

DISSERTATION

NEW METHODS FOR FIXED-MARGIN BINARY MATRIX SAMPLING, FRÉCHET  
COVARIANCE, AND MANOVA TESTS FOR RANDOM OBJECTS IN MULTIPLE METRIC  
SPACES

Submitted by

Alex M. Fout

Department of Statistics

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Summer 2022

Doctoral Committee:

Advisor: Bailey Fosdick

Andee Kaplan

Daniel Cooley

Henry Adams

Copyright by Alex M. Fout 2022

All Rights reserved

## ABSTRACT

### NEW METHODS FOR FIXED-MARGIN BINARY MATRIX SAMPLING, FRÉCHET COVARIANCE, AND MANOVA TESTS FOR RANDOM OBJECTS IN MULTIPLE METRIC SPACES

Many approaches to the analysis of network data essentially view the data as Euclidean and apply standard multivariate techniques. In this dissertation, we refrain from this approach, exploring two alternate approaches to the analysis of networks and other structured data. The first approach seeks to determine how unique an observed simple, directed network is by comparing it to like networks which share its degree distribution. Generating networks for comparison requires sampling from the space of all binary matrices with the prescribed row and column margins, since enumeration of all such matrices is often infeasible for even moderately sized networks with 20-50 nodes. We propose two new sampling methods for this problem. First, we extend two Markov chain Monte Carlo methods to sample from the space non-uniformly, allowing flexibility in the case that some networks are more likely than others. We show that non-uniform sampling could impede the MCMC process, but in certain special cases is still valid. Critically, we illustrate the differential conclusions that could be drawn from uniform vs. nonuniform sampling. Second, we develop a generalized divide and conquer approach which recursively divides matrices into smaller subproblems which are much easier to count and sample. Each division step reveals interesting mathematics involving the enumeration of integer partitions and points in convex lattice polytopes. The second broad approach we explore is comparing random objects in metric spaces lacking a coordinate system. Traditional definitions of the mean and variance no longer apply, and standard statistical tests have needed reconceptualization in terms of only distances in the metric space. We consider the multivariate setting where random objects exist in multiple metric spaces, which can be thought of as distinct views of the random object. We define the notion of Fréchet

covariance to measure dependence between two metric spaces, and establish consistency for the sample estimator. We then propose several tests for differences in means and covariance matrices among two or more groups in multiple metric spaces, and compare their performance on scenarios involving random probability distributions and networks with node covariates.

## ACKNOWLEDGEMENTS

A well known corollary to the Law of Large Numbers states: It takes a village to raise a statistician. I certainly have my own village to thank, and the significance of your contributions to my success became even more clear to me as I choked up on the thank you slide at the end of my defense. Here is a regrettably partial list of people to whom I owe thanks. To Elisa, Mom, Dad, Kris, and Katherine for your unwavering love and support. To Elaine, Reed, Charlie, Susan, Declan, and Edwin for adopting me so readily into your homes. To my extended family and friends, for your flexibility and understanding of my limited availability at times. To Bailey for your constant support, encouragement, and patience with me as I navigated the program, my research, and considered next steps after graduation. To my committee for your insightful questions and kind feedback on my dissertation, which have strengthened it considerably. To my colleagues, who inspire me with your curiosity, creativity, tenacity, and humility in both scholarship and teaching. To my instructors, who have challenged and encouraged me, even in my non-traditional academic journey. To my students, the teaching of whom was a welcome reprieve from the rigors of grad school, and who helped me rekindle my passion for teaching. Thank you all, I could not have done this without you!

This work was partially supported by grant IOS-1856229 from the National Science Foundation.

## TABLE OF CONTENTS

ABSTRACT . . . . .	ii
ACKNOWLEDGEMENTS . . . . .	iv
LIST OF TABLES . . . . .	viii
LIST OF FIGURES . . . . .	ix
Chapter 1    Introduction . . . . .	1
1.1        Preamble . . . . .	1
1.2        Random Sampling of Binary Matrices with Fixed Margins . . . . .	2
1.2.1    Motivating Examples . . . . .	2
1.2.2    Problem Variants . . . . .	5
1.2.3    Brief History of Sampling Methods . . . . .	6
1.2.4    Our Contribution . . . . .	9
1.3        Statistical Analysis of Random Objects in Metric Spaces . . . . .	9
1.3.1    Statistics In Metric Spaces . . . . .	10
1.3.2    Our Contribution . . . . .	10
Chapter 2    Mathematical Preliminaries . . . . .	12
2.1        Binary Matrices and Margins . . . . .	12
2.2        Integer Partitions and Compositions . . . . .	12
2.3        Majorization & The Gale-Ryser Theorem . . . . .	13
2.4        Navigating the Young’s Lattice with Box Shifts . . . . .	14
2.5        The Shift-Right Algorithm . . . . .	15
2.6        Metric Spaces . . . . .	17
2.7        Riemannian Manifolds . . . . .	18
2.8        Matrix Logarithm . . . . .	21
Chapter 3    Non-Uniform Sampling of Fixed Margin Binary Matrices . . . . .	23
3.1        Introduction . . . . .	23
3.1.1    Review of Sampling Methods . . . . .	24
3.1.2    Non-Uniform Sampling . . . . .	27
3.2        Non-Uniform Sampling Of Binary Matrices . . . . .	28
3.2.1    Probability Model . . . . .	28
3.2.2    Weighted Checkerboard Swap Algorithm . . . . .	29
3.2.3    Structural Zeros . . . . .	31
3.2.4    Weighted Curveball Algorithm . . . . .	33
3.3        Simulations . . . . .	33
3.3.1    Small Example . . . . .	35
3.3.2    Mixing Performance . . . . .	36
3.3.3    Effect of Heterogeneous Weights . . . . .	38
3.4        Example: New Hebrides Bird Species . . . . .	41
3.5        Conclusions . . . . .	44

Chapter 4	A Divide and Conquer Algorithm for Exact Uniform Sampling of Matrices with Fixed Margins . . . . .	47
4.1	Introduction . . . . .	47
4.2	Miller & Harrison’s Exact Sampling . . . . .	48
4.3	Divide and Conquer Counting . . . . .	50
4.3.1	Margin Spreading . . . . .	51
4.3.2	Base Cases . . . . .	60
4.3.3	Algorithm . . . . .	62
4.4	Divide and Conquer Sampling . . . . .	64
4.4.1	Algorithm . . . . .	65
4.5	Experiments . . . . .	65
4.5.1	Example Counts . . . . .	65
4.5.2	Example Sampling . . . . .	67
4.5.3	Lookup Table Size . . . . .	67
4.5.4	Cut Point Location . . . . .	70
4.5.5	Division Strategy . . . . .	74
4.5.6	Connection to Miller & Harrison’s Approach . . . . .	76
4.6	Discussion . . . . .	76
4.6.1	Computational Complexity . . . . .	76
4.6.2	Connection to symmetric polynomials . . . . .	77
4.7	Conclusions . . . . .	78
Chapter 5	Fréchet Covariance and MANOVA Tests for Random Objects in Multiple Metric Spaces . . . . .	79
5.1	Introduction . . . . .	79
5.2	Analysis of Variance in Bounded Metric Spaces . . . . .	80
5.3	Fréchet Covariance . . . . .	84
5.4	Properties of Fréchet Covariance . . . . .	88
5.5	Tests for Differences in Means and Variance / Covariance Matrices . . . . .	90
5.5.1	Adaptations of Classical MANOVA Statistics . . . . .	91
5.5.2	Statistics Based on Riemannian Geometry . . . . .	93
5.5.3	Statistics Based on the Fréchet ANOVA . . . . .	95
5.5.4	Testing the Null Hypothesis . . . . .	96
5.6	Simulation Studies . . . . .	98
5.6.1	Scenario 1: Normal Distributions in Two Metric Spaces, Two Groups . . . . .	99
5.6.2	Scenario 2: Random Networks with Node Covariates, Two Groups . . . . .	103
5.6.3	Type I Error Rate . . . . .	106
5.6.4	Effect of Sample Size and Group Balance . . . . .	107
5.7	Conclusions . . . . .	109
Chapter 6	Conclusion . . . . .	111
Bibliography	. . . . .	114
Appendix A	Proofs . . . . .	125

A.1	Nonuniform Sampling of Fixed Margin Binary Matrices . . . . .	125
A.2	Divide and Conquer Sampling of Fixed Margin Binary Matrices . . . . .	130
A.3	Fréchet Covariance and MANOVA Tests for Random Objects in Multiple Metric Spaces . . . . .	132
Appendix B	Enumerating Valid Spreads: The Naïve Approach . . . . .	134

## LIST OF TABLES

3.1	Theoretical vs. empirical probabilities for matrices in Figure 3.5 with binary matrix and weight matrix as defined in (3.6). Empirical probabilities also show Tukey-Hanning [1] MCMC standard errors. . . . .	36
4.1	Example margins used for examining cut point and cut strategy. Margins were sorted as displayed when feeding to the algorithm. . . . .	72
5.1	Summary of simulation parameters for groups 1 and 2 in each study of scenario 1. Bold sections highlight the differences between groups in each study. . . . .	101
5.2	Summary of simulation parameters for Groups 1 and 2 in each study of scenario 2. Bold sections highlight the differences between groups in each study. . . . .	104
5.3	Type I error rates for each test in the studies of scenarios 1 and 2. The Type I error is not a function of the alternative hypothesis so it is the same across all studies within each scenario. We simulated 6,000 data sets under each scenario, giving a standard error of approximately 0.003 for each estimate, and used 1,000 permutations for each permutation test. . . . .	107

## LIST OF FIGURES

2.1	Young diagrams for $(3, 2, 1)$ and $(2, 2, 1, 1)$ , both partitions of 6. . . . .	13
2.2	Young diagrams for two partitions of 9, $x = (3, 3, 2, 1)$ (left) and its conjugate $x' = (4, 3, 2)$ (right). The Young diagrams are diagonal reflections of one another. . . . .	13
2.3	Example of how partitions with “wider” Young diagrams (left) tend to majorize partitions with “taller” Young diagrams. . . . .	14
2.4	Example of a $v$ shift followed by an $h$ shift in a Young diagram, as described in [2]. These shifts, and their reverse equivalents, enable full navigation of the Young lattice for a particular integer $\lambda$ . . . . .	15
2.5	Initial fill of the shift-right algorithm (a), which satisfies the row margins $p$ , along with the Young diagram for $p$ (b). . . . .	16
2.6	(a)-(d) show the shifting steps of the shift-right algorithm, in order, which produce a valid binary matrix. Red arrows denote the moving of a 1 from the tail to the head. The target column whose margin is being satisfied is marked with $j$ , and the source column from which 1s are being shifted is marked with $j^*$ . . . . .	17
3.1	A $2 \times 2$ submatrix of $A$ from rows $i$ and $i'$ , and columns $j$ and $j'$ . which are not necessarily consecutive. Swapping the positions of ones and zeros preserves the row and column sums of $A$ . . . . .	29
3.2	Matrices $A$ and $B$ sharing the same margins and structural zeros, with structural zeros shaded grey. The difference $A - B$ has a single elementary circuit and the hamming distance between $A$ and $B$ is 6, so Brualdi’s method [3] gives a maximum of 2 swaps to transition between $A$ and $B$ . While this is true in the absence of structural zeros, this example demonstrates that structural zeros can cause the bound to be violated, as seven swaps are required to transition between $A$ and $B$ while honoring the structural zeros. . . . .	31
3.3	Matrices $A$ and $B$ sharing the same margins and structural zeros, with structural zeros shaded grey. No sequence of checkerboard swaps which honor the structural zeros can transition from $A$ to $B$ , demonstrating that structural zeros can render the Markov state space is reducible. . . . .	32
3.4	Example binary matrix $A$ and weight matrix $W$ . There are six binary matrices with the same row and column margins as $A$ , shown in Figure 3.5 . . . . .	35
3.5	All $3 \times 3$ matrices with row and column sums equal to one. . . . .	36
3.6	Average effective sample size for the simulations described in Section 3.3.2, for varying weighting schemes and percentage of structural zeros, where structural zeros are either randomly placed (blue), or arranged monotonically (red). Each line represents the average ESS from ten separate chains using the same starting matrix. . . . .	39
3.7	Detailed view of effective sample size for the simulations described in Section 3.3.2, using a weight matrix of all ones and with 50% structural zeros. Average ESS is shown in bold, with each iteration shown partially transparent. . . . .	40

3.8	Initial binary matrix $A$ and base weight matrix $W$ for the simulation described in Section 3.3.3. Both use a shading scale where white denotes zero and black denotes one. The minimum value in $W$ is 0.5, achieved in the top-right and bottom-left most elements. Note the unfavorable state of $A$ , since it contains ones corresponding to lower weighted regions of $W$ . . . . .	40
3.9	Sampling distribution of the diagonal divergence statistic, under different weighting matrices, defined by raising the base weight matrix to the power $p$ . . . . .	41
3.10	Observed New Hebrides bird species combination $A$ [4], and non-uniform weight matrix $W$ for the data analysis described in Section 3.4. For $A$ , white denotes zero and grey denotes one. For $W$ , light grey denotes one and grey denotes 15. With higher values indicating species (rows) preference for certain island (columns) habitats. . . .	43
3.11	Sampling distribution of the C-score statistic, under uniform (red) and non-uniform (blue) weights. The empirical p-value under non-uniform weights is no longer significant, demonstrating how incorporating species habitat preference could impact whether the observed combination is relatively unlikely. . . . .	44
4.1	Creating an initial margin spread. The initial margins $p$ and $q$ are used by Fulkerson & Ryser's <i>shift-right</i> algorithm [5] to produce an initial matrix (left left). The matrix is then divided at the division point (determined by $k = 3$ in this example), and the initial margin spread, $p^l$ and $p^r$ , are just the row margins of the left and right parts (right).	52
4.2	Using the Young diagram for $p = (2, 2, 1, 1)$ to visualize the margin spread $(p^l, p^r) = ((1, 1, 1, 1), (1, 1, 0, 0))$ . The thick border line separates the boxes belonging to $p^l$ on the left from those belonging to $p^r$ on the right. . . . .	52
4.3	Creating a new margin spread from an existing one by moving the border line in the Young diagram for $p$ . The initial border corresponding to $(p^l, p^r) = ((1, 1, 1, 1), (1, 1, 0, 0))$ (left) is moved according to the arrows to produce a new border line (right) without changing number of boxes to its left and right. The shifted border corresponds to the new margin spread $(p^l, p^r) = ((2, 1, 1, 0), (0, 1, 0, 1))$ . . . . .	53
4.4	Creating a new margin spread from an existing one by moving the border line in the Young diagram for $p$ . The initial border corresponding to $(p^l, p^r) = ((2, 1, 1, 0), (0, 1, 0, 1))$ (left) is moved according to the arrows to produce a new border line (right) without changing number of boxes to its left and right. The shifted border corresponds to the new margin spread $(p^l, p^r) = ((2, 2, 0, 0), (0, 0, 1, 1))$ . . . . .	53
4.5	Left: the row margin $p = (2, 2, 1, 1)$ contains as elements either two or one, so there are two blocks. The indices of block 1 are $s_1 = \{3, 4\}$ , and the indices of block 2 are $s_2 = \{1, 2\}$ . Right: The left fills $f_1^l$ and $f_2^l$ correspond to the number of boxes in blocks 1 and 2 respectively that are left of the border line in the Young diagram, so $f_1^l = f_2^l = 2$ . Similarly, the right fills $f_1^r$ and $f_2^r$ are the number of boxes in blocks 1 and 2 respectively that are right of the border line in the Young diagram, so $f_1^r = 0$ and $f_2^r = 2$ . $f_1^r$ can also be deduced from $f_1^l$ , since the total number of boxes in this block is $1 \times  s_1  = 2$ , so we only need to keep track of $f_1^l$ . . . . .	54
4.6	The only three left margins $(p^l)$ which are consistent with the fill $f^l = (2, 2)$ . Each margin is shown alongside the Young diagram for $p$ with a thick border separating the boxes belonging to $p^l$ on the left and $p^r$ on the right. . . . .	55

4.7 Finding a new margin spread candidate for the fill  $f^l = (2, 2)$ . An initial margin spread  $p^l = (1, 1, 1, 1)$ , in this case produced from the *shift-right* canonical matrix, is split into individual blocks. The top block is selected for “elevation”, meaning the left partition  $(1,1)$  is transformed into the partition  $(2,0)$  which is higher in the Young’s lattice partial order. By elevating only within this block, we maintain the same fill,  $f^l = (2, 2)$ . The blocks are recombined resulting in the left margin spread  $p^l = (2, 0, 1, 1)$ . The new margin will always majorize the starting margin, so we must check that the Gale-Ryser condition is still satisfied (for both  $p^l$  and  $p^r$ ). Repeating this process systematically to each block for each new margin until margins are no longer graphical allows every valid margin within fill  $f^l$  to be enumerated. . . . . 59

4.8 Beginning with the margins  $p = q = (1, 1, 1, 1)$ , the matrix is divided first by columns, with the only valid margin spread (up to permutation of the rows) being  $(p^l, p^r) = ((1, 1, 0, 0), (0, 0, 1, 1))$ . Each smaller matrix is then divided by rows, resulting in four  $2 \times 2$  matrices with regular margins, either all one or all zero. The zero-regular margin matrices are at a base case, with only one way of satisfying the margins (all zeros). The one-regular margin matrices are not yet at a base case, so more division is required, however this process only has to occur once, since both one margin matrices are identical and have the same count. . . . . 61

4.9 If row or column margins have value zero at any point in the counting process, those rows or columns can be excluded from subsequent steps. . . . . 61

4.10 We repeatedly sample from matrices with margins  $p = (2, 2, 2, 1)$  and  $q = (3, 2, 1, 1)$  using the divide and conquer sampling approach, drawing samples of size 100, 300, and 1000. The panels show the distribution of 100,000 p-values from a  $\chi^2$  lack of fit test against the uniform distribution for each sample size. P-values are estimated using Monte-Carlo resampling [6] with 10,000 Monte-Carlo replicates per sample. The p-values should follow a uniform distribution if sampling is uniform. . . . . 68

4.11 Number of unique counts stored in lookup table for all graphical margins with sums up to 22. There are roughly 1.02 million graphical margins, representing roughly  $7.4 \times 10^{21}$  unique matrices. Lighter colors indicate higher density matrices, i.e. a higher percentage of ones. Note the log scale on the horizontal axis. . . . . 69

4.12 . . . . . 71

4.13 Measurements from the counting process for examples A-D from Table 4.1: A: $p = q = (4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ ; B: $p = q = (1, \dots, 1)$  (20 ones); C: $p = (1, \dots, 1)$  (45 ones),  $q = (5, 5, 5, 4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ ; D: $p = q = (5, 5, 5, 4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ . Three metrics are shown for varying cut points: number of spreads in the first division (red), size of the lookup table (green), and median running time out of three trials (blue). Note the different vertical scales on each plot. The rapid jumps between 0.000 and 0.001 seconds for example B are a result of limited time resolution. . . . . 73

4.14 Measurements from the counting process for examples A-D from Table 4.1: A: $p = q = (4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ ; B: $p = q = (1, \dots, 1)$  (20 ones); C: $p = (1, \dots, 1)$  (45 ones),  $q = (5, 5, 5, 4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ ; D: $p = q = (5, 5, 5, 4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ . We compare four division strategies: always divide the longer margin (Long), always divide the shorter margin (Short), always alternate row/column divisions (Alternate), always divide the row margin (Row). Note the different scales on each plot, and the fine vertical time scale for example B. . . . . 75

5.1	Non-centered ( $\hat{\rho}_{12}$ ) and centered ( $\hat{\rho}'_{12}$ ) Fréchet correlation for 5 Euclidean bivariate scenarios. Scenarios 1 and 2 illustrate how both definitions show strong positive correlation when Euclidean correlation is either strongly positive or strongly negative. Scenarios 3 and 4 highlight how the centered Fréchet correlation is negative when $X_1$ and $X_2$ tend to move in opposite directions relative to their means. Scenario 5 shows that the centered Fréchet correlation near zero when $X_1$ and $X_2$ are independent. . . .	87
5.2	Data examples for the first four studies in scenario 1 . . . . .	100
5.3	Results from each simulation study in scenario 1. In each plot, a vertical black line indicates the null hypothesis, and a horizontal black line indicates the desired Type I error rate, $\alpha = 0.05$ . . . . .	102
5.4	Results from each simulation study in scenario 2. In each plot, a vertical black line indicates the null hypothesis, and a horizontal black line indicates the desired Type I error rate, $\alpha = 0.05$ . . . . .	105
5.5	Results of study 5 in scenario 1 at three different sample sizes (equal for each group). In each plot, a vertical black line indicates the null hypothesis, and a horizontal black line indicates the desired Type I error rate, $\alpha = 0.05$ . . . . .	108
5.6	Results of study 5 in scenario 1 with varying balance (unequal samples sizes). In each plot, a vertical black line indicates the null hypothesis, and a horizontal black line indicates the desired Type I error rate, $\alpha = 0.05$ . . . . .	108
A.1	Illustration of irreducibility proof Theorem 3. $A$ and $B$ have structural zeros (in black), and agree for elements in the hashed region. Their partials sums must agree in the red regions, and their partial sums must agree in the blue regions as well. Four cases are possible for the orange element. . . . .	130

# Chapter 1

## Introduction

### 1.1 Preamble

In this dissertation we develop methods for two statistical problems: determining how noteworthy is an observed binary matrix, and comparing groups of random objects observed in multiple metric spaces. These seemingly distinct problems have a unifying theme: the analysis of networks. Binary matrices are a popular representation of simple, directed networks, or equivalently simple, undirected bipartite networks. On the other hand, if a set of networks is imbued with a dissimilarity metric, then they can be viewed as random objects living in a metric space. These two views, though being applied to the same class of objects, lead to vastly different approaches to analysis.

While every binary matrix can be viewed as a network, not all networks can be represented as a binary matrix (e.g. networks with multiple or weighted edges). However, every collection of networks can be viewed as objects in a metric space, but certainly not every metric space object is a network. In this sense, the methods we develop for binary matrices (Chapters 3 and 4) are more limited in their application compared to the methods we develop for random objects in metric spaces (Chapter 5), which can be applied to any class of networks and extends well beyond networks to any objects living in metric spaces.

With this framing in mind, we assert that both the binary matrix and random object views of networks are actively taken by practitioners, motivating continued research in both. The remainder of this chapter introduces in more detail the problems of sampling binary matrices with fixed margins and the analysis of random objects in metric spaces, providing motivation and recent developments for both.

## 1.2 Random Sampling of Binary Matrices with Fixed Margins

A binary matrix has elements valued either zero or one, and we call the set of sums for each row and column the *marginal sums* or *margins*. In general, several binary matrices may share the same set of margins, and in many cases, we may be interested in sampling from the set of all matrices satisfying a given set of margins in order to perform a data analysis.

### 1.2.1 Motivating Examples

Data in the form of binary matrices occur in a variety of disciplines. The matrices are typically used to capture dyadic relationships between individuals, locations, or events in the rows and columns. We begin by highlighting three scenarios which give rise to a binary matrix and to the need to them with fixed margins.

#### Species Co-Occurrence

In 1975, Diamond [4] proposed rules which govern the combination of several bird communities on islands in New Guinea, such as “Some pairs of species never coexist, either by themselves or as a part of a larger combination”. However, Connor & Simberloff [7] disputed these rules, suggesting that the combination of species on the islands was not exceptionally different from random. Aside from the ensuing feud within the ecological community [8, 9, 10, 11, 12, 13, 14], Connor & Simberloff sparked interest in a novel form of “null model analysis” [15, 16], a general approach which compares an observed matrix statistic to the distribution of that statistic under a null model. The null model of choice in this case was the set of all random combinations of species on islands satisfying the following three constraints:

- The number of species on each island matches the number observed.
- The number of islands inhabited by each species matches the number observed.
- Each species inhabits islands of comparable size to the islands on which it is observed.

Connor & Simberloff reason that if the observed species combination is not substantially different from the random combinations with respect to some statistic of interest, then it becomes difficult to argue that the observed pattern is anything noteworthy.

To generate combinations which adhere to the above constraints, Connor & Simberloff represent the species / island combinations as a matrix, with rows representing species and columns representing islands. The element in a particular row and column is set to one if the corresponding species inhabits the corresponding island, and zero otherwise. In the matrix representation, the first two constraints imply fixing the marginal sums of the matrix, while the third constraint limits which elements are permitted to be one. Thus, the first step in the “null model” analysis is to generate a collection of matrices which possess the desired marginal sums and exclude the prohibited ones.

### **Rasch Item Response Modeling**

In 1960, Georg Rasch proposed a new “latent trait” model to explain the responses of a group of individuals on a set of test questions (or “items”) [17, 9]. In this new *Rasch model*, each individual  $i$  is assigned a latent ability parameter  $\alpha_i$ , and each test question  $j$  is assigned a difficulty  $\beta_j$ . The model assigns a probability that individual  $i$  correctly answers question  $j$  as

$$p_{ij} = \frac{e^{\alpha_i - \beta_j}}{1 + e^{\alpha_i + \beta_j}}. \quad (1.1)$$

The results of an experiment in which each individual attempts each test question can be arranged in a binary matrix, where each row represents an individual, each column represents a test question, and an element is one if the individual responded correctly or zero if the individual responded incorrectly.

Maximum likelihood methods exist for estimating the ability and difficulty parameters, and various statistical tests have been devised to assess lack of fit [18, 19]. One class of tests relies on the fact that the row and column sums of the observed matrix are sufficient statistics for the ability and difficulty parameters, and furthermore each binary matrix satisfying the given marginals is equally

likely. Thus given a statistic of interest, a null hypothesis can be tested by sampling uniformly from all matrices satisfying the margins, and deriving a rejection region from the sample [20].

## Networks

A network is a collection of nodes representing entities, along with a set of edges describing the relationships between pairs of entities. Mathematically, nodes can be represented as a set  $V = v_{i=1}^n$ , and edges as ordered or unordered pairs  $(v_i, v_{i'}) \in V \times V$ , representing directed or undirected relationships. A network can also be represented with a binary matrix, where a one in the matrix represents an edge between the row node and the column node, and a zero represents the lack of an edge. If the matrix is square, the rows and columns can be interpreted as describing the same set of nodes, and the non-zero entries can be interpreted as undirected edges, if the matrix is symmetric, or directed edges, in the general case. In this case the matrix is the adjacency matrix of the network. If the matrix is non-square, then the rows and columns necessarily represent two distinct sets of nodes. In this case the matrix is the bi-adjacency matrix of a bipartite network.

In the networks context, global statistics, such as degree assortativity, are commonly used to characterize the structural characteristics of a given network [21]. Assortativity quantifies the similarity in degree between connected node pairs, and is given by

$$r = \frac{\frac{1}{|E|} \sum_{(i,i') \in E} k_i k_{i'} - \bar{k}^2}{s_k^2}, \quad (1.2)$$

where  $|E|$  is the number of edges,  $k_i$  is the degree of node  $i$ , and  $\bar{k}$  and  $s_k^2$  are the observed mean and variance of the node degree, respectively. However, the interpretation of assortativity is challenging because of its complex dependence on the network's degree distribution. Hence it is common to compare the observed assortativity to its distribution under a uniform sample of networks having the same degree distribution, which is identical to the task of sampling binary matrices with fixed margins.

In any of the above examples, we may wish that our sample satisfies statistical properties, for example that each possible matrix has equal probability of being included in the sample. The

discrete matrix elements, the margin constraints, and the statistical requirements collectively make the sampling process challenging, and this challenge has led to a diversity of sampling approaches. We introduce new sampling methods for this problem which expand the set of circumstances under which sampling can be performed, and in keeping with the family of existing approaches, our methods are similarly diverse.

### 1.2.2 Problem Variants

The binary matrix sampling problem has some variants which naturally arise in practice, and are briefly listed here.

1. *Nonuniform Sampling*. In some scenarios, it may be desirable to consider some matrices more important or more likely than others, so they have a higher probability of inclusion in the sample. For example, Harrison & Miller [22] describe an extension to the Rasch latent trait model which incorporates item-specific covariate effects. When conditioning on the margins, the distribution is no longer uniform, but instead proportional to the product of weights from a weight matrix. Matrices with ones in locations corresponding to higher weight are more likely. Statistical testing under this model must take into account the non-uniform conditional distribution when sampling a null distribution.
2. *Structural Zeros*. Depending on the data generating process, it may be impossible for specific entries to be one. To reflect this in the analysis, we may wish to generate matrices while enforcing certain elements to be zero. This variant is a special case of non-uniform sampling, because some matrices are assigned probability zero. Specifically, some elements in the matrix are not allowed to be one and are denoted *structural zeros*. This can arise in network modeling if self-edges are not allowed, causing the diagonal of the adjacency matrix to be structurally zero. Structural zeros are also imposed by the third modeling constraint outlined by Connor & Simberloff when modeling species co-occurrence.

A related problem to sampling binary matrices is the sampling of positive integer valued matrices with specified margins, motivated from the analysis of contingency tables [22, 23]. The

problem is similar enough that methods that work for binary matrices can sometimes be adapted to work for positive integer valued matrices [24, 22]. We do not consider integer valued matrices in this work.

### 1.2.3 Brief History of Sampling Methods

Early interest in sampling binary matrices with fixed margins came from domain scientists, and most early attempts lacked proof that they sample uniformly from the space of valid matrices. These early attempts, however, inspired more carefully crafted algorithms which possessed nicer theoretical properties. Generally, matrix sampling occurs either by starting with an matrix of all zeros and incrementally filling in ones to satisfy the margins, or by starting with a matrix which satisfies the given margins and repeatedly altering it while preserving the desired margins. We call this first set of approaches *fill methods* and second set *swap methods*. Below we describe some swap and fill methods which have arisen in the literature.

