. Then, writing $\mathbf{d}_i := ((\tilde{\mathbf{d}}_k)_i)_{k=1}^m \in \mathbb{R}^m$, $\mathbf{R} = \operatorname{diag}(\rho)$, and $\mathbf{v}'_i = \mathbf{R}^{1/2} \mathbf{d}_i$, (7.103) implies that $M = \operatorname{Gram}(\mathbf{v}'_1, \dots, \mathbf{v}'_n)$. There then exists $\mathbf{Z} \in \mathbb{R}^{m \times r}$ such that $\mathbf{Z} \mathbf{v}_i = \mathbf{v}'_i$ and $\mathbf{Z}^\top \mathbf{Z} = \mathbf{I}_r$.

We let $\mathbf{D}_i := \operatorname{diag}(\mathbf{d}_i)$, and define $\mathbf{A} \in \mathbb{R}^{rn \times rn}$ to have blocks

$$\mathbf{A}_{[ij]} := \mathbf{Z}^\top \mathbf{D}_i \mathbf{D}_j \mathbf{Z} = (\mathbf{D}_i \mathbf{Z})^\top (\mathbf{D}_j \mathbf{Z}). \quad (7.104)$$

The last expression gives \mathbf{A} as a Gram matrix, so $\mathbf{A} \geq \mathbf{0}$. Since $\mathbf{D}_i^2 = \mathbf{I}_r$ for each $i \in [n]$, $\mathbf{A}_{[ii]} = \mathbf{I}_r$, and since $\mathbf{D}_1, \dots, \mathbf{D}_n$ commute, $\mathbf{A}_{[ij]}$ is symmetric. Thus, $\mathbf{A} \in \mathcal{B}^{n,r}$. We also

have

$$\mathbf{v}^\top \mathbf{A} \mathbf{v} = \sum_{i=1}^n \sum_{j=1}^n \mathbf{v}_i'^\top \mathbf{D}_i \mathbf{D}_j \mathbf{v}_j' = \sum_{i=1}^n \sum_{j=1}^n \mathbf{d}_i^\top \mathbf{R}^{1/2} \mathbf{D}_i \mathbf{D}_j \mathbf{R}^{1/2} \mathbf{d}_j = \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^m \rho_k = n^2. \quad (7.105)$$

It only remains to check that \mathbf{A} is separable. To do this, let $\mathbf{z}_1, \dots, \mathbf{z}_m \in \mathbb{R}^r$ be the rows of \mathbf{Z} , then by rewriting (7.104) we have $\mathbf{A} = \sum_{i=1}^m (\tilde{\mathbf{d}}_i \otimes \mathbf{z}_i)(\tilde{\mathbf{d}}_i \otimes \mathbf{z}_i)^\top$. \square

7.4 EXAMPLES FROM EQUIANGULAR TIGHT FRAMES

We next use the tools developed above to analyze the highly-structured special case of extending Gram matrices of *equiangular tight frames*, where the constraints of the previous section successfully guide the search for a degree 4 pseudomoment matrix extending a given degree 2 pseudomoment matrix.

We first review some definitions of special types of *frames* in finite dimension, which are overcomplete collections of vectors with certain favorable geometric properties. A more thorough introduction, in particular for the more typical applications of these definitions in signal processing and harmonic analysis, may be found in [CK12]. In what follows, as before, let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$ be unit vectors and let $\mathbf{M} := \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n)$.

Definition 7.4.1. *The vectors \mathbf{v}_i form a unit norm tight frame (UNTF) if any of the following equivalent conditions hold:*

1. $\sum_{i=1}^n \mathbf{v}_i \mathbf{v}_i^\top = \frac{n}{r} \mathbf{I}_r$.
2. *The eigenvalues of \mathbf{M} all equal either zero or $\frac{n}{r}$.*
3. $\sum_{i=1}^n \sum_{j=1}^n \langle \mathbf{v}_i, \mathbf{v}_j \rangle^2 = \frac{n^2}{r}$.

(The equivalence of the final condition is elementary but less obvious; the quantity on its left-hand side is sometimes called the *frame potential* [BF03].)

Definition 7.4.2. The \mathbf{v}_i form an equiangular tight frame (ETF) if they form a UNTF, and there exists $\alpha \in [0, 1]$, called the coherence of the ETF, such that whenever $i \neq j$ then $|M_{ij}| = \alpha$.

The following remarkable result shows that ETFs are extremal among UNTFs in the sense of worst-case coherence. Moreover, when an ETF exists, α is determined by n and r .

Proposition 7.4.3 (Welch Bound [Wel74]). If $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$ with $\|\mathbf{v}_i\|_2 = 1$, then

$$\max_{\substack{1 \leq i, j \leq n \\ i \neq j}} |\langle \mathbf{v}_i, \mathbf{v}_j \rangle| \geq \sqrt{\frac{n-r}{r(n-1)}}, \quad (7.106)$$

with equality if and only if $\mathbf{v}_1, \dots, \mathbf{v}_n$ form an ETF.

ETFs usually arise from combinatorial constructions and should generally be understood as rigid and highly structured objects. For instance, there remain many open problems about the pairs of dimensions (n, r) for which ETFs do or do not exist. More comprehensive references on these aspects of the theory of ETFs include [STDHJ07, CRT08, FM15].

We also recall a classical result bounding the number of equiangular lines (not necessarily forming a tight frame) that may occur in a given dimension. We also include its elegant proof, since similar ideas will be involved in our arguments.

Proposition 7.4.4 (Gerzon Bound [LSG91]). If $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{S}^{r-1}$ and $|\langle \mathbf{v}_i, \mathbf{v}_j \rangle| = \alpha < 1$ for all $i, j \in [n]$ with $i \neq j$, then $n \leq \frac{r(r+1)}{2}$.

Proof. For all $i \neq j$, $\langle \mathbf{v}_i \mathbf{v}_i^\top, \mathbf{v}_j \mathbf{v}_j^\top \rangle = \alpha^2$. Thus,

$$\text{Gram}(\mathbf{v}_1 \mathbf{v}_1^\top, \dots, \mathbf{v}_n \mathbf{v}_n^\top) = (1 - \alpha^2) \mathbf{I}_n + \alpha^2 \mathbf{1}_n \mathbf{1}_n^\top, \quad (7.107)$$

which is non-singular. The $\mathbf{v}_i \mathbf{v}_i^\top$ are there linearly independent symmetric matrices, so $n \leq \dim(\mathbb{R}_{\text{sym}}^{r \times r}) = \frac{r(r+1)}{2}$. \square

The main reason that it is convenient to work with ETFs is that, when M is the Gram matrix of an ETF, then $|M_{ij}|$ takes only two values, 1 when $i = j$ and some $\alpha \in [0, 1]$ otherwise. Therefore, in particular, the matrix $M^{\circ 2}$ (occurring in the proof above as well as our earlier results) is very simple,

$$M^{\circ 2} = (1 - \alpha^2)I_n + \alpha^2 \mathbf{1}_n \mathbf{1}_n^\top. \quad (7.108)$$

As we have seen in Theorem 7.2.15, $M^{\circ 2}$ is intimately related to $\text{pert}_{\mathcal{E}_2^n}(M)$ and therefore to the possible degree 4 pseudomoment extensions of M . In the case of ETFs, its simple structure makes it possible to compute an explicit (albeit naive) guess for a degree 4 pseudomoment extension, which rather surprisingly turns out to be correct.

By such reasoning, we obtain a complete characterization of membership in \mathcal{E}_4^n for ETF Gram matrices M , which is quite simple in that it depends only on the dimension and rank of M . This result is as follows.

Theorem 7.4.5. *Let $v_1, \dots, v_n \in \mathbb{R}^r$ form an ETF, and let $M := \text{Gram}(v_1, \dots, v_n)$. Then, $M \in \mathcal{E}_4^n$ if and only if $n < \frac{r(r+1)}{2}$ or $r = 1$. If $r = 1$, then $M = \mathbf{x}\mathbf{x}^\top$ for $\mathbf{x} \in \{\pm 1\}^n$, and a degree 4 redundant pseudomoment matrix Y extending M is given by $Y = (\mathbf{x} \otimes \mathbf{x})(\mathbf{x} \otimes \mathbf{x})^\top$. If $r > 1$ and $n < \frac{r(r+1)}{2}$, then, letting $P_{\text{vec}(\text{pert}_{\mathcal{E}_2^n}(M))}$ be the orthogonal projection matrix to $\text{vec}(\text{pert}_{\mathcal{E}_2^n}(M)) \subset \mathbb{R}^{n^2}$, a degree 4 redundant pseudomoment matrix Y extending M is given by*

$$Y = \text{vec}(M)\text{vec}(M)^\top + \frac{n^2(1 - \frac{1}{r})}{\frac{r(r+1)}{2} - n} P_{\text{vec}(\text{pert}_{\mathcal{E}_2^n}(M))}, \text{ i.e.} \quad (7.109)$$

$$Y_{(ij)(k\ell)} = \frac{\frac{r(r-1)}{2}}{\frac{r(r+1)}{2} - n} (M_{ij}M_{k\ell} + M_{ik}M_{j\ell} + M_{i\ell}M_{jk}) - \frac{r^2 \left(1 - \frac{1}{n}\right)}{\frac{r(r+1)}{2} - n} \sum_{a=1}^n M_{ia}M_{ja}M_{ka}M_{\ell a}. \quad (7.110)$$

The *maximal* ETFs with $n = \frac{r(r+1)}{2}$ are notoriously elusive combinatorial objects; for instance, they are known to exist for only four values of n , and the question of their existence

is open for infinitely many values of n [FM15]. Our result invokes another regard in which these ETFs are extremal, which was in fact present but perhaps unnoticed in existing results (in particular in the proof of the Gerzon bound given above): maximal ETF Gram matrices are the only ETF Gram matrices that are extreme points of \mathcal{E}_2^n ; thus, by Theorem 7.2.15, these Gram matrices cannot belong to \mathcal{E}_4^n .

In our argument it will become clear that the case of ETFs (those non-maximal ones that do belong to \mathcal{E}_4^n) is perhaps the simplest possible situation for degree 4 pseudomoments over the hypercube: as shown in (7.109), the degree 4 pseudomoment matrix \mathbf{Y} will have only two distinct positive eigenvalues, and will equal of the sum of the rank one matrix $\text{vec}(\mathbf{M})\text{vec}(\mathbf{M})^\top$, which contributes the “naive” pseudomoment value $M_{ij}M_{k\ell}$, with a constant multiple of the projection matrix onto the subspace $\text{vec}(\text{pert}_{\mathcal{E}_2^n}(\mathbf{M}))$, which contributes the remaining “symmetrization” term appearing in (7.110).

In the proof, we will essentially show that $\mathbf{A} \in \mathcal{B}^{n,r}$ the Gram witness of membership in \mathcal{E}_4^n may be constructed from the orthogonal projection matrix to V'_{sym} from Lemma 7.2.7. Thus we compute this projection in advance below. Recall that we have

$$V'_{\text{sym}} = \left\{ \text{vec}(\mathbf{S}\mathbf{V}) : \mathbf{S} \in \mathbb{R}_{\text{sym}}^{r \times r}, \mathbf{v}_i^\top \mathbf{S} \mathbf{v}_i = 0 \text{ for } i \in [n] \right\}. \quad (7.111)$$

Proposition 7.4.6. *Suppose that the matrices $\mathbf{v}_i \mathbf{v}_i^\top$ are linearly independent, or equivalently that the matrix $\mathbf{M}^{\circ 2}$ is non-singular, and that the \mathbf{v}_i form a UNTF. Let $\mathbf{P}_{V'_{\text{sym}}}$ denote the orthogonal projection to V'_{sym} . Then, the blocks of $\mathbf{P}_{V'_{\text{sym}}}$ are given by*

$$(\mathbf{P}_{V'_{\text{sym}}})_{[ij]} = \frac{r}{n} \left(\frac{1}{2} \langle \mathbf{v}_i, \mathbf{v}_j \rangle \mathbf{I}_r + \frac{1}{2} \mathbf{v}_j \mathbf{v}_i^\top - \sum_{k=1}^n \sum_{\ell=1}^n ((\mathbf{M}^{\circ 2})^{-1})_{k\ell} M_{ik} M_{j\ell} \mathbf{v}_k \mathbf{v}_\ell^\top \right). \quad (7.112)$$

Proof. Suppose we are computing $\mathbf{P}_{V'_{\text{sym}}} \mathbf{y}$ for some $\mathbf{y} \in \mathbb{R}^{rn}$. We consider the associated

optimization over \mathbf{S} :

$$\begin{aligned} \text{obj}(\mathbf{S}; \mathbf{y}) &:= \frac{1}{2} \sum_{i=1}^n \|\mathbf{S}\mathbf{v}_i - \mathbf{y}_i\|^2 \\ &= \frac{1}{2} \|\mathbf{y}\|_2^2 + \frac{n}{2r} \text{tr}(\mathbf{S}^2) - \left\langle \mathbf{S}, \sum_{i=1}^n \frac{\mathbf{v}_i \mathbf{y}_i^\top + \mathbf{y}_i \mathbf{v}_i^\top}{2} \right\rangle, \end{aligned} \quad (7.113)$$

$$\mathbf{S}^*(\mathbf{y}) = \underset{\substack{\mathbf{S} \in \mathbb{R}_{\text{sym}}^{r \times r} \\ \mathbf{v}_i^\top \mathbf{S} \mathbf{v}_i = 0 \text{ for } i \in [n]}}{\text{arg min}} \text{obj}(\mathbf{S}; \mathbf{y}), \quad (7.114)$$

where we have used the tight frame property to simplify the quadratic term.²

We introduce the Lagrangian

$$L(\mathbf{S}, \boldsymbol{\gamma}; \mathbf{y}) := \text{obj}(\mathbf{S}; \mathbf{y}) - \left\langle \mathbf{S}, \sum_{i=1}^n \gamma_i \mathbf{v}_i \mathbf{v}_i^\top \right\rangle \quad (7.115)$$

and write the first-order condition $\frac{\partial L}{\partial \mathbf{S}}(\mathbf{S}^*, \boldsymbol{\gamma}; \mathbf{y}) = 0$, which gives

$$\mathbf{S}^* = \mathbf{S}^*(\mathbf{y}) = \frac{r}{n} \left(\sum_{j=1}^n \frac{\mathbf{v}_j \mathbf{y}_j^\top + \mathbf{y}_j \mathbf{v}_j^\top}{2} + \sum_{j=1}^n \gamma_j \mathbf{v}_j \mathbf{v}_j^\top \right). \quad (7.116)$$

The other first-order condition $\frac{\partial L}{\partial \boldsymbol{\gamma}}(\mathbf{S}^*, \boldsymbol{\gamma}; \mathbf{y}) = 0$ is equivalent to the constraints, $\langle \mathbf{S}^*, \mathbf{v}_i \mathbf{v}_i^\top \rangle = 0$ for all $i \in [n]$, which yields the system of linear equations for $\boldsymbol{\gamma}$,

$$\sum_{j=1}^n (\mathbf{M}^{\circ 2})_{ij} \gamma_j = - \sum_{j=1}^n M_{ij} \langle \mathbf{v}_i, \mathbf{y}_j \rangle \text{ for } i \in [n]. \quad (7.117)$$

Since $\mathbf{M}^{\circ 2}$ is invertible by assumption, this admits a unique solution which is given by

$$\gamma_j = - \sum_{k=1}^n \sum_{\ell=1}^n ((\mathbf{M}^{\circ 2})^{-1})_{jk} M_{k\ell} \langle \mathbf{v}_k, \mathbf{y}_\ell \rangle. \quad (7.118)$$

² Without the tight frame assumption, the matrix $\mathbf{V}\mathbf{V}^\top$ would appear and, upon differentiating with respect to \mathbf{S} , we would find a so-called *continuous matrix Lyapunov equation* giving $(\mathbf{V}\mathbf{V}^\top)\mathbf{S}^* + \mathbf{S}^*(\mathbf{V}\mathbf{V}^\top)$. Such an equation in principle admits an analytic solution by reducing to a linear equation in $\text{vec}(\mathbf{S}^*)$ (see, e.g., [Kuř74]), but this would further complicate the calculations.

Substituting into (7.116), we find

$$\mathbf{S}^* = \frac{r}{n} \left(\sum_{j=1}^n \frac{\mathbf{v}_j \mathbf{y}_j^\top + \mathbf{y}_j \mathbf{v}_j^\top}{2} - \sum_{j=1}^n \sum_{k=1}^n \sum_{\ell=1}^n ((\mathbf{M}^{\circ 2})^{-1})_{jk} M_{k\ell} \langle \mathbf{v}_k, \mathbf{y}_\ell \rangle \mathbf{v}_j \mathbf{v}_j^\top \right). \quad (7.119)$$

We then recover the blocks we are interested in,

$$\begin{aligned} (\mathbf{P}_{V'_{\text{sym}}} \mathbf{y})_{[i]} &= (\text{vec}(\mathbf{S}^* \mathbf{V}))_{[i]} \\ &= \mathbf{S}^* \mathbf{v}_i \\ &= \frac{r}{n} \left(\sum_{j=1}^n \frac{\langle \mathbf{v}_i, \mathbf{y}_j \rangle \mathbf{v}_j + \langle \mathbf{v}_i, \mathbf{v}_j \rangle \mathbf{y}_j}{2} - \sum_{j=1}^n \sum_{k=1}^n \sum_{\ell=1}^n ((\mathbf{M}^{\circ 2})^{-1})_{jk} M_{ij} M_{k\ell} \langle \mathbf{v}_k, \mathbf{y}_\ell \rangle \mathbf{v}_j \right) \\ &= \sum_{j=1}^n \frac{r}{n} \left(\frac{1}{2} \langle \mathbf{v}_i, \mathbf{v}_j \rangle \mathbf{I}_r + \frac{1}{2} \mathbf{v}_j \mathbf{v}_i^\top - \sum_{k=1}^n \sum_{\ell=1}^n ((\mathbf{M}^{\circ 2})^{-1})_{k\ell} M_{ik} M_{j\ell} \mathbf{v}_k \mathbf{v}_\ell^\top \right) \mathbf{y}_j, \end{aligned} \quad (7.120)$$

and the result follows. \square

Corollary 7.4.7. *Suppose that $\mathbf{v}_1, \dots, \mathbf{v}_n$ form an ETF with $r > 1$. Then, the blocks of $\mathbf{P}_{V'_{\text{sym}}}$ are given by*

$$(\mathbf{P}_{V'_{\text{sym}}})_{[ij]} = \frac{n-r}{n(r-1)} \mathbf{v}_i \mathbf{v}_j^\top + \frac{r}{2n} \mathbf{v}_j \mathbf{v}_i^\top + \frac{r}{2n} \langle \mathbf{v}_i, \mathbf{v}_j \rangle \mathbf{I}_r - \frac{r^2(n-1)}{n^2(r-1)} \sum_{k=1}^n M_{ik} M_{jk} \mathbf{v}_k \mathbf{v}_k^\top. \quad (7.121)$$

Proof. By Proposition 7.4.4, the conditions of Proposition 7.4.6 are satisfied, so it suffices to compute $(\mathbf{M}^{\circ 2})^{-1}$. The off-diagonal entries of \mathbf{M} all equal the coherence α , which by Proposition 7.4.3 is given by

$$\alpha = \sqrt{\frac{n-r}{r(n-1)}}. \quad (7.122)$$

Thus, we have

$$\mathbf{M}^{\circ 2} = (1 - \alpha^2) \mathbf{I}_n + \alpha^2 \mathbf{1}_n \mathbf{1}_n^\top = \frac{n(r-1)}{r(n-1)} \mathbf{I}_n + \frac{n-r}{r(n-1)} \mathbf{1}_n \mathbf{1}_n^\top. \quad (7.123)$$

This matrix may be inverted by the Sherman-Morrison formula, giving

$$(\mathbf{M}^{\circ 2})^{-1} = \frac{r(n-1)}{n(r-1)} \mathbf{I}_n - \frac{r(n-r)}{n^2(r-1)} \mathbf{1}_n \mathbf{1}_n^\top. \quad (7.124)$$

Thus, the entries are

$$(\mathbf{M}^{\circ 2})_{ij}^{-1} = \begin{cases} \mathbf{a} := \frac{r((n-1)^2+r-1)}{n^2(r-1)} & : i = j, \\ \mathbf{b} := -\frac{r(n-r)}{n^2(r-1)} & : i \neq j. \end{cases} \quad (7.125)$$

Substituting into the expression from Proposition 7.4.6, we find

$$\begin{aligned} & \sum_{k=1}^n \sum_{\ell=1}^n ((\mathbf{M}^{\circ 2})^{-1})_{k\ell} \langle \mathbf{v}_i, \mathbf{v}_k \rangle \langle \mathbf{v}_j, \mathbf{v}_\ell \rangle \mathbf{v}_k \mathbf{v}_\ell^\top \\ &= (\mathbf{a} - \mathbf{b}) \sum_{k=1}^n \langle \mathbf{v}_i, \mathbf{v}_k \rangle \langle \mathbf{v}_j, \mathbf{v}_k \rangle \mathbf{v}_k \mathbf{v}_k^\top + \mathbf{b} \sum_{k=1}^n \sum_{\ell=1}^n \langle \mathbf{v}_i, \mathbf{v}_k \rangle \langle \mathbf{v}_j, \mathbf{v}_\ell \rangle \mathbf{v}_k \mathbf{v}_\ell^\top \\ &= (\mathbf{a} - \mathbf{b}) \sum_{k=1}^n \langle \mathbf{v}_i, \mathbf{v}_k \rangle \langle \mathbf{v}_j, \mathbf{v}_k \rangle \mathbf{v}_k \mathbf{v}_k^\top + \mathbf{b} (\mathbf{V} \mathbf{V}^\top \mathbf{v}_i) (\mathbf{V} \mathbf{V}^\top \mathbf{v}_j)^\top \\ &= \frac{r(n-1)}{n(r-1)} \sum_{k=1}^n \langle \mathbf{v}_i, \mathbf{v}_k \rangle \langle \mathbf{v}_j, \mathbf{v}_k \rangle \mathbf{v}_k \mathbf{v}_k^\top - \frac{n-r}{r(r-1)} \mathbf{v}_i \mathbf{v}_j^\top. \end{aligned} \quad (7.126)$$

Combining with the full result of Proposition 7.4.6 then gives the claim. \square

Proof of Theorem 7.4.5. Let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$ form an ETF, let $\mathbf{V} \in \mathbb{R}^{r \times n}$ have the \mathbf{v}_i as its columns, let $\mathbf{v} = \text{vec}(\mathbf{V})$ be the concatenation of $\mathbf{v}_1, \dots, \mathbf{v}_n$, and let $\mathbf{M} = \mathbf{V}^\top \mathbf{V} = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n)$. Then, our result is that $\mathbf{M} \in \mathcal{E}_4^n$ if and only if $n < \frac{r(r+1)}{2}$ or $r = 1$. If $r = 1$, then each \mathbf{v}_i is a scalar equal to ± 1 , so $\mathbf{M} \in \mathbb{C}^n$. Thus, it suffices to restrict our attention to $r > 1$.

By Theorem 7.2.15, the negative direction immediately follows: if $n = \frac{r(r+1)}{2}$, then the $\mathbf{v}_i \mathbf{v}_i^\top$ span $\mathbb{R}_{\text{sym}}^{r \times r}$, so by Proposition 7.2.14 \mathbf{M} is an extreme point of \mathcal{E}_2^n , thus \mathbf{M} cannot belong to \mathcal{E}_4^n unless $\text{rank}(\mathbf{M}) = 1$, which is a contradiction if $r > 1$.

The positive direction with $r > 1$ is the more difficult part of the result. We proceed by explicitly constructing $\mathbf{A} \in \mathcal{B}^{n,r}$ with $\mathbf{v}^\top \mathbf{A} \mathbf{v} = n^2$. The construction is optimistic: we

consider the simplest possible choice for \mathbf{A} respecting the constraint of Lemma 7.2.7. The Lemma forces $\mathbf{M} = \mathbf{v}\mathbf{v}^\top + \mathbf{A}'$ where $\mathbf{A}' \geq 0$ with all of its eigenvectors with positive eigenvalue lying in the subspace V'_{sym} . We then simply choose \mathbf{A}' to equal a constant multiple of $\mathbf{P}_{V'_{\text{sym}}}$. Choosing the constant factor such that $\text{tr}(\mathbf{A}) = rn$, we obtain the candidate

$$\mathbf{A} := \mathbf{v}\mathbf{v}^\top + \frac{(r-1)n}{\frac{r(r+1)}{2} - n} \mathbf{P}_{V'_{\text{sym}}}. \quad (7.127)$$

If we could show that $\mathbf{A}_{[ii]} = \mathbf{I}_r$ and $\mathbf{A}_{[ij]}^\top = \mathbf{A}_{[ij]}$ for all $i, j \in [n]$, then the proof would be complete.

Surprisingly, the naive construction (7.127) does satisfy these properties, as may be verified by substituting in the explicit formulae for the blocks of $\mathbf{P}_{V'_{\text{sym}}}$ from Corollary 7.4.7, with which it is straightforward to check that $\mathbf{A} \in \mathcal{B}^{n,r}$.

Finally, using the relation (7.2) between the blocks of \mathbf{A} and the degree 4 pseudomoments, we recover the formula for the degree 4 pseudomoments:

$$Y_{(ij)(k\ell)} = \frac{\frac{r(r-1)}{2}}{\frac{r(r+1)}{2} - n} (M_{ij}M_{k\ell} + M_{ik}M_{j\ell} + M_{il}M_{jk}) - \frac{r^2 \left(1 - \frac{1}{n}\right)}{\frac{r(r+1)}{2} - n} \sum_{m=1}^n M_{im}M_{jm}M_{km}M_{\ell m}, \quad (7.128)$$

concluding the proof. □

This derivation at this point appears to be a rather egregious instance of “bookkeeping for a miracle” [Cla09], and it certainly remains an open question to provide an intuitive explanation for why any ETF Gram matrices ought to belong to \mathcal{E}_4^n at all—we first discovered this fact in numerical experiments. We will, however, give a more principled account of the remarkably symmetric formula (7.128) in Chapters 8, 9, and 10.

7.5 NEW INEQUALITIES CERTIFIABLE AT DEGREE 4

Finally, to close this chapter, we consider the dual question to that of the previous section. While there we found structured deterministic examples of Gram matrices extensible to degree 4 pseudomoments, here we ask: can we use structured deterministic examples to find new inequalities certified by degree 4 SOS over the hypercube?

Let us first motivate this question. There are many results in combinatorial optimization enumerating linear inequalities satisfied by \mathcal{C}^n (see, e.g., [DL09]). The practical purpose of this pursuit is that such linear inequalities may be included in LP relaxations of \mathcal{C}^n , which are typically more efficient than the SDP relaxations we work with here. The putative convenience of SDP relaxations is that they do not require their user to know specifically which inequalities will be relevant for a given problem; the psd constraint captures many relevant inequalities at once. For theoretical understanding, however, it is again important to know which specific inequalities over \mathcal{C}^n are satisfied at which degrees of SOS relaxation, since those inequalities may then be used as analytical tools in proving that SOS succeeds in various tasks. Given a specific valid inequality, the least degree of SOS at which that inequality is certified is also an intrinsically interesting measure of that inequality's “complexity.”

Yet, to the best of our knowledge, very few inequalities over \mathcal{C}^n are known to be satisfied in \mathcal{E}_4^n but not \mathcal{E}_2^n ; indeed, it appears that the only infinite such family known before this work was the *triangle inequalities*,

$$-s_i s_j M_{ij} - s_j s_k M_{jk} - s_i s_k M_{ik} \leq 1 \text{ for } \mathbf{M} \in \mathcal{E}_4^n, \mathbf{s} \in \{\pm 1\}^n. \quad (7.129)$$

Guided by the results from the previous section, we find a new family of similar but independent inequalities. First, from the negative result of Theorem 7.4.5, we obtain concrete examples of matrices $\mathbf{M} \in \mathcal{E}_2^n \setminus \mathcal{E}_4^n$, namely the Gram matrices of ETFs with $n = \frac{r(r+1)}{2}$.

As mentioned before, these are only known to exist for four specific dimensions, namely $r \in \{2, 3, 7, 23\}$. By convex duality, there must exist certificates that these matrices do not belong to \mathcal{E}_4^n , taking the form of linear inequalities that hold over \mathcal{E}_4^n but fail to hold for these matrices. Indeed, for the smallest two examples $r \in \{2, 3\}$, a triangle inequality is a valid certificate of infeasibility.

For $r = 7$, on the other hand, the absolute value of the off-diagonal entries of the Gram matrix is $\alpha = \frac{1}{3}$, so the triangle inequalities are satisfied, and the certificates of infeasibility must be new inequalities which cannot be obtained as linear combinations of triangle inequalities. We compute these certificates numerically and identify the constants that arise by hand to allow the certificates to be validated by symbolic computation (this amounts to checking that a certain $n^2 \times n^2$ matrix is psd, where in this case $n^2 = 28^2 = 784$).

For $r = 23$ the same appears to occur numerically and a similar argument shows that yet another independent family of inequalities must arise as the certificates of infeasibility, but the symbolic verification of such a certificate is a much larger problem which a naive software implementation does not solve in a reasonable time. We thus only present the verified result for $r = 7$ here as a proof of concept, leaving both further computational verification of exact inequalities and further theoretical analysis of these certificates to future work.

Theorem 7.5.1. *Let \mathbf{Z} be the Gram matrix of an ETF of 28 vectors in \mathbb{R}^7 . Then, for any $\mathbf{M} \in \mathcal{E}_4^n$ and any $\pi : [28] \rightarrow [n]$ injective,*

$$\sum_{1 \leq i < j \leq 28} \text{sgn}(Z_{ij}) M_{\pi(i)\pi(j)} \leq 112, \quad (7.130)$$

and this inequality cannot be obtained as a linear combination of the triangle inequalities

$$-s_i s_j M_{ij} - s_j s_k M_{jk} - s_i s_k M_{ik} \leq 1 \text{ for } \mathbf{s} \in \{\pm 1\}^n. \quad (7.131)$$

As a point of comparison, since $\mathbf{Z} \in \mathcal{E}_2^n$, the half-space parallel to that defined by (7.130) most tightly bounding \mathcal{E}_2^n must have the right-hand side at least

$$\sum_{1 \leq i < j \leq 28} \operatorname{sgn}(Z_{ij})Z_{ij} = \sum_{1 \leq i < j \leq 28} |Z_{ij}| = \frac{28(28-1)}{2} \cdot \frac{1}{3} = 126 > 112. \quad (7.132)$$

Thus, these inequalities describe directions in the vector space of symmetric matrices along which \mathcal{E}_4^n is strictly “narrower” than \mathcal{E}_2^n .

To better understand the structure of these inequalities, we refer to a general correspondence between ETFs and strongly regular graphs (SRGs) [FW15]. (In fact, there are two distinct correspondences between ETFs and SRGs: the one we will use applies to arbitrary ETFs and is described in [FW15], while the other applies only to ETFs with a certain additional symmetry and is described in [FJM⁺16].)

Definition 7.5.2. *A graph $G = (V, E)$ is a strongly regular graph with parameters (v, k, λ, μ) , abbreviated $\operatorname{srg}(v, k, \lambda, \mu)$, if $|V| = v$, G is k -regular, every $x, y \in V$ that are adjacent have λ common neighbors, and every $x, y \in V$ that are not adjacent have μ common neighbors.*

Proposition 7.5.3 (Theorem 3.1 of [FW15]). *Let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$ form an ETF with $n > r$, suppose that for all $i \in [n] \setminus \{1\}$ we have $\langle \mathbf{v}_1, \mathbf{v}_i \rangle > 0$, and let $\mathbf{M} = \operatorname{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n)$. Define the graph G on vertices in $[n] \setminus \{1\}$ where i and j are adjacent if and only if $\langle \mathbf{v}_i, \mathbf{v}_j \rangle > 0$. Then, G is an $\operatorname{srg}(v, k, \lambda, \mu)$ with parameters*

$$v = n - 1, \quad (7.133)$$

$$k = \frac{n}{2} - 1 + \left(\frac{n}{2r} - 1 \right) \sqrt{\frac{r(n-1)}{n-r}}, \quad (7.134)$$

$$\mu = \frac{k}{2}, \quad (7.135)$$

$$\lambda = \frac{3k - v - 1}{2}. \quad (7.136)$$

Note that the assumption that $\langle \mathbf{v}_1, \mathbf{v}_i \rangle > 0$ for all $i \neq 1$ is not a substantial restriction, since any vector in an ETF may be negated to produce another, essentially equivalent, ETF.

In our case, an ETF on 28 vectors in \mathbb{R}^7 corresponds to an $\text{srg}(27, 16, 10, 8)$. By the result of [Sei91], this graph is unique. Consequently, putting the ETF into the “canonical” form where $\langle \mathbf{v}_1, \mathbf{v}_i \rangle > 0$ for all $i \neq 1$, we obtain the following uniqueness result.

Proposition 7.5.4. *Let $\mathbf{v}_1, \dots, \mathbf{v}_{28}$ and $\mathbf{w}_1, \dots, \mathbf{w}_{28}$ be two ETFs in \mathbb{R}^7 . Then, there exist signs $1 = s_1, s_2, \dots, s_{28} \in \{\pm 1\}$ and $\mathbf{Q} \in \mathcal{O}(7)$ such that $\mathbf{w}_i = s_i \mathbf{Q} \mathbf{v}_i$ for each $i \in [28]$.*

The associated graph G is called the *Schläfli graph*, a remarkably symmetrical 16-regular graph on 27 vertices that describes, among other structures, the incidences of the 27 lines on a cubic surface. See, e.g., [Cam80, CS05] for further examples of its structure and significance in combinatorics. We thus propose referring to these inequalities as *Schläfli inequalities*.

We now describe the computer-assisted verification of the degree 4 SOS proof of these inequalities. As above, let $\mathbf{v}_1, \dots, \mathbf{v}_{28} \in \mathbb{R}^7$ form an ETF with $\langle \mathbf{v}_1, \mathbf{v}_i \rangle > 0$ for all $i \neq 1$, and let $\mathbf{Z} = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_{28})$. Such \mathbf{Z} is unique by the above. Since if $\mathbf{M} \in \mathcal{E}_4^n$ then $\mathbf{DMD} \in \mathcal{E}_4^n$ for any $\mathbf{D} = \text{diag}(\mathbf{d})$ with $\mathbf{d} \in \{\pm 1\}^n$, it suffices to fix this single ETF of 28 vectors in \mathbb{R}^7 and check (7.130), and the result will follow for all ETFs of the same dimensions. Let G be the graph on [28] where i and j are adjacent if $\langle \mathbf{v}_i, \mathbf{v}_j \rangle > 0$, so that G is the Schläfli graph with one extra vertex added that is attached to every other vertex. We will write $G|_S$ for the subgraph induced by G on the set of vertices S .

We show (7.130) by producing a $\mathbf{0} \leq \mathbf{A} \in \mathbb{R}^{28^2 \times 28^2}$ such that, for any \mathbf{Y} a degree 4 redundant pseudomoment matrix extending some \mathbf{M} a degree 2 pseudomoment matrix,

$$0 \leq \langle \mathbf{A}, \mathbf{Y} \rangle = 112 - \sum_{1 \leq i < j \leq 28} \text{sgn}(Z_{ij}) X_{ij}. \quad (7.137)$$

The construction of \mathbf{A} is based on studying the results of numerical experiments. We iden-

tify the constants appearing in \mathbf{A} as

$$\gamma_1 := \frac{1}{126}, \quad (7.138)$$

$$\gamma_2 := \frac{1}{36}, \quad (7.139)$$

$$\kappa_1 := \frac{2}{9}, \quad (7.140)$$

$$\kappa_2 := \frac{1}{28}. \quad (7.141)$$

With this, we define

$$A_{(ij)(k\ell)} := \begin{cases} 0 & : |\{i, j, k, \ell\}| = 4, \\ -\operatorname{sgn}(Z_{k\ell})\gamma_1 & : i = j, k \neq \ell, \\ \gamma_2 & : i = k, j \neq \ell, |E(G|_{\{i,j,\ell}\})| = 0, \\ \gamma_2 & : i = k, j \neq \ell, |E(G|_{\{i,j,\ell}\})| = 2, i \sim j, i \sim \ell, \\ -\gamma_2 & : i = k, j \neq \ell, |E(G|_{\{i,j,\ell}\})| = 2, j \sim \ell, \\ 0 & : i = k, j \neq \ell, |E(G|_{\{i,j,\ell}\})| \in \{1, 3\}, \\ -\operatorname{sgn}(Z_{i\ell})\gamma_1 & : i = j = k, i \neq \ell, \\ \kappa_1 & : i = k, j = \ell, i \neq j, \\ \kappa_2 & : i = j, k = \ell. \end{cases} \quad (7.142)$$

$$A_{(i_1 i_2)(i_3 i_4)} = A_{(i_{\pi(1)} i_{\pi(2)})(i_{\pi(3)} i_{\pi(4)})} \text{ for } \mathbf{i} \in [28]^4, \pi \in \operatorname{Sym}(4). \quad (7.143)$$

We then perform a computer verification that $\mathbf{A} \succeq \mathbf{0}$ using the SageMath software package for symbolic calculation of a Cholesky decomposition (more precisely, we first compute the row space symbolically, reduce to this space which has the effect of removing the kernel of \mathbf{A} , and then verify strict positivity with a Cholesky decomposition). Verifying that the equality of (7.137) holds is straightforward by counting the occurrences of various terms in $\langle \mathbf{A}, \mathbf{Y} \rangle$. Of course, this proof technique is rather unsatisfying, and it is an open prob-

lem to provide a more principled description of \mathbf{A} and a conceptual proof of its positive semidefiniteness (both for this specific case and for the general case of maximal ETFs for any dimensions they may exist in).

8 | SPECTRAL PSEUDOMOMENT EXTENSIONS

In the previous chapter, we obtained constraints on the spectrum of a degree 4 pseudomoment extension of a low-rank degree 2 pseudomoment matrix, and found that the Gram matrices of ETFs admit extensions by essentially the simplest possible construction that respects these spectral constraints. We now consider how these results might inform pseudomoment constructions for other Gram matrices and for higher degrees of SOS. To do this, we will give another justification for the construction we obtained for ETFs by more general probabilistic reasoning, constructing pseudomoments by introducing *surrogate random tensors* that behave like the “pseudo-random variable” tensors $x^{\otimes d}$ that pseudoexpectations evaluate. This will recover the degree 4 extension that the ETF examples suggest, and will also give an abstract description of a reasonable higher-degree construction. We derive that abstract description in this chapter, then digress to work through a deterministic example where it applies directly in Chapter 9, and finally apply a further heuristic argument to derive a concrete construction that we will apply to random problems in Chapter 10.

SUMMARY AND REFERENCES This chapter describes a reinterpretation of the degree 4 construction suggested in the previous chapter’s Theorem 7.4.5 that is presented in Section 4.2 of [KB20], and the generalization to higher degrees first proposed in Section 5 of the same. We also present part of the further elaboration of these ideas from [Kun20b]; in particular, while [KB20] proposed a construction in terms of symmetric tensors, [Kun20b] translated

it to the language of homogeneous polynomials and their Hilbert space structure under the apolar inner product, and showed how that construction relates to ideal-harmonic decompositions. The main results of this chapter are a construction summarized in two equivalent ways in Lemma 8.3.11 and Corollary 8.3.13.

8.1 NOTATIONS AND ASSUMPTIONS

We first introduce some notations for and assumptions on the degree 2 pseudomoment matrix that we will make in discussing our proposed pseudomoment extensions. These assumptions will also be in force (either approximately or exactly) in our applications in later chapters.

Suppose $\mathbf{M} \in \mathbb{R}_{\text{sym}}^{n \times n}$ with $\mathbf{M} \succeq \mathbf{0}$ and $M_{ii} = 1$ for all $i \in [n]$. We will be seeking to build pseudoexpectations $\tilde{\mathbb{E}}$ with $\tilde{\mathbb{E}}[\mathbf{x}\mathbf{x}^\top] = \mathbf{M}$. Since $\mathbf{M} \succeq \mathbf{0}$, we may further suppose that, for some $\mathbf{V} \in \mathbb{R}^{r \times n}$ with $r \leq n$ and having full row rank, $\mathbf{M} = \mathbf{V}^\top \mathbf{V}$. In particular then, $\text{rank}(\mathbf{M}) = r$. Since this number will come up repeatedly, we denote the ratio between the rank of \mathbf{M} and the ambient dimension by

$$\delta := \frac{r}{n}. \tag{8.1}$$

Writing $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$ for the columns of \mathbf{V} , we see that $\mathbf{M} = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n)$, and, since $\text{diag}(\mathbf{M}) = \mathbf{1}_n$, $\|\mathbf{v}_i\| = 1$ for all $i \in [n]$.

We now formulate our key assumptions on \mathbf{M} . For the purposes of our derivations in this section, it will suffice to leave the “approximate” statements below vague.

Assumption 8.1.1 (Informal). *The following spectral conditions on \mathbf{M} hold:*

1. *All non-zero eigenvalues of \mathbf{M} , of which there are r , are approximately equal (to $\text{tr}(\mathbf{M})/r = n/r = \delta^{-1}$).*

2. M is approximately equal to a projection matrix to an r -dimensional subspace of \mathbb{R}^n , multiplied by δ^{-1} .
3. $VV^\top \approx \delta^{-1}I_r$.
4. The vectors v_1, \dots, v_n approximately form a unit-norm tight frame (Definition 7.4.1).

(Note that, per our earlier discussion surrounding Definition 7.4.1, if we require that these conditions hold exactly, then they are all equivalent.)

Assumption 8.1.2 (Informal). *The following entrywise condition on M holds:*

5. For any $i \neq j$, $|M_{ij}| \approx \sqrt{(\delta^{-1} - 1)/r}$.

We will see that, to derive a pseudomoment extension of M , we may reason as if the approximate equalities are exact and obtain a sound result.

In light of Condition 3 above, it will be useful to define a normalized version of V , whose rows have approximately unit norm: we let $\widehat{V} := \delta^{1/2}V$, so that $\widehat{V}\widehat{V}^\top \approx I_r$. We note that this matrix can therefore be extended, by adding rows, to an orthogonal matrix (this is equivalent to the *Naimark complement* construction in frame theory; see Section 2.8 of [Wal18]).

8.2 GAUSSIAN CONDITIONING INTERPRETATION AT DEGREE 4

Recall that, in Section 7.4 and its Theorem 7.4.5, we showed that the Gram matrices of most ETFs admit an extension to degree 4 pseudomoment matrices. Matrices M satisfying Assumptions 8.1.1 and 8.1.2 are “approximate ETF Gram matrices,” in the sense that they are close to constant multiples of projection matrices and their off-diagonal entries are close to equal in magnitude. Therefore, taking the pseudomoment extension of Theorem 7.4.5 and simplifying its coefficients when $r = \delta n$ and $n \rightarrow \infty$ with δ fixed, we obtain the prediction

$$\widetilde{\mathbb{E}}[x_i x_j x_k x_\ell] = M_{ij}M_{kl} + M_{ik}M_{jl} + M_{il}M_{jk} - 2 \sum_{a=1}^n M_{ia}M_{ja}M_{ka}M_{la}. \quad (8.2)$$

We note in passing that the construction has “tuned” the coefficient -2 of the second term to produce a cancellation that is required for $\tilde{\mathbb{E}}$ to satisfy $\tilde{\mathbb{E}}[1] = 1$ and $\tilde{\mathbb{E}}[(x_i^2 - 1)p(\mathbf{x})] = 0$. Namely, if $i = j = k = \ell$, then only the term $a = i$ contributes in the latter sum, so the value is $\tilde{\mathbb{E}}[x_i^4] \approx 3 - 2 = 1$, as needed. On the other hand, if $i = j$ and $k = \ell$ but these two values are not equal, then the first term is 1, while all other terms are of sub-constant order, so $\tilde{\mathbb{E}}[x_i^2 x_k^2] \approx 1$ as well. We will discuss a broad generalization of these circumstances in Chapter 10.

This expression in hand, we could proceed to prove degree 4 lower bounds—this is what is done for the Grigoriev-Laurent pseudomoments (see Chapter 9) in [BK18], and no further insight is needed to reproduce the argument of [KB20] either. However, it seems difficult to generalize this idea to find pseudomoment constructions at higher degrees, since the constraints on pseudomoment extensions do not appear strong enough to immediately yield extensions of arbitrary ETFs in closed form. Thus below we give another, perhaps more principled argument through which we arrive at the same prediction of degree 4 pseudomoments. The remainder of this chapter will then be dedicated to showing that this latter construction in fact does generalize sensibly to higher degrees. Our discussion will be slightly redundant as the present derivation is a special case of the higher-degree derivation to come, but, as we will see, the degree 4 case is actually exceptionally simple and it is instructive to perform it with “bare hands” before introducing more machinery.

Let us suppose that the spectral constraints of Assumption 8.1.1 hold exactly; that is, \mathbf{M} is exactly the Gram matrix of a UNTF, so $\mathbf{M} = \delta^{-1} \mathbf{P}$ for \mathbf{P} a projection matrix satisfying $\text{diag}(\mathbf{P}) = \delta \mathbf{1}_n$. Note also that in this case $\mathbf{P} = \widehat{\mathbf{V}}^\top \widehat{\mathbf{V}}$. The key idea is to view the pseudomoments $\tilde{\mathbb{E}}[x_i x_j x_k x_\ell]$ as being given by the *actual* moments of the entries of an $n \times n$ Gaussian random matrix \mathbf{G} . That is, while \mathbf{x} is only a “pseudo-random variable,” we propose that $\mathbf{x}\mathbf{x}^\top$ may be identified with a genuine random matrix, albeit one of rank greater

than one:

$$“\mathbf{x}\mathbf{x}^\top = \mathbf{G}.” \quad (8.3)$$

To write such claims precisely in terms of pseudoexpectations, it is convenient to introduce a weakened notion.

Definition 8.2.1 (Bilinear pseudoexpectation). *For a bilinear operator $\tilde{\mathbb{E}} : \mathbb{R}[x_1, \dots, x_n]_{\leq d} \times \mathbb{R}[x_1, \dots, x_n]_{\leq d} \rightarrow \mathbb{R}$, we denote its action with parentheses, $\tilde{\mathbb{E}}(p(\mathbf{x}), q(\mathbf{x}))$. For a linear operator $\tilde{\mathbb{E}} : \mathbb{R}[x_1, \dots, x_n]_{\leq 2d} \rightarrow \mathbb{R}$, we denote its action with brackets, $\tilde{\mathbb{E}}[p(\mathbf{x})]$. We call a symmetric bilinear operator as above a bilinear pseudoexpectation if it satisfies the following properties:*

1. $\tilde{\mathbb{E}}(1, 1) = 1$;
2. $\tilde{\mathbb{E}}((x_i^2 - 1)p(\mathbf{x}), q(\mathbf{x})) = 0$ for all $p \in \mathbb{R}[x_1, \dots, x_n]_{\leq d-2}$, $q \in \mathbb{R}[x_1, \dots, x_n]_{\leq d}$, and $i \in [n]$;
3. $\tilde{\mathbb{E}}(p(\mathbf{x}), p(\mathbf{x})) \geq 0$ for all $p \in \mathbb{R}[x_1, \dots, x_n]_{\leq d}$.

We say a bilinear pseudoexpectation factors through multiplication if $\tilde{\mathbb{E}}(p(\mathbf{x}), q(\mathbf{x}))$ depends only on the product $p(\mathbf{x})q(\mathbf{x})$.

In terms of a pseudomoment matrix, factoring through multiplication is equivalent to forming a Hankel matrix (or a multidimensional generalization thereof); however, we adopt this language to emphasize the underlying algebraic property. Clearly, a bilinear pseudoexpectation that factors through multiplication in fact yields a true pseudoexpectation (Definition 6.1.2). However, it will be useful for us to work with bilinear pseudoexpectations that do *not* necessarily have this property.

Returning to our construction, given some $\mathbf{G} \in \mathbb{R}^{n \times n}$ with random jointly Gaussian

entries, we propose a bilinear pseudoexpectation

$$\tilde{\mathbb{E}}(\mathbf{x}_i \mathbf{x}_j, \mathbf{x}_k \mathbf{x}_\ell) := \mathbb{E}[G_{ij} G_{k\ell}]. \quad (8.4)$$

(To be fully precise, we also let $\tilde{\mathbb{E}}(\mathbf{x}_i, \mathbf{x}_j) = \tilde{\mathbb{E}}(1, \mathbf{x}_i \mathbf{x}_j) = M_{ij}$ as an extension of \mathbf{M} must satisfy, $\tilde{\mathbb{E}}(1, 1) = 1$, and extend by linearity.) We then design \mathbf{G} so that $\tilde{\mathbb{E}}$ automatically satisfies the constraints of Definition 8.2.1, not including factoring through multiplication. We will have no reason to expect *a priori* that $\tilde{\mathbb{E}}$ should factor through multiplication and yield a true pseudoexpectation, but, remarkably, this will still happen. In exchange for the difficulty of ensuring that $\tilde{\mathbb{E}}$ factors through multiplication, it will be easy to ensure that the other constraints are satisfied. Most importantly, a pseudoexpectation of the above form is, by construction, positive semidefinite.

To specify the law of \mathbf{G} , we first note that taking \mathbf{G} to be symmetric at least ensures that $\tilde{\mathbb{E}}$ will in fact be a well-defined operator on homogeneous degree 2 polynomials. We then begin with a matrix $\mathbf{G}^{(0)}$ having a canonical Gaussian distribution for symmetric matrices, the GOE, suitably rescaled to allow us a normalizing degree of freedom later: $G_{ii}^{(0)} \sim \mathcal{N}(0, 2\sigma^2)$ and $G_{ij}^{(0)} = G_{ji}^{(0)} \sim \mathcal{N}(0, \sigma^2)$. Next, we take \mathbf{G} to have the distribution of $\mathbf{G}^{(0)}$, conditional on the following two properties:

1. $(\mathbf{I}_n - \mathbf{P})\mathbf{G} = 0$.
2. $G_{ii} = 1$ for all $i \in [n]$.

Property 2 ensures that the second condition of Definition 8.2.1 holds, $\tilde{\mathbb{E}}(\mathbf{x}_i^2 - 1, \mathbf{x}_k \mathbf{x}_\ell) = 0$. Property 1 ensures that the similar condition $\tilde{\mathbb{E}}(((\mathbf{I}_n - \mathbf{P})\mathbf{x})_i \mathbf{x}_j, \mathbf{x}_k \mathbf{x}_\ell) = 0$ holds as well. Intuitively, this corresponds to the constraint that \mathbf{x} lies in the row space of \mathbf{P} (which equals that of \mathbf{M}), which is reasonable recalling that we seek to achieve $\tilde{\mathbb{E}}[\mathbf{x}\mathbf{x}^\top] = \mathbf{M}$. Formally, any $\tilde{\mathbb{E}}$ that factors through multiplication and (switching to linear operator notation) with

$\tilde{\mathbb{E}}[\mathbf{x}\mathbf{x}^\top] = \mathbf{M}$ must satisfy this property, since the matrix $\mathbf{F}^{(k,\ell)} = (\tilde{\mathbb{E}}[x_i x_j x_k x_\ell])_{i,j=1}^n$ formed by “freezing” one index pair (k, ℓ) and letting the other two indices vary must, by positivity of $\tilde{\mathbb{E}}$, satisfy

$$\begin{bmatrix} \mathbf{M} & \mathbf{F}^{(k,\ell)} \\ \mathbf{F}^{(k,\ell)} & \mathbf{M} \end{bmatrix} \geq \mathbf{0}, \quad (8.5)$$

as this is a submatrix of the degree 4 redundant pseudomoment matrix. Thus the row space of $\mathbf{F}^{(k,\ell)}$ is contained in that of \mathbf{M} .

What is the law of the resulting Gaussian matrix \mathbf{G} ? Conditioning on Property 1 yields the law of $\mathbf{P}\mathbf{G}^{(0)}\mathbf{P} = \widehat{\mathbf{V}}(\widehat{\mathbf{V}}^\top \mathbf{G}^{(0)} \widehat{\mathbf{V}}) \widehat{\mathbf{V}}^\top$. By rotational invariance of the GOE, the inner matrix $\widehat{\mathbf{V}}^\top \mathbf{G}^{(0)} \widehat{\mathbf{V}} =: \mathbf{G}^{(1)} \in \mathbb{R}_{\text{sym}}^{r \times r}$ has the same law as the upper left $r \times r$ block of $\mathbf{G}^{(0)}$, i.e., a smaller GOE matrix with the same variance scaling of σ^2 .

Next, we condition on Property 2, or equivalently condition $\mathbf{G}^{(1)}$ on having $\mathbf{v}_i^\top \mathbf{G}^{(1)} \mathbf{v}_i = \langle \mathbf{v}_i \mathbf{v}_i^\top, \mathbf{G}^{(1)} \rangle = 1$. To work with these conditions, it is useful to define an isometry between $\mathbb{R}_{\text{sym}}^{r \times r}$ endowed with the Frobenius inner product and $\mathbb{R}^{r(r+1)/2}$ endowed with the ordinary Euclidean inner product.

Definition 8.2.2 (Isometric vectorization). *Define* $\text{isovec} : \mathbb{R}_{\text{sym}}^{r \times r} \rightarrow \mathbb{R}^{r(r+1)/2}$ *by*

$$\text{isovec}(\mathbf{A}) := \begin{bmatrix} \text{diag}(\mathbf{A}) \\ \sqrt{2} \cdot \text{offdiag}(\mathbf{A}) \end{bmatrix}. \quad (8.6)$$

This indeed satisfies $\langle \text{isovec}(\mathbf{A}), \text{isovec}(\mathbf{B}) \rangle = \langle \mathbf{A}, \mathbf{B} \rangle = \text{tr}(\mathbf{A}\mathbf{B})$. Then, $\mathbf{G}^{(1)}$ has the law of $\text{isovec}^{-1}(\mathbf{g})$ for a Gaussian vector $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, 2\sigma^2 \mathbf{I}_{r(r+1)/2})$. Since isovec is an isometry, we may equivalently condition \mathbf{g} on $\langle \mathbf{g}, \text{isovec}(\mathbf{v}_i \mathbf{v}_i^\top) \rangle = 1$ for each $i \in [n]$. By basic properties of Gaussian conditioning, the resulting law is

$$\mathcal{N} \left(\sum_{i=1}^n ((\mathbf{P}^{\circ 2})^{-1} \mathbf{1}_n)_i \text{isovec}(\mathbf{v}_i \mathbf{v}_i^\top), 2\sigma^2 (\mathbf{I} - \tilde{\mathbf{P}}) \right), \quad (8.7)$$

where $\mathbf{P}^{\circ 2}$ is the Gram matrix of the isovec($\mathbf{v}_i \mathbf{v}_i^\top$) or equivalently the entrywise square of \mathbf{P} , and $\tilde{\mathbf{P}}$ is the orthogonal projector to the span of the isovec($\mathbf{v}_i \mathbf{v}_i^\top$). Let $\mathbf{G}^{(2)}$ be a matrix with the law of isovec $^{-1}$ applied to the law in (8.7).

Having finished the conditioning calculations, we may now obtain the statistics of \mathbf{G} . Recall that $G_{ij} = \mathbf{v}_i^\top \mathbf{G}^{(2)} \mathbf{v}_j = \langle \frac{1}{2}(\mathbf{v}_i \mathbf{v}_j^\top + \mathbf{v}_j \mathbf{v}_i^\top), \mathbf{G}^{(2)} \rangle$. Applying isovec to each matrix and using the expression derived above, we find the mean and covariance

$$\mathbb{E}[G_{ij}] = \sum_{a=1}^n ((\mathbf{P}^{\circ 2})^{-1} \mathbf{1}_n)_k P_{ia} P_{ja}, \quad (8.8)$$

$$\text{Cov}[G_{ij}, G_{k\ell}] = \frac{\sigma^2}{2} \text{isovec}(\mathbf{v}_i \mathbf{v}_j^\top + \mathbf{v}_j \mathbf{v}_i^\top)^\top (\mathbf{I} - \tilde{\mathbf{P}}) \text{isovec}(\mathbf{v}_k \mathbf{v}_\ell^\top + \mathbf{v}_\ell \mathbf{v}_k^\top). \quad (8.9)$$

Next, we make two simplifying approximations. For the means, we approximate

$$\mathbf{P}^{\circ 2} \approx \delta \mathbf{I} + \frac{\delta}{r} \mathbf{1}_n \mathbf{1}_n^\top, \quad (8.10)$$

which gives

$$\mathbb{E}[A_{ij}] \approx \delta^{-1} P_{ij} = M_{ij}. \quad (8.11)$$

For the covariances, since under our assumptions we have $\|\text{isovec}(\mathbf{v}_i \mathbf{v}_i^\top)\|_2 = \|\mathbf{v}_i \mathbf{v}_i^\top\|_F = \|\mathbf{v}_i\|_2^2 = \delta$, we approximate

$$\tilde{\mathbf{P}} \approx \delta^{-2} \sum_{i=1}^n \text{isovec}(\mathbf{v}_i \mathbf{v}_i^\top) \text{isovec}(\mathbf{v}_i \mathbf{v}_i^\top)^\top, \quad (8.12)$$

which gives

$$\text{Cov}[A_{ij}, A_{k\ell}] \approx \sigma^2 \left(P_{ik} P_{j\ell} + P_{i\ell} P_{jk} - 2\delta^{-2} \sum_{a=1}^n P_{ia} P_{ja} P_{ka} P_{\ell a} \right). \quad (8.13)$$

Finally, to recover what this prediction implies for the pseudoexpectation values, we

compute

$$\begin{aligned}
\tilde{\mathbb{E}}(x_i x_j, x_k x_\ell) &= \mathbb{E}[G_{ij} G_{k\ell}] \\
&= \mathbb{E}[G_{ij}] \mathbb{E}[G_{k\ell}] + \text{Cov}[G_{ij}, G_{k\ell}] \\
&= M_{ij} M_{k\ell} + \sigma^2 \left(P_{ik} P_{j\ell} + P_{i\ell} P_{jk} - 2\delta^{-2} \sum_{a=1}^n P_{ia} P_{ja} P_{ka} P_{\ell a} \right). \tag{8.14}
\end{aligned}$$

We can then choose σ^2 either such that $\tilde{\mathbb{E}}(x_i^2, x_i^2) = 1$ or such that $\tilde{\mathbb{E}}$ factors through multiplication, which turn out to be the same choice $\sigma^2 = \delta^{-2}$. Thus, writing $\tilde{\mathbb{E}}[x_i x_j x_k x_\ell] = \tilde{\mathbb{E}}(x_i x_j, x_k x_\ell)$ and restricting to $i \neq j$ and $k \neq \ell$, we recover the same formula as (8.2):

$$\text{“}\tilde{\mathbb{E}}[x_i x_j x_k x_\ell] = M_{ij} M_{k\ell} + M_{ik} M_{j\ell} + M_{i\ell} M_{jk} - 2 \sum_{a=1}^n M_{ia} M_{ja} M_{ka} M_{\ell a}.\text{”} \tag{8.15}$$

Remark 8.2.3. *It is worth noting the intriguing geometric interpretation of the random matrix \mathbf{G} we have constructed: we have $\mathbb{E}\mathbf{G} = \mathbf{M}$, $\text{diag}(\mathbf{G}) = \mathbf{1}_n$ deterministically, and \mathbf{G} fluctuates in the linear subspace $\text{pert}_{\mathcal{E}_2^n}(\mathbf{M})$ (as may be verified from the covariance formula (8.9) and is intuitive by analogy with the ETF case of Chapter 7; recall Definition 7.2.12 and the following results). Thus, \mathbf{G} behaves, roughly speaking, like a random element of \mathcal{E}_2^n , which lies on the same face of \mathcal{E}_2^n as \mathbf{M} and fluctuates as an isotropic Gaussian centered at \mathbf{M} along this face (but is allowed to fluctuate “off the edges” of the face). In this regard our construction is an enhanced version of the naive attempt $\tilde{\mathbb{E}}[x_i x_j x_k x_\ell] = M_{ij} M_{k\ell}$, which of course does not satisfy the necessary symmetries. Instead of merely extending a single feasible point \mathbf{M} in this way, we instead extend a random ensemble fluctuating along the same face; if \mathbf{M} is optimal for some degree 2 SOS relaxation, then the other elements of its face should also be optimal, so this is a case of making pseudomoment constructions “as random as possible,” to borrow the phrase of [BHK⁺ 19].*

8.3 GENERALIZING TO HIGHER DEGREE

Now, motivated by the second derivation of the degree 4 extension in the previous section, we propose a generalization to higher degrees. This is performed by replacing the random symmetric matrix \mathbf{G} above with a random symmetric tensor of higher order. Thus before proceeding we review some background on symmetric tensors and the canonical Gaussian distributions they admit. We will also find it useful to take advantage of the equivalence between symmetric tensors and homogeneous polynomials, so we review this below as well. The material on symmetric tensors is standard and may be found in references such as [BS84, KM89]. The Hilbert space structure for homogeneous polynomials that we discuss is apparently a more obscure topic, to which the reader may wish to pay particular attention; we will give some references below.

8.3.1 SYMMETRIC TENSORS AND HOMOGENEOUS POLYNOMIALS

The vector space of *symmetric d -tensors* $\text{Sym}^d(\mathbb{R}^n) \subset (\mathbb{R}^n)^{\otimes d}$ is the subspace of d -tensors whose entries are invariant under permutations of the indices. The vector space of *homogeneous degree d polynomials* $\mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]_d^{\text{hom}}$ is the subspace of degree d polynomials whose monomials all have total degree d . Having the same dimension $\binom{n+d-1}{d}$, these two vector spaces are isomorphic; a natural correspondence $\Phi : \text{Sym}^d(\mathbb{R}^n) \rightarrow \mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]_d^{\text{hom}}$ is

$$\Phi(\mathbf{A}) = \mathbf{A}[\mathbf{y}, \dots, \mathbf{y}] = \sum_{s \in [n]^d} A_s \mathbf{y}^s, \quad (8.16)$$

$$\Phi^{-1}(p(\mathbf{y}))_s = \binom{d}{\text{freq}(\mathbf{s})}^{-1} \cdot [\mathbf{y}^s](p), \quad (8.17)$$

where $\text{freq}(\mathbf{s})$ is the sequence of integers giving the number of times different indices occur in \mathbf{s} (sometimes called the *derived partition*) and $[\mathbf{y}^s](p)$ denotes the extraction of the

coefficient of \mathbf{y}^s from a polynomial.

The map $\text{sym} : (\mathbb{R}^n)^{\otimes d} \rightarrow \text{Sym}^d(\mathbb{R}^n)$ is given by averaging the entries over permutations:

$$\text{sym}(\mathbf{A})_s := \frac{1}{d!} \sum_{\pi \in S_d} A_{s_{\pi(1)} \cdots s_{\pi(d)}}. \quad (8.18)$$

The *symmetric product* is defined by composing the tensor product with symmetrization:

$$\mathbf{A} \odot \mathbf{B} := \text{sym}(\mathbf{A} \otimes \mathbf{B}). \quad (8.19)$$

This is easily seen to coincide with multiplication of polynomials through Φ ,

$$\Phi(\mathbf{A} \odot \mathbf{B}) = \Phi(\mathbf{A})\Phi(\mathbf{B}) \quad (8.20)$$

The general d -tensors $(\mathbb{R}^n)^{\otimes d}$ may be made into a Hilbert space by equipping them with the *Frobenius inner product*,

$$\langle \mathbf{A}, \mathbf{B} \rangle := \sum_{s \in [n]^d} A_s B_s. \quad (8.21)$$

The symmetric d -tensors inherit this inner product. To account for the permutation symmetry, it is useful to introduce the following notation for multisets.

Definition 8.3.1 (Multisets). *For a set S , let $\mathcal{M}(S)$ be the set of multisets with elements in S (equivalently, a function $S \rightarrow \mathbb{N}$), $\mathcal{M}_d(S)$ the multisets of size exactly d , and $\mathcal{M}_{\leq d}(S)$ the multisets of size at most d .*

Then, when $\mathbf{A}, \mathbf{B} \in \text{Sym}^d(\mathbb{R}^n)$, the Frobenius inner product may be written

$$\langle \mathbf{A}, \mathbf{B} \rangle = \sum_{S \in \mathcal{M}_d([n])} \binom{d}{\text{freq}(S)} A_S B_S, \quad (8.22)$$

Perhaps less well-known¹ is the inner product induced on homogeneous degree d poly-

¹[ER93] write: "...the notion of apolarity has remained sealed in the well of oblivion."

nomials by the Frobenius inner product pulled back through the mapping Φ , which is called the *apolar inner product* [ER93, Rez96, Veg00].² For the sake of clarity, we distinguish this inner product with a special notation:

$$\langle p, q \rangle_{\circ} := \langle \Phi^{-1}(p), \Phi^{-1}(q) \rangle = \sum_{S \in \mathcal{M}_d([n])} \binom{d}{\text{freq}(S)}^{-1} \cdot [\mathbf{y}^S](p) \cdot [\mathbf{y}^S](q). \quad (8.23)$$

In the sequel we also follow the standard terminology of saying that “ p and q are apolar” when $\langle p, q \rangle_{\circ} = 0$; we also use this term more generally to refer to orthogonality under the apolar inner product, speaking of apolar subspaces, apolar projections, and so forth.

The most important property of the apolar inner product that we will use is that multiplication and differentiation are adjoint to one another. We follow here the expository note [Rez96], which presents applications of this idea to PDEs, a theme we will develop further below. The basic underlying fact is the following. For $q \in \mathbb{R}[\mathcal{y}_1, \dots, \mathcal{y}_n]$, write $q(\partial) = q(\partial_{\mathcal{y}_1}, \dots, \partial_{\mathcal{y}_n})$ for the associated differential operator.³

Proposition 8.3.2 (Theorem 2.11 of [Rez96]). *Suppose $p, q, r \in \mathbb{R}[\mathcal{y}_1, \dots, \mathcal{y}_n]^{\text{hom}}$, with degrees $\deg(p) = a, \deg(q) = b$, and $\deg(r) = a + b$. Then,*

$$\langle pq, r \rangle_{\circ} = \frac{a!}{(a+b)!} \langle p, q(\partial)r \rangle_{\circ}. \quad (8.24)$$

In particular, if $\deg(p) = \deg(q) = a$, then $\langle p, q \rangle_{\circ} = p(\partial)q/a!$.

In fact, it will later be useful for us to define the following rescaled version of the apolar inner product that omits the constant factor above.

Definition 8.3.3. *For $p, q \in \mathbb{R}[\mathcal{y}_1, \dots, \mathcal{y}_n]^{\text{hom}}$ with $\deg(p) = \deg(q)$, let $\langle p, q \rangle_{\partial} := p(\partial)q$.*

²Other names used in the literature for this inner product include the *Bombieri*, *Bombieri-Weyl*, *Fischer*, or *Sylvester* inner product. The term *apolar* itself refers to polarity in the sense of classical projective geometry; see [ER93] for a historical overview in the context of invariant theory.

³If, for instance, $q(\mathbf{y}) = \mathcal{y}_1^2 \mathcal{y}_2 + \mathcal{y}_3^3$, then $q(\partial)f = \frac{\partial^3 f}{\partial \mathcal{y}_1^2 \partial \mathcal{y}_2} + \frac{\partial^3 f}{\partial \mathcal{y}_3^3}$.

Using the preceding formula, we also obtain the following second important property, that of invariance under orthogonal changes of monomial basis.

Proposition 8.3.4. *Suppose $p, q \in \mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]_d^{\text{hom}}$ and $Q \in \mathcal{O}(n)$. Then,*

$$\langle p(\mathbf{y}), q(\mathbf{y}) \rangle_{\circ} = \langle p(Q\mathbf{y}), q(Q\mathbf{y}) \rangle_{\circ}. \quad (8.25)$$

Associated to these two Hilbert space structures, we may then define isotropic Gaussian “vectors” (tensors or polynomials).

Definition 8.3.5. *For $\sigma > 0$, $\mathcal{G}_d^{\text{tens}}(n, \sigma^2)$ is the unique centered Gaussian measure over $\text{Sym}^d(\mathbb{R}^n)$ such that, when $\mathbf{G} \sim \mathcal{G}_d^{\text{tens}}(n, \sigma^2)$, then for any $\mathbf{A}, \mathbf{B} \in \text{Sym}^d(\mathbb{R}^n)$,*

$$\mathbb{E}[\langle \mathbf{A}, \mathbf{G} \rangle \langle \mathbf{B}, \mathbf{G} \rangle] = \sigma^2 \langle \mathbf{A}, \mathbf{B} \rangle. \quad (8.26)$$

Equivalently, the entries of \mathbf{G} have laws $G_s \sim \mathcal{N}(0, \sigma^2 / \binom{d}{\text{freq}(s)})$ and are independent up to equality under permutations. Equivalently again, letting $\mathbf{G}^{(0)} \in (\mathbb{R}^n)^{\otimes d}$ have i.i.d. entries distributed as $\mathcal{N}(0, \sigma^2)$, $\mathbf{G} = \text{sym}(\mathbf{G}^{(0)})$.

For example, the Gaussian orthogonal ensemble with the scaling we have used previously is $\text{GOE}(n) = \mathcal{G}_2^{\text{tens}}(n, 2/n)$. The tensor ensembles have also been used by [RM14] and subsequent works on tensor PCA under the name “symmetric standard normal” tensors.

Definition 8.3.6. *For $\sigma > 0$, $\mathcal{G}_d^{\text{poly}}(n, \sigma^2)$ is the unique measure on polynomials with centered Gaussian coefficients in $\mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]_d^{\text{hom}}$ such that, when $g \sim \mathcal{G}_d^{\text{poly}}(n, \sigma^2)$, then for any $p, q \in \mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]_d^{\text{hom}}$,*

$$\mathbb{E}[\langle p, g \rangle_{\circ} \langle q, g \rangle_{\circ}] = \sigma^2 \langle p, q \rangle_{\circ}. \quad (8.27)$$

Equivalently, the coefficients of g are independent with laws $[\mathbf{y}^s](g) \sim \mathcal{N}(0, \sigma^2 \binom{d}{\text{freq}(s)})$.

See [Kos02] for references to numerous works and results on this distribution over polynomials, and justification for why it is “the most natural random polynomial.” Perhaps the main reason is that, as a corollary of Proposition 8.3.4, this polynomial is orthogonally invariant (unlike, say, a superficially simpler-looking random polynomial with i.i.d. coefficients).

Proposition 8.3.7. *If $g \sim \mathcal{G}_d^{\text{poly}}(n, \sigma^2)$ and $Q \in \mathcal{O}(n)$, then $g \stackrel{(d)}{=} g \circ Q$.*

(Likewise, though we will not use it, $\mathcal{G}_d^{\text{tens}}(n, \sigma^2)$ is invariant under contraction of each index with the same orthogonal matrix, generalizing the orthogonal invariance of the GOE.)

Finally, by the isotropy properties and the isometry of apolar and Frobenius inner products under Φ , we deduce that these two Gaussian laws are each other’s pullbacks under those correspondences.

Proposition 8.3.8. *If $G \sim \mathcal{G}_d^{\text{tens}}(n, \sigma^2)$, then $\Phi(G)$ has the law $\mathcal{G}_d^{\text{poly}}(n, \sigma^2)$. Conversely, if $g \sim \mathcal{G}_d^{\text{poly}}(n, \sigma^2)$, then $\Phi^{-1}(g)$ has the law $\mathcal{G}_d^{\text{tens}}(n, \sigma^2)$.*

8.3.2 SURROGATE RANDOM TENSOR CONSTRUCTION

We now proceed to generalize our degree 4 construction to higher degrees. As for degree 4, the initial idea is to build pseudoexpectation values as second moments of the entries of a Gaussian random symmetric tensor. That is, for degree $2d$, we build $\tilde{\mathbb{E}}$ using a random $G^{(d)} \in \text{Sym}^d(\mathbb{R}^n)$ and taking, for multisets of indices $S, T \in \mathcal{M}_d([n])$,

$$\text{“} \tilde{\mathbb{E}}(\mathbf{x}^S, \mathbf{x}^T) := \mathbb{E} \left[G_S^{(d)} G_T^{(d)} \right]. \text{”} \tag{8.28}$$

Again, at an intuitive level, if \mathbf{x} has the pseudodistribution encoded by $\tilde{\mathbb{E}}$, then one should think of identifying

$$\text{“} G^{(d)} = \mathbf{x}^{\otimes d}. \text{”} \tag{8.29}$$

The key point is that, while we cannot model the values of $\tilde{\mathbb{E}}$ as being the moments of an actual random vector \mathbf{x} , we can model them as being the moments of the tensors $\mathbf{G}^{(d)}$, which are surrogates for the tensor powers of \mathbf{x} .

Forcing $\mathbf{G}^{(d)}$ to be a *symmetric* tensor, as we have indicated above, makes $\tilde{\mathbb{E}}$ above a well-defined bilinear pseudoexpectation (Definition 8.2.1). We will again begin with $\mathbf{G}^{(d)} \sim \mathcal{G}_d^{\text{tens}}(n, \sigma_d^2)$ for some degree of freedom $\sigma_d > 0$ to be chosen later, and condition on constraints analogous to those given earlier for degree 4. To express these, we introduce the “slicing notation” (or “MATLAB notation”) that, for $S' \in \mathcal{M}_{d-1}([n])$ and $\mathbf{A} \in \text{Sym}^d(\mathbb{R}^n)$, $\mathbf{A}[S', :] := (A_{S'+\{i\}})_{i=1}^n \in \mathbb{R}^n$. Then, our construction is as follows.

PSEUDOEXPECTATION PREDICTION (TENSORS) Let $\mathbf{w}_1, \dots, \mathbf{w}_{n-r}$ be a basis of $\ker(\mathbf{M})$. Define a jointly Gaussian collection of tensors $\mathbf{G}^{(d)} \in \text{Sym}^d(\mathbb{R}^n)$ as follows.

1. $\mathbf{G}^{(0)} \in \text{Sym}^0(\mathbb{R}^n)$ is a scalar, with one entry $G_{\emptyset}^{(0)} = 1$.
2. For $d \geq 1$, $\mathbf{G}^{(d)}$ has the law of $\mathcal{G}_d^{\text{tens}}(n, \sigma_d^2)$, conditional on the following two properties:
 - (a) If $d \geq 2$, then for all $S' \in \mathcal{M}_{d-2}([n])$ and $i \in [n]$, $G_{S'+\{i,i\}}^{(d)} = G_{S'}^{(d-2)}$.
 - (b) For all $S' \in \mathcal{M}_{d-1}([n])$ and $i \in [n-r]$, $\langle \mathbf{w}_i, \mathbf{G}^{(d)}[S', :] \rangle = 0$.

Then, for $S, T \in \mathcal{M}_d([n])$, set

$$\tilde{\mathbb{E}}(\mathbf{x}^S, \mathbf{x}^T) := \mathbb{E} \left[G_S^{(d)} G_T^{(d)} \right]. \quad (8.30)$$

The choice of $G_{\emptyset}^{(0)}$ was implicit in our discussion of degree 4 where we did not explicitly introduce a scalar corresponding to degree 0. The first Property (a) we condition on is the suitable generalization of $G_{ii} = 1$ to higher-order surrogate tensors, and the second Property (b) is the same generalization of $(\mathbf{I} - \mathbf{P})\mathbf{G} = \mathbf{0}$. If it is possible to choose σ_d^2

such that this $\tilde{\mathbb{E}}$ factors through multiplication, then $\tilde{\mathbb{E}}$ is a degree $2d$ pseudoexpectation extending M .

However, when we try to actually carry out the computation of the law of $\mathbf{G}^{(d)}$ after performing this conditioning, matters are subtler once $d > 2$ (corresponding to total SOS degree $2d > 4$). Recall that, before, this amounted to computing the projection to the span of the matrices $\mathbf{v}_i \mathbf{v}_i^\top \in \mathbb{R}_{\text{sym}}^{r \times r}$, which arose from Property (a) fixing $G_{ii}^{(2)}$. Once $d > 2$, Property (a) will instead fix $G_S^{(d)}$ for any S containing a repeated index. Thus the corresponding projection will be to the subspace

$$\text{span}(\{\mathbf{v}_i \circ \mathbf{v}_i \circ \mathbf{v}_{j_1} \circ \cdots \circ \mathbf{v}_{j_{d-2}} : i, j_1, \dots, j_{d-2} \in [n]\}) \subseteq \text{Sym}^d(\mathbb{R}^r). \quad (8.31)$$

While for $d = 2$ there was a convenient basis $\mathbf{v}_i \circ \mathbf{v}_i$ of this subspace, here the above spanning set is linearly dependent (since $\mathbf{v}_1, \dots, \mathbf{v}_n$ have many linear dependences), and there does not appear to be a convenient adjustment of this set to a basis retaining the symmetries of the $d = 2$ case. We therefore turn to the reinterpretation in terms of homogeneous polynomials: in that language, this unusual subspace is nothing but an *ideal*, and we may at least obtain an intrinsic description of our construction using a connection between ideals and *multiharmonic polynomials* under the apolar inner product, which we review next.

8.3.3 HOMOGENEOUS IDEALS AND MULTIHARMONIC POLYNOMIALS

We return to homogeneous polynomials and the apolar inner product, and describe a crucial consequence of Proposition 8.3.2. Namely, for any homogeneous ideal, any polynomial uniquely decomposes into one part belonging to the ideal, and another part, apolar to the first, that is “multiharmonic” in that it satisfies a certain system of PDEs associated to the ideal.⁴

⁴We will use “harmonic” to abbreviate but “multiharmonic” when we wish to explicitly distinguish our case from the case of harmonic polynomials satisfying the single Laplace equation $\Delta q = 0$. Unfortunately,

Proposition 8.3.9. Let $p_1, \dots, p_m \in \mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]^{\text{hom}}$ and $d \geq \max_{i=1}^m \deg(p_i)$. Define two subspaces of $\mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]_d^{\text{hom}}$:

$$V_I := \left\{ \sum_{i=1}^m p_i q_i : q_i \in \mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]_{d-\deg p_i}^{\text{hom}} \right\}, \text{ the "ideal subspace," and} \quad (8.32)$$

$$V_{\mathcal{H}} := \{q : p_i(\partial)q = 0 \text{ for all } i \in [m]\}, \text{ the "harmonic subspace."} \quad (8.33)$$

Then, V_I and $V_{\mathcal{H}}$ are orthogonal complements under the apolar inner product. Consequently, $\mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]_d^{\text{hom}} = V_I \oplus V_{\mathcal{H}}$.

Perhaps the most familiar example is the special case of harmonic polynomials, for which this result applies as follows.

Example 8.3.10. Suppose $m = 1$, and $p_1(\mathbf{y}) = \|\mathbf{y}\|_2^2 = \mathcal{Y}_1^2 + \dots + \mathcal{Y}_n^2$. Then, $p_1(\partial) = \Delta$, so Proposition 8.3.9 implies that any $q \in \mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]_d^{\text{hom}}$ may be written uniquely as $p(\mathbf{y}) = q_d(\mathbf{y}) + \|\mathbf{y}\|_2^2 q_{d-2}(\mathbf{y})$ where q_d is harmonic, $\deg(q_d) = d$, and $\deg(q_{d-2}) = d - 2$. Repeating this inductively, we obtain the well-known fact from harmonic analysis that we may in fact expand

$$p(\mathbf{y}) = \sum_{a=0}^{\lfloor d/2 \rfloor} \|\mathbf{y}\|_2^{2a} q_{d-2a}(\mathbf{y}) \quad (8.34)$$

where each q_i is harmonic with $\deg(q_i) = i$ and the q_i are uniquely determined by p .

This is sometimes called the ‘‘Fischer decomposition,’’ see also the ‘‘Expansion Theorem’’ in [Rez96] for a generalization of this type of decomposition.

8.3.4 CONDITIONING BY TRANSLATING TO HOMOGENEOUS POLYNOMIALS

Equipped with these concepts, we proceed with our conditioning computation, after translating the construction to homogeneous polynomials. Passing each $\mathbf{G}^{(d)}$ through the isometry

 the term *multiharmonic function* is also sometimes used to refer to what is usually called a *pluriharmonic function*, the real or imaginary part of a holomorphic function of several variables, or to what is usually called a *polyharmonic function*, one that satisfies $\Delta^m q = 0$ for some $m \in \mathbb{N}$.

between $\text{Sym}^d(\mathbb{R}^n)$ and $\mathbb{R}[\gamma_1, \dots, \gamma_n]_d^{\text{hom}}$ described in Section 8.3.1, we find an equivalent construction in terms of random polynomials $g^{(d)} \in \mathbb{R}[\gamma_1, \dots, \gamma_n]_d^{\text{hom}}$. Though this may seem unnatural at the surface level—bizarrely, we will be defining $\tilde{\mathbb{E}}$ for each degree in terms of the correlations of various *coefficients* of a random polynomial—we recall that we expect viewing the extraction of coefficients in terms of the apolar inner product to bring forth a connection to multiharmonic polynomials per the previous section, allowing us to use a variant of the ideas there to complete the calculation.

PSEUDOEXPECTATION PREDICTION (POLYNOMIALS) Let w_1, \dots, w_{n-r} be a basis of $\ker(M)$. Define a jointly Gaussian collection of polynomials $g^{(d)} \in \mathbb{R}[\gamma_1, \dots, \gamma_n]_d^{\text{hom}}$ as follows.

1. $g^{(0)}(\mathbf{y}) = 1$.
2. For $d \geq 1$, $g^{(d)}$ has the law of $G_d^{\text{poly}}(n, \sigma_d^2)$, conditional on the following two properties:
 - (a) If $d \geq 2$, then for all $i \in [n]$ and $S' \in \mathcal{M}_{d-2}([n])$, $\langle g^{(d)}, \mathbf{y}^{S'} \gamma_i^2 \rangle_{\circ} = \langle g^{(d-2)}, \mathbf{y}^{S'} \rangle_{\circ}$.
 - (b) For all $S' \in \mathcal{M}_{d-1}([n])$ and $i \in [n-r]$, $\langle g^{(d)}, \mathbf{y}^{S'} \langle w_i, \mathbf{y} \rangle \rangle_{\circ} = 0$.

Then, for $S, T \in \mathcal{M}_d([n])$, set

$$\tilde{\mathbb{E}}(\mathbf{x}^S, \mathbf{x}^T) := \mathbb{E} \left[\langle g^{(d)}, \mathbf{y}^S \rangle_{\circ} \langle g^{(d)}, \mathbf{y}^T \rangle_{\circ} \right]. \quad (8.35)$$

The most immediate advantage of reframing our prediction in this way is that it gives us access to the clarifying concepts of “divisibility” and “differentiation,” whose role is obscured by the previous symmetric tensor language. Moreover, these are nicely compatible with the apolar inner product per Proposition 8.3.2.

Let us briefly outline the computation before giving a careful justification. Roughly speaking, conditioning on Property (b) above projects $g^{(d)}$ to the subspace of polynomi-

als depending only on $\widehat{\mathbf{V}}\mathbf{y}$. On the other hand, by Proposition 8.3.7, $g^{(d)}$ is invariant under compositions with orthogonal matrices, and by our assumptions on $\widehat{\mathbf{V}}$, it is the upper $r \times n$ block of some orthogonal matrix. From this, if $g_1^{(d)}$ has the law of $G_d^{\text{poly}}(n, \sigma_d^2)$ conditional on Property (b), then the collection of coefficients $(\langle g_1^{(d)}, \mathbf{y}^S \rangle_{\circ})_{S \in \mathcal{M}_d([n])}$ has the same law as $(\langle h^{(d)}, (\widehat{\mathbf{V}}^\top \mathbf{z})^S \rangle_{\circ})_{S \in \mathcal{M}_d([n])}$ for $h^{(d)}(\mathbf{z}) \sim G_d^{\text{poly}}(r, \sigma_d^2)$. (We use $\mathbf{y} = (y_1, \dots, y_n)$ for formal variables of dimension n and $\mathbf{z} = (z_1, \dots, z_r)$ for formal variables of dimension r .) Thus conditioning on Property (b) is merely a dimensionality reduction of the canonical Gaussian polynomial, in a suitable basis.

Conditioning $h^{(d)}$ as above on Property (a) brings in the ideal and harmonic subspaces discussed in Section 8.3.3. Let us define

$$V_I := \left\{ \sum_{i=1}^n \langle \mathbf{v}_i, \mathbf{z} \rangle^2 q_i(\mathbf{z}) : q_i \in \mathbb{R}[z_1, \dots, z_r]_d^{\text{hom}} \right\}, \quad (8.36)$$

$$V_{\mathcal{H}} := \left\{ \mathbf{q} \in \mathbb{R}[z_1, \dots, z_r]_d^{\text{hom}} : \langle \mathbf{v}_i, \boldsymbol{\partial} \rangle^2 \mathbf{q} = 0 \text{ for all } i \in [n] \right\}, \quad (8.37)$$

instantiations of the subspaces of Proposition 8.3.9 for the specific collection of polynomials $\{\langle \mathbf{v}_i, \mathbf{z} \rangle^2\}_{i=1}^n$. Conditioning on Property (a) fixes the component of h belonging to V_I , leaving a fluctuating part equal to the apolar projection of $h^{(d)}$ to $V_{\mathcal{H}}$. This reasoning yields the following recursion. We give a more careful proof, and then give a closed version of these formulae.

Lemma 8.3.11 (Pseudoexpectation recursion). *Suppose the conditions of Assumption 8.1.1 hold exactly. Let P_I and $P_{\mathcal{H}}$ be the apolar projections to V_I and $V_{\mathcal{H}}$, respectively. For each $S \in \mathcal{M}_d([n])$, let $r_S \in \mathbb{R}[x_1, \dots, x_n]_d^{\text{hom}}$ be a polynomial having*

$$P_I[(\mathbf{V}^\top \mathbf{z})^S] = r_S(\mathbf{V}^\top \mathbf{z}), \quad (8.38)$$

$$r_S(\mathbf{x}) = \sum_{i=1}^n x_i^{2d_i} r_{S,i}(\mathbf{x}) \text{ for } d_i \geq 1, r_{S,i} \in \mathbb{R}[x_1, \dots, x_n]_{d-2d_i}^{\text{hom}}, \quad (8.39)$$

and further define

$$r_S^\perp(\mathbf{x}) = \sum_{i=1}^n r_{S,i}(\mathbf{x}) \in \mathbb{R}[x_1, \dots, x_n]_{\leq d-2}, \quad (8.40)$$

where we emphasize that r_S^\perp is not necessarily homogeneous. Set $h_S(\mathbf{x}) := \mathbf{x}^S - r_S(\mathbf{x})$, whereby $P_{\mathcal{H}}[(\mathbf{V}^\top \mathbf{z})^S] = h_S(\mathbf{V}^\top \mathbf{z})$. Then, the right-hand side of (8.35) is

$$\tilde{\mathbb{E}}(\mathbf{x}^S, \mathbf{x}^T) = \underbrace{\tilde{\mathbb{E}}(r_S^\perp(\mathbf{x}), r_T^\perp(\mathbf{x}))}_{\text{"ideal" term}} + \underbrace{\sigma_d^2 \delta^d \cdot \langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle}_\text{"harmonic" term}. \quad (8.41)$$

Proof. We must compute the distribution of $g^{(d)}(\mathbf{y}) \sim G_d^{\text{poly}}(n, \sigma_d^2)$, given $g^{(0)}(\mathbf{y}) = 1$, conditional on (a) having, if $d \geq 2$, for all $i \in [n]$ and $S' \in \mathcal{M}_{d-2}([n])$ that $\langle g^{(d)}, \mathbf{y}^{S'} y_i^2 \rangle_\circ = \langle g^{(d-2)}, \mathbf{y}^{S'} \rangle_\circ$, and (b) having for all $S' \in \mathcal{M}_{d-1}([n])$ and $i \in [n-r]$ that $\langle g^{(d)}, \mathbf{y}^{S'} \langle \mathbf{w}_i, \mathbf{y} \rangle \rangle_\circ = 0$.

Working first with Property (a), we see after extending by linearity that it is equivalent to $\langle g^{(d)}, q(\mathbf{y}) y_i^2 \rangle_\circ = \langle g^{(d-2)}, q(\mathbf{y}) \rangle_\circ$ for all $i \in [n]$ and $q \in \mathbb{R}[y_1, \dots, y_n]_{d-2}^{\text{hom}}$. On the other hand, by the adjointness property from Proposition 8.3.2, we have $\langle g^{(d)}, q(\mathbf{y}) y_i^2 \rangle_\circ = \frac{1}{d(d-1)} \langle \partial_{y_i}^2 g^{(d)}, q(\mathbf{y}) \rangle_\circ$, and thus Property (a) is equivalent to the simpler property:

$$(a') \text{ If } d \geq 2, \text{ then for all } i \in [n], \partial_{y_i}^2 g^{(d)}(\mathbf{y}) = d(d-1) \cdot g^{(d-2)}(\mathbf{y}).$$

Similarly, we see after extending Property (b) by linearity that it is equivalent to having $\langle g^{(d)}, q(\mathbf{y}) \langle \mathbf{w}, \mathbf{y} \rangle \rangle_\circ = 0$ for all $q \in \mathbb{R}[y_1, \dots, y_n]_{d-1}^{\text{hom}}$ and $\mathbf{w} \in \ker(\mathbf{M})$. Again by Proposition 8.3.2, $\langle g^{(d)}, q(\mathbf{y}) \langle \mathbf{w}, \mathbf{y} \rangle \rangle_\circ = d \langle \langle \mathbf{w}, \partial \rangle g^{(d)}, q(\mathbf{y}) \rangle_\circ$, so Property (b) is equivalent to:

$$(b') \text{ For all } \mathbf{w} \in \ker(\mathbf{M}), \langle \mathbf{w}, \partial \rangle g^{(d)}(\mathbf{y}) = 0.$$

Polynomials $p \in \mathbb{R}[y_1, \dots, y_n]_d^{\text{hom}}$ with $\langle \mathbf{w}, \partial \rangle p = 0$ for all $\mathbf{w} \in \ker(\mathbf{M})$ form a linear subspace, which admits the following simple description. Changing monomial basis to one which extends $z_i = (\widehat{\mathbf{V}} \mathbf{y})_i$ for $i \in [r]$ and invoking Proposition 8.3.4, we may define

$$V_{\mathcal{B}} := \left\{ p \in \mathbb{R}[y_1, \dots, y_n]_d^{\text{hom}} : p(\mathbf{y}) = q(\widehat{\mathbf{V}} \mathbf{y}) \text{ for some } q \in \mathbb{R}[z_1, \dots, z_r]_d^{\text{hom}} \right\}, \quad (8.42)$$

and with this definition Property (b) is equivalent to:

$$(b'') \quad \mathbf{g}^{(d)} \in V_{\mathcal{B}}.$$

Now, let $\mathbf{Q} \in \mathcal{O}(n)$ be an orthogonal matrix formed by adding rows to $\widehat{\mathbf{V}}$ (see Section 8.1 for why this is possible under our assumptions on \mathbf{M}). Letting $a_S \sim \mathcal{N}(0, \sigma_d^2 \binom{d}{\text{freq}(S)})$ for each $S \in \mathcal{M}_d([n])$ independently, we set $\mathbf{g}_0^{(d)}(\mathbf{y}) := \sum_{S \in \mathcal{M}_d([n])} a_S \cdot (\mathbf{Q}\mathbf{y})^S$; then, the law of $\mathbf{g}_0^{(d)}$ is $\mathbf{G}_d^{\text{poly}}(n, \sigma_d^2)$ by Proposition 8.3.7. Thus we may form the law of $\mathbf{g}^{(d)}$ by conditioning $\mathbf{g}_0^{(d)}$ on Properties (a') and (b'').

Conveniently, conditioning $\mathbf{g}_0^{(d)}$ on Property (b'') amounts to merely setting those a_S with $S \cap \{r+1, \dots, n\} \neq \emptyset$ to equal zero. Denoting the resulting random polynomial by $\mathbf{g}_1^{(d)}$, we see that $\mathbf{g}_1^{(d)}(\mathbf{y}) = \sum_{S \in \mathcal{M}_d([r])} a_S \cdot (\mathbf{Q}\mathbf{y})^S = \sum_{S \in \mathcal{M}_d([r])} a_S \cdot (\widehat{\mathbf{V}}\mathbf{y})^S$. To extract the coefficients of $\mathbf{g}_1^{(d)}$ in the standard monomial basis, we compute

$$\begin{aligned} \langle \mathbf{g}_1^{(d)}, \mathbf{y}^S \rangle_\circ &= \sum_{T \in \mathcal{M}_d([r])} a_T \langle (\widehat{\mathbf{V}}\mathbf{y})^T, \mathbf{y}^S \rangle_\circ \\ &= \sum_{T \in \mathcal{M}_d([r])} a_T \langle \mathbf{y}^T, (\mathbf{Q}^\top \mathbf{y})^S \rangle_\circ. \end{aligned} \quad (\text{Proposition 8.3.4})$$

and noting that no term involving y_{r+1}, \dots, y_n will contribute, we may define the truncation $\mathbf{z} = (y_1, \dots, y_r)$ and continue

$$\begin{aligned} &= \sum_{T \in \mathcal{M}_d([r])} a_T \langle \mathbf{z}^T, (\widehat{\mathbf{V}}^\top \mathbf{z})^S \rangle_\circ \\ &= \left\langle \sum_{T \in \mathcal{M}_d([r])} a_T \mathbf{z}^T, (\widehat{\mathbf{V}}^\top \mathbf{z})^S \right\rangle_\circ. \end{aligned} \quad (8.43)$$

Letting $h_0^{(d)}(\mathbf{z}) := \sum_{T \in \mathcal{M}_d([r])} a_T \mathbf{z}^T$, we see that the law of $h_0^{(d)}$ is $\mathbf{G}_d^{\text{poly}}(r, \sigma_d^2)$.

Thus we may rewrite the result of the remaining conditioning on Property (a') by letting $h^{(d)}$ have the law of $h_0^{(d)}$ conditioned on $\langle \mathbf{v}_i, \boldsymbol{\theta} \rangle^2 h^{(d)}(\mathbf{z}) = \delta^{-1} d(d-1) \cdot h^{(d-2)}(\mathbf{z})$ for all

$i \in [n]$. Then,

$$\begin{aligned}\tilde{\mathbb{E}}(\mathbf{x}^S, \mathbf{x}^T) &= \mathbb{E} \left[\langle \mathbf{h}^{(d)}, (\widehat{\mathbf{V}}^\top \mathbf{z})^S \rangle_\circ \langle \mathbf{h}^{(d)}, (\widehat{\mathbf{V}}^\top \mathbf{z})^T \rangle_\circ \right] \\ &= \delta^d \cdot \mathbb{E} \left[\langle \mathbf{h}^{(d)}, (\mathbf{V}^\top \mathbf{z})^S \rangle_\circ \langle \mathbf{h}^{(d)}, (\mathbf{V}^\top \mathbf{z})^T \rangle_\circ \right]\end{aligned}\quad (8.44)$$

The result of this remaining conditioning is simple to write down in a more explicit form, since now we have just a single family of linear constraints to condition on and $\mathbf{h}_0^{(d)}$ is isotropic with respect to the apolar inner product (per Definition 8.3.6). We recall that P_I and $P_{\mathcal{H}}$ are the orthogonal projections to the ideal and harmonic subspaces V_I and $V_{\mathcal{H}}$, respectively, with respect to the polynomials $\langle \mathbf{v}_i, \mathbf{z} \rangle^2$. We also define the “least-squares raising” operator L_I by

$$L_I[h] := \operatorname{argmin} \left\{ \|f\|_\circ^2 : f \in V_I, \langle \mathbf{v}_i, \partial \rangle^2 f = h \text{ for all } i \in [n] \right\}, \quad (8.45)$$

for all h such that $\langle \mathbf{v}_i, \delta \rangle^2 h$ does not depend on i .

We then obtain that the law of $\mathbf{h}^{(d)}$ is

$$\mathbf{h}^{(d)} \stackrel{(d)}{=} \delta^{-1} d(d-1) \cdot L_I[h^{(d-2)}] + P_{\mathcal{H}}[\mathbf{h}_0^{(d)}], \quad (8.46)$$

where we view $\mathbf{h}_0^{(d)}$ as independent of $\mathbf{h}^{(d')}$ for all $d' < d$. Note that the first summand above belongs to V_I and the second to $V_{\mathcal{H}}$, so this is also an ideal-harmonic decomposition for $\mathbf{h}^{(d)}$ precisely of the kind provided by Proposition 8.3.9.

Now, we work towards substituting this into (8.44). To do that, we must compute inner products of the form $\langle \mathbf{h}^{(d)}, (\mathbf{V}^\top \mathbf{z})^S \rangle_\circ$. We introduce a few further observations to do this: for each S , the projection of $(\mathbf{V}^\top \mathbf{z})^S$ to V_I is a linear combination of multiples of the $\langle \mathbf{v}_i, \mathbf{z} \rangle^2$. Moreover, the $\langle \mathbf{v}_i, \mathbf{z} \rangle$ are an overcomplete system of linear polynomials, so *any* polynomial in \mathbf{z} may be written as a polynomial in these variables instead. Combining these facts, we

see that there exists a polynomial $r_S \in \mathbb{R}[x_1, \dots, x_n]_d^{\text{hom}}$ such that

$$P_J[(\mathbf{V}^\top \mathbf{z})^S] = r_S(\mathbf{V}^\top \mathbf{z}), \quad (8.47)$$

$$r_S(\mathbf{x}) = \sum_{i=1}^n x_i^{2d_i} r_{S,i}(\mathbf{x}) \text{ for } d_i \geq 1, r_{S,i} \in \mathbb{R}[x_1, \dots, x_n]_{d-2d_i}^{\text{hom}}. \quad (8.48)$$

The polynomial $r_S(\mathbf{x})$ is not unique, since the vectors \mathbf{v}_i and therefore the polynomials $\langle \mathbf{v}_i, \mathbf{z} \rangle$ are linearly dependent. Nor is the decomposition of r_S into the $x_i^{2d_i} r_{S,i}$ unique, since for instance if some $d_i \geq 2$ then we may move a factor of x_i^2 into $r_{S,i}$. However, any choice of r_S satisfying (8.47) and $r_{S,i}$ satisfying (8.48) suffices for our purposes. Note also that we reuse the pseudodistribution variables $\mathbf{x} = (x_1, \dots, x_n)$ intentionally here because of the role $r_S(\mathbf{x})$ will play below.

With this definition, we compute

$$\begin{aligned} \langle \mathbf{h}^{(d)}, (\mathbf{V}^\top \mathbf{z})^S \rangle_\circ &= \delta^{-1} d(d-1) \cdot \langle L_J[\mathbf{h}^{(d-2)}], P_J[(\mathbf{V}^\top \mathbf{z})^S] \rangle_\circ + \langle P_{\mathcal{H}}[\mathbf{h}_0^{(d)}], (\mathbf{V}^\top \mathbf{z})^S \rangle_\circ \\ &= \delta^{-1} d(d-1) \sum_{i=1}^n \langle L_J[\mathbf{h}^{(d-2)}], \langle \mathbf{v}_i, \mathbf{z} \rangle^{2d_i} r_{S,i}(\mathbf{V}^\top \mathbf{z}) \rangle_\circ + \langle \mathbf{h}_0^{(d)}, P_{\mathcal{H}}[(\mathbf{V}^\top \mathbf{z})^S] \rangle_\circ \\ &= \delta^{-1} \left\langle \mathbf{h}^{(d-2)}, \sum_{i=1}^n \langle \mathbf{v}_i, \mathbf{z} \rangle^{2d_i-2} r_{S,i}(\mathbf{V}^\top \mathbf{z}) \right\rangle_\circ + \langle \mathbf{h}_0^{(d)}, P_{\mathcal{H}}[(\mathbf{V}^\top \mathbf{z})^S] \rangle_\circ. \end{aligned} \quad (8.49)$$

We note that $\mathbf{h}^{(d-2)}$ and $\mathbf{h}_0^{(d)}$ are independent, and $\mathbf{h}_0^{(d)}$ is isotropic with variance σ_d^2 . Therefore, we may finally substitute into (8.44), obtaining

$$\begin{aligned} &\tilde{\mathbb{E}}(\mathbf{x}^S, \mathbf{x}^T) \\ &= \delta^d \cdot \mathbb{E} \left[\langle \mathbf{h}^{(d)}, (\mathbf{V}^\top \mathbf{z})^S \rangle_\circ \langle \mathbf{h}^{(d)}, (\mathbf{V}^\top \mathbf{z})^T \rangle_\circ \right] \\ &= \delta^{d-2} \cdot \mathbb{E} \left[\left\langle \mathbf{h}^{(d-2)}, \sum_{i=1}^n \langle \mathbf{v}_i, \mathbf{z} \rangle^{2d_i-2} r_{S,i}(\mathbf{V}^\top \mathbf{z}) \right\rangle_\circ \left\langle \mathbf{h}^{(d-2)}, \sum_{i=1}^n \langle \mathbf{v}_i, \mathbf{z} \rangle^{2d_i-2} r_{T,i}(\mathbf{V}^\top \mathbf{z}) \right\rangle_\circ \right] \\ &\quad + \delta^d \cdot \mathbb{E} \left[\langle \mathbf{h}_0^{(d)}, P_{\mathcal{H}}[(\mathbf{V}^\top \mathbf{z})^S] \rangle_\circ \langle \mathbf{h}_0^{(d)}, P_{\mathcal{H}}[(\mathbf{V}^\top \mathbf{z})^T] \rangle_\circ \right]. \end{aligned}$$

While this expression appears complicated, each term simplifies substantially: the first term is an evaluation of $\tilde{\mathbb{E}}$ at degree $2d - 4$, per (8.44), while the second term, by the isotropy of $h_0^{(d)}$, is an apolar inner product:

$$= \tilde{\mathbb{E}} \left(\sum_{i=1}^n x_i^{2d_i-2} r_{S,i}(\mathbf{x}), \sum_{i=1}^n x_i^{2d_i-2} r_{T,i}(\mathbf{x}) \right) + \sigma_d^2 \delta^d \cdot \left\langle P_{\mathcal{H}}[(\mathbf{V}^\top \mathbf{z})^S], P_{\mathcal{H}}[(\mathbf{V}^\top \mathbf{z})^T] \right\rangle_\circ.$$

Lastly, we note that the factors $x_i^{2d_i-2}$ are irrelevant in the evaluation of $\tilde{\mathbb{E}}$ (by the pseudoexpectation ideal property from Definition 6.1.2 at degree $2d - 4$), so we obtain

$$= \tilde{\mathbb{E}} \left(\sum_{i=1}^n r_{S,i}(\mathbf{x}), \sum_{i=1}^n r_{T,i}(\mathbf{x}) \right) + \sigma_d^2 \delta^d \cdot \left\langle P_{\mathcal{H}}[(\mathbf{V}^\top \mathbf{z})^S], P_{\mathcal{H}}[(\mathbf{V}^\top \mathbf{z})^T] \right\rangle_\circ, \quad (8.50)$$

which is the result in the statement. \square

Thus our prediction for the degree $2d$ pseudoexpectation values decomposes according to the ideal-harmonic decomposition of the input; the ideal term depends only on the pseudoexpectation values of strictly lower degree, while the harmonic term is a new contribution that is, in a suitable spectral sense, orthogonal to the ideal term. In this way, one may think of building up the spectral structure of the pseudomoment matrices of $\tilde{\mathbb{E}}$ by repeatedly “raising” the pseudomoment matrix two degrees lower into a higher-dimensional domain, and then adding a new component orthogonal to the existing one.

Remark 8.3.12 (Multiharmonic basis and block diagonalization). *We also mention a different way to view this result that will be more directly useful in our proofs later in Chapter 10. Defining $h_S^\perp(\mathbf{x}) := \mathbf{x}^S - r_S^\perp(\mathbf{x})$, note that we have, in the setting of Lemma 8.3.11 where Assumption 8.1.1 is exactly satisfied, $\tilde{\mathbb{E}}(h_S^\perp(\mathbf{x}), r_S^\perp(\mathbf{x})) = \tilde{\mathbb{E}}(h_S(\mathbf{x}), r_S(\mathbf{x})) = 0$ since the ideal and harmonic subspaces are apolar. Thus, we also have*

$$\tilde{\mathbb{E}}(h_S^\perp(\mathbf{x}), h_T^\perp(\mathbf{x})) = \sigma_d^2 \delta^d \cdot \langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle_\circ. \quad (8.51)$$

The $h_S^\perp(\mathbf{x})$ are a basis modulo the ideal generated by the constraint polynomials $x_i^2 - 1$, which we call the multiharmonic basis. Since the inner product on the right-hand side is zero unless $|S| = |T|$, this basis achieves a block diagonalization of the pseudomoment matrix. This idea turns out to be easier to use to give a proof of positivity of $\tilde{\mathbb{E}}$ than the full recursion of (8.41).

Finally, we also give a concrete closed form of the above recursion. This is in terms of a Fischer-like decomposition analogous to that given for harmonic polynomials in Example 8.3.10 above.

Corollary 8.3.13 (Pseudoexpectation closed form). *Suppose that the conditions of Assumption 8.1.1 hold exactly. Let $h_{S,T} \in V_{\mathcal{H}}^{(|S|-2|T|)}$ be such that*

$$(\mathbf{V}^\top \mathbf{z})^S = \sum_{k=0}^{\lfloor |S|/2 \rfloor} \sum_{T \in \mathcal{M}_k(\{n\})} ((\mathbf{V}^\top \mathbf{z})^T)^2 h_{S,T}(\mathbf{V}^\top \mathbf{z}), \quad (8.52)$$

which may be obtained by repeatedly expanding the ideal-harmonic decompositions of the $r_{S,i}$ from Lemma 8.3.11. Define

$$h_{S,k}(\mathbf{x}) := \sum_{T \in \mathcal{M}_{(|S|-k)/2}(\{n\})} h_{S,T}(\mathbf{x}) \in \mathbb{R}[x_1, \dots, x_n]_k^{\text{hom}} \quad (8.53)$$

if $k \leq |S|$ and k and $|S|$ are of equal parity, and $p_{S,k} = 0$ otherwise. Then, the right-hand side of (8.35) is

$$\tilde{\mathbb{E}}(\mathbf{x}^S, \mathbf{x}^T) = \sum_{d=0}^{|S| \wedge |T|} \sigma_d^2 \delta^d \cdot \langle h_{S,d}(\mathbf{V}^\top \mathbf{z}), h_{T,d}(\mathbf{V}^\top \mathbf{z}) \rangle_\circ. \quad (8.54)$$

We will not use this representation much, as it seems difficult to understand what transpires in the large summations of polynomials made to form $h_{S,d}$, but we present it to emphasize that, at least in principle, our construction proposes an explicit Gram matrix description of the pseudomoment matrix of $\tilde{\mathbb{E}}$. In one situation in the next chapter we will also be able to make use of this directly, since we will be studying a very structured deterministic case where the $h_{S,d}$ can be adequately understood.

9 | GRIGORIEV-LAURENT LOWER BOUND

While we proposed in the previous chapter a construction of pseudomoments with pleasant spectral properties and connections to the geometry of polynomials, the major mystery of why what we defined, *a priori* still only a bilinear pseudoexpectation, should factor through multiplication and define a true pseudoexpectation, still remains. Before we show a generic further derivation that will justify this, however, here we digress to consider one very special case—the simplest non-trivial one of the Gram matrices of ETFs we studied in Chapter 7—where this prediction in fact gives an exact description of a previously studied pseudomoment matrix.

That pseudomoment matrix was used by Laurent [Lau03b] to show the result we cited earlier in Chapter 6, that $\mathcal{E}_{2d}^n \cong \mathbb{C}^n$ if $2d < n$. Essentially the same construction appeared in a different guise less related to our results in a result of Grigoriev [Gri01a]. We mention Grigoriev’s result to respect the chronology of these results, but we will focus on Laurent’s formulation as it is much closer to our setting. In this chapter, we will give a new proof of this lower bound—though similar ideas to ours are already implicit in another alternative proof—and use some further reasoning suggested by our spectral extensions to establish a conjecture of Laurent’s on the eigenvalues of the associated pseudomoment matrix.

SUMMARY AND REFERENCES This chapter is based on a forthcoming note, a joint work with Jess Banks and Cristopher Moore. The following is a summary of our main results in this

chapter.

1. (Section 9.3) An alternative representation-theoretic proof of the positivity of the pseudomoments considered by Laurent, which we state in Theorem 9.1.1.
2. (Theorem 9.4.1) A recursive structure obeyed by the eigenvalues of this pseudomoment matrix, conjectured by Laurent in [Lau03b].

We also mention the following representation-theoretic result, perhaps of independent interest.

3. (Proposition 9.2.11) An intrinsic description of the irreducible representations of S_n corresponding to Young diagrams with two rows as certain spaces of multiharmonic polynomials associated to an equilateral simplex.

This is essentially equivalent to the well-known description as a Specht module, but our description gives a rather more natural definition not requiring an explicit basis.

PRIOR WORK At least three other proofs (besides the original) of the result of Grigoriev and Laurent have appeared in the literature [KLM16, BGP16, Pot17]. However, Laurent’s conjecture on the pseudomoment eigenvalues has remained, to the best of our knowledge, unverified since its observation in [Lau03b]. The general, if somewhat vague, question of finding “synthetic” descriptions of the irreducible representations of the symmetric group (in particular, basis-free descriptions) remains interesting. This is mentioned and discussed explicitly in, e.g., [Dia88]: “What one wants is a set of objects on which S_n acts that are comprehensible...as far as I know, a ‘concrete’ determination of the representations of S_n is an open problem” (p. 136). There is a rich literature on these matters applying much more sophisticated algebraic machinery than we do here (see the other discussion in Chapter 7B of the above reference as well as the more recent Okounkov-Vershik approach [OV96]), but we are not aware of any descriptions of the same flavor as we propose.

9.1 PSEUDOMOMENT CONSTRUCTION

The pseudomoment construction we will study is as follows.

Theorem 9.1.1 (Theorem 6 of [Lau03b]). *Let $n \geq 1$ be odd. Let $\tilde{\mathbb{E}} : \mathbb{R}[x_1, \dots, x_n]_{\leq n-1} \rightarrow \mathbb{R}$ have values on multilinear monomials given by*

$$\tilde{\mathbb{E}} \left[\prod_{i \in S} x_i \right] = \mathbb{1}_{\{|S| \text{ even}\}} \cdot (-1)^{|S|/2} \prod_{i=0}^{|S|/2-1} \frac{2i+1}{n-2i-1} =: \alpha_{|S|}, \quad (9.1)$$

and extend to a linear operator satisfying $\tilde{\mathbb{E}}[x_i^2 p(x)] = \tilde{\mathbb{E}}[p(x)]$ whenever $\deg(p) \leq n-3$. Then, $\tilde{\mathbb{E}}$ is a degree $(n-1)$ pseudoexpectation, and satisfies $\tilde{\mathbb{E}}[(\sum_{i=1}^n x_i)^2] = 0$.

Since n is odd, for any $x \in \mathcal{C}^n$ we have $(\sum_{i=1}^n x_i)^2 \geq 1$ by parity considerations. Therefore, the result shows that SOS requires degree $n+1$ to certify this simple-looking inequality, and in particular it follows that $\mathcal{E}_{n-1}^n \not\supseteq \mathcal{C}^n$, showing which was Laurent's purpose in proving this result. We also note that the degree 2 pseudomoment matrix is

$$\tilde{\mathbb{E}}[xx^\top] = \begin{bmatrix} 1 & -\frac{1}{n-1} & \cdots & -\frac{1}{n-1} \\ -\frac{1}{n-1} & 1 & \cdots & -\frac{1}{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ -\frac{1}{n-1} & -\frac{1}{n-1} & \cdots & 1 \end{bmatrix} = \frac{n}{n-1} \mathbf{I}_n - \frac{1}{n-1} \mathbf{1}_n \mathbf{1}_n^\top \in \mathbb{R}_{\text{sym}}^{n \times n}, \quad (9.2)$$

which is the Gram matrix of the *simplex ETF* of n unit vectors in \mathbb{R}^{n-1} pointing to the vertices of an equilateral simplex. In particular, our Theorem 7.4.5 immediately gives degree 4 pseudomoments achieving Laurent's result. Here, however, we will go much further and give an alternate proof of the full statement.

It is not difficult to arrive at the values of the α_k defining the pseudomoments: we assume by symmetry that $\tilde{\mathbb{E}}[x^S]$ depends only on $|S|$; by symmetrizing $\tilde{\mathbb{E}}'[p(x)] = \frac{1}{2}(\tilde{\mathbb{E}}[p(x)] +$

$\tilde{\mathbb{E}}[p(-x)]$ we may assume that $\tilde{\mathbb{E}}[x^S] = 0$ whenever $|S|$ is odd; and we assume that not only does $\tilde{\mathbb{E}}[(\sum_{i=1}^n x_i)^2] = 0$, but moreover $\tilde{\mathbb{E}}[(\sum_{i=1}^n x_i)p(x)] = 0$ whenever $\deg(p) \leq n - 2$ (sometimes called $\tilde{\mathbb{E}}$'s "strongly satisfying" the constraint $\sum_{i=1}^n x_i = 0$). Then, we must have $0 = \tilde{\mathbb{E}}[(\sum_{i=1}^n x_i)x^S] = |S|\alpha_{|S|-1} + (n - |S|)\alpha_{|S|+1}$, and starting with $\tilde{\mathbb{E}}[1] = 1$ the values in (9.1) follow recursively. We also note that it is impossible to continue this construction past $|S| = (n - 1)/2$ while retaining these properties, as solving the recursion will call for a division by zero.

The content of the theorem is the positivity of $\tilde{\mathbb{E}}$; clearly $\tilde{\mathbb{E}}$ satisfies the remaining conditions of Definition 6.1.2 by construction. Let us spell this out in terms of a concrete pseudomoment matrix, as we will return to the spectral properties of this matrix below. We set

$$d_{\max} = d_{\max}(n) := \left\lfloor \frac{n}{2} \right\rfloor. \quad (9.3)$$

(This is denoted " k " in [Lau03b].) Let $\mathbf{Y}^{(n)} \in \mathbb{R}^{\binom{[n]}{\leq d_{\max}} \times \binom{[n]}{\leq d_{\max}}}$ have entries

$$Y_{S,T}^{(n)} = \tilde{\mathbb{E}}[\mathbf{x}^S \mathbf{x}^T]. \quad (9.4)$$

Then, the content of the theorem is that $\mathbf{Y}^{(n)} \geq 0$. Actually, $\mathbf{Y}^{(n)}$ decomposes as the direct sum of two principal submatrices, those indexed by $\binom{[n]}{d}$ with d even and odd respectively, which Laurent considered separately, but it will be more natural in our calculations to avoid this decomposition.

Adjusting for this minor change, Laurent's proof may be seen as identifying a $\binom{n}{\leq d_{\max}-1}$ -dimensional kernel of $\mathbf{Y}^{(n)}$, and proving that the principal submatrix $\mathbf{Z}^{(n)}$ of $\mathbf{Y}^{(n)}$ indexed by $\binom{[n-1]}{d_{\max}} \cup \binom{[n-1]}{d_{\max}-1}$, which has total dimension $\binom{n}{d_{\max}}$, is positive definite. While identifying the kernel is straightforward, for the second part Laurent uses that each block of $\mathbf{Z}^{(n)}$ belongs to the Johnson association scheme, and applies formulae for the eigenvalues of the matrices spanning the Johnson scheme's Bose-Mesner algebra. In concrete terms, this

expresses the eigenvalues of $Z^{(n)}$ as certain combinatorial sums involving binomial coefficients and having alternating signs. To establish positivity, Laurent appeals to general identities for transforming hypergeometric series [PWZ96], which yield different expressions for the eigenvalues of $Z^{(n)}$ as sums of positive terms.

Both Laurent’s result and proof are similar to those of [Gri01a] that appeared earlier and concerned the same statement over $x \in \{0, 1\}^n$ in the context of the knapsack problem. At least three other, conceptually different, proofs of Laurent’s result (that we know of) have since appeared. First, [KLM16] showed that, for highly symmetric problems over subsets of the hypercube, the positivity of the “natural” pseudoexpectation constructed from symmetry considerations reduces to a small number of univariate polynomial inequalities. In the case of Laurent’s result, these inequalities simplify algebraically and yield a proof, while in other cases this machinery calls for analytic arguments. Second, [BGP16] produced an elegant proof of a stronger result, showing that the function $(\sum_{i=1}^n x_i)^2 - 1$ is not even a sum of squares of *rational* functions of degree at most d_{\max} . Their proof works in the dual setting, describing the decomposition of the space of functions on the hypercube into irreducible representations of the symmetric group and considering how a hypothetical sum of squares expression decomposes into associated components in these subspaces. We will present this decomposition below and apply it at the beginning of our computations. Lastly, [Pot17] showed that the result, in Grigoriev’s form concerning knapsack, follows from another general representation-theoretic reduction of positivity conditions to a lower-dimensional space of polynomials.

Laurent also made several further empirical observations about $Y^{(n)}$ and its “quite remarkable structural properties” in the appendix of [Lau03b]. Most notably, as n increases, the spectrum of $Y^{(n)}$ appears to “grow” in a recursive fashion, with the eigenvalues of $Y^{(n+2)}$ equaling those of $Y^{(n)}$ multiplied by $\frac{n+2}{n+1}$, along with a new largest eigenvalue. Unfortunately, the proof outlined above does not make use of this elegant structure and does not

give any indication of why it should hold. Our reasoning below will prove this observation and show how it fits into a representation-theoretic perspective on the Grigoriev-Laurent pseudomoments.¹

9.2 REPRESENTATION THEORY OF THE SYMMETRIC GROUP

We will use standard tools such as Schur’s lemma, character orthogonality, and characterizations of and formulae for characters of irreducible representations (henceforth *irreps*) of the symmetric group. See, e.g., the standard reference [FH04] or [Dia88] for more an explicit combinatorial perspective.

We briefly recall some standard aspects of this theory. A partition $\tau = (\tau_1, \dots, \tau_m)$ of $n \in \mathbb{N}_+$ is an ordered sequence $\tau_1 \geq \dots \geq \tau_m > 0$ such that $\sum_{i=1}^m \tau_i = n$. We write $\text{Part}(n)$ for the set of partitions of n (note that previously we wrote $\text{Part}(S)$ for the partitions of a set S , while $\text{Part}(n)$ here denotes the partitions of the *number* n). The associated *Young diagram* is an array of left-aligned boxes, where the k th row contains τ_k boxes. A *Young tableau* of *shape* τ is an assignment of the numbers from $[n]$ to these boxes (possibly with repetitions). A tableau is *standard* if the rows and columns are strictly increasing (left to right and top to bottom, respectively), and *semistandard* if the rows are non-decreasing but the columns are strictly increasing. We give a concrete and perhaps old-fashioned treatment of the topics below, since these constructions in terms of polynomials will be directly useful for us.

Definition 9.2.1 (Combinatorial representation). *The combinatorial representation associ-*

¹Laurent also observed that it appears plausible to prove $\mathbf{Y}^{(n)} \geq 0$ by repeatedly taking Schur complements with respect to blocks indexed by subsets of fixed size $\binom{[n]}{d}$, in the order $d = 0, 1, \dots, d_{\max}$. This would give a perhaps more conceptually-satisfying proof than the original one relying on eigenvalue interlacing—which offers no direct insight into the spectrum of $\mathbf{Y}^{(n)}$ itself—but appears too technical to carry out by hand. Our reasoning showing that the Grigoriev-Laurent pseudomoments are a special case of spectral extension will also implicitly verify that such an approach is tenable.

ated to $\tau \in \text{Part}(n)$, which we denote U^τ , is the module of polynomials in $\mathbb{R}[x_1, \dots, x_n]$ where, for each $k \in [m]$, exactly τ_k of the x_i appear raised to the power $(k - 1)$ in each monomial (thus these polynomials are homogeneous of degree $\sum_{k=1}^m (k - 1)\tau_k$).

Definition 9.2.2 (Specht module). *The Specht module associated to $\tau \in \text{Part}(n)$, which we denote W^τ , is the module of polynomials in $\mathbb{R}[x_1, \dots, x_n]$ spanned by, over all standard tableaux T of shape τ ,*

$$\prod_C \prod_{\substack{i, j \in C \\ i < j}} (x_i - x_j), \quad (9.5)$$

where the product is over columns C of T . We write χ_τ for the character of W^τ , and identify $\chi_{(n,0)} := \chi_{(n)}$ for the sake of convenience.

The key and classical fact concerning the Specht modules is that, over $\tau \in \text{Part}(n)$, they are all non-isomorphic and enumerate all irreps of S_n . The main extra fact we will use is the following, showing how to decompose a combinatorial representation in these irreps.

Proposition 9.2.3 (Young's rule). *For $\tau, \mu \in \text{Part}(n)$, the multiplicity of W^μ in U^τ is the number of semistandard Young tableaux of shape μ in which k occurs τ_k times for each $k = 1, \dots, m$.*

Finally, we will use the following decomposition of the representation consisting of polynomials over the hypercube $\{\pm 1\}^n$ given in a recent work. We correct a small typo present in the published version in the limits of the second direct sum below.

Proposition 9.2.4 (Theorem 3.2 of [BGP16]). *Let $\mathbb{R}[\{\pm 1\}^n] := \mathbb{R}[x_1, \dots, x_n]/\mathcal{I}$, where \mathcal{I} is the ideal generated by $\{x_i^2 - 1\}_{i=1}^n$. Then,*

$$\mathbb{R}[\{\pm 1\}^n] = \bigoplus_{d=0}^{d_{\max}} \bigoplus_{k=0}^{n-2d} \left(\sum_{i=1}^n x_i \right)^k W^{(n-d,d)}. \quad (9.6)$$

(Note that this is not merely a statement of the isomorphism type of the irreps occurring in $\mathbb{R}[\{\pm 1\}^n]$, but an actual direct sum decomposition of the space of polynomials, where the $W^{(n-d,d)}$ are meant as specific subspaces of polynomials, per Definition 9.2.2.)

9.2.1 COMBINATORIAL INTERPRETATIONS OF CHARACTERS

As we will be computing extensively with $\chi_{(n-d,d)}$ below, it will be useful to establish a concrete combinatorial description of the values of these characters.

Definition 9.2.5. For each $0 \leq d \leq n$, let

$$c_d(\pi) := \# \left\{ A \in \binom{[n]}{d} : \pi(A) = A \right\}. \quad (9.7)$$

Proposition 9.2.6. For all $1 \leq d \leq n$, $\chi_{(n-d,d)} = c_d - c_{d-1}$, and $\chi_{(n)} = c_0 = 1$.

We give two proofs, one using the Frobenius formula for irrep characters and another using combinatorial representations.

Proof 1. The Frobenius formula implies that, for π having cycles C_1, \dots, C_k ,

$$\chi_{(n-d,d)}(\pi) = [x^d] \left\{ (1-x) \prod_{i=1}^k (1+x^{|C_i|}) \right\}. \quad (9.8)$$

Since a subset fixed by π is a disjoint union of cycles, the product term is the generating function of the numbers of fixed subsets of all sizes:

$$\prod_{i=1}^k (1+x^{|C_i|}) = \sum_{d=0}^n c_d(\pi) x^d. \quad (9.9)$$

The result follows since multiplication by $(1-x)$ makes the coefficients precisely the claimed differences. □

Proof 2. The combinatorial representation $U^{(n-d,d)}$ is the subspace of multilinear polynomials in $\mathbb{R}[x_1, \dots, x_n]_d^{\text{hom}}$. By Young's rule, $U^{(n-d,d)} = \bigoplus_{i=0}^d W^{(n-i,i)}$. On the other hand, clearly the character of $U^{(n-d,d)}$ is c_d . Thus, $\sum_{i=0}^d \chi_{(n-d,d)} = c_d$, and the result follows by inverting this relation. \square

9.2.2 COMBINATORIAL CLASS FUNCTIONS

We will need to compute inner products of various functions on conjugacy classes of S_n (henceforth *class functions*) with $\chi_{(n-d,d)}$. We compute several such inner products in advance below.

Definition 9.2.7. For $\pi \in S_n$, $0 \leq a, b \leq n$, and $0 \leq k, \ell \leq a \wedge b$, we define

$$f_{a,k}(\pi) := \# \left\{ A \in \binom{[n]}{a} : |\pi(A) \cap A| = k \right\}, \quad (9.10)$$

$$g_{a,b,k,\ell}(\pi) := \# \left\{ A \in \binom{[n]}{a}, B \in \binom{[n]}{b} : |A \cap B| = k, |\pi(A) \cap B| = \ell \right\}. \quad (9.11)$$

We will ultimately be interested in inner products with the $g_{a,b,k,\ell}$, but the following shows that these reduce to linear combinations of the $f_{a,k}$.

Proposition 9.2.8. For all $0 \leq k, \ell \leq a \wedge b$,

$$g_{b,a,k,\ell} = g_{a,b,k,\ell} = g_{a,b,\ell,k} = \sum_{j=0}^a \left(\sum_{i=0}^j \binom{j}{i} \binom{a-j}{k-i} \binom{a-j}{\ell-i} \binom{n-2a+j}{b-k-\ell+i} \right) f_{a,j} \quad (9.12)$$

Proof. The first equality holds since $|\pi(A) \cap B| = |A \cap \pi^{-1}(B)|$, and so since inversion does not change the conjugacy class of π , we have $g_{a,b,k,\ell}(\pi) = g_{b,a,k,\ell}(\pi^{-1}) = g_{b,a,k,\ell}(\pi)$.

Suppose $A \in \binom{[n]}{a}$ with $|A \cap \pi(A)| = j$. Then, $B \in \binom{[n]}{b}$ with $|A \cap B| = k$ and $|\pi(A) \cap B| = \ell$ consists of some $0 \leq i \leq j$ elements of $A \cap \pi(A)$, $k - i$ elements of $A \setminus \pi(A)$, $\ell - i$ elements

of $\pi(A) \setminus A$, and $b - i - (k - i) - (\ell - i) = b - k - \ell + i$ elements of $[n] \setminus A \setminus \pi(A)$. Thus,

$$\begin{aligned}
& \mathcal{G}_{a,b,k,\ell}(\pi) \\
&= \sum_{A \in \binom{[n]}{a}} \# \left\{ B \in \binom{[n]}{b} : |A \cap B| = k, |\pi(A) \cap B| = \ell \right\} \\
&= \sum_{A \in \binom{[n]}{a}} \sum_{i=0}^{|A \cap \pi(A)|} \binom{|A \cap \pi(A)|}{i} \binom{a - |A \cap \pi(A)|}{k - i} \binom{a - |A \cap \pi(A)|}{\ell - i} \binom{n - 2a + |A \cap \pi(A)|}{b - k - \ell + i} \\
&= \sum_{j=0}^a \# \left\{ A \in \binom{[n]}{a} : |A \cap \pi(A)| = j \right\} \sum_{i=0}^j \binom{j}{i} \binom{a - j}{k - i} \binom{a - j}{\ell - i} \binom{n - 2a + j}{b - k - \ell + i}, \tag{9.13}
\end{aligned}$$

and the remaining cardinality is by definition $f_{a,j}(\pi)$. \square

The following is our key combinatorial lemma, computing the inner product of $\chi_{(n-d,d)}$ with the $\mathcal{G}_{a,b,k,\ell}$ so long as one of a and b is at most d .

Lemma 9.2.9. *For all $0 \leq a \leq d \leq n/2$, $a \leq b \leq n$, and $0 \leq k, \ell \leq a \wedge b$,*

$$\frac{1}{n!} \sum_{\pi \in S_n} \chi_{(n-d,d)}(\pi) \mathcal{G}_{a,b,k,\ell}(\pi) = \begin{cases} 0 & \text{if } a < d, \\ (-1)^{k+\ell} \binom{d}{k} \binom{d}{\ell} \binom{n-2d}{b-d} & \text{if } a = d. \end{cases} \tag{9.14}$$

Proof. We first compute the inner products with the $f_{a,k}$. To this end, we introduce the functions

$$F_{a,j} := \sum_{k=j}^a \binom{k}{j} f_{a,k}. \tag{9.15}$$

Then, we have

$$\begin{aligned}
F_{a,j}(\pi) &= \sum_{A \in \binom{[n]}{a}} \binom{|A \cap \pi(A)|}{j} \\
&= \sum_{A \in \binom{[n]}{a}} \sum_{C \in \binom{A}{j}} \mathbf{1}\{\pi(C) \subseteq A\} \\
&= \sum_{C \in \binom{[n]}{j}} \sum_{B \in \binom{[n] \setminus C}{a-j}} \mathbf{1}\{\pi(C) \subseteq C \cup B\} \\
&= \sum_{C \in \binom{[n]}{j}} \binom{n - 2j + |\pi(C) \cap C|}{a - 2j + |\pi(C) \cap C|} \\
&= \sum_{i=0}^j \binom{n - 2j + i}{a - 2j + i} f_{j,i}(\pi). \tag{9.16}
\end{aligned}$$

On the other hand, we may invert the relation (9.15) (this “inversion of Pascal’s triangle” follows from the binomial coefficients giving the coefficients of the polynomial transformation $p(x) \mapsto p(x + 1)$, whereby the inverse gives the coefficients of the transformation $p(x) \mapsto p(x - 1)$; it is also sometimes called the *Euler transform*) to obtain the closed recursion

$$f_{a,k} = \sum_{j=k}^a (-1)^{j+k} \binom{j}{k} F_{a,j} = \sum_{j=k}^a (-1)^{j+k} \binom{j}{k} \sum_{i=0}^j \binom{n - 2j + i}{a - 2j + i} f_{j,i}. \tag{9.17}$$

In particular, the only non-zero term with $j = a$ is $(-1)^{a+k} \binom{a}{k} f_{a,a}$. We know that

$$f_{a,a} = c_a = \sum_{d=0}^a \chi_{(n-d,d)}. \tag{9.18}$$

Thus, by induction it follows that, in the character expansion of $f_{a,k}$, $\chi_{(n-d,d)}$ appears only if $a \geq d$, and when $a = d$ it appears with coefficient $(-1)^{d+k} \binom{d}{k}$. Thus we have

$$\frac{1}{n!} \sum_{\pi \in S_n} \chi_{(n-d,d)}(\pi) f_{a,k}(\pi) = \begin{cases} 0 & \text{if } a < d, \\ (-1)^{d+k} \binom{d}{k} & \text{if } a = d. \end{cases} \tag{9.19}$$

The first case of our claim, with $a < d$, now follows immediately from Proposition 9.2.8. For the second case, with $a = d$, we proceed by induction on n . First, making a general manipulation, again by Proposition 9.2.8 we have

$$\begin{aligned} & \frac{1}{n!} \sum_{\pi \in S_n} \chi^{(n-d,d)}(\pi) \mathcal{G}_{a,b,k,\ell}(\pi) \\ &= \sum_{j=0}^d \sum_{i=0}^j \binom{j}{i} \binom{d-j}{k-i} \binom{d-j}{\ell-i} \binom{n-2d+j}{b-k-\ell+i} \frac{1}{n!} \sum_{\pi \in S_n} \chi^{(n-d,d)}(\pi) f_{d,j}(\pi) \\ &= \sum_{j=0}^d (-1)^{d+j} \binom{d}{j} \sum_{i=0}^j \binom{j}{i} \binom{d-j}{k-i} \binom{d-j}{\ell-i} \binom{n-2d+j}{b-k-\ell+i} \end{aligned}$$

We start to treat the remaining sum using that $\sum_{j=0}^d (-1)^j \binom{d}{j} f(j)$ gives the d th finite difference of the function f . In particular, for f a polynomial of degree smaller than d , any such sum is zero. Furthermore, $\sum_{j=0}^d (-1)^j \binom{d}{j} j^d = (-1)^d d!$. Therefore, we may continue, always applying the differencing Δ transformation to functions of the variable j ,

$$\begin{aligned} &= \sum_{i=0}^d \sum_{w+x+y+z=d} \binom{d}{w, x, y, z} \\ & \quad \frac{\Delta^w j^i \cdot \Delta^x (d-j)^{k-i} \cdot \Delta^y (d-j)^{\ell-i} \cdot \Delta^z (n-2d+j)^{b-k-\ell+i} \Big|_{j=0}}{i!(k-i)!(\ell-i)!(b-k-\ell+i)!} \end{aligned} \quad (9.20)$$

Here, in all cases the first factor, $\Delta^w j^i \Big|_{j=0}$, will only be nonzero when $w = i$.

Let us now first specialize to the base case $n = 2d$. In this case, the last factor, $\Delta^z (n-2d+j)^{b-k-\ell+i} \Big|_{j=0}$, will likewise only be nonzero when $b-k-\ell+w = z$. In that case, we must have $x+y = k+\ell-2w+d-b$. Since in all nonzero terms $x \leq k-w$ and $y \leq \ell-w$, and $d \leq b$, we will only have a nonzero result if $d = b$, $x = k-w$, and $y = \ell-w$. In this

case, we have

$$\begin{aligned} \frac{1}{n!} \sum_{\pi \in \mathcal{S}_n} \chi_{(n-d,d)}(\pi) \mathcal{G}_{d,d,k,\ell}(\pi) &= (-1)^{k+\ell} \sum_{w=0}^d \binom{d}{w, k-w, \ell-w, d-k-\ell+w} \\ &= (-1)^{k+\ell} \binom{d}{k} \binom{d}{\ell}, \end{aligned} \quad (9.21)$$

the final step following since the remaining sum counts the number of ways to choose a subset of size k and a subset of size ℓ from $[d]$, with w being the size of the intersection. Thus the result holds when $n = 2d$.

Suppose now that $n > 2d$ and the result holds for $n - 1$. Continuing from (9.20) above and completing the computation of the differences,

$$\begin{aligned} &\frac{1}{n!} \sum_{\pi \in \mathcal{S}_n} \chi_{(n-d,d)}(\pi) \mathcal{G}_{a,b,k,\ell}(\pi) \\ &= \sum_{w+x+y+z=d} (-1)^{x+y} \binom{d}{w, x, y, z} \frac{1}{(k-w)!(\ell-w)!(b-k-\ell+w)!} \\ &\quad (k-i)^x (d-w-x)^{k-i-x} (\ell-i)^y (d-w-x-y)^{\ell-w-y} (b-k-\ell+w)^z \\ &\quad (n-2d+w+x+y)^{\underline{b-k-\ell+w-z}} \\ &= \sum_{w+x+y+z=d} (-1)^{x+y} \binom{d}{w, x, y, z} \binom{d-w-x}{k-w-x} \binom{d-w-x-y}{\ell-w-y} \binom{n-2d+w+x+y}{b-k-\ell+w-z} \end{aligned}$$

Reindexing in terms of $x' := k - w - x, y' = \ell - w - y, z' = b - k - \ell + w - z$, which we note must be non-negative and satisfy $x' + y' + z' = b - d$, we find

$$\begin{aligned} &= \sum_{w=0}^d \sum_{x'+y'+z'=b-d} (-1)^{x+y} \binom{d}{w, k-w-x', \ell-w-y', d-k-\ell+w+x'+y'} \\ &\quad \binom{d-k+x'}{x'} \binom{d-k-\ell+w+x'+y'}{y'} \binom{n-2d+k+\ell-w-x'-y'}{z'}. \end{aligned} \quad (9.22)$$

We emphasize here first that b appears only in the summation bounds for the inner sum,

and second that we have rewritten to leave only one occurrence of z' , in the final factor.

We group the terms of the sum according to whether $z' = 0$ or $z' > 0$:

$$S_0(b, d, k, \ell) := \sum_{w=0}^d \sum_{x'+y'=b-d} (-1)^{x+y} \binom{d}{w, k-w-x', \ell-w-y', d-k-\ell+w+x'+y'} \binom{d-k+x'}{x'} \binom{d-k-\ell+w+x'+y'}{y'}, \quad (9.23)$$

$$S_1(n, b, d, k, \ell) := \sum_{w=0}^d \sum_{x'+y'+z'=b-d-1} (-1)^{x+y} \binom{d}{w, k-w-x', \ell-w-y', d-k-\ell+w+x'+y'} \binom{d-k+x'}{x'} \binom{d-k-\ell+w+x'+y'}{y'} \binom{n-2d+k+\ell-w-x-y}{z'+1}. \quad (9.24)$$

Then, the sum we are interested in, that given in (9.22), is $S(n, b, d, k, \ell) := S_0(b, d, k, \ell) + S_1(n, b, d, k, \ell)$. Now, applying the identity $\binom{m}{a} = \binom{m-1}{a} + \binom{m-1}{a-1}$ to the last factor involving z' in S_1 , we find that

$$S_1(n, b, d, k, \ell) = S_1(n-1, b, d, k, \ell) + S(n-1, b-1, d, k, \ell). \quad (9.25)$$

Thus we have

$$\begin{aligned} S(n, b, d, k, \ell) &= S_0(b, d, k, \ell) + S_1(n, b, d, k, \ell) \\ &= S_0(b, d, k, \ell) + S_1(n-1, b, d, k, \ell) + S(n-1, b-1, d, k, \ell) \quad (\text{by (9.25)}) \\ &= S(n-1, b, d, k, \ell) + S(n-1, b-1, d, k, \ell) \end{aligned}$$

and by the inductive hypothesis

$$\begin{aligned}
&= (-1)^{k+\ell} \binom{d}{k} \binom{d}{\ell} \left(\binom{n-2d-1}{b-d} + \binom{n-2d-1}{b-d-1} \right) \\
&= (-1)^{k+\ell} \binom{d}{k} \binom{d}{\ell} \binom{n-2d}{b-d},
\end{aligned} \tag{9.26}$$

completing the induction. \square

The following is a reformulation of this result, perhaps with a more intuitive combinatorial interpretation and which will be more directly useful in our calculations to come.

Corollary 9.2.10. *Let $0 \leq a, b \leq d$, $A \in \binom{[n]}{a}$, $B \in \binom{[n]}{b}$, and $0 \leq k \leq a \wedge b$. Then,*

$$\frac{1}{n!} \sum_{\substack{\pi \in S_n \\ |\pi(A) \cap B| = k}} \chi_{(n-d,d)}(\pi) = \begin{cases} 0, & \text{if } a \wedge b < d, \\ (-1)^{k+|A \cap B|} \frac{\binom{d}{k}}{\binom{d, d-|A \cap B|, n-2d+|A \cap B|}} & \text{if } a = b = d. \end{cases} \tag{9.27}$$

Proof. Let us write $\ell := |A \cap B|$. Then, using that $\chi_{(n-d,d)}$ is a class function, we may average over conjugations,

$$\begin{aligned}
\sum_{\substack{\pi \in S_n \\ |\pi(A) \cap B| = k}} \chi_{(n-d,d)}(\pi) &= \frac{1}{n!} \sum_{\sigma \in S_n} \sum_{\substack{\pi \in S_n \\ |\sigma^{-1} \pi \sigma(A) \cap B| = k}} \chi_{(n-d,d)}(\sigma^{-1} \pi \sigma) \\
&= \frac{1}{n!} \sum_{\sigma \in S_n} \sum_{\substack{\pi \in S_n \\ |\pi \sigma(A) \cap \sigma(B)| = k}} \chi_{(n-d,d)}(\pi) \\
&= \frac{1}{\binom{\ell, a-\ell, b-\ell, n-a-b+\ell}{n}} \sum_{\substack{A' \in \binom{[n]}{a} \\ B' \in \binom{[n]}{b} \\ |A' \cap B'| = \ell}} \sum_{\substack{\pi \in S_n \\ |\pi(A') \cap B'| = k}} \chi_{(n-d,d)}(\pi) \\
&= \frac{1}{\binom{\ell, a-\ell, b-\ell, n-a-b+\ell}{n}} \sum_{\pi \in S_n} \chi_{(n-d,d)}(\pi) \mathcal{G}_{a,b,k,\ell}(\pi),
\end{aligned} \tag{9.28}$$

and the result now follows from Lemma 9.2.9 upon simplifying. \square

We note also that the special case $a = b = k$ gives the summation of the character over all

π with a specified mapping $\pi(A) = B$.

9.2.3 THE SIMPLEX-HARMONIC REPRESENTATION

We now discuss how our definition of the harmonic subspace $V_{\mathcal{H}}^{(d)}$ relates to the representation theory of the symmetric group. Let us fix $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^{n-1}$ unit vectors pointing to the vertices of an equilateral simplex. These vectors are related to our pseudoexpectation by

$$\tilde{\mathbb{E}}[\mathbf{x}\mathbf{x}^\top] = \text{Gram}(\mathbf{v}_1, \dots, \mathbf{v}_n). \quad (9.29)$$

We have that S_n acts on \mathbb{R}^{n-1} by permuting the \mathbf{v}_i (since $\sum_{i=1}^n \mathbf{v}_i = 0$ this is well-defined); this is the irreducible “standard representation” of S_n . The symmetric powers of this irrep give actions of S_n on $\mathbb{R}[z_1, \dots, z_{n-1}]_d^{\text{hom}}$ by likewise permuting the $\langle \mathbf{v}_i, \mathbf{z} \rangle$, products of which form an overcomplete set of monomials. This contains the invariant subspace

$$V_{\mathcal{H}}^{(d)} := \left\{ \mathbf{p} \in \mathbb{R}[z_1, \dots, z_{n-1}]_d^{\text{hom}} : \langle \mathbf{v}_i, \boldsymbol{\partial} \rangle^2 \mathbf{p} = 0 \text{ for all } i \in [n] \right\}. \quad (9.30)$$

We call $V_{\mathcal{H}}^{(d)}$ the *simplex-harmonic* representation of S_n . The next result identifies the isomorphism type of this representation.

Proposition 9.2.11 (Isomorphism type). $V_{\mathcal{H}}^{(d)} \cong W^{(n-d,d)}$. *The map $\Psi : \mathbb{R}[x_1, \dots, x_n]_d^{\text{hom}} \rightarrow \mathbb{R}[z_1, \dots, z_{n-1}]_d^{\text{hom}}$ given by defining $\Psi(\mathbf{x}^S) = (\mathbf{V}^\top \mathbf{z})^S$ and extending by linearity is an isomorphism between $W^{(n-d,d)}$ and $V_{\mathcal{H}}^{(d)}$ when restricted to $W^{(n-d,d)}$.*

Proof. Let us abbreviate $W = W^{(n-d,d)}$ and $V = V_{\mathcal{H}}^{(d)}$. We first compute the dimensions of V and W and show that they are equal.

For W , by the hook length formula,

$$\begin{aligned}
\dim(W) &= \frac{n!}{d! \cdot (n-d+1) \cdots (n-2d+2) \cdot (n-2d)!} \\
&= \frac{n!(n-2d+1)}{d!(n-d+1)!} \\
&= \binom{n}{d} \cdot \frac{n-2d+1}{n-d+1} \\
&= \binom{n}{d} - \binom{n}{d-1}.
\end{aligned} \tag{9.31}$$

(The same also follows by evaluating $\chi_{(n-d,d)}$ on the identity using the formula from Proposition 9.2.6.)

For V , we note that V is isomorphic to the subspace of $\text{Sym}^d(\mathbb{R}^n)$ consisting of symmetric tensors that are zero at any position with a repeated index and have any one-dimensional slice summing to zero. The tensors satisfying the first constraint have dimension $\binom{n}{d}$, and there are $\binom{n}{d-1}$ one-dimensional slices. We verify that these slice constraints are linearly independent: they may be identified with the vectors $\mathbf{a}_S \in \mathbb{R}^{\binom{[n]}{d}}$ for $S \in \binom{[n]}{d-1}$ with entries $(\mathbf{a}_S)_T = \mathbb{1}\{S \subseteq T\}$. These vectors satisfy

$$\langle \mathbf{a}_S, \mathbf{a}_{S'} \rangle = \begin{cases} n-d+1 & \text{if } S = S', \\ 1 & \text{if } |S \cap S'| = d-2, \\ 0 & \text{otherwise.} \end{cases} \tag{9.32}$$

Therefore, their Gram matrix is equal to $(n-d+1)\mathbf{I}_{\binom{[n]}{d-1}} + \mathbf{A}$, where \mathbf{A} is the adjacency matrix of the Johnson graph $J(n, d-1)$. Its most negative eigenvalue is equal to $-\min(d-1, n-d+1) = -(d-1)$ (see, e.g., Section 1.2.2 of [BVM]), whereby the Gram matrix of the \mathbf{a}_S is positive definite. Thus the \mathbf{a}_S are linearly independent, and $\dim(V) = \binom{n}{d} - \binom{n}{d-1} = \dim(W)$.

Therefore, to show $V \cong W$ it suffices to show that one of V or W contains a copy of the other. We show that V contains a copy of W . Recall from Definition 9.2.2 that W is the

subspace of $\mathbb{R}[x_1, \dots, x_n]_d^{\text{hom}}$ spanned by

$$\prod_{a=1}^d (x_{i_a} - x_{j_a}) \text{ for } i_a, j_b \in [n] \text{ distinct; } i_1 < \dots < i_{n-d}; j_1 < \dots < j_d; i_a < j_a. \quad (9.33)$$

Note that $\ker(\Psi)$ is the ideal generated by $x_1 + \dots + x_n$, and therefore is an invariant subspace of the S_n action. Since W is also an invariant subspace, and is irreducible, if W intersected $\ker(\Psi)$ non-trivially then W would be contained in $\ker(\Psi)$, which is evidently not true (for instance, any of the basis elements in (9.33) do not map to zero). Thus Ψ is an isomorphism on W , so it suffices to show that $\Psi(W) \subseteq V$. Indeed, all polynomials of $\Psi(W)$ also belong to V : writing $M = \text{Gram}(v_1, \dots, v_n)$, for any basis element and $k \in [n]$,

$$\begin{aligned} & \langle v_k, \boldsymbol{\theta} \rangle^2 \prod_{a=1}^d (\langle v_{i_a}, \mathbf{z} \rangle - \langle v_{j_a}, \mathbf{z} \rangle) \\ &= \sum_{\{a,b\} \in \binom{[n]}{2}} (M_{k,i_a} - M_{k,j_a})(M_{k,i_b} - M_{k,j_b}) \prod_{c \in [d] \setminus \{a,b\}} (\langle v_{i_c}, \mathbf{z} \rangle - \langle v_{j_c}, \mathbf{z} \rangle), \end{aligned} \quad (9.34)$$

and since the i_a and j_a are all distinct while all off-diagonal entries of M are equal, one of the two initial factors in each term will be zero. Thus, $W \cong \Psi(W) \subseteq V$, and by counting dimensions $V \cong W$. \square

Proposition 9.2.12 (Multiplicity). $V_{\mathcal{H}}^{(d)}$ has multiplicity one in $\mathbb{R}[z_1, \dots, z_{n-1}]_d^{\text{hom}}$

Proof. We show the stronger statement that the multiplicity of $V_{\mathcal{H}}^{(d)}$ in $\mathbb{R}[x_1, \dots, x_n]_d^{\text{hom}}$, which contains a copy of $\mathbb{R}[z_1, \dots, z_{n-1}]_d^{\text{hom}}$ as the quotient by the ideal generated by $x_1 + \dots + x_n$. We use that $\mathbb{R}[x_1, \dots, x_n]_d^{\text{hom}}$ admits a decomposition into invariant subspaces \tilde{U}^τ over $\tau \in \text{Part}(d)$, $\mathbb{R}[x_1, \dots, x_n]_d^{\text{hom}} = \bigoplus_{\tau \in \text{Part}(d)} \tilde{U}^\tau$, where \tilde{U}^τ consists of polynomials whose monomials have their set of exponents equal to the numbers appearing in τ . Each \tilde{U}^τ with $\tau = (\tau_1, \dots, \tau_m)$ is isomorphic to the combinatorial representation $U^{(n-m, f_1, \dots, f_\ell)}$, where the f_i give, in descending order, the frequencies of numbers appearing among the τ_i . By Young's rule, among these representations, only $\tilde{U}^{(1, \dots, 1)}$ contains a copy of $V_{\mathcal{H}}^{(d)} \cong W^{(n-d, d)}$

(as for this to happen we must have $m = d$), and it contains exactly one copy of $V_{\mathcal{H}}^{(d)}$. \square

9.3 PROOF OF POSITIVITY

Proposition 9.2.4 gives us a means of showing that $\mathbf{Y}^{(n)} \succeq \mathbf{0}$, for which we will not need to reason with the simplex-harmonic representation. Viewing $\mathbf{Y}^{(n)}$ as operating on polynomials, since the ideal generated by $\sum_{i=1}^n x_i$ is in the kernel of $\mathbf{Y}^{(n)}$, the only possible eigenspaces with non-zero eigenvalue are $W^{(n,0)}, W^{(n-1,1)}, \dots, W^{(n-d_{\max}, d_{\max})}$. Thus we can choose a non-zero element of each isotypic component, $p_i \in \bigoplus_{k=0}^{n-2d} (\sum_{i=1}^n x_i)^k W^{(n-i,i)}$, and verify that $\tilde{\mathbb{E}}[p_i(\mathbf{x})^2] > 0$ for each i .

To identify such polynomials, we compute the isotypic projections of monomials.

Definition 9.3.1 (Isotypic projection). *For each $S \in \binom{[n]}{d}$, define $h_S \in \mathbb{R}[x_1, \dots, x_n]_d^{\text{hom}}$ by*

$$h_S(\mathbf{x}) = \frac{\binom{n}{d} - \binom{n}{d-1}}{n!} \sum_{\pi \in S_n} \chi_{(n-d,d)}(\pi) \mathbf{x}^{\pi(S)}, \quad (9.35)$$

It then follows that $h_S(\mathbf{x}) \in \bigoplus_{k=0}^{n-2d} (\sum_{i=1}^n x_i)^k W^{(n-|S|, |S|)}$, the isotypic component or direct sum of all irrep components isomorphic to $W^{(n-|S|, |S|)}$.

Proposition 9.3.2. *For any $S \in \binom{[n]}{d}$,*

$$\tilde{\mathbb{E}}[h_S(\mathbf{x})^2] = \frac{n-2d+1}{n-d+1} \prod_{i=0}^{d-1} \frac{n-2i}{n-2i-1} > 0. \quad (9.36)$$

Proof. First, since $h_S(\mathbf{x})$ is the sum of the projection of \mathbf{x}^S to one of the eigenspaces of $\tilde{\mathbb{E}}$ and an element of the kernel of $\tilde{\mathbb{E}}$, we have

$$\tilde{\mathbb{E}}[h_S(\mathbf{x})^2] = \tilde{\mathbb{E}}[h_S(\mathbf{x})\mathbf{x}^S]$$

and from here we may compute directly,

$$\begin{aligned}
&= \frac{\binom{n}{d} - \binom{n}{d-1}}{n!} \sum_{k=0}^d \alpha_{2d-2k} \sum_{\substack{\pi \in \mathcal{S}_n \\ |\pi(S) \cap S| = k}} \chi_{(n-d,d)}(\pi) \\
&= \frac{\binom{n}{d} - \binom{n}{d-1}}{\binom{n}{d}} \sum_{k=0}^d (-1)^{d-k} \binom{d}{k} \alpha_{2d-2k} \quad (\text{Corollary 9.2.10}) \\
&= \frac{n-2d+1}{n-d+1} \sum_{k=0}^d (-1)^k \binom{d}{k} \alpha_{2k}. \quad (9.37)
\end{aligned}$$

It remains to analyze the sum. We view such a sum as a d th order finite difference, in this case a forward finite difference of the sequence $f(k) = \alpha(2k)$ for $k = 0, \dots, d$. Let us write $\Delta^a f$ for the sequence that is the a th forward finite difference. We will show by induction that

$$\Delta^a f(k) = \alpha_{2k} \prod_{i=0}^{a-1} \frac{n-2i}{n-2k-2i-1}. \quad (9.38)$$

Clearly this holds for $a = 0$. If the result holds for $a - 1$, then we have

$$\begin{aligned}
\Delta^a f(k) &= \Delta^{a-1} f(k) - \Delta^{a-1} f(k+1) \\
&= \alpha_{2k} \prod_{i=0}^{a-2} \frac{n-2i}{n-2k-2i-1} - \alpha_{2k+2} \prod_{i=0}^{a-2} \frac{n-2i}{n-2k-2i-3} \\
&= \alpha_{2k} \prod_{i=0}^{a-2} \frac{n-2i}{n-2k-2i-1} + \alpha_{2k} \frac{2k+1}{n-2k-1} \prod_{i=0}^{a-2} \frac{n-2i}{n-2k-2i-3} \\
&= \alpha_{2k} \prod_{i=0}^{a-2} \frac{n-2i}{n-2k-2i-1} \left(1 + \frac{2k+1}{n-2k-1} \cdot \frac{n-2k-1}{n-2k-2a+1} \right) \\
&= \alpha_{2k} \prod_{i=0}^{a-2} \frac{n-2i}{n-2k-2i-1} \cdot \frac{n-2(a-1)}{n-2k-2(a-1)-1}, \quad (9.39)
\end{aligned}$$

completing the induction. Evaluating at $a = d$ then gives

$$\sum_{k=0}^d (-1)^k \binom{d}{k} \alpha_{2k} = \Delta^d f(0) = \prod_{i=0}^{d-1} \frac{n-2i}{n-2i-1}, \quad (9.40)$$

completing the proof. \square

It is then straightforward to check that, together with some representation-theoretic reasoning, this implies Laurent's result.

Proof of Theorem 9.1.1. Let $p(\mathbf{x}) \in \mathbb{R}[x_1, \dots, x_n]$. By Proposition 9.2.4, there exist $h_{d,k} \in W^{(n-d,d)}$ for $d \in \{0, \dots, d_{\max}\}$ and $k \in \{0, \dots, n - 2d + 1\}$ such that

$$p(\mathbf{x}) = \sum_{d=0}^{d_{\max}} \sum_{k=0}^{n-2d} \left(\sum_{i=1}^n x_i \right)^k h_{d,k}(\mathbf{x}). \quad (9.41)$$

Since $\tilde{\mathbb{E}}$ is zero on multiples of $\sum_{i=1}^n x_i$, its pseudomoment matrix $\mathbf{Y}^{(n)}$ acts as a scalar on each of the $W^{(n-d,d)}$ by Schur's lemma, and $h_{d,0}$ for different d have different degrees and thus orthogonal vectors of coefficients, we have

$$\tilde{\mathbb{E}}[p(\mathbf{x})^2] = \tilde{\mathbb{E}} \left[\left(\sum_{d=0}^{d_{\max}} h_{d,0}(\mathbf{x}) \right)^2 \right] = \sum_{d=0}^{d_{\max}} \tilde{\mathbb{E}}[h_{d,0}(\mathbf{x})^2] \geq 0 \quad (9.42)$$

by Proposition 9.3.2, completing the proof. \square

9.4 PSEUDOMOMENT SPECTRUM AND LAURENT'S CONJECTURE

We now would like to recover the actual eigenvalues of $\mathbf{Y}^{(n)}$. Specifically, in this section we will show the following.

Theorem 9.4.1. $\mathbf{Y}^{(n)}$ has $d_{\max} + 2$ distinct eigenvalues, $0 < \lambda_{n,d_{\max}} < \dots < \lambda_{n,1} < \lambda_{n,0}$. The multiplicity of the zero eigenvalue is $\binom{n}{\leq d_{\max}-1}$, while the $\lambda_{n,d}$ have the following multiplicities

and recursive description:

$$\lambda_{n,0} = \sum_{k=0}^{d_{\max}} \binom{n}{k} \alpha_k^2, \text{ with multiplicity } 1,$$

$$\lambda_{n,d} = \frac{n}{n-1} \lambda_{n-2,d-1} \text{ for } 1 \leq d \leq (n-1)/2 \text{ with multiplicity } \binom{n}{d} - \binom{n}{d-1}.$$

This includes a conjecture Laurent makes in the Appendix of [Lau03b], and also gives a formula for the “base case” of the recursion in terms of the coefficients α_k from the pseudoexpectation as well as the multiplicities of all eigenvalues.

It may seem that we are close to obtaining the eigenvalues: since the $W^{(n-d,d)}$ are the eigenspaces of $\tilde{\mathbb{E}}$, it suffices to just find any concrete polynomial $p \in W^{(n-d,d)}$ that it is convenient to compute with, whereupon we will have $\lambda_{n,d} = \tilde{\mathbb{E}}[p(x)^2] / \|p\|^2$, where the norm of a polynomial is the norm of the vector of coefficients (not the apolar norm). However, our computation above does not quite achieve this: crucially, $h_S(x)$ does *not* belong to $W^{(n-d,d)}$; rather, it equals the projection of x^S to *all* copies of this irrep in $\mathbb{R}[\{\pm 1\}^n]$, of which there are $n - 2d + 1$. Since those copies that are divisible by $\sum_{i=1}^n x_i$ are in the kernel of $\tilde{\mathbb{E}}$, we have actually computed $\tilde{\mathbb{E}}[h_S(x)^2] = \tilde{\mathbb{E}}[\hat{h}_S(x)^2]$ where $\hat{h}_S(x) \in W^{(n-d,d)}$ is the relevant component of x^S . However, not having an explicit description of $\hat{h}_S(x)$, we have no immediate way to compute $\|\hat{h}_S\|^2$.

Remark 9.4.2. *One possible approach to implement this direct strategy is to try to take $p \in W^{(n-d,d)}$ to be one of the basis polynomials given in Definition 9.2.2. However, computing the pseudoexpectation of the square of such a polynomial gives an unusual combinatorial sum to which the character-theoretic tools we have developed do not seem to apply.*

Instead, we will use the representation of $\tilde{\mathbb{E}}$ as a Gram matrix suggested by our general computations with spectral extensions of pseudomoments in the previous chapter, specifically Lemma 8.3.11 and Corollary 8.3.13. We first verify that $\tilde{\mathbb{E}}$ indeed admits this

kind of description.