#### Early Heuristics

Connor & Simberloff [7] proposed both a fill and swap method when analyzing species co-occurrence. The fill method involved sampling one row at a time, where sampling a row involved randomly selecting from the eligible columns the correct number of elements to be set to one. The column margins would then be updated and the sampling could proceed with the next row. Unfortunately, this method would fail to complete successfully if, for example, a species would be forced to inhabit the same island twice to satisfy the row and column margins, a so called “dead end”. In light of this, they also use a swap method which “scans” the matrix for submatrices involving a checkerboard (two opposing corners are one and the other corners are zero) and swapping the zeros to ones and the ones to zeros, thus leaving the matrix altered but the margins unchanged:

$$\begin{array}{|c|c|} \hline 1 & 0 \\ \hline 0 & 1 \\ \hline \end{array} \longrightarrow \begin{array}{|c|c|} \hline 0 & 1 \\ \hline 1 & 0 \\ \hline \end{array}$$

There is some ambiguity in Connor & Simberloff's description of both methods, so the exact properties of their sampling are unclear. Nevertheless these two attempts were the seeds of many methods to come. Subsequent authors have proposed additional methods for generating the desired matrices, which are generally either a fill [25, 10, 12, 14] or swap [9, 26, 10, 11, 27] type method. Focus eventually shifted from whether to use these approaches to which method of producing matrices is best [12]. This accompanied theoretical work in binary matrices with fixed marginal sums [3] and Markov chain Monte Carlo (MCMC) sampling [9], paving the way for swap and fill methods with proven statistical properties [28, 29].

### **Modern Swap Methods**

At their heart, swap methods repeatedly perform constraint preserving operations which enable a random walk in the space of valid binary matrices. The simplest and earliest form of operations is the *checkerboard swap* described above, which involves identifying a  $2 \times 2$  checkerboard (not necessarily congruous) and swapping the ones and zeros. If checkerboards are randomly chosen and swaps are performed with some probability less than unity, then the Markov chain comprised of repeated swaps will sample from the desired space uniformly. To be a valid method of sampling, it must be true that the Markov chain is irreducible and gives rise to a steady state corresponding to the uniform distribution. Fortunately, Brualdi established the former [3] and Kannan the latter [28]. As for the initial state, margins are often derived from some observed matrix which is valid by definition. However, if no such matrix is available, then one can be constructed using the *shift right* algorithm of Fulkerson & Ryser [5]. An obvious advantage of swap methods is the preserved satisfaction of the constraints, since the initial state is valid and the swaps leave margins unchanged. Further, the stationary state of the underlying Markov chain is the uniform distribution, so MCMC samples will approach the uniform distribution for long enough chains. A disadvantage is that for reasonably large matrices, checkerboard swaps are small local changes in the matrix, resulting in highly correlated samples. To address this problem, alternative swapping mechanisms involving more than four elements have been proposed [30, 31], which help reduce the sample correlation.

## Modern Fill Methods

Fill methods construct each matrix *de novo*, so there is no correlation between subsequent random matrices. In addition, whereas early methods sometimes encountered dead ends [7, 13] that required backtracking or restarting, modern approaches are able to guarantee a valid matrix every time. This has been made possible by one of two means: importance sampling or recursive counts.

One way around the dead end problem of fill methods is to forfeit sampling from the desired probability distribution in order to guarantee margin satisfaction. Chen & Small propose filling the matrix by sampling one column (or row) at a time using the conditional Poisson distribution [32], which happens to be the marginal distribution of a single column. To prevent dead ends, each step of the sampling algorithm is accompanied by the checking of conditions which are necessary and sufficient for the existence of a valid sampled matrix. This sampling scheme is guaranteed to produce a valid matrix, however the resulting probability distribution over valid matrices is not uniform. Therefore importance weights are used to adjust relevant estimators for the uniform distribution and estimate standard errors. In a follow up paper, Chen [24] showed how to adapt the sequential importance sampling approach to scenarios involving certain arrangements of structural zeros without backtracking by checking additional conditions at each sampling step.

Inspired by the column-by-column approach of Chen, Harrison & Miller [22, 33] proposed a fill method which also sampled one row (or column) at a time, but did so such that the resulting distribution is guaranteed to be uniform. Their algorithm involves an initial counting phase where for each possible choice of the first row, the number of valid matrices consistent with the choice is computed. This counting phase heavily leverages symmetries associated with the choice of row as well as with the counting of valid matrices, so that it is not necessary to enumerate all valid matrices. In the sampling phase, each row is sampled by selecting from the possible choices inversely proportional to the count associated with each. This approach enables fast (for moderate sized matrices) exact sampling from the uniform distribution.

## 1.2.4 Our Contribution

In this dissertation, we propose two novel approaches for sampling binary matrices with fixed margins, one of them is a swap method, the other is a fill method. In Chapter 3, we develop an adaptation to traditional swap methods in order to sample from the non-uniform model introduced by Harrison & Miller [22], involving a matrix of weights defining the relative likelihood of the corresponding elements in the binary matrix. We apply our adaptation to two existing MCMC swap methods and prove that it samples from the desired non-uniform distribution. We then address the natural consequence of non-uniform weights: structural zeros and matrices with probability zero. We prove that in special cases, the Markov space is still connected and we can still sample correctly. We investigate sampling efficiency under various configurations of weights and structural zeros, and provide an example data analysis utilizing our method.

In Chapter 4, we generalize the the divide and conquer fill method of Miller & Harrison [33]. Instead of dividing the matrix into the first row and the rest of the matrix as with Miller & Harrison's method, we allow dividing the matrix at an arbitrary point, resulting to two smaller matrices with compatible margins. Our method repeats the division process recursively until each resulting matrix is so small that sampling is trivially easy. In developing the algorithm, we solve the problem of enumerating all possible divided matrices, which amounts to enumerating points in a convex integer polytope, and enumerating the elements of an integer partition partial ordering associated with each point in the polytope. We prove the correctness of our algorithm, and investigate its runtime and space requirements through simulations.

## 1.3 Statistical Analysis of Random Objects in Metric Spaces

Virtually every introductory statistics course characterizes data as fundamentally *categorical* or *quantitative*, with finer distinctions being made under each classification. This long established taxonomy [34] has far reaching consequences in the appropriate summarization, visualization, and statistical analysis for each type of data [35, 36, 37, 38, 39]. Methods for analyzing quantitative, continuous data are particularly pervasive, owing to the ubiquity of multivariate numerical

data sets, and the mature theory of multivariate normal distributions. Many multivariate data have an imbedded structure, e.g. time series, functional data, spatial data, or networks, which have led to specialized methods which attempt to honor their structural characteristics and the statistical implications therein (e.g. dependence, stationarity) [40, 41, 42, 43, 44]. These domain-specific methods often avoid the curse of dimensionality which can arise when applying traditional multivariate methods. The common insight for these specialized approaches is that while observations may consist of many numbers, they are not necessarily “high dimensional” when accounting for the structure they contain.

The field of *object oriented data analysis* takes the view that while each observation in a data set may be comprised of several continuous and categorical measurements, these measurements ultimately represent a singular “object” with a certain structure. This approach ignores the classical distinction between categorical and continuous data, and instead seeks methods that properly account for the structure of a data set. In doing so, object oriented data analysis draws on such areas of mathematics as algebra, topology, and geometry [45, 46, 47, 48, 49].

### **1.3.1 Statistics In Metric Spaces**

One popular approach to object oriented data analysis treats observations as living in a metric space, where the pairwise distances among the observations play a fundamental role in the development of statistical theory and methods [50, 51, 52, 53]. If we take this view, then we must consider random objects in a metric space to develop a statistical theory and analysis methods. In general, metric spaces lack a coordinate system, which proves quite inconvenient since all traditional multivariate methods rely on one. As a result, researchers have re-conceptualized statistical concepts to work with random objects and observations in metric spaces, beginning with the Fréchet mean, and re-building methods analogous to analysis of variance [50, 51] and regression [54, 55].

### **1.3.2 Our Contribution**

In this dissertation we consider random objects being measured in multiple metric spaces, which may arise when random objects may be measured in multiple distinct ways. In this new

multivariate setting, define a Fréchet covariance and Fréchet correlation in two metric spaces, and a Fréchet covariance matrix or Fréchet correlation matrix in an arbitrary number of metric spaces. We prove consistency for the sample Fréchet covariance, and propose several tests to compare the means and covariance matrices between two or more groups. Lastly, we investigate the power and Type I error of each test under a variety of scenarios.

## Chapter 2

### Mathematical Preliminaries

This chapter introduces various mathematical concepts which are invoked later in this dissertation, and also establishes notation for these concepts.

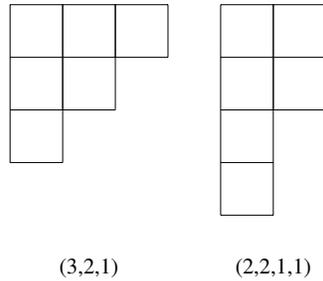
#### 2.1 Binary Matrices and Margins

For Chapters 3 and 4 involving sampling of binary matrices, we let  $p$  and  $q$  be vectors containing row and column margins, respectively, where  $p$  has length  $m$  and  $q$  has length  $n$ .  $N(p, q)$  is the number of valid matrices with row margins  $p$  and column margins  $q$ .  $N(p, q)$  is also called the *partition function*, and the set of margins  $(p, q)$  is *graphical* if  $N(p, q) > 0$ . Given a set of margins  $(p, q)$ , we may say that a matrix is *valid* if it has margins  $(p, q)$ , so that  $N(p, q)$  is also the number of valid matrices with those margins. Let  $\Omega(p, q)$  denote the space of binary matrices satisfying margins  $(p, q)$  so that the cardinality  $|\Omega(p, q)| = N(p, q)$ . Matrix elements will be indicated with subscripts, e.g.  $X_{1,2}$  indicates the element in row 1, column 2 of matrix  $X$ . A generic element is given by either  $X_{ij}$  or  $(i, j)$  if the matrix is obvious from context.

#### 2.2 Integer Partitions and Compositions

Chapter 4 will reference compositions, and partitions. A *composition* of an integer  $n$  is an ordered set of positive integers which sum to  $n$ . For example,  $(1, 3, 2)$ ,  $(3, 2, 1)$ , and  $(1, 1, 2, 2)$  are distinct compositions of 6. For our purposes, we will allow 0 to occur in a composition as well, so  $(3, 2, 1, 0)$  will still be considered a composition, which is distinct from the examples above. The purpose of the 0's is to provide padding. A *partition* of an integer  $n$  is an unordered set of positive integers which sum to  $n$ .  $(1, 2, 3)$  and  $(3, 2, 1)$  are indistinguishable as partitions, but still distinct from  $(1, 1, 2, 2)$ . By convention, partitions are listed in non-increasing order, as  $(3, 2, 1)$  and  $(2, 2, 1, 1)$  for example. Partitions can be represented as Young diagrams [56], which

are arrangements of boxes into rows, where each row is a part of the partition and the number of boxes in each row equals the size of that part.



**Figure 2.1:** Young diagrams for  $(3, 2, 1)$  and  $(2, 2, 1, 1)$ , both partitions of 6.

Figure 2.1 shows the Young diagrams for a few example partitions. Note the boxes are left justified, and all Young diagrams for an integer  $n$  will have exactly  $n$  boxes.

The Young diagram for  $x = (3, 3, 2, 1)$ , is shown in Figure 2.2a, and the Young diagram of its *conjugate partition*,  $x' = (4, 3, 2)$ , is just a diagonal reflection of the Young diagram for  $x$ , see Figure 2.2b.

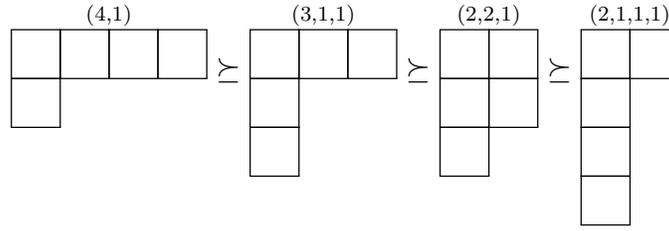
## 2.3 Majorization & The Gale-Ryser Theorem

Chapter 4 will make extensive use of the Gale-Ryser theorem to determine if there exists a valid binary matrix corresponding to a set of margins. Published in the same year by Gale [57] and Ryser [58], the Gale-Ryser theorem provides simple necessary and sufficient conditions for



(a) Young diagram for the partition  $x = (3, 3, 2, 1)$ . (b) Young diagram for the partition  $x' = (4, 3, 2)$ .

**Figure 2.2:** Young diagrams for two partitions of 9,  $x = (3, 3, 2, 1)$  (left) and its conjugate  $x' = (4, 3, 2)$  (right). The Young diagrams are diagonal reflections of one another.



**Figure 2.3:** Example of how partitions with “wider” Young diagrams (left) tend to majorize partitions with “taller” Young diagrams.

margins to be graphical (i.e. they correspond to at least one valid binary matrix). There are several equivalent ways to state the theorem, but here we choose the language of integer partitions and majorization.

Let  $x$  and  $y$  be integer partitions of  $\lambda$  of length  $l$ , possibly with zero elements, and assume by convention that  $x$  and  $y$  are sorted non-increasing. We say that  $x$  majorizes  $y$  and write  $x \succeq y$  if  $\sum_{i=1}^j x_i \geq \sum_{i=1}^j y_i$  for all  $1 \leq j \leq l$ . Further, denote the *conjugate partition* of  $x$  as  $x'$ , that is,  $x'_j = |\{i : x_i \geq j\}|$ . Again,  $x_i$  may terminate in a number of zero elements as necessary. When not obvious from context, we will specify the length of partitions and their conjugates.

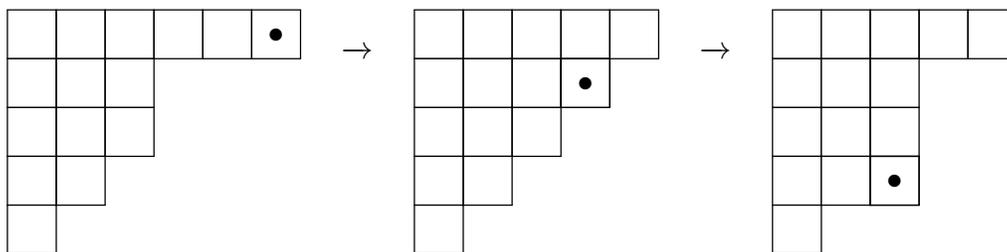
**Theorem 1** (Gale-Ryser). Integer partitions  $x$  and  $y$  are *graphical* if and only if  $x' \succeq y$ , equivalently, if  $y' \succeq x$ .

Incidentally, for integer partitions  $a$  and  $b$  satisfying  $a \succeq b$ , then  $b' \succeq a'$ .

The concept of majorization can be loosely understood using Young diagrams as well: Generally, partitions with “wider” diagrams will majorize partitions with “taller” diagrams. See Figure 2.3 for a few examples. Majorization induces a partial order on the set of integer partitions of integer  $\lambda$ , enabling Young diagrams can be arranged in a *lattice*, meaning any two elements have a unique least upper bound and unique greatest lower bound.

## 2.4 Navigating the Young’s Lattice with Box Shifts

Since partitions of  $\lambda$  produce Young diagrams with  $\lambda$  boxes, any diagram may be turned into any other by shifting the boxes. In fact, Goles et al. [2] describe two types of box shift,  $v$  and



**Figure 2.4:** Example of a  $v$  shift followed by an  $h$  shift in a Young diagram, as described in [2]. These shifts, and their reverse equivalents, enable full navigation of the Young lattice for a particular integer  $\lambda$ .

$h$ , which can be repeatedly applied to reconstruct the lattice, as proven by Brylawski [59]. Given partition  $x = (x_1, x_2, \dots, x_s)$ , consider shifting a box in the Young diagram from index  $i$  to  $j$ , where one of the following conditions hold:

1.  $j = i + 1$  and  $x_i - x_j > 2$  ( $v$  shift).
2.  $j > i$  and  $x_i - x_j = 2$  ( $h$  shift).

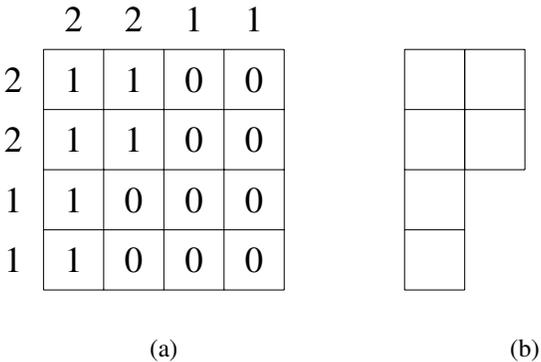
The resulting partition  $y = (x_1, \dots, x_i - 1, \dots, x_j + 1, \dots, x_s)$  is majorized by  $x$  and in fact,  $x$  is a *cover* (direct ancestor) of  $y$  in the majorization lattice. Figure 2.4 shows an example of a  $v$  shift ( $i = 1, j = 2$ ), followed by an  $h$  shift ( $i = 2, j = 4$ ) where the same box (marked with a dot) is involved in both shifts. The  $v$  and  $h$  notation comes from Goles' thinking of partitions as sandpiles, which are rotated counter clockwise 90 degrees from how they are displayed here. In this orientation a  $v$  shift would move a single box predominately vertically, and an  $h$  shift would move a single box predominately horizontally. While  $v$  and  $h$  shifts will always “lower” the diagram in the sense that the resulting partition is majorized by the starting partition, we can apply them in reverse, which will always “elevate” the diagram to a more majorizing partition. Reverse  $v$  and  $h$  shifts will be useful in the development of our divide and conquer algorithm in Chapter 4.

## 2.5 The Shift-Right Algorithm

In Chapter 4, we will require a way to produce an initial binary matrix which satisfies a given set of graphical margins. This can be accomplished using the algorithm developed by Fulkerson and

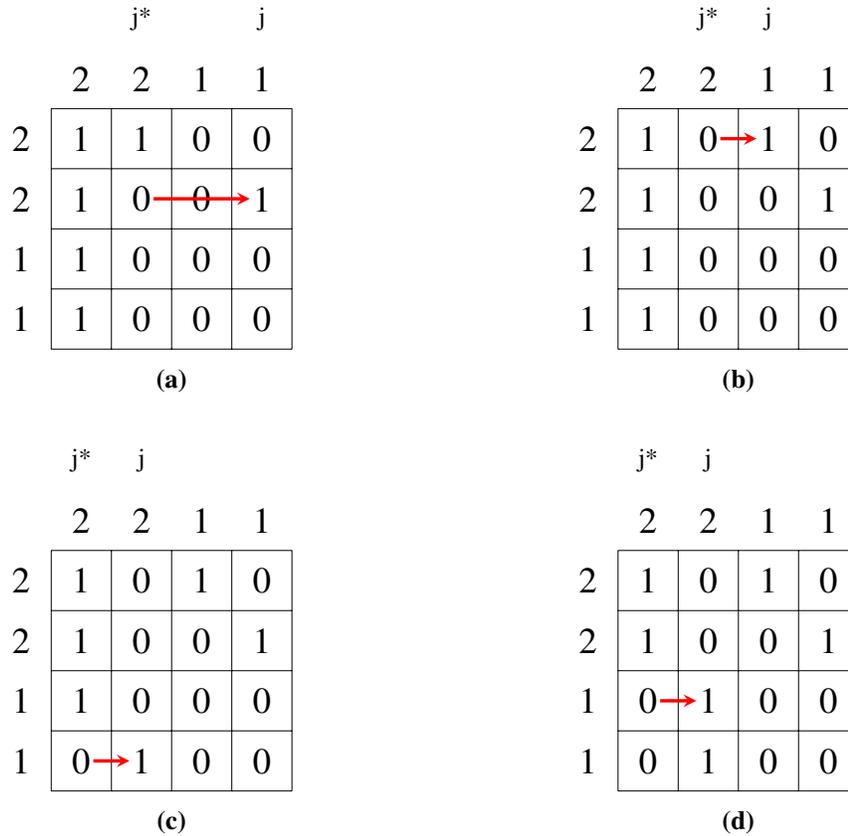
Ryser [5]. The algorithm works by first ignoring the column constraints and placing the required number of 1's in each row, left justified, then selectively shifting 1's to the right in order to satisfy the column constraints, right to left.

Fulkerson and Ryser's algorithm assumes the row margins are ordered non-increasing, from top to bottom, and the column margins are ordered non-increasing, from left to right. This way, the initial placement of 1's resembles the structure of the Young diagram.



**Figure 2.5:** Initial fill of the shift-right algorithm (a), which satisfies the row margins  $p$ , along with the Young diagram for  $p$  (b).

For instance, given row and column margins  $p = q = (2, 2, 1, 1)$  the initial placement, along with the Young diagram for  $p$  is given in figure 2.5. The algorithm proceeds by shifting 1's to the right in their respective rows in order to satisfy the columns margins, from right to left. Specifically, let  $j$  denote the target column which 1's are being shifted into. The 1's which get shifted are those in the rightmost column which is left of column  $j$ , which we call  $j^*$ . The lowest 1's are shifted first, and if column  $j^*$  is exhausted of 1's, then the column to the left of column  $j^*$  becomes the new column  $j^*$ .



**Figure 2.6:** (a)-(d) show the shifting steps of the shift-right algorithm, in order, which produce a valid binary matrix. Red arrows denote the moving of a 1 from the tail to the head. The target column whose margin is being satisfied is marked with  $j$ , and the source column from which 1s are being shifted is marked with  $j^*$ .

Figure 2.6 shows the sequence of shifts which produce a binary matrix satisfying margins  $(p, q)$ , where column  $j$  and  $j^*$  are annotated at each step.

## 2.6 Metric Spaces

Chapter 5 involves random objects existing in metric spaces. A *metric* on a set  $\Omega$  (not to be confused with  $\Omega(p, q)$ , the space of all binary matrices satisfying margins  $p$  and  $q$ ) is a function  $d : \Omega \times \Omega \rightarrow [0, \infty)$  such that, for all  $\omega_1, \omega_2, \omega_3 \in \Omega$ ,  $d$  satisfies the following:

1.  $d(\omega_1, \omega_2) = 0 \Leftrightarrow \omega_1 = \omega_2$
2.  $d(\omega_1, \omega_2) = d(\omega_2, \omega_1)$

$$3. d(\omega_1, \omega_2) \leq d(\omega_1, \omega_3) + d(\omega_3, \omega_2).$$

A *metric space* is a non-empty set endowed with a metric. When there are multiple metric spaces, we give them subscripts, for example  $\Omega_1, \dots, \Omega_D$ .

Some well known metrics that we use in this dissertation include the Wasserstein metric between probability spaces, and the Frobenius metric between matrices. Consider the space  $M$  of probability distributions for which every Borel probability measure on  $M$  is also a Radon measure, and assume a metric  $d$  is defined for this space. Then let  $P_p(M)$  consist of all probability measures on  $M$  with finite  $p$ th moment. The *Wasserstein- $p$  metric* between two probability measures  $\mu$  and  $\nu$  in  $P_p(M)$  is defined as

$$W_p(\mu, \nu) = \inf (E[d(X, Y)^p]^{1/p}), \quad (2.1)$$

where the infimum is taken over all joint distributions of  $X$  and  $Y$  with marginals  $\mu$  and  $\nu$ , respectively. For example, comparing univariate normal distributions with identical variance, say  $X \sim N(a, 1)$  and  $Y \sim N(b, 1)$ , the Wasserstein-2 metric is simply the absolute difference in their means,  $|a - b|$ .

The *Frobenius norm* of a real valued  $m \times n$  matrix  $A$  is

$$\|A\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n a_{ij}^2} = \sqrt{\sum_i^{\min(m,n)} \lambda_i}, \quad (2.2)$$

where  $\lambda_i$  are the eigenvalues of  $A^T A$ . This norm admits a metric between two  $m \times n$  matrices  $A$  and  $B$ :  $d(A, B) = \|A - B\|_F$ .

## 2.7 Riemannian Manifolds

In Chapter 5 we propose comparing symmetric positive definite matrices using a distance on a Riemannian manifold, which we introduce here, after some necessary preliminary definitions.

A *topological space* is an ordered pair  $(X, \tau)$  where  $X$  is a set of points and  $\tau$  is a collection of subsets of  $X$  satisfying:

1.  $X \in \tau$  and  $\emptyset \in \tau$ ,
2.  $\tau$  is closed under countable unions: that is, for  $A_1, \dots \in \tau$ ,  $\bigcup_i A_i \in \tau$ ,
3.  $\tau$  is closed under finite intersections, that is, for  $A_1, \dots, A_n \in \tau$ ,  $\bigcap_i A_i \in \tau$ .

The collection  $\tau$  is called a topology on  $X$ , and the elements of  $\tau$  are called open subsets. A *homeomorphism* between two topological spaces  $T_1$  and  $T_2$  is a function  $f : T_1 \rightarrow T_2$  such that  $f$  is a bijection,  $f$  is continuous, and  $f^{-1}$  is continuous. Loosely speaking, the existence of a homeomorphism implies that  $T_1$  can be continuously deformed into  $T_2$ , and vice versa.

A *manifold*  $M$  is a topological space where for each point in the topology, there exists a local neighborhood which is homeomorphic to an open subset of  $\mathbb{R}^n$  for some  $n$ . Manifolds may be described using a set of *charts*,  $A = \{(U_\alpha, \phi_\alpha)\}_\alpha$ , each of which consists of an open subset  $U_\alpha \in M$  and a homeomorphism  $\phi_\alpha$  from  $U_\alpha$  to a subset of  $\mathbb{R}^n$ . The subsets  $U_\alpha$  cover  $M$ , that is,  $\bigcup_\alpha U_\alpha = M$ , and  $A$  is called an *atlas*. Where any two open subsets  $U_\alpha$  and  $U_\beta$  overlap, we can define the *transition map*  $\phi_\beta^\alpha = \phi_\beta \phi_\alpha^{-1}$ . An atlas where all transition maps are differentiable is called a *differentiable atlas*. This means that overlapping charts have “similar” maps to  $\mathbb{R}^n$  in their overlap. A chart  $(U, \phi)$  is *compatible* with a differentiable atlas if its inclusion in the atlas results in a differentiable atlas, and a maximal differentiable atlas consists of all charts which are compatible with the given atlas.

A *differentiable manifold*  $M$  is a topological space with three additional properties: it is Hausdorff, second countable, and possesses a maximal differentiable atlas. To be Hausdorff, any two distinct points in  $M$  must have disjoint neighborhoods. To be second countable, the topology of  $M$  must have a countable base; that is, there exists some countable collection  $\{U_i\}_{i=1}^\infty$  with  $U_i \in \tau$  such that any element in  $\tau$  is a union of some set of  $U_i$ 's. Differentiable manifolds allow for the construction of a calculus on the topological space, since each chart is homeomorphic to  $\mathbb{R}^n$  and all transition maps are differentiable as well. This leads to the idea of tangent vectors at a point on

the manifold, and a *tangent space* at each point on the manifold, which is a vector space consisting of all tangent vectors at that point.

A *Riemannian manifold*  $M$  is a differentiable manifold in which the tangent space at each point  $p \in M$  is equipped with a positive definite inner product, and the inner product varies smoothly with the point  $p$ . Any vector space with an inner product admits a norm and a corresponding metric, so every tangent space in a Riemannian manifold is also equipped with a *Riemannian metric* that smoothly varies with  $p \in M$ . If  $M$  is connected (i.e. it cannot be expressed as a disjoint union of two nonempty open subsets), then it is possible to define the distance along a path between points  $p, q \in M$  by integrating the Riemannian metric along that path. The shortest possible path length is the *geodesic*, and defines a metric on  $M$ , so every connected Riemannian manifold is also a metric space. Different choices of Riemannian metric (at a point  $p$ ) will result in different metrics on the space  $M$ .

The space of symmetric positive definite (SPD) matrices can be thought of as a cone in Euclidean space, hence it is also a differentiable manifold. There are several choices for Riemannian metric, and each metric gives rise to a different shortest path and shortest path length between two symmetric positive definite matrices. Three common choices are the standard Euclidean metric (Euc), the affine invariant Riemannian metric (AIRM), and the log-Euclidean Riemannian metric (LERM) [60, 61, 62, 63]. The Euclidean Riemannian metric implies the distance between two SPD matrices  $A$  and  $B$  is

$$d_{Euc}(A, B) = \|A - B\|_F = \sqrt{\sum_i \lambda_i^2(A - B)}, \quad (2.3)$$

where  $\|\cdot\|_F$  is the Frobenius norm. As noted in [60, 64], this metric can lead to determinant swelling when used to compute means, that is, the Fréchet mean of a set of SPD matrices may have determinant larger than any of the individual matrices. They also note that shortest paths between two matrices of the same determinant will not necessarily preserve the determinant along

the path. This is used to motivate the AIRM and LERM Riemannian metrics, which do not suffer from the same deficiencies.

In the AIRM geometry, the distance between two SPD matrices  $A$  and  $B$  is given by

$$d_{AIRM}(A, B) = \|\text{Log}(A^{-1/2}BA^{-1/2})\|_F = \sqrt{\sum_i \log^2[\lambda_i(A^{-1/2}BA^{-1/2})]}, \quad (2.4)$$

where  $\text{Log}$  denotes the principal matrix logarithm (the inverse of the exponential map). The *affine invariant* naming comes from the fact that the underlying Riemannian metric is invariant to affine transformations (see [60] for details). Note that because  $A^{-1/2}BA^{-1/2}$  and  $A^{-1}B$  are similar matrices, they have the same eigenvalues, which allows us to avoid taking the square root of  $A^{-1}$  when computing  $d_{AIRM}$ .

In the LERM geometry, the distance between two SPD matrices  $A$  and  $B$  is given by

$$d_{LERM}(A, B) = \|\text{Log}(A) - \text{Log}(B)\|_F, \quad (2.5)$$

which is simply the Euclidean metric after  $A$  and  $B$  are mapped to the set of symmetric matrices via the principal matrix logarithm, hence the name.