Lemma 9.4.3 (Block diagonalization). $\tilde{\mathbb{E}}[h_S(\mathbf{x})h_T(\mathbf{x})] = 0$ if $|S| \neq |T|$. If $S, T \in \binom{[n]}{d}$, then

$$\tilde{\mathbb{E}}[h_S(\mathbf{x})h_T(\mathbf{x})] = \hat{\sigma}_d^2 \cdot \langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle. \quad (9.43)$$

where

$$\hat{\sigma}_d^2 = d! \left(\frac{n-1}{n} \right)^d \prod_{i=0}^{d-1} \frac{n-2i}{n-2i-1} > 0. \quad (9.44)$$

Proof. The first claim follows since if $|S| \neq |T|$ then $h_S(\mathbf{x})$ and $h_T(\mathbf{x})$ belong to orthogonal eigenspaces of $\tilde{\mathbb{E}}$. For the second claim, recall that $\Psi : \mathbb{R}[x_1, \dots, x_n]_d^{\text{hom}} \rightarrow \mathbb{R}[z_1, \dots, z_{n-1}]_d^{\text{hom}}$ as defined in Proposition 9.2.11 is an isomorphism on each eigenspace with non-zero eigenvalue of $\tilde{\mathbb{E}}$. Moreover, by Proposition 9.2.11, each such eigenspace is isomorphic to $V_{\mathcal{H}}^{(d)}$ and thus is irreducible. So, since the apolar inner product in $\mathbb{R}[z_1, \dots, z_{n-1}]_d^{\text{hom}}$ is invariant under the action of S_n (permuting the $\langle v_i, \mathbf{z} \rangle$) and $\Psi(h_S(\mathbf{x})) = h_S(\mathbf{V}^\top \mathbf{z})$, the result must hold with *some* $\hat{\sigma}_d^2 \geq 0$ (which must be non-negative by the positivity of $\tilde{\mathbb{E}}$).

It remains to compute $\hat{\sigma}_d^2$, which is

$$\hat{\sigma}_d^2 = \frac{\tilde{\mathbb{E}}[h_S(\mathbf{x})^2]}{\|h_S(\mathbf{V}^\top \mathbf{z})\|_2^2} \quad (9.45)$$

for any $S \in \binom{[n]}{d}$. We computed the numerator in Proposition 9.3.2, so we need only compute the denominator.

Define, for $0 \leq k \leq d$, $\beta_{d,k} := \langle (\mathbf{V}^\top \mathbf{z})^S, (\mathbf{V}^\top \mathbf{z})^T \rangle$ for any $S, T \in \binom{[n]}{d}$ with $|S \cap T| = k$ (as this value only depends on $|S \cap T|$). With this notation, since $h_S(\mathbf{V}^\top \mathbf{z})$ is the apolar projection of $(\mathbf{V}^\top \mathbf{z})^S$ to $V_{\mathcal{H}}^{(d)}$ (as it is by definition the isotypic projection and by Proposition 9.2.12

$V_{\mathcal{H}}^{(d)}$ has multiplicity one in $\mathbb{R}[z_1, \dots, z_{n-1}]_d^{\text{hom}}$,

$$\begin{aligned}
\|h_S(\mathbf{V}^\top \mathbf{z})\|_\circ^2 &= \langle h_S(\mathbf{V}^\top \mathbf{z}), (\mathbf{V}^\top \mathbf{z})^S \rangle_\circ \\
&= \frac{\binom{n}{d} - \binom{n}{d-1}}{n!} \sum_{k=0}^d \beta_{d,k} \sum_{\substack{\pi \in S_n \\ |\pi(S) \cap S| = k}} \chi^{(n-d,d)}(\pi) \\
&= (-1)^d \binom{d}{|S \cap T|} \frac{\binom{n}{d} - \binom{n}{d-1}}{\binom{n}{d}} \sum_{k=0}^d (-1)^k \binom{d}{k} \beta_{d,k}, \quad (\text{Corollary 9.2.10})
\end{aligned}$$

and we are left with a similar sum as in Proposition 9.3.2, but now a d th forward difference of the sequence $f(k) = \beta_{d,k}$. We note that, choosing a concrete S and T in the definition, we may write

$$\beta_{d,k} = \left\langle \prod_{i=1}^d \langle \mathbf{v}_i, \mathbf{z} \rangle, \prod_{i=1}^k \langle \mathbf{v}_i, \mathbf{z} \rangle \prod_{i=d+1}^{2d-k} \langle \mathbf{v}_i, \mathbf{z} \rangle \right\rangle_\circ. \quad (9.46)$$

Using this representation, it is straightforward to show, again by induction, that

$$\Delta^a f(k) = \left\langle \prod_{i=1}^d \langle \mathbf{v}_i, \mathbf{z} \rangle, \prod_{i=1}^k \langle \mathbf{v}_i, \mathbf{z} \rangle \prod_{i=d+a+1}^{2d-k} \langle \mathbf{v}_i, \mathbf{z} \rangle \prod_{j=1}^a \langle \mathbf{v}_{d+j} - \mathbf{v}_{k+j}, \mathbf{z} \rangle \right\rangle_\circ. \quad (9.47)$$

Therefore, we have

$$\begin{aligned}
\sum_{k=0}^d (-1)^k \binom{d}{k} \beta_{d,k} &= \Delta^a f(0) \\
&= \left\langle \prod_{i=1}^d \langle \mathbf{v}_i, \mathbf{z} \rangle, \prod_{i=1}^d \langle \mathbf{v}_i - \mathbf{v}_{d+i}, \mathbf{z} \rangle \right\rangle_\circ.
\end{aligned}$$

where the only contribution applying the product rule to the inner product is in the matching of the two products in their given order, whereby

$$\begin{aligned}
&= \frac{1}{d!} \left(-1 - \frac{1}{n-1} \right)^d \\
&= \frac{(-1)^d}{d!} \left(\frac{n}{n-1} \right)^d, \quad (9.48)
\end{aligned}$$

and substituting completes the proof. \square

The following then follows immediately since the $h_S(\mathbf{x})$ are a spanning set of the isotropic components of $\mathbb{R}[\{\pm 1\}^n]$; this is analogous to Corollary 8.3.13 for the general spectral extension.

Corollary 9.4.4 (Gram matrix expression). *Let $h_{S,T} \in V_{\mathcal{H}}^{(|S|-2|T|)}$ be such that*

$$(\mathbf{V}^\top \mathbf{z})^S = \sum_{k=0}^{\lfloor |S|/2 \rfloor} \sum_{T \in \mathcal{M}_k(\{n\})} ((\mathbf{V}^\top \mathbf{z})^T)^2 h_{S,T}(\mathbf{V}^\top \mathbf{z}), \quad (9.49)$$

Define

$$h_{S,k}(\mathbf{x}) := \sum_{T \in \mathcal{M}_{(|S|-k)/2}(\{n\})} h_{S,T}(\mathbf{x}) \in \mathbb{R}[x_1, \dots, x_n]_k^{\text{hom}} \quad (9.50)$$

if $k \leq |S|$ and k and $|S|$ are of equal parity, and $h_{S,k} = 0$ otherwise. Then,

$$\tilde{\mathbb{E}}[\mathbf{x}^S \mathbf{x}^T] = \sum_{d=0}^{|S| \wedge |T|} \hat{\sigma}_d^2 \cdot \langle h_{S,d}(\mathbf{V}^\top \mathbf{z}), h_{T,d}(\mathbf{V}^\top \mathbf{z}) \rangle_\circ. \quad (9.51)$$

Proof of Theorem 9.4.1. Let $\mathbf{A}^{(d)} \in \mathbb{R}^{\dim(V_{\mathcal{H}}^{(d)}) \times \binom{[n]}{\leq d_{\max}}}$ have an isometric embedding of the $h_{S,d}$ as its columns. Then, the expression in (9.51) says that

$$\mathbf{Y}^{(n)} = \sum_{d=0}^{d_{\max}} \hat{\sigma}_d^2 \mathbf{A}^{(d)\top} \mathbf{A}^{(d)}. \quad (9.52)$$

Define the matrix

$$\mathbf{A} := \begin{bmatrix} \hat{\sigma}_0 \mathbf{A}^{(0)} \\ \vdots \\ \hat{\sigma}_{d_{\max}} \mathbf{A}^{(d_{\max})} \end{bmatrix}. \quad (9.53)$$

Then, $\mathbf{Y}^{(n)} = \mathbf{A}^\top \mathbf{A}$, so the non-zero eigenvalues of $\mathbf{Y}^{(n)}$ are equal to those of $\mathbf{A} \mathbf{A}^\top$.

By Schur's lemma, whenever $d' \geq d$ and d and d' have the same parity, then we have that $\{h_{S,d}\}_{S \in \binom{[n]}{d'}} \subset V_{\mathcal{H}}^{(d)}$ forms a tight frame in $V_{\mathcal{H}}^{(d)}$, since the $h_{S,d}$ form a union of orbits under

the action of S_n and $V_{\mathcal{H}}^{(d)}$ is irreducible. Let $f_{d',d}$ denote the frame constant, so that, for all $p \in V_{\mathcal{H}}^{(d)}$, we have

$$\sum_{S \in \binom{[n]}{d'}} \langle h_{S,d}, p \rangle \circ h = f_{d',d} p. \quad (9.54)$$

Let $f_{d',d} = 0$ if $d > d'$ or d' and d have different parity. We then have

$$\mathbf{A}^{(d)} \mathbf{A}^{(d)\top} = \left(\sum_{d'=d}^{d_{\max}} f_{d',d} \right) \mathbf{I}_{\dim(V_{\mathcal{H}}^{(d)})}. \quad (9.55)$$

In particular, we may exhibit $\dim(V_{\mathcal{H}}^{(d)})$ many eigenvectors of $\mathbf{A}\mathbf{A}^\top$ with this positive eigenvalue supported on the d th block (in the same block decomposition as that of \mathbf{A}). So, $\mathbf{A}\mathbf{A}^\top \succ \mathbf{0}$ strictly, and its distinct eigenvalues are $\lambda_{n,0}, \dots, \lambda_{n,d_{\max}} > 0$ given by

$$\lambda_{n,d} = \hat{\sigma}_d^2 \sum_{d'=d}^{d_{\max}} f_{d',d} \text{ with multiplicity } \dim(V_{\mathcal{H}}^{(d)}) = \binom{n}{d} - \binom{n}{d-1}. \quad (9.56)$$

Thus these are also precisely the positive eigenvalues of $\mathbf{Y}^{(n)}$, and the multiplicity of the zero eigenvalue of $\mathbf{Y}^{(n)}$ is $\binom{n}{\leq d_{\max}} - \sum_{d=0}^{d_{\max}} \left(\binom{n}{d} - \binom{n}{d-1} \right) = \binom{n}{\leq d_{\max}-1}$, as claimed.

We now turn to the explicit computation of the eigenvalues. Let us write

$$\eta_{d',d}^2 := \|h_{S,d}\|_0^2 \text{ for any } S \in \binom{[n]}{d'}, \quad (9.57)$$

noting that these numbers are all equal by symmetry. Then, the frame constants from (9.54) are

$$f_{d',d} = \frac{\binom{n}{d'}}{\dim(V_{\mathcal{H}}^{(d)})} \eta_{d',d}^2 = \frac{\binom{n}{d'}}{\binom{n}{d} - \binom{n}{d-1}} \eta_{d',d}^2. \quad (9.58)$$

It remains to compute the $\eta_{d',d}$, which will yield the $f_{d',d}$ and thus the eigenvalues $\lambda_{n,d}$.

We first note that, by our earlier computation in Lemma 9.4.3, for any given $S \in \binom{[n]}{d}$,

$$\eta_{d,d}^2 = \|h_S\|_0^2 = \frac{\binom{n}{d} - \binom{n}{d-1}}{\binom{n}{d}} \frac{1}{d!} \left(\frac{n}{n-1}\right)^d. \quad (9.59)$$

Therefore,

$$f_{d,d} = \frac{\binom{n}{d}}{\binom{n}{d} - \binom{n}{d-1}} \eta_{d,d}^2 = \frac{1}{d!} \left(\frac{n}{n-1}\right)^d. \quad (9.60)$$

To compute the $\eta_{d',d}$ with $d' > d$, we use that $\tilde{\mathbb{E}}$ itself can be used to compute the following inner products, by Corollary 9.4.4:

$$\tilde{\mathbb{E}}[\mathbf{x}^S h_T(\mathbf{x})] = \hat{\sigma}_d^2 \langle h_{S,d}(\mathbf{x}), h_T(\mathbf{x}) \rangle_0. \quad (9.61)$$

Using that the $\{h_T(\mathbf{x})\}_{T \in \binom{[n]}{d}}$ form a tight frame with frame constant $f_{d,d}$, we have

$$\begin{aligned} \eta_{d',d}^2 &= \|h_{S,d}\|_0^2 \\ &= \frac{1}{f_{d,d}} \sum_{T \in \binom{[n]}{d}} \langle h_{S,d}(\mathbf{x}), h_T(\mathbf{x}) \rangle_0^2 \\ &= \frac{1}{\hat{\sigma}_d^4 f_{d,d}} \sum_{T \in \binom{[n]}{d}} (\tilde{\mathbb{E}}[\mathbf{x}^S h_T(\mathbf{x})])^2 \end{aligned} \quad (9.62)$$

We next expand these pseudoexpectations directly:

$$\begin{aligned} \tilde{\mathbb{E}}[\mathbf{x}^S h_T(\mathbf{x})] &= \frac{\binom{n}{d} - \binom{n}{d-1}}{n!} \sum_{\pi \in S_n} \chi_{(n-d,d)}(\pi) \mathbf{x}^{S+\pi(T)} \\ &= \frac{\binom{n}{d} - \binom{n}{d-1}}{n!} \sum_{\pi \in S_n} \chi_{(n-d,d)}(\pi) \alpha_{d+d'-2|S \cap \pi(T)|} \\ &= \frac{\binom{n}{d} - \binom{n}{d-1}}{n!} \sum_{k=0}^d \alpha_{d+d'-2k} \sum_{\substack{\pi \in S_n \\ |S \cap \pi(T)|=k}} \chi_{(n-d,d)}(\pi) \end{aligned}$$

Suppose now that $|S \cap T| = \ell$. Then, by Corollary 9.2.10 we have

$$= \frac{\binom{n}{d} - \binom{n}{d-1}}{\binom{n}{\ell, d-\ell, d'-\ell, n-d-d'+\ell}} \binom{n-2d}{d'-d} \binom{d}{\ell} (-1)^\ell \sum_{k=0}^d \binom{d}{k} (-1)^k \alpha_{d+d'-2k}$$

The remaining sum is one we evaluated in the course of our proof of Proposition 9.3.2 using finite differences. Substituting that result here then gives

$$= \frac{\binom{n}{d} - \binom{n}{d-1}}{\binom{n}{\ell, d-\ell, d'-\ell, n-d-d'+\ell}} \binom{n-2d}{d'-d} \binom{d}{\ell} \prod_{i=0}^{d-1} \frac{n-2i}{n-d'+d-2i-1} \cdot \alpha_{d'-d}. \quad (9.63)$$

Substituting this into the summation that occurs in our expression for $\eta_{d,d'}$, we then find

$$\begin{aligned} & \sum_{T \in \binom{[n]}{d}} (\tilde{\mathbb{E}}[\mathbf{x}^S h_T(\mathbf{x})])^2 \\ &= \sum_{\ell=0}^d \binom{d'}{\ell} \binom{n-d'}{d-\ell} \left(\frac{\binom{n}{d} - \binom{n}{d-1}}{\binom{n}{\ell, d-\ell, d'-\ell, n-d-d'+\ell}} \binom{n-2d}{d'-d} \binom{d}{\ell} \prod_{i=0}^{d-1} \frac{n-2i}{n-d'+d-2i-1} \cdot \alpha_{d'-d} \right)^2 \\ &= \alpha_{d'-d}^2 \left(\left(\binom{n}{d} - \binom{n}{d-1} \right) \binom{n-2d}{d'-d} \prod_{i=0}^{d-1} \frac{n-2i}{n-d'+d-2i-1} \right)^2 \\ & \quad \sum_{\ell=0}^d \binom{d'}{\ell} \binom{n-d'}{d-\ell} \frac{\binom{d}{\ell}^2}{\binom{n}{\ell, d-\ell, d'-\ell, n-d-d'+\ell}^2} \\ &= \alpha_{d'-d}^2 \left(\left(\binom{n}{d} - \binom{n}{d-1} \right) \binom{n-2d}{d'-d} \prod_{i=0}^{d-1} \frac{n-2i}{n-d'+d-2i-1} \right)^2 \\ & \quad \frac{d!^2 d'! (n-d')! (n-d-d')! (d'-d)!}{n!^2} \sum_{\ell=0}^d \binom{d'-\ell}{d'-d} \binom{n-d-d'+\ell}{n-d-d'} \end{aligned}$$

and the remaining sum evaluates by the Chu-Vandermonde identity to

$$\begin{aligned} &= \alpha_{d'-d}^2 \left(\left(\binom{n}{d} - \binom{n}{d-1} \right) \binom{n-2d}{d'-d} \prod_{i=0}^{d-1} \frac{n-2i}{n-d'+d-2i-1} \right)^2 \\ & \quad \frac{d!^2 d'! (n-d')! (n-d-d')! (d'-d)!}{n!^2} \binom{n-d+1}{d}. \end{aligned} \quad (9.64)$$

Having reached this expression, we may substitute for $\eta_{d',d}^2$ and find many cancellations, obtaining

$$\begin{aligned}
\eta_{d',d}^2 &= \frac{1}{\hat{\sigma}_d^4 f_{d,d}} \sum_{T \in \binom{[n]}{d}} (\tilde{\mathbb{E}}[\mathbf{x}^S h_T(\mathbf{x})])^2 \\
&= \alpha_{d'-d}^2 \left(\frac{n}{n-1} \right)^d \left(\prod_{i=0}^{d-1} \frac{n-2i-1}{n-2i-1-d'+d} \right)^2 \left(\binom{n}{d} - \binom{n}{d-1} \right)^2 \\
&\quad \frac{d!d'!(n-d'!(n-2d)!^2}{n!^2(n-d-d')!(d'-d)!} \binom{n-d+1}{d}.
\end{aligned} \tag{9.65}$$

Then we may again substitute for $\lambda_{n,d}$ and find more cancellations, obtaining

$$\begin{aligned}
\lambda_{n,d} &= \hat{\sigma}_d^2 \sum_{d'=d}^{d_{\max}} f_{d',d} \\
&= \hat{\sigma}_d^2 \sum_{d'=d}^{d_{\max}} \frac{\binom{n}{d'}}{\binom{n}{d} - \binom{n}{d-1}} \eta_{d',d}^2 \\
&= n! \sum_{d'=d}^{d_{\max}} \frac{\alpha_{d'-d}^2}{(n-d-d')!(d'-d)!} \prod_{i=0}^{d-1} \frac{1}{(n-2i-1-d'+d)^2}.
\end{aligned} \tag{9.66}$$

The formula for $\lambda_{n,0}$ then follows immediately. To obtain the recursion, we compute

$$\begin{aligned}
\lambda_{n+2,d+1} &= (n+2)! \sum_{d'=d+1}^{(n+1)/2} \frac{\alpha_{n+2,d'-d-1}^2}{(n-d-d'+1)!(d'-d-1)!} \prod_{i=0}^d \frac{1}{(n-2i+2-d'+d)^2} \\
&= (n+2)! \sum_{d'=d}^{(n-1)/2} \frac{\alpha_{n+2,d'-d}^2}{(n-d-d')!(d'-d)!} \prod_{i=0}^d \frac{1}{(n-2i+1-d'+d)^2} \\
&= (n+2)! \sum_{d'=d}^{(n-1)/2} \frac{1}{(n+1-d'+d)^2} \frac{\alpha_{n+2,d'-d}^2}{(n-d-d')!(d'-d)!} \prod_{i=0}^{d-1} \frac{1}{(n-2i-1-d'+d)^2}
\end{aligned}$$

and noting that $\alpha_{n+2,2k} = \frac{n-2k+1}{n+1} \alpha_{n,2k}$, we find

$$\begin{aligned}
 &= \frac{(n+2)!}{(n+1)^2} \sum_{d'=d}^{(n-1)/2} \frac{\alpha_{n,d'-d}^2}{(n-d-d')!(d'-d)!} \prod_{i=0}^{d-1} \frac{1}{(n-2i-1-d'+d)^2} \\
 &= \frac{n+2}{n+1} \lambda_{n,d},
 \end{aligned} \tag{9.67}$$

completing the proof. □

10 | SUM-OF-FORESTS PSEUDOMOMENTS AND LIFTING THEOREMS

Having established with a deterministic example that our spectral pseudomoment extension from Chapter 8 gives correct results at least sometimes, we now continue with our earlier plan of extending generic low-rank projection matrices.

SUMMARY AND REFERENCES This chapter is based on part of the reference [Kun20b]. The following is a summary of our main results in this chapter.

1. (Theorem 10.2.3) An extension to degree $\omega(1)$ (more precisely, an extension to degree $\Omega(\log n / \log \log n)$ for cases we will be interested in) that we expect to succeed for incoherent degree 2 pseudomoments of rank $(1 - o(1))n$.
2. (Theorem 10.11.2) An extension to degree 6 that we expect to succeed for incoherent degree 2 pseudomoments of rank δn for $\delta \in (0, 1)$.

We also emphasize the following important technical details of our proofs that may be interesting contribution beyond the lifting theorems themselves.

3. (Lemma 10.3.11) The derivation of the Möbius function of a set of forests under a “compositional” ordering, which appears as the coefficients of corresponding terms of pseudomoments.

4. (Remark 10.8.6) A combinatorial identity that, per our arguments, appears to be responsible for the possibility of pseudomoments satisfying both the positivity and entrywise constraints. The identity relates the Möbius function of a partially ordered set of forests to a sum over matrices associated to a pair of partitions similar to those appearing in the Robinson-Schensted-Knuth (RSK) correspondence of representation theory and the combinatorics of Young tableaux.

PRIOR WORK The idea of a “lifting” theorem that automatically applies to many low-degree pseudomoments is a relatively new one; our description of our results in this way (here and in [Kun20b]) is inspired by [MRX20], who appear to be the first to have given an explicit such statement. Our pseudomoment construction is inspired by on an old construction of harmonic polynomials dating back to Maxwell [Max73] and Sylvester [Syl76] and rediscovered many times since as well as a generalization due to Clerc [Cle00]; see Section 10.1.1.

10.1 CONCRETE PSEUDOMOMENT EXTENSION

Recall that our construction of spectral extensions in Chapter 8 left as an unspecified input the description of how a polynomial of the form $(\mathbf{V}^\top \mathbf{z})^S = \prod_{i \in S} \langle \mathbf{v}_i, \mathbf{z} \rangle$ decomposes into an ideal part of a linear combination of multiples of $\langle \mathbf{v}_i, \mathbf{z} \rangle^2$, belonging to the subspace V_I , and a harmonic part that is a zero of any linear combination of the differential operators $\langle \mathbf{v}_i, \partial \rangle^2$, belonging to the subspace $V_{\mathcal{H}}$. In this section, we develop a heuristic method to compute these projections. Since $(\mathbf{V}^\top \mathbf{z})^S$ is the sum of the two projections, it suffices to compute either one. We will work with the projection to the multiharmonic subspace $V_{\mathcal{H}}$. We warn in advance that this portion of the derivation is not mathematically rigorous; our goal is only to obtain a plausible prediction for the projections in question, which we will then analyze more precisely. Our construction will be based on a technique for computing these

projections exactly for harmonic polynomials and certain multiharmonic generalizations, which we review below.

10.1.1 GENERALIZED MAXWELL-SYLVESTER REPRESENTATIONS

We describe a line of work describing how the projection of a polynomial to $V_{\mathcal{H}}$ may sometimes be computed, rather surprisingly, by differentiating a *Green's function* associated to the defining PDE or system of PDEs. The following is the clearest instance of this idea, which concerns the case of harmonic polynomials.

Proposition 10.1.1 (Theorem 1.7 of [AR95]; Theorem 5.18 of [ABW13]). *Suppose $n \geq 3$.¹ Let $V_{\mathcal{H}} \subset \mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]_d^{\text{hom}}$ be the subspace of harmonic polynomials ($q(\mathbf{y})$ with $\Delta q = 0$), and let $P_{\mathcal{H}}$ be the apolar projection to $V_{\mathcal{H}}$. Define*

$$\varphi(\mathbf{y}) := \|\mathbf{y}\|_2^{2-n} \text{ (the Green's function of } \Delta), \text{ and} \quad (10.1)$$

$$K[u](\mathbf{y}) := \varphi(\mathbf{y})u(\mathbf{y}/\|\mathbf{y}\|_2^2) \text{ (the Kelvin transform),} \quad (10.2)$$

the latter defined for $u : \mathbb{R}^n \setminus \{\mathbf{0}\} \rightarrow \mathbb{R}$ a smooth function. Let $p \in \mathbb{R}[\mathcal{Y}_1, \dots, \mathcal{Y}_n]_d^{\text{hom}}$. Then,

$$P_{\mathcal{H}}[q] = \frac{1}{\prod_{i=0}^{d-1} (2 - n - 2i)} K[q(\partial)\varphi]. \quad (10.3)$$

Roughly speaking, the Kelvin transform is a generalization to higher dimensions of inversion across a circle, so this result says that apolar projections to harmonic polynomials may be computed by inverting corresponding derivatives of the Green's function of Δ . In other words again, the Green's function is a kind of generating function of the apolar projections of the monomials.

This result has a long history. At least for $n = 3$, the idea and its application to

¹A variant of this result also holds for $n = 2$; see Section 4 of [AR95].

the expansion of Example 8.3.10 were already present in classical physical reasoning of Maxwell [Max73]. Soon after, Sylvester [Syl76] gave a mathematical treatment, mentioning that an extension to other n is straightforward. See Section VII.5.5 of [CH89] on “The Maxwell-Sylvester representation of spherical harmonics” for a modern exposition. These ideas were rediscovered by [AR95]; there and in the later textbook treatment [ABW13] there is greater emphasis on $P_{\mathcal{H}}$ being a projection, though the fact that the apolar inner product makes it an *orthogonal* projection goes unmentioned. Further historical discussion and a presentation with all of the ideas relevant to us are given in an unpublished note of Gichev [Gic]. A note of Arnol’d [Arn96] and Appendix A of his lecture notes [Arn13] also discuss topological interpretations of these results and give historical commentary.

When we seek to apply these ideas in our setting, we will want to project to *multiharmonic* polynomials, which satisfy $p_i(\partial)q = 0$ for several polynomials p_1, \dots, p_m . In our case the polynomials will be quadratic, but a generalization to arbitrary polynomials is also sensible. This question has been studied much less. The main work we are aware of in this direction is due to Clerc [Cle00] (whose Green’s function construction was suggested earlier in Herz’s thesis [Her55]; see Lemma 1.6 of the latter), where the p_i are quadratic forms with matrices spanning a Jordan subalgebra of $\mathbb{R}_{\text{sym}}^{r \times r}$. The following is one, essentially trivial, instance of those results.

Proposition 10.1.2. *Let v_1, \dots, v_r be an orthonormal basis of \mathbb{R}^r . Define an associated Green’s function and Kelvin transform*

$$\varphi(z) := \prod_{i=1}^r \langle v_i, z \rangle, \tag{10.4}$$

$$K[f](z) := \varphi(z) f \left(\sum_{i=1}^r \langle v_i, z \rangle^{-1} v_i \right). \tag{10.5}$$

Then, for any $p \in \mathbb{R}[x_1, \dots, x_r] \setminus \{0\}$, the apolar projection of p to the harmonic subspace is

$K[p(\partial)\varphi]$.

In this case, it is easy to give a hands-on proof: one may write p in the monomial basis $\langle \mathbf{v}_i, \mathbf{z} \rangle$, and in this basis the desired projection is just the multilinear part of p . On the other hand, we have $p(\partial)\varphi = q/\varphi$, where q is the multilinear part of p , and the result follows. Though this is a simple derivation, we will see that extending it to overcomplete families of vectors \mathbf{v}_i in fact forms one of the key heuristic steps in our derivation.

10.1.2 CLOSING THE PREDICTION

We now use these ideas to heuristically project to our $V_{\mathcal{H}}$, and thus find a closed-form prediction of $\tilde{\mathbb{E}}$. The difference between our setting and the above is that $n > r$, and the \mathbf{v}_i form an overcomplete set. In particular, in the Kelvin transform, it is not guaranteed that, given some \mathbf{z} , there exists a \mathbf{z}' such that $\mathbf{V}^\top \mathbf{z}'$ is the coordinatewise reciprocal of $\mathbf{V}^\top \mathbf{z}$, so a genuine “inversion” like in the case of an orthonormal basis may be impossible. We also remark that it is not possible to apply the results of [Cle00] directly, since the $\mathbf{v}_i \mathbf{v}_i^\top$ typically do not span a Jordan subalgebra: $\frac{1}{2}(\mathbf{v}_i \mathbf{v}_i^\top \mathbf{v}_j \mathbf{v}_j^\top + \mathbf{v}_j \mathbf{v}_j^\top \mathbf{v}_i \mathbf{v}_i^\top) = \frac{1}{2} \langle \mathbf{v}_i, \mathbf{v}_j \rangle (\mathbf{v}_i \mathbf{v}_j^\top + \mathbf{v}_j \mathbf{v}_i^\top)$ is not always a linear combination of the $\mathbf{v}_k \mathbf{v}_k^\top$, since in fact the $\frac{1}{2}(\mathbf{v}_i \mathbf{v}_j^\top + \mathbf{v}_j \mathbf{v}_i^\top)$ span all of $\mathbb{R}_{\text{sym}}^{r \times r}$ when the \mathbf{v}_i are overcomplete. There *are* plenty of non-trivial Jordan subalgebras of $\mathbb{R}_{\text{sym}}^{r \times r}$, as discussed e.g. in [BES20]; they are just not generated in this way.

Nonetheless, let us continue with the first part of the calculation by analogy with Proposition 10.1.2. Define

$$\varphi(\mathbf{z}) := \prod_{i=1}^n \langle \mathbf{v}_i, \mathbf{z} \rangle. \quad (10.6)$$

Then, one may compute inductively by the product rule that

$$(\mathbf{V}^\top \partial)^S \varphi(\mathbf{z}) = \left(\sum_{\sigma \in \text{Part}(S)} \prod_{A \in \sigma} (-1)^{|A|-1} (|A| - 1)! \left\{ \sum_{a=1}^n \prod_{j \in A} M_{a,j} \cdot \langle \mathbf{v}_a, \mathbf{z} \rangle^{-|A|} \right\} \right) \varphi(\mathbf{z}). \quad (10.7)$$

(That various summations over partitions arise in such calculations is well-known; see, e.g., [Har06] for a detailed discussion.) We now take a leap of faith: despite the preceding caveats, let us suppose we could make a fictitious mapping $\tilde{F} : \mathbb{R}^r \rightarrow \mathbb{R}^r$ that would invert the values of each $\langle \mathbf{v}_i, \mathbf{z} \rangle$, i.e., $\langle \mathbf{v}_i, \tilde{F}(\mathbf{z}) \rangle = \langle \mathbf{v}_i, \mathbf{z} \rangle^{-1}$ for each $i \in [n]$. Then, we would define a Kelvin transform (also fictitious) by $\tilde{K}[f](\mathbf{z}) = f(\tilde{F}(\mathbf{z})) \cdot \varphi(\mathbf{z})$. Using this, and noting that $\varphi(\tilde{F}(\mathbf{z})) = \varphi(\mathbf{z})^{-1}$, we predict

$$\begin{aligned} P_{\mathcal{H}} \left[(\mathbf{V}^\top \mathbf{z})^S \right] &= \tilde{K} \left[(\mathbf{V}^\top \boldsymbol{\theta})^S \varphi \right] (\mathbf{z}) \\ &= \sum_{\sigma \in \text{Part}(S)} \prod_{A \in \sigma} (-1)^{|A|-1} (|A| - 1)! \left\{ \sum_{a=1}^n \prod_{j \in A} M_{aj} \cdot \langle \mathbf{v}_a, \mathbf{z} \rangle^{|A|} \right\}. \end{aligned} \quad (10.8)$$

We make one adjustment to this prediction: when $|A| = 1$ with $A = \{i\}$, then the inner summation is $\sum_{a=1}^n M_{ai} \langle \mathbf{v}_a, \mathbf{z} \rangle = (\mathbf{M}\mathbf{V}^\top \mathbf{z})_i = (\mathbf{V}^\top \mathbf{V}\mathbf{V}^\top \mathbf{z})_i \approx \delta^{-1} \langle \mathbf{v}_i, \mathbf{z} \rangle$. However, the factor of δ^{-1} here appears to be superfluous; one way to confirm this is to compare this prediction for $|S| = 2$ with the direct calculations of $P_{\mathcal{H}}$ for $d = 2$ in Chapter 7 for the case of ETFs. Thus we omit this factor in our final prediction.

We are left with the following prediction for the harmonic projection. First, it will be useful to set notation for the polynomials occurring inside the summation.

Definition 10.1.3. For $S \in \mathcal{M}([n])$ with $S \neq \emptyset$, $m \in \mathbb{N}$, and $\mathbf{x} \in \mathbb{R}^n$, define

$$q_{S,m}(\mathbf{x}) := \begin{cases} x_i^m & \text{if } |S| = 1 \text{ with } S = \{i\}, \\ \sum_{a=1}^n \prod_{j \in T} M_{aj} x_a^m & \text{otherwise,} \end{cases} \quad (10.9)$$

$$q_S(\mathbf{x}) := q_{S,|S|}(\mathbf{x}). \quad (10.10)$$

We then predict

$$P_{\mathcal{H}} \left[(\mathbf{V}^\top \mathbf{z})^S \right] \approx \sum_{\sigma \in \text{Part}(S)} \prod_{A \in \sigma} (-1)^{|A|-1} (|A| - 1)! q_A(\mathbf{V}^\top \mathbf{z}). \quad (10.11)$$

By the orthogonality of the ideal and harmonic subspaces, we also immediately obtain a prediction for the orthogonal projection to V_I :

$$P_I \left[(\mathbf{V}^\top \mathbf{z})^S \right] = (\mathbf{V}^\top \mathbf{z})^S - P_{\mathcal{H}} \left[(\mathbf{V}^\top \mathbf{z})^S \right] \approx - \sum_{\substack{\sigma \in \text{Part}(S) \\ |\sigma| < |S|}} \prod_{A \in \sigma} (-1)^{|A|-1} (|A|-1)! q_A(\mathbf{V}^\top \mathbf{z}). \quad (10.12)$$

We therefore obtain the following corresponding predictions for the polynomials $h_S(\mathbf{x})$ and $r_S(\mathbf{x})$ appearing in Lemma 8.3.11. We redefine these here for the duration of this chapter, though we emphasize that these do *not* give the actual projections $P_{\mathcal{H}}$ or P_I but rather are only heuristic approximations.

Definition 10.1.4 (h_S and r_S polynomials). *For $S \subseteq [n]$ and $\mathbf{x} \in \mathbb{R}^n$, define*

$$h_S(\mathbf{x}) := \sum_{\sigma \in \text{Part}(S)} \prod_{A \in \sigma} (-1)^{|A|-1} (|A|-1)! q_A(\mathbf{x}), \quad (10.13)$$

$$r_S(\mathbf{x}) := - \sum_{\substack{\sigma \in \text{Part}(S) \\ |\sigma| < |S|}} \prod_{A \in \sigma} (-1)^{|A|-1} (|A|-1)! q_A(\mathbf{x}). \quad (10.14)$$

The “lowered” polynomials $r_S^\downarrow(\mathbf{x})$ may also be defined by simply reducing the powers of x_i appearing in $q_A(\mathbf{x})$. Here again, however, we make a slight adjustment: when $|A| = 2$ with $A = \{i, j\}$, we would compute $q_{A,0}(\mathbf{x}) = \sum_{a=1}^n M_{ai} M_{aj} = (\mathbf{M}^2)_{ij} \approx \delta^{-1} M_{ij}$. This factor of δ^{-1} again appears to be superfluous, with the same justification as before. Removing it, we make the following definition.

Definition 10.1.5. *For $S \in \mathcal{M}([n])$ with $S \neq \emptyset$ and $|S|$ even, define*

$$q_S^\downarrow(\mathbf{x}) = q_S^\downarrow(\mathbf{x}; \mathbf{M}) := \begin{cases} q_{S,1}(\mathbf{x}) & \text{if } |S| \text{ is odd,} \\ M_{ij} & \text{if } |S| = 2 \text{ with } S = \{i, j\}, \\ q_{S,0}(\mathbf{x}) & \text{if } |S| \geq 4 \text{ is even.} \end{cases} \quad (10.15)$$

With this adjustment, we obtain the following values of r_S^\downarrow appearing in Lemma 8.3.11 and

h_S^\dagger appearing in Remark 8.3.12.

Definition 10.1.6 (h_S^\dagger and r_S^\dagger polynomials). For $S \subseteq [n]$ and $\mathbf{x} \in \mathbb{R}^n$, define

$$h_S^\dagger(\mathbf{x}) := \sum_{\sigma \in \text{Part}(S)} \prod_{A \in \sigma} (-1)^{|A|-1} (|A|-1)! q_A^\dagger(\mathbf{x}), \quad (10.16)$$

$$r_S^\dagger(\mathbf{x}) := - \sum_{\substack{\sigma \in \text{Part}(S) \\ |\sigma| < |S|}} \prod_{A \in \sigma} (-1)^{|A|-1} (|A|-1)! q_A^\dagger(\mathbf{x}). \quad (10.17)$$

Substituting the approximations (10.13) and (10.17) into the pseudoexpectation expression we obtained in Lemma 8.3.11, we are now equipped with a fully explicit heuristic recursion for $\tilde{\mathbb{E}}$, up to the choice of the constants σ_d^2 . Let us demonstrate how, with some further heuristic steps, this recovers our Definition 10.1.13 for $d = 1$ and $d = 2$. For $d = 1$ we expect a sanity check recovering that $\tilde{\mathbb{E}}[x_i x_j] = M_{ij}$, while for $d = 2$ we expect to recover the formula we obtained for ETFs in Chapter 7 and then rederived with the Gaussian conditioning interpretation in Section 8.2.

Example 10.1.7 ($d = 1$). If $|S| = 1$ with $S = \{i\}$, then there is no partition σ of S with $|\sigma| < |S|$, so $h_{\{i\}}(\mathbf{x}) = x_i$ and $r_{\{i\}}^\dagger(\mathbf{x}) = 0$. Thus, if $S = \{i\}$ and $T = \{j\}$, we are left with simply

$$\tilde{\mathbb{E}}(x_i, x_j) = \sigma_1^2 \delta \cdot \left\langle \langle \mathbf{v}_i, \mathbf{z} \rangle, \langle \mathbf{v}_j, \mathbf{z} \rangle \right\rangle_\circ = \sigma_1^2 \delta \cdot \langle \mathbf{v}_i, \mathbf{v}_j \rangle = \sigma_1^2 \delta \cdot M_{ij} \quad (10.18)$$

which upon taking $\sigma_1^2 = \delta^{-1}$ gives $\tilde{\mathbb{E}}(x_i, x_j) = M_{ij}$, as expected.

Example 10.1.8 ($d = 2$). If $S = \{i, j\}$ then the only partition σ of S with $|\sigma| < |S|$ is the partition $\sigma = \{\{i, j\}\}$. Therefore,

$$\begin{aligned} h_{\{i,j\}}(\mathbf{x}) &= x_i x_j - (2-1)! \cdot q_{\{i,j\}}^\dagger(\mathbf{x}) \\ &= x_i x_j - \sum_{a=1}^n M_{ai} M_{aj} x_a^2, \end{aligned} \quad (10.19)$$

$$r_{\{i,j\}}^\dagger(\mathbf{x}) = (2-1)! \cdot Z^{\{i,j\}} = M_{ij}. \quad (10.20)$$

If, furthermore, $T = \{k, \ell\}$, then we compute

$$\begin{aligned}
& \langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle. \\
&= \left\langle \langle \mathbf{v}_i, \mathbf{z} \rangle \langle \mathbf{v}_j, \mathbf{z} \rangle - \sum_{a=1}^n M_{ai} M_{aj} \langle \mathbf{v}_a, \mathbf{z} \rangle^2, \langle \mathbf{v}_k, \mathbf{z} \rangle \langle \mathbf{v}_\ell, \mathbf{z} \rangle - \sum_{a=1}^n M_{ak} M_{a\ell} \langle \mathbf{v}_a, \mathbf{z} \rangle^2 \right\rangle. \\
&= \frac{1}{2} M_{ik} M_{j\ell} + \frac{1}{2} M_{i\ell} M_{jk} - 2 \sum_{a=1}^n M_{ai} M_{aj} M_{ak} M_{a\ell} + \sum_{a=1}^n \sum_{b=1}^n M_{ai} M_{aj} M_{bk} M_{b\ell} M_{ab}^2
\end{aligned}$$

and if we make the approximation that the only important contributions from the final sum are when $a = b$, then we obtain

$$\approx \frac{1}{2} M_{ik} M_{j\ell} + \frac{1}{2} M_{i\ell} M_{jk} - \sum_{a=1}^n M_{ai} M_{aj} M_{ak} M_{a\ell}. \quad (10.21)$$

Substituting the above into Lemma 8.3.11, we compute

$$\begin{aligned}
& \tilde{\mathbb{E}}(x_i x_j, x_k x_\ell) \\
&= M_{ij} M_{k\ell} + \sigma_2^2 \delta^2 \cdot \left\langle \langle \mathbf{v}_i, \mathbf{z} \rangle \langle \mathbf{v}_j, \mathbf{z} \rangle - \sum_{a=1}^n M_{ai} M_{aj} x_a^2, \langle \mathbf{v}_k, \mathbf{z} \rangle \langle \mathbf{v}_\ell, \mathbf{z} \rangle - \sum_{a=1}^n M_{ak} M_{a\ell} x_a^2 \right\rangle. \\
&= M_{ij} M_{k\ell} + \sigma_2^2 \delta^2 \left(\frac{1}{2} M_{ik} M_{j\ell} + \frac{1}{2} M_{i\ell} M_{jk} - \sum_{a=1}^n M_{ai} M_{aj} M_{ak} M_{a\ell} \right)
\end{aligned}$$

where we see that the only value of σ_2^2 that will both make $\tilde{\mathbb{E}}$ factor through multiplication and achieve the normalization conditions $\tilde{\mathbb{E}}(x_i x_j, x_i x_j) \approx 1$ is $\sigma_2^2 = 2\delta^{-2}$, which gives

$$= M_{ij} M_{k\ell} + M_{ik} M_{j\ell} + M_{i\ell} M_{jk} - 2 \sum_{a=1}^n M_{ai} M_{aj} M_{ak} M_{a\ell}, \quad (10.22)$$

the formula discussed in Section 8.2.

The computation above is essentially the same as that in Section 8.2, only rewritten in the language of degree 2 homogeneous polynomials rather than symmetric matrices.

One may continue these (increasingly tedious) calculations for larger d to attempt to find a pattern in the resulting polynomials of M . This is how we arrive at the values given below, and we will sketch the basic idea of how to close the recursion. The more precise version of that calculation is rather more complicated and will be implicit in our proofs of the first lifting theorem in Sections 10.7 and 10.8. We mention for now that, in these heuristic calculations, it is important to be careful with variants of the step above where we restricted the double summation to indices $a = b$ (indeed, that step in the above derivation is not always valid; as we will detail in Section 10.10, this is actually the crux of the difficulty in applying our method to low-rank rather than high-rank M). This type of operation is valid only when the difference between the *matrices* containing the given summations as their entries has negligible operator norm—a subtle condition. We gloss this point for now, but give much attention to these considerations in the technical proof details below and in Section 10.13.

10.1.3 SUM-OF-FORESTS PSEUDOMOMENTS

The precise definition of our construction is as follows. Generalizing the degree 4 case, the pseudoexpectation is formed as a linear combination of a particular type of polynomial in the degree 2 pseudomoments, which we describe below.

Definition 10.1.9 (Contractive graphical scalar). *Suppose $G = (V, E)$ is a graph with two types of vertices, which we denote \bullet and \square visually and whose subsets we denote $V = V^\bullet \sqcup V^\square$. Suppose also that V^\bullet is equipped with a labelling $\kappa : V^\bullet \rightarrow [|V^\bullet|]$. For $\mathbf{s} \in [n]^{|V^\bullet|}$ and $\mathbf{a} \in [n]^{V^\square}$, let $f_{\mathbf{s}, \mathbf{a}} : V \rightarrow [n]$ have $f_{\mathbf{s}, \mathbf{a}}(v) = s_{\kappa(v)}$ for $v \in V^\bullet$ and $f_{\mathbf{s}, \mathbf{a}}(v) = a_v$ for $v \in V^\square$. Then, for $M \in \mathbb{R}_{\text{sym}}^{n \times n}$, we define*

$$Z^G(\mathbf{M}; \mathbf{s}) := \sum_{\mathbf{a} \in [n]^{V^\square}} \prod_{\{v, w\} \in E} M_{f_{\mathbf{s}, \mathbf{a}}(v), f_{\mathbf{s}, \mathbf{a}}(w)}. \quad (10.23)$$

We call this quantity a *contractive graphical scalar* (CGS) whose diagram is the graph G . When S is a set or multiset of elements of $[n]$ with $|S| = |V^*|$, we define $Z^G(\mathbf{M}; S) := Z^G(\mathbf{M}; \mathbf{s})$ where \mathbf{s} is the tuple of the elements of S in ascending order.

As an intuitive summary, the vertices of the underlying diagram G correspond to indices in $[n]$, and edges specify multiplicative factors given by entries of \mathbf{M} . The \bullet vertices are “pinned” to the indices specified by \mathbf{s} , while the \square vertices are “contracted” over all possible index assignments. CGSs are also a special case of existing formalisms, especially popular in the physics literature, of *trace diagrams* and *tensor networks* [BB17].

Remark 10.1.10. *Later, in Section 10.4, we will also study contractive graphical matrices (CGMs), set- or tuple-indexed matrices whose entries are CGSs with the set S varying according to the indices. CGMs are similar to graphical matrices as used in other work on SOS relaxations [AMP16, BHK⁺19, MRX20]. Aside from major but ultimately superficial notational differences, the main substantive difference is that graphical matrices require all indices labelling the vertices in the summation to be different from one another, while CGMs and CGSs do not. This restriction is natural in the combinatorial setting—if \mathbf{M} is an adjacency matrix then the entries of graphical matrices count occurrences of subgraphs—but perhaps artificial more generally. While the above works give results on the spectra of graphical matrices, and tensors formed with tensor networks have been studied at length elsewhere, the spectra of CGM-like matrix “flattenings” of tensor networks remain poorly understood.² We develop some further tools for working with such objects in Section 10.13.*

Next, we specify the fairly simple class of diagrams whose CGSs will actually appear in our construction.

Definition 10.1.11 (Good forest). *We call a forest good if it has the following properties:*

1. *no vertex is isolated, and*

²One notable exception is the calculations with the trace method in the recent work [MW19].

2. the degree of every internal (non-leaf) vertex is even and at least 4.

We count the empty forest as a good forest. Denote by $\mathcal{F}(m)$ the set of good forests on m leaves, equipped with a labelling κ of the leaves by the set $[m]$. We consider two labelled forests equivalent if they are isomorphic as partially labelled graphs; thus, the same underlying forest may appear in $\mathcal{F}(m)$ with some but not all of the $m!$ ways that it could be labelled. For $F \in \mathcal{F}(m)$, we interpret F as a diagram by calling V^\bullet the leaves of F and calling V^\square the internal vertices of F . Finally, we denote by $\mathcal{T}(m)$ the subset of $F \in \mathcal{F}(m)$ that are connected (and therefore trees).

We note that, for m odd, the constraints imply that $\mathcal{F}(m)$ is empty. We give some examples of these forests and the associated CGSs in Figure 10.1.

Finally, we define the coefficients that are attached to each forest diagram's CGS in our construction.

Definition 10.1.12 (Möbius function of good forests). For $F = (V^\bullet \sqcup V^\square, E) \in \mathcal{F}(m)$, define

$$\mu(F) := \prod_{v \in V^\square} (-(\deg(v) - 2)!) = (-1)^{|V^\square|} \prod_{v \in V^\square} (\deg(v) - 2)!. \quad (10.24)$$

For F the empty forest, we set $\mu(F) = 1$ by convention.

These constants have an important interpretation in terms of the combinatorics of $\mathcal{F}(m)$: as we will show in Section 10.3, when $\mathcal{F}(m)$ is endowed with a natural partial ordering, $\mu(F)$ is (up to sign) the Möbius function of the “interval” of forests lying below F in this ordering. In general, Möbius functions encode the combinatorics of inclusion-exclusion calculations under a partial ordering [Rot64]. In our situation, $\mu(F)$ ensures that, even if we allow repeated indices in the monomial index S in the definition below, a suitable cancellation occurs such that the pseudoexpectation of \mathfrak{x}^S still approximately satisfies the ideal annihilation constraint in Definition 6.1.2.




Figure 10.1: Forests, polynomials, and coefficients. We show three examples of good forests $F \in \mathcal{F}(6)$ with labelled leaves, together with the corresponding CGS terms $\mu(F) \cdot Z^F(\mathbf{M}; (i, j, k, \ell, m, p))$ appearing in the pseudoexpectation of Definition 10.1.13.

With these ingredients defined, we are prepared to define our pseudoexpectation.

Definition 10.1.13 (Sum-of-forests pseudoexpectation). *Suppose $\mathbf{M} \in \mathbb{R}_{\text{sym}}^{n \times n}$. We define $\tilde{\mathbb{E}}_{\mathbf{M}} : \mathbb{R}[x_1, \dots, x_n] \rightarrow \mathbb{R}$ to be a linear operator with $\tilde{\mathbb{E}}_{\mathbf{M}}[x_i^2 p(\mathbf{x})] = \tilde{\mathbb{E}}_{\mathbf{M}}[p(\mathbf{x})]$ for all $i \in [n]$ and $p \in \mathbb{R}[x_1, \dots, x_n]$, and values on multilinear monomials given by*

$$\tilde{\mathbb{E}}_{\mathbf{M}} \left[\prod_{i \in S} x_i \right] := \sum_{F \in \mathcal{F}(|S|)} \mu(F) \cdot Z^F(\mathbf{M}; S) \text{ for all } S \subseteq [n]. \quad (10.25)$$

Remark 10.1.14 (Pseudocumulant generating function). *One interesting intuition for this construction is in terms of the formal cumulant generating function associated to $\tilde{\mathbb{E}}_{\mathbf{M}}$. If we approximate by collapsing all forests to “star trees” on their connected components (i.e., by contracting all edges between \square vertices), then, by the calculations we present later in Section 10.3, we find*

$$\log \tilde{\mathbb{E}}_{\mathbf{M}} \exp(\langle \boldsymbol{\lambda}, \mathbf{x} \rangle) \approx \frac{1}{2} \boldsymbol{\lambda}^\top \mathbf{M} \boldsymbol{\lambda} + \sum_{i=1}^n \left(\log \cosh((\mathbf{M} \boldsymbol{\lambda})_i) - \frac{1}{2} (\mathbf{M} \boldsymbol{\lambda})_i^2 \right), \quad (10.26)$$

viewing the left-hand side as a formal power series. Since $\log \cosh(\lambda)$ is the cumulant generating function of $x \sim \text{Unif}(\{\pm 1\})$, we see that $\tilde{\mathbb{E}}_{\mathbf{M}}$ encodes some compromise among three desiderata: \mathbf{x} lying in the row space of \mathbf{M} , \mathbf{x} having approximately independent ± 1 -valued coordinates, and \mathbf{x} having covariance matrix \mathbf{M} .

Let us now discuss how one arrives at this definition from continuing in the vein of Examples 10.1.7 and 10.1.8 from the previous section. From those examples, it is plausible that the degree d pseudomoments might in general be given by linear combinations of CGSs whose diagrams are forests. We may make the *ansatz* that this is the case, and attempt to derive a recursion describing the coefficients of these forests.

Reasoning diagrammatically, increasing the degree generates new diagrams whose CGSs occur in the pseudomoments in two ways, corresponding to the two terms in Lemma 8.3.11, whose recursion we recall here:

$$\tilde{\mathbb{E}}(\mathbf{x}^S, \mathbf{x}^T) = \underbrace{\tilde{\mathbb{E}}(r_S^\perp(\mathbf{x}), r_T^\perp(\mathbf{x}))}_{\text{“ideal” term}} + \underbrace{\sigma_d^2 \delta^d \cdot \langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle}_\circ}_{\text{“harmonic” term}}. \quad (10.27)$$

First, it turns out that all CGSs that arise from the inner product $\langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle_\circ$ in the harmonic term may be fully “collapsed” to a CGS with only one summation (or a product of summations over subsets of indices), as we have done above. These contribute diagrams that are forests of *stars*: each connected component is either two \bullet vertices connected to one another, or a single \square vertex connected to several leaves. Second, in computing the ideal term $\tilde{\mathbb{E}}(r_S^\perp(\mathbf{x}), r_T^\perp(\mathbf{x}))$, we join the diagrams of odd partition parts to the leaves of existing diagrams at lower degree, a process we illustrate later in Figure 10.3. Thus our *ansatz* is closed, in that the recursion—assuming the collapsing step detailed above is sound—will only yield forest diagrams at higher degree, since any forest is either a forest of only stars, or stars on some subsets of leaves with some subsets attached to a smaller forest.

Thus we indeed expect the pseudomoments to be a linear combination of CGSs whose diagrams are forests. Moreover, more careful parity considerations show that we expect only good forests (Definition 10.1.11) to appear. Taking this for granted, if the pseudomoments are to be symmetric under permutations of the indices, then the coefficients $\mu(F)$ should depend only on the unlabelled isomorphism type of the graph F , not on the leaf labels

κ . Making this assumption, each successive $\mu(F)$ may be expressed in a cumbersome but explicit combinatorial recursion, eventually letting us predict the formula for $\mu(F)$. This computation will be implicitly carried out in Section 10.7, where we prove that the $h_S^\dagger(\mathbf{x})$ give a block diagonalization of our pseudomoments.

We emphasize the pleasant interplay of diagrammatic and linear-algebraic ideas here. As we mentioned after Lemma 8.3.11, the decomposition of the pseudomoments into the ideal and harmonic parts expresses the spectral structure of the pseudomoment matrices, which involves a sequence of alternating “lift lower-degree pseudomoment matrix” and “add orthogonal harmonic part” steps. These correspond precisely to the sequence of alternating “compose partitions with old forests” and “add new star forests” steps generating good forests recursively.

Remark 10.1.15 (Setting σ_d^2). *We have glossed above the remaining detail of choosing the constants σ_d^2 to make $\tilde{\mathbb{E}}$ factor through multiplication. The further calculations discussed above confirm the pattern in the examples $d = 1, 2$ that the correct choice of this remaining scaling factor is $\sigma_d^2 := d! \delta^{-d}$. With this choice, we note that the harmonic term may be written more compactly as*

$$\sigma_d^2 \delta^d \cdot \langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle_\circ = \langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle_\partial, \quad (10.28)$$

where $\langle \cdot, \cdot \rangle_\partial$ is the rescaled apolar inner product from Definition 8.3.3, given simply by $\langle p, q \rangle_\partial = p(\partial)q$.

10.2 LIFTING 1: HIGH RANK TO HIGH DEGREE

We now proceed to our first lifting theorem. We first introduce several quantities measuring favorable behavior of M . As a high-level summary, these quantities capture various

aspects of the “incoherence” of M with respect to the standard basis vectors e_1, \dots, e_n . To formulate the subtlest of the incoherence quantities precisely, we will require the following preliminary technical definition, whose relevance will only become clear in the course of our proof in Section 10.5. There, it will describe a residual error term arising from allowing repeated indices in S in Definition 10.1.13, after certain cancellations are taken into account.

Definition 10.2.1 (Maximal repetition-spanning forest). *For each $F \in \mathcal{F}(m)$ and $s \in [n]^m$, let $\text{MaxSpan}(F, s)$ be the subgraph of F formed by the following procedure. Let C_1, \dots, C_k be the connected components of F .*

```

Initialize with  $\text{MaxSpan}(F, s) = \emptyset$ .
for  $i = 1, \dots, n$  do
  for  $j = 1, \dots, k$  do
    if  $C_j$  has two leaves  $\ell_1 \neq \ell_2$  with  $s_{\kappa(\ell_1)} = s_{\kappa(\ell_2)} = i$  then
      Let  $T$  be the minimal spanning tree of all leaves  $\ell$  of  $C_j$  with  $s_{\kappa(\ell)} = i$ .
      if  $T$  is vertex-disjoint from  $\text{MaxSpan}(F, s)$  then
        Add  $T$  to  $\text{MaxSpan}(F, s)$ .
      end if
    end if
  end for
end for

```

We say that $\mathbf{a} \in [n]^{V^\square(F)}$ is (F, s) -tight if, for all connected components C of $\text{MaxSpan}(F, s)$ with $s_{\kappa(\ell)} = i$ for all leaves ℓ of C , for all $v \in V^\square(C)$, $a_v = i$. Otherwise, we say that \mathbf{a} is (F, s) -loose.

With this, we define the following functions of M . Below, $M^{\circ k}$ denotes the k th entrywise power of M , and $\text{set}(s)$ for a tuple s denotes the set of indices occurring in s .

Definition 10.2.2 (Incoherence quantities). For $M \in \mathbb{R}_{\text{sym}}^{n \times n}$, define the following quantities:

$$\epsilon_{\text{offdiag}}(M) := \max_{1 \leq i < j \leq n} |M_{ij}|, \quad (10.29)$$

$$\epsilon_{\text{corr}}(M) := \left(\max_{1 \leq i < j \leq n} \sum_{k=1}^n M_{ik}^2 M_{jk}^2 \right)^{1/2}, \quad (10.30)$$

$$\epsilon_{\text{pow}}(M) := \max_{k \geq 2} \|M^{\circ k} - I_n\|, \quad (10.31)$$

$$\epsilon_{\text{tree}}(M; 2d) := \max_{0 \leq d' \leq d} \max_{T \in \mathcal{T}(2d')} \max_{s \in [n]^{2d'}} \left| Z^T(M; s) - \mathbb{1}\{s_1 = \dots = s_n\} \right|, \quad (10.32)$$

$$\epsilon_{\text{err}}(M; 2d) := \max_{0 \leq d' \leq d} \max_{T \in \mathcal{T}(2d')} \max_{s \in [n]^{2d'}} n^{|\text{set}(s)|/2} \left| \sum_{\substack{\alpha \in [n]^{\vee \square} \\ \alpha \text{ (T,s)-loose}}} \prod_{(v,w) \in E(T)} M_{f_{s,\alpha}(v)f_{s,\alpha}(w)} \right|, \quad (10.33)$$

$$\epsilon(M; 2d) := \epsilon_{\text{offdiag}}(M) + \epsilon_{\text{corr}}(M) + \epsilon_{\text{pow}}(M) + \epsilon_{\text{tree}}(M; d) + \epsilon_{\text{err}}(M; 2d). \quad (10.34)$$

Our main result then states that M may be extended to a high-degree pseudoexpectation so long as its smallest eigenvalue is not too small compared to the sum of the incoherence quantities.

Theorem 10.2.3. Let $M \in \mathbb{R}_{\text{sym}}^{n \times n}$ with $M_{ii} = 1$ for all $i \in [n]$. Suppose that

$$\lambda_{\min}(M) \geq (12d)^{32} \|M\|^5 \epsilon(M; 2d)^{1/d}. \quad (10.35)$$

Then, $\tilde{\mathbb{E}}_M$ is a degree $2d$ pseudoexpectation with $\tilde{\mathbb{E}}_M[\mathbf{x}\mathbf{x}^\top] = M$.

In practice, Theorem 10.2.3 will not be directly applicable to the M we wish to extend, which, as mentioned earlier, will be rank-deficient and therefore have $\lambda_{\min}(M) = 0$ (or very small). This obstacle is easily overcome by instead extending $M' = (1 - \alpha)M + \alpha I_n$ for $\alpha \in (0, 1)$ a small constant, whereby $\lambda_{\min}(M') \geq \alpha$. Unfortunately, it seems difficult to make a general statement about how the more intricate quantities ϵ_{tree} and ϵ_{err} transform when M is replaced with M' ; however, we will show in our applications that directly analyzing these quantities for M' is essentially no more difficult than analyzing them for M . Indeed,

we expect these to only become smaller under this replacement since M' equals M with the off-diagonal entries multiplied by $(1 - \alpha)$.

Remark 10.2.4 (Different ways of nudging). *A similar “nudging” operation to the one we propose above, moving M towards the identity matrix, has been used before in [MRX20, KB20] for degree 4 SOS and in the earlier work [AU03] for LP relaxations.³ However, the way that this adjustment propagates through our construction is quite different: while [MRX20, KB20] consider, in essence, a convex combination of the form $(1 - \alpha)\tilde{\mathbb{E}}_M + \alpha\tilde{\mathbb{E}}_{I_n}$, we instead consider $\tilde{\mathbb{E}}_{(1-\alpha)M+\alpha I_n}$. The mapping $M \mapsto \tilde{\mathbb{E}}_M$ is highly non-linear, so this is a major difference, which indeed turns out to be crucial for the adjustment to effectively counterbalance the error terms in our analysis.*

We expect the following general quantitative behavior from this result. Typically, we will have $\epsilon(M; 2d) = O(n^{-\gamma})$ for some $\gamma > 0$. We will also have $\|M\| = O(1)$ and $\lambda_{\min}(M) = \tilde{\Omega}(1)$ after the adjustment discussed above. Therefore, Theorem 10.2.3 will ensure that M is extensible to degree $2d$ so long as $n^{-\gamma/d} \text{poly}(d) = O(1)$, whereby the threshold scaling at which the condition of Theorem 10.2.3 is no longer satisfied is slightly smaller than $d \sim \log n$; for instance, such M will be extensible to degree $d \sim \log n / \log \log n$. See the brief discussion after Proposition 10.12.11 for an explanation of why this scaling of the degree is likely the best our proof techniques can achieve.

10.3 PARTIAL ORDERING AND MÖBIUS FUNCTION OF $\mathcal{F}(m)$

In the previous section, we found a way to compute the coefficients $\mu(F)$ attached to each forest diagram F in Definition 10.1.13. Calculating examples, one is led to conjecture the formula given in Definition 10.1.12 for these quantities. We will eventually give the rather difficult justification for that equality in the course of proving Theorem 10.2.3 in Section 10.7.

³I thank Aida Khajavirad for bringing the reference [AU03] to my attention.

For now, we prove the simpler interpretation of these constants that will guide other parts of that proof: the $\mu(F)$ give the Möbius function of a certain partial ordering on the CGS terms of the pseudoexpectation.

Before proceeding, let us introduce some notations and clarifications concerning partitions that will be useful here and in the remainder of the chapter.

Definition 10.3.1 (Partitions). *For A a set or multiset, we write $\text{Part}(A)$ for the set or multiset, respectively, of partitions of A . Repeated elements in a multiset are viewed as distinct for generating partitions, making $\text{Part}(A)$ a multiset when A is a multiset. For example,*

$$\text{Part}(\{i, i, j\}) = \left\{ \{\{i\}, \{i\}, \{j\}\}, \{\{i, i\}, \{j\}\}, \underbrace{\{\{i\}, \{i, j\}\}, \{\{i\}, \{i, j\}\}}_{\text{repeated}}, \{\{i, i, j\}\} \right\}. \quad (10.36)$$

We write $\text{Part}(A; \text{even})$ and $\text{Part}(A; \text{odd})$ for partitions into only even or odd parts, respectively, and $\text{Part}(A; k)$, $\text{Part}(A; \geq k)$, and $\text{Part}(A; \leq k)$ for partitions into parts of size exactly, at most, and at least k , respectively. We also allow these constraints to be chained, so that, e.g., $\text{Part}(A; \text{even}; \geq k)$ is the set of partitions into even parts of size at least k . Similarly, for a specific partition $\pi \in \text{Part}(A)$, we write $\pi[\text{even}]$, $\pi[\text{odd}]$, $\pi[k]$, $\pi[\geq k]$, $\pi[\leq k]$, $\pi[\text{even}; \geq k]$ and so forth for the parts of π with the specified properties.

10.3.1 MÖBIUS FUNCTIONS OF PARTIALLY ORDERED SETS

First, we review some basic concepts of the combinatorics of partially ordered sets (henceforth *posets*). Recall that a poset is a set \mathcal{P} equipped with a relation \leq that satisfies reflexivity ($x \leq x$ for all $x \in \mathcal{P}$), antisymmetry (if $x \leq y$ and $y \leq x$ then $x = y$), and transitivity (if $x \leq y$ and $y \leq z$, then $x \leq z$). For the purposes of this paper, we will assume all posets are finite. The following beautiful and vast generalization of the classical Möbius function of number theory was introduced by Rota in [Rot64] (the reference's introduction gives a more nuanced discussion of the historical context at the time).

Definition 10.3.2 (Poset Möbius function). *Let \mathcal{P} be a poset. Then, the Möbius function of \mathcal{P} , denoted $\mu_{\mathcal{P}}(x, y)$, is defined over all pairs $x \leq y$ by the relations*

$$\mu_{\mathcal{P}}(x, x) = 1, \tag{10.37}$$

$$\sum_{x \leq y \leq z} \mu_{\mathcal{P}}(x, y) = 0 \text{ for all } x < z. \tag{10.38}$$

The key consequence of this definition is the following general inclusion-exclusion principle over posets, again a vast generalization of both the Möbius inversion formula of number theory and the ordinary inclusion-exclusion principle over the poset of subsets of a set.

Proposition 10.3.3 (Poset Möbius inversion). *If \mathcal{P} has a minimal element, $f : \mathcal{P} \rightarrow \mathbb{R}$ is given, and $g(x) := \sum_{y \leq x} f(y)$, then $f(x) = \sum_{y \leq x} \mu_{\mathcal{P}}(y, x)g(y)$. Similarly, if \mathcal{P} has a maximal element and $g(x) := \sum_{y \geq x} f(y)$, then $f(x) = \sum_{y \geq x} \mu_{\mathcal{P}}(x, y)g(y)$.*

In addition to [Rot64], the reader may consult, e.g., [BG75] for some consequences of this result in enumerative combinatorics.

We give three examples of Möbius functions of posets of partitions that will be useful in our calculations. The first concerns subsets and corresponds to the classical inclusion-exclusion principle, and the latter two concern partitions of a set.

Example 10.3.4 (Subsets). *Give $2^{[m]}$ the poset structure of $S \leq T$ whenever $S \subseteq T$. Write $\mu_{\text{Subset}}(\cdot, \cdot)$ for the Möbius function of $[m]$. Then,*

$$\mu_{\text{Subset}}(S, T) = (-1)^{|T|-|S|}. \tag{10.39}$$

Example 10.3.5 (Partitions [Rot64]). *Let $\text{Part}([m])$ denote the poset of partitions of $[m]$, where $\pi \leq \rho$ whenever π is a refinement of ρ . Write $\mu_{\text{Part}}(\cdot, \cdot)$ for the Möbius function of*

$\text{Part}([m])$, eliding m for the sake of brevity. Then,

$$\mu_{\text{Part}}(\pi, \rho) = \prod_{A \in \rho} (-1)^{\#\{B \in \pi : B \subseteq A\} - 1} (\#\{B \in \pi : B \subseteq A\} - 1)!. \quad (10.40)$$

In particular, letting $\emptyset := \{\{1\}, \dots, \{m\}\}$ be the unique minimal element of $\text{Part}([m])$, we have

$$\mu_{\text{Part}}(\emptyset, \rho) = \prod_{A \in \rho} (-1)^{|A| - 1} (|A| - 1)!. \quad (10.41)$$

Example 10.3.6 (Partitions into even parts [Syl76]). For $m \geq 2$ even, let $\text{EvenPart}([m])$ denote the poset of partitions of $[m]$ into even parts, where $\pi \leq \rho$ whenever π is a refinement of ρ , along with the additional formal element \emptyset with $\emptyset \leq \pi$ for all partitions π . Again write $\mu_{\text{EvenPart}}(\cdot, \cdot)$ for the associated Möbius function, eliding m for the sake of brevity. Let the sequence $\nu(k)$ for $k \geq 0$ be defined by the exponential generating function $\log \cosh(x) =: \sum_{k=0}^{\infty} \frac{\nu(k)}{k!} x^k$, or equivalently $\tanh(x) =: \sum_{k=0}^{\infty} \frac{\nu(k+1)}{k!} x^k$. Then,

$$\mu_{\text{EvenPart}}(\emptyset, \rho) = - \prod_{A \in \rho} \nu(|A|). \quad (10.42)$$

On the other hand, if $\pi > \emptyset$, the $[\pi, \rho]$ is isomorphic to a poset of ordinary partitions, so we recover

$$\mu_{\text{EvenPart}}(\pi, \rho) = \prod_{A \in \rho} (-1)^{\#\{B \in \pi : B \subseteq A\} - 1} (\#\{B \in \pi : B \subseteq A\} - 1)!. \quad (10.43)$$

There is no convenient closed form for $\nu(k)$, but a combinatorial interpretation (up to sign) is given by $(-1)^k \nu(2k)$ counting the number of alternating permutations of $2k + 1$ elements. This fact, as a generating function identity, is a classical result due to André [And81] who used it to derive the asymptotics of ν ; see also [Sta10] for a survey. The connection with Möbius functions was first observed in Sylvester's thesis [Syl76], and Stanley's subsequent work [Sta78] explored further situations where the Möbius function of a poset is given by an exponential generating function. Some of our calculations in Section 10.3 indicate that the

poset defined there, while not one of Stanley’s “exponential structures,” is still amenable to analysis via exponential generating functions, suggesting that the results of [Sta78] might be generalized to posets having more general self-similarity properties.

10.3.2 THE COMPOSITIONAL ORDERING

We first introduce the poset structure that is associated with $\mu(F)$. We call this a *compositional* structure because it is organized according to which forests are obtained by “composing” one forest with another by inserting smaller forests at each \square vertex.

Definition 10.3.7 (Compositional poset). *Suppose $F \in \mathcal{F}(m)$. For each $v \in V^\square(F)$, write $E(v)$ for the set of edges incident with v , and fix $\kappa_v : E(v) \rightarrow [|E(v)|]$ a labelling of $E(v)$. Suppose that for each $v \in V^\square(F)$, we are given $F_v \in \mathcal{F}(\deg(v))$. Write $F[(F_v)_{v \in V^\square(F)}] \in \mathcal{F}(m)$ for the forest formed as follows. Begin with the disjoint union of all F_v for $v \in V^\square(F)$ and all pairs in F . Denote the leaves of F_v in this disjoint union by $\ell_{v,1}, \dots, \ell_{v,\deg(v)}$. Then, merge the edges ending at $\ell_{v,i}$ and $\ell_{w,j}$ whenever $\kappa_v^{-1}(i) = \kappa_w^{-1}(j)$. Whenever $\kappa_v^{-1}(i)$ terminates in a leaf x of F , give $\ell_{v,i}$ the label that x has in F . Finally, whenever x belongs to a pair of F , give x in the disjoint union the same label that it has in F .*

Let the compositional relation \leq on $\mathcal{F}(m)$ be defined by setting $F' \leq F$ if, for each $v \in V^\square(F)$, there exists $F_v \in \mathcal{F}(\deg(v))$ such that $F' = F[(F_v)_{v \in V^\square(F)}]$.

It is straightforward to check that this relation does not depend on the auxiliary orderings κ_v used in the definition. While the notation used to describe the compositional relation above is somewhat heavy, we emphasize that it is conceptually quite intuitive, and give an illustration in Figure 10.2.

We give the following additional definition before continuing to the basic properties of the resulting poset.

Definition 10.3.8 (Star tree). *For $m \geq 4$ an even number, we denote by S_m the star tree on m*




Figure 10.2: Ordering in the compositional poset $\mathcal{F}(m)$. We give an example of the ordering relation between two forests in $\mathcal{F}(14)$, highlighting the “composing forests” at each \square vertex of the greater forest that witness this relation.

leaves, consisting of a single \square vertex connected to $m \bullet$ vertices. For $m = 2$, denote by S_2 the tree with no \square vertices and two \bullet vertices connected to one another. Note that all labellings of S_m are isomorphic, so there is a unique labelled star tree in $\mathcal{F}(m)$.

Proposition 10.3.9. $\mathcal{F}(m)$ endowed with the relation \leq forms a poset. The unique maximal element in $\mathcal{F}(m)$ is S_m , while any perfect matching in $\mathcal{F}(m)$ is a minimal element.

To work with the Möbius function, it will be more convenient to define a version of this poset augmented with a unique minimal element, as follows (this is the same manipulation as is convenient to use, for example, in the analysis of the poset of partitions into sets of even size; see [Sta78]).

Definition 10.3.10. Let $\bar{\mathcal{F}}(m)$ consist of $\mathcal{F}(m)$ with an additional formal element denoted \emptyset . We extend the poset structure of $\mathcal{F}(m)$ to $\bar{\mathcal{F}}(m)$ by setting $\emptyset \leq F$ for all $F \in \bar{\mathcal{F}}(m)$. When we wish to distinguish \emptyset from the elements of $\mathcal{F}(m)$, we call the latter proper forests.

The main result of this section obtains the Möbius function of this partial ordering.

Lemma 10.3.11. Let $\mu_{\bar{\mathcal{F}}}(\cdot, \cdot)$ be the Möbius function of $\bar{\mathcal{F}}(m)$ (where we elide m for the sake of brevity, as it is implied by the arguments). Then $\mu_{\bar{\mathcal{F}}}(\emptyset, \emptyset) = 1$, and for $F \in \mathcal{F}(m)$,

$$\mu_{\bar{\mathcal{F}}}(\emptyset, F) = (-1)^{|V^\square(F)|+1} \prod_{v \in V^\square(F)} (\deg(v) - 2)! = -\mu(F), \quad (10.44)$$

where $\mu(\cdot)$ on the right-hand side is the quantity from Definition 10.1.12.

We proceed in two steps: first, and what is the main part of the argument, we compute the Möbius function of a star tree. Then, we show that the Möbius function of a general forest factorizes into that of the star trees corresponding to each of its internal vertices. The following ancillary definition will be useful both here and in a later proof.

Definition 10.3.12 (Rooted odd tree). *For odd $\ell \geq 3$, define the set of rooted odd trees on ℓ leaves, denoted $\mathcal{T}^{\text{root}}(\ell)$, to be the set of rooted trees where the number of children of each internal vertex is odd and at least 3, and where the leaves are labelled by $[\ell]$. Define a map $e : \mathcal{T}^{\text{root}}(m-1) \rightarrow \mathcal{T}(m)$ that attaches the leaf labelled m to the root.*

While it is formally easier to express this definition in terms of rooted trees, it may be intuitively clearer to think of a rooted odd tree as still being a good tree, only having one distinguished “stub” leaf, whose lone neighbor is viewed as the root.

Proposition 10.3.13. $\mu_{\bar{f}}(\emptyset, S_2) = -1$. For all even $m \geq 4$, $\mu_{\bar{f}}(\emptyset, S_m) = (m-2)!$.

Proof. We first establish the following preliminary identity.

$$\text{Claim: } \sum_{T \in \mathcal{T}(m)} \mu_{\bar{f}}(\emptyset, T) = -\nu(m) = \mu_{\text{EvenPart}}(\emptyset, \{[m]\}). \quad (10.45)$$

We proceed using a common idiom of Möbius inversion arguments, similar to, e.g., counting labelled connected graphs (see Example 2 in Section 4 of [BG75]). For $F \in \mathcal{F}(m)$, let $\text{conn}(F) \in \text{Part}([m]; \text{even})$ denote the partition of leaves into those belonging to each connected component of F . For $\pi \in \text{Part}([m]; \text{even})$, define

$$b(\pi) := - \sum_{\substack{F \in \mathcal{F}([m]) \\ \text{conn}(F) = \pi}} \mu(\emptyset, F), \quad (10.46)$$

and $b(\emptyset) = 0$. Then, the quantity we are interested in is $-b(\{[m]\})$. By Möbius inversion,

$$b(\{[m]\}) = \sum_{\pi \in \text{Part}(m; \text{even}) \cup \{\emptyset\}} \left(\sum_{\emptyset \leq \rho \leq \pi} b(\rho) \right) \mu_{\text{EvenPart}}(\pi, \{[m]\}). \quad (10.47)$$

The inner summation is zero if $\pi = \emptyset$, and otherwise equals

$$\sum_{\emptyset \leq \rho \leq \pi} b(\rho) = \prod_{A \in \pi} \left(\sum_{F \in \mathcal{F}(|A|)} -\mu(F) \right) = 1 \text{ if } \pi \neq \emptyset. \quad (10.48)$$

Therefore, we may continue

$$b(\{[m]\}) = \sum_{\pi \in \text{Part}(m; \text{even})} \mu_{\text{EvenPart}}(\pi, \{[m]\}) = \sum_{\pi \in \text{Part}(m; \text{even})} (-1)^{|\pi|-1} (|\pi| - 1)!. \quad (10.49)$$

By the composition formula for exponential generating functions, this means

$$\sum_{k \geq 0} \frac{b(\{[2k]\})}{(2k)!} x^{2k} = \log(1 + (\cosh(x) - 1)) = \log \cosh(x), \quad (10.50)$$

and the result follows by equating coefficients.

Next, we relate the trees of $\mathcal{T}([m])$ that we sum over in this identity to the rooted odd trees introduced in Definition 10.3.12. We note that the map e defined there is a bijection between $\mathcal{T}^{\text{root}}(m - 1)$ and $\mathcal{T}(m)$ (the inverse map removes the leaf labelled m and sets its single neighbor to be the root). The Möbius function composed with this bijection is

$$\mu(\emptyset, e(T)) = (-1)^{|V^\square(T)|+1} \prod_{v \in V^\square(T)} \mu(\emptyset, S_{|c(v)|+1}),$$

where $c(v)$ gives the number of children of an internal vertex.

Finally, we combine the recursion associated to the rooted structure of $\mathcal{T}^{\text{root}}(m - 1)$ (whereby a rooted tree is, recursively, the root and a collection of rooted trees attached to the root) and the identity of the Claim to derive a generating function identity that completes

the proof. Namely, we may manipulate, for $m \geq 4$,

$$\begin{aligned}
\nu(m) &= - \sum_{T \in \mathcal{T}(m)} \mu(\emptyset, T) \\
&= - \sum_{T \in \mathcal{T}^{\text{root}}(m-1)} \mu(\emptyset, e(T)) \\
&= \sum_{T \in \mathcal{T}^{\text{root}}(m-1)} (-1)^{|V^\square(T)|} \prod_{v \in V^\square(G)} \mu(\emptyset, S_{|c(v)|+1}) \\
&= \sum_{\substack{\pi \in \text{Part}([m-1]; \text{odd}) \\ |\pi| > 1}} (-\mu(\emptyset, S_{|\pi|+1})) \prod_{S \in \pi} \left(\sum_{T \in \mathcal{T}^{\text{root}}(S)} (-1)^{|V^\square(T)|} \prod_{v \in V^\square(T)} \mu(\emptyset, S_{|c(v)|+1}) \right) \\
&= \sum_{\substack{\pi \in \text{Part}([m-1]; \text{odd}) \\ |\pi| > 1}} (-\mu(\emptyset, S_{|\pi|+1})) \prod_{S \in \pi} \left(- \sum_{T \in \mathcal{T}^{\text{root}}(|S|)} \mu(\emptyset, e(T)) \right) \\
&= \sum_{\substack{\pi \in \text{Part}([m-1]; \text{odd}) \\ |\pi| > 1}} (-\mu(\emptyset, S_{|\pi|+1})) \prod_{S \in \pi} \nu(|S| + 1) \tag{10.51}
\end{aligned}$$

We now have a relatively simple identity connecting $\mu(\emptyset, S_m)$ with $\nu(m)$. To translate this into a relation of generating functions, we remove the condition $|\pi| > 1$ and correct to account for the case $m = 2$, obtaining, for any even $m \geq 2$,

$$2\nu(m) = \mathbb{1}\{m = 2\} + \sum_{\pi \in \text{Part}([m-1]; \text{odd})} (-\mu(\emptyset, S_{|\pi|+1})) \prod_{S \in \pi} \nu(|S| + 1). \tag{10.52}$$

Now, let $F(x) := \sum_{k \geq 1} \frac{\mu(\emptyset, S_{2k})}{(2k)!} x^{2k}$. Multiplying by $x^{2k-1}/(2k-1)!$ on either side of (10.52) and summing over all $k \geq 1$, we find, by the composition formula for exponential generating functions,

$$2 \tanh(x) = x - F'(\tanh(x)). \tag{10.53}$$

Equivalently, taking $y = \tanh(x)$, we have

$$F'(y) = \tanh^{-1}(y) - 2y. \tag{10.54}$$

Recalling

$$\tanh^{-1}(y) = \frac{1}{2} (\log(1 + y) - \log(1 - y)) = \sum_{k \geq 0} \frac{y^{2k+1}}{2k+1}, \quad (10.55)$$

we have

$$F(y) = -\frac{1}{2}x^2 + \sum_{k \geq 2} \frac{x^{2k}}{2k(2k-1)} = -\frac{1}{2!}x^2 + \sum_{k \geq 2} \frac{(2k-2)!}{(2k)!}x^{2k}, \quad (10.56)$$

and the result follows. \square

Before completing the proof of Lemma 10.3.11, we give the following preliminary result, describing the interval lying below a forest as a product poset. This follows immediately from the definition of the compositional relation, since the set of forests smaller than F corresponds to a choice of a local “composing forest” at each $v \in V^\square(F)$.

Proposition 10.3.14. *Let $F = ((V^\bullet, V^\square), E) \in \mathcal{F}(m)$. Then, we have the isomorphism of posets*

$$(\emptyset, F] \cong \prod_{v \in V^\square(F)} (\emptyset, S_{\deg(v)}]. \quad (10.57)$$

We now complete the proof of the main Lemma.

Proof of Lemma 10.3.11. Let $\hat{\mu}(\emptyset, \cdot)$ be the putative Möbius function from the statement,

$$\hat{\mu}(\emptyset, F) = (-1)^{|V^\square(F)|+1} \prod_{v \in V^\square(F)} (\deg(v) - 2)!. \quad (10.58)$$

We proceed by induction on m . For $m = 2$, the result holds by inspection. Suppose $m \geq 4$. Since $\hat{\mu}(\emptyset, \emptyset) = 1$ by definition, and since by Proposition 10.3.13 we know that $\hat{\mu}(\emptyset, S_m) = \mu(\emptyset, S_m)$, it suffices to show that, for all $F \in \mathcal{F}(m)$ with $F \neq S_m$, we have

$$\sum_{F' \in [\emptyset, F]} \hat{\mu}(\emptyset, F') = 0. \quad (10.59)$$

Let $F \in \mathcal{F}(m)$ with $F \neq S_m$. We then compute:

$$\begin{aligned}
\sum_{F' \in [\emptyset, F]} \hat{\mu}(\emptyset, F') &= 1 - \sum_{F' \in (\emptyset, F]} \prod_{v \in V^\square(F')} (-(\deg(v) - 2)!) \\
&= 1 - \prod_{v \in V^\square(F)} \left(\sum_{F' \in (\emptyset, S_{\deg(v)}]} \prod_{w \in V^\square(F')} (-(\deg(w) - 2)!) \right) \quad (\text{Proposition 10.3.14}) \\
&= 1 - \prod_{v \in V^\square(F)} \left(- \sum_{F' \in (\emptyset, S_{\deg(v)}]} \mu(\emptyset, F') \right) \\
&= 1 - \prod_{v \in V^\square(F)} \mu(\emptyset, \emptyset) \quad (\text{inductive hypothesis}) \\
&= 0, \tag{10.60}
\end{aligned}$$

completing the proof. □

10.4 PSEUDOMOMENT AND CONTRACTIVE GRAPHICAL MATRICES

We now proceed to the proof of Theorem 10.2.3. We first outline the general approach of our proof and introduce the main objects involved. By construction, $\tilde{\mathbb{E}}$ as given in Definition 10.1.13 satisfies Conditions 1 and 2 of the pseudoexpectation properties from Definition 6.1.2 (normalization and ideal annihilation); therefore, it suffices to prove positivity. Moreover, positivity may be considered in any suitable basis modulo the ideal generated by the constraint polynomials, and given any fixed basis positivity may be written in linear-algebraic terms as the positive semidefiniteness of the associated *pseudomoment matrix*. We state this explicitly below, in an application of standard reasoning in the SOS literature (see, e.g., [Lau09]).

Proposition 10.4.1. *Let $\tilde{\mathbb{E}} : \mathbb{R}[x_1, \dots, x_n]_{\leq 2d} \rightarrow \mathbb{R}$ be a linear operator satisfying the normalization and ideal annihilation properties of Definition 6.1.2. Let $\mathcal{I} \subset \mathbb{R}[x_1, \dots, x_n]$ be the ideal generated by $x_i^2 - 1$ for $i = 1, \dots, n$, and let $p_1, \dots, p_{\binom{n}{\leq d}} \in \mathbb{R}[x_1, \dots, x_n]_{\leq d}$ be a collection of*

coset representatives for a basis of $\mathbb{R}[x_1, \dots, x_n]_{\leq d} / \mathcal{I}$. Define the associated pseudomoment matrix $\mathbf{Z} \in \mathbb{R}^{\binom{n}{\leq d} \times \binom{n}{\leq d}}$ with entries

$$Z_{s,t} = \tilde{\mathbb{E}}[p_s(\mathbf{x})p_t(\mathbf{x})]. \quad (10.61)$$

Then, $\tilde{\mathbb{E}}$ satisfies the positivity property of Definition 6.1.2 if and only if $\mathbf{Z} \succeq \mathbf{0}$.

If we were to take the standard multilinear monomial basis for the $p_s(\mathbf{x})$, we would wind up with \mathbf{Z} being a sum of CGSs of different diagrams in each entry, with the CGS indices corresponding to the set indexing of \mathbf{Z} . While we will ultimately work in a different basis, this general observation will still hold, so we define the following broad formalism for the matrices that will arise.

The following enhancement of the diagrams introduced in Definition 10.1.9 is the analogous object to what is called a *shape* in the literature on graphical matrices [AMP16, BHK⁺19]. We prefer to reserve the term *diagram* for any object specifying some contractive calculation, to use that term unadorned for the scalar version, and to add *ribbon* to indicate the specification of “sidedness” that induces a matrix structure.

Definition 10.4.2 (Ribbon diagram). *Suppose $G = (V, E)$ is a graph with two types of vertices, which we denote \bullet and \square visually and whose subsets we denote $V = V^\bullet \sqcup V^\square$. Suppose also that V^\bullet is further partitioned into two subsets, which we call “left” and “right” and denote $V^\bullet = \mathcal{L} \sqcup \mathcal{R}$. Finally, suppose that each of \mathcal{L} and \mathcal{R} is equipped with a labelling $\kappa_{\mathcal{L}} : \mathcal{L} \rightarrow [|\mathcal{L}|]$ and $\kappa_{\mathcal{R}} : \mathcal{R} \rightarrow [|\mathcal{R}|]$. We call such G together with the labellings $\kappa_{\mathcal{L}}$ and $\kappa_{\mathcal{R}}$ a ribbon diagram.*

Definition 10.4.3 (Good forest ribbon diagram). *We write $\mathcal{F}(\ell, m)$ for the set of good forests on $\ell + m$ vertices, equipped with a partition of the leaves $V^\bullet = \mathcal{L} \sqcup \mathcal{R}$ with $|\mathcal{L}| = \ell$ and $|\mathcal{R}| = m$.*

Definition 10.4.4 (Contractive graphical matrix). *Suppose G is a ribbon diagram with labellings $\kappa_{\mathcal{L}}$ and $\kappa_{\mathcal{R}}$. Define $\kappa : V^\bullet \rightarrow [|V^\bullet|]$ by $\kappa(\ell) = \kappa_{\mathcal{L}}(\ell)$ for $\ell \in \mathcal{L}$ and $\kappa(r) = |\mathcal{L}| + \kappa_{\mathcal{R}}(r)$ for $r \in \mathcal{R}$. With this labelling, we interpret G as a CGS diagram.*

For $M \in \mathbb{R}_{\text{sym}}^{n \times n}$, we then define the contractive graphical matrix (CGM) of G to be the matrix $Z^G \in \mathbb{R}^{\binom{[n]}{|\mathcal{L}|} \times \binom{[n]}{|\mathcal{R}|}}$ with entries

$$Z_{S,T}^G = Z_{S,T}^G(\mathbf{M}) := Z^G(\mathbf{M}; (s_1, \dots, s_{|\mathcal{L}|}, t_1, \dots, t_{|\mathcal{R}|})) \quad (10.62)$$

where $S = \{s_1, \dots, s_{|\mathcal{L}|}\}$ and $T = \{t_1, \dots, t_{|\mathcal{R}|}\}$ with $s_1 < \dots < s_{|\mathcal{L}|}$ and $t_1 < \dots < t_{|\mathcal{R}|}$.

We note that the restriction to set-valued indices in this definition is rather artificial; the most natural indexing would be by $[n]^{|\mathcal{L}|} \times [n]^{|\mathcal{R}|}$. However, as the set-indexed submatrix of this larger matrix is most relevant for our application, we use this definition in the main text; we present several technical results with the more general tuple-indexed CGMs in Section 10.13.

Remark 10.4.5 (Multiscale spectrum). *As in calculations involving graphical matrices in the pseudocalibration approach [AMP16, RSS18, BHK⁺19], the scale of the norm of a CGM may be read off of its ribbon diagram. We emphasize the following general principle: if $\|M\| = O(1)$ and $F \in \mathcal{F}(2d)$, then $\|Z^F\| = \omega(1)$ **if and only if some connected components of F have leaves in only \mathcal{L} or only \mathcal{R}** . We call such components sided. CGMs tensorize over connected components (Proposition 10.13.3), so the norm of a CGM is the product of the norms of CGMs of its diagram's components. In the case of M a rescaled random low-rank projection matrix, where $\|M\| = O(1)$ and $\|M\|_F = \Theta(n^{1/2})$, components that are not sided give norm $O(1)$ (Proposition 10.13.13), while each sided component gives norm roughly $\tilde{\Theta}(n^{1/2})$, which follows from calculating the sizes of the individual CGM entries assuming square root cancellations. Thus the norm of a CGM is $\tilde{\Theta}(n^{\#\{\text{sided components}\}/2})$.*

In particular, we will encounter the same difficulty as in other SOS lower bounds that the

pseudomoment matrices we work with have a multiscale spectrum, meaning simply that the scaling of different Z^G with n can be very different. For the main part of our analysis we will be able to ignore this, since by working in the multiharmonic basis from Remark 8.3.12, we will be able to eliminate all ribbon diagrams with sided components, leaving us with only terms of norm $O(1)$. Unfortunately, this issue returns when handling various error terms, so some of our choices below will still be motivated by handling the multiscale difficulty correctly.

10.5 MAIN AND ERROR TERMS

Recall that our pseudoexpectation was constructed in Section 10.1.3 as a sum of $Z^F(M; S)$ for S a *multiset*, and had the property of being approximately unchanged by adding pairs of repeated indices to S . While in Definition 10.1.13 we have forced these to be exact equalities to produce a pseudoexpectation satisfying the ideal annihilation constraints exactly, the approximate version of the pseudoexpectation, which is better suited for the diagrammatic reasoning that will justify positivity, will still be very useful. Therefore, we decompose $\tilde{\mathbb{E}}$ into a “main term,” which is the actual result of our heuristic calculation but only approximately satisfies the hypercube constraints, and an “error term” that implements the remaining correction, as follows.

Definition 10.5.1 (Main and error pseudoexpectations). *Define $\tilde{\mathbb{E}}^{\text{main}}, \tilde{\mathbb{E}}^{\text{err}} : \mathbb{R}[x_1, \dots, x_n] \rightarrow \mathbb{R}$ to be linear operators with values on monomials given by*

$$\tilde{\mathbb{E}}^{\text{main}}[\mathbf{x}^S] := \sum_{F \in \mathcal{F}([S])} \mu(F) \cdot Z^F(M; S), \quad (10.63)$$

$$\tilde{\mathbb{E}}^{\text{err}}[\mathbf{x}^S] := \tilde{\mathbb{E}}[\mathbf{x}^S] - \tilde{\mathbb{E}}^{\text{main}}[\mathbf{x}^S], \quad (10.64)$$

for all multisets $S \in \mathcal{M}([n])$. Note that for S a multiset, $\tilde{\mathbb{E}}[\mathbf{x}^S] = \tilde{\mathbb{E}}[\mathbf{x}^{S'}]$ where S' is the (non-multi) set of indices occurring an odd number of times in S .

In the remainder of this section, we show how the presence of the Möbius function in our pseudomoment values implies that $\tilde{\mathbb{E}}^{\text{err}}[\mathbf{x}^S]$ is small. It is not difficult to see that $\tilde{\mathbb{E}}$ and $\tilde{\mathbb{E}}^{\text{main}}$ are equal to leading order, since if, for instance, only one index is repeated two times in S , then the dominant terms of $\tilde{\mathbb{E}}^{\text{main}}$ will be those from diagrams where the two occurrences of this index are paired and there is an arbitrary forest on the remaining indices; various generalizations thereof hold for greater even and odd numbers of repetitions. This kind of argument shows, for example, that for M a rescaled random low-rank projection matrix, we have $\tilde{\mathbb{E}}^{\text{main}}[\mathbf{x}^S] = (1 + O(n^{-1/2}))\tilde{\mathbb{E}}[\mathbf{x}^S]$ as $n \rightarrow \infty$. However, due to the multiscale spectrum of the pseudomoments as discussed in Remark 10.4.5, it turns out that this does not give sufficient control of $\tilde{\mathbb{E}}^{\text{err}}$.

We must go further than this initial analysis and take advantage of cancellations among even the sub-leading order terms of $\tilde{\mathbb{E}}^{\text{err}}$, a fortunate side effect of the Möbius function coefficients. These cancellations generalize the following observation used in [KB20] for the degree 4 case. If we take $S = \{i, i, j, k\}$ (the simple situation mentioned above), then we have

$$\begin{aligned} \tilde{\mathbb{E}}^{\text{err}}[\mathbf{x}^S] &= \underbrace{M_{jk}}_{\tilde{\mathbb{E}}[\mathbf{x}^S]} - \underbrace{\left(M_{ii}M_{jk} + M_{ij}M_{ik} + M_{ik}M_{ij} - 2 \sum_{a=1}^n M_{ai}^2 M_{aj} M_{ak} \right)}_{\tilde{\mathbb{E}}^{\text{main}}[\mathbf{x}^S]} \\ &= 2 \sum_{a \in [n] \setminus \{i\}} M_{ai}^2 M_{aj} M_{ak}, \end{aligned} \tag{10.65}$$

where the term $2M_{ij}M_{ik}$ in $\tilde{\mathbb{E}}^{\text{main}}[\mathbf{x}^S]$ has cancelled. For M a rescaled random low-rank projection matrix, this makes a significant difference: the term that cancels is $\Theta(n^{-1})$, while the remaining error term after the cancellation is only $\Theta(n^{-3/2})$ (assuming square root cancellations).

Surprisingly, a similar cancellation obtains at all degrees and for any combination of repeated indices. The general character of the remaining error terms is that, as in the above simple example the \square vertex connecting two equal leaves labelled i was not allowed to have

its index equal i , so in general the minimal spanning subtree of a collection of leaves with the same label cannot “collapse” by having all of its internal vertices have that same label.

The collections of spanning subtrees with respect to which we will study this cancellation are precisely the forests $\text{MaxSpan}(F, \mathbf{s})$, as defined earlier in Definition 10.2.1. Below we record the important properties of the subgraphs that result from this construction.

Proposition 10.5.2 (Properties of MaxSpan). *For any $F \in \mathcal{F}(m)$ and $\mathbf{s} \in [n]^m$, $\text{MaxSpan}(F, \mathbf{s})$ satisfies the following.*

1. *(Components) For every connected component C of $\text{MaxSpan}(F, \mathbf{s})$, there is some $i \in [n]$ and C_j a connected component of F such that $|\kappa^{-1}(i) \cap V^\bullet(C_j)| \geq 2$ and C is the minimal spanning tree of $\kappa^{-1}(i) \cap V^\bullet(C_j)$.*
2. *(Maximality) $\text{MaxSpan}(F, \mathbf{s})$ is the union of a maximal collection of vertex-disjoint spanning trees of the above kind.*
3. *(Independence over connected components) If C_1, \dots, C_k are the connected components of F , then $\text{MaxSpan}(F, \mathbf{s}) = \text{MaxSpan}(C_1, \mathbf{s}|_{C_1}) \sqcup \dots \sqcup \text{MaxSpan}(C_k, \mathbf{s}|_{C_k})$. (We write $\mathbf{s}|_{C_i}$ for the restriction of \mathbf{s} to the indices that appear as labels of the leaves of C_i .)*
4. *(Priority of small indices) Whenever $i < j$, $|\kappa^{-1}(i) \cap V^\bullet(C_k)| \geq 2$, $|\kappa^{-1}(j) \cap V^\bullet(C_k)| \geq 2$, and $\text{MaxSpan}(F, \mathbf{s})$ contains the minimal spanning tree of $\kappa^{-1}(j) \cap V^\bullet(C_k)$, then it also contains the minimal spanning tree of $\kappa^{-1}(i) \cap V^\bullet(C_k)$.*

We are now prepared to express our generalization of the cancellation that we observed above in (10.65), which amounts to the cancellation of all summation terms where the entire subgraph $\text{MaxSpan}(F, \mathbf{s})$ collapses in the sense discussed previously.

Definition 10.5.3 (Graphical error terms). *Let $F \in \mathcal{F}(m)$ and $\mathbf{s} \in [n]^m$. Recall that we say $\mathbf{a} \in [n]^{V^\square(F)}$ is (F, \mathbf{s}) -tight if, for all connected components C of $\text{MaxSpan}(F, \mathbf{s})$, if $s_{\kappa(x)} = i$*

for all leaves x of C , then $a_v = i$ for all $v \in V^\square(C)$ as well. Otherwise, we say that \mathbf{a} is (F, \mathbf{s}) -loose. With this, we define

$$\Delta^F(\mathbf{M}; \mathbf{s}) := \sum_{\substack{\mathbf{a} \in [n]^{V^\square} \\ \mathbf{a} \text{ is } (F, \mathbf{s})\text{-loose}}} \prod_{\{v, w\} \in E} M_{f_{\mathbf{s}, \mathbf{a}}(v), f_{\mathbf{s}, \mathbf{a}}(w)}. \quad (10.66)$$

As in Definition 10.1.9, we also extend the definition to allow sets or multisets in the second argument of Δ^F by replacing them with the corresponding tuple of elements in ascending order.

The following preliminary definition, building on the rooted odd trees introduced in Definition 10.3.12, will be useful in the argument.

Definition 10.5.4 (Good forest with rooted components). For $m \geq 2$, let $\mathcal{F}^{\text{root}}(m)$ be the set of forests on m leaves where every connected component is either a good tree (per Definition 10.1.11) or a rooted odd tree (per Definition 10.3.12), and where the leaves are labelled by $[m]$. Note that some but not all components of such a forest may have distinguished roots. For $F \in \mathcal{F}^{\text{root}}(m)$, let $\text{odd}(F)$ denote the set of rooted odd tree components of F , and let $\mu(F) := \mu(F')$ for F' formed by attaching an extra leaf to the root of every tree in $\text{odd}(F)$.

Lemma 10.5.5 (Graphical error pseudomoments). For any $S \in \mathcal{M}([n])$,

$$\tilde{\mathbb{E}}^{\text{err}}[\mathbf{x}^S] = - \sum_{F \in \mathcal{F}(m)} \mu(F) \cdot \Delta^F(\mathbf{M}; S). \quad (10.67)$$

Proof. Our result will follow from the following, purely combinatorial, result. For $F \in \mathcal{F}(m)$ and $A \subseteq [m]$, let us say that F is A -dominated if, for every connected component C of F , every \square vertex of C is contained in the minimal spanning tree of the leaves $\kappa^{-1}(A) \cap V^\bullet(C)$.

Then,

$$\text{Claim: } \sum_{\substack{F \in \mathcal{F}(m) \\ F \text{ is } A\text{-dominated}}} \mu(F) = \begin{cases} 1 & \text{if } m \in \{|A|, |A| + 1\}, \\ 0 & \text{otherwise.} \end{cases} \quad (10.68)$$

We first prove the Claim. Let $\ell = |A|$; without loss of generality we take $A = [\ell]$. Let us write

$$c(\ell, m) := \sum_{\substack{F \in \mathcal{F}(m) \\ F \text{ is } [\ell]\text{-dominated}}} \mu(F). \quad (10.69)$$

For each fixed ℓ , we will proceed by induction on $m \geq \ell$. For the base case, we have $c(\ell, \ell) = 1$ by the defining property of the Möbius function, since in this case the summation is over all $F \in \mathcal{F}(\ell)$.

Let $r_\ell : \mathcal{F}(m) \rightarrow \mathcal{F}^{\text{root}}(m - \ell)$ return the rooted forest formed by deleting the minimal spanning trees of the elements of $[\ell]$ in each connected component, where upon deleting part of a tree, we set any vertex with a deleted neighbor to be the root of the new odd tree connected component thereby formed. Then, we have

$$\begin{aligned} 1 &= \sum_{F \in \mathcal{F}(m)} \mu(F) \\ &= \sum_{R \in \mathcal{F}^{\text{root}}(m-\ell)} \sum_{\substack{F \in \mathcal{F}(m) \\ r_\ell(F) = R}} \mu(F) \end{aligned}$$

and, factoring out $\mu(R)$ from $\mu(F)$ with $r_\ell(F) = R$, we note that what is left is a sum of $\mu(F)$ over $[\ell]$ -dominated forests F on $[\ell + |\text{odd}(R)|]$, whereby

$$= \sum_{R \in \mathcal{F}^{\text{root}}(m-\ell)} \mu(R) c(\ell, \ell + |\text{odd}(R)|). \quad (10.70)$$

Now, we consider two cases. First, if $m = \ell + 1$ for ℓ odd, then there is only one R in the above summation, having one leaf connected to a root, which has $\mu(R) = 1$. Therefore,

$$c(\ell, \ell + 1) = 1.$$

Otherwise, supposing $m \geq \ell + 2$ and continuing the induction, if we assume the Claim holds for all smaller m , then we find

$$1 = \sum_{\substack{R \in \mathcal{F}^{\text{root}}(m-\ell) \\ |\text{odd}(R)| = \mathbb{1}\{\ell \text{ odd}\}}} \mu(R) + c(\ell, m). \quad (10.71)$$

If ℓ is even, then the first term is a sum over $R \in \mathcal{F}(m - \ell)$. If ℓ is odd, then the sum may be viewed as a sum over $R \in \mathcal{F}(m - \ell + 1)$ by viewing the single root vertex as an additional leaf. In either case, this sum equals 1 by the definition of the Möbius function, whereby $c(\ell, m) = 0$, completing the proof.

We now return to the proof of the statement. Suppose $S \in \mathcal{M}([n])$, and let $\mathbf{s} \in [n]^{|S|}$ be the tuple of the elements of S in ascending order. Given $F \in \mathcal{F}(|S|)$, let us write $\text{ind}(F, S)$ for the multiset with one occurrence of each $i \in [n]$ for each connected component C of $\text{MaxSpan}(F, \mathbf{s})$ with $s_{\kappa(\ell)} = i$ for all leaves ℓ in C , and a further occurrence of each $i \in [n]$ for each leaf ℓ not belonging to $\text{MaxSpan}(F; \mathbf{s})$ with $s_{\kappa(\ell)} = i$. And, write $\text{coll}(F, S) \in \mathcal{F}(|\text{ind}(F, S)|)$ for the good forest obtained by deleting from F each component of $\text{MaxSpan}(F; \mathbf{s})$ and replacing each incidence between $\text{MaxSpan}(F; \mathbf{s})$ and $F \setminus \text{MaxSpan}(F; \mathbf{s})$ with a new leaf, labelled such that

$$Z^{\text{coll}(F, S)}(\mathbf{M}; \text{ind}(F, S)) = \sum_{\substack{\mathbf{a} \in [n]^{\text{V}^\square(F)} \\ \mathbf{a} \text{ (F,S)-tight}}} \prod_{\{v, w\} \in E(F)} M_{f_{S, \mathbf{a}(v)} f_{S, \mathbf{a}(w)}}. \quad (10.72)$$

Intuitively, these definitions describe the forest obtained from F by collapsing all tight subtrees in $\text{MaxSpan}(F, \mathbf{s})$, with extra occurrences of their indices added as labels on leaves in the new “fragmented” tree.

Using these definitions, we may rewrite the quantity that we need to compute as follows.

$$\begin{aligned}
\tilde{\mathbb{E}}^{\text{main}}[\mathbf{x}^S] &= \sum_{F \in \mathcal{F}(|S|)} \mu(F) \cdot \Delta^F(\mathbf{M}; S) \\
&= \sum_{F \in \mathcal{F}(|S|)} \mu(F) \sum_{\substack{\mathbf{a} \in [n]^{V^\square(F)} \\ \mathbf{a} \text{ (F,S)-tight}}} \prod_{\{v,w\} \in E(F)} M_{f_{S,\mathbf{a}}(v), f_{S,\mathbf{a}}(w)} \\
&= \sum_{F \in \mathcal{F}(|S|)} \mu(F) Z^{\text{coll}(F,S)}(\mathbf{M}; \text{ind}(F,S)) \\
&= \sum_{S' \in \mathcal{M}([n])} \sum_{F' \in \mathcal{F}(|S'|)} \underbrace{\left(\sum_{\substack{F \in \mathcal{F}(|S|) \\ \text{coll}(F,S)=F' \\ \text{ind}(F,S)=S'}} \mu(F) \right)}_{=: \zeta(F', S', S)} Z^{F'}(\mathbf{M}; S'). \tag{10.73}
\end{aligned}$$

We claim that the inner coefficient $\zeta(F', S', S)$ is zero unless S' is the (non-multi) set of indices occurring an odd number of times in S , in which case it is $\mu(F')$. This will complete the proof, since we will then have that the above equals $\tilde{\mathbb{E}}[\mathbf{x}^S]$ (by definition of the latter).

Since $\text{ind}(F, S)$ for any F only contains indices occurring in S , we will have $\zeta(F', S', S) = 0$ unless S' only contains indices also occurring in S . In other words, we have $\text{set}(S') \subseteq \text{set}(S)$; note, however, that a given index can occur more times in S' than in S .

Let C'_1, \dots, C'_m be the connected components of F' , let κ' be the function labelling the leaves of F' , and let \mathbf{s}' be the tuple of elements of S' in ascending order. Let $S'_i := \{s'_{\kappa'(\ell)} : \ell \in V^\bullet(C'_i)\}$, *a priori* a multiset. In fact, no index can occur twice in any S'_i : if j is the least such index, then by construction the minimal spanning tree on all $\ell \in V^\bullet(C'_i)$ with $s'_{\kappa'(\ell)} = j$ would have been included in $\text{MaxSpan}(F, \mathbf{s})$ and would have been collapsed in forming F' . Therefore, each S'_i is a set, and $S = S'_1 + \dots + S'_m$.

Now, we define the subsets of connected components containing a leaf labelled by each index: for $j \in [n]$, let $A_j = \{i \in [m] : j \in S'_i\}$. Also, let n_j equal the number of occurrences of j in S . Then, every F with $\text{coll}(F, S) = F'$ and $\text{ind}(F, S) = S'$ is obtained by composing with F' forests F_j for $j \in [n]$ whose leaves are $\kappa'^{-1}(j)$, together with some $n_j - |A_j|$ further

leaves $\ell_{j,1}, \dots, \ell_{j,n_j}$, such that F_j is dominated (in the sense above) by these further leaves, which all have $s_{\kappa(\ell_{j,k})} = j$, and such that F_j does not connect any C'_{i_1}, C'_{i_2} for $i_1, i_2 \in A_{j'}$ with $j' < j$. It is easier to understand the description in reverse: the F_j are precisely the forests added to $\text{MaxSpan}(F, \mathbf{s})$ for index j , if F collapses to F' .

Using this description of the F appearing in $\zeta(F', S', S)$ and the fact that $\mu(F)$ factorizes over \square vertices, we may factorize

$$\zeta(F', S', S) = \mu(F') \prod_{j=1}^n \left(\sum_{\substack{F \in \mathcal{F}(n_j) \\ F \text{ } [n_j - |A_j|]\text{-dominated} \\ F \text{ does not connect } C'_{i_1}, C'_{i_2} \\ \text{for } i_1, i_2 \in A_{j'}, j' < j}} \mu(F) \right). \quad (10.74)$$

Now, suppose for the sake of contradiction that $\zeta(F', S', S) \neq 0$ for some F' , and $|A_j| \neq \mathbb{1}\{n_j \text{ odd}\}$ for some j (remembering that A_j are defined in terms of F'). Choose the smallest such j . Then, the connectivity property on F in the j th factor above is vacuous since $|A_{j'}| \leq 1$ for all $j' < j$, so it may be removed in the summation. By the Claim, that factor is then zero, whereby $\zeta(F', S', S) = 0$, unless $|A_j| = \mathbb{1}\{n_j \text{ odd}\}$, so we reach a contradiction. Finally, if indeed $|A_j| = \mathbb{1}\{n_j \text{ odd}\}$ for all j , then the connectivity condition is vacuous for all terms, so it may always be removed, whereupon by the Claim the product above is 1 and $\zeta(F', S', S) = \mu(F')$ as desired. \square

Lastly, we prove the following additional result on $\tilde{\mathbb{E}}^{\text{err}}$ that will be useful later, showing that it decomposes into a sum over a choice of some “main term trees” and some “error trees” to apply to subsets of S .

Proposition 10.5.6 (Error term factorizes over connected components). *For any $S \in \mathcal{M}([n])$,*

$$\tilde{\mathbb{E}}^{\text{err}}[\mathbf{x}^S] = \sum_{\substack{A \subseteq S \\ A \neq \emptyset}} \tilde{\mathbb{E}}^{\text{main}}[\mathbf{x}^{S-A}] \sum_{\pi \in \text{Part}(A; \text{even})} \prod_{R \in \pi} \left(- \sum_{T \in \mathcal{T}(|R|)} \mu(T) \cdot \Delta^T(\mathbf{M}; R) \right). \quad (10.75)$$

Proof. We begin from the definition,

$$\Delta^F(\mathbf{M}; \mathbf{s}) = \sum_{\substack{\mathbf{a} \in [n]^{V^\square} \\ \mathbf{a} \text{ is } (F, \mathbf{s})\text{-loose}}} \prod_{\{v, w\} \in E} M_{f_{\mathbf{s}, \mathbf{a}}(v) f_{\mathbf{s}, \mathbf{a}}(w)}$$

Now, we observe from Proposition 10.5.2 that \mathbf{a} is (F, \mathbf{s}) -loose if and only if $\mathbf{a}|_T$ is $(F, \mathbf{s}|_T)$ -loose for some $T \in \text{conn}(F)$. Therefore, by the inclusion-exclusion principle, we may write

$$\begin{aligned} &= \sum_{\substack{A \subseteq \text{conn}(F) \\ A \neq \emptyset}} (-1)^{|A|-1} \sum_{\substack{\mathbf{a} \in [n]^{V^\square} \\ \mathbf{a}|_T \text{ is } (F, \mathbf{s}|_T)\text{-loose for } T \in A}} \prod_{\{v, w\} \in E} M_{f_{\mathbf{s}, \mathbf{a}}(v) f_{\mathbf{s}, \mathbf{a}}(w)} \\ &= - \sum_{\substack{A \subseteq \text{conn}(F) \\ A \neq \emptyset}} \prod_{T \in A} (-\Delta^T(\mathbf{M}; \mathbf{s}|_T)) \prod_{T \notin A} Z^T(\mathbf{M}; \mathbf{s}|_T). \end{aligned} \quad (10.76)$$

Now, we use that $\mu(F) = \prod_{i=1}^k \mu(T_i)$, so by definition of $\tilde{\mathbb{E}}^{\text{err}}$, we have

$$\begin{aligned} \tilde{\mathbb{E}}^{\text{err}}[\mathbf{x}^{\mathbf{s}}] &= - \sum_{F \in \mathcal{F}(m)} \mu(F) \cdot \Delta^F(\mathbf{M}; \mathbf{s}) \\ &= \sum_{F \in \mathcal{F}(m)} \sum_{\substack{A \subseteq \text{conn}(F) \\ A \neq \emptyset}} \prod_{T \in A} (-\mu(T) \cdot \Delta^T(\mathbf{M}; \mathbf{s}|_T)) \prod_{T \notin A} (\mu(T) \cdot Z^T(\mathbf{M}; \mathbf{s}|_T)). \end{aligned} \quad (10.77)$$

Reversing the order of summation and then reorganizing the inner sum according to the partition π of the leaves of F lying in each connected component then gives the result. \square

10.6 SPECTRAL ANALYSIS IN THE HARMONIC BASIS: OUTLINE OF THEOREM 10.2.3

Our basic strategy for proving positivity is to invoke Proposition 10.4.1 with the multiharmonic basis discussed in Remark 8.3.12. As our heuristic calculations there suggested, this will attenuate the multiscale spectrum of the pseudomoment matrix written in the standard

monomial basis, making the analysis of the spectrum much simpler. It will also let us use the heuristic Gram matrix expression (8.51) as a tool for proving positivity.

In this section, we describe the objects that arise after writing the pseudomoments in this basis, and state the main technical results that lead to the proof of Theorem 10.2.3. First, we recall the definition of the basis. For $S \subseteq [n]$, we have

$$q_S^\perp(\mathbf{x}; \mathbf{M}) := \begin{cases} x_i & \text{if } |S| = 1 \text{ with } S = \{i\}, \\ M_{ij} & \text{if } |S| = 2 \text{ with } S = \{i, j\}, \\ \sum_{a=1}^n \prod_{i \in S} M_{ia} \cdot x_a & \text{if } |S| \geq 3 \text{ is odd,} \\ \sum_{a=1}^n \prod_{i \in S} M_{ia} & \text{if } |S| \geq 4 \text{ is even,} \end{cases} \quad (10.78)$$

$$h_S^\perp(\mathbf{x}; \mathbf{M}) := \sum_{\sigma \in \text{Part}(S)} \prod_{A \in \sigma} (-1)^{|A|-1} (|A| - 1)! q_A^\perp(\mathbf{x}; \mathbf{M}). \quad (10.79)$$

For the sake of brevity, we will usually omit the explicit dependence on \mathbf{M} below, abbreviating $h_S^\perp(\mathbf{x}) = h_S^\perp(\mathbf{x}; \mathbf{M})$.

Next, we write the pseudomoments in this basis, separating the contributions of the main and error terms.

Definition 10.6.1 (Main and error pseudomoments). *Define $\mathbf{Z}^{\text{main}}, \mathbf{Z}^{\text{err}}, \mathbf{Z} \in \mathbb{R}^{\binom{[n]}{\leq d} \times \binom{[n]}{\leq d}}$ to have entries*

$$\mathbf{Z}_{S,T}^{\text{main}} := \tilde{\mathbb{E}}^{\text{main}}[h_S^\perp(\mathbf{x})h_T^\perp(\mathbf{x})], \quad (10.80)$$

$$\mathbf{Z}_{S,T}^{\text{err}} := \tilde{\mathbb{E}}^{\text{err}}[h_S^\perp(\mathbf{x})h_T^\perp(\mathbf{x})], \quad (10.81)$$

$$\mathbf{Z}_{S,T} := \tilde{\mathbb{E}}[h_S^\perp(\mathbf{x})h_T^\perp(\mathbf{x})] \quad (10.82)$$

$$= \mathbf{Z}_{S,T}^{\text{main}} + \mathbf{Z}_{S,T}^{\text{err}}. \quad (10.83)$$

We assign a technical lemma to the analysis of each of the two terms.

Lemma 10.6.2 (Positivity of main term). *Under the assumptions of Theorem 10.2.3,*

$$\lambda_{\min}(\mathbf{Z}^{\text{main}}) \geq \lambda_{\min}(\mathbf{M})^d - (6d)^{10d} \|\mathbf{M}\|^{3d} (\epsilon_{\text{tree}}(\mathbf{M}; d) + \epsilon_{\text{pow}}(\mathbf{M}) + \epsilon_{\text{offdiag}}(\mathbf{M}) + \epsilon_{\text{corr}}(\mathbf{M})). \quad (10.84)$$

Lemma 10.6.3 (Bound on error term). *Under the assumptions of Theorem 10.2.3,*

$$\|\mathbf{Z}^{\text{err}}\| \leq (12d)^{32d} \|\mathbf{M}\|^{5d} \epsilon_{\text{err}}(\mathbf{M}; 2d). \quad (10.85)$$

Given these statements, it is straightforward to prove our main theorem.

Proof of Theorem 10.2.3. Since the only multilinear monomial in $h_{\zeta}^{\perp}(\mathbf{x})$ is \mathbf{x}^S , the $h_{\zeta}^{\perp}(\mathbf{x})$ for $S \in \binom{[n]}{\leq d}$ form a basis for $\mathbb{R}[x_1, \dots, x_n]_{\leq d} / \mathcal{I}$ for \mathcal{I} the ideal generated by $\{x_i^2 - 1\}_{i=1}^n$. Thus by Proposition 10.4.1 it suffices to show $\mathbf{Z} \geq \mathbf{0}$. Since $\mathbf{Z} = \mathbf{Z}^{\text{main}} + \mathbf{Z}^{\text{err}}$, we have $\lambda_{\min}(\mathbf{Z}) \geq \lambda_{\min}(\mathbf{Z}^{\text{main}}) - \|\mathbf{Z}^{\text{err}}\|$. Substituting the results of Lemmata 10.6.2 and 10.6.3 then gives the result. \square

Before proceeding to the proof details, we note that we will use tools given later in Section 10.12 (various miscellaneous combinatorial and linear-algebraic bounds) and Section 10.13 (tools for working with CGMs).

10.7 APPROXIMATE BLOCK DIAGONALIZATION: TOWARDS

LEMMA 10.6.2

As a first step towards showing the positivity of \mathbf{Z}^{main} , we show that our choice of writing the pseudomoments of $\tilde{\mathbb{E}}^{\text{main}}$ in the multiharmonic basis makes \mathbf{Z}^{main} approximately block diagonal. This verifies what we expect based on the informal argument leading up to Re-

mark 8.3.12.

10.7.1 STRETCHED FOREST RIBBON DIAGRAMS

We first describe an important cancellation in Z^{main} . Writing the pseudomoments in the multiharmonic basis in fact leaves only the following especially well-behaved type of forest ribbon diagram. Below we call a \square vertex *terminal* if it is incident to any leaves.

Definition 10.7.1 (Stretched forest ribbon diagram). *We say that $F \in \mathcal{F}(\ell, m)$ is stretched if it satisfies the following properties:*

1. Every terminal \square vertex of F has a neighbor in both \mathcal{L} and \mathcal{R} .
2. No connected component of F is a sided pair: a pair of connected \bullet vertices both lying in \mathcal{L} or both lying in \mathcal{R} .
3. No connected component of F is a skewed star: a star with one vertex in \mathcal{L} and more than one vertex in \mathcal{R} , or one vertex in \mathcal{R} and more than one vertex in \mathcal{L} .

A fortunate combinatorial cancellation shows that, in the multiharmonic basis, the pseudomoment terms of stretched forest ribbon diagrams retain their initial coefficients, while non-stretched forest ribbon diagrams are eliminated.

Proposition 10.7.2. *For any $S, T \subseteq [n]$,*

$$Z_{S,T}^{\text{main}} = \tilde{\mathbb{E}}^{\text{main}}[h_S^\dagger(\mathbf{x})h_T^\dagger(\mathbf{x})] = \sum_{\substack{F \in \mathcal{F}(|S|, |T|) \\ F \text{ stretched}}} \mu(F) \cdot Z_{S,T}^F(\mathbf{M}). \quad (10.86)$$

Proof. We expand directly:

$$\begin{aligned} & \tilde{\mathbb{E}}^{\text{main}}[h_S^\dagger(\mathbf{x})h_T^\dagger(\mathbf{x})] \\ &= \tilde{\mathbb{E}}^{\text{main}} \left[\left(\sum_{\sigma \in \text{Part}(S)} \prod_{A \in \sigma} (-1)^{|A|-1} (|A|-1)! q_A^\dagger(\mathbf{x}) \right) \left(\sum_{\tau \in \text{Part}(T)} \prod_{B \in \tau} (-1)^{|B|-1} (|B|-1)! q_B^\dagger(\mathbf{x}) \right) \right] \end{aligned}$$

$$\begin{aligned}
&= \sum_{\substack{\sigma \in \text{Part}(S) \\ \tau \in \text{Part}(T)}} \prod_{R \in \sigma + \tau} (-1)^{|R|-1} (|R| - 1)! \prod_{R \in \sigma[\text{even}] + \tau[\text{even}]} a_R^\dagger \\
&\quad \sum_{\substack{\mathbf{a} \in [n]^{\sigma[\text{odd}; \geq 3]} \\ \mathbf{b} \in [n]^{\tau[\text{odd}; \geq 3]}}} \prod_{A \in \sigma[\text{odd}; \geq 3]} \prod_{i \in A} M_{a(A), i} \cdot \prod_{B \in \tau[\text{odd}; \geq 3]} \prod_{j \in B} M_{b(B), j} \cdot \\
&\quad \tilde{\mathbb{E}}^{\text{main}} \left[\prod_{\{i\} \in \sigma[1]} x_i \prod_{A \in \sigma[\text{odd}; \geq 3]} x_{a(A)} \prod_{\{j\} \in \tau[1]} x_j \prod_{B \in \tau[\text{odd}; \geq 3]} x_{b(B)} \right] \quad (10.87)
\end{aligned}$$

Let us write $f_a : \sigma[\text{odd}] \rightarrow [n]$ to map $A = \{i\} \mapsto i$ when $|A| = 1$ and to map $A \mapsto a(A)$ when $|A| \geq 3$, and likewise $g_b : \tau[\text{odd}] \rightarrow [n]$. Then, expanding the pseudoexpectation, we have

$$\begin{aligned}
&= \sum_{\substack{\sigma \in \text{Part}(S) \\ \tau \in \text{Part}(T)}} \prod_{R \in \sigma + \tau} (-1)^{|R|-1} (|R| - 1)! \prod_{R \in \sigma[\text{even}] + \tau[\text{even}]} a_R^\dagger \\
&\quad \sum_{\substack{\mathbf{a} \in [n]^{\sigma[\text{odd}; \geq 3]} \\ \mathbf{b} \in [n]^{\tau[\text{odd}; \geq 3]}}} \prod_{A \in \sigma[\text{odd}; \geq 3]} \prod_{i \in A} M_{a(A), i} \cdot \prod_{B \in \tau[\text{odd}; \geq 3]} \prod_{j \in B} M_{b(B), j} \cdot \\
&\quad \sum_{F \in \mathcal{F}(|\sigma[\text{odd}|], |\tau[\text{odd}|])} \mu(F) \cdot Z_{(f_a(A))_{A \in \sigma[\text{odd}], (g_b(B))_{B \in \tau[\text{odd}]}}^F
\end{aligned}$$

We say that $F \in \mathcal{F}(|S|, |T|)$ is an *odd merge of (σ, τ) through $F' \in \mathcal{F}(|\sigma[\text{odd}|], |\tau[\text{odd}|])$* if F consists of even stars on the even parts of σ and τ , and even stars on the odd parts of σ and τ with one extra leaf added to each, composed in the sense of the compositional ordering of Definition 10.3.7 with F' . See Figure 10.3 for an example. When F is an odd merge of (σ, τ) through F' , then F' is uniquely determined by σ, τ , and F . Using this notion, we may rewrite the above as

$$= \sum_{F \in \mathcal{F}(|S|, |T|)} \left(\sum_{\substack{\sigma \in \text{Part}(S) \\ \tau \in \text{Part}(T) \\ F' \in \mathcal{F}(|\sigma[\text{odd}|], |\tau[\text{odd}|]) \\ F \text{ is an odd merge} \\ \text{of } (\sigma, \tau) \text{ through } F'}} \prod_{R \in \sigma + \tau} (-1)^{|R|-1} (|R| - 1)! \cdot \mu(F') \right) Z_{S, T}^F$$




Figure 10.3: Odd merge of partitions. We illustrate an odd merge of two partitions σ and τ through a forest ribbon diagram F , as used in the proof of Proposition 10.7.2. The gray boxes show the components of the resulting diagram arising from the partitions (and, one step before, arising from terms in the multiharmonic basis polynomials $h_S^1(x)$), while the remainder is the forest ribbon diagram that merges the odd parts of the partitions. The odd parts that are merged by F are highlighted with bold brackets.

We make two further simplifying observations. First, the factors of $(-1)^{|R|}$ multiply to $(-1)^{|S|+|T|}$, and $|S| + |T|$ must be even in order for $\mathcal{F}(|S|, |T|)$ to be non-empty, so we may omit the $(-1)^{|R|}$ factors. Second, by the factorization $\mu(F) = \prod_{v \in V^\square(F)} (-(\deg(v) - 2)!)$, when F is an odd merge of (σ, τ) through F' then we have $\mu(F) = \mu(F') \prod_{|R| \geq 3 \text{ is odd}} (-(|R| - 1)!) \prod_{|R| \geq 4 \text{ is even}} (-(|R| - 2)!)$, where both products are over $R \in \sigma + \tau$ satisfying the given conditions. It may again be helpful to consult Figure 10.3 to see why this formula holds. Using this, we may extract $\mu(F)$ and continue

$$= \sum_{F \in \mathcal{F}(|S|, |T|)} \underbrace{\left(\sum_{\substack{\sigma \in \text{Part}(S) \\ \tau \in \text{Part}(T) \\ F \text{ odd merge} \\ \text{of } (\sigma, \tau)}} (-1)^{|\sigma_{[\leq 2]}| + |\tau_{[\leq 2]}|} \prod_{R \in \sigma[\text{even}; \geq 4] + \tau[\text{even}; \geq 4]} (R - 1) \right)}_{=: \eta(F)} \mu(F) \cdot Z_{S, T}^F. \quad (10.88)$$

It remains to analyze the inner coefficient $\eta(F)$. To do this, it suffices to enumerate the

pairs of partitions (σ, τ) of which F is an odd merge. We describe the possible partitions below.

- If $v \in V^\square(F)$ is the only \square vertex of a skewed star connected component with leaves $i_1, \dots, i_k \in \mathcal{L}$ (for k odd) and $j \in \mathcal{R}$, then j must be a singleton in τ while i_1, \dots, i_k can either (1) all be singletons in σ or (2) constitute one part $\{i_1, \dots, i_k\}$ of σ . A symmetric condition holds if there is more than one leaf in \mathcal{R} and one leaf in \mathcal{L} .
- If $v \in V^\square(F)$ is the only \square vertex of a sided star connected component with leaves $i_1, \dots, i_k \in \mathcal{L}$ (for k even), then the i_1, \dots, i_k can either (1) all be singletons in σ , (2) constitute one part $\{i_1, \dots, i_k\}$ of σ , or (3) be divided into an odd part $\{i_1, \dots, i_k\} \setminus \{i_{k^*}\}$ and a singleton $\{i_{k^*}\}$ for any choice of $k^* \in 1, \dots, k$. A symmetric condition holds if the leaves are all in \mathcal{R} .
- If $i_1, i_2 \in \mathcal{L}$ form a sided pair in F , then i_1, i_2 can either (1) both be singletons in σ , or (2) constitute one part $\{i_1, i_2\}$ of σ . A symmetric condition holds if the two leaves are in \mathcal{R} .
- If $v \in V^\square(F)$ is terminal, is not the only \square vertex of its connected component, and has leaf neighbors $i_1, \dots, i_k \in \mathcal{L}$ (for k odd), then the i_1, \dots, i_k can either (1) all be singletons in σ , or (2) constitute one part $\{i_1, \dots, i_k\}$ of σ . A symmetric condition holds if the leaves are all in \mathcal{R} .
- If $v \in V^\square(F)$ is terminal, is not the only \square vertex of its connected component, and has leaf neighbors in both \mathcal{L} and \mathcal{R} , then all leaves attached to v must be singletons in σ and τ (according to whether they belong to \mathcal{L} or \mathcal{R} , respectively).
- If $i \in \mathcal{L}$ and $j \in \mathcal{R}$ form a non-sided pair in F , then i and j must be singletons in σ and τ , respectively.

Factorizing $\eta(F)$ according to which terminal \square vertex or pair connected component each leaf of F is attached to using these criteria, we find

$$\begin{aligned}
\eta(F) &= \prod_{\substack{C \in \text{conn}(F) \\ C \text{ sided pair}}} (1 - 1) \prod_{\substack{v \text{ terminal in } V^\square(F), \\ \text{all leaf neighbors of } v \text{ in } \mathcal{L} \\ \text{or all leaf neighbors of } v \text{ in } \mathcal{R}}} (1 - 1) \prod_{\substack{C \in \text{conn}(F) \\ C \text{ sided even star} \\ \text{on } k \geq 4 \text{ leaves}}} ((k - 1) - k + 1) \cdot \\
&\quad \prod_{\substack{C \in \text{conn}(F) \\ C \text{ skewed star} \\ \text{on } \geq 4 \text{ leaves}}} (1 - 1) \\
&= \mathbb{1}\{F \text{ is stretched}\}, \tag{10.89}
\end{aligned}$$

completing the proof. □

10.7.2 TYING BOUND AND BOWTIE FOREST RIBBON DIAGRAMS

The above does not appear to give the block diagonalization we promised—there exist stretched ribbon diagrams in $\mathcal{F}(\ell, m)$ even when $\ell \neq m$, so the off-diagonal blocks of \mathbf{Z}^{main} are non-zero. To find that this is an *approximate* block diagonalization, we must recognize that the CGMs of many stretched ribbon diagrams are approximately equal (up to a small error in operator norm) and then observe another combinatorial cancellation in these “standardized” diagrams.

Specifically, we will show that all stretched forest ribbon diagrams’ CGMs can be reduced to the following special type of stretched forest ribbon diagram.

Definition 10.7.3 (Bowtie forest ribbon diagram). *We call $F \in \mathcal{F}(\ell, m)$ a bowtie forest if every connected component of F is a pair or star with at least one leaf in each of \mathcal{L} and \mathcal{R} . Each connected component of such F is a bowtie. We call a bowtie or bowtie forest balanced if all components have an equal number of leaves in \mathcal{L} and in \mathcal{R} .*

Note that there are no balanced bowtie forests in $\mathcal{F}(\ell, m)$ unless $\ell = m$; thus, since in our

final expression for the pseudomoments below only balanced bowtie forests will remain, we will indeed have an approximate block diagonalization.

Next, we show that any stretched forest ribbon diagram can be “tied” to form a bowtie forest ribbon diagram by collapsing every non-pair connected component to have a single \square vertex, while incurring only a small error in the associated CGMs.

Lemma 10.7.4 (Stretched forest ribbon diagrams: tying bound). *Suppose that $\epsilon_{\text{tree}}(\mathbf{M}; (\ell + m)/2) \leq 1$. Let $F \in \mathcal{F}(\ell, m)$ be stretched. Let $\text{tie}(F) \in \mathcal{F}(\ell, m)$ be formed by replacing each connected component of F that is not a pair with the bowtie of a single \square vertex attached to all of the leaves of that connected component. Then, $\text{tie}(F)$ is a bowtie forest, and*

$$\|\mathbf{Z}^F - \mathbf{Z}^{\text{tie}(F)}\| \leq (\ell + m)(2\|\mathbf{M}\|)^{\frac{3}{2}(\ell+m)} \epsilon_{\text{tree}}(\mathbf{M}; (\ell + m)/2).$$

The basic intuition behind the result is that, since every terminal \square vertex of a stretched ribbon diagram is connected to both \mathcal{L} and \mathcal{R} , the corresponding CGM may be factorized into $\mathbf{Z}^F = \mathbf{A}^{\mathcal{L}} \mathbf{D} \mathbf{A}^{\mathcal{R}}$, where $\mathbf{A}^{\mathcal{L}}$ and $\mathbf{A}^{\mathcal{R}}$ correspond to the contributions of edges attaching \mathcal{L} and \mathcal{R} respectively to the internal vertices, while \mathbf{D} is CGM of the induced “inner” ribbon diagram on the \square vertices. Thanks to the diagram being stretched, \mathbf{D} is actually diagonal, and the result essentially states that its only significant entries are those corresponding to all \square vertices in each connected component having the same index. That is the origin of the ϵ_{tree} incoherence quantity here.

Proof of Lemma 10.7.4. We recall the statement of the result. Let $F \in \mathcal{F}(\ell, m)$ be a stretched forest ribbon diagram, with all edges labelled with \mathbf{M} . Let $\text{tie}(F)$ be the forest ribbon diagram constructed by tying every connected component of F that is not a pair into a bowtie. By construction, $\text{tie}(F)$ is a bowtie forest ribbon diagram. Our goal is then to show the bound

$$\|\mathbf{Z}^F - \mathbf{Z}^{\text{tie}(F)}\| \leq (\ell + m)(2\|\mathbf{M}\|)^{\frac{3}{2}(\ell+m)} \epsilon_{\text{tree}}(\mathbf{M}; (\ell + m)/2). \quad (10.90)$$

Let us first suppose that $T \in \mathcal{T}(\ell, m)$ is connected. We will then show the bound

$$\|\mathbf{Z}^T - \mathbf{Z}^{\text{tie}(T)}\| \leq 2^{\frac{3}{2}(\ell+m)} \epsilon_{\text{tree}}(\mathbf{M}; (\ell+m)/2), \quad (10.91)$$

the same as the above but without the leading factor of $(\ell+m)$. Since the norm of the CGM of a connected component of F that has k leaves is at most $\|\mathbf{M}\|^{\frac{3}{2}k}$ by Proposition 10.13.13 and Corollary 10.12.10, the bound on arbitrary F will then follow by applying Proposition 10.13.4.

Write $U \subseteq V^\square(T)$ for the set of terminal vertices of T . By Corollary 10.12.10, $|U| \leq |V^\square(T)| \leq (\ell+m)/2$. Recall that, since T is stretched, every vertex of U is adjacent to some vertex of \mathcal{L} and some vertex of \mathcal{R} . Let A be the ribbon diagram on vertex triplet $((\mathcal{L}, U), \emptyset)$ induced by T , let B be the ribbon diagram on vertex triplet $((U, U), \emptyset)$ obtained by deleting \mathcal{L} and \mathcal{R} from T , and let C be the ribbon diagram on vertex triplet $((U, \mathcal{R}), \emptyset)$ induced by T . Then, by Proposition 10.13.11, we may factorize $\mathbf{Z}^T = \mathbf{Z}^{G[A]} \mathbf{Z}^{G[B]} \mathbf{Z}^{G[C]}$. Moreover, let $G'[B]$ be the ribbon diagram obtained by relabelling every edge in $G[B]$ with the identity matrix. Then we have $\mathbf{Z}^{\text{tie}(T)} = \mathbf{Z}^{G[A]} \mathbf{Z}^{G'[B]} \mathbf{Z}^{G[C]}$. Therefore, by norm submultiplicativity and Proposition 10.13.13 applied to $G[A]$ and $G[C]$, we may bound

$$\|\mathbf{Z}^T - \mathbf{Z}^{\text{tie}(T)}\| \leq \|\mathbf{Z}^{G[A]}\| \cdot \|\mathbf{Z}^{G[C]}\| \cdot \|\mathbf{Z}^{G[B]} - \mathbf{Z}^{G'[B]}\| \leq \|\mathbf{M}\|^{\ell+m} \|\mathbf{Z}^{G[B]} - \mathbf{Z}^{G'[B]}\|. \quad (10.92)$$

Since $\mathcal{L}(G[B]) = \mathcal{R}(G[B]) = \mathcal{L}(G'[B]) = \mathcal{R}(G'[B]) = U$, the matrices $\mathbf{Z}^{G[B]}$ and $\mathbf{Z}^{G'[B]}$ are diagonal. Let us view $G[B]$ as being partitioned into edge-disjoint (but not necessarily vertex disjoint) subtrees B_1, \dots, B_n , such that every leaf of every B_i belongs to U . Since the B_i are edge-disjoint, n is at most the number of edges in T , which by Corollary 10.12.10 is at most $\frac{3}{2}(\ell+m)$. Since $\mathcal{L}(G[B]) = \mathcal{R}(G[B])$, the labellings $\kappa_{\mathcal{L}}$ and $\kappa_{\mathcal{R}}$ must be equal, so let us simply write κ for this single labelling. Suppose that the leaves of B_i are $\ell_{i,1}, \dots, \ell_{i,a_i} \in U$.

Then, we have

$$Z_{s,s}^{G[B]} = \prod_{i=1}^n Z^{B_i} \left(\mathbf{M}; (s(\kappa(\ell_{i,1})), \dots, s(\kappa(\ell_{i,a_i}))) \right). \quad (10.93)$$

By the definition of ϵ_{tree} , we have

$$\begin{aligned} & \left| \mathbb{1}\{s(\kappa(\ell_{i,1})) = \dots = s(\kappa(\ell_{i,a_i}))\} - Z^{B_i} \left(\mathbf{M}; (s(\kappa(\ell_{i,1})), \dots, s(\kappa(\ell_{i,a_i}))) \right) \right| \\ & \leq \epsilon_{\text{tree}}(\mathbf{M}; a_i) \\ & \leq \epsilon_{\text{tree}}(\mathbf{M}; (\ell + m)/2). \end{aligned} \quad (10.94)$$

Since the B_i form an edge partition of T into subtrees, we have

$$\prod_{i=1}^n \mathbb{1}\{s(\kappa(\ell_{i,1})) = \dots = s(\kappa(\ell_{i,a_i}))\} = \mathbb{1}\{s(i) = s(j) \text{ for all } i, j \in [U]\} = Z_{s,s}^{G'[B]}. \quad (10.95)$$

Therefore, substituting and expanding, we find

$$\begin{aligned} & \left| Z_{s,s}^{G[B]} - Z_{s,s}^{G'[B]} \right| \\ & = \left| \prod_{i=1}^n Z^{B_i} \left(\mathbf{M}; (s(\kappa(\ell_{i,1})), \dots, s(\kappa(\ell_{i,a_i}))) \right) - \prod_{i=1}^n \mathbb{1}\{s(\kappa(\ell_{i,1})) = \dots = s(\kappa(\ell_{i,a_i}))\} \right| \\ & = \left| \sum_{\substack{A \subseteq [n] \\ A \neq \emptyset}} \prod_{i \in A} \left(Z^{B_i} \left(\mathbf{M}; (s(\kappa(\ell_{i,1})), \dots, s(\kappa(\ell_{i,a_i}))) \right) - \mathbb{1}\{s(\kappa(\ell_{i,1})) = \dots = s(\kappa(\ell_{i,a_i}))\} \right) \right. \\ & \quad \left. \prod_{i \notin A} \mathbb{1}\{s(\kappa(\ell_{i,1})) = \dots = s(\kappa(\ell_{i,a_i}))\} \right| \\ & \leq \sum_{\substack{A \subseteq [n] \\ A \neq \emptyset}} \epsilon_{\text{tree}}(\mathbf{M}; (\ell + m)/2)^{|A|} \\ & \leq 2^n \epsilon_{\text{tree}}(\mathbf{M}; (\ell + m)/2) \\ & \leq 2^{\frac{3}{2}(\ell+m)} \epsilon_{\text{tree}}(\mathbf{M}; (\ell + m)/2), \end{aligned} \quad (10.96)$$

completing the proof. □

10.7.3 SIMPLIFICATION OF COEFFICIENTS

We next define the result of tying all of the ribbon diagrams in $\mathcal{Z}^{\text{main}}$:

$$Z_{S,T}^{\text{tied}} := \sum_{\substack{F \in \mathcal{F}(|S|,|T|) \\ F \text{ stretched}}} \mu(F) \cdot Z_{S,T}^{\text{tie}(F)}, \quad (10.97)$$

where since each bowtie forest can be formed by tying multiple stretched forests, we rewrite to isolate the resulting coefficient of each bowtie forest,

$$= \sum_{\substack{F \in \mathcal{F}(|S|,|T|) \\ F \text{ bowtie forest}}} \left(\underbrace{\sum_{\substack{F' \in \mathcal{F}(|S|,|T|) \\ F' \text{ stretched} \\ \text{tie}(F')=F}} \mu(F')}_{=:\xi(F)} \right) Z_{S,T}^F. \quad (10.98)$$

The following result gives the combinatorial analysis of the coefficients $\xi(F)$ appearing here, which yields a surprising cancellation that verifies that $\mathcal{Z}^{\text{main}}$ is approximately block diagonal.

Lemma 10.7.5 (Stretched forest ribbon diagrams: combinatorial reduction). *Let $F \in \mathcal{F}(\ell, m)$ be a bowtie forest. Then,*

$$\xi(F) := \sum_{\substack{F' \in \mathcal{F}(|S|,|T|) \\ F' \text{ stretched} \\ \text{tie}(F')=F}} \mu(F') = \mathbb{1}\{F \text{ balanced}\} \prod_{\substack{C \in \text{conn}(F) \\ C \text{ balanced bowtie} \\ \text{on } 2k \text{ leaves}}} (-1)^{k-1} (k-1)! k!. \quad (10.99)$$

We give the proof, a rather involved calculation with exponential generating functions, below. We leave open the interesting problem of finding a more conceptual combinatorial proof of this result, especially in light of the appearance of $\xi(F)$ again in Lemma 10.8.5 later.

Proof of Lemma 10.7.5. It suffices to consider connected forests, since both the left- and right-hand sides of the statement factorize over connected components. Thus, we want to show

$$\sum_{\substack{T \in \mathcal{T}(\ell, m) \\ T \text{ stretched}}} \mu(F) = \mathbb{1}\{\ell = m\} (-1)^{m-1} (m-1)! m!. \quad (10.100)$$

It will be slightly easier to work with a less stringent definition of “stretched” which removes the exceptions for skewed stars, and also allows sided stars if the other side has no \bullet vertices. Let us call F *weakly stretched* if every terminal \square vertex has a neighbor both in \mathcal{L} and in \mathcal{R} , or if there is only one \square vertex and one of \mathcal{L} and \mathcal{R} is empty. Then, our task is equivalent to showing

$$\sum_{\substack{T \in \mathcal{T}(\ell, m) \\ T \text{ weakly stretched}}} \mu(F) = \begin{cases} (-1)^{m-1} (m-1)! m! & \text{if } \ell = m \geq 1, \\ -(m-2)! & \text{if } \ell = 0, m \geq 4 \text{ is even,} \\ -(\ell-2)! & \text{if } m = 0, \ell \geq 4 \text{ is even,} \\ -(m-1)! & \text{if } \ell = 1, m \geq 3 \text{ is odd,} \\ -(\ell-1)! & \text{if } m = 1, \ell \geq 3 \text{ is odd,} \\ 0 & \text{otherwise.} \end{cases} \quad (10.101)$$

We use the convention that lowercase functions of combinatorial variables, like $f(a, b)$, give coefficients, and uppercase functions of analytic variables, like $F(x, y)$, give the corresponding exponential generating functions.

Define the coefficients

$$c(k) = \begin{cases} -(k-2)! & \text{if } k \geq 4 \text{ is even,} \\ 0 & \text{otherwise.} \end{cases} \quad (10.102)$$

Our goal is to compute the coefficients

$$f(\ell, m) = \sum_{\substack{T \in \mathcal{T}(\ell, m) \\ T \text{ weakly stretched}}} \prod_{v \in V^\square(T)} c(\deg(v)), \quad (10.103)$$

Equivalently, separating the terminal and non-terminal vertices, we may rewrite with the following intermediate quantities:

$$g_T(m) := \sum_{\phi: [m] \rightarrow V^\square(T)} \prod_{v \in V^\square(T)} c(\deg(v) + |\phi^{-1}(v)|) \text{ for a given } T \in \mathcal{T}(m), \quad (10.104)$$

$$g(\ell, m) := \sum_{T \in \mathcal{T}(\ell)} g_T(m), \quad (10.105)$$

$$h(\ell, m, n, p) := \sum_{\substack{\pi \in \text{Part}([\ell + m]; \text{odd}) \\ S \cap \{1, \dots, \ell\} \neq \emptyset \text{ for all } S \in \pi \\ S \cap \{\ell + 1, \dots, m\} \neq \emptyset \text{ for all } S \in \pi}} \prod_{S \in \pi} (-(|S| - 1)!) g(n + |\pi|, p), \quad (10.106)$$

$$f(\ell, m) = c(\ell + m) + \sum_{a=0}^{\ell} \sum_{b=0}^m \binom{\ell}{a} \binom{m}{b} h(a, b, 0, \ell + m - a - b). \quad (10.107)$$

The first term in the final expression counts the star tree on $[a + b]$ having a single \square vertex. In any other tree, every terminal \square vertex must be adjacent to an odd number of leaves, giving the remaining recursion. We will calculate the exponential generating functions of the sums appearing, in the same order as they are given above. We have introduced a needlessly general version of $h(\cdot, \cdot, \cdot, \cdot)$ in order to make it simpler to close a recursion to come.

Before proceeding, we compute the exponential generating function of the $c(k)$:

$$\begin{aligned} C(x) &= \sum_{k \geq 0} \frac{x^k}{k!} c(k) \\ &= - \sum_{k=2}^{\infty} \frac{x^{2k}}{2k(2k-1)} \\ &= \frac{1}{2} x^2 - \frac{1}{2} (1+x) \log(1+x) - \frac{1}{2} (1-x) \log(1-x). \end{aligned} \quad (10.108)$$

Next, to compute the exponential generating function of the $g_T(m)$, note that, grouping by the values of $|\phi^{-1}(v)|$ for each v , we may rewrite

$$g_T(m) = \sum_{\substack{z \in \mathbb{N}^{V^\square(T)} \\ |z|=m}} \binom{m}{z} \prod_{v \in V^\square(T)} c(\deg(v) + z_v). \quad (10.109)$$

Thus, the generating function factorizes as

$$\begin{aligned} G_T(x) &= \sum_{m \geq 0} \frac{x^m}{m!} g_T(m) \\ &= \prod_{v \in V^\square(T)} \left(\sum_{m \geq 0} \frac{x^m}{m!} f(\deg(v) + m) \right) \\ &= \prod_{v \in V^\square(T)} \frac{d^{\deg(v)}}{dx^{\deg(v)}} \left(\sum_{m \geq 0} \frac{x^m}{m!} c(m) \right) \\ &= \prod_{v \in V^\square(T)} C^{(\deg(v))}(x). \end{aligned} \quad (10.110)$$

Next, for the $g(\ell, m)$, we have

$$G(x, y) = \sum_{\ell \geq 0} \frac{x^\ell}{\ell!} \sum_{\substack{T \in \mathcal{F}(\ell) \\ T \text{ connected}}} G_T(y). \quad (10.111)$$

Let us define

$$G_\ell(y) = \sum_{\substack{T \in \mathcal{F}(\ell) \\ T \text{ connected}}} G_T(y). \quad (10.112)$$

A tree on $\ell > 1$ leaves can either be a single edge between two leaves (if $\ell = 2$), or will have every leaf connected to an internal vertex. In the latter case, let us think of the tree as being rooted as the internal vertex that the leaf labelled ℓ is attached to. Then, recursively, the tree consists of several rooted trees, whose leaves form a partition of $[\ell - 1]$, attached to the single new root. (This is similar to our formalism of “rooted odd trees” from

Definition 10.3.12.) Writing this recursive structure in terms of generating functions,

$$G_\ell(\mathbf{y}) = \mathbb{1}\{\ell = 2\} + \sum_{\pi \in \text{Part}([\ell-1])} C^{(|\pi|+1)}(\mathbf{y}) \prod_{S \in \pi} G_{|S|+1}(\mathbf{y}). \quad (10.113)$$

Now, noting that $G(x, \mathbf{y})$ is the exponential generating function of the $G_k(\mathbf{y})$, we use the composition formula. This calculation is clearer if we reindex, defining $\tilde{G}_\ell(\mathbf{y}) = G_{\ell+1}(\mathbf{y})$, which satisfy

$$\tilde{G}_\ell(\mathbf{y}) = \mathbb{1}\{\ell = 1\} + \sum_{\pi \in \text{Part}([\ell])} C^{(|\pi|+1)}(\mathbf{y}) \prod_{S \in \pi} \tilde{G}_{|S|}(\mathbf{y}). \quad (10.114)$$

Note that

$$\sum_{\ell \geq 0} \frac{x^\ell}{\ell!} C^{(\ell)}(\mathbf{y}) = C(x + \mathbf{y}) \quad (10.115)$$

since this is just a Taylor expansion of C about \mathbf{y} . The generating function of $C^{(k+1)}(\mathbf{y})$ is then the derivative in x , which is $C'(x + \mathbf{y})$. However, we must subtract off the term that is constant in x before using this in the composition formula, giving $C'(x + \mathbf{y}) - C'(\mathbf{y})$. Thus, we find by the composition formula

$$\tilde{G}(x, \mathbf{y}) := \sum_{\ell=0}^{\infty} \frac{x^\ell}{\ell!} \tilde{G}_\ell(\mathbf{y}) = x + C'(\tilde{G}(x, \mathbf{y}) + \mathbf{y}) - C'(\mathbf{y}). \quad (10.116)$$

We have

$$C'(x) = x - \frac{1}{2} \log(1+x) + \frac{1}{2} \log(1-x) \quad (10.117)$$

whereby the above functional equation is

$$\begin{aligned} \tilde{G}(x, \mathbf{y}) &= x + \tilde{G}(x, \mathbf{y}) + \mathbf{y} - \frac{1}{2} \log(1 + \tilde{G}(x, \mathbf{y}) + \mathbf{y}) + \frac{1}{2} \log(1 - \tilde{G}(x, \mathbf{y}) - \mathbf{y}) - \mathbf{y} \\ &\quad + \frac{1}{2} \log(1 + \mathbf{y}) - \frac{1}{2} \log(1 - \mathbf{y}), \end{aligned}$$

which, after cancellations, gives

$$0 = x - \frac{1}{2} \log(1 + \tilde{G}(x, y) + y) + \frac{1}{2} \log(1 - \tilde{G}(x, y) - y) + \frac{1}{2} \log(1 + y) - \frac{1}{2} \log(1 - y), \quad (10.118)$$

and exponentiating we have

$$e^{-2x} = \frac{(1 - \tilde{G}(x, y) - y)(1 + y)}{(1 + \tilde{G}(x, y) + y)(1 - y)} = \frac{1 - \frac{\tilde{G}(x, y)}{1-y}}{1 + \frac{\tilde{G}(x, y)}{1+y}} \quad (10.119)$$

solving which we find

$$\tilde{G}(x, y) = \frac{1 - e^{-2x}}{\frac{1}{1-y} + \frac{e^{-2x}}{1+y}} = \frac{(1 - y^2)(1 - e^{-2x})}{1 + e^{-2x} + y(1 - e^{-2x})} = \frac{1 - y^2}{y + \coth(x)}. \quad (10.120)$$

Finally, $G(x, y)$ is the integral with respect to x , with the boundary condition $G(0, y) = 0$.

This gives

$$G(x, y) = \log(\cosh(x)) - xy + \log(1 + y \tanh(x)) \quad (10.121)$$

Note that, when $y = 0$, we recover the result from the proof of Lemma 10.3.11 that the sum over trees of our Mobius function gives the alternating tangent numbers from Example 10.3.6, whose generating function is $\log(\cosh(x))$.

We next compute the exponential generating function of the $h(\ell, m, n, p)$. Define the simpler version of these coefficients, without the condition that each subset of the partition π intersect both $\{1, \dots, \ell\}$ and $\{\ell + 1, \dots, m\}$:

$$h(\ell, m, n, p) := \sum_{\pi \in \text{Part}([\ell+m]; \text{odd})} \prod_{S \in \pi} (-(|S| - 1)!) g(n + |\pi|, p). \quad (10.122)$$

Note that, by decomposing an odd partition into the parts that are contained in k , the parts

that are contained in ℓ , and all the other parts, we have

$$\tilde{h}(\ell, m, n, p) = \sum_{q=0}^{\ell} \sum_{r=0}^m \binom{\ell}{q} \binom{m}{r} \sum_{\substack{\pi \in \text{Part}([q]; \text{odd}) \\ \rho \in \text{Part}([r]; \text{odd})}} \prod_{S \in \pi + \rho} (-(|S|-1)!) h(\ell - q, m - r, n + |\pi| + |\rho|, p). \quad (10.123)$$

Now, by the composition formula this implies

$$\tilde{H}(w, x, y, z) = H(w, x, y - \tanh^{-1}(w) - \tanh^{-1}(x), z) \quad (10.124)$$

and therefore we can conversely recover H from \tilde{H} by

$$H(w, x, y, z) = \tilde{H}(w, x, y + \tanh^{-1}(w) + \tanh^{-1}(x), z). \quad (10.125)$$

On the other hand, again by the composition formula and the addition formula,

$$\tilde{H}(w, x, y, z) = \sum_{m \geq 0} \frac{y^m}{m!} \frac{\partial^m}{\partial t^m} [G(t, z)]_{t = -\tanh^{-1}(w+x)} = G(y - \tanh^{-1}(w+x), z) \quad (10.126)$$

and thus

$$H(w, x, y, z) = G(y + \tanh^{-1}(w) + \tanh^{-1}(x) - \tanh^{-1}(w+x), z). \quad (10.127)$$

Lastly, by the addition formula we have

$$\begin{aligned} F(x, y) &= C(x+y) + H(x, y, 0, x+y) \\ &= C(x+y) + G(\tanh^{-1}(x) + \tanh^{-1}(y) - \tanh^{-1}(x+y), x+y) \end{aligned}$$

and substituting in for C and G ,

$$\begin{aligned}
&= \frac{1}{2}(x+y)^2 - \frac{1}{2}(1+x+y)\log(1+x+y) - \frac{1}{2}(1-x-y)\log(1-x-y) \\
&\quad + \log(\cosh(\tanh^{-1}(x) + \tanh^{-1}(y) - \tanh^{-1}(x+y))) \\
&\quad - (x+y)(\tanh^{-1}(x) + \tanh^{-1}(y) - \tanh^{-1}(x+y)) \\
&\quad + \log(1 + (x+y)\tanh(\tanh^{-1}(x) + \tanh^{-1}(y) - \tanh^{-1}(x+y))), \quad (10.128)
\end{aligned}$$

which after some algebra is equivalent to

$$= C(x) + C(y) - x \tanh^{-1}(y) - y \tanh^{-1}(x) + xy + \log(1 + xy). \quad (10.129)$$

Expanding the exponential generating function coefficients of the final line then gives the result. \square

Equipped with these results, let us summarize our analysis below, giving a combined error bound between \mathbf{Z}^{main} and \mathbf{Z}^{tied} as well as the final form of the latter.

Corollary 10.7.6. *Suppose that $\epsilon_{\text{tree}}(\mathbf{M}; d) \leq 1$. We have*

$$\mathbf{Z}_{S,T}^{\text{tied}} = \sum_{\substack{F \in \mathcal{F}(|S|, |T|) \\ F \text{ balanced bowtie forest}}} \xi(F) \cdot \mathbf{Z}_{S,T}^F, \quad (10.130)$$

and this matrix satisfies

$$\|\mathbf{Z}^{\text{main}} - \mathbf{Z}^{\text{tied}}\| \leq (6d)^{10d} \|\mathbf{M}\|^{3d} \epsilon_{\text{tree}}(\mathbf{M}; d). \quad (10.131)$$

Proof. The first formula follows from Lemma 10.7.5. For the norm bound, we apply the result of Lemma 10.7.4 to each CGM term in each block of \mathbf{Z}^{main} and use the norm bound of

Proposition 10.12.4, which gives

$$\begin{aligned} \|\mathbf{Z}^{\text{main}} - \mathbf{Z}^{\text{tied}}\| &\leq \sum_{\ell, m=0}^d (\ell + m) (2\|\mathbf{M}\|)^{\frac{3}{2}(\ell+m)} \epsilon_{\text{tree}}(\mathbf{M}; (\ell + m)/2) \sum_{F \in \mathcal{F}(\ell, m)} |\mu(F)| \\ &\leq (2d)(2\|\mathbf{M}\|)^{3d} \epsilon_{\text{tree}}(\mathbf{M}; d) \sum_{\ell, m=0}^d |\mathcal{F}(\ell, m)| \max_{F \in \mathcal{F}(\ell, m)} |\mu(F)| \end{aligned}$$

and using Propositions 10.12.11 and 10.12.5 to bound $|\mathcal{F}(\ell, m)|$ and $|\mu(F)|$ respectively, we finish

$$\leq (2d)(2\|\mathbf{M}\|)^{3d} \epsilon_{\text{tree}}(\mathbf{M}; d) \cdot d^2 (6d)^{3d} (6d)^{6d}, \quad (10.132)$$

and the remaining bound follows from elementary manipulations. \square

10.8 APPROXIMATE GRAM FACTORIZATION: PROOF OF LEMMA 10.6.2

To prove a lower bound on $\lambda_{\min}(\mathbf{Z}^{\text{main}})$, our strategy will be to justify the equality

$$\tilde{\mathbb{E}}^{\text{main}}[h_S^\dagger(\mathbf{x})h_T^\dagger(\mathbf{x})] \approx \langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle_{\partial}, \quad (10.133)$$

that was suggested by our calculations in Chapter 8. The right-hand side is block diagonal (since homogeneous polynomials of different degrees are apolar), so our block diagonalization of the left-hand side is a useful start. To continue, we follow the same plan for the right-hand side in this section as we did for the left-hand side in the previous section: we (1) express the Gram matrix as a linear combination of CGMs, (2) show that the corresponding ribbon diagrams may be simplified to the same bowtie forests from Definition 10.7.3, and (3) perform a combinatorial analysis of the coefficients attached to each bowtie forest after

the simplification.

10.8.1 PARTITION TRANSPORT PLANS AND RIBBON DIAGRAMS

We first describe the class of ribbon diagrams that will arise in expanding the inner products above, which may be viewed as encoding the following kind of combinatorial object.