## 2.8 Matrix Logarithm

The LERM and AIRM Riemannian metrics considered in Chapter 5 lead to distances on the space of symmetric positive definite matrices which rely on the matrix logarithm. The exponential map of an  $n \times n$  matrix  $A$  is given by the power series

$$\exp A = \sum_{k=0}^{\infty} \frac{A^k}{k!}. \quad (2.6)$$

Matrix  $B$  is a matrix logarithm of  $A$  if  $\exp B = A$ . Since positive definite matrix  $A$  admits an eigenvalue decomposition,  $A = UDU^T$ , we can define the *principal matrix logarithm* as

$$\text{Log}(A) = U\text{Log}(D)U^T, \quad (2.7)$$

Where  $\text{Log}(D)$  is a diagonal matrix where each element is the (scalar) natural log of the corresponding eigenvalue of  $A$ .

## Chapter 3

# Non-Uniform Sampling of Fixed Margin Binary

## Matrices

### 3.1 Introduction

In 1975, Diamond [4] proposed rules which govern the combination of several bird communities on islands in New Guinea, such as “Some pairs of species never coexist, either by themselves or as a part of a larger combination”. However, Connor & Simberloff [7] disputed these rules, suggesting that the combination of species on the islands was not exceptionally different from random. Aside from the ensuing feud within the ecological community [8, 9, 10, 11, 12, 13, 14], Connor & Simberloff sparked interest in a novel form of "null model analysis" [15, 16], a general approach which compares an observed matrix statistic to the distribution of that statistic under a null model. The null model of choice in this case was the set of all random combinations of species on islands satisfying the following three constraints:

- The number of species on each island matches the number observed.
- The number of islands inhabited by each species matches the number observed.
- Each species inhabits islands of comparable size to the islands on which it is observed.

Connor & Simberloff reason that if the observed species combination is not substantially different from the random combinations with respect to some statistic of interest, then it becomes difficult to argue that the observed pattern is anything noteworthy.

To generate combinations which adhere to the above constraints, they represent the species/island combinations as a matrix, with rows representing species and columns representing islands. The element in a particular row and column is set to one if the corresponding species inhabits the corresponding island, and zero otherwise. In the matrix representation, the first two constraints imply

fixing the marginal sums of the matrix, while the third constraint limits which elements are permitted to be one. Thus, the aim in the "null model" analysis is to generate a collection of matrices which possess the desired marginal sums and exclude the prohibited ones. Connor & Simberloff propose two methods, one which begins with a matrix of all zeros matrix and sequentially *fills* in ones in each row randomly, being careful to adhere to the constraints, and another which begins with the observed matrix, and repeatedly *swaps* ones such that the constraints are maintained. Subsequent authors have proposed additional methods for generating the desired matrices, which are generally either a "fill" [25, 10, 12, 14] or "swap" [9, 26, 10, 11, 27] type method.

Early debates surrounding the use of this particular null model focused on the appropriateness of imposing Connor & Simberloff's constraints, the choice of statistic upon which to base comparison, and the interpretability of null model tests [8, 26]. However, focus eventually shifted from whether to use this approach to which method of producing matrices is best [12]. This accompanied theoretical work in binary matrices with fixed marginal sums [3] and Markov Chain Monte Carlo (MCMC) sampling [9], paving the way for swap and fill methods with proven statistical properties [28, 29].

Generally speaking, sampling from a null model can be useful for quantifying the degree of "non-randomness" in an observed matrix (although care should be exercised in the selection of a test statistic and mechanistic interpretation of results). This approach has been used for testing for significance of species co-occurrence [65], evaluating goodness of fit for the Rasch item-response model [9, 29], and in the analysis of contingency tables [29].

In this chapter, we extend the capability of null model analysis by allowing for non-uniform sampling. This permits the construction of more sophisticated "null" models corresponding to non-uniform probability distributions.

### **3.1.1 Review of Sampling Methods**

Here we briefly review the variety of existing swap and fill methods that have been employed for sampling binary matrices with fixed margins. To highlight the contribution of this work, we

point out that most efforts at sampling these matrices thus far have focused on uniform sampling, where each matrix has an equal probability of being included in the sample. Early work gave little attention to actually proving uniformity and few authors address situations in which some matrix elements must be zero, so called *structural zeros*, with exceptions noted below.

As mentioned, these sampling methods can be categorized into *fill* and *swap* methods. Fill methods are typically faster than swap methods, but risk hitting "dead ends", where the process must restart or backtrack in order to avoid violating row and column margin constraints. Swap methods, on the other hand, are arguably simpler to understand and implement, but rely on local changes, so sampled matrices tend to be highly correlated. For both categories, initial methods lacked much theoretical treatment, but more recent work has considered the statistical properties of the algorithms.

### **Fill Methods**

Fill methods begin with a matrix of all zeros, and randomly *fill* in ones in such a way that adheres to the margin (row and column) constraints.

Connor & Simberloff [7] first placed the row sums in decreasing order and column sums in increasing order ("canonical form"), then proceeded row by row, randomly selecting columns to fill with ones, until the desired row sum is reached. If the column is already one, or it has met its column sum, another random selection is made. If the procedure reaches a point where no one can be placed without violating some constraint (a dead end), the process restarts with an empty matrix. No mention is made of the probability of generating a particular matrix. In their "Milne method", Stone & Roberts [10] adapt the fill method of Connor & Simberloff to check whether the proposed fill will lead to a dead end.

Sanderson et al. [12] propose a "knight's tour" algorithm which randomly samples a row and column index for filling, rather than filling one row at a time, and backtracks rather than restarts if the algorithm reaches a dead end. They claim that since rows and columns are sampled randomly, then each valid matrix occurs with equal probability. However, Gotelli & Entsminger show that the "knight's tour" algorithm does not sample uniformly, and propose a "random knight's tour"

which only samples a subset of available elements before backtracking, and backtracks randomly rather than sequentially. They provide no proof of uniform sampling, instead showing on a small example that a sample generated by the random knight's tour does not reject a lack of fit test.

Chen et al. [29] propose a sequential importance sampling approach which randomly fills in each row, according to a conditional Poisson proposal distribution that depends on the previously filled rows. The matrices in the sample are then re-weighted to approximate a uniform distribution. This method also approximates the size of the space of valid matrices.

Taking a different approach, Miller & Harrison [33] devise a recursive sampling algorithm based on the Gale-Ryser Theorem to generate matrices from the uniform distribution exactly, and compute the size of the space exactly. The recursion is over row and column sums, so this method is limited to matrices of moderate size.

Finally, Harrison & Miller [22] adopt the sequential importance sampling approach of Chen et al. for non-uniform sampling. This is, to our knowledge, the only work dedicated towards deliberate non-uniform sampling. They propose a multiplicative weights model that forms the basis of our weights model below, and permit specifying structural zeros by setting a corresponding weight to zero.

## **Swap Methods**

Swap methods have developed in parallel to fill methods, beginning with the original paper of Connor & Simberloff [7]. These methods involve identifying “checkerboard” patterns of zeros and ones in an existing valid matrix, and swapping the ones for the zeros to produce a new matrix which still adheres to the constraints (see Figure 3.1). Repeated swaps result in a Markov chain Monte Carlo algorithm that produces a dependent sample of valid matrices. Of particular importance is the method used to identify potential swaps, and what happens if a proposed swap is invalid, which have computational as well as theoretical implications. Initially, there was some question about whether swapping methods could produce all possible valid matrices by starting at the observed matrix; in other words, if the space of valid matrices is connected under checkerboard swaps.

However, Brualdi [3] showed this to be true using elementary circuits. This paved the way for more development in swap methods [10, 11].

As with fill methods, early swap methods gave no proof that matrices are sampled uniformly, until Kannan et al. [28] proved this for the following procedure: Select two rows and two columns uniformly at random. If they form a checkerboard, then perform a swap, and if not, keep the previous matrix. This result influenced the development of subsequent swapping algorithms. Gotelli & Entsminger’s [14] version of the swapping algorithm importantly picks a swap randomly from the set of eligible swaps with equal probability, not as prescribed by Kannan et al., and consequently demonstrate that the resulting distribution which this samples from is not uniform. This approach is also taken by Zaman & Simberloff [27], who correct for the imbalance by reweighting each matrix when computing summary statistics. These approaches require identifying all possible swaps at each step in order to sample uniformly from them, which is prohibitive for large matrices.

Finally, in the “curveball” algorithm proposed by Strona et al. [66] multiple swaps are performed at once by first selecting two rows (or columns), and permuting the elements in their symmetric difference. Following the proof technique of Kannan et al., Carstens [30] proves that curveball samples uniformly and demonstrates that multiple swaps can lead to faster mixing compared to Kannan et al.’s algorithm.

### **3.1.2 Non-Uniform Sampling**

Most existing null matrix sampling algorithms focus on the uniform distribution, the goal being to compare the observed matrix to other matrices with the same marginal properties, without preference for particular matrices over others. We consider, however, that for some data analyses we may wish to lend more importance to some matrices compared to others, when testing for a specific effect. For instance, Diamond [8] mentions several factors which may give rise to certain species combinations, such as resource overlap, dispersal ability, proneness to local extinction, and distributional strategy, and criticised Connor & Simberloff [7] for not accounting for such “signif-

icant structuring forces”. We argue that it may be possible to control for such external effects in a modified null model, which does not give equal weight to all possible matrices.

Here we consider the task of non-uniform sampling of matrices in a general setting. Following the formulation of Harrison & Miller [22], we assume the existence of a weight matrix indicating the relative likelihood of a one in each element and define the probability of a matrix as a function of these weights. Careful attention must be paid to the consequences of weights equal to zero, which further restrict the set of matrices with positive probability, and affect the validity of the proposed sampling methods.

In the rest of this chapter, we describe a weight-based probability model, then present a novel MCMC checkerboard swap algorithm which incorporates the weights and samples according to the given model. The weights introduce the notion of a *structural zero*, which can impact the efficiency and even correctness of sampling. We then present a novel weighted curveball MCMC algorithm, which in most circumstances is a faster alternative to the weighted checkerboard swap algorithm. We investigate the correctness and efficiency under different weighting and structural zero schemes via simulation. Finally, we revisit the bird species-island dataset first discussed by Diamond [4] and Connor & Simberloff [7] and consider incorporating a weight matrix in the null distribution to show the impact that non-uniform sampling can have on scientific conclusions.

## 3.2 Non-Uniform Sampling Of Binary Matrices

### 3.2.1 Probability Model

Let  $\Omega(p, q)$  be the set of all  $m \times n$  binary matrices satisfying row sums  $p = (p_1, p_2, \dots, p_m)$  and column sums  $q = (q_1, q_2, \dots, q_n)$ . Furthermore let  $W = (w_{ij}) \in [0, \infty)^{m \times n}$  be a non-negative weight matrix representing the relative likelihood of a one in element  $(i, j)$ . Following Harrison & Miller [22], define the probability of observed matrix  $A \in \Omega(p, q)$  as

$$\mathbb{P}(A) = \frac{1}{\kappa} \prod_{ij} w_{ij}^{a_{ij}}, \quad \kappa = \sum_{A \in \Omega(p, q)} \prod_{ij} w_{ij}^{a_{ij}}. \quad (3.1)$$

$$\begin{array}{c}
j & j' \\
i & \begin{array}{|c|c|} \hline 1 & 0 \\ \hline \end{array} \\
i' & \begin{array}{|c|c|} \hline 0 & 1 \\ \hline \end{array}
\end{array}
\longrightarrow
\begin{array}{c}
j & j' \\
i & \begin{array}{|c|c|} \hline 0 & 1 \\ \hline \end{array} \\
i' & \begin{array}{|c|c|} \hline 1 & 0 \\ \hline \end{array}
\end{array}$$

**Figure 3.1:** A  $2 \times 2$  submatrix of  $A$  from rows  $i$  and  $i'$ , and columns  $j$  and  $j'$ . which are not necessarily consecutive. Swapping the positions of ones and zeros preserves the row and column sums of  $A$ .

We call the set of elements  $\{(i, j) : w_{ij} = 0\}$  *structural zeros*, because any matrix with positive probability must have  $a_{ij} = 0$  if  $w_{ij} = 0$ . Let  $\Omega'(p, q) = \{A \in \Omega(p, q) : \mathbb{P}(A) > 0\}$  be the set of matrices with positive probability. We interpret the weights by considering the relative probability between two matrices  $A, B \in \Omega'(p, q)$ :

$$\frac{\mathbb{P}(A)}{\mathbb{P}(B)} = \frac{\prod_{\substack{ij: a_{ij}=1 \\ ij: b_{ij}=0}} w_{ij}}{\prod_{\substack{ij: b_{ij}=1 \\ ij: a_{ij}=0}} w_{ij}}, \quad (3.2)$$

which is governed by the weights of the elements not shared between them. Furthermore, we can write the conditional probability:

$$\mathbb{P}(B|\{A, B\}) = \frac{\prod_{\substack{ij: b_{ij}=1 \\ ij: a_{ij}=0}} w_{ij}}{\prod_{\substack{ij: a_{ij}=1 \\ ij: b_{ij}=0}} w_{ij} + \prod_{\substack{ij: b_{ij}=1 \\ ij: a_{ij}=0}} w_{ij}}. \quad (3.3)$$

This forms the basis for the swapping probability defined in the next section.

### 3.2.2 Weighted Checkerboard Swap Algorithm

We propose a swapping algorithm similar to that of Brualdi [3], which identifies  $2 \times 2$  submatrices with diagonal ones and zeros, as in Figure 3.1, where the rows and columns are not necessarily consecutive. The positions of the ones and zeros can be swapped without altering the row or column sums. This swap operation produces a matrix  $B \in \Omega(p, q)$  which differs from  $A$  in four elements. In the uniform sampling case, all matrices have equal probability, so swaps are

performed with a fixed probability. In our version, the swap is performed with probability

$$\mathbb{P}(B|\{A, B\}) = \frac{w_{i'j}w_{ij'}}{w_{ij}w_{i'j'} + w_{i'j}w_{ij'}} := p_{ij;i'j'}. \quad (3.4)$$

This modification implies matrices with zero probability will be visited with probability zero.

In Algorithm 1, we present an MCMC procedure based on weighted checkerboard swap. As with any MCMC algorithm, this algorithm can be modified to incorporate burn-in and thinning to reduce autocorrelation between saved samples.

---

**Algorithm 1:** Sampling Via Weighted Checkerboard Swaps

---

```

1 Let  $A^{(0)}$  be the observed matrix;
2 for  $k$  in  $1:N$  do
3   Set  $A^{(k)} = A^{(k-1)}$ ;
4   Sample two elements  $(i, j)$  and  $(i', j')$  uniformly at random from the non-zero
     elements of  $A^{(k)}$ ;
5   if  $A_{i'j}^{(k)} = 1$  or  $A_{ij'}^{(k)} = 1$  then
6     | next  $k$  (no checkerboard);
7   end
8   Sample  $s \sim \text{Bernoulli}(p_{ij;i'j'})$ ;
9   if  $s = 1$  then
10    | set  $A_{ij}^{(k)} = A_{i'j'}^{(k)} = 0$  and  $A_{i'j}^{(k)} = A_{ij}^{(k)} = 1$  (perform swap);
11  end
12 end

```

---

If all weights are positive, then we have the following result:

**Theorem 2.** Given binary matrix  $A$  and weight matrix  $W$  with  $w_{ij} > 0$ , then Algorithm 1 generates a Markov chain with stationary distribution given by (3.1).

Proof of Theorem 2 is given in Appendix A. If weights are allowed to equal zero, the Markov chains may mix more slowly and even become reducible as we demonstrate in the next section.

$$A = \begin{array}{|c|c|c|c|c|} \hline 1 & 0 & 0 & 0 & 0 \\ \hline 0 & 1 & 0 & 0 & 0 \\ \hline 0 & 0 & 1 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 & 0 \\ \hline 0 & 0 & 0 & 0 & 1 \\ \hline \end{array}
\qquad
B = \begin{array}{|c|c|c|c|c|} \hline 1 & 0 & 0 & 0 & 0 \\ \hline 0 & 1 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 & 0 \\ \hline 0 & 0 & 0 & 0 & 1 \\ \hline 0 & 0 & 1 & 0 & 0 \\ \hline \end{array}$$

**Figure 3.2:** Matrices  $A$  and  $B$  sharing the same margins and structural zeros, with structural zeros shaded grey. The difference  $A - B$  has a single elementary circuit and the hamming distance between  $A$  and  $B$  is 6, so Brualdi’s method [3] gives a maximum of 2 swaps to transition between  $A$  and  $B$ . While this is true in the absence of structural zeros, this example demonstrates that structural zeros can cause the bound to be violated, as seven swaps are required to transition between  $A$  and  $B$  while honoring the structural zeros.

### 3.2.3 Structural Zeros

As mentioned, structural zeros have received little attention since originally introduced by Connor & Simberloff [7]. We show here that structural zeros can be problematic for swapping algorithms, first addressing sampling efficiency then correctness.

Structural zeros can result in some checkerboard swaps having zero probability, which removes transitions from the Markov chain. This can increase the number of swaps necessary to transition between two states, in turn reducing the sampling efficiency, and is reflected in a larger diameter of the Markov state space. Here we show that Brualdi’s [3] upper bound on the diameter of the Markov space no longer holds with structural zeros.

Given matrices  $A, B \in \Omega(p, q)$  the upper bound on the maximum number of swaps required to transition from one matrix to the other is given by  $d_H(A, B)/2 - k$ , where  $d_H(A, B)$  is the Hamming distance between  $A$  and  $B$ , and  $k$  is the largest number of non-overlapping “elementary circuits” in  $A - B$ , which we briefly explain here. Consider  $A - B$  as a weighted adjacency matrix of a bipartite graph, where the rows and columns form separate partitions of a vertex set. Suppose an entry of  $+1$  indicates an edge from the row vertex to the column vertex, and an element of  $-1$  indicates an edge from column vertex to the row vertex. An elementary circuit is then defined as a cycle of edges in this graph.

To show how the bound can be violated, consider the matrices in Figure 3.2, which have unit row and column sums and we have indicated structural zeros using grey fill. The Hamming distance

$$A = \begin{array}{|c|c|c|} \hline 0 & 1 & 0 \\ \hline 0 & 0 & 1 \\ \hline 1 & 0 & 0 \\ \hline \end{array} \quad B = \begin{array}{|c|c|c|} \hline 0 & 0 & 1 \\ \hline 1 & 0 & 0 \\ \hline 0 & 1 & 0 \\ \hline \end{array}.$$

**Figure 3.3:** Matrices  $A$  and  $B$  sharing the same margins and structural zeros, with structural zeros shaded grey. No sequence of checkerboard swaps which honor the structural zeros can transition from  $A$  to  $B$ , demonstrating that structural zeros can render the Markov state space is reducible.

between them is  $d_H(A, B) = 6$  and there is a single elementary circuit, giving a bound of  $\frac{6}{2} - 1 = 2$ . However, seven swaps are required to transition from  $A$  to  $B$ , which can be described as a sequence of column pairs, since there is only one one per column:  $((1, 2), (1, 3), (1, 4), (1, 5), (1, 2), (2, 3))$ .

Additionally, the example in Figure 3.3 illustrates how structural zeros can render it impossible to transition between two matrices via checkerboard swaps. Because no sequence of swaps transforms  $A$  into  $B$ , the Markov chain described in Algorithm 1 is reducible and hence there is no unique stationary distribution.

In many cases, the presence of a few structural zeros will not result in a reducible chain, but in the general case structural zeros can be fatal to checkerboard swap based MCMC sampling algorithms. However, let us consider a special case where they are not problematic, specifically where structural zeros satisfy the following:

**Definition 3.2.1.** The structural zeros of a matrix  $A$  are *monotonic* if there exists a permutation of the rows and columns of  $A$ , such that  $w_{ij} = 0$  implies  $w_{i'j} = w_{ij'} = 0$  for  $i' < i, j' < j$ .

This situation can arise when the rows and columns reflect different moments in time. For example, let  $A$  represent a citation network among academic papers, where a one indicates that the row article cites the column article. Suppose the rows and columns are arranged chronologically with increasing row/column index. Clearly, articles cannot cite into the future, so we prohibit some ones using structural zeros. In this example, the upper triangle (including the diagonal) of  $A$  are structural zeros. Furthermore these structural zeros are monotonic as  $A$  will satisfy the definition if the order of the columns of  $A$  are reversed.

When structural zeros are monotonic, Algorithm 1 produces a valid sampler, formalized in the following theorem.

**Theorem 3.** Given binary matrix  $A$  and weight matrix  $W$  where any structural zeros are monotonic, then Algorithm 1 generates a Markov chain with stationary distribution given by (3.1).

The only modification of the proof for Theorem 2 needed to prove Theorem 3 is to show irreducibility of the chain given structural zeros (see Appendix A).

### 3.2.4 Weighted Curveball Algorithm

Strona et al. [66] proposed a faster version of checkerboard swapping called the curveball, which imagines two rows trading their ones like baseball cards. Two rows are randomly selected, and the ones in their symmetric difference become candidates for the trade, where several ones may move between rows in a trade. In the uniform setting, the proposed trade is always performed. In the non-uniform setting, we modify the algorithm in the same way as for weighted checkerboard swapping, and introduce a probabilistic trade, with probability dependent on the weights of the elements involved in the trade. The entire procedure is presented in Algorithm 2.

Carstens [30] proves that the unweighted curveball algorithm can be used to sample uniformly, and we contribute the following theorem for the weighted case:

**Theorem 4.** Given binary matrix  $A$  and weight matrix  $W$  where any structural zeros are monotonic, then Algorithm 2 generates a Markov chain with stationary distribution given by (3.1).

In Appendix A, we provide a proof of this theorem when there are no structural zeros and extend to the case where monotonic structural zeros are present.

## 3.3 Simulations

We performed simulations to investigate the behavior of the weighted swap and weighted curveball algorithms. In the first simulation, we verify the weighted checkerboard swap algorithm samples correctly from a non-uniform distribution for a small example. In the second simulation, we compare the mixing performance of both algorithms for different weighting and structural

---

**Algorithm 2: Sampling Via Weighted Curveball**

---

```
1 Let  $A^{(0)}$  be the observed matrix;
2 for  $k$  in  $1:N$  do
3   Set  $A^{(k)} = A^{(k-1)}$ ;
4   Sample two rows  $i$  and  $i'$  (such that  $i \neq i'$ ) uniformly at random from the non-zero
   elements of  $A^{(k-1)}$ ;
5   Let  $A_{i \setminus i'} = \{j : A_{ij} = 1, A_{i'j} = 0, w_{i'j} > 0\}$ ;
6   Let  $A_{i' \setminus i} = \{j : A_{i'j} = 1, A_{ij} = 0, w_{ij} > 0\}$ ;
7   if  $|A_{i' \setminus i}| = 0$  or  $|A_{i \setminus i'}| = 0$  then
8     | next  $k$  (no trade);
9   end
10  Let  $C$  be the ordered list resulting from concatenating  $A_{i \setminus i'}$  and  $A_{i' \setminus i}$ ;
11  Permute the elements of  $C$  uniformly at random;
12  Let  $B_{i \setminus i'}$  be the first  $|A_{i \setminus i'}|$  elements of  $C$ , and let  $B_{i' \setminus i}$  be the last  $|A_{i' \setminus i}|$  elements of  $C$ 
   ;
13  Let  $\mathbf{j}_i = \{j : j \notin A_{i \setminus i'}, j \in B_{i \setminus i'}\}$  (0 to 1 in row  $i$ );
14  Let  $\mathbf{j}_{i'} = \{j : j \notin A_{i' \setminus i}, j \in B_{i' \setminus i}\}$  (0 to 1 in row  $i'$ );
15  Compute  $p_{ii'} = \frac{\prod_{j \in \mathbf{j}_i} w_{ij} \prod_{j \in \mathbf{j}_{i'}} w_{i'j}}{\prod_{j \in \mathbf{j}_i} w_{ij} \prod_{j \in \mathbf{j}_{i'}} w_{i'j} + \prod_{j \in \mathbf{j}_i} w_{i'j} \prod_{j \in \mathbf{j}_{i'}} w_{ij}}$ ;
16  Sample  $s \sim \text{Bernoulli}(p_{ii'})$ ;
17  if  $s = 1$  then
18    | Set  $A_{ij}^{(k)} = 1$  and  $A_{i'j}^{(k)} = 0$  for  $j \in \mathbf{j}$ ;
19    | Set  $A_{ij}^{(k)} = 0$  and  $A_{i'j}^{(k)} = 1$  for  $j \in \mathbf{j}'$ ;
20  end
21 end
```

---

$$A = \begin{array}{|c|c|c|} \hline 1 & 0 & 0 \\ \hline 0 & 1 & 0 \\ \hline 0 & 0 & 1 \\ \hline \end{array} \quad W = \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 2 & 1 & 2 \\ \hline 1 & 2 & 1 \\ \hline \end{array}. \quad (3.6)$$

**Figure 3.4:** Example binary matrix  $A$  and weight matrix  $W$ . There are six binary matrices with the same row and column margins as  $A$ , shown in Figure 3.5

zero schemes. Finally, we investigate the effect of different weighting schemes on the resultant sampling distribution of a summary statistic.

For the second and third simulations, we will make use of a global statistic termed *diagonal divergence*, which quantifies how far the ones of a square  $(n \times n)$  matrix  $A$  are from its diagonal:

$$T(A) = \frac{1}{|A|n} \sum_{ij} |i - j| \mathbb{I}(A_{ij} = 1) \quad (3.5)$$

where  $|A|$  is the number of ones in  $A$  and  $\mathbb{I}(\cdot)$  is the indicator function.

### 3.3.1 Small Example

We first consider a small example, such that every Markov state can be enumerated and exact probabilities can be computed.

Consider the binary matrix  $A$ , and weight matrix  $W$  shown in Figure 3.4. Including  $A$ , there are six possible states with row and column sums equal to one (see Figure 3.5). Under a uniform distribution, all matrices would have probability  $\frac{1}{6}$ .  $W$  induces a non-uniform probability distribution among the matrices (see Table 3.1), where matrices  $A$  and  $D$  have lower probability because their non-zero elements correspond to ones in the weight matrix. We compare the true distribution to the distribution of a sample produced by weighted swapping. To sample, we used a burn-in of 10,000 swaps before retaining every 10,000th matrix, to generate a sample of 1,000 matrices. Table 3.1 shows agreement between the true and empirical probability distributions. The KL-Divergence between the empirical and true distributions is  $9.0 \times 10^{-4}$ .

$$\begin{array}{l}
A = \begin{array}{|c|c|c|} \hline 1 & 0 & 0 \\ \hline 0 & 1 & 0 \\ \hline 0 & 0 & 1 \\ \hline \end{array} \\
D = \begin{array}{|c|c|c|} \hline 0 & 0 & 1 \\ \hline 0 & 1 & 0 \\ \hline 1 & 0 & 0 \\ \hline \end{array} \\
B = \begin{array}{|c|c|c|} \hline 0 & 1 & 0 \\ \hline 1 & 0 & 0 \\ \hline 0 & 0 & 1 \\ \hline \end{array} \\
E = \begin{array}{|c|c|c|} \hline 0 & 1 & 0 \\ \hline 0 & 0 & 1 \\ \hline 1 & 0 & 0 \\ \hline \end{array} \\
C = \begin{array}{|c|c|c|} \hline 1 & 0 & 0 \\ \hline 0 & 0 & 1 \\ \hline 0 & 1 & 0 \\ \hline \end{array} \\
F = \begin{array}{|c|c|c|} \hline 0 & 0 & 1 \\ \hline 1 & 0 & 0 \\ \hline 0 & 1 & 0 \\ \hline \end{array}
\end{array}$$

**Figure 3.5:** All  $3 \times 3$  matrices with row and column sums equal to one.

State	Probability	Empirical Probability
<i>A</i>	0.056	0.051 (+/- 0.007)
<i>B</i>	0.222	0.237 (+/- 0.013)
<i>C</i>	0.222	0.217 (+/- 0.013)
<i>D</i>	0.056	0.058 (+/- 0.008)
<i>E</i>	0.222	0.222 (+/- 0.013)
<i>F</i>	0.222	0.215 (+/- 0.014)

**Table 3.1:** Theoretical vs. empirical probabilities for matrices in Figure 3.5 with binary matrix and weight matrix as defined in (3.6). Empirical probabilities also show Tukey-Hanning [1] MCMC standard errors.

### 3.3.2 Mixing Performance

In MCMC sampling, efficient mixing is desirable in order to maximize inferential capability and minimize computation time. Due to autocorrelation in sequential observations, the precision of a statistic from MCMC sampling is less than it would be for independent, identically distributed (iid) sampling. The difference in precision can be measured by effective sample size (ESS), which gives the number of iid samples required to match the precision of the MCMC sample [67]. We simulated matrices using weighted checkerboard swapping and weighted curveball under different weight schemes and structural zero configurations in order to assess their impact on mixing, as measured by ESS. ESS was computed for the diagonal divergence statistic using the *coda* package [68].

To assess the impact of weights, we constructed  $20 \times 20$  weight matrices by sampling each element independently from one of the distributions: Exponential(1), Uniform(0,1), and Uniform(0.5, 1), and also a weight matrix of all ones, corresponding to uniform sampling. After constructing the weight matrices, we introduced structural zeros with one of two methods. The first method was to

randomly set elements equal to zero with fixed probability  $p$ , where  $p$  was either 0.1, 0.25, or 0.5. The resulting matrices will likely correspond with irreducible Markov chains, given the relatively small probabilities. The second method was to designate a lower triangular region of the weight matrix as all zeros, so that they would be monotonic. Triangles covered either 10%, 25%, or 50% of the weight matrix. We generated a  $20 \times 20$  starting binary matrix by setting each element to one independently with probability 0.25, resulting in a density of 23.5% before enforcing structural zeros.