Definition 10.8.1 (Partition transport plan). *For a pair of partitions $\sigma, \tau \in \text{Part}([d])$, we write $\text{Plans}(\sigma, \tau)$ for the set of matrices $\mathbf{D} \in \mathbb{N}^{\sigma \times \tau}$ where the sum of each row indexed by $A \in \sigma$ is $|A|$, and the sum of each column indexed by $B \in \tau$ is $|B|$.*

We borrow the terms “transport” and “plan” from the theory of optimal transport [Vil08], since \mathbf{D} may be seen as specifying a protocol for moving masses corresponding to the part sizes of σ and τ . These same matrices also play a crucial role in the Robinson-Schensted-Knuth correspondence of representation theory and the combinatorics of Young tableaux (where they are sometimes called *generalized permutations*) [Ful97]. It is an intriguing question for future investigation whether this connection can shed light on our use of $\text{Plans}(\sigma, \tau)$.

We encode a pair of partitions and a partition transport plan between them into a ribbon diagram in the following way.

Definition 10.8.2 (Partition transport ribbon diagram). *Suppose that $\sigma, \tau \in \text{Part}([d])$ and $\mathbf{D} \in \text{Plans}(\sigma, \tau)$. Then, let $G = G(\sigma, \tau, \mathbf{D})$ be the associated partition transport ribbon diagram, with graphical structure defined as follows:*

- $\mathcal{L}(G)$ and $\mathcal{R}(G)$ are two disjoint sets, each labelled by $1, \dots, d$.
- $V^\square(G)$ contains one vertex for each part of $\sigma[\geq 2]$ and each part of $\tau[\geq 2]$.
- Whenever $i \in A \in \sigma$, then the vertex labelled i in \mathcal{L} and the \square vertex corresponding to

A have an edge between them. Likewise, whenever $j \in B \in \tau$, then the vertex labelled j in \mathcal{R} and the \square vertex corresponding to B have an edge between them.

- When $A \in \sigma[\geq 2]$ and $B \in \tau[\geq 2]$, then there are $D_{A,B}$ parallel edges between the corresponding \square vertices.
- When $A = \{i\} \in \sigma[1]$, $B \in \tau[\geq 2]$, and $D_{A,B} = 1$, then there is an edge between the vertex labelled i in \mathcal{L} and the \square vertex corresponding to B . Likewise, when $B = \{j\} \in \tau[1]$, $A \in \sigma[\geq 2]$, and $D_{A,B} = 1$, then there is an edge between the vertex labelled j in \mathcal{R} and the \square vertex corresponding to A .
- When $A = \{i\} \in \sigma[1]$, $B = \{j\} \in \tau[1]$, and $D_{A,B} = 1$, then there is an edge between the vertex labelled i in \mathcal{L} and the vertex labelled j in \mathcal{R} .

See Figure 10.4 for an example featuring all of these situations that may be clearer than the written description.

This formalism allows us to make the following compact CGM description of the Gram matrix of the non-lowered multiharmonic polynomials.

Proposition 10.8.3 (CGM expansion of multiharmonic Gram matrix). *Define $\mathbf{Y} \in \mathbb{R}^{\binom{[n]}{\leq d} \times \binom{[n]}{\leq d}}$ by*

$$Y_{S,T} = \langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle_{\partial}. \quad (10.134)$$

Then, \mathbf{Y} is block diagonal, with diagonal blocks $\mathbf{Y}^{[d,d]} \in \mathbb{R}^{\binom{[n]}{d} \times \binom{[n]}{d}}$ given by

$$\mathbf{Y}^{[d,d]} = \sum_{\sigma, \tau \in \text{Part}([d])} \prod_{R \in \sigma + \tau} (-1)^{|R|-1} (|R|-1)! (|R|)! \sum_{D \in \text{Plans}(\sigma, \tau)} \frac{1}{D!} \mathbf{Z}^{G(\sigma, \tau, D)}(\mathbf{M}), \quad (10.135)$$

where $D! := \prod_{A \in \sigma} \prod_{B \in \tau} D_{A,B}!$.




Figure 10.4: Partition transport plan and ribbon diagram. We illustrate an example of two partitions $\sigma, \tau \in \text{Part}([12])$, a partition transport plan $D \in \text{Plans}(\sigma, \tau)$, and the associated partition transport ribbon diagram $G^{(\sigma, \tau, D)}$.

Proof. We expand directly by linearity:

$$\begin{aligned}
& \langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle_\partial \\
&= \left\langle \sum_{\sigma \in \text{Part}(S)} \prod_{A \in \sigma} (-1)^{|A|-1} (|A|-1)! q_A(\mathbf{V}^\top \mathbf{z}), \sum_{\tau \in \text{Part}(T)} \prod_{B \in \tau} (-1)^{|B|-1} (|B|-1)! q_B(\mathbf{V}^\top \mathbf{z}) \right\rangle_\partial \\
&= \sum_{\substack{\sigma \in \text{Part}(S) \\ \tau \in \text{Part}(T)}} \prod_{R \in \sigma + \tau} (-1)^{|R|-1} (|R|-1)! \left\langle \prod_{A \in \sigma} q_A(\mathbf{V}^\top \mathbf{z}), \prod_{B \in \tau} q_B(\mathbf{V}^\top \mathbf{z}) \right\rangle_\partial \tag{10.136}
\end{aligned}$$

The remaining polynomials we may further expand

$$\begin{aligned}
\prod_{A \in \sigma} q_A(\mathbf{V}^\top \mathbf{z}) &= \prod_{A = \{i\} \in \sigma[1]} \langle \mathbf{v}_i, \mathbf{z} \rangle \cdot \prod_{A \in \sigma[\geq 2]} \left\{ \sum_{a=1}^n \prod_{i \in A} M_{ai} \cdot \langle \mathbf{v}_a, \mathbf{z} \rangle^{|A|} \right\} \\
&= \prod_{A = \{i\} \in \sigma[1]} \langle \mathbf{v}_i, \mathbf{z} \rangle \sum_{a \in [n]^{\sigma[\geq 2]}} \prod_{A \in \sigma[\geq 2]} \prod_{i \in A} M_{i, a(A)} \cdot \langle \mathbf{v}_{a(A)}, \mathbf{z} \rangle^{|A|} \\
&= \sum_{a \in [n]^{\sigma[\geq 2]}} \prod_{A \in \sigma} \prod_{i \in A} M_{i, f_a(A)} \cdot \langle \mathbf{v}_{f_a(A)}, \mathbf{z} \rangle^{|A|}, \tag{10.137}
\end{aligned}$$

where we define $f_a(A) = a(A)$ if $|A| \geq 2$, and $f_a(A) = i$ if $A = \{i\}$. Likewise, for $\mathbf{b} \in [n]^{\tau[\geq 2]}$, as will arise in q_B , we set $g_b(B) = b(B)$ if $|B| \geq 2$, and $g_b(B) = j$ if $B = \{j\}$. Thus we may expand the remaining inner product from before as

$$\begin{aligned} & \left\langle \prod_{A \in \sigma} q_A(\mathbf{V}^\top \mathbf{z}), \prod_{B \in \tau} q_B(\mathbf{V}^\top \mathbf{z}) \right\rangle_{\partial} \\ &= \sum_{\substack{\mathbf{a} \in [n]^{\sigma[\geq 2]} \\ \mathbf{b} \in [n]^{\tau[\geq 2]}}} \prod_{A \in \sigma} \prod_{i \in A} M_{i, f_a(A)} \cdot \prod_{B \in \tau} \prod_{j \in B} M_{j, g_b(B)} \cdot \left\langle \prod_{A \in \sigma} \langle \mathbf{v}_{f_a(A)}, \mathbf{z} \rangle^{|A|}, \prod_{B \in \tau} \langle \mathbf{v}_{f_b(B)}, \mathbf{z} \rangle^{|B|} \right\rangle_{\partial}. \end{aligned} \quad (10.138)$$

Finally, this remaining inner product we expand by the product rule, executing which gives rise to the summation over partition transport plans that arises in our result (this calculation is easy to verify by induction, or may be seen as an application of the more general Faà di Bruno formula; see, e.g., [Har06]):

$$\begin{aligned} & \left\langle \prod_{A \in \sigma} \langle \mathbf{v}_{f(A)}, \mathbf{z} \rangle^{|A|}, \prod_{B \in \tau} \langle \mathbf{v}_{f(B)}, \mathbf{z} \rangle^{|B|} \right\rangle_{\partial} \\ &= \prod_{A \in \sigma} \langle \mathbf{v}_{f(A)}, \boldsymbol{\partial} \rangle^{|A|} \prod_{B \in \tau} \langle \mathbf{v}_{f(B)}, \mathbf{z} \rangle^{|B|} \\ &= \prod_{B \in \tau} (|B|)! \sum_{D \in \text{Plans}(\sigma, \tau)} \prod_{A \in \sigma} \left(D_{A, B_1}^{|A|} \cdots D_{A, B_{|\tau|}}^{|A|} \right) \prod_{A \in \sigma} \prod_{B \in \tau} (M_{f(A), f(B)})^{D_{A, B}} \\ &= \prod_{R \in \sigma + \tau} (|R|)! \sum_{D \in \text{Plans}(\sigma, \tau)} \prod_{A \in \sigma} \prod_{B \in \tau} \frac{(M_{f(A), f(B)})^{D_{A, B}}}{D_{A, B}!}, \end{aligned} \quad (10.139)$$

where we remark that in the final expression here we restore the symmetry between σ and τ , which was briefly broken to perform the calculation. The final result then follows from combining the preceding equations (10.136), (10.138), and (10.139), identifying the summation occurring as precisely that defined by the partition transport ribbon diagram corresponding to (σ, τ, D) . \square




Figure 10.5: Tying stretched forest and partition transport ribbon diagrams. We illustrate the key diagrammatic idea of the argument proving Lemma 10.6.2, that both the stretched forest ribbon diagrams appearing in $\mathcal{Z}^{\text{main}}$ and the partition transport ribbon diagrams appearing in \mathcal{Y} may be “tied” to form the same bowtie forest ribbon diagrams.

10.8.2 TYING BOUND

We now describe the “tied” version of a partition transport ribbon diagram and bound the error in operator norm incurred by the tying procedure.

Lemma 10.8.4 (Partition transport ribbon diagrams: tying bound). *Let $\sigma, \tau \in \text{Part}([d])$ and $\mathbf{D} \in \text{Plans}(\sigma, \tau)$. Let $G = G(\sigma, \tau, \mathbf{D}) \in \mathcal{F}(d, d)$ be the associated partition transport ribbon diagram. Let $\text{tie}(G)$ denote the diagram obtained from G by contracting all connected components that are not pairs to a bowtie. Then, $\text{tie}(G)$ is a balanced bowtie forest, and*

$$\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| \leq d^2 \|\mathbf{M}\|^{3d} (\epsilon_{\text{pow}}(\mathbf{M}) + \epsilon_{\text{offdiag}}(\mathbf{M}) + \epsilon_{\text{corr}}(\mathbf{M})). \quad (10.140)$$

The proof considers various cases depending on the graph structure of G , but is similar in principle to the proof of Lemma 10.7.4, the tying bound for stretched forest diagrams—we again factorize CGMs and argue that the “inner” ribbon diagrams may be collapsed without incurring a substantial error in operator bound.

Proof of Lemma 10.8.4. We recall the statement: let $\sigma, \tau \in \text{Part}([d])$, let $\mathbf{D} \in \text{Plans}(\sigma, \tau)$, and let $G = G(\sigma, \tau, \mathbf{D})$ be the associated partition transport ribbon diagram. We will show

the slightly stronger bound,

$$\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| \leq d \|\mathbf{M}\|^{3d} \left(\epsilon_{\text{pow}}(\mathbf{M}) + d\epsilon_{\text{offdiag}}(\mathbf{M}) + d\epsilon_{\text{corr}}(\mathbf{M}) \right). \quad (10.141)$$

As in the previous proof, let us first suppose that G is connected. We will then show the bound

$$\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| \leq \|\mathbf{M}\|^{3d} \left(\epsilon_{\text{pow}}(\mathbf{M}) + d\epsilon_{\text{offdiag}}(\mathbf{M}) + d\epsilon_{\text{corr}}(\mathbf{M}) \right), \quad (10.142)$$

the same as the above but without the leading factor of d . Since there are exactly $3d$ edges in any connected component of a partition transport ribbon diagram such that the component has $2d$ leaves, the final result will then follow by applying Propositions 10.13.13 and 10.13.4.

Let us write $\partial^\square \mathcal{L}, \partial^\square \mathcal{R} \subset V^\square$ for the sets of \square vertices with a neighbor in \mathcal{L} and \mathcal{R} , respectively. We have $\partial^\square \mathcal{L} \cup \partial^\square \mathcal{R} = V^\square$, and we observe that $\partial^\square \mathcal{L} \cap \partial^\square \mathcal{R} = \emptyset$ if and only if both σ and τ contain no singletons.

Case 1: $d = 1$. This is only possible if $\sigma = \tau = \{\{1\}\}$ and $\mathbf{D} = [1]$. In this case, $G = \text{tie}(G)$ (both consist of two leaves connected to one another), so $\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| = 0$.

Case 2: $\partial^\square \mathcal{L} \cap \partial^\square \mathcal{R} \neq \emptyset$ and $d \geq 2$. We argue by induction on $|V^\square(G)|$ that the following bound holds in this case:

$$\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| \leq \|\mathbf{M}\|^{3d} \cdot |V^\square(G)| \cdot \epsilon_{\text{offdiag}}(\mathbf{M}) \quad (10.143)$$

If $|V^\square(G)| = 1$, then $G = \text{tie}(G)$, so $\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| = 0$. Now, suppose we have established the result for all partition transport ribbon diagrams satisfying the assumptions of this case with $|V^\square| \leq m$, and have $|V^\square(G)| = m + 1 > 1$. Let $v \in \partial^\square \mathcal{L} \cap \partial^\square \mathcal{R}$. We apply the factorization of Proposition 10.13.11 with $A = \mathcal{L}$, $B = V^\square$, and $C = \mathcal{R}$, which factor-

izes $\mathbf{Z}^G = \mathbf{Z}^{G[A]} \mathbf{Z}^{G[B]} \mathbf{Z}^{G[C]}$. Since $v \in \mathcal{L}(G[B]) \cap \mathcal{R}(G[B])$, $\mathbf{Z}^{G[B]}$ is the direct sum of n further CGMs where v is pinned to each possible value in $[n]$. Let $G'[B]$ denote the ribbon diagram formed from $G[B]$ by relabelling all edges between v and all of its neighbors with the identity matrix. When v is pinned, the factors contributed by these edges are constant factors in each direct summand CGM, so we have $\|\mathbf{Z}^{G[B]} - \mathbf{Z}^{G'[B]}\| \leq \|\mathbf{M}\|^{|E(G[B])|} \epsilon_{\text{offdiag}}(\mathbf{M})$. Now, let G' denote the ribbon diagram formed from G by contracting all edges between v and all of its neighbors in V^\square . Then, we have $\mathbf{Z}^{G'} = \mathbf{Z}^{G[A]} \mathbf{Z}^{G'[B]} \mathbf{Z}^{G[C]}$, so $\|\mathbf{Z}^{G'} - \mathbf{Z}^G\| \leq \|\mathbf{M}^{|E(G)|}\| \epsilon_{\text{offdiag}}(\mathbf{M}) \leq \|\mathbf{M}\|^{3d} \epsilon_{\text{offdiag}}(\mathbf{M})$ by Proposition 10.13.13 and Corollary 10.12.10. Since $\text{tie}(G') = \text{tie}(G)$, we find

$$\begin{aligned} & \|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| \\ & \leq \|\mathbf{Z}^G - \mathbf{Z}^{G'}\| + \|\mathbf{Z}^{G'} - \mathbf{Z}^{\text{tie}(G')}\| \\ & \leq \|\mathbf{M}\|^{3d} \epsilon_{\text{offdiag}}(\mathbf{M}) + \|\mathbf{M}^{3d}\| \cdot |V^\square(G')| \cdot \epsilon_{\text{offdiag}}(\mathbf{M}) \quad (\text{inductive hypothesis}) \end{aligned}$$

Since $|V^\square(G')| < |V^\square(G)|$ by construction, the proof of (10.143) is complete.

Lastly, since each \square vertex of G corresponds to a part of σ or τ having size at least 2, we have $|V^\square(G)| \leq d$, so we obtain the simpler version

$$\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| \leq \|\mathbf{M}\|^{3d} d \epsilon_{\text{offdiag}}(\mathbf{M}). \quad (10.144)$$

Case 3: $\partial^\square \mathcal{L} \cap \partial^\square \mathcal{R} = \emptyset$, $d \geq 2$, and a row or column of \mathbf{D} has only one non-zero entry. Let us suppose without loss of generality that it is a row of \mathbf{D} that has the specified property, which corresponds to a part $S \in \sigma$. Let $v \in V^\square$ be the associated \square vertex in G . The given condition means that v has only one neighbor in $\partial^\square \mathcal{R}$, which we call w , and that there are $|S| \geq 2$ parallel edges between v and w . Let G' denote the diagram where v and w are

identified. Then, by Proposition 10.13.13, we have

$$\|\mathbf{Z}^G - \mathbf{Z}^{G'}\| \leq \|\mathbf{M}\|^{3d} \epsilon_{\text{pow}}(\mathbf{M}). \quad (10.145)$$

We note that $\text{tie}(G') = \text{tie}(G)$, and Case 2 applies to G' (indeed, G' is the partition transport ribbon diagram formed by replacing the part S of σ with $|S|$ singletons that are all transported to the part of τ corresponding to w). Therefore, using that result, we conclude

$$\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| \leq \|\mathbf{M}\|^{3d} (\epsilon_{\text{pow}}(\mathbf{M}) + d\epsilon_{\text{offdiag}}(\mathbf{M})). \quad (10.146)$$

Case 4: $\partial^\square \mathcal{L} \cap \partial^\square \mathcal{R} = \emptyset$, $d \geq 2$, and no row or column of \mathbf{D} has only one non-zero entry.

We argue by induction on $|V^\square(G)|$ that the following bound holds:

$$\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| \leq \|\mathbf{M}\|^{3d} (\epsilon_{\text{pow}}(\mathbf{M}) + |V^\square(G)| \epsilon_{\text{corr}}(\mathbf{M})). \quad (10.147)$$

We cannot have $|V^\square(G)| = 1$, since then the single \square vertex would need to belong to $\partial^\square \mathcal{L} \cap \partial^\square \mathcal{R}$. Thus the base case is $|V^\square(G)| = 2$. In this case, G consists of two \square vertices, each connected to d leaves, and with d parallel edges between them. We apply the factorization of Proposition 10.13.11 with $A = \mathcal{L}$, $B = V^\square$, and $C = \mathcal{R}$. Then, $G[B]$ consists of two vertices, one in $\mathcal{L}(G[B])$ and one in $\mathcal{R}(G[B])$, connected by d parallel edges. Writing $G'[B]$ for the diagram where these edges are replaced by a single one labelled with the identity, we then have $\|\mathbf{Z}^{G[B]} - \mathbf{Z}^{G'[B]}\| = \|\mathbf{M}^{\circ d} - \mathbf{I}_n\| \leq \epsilon_{\text{pow}}(\mathbf{M})$. And, since $\mathbf{Z}^{\text{tie}(G)} = \mathbf{Z}^{G[A]} \mathbf{Z}^{G'[B]} \mathbf{Z}^{G[C]}$, using Proposition 10.13.13 we find

$$\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| \leq \|\mathbf{M}\|^{3d} \epsilon_{\text{pow}}(\mathbf{M}). \quad (10.148)$$

(We could reduce the exponent to $2d$ here, but accept the insignificant additional slack to

make the final expression cleaner.)

Now, suppose we have established the result for all partition transport ribbon diagrams G with $|V^\square(G)| \leq m$, and have G with $|V^\square(G)| = m + 1 > 2$. Since G is connected, has $\partial^\square \mathcal{L} \cap \partial^\square \mathcal{R} = \emptyset$, and $d > 1$, all parts of σ and τ must have size at least 2. Moreover, since $|V^\square(G)| > 2$, we must have either $|\sigma| > 1$ or $|\tau| > 1$. Let us suppose, without loss of generality, that $|\sigma| > 1$ (otherwise we may reverse the roles of σ and τ , which amounts to transposing the ribbon diagram G and the associated CGM).

Choose any part $S \in \sigma$, and let $v \in V^\square(G)$ be the associated \square vertex. We apply the factorization of Proposition 10.13.11 with $A = \mathcal{L} \cup (\partial^\square \mathcal{L} \setminus \{v\})$, $B = \partial^\square \mathcal{R} \cup \{v\}$, and $C = \mathcal{R}$. Consider the diagram $G[B]$. We have $\mathcal{R}(G[B]) = \partial^\square \mathcal{R}$ and $V^\square(G[B]) = \emptyset$. Moreover, by the assumption of this case, every vertex of $\partial^\square \mathcal{R}$ has a neighbor in $\mathcal{L} \setminus \{v\}$. Therefore, $\mathcal{L}(G[B]) = \partial^\square \mathcal{R} \cup \{v\}$. In particular, $\mathcal{R}(G[B]) \subseteq \mathcal{L}(G[B])$, so $\mathcal{R}(G[B]) = \mathcal{L}(G[B]) \cap \mathcal{R}(G[B]) = \partial^\square \mathcal{R}$.

Following the pinning transformation of Proposition 10.13.10, after a suitable permutation, the CGM of $G[B]$ will then be the direct sum of column vectors $\mathbf{v}_s \in \mathbb{R}^n$, indexed by $s \in [n]^{\partial^\square \mathcal{R}}$, where

$$(\mathbf{v}_s)_i = \prod_{x \in \partial^\square v} M_{i,s(x)}, \quad (10.149)$$

where the product is understood to repeat vertices x when v is connected to x with multiple edges. In particular, we have

$$\|\mathbf{v}_s\|_2^2 = \sum_{i=1}^n \prod_{x \in \partial^\square v} M_{i,s(x)}^2, \quad (10.150)$$

and since, again by the assumption of this case, v has at least two different neighbors in $\partial^\square \mathcal{R}$, we have $\|\mathbf{v}_s\|_2 \leq \epsilon_{\text{corr}}(\mathbf{M})$ whenever s is not constant on $\partial^\square v$. Thus letting G' be the diagram formed from G by identifying all neighbors of v , and applying Proposition 10.13.13, we find that

$$\|\mathbf{Z}^G - \mathbf{Z}^{G'}\| \leq \|\mathbf{M}\|^{3d} \epsilon_{\text{corr}}(\mathbf{M}), \quad (10.151)$$

so by the inductive hypothesis,

$$\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| \leq \|\mathbf{M}\|^{3d} \epsilon_{\text{corr}}(\mathbf{M}) + \|\mathbf{M}\|^{3d} (\epsilon_{\text{pow}}(\mathbf{M}) + |V^\square(G')| \epsilon_{\text{corr}}(\mathbf{M})), \quad (10.152)$$

and since $|V^\square(G')| < |V^\square(G)|$, the induction is complete. Finally, using again that $|V^\square(G)| \leq d$ since all parts of σ and τ have size at least 2 in this case, we find the looser bound

$$\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\| \leq \|\mathbf{M}\|^{3d} (\epsilon_{\text{pow}}(\mathbf{M}) + d \epsilon_{\text{corr}}(\mathbf{M})). \quad (10.153)$$

Conclusion. In the four cases considered above, we have shown that either $\|\mathbf{Z}^G - \mathbf{Z}^{\text{tie}(G)}\|$ is zero, or is bounded as in (10.144), (10.146), and (10.153). We then observe that the “master bound” in the statement is larger than each of these, completing the proof. \square

10.8.3 SIMPLIFICATION OF COEFFICIENTS

Next, we describe the result of tying every ribbon diagram in \mathbf{Y} . As before, this involves a combinatorial calculation to sum over all diagrams that produce a given bowtie forest upon tying. As we would hope, the resulting coefficients are the same as those arising in \mathbf{Z}^{tied} .

Lemma 10.8.5 (Partition transport ribbon diagrams: combinatorial reduction). *For $d \in \mathbb{N}$, define $\mathbf{Y}^{\text{tied}[d,d]} \in \mathbb{R}^{\binom{[n]}{d} \times \binom{[n]}{d}}$ by*

$$\begin{aligned} \mathbf{Y}^{\text{tied}[d,d]} := & \sum_{\sigma, \tau \in \text{Part}([d])} \prod_{A \in \sigma + \tau} (-1)^{|A|-1} (|A|-1)! (|A|)! \\ & \sum_{D \in \text{Plans}(\sigma, \tau)} \frac{1}{D!} \mathbf{Z}^{\text{tie}(G(\sigma, \tau, D))}(\mathbf{M}). \end{aligned} \quad (10.154)$$

Then,

$$\mathbf{Y}^{\text{tied}[d,d]} = \sum_{\substack{F \in \mathcal{F}(d,d) \\ F \text{ balanced bowtie forest}}} \xi(F) \cdot \mathbf{Z}^F, \quad (10.155)$$

where $\xi(F)$ is the same coefficient from Lemma 10.7.5, given by

$$\xi(F) = \mathbb{1}\{F \text{ balanced}\} \prod_{\substack{C \in \text{conn}(F) \\ C \text{ balanced bowtie} \\ \text{on } 2k \text{ leaves}}} (-1)^{k-1} (k-1)! k!. \quad (10.156)$$

In particular, the direct sum of $\mathbf{Y}^{\text{tied}[d', d']}$ over $0 \leq d' \leq d$ equals \mathbf{Z}^{tied} as defined in (10.97).

Remark 10.8.6. We note that the fact that the combinatorial quantities in Lemma 10.7.5 earlier, sums of Möbius functions of stretched forests, and those in Lemma 10.8.5 above, sums of combinatorial coefficients of partition transport plans, are equal is quite surprising. We isolate this fact to emphasize its unusual form:

$$\sum_{\substack{F \in \mathcal{F}(\ell, m) \\ F \text{ stretched}}} \mu(F) = \sum_{\substack{\sigma \in \text{Part}([\ell]) \\ \tau \in \text{Part}([m])}} \prod_{A \in \sigma + \tau} (-1)^{|A|-1} (|A|-1)! (|A|)! \sum_{D \in \text{Plans}(\sigma, \tau)} \frac{1}{D!}. \quad (10.157)$$

Our proofs, unfortunately, give little insight as to why this should be the case, instead showing, in an especially technical manner for the left-hand side, that both sides equal another quantity. It would be interesting to find a combinatorial or order-theoretic argument explaining this coincidence, which is crucial to our argument, more directly.

Proof of Lemma 10.8.5. Let us say that a partition transport plan $D \in \text{Plans}(\sigma, \tau)$ is *connected* if its diagram $G^{(\sigma, \tau, D)}$ is connected, and refer to the connected components of D , denoted $\text{conn}(D)$, as the subsets of $\sigma + \tau$ that belong to each connected component of the diagram. As in the previous Lemma, both sides of the result factorize over connected components, so it suffices to show that

$$\sum_{\sigma, \tau \in \text{Part}([d])} (-1)^{|\sigma|+|\tau|} \prod_{A \in \sigma + \tau} (|A|-1)! (|A|)! \sum_{\substack{D \in \text{Plans}(\sigma, \tau) \\ D \text{ connected}}} \frac{1}{D!} = (-1)^{d-1} (d-1)! d!, \quad (10.158)$$

Let us first work with the innermost sum. For σ, τ arbitrary partitions, writing $\|\sigma\| =$

$\sum_{S \in \sigma} |S|$ and likewise for $\|\tau\|$, by the inner product calculation from Proposition 10.8.3 we have

$$\sum_{D \in \text{Plans}(\sigma, \tau)} \frac{1}{D!} = \mathbb{1}\{\|\sigma\| = \|\tau\|\} \frac{d!}{\prod_{A \in \sigma + \tau} (|A|)!}. \quad (10.159)$$

We now compute the restriction to connected D using Möbius inversion (a similar calculation to the proof of Lemma 10.3.11). For $\pi \in \text{Part}(\sigma + \tau)$, let us define

$$b(\pi) := \sum_{\substack{D \in \text{Plans}(\sigma, \tau) \\ \text{conn}(D) = \pi}} \frac{1}{D!}. \quad (10.160)$$

Then, the quantity we are interested in is $b(\{\sigma + \tau\})$. The downward sums of $b(\pi)$ are

$$\begin{aligned} c(\pi) &:= \sum_{\pi' \leq \pi} b(\pi') \\ &= \sum_{\substack{D \in \text{Plans}(\sigma, \tau) \\ \text{conn}(D) \leq \pi}} \frac{1}{D!} \\ &= \prod_{\rho \in \pi} \left(\sum_{D \in \text{Plans}(\sigma \cap \rho, \tau \cap \rho)} \frac{1}{D!} \right) \\ &= \mathbb{1}\{\|\sigma \cap \rho\| = \|\tau \cap \rho\| \text{ for all } \rho \in \pi\} \frac{\prod_{\rho \in \pi} (\|\rho\|/2)!}{\prod_{A \in \sigma + \tau} (|A|)!}. \end{aligned} \quad (10.161)$$

Therefore, by Möbius inversion (in the poset $\text{Part}(\sigma + \tau)$),

$$\begin{aligned} &\sum_{\substack{D \in \text{Plans}(\sigma, \tau) \\ D \text{ connected}}} \frac{1}{D!} \\ &= b(\{\sigma + \tau\}) \\ &= \sum_{\pi \in \text{Part}(\sigma + \tau)} \mu_{\text{Part}}(\pi, \{\sigma + \tau\}) c(\pi) \\ &= \frac{1}{\prod_{A \in \sigma + \tau} (|A|)!} \sum_{\substack{\pi \in \text{Part}(\sigma + \tau) \\ \|\sigma \cap \rho\| = \|\tau \cap \rho\| \text{ for all } \rho \in \pi}} (-1)^{|\pi|-1} (|\pi| - 1)! \prod_{\rho \in \pi} (\|\rho\|/2)! \end{aligned} \quad (10.162)$$

Now, we substitute this into the left-hand side in the initial statement:

$$\begin{aligned}
& \sum_{\sigma, \tau \in \text{Part}([d])} (-1)^{|\sigma|+|\tau|} \prod_{A \in \sigma + \tau} (|A| - 1)! (|A|)! \sum_{\substack{D \in \text{Plans}(\sigma, \tau) \\ D \text{ connected}}} \frac{1}{D!} \\
&= \sum_{\sigma, \tau \in \text{Part}([d])} (-1)^{|\sigma|+|\tau|} \prod_{A \in \sigma + \tau} (|A| - 1)! \\
& \quad \sum_{\substack{\pi \in \text{Part}(\sigma + \tau) \\ \|\sigma \cap \rho\| = \|\tau \cap \rho\| \text{ for all } \rho \in \pi}} (-1)^{|\pi|-1} (|\pi| - 1)! \prod_{\rho \in \pi} (\|\rho\|/2)!
\end{aligned}$$

and exchanging the order of summation,

$$\begin{aligned}
&= \sum_{\substack{\pi \in \text{Part}([2d]) \\ |R \cap \{1, \dots, d\}| = |R \cap \{d+1, \dots, 2d\}| \\ \text{for all } R \in \pi}} (-1)^{|\pi|-1} (|\pi| - 1)! \prod_{R \in \pi} (|R|/2)! \\
& \quad \sum_{\substack{\sigma \in \text{Part}(\{1, \dots, d\}) \\ \tau \in \text{Part}(\{d+1, \dots, 2d\}) \\ \sigma + \tau \leq \pi}} (-1)^{|\sigma|+|\tau|} \prod_{A \in \sigma + \tau} (|A| - 1)!. \tag{10.163}
\end{aligned}$$

We again think in terms of Möbius functions, but now on a different poset: on the product poset $\text{Part}(\{1, \dots, d\}) \times \text{Part}(\{d+1, \dots, 2d\})$, the Möbius function is (see [Rot64] for this fact)

$$\mu_{\text{Part} \times \text{Part}}(\{(\{1\}, \dots, \{\ell\}), \{(\ell+1), \dots, \{2\ell\})\}, (\sigma, \tau)) = (-1)^{|\sigma|+|\tau|} \prod_{A \in \sigma + \tau} (|A| - 1)!, \tag{10.164}$$

and $\{(\sigma, \tau) : \sigma + \tau \leq \pi\}$ is an interval. Therefore, the inner sum above is zero unless every part of π has size two, so that π is a matching of $\{1, \dots, d\}$ with $\{d+1, \dots, 2d\}$, in which case the inner sum is 1. There are $d!$ such matchings π , so we find that the above expression equals $(-1)^{d-1} (d-1)! d!$, giving the result. \square

We again summarize our findings below.

Corollary 10.8.7. $\|Y - Z^{\text{tied}}\| \leq d^{5d} \|M\|^{3d} (\epsilon_{\text{pow}}(M) + \epsilon_{\text{offdiag}}(M) + \epsilon_{\text{corr}}(M)).$

Proof. Since \mathbf{Y} and \mathbf{Z}^{tied} are both block diagonal, it suffices to consider a diagonal block indexed by $\binom{[n]}{d}$. Since by Lemma 10.8.5 this block of \mathbf{Z}^{tied} is $\mathbf{Y}^{\text{tied}[d,d]}$, this amounts to bounding $\|\mathbf{Y}^{[d,d]} - \mathbf{Y}^{\text{tied}[d,d]}\|$. Applying triangle inequality and Lemma 10.8.4, we find

$$\begin{aligned} & \|\mathbf{Y}^{[d,d]} - \mathbf{Y}^{\text{tied}[d,d]}\| \\ & \leq \sum_{\sigma, \tau \in \text{Part}([d])} \prod_{A \in \sigma + \tau} (|A| - 1)! (|A|)! \sum_{D \in \text{Plans}(\sigma, \tau)} \frac{1}{D!} \|\mathbf{Z}^{G(\sigma, \tau, D)} - \mathbf{Z}^{\text{tie}(G(\sigma, \tau, D))}\| \\ & \leq d^2 \|\mathbf{M}\|^{3d} (\epsilon_{\text{pow}}(\mathbf{M}) + \epsilon_{\text{offdiag}}(\mathbf{M}) + \epsilon_{\text{corr}}(\mathbf{M})) \cdot (d - 1)! d! \cdot \sum_{\sigma, \tau \in \text{Part}([d])} |\text{Plans}(\sigma, \tau)| \end{aligned}$$

and bounding $|\text{Part}([d])|$ and $|\text{Plans}(\sigma, \tau)|$ using Propositions 10.12.7 and 10.12.8 respectively and noting that $d \cdot d! \leq d^d$ we have

$$\leq \|\mathbf{M}\|^{3d} (\epsilon_{\text{pow}}(\mathbf{M}) + \epsilon_{\text{offdiag}}(\mathbf{M}) + \epsilon_{\text{corr}}(\mathbf{M})) \cdot d^{2d} \cdot d^{2d} d^d, \quad (10.165)$$

and the result follows. \square

Finally, what will be the crucial feature of \mathbf{Y} for us is that its spectrum is bounded below. Since \mathbf{Y} is formed by definition as a Gram matrix we will certainly have $\mathbf{Y} \geq 0$; however, as we show below, we moreover have $\mathbf{Y} \geq \lambda \mathbf{I}$ for some small but positive $\lambda > 0$, depending on the smallest eigenvalue of \mathbf{M} . Intuitively, before any technical reasoning we would already expect some quantitative statement of this kind, since whenever \mathbf{M} is non-singular the $h_S(\mathbf{V}^\top \mathbf{z})$ are linearly independent and their conditioning should depend mostly on the conditioning of $(\mathbf{V}^\top \mathbf{z})^S$.

Proposition 10.8.8. $\lambda_{\min}(\mathbf{Y}) \geq \lambda_{\min}(\mathbf{M})^d$.

Proof. Since \mathbf{Y} is block diagonal, it suffices to show the result for each $\mathbf{Y}^{[d,d]}$. Let $\mathbf{U} \in \mathbb{R}^{\mathcal{M}_d([n]) \times \binom{[n]}{d}}$ have as its columns the monomial coefficients of the polynomials $h_S(\mathbf{x})$ (noting that the entries of the columns are indexed by multisets in $[n]$, compatible with this

interpretation), and let $\mathbf{A} \in \mathbb{R}^{\mathcal{M}_d([n]) \times \mathcal{M}_d([n])}$ have as its entries $A_{S,T} = \langle (\mathbf{V}^\top \mathbf{z})^S, (\mathbf{V}^\top \mathbf{z})^T \rangle_\partial$. We then have $\mathbf{Y}^{[d,d]} = \mathbf{U}^\top \mathbf{A} \mathbf{U}$.

Write the singular value decomposition $\mathbf{V} = \mathbf{Q}_1 \boldsymbol{\Sigma} \mathbf{Q}_2^\top$ for $\mathbf{Q}_i \in \mathcal{O}(n)$ and $\boldsymbol{\Sigma}$ diagonal containing the singular values of \mathbf{V} . Then, applying the orthogonal invariance of Proposition 8.3.4, we have $A_{S,T} = \langle (\mathbf{Q}_2 \boldsymbol{\Sigma} \mathbf{z})^S, (\mathbf{Q}_2 \boldsymbol{\Sigma} \mathbf{z})^T \rangle_\partial$. Now, letting $\mathbf{U}' \in \mathbb{R}^{\mathcal{M}_d([n]) \times \mathcal{M}_d([n])}$ have as its columns the monomial coefficients of $(\mathbf{Q}_2 \mathbf{z})^S$ and letting $\mathbf{A}' \in \mathbb{R}^{\mathcal{M}_d([n]) \times \mathcal{M}_d([n])}$ have entries $A'_{S,T} = \langle (\boldsymbol{\Sigma} \mathbf{z})^S, (\boldsymbol{\Sigma} \mathbf{z})^T \rangle_\partial$, we have $\mathbf{A} = \mathbf{U}'^\top \mathbf{A}' \mathbf{U}'$, and \mathbf{A}' is a diagonal matrix with entries equal to $d!$ multiplied by various products of $2d$ of the singular values of \mathbf{V} . In particular, letting σ_{\min} be the smallest such singular value, we have $\mathbf{A}' \succeq \sigma_{\min}^{2d} \mathbf{I}_{\mathcal{M}_d([n])}$. Therefore, $\mathbf{A} \succeq \sigma_{\min}^{2d} \cdot d! \mathbf{U}'^\top \mathbf{U}'$. But the entries of the matrix $d! \mathbf{U}'^\top \mathbf{U}'$ are the inner products $\langle (\mathbf{Q}_2 \mathbf{z})^S, (\mathbf{Q}_2 \mathbf{z})^T \rangle_\partial = \langle \mathbf{z}^S, \mathbf{z}^T \rangle_\partial$, again by orthogonal invariance. Thus this matrix is merely $d! \mathbf{I}_{\mathcal{M}_d([n])}$, so we find $\mathbf{A} \succeq \sigma_{\min}^{2d} d! \mathbf{I}_{\mathcal{M}_d([n])}$.

Therefore, $\mathbf{Y}^{[d,d]} \succeq \sigma_{\min}^{2d} d! \mathbf{U}^\top \mathbf{U}$. Since the only multilinear monomial appearing in $h_S(\mathbf{x})$ is \mathbf{x}^S , and this occurs with coefficient 1, the block of \mathbf{U}^\top indexed by all columns and rows corresponding to multisets with no repeated elements is the identity. In particular then, $\mathbf{U}^\top \mathbf{U} \succeq \mathbf{I}_{\binom{[n]}{d}}$, so $\mathbf{Y}^{[d,d]} \succeq \sigma_{\min}^{2d} d! \mathbf{I}_{\binom{[n]}{d}}$. Finally, since $\mathbf{M} = \mathbf{V}^\top \mathbf{V}$, we have $\lambda_{\min}(\mathbf{M}) = \sigma_{\min}^2$, and the result follows. \square

Combining our results, we are now prepared to prove Lemma 10.6.2.

Proof of Lemma 10.6.2. We need only recall the main results from the last two sections:

$$\|\mathbf{Z}^{\text{main}} - \mathbf{Z}^{\text{tied}}\| \leq (6d)^{10d} \|\mathbf{M}\|^{3d} \epsilon_{\text{tree}}(\mathbf{M}; d), \quad (\text{Corollary 10.7.6})$$

$$\|\mathbf{Y} - \mathbf{Z}^{\text{tied}}\| \leq d^{5d} \|\mathbf{M}\|^{3d} (\epsilon_{\text{pow}}(\mathbf{M}) + \epsilon_{\text{offdiag}}(\mathbf{M}) + \epsilon_{\text{corr}}(\mathbf{M})), \quad (\text{Corollary 10.8.7})$$

$$\lambda_{\min}(\mathbf{Y}) \geq \lambda_{\min}(\mathbf{M})^d, \quad (\text{Proposition 10.8.8})$$

where we note that the assumption $\epsilon_{\text{tree}}(\mathbf{M}; d) \leq 1$ in Corollary 10.7.6 follows from the

condition of Theorem 10.2.3, which is assumed in the statement. The result follows then follows by the eigenvalue inequality $\lambda_{\min}(\mathbf{Z}^{\text{main}}) \geq \lambda_{\min}(\mathbf{Y}) - \|\mathbf{Z}^{\text{main}} - \mathbf{Z}^{\text{tied}}\| - \|\mathbf{Z}^{\text{tied}} - \mathbf{Y}\|$. \square

10.9 BOUND ON THE ERROR TERM: PROOF OF LEMMA 10.6.3

Our first step in analyzing the error term is, as for the main term, to evaluate it in the multiharmonic basis, giving the entries of \mathbf{Z}^{err} . We recall that, in Proposition 10.5.6, we found that $\tilde{\mathbb{E}}^{\text{err}}$ decomposes as an application of $\tilde{\mathbb{E}}^{\text{main}}$ to some of the input indices and a combination of error trees applied to the other indices. Thus, part of the result of this calculation will be the familiar stretched forest terms from the calculations in Proposition 10.7.2, while the remainder will consist of the Δ^F error components from Section 10.5, applied to some of the partition components of the multiharmonic basis polynomials. To make it easier to describe and eventually bound the latter, we define the following type of error matrix.

Definition 10.9.1 (Partition-error matrices). *Suppose $\sigma = \{A_1, \dots, A_n\} \in \text{Part}([\ell]; \text{odd})$, $\tau = \{B_1, \dots, B_p\} \in \text{Part}([m]; \text{odd})$, and $T \in \mathcal{T}(n + p)$. We then define the partition-error matrix $\Delta^{(\sigma, \tau, T)} \in \mathbb{R}^{\binom{[n]}{\ell} \times \binom{[n]}{m}}$ associated to this triplet to have entries*

$$\Delta_{st}^{(\sigma, \tau, T)} := \sum_{\substack{a \in [n]^{\sigma[\geq 3]} \\ b \in [n]^{\tau[\geq 3]}}} \prod_{A \in \sigma} \prod_{i \in A} M_{i, f_{s,a}(A)} \cdot \prod_{B \in \tau} \prod_{j \in B} M_{j, g_{t,b}(B)} \cdot \Delta^T(\mathbf{M}; (f_{s,a}(A_1), \dots, f_{s,a}(A_n), g_{t,b}(B_1), \dots, g_{t,b}(B_p))). \quad (10.166)$$

Proposition 10.9.2. For any $S, T \subseteq [n]$,

$$\begin{aligned}
Z_{S,T}^{\text{err}} &= \tilde{\mathbb{E}}^{\text{err}}[h_S^\dagger(\mathbf{x})h_T^\dagger(\mathbf{x})] \\
&= \sum_{\substack{A \subseteq S \\ B \subseteq T \\ A+B \neq S+T}} \left(\sum_{\substack{F \in \mathcal{F}(|A|,|B|) \\ F \text{ stretched}}} \mu(F) \cdot Z_{A,B}^F \right) \left(\sum_{\pi \in \text{Part}((S \setminus A) + (T \setminus B); \text{even})} (-1)^{|\pi|} \prod_{R \in \pi} \sum_{\substack{\sigma \in \text{Part}([|R \cap S|]; \text{odd}) \\ \tau \in \text{Part}([|R \cap T|]; \text{odd})}} \right. \\
&\quad \left. \prod_{A \in \sigma + \tau} (-1)^{|A|-1} (|A|-1)! \sum_{F \in \mathcal{T}(|\sigma|+|\tau|)} \mu(F) \cdot \Delta_{R \cap S, R \cap T}^{(\sigma, \tau, F)} \right). \tag{10.167}
\end{aligned}$$

Proof. As in Proposition 10.7.2, we begin by expanding directly:

$$\begin{aligned}
&\tilde{\mathbb{E}}^{\text{err}}[h_S^\dagger(\mathbf{x})h_T^\dagger(\mathbf{x})] \\
&= \sum_{\substack{\sigma \in \text{Part}(S) \\ \tau \in \text{Part}(T)}} \prod_{R \in \sigma + \tau} (-1)^{|R|-1} (|R|-1)! \prod_{R \in \sigma[\text{even}] + \tau[\text{even}]} q_R^\dagger \cdot \\
&\quad \sum_{\substack{\mathbf{a} \in [n]^{\sigma[\text{odd}; \geq 3]} \\ \mathbf{b} \in [n]^{\tau[\text{odd}; \geq 3]}}} \prod_{A \in \sigma[\text{odd}]} \prod_{i \in A} M_{f_a(A), i} \prod_{B \in \tau[\text{odd}]} \prod_{j \in B} M_{g_b(B), j} \cdot \tilde{\mathbb{E}}^{\text{err}} \left[\prod_{A \in \sigma[\text{odd}]} \mathcal{X}_{f_a(A)} \prod_{B \in \tau[\text{odd}]} \mathcal{X}_{g_b(B)} \right]
\end{aligned}$$

where f_a, g_b are defined as in Proposition 10.7.2. Now, expanding the pseudoexpectation according to Proposition 10.5.6, we have

$$\begin{aligned}
&= \sum_{\substack{\sigma \in \text{Part}(S) \\ \tau \in \text{Part}(T)}} \prod_{R \in \sigma + \tau} (-1)^{|R|-1} (|R|-1)! \prod_{R \in \sigma[\text{even}] + \tau[\text{even}]} q_R^\dagger \cdot \\
&\quad \sum_{\substack{\mathbf{a} \in [n]^{\sigma[\text{odd}; \geq 3]} \\ \mathbf{b} \in [n]^{\tau[\text{odd}; \geq 3]}}} \prod_{A \in \sigma[\text{odd}]} \prod_{i \in A} M_{f_a(A), i} \prod_{B \in \tau[\text{odd}]} \prod_{j \in B} M_{g_b(B), j} \sum_{\substack{\pi \subseteq \sigma[\text{odd}] \\ \rho \subseteq \tau[\text{odd}] \\ |\pi| + |\rho| < |\sigma[\text{odd}]| + |\tau[\text{odd}]|}} \tilde{\mathbb{E}}^{\text{main}} \left[\prod_{A \in \pi} \mathcal{X}_{f_a(A)} \prod_{B \in \rho} \mathcal{X}_{g_b(B)} \right] \\
&\quad \sum_{\beta \in \text{Part}((\sigma[\text{odd}] \setminus \pi) + (\tau[\text{odd}] \setminus \rho); \text{even})} \prod_{\gamma \in \beta} \left(- \sum_{T \in \mathcal{T}(|R|)} \mu(T) \cdot \Delta^T(\mathbf{M}; (f_a(A))_{A \in \gamma \cap \sigma} \circ (g_b(B))_{B \in \gamma \cap \tau}) \right)
\end{aligned}$$

Here, we swap the order of summation and reorganize the sum according to the union of all parts of π and $\sigma[\text{even}]$, which we call J , and the union of all parts of ρ and $\tau[\text{even}]$, which we call K . Recognizing after this manipulation the intermediate result from Proposition 10.7.2, we continue

$$\begin{aligned}
&= \sum_{\substack{J \subseteq S \\ K \subseteq T \\ J+K \neq S+T}} \tilde{\mathbb{E}}^{\text{main}}[h_J^\downarrow(\mathbf{x})h_K^\downarrow(\mathbf{x})] \sum_{\substack{\sigma \in \text{Part}(S \setminus J; \text{odd}) \\ \tau \in \text{Part}(T \setminus K; \text{odd})}} \prod_{R \in \sigma + \tau} (-1)^{|R|-1} (|R| - 1)! \\
&\quad \sum_{\substack{\mathbf{a} \in [n]^{\sigma[\text{odd}; \geq 3]} \\ \mathbf{b} \in [n]^{\tau[\text{odd}; \geq 3]}}} \prod_{A \in \sigma} \prod_{i \in A} M_{f_a(A), i} \prod_{B \in \tau} \prod_{j \in B} M_{f_b(B), i} \\
&\quad \sum_{\beta \in \text{Part}(\sigma + \tau; \text{even})} \prod_{\gamma \in \beta} \left(- \sum_{T \in \mathcal{T}(|R|)} \mu(T) \cdot \Delta^T(\mathbf{M}; (f_a(A))_{A \in \gamma \cap \sigma} \circ (g_b(B))_{B \in \gamma \cap \tau}) \right)
\end{aligned}$$

and again exchanging the order of summation and letting π be the partition formed by taking the union of the sets in every part of β , by a similar manipulation to that in Proposition 10.7.2 we complete the proof. \square

We now develop a few tools to bound the norms of partition-error matrices. The following is a minor variant of Proposition 10.13.11, a diagrammatic factorization of CGMs that we use at length in the deferred technical proofs. This shows how $\Delta^{(\sigma, \tau, T)}$ can be factorized into two outer factors that are similar to CGMs with no \square vertices, and an inner factor that consists of values of Δ^T arranged in a matrix of suitable shape.

Proposition 10.9.3. *Let $\sigma = \{A_1, \dots, A_q\} \in \text{Part}([\ell]; \text{odd})$, $\tau = \{B_1, \dots, B_p\} \in \text{Part}([m]; \text{odd})$, and $T \in \mathcal{T}(q + p)$. Define $\mathbf{Z}^\sigma \in \mathbb{R}^{\binom{[n]}{\ell} \times [n]^q}$ to have entries*

$$\mathbf{Z}_{sa}^\sigma = \prod_{A_q = \{i\} \in \sigma[1]} \mathbb{1}\{s_i = a_q\} \prod_{A_q \in \sigma[\geq 3]} \prod_{i \in A_q} M_{s_i, a_q}, \tag{10.168}$$

and similarly $\mathbf{Z}^\tau \in \mathbb{R}^{\binom{[n]}{m} \times [n]^p}$. Let $\mathbf{F}^{(T,q,p)} \in \mathbb{R}^{[n]^q \times [n]^p}$ have entries

$$F_{ab}^{(T,q,p)} = \Delta^T(\mathbf{M}; (a_1, \dots, a_q, b_1, \dots, b_p)). \quad (10.169)$$

Then, $\Delta^{(\sigma,\tau,T)} = \mathbf{Z}^\sigma \mathbf{F}^{(T,q,p)} \mathbf{Z}^{\tau^\top}$.

Proof. The result follows from expanding the definition of the matrix multiplication and comparing with Definition 10.9.1. \square

Next, we show how the norm of the inner matrix can be controlled; in fact, we give a stronger statement bounding the Frobenius norm.

Proposition 10.9.4 (Error matrix Frobenius norm bound). *For any $d' \leq d$ and $T \in \mathcal{T}(2d')$,*

$$\left(\sum_{\mathbf{s} \in [n]^{2d'}} (\Delta^T(\mathbf{M}; \mathbf{s}))^2 \right)^{1/2} \leq (2d)^d \epsilon_{\text{err}}(\mathbf{M}; 2d) \quad (10.170)$$

Proof. Recall that we denote by $\text{set}(\mathbf{s})$ the set of distinct indices appearing in \mathbf{s} . By definition of ϵ_{err} , we have

$$|\Delta^T(\mathbf{M}; \mathbf{s})| \leq n^{-|\text{set}(\mathbf{s})|/2} \epsilon_{\text{err}}(\mathbf{M}; 2d). \quad (10.171)$$

Therefore, we find

$$\begin{aligned} \left(\sum_{\mathbf{s} \in [n]^{2d'}} (\Delta^T(\mathbf{M}; \mathbf{s}))^2 \right)^{1/2} &\leq \left(\sum_{\mathbf{s} \in [n]^{2d'}} n^{-|\text{set}(\mathbf{s})|} \right)^{1/2} \epsilon_{\text{err}}(\mathbf{M}; 2d) \\ &\leq \left(\sum_{k=1}^{2d'} n^{-k} \cdot \#\{\mathbf{s} \in [n]^{2d'} : |\text{set}(\mathbf{s})| = k\} \right)^{1/2} \epsilon_{\text{err}}(\mathbf{M}; 2d) \\ &\leq \left(\sum_{k=1}^{2d'} \frac{k^{2d'}}{k!} \right)^{1/2} \epsilon_{\text{err}}(\mathbf{M}; 2d) \\ &\leq (2d')^{d'} \epsilon_{\text{err}}(\mathbf{M}; 2d), \end{aligned} \quad (10.172)$$

and the result follows since $d' \leq d$. \square

Combining these results with an ancillary result gives the following bound.

Corollary 10.9.5. $\|\Delta^{(\sigma,\tau,T)}\| \leq (2(\ell + m))^{(\ell+m)} \|\mathbf{M}\|^{\ell+m} \epsilon_{\text{err}}(\mathbf{M}; 2d)$.

Proof. By norm submultiplicativity, $\|\Delta^{(\sigma,\tau,T)}\| \leq \|\mathbf{Z}^\sigma\| \cdot \|\mathbf{F}^{(T,n,p)}\| \cdot \|\mathbf{Z}^\tau\|$. By Proposition 10.13.13, we have $\|\mathbf{Z}^\sigma\| \leq \|\mathbf{M}\|^\ell$ and $\|\mathbf{Z}^\tau\| \leq \|\mathbf{M}\|^m$, and by Proposition 10.9.4, we have $\|\mathbf{F}^{(T,n,p)}\| \leq \|\mathbf{F}^{(T,n,p)}\|_F \leq (2(n+p))^{n+p} \epsilon_{\text{err}}(\mathbf{M}; n+p)$, and the result then follows after noting $n+p \leq \ell+m$ since $n = |\sigma|$ and $p = |\tau|$. \square

Proof of Lemma 10.6.3. First, we note that under the assumptions of Theorem 10.2.3, which we have also assumed in the statement of the Lemma, we have $\epsilon_{\text{err}}(\mathbf{M}; 2d) \leq 1$.

We then follow the same manipulations as in Lemma 10.6.2, using Proposition 10.12.4 to bound the norm of \mathbf{Z}^{err} by the sum of all block norms. Also, since products over subsets of indices correspond to tensorizations of terms in this sum (see Proposition 10.13.3), we may bound such products by corresponding products of matrix norms. We therefore find

$$\begin{aligned} \|\mathbf{Z}^{\text{err}}\| \leq & \sum_{\ell,m=0}^d \sum_{\substack{a \in [\ell] \\ b \in [m] \\ a+b < \ell+m}} \left(\sum_{\substack{F \in \mathcal{F}(a,b) \\ F \text{ stretched}}} |\mu(F)| \cdot \|\mathbf{Z}^F\| \right) \\ & \sum_{\pi \in \text{Part}([\ell+m-a-b]; \text{even})} \prod_{R \in \pi} \left(\sum_{\substack{\sigma \in \text{Part}(R \cap [\ell-a]; \text{odd}) \\ \tau \in \text{Part}(R \cap \{\ell-a+1, \dots, \ell+m-a-b\}; \text{odd})}} \prod_{A \in \sigma + \tau} (|A| - 1)! \right. \\ & \left. \sum_{T \in \mathcal{T}(|\sigma|, |\tau|)} |\mu(T)| \cdot \|\Delta^{(\sigma,\tau,T)}\| \right). \end{aligned} \quad (10.173)$$

In the sum over stretched forest ribbon diagrams, by Proposition 10.12.5 we have $|\mu(F)| \leq (3(a+b))! \leq (3(a+b))^{3(a+b)} \leq 3(\ell+m)^{3(\ell+m)} \leq (6d)^{6d}$, by Proposition 10.12.11 we have $|\mathcal{F}(a,b)| \leq (a+b)^{3(a+b)} \leq (2d)^{6d}$, and by Proposition 10.13.13 and Corollary 10.12.10 we have $\|\mathbf{Z}^F\| \leq \|\mathbf{M}\|^{3d}$. In the second term, by Corollary 10.9.5 we have $\|\Delta^{(\sigma,\tau,T)}\| \leq (2|R|)^{|R|} \|\mathbf{M}\|^{|R|} \epsilon_{\text{err}}(\mathbf{M}; 2d)$, by Proposition 10.12.5 we have $|\mu(T)| \leq (3|R|)! \leq (3|R|)^{3|R|}$, and by Proposition 10.12.11 we have $|\mathcal{T}(|\sigma|, |\tau|)| \leq |\mathcal{F}(|\sigma|, |\tau|)| \leq (|\sigma| + |\tau|)^{3(|\sigma| + |\tau|)} \leq$

$(|R|)^{3|R|}$. We also have $\prod_{A \in \sigma + \tau} (|A| - 1)! \leq (|\sigma| + |\tau|)! \leq (|R|)! \leq |R|^{|R|}$. Finally, by Proposition 10.12.7, viewing σ and τ taken together as a partition of R , the number of choices of σ and τ is at most $|R|^{|R|}$. Combining these observations, we continue

$$\|\mathbf{Z}^{\text{err}}\| \leq (6d)^{12d} \|\mathbf{M}\|^{3d} \epsilon_{\text{err}}(\mathbf{M}; 2d) \sum_{\ell, m=0}^d \sum_{\substack{a \in [\ell] \\ b \in [m] \\ a+b < \ell+m}} \sum_{\pi \in \text{Part}([\ell+m-a-b]; \text{even})} \prod_{R \in \pi} (3|R|)^{9|R|} \|\mathbf{M}\|^{|R|}$$

where since $\sum_{R \in \pi} |R| \leq 2d$ we continue

$$\leq (6d)^{30d} \|\mathbf{M}\|^{5d} \epsilon_{\text{err}}(\mathbf{M}; 2d) \cdot d^2 \cdot d^2 \cdot |\text{Part}(2d)|$$

and by Proposition 10.12.7 again we may finish

$$\leq (12d)^{32d} \|\mathbf{M}\|^{5d} \epsilon_{\text{err}}(\mathbf{M}; 2d), \tag{10.174}$$

completing the proof. □

10.10 OBSTACLES TO LOW-RANK EXTENSIONS

Unfortunately, our approach above does not appear to extend directly to the low-rank setting. There are two major obstacles to attempting to apply Theorem 10.2.3 to \mathbf{M} equal to a rescaled low-rank projection matrix. These correspond to two incoherence quantities that are no longer $o(1)$ as $n \rightarrow \infty$ once $\text{rank}(\mathbf{M}) = n - \Theta(n)$: ϵ_{pow} and ϵ_{err} . As we describe below, the former seems to represent a fundamental barrier requiring us to reconsider our derivation of the pseudomoments, while the latter can probably be handled without much difficulty.

10.10.1 ORDER-2 TENSOR ORTHONORMALITY: THE ϵ_{pow} OBSTACLE

The more substantial issue arises, in the course of our proof of Theorem 10.2.3, in the collapse of partition transport diagrams. Put simply, our construction relies on the $\mathbf{v}_i^{\otimes k}$ for all $k \geq 2$ behaving like a nearly-orthonormal set. Once $\text{rank}(M) = n - \Theta(n)$, this is no longer the case: for $k \geq 3$ the $\mathbf{v}_i^{\otimes k}$ still behave like an orthonormal set, but the $\mathbf{v}_i^{\otimes 2}$, which equivalently may be viewed as the matrices $\mathbf{v}_i \mathbf{v}_i^\top$, are too “crowded” in $\mathbb{R}_{\text{sym}}^{r \times r}$ and have an overly significant collective bias in the direction of the identity matrix.

More precisely, as illustrated in Figure 10.6, when $\text{rank}(M) = n - \Theta(n)$ then we no longer have merely $M^{\circ 2} \approx I_n$, but rather $M^{\circ 2} \approx I_n + t \mathbf{1}_n \mathbf{1}_n^\top$ for $t = \Theta(n^{-1})$, whereby one particular diagram contributes a non-negligible new term. (To see explicitly that this occurs, we may compute $n^{-1} \mathbf{1}_n^\top M^{\circ 2} \mathbf{1}_n = \|M\|_F^2 / n = \Theta(1)$ in this scaling.) This happens already at degree $2d = 4$, where this diagram has coefficient $+2$, and therefore our approximate Gramian factorization \mathbf{Y} of \mathbf{Z}^{main} has an extra positive term of the form of the right-most diagram in the Figure, two sided pairs with edges labelled by M^2 instead of M . The associated CGM, after multiplying by t , has spectral norm $O(1)$ (see the scaling discussed in Remark 10.4.5), so we have “ $\mathbf{Y} \gg \mathbf{Z}^{\text{main}}$ ” in psd ordering—in this case, our approximate Gramian factorization is simply false. Moreover, when $M^2 \approx \lambda M$ (as for M a rescaled projector), this additional diagram is the same as the diagram that is “orthogonalized away” by writing the pseudomoments in the multiharmonic basis. Therefore, the other diagrams in \mathbf{Z}^{main} cannot compensate for this negative term, whereby $\mathbf{Z}^{\text{main}} \not\approx \mathbf{0}$ and it is not merely the Gramian approximation that fails but rather the pseudomoment construction itself.

Some technical tricks can work around this issue at low degrees: in [KB20], we made an adjustment to the pseudomoments before passing to the multiharmonic basis, and also adjusted the basis so that \mathbf{Z}^{main} written in this basis has a similar extra term, whereby we can




Figure 10.6: Extra term from $D = [2]$ partition transport ribbon diagram. We illustrate the issue discussed in Section 10.10.1, that a partition transport ribbon diagram with associated plan $D = [2]$ produces an extra term consisting of two sided pairs when collapsed. The two parallel edges in the diagram on the left represent the matrix $M^{\circ 2}$, and this expansion corresponds to an approximation $M^{\circ 2} \approx I_n + t \mathbf{1}_n \mathbf{1}_n^\top$.

restore the equality $Z^{\text{main}} \approx Y$.⁴ In Section 10.11 below we will take a different approach, simply adding extra terms to $\tilde{\mathbb{E}}_M$ itself that achieve the same thing. This is slightly more flexible, working up to $2d = 6$.

It seems, however, that to resolve this issue for arbitrarily high degrees would require rethinking much of our derivation to include a second-order correction. Namely, our very initial construction of the approximate Green's function to the system of PDEs $\langle \mathbf{v}_i, \partial \rangle^2 p = 0$ in Section 10.1.2 already has built in the premise that the matrices $\mathbf{v}_i \mathbf{v}_i^\top$ are nearly orthonormal, when compared to the actual Green's function derivations in similar situations by [Cle00]. Specifically, we took $\varphi(\mathbf{z}) = \prod_{i=1}^n \langle \mathbf{v}_i, \mathbf{z} \rangle$, while Clerc's derivation suggests that the more principled choice would be $\varphi(\mathbf{z}) = \det(\sum_{i=1}^n \langle \mathbf{v}_i, \mathbf{z} \rangle^2 \mathbf{v}_i \mathbf{v}_i^\top)^{1/2}$. Unfortunately, we have not been able to achieve the symmetries required of the pseudomoments building harmonic projections by starting with such a Green's function. Nonetheless, since this same assumption eventually leads our construction astray, it seems plausible that the correct resolution will come from building a more nuanced approximate Green's function and deriving the prediction anew.

⁴The change of basis is expressed there as a Schur complement, which turns out to be equivalent.

10.10.2 SIDED ERROR TERMS: THE ϵ_{err} OBSTACLE

Another, milder difficulty that arises when $\text{rank}(\mathbf{M}) = n - \Theta(n)$ is that we have $\epsilon_{\text{err}}(\mathbf{M}) = \Theta(1)$. This is less of a problem because we have used a very coarse approach to bounding the error contribution \mathbf{Z}^{err} , bounding related matrix norms by Frobenius norms in the proof of Lemma 10.6.3. Since, assuming randomly behaving entries, we only expect such a bound to be tight when a matrix is actually a vector, this should not be a problem except in situations where the diagram-like objects of error terms appearing in the proof of Lemma 10.6.3 are *sided*, applying only to left or right indices in a diagram. To be more specific, we recall that, per Proposition 10.5.6, the full pseudoexpectation may be written as a sum over diagrams where “main term trees” are applied to some subsets of indices and “error term trees” are applied to others. The multiharmonic basis eliminates sided main term trees, but does not have any useful effect on sided error term trees.

Below, we give a simple result showing that, in quite general pseudoexpectations like this, one may build a basis that eliminates sided terms of any kind. It is possible to combine this construction with the multiharmonic basis to eliminate sided error terms, though this complicates other parts of the analysis and has no direct benefit in our applications, so we do not pursue it here.

Proposition 10.10.1. *Let $\tilde{\mathbb{E}} : \mathbb{R}[x_1, \dots, x_n] \rightarrow \mathbb{R}$ be linear and given by*

$$\tilde{\mathbb{E}}[\mathbf{x}^S] = \sum_{\pi \in \text{Part}(S)} \prod_{A \in \pi} E_A \quad (10.175)$$

for all $S \in \mathcal{M}([n])$, where E_A are arbitrary multiset-indexed quantities. Define the polynomials

$$p_S(\mathbf{x}) := \sum_{T \subseteq S} \left(\sum_{\pi \in \text{Part}(S-T)} (-1)^{|\pi|} \prod_{A \in \pi} E_A \right) \mathbf{x}^T. \quad (10.176)$$

for all $S \in \mathcal{M}([n])$. The pseudomoments written in this basis are then

$$\tilde{\mathbb{E}}[p_S(\mathbf{x})p_T(\mathbf{x})] = \sum_{\substack{\pi \in \text{Part}(S+T) \\ A \cap S \neq \emptyset, A \cap T \neq \emptyset \\ \text{for all } A \in \pi}} \prod_{A \in \pi} E_A. \quad (10.177)$$

Proof. We calculate directly:

$$\begin{aligned} & \tilde{\mathbb{E}}[p_S(\mathbf{x})p_T(\mathbf{x})] \\ &= \sum_{\substack{S' \subseteq S \\ T' \subseteq T}} \sum_{\substack{\sigma \in \text{Part}(S') \\ \tau \in \text{Part}(T') \\ \pi \in \text{Part}((S-S')+(T-T')) \\ A \cap S \neq \emptyset, A \cap T \neq \emptyset \\ \text{for all } A \in \pi}} \left(\sum_{\substack{\sigma' \subseteq \sigma \\ \tau' \subseteq \tau}} (-1)^{|\sigma'|+|\tau'|} \right) \prod_{A \in \sigma+\tau+\pi} E_A \\ &= \sum_{\substack{S' \subseteq S \\ T' \subseteq T}} \sum_{\substack{\sigma \in \text{Part}(S') \\ \tau \in \text{Part}(T') \\ \pi \in \text{Part}((S-S')+(T-T')) \\ A \cap S \neq \emptyset, A \cap T \neq \emptyset \\ \text{for all } A \in \pi}} \left(\sum_{\sigma' \subseteq \sigma} (-1)^{|\sigma'|} \right) \left(\sum_{\tau' \subseteq \tau} (-1)^{|\tau'|} \right) \prod_{A \in \sigma+\tau+\pi} E_A, \end{aligned} \quad (10.178)$$

and the remaining coefficients are 1 if $|S'| = |T'| = 0$ and 0 otherwise, completing the proof. \square

Here we have used implicitly the simple Möbius function of the subset poset from Example 10.3.4 in building our basis. There would appear to be an analogy between this feature and the appearance of the Möbius function of partitions from Example 10.3.5 in the multi-harmonic basis. It would be interesting to develop more general techniques for “orthogonalizing away” the terms of pseudoexpectations that contribute to a multiscale spectrum in the pseudomoment matrix using bases that incorporate poset combinatorics.

10.11 LIFTING 2: LOW RANK TO LOW DEGREE

As discussed above, Theorem 10.2.3 does not apply directly to low-rank M , as we sought for our applications to the SK Hamiltonian and related problems. However, for low degrees

of SOS, we can still make a manual correction and obtain a lower bound.

We present our result in terms of another, modified lifting theorem for arbitrary degree 2 pseudomoments. This extension only reaches degree 6, but allows the flexibility we sought above in ϵ_{pow} . We obtain it by inelegant means, using simplifications specific to the diagrams appearing at degree 6 to make some technical improvements in the argument of Theorem 10.2.3.

Definition 10.11.1 (Additional incoherence quantities). *For $M \in \mathbb{R}_{\text{sym}}^{n \times n}$ and $t > 0$, define the following quantities:*

$$\tilde{\epsilon}_{\text{pow}}(M, t) := \max \left\{ \|M^{\circ 2} - I_n - t \mathbf{1}_n \mathbf{1}_n^\top\|, \max_{k \geq 3} \|M^{\circ k} - I_n\| \right\}, \quad (10.179)$$

$$\tilde{\epsilon}(M, t) := \epsilon_{\text{offdiag}}(M) + \tilde{\epsilon}_{\text{pow}}(M, t) + n^{-1/2} \epsilon_{\text{err}}(M; 6). \quad (10.180)$$

Theorem 10.11.2. *Let $M \in \mathbb{R}_{\text{sym}}^{n \times n}$ with $M_{ii} = 1$ for all $i \in [n]$, and suppose $t_{\text{pow}} > 0$. Suppose that*

$$\lambda_{\min}(M) \geq 10^6 \|M\|^5 \tilde{\epsilon}(M, t_{\text{pow}})^{1/3}. \quad (10.181)$$

Define the constant

$$c := 250 t_{\text{pow}} (\|M\|^6 \|M^2\|_F + n \epsilon_{\text{offdiag}}(M^2) + n^2 \epsilon_{\text{offdiag}}(M^2)^3). \quad (10.182)$$

Then, there exists a degree 6 pseudoexpectation $\tilde{\mathbb{E}}$ with $\tilde{\mathbb{E}}[\mathbf{x}\mathbf{x}^\top] = (1 - c)M + cI_n$.

We show as part of the proof that a pseudoexpectation achieving this can be built by adding a correction of sub-leading order to those terms of the pseudoexpectation in Definition 10.1.13 where F is a perfect matching. As mentioned above, it is likely that to extend this result to degree $\omega(1)$ using our ideas would require somewhat rethinking our construction and the derivation we give in Section 10.1.2 to take into account the above “ ϵ_{pow}

obstacle,” but this makes it plausible that the result will be some form of small correction added to $\tilde{\mathbb{E}}_M$.

Proof of Theorem 10.11.2. The construction of $\tilde{\mathbb{E}}$ is a minor variation on that of $\tilde{\mathbb{E}}_M$ from Theorem 10.2.3, modified to counteract the negative terms discussed in Section 10.10.1.

$$\tilde{\mathbb{E}}_M^{\text{pairs}} \left[\prod_{i \in S} x_i \right] := \mathbb{1}\{|S| \in \{4, 6\}\} \sum_{\substack{F \in \mathcal{F}(|S|) \\ F \text{ all pairs}}} Z^F(M^2; S) \text{ for all } S \subseteq [n], \quad (10.183)$$

$$\tilde{\mathbb{E}}^{\text{id}} \left[\prod_{i \in S} x_i \right] := \mathbb{1}\{S = \emptyset\}, \quad (10.184)$$

$$\tilde{\mathbb{E}} := (1 - c) \left(\tilde{\mathbb{E}}_M + 2t_{\text{pow}} \tilde{\mathbb{E}}_M^{\text{pairs}} \right) + c \tilde{\mathbb{E}}^{\text{id}}. \quad (10.185)$$

We remark that here we use a combination of the two strategies for attenuating the spectrum of the error terms that were discussed in Remark 10.2.4. We also emphasize the detail that the matrix used in the CGSs in $\tilde{\mathbb{E}}^{\text{pairs}}$ is the *square* of M . (For M a rescaled random projection we expect $M^2 \approx \lambda M$ for some λ , but we do not require such a relation to hold, nor is this taken into account in the incoherence quantities used in the statement.)

Let us moreover decompose $\tilde{\mathbb{E}}_M^{\text{pairs}}$ into three terms, as follows. Note that the first two are bilinear operators on polynomials of degree at most d , in the sense of Definition 8.2.1, while the last has the additional symmetry making it a linear operator on polynomials of degree

at most $2d$.

$$\begin{aligned} & \tilde{\mathbb{E}}_M^{\text{pairs:main:1}}(\mathbf{x}^S, \mathbf{x}^T) \\ & := \mathbb{1}\{|S| + |T| \in \{4, 6\}\} \sum_{\substack{F \in \mathcal{F}(|S|, |T|) \\ F \text{ all pairs} \\ F \text{ has 2 sided pairs}}} Z_{S,T}^F(\mathbf{M}^2) \text{ for all } S, T \in \mathcal{M}([n]), \end{aligned} \quad (10.186)$$

$$\begin{aligned} & \tilde{\mathbb{E}}_M^{\text{pairs:main:2}}(\mathbf{x}^S, \mathbf{x}^T) \\ & := \mathbb{1}\{|S| + |T| \in \{4, 6\}\} \sum_{\substack{F \in \mathcal{F}(|S|, |T|) \\ F \text{ all pairs} \\ F \text{ has } \leq 1 \text{ sided pair}}} Z_{S,T}^F(\mathbf{M}^2) \text{ for all } S \in \mathcal{M}([n]), \end{aligned} \quad (10.187)$$

$$\begin{aligned} & \tilde{\mathbb{E}}_M^{\text{pairs:err}}(\mathbf{x}^S, \mathbf{x}^T) \\ & := \tilde{\mathbb{E}}_M^{\text{pairs}}[\mathbf{x}^{S+T}] - \tilde{\mathbb{E}}_M^{\text{pairs:main:1}}(\mathbf{x}^S, \mathbf{x}^T) - \tilde{\mathbb{E}}_M^{\text{pairs:main:2}}(\mathbf{x}^S, \mathbf{x}^T) \\ & = \tilde{\mathbb{E}}_M^{\text{pairs:err}}[\mathbf{x}^{S+T}]. \end{aligned} \quad (10.188)$$

The point here is that, since we expect $t_{\text{pow}} \|\mathbf{M}^2\|_F^2 = O(1)$, only the ribbon diagrams with two sided pairs, those in $\tilde{\mathbb{E}}^{\text{pairs:main:1}}$, will contribute significantly. The further decomposition between $\tilde{\mathbb{E}}^{\text{pairs:main:1}} + \tilde{\mathbb{E}}^{\text{pairs:main:2}}$ and $\tilde{\mathbb{E}}^{\text{pairs:err}}$ is precisely the same as that between $\tilde{\mathbb{E}}^{\text{main}}$ and $\tilde{\mathbb{E}}^{\text{err}}$ in Theorem 10.2.3, the former being simpler to work with in terms of CGMs and the latter being a small correction.

Our result will then follow from the following three claims:

$$\tilde{\mathbb{E}}_M + 2t_{\text{pow}} \tilde{\mathbb{E}}_M^{\text{pairs:main:1}} \geq \mathbf{0}, \quad (10.189)$$

$$2(1 - c)t_{\text{pow}} \tilde{\mathbb{E}}_M^{\text{pairs:main:2}} + \frac{c}{2} \tilde{\mathbb{E}}^{\text{id}} \geq \mathbf{0}, \quad (10.190)$$

$$2(1 - c)t_{\text{pow}} \tilde{\mathbb{E}}_M^{\text{pairs:err}} + \frac{c}{2} \tilde{\mathbb{E}}^{\text{id}} \geq \mathbf{0}. \quad (10.191)$$

For (10.189) we will argue by adjusting the proof of Theorem 10.2.3, arguing for positivity in the harmonic basis, and using that the additional term counteracts the negative terms

discussed in Section 10.10.1. For (10.190) and (10.191), we will make simpler arguments in the standard monomial basis.

Proof of (10.189): We will be quite explicit about the calculations in this section, essentially recapitulating this special case of Theorem 10.2.3 with adjustments as needed. We notice first that the only cases where $\tilde{\mathbb{E}}_M^{\text{pairs:main:1}}(\mathbf{x}^S, \mathbf{x}^T) \neq 0$ are where either $|S| = |T| = 2$ or $|S| = |T| = 3$. In the former case there is only a single diagram, with one sided pair in each \mathcal{L} and \mathcal{R} , while in the latter case there are 9 such diagrams, with one additional non-sided pair (there are $3 \cdot 3 = 9$ ways to choose the leaves belonging to this pair).

Let us enumerate explicitly the multiharmonic basis polynomials $h_S^\perp(\mathbf{x})$ for $|S| \leq 3$:

$$h_\emptyset^\perp(\mathbf{x}) = 1, \tag{10.192}$$

$$h_{\{i\}}^\perp(\mathbf{x}) = x_i, \tag{10.193}$$

$$h_{\{i,j\}}^\perp(\mathbf{x}) = x_i x_j - M_{ij}, \tag{10.194}$$

$$h_{\{i,j,k\}}^\perp(\mathbf{x}) = x_i x_j x_k - M_{ij} x_k - M_{ik} x_j - M_{jk} x_i + 2 \sum_{a=1}^n M_{ai} M_{aj} M_{ak} x_a. \tag{10.195}$$

We see therefore that the only cases with $\tilde{\mathbb{E}}^{\text{pairs:main:1}}(h_S^\perp(\mathbf{x}), h_T^\perp(\mathbf{x})) \neq 0$ will again be those with either $|S| = |T| = 2$ or $|S| = |T| = 3$, and in these two cases $\tilde{\mathbb{E}}^{\text{pairs:main:1}}(h_S^\perp(\mathbf{x}), h_T^\perp(\mathbf{x})) = \tilde{\mathbb{E}}^{\text{pairs:main:1}}(\mathbf{x}^S, \mathbf{x}^T)$; thus, the more complicated terms in the harmonic basis polynomials are in fact entirely “invisible” to the corrective term $\tilde{\mathbb{E}}^{\text{pairs:main:1}}$. That this does not happen anymore once $2d \geq 8$ seems to be one of the main obstructions to applying a similar adjustment technique there.

Following the proof of Theorem 10.2.3 but adding an extra detail, we define the pseudo-

moment matrices $\mathbf{Z}^{\text{main:1}}, \mathbf{Z}^{\text{main:2}}, \mathbf{Z}^{\text{err}}, \mathbf{Z} \in \mathbb{R}^{\binom{[n]}{\leq 3} \times \binom{[n]}{\leq 3}}$ to have entries

$$\mathbf{Z}_{S,T}^{\text{main:1}} := \tilde{\mathbb{E}}_M^{\text{main}}[h_S^\perp(\mathbf{x})h_T^\perp(\mathbf{x})], \quad (10.196)$$

$$\mathbf{Z}_{S,T}^{\text{main:2}} := 2t_{\text{pow}} \tilde{\mathbb{E}}_M^{\text{pairs:main:1}}[h_S^\perp(\mathbf{x})h_T^\perp(\mathbf{x})], \quad (10.197)$$

$$\mathbf{Z}_{S,T}^{\text{err}} := \tilde{\mathbb{E}}_M^{\text{err}}[h_S^\perp(\mathbf{x})h_T^\perp(\mathbf{x})], \quad (10.198)$$

$$\mathbf{Z} := \mathbf{Z}^{\text{main:1}} + \mathbf{Z}^{\text{main:2}} + \mathbf{Z}^{\text{err}}. \quad (10.199)$$

It then suffices to prove $\mathbf{Z} \geq \mathbf{0}$.