For each combination of weights and structural zeros, we ran each sampling algorithm for 10,000 iterations (with no burn-in or thinning). Figure 3.6 shows ESS for the different conditions, where each line is the average over ten individual chains with different random seeds. For illustration, Figure 3.7 shows ESS for all chains for the condition in the bottom-right panel of Figure 3.6.

Comparing the weighted swap algorithm to the weighted curveball algorithm, we see that the curveball algorithm tends to mix more rapidly. This reflects the fact that weighted curveball can perform multiple trades simultaneously, whereas weighted swap only swaps two ones at a time, increasing the autocorrelation between sequential observations. This difference is less pronounced for the Exponential(1) and Uniform(0,1) weighting schemes compared to the Uniform(0.5, 1) and “All ones” weighting schemes. This is because the former weight schemes contain weights near zero, making one in the corresponding position of  $A$  possible but unlikely. Since curveball attempts to make several trades simultaneously, this increases the chances of attempting to trade a one into an unlikely position. This makes the trade itself unlikely, so fewer trades take place overall, removing the advantage that weighted curveball has over weighted swap. The latter two weight schemes do not produce weights near zero (ignoring structural zeros), so more curveball trades are successful, restoring the ESS gap between the algorithms. One instance where weighted swap may have an advantage over weighted curveball is in sparse matrices with weights near zero. The sparsity would limit the ability of weighted curveball to perform multiple trades, and the low weights would favor the “smaller steps” taken by weighted swap.

Comparing random structural zeros to monotonic structural zeros, we see little difference for the top two weighting schemes. As mentioned before, these two schemes already tend to produce matrices with weights near zero, so setting some elements equal to zero does not change the weight matrices much, whether done randomly or monotonically. However, the bottom two weight schemes produce weights away from zero, so setting some as structural zeros is a bigger change. In these cases, mixing is better when structural zeros are monotonic compared to random, with the difference being more pronounced when there are more structural zeros. This is unsurprising given the discussion in section 3.2.3 showing that structural zeros may give rise to reducible chains, whereas for monotonic structural zeros irreducibility is preserved, and the space has bounded diameter (see Appendix A).

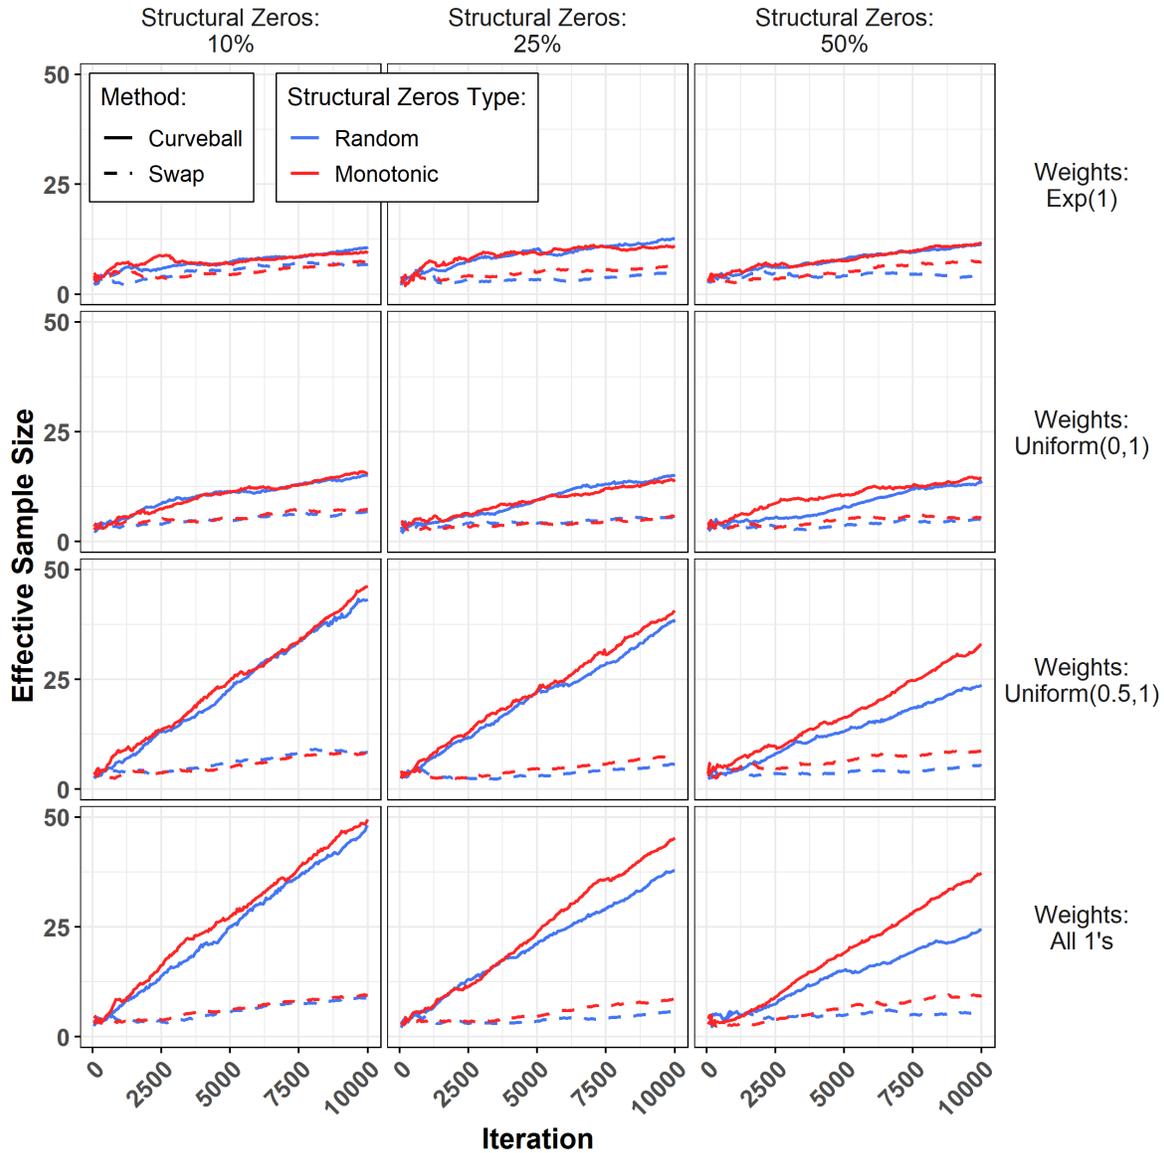
Finally, we acknowledge that “well mixing” may be a questionable label for any of the scenarios shown, since ESS is relatively paltry compared to the number of iterations. In practice, this is ameliorated by careful algorithm implementation (which we do not discuss here) and aggressive thinning.

### 3.3.3 Effect of Heterogeneous Weights

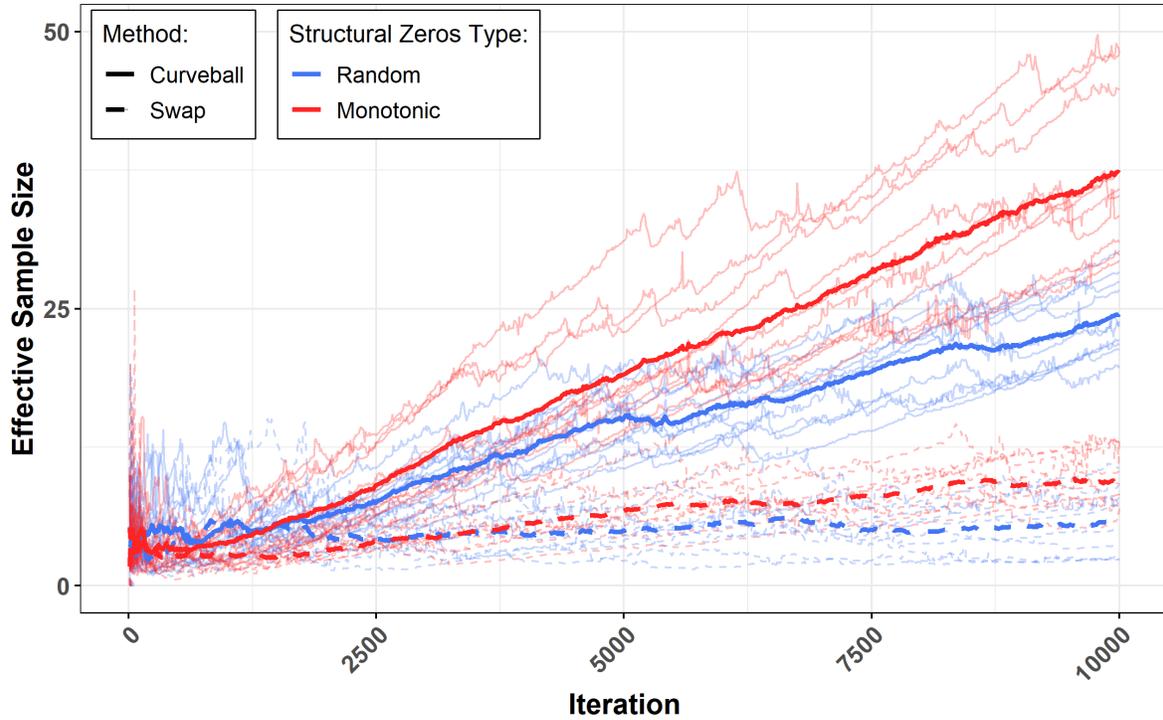
The matrix  $W$  adds flexibility to the sampling algorithms by allowing for non-uniform stationary distributions, however it is unclear how much influence  $W$  has over the resultant sample given the fixed row and column margins. To investigate this issue, we compare the sampling distribution for a statistic under different “strength” weight matrices, which are constructed by taking the elements of a base weight matrix to a positive or negative power. In this case, we let  $A$  have dimension  $50 \times 50$ , and define the base weight matrix as:

$$W_{ij} = \frac{100 - |i - j|}{100}, \quad (3.7)$$

which has value one on the diagonal and value near 0.5 in the corners (see Figure 3.8). Sampling proceeded with a burn-in period of 5,000 proposed swaps, before retaining every 1,000th matrix, to generate a sample of 5,000 retained matrices.



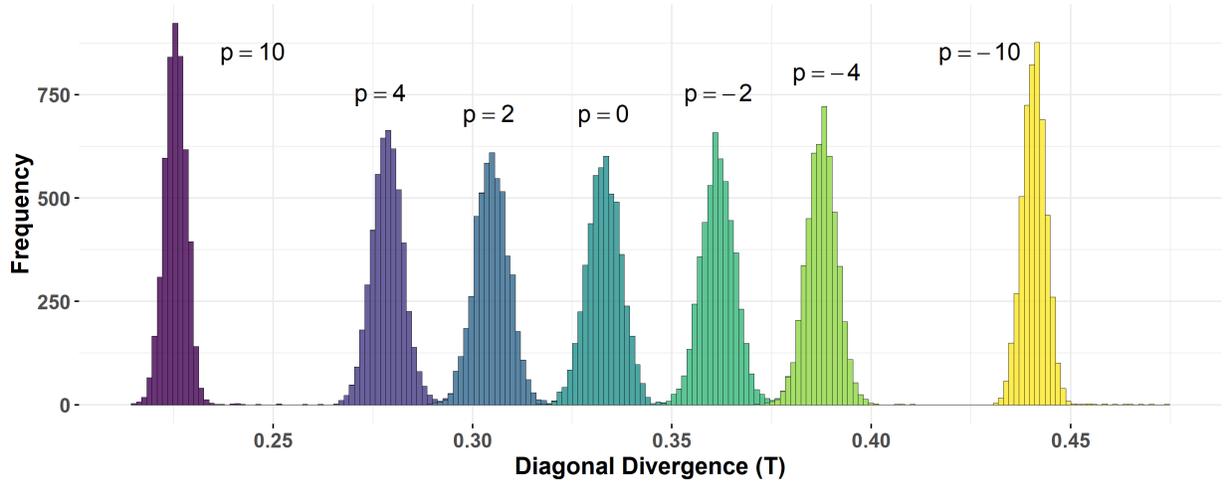
**Figure 3.6:** Average effective sample size for the simulations described in Section 3.3.2, for varying weighting schemes and percentage of structural zeros, where structural zeros are either randomly placed (blue), or arranged monotonically (red). Each line represents the average ESS from ten separate chains using the same starting matrix.



**Figure 3.7:** Detailed view of effective sample size for the simulations described in Section 3.3.2, using a weight matrix of all ones and with 50% structural zeros. Average ESS is shown in bold, with each iteration shown partially transparent.



**Figure 3.8:** Initial binary matrix  $A$  and base weight matrix  $W$  for the simulation described in Section 3.3.3. Both use a shading scale where white denotes zero and black denotes one. The minimum value in  $W$  is 0.5, achieved in the top-right and bottom-left most elements. Note the unfavorable state of  $A$ , since it contains ones corresponding to lower weighted regions of  $W$ .



**Figure 3.9:** Sampling distribution of the diagonal divergence statistic, under different weighting matrices, defined by raising the base weight matrix to the power  $p$ .

Figure 3.9 shows the approximate sampling distributions of the diagonal divergence  $T$  for selected powers of the weight matrix. Since  $W$  favors matrices with ones near the diagonal, larger powers of  $W$  favor lower values for  $T$ , and negative powers of  $W$  favor higher values for  $T$ . Larger values of  $|p|$  emphasize having ones close to the diagonal ( $p$  positive) or packed in the corners ( $p$  negative), where there are relatively few distinct matrices, and there is a corresponding decrease in the variance of the samples in those cases. This figure illustrates how vast the space  $\Omega'(p, q)$  is as there is virtually no overlap in the sampling distributions under different weighting schemes.

### 3.4 Example: New Hebrides Bird Species

We revisit the original Vanuatu (formerly New Hebrides) bird species data provided by Diamond in 1975 [4], to illustrate the impact weighting can have on the analysis of real data. Of interest is whether or not there exists competitive exclusion between different bird species, or groups of species, on these islands. That is, do certain species or groups exclude one another from islands due to competition over resources? The data are represented in a matrix  $A$ , where rows are species, and columns are islands (see Figure 3.10a). The elements of the matrix are one if the corresponding species is found on the corresponding island, and zero otherwise. If a checkerboard pattern exists between two species on two islands, this can indicate competitive exclusion (see

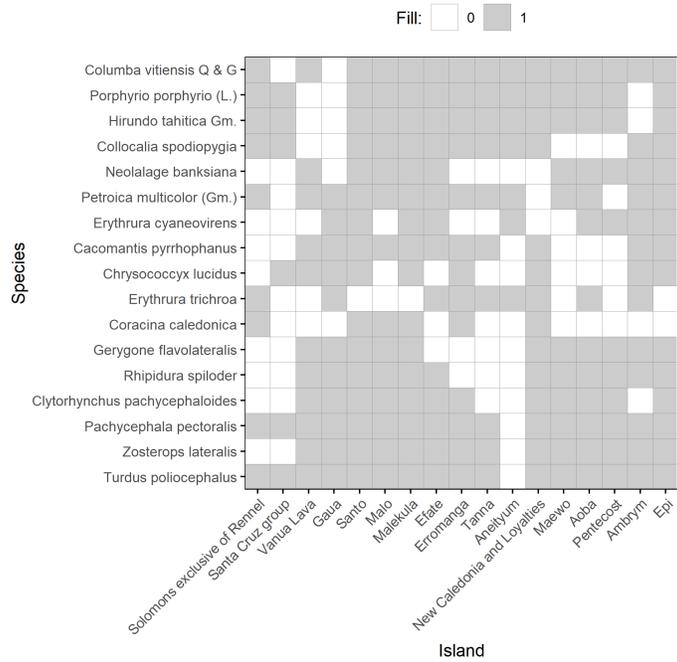
[26]). Accordingly, Stone & Roberts [10] propose the *C-score* to measure competitive exclusion by counting how many checkerboards exist between species pairs in the matrix:

$$C(A) = \frac{2}{m(m-1)} \sum_{j=2}^m \sum_{i<j} (p_i - S_{ij})(p_j - S_{ij}). \quad (3.8)$$

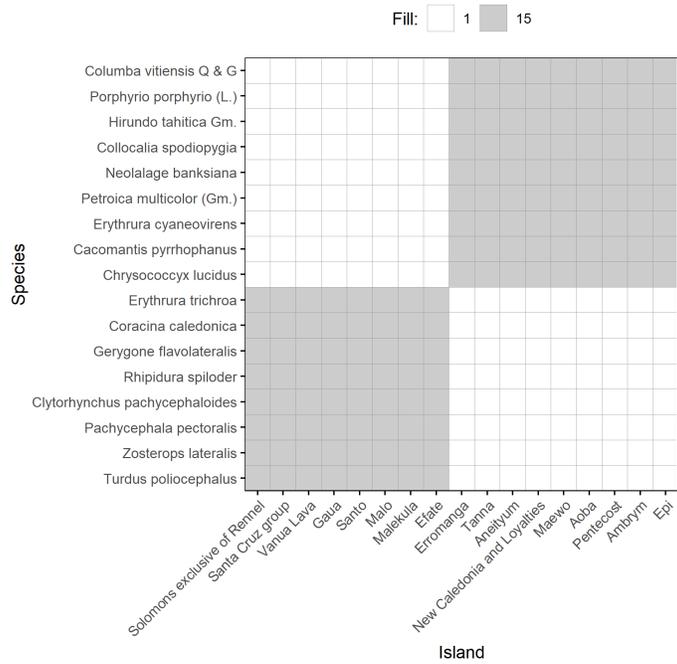
Here,  $S_{ij}$  is the number of shared islands between species  $i$  and  $j$ ,  $p_i$  is the number of islands inhabited by species  $i$ , and  $m$  is the number of species. A high C-score is interpreted as evidence of competitive exclusion. We will analyze the data using a null model analysis twice, once with equal weights (as has been done historically), and once with heterogeneous-weights. Though we are using real data, this is an illustrative example only and we select a weight matrix to that end.

We define a weight matrix  $W$  to encode the hypothesis that island preference by species is higher for certain islands than others (see Figure 3.10b). This could result, for example, from several environmental and biological factors. We then proceed, as did Connor & Simberloff [7], to generate a sample of matrices under the null distribution, assuming equal weights (all ones) and then under  $W$ . We generate 5,000 samples from each weighting regime, after a burn-in of 1,000, and thinning of 500, using the weighted curveball algorithm.

The resulting distributions of C-scores are shown in Figure 3.11, along with the empirical upper-tailed p-values of the observed matrix under each model. We see significance at the  $\alpha = 0.05$  level under uniform sampling, but not after accounting for the weights. The blocked structure of species-island preferences represented in  $W$  favors more checkerboard structures between species in the top half and bottom half, such that the observed matrix no longer appears so exceptional. The lesson here is that by incorporating more information about the bird species into our null model, the evidence in favor of competitive exclusion changes, and thus researchers should carefully consider the appropriate weight matrix when performing such analyses.

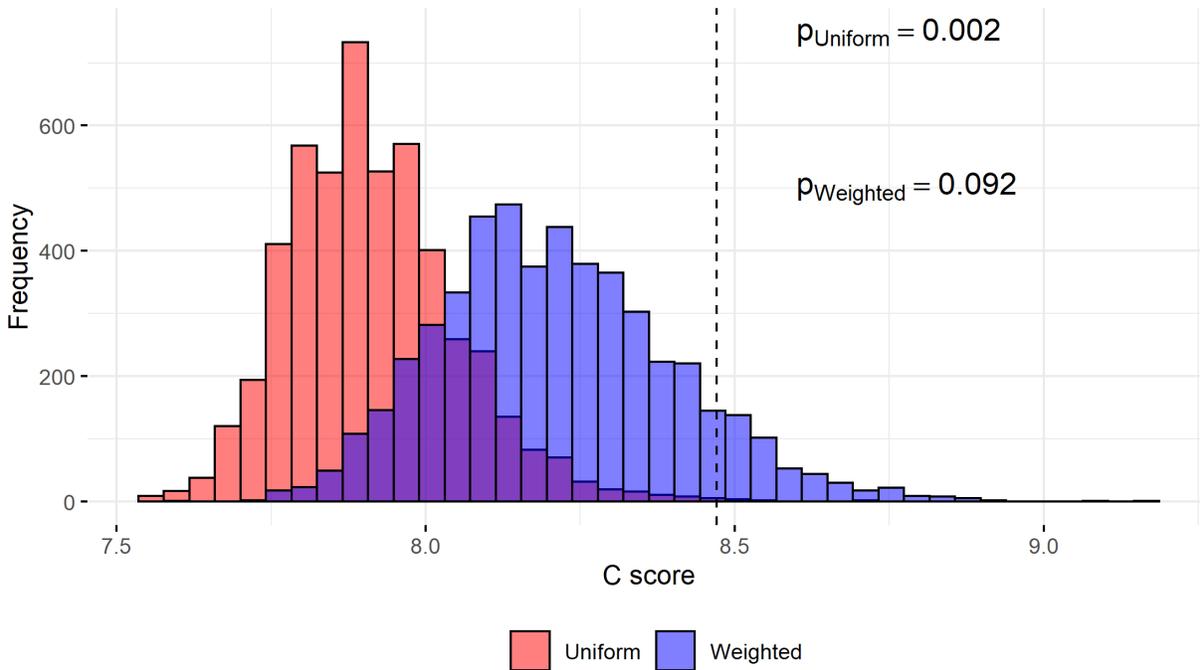


(a) Species-Island occurrence,  $A$



(b) Weight matrix,  $W$

**Figure 3.10:** Observed New Hebrides bird species combination  $A$  [4], and non-uniform weight matrix  $W$  for the data analysis described in Section 3.4. For  $A$ , white denotes zero and grey denotes one. For  $W$ , light grey denotes one and grey denotes 15. With higher values indicating species (rows) preference for certain island (columns) habitats.



**Figure 3.11:** Sampling distribution of the C-score statistic, under uniform (red) and non-uniform (blue) weights. The empirical p-value under non-uniform weights is no longer significant, demonstrating how incorporating species habitat preference could impact whether the observed combination is relatively unlikely.

### 3.5 Conclusions

Given the healthy debate surrounding the appropriate use of null matrix model analysis methods, we recommend that the practitioner weigh this approach against their particular situation and needs. One advantage of this approach is that by constraining, equivalently conditioning on, the row and column sums of the matrix, we remove the need to explicitly model them. However, without an estimation procedure for the weights, these values must be specified *a priori*, and reasonable sensitivity analyses are warranted. Fixing row and column sums is only one way to perform null model analysis, and as Gotelli stated [13]: "There is great value in exploring the results of several null models that incorporate different degrees of randomness". These types of null models provide information about mechanism through the careful consideration of a summary statistic, but it often can be difficult to argue for the correspondence between mechanism and statistic directly. Therefore, if mechanism is of importance, a direct modeling approach may be more appropriate.

We have introduced weighted versions of the checkerboard swap and curveball algorithms, which extend binary matrix sampling to non-uniform probability distributions. Simulations demonstrate that the choice of weight matrix can impact the efficiency of sampling. In fact, when weights are equal to zero, this can slow the mixing process of the MCMC chain, or even prohibit sampling from the desired distribution. However, monotonic structural zeros can arise in time-evolved networks and are better behaved.

Both the weighted swap and weighted curveball algorithms rely on local changes to perform an MCMC walk, which has implications for sampling efficiency (as seen in Figure 3.6). An unavoidable consequence is that large numbers of swaps are required to generate near independent samples. Of course, the computational cost of each swap is very small, and several chains can be run in parallel to produce larger samples. In general, the weighted checkerboard swaps is preferred when the matrix is very sparse or very dense, and weighted curveball is preferred when the matrix is closer to 50% dense (which will produce the most trade candidates). The choice of summary statistic also has implications for mixing. A single “global” statistic, like the C-score, mixes faster than a local statistic, like, for example, the marginal probability for an element of the matrix.

Another possible area of application is directed graphs, which are naturally represented as square binary adjacency matrices. Fixing marginal sums of these matrices is equivalent to fixing in-degree and out-degree for each vertex. Unweighted checkerboard swapping has been used to sample uniformly for vertex-labelled networks with self-loops, also known as the Configuration Model [21]. However, we know of no implementation of graph sampling where edges are weighted, as in this chapter.

In this work, we do not address how to select appropriate weights, as this should be informed by the scientific setting and specific hypothesis being tested. However, the following observation may help guide the construction of  $W$ . Consider binary matrices  $A$  and  $B$  which differ only by a checkerboard swap in rows  $i$  and  $i'$ , and columns  $j$  and  $j'$ :

$$\frac{\mathbb{P}(A)}{\mathbb{P}(B)} = \frac{w_{ij}w_{i'j'}}{w_{i'j}w_{ij}}. \quad (3.9)$$

The ratio of weights communicates the relative probability of the differing elements of two matrices, conditioning on all other elements. Therefore eliciting or estimating the relative propensity of matrix elements could inform the weights specification.

Classic (unweighted) checkerboard swapping and curveball trading, while intuitive and simple to implement, lack the ability to incorporate heterogeneity of probabilities when sampling, rendering them useless in situations where such heterogeneity is known to exist. With our contribution, these algorithms are now equipped to incorporate easy to interpret weights, and to consider more plausible null models, which provides a valuable new tool for matrix analysis.

## Chapter 4

# A Divide and Conquer Algorithm for Exact Uniform Sampling of Matrices with Fixed Margins

### 4.1 Introduction

In this chapter we introduce an approach to uniform sampling of matrices with fixed margins, driven by the intuition that smaller matrices are easier to sample than larger matrices. For small matrices it is often possible to enumerate all matrices satisfying the margins, and so uniform sampling is trivial. Indeed, if all matrices can be enumerated, then moments, distributions, and p-values can all be computed exactly. However the number of feasible matrices increases dramatically as the matrix size increases (factorial in the worst case), consider the following examples:

Row Margins - $p$	Column Margins - $q$	Number of Matrices - $N(p, q)$
(1)	(1)	1
(1, 1)	(1, 1)	2
(2, 2, 1)	(2, 1, 1, 1)	12
(5, 5, 4, 4, 3, 3, 2, 2, 1, 1)	(5, 5, 4, 4, 3, 3, 2, 2, 1, 1)	$\approx 2 \times 10^{13}$
(1, 1, \dots, 1) (k times)	(1, 1, \dots, 1) (k times)	$k!$

The approach of this chapter is to sample arbitrarily sized matrices by recursively dividing them into smaller matrices until sampling is trivial. The challenge in developing such a method is in the mechanics of the division process and careful bookkeeping when doing so. Before describing our method in detail, we first describe a related method which partially inspired ours, and of which our method can be viewed as a generalization.

## 4.2 Miller & Harrison's Exact Sampling

Miller & Harrison [33] devised a method that samples a matrix one row at a time, similar to the importance sampling in [29]. Instead of approximating the row's distribution with a conditional Poisson as in [29], however, they are able to sample from the target distribution exactly, facilitated by an initial recursive counting step.

The key insight is to recognize that specifying the first row of elements results in a smaller (by one row) version of the same problem, where the margins are dependent on the selection for the first row. Since there are generally multiple ways of specifying the first row, there are multiple smaller versions of the problem. Summing over all possible choices for the first row, one is able to equate the number of valid matrices for the full problem with the number of valid matrices across the smaller versions after specifying the first column. The simplest form of the recurrence relation is

$$N(p, q) = \sum_{u \in \{0,1\}^n} N(Lp, q - u) = \sum_{u \in \{0,1\}^n : \{u_i \leq q_i \forall i, \sum_i u_i = p_1\}} N(Lp, q - u), \quad (4.1)$$

where the vector  $u$  with  $u_i \in \{0, 1\}$  represents the selection of the first row, and  $Lp = (p_2, p_3, \dots)$  represents the "left shifted" version of  $p$ , that is,  $p$  with the first element removed.  $N(Lp, q - u)$  is the number of matrices in the reduced problem, where the first row has been removed, and the column margins have been updated to account for the choice of  $u$ . The first expression in (4.1) sums over all  $2^n$  possible choices of  $u$ , but the summand will be zero whenever  $u$  is not compatible with the margins  $q$  and  $p_1$ . The second term excludes these zero terms by restricting the sum to only those  $u$  such that  $u_i$  is not larger than the column margin  $q_i$  for each  $i$ , and  $u$  sums to the required row margin  $p_1$ . Even after this restriction, some terms may be zero if the associated choice of  $u$  cannot lead to a valid binary matrix, that is, if  $(Lp, q - u)$  is not graphical (see Chapter 2).

Initially, this recurrence relation is of little use, as blindly expanding it would result in enumerating all valid matrices, which is undesirable. Thankfully, the choices of  $u$  can be grouped into equivalence classes with the summand being identical for all  $u$  in that class. To see how, imag-

ine  $q$  is sorted non-increasing, so that it can be divided into contiguous blocks of constant value. When specifying  $u$ , placing a one into any position of a block results in the same margin  $q - u$ , just permuted. For example, if  $q = (2, 2, 1, 1)$ , then  $u = (1, 0, 1, 0)$  results in the updated column margins  $q - u = (1, 2, 0, 1)$ , whereas the choice  $u^* = (0, 1, 0, 1)$  results in  $q - u^* = (2, 1, 1, 0)$ . Importantly,  $N(Lp, q - u) = N(Lp, q - u^*)$ , so we may group these terms together. If block  $i$  has  $r_i$  elements and  $f_i$  elements are placed into it, then there are  $\prod_i \binom{r_i}{f_i}$  placements of ones giving the same count. Accounting for permutations within each block gives an alternate recurrence relation with fewer terms:

$$N(p, q) = \sum_{u_f} \prod_i \binom{r_i}{f_i} N(Lp, q - u_f), \quad (4.2)$$

where  $u_f$  is now a single representative from each equivalence class,  $r_i$  determines the size of block  $i$  in  $q$ ,  $f_i$  is less or equal to  $r_i$ , and  $f$  (the vector of  $f_i$ s) sums to  $p_1$ . One more version of the equation is possible: since the equivalence classes are distinguished only in how they allocate ones to the blocks of  $q$ , one can deal with the blocks directly. Defining  $\bar{N}(p, r) = N(p, q)$  where  $r$  is the vector of  $r_i$ s, then the sum can be written as

$$\bar{N}(p, r) = \sum_{f \in C^*} \prod_i \binom{r_i}{f_i} \bar{N}(Lp, r - f + Lf), \quad (4.3)$$

where  $C^*$  is the set of vectors  $f$  such that  $f_i \leq r_i$  for all  $i$  and  $\sum_i f_i = p_1$ . Note that in (4.3) the  $u$ 's need not be specified explicitly.

The recurrence relation in (4.3) allows for  $N(p, q)$  to be computed using dynamic programming, where summands which arise several times when expanding the recurrence only need to be computed once and stored in a lookup table to avoid redundant computation. Because each combination of  $(p, r)$  corresponds with multiple distinct matrices, the count can be produced without enumerating all valid matrices. The lookup table of counts can then be employed in sampling each row, by first selecting an equivalence class  $f$  with probability  $\frac{1}{N(p, q)} \prod_i \binom{r_i}{f_i} \bar{N}(Lp, r - f + Lf)$ ,

then selecting one of the  $u$  in the equivalence class uniformly at random, and repeating until all rows have been sampled.

### 4.3 Divide and Conquer Counting

The approach of Miller & Harrison can be viewed as a division of the matrix into two parts: the top row, and the remaining matrix (or equivalently, the left column, and the remaining matrix, if sampling one column at a time). In this chapter, we generalize this process so division can take place at any row or column in the matrix. Without loss of generality, we consider dividing the *columns* of the matrix into a left part and right part, though the same procedure can be applied if dividing the rows into a top part and bottom part.

Dividing the matrix amounts to selecting an integer  $k$  with  $1 < k \leq n$  which divides the column margins  $q$  into a left part  $q^l = (q_1, \dots, q_{k-1})$  and a right part  $q^r = (q_k, \dots, q_n)$ . We call  $(q^l, q^r)$  a *division*. The associated row margins  $p^l$  and  $p^r$  are undetermined, but their sums must agree with the sums of  $q^l$  and  $q^r$ , respectively, and they must add together to produce the original row margins  $p$ . We also restrict  $p^l$  and  $p^r$  to be graphical, since non-graphical margins contribute nothing to the partition function. These considerations imply the following conditions:

$$\sum_{i=1}^m p_i^l = \sum_{i=1}^{k-1} q_i^l := n^l \quad (4.4)$$

$$\sum_{i=1}^m p_i^r = \sum_{i=k}^n q_i^r := n^r \quad (4.5)$$

$$p^l + p^r = p \quad (4.6)$$

$$p^{l'} \succeq q^l \quad (\text{equivalently } q^{l'} \succeq p^l) \quad (4.7)$$

$$p^{r'} \succeq q^r \quad (\text{equivalently } q^{r'} \succeq p^r), \quad (4.8)$$

Where we introduce  $n^l$  and  $n^r$  as the number of ones in the left and right part respectively, with  $n^l + n^r$ . Conditions (4.7) and (4.8) require that the left and right sets of margins each satisfy the Gale Ryser condition, ensuring that both the left and right submatrices are valid. This gives rise to

the recurrence relation

$$N(p, q) = \sum_{(p^l, p^r) \in C} N(p^l, q^l) N(p^r, q^r), \quad (4.9)$$

where  $C$  is the class of pairs  $(p^l, p^r)$  which satisfy (4.4)-(4.8). Recursive counting and sampling is conceptually similar to the approach of Miller & Harrison [33], where counts are stored in a lookup table for later use when sampling. As in the previous section, some terms in (4.9) will be zero, but in this case if *either* the left margins *or* the right margins are not graphical, the summand will be zero. The central challenge in this divide and conquer approach is in the enumeration of all terms in the set  $C$ , which we call *margin spreading*, while referring to a single proposed  $(p^l, p^r)$  as a *spread*. We will discuss this task in more detail in the next section.

### 4.3.1 Margin Spreading

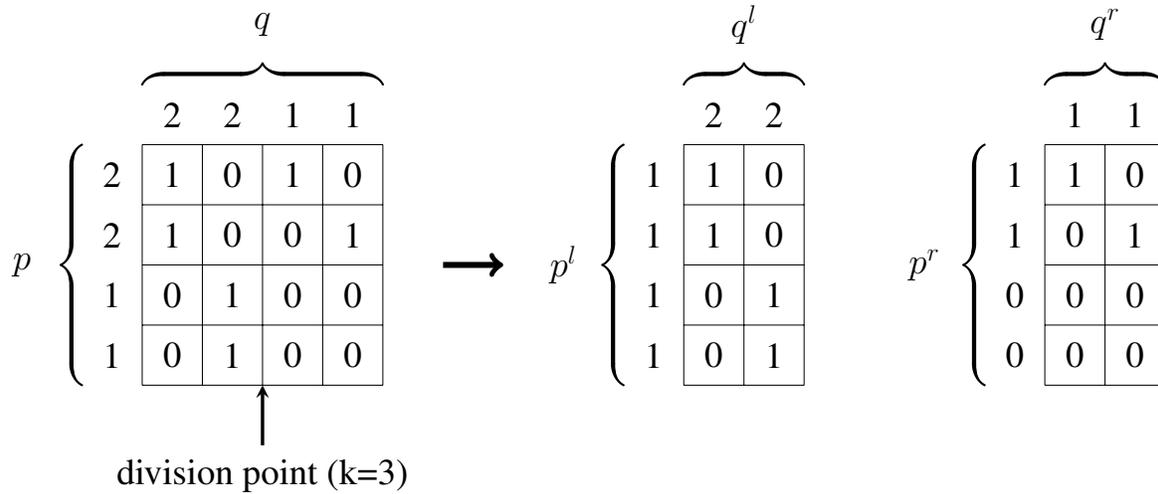
Requiring that a spread  $(p^l, p^r)$  satisfies (4.4) restricts  $p^l$  to the set of length  $m$  compositions of  $n^l$  (allowing the composition to include zeros). Similarly, (4.5) restricts  $p^r$  to the set of length  $m$  compositions of  $n^r$ , permitting zeros in both compositions (again with zeros). Constraint (4.6) resembles a NP-hard 3D-matching problem, where some permutation of vectors  $x$  and  $y$  must sum to vector  $z$ . In our case the matching is even more difficult since  $p^l$  and  $p^r$  are initially unknown. Appendix B outlines an approach to margin spreading using 3D-matching, however it turns out that spread enumeration can be reduced to enumerating lattice points in a convex polytope (described below), which avoids the matching problem.

#### Enumerating Valid Spreads

To enumerate all margin spreads which satisfy (4.4)-(4.6), we begin by trivially considering *one* valid spread, and applying iterative changes to enumerate all the other valid spreads. This avoids any matching, and as will be shown, avoids considering many non-graphical spreads.

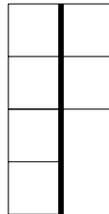
First we produce an initial spread using the *canonical matrix* resulting from Fulkerson & Ryser's *shift-right* algorithm [5] (see Chapter 2 for more details), then observe the implied margins

for each part in the division. As an example, consider the margins  $p = q = (2, 2, 1, 1)$ . Applying the shift-right algorithm produces the left matrix shown in Figure 4.1.



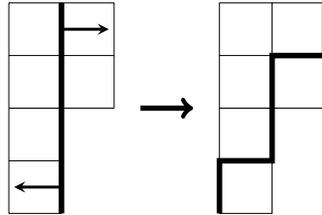
**Figure 4.1:** Creating an initial margin spread. The initial margins  $p$  and  $q$  are used by Fulkerson & Ryser’s *shift-right* algorithm [5] to produce an initial matrix (left left). The matrix is then divided at the division point (determined by  $k = 3$  in this example), and the initial margin spread,  $p^l$  and  $p^r$ , are just the row margins of the left and right parts (right).

If we choose the center as the division point ( $k = 3$ ), the left two columns imply  $p^l = (1, 1, 1, 1)$ ,  $q^l = (2, 2)$  and the right two columns imply  $p^r = (1, 1, 0, 0)$ ,  $q^r = (1, 1)$ . The  $(p^l, p^r)$  spread can be visualized using the Young diagram for  $p$  (see Chapter 2 for more details on Young diagrams), as in Figure 4.2,



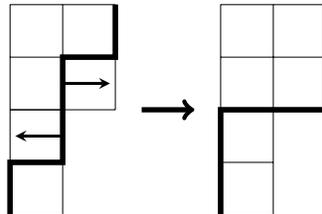
**Figure 4.2:** Using the Young diagram for  $p = (2, 2, 1, 1)$  to visualize the margin spread  $(p^l, p^r) = ((1, 1, 1, 1), (1, 1, 0, 0))$ . The thick border line separates the boxes belonging to  $p^l$  on the left from those belonging to  $p^r$  on the right.

where the thick border line divides the boxes belonging to  $p^l$  on the left from those belonging to  $p^r$  on the right. Now imagine moving the border so that boxes are exchanged between the left and right sides, where the number of boxes on either side is unchanged, i.e.  $n^l$  and  $n^r$  are held constant, as shown in Figure 4.3.



**Figure 4.3:** Creating a new margin spread from an existing one by moving the border line in the Young diagram for  $p$ . The initial border corresponding to  $(p^l, p^r) = ((1, 1, 1, 1), (1, 1, 0, 0))$  (left) is moved according to the arrows to produce a new border line (right) without changing number of boxes to its left and right. The shifted border corresponds to the new margin spread  $(p^l, p^r) = ((2, 1, 1, 0), (0, 1, 0, 1))$ .

This produces an alternative spread,  $p^l = (2, 1, 1, 0)$ ,  $p^r = (0, 1, 0, 1)$ . Repeating the process creates a third spread, shown in Figure 4.4, where  $p^l = (2, 2, 0, 0)$ ,  $p^r = (0, 0, 1, 1)$ .



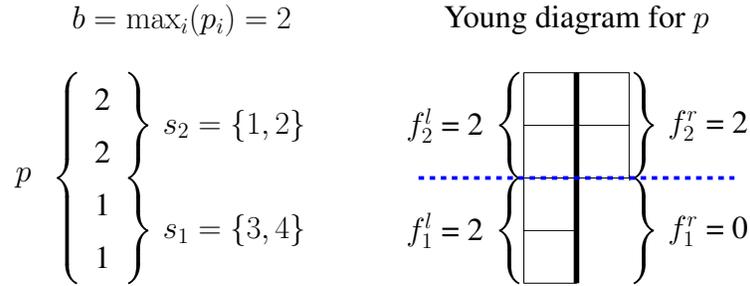
**Figure 4.4:** Creating a new margin spread from an existing one by moving the border line in the Young diagram for  $p$ . The initial border corresponding to  $(p^l, p^r) = ((2, 1, 1, 0), (0, 1, 0, 1))$  (left) is moved according to the arrows to produce a new border line (right) without changing number of boxes to its left and right. The shifted border corresponds to the new margin spread  $(p^l, p^r) = ((2, 2, 0, 0), (0, 0, 1, 1))$ .

The general idea of this approach is to enumerate all valid spreads by moving the border in this way. Ideally, we would do so using the fewest number of border moves. For this purpose, we decompose the enumeration process by partitioning the space of spreads into *fills* (analogous to  $f$  from Miller & Harrison's method). We then enumerate margin spreads within each fill which are

unique up to permutation (analogous to Miller & Harrison's  $u_s$ ), and account for the permutations allowed for each spread. First, we describe the partitioning of margins into fills.

### Partitioning Margins Into Fills

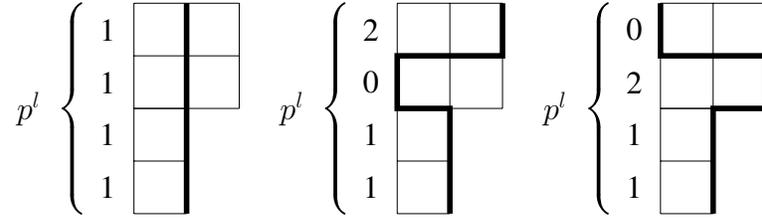
As in [33], we segment the elements of  $p$  into  $b$  blocks  $(s_1, \dots, s_b)$  such that  $s_i = \{j : p_j = i\}$  for  $i \in (1, \dots, b)$  gives the index location of all elements in  $p$  which are equal to  $i$ , where  $b = \max_i(p_i)$ . Each block represents a *level set* of  $p$ , within which permutations can be easily counted. Define the  $b$ -dimensional *left fill* vector  $f^l$  such that  $f_i^l$  is the number of boxes in the Young diagram in the  $i$ th block that are *left* of the border, in other words,  $f_i^l = \sum_{j \in s_i} p_j^l$ . Define the  $b$ -dimensional *right fill* vector analogously as  $f_i^r = \sum_{j \in s_i} p_j^r$ , the number of boxes to the right of the border. Figure 4.5 continues the above example, showing the blocks associated with  $p$  along with the elements of the fill vectors  $f^l$  and  $f^r$ .



**Figure 4.5:** Left: the row margin  $p = (2, 2, 1, 1)$  contains as elements either two or one, so there are two blocks. The indices of block 1 are  $s_1 = \{3, 4\}$ , and the indices of block 2 are  $s_2 = \{1, 2\}$ . Right: The left fills  $f_1^l$  and  $f_2^l$  correspond to the number of boxes in blocks 1 and 2 respectively that are left of the border line in the Young diagram, so  $f_1^l = f_2^l = 2$ . Similarly, the right fills  $f_1^r$  and  $f_2^r$  are the number of boxes in blocks 1 and 2 respectively that are right of the border line in the Young diagram, so  $f_1^r = 0$  and  $f_2^r = 2$ .  $f_1^r$  can also be deduced from  $f_1^l$ , since the total number of boxes in this block is  $1 \times |s_1| = 2$ , so we only need to keep track of  $f_1^l$ .

Note that the left and right fills for a given block must together account for all of the ones in that block, that is,  $f_i^l + f_i^r = |s_i| \times i$ . In other words, the value of  $f_i^r$  can be deduced knowing the value of  $f_i^l$ , so we need only consider unique values of  $f_i^l$ . There may be several margin spreads consistent with a specified fill. Continuing the example, we have  $b = 2$ ,  $s_1 = \{3, 4\}$  and  $s_2 = \{1, 2\}$ . The

shift-right construction produced the left fill  $f^l = (2, 2)$ , and there are three candidate left margin spreads  $p^l$  consistent with this fill:  $(1, 1, 1, 1)$ ,  $(2, 0, 1, 1)$ , and  $(0, 2, 1, 1)$ , see Figure 4.6.



**Figure 4.6:** The only three left margins ( $p^l$ ) which are consistent with the fill  $f^l = (2, 2)$ . Each margin is shown alongside the Young diagram for  $p$  with a thick border separating the boxes belonging to  $p^l$  on the left and  $p^r$  on the right.

These margin spreads are *candidates* because the Gale-Ryser conditions (4.7) and (4.8) have not yet been checked. Checking the conditions for each candidate spread involves sorting  $p^l$  descending and checking that its conjugate majorizes  $q^l$ , then doing the same for  $p^r$  and  $q^r$ . The sorting step means that the second and third candidates can be checked at the same time, since their  $p^l$ 's (respectively  $p^r$ 's) are permutations of each other. Furthermore, these permutations don't affect the number of valid binary matrices, so we only have to consider margin spreads which are unique up to permutation within each block (permuting between blocks may change the fill,  $f^l$ , but since we already intend to sum over all possible  $f^l$ , we restrict permutations to within each block), and account for the number of permutations separately. In the language of group theory, we consider the group action of permuting the elements of  $p^l$  within each block, resulting in a space of orbits, with each member of an orbit sharing the same number of valid binary matrices.

Let  $\tau(p^l)$  represent the orbit of  $p^l$  under the group action of permuting within the blocks of  $p$ . We may also refer to the orbit of a margin spread,  $\tau(p^l, p^r)$ , where for each point in the orbit, the same permutation is applied to both  $p^l$  and  $p^r$ . We also extend the notation to incorporate majorization, so that  $\tau(a) \succeq \tau(b)$  means that  $a$  majorizes  $b$  after both have been sorted non-increasing. The sorting step implies that the relationship holds irrespective of which representatives from  $\tau(a)$  and  $\tau(b)$  are considered. Furthermore, if the Gale-Ryser conditions hold for any members of  $\tau(p^l, p^r)$ ,

then they hold for all members in the orbit, so it is necessary only to check the Gale Ryser condition once per orbit.

Let  $\mathcal{T}(f^l)$  be the class of orbits consistent with fill  $f^l$  which satisfy the Gale-Ryser condition in both the left and right parts. In the example, there are two orbits consistent with  $f^l = (2, 2)$ , represented by  $(p^l, p^r) = ((1, 1, 1, 1), (1, 1, 0, 0))$  and  $(p^l, p^r) = ((2, 0, 1, 1), (0, 2, 0, 0))$ . Spread enumeration proceeds by carefully enumerating fills  $f^l$  for which  $\mathcal{T}(f^l)$  is non-empty, then enumerating the orbits of each fill class and accounting for permutations. To determine if a fill class is non-empty, it is sufficient to check only the minimum orbit (by majorization) from the class against the Gale-Ryser condition, since orbits can be partially ordered by majorization. The existence of the minimum is guaranteed by the following lemma:

**Lemma 1.** Given fill  $f^l$ , assume  $\mathcal{T}(f^l)$  is non-empty. Then there exists a unique minimum orbit  $\tau(p_*^l, p_*^r) \in \mathcal{T}(f^l)$  such that  $\tau(p_*^l) \preceq \tau(p^l)$  and  $\tau(p_*^r) \preceq \tau(p^r)$  for all  $\tau(p^l, p^r) \in \mathcal{T}(f^l)$

Proof of Lemma 1 is given in Appendix A. An immediate consequence of Lemma 1 is that if any orbit  $\tau(p^l, p^r)$  satisfies the Gale-Ryser conditions,  $\tau(p_*^l, p_*^r)$  will also, so non-emptiness of  $\mathcal{T}(f^l)$  can be checked by checking the Gale-Ryser condition on only  $\tau(p_*^l, p_*^r)$ . Moreover, the minimum element can be trivially found by arranging boxes as flat as possible within each block, that is, consider the integer partition of each block separately, and take the minimum partition (by majorization) for each one. In the example of this section, the minimum element corresponding with  $f^l = (2, 2)$  is  $\tau((1, 1, 1, 1))$ .

It turns out that enumerating the set of  $f^l$  with non-empty  $\mathcal{T}(f^l)$  amounts to local search. We will show that these fills are arranged in a convex integral polytope in dimension  $b$ , and that the spread orbit implied by the shift-right canonical matrix provides an appropriate starting point for enumerating all points in the polytope.

The vector  $f^l$  exists in  $\mathbb{Z}_{\geq 0}^b$  and we have the constraint  $\sum_i f_i^l = n^l$ , since the left part must contain  $n^l$  ones, which restricts the set of  $f^l$  to a hyperplane of dimension  $b - 1$ . Each  $f_i^l$  is non-negative,  $0 \leq f_i^l$ , and each block of  $p$  defines a maximum capacity for  $f_i$ :  $f_i \leq |s_i| \times i$ . The intersection of these half space constraints is convex, so the hyperplane is restricted to a (convex)

integral polytope. Applying similar constraints for the right side produces another (convex) integral polytope. Let  $F \subset \mathbb{Z}_{\geq 0}^b$  be the intersection of the left and right polytopes, so  $F$  is also convex. Even though we are applying constraints to the left and right sides, recall that a fill pair  $(f^l, f^r)$  is fully specified with just  $f^l$  or  $f^r$ , and we use  $f^l$  by convention.

Consider fill  $f^l$ , and spread orbit  $\tau(p^l, p^r) \in \mathcal{T}(f^l)$ . Define sets  $F_1 \subseteq F$  and  $F_2 \subseteq F$  according to the following:

1.  $f^l \in F_1$  if and only if  $\tau(p^l, p^r)$  satisfies the left Gale-Ryser condition.
2.  $f^l \in F_2$  if and only if  $\tau(p^l, p^r)$  satisfies the right Gale-Ryser condition.

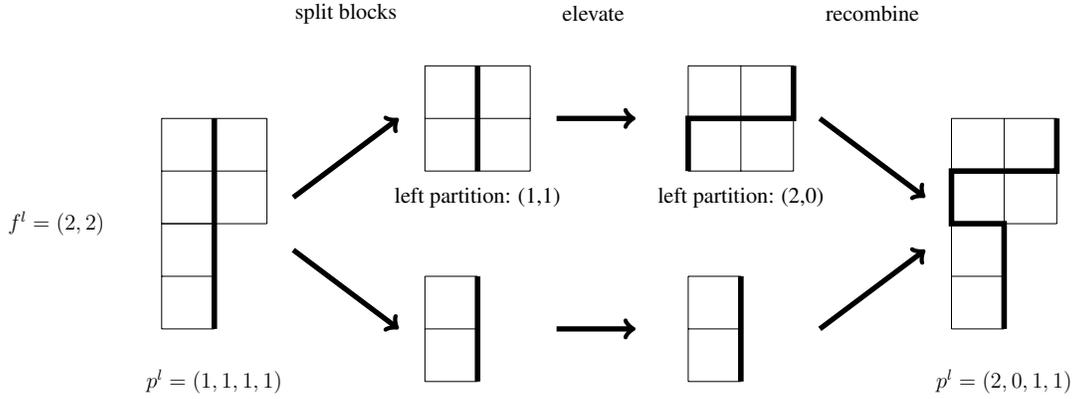
**Theorem 5.**  $F_1 \cap F_2$  is convex.

The proof of Theorem 5 is in Appendix A. For the purposes of enumerating all  $f^l$  with non-empty  $\mathcal{T}(f^l)$ , the important implication of Theorem 5 is that the space is connected, and hence can be enumerated using local search. Convexity is not required for enumeration, but does make the process simpler. Essentially, the shift-right construction provides an initial margin spread, and hence an initial non-empty  $\mathcal{T}(f^l)$  (which satisfied the Gale-Ryser conditions by construction). The remaining fills can be found by expanding outward from the initial lattice point in  $b$  dimensions, checking for non-emptiness at each point using Lemma 1.

In our example, this process starts with the initial fill  $f^l = (2, 2)$  produced by the shift-right construction. We can explore locally to find other fills, while keeping the constraint  $f_1^l + f_2^l = 4$  to ensure  $f^l \in F$ . The neighboring point  $(3, 1)$  doesn't work because  $f_1^l = 3$  exceeds the space available in block 1, given by  $|s_1| \times 1 = 2$ . That is,  $(3, 1) \notin F$ . Further in the same direction is the point  $(4, 0)$ , but Theorem 5 guarantees the space of valid fills is convex, so we don't need to continue in that direction after finding an invalid fill. Exploring in the opposite direction, we consider the point  $f^l = (1, 3)$ , which does satisfy the block size constraints, so  $(1, 3) \in F$ . The next step is to check whether the minimum spread orbit in this fill satisfies the Gale-Ryser conditions for the left and right margins. The minimum orbit is the flattest configuration of boxes on left and right sides:  $\tau(p^l, p^r) = \tau((2, 1, 1, 0), (0, 1, 0, 1))$ , and inspection shows that both the left and

right sides satisfy the Gale-Ryser conditions, so  $(1, 3) \in F_1 \cap F_2$ , and  $\mathcal{T}((1, 3))$  is not empty. The next possible point is  $f^l = (0, 4)$ , with minimum element  $\tau(p^l, p^r) = \tau(((2, 2, 0, 0), (0, 0, 1, 1)))$  that still satisfies the Gale-Ryser conditions. So  $(0, 4) \in F_1 \cap F_2$ , and  $\mathcal{T}((0, 4))$  is not empty. We cannot continue in that direction anymore, so the search has stopped, giving us the convex integral polytope  $\{(2, 2), (1, 3), (2, 4)\}$  which exists on a one-dimensional hyperplane embedded in  $\mathbb{Z}_{\geq 0}^2$ .

Next we enumerate all valid spread orbits within each fill class. In the example above where  $f^l = (2, 2)$ , the shift-right canonical matrix produced spread  $(p_*^l, p_*^r) = ((1, 1, 1, 1), (1, 1, 0, 0))$ , which is the flattest configuration of boxes in each block of  $p$ , so  $\tau(p_*^l, p_*^r)$  is the minimum orbit. To find other orbits, we will “elevate” the left and right margins using the  $v$  and  $h$  shifts described in [2] (more detail in Chapter 2), without changing fill. Preserving the fill requires that each shift take place entirely within a block. Within the  $i$ th block, permutations are ignored (and accounted for separately), so we only need to enumerate the unique integer partitions within that block, for the left and right parts. The block size implies the partitions can have no more than  $|s_i|$  “parts” (in the partition sense of the word) and no part can exceed  $i$ . New candidate orbits can be constructed by picking one partition from each block to form a spread  $(p^l, p^r)$ , and each candidate can be checked against the Gale-Ryser conditions to see if the margins are graphical. Figure 4.7 shows how this step would look for our example. Block 1 cannot be elevated at all, so the only option is to elevate in block 2, taking the left partition from  $(1, 1)$  to  $(2, 0)$ , producing the new spread orbit  $\tau(p^l, p^r) = \tau(((2, 0, 1, 1), (0, 2, 0, 0)))$ , which upon inspection satisfies the left and right Gale-Ryser conditions. No further elevation can take place, so we have found all spread orbits consistent with fill  $f^l = (2, 2)$ . Repeating this process for the other fills produces all possible spread orbits for this division.



**Figure 4.7:** Finding a new margin spread candidate for the fill  $f^l = (2, 2)$ . An initial margin spread  $p^l = (1, 1, 1, 1)$ , in this case produced from the *shift-right* canonical matrix, is split into individual blocks. The top block is selected for “elevation”, meaning the left partition  $(1,1)$  is transformed into the partition  $(2,0)$  which is higher in the Young’s lattice partial order. By elevating only within this block, we maintain the same fill,  $f^l = (2, 2)$ . The blocks are recombined resulting in the left margin spread  $p^l = (2, 0, 1, 1)$ . The new margin will always majorize the starting margin, so we must check that the Gale-Ryser condition is still satisfied (for both  $p^l$  and  $p^r$ ). Repeating this process systematically to each block for each new margin until margins are no longer graphical allows every valid margin within fill  $f^l$  to be enumerated.

Since the partitions within each block can be partially ordered by majorization, then there exists a partial order on the spread candidates. Checking candidates in accordance with the partial order prevents checking too many non-graphical spreads, since once a non-graphical candidate is found, no subsequent candidates in the ordering need be considered. The following lemma summarizes this process and guarantees its success.

**Lemma 2.** Given fill class  $\mathcal{T}(f^l)$  and minimum spread orbit  $\tau(p_*^l, p_*^r)$ , all valid spread orbits in  $\mathcal{T}(f^l)$  can be enumerated using sequences of reverse  $v$  and  $h$  shifts within the blocks of  $p$ .

Proof of Lemma 2 is in Appendix A.

The last step is to account for the allowed permutations within each block of  $p$  (that is, the number of spreads in the orbit  $\tau(p^l, p^r)$ ), which we denote with  $\sigma(p^l, p^r)$ . This number depends on the number of distinct elements of  $p^l$  (or  $p^r$ ) within each block of  $p$  and is easy to compute. For the spread orbit  $\tau((2, 0, 1, 1), (0, 2, 0, 0))$ , there are two permutations in block 1 and two permutations in block 2, giving  $\sigma((2, 0, 1, 1), (0, 2, 0, 0)) = 4$ . Since the permutations don’t affect the number

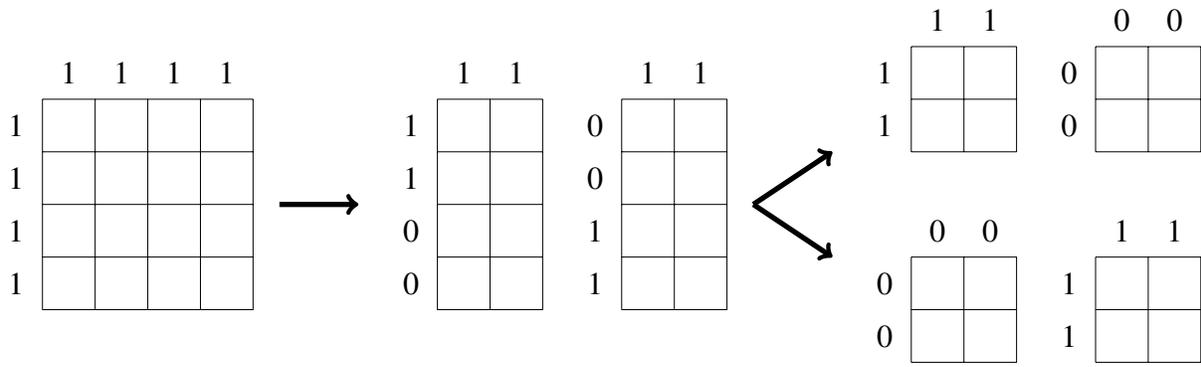
of consistent binary matrices, we only have to consider one representative from each orbit when recursing in the algorithm. Using the approach to enumerating margin spreads outlined above, we revise the recurrence relation one more time:

$$N(p, q) = \sum_{\{f^l: |\mathcal{T}(f^l)| > 0\}} \sum_{\{\tau(p^l, p^r) \in \mathcal{T}(f^l)\}} \sigma(p^l, p^r) N(p^l, q^l) N(p^r, q^r) \quad (4.10)$$

This final version can be used to compute the partition function for any set of margins  $(p, q)$ , which can be stored along with all the intermediate counts in a lookup table to be used for sampling, analogous to Miller & Harrison [33]. Before formally outlining the algorithm in full, we will briefly discuss the base cases.