By Proposition 10.7.2, $\mathbf{Z}^{\text{main:1}}$ is block diagonal (note that in this small case there are no stretched forest ribbon diagrams with unequal numbers of leaves in \mathcal{L} and \mathcal{R} , so the block diagonalization is exact). Define \mathbf{Z}^{tied} as in Corollary 10.7.6. We apply Corollary 10.7.6, but improve the constants (to avoid truly astronomical values) by adjusting the proof to $d = 6$, noting that $|\mathcal{F}(2)| = 1$, $|\mathcal{F}(4)| = 4$, and $|\mathcal{F}(6)| = 51$, and for all $2d' \leq 6$ and $F \in \mathcal{F}(2d')$ we have $|\mu(F)| \leq 24$. This gives

$$\|\mathbf{Z}^{\text{main:1}} - \mathbf{Z}^{\text{tied}}\| \leq 10^7 \|\mathbf{M}\|^9 \epsilon_{\text{tree}}(\mathbf{M}; 3),$$

where we can note that $\epsilon_{\text{tree}}(\mathbf{M}; 3) = \epsilon_{\text{tree}}(\mathbf{M}; 2) = \epsilon_{\text{offdiag}}(\mathbf{M})$, since the only good tree on two leaves is a pair, allowing us to eliminate ϵ_{tree} ,

$$= 10^7 \|\mathbf{M}\|^9 \epsilon_{\text{offdiag}}(\mathbf{M}). \quad (10.200)$$

We claim that the same approximate Gram factorization that held for \mathbf{Z}^{tied} in Theorem 10.2.3 holds for $\mathbf{Z}^{\text{tied}} + \mathbf{Z}^{\text{main:2}}$ in this case. Namely, as in Theorem 10.2.3, we define $\mathbf{Y} \in \mathbb{R}^{\binom{[n]}{\leq 3} \times \binom{[n]}{\leq 3}}$ to have entries $Y_{S,T} = \langle h_S(\mathbf{V}^\top \mathbf{z}), h_T(\mathbf{V}^\top \mathbf{z}) \rangle_\partial$. By Proposition 10.8.8, we have $\mathbf{Y} \geq \lambda_{\min}(\mathbf{M})^3$, and we will bound $\|\mathbf{Z}^{\text{tied}} + \mathbf{Z}^{\text{main:2}} - \mathbf{Y}\|$ below.

By construction \mathbf{Y} is block diagonal. We let $\mathbf{Y}^{[d,d]} \in \mathbb{R}^{\binom{[n]}{d} \times \binom{[n]}{d}}$ be the diagonal block

indexed by sets of size d . Likewise, \mathbf{Z}^{tied} and $\mathbf{Z}^{\text{main:2}}$ are block diagonal, by Corollary 10.7.6 and our remark above, respectively. Denote $\mathbf{Z}^{\text{tied}[d,d]}$ and $\mathbf{Z}^{\text{main:2}[d,d]}$ for their respective diagonal blocks. We have $\mathbf{Z}^{\text{main:2}[0,0]}$ and $\mathbf{Z}^{\text{main:2}[1,1]}$ both identically zero, and $\mathbf{Z}^{\text{tied}[0,0]} = \mathbf{Y}^{[0,0]} = [1]$ and $\mathbf{Z}^{\text{tied}[1,1]} = \mathbf{Y}^{[1,1]} = \mathbf{M}$, so it suffices to bound $\|\mathbf{Z}^{\text{tied}[d,d]} + \mathbf{Z}^{\text{main:2}[d,d]} - \mathbf{Y}^{[d,d]}\|$ for $d \in \{2, 3\}$.

For $d = 2$, as mentioned above, there is only one diagram consisting of two sided pairs occurring in $\mathbf{Z}^{\text{main:2}[2,2]}$. There are three bowtie forest ribbon diagrams in $\mathbf{Z}^{\text{tied}[2,2]}$: two consist of two pairs, while the third consists of all four leaves connected to a \square vertex. For a given choice of $\sigma, \tau \in \text{Part}([2])$, if either $|\sigma| = 1$ or $|\tau| = 1$ then there is a unique $D \in \text{Plans}(\sigma, \tau)$; otherwise, there are two plans corresponding to the two matchings of two copies of $\{1, 2\}$. Enumerating all of these terms, we find that most of them cancel in $\mathbf{Z}^{\text{tied}[2,2]} + \mathbf{Z}^{\text{main:2}[2,2]} - \mathbf{Y}^{[2,2]}$: see Figure 10.7 for a graphical depiction of the calculation. We are left only with the final diagrams illustrated there. The sum of their CGMs' norm we can bound using the factorization of Proposition 10.13.11 and the norm bound of Proposition 10.13.13:

$$\|\mathbf{Z}^{\text{tied}[2,2]} + \mathbf{Z}^{\text{main:2}[2,2]} - \mathbf{Y}^{[2,2]}\| \leq 2\|\mathbf{M}\|^4\|\mathbf{M} + t_{\text{pow}}\mathbf{1}_n\mathbf{1}_n^\top - \mathbf{M}^{\circ 2}\| = 2\|\mathbf{M}\|^4\tilde{\epsilon}_{\text{pow}}(\mathbf{M}). \quad (10.201)$$

For $d = 3$, again as mentioned before, there are 9 diagrams in $\mathbf{Z}^{\text{main:2}[3,3]}$, each consisting of two sided pairs and one non-sided pair. There are 16 bowtie forest ribbon diagrams in $\mathbf{Z}^{\text{tied}[3,3]}$: 6 where every connected component is a pair, 9 where one connected component is a pair and another is a star on 4 leaves, and 1 star on all 6 leaves. We first apply Lemma 10.8.4, which controls the norm error incurred in tying any partition transport ribbon diagram. We use the following minor variant of this bound, which follows upon examining the proof of Lemma 10.8.4 for the particular case $d = 3$: for $\sigma, \tau \in \text{Part}([3])$ and

$$\begin{aligned}
& \left[\begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} + \begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} - 2 \begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} \right] \\
+ & \left[2t_{\text{pow}} \begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} \right] \\
- & \left[\begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} + \begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} - 2 \begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} - 2 \begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} \right] \\
& + 2 \left[\begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} \right] \\
= & 2 \left[\begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} + t_{\text{pow}} \begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} - \begin{array}{c} \text{---} \quad \text{---} \\ \bullet \quad \bullet \\ \bullet \quad \bullet \end{array} \right]
\end{aligned}$$

Figure 10.7: Diagrammatic manipulation of degree 4 adjustment. We show the calculation of $Z^{\text{tied}[2,2]} + Z^{\text{main}:2[2,2]} - Y^{[2,2]}$ (listed here in the same order they appear in the figure) and the role of the additional diagram of $Z^{\text{main}:2[2,2]}$ in adjusting the result. The reader may compare with Figure 10.6 to see why the right-hand side is a desirable outcome whose norm we are able to bound.

$D \in \text{Plans}(\sigma, \tau)$, if D does not equal a permutation of the matrix $D^* := \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$, then

$$\|Z^{G(\sigma, \tau, D)} - Z^{\text{tie}(G(\sigma, \tau, D))}\| \leq 9\|M\|^9(\epsilon_{\text{offdiag}}(M) + \tilde{\epsilon}_{\text{pow}}(M)). \quad (10.202)$$

This follows simply by observing that, in this case, either σ or τ contains a singleton whereby Case 2 of that proof applies (giving the first term above), or $\sigma = \tau = \{\{1, 2, 3\}\}$, in which case $D = [3]$ and Case 3 applies and immediately yields the tied diagram after the first step, incurring error of only $\|M\|^6\|M^{\circ 3} - I_n\|$. (We do this carefully to avoid incurring a cost of $\|M^{\circ 2} - I_n\|$, which we are no longer assuming we have good control over.)

Applying this bound to all partition transport ribbon diagrams in Y whose matrix D is not a permutation of D^* as above, we may form $Y^{\text{tied}[3,3]}$ where all of these diagrams are

tied which, by the same counting as in Corollary 10.8.7, will satisfy

$$\|\mathbf{Y}^{[3,3]} - \mathbf{Y}^{\text{tied}[3,3]}\| \leq 10^8 \|\mathbf{M}\|^9 (\epsilon_{\text{offdiag}}(\mathbf{M}) + \tilde{\epsilon}_{\text{pow}}(\mathbf{M})). \quad (10.203)$$

Now, in $\mathbf{Z}^{\text{tied}[3,3]} + \mathbf{Z}^{\text{main:2}[3,3]} - \mathbf{Y}^{\text{tied}[3,3]}$, all ribbon diagrams cancel (as in the $d = 2$ case) except for 9 copies of the situation illustrated in Figure 10.7, each with an extra non-sided pair. Therefore, applying the same argument, we obtain the same bound, multiplied by 9 for the number of diagrams and by $\|\mathbf{M}\|$ for the extra non-sided pair. Thus:

$$\|\mathbf{Z}^{\text{tied}[3,3]} + \mathbf{Z}^{\text{main:2}[3,3]} - \mathbf{Y}^{\text{tied}[3,3]}\| \leq 18 \|\mathbf{M}\|^5 \tilde{\epsilon}_{\text{pow}}(\mathbf{M}), \quad (10.204)$$

and by triangle inequality, combining this with the previous inequality we find

$$\|\mathbf{Z}^{\text{tied}[3,3]} + \mathbf{Z}^{\text{main:2}[3,3]} - \mathbf{Y}^{[3,3]}\| \leq 10^9 \|\mathbf{M}\|^9 (\epsilon_{\text{offdiag}}(\mathbf{M}) + \tilde{\epsilon}_{\text{pow}}(\mathbf{M})). \quad (10.205)$$

Combining (10.200), (10.201), and (10.204), we have:

$$\begin{aligned} & \lambda_{\min}(\mathbf{Z}^{\text{main:1}} + \mathbf{Z}^{\text{main:2}}) \\ & \geq \lambda_{\min}(\mathbf{Y}) - \lambda_{\max}(\mathbf{Z}^{\text{tied}} + \mathbf{Z}^{\text{main:2}} - \mathbf{Y}) - \lambda_{\max}(\mathbf{Z}^{\text{main:1}} - \mathbf{Z}^{\text{tied}}) \\ & \geq \lambda_{\min}(\mathbf{M})^3 - 10^{10} \|\mathbf{M}\|^9 (\tilde{\epsilon}_{\text{pow}}(\mathbf{M}) + \epsilon_{\text{offdiag}}(\mathbf{M})). \end{aligned} \quad (10.206)$$

It remains to bound $\|\mathbf{Z}^{\text{err}}\|$. Here we again use a small improvement on the general strategy, this time that of Corollary 10.9.5 for bounding the inner error matrices $\Delta^{(\sigma, \tau, T)}$, specific to degree 6. Recall that here $\sigma \in \text{Part}([\ell]; \text{odd})$, $\tau \in \text{Part}([m]; \text{odd})$, and $T \in \mathcal{T}(|\sigma| + |\tau|)$, for $0 \leq \ell, m \leq 3$. $\Delta^{(\sigma, \tau, T)}$ is indexed by $\binom{[n]}{\ell} \times \binom{[n]}{m}$, and contains terms $\Delta^T(\mathbf{M}; \cdot)$. We note that if $|\sigma| + |\tau| < 4$, then $\Delta^T(\mathbf{M}; \mathbf{s}) = 0$ identically, so $\Delta^{(\sigma, \tau, T)} = \mathbf{0}$ for any T in this case. But, when $\ell, m \leq 3$, the only way this can be avoided and $|\sigma|$ and $|\tau|$ can have the

same parity is in one of three cases: (1) σ and τ both consist of two singletons, (2) σ and τ both consist of three singletons, or (3) one has a single part of size 3 and the other consists of three singletons. In any case, the matrix \mathbf{F} from Proposition 10.9.3 may be viewed as *sparse*, along whichever of the rows or columns is indexed by a partition consisting only of singletons. Therefore, our norm bound which naively bounded $\|\mathbf{F}\| \leq \|\mathbf{F}\|_F$ can be improved by applying Proposition 10.12.1. Repeating the argument of Corollary 10.9.5 in this way, we find

$$\begin{aligned} \|\Delta^{(\sigma,\tau,T)}\| &\leq \|\mathbf{M}\|^6 \sqrt{(n^{-1}\epsilon_{\text{err}}(\mathbf{M};6) + n^{-3/2}\epsilon_{\text{err}}(\mathbf{M};6) \cdot n)^2} \\ &\leq 2\|\mathbf{M}\|^6 n^{-1/2}\epsilon_{\text{err}}(\mathbf{M};6), \text{ if } \sigma = \tau = \{\{1\}, \{2\}\}, \end{aligned} \quad (10.207)$$

$$\begin{aligned} \|\Delta^{(\sigma,\tau,T)}\| &\leq \|\mathbf{M}\|^6 \sqrt{(n^{-3/2}\epsilon_{\text{err}}(\mathbf{M};6) + n^{-4/2}\epsilon_{\text{err}}(\mathbf{M};6) \cdot n + n^{-5/2}\epsilon_{\text{err}}(\mathbf{M};6) \cdot n^2)^2} \\ &\leq 3\|\mathbf{M}\|^6 n^{-1/2}\epsilon_{\text{err}}(\mathbf{M};6), \text{ if } \sigma = \tau = \{\{1\}, \{2\}, \{3\}\}, \end{aligned} \quad (10.208)$$

$$\begin{aligned} \|\Delta^{(\sigma,\tau,T)}\| &\leq \sqrt{(n^{-3/2}\epsilon_{\text{err}}(\mathbf{M};6) \cdot n)(n^{-3/2}\epsilon_{\text{err}}(\mathbf{M};6) \cdot 3)} \\ &\leq 3\|\mathbf{M}\|^6 n^{-1/2}\epsilon_{\text{err}}(\mathbf{M};6), \text{ if } \sigma = \{\{1, 2, 3\}\}, \tau = \{\{1\}, \{2\}, \{3\}\} \\ &\text{or vice-versa.} \end{aligned} \quad (10.209)$$

In effect, we are able to scale $\epsilon_{\text{err}}(\mathbf{M};6)$ down by an additional factor of $n^{-1/2}$ using this technique. Following the remainder of the proof of Lemma 10.6.3 with this improvement and again slightly improving the constants as above for the specific case $2d = 6$ then gives

$$\|\mathbf{Z}^{\text{err}}\| \leq 10^{18}\|\mathbf{M}\|^{15}n^{-1/2}\epsilon_{\text{err}}(\mathbf{M};6). \quad (10.210)$$

Combining this with (10.206) then gives

$$\begin{aligned}
\lambda_{\min}(\mathbf{Z}) &\geq \lambda_{\min}(\mathbf{Z}^{\text{main:1}} + \mathbf{Z}^{\text{main:2}}) - \|\mathbf{Z}^{\text{err}}\| \\
&\geq \lambda_{\min}(\mathbf{M})^3 - 10^{18} \|\mathbf{M}\|^{15} (\tilde{\epsilon}_{\text{pow}}(\mathbf{M}) + \epsilon_{\text{offdiag}}(\mathbf{M}) + n^{-1/2} \epsilon_{\text{err}}(\mathbf{M}; 6)) \\
&\geq 0
\end{aligned} \tag{10.211}$$

by our assumption in the statement. Thus, $\mathbf{Z} \succeq \mathbf{0}$.

Proof of (10.190): We consider the pseudomoment matrix of the left-hand side, written in the standard monomial basis. The term arising from $\tilde{\mathbb{E}}^{\text{id}}$ is then simply the identity matrix. Let us write $\hat{\mathbf{Z}}^{\text{pairs:main:2}} \in \mathbb{R}^{\binom{[n]}{\leq 3} \times \binom{[n]}{\leq 3}}$ for the matrix arising from $\tilde{\mathbb{E}}^{\text{pairs:main:2}}$ (the hat serving as a reminder that this is a pseudomoment in the standard basis rather than the harmonic basis). Then, $\hat{\mathbf{Z}}^{\text{pairs:main:2}}$ is blockwise a sum of CGM terms corresponding to diagrams F , each of which consists only of pairs and has at most one sided pair. By and Propositions 10.13.12 and 10.13.13, the norm of such a CGM is at most $\|\mathbf{M}^2\|^3 \|\mathbf{M}^2\|_F \leq \|\mathbf{M}\|^6 \|\mathbf{M}^2\|_F$, the operator norm terms accounting for the non-sided pairs and the Frobenius norm term for the sided pair. The total number of such CGM terms across all blocks is 14: $2! = 2$ from the block with $|S| = |T| = 2$, 3 from the block with $|S| = 1, |T| = 3$, 3 from the block with $|S| = 3, |T| = 1$, $3! = 6$ from the block with $|S| = |T| = 3$. Therefore, we have

$$\|2(1 - c)t_{\text{pow}} \hat{\mathbf{Z}}^{\text{pairs:main:2}}\| \leq 28t_{\text{pow}} \|\mathbf{M}\|^6 \|\mathbf{M}^2\|_F \leq \frac{c}{2} \tag{10.212}$$

by our choice of c , concluding the proof of the claim.

Proof of (10.191): We again consider the pseudomoment matrix of the left-hand side, written in the standard monomial basis, and the arising from $\tilde{\mathbb{E}}^{\text{id}}$ is again the identity matrix. Let us write $\hat{\mathbf{Z}}^{\text{pairs:err}} \in \mathbb{R}^{\binom{[n]}{\leq 3} \times \binom{[n]}{\leq 3}}$ for the matrix arising from $\tilde{\mathbb{E}}^{\text{pairs:err}}$. The entries of this

matrix are as follows. First, $\widehat{Z}_{S,T}^{\text{pairs:err}} = 0$ whenever $|S| + |T| < 4$, $|S|$ and $|T|$ have different parity, or $|S \cap T| = 0$. Also, $\widehat{Z}^{\text{pairs:err}}$ is symmetric. The remaining entries are given, writing $\mathbf{H} = \mathbf{M}^2$ here to lighten the notation, by

$$\widehat{Z}_{\{i\}\{i,j,k\}}^{\text{pairs:err}} = \widehat{Z}_{\{i,j\}\{i,k\}}^{\text{pairs:err}} = -H_{jk} - 2H_{ij}H_{ik}, \quad (10.213)$$

$$\begin{aligned} \widehat{Z}_{\{i,j,k\}\{i,\ell,m\}}^{\text{pairs:err}} &= -2H_{ij}H_{ik}H_{\ell m} - 2H_{ij}H_{i\ell}H_{km} - 2H_{ij}H_{im}H_{k\ell} \\ &\quad - 2H_{ik}H_{i\ell}H_{jm} - 2H_{ik}H_{im}H_{j\ell} - 2H_{i\ell}H_{im}H_{jk} \\ &\quad \text{for } j, k, \ell, m \text{ distinct,} \end{aligned} \quad (10.214)$$

$$\begin{aligned} \widehat{Z}_{\{i,j,k\}\{i,j,\ell\}}^{\text{pairs:err}} &= -H_{k\ell} - 2H_{ij}^2H_{k\ell} - 2H_{ik}H_{i\ell} - 2H_{jk}H_{j\ell} \\ &\quad - 4H_{ij}H_{ik}H_{j\ell} - 4H_{ij}H_{i\ell}H_{jk}. \end{aligned} \quad (10.215)$$

Accordingly, we find the entrywise bounds

$$|\widehat{Z}_{\{i\}\{i,j,k\}}^{\text{pairs:err}}| \leq 3\epsilon_{\text{offdiag}}(\mathbf{M}^2) \quad (10.216)$$

$$|\widehat{Z}_{\{i,j\}\{i,k\}}^{\text{pairs:err}}| \leq 3\epsilon_{\text{offdiag}}(\mathbf{M}^2) \text{ if } j, k \text{ distinct,} \quad (10.217)$$

$$|\widehat{Z}_{\{i,j\}\{i,j\}}^{\text{pairs:err}}| \leq 3 \quad (10.218)$$

$$|\widehat{Z}_{\{i,j,k\}\{i,\ell,m\}}^{\text{pairs:err}}| \leq 12\epsilon_{\text{offdiag}}(\mathbf{M}^2)^3 \text{ if } j, k, \ell, m \text{ distinct,} \quad (10.219)$$

$$|\widehat{Z}_{\{i,j,k\}\{i,j,\ell\}}^{\text{pairs:err}}| \leq 15\epsilon_{\text{offdiag}}(\mathbf{M}^2) \text{ if } k, \ell \text{ distinct,} \quad (10.220)$$

$$|\widehat{Z}_{\{i,j,k\}\{i,j,k\}}^{\text{pairs:err}}| \leq 15. \quad (10.221)$$

Let us write $\widehat{Z}^{\text{pairs:err}[\ell,m]}$ for the submatrix of $\widehat{Z}^{\text{pairs:err}}$ indexed by $|S| = \ell$ and $|T| = m$. By the Gershgorin circle theorem, we then find the bounds

$$\|\widehat{Z}^{\text{pairs:err}[2,2]}\| \leq 3 + 6n\epsilon_{\text{offdiag}}(\mathbf{M}^2) \quad (10.222)$$

$$\|\widehat{Z}^{\text{pairs:err}[3,3]}\| \leq 15 + 45n\epsilon_{\text{offdiag}}(\mathbf{M}^2) + 36n^2\epsilon_{\text{offdiag}}(\mathbf{M}^2)^3, \quad (10.223)$$

and by the “rectangular Gershgorin” bound of Proposition 10.12.1

$$\|\hat{\mathbf{Z}}^{\text{pairs:err}[1,3]}\| \leq \sqrt{3n^2\epsilon_{\text{offdiag}}(\mathbf{M}^2) \cdot 3\epsilon_{\text{offdiag}}(\mathbf{M}^2)} = 3n\epsilon_{\text{offdiag}}(\mathbf{M}^2). \quad (10.224)$$

Finally, by Proposition 10.12.4, we combine these bounds to find

$$\begin{aligned} \|\hat{\mathbf{Z}}^{\text{pairs:err}}\| &\leq \|\hat{\mathbf{Z}}^{\text{pairs:err}[2,2]}\| + \|\hat{\mathbf{Z}}^{\text{pairs:err}[3,3]}\| + 2\|\hat{\mathbf{Z}}^{\text{pairs:err}[2,2]}\| \\ &\leq 18 + 57n\epsilon_{\text{offdiag}}(\mathbf{M}^2) + 36n^2\epsilon_{\text{offdiag}}(\mathbf{M}^2)^3. \end{aligned} \quad (10.225)$$

Therefore,

$$\|2(1-c)t_{\text{pow}}\hat{\mathbf{Z}}^{\text{pairs:err}}\| \leq 114t_{\text{pow}}(1+n\epsilon_{\text{offdiag}}(\mathbf{M}^2)+n^2\epsilon_{\text{offdiag}}(\mathbf{M}^2)^3) \leq \frac{c}{2} \quad (10.226)$$

by the definition of c , concluding the proof of the claim. \square

10.12 MISCELLANEOUS BOUNDS

We collect here some simple technical results we have used in the proofs above.

10.12.1 MATRIX INEQUALITIES

We first give the following matrix norm inequality that is effective for sparse matrices. Recall that the ∞ -norm of a matrix is defined as $\|\mathbf{A}\|_\infty = \max_{\mathbf{x} \neq 0} \|\mathbf{A}\mathbf{x}\|_\infty / \|\mathbf{x}\|_\infty = \max_i \sum_j |A_{ij}|$.

Proposition 10.12.1. *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$. Then,*

$$\|\mathbf{A}\| \leq \sqrt{\|\mathbf{A}\|_\infty \|\mathbf{A}^\top\|_\infty} = \sqrt{\left(\max_{i=1}^m \sum_{j=1}^n |A_{ij}| \right) \left(\max_{j=1}^n \sum_{i=1}^m |A_{ij}| \right)}. \quad (10.227)$$

Proof. We have $\|\mathbf{A}\| = \sigma_{\max}(\mathbf{A}) = \sqrt{\lambda_{\max}(\mathbf{A}\mathbf{A}^\top)}$. By the Gershgorin circle theorem, we may bound $\lambda_{\max}(\mathbf{A}\mathbf{A}^\top) \leq \|\mathbf{A}\mathbf{A}^\top\|_\infty$. Since the ∞ -norm on matrices is induced as an operator norm by the ℓ^∞ vector norm, it is submultiplicative, whereby $\|\mathbf{A}\mathbf{A}^\top\|_\infty \leq \|\mathbf{A}\|_\infty \|\mathbf{A}^\top\|_\infty$. \square

Remark 10.12.2. *This inequality is tight for any $\mathbf{A} \in \{0, 1\}^{m \times n}$ where every row has exactly one 1, but every column can have arbitrary numbers of 1's. The extremes in this class of matrices are the identity on the one hand, and $\mathbf{1}e_k^\top$ on the other (for a standard basis vector e_k).*

The following result on block matrix norms will also be useful.

Proposition 10.12.3. *Suppose $\mathbf{A} \in \mathbb{R}^{m \times n}$ is partitioned into $a \times b$ blocks $\mathbf{A}^{[k, \ell]} \in \mathbb{R}^{m_k \times n_\ell}$ for $k \in [a]$ and $\ell \in [b]$, where $\sum_{k=1}^a m_k = m$ and $\sum_{\ell=1}^b n_\ell = n$. Let $\tilde{\mathbf{A}} \in \mathbb{R}^{a \times b}$ have entries $\tilde{A}_{k\ell} = \|\mathbf{A}^{[k, \ell]}\|$. Then, $\|\mathbf{A}\| \leq \|\tilde{\mathbf{A}}\|$.*

Proof. If $\mathbf{x} \in \mathbb{R}^m$ and $\mathbf{y} \in \mathbb{R}^n$ are partitioned into vectors \mathbf{x}_k and \mathbf{y}_ℓ with compatible sizes to the blocks of \mathbf{A} , then we have

$$|\mathbf{x}^\top \mathbf{A} \mathbf{y}| = \left| \sum_{k, \ell} \mathbf{x}_k^\top \mathbf{A}^{[k, \ell]} \mathbf{y}_\ell \right| \leq \sum_{k, \ell} \|\mathbf{x}_k\|_2 \|\mathbf{y}_\ell\|_2 \|\mathbf{A}^{[k, \ell]}\|. \quad (10.228)$$

If $\|\mathbf{x}\| = \|\mathbf{y}\| = 1$, then the right-hand side is a bilinear form of unit vectors with $\tilde{\mathbf{A}}$, and the result follows. \square

The following other relative of the Gershgorin circle theorem gives a straightforward bound on block matrix norms.

Proposition 10.12.4. *Suppose $\mathbf{A} \in \mathbb{R}^{m \times n}$ is partitioned into blocks $\mathbf{A}^{[k, \ell]} \in \mathbb{R}^{m_k \times n_\ell}$ where $\sum m_k = m$ and $\sum n_\ell = n$. Then, $\|\mathbf{A}\| \leq \sum_{k, \ell} \|\mathbf{A}^{[k, \ell]}\|$.*

Proof. If $\mathbf{x} \in \mathbb{R}^m$ and $\mathbf{y} \in \mathbb{R}^n$ are partitioned into vectors \mathbf{x}_k and \mathbf{y}_ℓ with compatible sizes

to the blocks of \mathbf{A} , then we have

$$|\mathbf{x}^\top \mathbf{A} \mathbf{y}| = \left| \sum_{k,\ell} \mathbf{x}_k^\top \mathbf{A}^{[k,\ell]} \mathbf{y}_\ell \right| \leq \sum_{k,\ell} \|\mathbf{x}_k\|_2 \|\mathbf{y}_\ell\|_2 \|\mathbf{A}^{[k,\ell]}\|. \quad (10.229)$$

Thus, letting \mathbf{A}' be the matrix of norms of the blocks of \mathbf{A} , we have $\|\mathbf{A}\| \leq \|\mathbf{A}'\|$. The result then follows since if $\|\mathbf{v}\| = \|\mathbf{w}\| = 1$, then $\mathbf{v}^\top \mathbf{A}' \mathbf{w} \leq (\max_k |v_k|)(\max_\ell |w_\ell|)(\sum_{k,\ell} |A'_{k\ell}|) \leq \sum_{k,\ell} |A'_{k\ell}|$. \square

10.12.2 COMBINATORIAL BOUNDS

We next prove several coarse bounds on combinatorial quantities arising in our arguments. We begin with bounds on the coefficients of forests that arise in our calculations. The only tool required for these is that $(a + b)! \geq a! b!$, which follows from observing that $\binom{a+b}{a} \geq 1$.

Proposition 10.12.5. *For any $F \in \mathcal{F}(d)$, $|\mu(F)| \leq (3d)!$.*

Proof. We have

$$\begin{aligned} |\mu(F)| &= \prod_{v \in V^\square} (\deg(v) - 2)! \\ &\leq \left(\sum_{v \in V} \deg(v) \right)! \\ &\leq \left(2 \cdot \frac{3}{2} d \right)!, \end{aligned} \quad (10.230)$$

the last step following by Corollary 10.12.10. \square

Proposition 10.12.6. *For any $F \in \mathcal{F}(d, d)$ a balanced bowtie forest ribbon diagram, $|\xi(F)| \leq (d - 1)! d! \leq d^{2d}$.*

Proof. We have

$$\begin{aligned}
|\xi(F)| &= \prod_{\substack{C \in \text{conn}(F) \\ C \text{ balanced bowtie} \\ \text{on } 2k \text{ leaves}}} (k-1)!k! \\
&\leq \left(\sum_{\substack{C \in \text{conn}(F) \\ C \text{ balanced bowtie} \\ \text{on } 2k \text{ leaves}}} (k-1) \right)! \left(\sum_{\substack{C \in \text{conn}(F) \\ C \text{ balanced bowtie} \\ \text{on } 2k \text{ leaves}}} k \right)! \\
&= (d - |\text{conn}(F)|)! d! \\
&\leq (d-1)! d!, \tag{10.231}
\end{aligned}$$

completing the proof. \square

We next give some bounds on the cardinalities of various sets of combinatorial objects arising in our analysis.

Proposition 10.12.7. $|\text{Part}([d])| \leq d^d$.

Proof. For $d \leq 5$, the inequality may be verified directly. Assuming $d \geq 6$, we begin with a “stars and bars” argument: all partitions of $[d]$ may be obtained by writing the numbers $1, \dots, d$ in some order, and then choosing some subset of the $d-1$ possible positions between two numbers where a boundary between parts of the partition may be placed. Therefore, $|\text{Part}([d])| \leq 2^d d!$. For $d \geq 6$, we have $d! \leq (d/2)^d$, and the result follows. \square

We note that the numbers $|\text{Part}([d])|$ are known as the *Bell numbers*, for which many more precise asymptotics are known [BT10]. We prefer to give a simple hands-on proof here, which matches the correct coarse scaling $\log |\text{Part}([d])| = \Theta(d \log d)$.

Proposition 10.12.8. For any $\sigma, \tau \in \text{Part}([d])$, $|\text{Plans}(\sigma, \tau)| \leq d!$.

Proof. For every $\mathbf{D} \in \text{Plans}(\sigma, \tau)$, there exists a bijection $f : [d] \rightarrow [d]$ for which $D_{A,B} = \#\{i \in A : f(i) \in B\}$. Therefore, the total number of such \mathbf{D} is at most the total number of bijections of $[d]$ with itself, which is $d!$. \square

The following simple and general result gives a bound on the number of vertices in a tree if the degrees of internal vertices are bounded below.

Proposition 10.12.9. *Suppose $T = (V, E)$ is a tree with ℓ leaves. Assume that, for any internal vertex v of T , $\deg(v) \geq k \geq 3$. Then,*

$$|V| < \frac{k-1}{k-2}\ell. \quad (10.232)$$

Proof. We count the number of edges $|E|$ in two ways, and then apply the degree bound:

$$\begin{aligned} |E| &= |V| - 1 \\ &= \frac{1}{2} \sum_{v \in V} \deg(v) \\ &\geq \frac{1}{2} (\ell + k(|V| - \ell)) \\ &= \frac{k}{2}|V| - \frac{k-1}{2}\ell. \end{aligned} \quad (10.233)$$

Solving for $|V|$, we have

$$|V| \leq \frac{k-1}{k-2}\ell - \frac{2}{k-2}, \quad (10.234)$$

and the result follows. □

Corollary 10.12.10. *For any $F \in \mathcal{F}(d)$, the number of vertices and edges in F are both at most $\frac{3}{2}d$, and the number of \square vertices is at most $\frac{1}{2}d$.*

This allows us to bound the number of good forests, as follows.

Proposition 10.12.11. *For d even, $(d/2)! \leq |\mathcal{F}(d)| \leq 2(\frac{3}{2}d)^{\frac{3}{2}d}$.*

Proof. For the lower bound, we simply note that any matching of $1, \dots, d/2$ with $d/2 + 1, \dots, d$ corresponds to a distinct element of $\mathcal{F}(d)$ consisting of the corresponding forest

all of whose connected components are pairs. Since there are $(d/2)!$ such matchings, the bound follows.

For the upper bound, Theorem 4.1 of [Moo70], attributed there to Rényi, gives the following explicit formula for the number of labelled forests on n nodes with k connected components, which we denote $f_{n,k}$ (this is a generalization of Cayley's well-known formula counting labelled trees, which is the case $k = 1$):

$$f_{n,k} = \binom{n}{k} \sum_{i=0}^k \left(-\frac{1}{2}\right)^i (k+i) i! \binom{k}{i} \binom{n-k}{i} n^{n-k-i-1} \quad (10.235)$$

from which by coarse bounds we find

$$\leq \frac{n^k}{k!} \sum_{i=0}^k \frac{k+i}{2^i} \frac{k!}{(k-i)!} \frac{n^i}{i!} n^{n-k-i-1} \quad (10.236)$$

$$= n^{n-1} \sum_{i=0}^k \frac{k+i}{2^i (k-i)! i!} \quad (10.237)$$

$$\leq 2n^{n-1}, \quad (10.238)$$

where the final inequality may be checked by verifying by hand for $k \leq 4$, and for $k \geq 5$ bounding the inner term by $2k/(\lfloor k/2 \rfloor! \lceil k/2 \rceil!) \leq 2$. Thus the number of labelled forests on n nodes with any number of connected components, which equals $\sum_{k=1}^n f_{n,k}$, is at most $2n^n$. By Corollary 10.12.10, the total number of vertices in $F \in \mathcal{F}(d)$ is at most $\frac{3}{2}d$. Since there are no isolated vertices in such F , we may always view F as embedded uniquely into a fully labelled forest on exactly $\frac{3}{2}d$ vertices by adding isolated vertices and labelling internal vertices with any deterministic procedure. Thus, $|\mathcal{F}(d)|$ is at most the number of labelled forests on $\frac{3}{2}d$ vertices, and the result follows. \square

We include the lower bound above to emphasize that the upper bound correctly identifies the coarse behavior $\log |\mathcal{F}(d)| = \Theta(d \log d)$. This suggests that, without much more careful

proof techniques, the d^d behavior of the leading factor in the condition of Theorem 10.2.3 cannot be improved.

10.13 CALCULUS OF CONTRACTIVE GRAPHICAL MATRICES

Finally, we present some general tools for working with contractive graphical matrices (henceforth CGMs, as in the main text). In fact, to make some technicalities easier to work around, we use a more general definition, which allows different edges of a ribbon diagram to be labelled with different matrices and also allows for the left and right index subsets to overlap.

Definition 10.13.1 (Generalized ribbon diagram). *Let $G = ((\mathcal{L} \cup \mathcal{R}) \sqcup V^\square, E)$ be a graph with two types of vertices, \bullet and \square , whose subsets are $V^\bullet = \mathcal{L} \cup \mathcal{R}$ and V^\square . Suppose that G is equipped with labellings $\kappa_{\mathcal{L}} : \mathcal{L} \rightarrow [|\mathcal{L}|]$ and $\kappa_{\mathcal{R}} : \mathcal{R} \rightarrow [|\mathcal{R}|]$. Suppose that for every $e = \{x, y\} \in E$, we have a matrix $M^{(x,y)} \in \mathbb{R}^{n(x) \times n(y)}$, for some $n : V \rightarrow \mathbb{N}$, which satisfies $M^{(y,x)} = M^{(x,y)\top}$. Note that n is determined by the collection of matrices $M^{(x,y)}$, provided that their dimensions satisfy the appropriate equalities. We call such a collection of matrices a compatible matrix labelling of the edges of G , noting that a matrix is associated to each oriented edge in such a labelling. For the course of this appendix, we call such labelled G a ribbon diagram, instead of the weaker definition from the main text.*

Definition 10.13.2 (Generalized CGM). *Let G be a ribbon diagram. Given $\mathbf{a} \in \prod_{v \in V^\square} [n(v)]$, $\mathbf{s} \in \prod_{i \in [|\mathcal{L}|]} [n(\kappa_{\mathcal{L}}^{-1}(i))]$, and $\mathbf{t} \in \prod_{j \in [|\mathcal{R}|]} [n(\kappa_{\mathcal{R}}^{-1}(j))]$ such that, for all $i \in \mathcal{L} \cap \mathcal{R}$, $s(\kappa_{\mathcal{L}}(i)) = t(\kappa_{\mathcal{R}}(j))$, define $f_{\mathbf{a}, \mathbf{s}, \mathbf{t}} : V \rightarrow \mathbb{N}$ by*

$$f_{\mathbf{a}, \mathbf{s}, \mathbf{t}}(x) = \begin{cases} s(\kappa_{\mathcal{L}}(x)) & \text{if } x \in \mathcal{L}, \\ t(\kappa_{\mathcal{R}}(x)) & \text{if } x \in \mathcal{R}, \\ \mathbf{a}(x) & \text{if } x \in V^\square. \end{cases} \quad (10.239)$$

Note that we always have $f_{a,s,t}(x) \in [n(x)]$. Then, the contractive graphical matrix associated to the ribbon diagram G labelled by the matrices $M^{(x,y)}$ has entries

$$Z_{s,t}^G = \prod_{i \in \mathcal{L} \cap \mathcal{R}} \mathbb{1}\{s(\kappa_{\mathcal{L}}(i)) = t(\kappa_{\mathcal{R}}(i))\} \sum_{a \in \prod_{v \in V \square} [n(v)]} \prod_{\{x,y\} \in E} M_{f_{a,s,t}(x), f_{a,s,t}(y)}^{(x,y)}. \quad (10.240)$$

We note that, for the purposes of this appendix, we always will think of CGMs as being labelled by tuples rather than sets. Since set-indexed CGMs are submatrices of tuple-indexed ones, all norm bounds will immediately be inherited by the set-indexed CGMs encountered in the main text.

The general goal we pursue in the following sections is to develop some general tools for connecting the graphical structure of a ribbon diagram and the matrix structure of its CGM.

10.13.1 CONNECTED COMPONENTS AND TENSORIZATION

We first consider the effect of a diagram being disconnected on the CGM. In this case, it is simple to see that the expression (10.240) factorizes, and therefore the CGM decomposes as a tensor product. We give a precise statement below, taking into account the ordering of indices.

Proposition 10.13.3. *Let $G = ((\mathcal{L} \cup \mathcal{R}) \sqcup V \square, E)$ be a ribbon diagram with connected components G_1, \dots, G_m , where $V(G_\ell) \cap \mathcal{L} =: \mathcal{L}_\ell$ and $V(G_\ell) \cap \mathcal{R} =: \mathcal{R}_\ell$. Define $\kappa_{\mathcal{L}_\ell} : \mathcal{L}_\ell \rightarrow [|\mathcal{L}_\ell|]$ and $\kappa_{\mathcal{R}_\ell} : \mathcal{R}_\ell \rightarrow [|\mathcal{R}_\ell|]$ to be the labellings inherited from $\kappa_{\mathcal{L}}$ and $\kappa_{\mathcal{R}}$, i.e.,*

$$\kappa_{\mathcal{L}_\ell}(i) := \#\{i' \in \mathcal{L}_\ell : \kappa_{\mathcal{L}}(i') \leq \kappa_{\mathcal{L}}(i)\}, \quad (10.241)$$

$$\kappa_{\mathcal{R}_\ell}(j) := \#\{j' \in \mathcal{R}_\ell : \kappa_{\mathcal{R}}(j') \leq \kappa_{\mathcal{R}}(j)\}. \quad (10.242)$$

Equipped with these labellings, view G_1, \dots, G_m as ribbon diagrams. Let $\pi_{\mathcal{L}} \in \text{Sym}(|\mathcal{L}|)$ and

$\pi_{\mathcal{R}} \in \text{Sym}(|\mathcal{R}|)$ be the permutations with

$$\begin{aligned} & (\pi_{\mathcal{L}}(1), \dots, \pi_{\mathcal{L}}(|\mathcal{L}|)) \\ &= \left(\kappa_{\mathcal{L}}(\kappa_{\mathcal{L}_1}^{-1}(1)), \dots, \kappa_{\mathcal{L}}(\kappa_{\mathcal{L}_1}^{-1}(|\mathcal{L}_1|)), \dots, \kappa_{\mathcal{L}}(\kappa_{\mathcal{L}_m}^{-1}(1)), \dots, \kappa_{\mathcal{L}}(\kappa_{\mathcal{L}_m}^{-1}(|\mathcal{L}_m|)) \right) \\ & (\pi_{\mathcal{R}}(1), \dots, \pi_{\mathcal{R}}(|\mathcal{R}|)) \\ &= \left(\kappa_{\mathcal{R}}(\kappa_{\mathcal{R}_1}^{-1}(1)), \dots, \kappa_{\mathcal{R}}(\kappa_{\mathcal{R}_1}^{-1}(|\mathcal{R}_1|)), \dots, \kappa_{\mathcal{R}}(\kappa_{\mathcal{R}_m}^{-1}(1)), \dots, \kappa_{\mathcal{R}}(\kappa_{\mathcal{R}_m}^{-1}(|\mathcal{R}_m|)) \right) \end{aligned}$$

Let $\sigma_{\mathcal{L}} \in \text{Sym}([n]^{|\mathcal{L}|})$ map $(a_1, \dots, a_{|\mathcal{L}|}) \mapsto (a_{\pi_{\mathcal{L}}^{-1}(1)}, \dots, a_{\pi_{\mathcal{L}}^{-1}(|\mathcal{L}|)})$ and $\sigma_{\mathcal{R}} \in \text{Sym}([n]^{|\mathcal{R}|})$ map $(a_1, \dots, a_{|\mathcal{R}|}) \mapsto (a_{\pi_{\mathcal{R}}^{-1}(1)}, \dots, a_{\pi_{\mathcal{R}}^{-1}(|\mathcal{R}|)})$. Finally, let $\Pi_{\mathcal{L}} \in \mathbb{R}^{[n]^{|\mathcal{L}|} \times [n]^{|\mathcal{L}|}}$ and $\Pi_{\mathcal{R}} \in \mathbb{R}^{[n]^{|\mathcal{R}|} \times [n]^{|\mathcal{R}|}}$ be the permutation matrices of $\sigma_{\mathcal{L}}$ and $\sigma_{\mathcal{R}}$, respectively. Then,

$$\mathbf{Z}^G = \Pi_{\mathcal{L}} \left(\bigotimes_{\ell=1}^m \mathbf{Z}^{G_{\ell}} \right) \Pi_{\mathcal{R}}^{\top}.$$

This fact will be most useful when bounding the difference in operator norm incurred by replacing each connected component G_i by some other diagram in terms of the differences of the smaller CGMs corresponding to each connected component taken in isolation.

Proposition 10.13.4. *Let G be a ribbon diagram with connected components G_1, \dots, G_m , where $V(G_{\ell}) \cap \mathcal{L} = \mathcal{L}_{\ell}$ and $V(G_{\ell}) \cap \mathcal{R} = \mathcal{R}_{\ell}$. Suppose H_1, \dots, H_m are other ribbon diagrams on $(\mathcal{L}_{\ell}, \mathcal{R}_{\ell}, V^{\square}(G_{\ell}))$, and write H for the union diagram of the H_i . Then,*

$$\|\mathbf{Z}^G - \mathbf{Z}^H\| \leq \sum_{\ell=1}^m \|\mathbf{Z}^{G_{\ell}} - \mathbf{Z}^{H_{\ell}}\| \prod_{\ell'=1}^{\ell-1} \|\mathbf{Z}^{H_{\ell'}}\| \prod_{\ell'=\ell+1}^m \|\mathbf{Z}^{G_{\ell'}}\|. \quad (10.243)$$

Proof. Following the notation of Proposition 10.13.3, we can write a telescoping sum,

$$\mathbf{Z}^G - \mathbf{Z}^H = \Pi_{\mathcal{L}} \left(\sum_{\ell=1}^m \bigotimes_{\ell'=1}^{\ell-1} \mathbf{Z}^{H_{\ell'}} \otimes (\mathbf{Z}^{G_{\ell}} - \mathbf{Z}^{H_{\ell}}) \otimes \bigotimes_{\ell'=\ell+1}^m \mathbf{Z}^{G_{\ell'}} \right) \Pi_{\mathcal{R}} \quad (10.244)$$

The bound then follows by the triangle inequality and the tensorization of the operator

norm. □

10.13.2 SPLITTING

We describe two operations on a ribbon diagram, which we call “splittings,” that add edges to the diagram without changing the associated CGM. Later this will allow us to perform some regularizing operations on a ribbon diagram’s graph structure when making arguments about its CGM.

The first type of splitting lets us expand a diagram and thereby eliminate the intersection of \mathcal{L} and \mathcal{R} by adding redundant vertices and suitable adjacencies.

Proposition 10.13.5 (Intersection splitting). *Let $G = ((\mathcal{L} \cup \mathcal{R}) \sqcup V^\square, E)$ be a ribbon diagram with $(M^{(x,y)})$ a compatible labelling of its edges. Write $\mathcal{L} \cap \mathcal{R} = \{v_1, \dots, v_n\}$. Let G' be another labelled ribbon diagram, formed by adding new vertices $\{v'_1, \dots, v'_n\}$ to G , setting $V^\square(G') = V^\square(G)$, $\mathcal{L}(G') = \mathcal{L}(G)$, and $\mathcal{R}(G') = \mathcal{R}(G) \setminus \mathcal{L}(G) \cup \{v'_1, \dots, v'_n\}$, and adding edges $\{v_i, v'_i\}$ labelled by the matrix $I_{n(v_i)}$ for each $i \in [n]$. Then, $Z^G = Z^{G'}$.*

The second type of splitting lets us represent a factorization of a matrix labelling an edge by subdividing that edge with intermediate vertices.

Proposition 10.13.6 (Edge splitting). *Let $G = ((\mathcal{L} \cup \mathcal{R}) \sqcup V^\square, E)$ be a ribbon diagram with $(M^{(x,y)})$ a compatible labelling of its edges. Suppose $x \sim z$ in G , and there exist matrices $A \in \mathbb{R}^{n(x) \times n}$, $B \in \mathbb{R}^{n \times n(z)}$ such that $M^{(x,z)} = AB$. Let G' be another labelled ribbon diagram, with $\mathcal{L}(G') = \mathcal{L}(G)$, $\mathcal{R}(G') = \mathcal{R}(G)$, $V^\square(G') = V^\square(G) \cup \{y\}$ for a new vertex y , and $E(G') = E(G) \cup \{\{x, y\}, \{y, z\}\}$. Let $\{x, y\}$ in G' be labelled with the matrix $M^{(x,y)} = A$, and $\{y, z\}$ be labelled with the matrix $M^{(y,z)} = B$. Then, $Z^G = Z^{G'}$.*

Note that, as an especially useful special case, we may always take $n = n(x)$, $A = I_{n(x)}$, and $B = M^{(x,z)}$. This particular technique allows us to adjust the graph of the ribbon diagram without needing to find any special factorization of $M^{(x,z)}$.

10.13.3 PINNING, CUTTING, AND DIRECT SUM DECOMPOSITION

We next explore special operations that may be performed on the following special type of \square vertex in a ribbon diagram.

Definition 10.13.7 (Pinned vertex). *In a ribbon diagram $G = ((\mathcal{L} \cup \mathcal{R}) \sqcup V^\square, E)$ with dimension labels $n : V \rightarrow \mathbb{N}$, we call a vertex $v \in V^\square$ pinned if $n(v) = 1$.*

Any edge one of whose endpoints is a pinned vertex must be labelled with a vector, and in the formula (10.240) there is effectively no summation corresponding to a pinned vertex, since $[n(v)] = \{1\}$ whereby the vertex's assigned index is always the same—this is the reason for the term “pinned.”

In terms of manipulations of the ribbon diagram G , the important special property of a pinned vertex is that it allows the diagram to be “cut” at that vertex without changing the resulting CGM.

Proposition 10.13.8 (Cutting). *Let $G = ((\mathcal{L} \cup \mathcal{R}) \sqcup V^\square, E)$ be a ribbon diagram, and let $v \in V^\square$ be pinned. Suppose $\deg(v) = m$, enumerate the neighbors of v as w_1, \dots, w_m , and suppose these edges are labelled with vectors $\mathbf{m}_i \in \mathbb{R}^{n(w_i)}$ for $i = 1, \dots, m$. Let G' be another ribbon diagram, formed by removing v from G and adding new vertices v'_1, \dots, v'_m to V^\square , with $n(v'_i) = 1$, v'_i adjacent to only w_i , and this edge labelled by \mathbf{m}_i . Then, $\mathbf{Z}^G = \mathbf{Z}^{G'}$.*

Note that, after splitting, every pinned vertex has degree 1. In our case, when we work with tree ribbon diagrams, this means that every pinned vertex is a leaf of the resulting forest, a property that will be important in our analysis.

Finally, we show two ways that pinned vertices arise naturally from matrix-labelled ribbon diagrams where no vertex has dimension label 1 to begin with. The first, simpler situation is where an edge is labelled with a rank 1 matrix.

Proposition 10.13.9. *Let $G = ((\mathcal{L} \cup \mathcal{R}) \sqcup V^\square, E)$ be a ribbon diagram, and suppose $\{v, w\} \in E$ is labelled with $M^{(v,w)} = \mathbf{x}\mathbf{y}^\top$. Let G' be formed by adding a vertex x along this edge between v and w , setting $n(x) = 1$, and setting $M^{(v,x)} = \mathbf{x}$ and $M^{(x,w)} = \mathbf{y}$. Then, $\mathbf{Z}^G = \mathbf{Z}^{G'}$.*

Proof. The result follows from applying Proposition 10.13.6 to the edge $\{v, w\}$ using the given rank one factorization. □

The second, perhaps more natural, situation is that the CGM of any ribbon diagram with $\mathcal{L} \cap \mathcal{R} \neq \emptyset$ may be written as a direct sum of CGMs with those vertices pinned (that is, as a block diagonal matrix with these CGMs as the diagonal blocks).

Proposition 10.13.10 (Direct sum decomposition). *Let $G = ((\mathcal{L} \cup \mathcal{R}) \sqcup V^\square, E)$ be a ribbon diagram, and enumerate $\mathcal{L} \cap \mathcal{R} = \{v_1, \dots, v_m\}$. Given $\mathbf{a} \in \prod_{i=1}^m [n(v_i)]$, let $G[\mathbf{a}]$ be the diagram formed by moving each v_i to V^\square , letting $n(v_i) = 1$, and labelling each edge incident with v_i , say $\{v_i, x\}$, with the vector ${}^{(v_i,x)}$ equal to the a_i th row of $M^{(v_i,x)}$. If there is an edge between v_i and v_j , then in $G[\mathbf{a}]$ it is labelled with the constant $M_{a_i a_j}^{(v_i, v_j)}$. Then, there exist permutation matrices $\Pi_{\mathcal{L}}$ and $\Pi_{\mathcal{R}}$ such that*

$$\mathbf{Z}^G = \Pi_{\mathcal{L}} \left(\bigoplus_{\mathbf{a} \in \prod_{i=1}^m [n(v_i)]} \mathbf{Z}^{G(\mathbf{a})} \right) \Pi_{\mathcal{R}}^\top. \quad (10.245)$$

While formally a pinned vertex is a \square vertex, in this setting we see that it behaves more like a \bullet vertex whose index is fixed instead of varying with the matrix indices.

10.13.4 FACTORIZATION

We now arrive at perhaps the most useful and important manipulation of CGMs via ribbon diagrams. Namely, certain graphical decompositions of ribbon diagrams correspond to factorizations of CGMs into products of simpler CGMs.

Proposition 10.13.11. *Let $G = (V, E)$ be a ribbon diagram with $V = (\mathcal{L} \cup \mathcal{R}) \sqcup V^\square$. Suppose that V admits a partition $V = A \sqcup B \sqcup C$, such that the following properties hold:*

1. $\mathcal{L} \subseteq A$.
2. $\mathcal{R} \subseteq C$.
3. $\partial_{\text{out}}A, \partial_{\text{out}}C \subseteq B$.

(Here ∂_{out} denotes the “outer boundary” of a set, those vertices not in the set but with a neighbor in the set.) Suppose also that the edges within B admit a partition $E(B, B) = E^A \sqcup E^B \sqcup E^C$, where $E^A \subseteq E(\partial_{\text{out}}A, \partial_{\text{out}}A)$ and $E^C \subseteq E(\partial_{\text{out}}C, \partial_{\text{out}}C)$. Define the following ancillary ribbon diagrams:

1. $G[A]$ with vertex triple $(\mathcal{L}, \partial_{\text{out}}A, A \setminus \mathcal{L})$ and edges $E(A, A \cup \partial_{\text{out}}A) \cup E^A$;
2. $G[B]$ with vertex triple $(\partial_{\text{out}}A, \partial_{\text{out}}C, B \setminus \partial_{\text{out}}A \setminus \partial_{\text{out}}C)$ and edges E^B ; and
3. $G[C]$ with vertex triple $(\partial_{\text{out}}C, \mathcal{R}, C \setminus \mathcal{R})$ and edges $E(C, C \cup \partial_{\text{out}}C) \cup E^C$.

In these diagrams, $\partial_{\text{out}}A$ and $\partial_{\text{out}}C$ are given arbitrary labellings, but the same labelling each time that they appear in different diagrams. Then,

$$\mathbf{Z}^G = \mathbf{Z}^{G[A]} \mathbf{Z}^{G[B]} \mathbf{Z}^{G[C]}. \quad (10.246)$$

Proof. The proof is a direct verification by expanding the matrix multiplications and definitions of the CGMs involved. Note that, by assumption, since A and C are disjoint and $\mathcal{L} \subseteq A$ and $\mathcal{R} \subseteq C$, we must have $\mathcal{L} \cap \mathcal{R} = \emptyset$. We have

$$(\mathbf{Z}^{G[A]} \mathbf{Z}^{G[B]} \mathbf{Z}^{G[C]})_{s,t} = \sum_{\substack{a \in [n]^{\partial_{\text{out}}A} \\ c \in [n]^{\partial_{\text{out}}C}}} \mathbf{Z}_{s,a}^{G[A]} \mathbf{Z}_{a,c}^{G[B]} \mathbf{Z}_{c,t}^{G[C]}. \quad (10.247)$$

Given $\mathbf{a} \in [n]^{\partial_{\text{out}}A}$, $\mathbf{a}' \in [n]^{A \setminus \mathcal{L}}$, $\mathbf{b} \in [n]^{B \setminus \partial_{\text{out}}A \setminus \partial_{\text{out}}C}$, $\mathbf{c} \in [n]^{\partial_{\text{out}}C}$, $\mathbf{c}' \in [n]^{C \setminus \mathcal{R}}$ such that for $x \in \partial_{\text{out}}A \cap \partial_{\text{out}}C$ we have $a(x) = c(x)$, let us define $g = g_{\mathbf{a}, \mathbf{a}', \mathbf{b}, \mathbf{c}, \mathbf{c}', s, t} : V^\square \rightarrow [n]$ by

$$g(x) := \begin{cases} s(\kappa_{\mathcal{L}}(x)) & \text{if } x \in \mathcal{L}, \\ t(\kappa_{\mathcal{R}}(x)) & \text{if } x \in \mathcal{R}, \\ a(x) & \text{if } x \in \partial_{\text{out}}A, \\ a'(x) & \text{if } x \in A \setminus \mathcal{L}, \\ b(x) & \text{if } x \in B \setminus \partial_{\text{out}}A \setminus \partial_{\text{out}}C, \\ c(x) & \text{if } x \in \partial_{\text{out}}C, \\ c'(x) & \text{if } x \in C \setminus \mathcal{R}. \end{cases} \quad (10.248)$$

We then compute

$$\begin{aligned} (\mathbf{Z}^{G[A]} \mathbf{Z}^{G[B]} \mathbf{Z}^{G[C]})_{s,t} &= \sum_{\substack{\mathbf{a} \in [n]^{\partial_{\text{out}}A} \\ \mathbf{a}' \in [n]^{A \setminus \mathcal{L}} \\ \mathbf{c} \in [n]^{\partial_{\text{out}}C} \\ \mathbf{c}' \in [n]^{C \setminus \mathcal{R}}} \prod_{\{x,y\} \in E(A, A \cup \partial_{\text{out}}A) \cup E^A} M_{g(x), g(y)}^{(x,y)} \prod_{\{x,y\} \in E(C, C \cup \partial_{\text{out}}C) \cup E^C} M_{g(x), g(y)}^{(x,y)} \\ &\quad \prod_{x \in \partial_{\text{out}}A \cap \partial_{\text{out}}C} \mathbb{1}\{a(x) = c(x)\} \prod_{\{x,y\} \in E^B} M_{g(x), g(y)}^{(x,y)} \\ &= \sum_{\mathbf{a} \in [n]^{V^\square}} \prod_{\{x,y\} \in E} M_{f_{\mathbf{a}, s, t}(x), f_{\mathbf{a}, s, t}(y)}^{(x,y)} \\ &= Z_{s,t}^G, \end{aligned} \quad (10.249)$$

completing the proof. □

10.13.5 GENERAL-PURPOSE NORM BOUNDS

Our first application is to prove general-purpose bounds on the norms of CGMs based on ribbon diagram structure and the norms of constituent labelling matrices.

First, we show that norms multiply over connected components, as we have alluded to

in Remark 10.4.5 in the main text. This is a direct application of Proposition 10.13.3.

Proposition 10.13.12. *Let G be a ribbon diagram with connected components G_1, \dots, G_m , as in Proposition 10.13.3. Then, $\mathbf{Z}^G = \prod_{\ell=1}^m \|\mathbf{Z}^{G_\ell}\|$.*

The following bound is less trivial and is used repeatedly in our arguments.

Proposition 10.13.13. *Let $G = (V, E)$ be a ribbon diagram with $V = (\mathcal{L} \cup \mathcal{R}) \sqcup V^\square$. Suppose that V admits a partition $V = V_1 \sqcup \dots \sqcup V_m$ with $m \geq 2$, such that $V_1 = \mathcal{L}$, $V_m = \mathcal{R}$, and the following properties hold:*

1. *For every $v \in V_1 = \mathcal{L}$, there exists some $k > 1$ such that v has a neighbor in V_k .*
2. *Every $v \in V_m = \mathcal{R}$, there exists some $k < m$ such that v has a neighbor in V_k .*
3. *For every $1 < j < m$ and every $v \in V_j$, there exist $i < j < k$ such that v has a neighbor in V_i and a neighbor in V_k .*

Then,

$$\|\mathbf{Z}^G\| \leq \prod_{\{x, y\} \in E} \|\mathbf{M}^{(x, y)}\|. \quad (10.250)$$

Proof. Note that, by repeatedly applying Proposition 10.13.6 with edges labelled by an identity matrix, we may furthermore assume without loss of generality that every edge of G is either between two vertices of V_i , or between one vertex of V_i and one vertex of V_{i+1} for some i . Under this assumption, the three conditions in the statement may be rewritten as follows:

1. For every $v \in V_1 = \mathcal{L}$, v has a neighbor in V_2 .
2. Every $v \in V_m = \mathcal{R}$, v has a neighbor in V_{m-1} .
3. For every $1 < j < m$ and every $v \in V_j$, v has a neighbor in V_{j-1} and a neighbor in V_{j+1} .

Next, we proceed by induction on m . Suppose first that $m = 2$. Then, the assumptions imply that $V^\square = \emptyset$ and $\mathcal{L} \cap \mathcal{R} = \emptyset$. Let us enumerate the edges within \mathcal{L} , within \mathcal{R} , and between \mathcal{L} and \mathcal{R} as follows:

$$E(\mathcal{L}, \mathcal{L}) = \{\{i_1^{(1)}, i_1^{(2)}\}, \dots, \{i_a^{(1)}, i_a^{(2)}\}\} \text{ for } i_\ell^{(k)} \in \mathcal{L}, \quad (10.251)$$

$$E(\mathcal{R}, \mathcal{R}) = \{\{j_1^{(1)}, j_1^{(2)}\}, \dots, \{j_b^{(1)}, j_b^{(2)}\}\} \text{ for } j_\ell^{(k)} \in \mathcal{R}, \quad (10.252)$$

$$E(\mathcal{L}, \mathcal{R}) = \{\{i_1, j_1\}, \dots, \{i_c, j_c\}\} \text{ for } i_\ell \in \mathcal{L} \text{ and } j_\ell \in \mathcal{R}. \quad (10.253)$$

Then, we have

$$Z_{s,t}^G = \prod_{\ell=1}^a M_{s(\kappa_{\mathcal{L}}(i_\ell^{(1)})), s(\kappa_{\mathcal{L}}(i_\ell^{(2)}))}^{(i_\ell^{(1)}, i_\ell^{(2)})} \prod_{\ell=1}^b M_{t(\kappa_{\mathcal{R}}(j_\ell^{(1)})), t(\kappa_{\mathcal{R}}(j_\ell^{(2)}))}^{(j_\ell^{(1)}, j_\ell^{(2)})} \prod_{\ell=1}^c M_{s(\kappa_{\mathcal{L}}(i_\ell)), t(\kappa_{\mathcal{R}}(j_\ell))}^{(i_\ell, j_\ell)}. \quad (10.254)$$

Let us define an ancillary matrix

$$\hat{Z}_{s,t}^G := \prod_{\ell=1}^c M_{s(\kappa_{\mathcal{L}}(i_\ell)), t(\kappa_{\mathcal{R}}(j_\ell))}^{(i_\ell, j_\ell)}. \quad (10.255)$$

Then, we may write $Z^G = D^{\mathcal{L}} \hat{Z}^G D^{\mathcal{R}}$, for suitable diagonal matrices $D^{\mathcal{L}}$ and $D^{\mathcal{R}}$ having entries equal to the first two products above, respectively. Since every entry of a matrix is bounded by the matrix norm, we then have

$$\|Z^G\| \leq \|D^{\mathcal{L}}\| \cdot \|\hat{Z}^G\| \cdot \|D^{\mathcal{R}}\| \leq \|\hat{Z}^G\| \prod_{\{x,y\} \in E(\mathcal{L}, \mathcal{L}) \cup E(\mathcal{R}, \mathcal{R})} \|M^{(x,y)}\|. \quad (10.256)$$

For the remaining factor, by taking a singular value decomposition, we can factorize each labelling matrix as $M^{(x,y)} = U^{(x,y)\top} V^{(x,y)}$ such that $\|M^{(x,y)}\| = \|U^{(x,y)}\| \cdot \|V^{(x,y)}\|$ (by including the factor of the singular values in either of the singular vector matrices). Writing $u^{(x,y,i)}$ for the columns of $U^{(x,y)}$ and $v^{(x,y,i)}$ for the columns of $V^{(x,y)}$, we then

have $M_{ij}^{(x,y)} = \langle \mathbf{u}^{(x,y,i)}, \mathbf{v}^{(x,y,j)} \rangle$. Therefore,

$$\hat{\mathbf{Z}}_{s,t}^G = \prod_{\ell=1}^c \left\langle \mathbf{u}^{(i_\ell, j_\ell, s(\kappa_{\mathcal{L}}(i_\ell)))}, \mathbf{v}^{(i_\ell, j_\ell, t(\kappa_{\mathcal{R}}(j_\ell)))} \right\rangle = \left\langle \bigotimes_{\ell=1}^c \mathbf{u}^{(i_\ell, j_\ell, s(\kappa_{\mathcal{L}}(i_\ell)))}, \bigotimes_{\ell=1}^c \mathbf{v}^{(i_\ell, j_\ell, t(\kappa_{\mathcal{R}}(j_\ell)))} \right\rangle. \quad (10.257)$$

This writes $\hat{\mathbf{Z}}^G = \mathbf{U}^{G^\top} \mathbf{V}^G$, so we have $\|\hat{\mathbf{Z}}^G\| \leq \|\mathbf{U}^G\| \cdot \|\mathbf{V}^G\|$. We then compute

$$\begin{aligned} (\mathbf{U}^{G^\top} \mathbf{U}^G)_{s,s'} &= \prod_{\ell=1}^c (\mathbf{U}^{(i_\ell, j_\ell)^\top} \mathbf{U}^{(i_\ell, j_\ell)})_{s(\kappa_{\mathcal{L}}(i_\ell)), s'(\kappa_{\mathcal{L}}(i_\ell))} \\ &= \prod_{i \in \mathcal{L}} \prod_{j \sim i} (\mathbf{U}^{(i,j)^\top} \mathbf{U}^{(i,j)})_{s(\kappa_{\mathcal{L}}(i)), s'(\kappa_{\mathcal{L}}(i))}. \end{aligned} \quad (10.258)$$

Thus, $\mathbf{U}^{G^\top} \mathbf{U}^G$ is the tensor product over $i \in \mathcal{L}$ of the Hadamard products over $j \sim i$ of $\mathbf{U}^{(i,j)^\top} \mathbf{U}^{(i,j)}$. Since the operator norm is multiplicative over tensor products and submultiplicative over Hadamard products, and every $i \in \mathcal{L}$ has a neighbor $j \in \mathcal{R}$, we find

$$\|\mathbf{U}^{G^\top} \mathbf{U}^G\| \leq \prod_{i \in \mathcal{L}} \prod_{j \sim i} \|\mathbf{U}^{(i,j)^\top} \mathbf{U}^{(i,j)}\|, \quad (10.259)$$

whereby

$$\|\mathbf{U}^G\| \leq \prod_{i \in \mathcal{L}} \prod_{j \sim i} \|\mathbf{U}^{(i,j)}\| = \prod_{\{i,j\} \in E} \|\mathbf{U}^{(i,j)}\|. \quad (10.260)$$

Repeating the same argument for \mathbf{V}^G and multiplying the results together, we have

$$\|\mathbf{Z}^G\| \leq \|\mathbf{U}^G\| \cdot \|\mathbf{V}^G\| \leq \prod_{\{i,j\} \in E} \|\mathbf{U}^{(i,j)}\| \cdot \|\mathbf{V}^{(i,j)}\| = \prod_{\{i,j\} \in E} \|\mathbf{M}^{(i,j)}\|, \quad (10.261)$$

completing the argument for $m = 2$.

For the inductive step, if we have the result for m and are given a decomposition of G into $m + 1$ sets, the result follows by applying the factorization of Proposition 10.13.11 with $A = V_1$, $B = V_2$, and $C = V_3 \sqcup \dots \sqcup V_m$, and using that $\|\mathbf{Z}^{G[A]}\|$ and $\|\mathbf{Z}^{G[B]}\|$ may be bounded using the $m = 2$ case. \square

To see that the connectivity requirements are important for this argument, one may consider the simple case where G has isolated vertices in \mathcal{L} or \mathcal{R} : if so, then the associated CGM is the tensor product of an all-ones matrix with the CGM associated to G with the isolated vertices removed. The norm of this all-ones matrix is polynomial in n , whereby the best bound of the type (10.250) that we could hope for would depend on n , spoiling many of our applications. Other cases where the connectivity requirements fail reduce to a similar situation after sufficiently many applications of the factorization of Proposition 10.13.11.

11 | APPLICATIONS OF LIFTING THEOREMS

Finally, we give applications of our lifting theorems. As we expect from the previous chapter, these results do not achieve the “gold standard” of proving SOS lower bounds of arbitrary constant degree for the SK Hamiltonian, where we wish to lift low-rank matrices. However, we provide some ancillary results suggesting that our construction is correct in other high-rank settings, and take this as additional evidence that it gives a correct first-order approximation to a pseudomoment construction, which must likely be corrected with more detailed considerations to work for the low-rank case.

SUMMARY AND REFERENCES This chapter is based on the parts not discussed earlier of the reference [Kun20b]. The following is a summary of our main results in this chapter.

1. (Theorem 11.2.1) The sum-of-forests pseudomoments built from the Gram matrix of the simplex ETF approximately recover the Grigoriev-Laurent pseudomoments, and give an enumerative combinatorial interpretation of their leading-order behavior.
2. (Theorem 11.3.1) The sum-of-forests pseudomoments give an extension of random high-rank projection matrices to pseudomoments of arbitrary constant degree.
3. (Theorem 11.4.1) A tight degree 6 SOS lower bound for the SK Hamiltonian.

PRIOR WORK The work [MRX20], concurrent with [KB20], also proved a degree 4 lower bound, while the work [GJJ⁺20], concurrent with [Kun20b], also proved a stronger de-

gree $\Omega(n^\epsilon)$ lower bound. Both other results used the pseudocalibration construction of [BHK⁺19]. The master's thesis [dB19] explored techniques for degree 4 lower bounds similar to our own, avoiding pseudocalibration.

11.1 PATTERN DIAGRAMS

We first introduce the following device for counting arguments involving the patterns of equal and unequal entries labelling a CGS diagram in the summations we will encounter.

Definition 11.1.1 (Pattern diagram). *Suppose $F = (V^\bullet \sqcup V^\square, E)$ is a diagram, $s \in [n]^{|V^\bullet|}$, and $\mathbf{a} \in [n]^{V^\square}$. Let the associated pattern diagram, denoted $\text{pat}(F, s, \mathbf{a})$, be the graph G with two types of vertices, \bullet and \square (as for diagrams), formed by starting with G and identifying all v whose value of $f_{s,\mathbf{a}}(v)$ is equal, where if we identify a \bullet vertex with either a \bullet or a \square vertex then the result is a \bullet vertex, but if we identify two \square vertices then the result is again a \square vertex. We then remove all self-loops from $\text{pat}(F, s, \mathbf{a})$ (but allow parallel edges). The graph G is also equipped with a natural labelling inherited from s and \mathbf{a} , which we denote f , sometimes writing $(G, f) = \text{pat}(F, s, \mathbf{a})$.*

Finally, if F_1, \dots, F_m are diagrams, and $s_i \in [n]^{|V^\bullet(F_i)|}$ and $\mathbf{a}_i \in [n]^{V^\square(F_i)}$, then we let $\text{pat}((F_1, s_1, \mathbf{a}_1), \dots, (F_m, s_m, \mathbf{a}_m))$ be the graph formed by applying the above identification procedure to the disjoint union of the F_i , each labelled by s_i and \mathbf{a}_i .

Definition 11.1.2. *Let $\text{Pat}^{\leq 2d}(m)$ be the set of unlabelled $\text{pat}((T_1, s, \mathbf{a}_1), \dots, (T_m, s, \mathbf{a}_m))$ that occur for $T_i \in \mathcal{T}(2d')$, $s \in [n]^{2d'}$, and $\mathbf{a}_i \in [n]^{V^\square(T)}$ for some choice of $1 \leq d' \leq d$. We emphasize that we force s to be the same in all inputs here.*

The way we will use this is by considering the pattern diagram G of any term in any CGS quantity, whose magnitude scales, depending on the behavior of the entries of M , as $n^{-\gamma|E(G)|}$. On the other hand, the number of terms sharing a given pattern diagram is es-

essentially $n^{|V^\square(G)|}$. Grouping terms by pattern diagram allows us to take advantage of the tradeoff between these two quantities.

In particular, we will want to use this to analyze the quantities ϵ_{tree} and ϵ_{err} , so we define the following subsets of pattern diagrams.

Definition 11.1.3. Let $\text{Pat}_{\text{tree}}^{\leq 2d}(m) \subseteq \text{Pat}^{\leq 2d}(m)$ be the set of unlabelled pattern diagrams $\text{pat}((T_1, \mathbf{s}, \mathbf{a}_1), \dots, (T_m, \mathbf{s}, \mathbf{a}_m))$ that occur for $T_i \in \mathcal{T}(2d')$, $\mathbf{s} \in [n]^{2d'}$, and $\mathbf{a}_i \in [n]^{V^\square(T)}$, such that either the entries of \mathbf{s} are not all equal, or $s_1 = \dots = s_{2d'} = j$ but not all of the entries of \mathbf{a}_i equal j for all i , for some $1 \leq d' \leq d$.

Definition 11.1.4. Let $\text{Pat}_{\text{err}}^{\leq 2d}(m) \subseteq \text{Pat}^{\leq 2d}(m)$ be the set of unlabelled pattern diagrams $\text{pat}((T_1, \mathbf{s}, \mathbf{a}_1), \dots, (T_m, \mathbf{s}, \mathbf{a}_m))$ that occur for $T_i \in \mathcal{T}(2d')$, $\mathbf{s} \in [n]^{2d'}$, and $\mathbf{a}_i \in [n]^{V^\square(T)}$, such that \mathbf{a}_i is (T_i, \mathbf{s}) -loose for all i , for some choice of $1 \leq d' \leq d$.

The following two simple facts will be useful throughout; we will introduce other combinatorial properties as needed in our arguments.

Proposition 11.1.5. All diagrams in $\text{Pat}^{\leq 2d}(m)$ are connected for any $d \geq 1$ and $m \geq 1$.

Proposition 11.1.6. $|\text{Pat}^{\leq 2d}(m)| \leq (3md)^{9md}$.

Proof. Every $G \in \text{Pat}^{\leq 2d}(m)$ is connected, and has at most $3md$ vertices and $3md$ edges by Corollary 10.12.10 since this holds for each $T_i \in \mathcal{T}(2d')$ for any $d' \leq d$ and G can have only fewer vertices and edges than the disjoint union of the T_i . Generally, the number of connected graphs on at most $m \geq 2$ vertices, with at most n edges for $n \geq m$, and equipped with a partition of the vertices into two parts is at most $2^m \cdot (m^2)^n \leq m^{3n}$, where we ignore that there may be fewer vertices or edges by allowing “excess” vertices and edges to be added to different connected components that we may ignore. \square

11.2 WARMUP 1: GRIGORIEV-LAURENT PSEUDOMOMENTS

As a first application of Theorem 10.2.3, we show that we can recover a “soft version” of the pseudomoments studied in Chapter 9.

Theorem 11.2.1. *Let $\alpha = \alpha(n) = (\log \log n)^{-50}$. Then, for all n sufficiently large, exists there $\tilde{\mathbb{E}}$ a degree $\frac{1}{100} \log n / \log \log n$ pseudoexpectation satisfying*

$$\tilde{\mathbb{E}} \left[\prod_{i \in S} x_i \right] = \mathbb{1}_{\{|S| \text{ even}\}} \cdot \left(\frac{(-1)^{|S|/2} (|S| - 1)!!}{(n/(1 - \alpha))^{|S|/2}} + O_{|S|} \left(\frac{1}{n^{|S|/2+1}} \right) \right), \quad (11.1)$$

$$\tilde{\mathbb{E}}[\mathbf{x}\mathbf{x}^\top] = \left(1 + \frac{1 - \alpha}{n - 1}\right) \mathbf{I}_n - \frac{1 - \alpha}{n - 1} \mathbf{1}_n \mathbf{1}_n^\top. \quad (11.2)$$

This is weaker than the original statement; most importantly, it only gives a pseudoexpectation $\tilde{\mathbb{E}}$ with $\tilde{\mathbb{E}}[(\mathbf{1}_n^\top \mathbf{x})^2] \approx \alpha n$, and thus does not show that the parity inequality above fails for $\tilde{\mathbb{E}}$. However, it has two important qualitative features: (1) it implies that we need only add to $\tilde{\mathbb{E}}[\mathbf{x}\mathbf{x}^\top]$ an adjustment with operator norm $o(1)$ to obtain an automatically-extensible degree 2 pseudomoment matrix, and (2) it gives the correct leading-order behavior of the pseudomoments. Elaborating on the latter point, our derivation in fact shows how the combinatorial interpretation of $(|S| - 1)!!$ as the number of perfect matchings of a set of $|S|$ objects is related to the appearance of this quantity in the Grigoriev-Laurent construction. While in the original derivation this arises from a somewhat technical induction, in our derivation, this coefficient simply comes from counting the diagrams of $\mathcal{F}(|S|)$ making leading-order contributions, which are the diagrams of perfect matchings.

In the proof we will use the following more detailed bounds on pattern diagrams.

Proposition 11.2.2. *If $G = (V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\text{tree}}^{\leq 2d}(1)$, then $|E| \geq |V^\bullet| + |V^\square| - \mathbb{1}_{\{|V^\bullet| > 1\}}$.*

Proof. We consider two cases. If $|V^\bullet| > 1$, then the result follows since G is connected by Proposition 11.1.5. If $|V^\bullet| = 1$, then since the initial diagram G is formed from by identifying

vertices is a tree, all leaves of that tree are identified in forming G , and G has more than one vertex, G must have a cycle. Therefore, in this case, $|E| \geq |V^\bullet| + |V^\square|$, completing the proof. \square

Proposition 11.2.3. *If $G = (V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\text{err}}^{\leq 2d}(1)$, then $|E| \geq |V^\bullet| + |V^\square|$.*

Proof. Suppose for the sake of contradiction that this is not the case. Since T is connected, by Proposition 11.1.5 G is connected as well, and if $|E| \leq |V^\bullet| + |V^\square| - 1$ then in fact equality holds and G is a tree, in particular having no parallel edges. On the other hand, if \mathbf{a} is (T, \mathbf{s}) -loose, then there exists some index $i \in [n]$ and a \square vertex v in the minimal spanning subtree of $\{w \in V^\bullet : s_{\kappa(w)} = i\}$ such that $a_v \neq i$. Thus there must exist some \square vertex v' in this minimal spanning subtree with at least two neighbors w_1, w_2 such that $f_{s, \mathbf{a}}(w_1) = f_{s, \mathbf{a}}(w_2) = i$ but $a_{v'} \neq i$. The \square vertex of G to which v' is identified will then be incident with a pair of parallel edges, giving a contradiction. \square

Proof of Theorem 11.2.1. We will set

$$\mathbf{M} := \left(1 + \frac{1 - \alpha}{n - 1}\right) \mathbf{I}_n - \frac{1 - \alpha}{n - 1} \mathbf{1}_n \mathbf{1}_n^\top \quad (11.3)$$

and take $\tilde{\mathbb{E}} = \tilde{\mathbb{E}}_{\mathbf{M}}$. We first use Theorem 10.2.3 to show that $\tilde{\mathbb{E}}$ is a degree $2d$ pseudoexpectation.