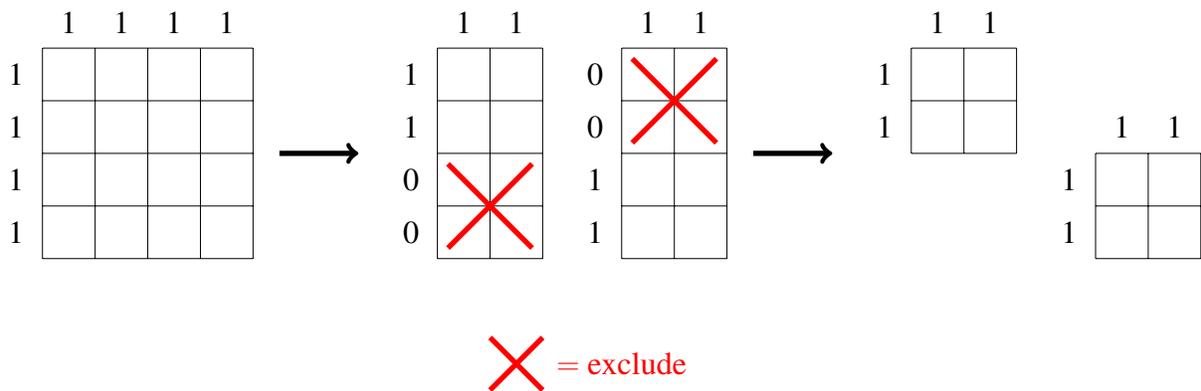
### 4.3.2 Base Cases

The recursion terminates when one of two base cases is encountered. The first case is when the margins fully define the matrix, and the count is one. This occurs when either  $p$  or  $q$  (or both) have length one, or the margins are all zero. The second case is when the count is not necessarily one, but is already known. This is possible because the divide and conquer nature of the algorithm presents the opportunity for margins to be encountered more than once in the recursion process, in which case the count can be stored in the lookup table the first time it is encountered, and retrieved from the lookup table on subsequent encounters. As a simple example, consider margins  $p = q = (1, 1, 1, 1)$  and suppose the cut point  $k = 3$ , so that  $q^l = q^r = (1, 1)$ . The only valid margin spreads are permutations of  $(p^l, p^r) = \left( (1, 1, 0, 0), (0, 0, 1, 1) \right)$ . Dividing the row margins next produces  $2 \times 2$  matrices which are either zero-regular (having margins all zero) or one-regular (having margins all one), as seen in Figure 4.8



**Figure 4.8:** Beginning with the margins  $p = q = (1, 1, 1, 1)$ , the matrix is divided first by columns, with the only valid margin spread (up to permutation of the rows) being  $(p^l, p^r) = ((1, 1, 0, 0), (0, 0, 1, 1))$ . Each smaller matrix is then divided by rows, resulting in four  $2 \times 2$  matrices with regular margins, either all one or all zero. The zero-regular margin matrices are at a base case, with only one way of satisfying the margins (all zeros). The one-regular margin matrices are not yet at a base case, so more division is required, however this process only has to occur once, since both one margin matrices are identical and have the same count.

The count for both of the zero-regular margin matrices is one, and there is no need to subdivide further. The one-regular margin matrices need to be divided further, but this process only needs to be done once since both matrices will have the same count. This example shows another situation which may be exploited, namely when a row or column has margin zero after dividing. There is exactly one way to satisfy such a margin (leave it empty), so it can be excluded when performing the count, leaving reduced versions of the left and right parts, as shown in Figure 4.9.



**Figure 4.9:** If row or column margins have value zero at any point in the counting process, those rows or columns can be excluded from subsequent steps.

The same simplification can be applied if a margin implies that a row or column must be populated with all ones rather than all zeros: that row or column can be excluded from subsequent steps.

### 4.3.3 Algorithm

Here we formally state the algorithm for divide and conquer counting of fixed margin binary matrices, see Algorithm 3. We use the functions `Contains( $LT, x$ )`, `Get( $LT, x$ )`, and `Set( $LT, x, y$ )` to check whether lookup table  $LT$  contains object  $x$ , get the value associated with  $x$ , or set the stored lookup value of  $x$  to  $y$ . The algorithm uses multiple lookup tables, one to store matrix counts,  $LT_C$ , and one to store valid margin spreads,  $LT_S$ , and a few more for bookkeeping when computing margin spreads ( $LT_1$  and  $LT_2$ ). The spread enumeration process involves first enumerating all of the fills valid for a division, then enumerating all of the margins consistent with each fill. Enumerating fills involves a local search which processes fills one at a time, and we use a queue *todo* to store fills waiting to be processed. We use functions `Push(todo,  $x$ )` and `Pop(todo)` to emplace and retrieve fills from the queue. Valid spreads are collected in a list  $S$ , to which we append an element  $x$  using `Append( $S, x$ )`. When enumerating spreads consistent with fills, we use the function `EnumerateOrbits( $f^l$ )` function which iteratively applies the *split-elevate-combine* idea to enumerate representatives for all orbits consistent with the given fill.

We first define the `Spreads` function which enumerates all margin spreads, then the `DAC-Count` function which performs the counting. In this implementation we alternate between dividing rows and columns, and always divide closest to the midpoint, but these choices can be made differently, and the ramifications are discussed in Section 4.5.4.

Since the algorithm populates lookup tables of counts and spreads, the `DAC-Count` function is endowed with a form of “idempotency”: after the initial run, no additional counting is performed, and subsequent calls to the function are composed entirely of lookups. Also, we populate the lookup table of counts and margin spreads at each level of the recurrent hierarchy, which is pivotal to the sampling algorithm described in the next section and ensures that no margin pair is counted more than once.

---

**Algorithm 3:** Divide and Conquer Counting of Fixed Margin Binary Matrices.

---

```
1 Function Spreads ( $p, q^l, q^r$ ):
2   S = empty list;
3   todo = empty queue, contains fills waiting to be processed;
4    $LT_1$  = empty lookup table;
5    $LT_2$  = empty lookup table;
6    $f_0^l$  = left fill implied by shift-right canonical matrix of  $(p, q)$ ;
   // Enumerate convex integral polytope of fills:
7   Push(todo,  $f_0^l$ );
8   while todo is not empty do
9      $f$  = Pop(todo);
10    Set( $LT_1$ ,  $f$ , true);
11     $(p_*^l, p_*^r)$  = flattest partition consistent with  $f$ ;
   // Check if  $\mathcal{T}(f^l)$  is nonempty:
12    if  $(p_*^l, p_*^r)$  satisfy Gale-Ryser conditions then
13      // Add spreads to spreads list:
14       $\sigma$  = Number of possible permutations of  $p_*^l$  within blocks  $p$ .;
15      Append(S,  $(p_*^l, p_*^r, \sigma)$ );
16      for  $(p^l, p^r)$  in EnumerateOrbits( $f$ ) do
17        if  $(p^l, p^r)$  satisfy Gale-Ryser then
18           $\sigma$  = Number of possible permutations of  $p^l$  within blocks of  $p$ .;
19          Append(S,  $(p^l, p^r, \sigma)$ );
20        end
21      end
   // Explore nearby fills:
22      for  $f'$  in fills hamming distance one from  $f^l$  do
23        if Not Contains( $LT_1, f'$ ) then
24          Push(todo,  $f'$ );
25        end
26      end
27    end
28  return S;
29 end
```

---

---

```

30 Function DAC-Count ( $p, q$ ) :
31   if Contains ( $LT_C, (p, q)$ ) then
32     | return Get ( $LT_C, (p, q)$ )
33   end
34    $m = \text{length}(p)$ ;
35    $n = \text{length}(q)$ ;
36   if  $\sum_i(p_i) = 0$  or  $n = 1$  or  $m = 1$  then
37     |  $N = 1$ ;
38   else if ( $p, q$ ) do not satisfy Gale-Ryser conditions then
39     |  $N = 0$ ;
40   else
41     |  $q^l = (q_1, \dots, q_{\lfloor m/2 \rfloor})$ ;
42     |  $q^r = (q_{\lfloor m/2 \rfloor + 1}, \dots, q_m)$ ;
43     |  $S = \text{Spreads}(p, q^l, q^r)$ ;
44     |  $N = 0$ ;
45     | for ( $p^l, p^r, \sigma$ ) in  $S$  do
46       |  $N = N + \sigma \times \text{DAC-Count}(q^l, p^l) \times \text{DAC-Count}(q^r, p^r)$ ;
47     | end
48   Set ( $LT_C, (p, q), N$ );
49   return  $N$ ;
50 end

```

---

## 4.4 Divide and Conquer Sampling

In this section, we describe how to use the counts and margin spreads produced in the previous section to sample from the space of binary matrices with specified margins uniformly. Following the divide and conquer approach to counting, sampling involves randomly sampling from the available margin spreads at each division, and the specification of the actual matrix elements is completed once matrices are at a base case and there is only one option for filling in the matrix. Sampling uniformly from the space of binary matrices does not translate to sampling uniformly from the space of spreads at each division, because some spreads may correspond with more matrices than others, not only because the number of matrices for the left and right margins ( $N(p^l, q^l)$  and  $N(p^r, q^r)$ ) may be greater or lesser, but also the allowable permutations within each block ( $\sigma(p^l, p^r)$ ) may vary as well. The key is to sample from spreads in proportion to the number of matrices corresponding to that choice. For this reason lookup tables are created not only of all the spreads, but also the matrix counts at each division in the hierarchy.

### 4.4.1 Algorithm

Here we describe Algorithm 4 for sampling binary matrices with fixed margins explicitly. These algorithms make use of the `DAC-Count` function, but it is not strictly necessary that counting be performed before sampling starts: if counting has been performed already, then all calls to `DAC-Count` will amount to lookups; if counting has not been performed already, then the first sampled matrix will involve the counting process, but subsequently sampled matrices will use the lookup tables. In our description of the algorithm, we use the function `Matrix( $m, n, x$ )` to denote a  $m \times n$  matrix with elements given by  $x$ . If  $x$  is a vector, then the matrix will have only one column so the meaning should be unambiguous, if  $x$  is a scalar, then all of the elements of the matrix should be given that value. For sampling, we assemble the matrix counts  $c_k$  associated with margin spread  $k$ , then normalize to produce probabilities for sampling,  $P_k$ . The function `Hcat` is horizontal concatenation, `⊤` is the transpose function, and `Sample( $S, P$ )` samples a spread given weights vector  $P$ . As before, we choose to alternate between dividing rows and columns, and divide as close to the midpoint as possible.

## 4.5 Experiments

### 4.5.1 Example Counts

To verify our implementation of divide and conquer counting, we repeat a few example counts performed by previous authors. The first example comes from finch species observations made by Charles Darwin in the Galapagos islands, where the rows represent species and columns represent islands. Each element in the matrix is either one indicating the presence of the row species on the column island, or zero indicating the absence of the species on the island (see [29]). The matrix margins are  $p = (14, 13, 14, 10, 12, 2, 10, 1, 10, 11, 6, 2, 17)$  and  $q = (4, 4, 11, 10, 10, 8, 9, 10, 8, 9, 3, 10, 4, 7, 9, 3, 3)$ . We compute the total number of matrices as  $N(p, q) = 67, 149, 106, 137, 567, 626$  which agrees with Miller & Harrison [33].

In the special case of  $k \times k$  square matrices with one-regular margins, the number of valid matrices is just the number of permutation matrices, or  $k!$ . In the case of  $k \times k$  square matrices

---

**Algorithm 4:** Divide and Conquer Sampling of Fixed Margin Binary Matrices

---

```
1 Function DAC-Sample ( $p, q$ ) :
2    $m = \text{length}(p)$ ;
3    $n = \text{length}(q)$ ;
4   if  $m = 1$  then
5      $M = \text{Matrix}(1, n, q)$ ;
6   else if  $n = 1$  then
7      $M = \text{Matrix}(m, 1, p)$ ;
8   else if  $\sum_i p_i = 0$  then
9      $M = \text{Matrix}(m, n, 0)$ ;
10  else
11     $q^l = (q_1, \dots, q_{\lfloor m/2 \rfloor})$ ;
12     $q^r = (q_{\lfloor m/2 \rfloor + 1}, \dots, q_m)$ ;
13     $S = \text{Spreads}(p, q^l, q^r)$ ;
14     $N = 0$ ;
15     $P = \text{vector of length } |S|$ ;
16    for  $k$  in  $1, \dots, |S|$  do
17       $(p^l, p^r, \sigma) = k^{\text{th}} \text{ spread of } S$ ;
18       $N_k = \sigma \times \text{DAC-Count}(q^l, p^l) \times \text{DAC-Count}(q^r, p^r)$ ;
19       $N = N + N_k$ ;
20    end
21    for  $k$  in  $1, \dots, |S|$  do
22       $P_k = N_k / N$ ;
23    end
24     $(p^l, p^r, \sigma) = \text{Sample}(S, (P_1, \dots, P_{|S|}))$ ;
25     $M^l = \text{DAC-Sample}(q^l, p^l)$ ;
26     $M^r = \text{DAC-Sample}(q^r, p^r)$ ;
27     $M = \text{Hcat}(\text{t}(M^l), \text{t}(M^r))$ ;
28  return  $M$ ;
29 end
```

---

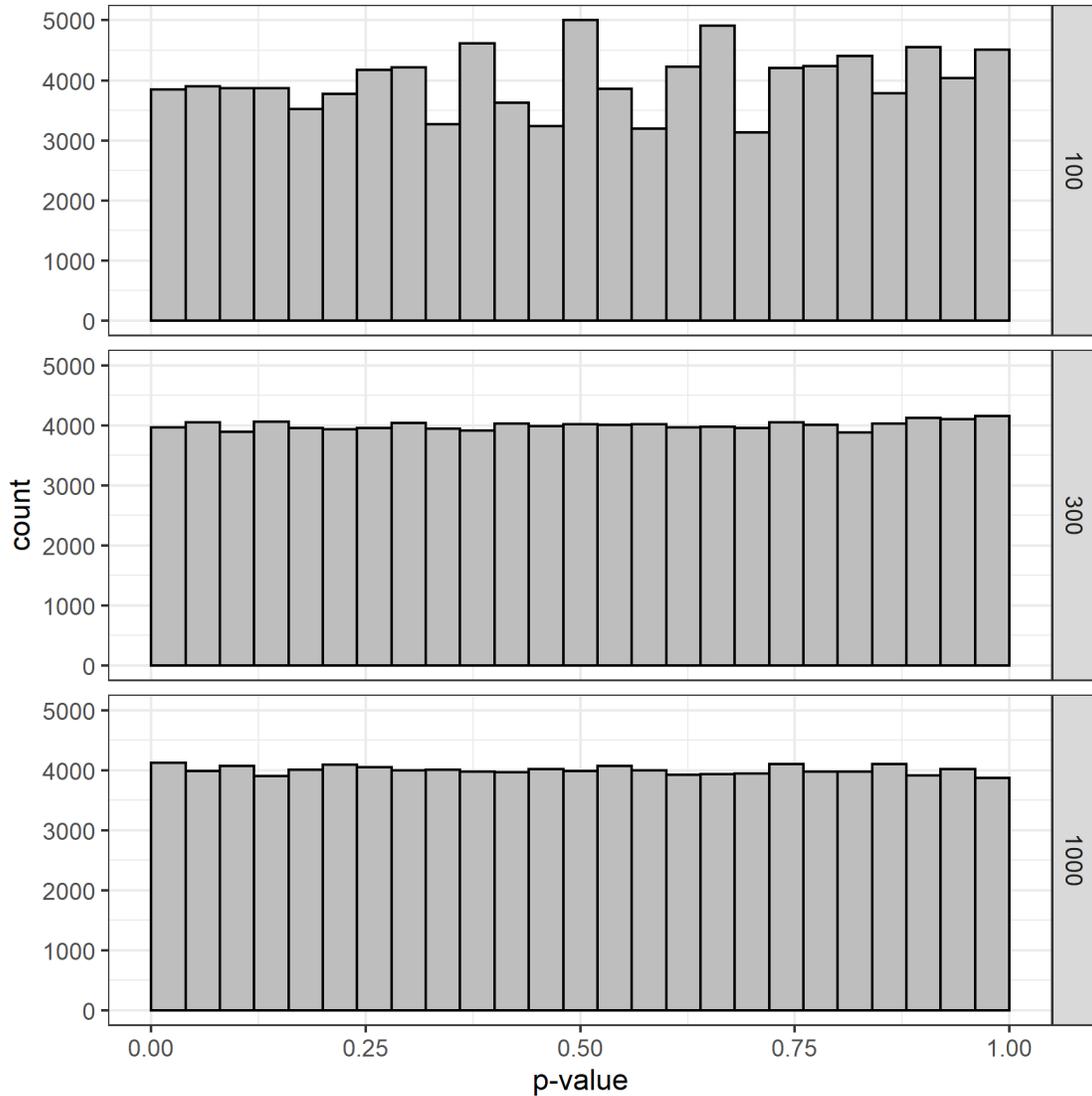
with  $(k - 1)$ -regular margins, the number of valid matrices is also  $k!$  (imagine starting with a permutation matrix and trading out ones for zeros and zeros for ones). We verified our counts of one-regular and  $(k - 1)$ -regular matrices are correct up to  $k = 10,000$ .

### 4.5.2 Example Sampling

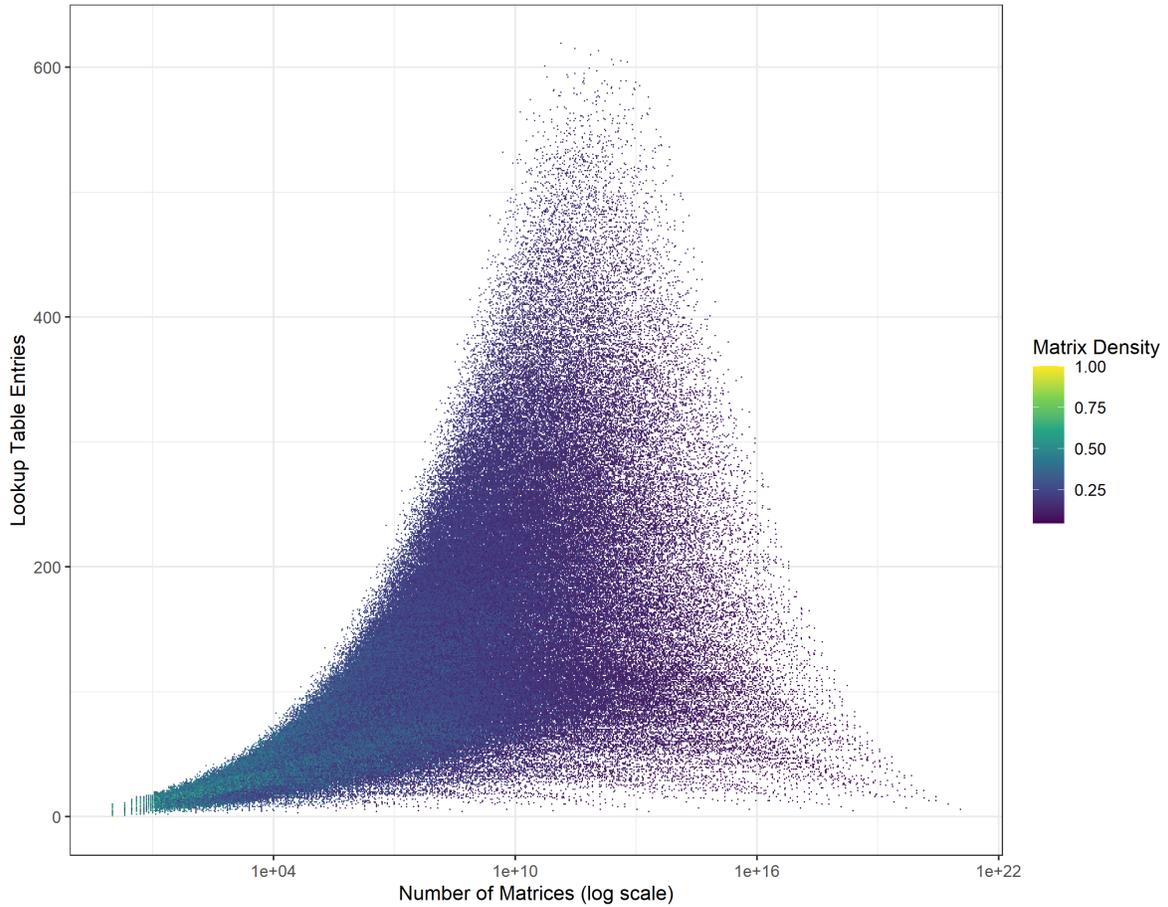
To verify our implementation of divide and conquer sampling, we consider drawing a sample and performing a lack of fit test against a uniform null distribution. We expect the distribution of p-values in such case to be uniform. We draw samples of size 100, 300, and 1,000 for the margins  $p = (2, 2, 2, 1)$ ,  $q = (3, 2, 1, 1)$ , and perform a  $\chi^2$  lack of fit test against the uniform null distribution. There are 24 valid matrices with these margins, so the expected frequency of each matrix under each sample size is approximately 4.1, 12.5, and 41.7, arguably too low for the asymptotic  $\chi^2$  lack of fit test to be appropriate. Hence p-values are estimated using Monte Carlo resampling as in [6] with 10,000 replicates. Repeating this process 100,000 times for each sample size generates distributions of p-values, as shown in Figure 4.10. The uniformity of these distributions lacks evidence of incorrect sampling.

### 4.5.3 Lookup Table Size

By using a lookup table, our algorithm is designed to only compute the counts for a set of margins once and in the course of recursively counting, a count may be used several times. To investigate the degree of re-use, we compare the size of the lookup table (the number of unique counts it contains) to the number of valid matrices, and the density of the matrix. For this comparison, we examine all graphical margin pairs with sums from one up to 22. This comprises about 1.02 million margin pairs, representing roughly  $7.4 \times 10^{21}$  unique matrices. The dimensions of each matrix are implied by the length of each margin, assuming no zero margins. In each case, we employ Algorithm 3 to count the number of compatible matrices and the number of elements in the lookup table after counting. We also compute the density of ones in the matrix, which is influenced by both the sum and length of each margin. Figure 4.11 illustrates the results from counting. The left-most region of the figure corresponds to margins of small, dense matrices where the number



**Figure 4.10:** We repeatedly sample from matrices with margins  $p = (2, 2, 2, 1)$  and  $q = (3, 2, 1, 1)$  using the divide and conquer sampling approach, drawing samples of size 100, 300, and 1000. The panels show the distribution of 100,000 p-values from a  $\chi^2$  lack of fit test against the uniform distribution for each sample size. P-values are estimated using Monte-Carlo resampling [6] with 10,000 Monte-Carlo replicates per sample. The p-values should follow a uniform distribution if sampling is uniform.



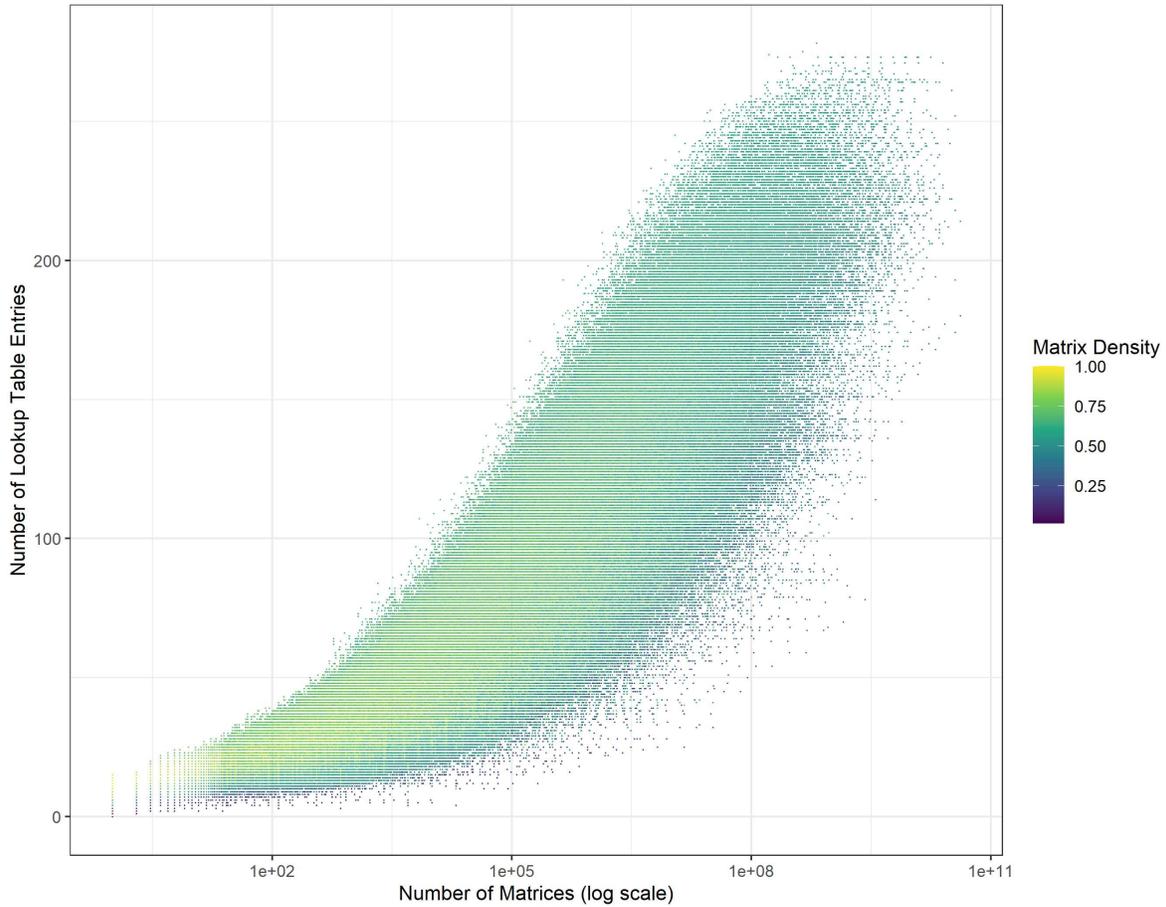
**Figure 4.11:** Number of unique counts stored in lookup table for all graphical margins with sums up to 22. There are roughly 1.02 million graphical margins, representing roughly  $7.4 \times 10^{21}$  unique matrices. Lighter colors indicate higher density matrices, i.e. a higher percentage of ones. Note the log scale on the horizontal axis.

of valid matrices is low, and the required number of lookup table entries is also small. As margin length ( $m$  or  $n$ ) increases, the number of valid matrices generally increases as well, and we see an increase in the lookup table entries also. This reflects the higher number of divisions before reaching the base case for larger matrix dimensions, and the greater number of spreads to consider as the margins get longer and there are more valid matrices. However, the margins with the greatest number of valid matrices tend to require fewer lookup table entries due to increasing information reuse. For example, the largest count corresponds with one-regular margins, where symmetry implies that every margin pair considered in the left part of the split will also be considered in the right part of the split.

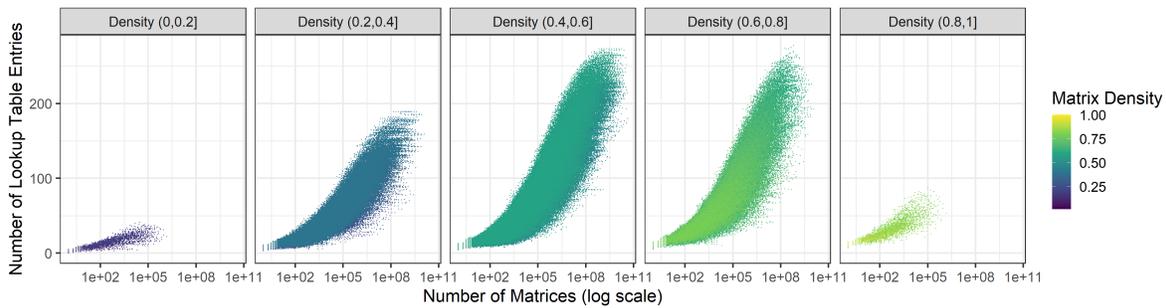
Instead of allowing matrix dimensions to vary for a given number of ones, we can instead consider matrices of fixed size and examine all margin pairs consistent with that size, allowing some margins to be zero now. Figure 4.12 compares the number of lookup table entries to the number of matrices consistent with the given margins for all matrices of size  $9 \times 7$ , where brighter colors indicate more dense matrices. This figure illustrates how lookup table size generally increases for higher count matrices. This figure excludes the right tail present in Figure 4.11 because the dimensions are constrained. To show the influence of density, Figure 4.12b separates points into five density ranges, with less dense on the left and more dense on the right. Densities near 0.5 correspond with the highest matrix counts and the largest lookup tables as well. Though the figure focuses on one size, the same trends can be seen in other sizes of matrices as well.

#### 4.5.4 Cut Point Location

In Algorithms 3 and 4, the cut point is chosen closest to the midpoint of  $q$ , so that the left and right parts are roughly half the original size ( $\lfloor m/2 \rfloor$ ), but this is not strictly necessary. In this section, we explore the effect of moving the cut point, using four example sets of margins. For each example, we perform divide and conquer counting several times for different cut points. Specifically, we define the split margins  $q^l = (q_1, \dots, q_{\lfloor m \times c \rfloor})$  and  $q^r = (q_{\lfloor m \times c \rfloor + 1}, \dots, m)$  where  $c$  is allowed to vary between 0 and 1, and we also ensure that  $p^l$  and  $p^r$  each have length at least



(a) Number of unique counts stored in lookup table for all possible margins of size  $9 \times 7$ . There are roughly 1.29 million graphical margins, representing roughly  $1.99 \times 10^{13}$  unique matrices. Lighter colors indicate higher density matrices, i.e. a higher percentage of ones. Note the log scale on the horizontal axis.