For the simpler incoherence quantities, we directly bound

$$\epsilon_{\text{offdiag}}(\mathbf{M}) \leq \frac{1}{n - 1}, \quad (11.4)$$

$$\epsilon_{\text{corr}}(\mathbf{M}) \leq \left(2 \left(\frac{1}{n - 1}\right)^2 + (n - 2) \left(\frac{1}{n - 1}\right)^4\right)^{1/2} \leq \frac{2}{n - 1}, \quad (11.5)$$

$$\epsilon_{\text{pow}}(\mathbf{M}) \leq (n - 1) \left(\frac{1}{n - 1}\right)^2 = \frac{1}{n - 1}. \quad (11.6)$$

For ϵ_{tree} , we group terms according to their pattern diagram. A given $G = ((V^\bullet, V^\square), E)$

can occur in at most $n^{|V^\square|}$ terms, and each term contributes at most $(n-1)^{-|E|}$. We then have

$$\begin{aligned}
\epsilon_{\text{tree}}(\mathbf{M}; 2d) &= \max_{0 \leq d' \leq d} \max_{T \in \mathcal{T}(2d')} \max_{s \in [n]^{2d'}} \left| Z^T(\mathbf{M}; s) - \mathbb{1}\{s_1 = \dots = s_n\} \right| \\
&\leq \sum_{G = ((V^\bullet, V^\square), E) \in \text{Pat}_{\text{tree}}^{\leq 2d}(1)} n^{|V^\square|} (n-1)^{-|E|} \\
&\leq \sum_{G = ((V^\bullet, V^\square), E) \in \text{Pat}_{\text{tree}}^{\leq 2d}(1)} n^{|V^\square|} (n-1)^{-|V^\square| - |V^\bullet| + \mathbb{1}\{|V^\bullet| > 1\}} \quad (\text{Proposition 11.2.2}) \\
&\leq \frac{2^{3d}}{n-1} \sum_{G = ((V^\bullet, V^\square), E) \in \text{Pat}_{\text{tree}}^{\leq 2d}(1)} 1 \quad (\text{Corollary 10.12.10}) \\
&\leq \frac{(3d)^{12d}}{n-1}. \quad (\text{Proposition 11.1.6})
\end{aligned}$$

For ϵ_{err} , we follow the same strategy. The main additional observation is that $|\text{set}(s)|$ is always simply the number of \bullet vertices in $\text{pat}(T, s, \mathbf{a})$, regardless of T or \mathbf{a} . Thus we find

$$\begin{aligned}
\epsilon_{\text{err}}(\mathbf{M}; 2d) &= \max_{0 \leq d' \leq d} \max_{T \in \mathcal{T}(2d')} \max_{s \in [n]^{2d'}} n^{|\text{set}(s)|/2} \left| \sum_{\substack{\mathbf{a} \in [n]^{V^\square} \\ \mathbf{a} \text{ } (T, s)\text{-loose}}} \prod_{(v, w) \in E(T)} M_{f_{s, \mathbf{a}}(v), f_{s, \mathbf{a}}(w)} \right| \\
&\leq \sum_{G = (V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\text{err}}^{\leq 2d}(1)} n^{|V^\bullet|/2} \cdot n^{|V^\square|} (n-1)^{-|E|} \\
&\leq \sum_{G = (V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\text{err}}^{\leq 2d}(1)} n^{|V^\bullet|/2 + |V^\square|} (n-1)^{-|V^\bullet| - |V^\square|} \quad (\text{Proposition 11.2.3}) \\
&\leq \frac{2^{6d}}{\sqrt{n}} \sum_{G = (V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\text{err}}^{\leq 2d}(1)} 1 \quad (\text{Corollary 10.12.10}) \\
&\leq \frac{(3d)^{15d}}{\sqrt{n}}. \quad (\text{Proposition 11.1.6})
\end{aligned}$$

(This may be sharpened to $O(n^{-1})$ by considering the case $k = 1$ separately, but that would not change the final result significantly.)

Combining these results, we find $\epsilon(\mathbf{M}; 2d) \leq 2(3d)^{15d} n^{-1/2}$. Thus, since $\lambda_{\min}(\mathbf{M}) \geq \alpha$ and $\|\mathbf{M}\| \leq 2$, the result follows so long as $\alpha \geq 64(12d)^{47} n^{-1/2d}$. If $d = \log n / 100 \log \log n$,

then we have

$$64(12d)^{47}n^{-1/2d} = 64 \left(\frac{12 \log n}{100 \log \log n} \right)^{47} (\log n)^{-100/2} \leq 64 \left(\frac{3}{20 \log \log n} \right)^{50}, \quad (11.7)$$

so with $\alpha = (\log \log n)^{-50}$ it follows that $\tilde{\mathbb{E}}_M$ is a degree $2d$ pseudoexpectation.

It remains to verify (11.1), which gives the leading order behavior of the pseudomoments:

$$\tilde{\mathbb{E}} \left[\prod_{i \in S} x_i \right] = \mathbb{1}_{\{|S| \text{ even}\}} \cdot \left(\frac{(-1)^{|S|/2} (|S| - 1)!!}{(n/(1 - \alpha))^{|S|/2}} + O_{|S|} \left(\frac{1}{n^{|S|/2+1}} \right) \right). \quad (11.8)$$

We claim that the leading order part is exactly the sum of the terms $\mu(F) \cdot Z^F(\mathbf{M}; S)$ for F a forest where every connected component is two \bullet vertices connected by an edge (a “pair”), since there are $(|S| - 1)!!$ such perfect matchings. Thus it suffices to bound the contributions of all other terms.

We use pattern graphs once again, now noting that, since we are assuming S is a *set*, no \bullet vertices will be identified with one another. Suppose F is a good forest and $G = \text{pat}(F, \mathbf{s}, \mathbf{a})$ where all indices of \mathbf{s} are distinct. We then have $|E(G)| \geq 2|V^\square(G)| + \frac{1}{2}|V^\bullet(G)|$, because $|E(G)| = \frac{1}{2} \sum_v \deg(v)$, and every \square vertex in G after the identification procedure will still have degree at least 4 and every \bullet vertex will still have degree at least 1 since no \bullet vertices are identified. Moreover, if $|V^\square(G)| = 0$, then the above inequality is tight if and only if F is a perfect matching to begin with. Therefore, if F is *not* a perfect matching, then, writing for the moment $\text{Pat}_{\mathcal{F}}^{2d}$ for the pattern graphs arising as any $\text{pat}(F, \mathbf{s}, \mathbf{a})$ for $F \in \mathcal{F}(2d)$ (with

forests rather than trees), we find

$$\begin{aligned}
|Z^F(\mathbf{M}; S)| &\leq \sum_{\substack{G=(V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\mathcal{F}}^{|S|} \\ |V^\bullet|=|S|}} n^{|V^\square|} (n-1)^{-|E|} \\
&\leq \sum_{\substack{G=(V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\mathcal{F}}^{|S|} \\ |V^\bullet|=|S|}} n^{|V^\square|} (n-1)^{-2|V^\square| - |V^\bullet|/2 - \mathbb{1}\{|V^\square|=0\}} \\
&\leq n^{-|V^\bullet|/2-1} \sum_{\substack{G=(V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\mathcal{F}}^{|S|} \\ |V^\bullet|=|S|}} 1. \tag{11.9}
\end{aligned}$$

Finally, since the remaining counting coefficient, $\max_{F \in \mathcal{F}(|S|)} |\mu(F)|$, and $|\mathcal{F}(|S|)|$ all depend only on $|S|$, the result follows. (Bounding the remaining combinatorial coefficient and using Propositions 10.12.5 and 10.12.11 here can give a weak quantitative dependence on $|S|$ as well.) \square

11.3 WARMUP 2: LIFTING RANDOM HIGH-RANK PROJECTORS

We also consider a random variant of the setting of Laurent’s theorem, where the special subspace spanned by $\mathbf{1}_n$ is replaced with a random low-dimensional subspace. This is also essentially identical to the setting we would like to treat to give SOS lower bounds for the SK Hamiltonian, except for the dimensionality of the subspace.

Theorem 11.3.1. *Suppose $m : \mathbb{N} \rightarrow \mathbb{N}$ is an increasing function with $\log(n) \ll m(n) \ll n / \log n$ as $n \rightarrow \infty$. Let V be a uniformly random $(n - m)$ -dimensional subspace of \mathbb{R}^n . Then, with high probability as $n \rightarrow \infty$, there exists $\tilde{\mathbb{E}}$ a degree $\frac{1}{300} \log(n/m) / \log \log n$ pseudoexpect-*

tation satisfying

$$\frac{\tilde{\mathbb{E}}[\langle \mathbf{x}, \mathbf{v} \rangle^2]}{\|\mathbf{v}\|^2} \in \left[1, 1 + 4 \frac{m}{n} \right] \quad \text{for all } \mathbf{v} \in V \setminus \{\mathbf{0}\}, \quad (11.10)$$

$$\frac{\tilde{\mathbb{E}}[\langle \mathbf{x}, \mathbf{v} \rangle^2]}{\|\mathbf{v}\|^2} \in \left[0, \frac{1}{(\log \log n)^{32}} + 4 \frac{m}{n} \right] \quad \text{for all } \mathbf{v} \in V^\perp \setminus \{\mathbf{0}\}. \quad (11.11)$$

As in the case of our version of Laurent’s theorem, this result does not imply an SOS integrality gap that is in itself particularly interesting. Indeed, results in discrepancy theory have shown that hypercube vectors can avoid random subspaces of sub-linear dimension (V^\perp , in our case) unusually effectively; see, e.g., [TMR20] for the recent state-of-the-art. Rather, we present this example as another qualitative demonstration of our result, showing that it is possible to treat the random case in the same way as the deterministic case above, and that we can again obtain an automatic higher-degree extension after an adjustment with operator norm $o(1)$ of $\tilde{\mathbb{E}}[\mathbf{x}\mathbf{x}^\top]$ from a random projection matrix.

To handle this random case, we will need some more involved tools that we introduce now. Our main probabilistic tool for controlling the more complicated incoherence quantities will be the following family of *hypercontractive* concentration inequalities, which state (in the case we will use) that low-degree polynomials of independent gaussian random variables concentrate well.

Proposition 11.3.2 (Theorem 5.10 of [Jan97]). *Let $p \in \mathbb{R}[x_1, \dots, x_n]$ be a polynomial with $\deg(p) \leq D$. Then, for all $q \geq 2$,*

$$(\mathbb{E}|p(\mathbf{g})|^q)^{1/q} \leq (q-1)^{D/2} \cdot (\mathbb{E}|p(\mathbf{g})|^2)^{1/2} \quad (11.12)$$

The consequence we will be interested in is the following convenient tail bound, which reduces analyzing the concentration of a polynomial to computing its second moment.

Corollary 11.3.3. *Let $p \in \mathbb{R}[x_1, \dots, x_n]$ be a polynomial with $\deg(p) \leq D$. Then, for all*

$t \geq (2e^2)^{D/2}$,

$$\mathbb{P}\left[\|\mathbf{p}(\mathbf{g})\|_2^2 \geq t \cdot \mathbb{E}\|\mathbf{p}(\mathbf{g})\|_2^2\right] \leq \exp\left(-\frac{D}{e^2}t^{2/D}\right). \quad (11.13)$$

Proof. By Proposition 11.3.2, for any $q \geq 2$,

$$\begin{aligned} \mathbb{P}[|\mathbf{p}(\mathbf{g})| \geq t(\mathbb{E}|\mathbf{p}(\mathbf{g})|^2)^{1/2}] &= \mathbb{P}[|\mathbf{p}(\mathbf{g})|^q \geq t^q(\mathbb{E}|\mathbf{p}(\mathbf{g})|^2)^{q/2}] \\ &\leq t^{-q}(\mathbb{E}|\mathbf{p}(\mathbf{g})|^2)^{-q/2}\mathbb{E}|\mathbf{p}(\mathbf{g})|^q \\ &\leq t^{-q}(q-1)^{qD/2} \\ &\leq (q^{D/2}/t)^q, \end{aligned}$$

and setting $q := t^{2/D}/e^2 \geq 2$ we have

$$= \exp\left(-\frac{D}{e^2}t^{2/D}\right), \quad (11.14)$$

completing the proof. □

We will also use the following standard tail bounds on χ^2 random variables and rectangular Gaussian random matrices.

Proposition 11.3.4 (Lemma 1 of [LM00]). *Let $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$. Then,*

$$\mathbb{P}\left[|\|\mathbf{g}\|_2 - \sqrt{n}| \geq t\right] \leq 2 \exp\left(-\frac{t^2}{2}\right). \quad (11.15)$$

Proposition 11.3.5 (Corollary 7.3.3 and Exercise 7.3.4 of [Ver18]). *Let $\mathbf{G} \in \mathbb{R}^{m \times n}$ with $m \geq n$ have i.i.d. entries distributed as $\mathcal{N}(0, 1)$. Let $\sigma_1(\mathbf{G}) \geq \dots \geq \sigma_n(\mathbf{G}) \geq 0$ denote the ordered singular values of \mathbf{G} . Then,*

$$\mathbb{P}\left[\sqrt{m} - \sqrt{n} - t \leq \sigma_n(\mathbf{G}) \leq \sigma_1(\mathbf{G}) \leq \sqrt{m} + \sqrt{n} + t\right] \geq 1 - 4 \exp(-Ct^2) \quad (11.16)$$

for a universal constant $C > 0$.

We will also need some more specific combinatorial preliminaries, which describe how to compute expectations of gaussian polynomials like those that will wind up associated with pattern graphs in our calculations.

Definition 11.3.6. A cycle cover of a graph G is a partition of the edges into edge-disjoint cycles. We denote the number of cycles in the largest cycle cover of G by $c_{\max}(G)$.

Proposition 11.3.7. Let $G = (V, E)$ be a graph, and for each $v \in V$ draw $g_v \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$ independently. Then,

$$\mathbb{E} \left[\prod_{(v,w) \in E} \langle g_v, g_w \rangle \right] = \sum_{C \text{ cycle cover of } G} k^{|C|} \leq |E|^{|E|} k^{c_{\max}(G)}. \quad (11.17)$$

Proof. The first equality is proved in Section 4 of [MR11]. The inequality follows from the fact that the number of cycle covers of G is at most the number of partitions of the edges of G , which is at most $|E|^{|E|}$ by Proposition 10.12.7. \square

As a historical remark, we note that essentially the same results, though in somewhat different language, are given in the earlier paper [LR01].

Proposition 11.3.8. Suppose $G = (V, E)$ is a connected graph with no self-loops, but possibly with parallel edges. Then, $|V| + c_{\max}(G) - 1 \leq |E|$, with equality if and only if G is an inflated tree—a tree where every edge has been replaced with a cycle.

Proof. Let C be a maximum cycle cover of G . Let G' be the graph formed by removing an arbitrary edge from every cycle in C . Then, $|E(G')| = |E| - c_{\max}(G)$ and $|V(G')| = |V|$. Moreover, G' is connected, since there is a path in G' between the endpoints of each edge that was removed (along the remaining edges of the corresponding cycle). Thus, $|E(G')| \geq |V(G')| - 1$, and substituting gives $|E| - c_{\max}(G) \geq |V| - 1$.

Equality holds if and only if G' is a tree. If G is an inflated tree, this will clearly be the case. Suppose now that G' is a tree; we want to show that G is an inflated tree. If two edge-disjoint cycles intersect in more than one vertex, then after one edge is removed from each cycle, there still exists a cycle among their edges. Therefore, if G' is a tree, then any two cycles of C can intersect in at most one vertex. Moreover, again because G' is a tree, there can exist no further cycle of G including edges from more than one of the cycles of C . Therefore, the graph formed by collapsing each cycle of C to an edge must be a tree, whereby G is an inflated tree. \square

Proof of Theorem 11.3.1. Let $\mathbf{g}_1, \dots, \mathbf{g}_m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ be a collection of independent gaussian vectors coupled to V^\perp such that $V^\perp = \text{span}(\mathbf{g}_1, \dots, \mathbf{g}_m)$. Let us define

$$\alpha := (\log \log n)^{-32}, \quad (11.18)$$

which will play a similar role here to that of α in the proof of Theorem 11.2.1. Define $\mathbf{M}^{(0)} := (1 - \alpha/2) \frac{1}{n} \sum_{i=1}^m \mathbf{g}_i \mathbf{g}_i^\top$, let \mathbf{D} be the diagonal matrix with $\text{diag}(\mathbf{D}) = \text{diag}(\mathbf{M}^{(0)})$, and define $\mathbf{M} := \mathbf{I} + \mathbf{D} - \mathbf{M}^{(0)}$. We will then take $\tilde{\mathbb{E}} = \tilde{\mathbb{E}}_M$ for this choice of \mathbf{M} (which we note satisfies $\text{diag}(\mathbf{M}) = \mathbf{1}_n$ by construction).

We first establish a preliminary asymptotic on the eigenvalues of $\mathbf{M}^{(0)}$. Let $\lambda_1(\mathbf{M}^{(0)}) \geq \lambda_2(\mathbf{M}^{(0)}) \geq \dots \geq \lambda_n(\mathbf{M}^{(0)}) \geq 0$ be the ordered eigenvalues of $\mathbf{M}^{(0)}$. Then, $\lambda_{m+1}(\mathbf{M}^{(0)}) = \dots = \lambda_n(\mathbf{M}^{(0)}) = 0$ almost surely. We note that $\sqrt{m/n} \ll 1/\sqrt{\log n} \ll \alpha$, whereby the concentration inequality of Proposition 11.3.5 implies that, with high probability as $n \rightarrow \infty$,

$$1 - \alpha \leq \lambda_m(\mathbf{M}^{(0)}) \leq \dots \leq \lambda_1(\mathbf{M}^{(0)}) \leq 1 - \frac{1}{3}\alpha. \quad (11.19)$$

We next control the entries of \mathbf{D} . These are

$$D_{ii} = \left(1 - \frac{\alpha}{2}\right) \frac{1}{n} \sum_{j=1}^m (\mathbf{g}_j)_i^2, \quad (11.20)$$

where the law of the inner sum is $\chi^2(m)$. By the concentration inequality of Proposition 11.3.4, we have that $\mathbb{P}[D_{ii} \geq 4\frac{m}{n}] \leq \exp(-n)$, and since $m \gg \log n$ by assumption, upon taking a union bound we have that, with high probability, $\mathbf{0} \leq \mathbf{D} \leq 4\frac{m}{n}\mathbf{I}_n$.

We next establish the projection-like behavior of \mathbf{M} . Suppose first that $\mathbf{v} \in V$. Since the row space of $\mathbf{M}^{(0)}$ is V^\perp , we have, on the event that the bound for \mathbf{D} above holds,

$$\|\mathbf{v}\|^2 \leq \mathbf{v}^\top \mathbf{M} \mathbf{v} = \|\mathbf{v}\|^2 + \mathbf{v}^\top \mathbf{D} \mathbf{v} \leq \left(1 + 4\frac{m}{n}\right) \|\mathbf{v}\|^2 \quad (11.21)$$

Now, suppose $\mathbf{v} \in V^\perp$. Then, on the event that the bound for $\mathbf{M}^{(0)}$ above holds, we have that since \mathbf{v} is in the subspace spanned by the top m eigenvectors of $\mathbf{M}^{(0)}$,

$$\mathbf{v}^\top \mathbf{M} \mathbf{v} = \|\mathbf{v}\|^2 + \mathbf{v}^\top \mathbf{D} \mathbf{v} - \mathbf{v}^\top \mathbf{M}^{(0)} \mathbf{v} \leq \|\mathbf{v}\|^2 + 4\frac{m}{n} \|\mathbf{v}\|^2 - (1 - \alpha) \|\mathbf{v}\|^2 = \left(4\frac{m}{n} + \alpha\right) \|\mathbf{v}\|^2. \quad (11.22)$$

We now take up the main task of showing that $\tilde{\mathbb{E}}_M$ is a pseudoexpectation of the required degree. Note that the above results imply that $\lambda_{\min}(\mathbf{M}) \geq 1 - \lambda_{\max}(\mathbf{M}^{(0)}) \geq \alpha/3$, giving the necessary control of the smallest eigenvalue. It remains to control the incoherence quantities.

Writing $\mathbf{h}_1, \dots, \mathbf{h}_n \in \mathbb{R}^m$ for the vectors $\mathbf{h}_i = ((\mathbf{g}_j)_i)_{j=1}^n$, we note that, for $i \neq j$, we have $M_{ij} = -(1 - \alpha)\frac{1}{n} \langle \mathbf{h}_i, \mathbf{h}_j \rangle$, and the \mathbf{h}_i are independent and identically distributed with law $\mathcal{N}(\mathbf{0}, \mathbf{I}_m)$. Since for any fixed $i \neq j$ the law of $\langle \mathbf{h}_i, \mathbf{h}_j \rangle$ is the same as that of $\|\mathbf{h}_i\|_2 \langle \mathbf{h}_j \rangle_1$ (by

orthogonal invariance of gaussian vectors), using Proposition 11.3.4 again we may bound

$$\begin{aligned} \mathbb{P} \left[\frac{1}{n} |\langle \mathbf{h}_i, \mathbf{h}_j \rangle| \geq t \right] &\leq \mathbb{P}[\|\mathbf{h}_i\|_2 \geq 2\sqrt{m}] + \mathbb{P} \left[|(\mathbf{h}_j)_1| \geq \frac{tn}{2\sqrt{m}} \right] \\ &\leq \exp \left(-\frac{m}{2} \right) + \exp \left(-\frac{t^2 n^2}{8m} \right) \end{aligned} \quad (11.23)$$

Recall that we have assumed $m \gg \log n$. Therefore, taking a union bound over these events for $\{i, j\} \in \binom{[n]}{2}$ we find that, with high probability as $n \rightarrow \infty$, the simpler incoherence quantities will satisfy

$$\epsilon_{\text{offdiag}}(\mathbf{M}) \leq 5\sqrt{\frac{m \log n}{n^2}}, \quad (11.24)$$

$$\epsilon_{\text{corr}}(\mathbf{M}) \leq 25 \left(2\frac{m \log n}{n^2} + (n-2)\frac{m^2 \log^2 n}{n^4} \right)^{1/2} \leq 50\sqrt{\frac{m \log n}{n^2}}. \quad (11.25)$$

For ϵ_{pow} , we observe that by the above reasoning with high probability $\mathbf{M} \geq \mathbf{0}$, and thus $|M_{ij}| \leq 1$ for all $i \neq j$. On this event, we have

$$\epsilon_{\text{pow}}(\mathbf{M}) \leq \max_{i \in [n]} \sum_{j \neq i} M_{ij}^2 \leq \frac{1}{n^2} \max_{i \in [n]} \mathbf{h}_i^\top \left(\sum_{j \neq i} \mathbf{h}_j \mathbf{h}_j^\top \right) \mathbf{h}_i \leq \frac{1}{n^2} \left\| \sum_{i=1}^n \mathbf{h}_i \mathbf{h}_i^\top \right\| \max_{i \in [n]} \|\mathbf{h}_i\|_2^2. \quad (11.26)$$

By the calculations above, with high probability we have both $\|\mathbf{h}_i\|_2^2 \leq 4m$ for all $i \in [n]$ and $\left\| \sum_{i=1}^n \mathbf{h}_i \mathbf{h}_i^\top \right\| = \left\| \sum_{i=1}^m \mathbf{g}_i \mathbf{g}_i^\top \right\| = \frac{n}{1-\alpha} \|\mathbf{M}^{(0)}\| \leq 2n$. Thus we find that, with high probability,

$$\epsilon_{\text{pow}}(\mathbf{M}) \leq 8\frac{m}{n}. \quad (11.27)$$

Finally, for ϵ_{tree} and ϵ_{err} we will use pattern diagrams together with hypercontractivity. We begin with ϵ_{tree} . Examining one term in the maximization, for a given $T \in \mathcal{T}(2d')$ and

$\mathbf{s} \in [n]^{2d'}$, we have

$$\begin{aligned}
Z^T(\mathbf{M}; \mathbf{s}) - \mathbb{1}\{s_1 = \dots = s_n\} \\
&= \sum_{\mathbf{a} \in [n]^{V^\square} \text{ with } s, \mathbf{a} \text{ not all equal}} \prod_{(v,w) \in E(T)} M_{f_{s,\mathbf{a}}(v), f_{s,\mathbf{a}}(w)} \\
&= \sum_{\substack{\mathbf{a} \in [n]^{V^\square} \text{ with } s, \mathbf{a} \text{ not all equal} \\ (G,f) = \text{pat}(T, s, \mathbf{a})}} \prod_{(v,w) \in E(G)} M_{f(v), f(w)}
\end{aligned}$$

and since the pattern diagram is constructed to have all edges between vertices with different indices, we may expand this in terms of the \mathbf{h}_i ,

$$= \sum_{\substack{\mathbf{a} \in [n]^{V^\square} \text{ with } s, \mathbf{a} \text{ not all equal} \\ (G,f) = \text{pat}(T, s, \mathbf{a})}} \left(-\frac{1-\alpha}{n} \right)^{|E(G)|} \prod_{(v,w) \in E(G)} \langle \mathbf{h}_{f(v)}, \mathbf{h}_{f(w)} \rangle. \quad (11.28)$$

Towards applying the hypercontractive inequality, we compute the second moment:

$$\begin{aligned}
&\mathbb{E}[(Z^T(\mathbf{M}; \mathbf{s}) - \mathbb{1}\{s_1 = \dots = s_n\})^2] \\
&= \sum_{\substack{\mathbf{a}_1, \mathbf{a}_2 \in [n]^{V^\square} \text{ with } s, \mathbf{a}_1 \text{ not all equal} \\ \text{and } s, \mathbf{a}_2 \text{ not all equal} \\ (G,f) = \text{pat}((T, s, \mathbf{a}_1), (T, s, \mathbf{a}_2))}} \left(-\frac{1-\alpha}{n} \right)^{|E(G)|} \mathbb{E} \left[\prod_{(v,w) \in E(G)} \langle \mathbf{h}_{f(v)}, \mathbf{h}_{f(w)} \rangle \right]
\end{aligned}$$

and simplifying the remaining expectation using Proposition 11.3.7 and bounding the first term,

$$\leq (6d)^{6d} \sum_{\substack{\mathbf{a}_1, \mathbf{a}_2 \in [n]^{V^\square} \text{ with } s, \mathbf{a}_1 \text{ not all equal} \\ \text{and } s, \mathbf{a}_2 \text{ not all equal} \\ G = \text{pat}((T, s, \mathbf{a}_1), (T, s, \mathbf{a}_2))}} \mathbb{1}\{G \text{ has a cycle cover}\} n^{-|E(G)|} m^{c_{\max}(G)},$$

where we note that the expression does not depend on the labelling f of the vertices of G anymore. Now, as before, we group terms according to the graph G , using that each occurs at most $n^{|V^\square(G)|}$ times in the sum, and that each G arising is connected, contains at most $6d$ vertices and $6d$ edges, and at least one \bullet vertex:

$$\begin{aligned} &\leq (6d)^{6d} \sum_{G=(V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\text{tree}}^{\leq 2d}(2)} \mathbb{1}\{G \text{ has a cycle cover}\} n^{|V^\square| - |E|} m^{c_{\max}(G)} \\ &\leq (6d)^{6d} \frac{m}{n} \sum_{G=(V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\text{tree}}^{\leq 2d}(2)} n^{|V^\square| + c_{\max}(G) - |E|}. \end{aligned}$$

Now, by Proposition 11.3.8 we have $|V^\bullet| + |V^\square| + c_{\max}(G) - |E| - 1 \leq 0$. If $|V^\bullet| \geq 2$, then this yields $|V^\square| + c_{\max}(G) - |E| \leq -1$. If $|V^\bullet| = 1$, we argue slightly more carefully and note that in this case, since all \bullet vertices in both underlying trees collapsed to a single vertex, in fact all vertices in G have degree at least 4, so G cannot be an inflated tree as in the only case of equality for Proposition 11.3.8. Therefore, in this case we have $|V^\square| + c_{\max}(G) - |E| \leq -1$ again, whereby

$$\leq (6d)^{6d} \frac{m}{n^2} \sum_{G=(V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\text{tree}}^{\leq 2d}(2)} 1$$

and concluding with Proposition 11.1.6, we find

$$\begin{aligned} &\leq (6d)^{24d} \frac{m}{n^2}. \\ &\leq \frac{(6d)^{24d}}{n}. \end{aligned} \tag{11.29}$$

(Here we have been slightly more precise than strictly necessary, in anticipation of referring to our results when discussing the SK Hamiltonian below.)

Now, we observe that $Z^T(\mathbf{M}; \mathbf{s}) - \mathbb{1}\{s_1 = \dots = s_{2d'}\}$ is a polynomial of degree at most $2|E(T)| \leq 6d$ (by Corollary 10.12.10) in the entries of the \mathbf{h}_i , which are i.i.d. standard gaus-

sians. Thus we can apply the hypercontractive tail bound of Corollary 11.3.3 to find, taking $t = n^{1/4} \geq (2e^2)^{3d}$ for n sufficiently large,

$$\mathbb{P} \left[|Z^T(\mathbf{M}; \mathbf{s}) - \mathbb{1}\{s_1 = \dots = s_{2d'}\}| \geq (6d)^{12d} n^{-1/4} \right] \leq \exp(-6dn^{1/12d}). \quad (11.30)$$

Taking a union bound, since the number of choices of d' , T , and \mathbf{s} is at most $d \cdot 2(3d)^{3d} \cdot n^{2d} \leq n^{6d}$ for n sufficiently large, we have

$$\begin{aligned} \mathbb{P} \left[\epsilon_{\text{tree}}(\mathbf{M}; 2d) \geq (6d)^{12d} n^{-1/4} \right] &\leq n^{6d} \exp(-6dn^{1/12d}) \\ &\leq \exp \left(6d(\log n - n^{1/12d}) \right), \end{aligned} \quad (11.31)$$

and recalling that $d \leq \frac{1}{300} \log(n/m) / \log \log n$ from our assumption we find that the event above holds with high probability. Also, from this same assumption we find $(6d)^{12d} \leq n^{1/8}$ for n sufficiently large, whereby with high probability

$$\epsilon_{\text{tree}}(\mathbf{M}; 2d) \leq n^{-1/8}. \quad (11.32)$$

We now perform the same analysis for $\epsilon_{\text{err}}(\mathbf{M}; 2d)$. Again examining one term with a given $T \in \mathcal{T}(2d')$ and $\mathbf{s} \in [n]^{2d'}$, manipulating as before, and computing the second

moment, we find

$$\begin{aligned}
& \mathbb{E} \left[\left(\sum_{\substack{\mathbf{a} \in [n]^{V^\square} \\ \mathbf{a} \text{ } (T, \mathbf{s})\text{-loose}}} \prod_{(v, w) \in E(T)} M_{f_{\mathbf{s}, \mathbf{a}}(v), f_{\mathbf{s}, \mathbf{a}}(w)} \right)^2 \right] \\
&= \sum_{\substack{\mathbf{a}_1, \mathbf{a}_2 \in [n]^{V^\square} \\ \mathbf{a}_1, \mathbf{a}_2 \text{ } (T, \mathbf{s})\text{-loose} \\ (G, f) = \text{pat}((T, \mathbf{s}, \mathbf{a}_1), (T, \mathbf{s}, \mathbf{a}_2))}} \left(-\frac{1 - \alpha}{n} \right)^{|E(G)|} \mathbb{E} \left[\prod_{(v, w) \in E(G)} \langle \mathbf{h}_{f(v)}, \mathbf{h}_{f(w)} \rangle \right] \\
&\leq (6d)^{6d} \sum_{\substack{\mathbf{a}_1, \mathbf{a}_2 \in [n]^{V^\square} \\ \mathbf{a}_1, \mathbf{a}_2 \text{ } (T, \mathbf{s})\text{-loose} \\ G = \text{pat}((T, \mathbf{s}, \mathbf{a}_1), (T, \mathbf{s}, \mathbf{a}_2))}} \mathbb{1}\{G \text{ has a cycle cover}\} n^{-|E(G)|} m^{c_{\max}(G)} \\
&\leq (6d)^{6d} \sum_{G = (V^\bullet \sqcup V^\square, E) \in \text{Pat}_{\text{err}}^{\leq 2d}(2)} \mathbb{1}\{G \text{ has a cycle cover}\} n^{|V^\square| - |E|} m^{c_{\max}(G)}
\end{aligned}$$

We recall that $|V^\bullet| = |\text{set}(\mathbf{s})|$ and $|V^\bullet| + |V^\square| = |V|$, so we may rewrite this by “forgetting” the vertex types as

$$\begin{aligned}
&= n^{-|\text{set}(\mathbf{s})|} (6d)^{6d} \sum_{G = (V, E) \in \text{Pat}_{\text{err}}^{\leq 2d}(2)} \mathbb{1}\{G \text{ has a cycle cover}\} n^{|V| - |E|} m^{c_{\max}(G)} \\
&\leq n^{-|\text{set}(\mathbf{s})|} (6d)^{6d} \frac{m}{n} \sum_{G = (V, E) \in \text{Pat}_{\text{err}}^{\leq 2d}(2)} n^{|V| + c_{\max}(G) - |E|}. \tag{11.33}
\end{aligned}$$

Now, by Proposition 11.3.8, the inner term is at most 1 unless G is an inflated tree. We claim that, when $G = \text{pat}((T, \mathbf{s}, \mathbf{a}_1), (T, \mathbf{s}, \mathbf{a}_2))$ where \mathbf{a}_1 and \mathbf{a}_2 are both (T, \mathbf{s}) -loose, then G cannot be an inflated tree. To prove this, we consider two cases.

Case 1: $|\text{set}(\mathbf{s})| = 1$. In this case, as we have argued above, since the total number of vertices in the two initial trees taken together is at least 4 and neither of these trees is a pair, every vertex in G will have degree at least 4, whereby G cannot be an inflated tree.

Case 2: $|\text{set}(\mathbf{s})| > 1$. Suppose, more specifically, that $G = \text{pat}((T_1, \mathbf{s}, \mathbf{a}_1), (T_2, \mathbf{s}, \mathbf{a}_2))$. Since \mathbf{a}_1 is (T_1, \mathbf{s}) -loose, there exists some $i \in [n]$ and some $v \in V^\square(T_1)$ such that v belongs to the minimal spanning tree of leaves ℓ with $s_{\kappa_{T_1}(\ell)} = i$, but $(\mathbf{a}_1)_v \neq i$. In particular, there

must exist two such leaves ℓ_1, ℓ_2 such that v is along the path from ℓ_1 to ℓ_2 . In G , the vertex that v is identified to—call it x —is different from the vertex that ℓ_1 and ℓ_2 are identified to—call it y . Since $|\text{set}(s)| > 1$, there is some $j \neq i$ and a leaf ℓ' with $s_{\kappa_{T_1}(\ell')} = j$. Suppose ℓ' is identified to a vertex z in G . Then, there is a path from x to z in G , so there are two different paths from y to z consisting only of edges coming from T_1 .

On the other hand, since T_2 is a tree and a_2 is (T_2, s) -loose, there is another path in G from y to z and consisting of edges different from the first two paths, coming from T_2 . Therefore, in G there exist three different paths between y and z ; put differently, a triple edge can be obtained as a minor of G (after discarding self-loops). On the other hand, any minor of an inflated tree is still an inflated tree (again after discarding self-loops), and a triple edge is not an inflated tree. Thus, G cannot be an inflated tree. This concludes the proof of our intermediate claim.

We then conclude the main argument using Proposition 11.1.6:

$$\begin{aligned} \mathbb{E} \left[\left(\sum_{\substack{a \in [n]^{V \square} \\ a \text{ } (T,s)\text{-loose}}} \prod_{(v,w) \in E(T)} M_{f_{s,a}(v), f_{s,a}(w)} \right)^2 \right] &\leq n^{-|\text{set}(s)|} (6d)^{6d} \frac{m}{n} \sum_{G=(V,E) \in \text{Pat}_{\text{err}}^{\leq 2d}(2)} 1 \\ &\leq n^{-|\text{set}(s)|} (6d)^{24d} \frac{m}{n}. \end{aligned} \quad (11.34)$$

Similarly to before, we apply Corollary 11.3.3 with $t = (n/m)^{1/4}$, finding

$$\begin{aligned} \mathbb{P} \left[n^{|\text{set}(s)|/2} \left| \sum_{\substack{a \in [n]^{V \square} \\ a \text{ } (T,s)\text{-loose}}} \prod_{(v,w) \in E(T)} M_{f_{s,a}(v), f_{s,a}(w)} \right| \geq (6d)^{12d} (m/n)^{1/4} \right] \\ \leq \exp(-6d(n/m)^{1/12d}), \end{aligned} \quad (11.35)$$

and performing the same union bound calculation over all choices of d' , T , and s shows that, with high probability,

$$\epsilon_{\text{err}}(\mathbf{M}; 2d) \leq (m/n)^{1/8}. \quad (11.36)$$

Thus, combining the results on the incoherence quantities, with high probability we have

$$\epsilon(\mathbf{M}; 2d) \leq 55\sqrt{\frac{m \log n}{n^2}} + 8\frac{m}{n} + \left(\frac{1}{n}\right)^{1/8} + \left(\frac{m}{n}\right)^{1/8} \leq 65\left(\frac{m}{n}\right)^{1/8}. \quad (11.37)$$

On this event, we work with the condition of Theorem 10.2.3, for n sufficiently large:

$$(12d)^{32} \|\mathbf{M}\|^5 \epsilon(\mathbf{M}; 2d)^{1/d} \leq 64 \left(\frac{12d}{\log n}\right)^{32} \leq \frac{1}{3} (\log \log n)^{-32} = \frac{1}{3} \alpha \leq \lambda_{\min}(\mathbf{M}), \quad (11.38)$$

concluding the proof. □

11.4 DEGREE 6 LOWER BOUND FOR SHERRINGTON-KIRKPATRICK HAMILTONIAN

We now prove our SOS lower bound for the SK Hamiltonian.

Theorem 11.4.1. *For any $\epsilon > 0$, for $\mathbf{W} \sim \text{GOE}(n)$, $\lim_{n \rightarrow \infty} \mathbb{P}[\text{SOS}_6(\mathbf{W}) \geq (2 - \epsilon)n] = 1$.*

The new technical lemma we need here compared to previously controls the ϵ_{pow} incoherence quantity from Definition 10.2.2 in the previous chapter. Recall that, in this setting, this amounts to controlling the spectrum of $\mathbf{M}^{\circ 2}$ for \mathbf{M} a rescaled low-rank projection matrix.

Lemma 11.4.2. *Let $\delta \in (0, 1)$ and $r = \delta n$. Then, for any $K > 0$ there exist constants $C_1, C_2 > 0$ depending only on K and δ such that, letting $\mathbf{h}_1, \dots, \mathbf{h}_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$ be independent, $\mathbf{a}_i := \text{isovec}(\mathbf{h}_i \mathbf{h}_i^\top - \mathbf{I}_r)$, and $\mathbf{A} \in \mathbb{R}^{r(r+1)/2 \times n}$ have the \mathbf{a}_i as its columns,*

$$\mathbb{P}\left[\left\|\frac{1}{r^2} \mathbf{A}^\top \mathbf{A} - \mathbf{I}_n\right\| \leq C_1 \frac{\log n}{\sqrt{n}}\right] \geq 1 - \frac{C_2}{n^K}. \quad (11.39)$$

In essence, this says that the $\frac{1}{r} \mathbf{a}_i$ are approximately orthonormal. We give the proof in Section 11.5 below using powerful general results of [ALPTJ11], as well as a related result used

in [KB20] that gives a more direct proof for the case of M a rescaling of a genuine random projection matrix—not just the Gaussian approximation thereof—showing how such results can be proved even without independence assumptions.¹

Proof of Theorem 11.4.1. We start out following similar steps to Theorem 11.3.1. Fix some small $\delta > 0$, and let V be the eigenspace spanned by the δn leading eigenvectors of W . Let us define $r := \delta n$, following the notation from earlier. Let $\mathbf{g}_1, \dots, \mathbf{g}_r \in \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ be a collection of independent gaussian vectors coupled to V such that $V = \text{span}(\mathbf{g}_1, \dots, \mathbf{g}_r)$. Define $M^{(0)} := (1 - \alpha/2) \frac{\delta^{-1}}{n} \sum_{i=1}^r \mathbf{g}_i \mathbf{g}_i^\top$. Let D be the diagonal matrix with $\text{diag}(D) = \text{diag}(M^{(0)})$, and define $M := I - D + M^{(0)}$.

We first control the entries of D . These are

$$D_{ii} = \left(1 - \frac{\alpha}{2}\right) \frac{\delta^{-1}}{n} \sum_{j=1}^r (\mathbf{g}_j)_i^2 = \left(1 - \frac{\alpha}{2}\right) \frac{1}{r} \sum_{j=1}^r (\mathbf{g}_j)_i^2. \quad (11.40)$$

Applying the concentration inequality of Proposition 11.3.4 and a union bound, we find that with high probability $(1 - \alpha)\mathbf{I}_n \leq D \leq (1 - \alpha/3)\mathbf{I}_n$. Since $M^{(0)} \geq \mathbf{0}$, on this event we have $M \geq (\alpha/3)\mathbf{I}_n$, and also $\|M - M^{(0)}\| = \|\mathbf{I}_n - D\| \leq \alpha$.

We now show that M satisfies the conditions of Theorem 10.11.2. We note that, again writing $\mathbf{h}_1, \dots, \mathbf{h}_n \in \mathbb{R}^r$ for the vectors $\mathbf{h}_i = ((\mathbf{g}_j)_i)_{j=1}^r$, for $i \neq j$ we have $M_{ij} = (1 - \alpha/2) \frac{1}{r} \langle \mathbf{h}_i, \mathbf{h}_j \rangle$. Following the same calculations as in Theorem 11.3.1 (noting that we may follow them truly verbatim, since the setting is identical except for the constant in front of each M_{ij} with $i \neq j$), we find that we have, with high probability

$$\epsilon_{\text{offdiag}}(M) \leq K \sqrt{\frac{\log n}{n}}, \quad (11.41)$$

$$\epsilon_{\text{err}}(M; 6) \leq K. \quad (11.42)$$

¹I thank Ramon van Handel for suggesting that the line of work in [ALPTJ11] would be applicable here.

Here and in the rest of this proof, we adopt the convention that $K = K(\delta) > 0$ is a constant that may change from line to line.

We now control $\epsilon_{\text{pow}}(\mathbf{M})$. For $k \geq 3$, by the Gershgorin circle theorem and substituting in $\epsilon_{\text{offdiag}}(\mathbf{M})$, we note that we have

$$\|\mathbf{M}^{\circ k} - \mathbf{I}_n\| \leq n \left(K \sqrt{\frac{\log n}{n}} \right)^3 \leq K^3 \frac{\log^2 n}{\sqrt{n}}. \quad (11.43)$$

We will take $t_{\text{pow}} = (1 - \alpha/2)^2 \frac{1}{r}$. Let us define vectors $\mathbf{a}_i := \text{isovec}(\mathbf{h}_i \mathbf{h}_i^\top - \mathbf{I}_r)$, where the mapping $\text{isovec} : \mathbb{R}_{\text{sym}}^{r \times r} \rightarrow \mathbb{R}^{r(r+1)/2}$ is such that $\langle \mathbf{S}, \mathbf{T} \rangle = \langle \text{isovec}(\mathbf{S}), \text{isovec}(\mathbf{T}) \rangle$. We note that

$$\langle \mathbf{a}_i, \mathbf{a}_j \rangle = \langle \mathbf{h}_i, \mathbf{h}_j \rangle^2 - \|\mathbf{h}_i\|^2 - \|\mathbf{h}_j\|^2 + r. \quad (11.44)$$

Thus, writing \mathbf{A} for the matrix with the \mathbf{a}_i as its columns and \mathbf{b} for the vector with entries $\|\mathbf{h}_i\|^2$, we have

$$\mathbf{M}^{\circ 2} - \mathbf{I}_n = \frac{(1 - \alpha/2)^2}{r^2} \left(\mathbf{A}^\top \mathbf{A} + \mathbf{b} \mathbf{1}_n^\top + \mathbf{1}_n \mathbf{b}^\top - r \mathbf{1}_n \mathbf{1}_n^\top - \text{diag}(\mathbf{b})^2 \right) \quad (11.45)$$

$$= \frac{(1 - \alpha/2)^2}{r^2} \left(\mathbf{A}^\top \mathbf{A} + (\mathbf{b} - r \mathbf{1}_n) \mathbf{1}_n^\top + \mathbf{1}_n (\mathbf{b} - r \mathbf{1}_n)^\top + r \mathbf{1}_n \mathbf{1}_n^\top - \text{diag}(\mathbf{b})^2 \right), \quad (11.46)$$

whereby with our choice of t_{pow} we may bound

$$\begin{aligned} & \|\mathbf{M}^{\circ 2} - \mathbf{I}_n - t_{\text{pow}} \mathbf{1}_n \mathbf{1}_n^\top\| \\ & \leq \frac{(1 - \alpha/2)^2}{r^2} \left(\|\mathbf{A}^\top \mathbf{A} - r^2 \mathbf{I}_n\| + 2n^{1/2} \|\mathbf{b} - r \mathbf{1}_n\|_2 + \|r^2 \mathbf{I}_n - \text{diag}(\mathbf{b})^2\| \right) \end{aligned}$$

By Proposition 11.3.4, with high probability for all $i \in [n]$ we have $|\|\mathbf{h}\|_i^2 - r| \leq r^{3/4}$, on which event we have $\|\mathbf{b} - r\mathbf{1}_n\|_2 \leq r^{3/4}n^{1/2}$ and $\|r^2\mathbf{I}_n - \text{diag}(\mathbf{b})^2\| \leq 3r^{7/4}$, which leaves only the more interesting term,

$$\leq \left\| \frac{1}{r^2} \mathbf{A}^\top \mathbf{A} - \mathbf{I}_n \right\|^2 + (3 + \delta^{-1})r^{-1/4} \quad (11.47)$$

Finally, by Lemma 11.4.2 the remaining term is at most $K\sqrt{\log^2 n/n}$ with high probability, whereby we find with high probability

$$\leq Kn^{-1/4}. \quad (11.48)$$

Combining these results, we see that $\tilde{\epsilon}(\mathbf{M}, t_{\text{pow}}) \leq Kn^{-1/4}$. Since by our earlier calculations $\lambda_{\min}(\mathbf{M}) \geq \alpha/3$, the condition of Theorem 10.11.2 will hold with high probability.

Now, we consider the constant c appearing in the theorem. Recall that we chose $t_{\text{pow}} \leq Kn^{-1}$. We have $\|\mathbf{M}\| \leq 1 + \|\mathbf{M}^{(0)}\| \leq K$ with high probability by Proposition 11.3.5, and $\|\mathbf{M}^2\|_F \leq \sqrt{r\|\mathbf{M}^2\|^2} \leq r^{1/2}\|\mathbf{M}\|^2 \leq Kn^{1/2}$. The only quantity it remains to control is $\epsilon_{\text{offdiag}}(\mathbf{M}^2)$. We have

$$\mathbf{M}^2 = (\mathbf{I} - \mathbf{D})^2 + (\mathbf{I} - \mathbf{D})\mathbf{M}^{(0)} + \mathbf{M}^{(0)}(\mathbf{I} - \mathbf{D}) + \mathbf{M}^{(0)2}, \quad (11.49)$$

and thus, for $i \neq j$,

$$\begin{aligned} |(\mathbf{M}^2)_{ij}| &\leq |2 - D_{ii} - D_{jj}| |M_{ij}^{(0)}| + |(\mathbf{M}^{(0)2})_{ij}| \\ &\leq K\sqrt{\frac{\log n}{n}} + |(\mathbf{M}^{(0)2})_{ij}| \end{aligned} \quad (11.50)$$

with high probability for all $i \neq j$ by our previous reasoning. For the remaining term, let $\mathbf{G} \in \mathbb{R}^{n \times r}$ have the \mathbf{g}_i as its columns and the \mathbf{h}_i as its rows, so that $\mathbf{M}^{(0)} = (1 - \alpha/2)\frac{\delta^{-1}}{n}\mathbf{G}\mathbf{G}^\top$.

Then, we may write

$$\begin{aligned}
& (\mathbf{M}^{(0)^2})_{ij} \\
&= (1 - \alpha/2)^2 \frac{\delta^{-2}}{n^2} \mathbf{e}_i^\top \mathbf{G} \mathbf{G}^\top \mathbf{G} \mathbf{G}^\top \mathbf{e}_j \\
&= (1 - \alpha/2)^2 \delta^{-2} \frac{1}{n^2} \mathbf{h}_i^\top \left(\sum_{k=1}^n \mathbf{h}_k \mathbf{h}_k^\top \right) \mathbf{h}_j \\
&= (1 - \alpha/2)^2 \delta^{-2} \frac{1}{n^2} \left((\|\mathbf{h}_i\|_2^2 + \|\mathbf{h}_j\|_2^2) \langle \mathbf{h}_i, \mathbf{h}_j \rangle + \mathbf{h}_i^\top \left(\sum_{k \in [n] \setminus \{i,j\}} \mathbf{h}_k \mathbf{h}_k^\top \right) \mathbf{h}_j \right). \quad (11.51)
\end{aligned}$$

In the last factor, by our previous reasoning with high probability the first summand is, in magnitude, at most $Kn^{3/2} \log n$ for all $i \neq j$. In the second summand, note that for each fixed $i \neq j$, the vectors $\mathbf{h}_i, \mathbf{h}_j$, and the matrix $\sum_{k \in [n] \setminus \{i,j\}} \mathbf{h}_k \mathbf{h}_k^\top =: \mathbf{B}^{\sim\{i,j\}}$ are independent. By Proposition 11.3.5, for all $i \neq j$, $\|\mathbf{B}^{\sim\{i,j\}}\| \leq Kn$, and therefore also $\|\mathbf{B}^{\sim\{i,j\}}\|_F^2 \leq n^3$, with high probability. Conditioning on the value of this matrix and applying the Hanson-Wright inequality [RV13] for a fixed $i \neq j$ then shows, after a union bound, that with high probability, $|\mathbf{h}_i^\top \mathbf{B}^{\sim\{i,j\}} \mathbf{h}_j| \leq Kn^{13/8}$ for all $i \neq j$ (indeed, this will hold with any exponent larger than $3/2$). Thus we find $\epsilon_{\text{offdiag}}(\mathbf{M}^2) \leq Kn^{-3/8}$ with high probability (this, in turn, will hold with any exponent smaller than $1/2$, which coincides with our expectation that we should have $\epsilon_{\text{offdiag}}(\mathbf{M}^2) \lesssim K\epsilon_{\text{offdiag}}(\mathbf{M})$ for \mathbf{M} close to a rescaled projection matrix).

Therefore, the constant c from the statement of Theorem 10.11.2 will satisfy

$$c \leq \frac{K}{n} \left(\sqrt{n} + n \cdot n^{-3/8} + n^2 \cdot n^{-9/8} \right) \leq Kn^{-1/8}. \quad (11.52)$$

The theorem produces $\tilde{\mathbb{E}}$ a degree 6 pseudoexpectation with $\tilde{\mathbb{E}}[\mathbf{x}\mathbf{x}^\top] = (1 - c)\mathbf{M} + c\mathbf{I}_n$. Suppose we have chosen δ small enough that, with high probability, the largest δn eigenvalues

of \mathbf{W} are at least $2 - \epsilon/2$. Then we have, with high probability,

$$n^{-1}\text{SOS}_6(\mathbf{W}) \geq n^{-1}\langle \mathbf{W}, (1 - c)\mathbf{M} + c\mathbf{I}_n \rangle \quad (11.53)$$

$$\geq (1 - c)n^{-1}\langle \mathbf{W}, \mathbf{M}^{(0)} \rangle - \alpha\|\mathbf{W}\| - c|\text{tr}(\mathbf{W})| \quad (11.54)$$

and since we have, with high probability, $\|\mathbf{W}\| \leq 2 + \alpha$, $|\text{tr}(\mathbf{W})| \leq \log n$, and $\langle \mathbf{W}, \mathbf{M}^{(0)} \rangle \geq \lambda_r(\mathbf{M}^{(0)})\lambda_{\delta n}(\mathbf{W}) \geq (1 - \alpha)(2 - \epsilon/2)$, we find on this event that

$$\geq (1 - Kn^{-1/8})(1 - \alpha)(2 - \epsilon/2) - \alpha(2 + \alpha) - Kn^{-1/8} \log n \quad (11.55)$$

and choosing α sufficiently small, depending on our choice of δ above, we will have for sufficiently large n

$$\geq 2 - \epsilon, \quad (11.56)$$

completing the proof. □

11.5 TENSORIAL SAMPLE COVARIANCE MATRICES

In this section we give the proof of Lemma 11.4.2 controlling a covariance-like matrix consisting of Gaussian tensors, as well as a variant for tensors built from a matrix distributed uniformly on the Stiefel manifold, used in [KB20] and involving a different and independently interesting proof technique. Related questions have been studied recently in, e.g., [Ver20]; see also our references below.

11.5.1 GAUSSIAN TENSORS: PROOF OF LEMMA 11.4.2

The proof will use powerful concentration inequalities developed for low-rank sample covariance matrices of this kind for random vectors having a certain concentration property.

Definition 11.5.1 (Orlicz norm). *Let $\mathbf{x} \in \mathbb{R}^d$ be a random vector. For $\rho \geq 1$, we define its ψ_ρ norm as*

$$\|\mathbf{x}\|_{\psi_\rho} := \sup_{\substack{\mathbf{y} \in \mathbb{R}^d \\ \|\mathbf{y}\|_2=1}} \inf \left\{ C > 0 : \mathbb{E} \left[\exp \left(\left(\frac{|\langle \mathbf{x}, \mathbf{y} \rangle|}{C} \right)^\rho \right) \right] \leq 2 \right\}. \quad (11.57)$$

(The specific constant 2 is not essential, but is a common convention. Note also that ψ_r is the standard notation, but we switch the index variable to ρ to avoid confusion with our dimension variable r .)

Proposition 11.5.2 (Theorem 3.3 of [ALPTJ11]). *Let $\rho \in [1, 2]$. Let $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^d$ be independent centered random vectors, each having finite ψ_ρ norm, and let $\psi := \max_{i \in [n]} \|\mathbf{x}_i\|_{\psi_\rho}$. Let $\mathbf{A} \in \mathbb{R}^{d \times n}$ have the $\mathbf{x}_1, \dots, \mathbf{x}_n$ as its columns. There exist universal constants $C_1, C_2 > 0$ such that, for any $\theta \in (0, 1)$,*

$$\left| \|\mathbf{A}\mathbf{y}\|_2^2 - 1 \right| \leq C_1 (\psi\sqrt{d} + \sqrt{1 + \theta})^2 \sqrt{\frac{n}{d}} \left[1 + \log \left(\sqrt{\frac{d}{n}} \right) \right]^{1/\rho} + \theta \text{ for all } \mathbf{y} \in \mathbb{S}^{d-1} \quad (11.58)$$

with probability at least

$$1 - C_1 \exp \left(-C_2 \sqrt{n} \left[1 + \log \left(\sqrt{\frac{d}{n}} \right) \right] \right) - 2\mathbb{P} \left[\max_{i \in [n]} \left| \|\mathbf{x}_i\|_2^2 - 1 \right| \geq \theta \right]. \quad (11.59)$$

To apply Proposition 11.5.2, we must show that the distribution of the vectors $\frac{1}{r}\mathbf{a}_i$ has bounded ψ_ρ norm for some $\rho \in [1, 2]$. The following achieves this for $\rho = 1$.

Proposition 11.5.3. *Let $\mathbf{h} \sim \mathcal{N}(0, \mathbf{I}_r)$ and $\mathbf{a} = \text{isovec}(\mathbf{h}\mathbf{h}^\top - \mathbf{I}_r)$. Then, $\|\mathbf{a}\|_{\psi_1} \leq C$ for a universal constant $C > 0$ (concretely, $C = 30$ suffices).*

Let us give the intuition behind this claim. Bounded ψ_1 norm is a certain characterization of subexponential distribution of the linear forms of a random vector. After centering, a vector of squared gaussians exhibits tail decay of this order by well-known concentration inequalities [LM00]. Thus, assuming that the highest variance components of \mathbf{a} correspond to the diagonal entries of $\mathbf{h}\mathbf{h}^\top - \mathbf{I}_r$, our result is not surprising.

Proof of Proposition 11.5.3. Given $\mathbf{y} \in \mathbb{R}^{r(r+1)/2}$ with $\|\mathbf{y}\|_2 = 1$, let us view $\mathbf{y} = \text{isovec}(\mathbf{Y})$ for $\mathbf{Y} \sim \mathbb{R}_{\text{sym}}^{r \times r}$ with $\|\mathbf{Y}\|_F = 1$. By the spectral theorem, there exist $\mathbf{U} \in \mathcal{O}(r)$ and $\boldsymbol{\lambda} \in \mathbb{R}^r$ such that $\mathbf{Y} = \mathbf{U} \text{diag}(\boldsymbol{\lambda}) \mathbf{U}^\top$, $\|\boldsymbol{\lambda}\|_2 = \|\mathbf{Y}\|_F = 1$, and $\|\boldsymbol{\lambda}\|_\infty = \|\mathbf{Y}\| \leq 1$. We have

$$\langle \mathbf{a}, \mathbf{y} \rangle = \mathbf{h}^\top \mathbf{Y} \mathbf{h} - \text{tr}(\mathbf{Y}) = \sum_{i=1}^r \lambda_i \left(\langle \mathbf{h}, \mathbf{u}_i \rangle^2 - 1 \right). \quad (11.60)$$

Let $\boldsymbol{\lambda}^+, \boldsymbol{\lambda}^- \in \mathbb{R}^r$ have $\lambda_i^+ = \max(0, \lambda_i)$ and $\lambda_i^- = -\min(0, \lambda_i)$, so that $\lambda_i^\pm \geq 0$ and $\boldsymbol{\lambda} = \boldsymbol{\lambda}^+ - \boldsymbol{\lambda}^-$. Note that $\|\boldsymbol{\lambda}^\pm\|_2 \leq \|\boldsymbol{\lambda}\|_2 = 1$ and $\|\boldsymbol{\lambda}^\pm\|_\infty \leq \|\boldsymbol{\lambda}\|_\infty \leq 1$.

We bound

$$|\langle \mathbf{a}, \mathbf{y} \rangle| \leq \underbrace{\left| \sum_{i=1}^r \lambda_i^+ \left(\langle \mathbf{h}, \mathbf{u}_i \rangle^2 - 1 \right) \right|}_{K^+(\mathbf{y})} + \underbrace{\left| \sum_{i=1}^r \lambda_i^- \left(\langle \mathbf{h}, \mathbf{u}_i \rangle^2 - 1 \right) \right|}_{K^-(\mathbf{y})}. \quad (11.61)$$

Note that $\langle \mathbf{h}, \mathbf{u}_i \rangle$ for $i \in [r]$ are r i.i.d. random variables distributed as $\mathcal{N}(0, 1)$, since the \mathbf{u}_i form an orthonormal basis. Suppose $t \geq 1$. Then, using Lemma 1 of [LM00],

$$\begin{aligned} & \mathbb{P} [|\langle \mathbf{a}, \mathbf{y} \rangle| \geq 8t] \\ & \leq \mathbb{P} \left[|\langle \mathbf{a}, \mathbf{y} \rangle| \geq 4 \left(\|\boldsymbol{\lambda}\|_2 \sqrt{t} + \|\boldsymbol{\lambda}\|_\infty t \right) \right] \\ & \leq \mathbb{P} \left[K^+(\mathbf{a}) \geq 2 \left(\|\boldsymbol{\lambda}\|_2 \sqrt{t} + \|\boldsymbol{\lambda}\|_\infty t \right) \right] + \mathbb{P} \left[K^-(\mathbf{a}) \geq 2 \left(\|\boldsymbol{\lambda}\|_2 \sqrt{t} + \|\boldsymbol{\lambda}\|_\infty t \right) \right] \\ & \leq 4e^{-t}. \end{aligned} \quad (11.62)$$

Fix $C > 0$, then changing variables in (11.62) we find that for any $b \geq \exp(\frac{8}{C})$,

$$\mathbb{P} \left[\exp \left(\frac{|\langle \mathbf{a}, \mathbf{y} \rangle|}{C} \right) \geq b \right] \leq 4b^{-C/8}. \quad (11.63)$$

Integrating (11.63), we find that for any $b_0 \geq \exp(\frac{8}{C})$ and $C > 8$,

$$\begin{aligned} \mathbb{E} \left[\exp \left(\frac{|\langle \mathbf{a}, \mathbf{y} \rangle|}{C} \right) \right] &\leq b_0 + \int_{b_0}^{\infty} 4b^{-C/8} db \\ &= b_0 + \frac{32b_0^{1-C/8}}{C-8} \end{aligned}$$

Optimizing over b_0 , we find the optimal value $b_0^* = b^{8/C}$, whereby

$$\mathbb{E} \left[\exp \left(\frac{|\langle \mathbf{s}, \mathbf{a} \rangle|}{C} \right) \right] \leq 4^{8/C} \left(1 + \frac{8}{C-8} \right).$$

We see that the above expression tends to 1 as $C \rightarrow \infty$, and in particular is smaller than 2 for sufficiently large C . One may verify numerically that $C = 30$ suffices. \square

To control the other term appearing in (11.59), we will also need to control the norms of the \mathbf{a}_i .

Proposition 11.5.4. *For any $K > 0$, there exist $C_1, C_2 > 0$ such that, for any $r \geq 1$, letting $\mathbf{h} \sim \mathcal{N}(0, \mathbf{I}_r)$ and $\mathbf{a} = \text{isovec}(\mathbf{h}\mathbf{h}^\top - \mathbf{I}_r)$,*

$$\mathbb{P} \left[\left| \frac{1}{r^2} \|\mathbf{a}\|_2^2 - 1 \right| \geq C_1 \sqrt{\frac{\log r}{r}} \right] \leq \frac{C_2}{r^K}. \quad (11.64)$$

Proof. We compute

$$\frac{1}{r^2} \|\mathbf{a}\|_2^2 = \frac{1}{r^2} \|\mathbf{h}\mathbf{h}^\top - \mathbf{I}_r\|_F^2 = \left(\frac{\|\mathbf{h}\|_2}{\sqrt{r}} \right)^4 - 2 \left(\frac{\|\mathbf{h}\|_2}{r} \right)^2 + \frac{1}{r}. \quad (11.65)$$

The result then follows by applying Proposition 11.3.4 to the first and second terms. \square

Proof of Lemma 11.4.2. We apply Proposition 11.5.2 to the vectors $\frac{1}{r}\mathbf{a}_i$. Note that in our case $d = \frac{r(r+1)}{2}$ while $\psi = O(1/r)$ by Proposition 11.5.3, whereby $\psi\sqrt{d} = O(1)$. Taking $\theta = C\sqrt{\frac{\log n}{n}}$ for C sufficiently large and applying Proposition 11.5.4 then gives the result. \square

A very similar situation to ours is also treated as an application of [ALPTJ11] by [FJ19]. Analogous results are mentioned without proof in [AHH12] for the distribution $\mathbf{a}' = \mathbf{h} \otimes \mathbf{h}'$ where \mathbf{h} and \mathbf{h}' are i.i.d. standard gaussian vectors. Thus we show that the ψ_1 norm does not “see” the additional weak dependences present in our distribution. This is not a new idea; for instance, it was shown in [Cun14] that in a suitable parameter regime the sample covariance matrices of either of these distributions of vectors have empirical spectral distribution converging to the same Marcenko-Pastur limit. Unfortunately, these results are obtained with free probability techniques and thus, unlike a moment calculation, do not directly assist in controlling the largest eigenvalue. In [AHH12] a moment calculation is carried out for the $\mathbf{h} \otimes \mathbf{h}'$ distribution; extending this to the $\mathbf{h} \otimes \mathbf{h}$ distribution is an interesting challenge.

11.5.2 HAAR TENSORS

We also consider the analogous situation where the law of the vectors in question is associated to the uniform or Haar measure on the Stiefel manifold. The Stiefel manifolds are defined as follows:

$$\text{Stief}(n, r) := \{\mathbf{V} \in \mathbb{R}^{r \times n} : \mathbf{V}\mathbf{V}^\top = \mathbf{I}_r\}. \quad (11.66)$$

In words, $\text{Stief}(n, r)$ consists of the $r \times n$ matrices with orthonormal rows. The Haar measure $\text{Haar}(\text{Stief}(n, r))$ is the unique measure on $\text{Stief}(n, r)$ that is invariant under the action of $\mathcal{O}(n)$ on $\text{Stief}(n, r)$ by multiplication on the right. Equivalently, $\text{Haar}(\text{Stief}(n, r))$ is the measure obtained by restricting $\text{Haar}(\mathcal{O}(n))$ (defined in the usual way) to the upper $r \times n$ matrix block.

These measures enjoy the following concentration inequality when $r < n$, obtained by standard arguments from logarithmic Sobolev or isoperimetric inequalities for the special orthogonal group $S\mathcal{O}(n)$, of which $\text{Stief}(n, r)$ is a quotient when $r < n$ (see, e.g., the discussion following Theorem 2.4 of [Led01]).

Proposition 11.5.5. *Suppose $1 \leq r < n$, and $F : \text{Stief}(n, r) \rightarrow \mathbb{R}$ has Lipschitz constant at most L when $\text{Stief}(n, r)$ is endowed with the metric of the Frobenius matrix norm. Then, for an absolute constant $C > 0$,*

$$\mathbb{P}_{\mathbf{V} \sim \text{Haar}(\text{Stief}(n, r))} [|F(\mathbf{V}) - \mathbb{E}F(\mathbf{V})| \geq t] \leq 2 \exp\left(-\frac{C n t^2}{L^2}\right). \quad (11.67)$$

We also register the following preliminary result on the moments of Haar-distributed orthogonal matrices. The following gives the low-degree moments of Haar-distributed orthogonal matrices.

Proposition 11.5.6 (Lemma 9 of [CM07]). *Let $\mathbf{Q} \sim \text{Haar}(\mathcal{O}(n))$. The moment $\mathbb{E} \prod_{k=1}^d Q_{i_k j_k}$ is zero if any index occurs an odd number of times among either the i_k or j_k . The non-zero degree 2 and 4 moments are given by*

$$\mathbb{E} Q_{11}^2 = \frac{1}{n}, \quad (11.68)$$

$$\mathbb{E} Q_{11}^4 = \frac{3}{n(n+2)}, \quad (11.69)$$

$$\mathbb{E} Q_{11}^2 Q_{12}^2 = \frac{1}{n(n+2)}, \quad (11.70)$$

$$\mathbb{E} Q_{11}^2 Q_{22}^2 = \frac{n+1}{(n-1)n(n+2)}, \quad (11.71)$$

$$\mathbb{E} Q_{11} Q_{12} Q_{21} Q_{22} = -\frac{1}{(n-1)n(n+2)}. \quad (11.72)$$

Our result is then as follows.

Lemma 11.5.7. *Let $r = \delta n$, $\mathbf{V} \sim \text{Haar}(\text{Stief}(n, r))$, and $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^r$ be the columns*

of \mathbf{V} . Let \mathbf{A}^{orth} have $\text{isovec}(\delta^{-1}\mathbf{v}_i\mathbf{v}_i^\top - \frac{1}{r}\mathbf{I}_r)$ as its columns (recall $\text{isovec}(\cdot)$ is the isometric vectorization of Definition 8.2.2). Let $\mathbf{P}_{\mathbf{1}_n^\perp} := \mathbf{I}_n - \frac{1}{n}\mathbf{1}_n\mathbf{1}_n^\top$, the orthogonal projector to the subspace orthogonal to $\mathbf{1}_n$. Then,

$$\mathbb{P}\left[\left\|\mathbf{A}^{\text{orth}\top}\mathbf{A}^{\text{orth}} - \mathbf{P}_{\mathbf{1}_n^\perp}\right\| \leq O_\delta\left(\frac{\log n}{n^{1/4}}\right)\right] \geq 1 - \exp\left(-\Omega_\delta(n^{1/2})\right). \quad (11.73)$$

Proof. Note that $\mathbf{A}^{\text{orth}}\mathbf{1}_n = \text{isovec}(\frac{n}{r}\mathbf{V}\mathbf{V}^\top - \frac{n}{r}\mathbf{I}_r) = \mathbf{0}$; thus it is impossible for \mathbf{A}^{orth} to act on \mathbb{R}^n as an approximate isometric embedding, as we might naively expect from its weakly dependent columns. Our argument is more natural to carry out if we remove this caveat; therefore, let us define $\mathbf{A}_0^{\text{orth}}$ to have columns $\text{isovec}(\frac{n}{r}\mathbf{v}_i\mathbf{v}_i^\top - \frac{1-\sqrt{\delta}}{r}\mathbf{I}_r)$. One may check that $\|\mathbf{A}_0^{\text{orth}}\mathbf{1}_n\|_2 = \|\mathbf{1}_n\|_2 = \sqrt{n}$, and that

$$\mathbf{A}_0^{\text{orth}\top}\mathbf{A}_0^{\text{orth}} = \mathbf{A}^{\text{orth}\top}\mathbf{A}^{\text{orth}} + \frac{1}{n}\mathbf{1}_n\mathbf{1}_n^\top. \quad (11.74)$$

In particular, $\mathbf{A}_0^{\text{orth}\top}\mathbf{A}_0^{\text{orth}} - \mathbf{I}_n = \mathbf{A}^{\text{orth}\top}\mathbf{A}^{\text{orth}} - \mathbf{P}_{\mathbf{1}_n^\perp}$, so it suffices to show the operator norm bound of (11.67) for $\mathbf{A}_0^{\text{orth}\top}\mathbf{A}_0^{\text{orth}} - \mathbf{I}_n$.

For $\mathbf{x} \in \mathbb{R}^n$, let us denote $\mathbf{D}_x := \text{diag}(\mathbf{x})$ for the course of this proof. Then,

$$\begin{aligned} \mathbf{A}_0^{\text{orth}}\mathbf{x} &= \text{isovec}\left(\delta^{-1}\sum_{i=1}^n x_i\mathbf{v}_i\mathbf{v}_i^\top - \frac{1-\sqrt{\delta}}{r}\langle\mathbf{1}_n, \mathbf{x}\rangle\mathbf{I}_r\right) \\ &= \delta^{-1}\text{isovec}\left(\mathbf{V}\mathbf{D}_x\mathbf{V}^\top - \frac{1-\sqrt{\delta}}{n}\langle\mathbf{1}_n, \mathbf{x}\rangle\mathbf{I}_r\right). \end{aligned} \quad (11.75)$$

For $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, define

$$F_{\mathbf{x}, \mathbf{y}}(\mathbf{V}) := \langle \mathbf{A}_0^{\text{orth}}\mathbf{x}, \mathbf{A}_0^{\text{orth}}\mathbf{y} \rangle. \quad (11.76)$$

Then, recalling that $\mathbf{P} := \mathbf{V}^\top \mathbf{V}$ is the orthogonal projector to the row space of \mathbf{V} ,

$$\begin{aligned}
F_{x,y}(\mathbf{V}) &= \delta^{-2} \left\langle \mathbf{V} \mathbf{D}_x \mathbf{V}^\top - \frac{1 - \sqrt{\delta}}{n} \langle \mathbf{1}_n, \mathbf{x} \rangle \mathbf{I}_r, \mathbf{V} \mathbf{D}_y \mathbf{V}^\top - \frac{1 - \sqrt{\delta}}{n} \langle \mathbf{1}_n, \mathbf{y} \rangle \mathbf{I}_r \right\rangle \\
&= \delta^{-2} \left[\text{tr}(\mathbf{D}_x \mathbf{P} \mathbf{D}_y \mathbf{P}) \right. \\
&\quad \left. - \frac{1 - \sqrt{\delta}}{n} (\langle \mathbf{1}_n, \mathbf{x} \rangle \text{tr}(\mathbf{P} \mathbf{D}_y) + \langle \mathbf{1}_n, \mathbf{y} \rangle \text{tr}(\mathbf{P} \mathbf{D}_x)) \right. \\
&\quad \left. + \frac{\delta(1 - \sqrt{\delta})^2}{n} \langle \mathbf{1}_n, \mathbf{x} \rangle \langle \mathbf{1}_n, \mathbf{y} \rangle \right]. \tag{11.77}
\end{aligned}$$

Let us denote balls in Euclidean space by

$$B(\mathbf{x}, r) := \{\mathbf{y} \in \mathbb{R}^n : \|\mathbf{x} - \mathbf{y}\|_2 \leq r\}. \tag{11.78}$$

Our first goal will be to obtain concentration bounds on $F_{x,y}(\mathbf{V})$ when $\mathbf{V} \sim \text{Haar}(\text{Stief}(n, r))$ for each fixed pair $(\mathbf{x}, \mathbf{y}) \in B(\mathbf{0}, 1)^2$, by applying the Lipschitz concentration inequality.

Claim 1. *Let $\mathbf{x}, \mathbf{y} \in B(\mathbf{0}, 1)$. Then,*

$$\text{Lip}(F_{x,y}) \leq 4\delta^{-2} \left[\min\{\|\mathbf{x}\|_\infty, \|\mathbf{y}\|_\infty\} + \frac{1}{\sqrt{n}} \right]. \tag{11.79}$$

Proof. For $\mathbf{V}_1, \mathbf{V}_2 \in \text{Stief}(n, r)$, letting $\mathbf{P}_i = \mathbf{V}_i^\top \mathbf{V}_i$, we have using (11.77) and the triangle inequality

$$\begin{aligned}
&\delta^2 |F_{x,y}(\mathbf{V}_1) - F_{x,y}(\mathbf{V}_2)| \\
&= |\text{tr}(\mathbf{D}_x \mathbf{P}_1 \mathbf{D}_y \mathbf{P}_1) - \text{tr}(\mathbf{D}_x \mathbf{P}_2 \mathbf{D}_y \mathbf{P}_2)| \\
&\quad + \frac{(1 - \sqrt{\delta}) |\langle \mathbf{1}_n, \mathbf{x} \rangle|}{n} |\text{tr}(\mathbf{P}_1 \mathbf{D}_y) - \text{tr}(\mathbf{P}_2 \mathbf{D}_y)| \\
&\quad + \frac{(1 - \sqrt{\delta}) |\langle \mathbf{1}_n, \mathbf{y} \rangle|}{n} |\text{tr}(\mathbf{P}_1 \mathbf{D}_x) - \text{tr}(\mathbf{P}_2 \mathbf{D}_x)|,
\end{aligned}$$

then using that $|\langle \mathbf{1}_n, \mathbf{x} \rangle| \leq \|\mathbf{x}\|_1 \leq \sqrt{n}$ and likewise for \mathbf{y} ,

$$\begin{aligned}
&\leq |\operatorname{tr}(\mathbf{D}_x \mathbf{P}_1 \mathbf{D}_y (\mathbf{P}_1 - \mathbf{P}_2))| + |\operatorname{tr}(\mathbf{D}_x (\mathbf{P}_1 - \mathbf{P}_2) \mathbf{D}_y \mathbf{P}_2)| \\
&\quad + \frac{1}{\sqrt{n}} (|\operatorname{tr}((\mathbf{P}_1 - \mathbf{P}_2) \mathbf{D}_x)| + |\operatorname{tr}((\mathbf{P}_1 - \mathbf{P}_2) \mathbf{D}_y)|) \\
&\leq (\|\mathbf{D}_x \mathbf{P}_1\|_F + \|\mathbf{D}_x \mathbf{P}_2\|_F) \|\mathbf{D}_y (\mathbf{P}_1 - \mathbf{P}_2)\|_F \\
&\quad + \frac{2}{\sqrt{n}} \|\mathbf{P}_1 - \mathbf{P}_2\|_F. \tag{11.80}
\end{aligned}$$

Since \mathbf{P}_i is an orthogonal projector for $i \in \{1, 2\}$,

$$\|\mathbf{D}_x \mathbf{P}_i\|_F = \langle \mathbf{D}_x^2, \mathbf{P}_i \rangle^{1/2} \leq (\operatorname{tr}[\mathbf{D}_x^2])^{1/2} \leq 1. \tag{11.81}$$

We bound the other term by

$$\|\mathbf{D}_y (\mathbf{P}_1 - \mathbf{P}_2)\|_F = \langle \mathbf{D}_y^2, (\mathbf{P}_1 - \mathbf{P}_2)^2 \rangle^{1/2} \leq \|\mathbf{y}\|_\infty \|\mathbf{P}_1 - \mathbf{P}_2\|_F. \tag{11.82}$$

Combining these observations and a symmetric argument with \mathbf{x} and \mathbf{y} in opposite roles gives

$$|F_{x,y}(\mathbf{V}_1) - F_{x,y}(\mathbf{V}_2)| \leq 2\delta^{-2} \left[\min\{\|\mathbf{x}\|_\infty, \|\mathbf{y}\|_\infty\} + \frac{1}{\sqrt{n}} \right] \|\mathbf{P}_1 - \mathbf{P}_2\|_F. \tag{11.83}$$

Lastly, we bound

$$\begin{aligned}
\|\mathbf{P}_1 - \mathbf{P}_2\|_F &= \|\mathbf{V}_1^\top \mathbf{V}_1 - \mathbf{V}_2^\top \mathbf{V}_2\|_F \\
&= \|(\mathbf{V}_1 - \mathbf{V}_2)^\top \mathbf{V}_1 + \mathbf{V}_2^\top (\mathbf{V}_1 - \mathbf{V}_2)\|_F
\end{aligned}$$

where by triangle inequality

$$\begin{aligned}
&\leq \|(\mathbf{V}_1 - \mathbf{V}_2)^\top \mathbf{V}_1\|_F + \|\mathbf{V}_2^\top (\mathbf{V}_1 - \mathbf{V}_2)\|_F \\
&= (\text{tr}[(\mathbf{V}_1 - \mathbf{V}_2)^\top \mathbf{V}_1 \mathbf{V}_1^\top (\mathbf{V}_1 - \mathbf{V}_2)])^{1/2} \\
&\quad + (\text{tr}[(\mathbf{V}_1 - \mathbf{V}_2)^\top \mathbf{V}_2 \mathbf{V}_2^\top (\mathbf{V}_1 - \mathbf{V}_2)])^{1/2} \\
&= 2\|\mathbf{V}_1 - \mathbf{V}_2\|_F,
\end{aligned} \tag{11.84}$$

where we have used that $\mathbf{V}_1 \mathbf{V}_1^\top = \mathbf{V}_2 \mathbf{V}_2^\top = \mathbf{I}_r$, and the result follows. \square

Therefore, and what is crucial to our argument, while for the worst-case $\mathbf{x} \in B(\mathbf{0}, 1)$, namely $\mathbf{x} = \mathbf{e}_i$ a standard basis vector, $F_{\mathbf{x}, \mathbf{x}}$ will have Lipschitz constant $O(1)$, for typical $\mathbf{x} \in B(\mathbf{0}, 1)$, $F_{\mathbf{x}, \mathbf{x}}$ will rather have Lipschitz constant $\tilde{O}(n^{-1/2})$. Moreover, the Lipschitz constant of $F_{\mathbf{x}, \mathbf{y}}$ is comparable to the *smaller* of the Lipschitz constants of $F_{\mathbf{x}, \mathbf{x}}$ and $F_{\mathbf{y}, \mathbf{y}}$.

Claim 2. For $\mathbf{x}, \mathbf{y} \in B(\mathbf{0}, 1)$, $\mathbb{E}_{\mathbf{V} \sim \text{Haar}(\text{Stief}(n, r))} F_{\mathbf{x}, \mathbf{y}}(\mathbf{V}) = \langle \mathbf{x}, \mathbf{y} \rangle + O_\delta(n^{-1})$.

Proof. We have

$$\begin{aligned}
\delta^2 \mathbb{E} F_{\mathbf{x}, \mathbf{y}}(\mathbf{V}) &= \mathbb{E} \text{tr}(\mathbf{V} \mathbf{D}_x \mathbf{V}^\top \mathbf{V} \mathbf{D}_y \mathbf{V}^\top) \\
&\quad - \frac{1 - \sqrt{\delta}}{n} \langle \mathbf{1}_n, \mathbf{x} \rangle \mathbb{E} \text{tr}(\mathbf{V}^\top \mathbf{V} \mathbf{D}_y) \\
&\quad - \frac{1 - \sqrt{\delta}}{n} \langle \mathbf{1}_n, \mathbf{y} \rangle \mathbb{E} \text{tr}(\mathbf{V}^\top \mathbf{V} \mathbf{D}_x) \\
&\quad + \frac{\delta(1 - \sqrt{\delta})^2}{n} \langle \mathbf{1}_n, \mathbf{x} \rangle \langle \mathbf{1}_n, \mathbf{y} \rangle,
\end{aligned} \tag{11.85}$$

and by either the moment formulae of Proposition 11.5.6 or an argument from orthogonal invariance of Haar measure, we have $\mathbb{E} \mathbf{V}^\top \mathbf{V} = \delta \mathbf{I}_n$, whereby

$$= \mathbb{E} \text{tr}(\mathbf{V} \mathbf{D}_x \mathbf{V}^\top \mathbf{V} \mathbf{D}_y \mathbf{V}^\top) - \frac{\delta(1 - \delta)}{n} \langle \mathbf{1}_n, \mathbf{x} \rangle \langle \mathbf{1}_n, \mathbf{y} \rangle. \tag{11.86}$$

View $\mathbf{V} \sim \text{Haar}(\text{Stief}(n, r))$ as the top $r \times n$ block of $\mathbf{Q} \sim \text{Haar}(\mathcal{O}(n))$. Then, expanding the first term with the moment formulae of Proposition 11.5.6,

$$\begin{aligned}
& \mathbb{E} \text{tr}(\mathbf{V} \mathbf{D}_x \mathbf{V}^\top \mathbf{V} \mathbf{D}_y \mathbf{V}^\top) \\
&= \sum_{i,j=1}^n x_i y_j \left(\sum_{a,b=1}^r \mathbb{E}[Q_{ai} Q_{bi} Q_{aj} Q_{bj}] \right) \\
&= \sum_{i=1}^n x_i y_i \left(\frac{3r}{n(n+2)} + \frac{r(r-1)}{n(n+2)} \right) \\
&\quad + \sum_{1 \leq i < j \leq n} x_i y_j \left(\frac{r}{n(n+2)} - \frac{r(r-1)}{(n-1)n(n+2)} \right) \\
&= \frac{\delta}{n+2} \left(r+1 + \frac{r-1}{n-1} \right) \sum_{i=1}^n x_i y_i \\
&\quad + \frac{\delta}{n+2} \left(1 - \frac{r-1}{n-1} \right) \left(\sum_{i=1}^n x_i \right) \left(\sum_{i=1}^n y_i \right) \\
&= \left(\delta^2 + O_\delta(n^{-1}) \right) \langle \mathbf{x}, \mathbf{y} \rangle + \left(1 + O_\delta(n^{-1}) \right) \frac{\delta(1-\delta)}{n} \langle \mathbf{1}_n, \mathbf{x} \rangle \langle \mathbf{1}_n, \mathbf{y} \rangle, \tag{11.87}
\end{aligned}$$

and since $|\langle \mathbf{x}, \mathbf{y} \rangle| \leq \|\mathbf{x}\|_2 \|\mathbf{y}\|_2 \leq 1$ and $|\langle \mathbf{1}_n, \mathbf{x} \rangle| \cdot |\langle \mathbf{1}_n, \mathbf{y} \rangle| \leq n$, the result follows. \square

Combining Claim 1, Claim 2, and the concentration result Proposition 11.5.5, we find the following corollary on pointwise concentration of $F_{x,y}(\mathbf{V})$.

Claim 3. *There exist constants $C_1, C_2 > 0$ depending only on δ such that, for any $\mathbf{x}, \mathbf{y} \in B(\mathbf{0}, 1)$,*

$$\begin{aligned}
& \mathbb{P}_{\mathbf{V} \sim \text{Haar}(\text{Stief}(n, r))} \left[\left| \langle \mathbf{A}_0^{\text{orth}} \mathbf{x}, \mathbf{A}_0^{\text{orth}} \mathbf{y} \rangle - \langle \mathbf{x}, \mathbf{y} \rangle \right| \geq \frac{C_1}{n} + t \right] \\
& \leq 2 \exp \left(- \frac{C_2 n t^2}{(\min\{\|\mathbf{x}\|_\infty, \|\mathbf{y}\|_\infty\} + n^{-1/2})^2} \right). \tag{11.88}
\end{aligned}$$

This concludes the first part of the argument.

The remaining part of the argument is to apply a union bound of the probabilities controlled in Claim 3 over suitable nets of $B(\mathbf{0}, 1)$. We divide our task into a bound over

sparse vectors and vectors with bounded largest entry, very similar to the technique in [Rud08, RV08] and especially [Ver11]. Introduce a parameter $\rho \in (0, 1)$ to be chosen later. Define

$$B_s := \{\mathbf{y} \in B(\mathbf{0}, 1) : \|\mathbf{y}\|_0 \leq \rho n\}, \quad (11.89)$$

$$B_b := \left\{ \mathbf{z} \in B(\mathbf{0}, 1) : \|\mathbf{z}\|_\infty \leq \frac{1}{\sqrt{\rho n}} \right\}. \quad (11.90)$$

For any $\mathbf{x} \in B(\mathbf{0}, 1)$, we define $\mathbf{y} = \mathbf{y}(\mathbf{x})$ and $\mathbf{z} = \mathbf{z}(\mathbf{x})$ by thresholding the entries of \mathbf{x} , setting $y_i := x_i \mathbb{1}\{|x_i| > \frac{1}{\sqrt{\rho n}}\}$ and $z_i := x_i \mathbb{1}\{|x_i| \leq \frac{1}{\sqrt{\rho n}}\}$. Then, $\mathbf{x} = \mathbf{y} + \mathbf{z}$, $\mathbf{y} \in B_s$, and $\mathbf{z} \in B_b$.

Introduce another parameter $\gamma \in (0, 1)$ to be chosen later. Let $\mathcal{N}_s \subset B_s$ and $\mathcal{N}_b \subset B_b$ be γ -nets. By a standard bound (see, e.g., Lemma 9.5 of [LT13]), we may choose $|\mathcal{N}_b| \leq \exp(2n/\gamma)$, and by the same bound applied to each choice of ρn support coordinates for an element of B_s , we may choose

$$|\mathcal{N}_s| \leq \binom{n}{\lfloor \rho n \rfloor} \exp\left(\frac{2\rho n}{\gamma}\right) \leq \exp\left(\frac{2\rho n}{\gamma} + \rho n + \log\left(\frac{1}{\rho}\right) \rho n\right). \quad (11.91)$$

To lighten the notation, let us set $\mathbf{S} := \mathbf{A}_0^{\text{orth}\top} \mathbf{A}_0^{\text{orth}} - \mathbf{I}_n$. The following is an adaptation to our setting of a standard technique for estimating a matrix norm over a net: we first bound

$$\begin{aligned} \|\mathbf{S}\| &= \max_{\mathbf{x} \in B(\mathbf{0}, 1)} |\mathbf{x}^\top \mathbf{S} \mathbf{x}| \\ &\leq \max_{\substack{\mathbf{y} \in B_s \\ \mathbf{z} \in B_b}} |(\mathbf{y} + \mathbf{z})^\top \mathbf{S}(\mathbf{y} + \mathbf{z})| \\ &\leq \max_{\mathbf{y} \in B_s} |\mathbf{y}^\top \mathbf{S} \mathbf{y}| + \max_{\mathbf{z} \in B_b} |\mathbf{z}^\top \mathbf{S} \mathbf{z}| + 2 \max_{\substack{\mathbf{y} \in B_s \\ \mathbf{z} \in B_b}} |\mathbf{y}^\top \mathbf{S} \mathbf{z}| \\ &\leq \max_{\mathbf{y} \in \mathcal{N}_s} |\mathbf{y}^\top \mathbf{S} \mathbf{y}| + \max_{\mathbf{z} \in \mathcal{N}_b} |\mathbf{z}^\top \mathbf{S} \mathbf{z}| + 2 \max_{\substack{\mathbf{y} \in \mathcal{N}_s \\ \mathbf{z} \in \mathcal{N}_b}} |\mathbf{y}^\top \mathbf{S} \mathbf{z}| + 12\gamma \|\mathbf{S}\|. \end{aligned} \quad (11.92)$$

Rearranging this, we obtain

$$\|\mathbf{S}\| \leq \frac{1}{1-12\gamma} \left[\max_{\mathbf{y} \in \mathcal{N}_s} |\mathbf{y}^\top \mathbf{S} \mathbf{y}| + \max_{\mathbf{z} \in \mathcal{N}_b} |\mathbf{z}^\top \mathbf{S} \mathbf{z}| + 2 \max_{\substack{\mathbf{y} \in \mathcal{N}_s \\ \mathbf{z} \in \mathcal{N}_b}} |\mathbf{y}^\top \mathbf{S} \mathbf{z}| \right]. \quad (11.93)$$

Using Claim 3 and a union bound, we have that

$$\begin{aligned} & \mathbb{P} \left[\|\mathbf{S}\| \geq \frac{4}{1-12\gamma} \left(\frac{C_1}{n} + t \right) \right] \\ & \leq 2(|\mathcal{N}_b| + |\mathcal{N}_s| \cdot |\mathcal{N}_b|) \exp \left(-C_2 \frac{\rho}{(1+\sqrt{\rho})^2} n^2 t^2 \right) \\ & \quad + 2|\mathcal{N}_s| \exp \left(-\frac{C_2}{2} n t^2 \right) \\ & \leq 3 \exp \left(n \left[\frac{2}{\gamma} (1+\rho) + \rho + \log \left(\frac{1}{\rho} \right) \rho - C_2 \frac{\rho}{(1+\sqrt{\rho})^2} n t^2 \right] \right) \\ & \quad + 2 \exp \left(n \left[\frac{2\rho}{\gamma} + \rho + \log \left(\frac{1}{\rho} \right) \rho - \frac{C_2}{2} t^2 \right] \right). \end{aligned} \quad (11.94)$$

Taking $\rho = n^{-1/2}$, $t = C_3 n^{-1/4} \log n$ for a large constant C_3 , and $\gamma < \frac{1}{12}$ a small constant, we obtain the result. \square

A | OPEN PROBLEMS

We briefly present several open problems motivated by the various topics discussed in the main text.

A.1 LOW-DEGREE METHOD BEYOND INTEGRABLE MODELS

Our analysis of the low-degree likelihood ratio in Chapter 4 depended on elegant “integrability” properties of models where we make observations drawn from convenient distributions. In particular, we used at length the algebraic and combinatorial properties of orthogonal polynomials of these distributions, drawing on the “umbral calculus” of Hermite polynomials and various related systems of identities for other polynomial families. There are two directions in which these distributional assumptions may be weakened. First, we may consider well-behaved entrywise distributions with complicated correlations across entries, such as in the random regular graph ensemble $\text{Reg}(n, d)$. Second, we may consider i.i.d. distributions whose entrywise distributions do not belong to convenient exponential families. Is it possible to develop techniques for the computations called for by the low-degree heuristic for such situations, techniques that do not rely on detailed knowledge of the associated orthogonal polynomials?

A.2 OVERLAPS, LOW-DEGREE POLYNOMIALS, AND STATISTICAL PHYSICS

We showed in Chapter 4 that the efficacy of low-degree polynomials for problems like spiked matrix models and many generalizations thereof are governed by the *overlap distribution* of the spike or signal prior, the law of $\langle \mathbf{x}^1, \mathbf{x}^2 \rangle$ for $\mathbf{x}^1, \mathbf{x}^2 \sim \mathcal{P}_n$ independently in the case of a rank-one spiked matrix model. Similar distributions with overlaps of two independent draws from a *posterior* distribution also arise naturally in Bayesian analyses of problems like these, especially in their treatment with the methods of statistical physics. Can this formal resemblance be used to show that the predictions made by statistical physics methods and predictions based on limitations of low-degree polynomials may, for some broad class of problems, be equivalent?

A.3 CHANNEL MONOTONICITY

The result of Theorem 4.3.14 suggests two intriguing open problems further probing the low degree method.

1. Are the channel monotonicity predictions accurate, i.e., can they be corroborated with any other form of evidence of computational hardness? (One intriguing possibility is average-case reductions in the style of [BR13, BB20] between different NEF-QVFs.)
2. If these predictions are accurate, then does *strict* inequality hold in computational cost between any of these versions of a given problem, or does *channel universality* hold (we borrow the term from [LKZ15a] but use it in a slightly different sense), where in fact computational complexity of testing does not depend on the NEF-QVF through which the data are observed?

A.4 FISHER INFORMATION-ORTHOGONAL POLYNOMIAL

IDENTITIES

As we have mentioned, the original argument of [PWBM18] derives the critical value λ^* for the non-Gaussian spiked matrix model we treated in Section 5.4.2 in terms of the Fisher information in the family of translates of the distribution ρ^{sech} , while our calculation, if we consider $D = D(n)$ growing slowly, obtains the same predicted value using orthogonal polynomials. It appears that the connection between these derivations lies in the summation identity $\sum_{\ell \geq 0} \frac{1}{(2\ell+1)^2} = \frac{\pi^2}{8}$. We suspect that there are similar identities associated to these two approaches to calculating the critical signal-to-noise ratio in spiked matrix models for other algebraically-convenient noise measures ρ . It would be interesting to understand what class of summation identities arises in this way, and whether equating these two derivations can give novel proofs of such identities.

A.5 CAN LOW-DEGREE POLYNOMIALS EXPLORE?

The discussion in Section 5.4.2 on a non-Gaussian spiked matrix model suggests that, for algorithms computing low-degree polynomials, there is a tension between robustness to heavy-tailed noise distributions and optimality for specific rapidly-decaying (and, in that case, non-Gaussian) noise distributions. In particular, we might expect low-degree polynomials to have difficulty performing optimally on problems with an *unknown* noise distribution. For example, it was shown in [MRY18] that there is a distribution-agnostic algorithm achieving optimal performance in a wide range of Wigner spiked matrix models, that first identifies the noise distribution using kernel density estimation and then applies a recovery algorithm tailored to that distribution. Since the initial phase of such an algorithm is often

said to fall under the rubric of “exploratory data analysis,” we call this algorithmic strategy *exploration before inference*. Can low-degree polynomials explore in this way?

It is easy to show that, for instance, if we draw observations from a model where the noise is drawn with probability $\frac{1}{2}$ from ρ^{sech} for all entries and with probability $\frac{1}{2}$ from some heavy-tailed ρ^{heavy} with only finitely many moments for all entries, then low-degree polynomials will be suboptimal in the sense we have been considering, simply because $L^2(\mathbb{Q}_n)$ will only contain polynomials of bounded entrywise degree. Is this an artificial example that can be repaired, or does it indicate a more fundamental weakness of low-degree algorithms?

A.6 SUM-OF-SQUARES AND ENTANGLEMENT

We have seen that the family of matrices $\mathcal{B}^{n,r}$ which is, through the results of Chapter 7, essentially equivalent to the feasible set of degree 4 SOS over the hypercube, is also closely related to the partial transpose operation. We explained in that chapter how this operation is used as a test to detect entanglement of bipartite quantum states, and showed in Theorem 7.3.3 that, roughly speaking, separability of matrices in $\mathcal{B}^{n,r}$ is equivalent to integrality of pseudomoments over the hypercube. In [DPS04], the authors show that the partial transpose test may be viewed as only the first of a family of entanglement criteria, and that these criteria eventually detect any entangled state. Is there an equivalence between this “DPS hierarchy” of entanglement criteria (or a real-valued analog thereof) and higher degrees of SOS over the hypercube, analogous to the equivalence we have shown between “passing” the first level of DPS and being feasible for degree 4 of SOS?¹

¹Not to be confused with the different and previously-known equivalence between the DPS hierarchy and SOS over the *sphere* [FF20].

A.7 GENERALIZED MAXWELL-SYLVESTER REPRESENTATIONS

We have seen that even a heuristic choice of a Green’s function for an approximate Maxwell-Sylvester representation of multiharmonic polynomials (Section 10.1.1) is a powerful tool for deriving predictions of SOS pseudomoments. Alas, even the original Maxwell-Sylvester representation of spherical harmonics is a somewhat obscure topic, and Clerc’s generalizations to certain multiharmonic settings is also little-known. Thus we ask: in what generality do Maxwell-Sylvester representations of multiharmonic polynomials exist? What is the structure of the associated Green’s functions and Kelvin transforms? As a concrete example from the results obtained here, can the isotypic projection expressing $h_S(\mathbf{V}^\top \mathbf{z})$ for $\mathbf{V} \in \mathbb{R}^{(n-1) \times n}$ the simplex frame studied in Chapter 9 (see Definition 9.3.1) be expressed through differentiation of some Green’s function?

A.8 PSEUDOCALIBRATION RECONCILIATION

Both our results using the spectral pseudomoment extensions of Chapter 8 and those of [MRX20, GJJ⁺20] using pseudocalibration prove SOS lower bounds for the SK Hamiltonian. Moreover, both constructions may be reasonably viewed as the “simplest possible” extensions of degree 2 pseudomoments given by a rescaled low-rank projection matrix—spectral extensions give the simplest extension from the point of view of a Gram factorization, while pseudocalibration gives the simplest extension “calibrated” to the individual moments of a planted distribution (see Section 3.2.2). It is therefore natural to conjecture that the constructions are closely related, perhaps with spectral extensions, which appear at least superficially simpler, giving some form of first-order approximation of pseudocalibration. However, it remains unclear how to draw such a connection, as quite different combinatorial objects appear in either construction: the forest poset $\mathcal{F}(m)$ and its Möbius function

drive the spectral extension construction, while evaluations of Hermite polynomials (whose coefficients may be related to combinatorics of matchings) appear in pseudocalibration. Can these seemingly different constructions be reconciled and given a unified interpretation?

A.9 REPLICATED SUM-OF-SQUARES

A great deal of the theory surrounding the SK model and spin glass models more generally takes advantage of studying *replicas*, independent copies from a Gibbs measure, and quantities such as their *overlap* or inner product, whose distribution and in particular its support describes aspects of the geometry of the optimization landscape (see [MPV87, Tal10, Pan13]). In [RS00], the authors proposed a variant of Hilbert’s seventeenth problem (on expressing non-negative rational functions as sums-of-squares of rational functions) related to “umbral polynomials.” In simple terms, this amounts to a variant of SOS reasoning with independent copies of the pseudo-random variable x allowed. The authors refer to a conjecture of Rota’s that any polynomial of moments that is non-negative for all real-valued random variables can be written as the expectation of a sum-of-squares involving independent copies. (For example, $\text{Var}[X] = \mathbb{E}X^2 - (\mathbb{E}X)^2 \geq 0$ for all real-valued X , and $\text{Var}[X] = \frac{1}{2}\mathbb{E}(X^1 - X^2)^2$ for independent copies X^i .) Do there exist low-degree proofs in this extended proof system that give non-trivial bounds on the SK Hamiltonian?

BIBLIOGRAPHY

- [ABAČ13] Antonio Auffinger, Gérard Ben Arous, and Jiří Černý. Random matrices and complexity of spin glasses. *Communications on Pure and Applied Mathematics*, 66(2):165–201, 2013.
- [Abb17] Emmanuel Abbe. Community detection and stochastic block models: recent developments. *The Journal of Machine Learning Research*, 18(1):6446–6531, 2017.
- [Abi19] Aida Abiad. A characterization and an application of weight-regular partitions of graphs. *Linear Algebra and its Applications*, 569:162–174, 2019.
- [ABM20] Louigi Addario-Berry and Pascal Maillard. The algorithmic hardness threshold for continuous random energy models. *Mathematical Statistics and Learning*, 2(1):77–101, 2020.
- [ABW13] Sheldon Axler, Paul Bourdon, and Ramey Wade. *Harmonic function theory*. Springer Science & Business Media, 2013.
- [ACO08] Dimitris Achlioptas and Amin Coja-Oghlan. Algorithmic barriers from phase transitions. In *49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2008)*, pages 793–802. IEEE, 2008.

- [ADGM17] Anima Anandkumar, Yuan Deng, Rong Ge, and Hossein Mobahi. Homotopy analysis for tensor PCA. In *30th Annual Conference on Learning Theory (COLT 2017)*, pages 79–104. PMLR, 2017.
- [AG84] Bengt Aspvall and John R Gilbert. Graph coloring using eigenvalue decomposition. *SIAM Journal on Algebraic Discrete Methods*, 5(4):526–538, 1984.
- [AGZ10] Greg W Anderson, Alice Guionnet, and Ofer Zeitouni. *An introduction to random matrices*. Cambridge University Press, 2010.
- [AHH12] Andris Ambainis, Aram W Harrow, and Matthew B Hastings. Random tensor theory: extending random matrix theory to mixtures of random product states. *Communications in Mathematical Physics*, 310(1):25–74, 2012.
- [AK97] Noga Alon and Nabil Kahale. A spectral technique for coloring random 3-colorable graphs. *SIAM Journal on Computing*, 26(6):1733–1748, 1997.
- [AKS98] Noga Alon, Michael Krivelevich, and Benny Sudakov. Finding a large hidden clique in a random graph. *Random Structures & Algorithms*, 13(3-4):457–466, 1998.
- [ALPTJ11] Radoslaw Adamczak, Alexander E Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann. Restricted isometry property of matrices with independent columns and neighborly polytopes by random sampling. *Constructive Approximation*, 34(1):61–88, 2011.
- [ALR87] Michael Aizenman, Joel L Lebowitz, and David Ruelle. Some rigorous results on the Sherrington-Kirkpatrick spin glass model. *Communications in Mathematical Physics*, 112(1):3–20, 1987.

- [AM04] Dimitris Achlioptas and Cristopher Moore. The chromatic number of random regular graphs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 219–228. Springer, 2004.
- [AM20] Ahmed El Alaoui and Andrea Montanari. Algorithmic thresholds in mean field spin glasses. *arXiv preprint arXiv:2009.11481*, 2020.
- [AMP16] Kwangjun Ahn, Dhruv Medarametla, and Aaron Potechin. Graph matrices: norm bounds and applications. *arXiv preprint arXiv:1604.03423*, 2016.
- [AMS20] Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Optimization of mean-field spin glasses. *arXiv preprint arXiv:2001.00904*, 2020.
- [AN04] Dimitris Achlioptas and Assaf Naor. The two possible values of the chromatic number of a random graph. In *36th Annual ACM Symposium on Theory of Computing (STOC 2004)*, pages 587–593, 2004.
- [And81] Désiré André. Sur les permutations alternées. *Journal de mathématiques pures et appliquées*, 7:167–184, 1881.
- [AR95] Sheldon Axler and Wade Ramey. Harmonic polynomials and Dirichlet-type problems. *Proceedings of the American Mathematical Society*, pages 3765–3773, 1995.
- [Arn96] Vladimir I Arnol'd. Topological content of the Maxwell theorem on multiple representation of spherical functions. *Topological Methods in Nonlinear Analysis*, 7(2):205–217, 1996.
- [Arn13] Vladimir I Arnol'd. *Lectures on partial differential equations*. Springer Science & Business Media, 2013.

- [Art27] Emil Artin. Über die zerlegung definiter funktionen in quadrate. In *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, volume 5, pages 100–115, 1927.
- [ART06] Dimitris Achlioptas and Federico Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. In *38th Annual ACM Symposium on Theory of Computing (STOC 2006)*, pages 130–139, 2006.
- [AS17] Guillaume Aubrun and Stanisław J Szarek. *Alice and Bob Meet Banach: The Interface of Asymptotic Geometric Analysis and Quantum Information Theory*. American Mathematical Society, 2017.
- [AU03] David Avis and Jun Umemoto. Stronger linear programming relaxations of max-cut. *Mathematical Programming*, 97(3):451–469, 2003.
- [BB17] Jacob Biamonte and Ville Bergholm. Tensor networks in a nutshell. *arXiv preprint arXiv:1708.00006*, 2017.
- [BB19a] Matthew Brennan and Guy Bresler. Average-case lower bounds for learning sparse mixtures, robust estimation and semirandom adversaries. *arXiv preprint arXiv:1908.06130*, 2019.
- [BB19b] Matthew Brennan and Guy Bresler. Optimal average-case reductions to sparse PCA: From weak assumptions to strong hardness. In *32nd Annual Conference on Learning Theory (COLT 2019)*, pages 469–470. PMLR, 2019.
- [BB20] Matthew Brennan and Guy Bresler. Reducibility and statistical-computational gaps from secret leakage. In *33rd Annual Conference on Learning Theory (COLT 2020)*, pages 648–847. PMLR, 2020.

- [BBAP05] Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697, 2005.
- [BBH18] Matthew Brennan, Guy Bresler, and Wasim Huleihel. Reducibility and computational lower bounds for problems with planted sparse structure. In *31st Annual Conference On Learning Theory (COLT 2018)*, pages 48–166. PMLR, 2018.
- [BBH⁺20] Matthew Brennan, Guy Bresler, Samuel B Hopkins, Jerry Li, and Tselil Schramm. Statistical query algorithms and low-degree tests are almost equivalent. *arXiv preprint arXiv:2009.06107*, 2020.
- [BBK⁺20] Afonso S Bandeira, Jess Banks, Dmitriy Kunisky, Cristopher Moore, and Alexander S Wein. Spectral planting and the hardness of refuting cuts, colorability, and communities in random graphs. *arXiv preprint arXiv:2008.12237*, 2020.
- [BDM⁺16] Jean Barbier, Mohamad Dia, Nicolas Macris, Florent Krzakala, Thibault Lesieur, and Lenka Zdeborová. Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula. *arXiv preprint arXiv:1606.04142*, 2016.
- [Ber18] Christoph Berkholz. The relation between polynomial calculus, Sherali-Adams, and sum-of-squares proofs. In *35th Symposium on Theoretical Aspects of Computer Science (STACS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [BES20] Arthur Bik, Henrik Eisenmann, and Bernd Sturmfels. Jordan algebras of symmetric matrices. *arXiv preprint arXiv:2010.00277*, 2020.
- [BF03] John J Benedetto and Matthew Fickus. Finite normalized tight frames. *Advances in Computational Mathematics*, 18(2-4):357–385, 2003.

- [BG75] Edward A Bender and Jay R Goldman. On the applications of Möbius inversion in combinatorial analysis. *The American Mathematical Monthly*, 82(8):789–803, 1975.
- [BGG⁺16] Vijay Bhattiprolu, Mrinalkanti Ghosh, Venkatesan Guruswami, Euiwoong Lee, and Madhur Tulsiani. Multiplicative approximations for polynomial optimization over the unit sphere. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:185, 2016.
- [BGJ20] Gérard Ben Arous, Reza Gheissari, and Aukosh Jagannath. Algorithmic thresholds for tensor PCA. *Annals of Probability*, 48(4):2052–2087, 2020.
- [BGL16] Vijay Bhattiprolu, Venkatesan Guruswami, and Euiwoong Lee. Sum-of-squares certificates for maxima of random tensors on the sphere. *arXiv preprint arXiv:1605.00903*, 2016.
- [BGN11] Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011.
- [BGP16] Grigoriy Blekherman, João Gouveia, and James Pfeiffer. Sums of squares on the hypercube. *Mathematische Zeitschrift*, 284(1-2):41–54, 2016.
- [BGT10] Mohsen Bayati, David Gamarnik, and Prasad Tetali. Combinatorial approach to the interpolation method and scaling limits in sparse random graphs. In *42nd ACM Symposium on Theory of Computing (STOC 2010)*, pages 105–114, 2010.
- [BHK⁺19] Boaz Barak, Samuel Hopkins, Jonathan Kelner, Pravesh K Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM Journal on Computing*, 48(2):687–735, 2019.

- [BIK⁺96] Paul Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, and Pavel Pudlák. Lower bounds on Hilbert’s Nullstellensatz and propositional proofs. *Proceedings of the London Mathematical Society*, 3(1):1–26, 1996.
- [BK18] Afonso S Bandeira and Dmitriy Kunisky. A Gramian description of the degree 4 generalized ellipsope. *arXiv preprint arXiv:1812.11583*, 2018.
- [BK19a] Afonso S Bandeira and Dmitriy Kunisky. Connections between structured tight frames and sum-of-squares optimization. In *Wavelets and Sparsity XVIII (SPIE 2019)*, volume 11138, page 111381F. International Society for Optics and Photonics, 2019.
- [BK19b] Afonso S Bandeira and Dmitriy Kunisky. Sum-of-squares optimization and the sparsity structure of equiangular tight frames. In *13th International Conference on Sampling Theory and Applications (SampTA 2019)*, pages 1–4. IEEE, 2019.
- [BKM19] Jess Banks, Robert Kleinberg, and Cristopher Moore. The Lovász theta function for random regular graphs and community detection in the hard regime. *SIAM Journal on Computing*, 48(3):1098–1119, 2019.
- [BKS16] Afonso S Bandeira, Christopher Kennedy, and Amit Singer. Approximating the little Grothendieck problem over the orthogonal and unitary groups. *Mathematical Programming*, 160(1-2):433–475, 2016.
- [BKW20a] Afonso S Bandeira, Dmitriy Kunisky, and Alexander S Wein. Average-case integrality gap for non-negative principal component analysis. *arXiv preprint arXiv:2012.02243*, 2020.
- [BKW20b] Afonso S Bandeira, Dmitriy Kunisky, and Alexander S Wein. Computational hardness of certifying bounds on constrained PCA problems. In *11th Inno-*

- vations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151, pages 78:1–78:29, 2020.
- [BM03] Samuel Burer and Renato DC Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003.
- [BM17] Debapratim Banerjee and Zongming Ma. Optimal hypothesis testing for stochastic block models with growing degrees. *arXiv preprint arXiv:1705.05305*, 2017.
- [BMNN16] Jess Banks, Cristopher Moore, Joe Neeman, and Praneeth Netrapalli. Information-theoretic thresholds for community detection in sparse networks. In *29th Annual Conference on Learning Theory (COLT 2016)*, pages 383–416. PMLR, 2016.
- [BMR21] Jess Banks, Sidhanth Mohanty, and Prasad Raghavendra. Local statistics, semidefinite programming, and community detection. In *32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2021)*, pages 1298–1316. SIAM, 2021.
- [BMV⁺18] Jess Banks, Cristopher Moore, Roman Vershynin, Nicolas Verzelen, and Jiaming Xu. Information-theoretic bounds and phase transitions in clustering, sparse PCA, and submatrix localization. *IEEE Transactions on Information Theory*, 64(7):4872–4894, 2018.
- [Bor82] Karl Heinx Borgwardt. The average number of pivot steps required by the simplex-method is polynomial. *Zeitschrift für Operations Research*, 26(1):157–177, 1982.
- [Bor88] Karl Heinx Borgwardt. Probabilistic analysis of the simplex method. In *DGOR/N-SOR*, pages 564–575. Springer, 1988.

- [BPT12] Grigoriy Blekherman, Pablo A Parrilo, and Rekha R Thomas. *Semidefinite optimization and convex algebraic geometry*. SIAM, 2012.
- [BPW18] Afonso S Bandeira, Amelia Perry, and Alexander S Wein. Notes on computational-to-statistical gaps: predictions using statistical physics. *arXiv preprint arXiv:1803.11132*, 2018.
- [BR13] Quentin Berthet and Philippe Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In *26th Annual Conference on Learning Theory (COLT 2013)*, pages 1046–1066, 2013.
- [Bro87] W Dale Brownawell. Bounds for the degrees in the Nullstellensatz. *Annals of Mathematics*, 126(3):577–591, 1987.
- [BRS15] Jop Briët, Oded Regev, and Rishi Saket. Tight hardness of the non-commutative Grothendieck problem. In *56th Annual Symposium on Foundations of Computer Science (STOC 2015)*, pages 1108–1122. IEEE, 2015.
- [BS84] Aart Blokhuis and Johan J Seidel. Introduction to multilinear algebra and some applications. *Philips J. Res.*, 39(4):111–120, 1984.
- [BS06] Jinho Baik and Jack W Silverstein. Eigenvalues of large sample covariance matrices of spiked population models. *Journal of Multivariate Analysis*, 97(6):1382–1408, 2006.
- [BS14] Boaz Barak and David Steurer. Sum-of-squares proofs and the quest toward optimal algorithms. *arXiv preprint arXiv:1404.5236*, 2014.
- [BS16] Boaz Barak and David Steurer. Proofs, beliefs, and algorithms through the lens of sum-of-squares, 2016.

- [BSW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow - resolution made simple. *Journal of the ACM (JACM)*, 48(2):149-169, 2001.
- [BT10] Daniel Berend and Tamir Tassa. Improved bounds on Bell numbers and on moments of sums of random variables. *Probability and Mathematical Statistics*, 30(2):185-205, 2010.
- [BTN01] Ahron Ben-Tal and Arkadi Nemirovski. *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*. SIAM, 2001.
- [BV04] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [BVM] Andries E Brouwer and Hendrik Van Maldeghem. Strongly regular graphs.
- [BŻ17] Ingemar Bengtsson and Karol Życzkowski. *Geometry of quantum states: an introduction to quantum entanglement*. Cambridge University Press, 2017.
- [Cam80] Peter J Cameron. 6-transitive graphs. *Journal of Combinatorial Theory, Series B*, 28(2):168-179, 1980.
- [CD06] Charles J Colbourn and Jeffrey H Dinitz. *Handbook of combinatorial designs*. CRC Press, 2006.
- [CD12] Lin Chen and Dragomir Ž Djoković. Qubit-qudit states with positive partial transpose. *Physical Review A*, 86(6):062332, 2012.
- [CDMF09] Mireille Capitaine, Catherine Donati-Martin, and Delphine Féral. The largest eigenvalues of finite rank deformation of large Wigner matrices: convergence and nonuniversality of the fluctuations. *The Annals of Probability*, 37(1):1-47, 2009.

- [CGPR19] Wei-Kuo Chen, David Gamarnik, Dmitry Panchenko, and Mustazee Rahman. Sub-optimality of local algorithms for a class of max-cut problems. *The Annals of Probability*, 47(3):1587–1618, 2019.
- [CH89] Richard Courant and David Hilbert. *Methods of mathematical physics*, volume 1. Wiley, New York, 1989.
- [Cha05] Sourav Chatterjee. An error bound in the Sudakov-Fernique inequality. *arXiv preprint math/0510424*, 2005.
- [CHK⁺20] Yeshwanth Cherapanamjeri, Samuel B Hopkins, Tarun Kathuria, Prasad Raghavendra, and Nilesch Tripuraneni. Algorithms for heavy-tailed statistics: Regression, covariance estimation, and beyond. In *52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2020)*, pages 601–609, 2020.
- [Chv73] Vasek Chvatal. Edmonds polytopes and a hierarchy of combinatorial problems. *Discrete Mathematics*, 4(4):305–337, 1973.
- [CK12] Peter G Casazza and Gitta Kutyniok. *Finite frames: Theory and applications*. Springer, 2012.
- [CL19] Hye Won Chung and Ji Oon Lee. Weak detection of signal in the spiked Wigner model. In *International Conference on Machine Learning*, pages 1233–1241. PMLR, 2019.
- [Cla09] Pete L Clark. Quadratic Reciprocity II: The Proofs (Lecture Notes), 2009. URL: <http://math.uga.edu/~pete/NT2009qrproof.pdf>. Last visited on 2018/05/21. Quotation: “Working through this proof feels a little bit like being an accountant who has been assigned to carefully document a miracle”.

- [Cle00] Jean-Louis Clerc. Kelvin transform and multi-harmonic polynomials. *Acta Mathematica*, 185(1):81-99, 2000.
- [CM07] Sourav Chatterjee and Elizabeth Meckes. Multivariate normal approximation using exchangeable pairs. *arXiv preprint math/0701464*, 2007.
- [CMW15] Tony Cai, Zongming Ma, and Yihong Wu. Optimal estimation and rank detection for sparse spiked covariance matrices. *Probability Theory and Related Fields*, 161(3-4):781-815, 2015.
- [CO03] Amin Coja-Oghlan. The Lovász number of random graphs. In *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques*, pages 228-239. Springer, 2003.
- [COE15] Amin Coja-Oghlan and Charilaos Efthymiou. On independent sets in random graphs. *Random Structures & Algorithms*, 47(3):436-486, 2015.
- [COEH16] Amin Coja-Oghlan, Charilaos Efthymiou, and Samuel Hetterich. On the chromatic number of random regular graphs. *Journal of Combinatorial Theory, Series B*, 116:367-439, 2016.
- [COKV07a] Amin Coja-Oghlan, Michael Krivelevich, and Dan Vilenchik. Why almost all k -CNF formulas are easy. In *13th International Conference on Analysis of Algorithms*, 2007.
- [COKV07b] Amin Coja-Oghlan, Michael Krivelevich, and Dan Vilenchik. Why almost all k -colorable graphs are easy. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 121-132. Springer, 2007.
- [Coo71] Stephen A Cook. The complexity of theorem-proving procedures. In *3rd Annual ACM Symposium on Theory of Computing (STOC 1971)*, pages 151-158, 1971.

- [CR79] Stephen A Cook and Robert A Reckhow. The relative efficiency of propositional proof systems. *The Journal of Symbolic Logic*, 44(1):36–50, 1979.
- [CR02] Andrea Crisanti and Tommaso Rizzo. Analysis of the ∞ -replica symmetry breaking solution of the Sherrington-Kirkpatrick model. *Physical Review E*, 65(4):046137, 2002.
- [CRT08] Peter G Casazza, Dan Redmond, and Janet C Tremain. Real equiangular frames. In *42nd Annual Conference on Information Sciences and Systems (CISS 2008)*, pages 715–720. IEEE, 2008.
- [CS05] Maria Chudnovsky and Paul D Seymour. The structure of claw-free graphs. *Surveys in Combinatorics*, 327:153–171, 2005.
- [CT12] Eden Chlamtac and Madhur Tulsiani. Convex relaxations and integrality gaps. In *Handbook on semidefinite, conic and polynomial optimization*, pages 139–169. Springer, 2012.
- [Cun14] Fabio Deelan Cunden. *Spectral methods in random matrix theory: from classical ensembles to quantum random tensors*. PhD thesis, University of Bari Aldo Moro, 2014.
- [Dan65] George Bernard Dantzig. *Linear programming and extensions*. Princeton University Press, 1965.
- [dB19] Corwin de Boer. In search of degree-4 sum-of-squares lower bounds for MaxCut. Master’s thesis, Carnegie Mellon University, 2019.
- [DCS03] Giuseppe Dattoli, Clemente Cesarano, and Dario Sacchetti. A note on truncated polynomials. *Applied Mathematics and Computation*, 134(2-3):595–605, 2003.

- [DFJ02] Martin Dyer, Alan Frieze, and Mark Jerrum. On counting independent sets in sparse graphs. *SIAM Journal on Computing*, 31(5):1527–1541, 2002.
- [DHS20] Jingqiu Ding, Samuel B Hopkins, and David Steurer. Estimating rank-one spikes from heavy-tailed noise via self-avoiding walks. *arXiv preprint arXiv:2008.13735*, 2020.
- [Dia88] Persi Diaconis. Group representations in probability and statistics. *Lecture Notes Monograph Series*, 11, 1988.
- [DKMZ11a] Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6):066106, 2011.
- [DKMZ11b] Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Inference and phase transitions in the detection of modules in sparse networks. *Physical Review Letters*, 107(6):065701, 2011.
- [DKS17] Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. In *58th Annual Symposium on Foundations of Computer Science (FOCS 2017)*, pages 73–84. IEEE, 2017.
- [DKWB19] Yunzi Ding, Dmitriy Kunisky, Alexander S Wein, and Afonso S Bandeira. Subexponential-time algorithms for sparse PCA. *arXiv preprint arXiv:1907.11635*, 2019.
- [DKWB20] Yunzi Ding, Dmitriy Kunisky, Alexander S Wein, and Afonso S Bandeira. The average-case time complexity of certifying the restricted isometry property. *arXiv preprint arXiv:2005.11270*, 2020.

- [DL09] Michel Marie Deza and Monique Laurent. *Geometry of cuts and metrics*. Springer, 2009.
- [DM14] Yash Deshpande and Andrea Montanari. Sparse PCA via covariance thresholding. In *Advances in Neural Information Processing Systems*, pages 334–342, 2014.
- [DM15a] Yash Deshpande and Andrea Montanari. Finding hidden cliques of size $\sqrt{N/e}$ in nearly linear time. *Foundations of Computational Mathematics*, 15(4):1069–1128, 2015.
- [DM15b] Yash Deshpande and Andrea Montanari. Improved sum-of-squares lower bounds for hidden clique and hidden submatrix problems. In *28th Annual Conference on Learning Theory (COLT 2015)*, pages 523–562, 2015.
- [DMO⁺19] Yash Deshpande, Andrea Montanari, Ryan O’Donnell, Tselil Schramm, and Subhabrata Sen. The threshold for SDP-refutation of random regular NAE-3SAT. In *30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2019)*, pages 2305–2321. SIAM, 2019.
- [DMR14] Yash Deshpande, Andrea Montanari, and Emile Richard. Cone-constrained principal component analysis. In *Advances in Neural Information Processing Systems*, pages 2717–2725, 2014.
- [DMS17] Amir Dembo, Andrea Montanari, and Subhabrata Sen. Extremal cuts of sparse random graphs. *The Annals of Probability*, 45(2):1190–1217, 2017.
- [DPS04] Andrew C Doherty, Pablo A Parrilo, and Federico M Spedalieri. Complete family of separability criteria. *Physical Review A*, 69(2):022308, 2004.

- [DSW07] Josep Díaz, Maria J Serna, and Nicholas C Wormald. Bounds on the bisection width for random d -regular graphs. *Theoretical Computer Science*, 382(2):120–130, 2007.
- [EAKJ20] Ahmed El Alaoui, Florent Krzakala, and Michael Jordan. Fundamental limits of detection in the spiked Wigner model. *Annals of Statistics*, 48(2):863–885, 2020.
- [ER93] Richard Ehrenborg and Gian-Carlo Rota. Apolarity and canonical forms for homogeneous polynomials. *European Journal of Combinatorics*, 14(3):157–181, 1993.
- [ES83] David Elderfield and David Sherrington. The curious case of the Potts spin glass. *Journal of Physics C: Solid State Physics*, 16(15):L497, 1983.
- [Fei02] Uriel Feige. Relations between average case complexity and approximation complexity. In *34th Annual ACM Symposium on Theory of Computing (STOC 2002)*, pages 534–543. ACM, 2002.
- [FF20] Kun Fang and Hamza Fawzi. The sum-of-squares hierarchy on the sphere and applications in quantum information theory. *Mathematical Programming*, pages 1–30, 2020.
- [FGR⁺17] Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh S Vempala, and Ying Xiao. Statistical algorithms and a lower bound for detecting planted cliques. *Journal of the ACM (JACM)*, 64(2):8, 2017.
- [FH04] William Fulton and Joe Harris. *Representation theory: a first course*. Springer Science & Business Media, 2004.
- [Fis13] Matthias J Fischer. *Generalized hyperbolic secant distributions: with applications to finance*. Springer Science & Business Media, 2013.

- [FJ19] Alexander Fengler and Peter Jung. On the restricted isometry property of centered self Khatri-Rao products. *arXiv preprint arXiv:1905.09245*, 2019.
- [FJM⁺16] Matthew Fickus, John Jasper, Dustin G Mixon, Jesse D Peterson, and Cody E Watson. Equiangular tight frames with centroidal symmetry. *Applied and Computational Harmonic Analysis*, 2016.
- [FK00] Uriel Feige and Robert Krauthgamer. Finding and certifying a large hidden clique in a semirandom graph. *Random Structures & Algorithms*, 16(2):195–208, 2000.
- [FM15] Matthew Fickus and Dustin G Mixon. Tables of the existence of equiangular tight frames. *arXiv preprint arXiv:1504.00253*, 2015.
- [FP07] Delphine Féral and Sandrine Péché. The largest eigenvalue of rank one deformation of large Wigner matrices. *Communications in Mathematical Physics*, 272(1):185–228, 2007.
- [FP16] Hamza Fawzi and Pablo A Parrilo. Self-scaled bounds for atomic cone ranks: applications to nonnegative rank and cp-rank. *Mathematical Programming*, 158(1-2):417–465, 2016.
- [FPV18] Vitaly Feldman, Will Perkins, and Santosh Vempala. On the complexity of random satisfiability problems with planted solutions. *SIAM Journal on Computing*, 47(4):1294–1338, 2018.
- [Fri03] Joel Friedman. A proof of Alon’s second eigenvalue conjecture. In *35th Annual ACM Symposium on Theory of Computing (STOC 2003)*, pages 720–724. ACM, 2003.
- [FS09] Philippe Flajolet and Robert Sedgewick. *Analytic combinatorics*. Cambridge University Press, 2009.

- [FSP16] Hamza Fawzi, James Saunderson, and Pablo A Parrilo. Sparse sums of squares on finite abelian groups and improved semidefinite lifts. *Mathematical Programming*, 160(1-2):149–191, 2016.
- [Ful97] William Fulton. *Young tableaux: with applications to representation theory and geometry*. Cambridge University Press, 1997.
- [FV62] David G Feingold and Richard S Varga. Block diagonally dominant matrices and generalizations of the Gerschgorin circle theorem. *Pacific Journal of Mathematics*, 12(4):1241–1250, 1962.
- [FW15] Matthew Fickus and Cody E Watson. Detailing the equivalence between real equiangular tight frames and certain strongly regular graphs. In *Wavelets and Sparsity XVI*, volume 9597. International Society for Optics and Photonics, 2015.
- [Gem80] Stuart Geman. A limit theorem for the norm of random matrices. *The Annals of Probability*, pages 252–261, 1980.
- [GHP02] Dima Grigoriev, Edward A Hirsch, and Dmitrii V Pasechnik. Complexity of semi-algebraic proofs. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 419–430. Springer, 2002.
- [Gic] Victor Gichev. The harmonic component of a homogeneous polynomial. <https://www.puremath.no/wp-content/uploads/2019/02/EsseGichev.pdf>.
- [GJJ⁺20] Mrinalkanti Ghosh, Fernando Granha Jeronimo, Chris Jones, Aaron Potechin, and Goutham Rajendran. Sum-of-squares lower bounds for Sherrington-Kirkpatrick via planted affine planes. *arXiv preprint arXiv:2009.01874*, 2020.

- [GJW20] David Gamarnik, Aukosh Jagannath, and Alexander S Wein. Low-degree hardness of random optimization problems. *arXiv preprint arXiv:2004.12063*, 2020.
- [GKS85] David J Gross, Ido Kanter, and Haim Sompolinsky. Mean-field theory of the Potts glass. *Physical Review Letters*, 55(3):304, 1985.
- [GL18] David Gamarnik and Quan Li. On the max-cut of sparse random graphs. *Random Structures & Algorithms*, 52(2):219–262, 2018.
- [GM75] Geoffrey R Grimmett and Colin JH McDiarmid. On colouring random graphs. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 77, pages 313–324. Cambridge University Press, 1975.
- [Gom10] Ralph E Gomory. Outline of an algorithm for integer solutions to linear programs and an algorithm for the mixed integer problem. In *50 Years of Integer Programming 1958-2008*, pages 77–103. Springer, 2010.
- [Gri01a] Dima Grigoriev. Complexity of Positivstellensatz proofs for the knapsack. *Computational Complexity*, 10(2):139–154, 2001.
- [Gri01b] Dima Grigoriev. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theoretical Computer Science*, 259(1-2):613–622, 2001.
- [GS14] David Gamarnik and Madhu Sudan. Limits of local algorithms over sparse random graphs. In *5th Conference on Innovations in Theoretical Computer Science (ITCS 2014)*, pages 369–376. ACM, 2014.
- [Gue01] Francesco Guerra. Sum rules for the free energy in the mean field spin glass model. *Fields Institute Communications*, 30(11), 2001.
- [Gue03] Francesco Guerra. Broken replica symmetry bounds in the mean field spin glass model. *Communications in Mathematical Physics*, 233(1):1–12, 2003.

- [GV01] Dima Grigoriev and Nicolai Vorobjov. Complexity of Null- and Positivstellensatz proofs. *Annals of Pure and Applied Logic*, 113(1-3):153–160, 2001.
- [GW95] Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.
- [GZ19] David Gamarnik and Ilias Zadik. The landscape of the planted clique problem: Dense subgraphs and the overlap gap property. *arXiv preprint arXiv:1904.07174*, 2019.
- [Har06] Michael Hardy. Combinatorics of partial derivatives. *Electronic Journal of Combinatorics*, 2006.
- [Her55] Carl S Herz. Bessel functions of matrix argument. *Annals of Mathematics*, pages 474–523, 1955.
- [Her98] Grete Hermann. The question of finitely many steps in polynomial ideal theory. *ACM SIGSAM Bulletin*, 32(3):8–30, 1998.
- [HHH96] Michal Horedecki, Pawel Horodecki, and Ryszard Horodecki. Separability of mixed states: necessary and sufficient conditions. *Physical Letters. A*, 223:1–8, 1996.
- [Hil93] David Hilbert. Über die vollen invariantensysteme. *Mathematische Annalen*, 42(3):313–373, 1893.
- [Hil07] Roland Hildebrand. Positive partial transpose from spectra. *Physical Review A*, 76(5):052325, 2007.
- [HKP⁺17] Samuel B Hopkins, Pravesh K Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. The power of sum-of-squares for detect-

- ing hidden structures. In *58th Annual Symposium on Foundations of Computer Science (FOCS 2017)*, pages 720–731. IEEE, 2017.
- [HKP⁺18] Samuel B Hopkins, Pravesh Kothari, Aaron Henry Potechin, Prasad Raghavendra, and Tselil Schramm. On the integrality gap of degree-4 sum of squares for planted clique. *ACM Transactions on Algorithms (TALG)*, 14(3):1–31, 2018.
- [Hof70] Alan J Hoffman. On eigenvalues and colorings of graphs. In Bernard Harris, editor, *Graph theory and its applications*. 1970.
- [Hop18] Samuel Hopkins. *Statistical inference and the sum of squares method*. PhD thesis, Cornell University, 2018.
- [HS17] Samuel B Hopkins and David Steurer. Efficient Bayesian estimation from few samples: community detection and related problems. In *58th Annual Symposium on Foundations of Computer Science (FOCS 2017)*, pages 379–390. IEEE, 2017.
- [HSS15] Samuel B Hopkins, Jonathan Shi, and David Steurer. Tensor principal component analysis via sum-of-square proofs. In *28th Annual Conference on Learning Theory (COLT 2015)*, pages 956–1006, 2015.
- [HSS16] Samuel B Hopkins, Tselil Schramm, Jonathan Shi, and David Steurer. Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors. In *48th Annual ACM Symposium on Theory of Computing (STOC 2016)*, pages 178–191, 2016.
- [HSV20] Guy Holtzman, Adam Soffer, and Dan Vilenchik. A greedy anytime algorithm for sparse PCA. In *33rd Annual Conference on Learning Theory (COLT 2020)*, pages 1939–1956. PMLR, 2020.

- [HWX15] Bruce Hajek, Yihong Wu, and Jiaming Xu. Computational lower bounds for community detection on random graphs. In *28th Annual Conference on Learning Theory (COLT 2015)*, pages 899–928, 2015.
- [IKKM12] Morteza Ibrahimi, Yashodhan Kanoria, Matt Kranning, and Andrea Montanari. The set of solutions of random XORSAT formulae. In *23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)*, pages 760–779. SIAM, 2012.
- [Jae07] Gregg Jaeger. *Quantum information*. Springer, 2007.
- [Jan97] Svante Janson. *Gaussian Hilbert spaces*. Cambridge University Press, 1997.
- [Jer92] Mark Jerrum. Large cliques elude the Metropolis process. *Random Structures & Algorithms*, 3(4):347–359, 1992.
- [JKR19] Vishesh Jain, Frederic Koehler, and Andrej Risteski. Mean-field approximation, convex hierarchies, and the optimality of correlation rounding: a unified perspective. In *51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)*, pages 1226–1236, 2019.
- [JL09] Iain M Johnstone and Arthur Yu Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693, 2009.
- [JLM20] Aukosh Jagannath, Patrick Lopatto, and Leo Miolane. Statistical thresholds for tensor PCA. *Annals of Applied Probability*, 30(4):1910–1933, 2020.
- [Joh01] Iain M Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Annals of Statistics*, pages 295–327, 2001.
- [JP00] Ari Juels and Marcus Peinado. Hiding cliques for cryptographic security. *Designs, Codes and Cryptography*, 20(3):269–280, 2000.

- [JP18] Nathaniel Johnston and Everett Patterson. The inverse eigenvalue problem for entanglement witnesses. *Linear Algebra and its Applications*, 550:1–27, 2018.
- [Kan39] Leonid V Kantorovich. The mathematical method of production planning and organization. *Management Science*, 6(4):363–422, 1939.
- [Kar72] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- [Kar76] Richard M. Karp. The probabilistic analysis of some combinatorial search algorithms. In *Algorithms and complexity: New Directions and Recent Results*, 1976.
- [Kar86] Richard M Karp. Combinatorics, complexity, and randomness. *Communications of the ACM*, 29(2):98–109, 1986.
- [KB20] Dmitriy Kunisky and Afonso S Bandeira. A tight degree 4 sum-of-squares lower bound for the Sherrington-Kirkpatrick Hamiltonian. *Mathematical Programming*, 2020.
- [Kea98] Michael Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM (JACM)*, 45(6):983–1006, 1998.
- [KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007.
- [KLM16] Adam Kurpisz, Samuli Leppänen, and Monaldo Mastrolilli. Sum-of-squares hierarchy lower bounds for symmetric formulations. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 362–374. Springer, 2016.

- [KM72] Victor Klee and George J Minty. How good is the simplex algorithm? *Inequalities*, 3(3):159–175, 1972.
- [KM89] Alexei I Kostrikin and Yuri I Manin. *Linear algebra and geometry*. CRC Press, 1989.
- [KMOW17] Pravesh K Kothari, Ryuhei Mori, Ryan O’Donnell, and David Witmer. Sum of squares lower bounds for refuting any CSP. In *49th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2017)*, pages 132–145, 2017.
- [KMRT⁺07] Florent Krzakala, Andrea Montanari, Federico Ricci-Tersenghi, Guilhem Semerjian, and Lenka Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proceedings of the National Academy of Sciences*, 104(25):10318–10323, 2007.
- [Kol88] János Kollár. Sharp effective Nullstellensatz. *Journal of the American Mathematical Society*, pages 963–975, 1988.
- [Kos02] Eric Kostlan. On the expected number of real roots of a system of random polynomial equations. In *Foundations of Computational Mathematics*, pages 149–188. World Scientific, 2002.
- [KPGW10] Graeme Kemkes, Xavier Pérez-Giménez, and Nicholas Wormald. On the chromatic number of random d -regular graphs. *Advances in Mathematics*, 223(1):300–328, 2010.
- [KS07] Adam R Klivans and Alexander A Sherstov. Unconditional lower bounds for learning intersections of halfspaces. *Machine Learning*, 69(2-3):97–114, 2007.
- [Kuč74] Vladimír Kučera. The matrix equation $ax + xb = c$. *SIAM Journal on Applied Mathematics*, 26(1):15–25, 1974.

- [Kuč95] Luděk Kučera. Expected complexity of graph partitioning problems. *Discrete Applied Mathematics*, 57(2-3):193–212, 1995.
- [Kun20a] Dmitriy Kunisky. Hypothesis testing with low-degree polynomials in the Morris class of exponential families. *arXiv preprint arXiv:2011.03693*, 2020.
- [Kun20b] Dmitriy Kunisky. Positivity-preserving extensions of sum-of-squares pseudomoments over the hypercube. *arXiv preprint arXiv:2009.07269*, 2020.
- [KV05] Subhash Khot and Nisheeth K Vishnoi. On the unique games conjecture. In *46th Annual Symposium on Foundations of Computer Science (FOCS 2005)*, 2005.
- [KV16] Ravi Kannan and Santosh Vempala. Beyond spectral: Tight bounds for planted gaussians. *arXiv preprint arXiv:1608.03643*, 2016.
- [KWB19] Dmitriy Kunisky, Alexander S Wein, and Afonso S Bandeira. Notes on computational hardness of hypothesis testing: predictions using the low-degree likelihood ratio. *arXiv preprint arXiv:1907.11636*, 2019.
- [KXZ16] Florent Krzakala, Jiaming Xu, and Lenka Zdeborová. Mutual information in rank-one matrix estimation. In *2016 IEEE Information Theory Workshop (ITW)*, pages 71–75. IEEE, 2016.
- [KZ09] Florent Krzakala and Lenka Zdeborová. Hiding quiet solutions in random constraint satisfaction problems. *Physical Review Letters*, 102(23):238701, 2009.
- [Lan75] Henry Oliver Lancaster. Joint probability distributions in the Meixner classes. *Journal of the Royal Statistical Society: Series B (Methodological)*, 37(3):434–443, 1975.
- [Las01] Jean B Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.

- [Lau03a] Monique Laurent. A comparison of the Sherali-Adams, Lovász-Schrijver, and Lasserre relaxations for 0-1 programming. *Mathematics of Operations Research*, 28(3):470–496, 2003.
- [Lau03b] Monique Laurent. Lower bound for the number of iterations in semidefinite hierarchies for the cut polytope. *Mathematics of Operations Research*, 28(4):871–883, 2003.
- [Lau09] Monique Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging applications of algebraic geometry*, pages 157–270. Springer, 2009.
- [LCY12] Lucien Le Cam and Grace Lo Yang. *Asymptotics in statistics: some basic concepts*. Springer Science & Business Media, 2012.
- [Led01] Michel Ledoux. *The concentration of measure phenomenon*. Number 89 in Mathematical Surveys & Monographs. American Mathematical Society, 2001.
- [Lev73] Leonid Anatolevich Levin. Universal sequential search problems. *Problemy peredachi informatsii*, 9(3):115–116, 1973.
- [Liu] Ya-Feng Liu. Set-completely-positive representations and cuts for the max-cut polytope and the unit modulus lifting.
- [LKCH00] Maciej Lewenstein, Barbara Kraus, J Ignacio Cirac, and Pawel Horodecki. Optimization of entanglement witnesses. *Physical Review A*, 62(5):052310, 2000.
- [LKZ15a] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. MMSE of probabilistic low-rank matrix estimation: Universality with respect to the output channel. In *53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton 2015)*, pages 680–687. IEEE, 2015.

- [LKZ15b] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. Phase transitions in sparse PCA. In *IEEE International Symposium on Information Theory (ISIT 2015)*, pages 1635–1639. IEEE, 2015.
- [LM00] Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- [LM08] Gérard Letac and H el ene Massam. The noncentral Wishart as an exponential family, and its moments. *Journal of Multivariate Analysis*, 99(7):1393–1417, 2008.
- [LM19] Marc Lelarge and L eo Miolane. Fundamental limits of symmetric low-rank matrix estimation. *Probability Theory and Related Fields*, 173(3):859–929, 2019.
- [LML⁺17] Thibault Lesieur, L eo Miolane, Marc Lelarge, Florent Krzakala, and Lenka Zdeborov a. Statistical and computational phase transitions in spiked tensor estimation. In *IEEE International Symposium on Information Theory (ISIT 2017)*, pages 511–515. IEEE, 2017.
- [LMS10] Jon Magne Leinaas, Jan Myrheim, and Per  yvind Sollid. Numerical studies of entangled positive-partial-transpose states in composite quantum systems. *Physical Review A*, 81(6):062329, 2010.
- [LP96] Monique Laurent and Svatopluk Poljak. On the facial structure of the set of correlation matrices. *SIAM Journal on Matrix Analysis and Applications*, 17(3):530–547, 1996.
- [LR01] I-Li Lu and Donald St P Richards. MacMahon’s master theorem, representation theory, and moments of Wishart distributions. *Advances in Applied Mathematics*, 27(2-3):531–547, 2001.

- [LR06] Erich L Lehmann and Joseph P Romano. *Testing statistical hypotheses*. Springer Science & Business Media, 2006.
- [LRS15] James R Lee, Prasad Raghavendra, and David Steurer. Lower bounds on the size of semidefinite programming relaxations. In *47th Annual ACM Symposium on Theory of Computing (STOC 2015)*, pages 567–576, 2015.
- [LS91] László Lovász and Alexander Schrijver. Cones of matrices and set-functions and 0–1 optimization. *SIAM Journal on Optimization*, 1(2):166–190, 1991.
- [LSG91] Petrus WH Lemmens, Johan J Seidel, and JA Green. Equiangular lines. In *Geometry and Combinatorics*, pages 127–145. Elsevier, 1991.
- [LT94] Chi-Kwong Li and Bit-Shun Tam. A note on extreme correlation matrices. *SIAM Journal on Matrix Analysis and Applications*, 15(3):903–908, 1994.
- [LT13] Michel Ledoux and Michel Talagrand. *Probability in Banach spaces: isoperimetry and processes*. Springer Science & Business Media, 2013.
- [LWB20] Matthias Löffler, Alexander S Wein, and Afonso S Bandeira. Computationally efficient sparse clustering. *arXiv preprint arXiv:2005.10817*, 2020.
- [Max73] James Clerk Maxwell. *A treatise on electricity and magnetism*, volume 1. 1873.
- [McK81] Brendan D McKay. The expected eigenvalue distribution of a large regular graph. *Linear Algebra and its Applications*, 40:203–216, 1981.
- [Mei34] Joseph Meixner. Orthogonale polynomsysteme mit einer besonderen gestalt der erzeugenden funktion. *Journal of the London Mathematical Society*, 1(1):6–13, 1934.

- [MKUZ19] Stefano Sarao Mannelli, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborová. Passed & spurious: Descent algorithms and local minima in spiked matrix-tensor models. In *International Conference on Machine Learning*, pages 4333–4342, 2019.
- [MM82] Ernst W Mayr and Albert R Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Advances in Mathematics*, 46(3):305–329, 1982.
- [MM09] Marc Mézard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.
- [MMZ05] Marc Mézard, Thierry Mora, and Riccardo Zecchina. Clustering of solutions in the random satisfiability problem. *Physical Review Letters*, 94(19):197205, 2005.
- [MNS15] Elchanan Mossel, Joe Neeman, and Allan Sly. Reconstruction and estimation in the planted partition model. *Probability Theory and Related Fields*, 162(3-4):431–461, 2015.
- [Moi20] Ankur Moitra. Sum of squares in theoretical computer science. In *Sum of Squares: Theory and Applications*, volume 77 of *Proceedings of Symposia in Applied Mathematics*. American Mathematical Society, 2020.
- [Mon18] Andrea Montanari. Optimization of the Sherrington-Kirkpatrick Hamiltonian. *arXiv preprint arXiv:1812.10897*, 2018.
- [Moo70] John W Moon. *Counting labelled trees*. Number 1 in Canadian Mathematical Monographs. Canadian Mathematical Congress, 1970.

- [Moo17] Cristopher Moore. The computer science and physics of community detection: landscapes, phase transitions, and hardness. *arXiv preprint arXiv:1702.00467*, 2017.
- [Mor82] Carl N Morris. Natural exponential families with quadratic variance functions. *Annals of Statistics*, pages 65–80, 1982.
- [Mor83] Carl N Morris. Natural exponential families with quadratic variance functions: statistical theory. *Annals of Statistics*, 11(2):515–529, 1983.
- [Mot67] Theodore Samuel Motzkin. The arithmetic-geometric inequality. *Inequalities (Proc. Sympos. Wright-Patterson Air Force Base)*, pages 205–224, 1967.
- [MPV87] Marc Mézard, Giorgio Parisi, and Miguel Virasoro. *Spin glass theory and beyond: an introduction to the replica method and its applications*. World Scientific Publishing Company, 1987.
- [MPW15] Raghu Meka, Aaron Potechin, and Avi Wigderson. Sum-of-squares lower bounds for planted clique. In *47th Annual ACM Symposium on Theory of Computing (STOC 2015)*, pages 87–96. ACM, 2015.
- [MR11] Cristopher Moore and Alexander Russell. A graph integral formulation of the circuit partition polynomial. *Combinatorics, Probability & Computing*, 20(6):911, 2011.
- [MR15] Andrea Montanari and Emile Richard. Non-negative principal component analysis: message passing algorithms and sharp asymptotics. *IEEE Transactions on Information Theory*, 62(3):1458–1484, 2015.

- [MRX20] Sidhanth Mohanty, Prasad Raghavendra, and Jeff Xu. Lifting sum-of-squares lower bounds: degree-2 to degree-4. In *52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2020)*, pages 840–853, 2020.
- [MRY18] Andrea Montanari, Feng Ruan, and Jun Yan. Adapting to unknown noise distribution in matrix denoising. *arXiv preprint arXiv:1810.02954*, 2018.
- [MRZ15] Andrea Montanari, Daniel Reichman, and Ofer Zeitouni. On the limitation of spectral methods: From the gaussian hidden clique problem to rank-one perturbations of gaussian tensors. In *Advances in Neural Information Processing Systems*, pages 217–225, 2015.
- [MS16] Andrea Montanari and Subhabrata Sen. Semidefinite programs on sparse random graphs and their application to community detection. In *48th Annual ACM Symposium on Theory of Computing (STOC 2016)*, pages 814–827. ACM, 2016.
- [MW13] Raghu Meka and Avi Wigderson. Association schemes, non-commutative polynomial concentration, and sum-of-squares lower bounds for planted clique. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 20, page 10, 2013.
- [MW15] Tengyu Ma and Avi Wigderson. Sum-of-squares lower bounds for sparse PCA. In *Advances in Neural Information Processing Systems*, pages 1612–1620, 2015.
- [MW19] Ankur Moitra and Alexander S Wein. Spectral methods from tensor networks. In *51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)*, pages 926–937, 2019.
- [Nes98] Yurii Nesterov. Semidefinite relaxation and nonconvex quadratic optimization. *Optimization Methods and Software*, 9(1-3):141–160, 1998.

- [Nes00] Yurii Nesterov. Squared functional systems and optimization problems. In *High performance optimization*, pages 405–440. Springer, 2000.
- [NP33] Jerzy Neyman and Egon Sharpe Pearson. IX. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231(694-706):289–337, 1933.
- [NRV13] Assaf Naor, Oded Regev, and Thomas Vidick. Efficient rounding for the noncommutative Grothendieck inequality. In *45th Annual ACM Symposium on Theory of Computing (STOC 2013)*, pages 71–80, 2013.
- [O'D17] Ryan O'Donnell. SOS is not obviously automatizable, even approximately. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- [OV96] Andrei Okounkov and Anatoly Vershik. A new approach to representation theory of symmetric groups. *Selecta Mathematica New Series*, 2(4):581–606, 1996.
- [Pan13] Dmitry Panchenko. *The Sherrington-Kirkpatrick model*. Springer Science & Business Media, 2013.
- [Pan18] Dmitry Panchenko. Free energy in the Potts spin glass. *The Annals of Probability*, 46(2):829–864, 2018.
- [Par79] Giorgio Parisi. Infinite number of order parameters for spin-glasses. *Physical Review Letters*, 43(23):1754, 1979.
- [Par80] Giorgio Parisi. A sequence of approximated solutions to the SK model for spin glasses. *Journal of Physics A: Mathematical and General*, 13(4):L115, 1980.

- [Par00] Pablo A Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, 2000.
- [Par03] Pablo A Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming*, 96(2):293–320, 2003.
- [Pat98] Gábor Pataki. On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues. *Mathematics of Operations Research*, 23(2):339–358, 1998.
- [Per96] Asher Peres. Separability criterion for density matrices. *Physical Review Letters*, 77(8):1413, 1996.
- [Pot17] Aaron Potechin. Sum of squares lower bounds from symmetry and a good story. *arXiv preprint arXiv:1711.11469*, 2017.
- [Put93] Mihai Putinar. Positive polynomials on compact semi-algebraic sets. *Indiana University Mathematics Journal*, 42(3):969–984, 1993.
- [PWB16] Amelia Perry, Alexander S Wein, and Afonso S Bandeira. Statistical limits of spiked tensor models. *arXiv preprint arXiv:1612.07728*, 2016.
- [PWBM18] Amelia Perry, Alexander S Wein, Afonso S Bandeira, and Ankur Moitra. Optimality and sub-optimality of PCA I: Spiked random matrix models. *Annals of Statistics*, 46(5):2416–2451, 2018.
- [PWZ96] Marko Petkovšek, Herbert S Wilf, and Doron Zeilberger. *A = B*. 1996.
- [Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *40th Annual ACM Symposium on Theory of Computing (STOC 2008)*, pages 245–254. ACM, 2008.

- [Raz01] Alexander A Razborov. Proof complexity of pigeonhole principles. In *International Conference on Developments in Language Theory*, pages 100–116. Springer, 2001.
- [Raz02] Ran Raz. Resolution lower bounds for the weak pigeon hole principle. In *17th IEEE Annual Conference on Computational Complexity*, page 3. IEEE, 2002.
- [Raz03] Alexander Razborov. Propositional proof complexity. *Journal of the ACM*, 50(1):80–82, 2003.
- [Rez95] Bruce Reznick. Uniform denominators in Hilbert’s seventeenth problem. *Mathematische Zeitschrift*, 220(1):75–97, 1995.
- [Rez96] Bruce Reznick. Homogeneous polynomial solutions to constant coefficient PDE’s. *Advances in Mathematics*, 117(2):179–192, 1996.
- [Rez00] Bruce Reznick. Some concrete aspects of Hilbert’s 17th problem. *Contemporary Mathematics*, 253:251–272, 2000.
- [RH17] Phillippe Rigollet and Jan-Christian Hütter. Lecture notes on high dimensional statistics. 2017.
- [RM14] Emile Richard and Andrea Montanari. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems*, pages 2897–2905, 2014.
- [Rom05] Steven Roman. *The umbral calculus*. Springer, 2005.
- [Ros10] Benjamin Rossman. *Average-case complexity of detecting cliques*. PhD thesis, Massachusetts Institute of Technology, 2010.
- [Ros14] Benjamin Rossman. The monotone complexity of k -clique on random graphs. *SIAM Journal on Computing*, 43(1):256–279, 2014.

- [Rot64] Gian-Carlo Rota. On the foundations of combinatorial theory: I. Theory of Möbius functions. *Probability Theory and Related Fields*, 2(4):340–368, 1964.
- [RRS17] Prasad Raghavendra, Satish Rao, and Tselil Schramm. Strongly refuting random CSPs below the spectral threshold. In *49th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2017)*, pages 121–131. ACM, 2017.
- [RS00] Gian-Carlo Rota and Jianhong Shen. On the combinatorics of cumulants. 2000.
- [RS15] Prasad Raghavendra and Tselil Schramm. Tight lower bounds for planted clique in the degree-4 SOS program. *arXiv preprint arXiv:1507.05136*, 2015.
- [RSS18] Prasad Raghavendra, Tselil Schramm, and David Steurer. High-dimensional estimation via sum-of-squares proofs. *arXiv preprint arXiv:1807.11419*, 2018.
- [Rud08] Mark Rudelson. Invertibility of random matrices: norm of the inverse. *Annals of Mathematics*, pages 575–600, 2008.
- [RV08] Mark Rudelson and Roman Vershynin. The Littlewood-Offord problem and invertibility of random matrices. *Advances in Mathematics*, 218(2):600–633, 2008.
- [RV13] Mark Rudelson and Roman Vershynin. Hanson-Wright inequality and sub-gaussian concentration. *Electronic Communications in Probability*, 18, 2013.
- [RW92] Robert W Robinson and Nicholas C Wormald. Almost all cubic graphs are Hamiltonian. *Random Structures & Algorithms*, 3(2):117–125, 1992.
- [RW94] Robert W Robinson and Nicholas C Wormald. Almost all regular graphs are Hamiltonian. *Random Structures & Algorithms*, 5(2):363–374, 1994.
- [RW17] Prasad Raghavendra and Benjamin Weitz. On the bit complexity of sum-of-squares proofs. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and

Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 80 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 80:1–80:13, Dagstuhl, Germany, 2017. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

- [SA90] Hanif D Sherali and Warren P Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3(3):411–430, 1990.
- [SBŻ06] Stanisław J Szarek, Ingemar Bengtsson, and Karol Życzkowski. On the structure of the body of states with positive partial transpose. *Journal of Physics A: Mathematical and General*, 39(5):L119, 2006.
- [Sch91] Konrad Schmüdgen. The k -moment problem for compact semi-algebraic sets. *Mathematische Annalen*, 289(1):203–206, 1991.
- [Sch05] Alexander Schrijver. On the history of combinatorial optimization (till 1960). *Handbooks in operations research and management science*, 12:1–68, 2005.
- [Sch08] Grant Schoenebeck. Linear level Lasserre lower bounds for certain k -CSPs. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 593–602. IEEE, 2008.
- [Sch17] Konrad Schmüdgen. *The moment problem*, volume 9. Springer, 2017.
- [Sei91] Johan Jacob Seidel. Strongly regular graphs with $(-1, 1, 0)$ adjacency matrix having eigenvalue 3. In *Geometry and Combinatorics*, pages 26–43. Elsevier, 1991.
- [Sen18] Subhabrata Sen. Optimization on sparse random hypergraphs and spin glasses. *Random Structures & Algorithms*, 53(3):504–536, 2018.

- [Sho87] Naum Zuselevich Shor. An approach to obtaining global extremums in polynomial mathematical programming problems. *Cybernetics*, 23(5):695–700, 1987.
- [Sin11] Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30(1):20–36, 2011.
- [SK75] David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Physical Review Letters*, 35(26):1792, 1975.
- [Sma83] Steve Smale. On the average number of steps of the simplex method of linear programming. *Mathematical Programming*, 27(3):241–262, 1983.
- [Sma98] Steve Smale. Mathematical problems for the next century. *The Mathematical Intelligencer*, 20(2):7–15, 1998.
- [SS94] Michael Shub and Steve Smale. On the intractability of Hilbert’s Nullstellensatz and an algebraic version of “NP = P?”. 1994.
- [ST04] Daniel A Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.
- [Sta78] Richard P Stanley. Exponential structures. *Studies in Applied Mathematics*, 59(1):73–82, 1978.
- [Sta10] Richard P Stanley. A survey of alternating permutations. *Contemporary Mathematics*, 531:165–196, 2010.
- [STDHJ07] Mátyás A Sustik, Joel A Tropp, Inderjit S Dhillon, and Robert W Heath Jr. On the existence of equiangular tight frames. *Linear Algebra and its Applications*, 426(2-3):619–635, 2007.

- [Ste74] Gilbert Stengle. A Nullstellensatz and a Positivstellensatz in semialgebraic geometry. *Mathematische Annalen*, 207(2):87-97, 1974.
- [Sub18] Eliran Subag. Following the ground-states of full-RSB spherical spin glasses. *arXiv preprint arXiv:1812.04588*, 2018.
- [SW07] Lingsheng Shi and Nicholas Wormald. Colouring random regular graphs. *Combinatorics, Probability and Computing*, 16(3):459-494, 2007.
- [SW20] Tselil Schramm and Alexander S Wein. Computational barriers to estimation from low-degree polynomials. *arXiv preprint arXiv:2008.02269*, 2020.
- [Syl76] James Joseph Sylvester. XXXVII. Note on spherical harmonics. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):291-307, 1876.
- [Syl76] Garrett Smith Sylvester. *Continuous-spin Ising ferromagnets*. PhD thesis, Massachusetts Institute of Technology, 1976.
- [Sze39] Gabor Szegő. *Orthogonal polynomials*. American Mathematical Society, 1939.
- [Tal06] Michel Talagrand. The Parisi formula. *Annals of Mathematics*, pages 221-263, 2006.
- [Tal10] Michel Talagrand. *Mean field models for spin glasses: Volume I: Basic examples*. Springer Science & Business Media, 2010.
- [TMR20] Paxton Turner, Raghu Meka, and Philippe Rigollet. Balancing Gaussian vectors in high dimension. In *33rd Annual Conference on Learning Theory (COLT 2020)*, pages 3455-3486, 2020.

- [Tre12a] Luca Trevisan. Max cut and the smallest eigenvalue. *SIAM Journal on Computing*, 41(6):1769–1786, 2012.
- [Tre12b] Luca Trevisan. On Khot’s Unique Games Conjecture. *Bulletin (New Series) of the American Mathematical Society*, 49(1), 2012.
- [Tul10] Madhur Tulsiani. Lovász-Schrijver reformulation. *Wiley Encyclopedia of Operations Research and Management Science*, 2010.
- [Tur37] Alan Mathison Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 2(1):230–265, 1937.
- [Veg00] Gert Vegter. The apolar bilinear form in geometric modeling. *Mathematics of Computation*, 69(230):691–720, 2000.
- [Ver11] Roman Vershynin. Spectral norm of products of random and deterministic matrices. *Probability Theory and Related Fields*, 150(3-4):471–509, 2011.
- [Ver18] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*. Cambridge University Press, 2018.
- [Ver20] Roman Vershynin. Concentration inequalities for random tensors. *Bernoulli*, 26(4):3139–3162, 2020.
- [Vil08] Cédric Villani. *Optimal transport: old and new*. Springer Science & Business Media, 2008.
- [Wal18] Shayne FD Waldron. *An introduction to finite tight frames*. Springer, 2018.
- [WBP16] Tengyao Wang, Quentin Berthet, and Yaniv Plan. Average-case hardness of RIP certification. In *Advances in Neural Information Processing Systems*, pages 3819–3827, 2016.

- [WBS16] Tengyao Wang, Quentin Berthet, and Richard J Samworth. Statistical and computational trade-offs in estimation of sparse principal components. *Annals of Statistics*, 44(5):1896–1930, 2016.
- [Wei20] Alexander S Wein. Optimal low-degree hardness of maximum independent set. *arXiv preprint arXiv:2010.06563*, 2020.
- [Wel74] Lloyd Welch. Lower bounds on the maximum cross correlation of signals (corresp.). *IEEE Transactions on Information Theory*, 20(3):397–399, 1974.
- [WEM19] Alexander S Wein, Ahmed El Alaoui, and Cristopher Moore. The Kikuchi hierarchy and tensor PCA. *arXiv preprint arXiv:1904.03858*, 2019.
- [Wig93] Eugene P Wigner. Characteristic vectors of bordered matrices with infinite dimensions I. In *The Collected Works of Eugene Paul Wigner*, pages 524–540. Springer, 1993.
- [Wor99] Nicholas C Wormald. Models of random regular graphs. *London Mathematical Society Lecture Note Series*, pages 239–298, 1999.
- [ZB10] Lenka Zdeborová and Stefan Boettcher. A conjecture on the maximum cut and bisection width in random regular graphs. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(02):P02020, 2010.
- [ZHT06] Hui Zou, Trevor Hastie, and Robert Tibshirani. Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15(2):265–286, 2006.
- [ZK11] Lenka Zdeborová and Florent Krzakala. Quiet planting in the locked constraint satisfaction problems. *SIAM Journal on Discrete Mathematics*, 25(2):750–770, 2011.

- [ZK16] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.
- [Zwi98] Uri Zwick. Approximation algorithms for constraint satisfaction problems involving at most three variables per constraint. In *9th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 1998)*, pages 201–210, 1998.