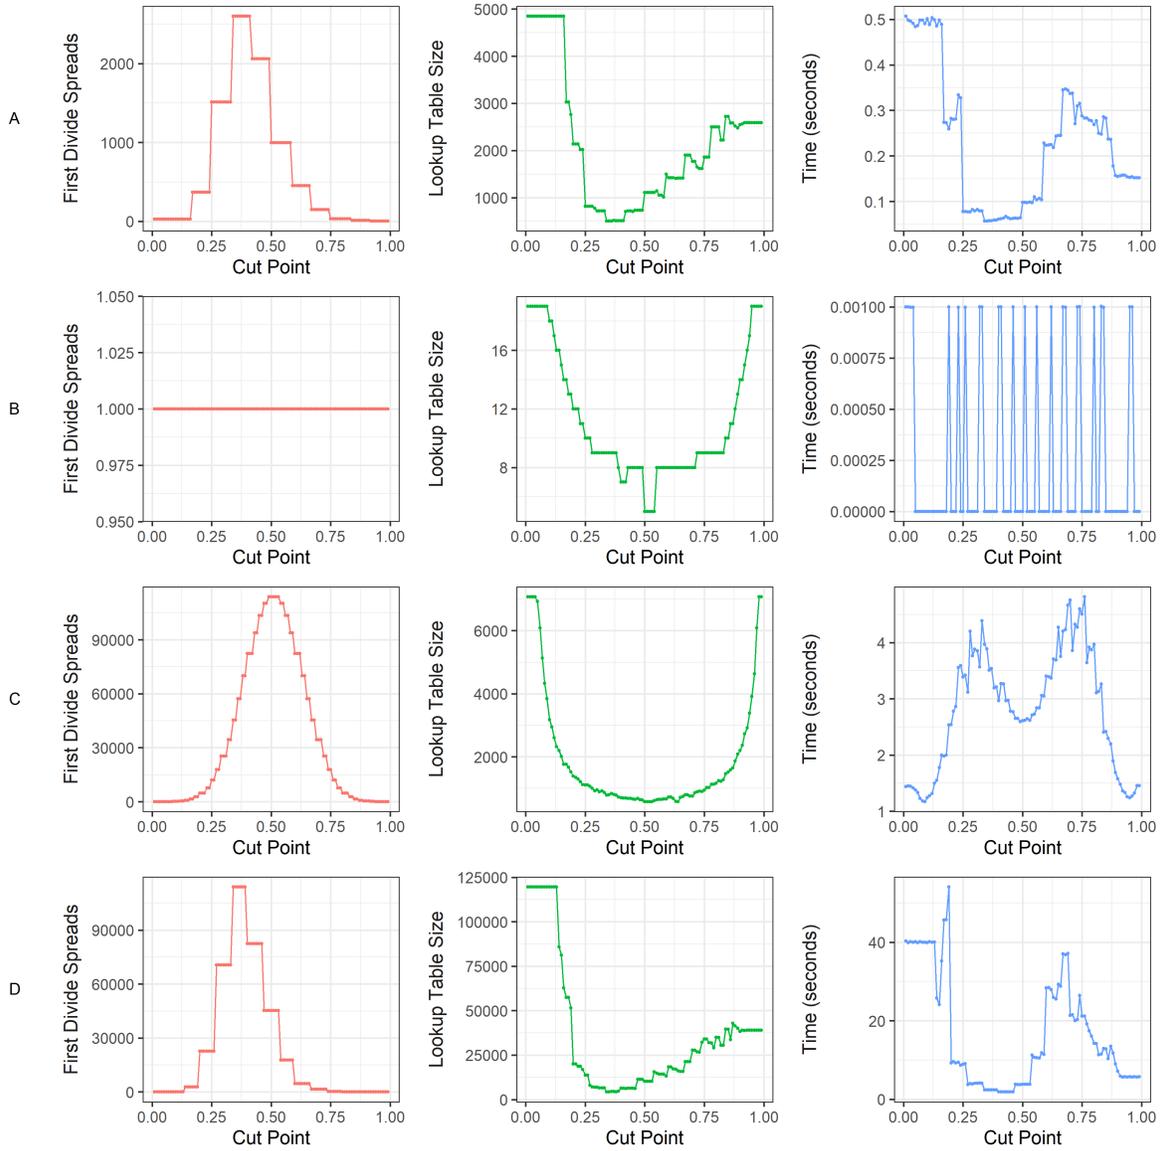
(b) The same results as panel (a), separated by density to more clearly show its effect.

**Figure 4.12**

Label	Row Margins ( $p$ )	Column Margins ( $q$ )
A	(4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)	(same as $p$ )
B	(1, . . . , 1) (20 ones)	(same as $p$ )
C	(1, . . . , 1) (45 ones)	(5, 5, 5, 4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)
D	(5, 5, 5, 4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)	(same as $p$ )

**Table 4.1:** Example margins used for examining cut point and cut strategy. Margins were sorted as displayed when feeding to the algorithm.

one. The division point is applied at all levels of the recursive algorithm (not just the initial division). Varying the cut point can affect the number of margin spreads, which directly increases or decreases the number of recursive calls and the running time. On the other hand, some divisions may be more likely to produce repeated margin lookups, reducing run time. The first division impacts run time the most, as the lengths of  $p$  and  $q$ , the marginal sums, and the number of blocks in  $p$  produce more margin spreads than at lower levels. Increased recursive calls at the first level will create a larger recursion tree than the same number of calls lower in the tree. We use our R/C++ implementation of Algorithm 3 on a PC with an AMD Ryzen 3700X processor and 48GB of RAM (which was well in excess of that required by the software). Figure 4.13 summarizes the simulation results, across three metrics: number of spreads in the first division (red), size of the lookup table (green), and median running time out of three trials (blue). Even though the algorithm is deterministic, run time will vary slightly from run to run because of other computer processes, hence we take the median time for three runs. When dividing near the center of the margins, the number of spreads in the first division tends to be larger compared to when dividing near the edges. Allowing half of  $q$  in each divided part reduces the restrictions imposed by the Gale-Ryser condition, which require  $q^l$  and  $q^r$  to majorize the potential spreads. In cases A and D, where  $q$  is asymmetric, the worst run time occurs left of center, likely because that more evenly divides the ones between  $q^l$  and  $q^r$  (compared to the symmetry seen in example C, which has a symmetric  $q$ ). The exception to this rule is example B, where the regular margins give a single spread no matter what the division point. This illustrates the value of the permutation factor  $\sigma$ , which exploits the symmetry within the margin blocks. The lesson from this metric alone is that perhaps division should take place at the left or right extremes, as with Miller & Harrison’s algorithm.



**Figure 4.13:** Measurements from the counting process for examples A-D from Table 4.1: A:  $p = q = (4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ ; B:  $p = q = (1, \dots, 1)$  (20 ones); C:  $p = (1, \dots, 1)$  (45 ones),  $q = (5, 5, 5, 4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ ; D:  $p = q = (5, 5, 5, 4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ . Three metrics are shown for varying cut points: number of spreads in the first division (red), size of the lookup table (green), and median running time out of three trials (blue). Note the different vertical scales on each plot. The rapid jumps between 0.000 and 0.001 seconds for example B are a result of limited time resolution.

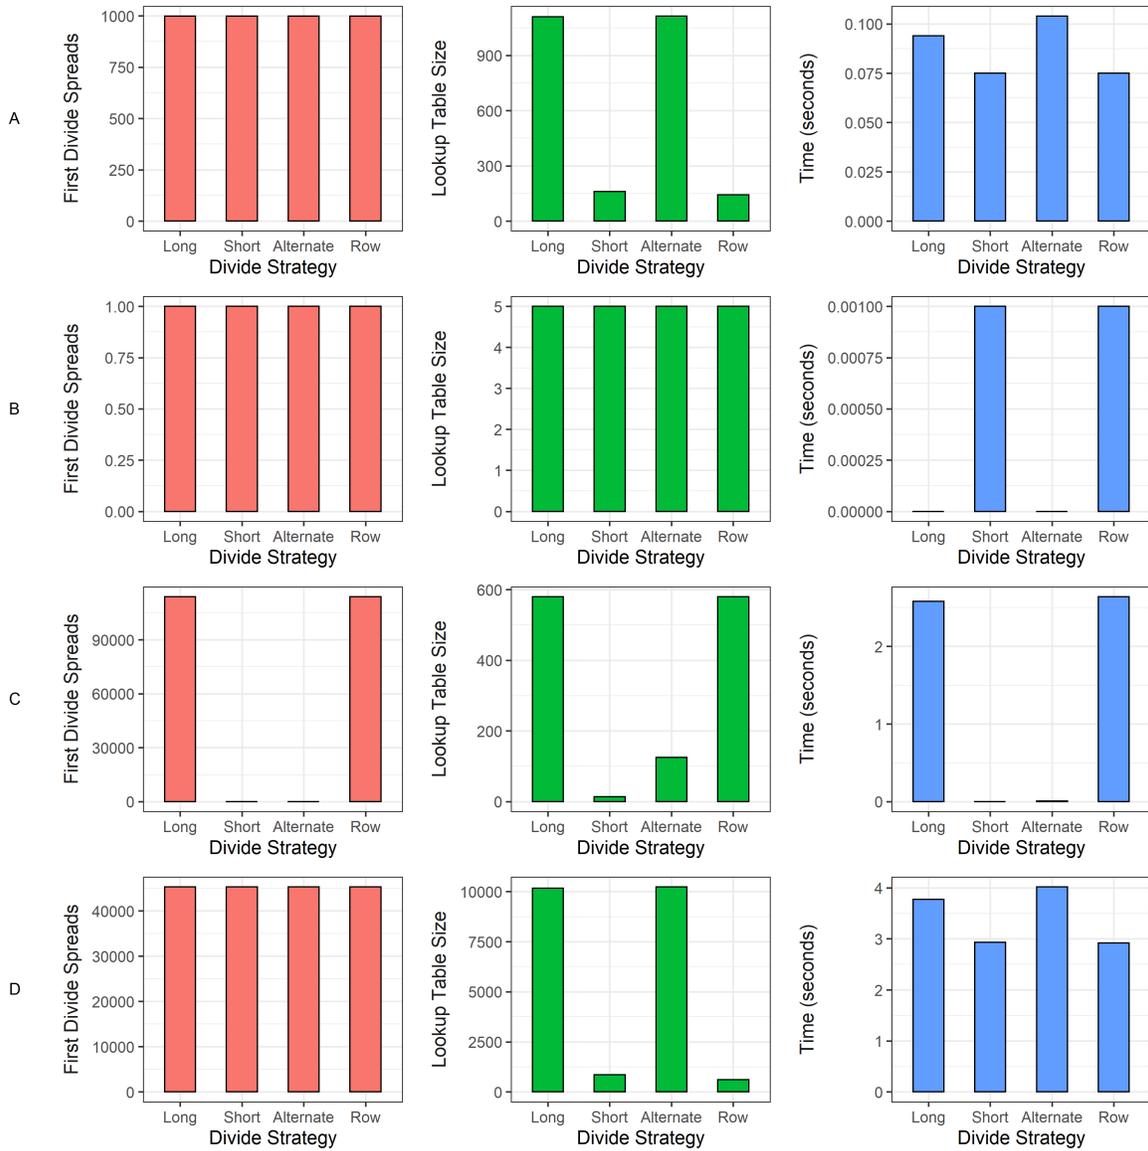
Inspecting lookup table size, however, shows that the number of lookup table entries is lowest when division occurs midway through  $q$ . In this case, the left and right parts are of comparable, if not equal, size, so the likelihood of encountering the same margins more than once is greatly increased. Note that even though the number of recursive calls may be greater, caused by a greater number of spreads, the number of unique margins encountered can still be far less, which is captured by lookup table size.

Finally, the running times for each example show the interplay between increasing spreads and decreasing lookups near a cut point of 0.5 particularly for example C. The shifting between two values in example B is a result of the limited measurement precision for short time intervals.

#### **4.5.5 Division Strategy**

Algorithms 3 and 4 alternate between dividing row and column margins and at each level of the recursion, the rationale being that the total margin length for both parts would be minimized. The divide and conquer approach to counting and sampling need not adhere to this division strategy, however. Indeed, we could choose to always divide the column margins, resulting in narrower and narrower matrices with the same height as the original matrix. Alternately, we could always divide the row margins, always divide the longer margin, or always divide the shorter margin. To examine the impact of different division strategies, we take examples A-D from the previous section, a cut point of 0.5, and apply each division strategy. In each case, we employ the division strategy at all recursive levels.

Figure 4.14 shows measurements in the same three metrics as in the previous section, where the division strategies are: always divide the longer margin (Long), always divide the shorter margin (Short), alternate between row and column margins (Alternate), always divide the row margin (Horizontal). For symmetric matrices like examples A, B, and D, the Long and Alternate strategies are very similar, differing only in how they handle equal length row and column margins. This is reflected in similar first divide spreads, lookup table size, and time between those two strategies on those examples. Surprisingly, the Short strategy has reduced run time for examples A, C, and



**Figure 4.14:** Measurements from the counting process for examples A-D from Table 4.1: A:  $p = q = (4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ ; B:  $p = q = (1, \dots, 1)$  (20 ones); C:  $p = (1, \dots, 1)$  (45 ones),  $q = (5, 5, 5, 4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ ; D:  $p = q = (5, 5, 5, 4, 4, 4, 3, 3, 3, 2, 2, 2, 1, 1, 1)$ . We compare four division strategies: always divide the longer margin (Long), always divide the shorter margin (Short), always alternate row/column divisions (Alternate), always divide the row margin (Row). Note the different scales on each plot, and the fine vertical time scale for example B.

D. The insight gained here is that reducing the size of one of the margins as fast as possible limits the valid spreads for the other margin due to the Gale-Ryser constraints, so repeatedly dividing the shorter margin results in smaller lookup tables compared to alternating or always choosing the longer side. For the square examples (A, B, and D), the Shorter and Horizontal strategies are effectively doing the same thing, so we see very similar performance. Because example C is much taller than it is wide, the Horizontal and Long strategies are the same for the first few steps. Despite the row margins eventually becoming shorter, the Horizontal strategy suffers because of the high number of spreads in the first few recursive levels. These examples suggest that always dividing the shorter margin is preferable to the alternating strategy outlined in Algorithms 3 and 4.

#### 4.5.6 Connection to Miller & Harrison's Approach

We stated earlier that our approach is a generalization of the one-column-at-a-time approach of Miller & Harrison. Indeed, (4.10) reduces to (4.3) if we set the division point  $k = 2$ , so the left part of the division is a single column. More specifically, the sum over non-empty fill classes in (4.10) reduces to the sum over choices of vector  $s$  with  $s_i \leq r_i$  and summing to  $q_i$ . The inner sum of (4.10) reduces to a single term (no sum) since the partitions in each block of  $p$  are limited to size one. The combinatorial factor  $\sigma(p^l, p^r)$  permuting within each block of  $p$  reduces to counting the combinations associated with placing  $r_i$  ones in each block of  $p$  in the column  $q_i$ , which is given by the product of choose functions in (4.3). Finally, the product of left and right recursive calls to  $N$  in (4.10) reduces to a single call in (4.3) because the left count is always one.

## 4.6 Discussion

### 4.6.1 Computational Complexity

Counting and sampling via our divide and conquer approach involves recursive evaluations of (4.10), each of which involves sums over lattice structures and computation of a combinatorial factor. Computing the theoretical worst-case complexity of the algorithm is challenging because of the need to bound the number of terms being summed at each level in the algorithm. The

outer sum covers the space of fills which have at least one valid spread orbit, whose boundaries are hyperplanes arising from two types of constraints: the fill in each block is bounded by that block's size, so  $0 \leq f_i^l \leq |s_i| \times i$  and at least one margin consistent with the fill must satisfy the Gale Ryser conditions. These fills are lattice points in an integral polytope with dimension equal to the number of blocks in  $p$  (the margin being spread). Short of specialized cases (e.g. Pick's formula for dimension 2), few closed form expressions exist for the number of lattice in an integral polytope [69]. The counting problem in general is #P-Hard, but in fixed dimension (as is our case) there exists an algorithm for counting which is polynomial in the number of vertices, or even linear in the number of vertices if special conditions are met [70]. However, these complexities bound how long it takes to *determine the number terms in the outer sum of (4.10)*, not the number of terms themselves, so the polynomial and linear algorithms imply nothing about the computation complexity of Algorithms 3 and 4.

The inner sum in (4.10) is over spread orbits consistent with a given fill. Given a fill  $f^l$ , each block of  $p$  can be viewed as a partition with bounded number of parts and size. Each spread orbit is then a point in the  $d$ -dimensional product lattice, whose boundaries are imposed again by the block and Gale-Ryser constraints. The task of counting integer partitions dates at least to Hardy & Ramanujan [71], who give asymptotic expressions for the number of partitions as the integer approaches infinity. Rademacher later proved an exact formula, but which involves an infinite series [72].

Miller & Harrison [73] gave expressions for computational complexity for their algorithm by bounding the total number of nodes in the directed acyclic graph arising from the recursive evaluation of (4.3). Unfortunately we have not found a straightforward application of their approach to our algorithm.

## 4.6.2 Connection to symmetric polynomials

There is a connection between our approach for counting and taking products of elementary symmetric polynomials. First identify the  $i$ th row with a variable  $x_i$ , for  $i \in \{1, \dots, m\}$ . The

$j$ th column margin,  $q_j$ , indicates a particular elementary symmetric polynomial, i.e. if  $q_j = b$ , then each term in  $e_b(x_1, \dots, x_m)$  indicates a unique way of selecting  $b$  out of  $m$  rows in column  $j$ . The product of the polynomials associated with the columns gives a collection of monomial terms, each of which represents a resulting set of row margins, and the coefficient on each term represents the number of unique choices for the columns that give rise to the row margins. Our objective of counting the number of binary matrices with margins  $p$  and  $q$  maps to identifying the coefficient on the term that corresponds to the desired row margins,  $p$ . Naively taking this approach would be enormously wasteful, because the number of resulting terms could be quite large (bounded above by  $\binom{m}{q_1} \times \dots \times \binom{m}{q_n}$ ), and we would discard all but one of them. However, the concept of dividing the matrix into two smaller parts can be considered in this context, where the product of column-based polynomials is computed separately for the left and right parts of the division. At this stage, some terms could be eliminated from the left and right parts because they don't contribute to the term whose coefficient contains the desired count (e.g. if the exponent on  $x_i$  exceeds  $p_i$ ).

## 4.7 Conclusions

At first glance, the notion of dividing a problem into smaller independent pieces seems quite promising, as the resulting sub-parts are decoupled from each other. The challenge in our approach is in the decoupling of the two sub-parts, which gives rise to a large number of sub-problems to consider, not just two as intuition might suggest. Nevertheless our approach demonstrates that a generalized divide and conquer approach to binary matrix sampling is achievable, though open questions about computational complexity and optimal algorithm design remain.

# Chapter 5

## Fréchet Covariance and MANOVA Tests for Random Objects in Multiple Metric Spaces

### 5.1 Introduction

Much networks research has traditionally addressed research questions pertaining to a single observed network [42, 74], focusing on either global characteristics (e.g. degree distribution, number of triadic relationships) or local characteristics (e.g. node attributes, individual dyadic relationships). Statistics in this context has often been applied at the local scale by treating each node or dyadic relationship as an observation and addressing the dependencies between observations implied by the network structure [75, 76]. At the global scale, however, there is only one network, which is a single observation, so there is little opportunity for statistical modeling without placing assumptions on the generating process [77, 78, 79].

However, the advent of network data comprised of multiple networks led researchers to readdress basic statistical concepts in this new setting, such as the sample mean network, population mean network, and measurement errors [52, 80]. With observed networks themselves now being viewed as realizations of random networks having some probability measure, tools from shape analysis, non-Euclidean data sets, and object oriented data analysis were leveraged [45, 49, 46, 81]. These tools, such as the notion of a Fréchet mean [82], have been used to redefine the concepts of consistency and convergence for random objects. This has led to the development of methods analogous to traditional two sample testing, ANOVA, and regression in service of practitioners dealing with random objects data sets [83].

The purpose of this chapter is to continue in the vein of building basic analysis tools for these more complicated random objects in bounded metric spaces. The appeal of metric spaces is that they encompass several examples of non-Euclidean random objects, including networks, probabil-

ity distributions, and other multivariate objects [50]. On the other hand, metric spaces do not admit an arithmetic mean, voiding many routine statistical models and testing procedures. In contrast to the Euclidean setting, where the observations in a sample each have an absolute position, a sample in metric space can produce only a set of pairwise distances between observations, with no other information about the underlying geometry, if there is one. Without an arithmetic mean, statisticians have turned to the Fréchet mean for a basis from which to build anew classical statistical methods [50, 84]. The next section reviews one such development, the analysis of variance, in preparation for our introduction of a multivariate extension.

In this chapter we propose a definition for Fréchet covariance and establish consistency of the sample estimator. We also discuss various definitions of the Fréchet correlation, with particular attention paid to the interpretation of each. We then propose several statistical tests for differences in group mean and covariance matrices among two or more groups of random objects in bounded metric spaces. Some statistics draw inspiration from existing MANOVA tests, others relying on the Riemannian geometry of symmetric positive definite matrices, and others which use our CLT results. In two scenarios, we perform simulation studies to assess the Type I error of each statistic, as well as the power under various departures from the null hypothesis.

## 5.2 Analysis of Variance in Bounded Metric Spaces

Our work builds upon that of Dubey & Müller [50], which studies Fréchet variance and an ANOVA procedure for differences in means and variances between two or more groups. In this section we briefly review their work, though we have adjusted the notation to be consistent with the rest of this chapter. We begin with the following definitions. Let  $\Omega_s$  be a non-empty metric space with metric  $d_s$  and generic point  $\omega \in \Omega_s$ . We have added the  $s$  subscripts in preparation for adding more metric spaces later. Suppose random object  $X_s \in \Omega_s$  has probability measure  $P$ . Then the population Fréchet mean is given by

$$\mu_s = \arg \min_{\omega \in \Omega_s} E[d_s^2(\omega, X_s)], \quad (5.1)$$

and for an independent and identically distributed sample  $\{X_{si}\}$  with  $i \in \{1, \dots, n\}$ , each  $X_{si}$  having probability measure  $P$ , the corresponding sample mean is

$$\hat{\mu}_s = \arg \min_{\omega \in \Omega_s} \frac{1}{n} \sum_{i=1}^n d_s^2(\omega, X_{si}). \quad (5.2)$$

If the metric space happens to be Euclidean space equipped with the Euclidean  $L^2$  metric, then the Fréchet mean is the same as the arithmetic mean. The Fréchet variance is given by

$$\sigma_s^2 = E[d_s^2(\mu_s, X_s)] \quad (5.3)$$

and the corresponding sample variance is

$$\hat{\sigma}_s^2 = \frac{1}{n} \sum_{i=1}^n d_s^2(\hat{\mu}_s, X_{si}). \quad (5.4)$$

Petersen & Müller established the consistency of the sample Fréchet mean [84], and Dubey & Müller established the consistency and asymptotic normality of the sample Fréchet variance [50]. From these results, Dubey & Müller developed an ANOVA procedure to test for equality of variance and means simultaneously among several groups for a random variable defined on a metric space.

Dubey & Müller's test combines two statistics,  $F_s$  and  $U_s$ , where  $F_s$  is analogous to the  $F_s$  statistic from classical ANOVA tests, and  $U_s$  (which is a "U-statistic" in the sense of [85]) tests for equality of variances. While typically equality of group variances is an assumption to be checked before performing the ANOVA procedure (for example, see [86]), the metric space ANOVA tests against a null of the means and variances being equal among the groups. One advantage of this approach is that it produces an asymptotic test thanks to the asymptotic distribution of  $U_s$  and the fact that  $F_s$  is  $o_p(n^{-1/2})$ . Their proposed test also does not require normality assumptions of the groups, which seems appropriate given the diversity of random objects which are found in the

setting of bounded metric spaces. We outline the major components of Dubey & Müller's test below insofar as they are relevant to our approach later.

Given  $j$  groups denoted by  $j \in \{1, \dots, J\}$ , let observation  $X_{jsi}$  be the  $i$ th observation from group  $j$ , where  $i \in \{1, \dots, n_j\}$ , and assume  $\{X_{jsi}\}$  are independent across all  $i$  and  $j$ , and identically distributed within each group  $j$ . Define the group specific means and variances as

$$\mu_{js} = \arg \min_{\omega \in \Omega_s} E[d_s^2(\omega, X_{jsi})], \quad j \in \{1, \dots, J\} \quad (5.5)$$

$$\sigma_{js}^2 = E[d_s^2(\mu, X_{jsi})], \quad (5.6)$$

and corresponding sample estimators as

$$\hat{\mu}_{js} = \arg \min_{\omega \in \Omega_s} \frac{1}{n_j} \sum_{i=1}^{n_j} d_s^2(\omega, X_{jsi}), \quad j \in \{1, \dots, J\}, \quad (5.7)$$

$$\widehat{\sigma}_{js}^2 = \frac{1}{n_j} \sum_{i=1}^{n_j} d_s^2(\hat{\mu}_{js}, X_{jsi}) \quad j \in \{1, \dots, J\}. \quad (5.8)$$

The hypothesis being tested is

$$H_0 : \quad \mu_{1s} = \dots = \mu_{Js}, \quad \sigma_{1s}^2 = \dots = \sigma_{Js}^2. \quad (5.9)$$

Define the pooled sample mean and variance as

$$\hat{\mu}_{ps} = \arg \min_{\omega \in \Omega_s} \frac{1}{n_j} \sum_{j=1}^J \sum_{i=1}^{n_j} d_s^2(\omega, X_{jsi}), \quad (5.10)$$

$$\widehat{\sigma}_{ps}^2 = \sum_{j=1}^J \frac{1}{n_j} \sum_{i=1}^{n_j} d_s^2(\hat{\mu}_{js}, X_{jsi}). \quad (5.11)$$

Since the sample sizes are not necessarily equal, let  $\gamma_j = n_j/n$  represent the proportion of observations belonging to group  $j$ . Then define the statistics

$$F_s = \widehat{\sigma}_{ps}^2 - \sum_{j=1}^J \gamma_j \widehat{\sigma}_{js}^2 \quad (5.12)$$

$$U_s = \sum_{j < j'} \frac{\gamma_j \gamma_{j'}}{\widehat{\text{Var}}(\widehat{\sigma}_{js}^2) \widehat{\text{Var}}(\widehat{\sigma}_{j's}^2)} \left( \widehat{\sigma}_j^2 - \widehat{\sigma}_{j'}^2 \right)^2, \quad (5.13)$$

where

$$\widehat{\text{Var}}(\widehat{\sigma}_{js}^2) = \frac{1}{n_j} \sum_{i=1}^{n_j} d_s^4(\hat{\mu}_{js}, X_{jsi}) - \left\{ \frac{1}{n_j} \sum_{i=1}^{n_j} d_s^2(\hat{\mu}_{js}, X_{jsi}) \right\}^2, \quad j \in \{1, \dots, J\} \quad (5.14)$$

is an estimate of the variance of  $\widehat{\sigma}_{js}^2$ . The  $F_s$  statistic can be seen as estimating the between group variances by subtracting the average within group variance from the the total variance, and is analogous to the numerator of the Euclidean ANOVA statistic. The  $U_s$  statistic is similar to the test statistic underlying Levene's test, and when appropriately scaled, is asymptotically  $\chi_{(J-1)}^2$  as  $n \rightarrow \infty$ , assuming each  $\gamma_j$  converges as well. To perform the metric space ANOVA, Dubey & Müller propose a test statistic  $T_s$  defined as

$$T_s = \frac{nU_s}{\sum_{j=1}^J \frac{\gamma_j}{\widehat{\text{Var}}(\widehat{\sigma}_{js}^2)}} + \frac{nF_s^2}{\sum_{j=1}^J \gamma_j^2 \widehat{\text{Var}}(\widehat{\sigma}_{js}^2)}. \quad (5.15)$$

Under some existence and uniqueness assumptions for the population means and variances,  $T_s$  converges in distribution to  $\chi_{(J-1)}^2$  under the null hypothesis as  $n \rightarrow \infty$ .

Simulation studies show that Dubey & Müller's metric space ANOVA performs well when considering scale-free networks, probability distributions, and even multivariate Euclidean data (each object being an observation in  $\mathbb{R}^5$ ). Our purpose in the rest of this chapter is to extend this approach to the setting of two or more metric spaces, to test for differences in two or more groups measured in multiple metric spaces, analogous to a metric space MANOVA procedure. We begin this with the introduction of Fréchet covariance in the next section.

### 5.3 Fréchet Covariance

Consider now two distinct metric spaces,  $\Omega_s$  with metric  $d_s$  and  $\Omega_{s'}$  with metric  $d_{s'}$ , and random object vector  $(X_s, X_{s'}) \in \Omega_s \times \Omega_{s'}$  with joint distribution  $P_{ss'}$ . Strictly speaking,  $(X_s, X_{s'})$  is an ordered pair not necessarily belonging to a vector space, but we use the term “vector” for convenience. For compactness we will omit the subscripts of the metrics  $d_s$  and  $d_{s'}$  and use simply  $d$  when it is obvious from context which one we mean. In particular,  $d_s$  measures distances in  $\Omega_s$  and  $d_{s'}$  measures distances in  $\Omega_{s'}$ . We define the population mean vector as  $(\mu_s, \mu_{s'})$ , with elements given by:

$$\mu_* = \arg \min_{\omega \in \Omega_*} E[d^2(\omega, X_*)], \quad * \in \{s, s'\}. \quad (5.16)$$

Given a set of independent and identically distributed observations  $\{(X_{si}, X_{s'i})\}$  with  $i \in \{1, \dots, n\}$ , the sample mean vector  $(\hat{\mu}_s, \hat{\mu}_{s'})$  has elements

$$\hat{\mu}_* = \arg \min_{\omega \in \Omega_*} \frac{1}{n} \sum_{i=1}^n d^2(\omega, X_{*i}), \quad * \in \{s, s'\}. \quad (5.17)$$

The population and sample variance vectors,  $(\sigma_s^2, \sigma_{s'}^2)$  and  $(\widehat{\sigma}_s^2, \widehat{\sigma}_{s'}^2)$ , are defined analogously to (5.3)-(5.4) with elements

$$\sigma_* = \arg \min_{\omega \in \Omega_*} E[d^2(\omega, X_*)] \quad * \in \{s, s'\}, \quad (5.18)$$

$$\widehat{\sigma}_*^2 = \frac{1}{n} \sum_{i=1}^n d^2(\hat{\mu}_*, X_{*i}) \quad * \in \{s, s'\}. \quad (5.19)$$

Now that there are multiple metric spaces, we may consider the dependency which may exist between  $X_s$  and  $X_{s'}$ , the two elements of the random object vector. Before proposing how to do so, however, we first address exactly what dependence means in a two-metric-space context.

Because general metric spaces lack a coordinate system which endows direction, it does not make sense to talk about the tendency for  $X_s$  to *increase* or *decrease* as  $X_{s'}$  *increases*. However,

$X_s$  and  $X_{s'}$  can be characterized by their distance from their respective means, and one sensible notion of dependence would be whether  $X_s$  tends to move away from or toward its mean as  $X_{s'}$  moves away from its mean. This suggests the following definitions for the population and sample Fréchet covariance:

$$\sigma_{ss'} = E[d(X_s, \mu_s)d(X_{s'}, \mu_{s'})] \quad (5.20)$$

$$\hat{\sigma}_{ss'} = \frac{1}{n} \sum_{i=1}^n d(X_{si}, \hat{\mu}_s)d(X_{s'i}, \hat{\mu}_{s'}), \quad (5.21)$$

which we expect to be large if  $X_s$  and  $X_{s'}$  tend to deviate from their mean together, and small if not. This definition, however, always produces a positive covariance, since distances are always positive, so it does not clearly indicate when  $X_s$  and  $X_{s'}$  both tend to move away from their respective means at the same time, or if one tends to move away from its mean as the other approaches its mean. Indeed, unlike the Fréchet variance, which reduces to the classical definition when applied to Euclidean objects with the Euclidean metric, the definition in (5.20) reduces to  $E[|X_s - \mu_s||X_{s'} - \mu_{s'}|]$ , which is not the Euclidean covariance.

Using the definition for covariance in (5.20) allows us to express the covariance matrix of  $(X_s, X_{s'})$  as  $\Sigma = E[d_\mu(X_s, X_{s'})d_\mu(X_s, X_{s'})^T]$ , where the distance vector  $d_\mu(X_s, X_{s'})^T = (d(X_s, \mu_s), d(X_{s'}, \mu_{s'}))$ . Both the population and sample versions are positive semi-definite, positive definite if  $X_s$  and  $X_{s'}$  are non-degenerate, as shown later in Lemma 3.

An alternative definition of covariance is motivated by the desire that the sign of the covariance be meaningful, which we call the *centered Fréchet covariance*:

$$\sigma'_{ss'} = E \{ [d(X_s, \mu_s) - E(d(X_s, \mu_s))] [d(X_{s'}, \mu_{s'}) - E(d(X_{s'}, \mu_{s'}))] \} \quad (5.22)$$

$$\hat{\sigma}'_{ss'} = \frac{1}{n} \sum_{i=1}^n [d(X_{si}, \hat{\mu}_s) - \bar{d}_s] [d(X_{s'i}, \hat{\mu}_{s'}) - \bar{d}_{s'}] \quad (5.23)$$

where  $\bar{d}_s = \frac{1}{n} \sum_{i=1}^n d(X_{si}, \hat{\mu}_s)$  and  $\bar{d}_{s'} = \frac{1}{n} \sum_{i=1}^n d(X_{s'i}, \hat{\mu}_{s'})$ . Equation (5.22) is simply the Euclidean covariance between  $d(X_s, \mu_s)$  and  $d(X_{s'}, \mu_{s'})$ , and (5.23) is the sample analog. If both the

$X_s$  and  $X_{s'}$  distances increase / decrease together, this will be positive. Unlike (5.20) and (5.21), equations (5.22) and (5.23) can be negative, when the  $X_s$  distances and  $X_{s'}$  distances have an inverse relationship. This is arguably the nicer interpretation, though it does not align as well with Dubey & Müller's definition of Fréchet variance in (5.3).

For both covariance definitions, we can define correlation by normalizing by the appropriate standard deviations. For the definitions in (5.20) and (5.21), we divide by the square root of Dubey & Müller's variance to produce the *non-centered Fréchet correlation*:

$$\rho_{ss'} = \frac{E[d(X_s, \mu_s)d(X_{s'}, \mu_{s'})]}{\sqrt{E[d^2(X_s, \mu_s)]E[d^2(X_{s'}, \mu_{s'})]}}, \quad (5.24)$$

$$\hat{\rho}_{ss'} = \frac{\frac{1}{n} \sum_{i=1}^n d(X_{si}, \hat{\mu}_s)d(X_{s'i}, \hat{\mu}_{s'})}{\sqrt{\frac{1}{n} \sum_{i=1}^n d^2(X_{si}, \hat{\mu}_s) \frac{1}{n} \sum_{i=1}^n d^2(X_{s'i}, \hat{\mu}_{s'})}}. \quad (5.25)$$

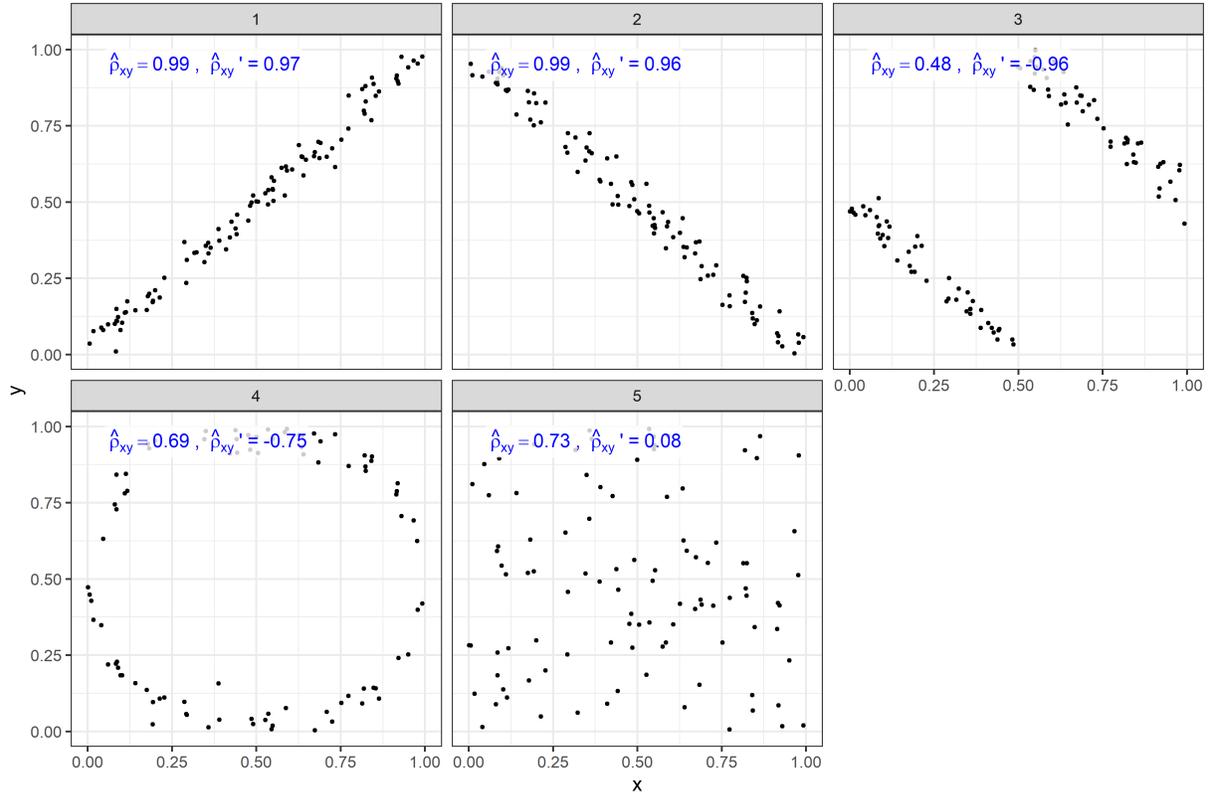
For the definitions in (5.22) and (5.23), we simply use the Euclidean version of correlation on the distances to produce the *centered Fréchet correlation*:

$$\rho'_{ss'} = \frac{E[(d(X_s, \mu_s) - Ed(X_s, \mu_s))(d(X_{s'}, \mu_{s'}) - Ed(X_{s'}, \mu_{s'}))]}{\sqrt{E[(d(X_s, \mu_s) - Ed(X_s, \mu_s))^2]E[(d(X_{s'}, \mu_{s'}) - Ed(X_{s'}, \mu_{s'}))^2]}}, \quad (5.26)$$

$$\hat{\rho}'_{ss'} = \frac{\frac{1}{n} \sum_{i=1}^n (d(X_{si}, \hat{\mu}_s) - \bar{d}_s)(d(X_{s'i}, \hat{\mu}_{s'}) - \bar{d}_{s'})}{\sqrt{\frac{1}{n} \sum_{i=1}^n (d(X_{si}, \hat{\mu}_s) - \bar{d}_s)^2 \frac{1}{n} \sum_{i=1}^n (d(X_{s'i}, \hat{\mu}_{s'}) - \bar{d}_{s'})^2}}. \quad (5.27)$$

The non-centered correlation lies in  $[0, 1]$ , and the centered version is in  $[-1, 1]$ .

To illustrate the behavior of each definition of correlation, we simulated data from two metric spaces, both Euclidean, and each equipped with the  $L^1$  metric. The first random variable is  $X_1 \in \Omega_1 = \mathbb{R}$  and the second random variable is  $X_2 \in \Omega_2 = \mathbb{R}$ . In a Euclidean sense, we are sampling bivariate data in  $\mathbb{R}^2$ . We simulate data under 5 different scenarios. For all scenarios the  $X_1$  values are generated from a Uniform(0, 1), distribution. The  $X_2$  values are generated either as a noisy function of  $X_1$  (scenarios 1-4) or from a Uniform(0, 1) distribution independent of  $X_1$  (scenario 5). For each scenario, we compute both the non-centered and centered sample Fréchet correlation, treating each bivariate observation as an object vector,  $(X_{1i}, X_{2i})$ . Figure 5.1 shows a scatter plot



**Figure 5.1:** Non-centered ( $\hat{\rho}_{12}$ ) and centered ( $\hat{\rho}'_{12}$ ) Fréchet correlation for 5 Euclidean bivariate scenarios. Scenarios 1 and 2 illustrate how both definitions show strong positive correlation when Euclidean correlation is either strongly positive or strongly negative. Scenarios 3 and 4 highlight how the centered Fréchet correlation is negative when  $X_1$  and  $X_2$  tend to move in opposite directions relative to their means. Scenario 5 shows that the centered Fréchet correlation near zero when  $X_1$  and  $X_2$  are independent.

of observations for each scenario, along with the non-centered and centered Fréchet correlations of each. In terms of Euclidean correlation, scenarios 1 and 2 show strong positive and negative correlation respectively, scenario 3 shows weaker positive correlation, and scenarios 4 and 5 have correlation near zero.

The non-centered Fréchet correlation measures how much the  $X_1$  and  $X_2$  deviate from their mean at the same time. In this regard, scenarios 1 and 2 both exhibit strong positive non-centered Fréchet correlation, as the points deviate from  $(0.5, 0.5)$  together. In scenarios 3 and 4,  $X_1$  values far from their mean are paired with  $X_2$  values near their mean, and vice versa, resulting in smaller (but still positive) non-centered Fréchet correlations. Finally, scenario 5 shows how the non-centered Fréchet correlation disagrees with the classical Euclidean definition.

The centered Fréchet correlation focuses on the degree to which the distance from  $X_1$  to its mean is above or below average, and similarly for  $X_2$ . For scenarios 1 and 2,  $X_1$  and  $X_2$  are both above and below average distance simultaneously, leading to strong positive correlation. In scenario 3,  $X_1$  is near its mean when  $X_2$  is far away from its and vice versa, leading to a strong negative centered Fréchet correlation, similarly for scenario 4. In scenario 5, when  $X_1$  and  $X_2$  are generated independently, the centered correlation is near zero. Scenarios 3-5 in particular show how the centered version of correlation retains some desirable characteristics of Euclidean correlation and is arguably easier to interpret than the non-centered version.

The examples in Figure 5.1 illustrate how both definitions of Fréchet correlation differ from the Euclidean notion when looking at Euclidean data, but this is expected since metric spaces force us to redefine what it means to be correlated using only distances. In the rest of this paper, we will consider the non-centered version of Fréchet covariance, as it integrates nicely with Dubey & Müller's Fréchet variance.

## 5.4 Properties of Fréchet Covariance

In this section we define the (non-centered) Fréchet covariance matrix, and establish consistency for the (non-centered) Fréchet covariance defined in (5.21). We begin by considering the Fréchet covariance matrix constructed using Dubey & Müller's Fréchet variances on the diagonal and our non-centered Fréchet covariance on the off-diagonals. Given  $S$  metric spaces  $\{\Omega_s\}$  for  $s \in \{1, \dots, S\}$ , and random object vector  $\mathbf{X} = (X_1, \dots, X_S) \in \Omega_1 \times \dots \times \Omega_S$  with probability measure  $P : \Omega_1 \times \dots \times \Omega_S \rightarrow \mathbb{R}_{\geq 0}$  and Fréchet mean vector  $(\mu_1, \dots, \mu_S)$ . Define the  $(s, s')$  element of the population covariance matrix  $\Sigma$  and the sample covariance matrix  $\hat{\Sigma}$  as

$$\Sigma(s, s') = \begin{cases} \sigma_s^2 & s = s' \\ \sigma_{ss'} & s \neq s' \end{cases}, \quad (5.28)$$

$$\widehat{\Sigma}(s, s') = \begin{cases} \widehat{\sigma}_s^2 & s = s' \\ \widehat{\sigma}_{ss'} & s \neq s' \end{cases}. \quad (5.29)$$

Defining the vector of distances  $d_\mu(\mathbf{X}) = (d(X_1, \mu_1), \dots, d(X_S, \mu_S))^T$ , we can write  $\Sigma = E[d_\mu d_\mu^T]$ . Similarly, if  $X_{si}$  is the  $i$ th observation in the  $s$ th metric space and  $\hat{\mu}_s$  the sample Fréchet mean in the  $s$ th metric space, define  $d_{\hat{\mu}i} = (d(X_{1i}, \hat{\mu}_1), \dots, d(X_{Si}, \hat{\mu}_S))^T$  so that  $\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n d_{\hat{\mu}i} d_{\hat{\mu}i}^T$ . This makes it clear the both  $\Sigma$  and  $\widehat{\Sigma}$  are positive semi-definite:

**Lemma 3.** The Fréchet population covariance matrix  $\Sigma$  and sample covariance matrix  $\widehat{\Sigma}$  are both positive semi-definite. If the elements of  $d_\mu$  are non-degenerate, then  $\Sigma$  is positive definite. If the vectors  $\{d_{\hat{\mu}1}, \dots, d_{\hat{\mu}n}\}$  span  $\mathbb{R}^S$ , then  $\widehat{\Sigma}$  is positive definite.

*Proof of Lemma 3.* Given vector  $z \in \mathbb{R}^S$ , then

$$z^T \Sigma z = z^T E[d_\mu d_\mu^T] z = E[z^T d_\mu d_\mu^T z] = E[|d_\mu^T z|^2] \geq 0. \quad (5.30)$$

The quantity is strictly greater than zero if the elements of  $d_\mu$  are non-degenerate. Similarly for  $\widehat{\Sigma}$ ,

$$z^T \widehat{\Sigma} z = \frac{1}{n} \sum_{i=1}^n z^T d_{\hat{\mu}i} d_{\hat{\mu}i}^T z = \frac{1}{n} \sum_{i=1}^n |d_{\hat{\mu}i}^T z|^2 \geq 0. \quad (5.31)$$

If the vectors  $\{d_{\hat{\mu}1}, \dots, d_{\hat{\mu}n}\}$  span  $\mathbb{R}^D$ , then any vector  $z$  can be written as  $a_1 d_{\hat{\mu}1} + \dots + a_n d_{\hat{\mu}n}$  so that  $z^T z = a_1 d_{\hat{\mu}1}^T z + \dots + a_n d_{\hat{\mu}n}^T z$ . Hence if (5.31) is equal to zero, then  $d_{\hat{\mu}i}^T z = 0$  for each  $i$ , implying  $z^T z = 0$ , so  $z = 0$ .  $\square$

From now on we assume that all population and sample covariances are strictly positive definite. The remaining results of this section will focus on the the covariance of  $(X_s, X_{s'}) \in \Omega_s \times \Omega_{s'}$ , and require the following assumptions:

A.1 The population and sample Fréchet means in each metric space,  $(\mu_s, \mu_{s'})$  and  $(\hat{\mu}_s, \hat{\mu}_{s'})$ , exist and are unique, and for any  $\epsilon > 0$ ,  $\inf_{d(\omega, \mu_s) > \epsilon} E[d^2(\omega, X_s)] > E[d^2(\mu_s, X_s)]$  and

$\inf_{d(\omega, \mu_{s'}) > \epsilon} E[d^2(\omega, X_{s'})] > E[d^2(\mu_{s'}, X_{s'})]$ . In other words, in each metric space, the expected distance from the mean is uniquely minimum compared to the expected distance from any other point  $\omega$  which is further than  $\epsilon$  from the mean.

This assumption is analogous to one made by Dubey & Müller in [50], and Petersen & Müller show in [84] that it is satisfied by the following two metric spaces: univariate probability distributions with compact support and finite second moment equipped with the  $L^2$ -Wasserstein metric, the set of correlation matrices of a fixed dimension equipped with the Frobenius metric. Dubey & Müller additionally claim that the set of graph Laplacians of connected, undirected, simple graphs of a fixed dimension also satisfy the assumptions [50]. Under Assumptions A.1, we establish consistency of the sample Fréchet covariance  $\widehat{\sigma}_{ss'}$ .

**Theorem 6 (Consistency).** Under Assumption 1 above, the sample covariance converges to the population covariance in probability, that is  $\widehat{\sigma}_{ss'} \xrightarrow{p} \sigma_{ss'}$  as  $n \rightarrow \infty$ .

The proof of Theorem 6 is provided in Appendix A.

## 5.5 Tests for Differences in Means and Variance / Covariance Matrices

Our goal is to investigate various tests for differences between  $J$  groups in  $S$  metric spaces, as an analog to the MANOVA procedure. We will consider various tests based on several test statistics, but in all cases the null hypothesis we wish to test is simultaneous equality of the mean vector across groups and equality of covariance matrix across groups. Letting  $\boldsymbol{\mu}_j = (\mu_{j1}, \dots, \mu_{jS})^T$  and  $\Sigma_j$  be the mean vector and covariance matrix for group  $j$  respectively, then the null hypothesis can be expressed as

$$H_0 : \quad \boldsymbol{\mu}_1 = \dots = \boldsymbol{\mu}_J, \quad \Sigma_1 = \dots = \Sigma_J. \quad (5.32)$$

If the groups deviate from the null hypothesis, we expect the *sample* means and *sample* covariance matrices to deviate from one another, so we seek test statistics which quantify these differences.

### 5.5.1 Adaptations of Classical MANOVA Statistics

Several Euclidean MANOVA tests, including Wilks' lambda [87], Pillai & Bartlett's trace [88], Lawley & Hotelling's trace [89], and Roy's Root [90], test for differences in means by relying on a partition of the total sums of squares and cross products matrix into a between group component (the treatment) and a within group component (the residuals). Such a partition is possible in Euclidean space, thanks to the Pythagorean theorem, but not always possible in generic metric spaces. We can circumvent this problem, however, following the approach of Dubey & Müller's  $F_s$  statistic, which compares the pooled covariance matrix to the weighted mean of group covariance matrices.

Let  $\hat{\sigma}_{jss'}$  be the covariance of group  $j$ , and  $\hat{\sigma}_{pss'}$  be the pooled covariance. Define the  $(s, s')$  element of the population and sample pooled covariance matrices as

$$\Sigma_p(s, s') = \begin{cases} \sigma_{ps}^2 & s = s' \\ \sigma_{pss'} & s \neq s' \end{cases}, \quad (5.33)$$

$$\hat{\Sigma}_p(s, s') = \begin{cases} \hat{\sigma}_{ps}^2 & s = s' \\ \hat{\sigma}_{pss'} & s \neq s' \end{cases}. \quad (5.34)$$

Similarly, define the  $(s, s')$  element of the population and sample group weighted mean covariance matrices as

$$\Sigma_g(s, s') = \begin{cases} \sum_{j=1}^J \gamma_j \sigma_{js}^2 & s = s' \\ \sum_{j=1}^J \gamma_j \sigma_{jss'} & s \neq s' \end{cases}, \quad (5.35)$$

$$\hat{\Sigma}_g(s, s') = \begin{cases} \sum_{j=1}^J \gamma_j \hat{\sigma}_{js}^2 & s = s' \\ \sum_{j=1}^J \gamma_j \hat{\sigma}_{jss'} & s \neq s' \end{cases}. \quad (5.36)$$

Whereas  $\Sigma_p$  and  $\widehat{\Sigma}_p$  capture the total variation across all groups,  $\Sigma_g$  and  $\widehat{\Sigma}_g$  capture the within group variation. When all groups share the same mean vector and covariance matrix, then  $\Sigma_p = \Sigma_g$ , so  $\widehat{\Sigma}_p$  will be close to  $\widehat{\Sigma}_g$ . As the group mean vectors deviate from one another,  $\widehat{\Sigma}_p$  will inflate while  $\widehat{\Sigma}_g$  will remain more or less the same.

It turns out that we can mimic Euclidean MANOVA statistics using  $\widehat{\Sigma}_p$  and  $\widehat{\Sigma}_g$ , without having an explicit residuals term. For instance, the Pillai-Bartlett trace [88] can be expressed as the trace of  $H(H+E)$ , where  $H$  is the between-group sums of squares and cross product matrix, and  $E$  is the residual sums of squares and cross product matrix. Consider starting with the pooled covariance,  $\widehat{\Sigma}_p$ , and subtracting the mean within group covariance,  $\widehat{\Sigma}_g$ . This is similar to the idea of subtracting the within group sum of squares from the total sum of squares, leaving only a between group effect. Therefore we may loosely equate  $\widehat{\Sigma}_p - \widehat{\Sigma}_g$  to  $H$ , and similarly loosely equate  $\widehat{\Sigma}_g$  to  $E$ , ignoring obvious normalizing constants. Translating the Pillai-Bartlett trace into our context then produces:

$$\Lambda_{pillai} = \text{tr}[H(H + E)^{-1}] = \text{tr} \left[ \left( \widehat{\Sigma}_p - \widehat{\Sigma}_g \right) \widehat{\Sigma}_p^{-1} \right] \quad (5.37)$$

$$= \text{tr}(I - \widehat{\Sigma}_g \widehat{\Sigma}_p^{-1}) = \sum_i \lambda_i \left( I - \widehat{\Sigma}_g \widehat{\Sigma}_p^{-1} \right), \quad (5.38)$$

Where  $\lambda_i(A)$  denotes the  $i^{\text{th}}$  eigenvalue of the matrix  $A$ . Under the null hypothesis,  $\widehat{\Sigma}_g \widehat{\Sigma}_p^{-1}$  is close to the identity matrix  $I$ , and  $\Lambda_{pillai}$  will be close to zero.

Another possible adaptation of a classical MANOVA test would be to perform a Euclidean MANOVA on distances between observations and sample means. That is, for group  $j$  and observation  $i$ , define the vector of distances  $d_{\hat{\mu}_j i} = (d(X_{j1i}, \hat{\mu}_{j1}), \dots, d(X_{jSi}, \hat{\mu}_{jS}))^T$ . each vector summarizes an observation based on its distance to its respective group sample mean. We can then apply the Pillai-Bartlett trace directly to the distances to compare the groups. We call this statistic  $\Lambda_{pillai,d}$ . Clearly, if any of the group variances changes, the average distance must be changing as well, which should be detected by  $\Lambda_{pillai,d}$ .

## 5.5.2 Statistics Based on Riemannian Geometry

As was mentioned above, as the group mean vectors deviate from one another,  $\hat{\Sigma}_p$  will inflate while  $\hat{\Sigma}_g$  will remain more or less the same, so a natural test for difference in means would be to compare them to one another.

Since  $\hat{\Sigma}_g$  is the sum of positive-definite matrices, it is positive definite, whereas  $\hat{\Sigma}_p$  is positive definite by Lemma 3. The space of symmetric positive definite (SPD) matrices is a manifold, so the so  $\hat{\Sigma}_p$  and  $\hat{\Sigma}_g$  can be thought of as points on the manifold. This manifold can be endowed with a Riemannian metric which defines the distance between points in any *tangent space* incident to the manifold. To avoid confusion, note that the Riemannian metric does not measure distances between points on the manifold (i.e. SPD matrices) directly, but the Riemannian metric can be integrated along a path on the manifold to determine the length on the manifold, and the distance between two points on the manifold can be taken as the length of the shortest path between them. For more details on manifolds and Riemannian metrics, refer to Chapter 2.7. The shortest path length between two SPD matrices provides a means to compare  $\hat{\Sigma}_g$  to  $\hat{\Sigma}_p$  and ultimately test for differences between the groups.

We define the following statistics which compare the within group variation  $\hat{\Sigma}_g$  to the total variation  $\hat{\Sigma}_p$ :

$$R_{\mu,Euc} = d_{Euc}(\hat{\Sigma}_p, \hat{\Sigma}_g) = \sqrt{\sum_i \lambda_i^2 (\hat{\Sigma}_p - \hat{\Sigma}_g)} \quad (5.39)$$

$$R_{\mu,AIRM} = d_{AIRM}(\hat{\Sigma}_p, \hat{\Sigma}_g) = \sqrt{\sum_i \log^2 [\lambda_i(\hat{\Sigma}_p^{-1}\hat{\Sigma}_g)]} \quad (5.40)$$

$$R_{\mu,LERM} = d_{LERM}(\hat{\Sigma}_p, \hat{\Sigma}_g) = \sqrt{\sum_i \lambda_i^2 (\text{Log}(\hat{\Sigma}_p) - \text{Log}(\hat{\Sigma}_g))} \quad (5.41)$$

All of these statistics are expected to be small if all groups have the same population mean vector and will increase as the mean vectors deviate from each other. The  $d_{LERM}$  and  $d_{AIRM}$  distances are particularly sensitive to changes in determinant, which will certainly change for  $\Sigma_p$  as the mean vectors grow further apart, so we anticipate that  $R_{\mu,LERM}$  and  $R_{\mu,AIRM}$  will be particularly

sensitive to departures from the null hypothesis in (5.32). Later we will integrate each of the statistics from this section into a global test for differences in means or covariance matrices.

The distances  $d_{Euc}$ ,  $d_{AIRM}$  and  $d_{LERM}$  defined respectively in (2.3), (2.4), and (2.5) can also be used to compare the group covariance matrices in a pairwise fashion, in order to determine if any covariance matrix differs from the others. Since the sample covariance matrices are each SPD, the use of these Riemannian metric based distances is appropriate. This is similar in spirit to Dubey & Müller's  $U_s$  statistic, which performs pairwise comparisons of group Fréchet variances. Like  $U_s$ , we consider the weighted average of all pairwise distances between group covariance matrices, where the weights are the product of the proportions of observations in the respective groups:

$$R_{\Sigma, Euc} = \sum_{j < j'} \gamma_j \gamma_{j'} d_{Euc} \left( \widehat{\Sigma}_j, \widehat{\Sigma}_{j'} \right) = \sum_{j < j'} \gamma_j \gamma_{j'} \sqrt{\sum_i \lambda_i^2 \left( \widehat{\Sigma}_j - \widehat{\Sigma}_{j'} \right)} \quad (5.42)$$

$$R_{\Sigma, AIRM} = \sum_{j < j'} \gamma_j \gamma_{j'} d_{AIRM} \left( \widehat{\Sigma}_j, \widehat{\Sigma}_{j'} \right) = \sum_{j < j'} \gamma_j \gamma_{j'} \sqrt{\sum_i \log^2 \left[ \lambda_i \left( \widehat{\Sigma}_j^{-1} \widehat{\Sigma}_{j'} \right) \right]} \quad (5.43)$$

$$R_{\Sigma, LERM} = \sum_{j < j'} \gamma_j \gamma_{j'} d_{LERM} \left( \widehat{\Sigma}_j, \widehat{\Sigma}_{j'} \right) = \sum_{j < j'} \gamma_j \gamma_{j'} \sqrt{\sum_i \lambda_i^2 \left( \text{Log}(\widehat{\Sigma}_j) - \text{Log}(\widehat{\Sigma}_{j'}) \right)} \quad (5.44)$$

Each of these quantities is expected to be small when the population covariance matrices of all groups are equal, and grow as the covariance matrices differ from each other. Since each type of distance is rooted in a different underlying geometry, covariance matrices which appear close in one geometry may appear further apart in another geometry, and vice versa. Therefore we expect the three statistics introduced in this section to have relative strengths and weaknesses for the types of departures from equality of group covariance matrices they can detect well. We anticipate that the statistics using  $d_{AIRM}$  and  $d_{LERM}$  will be more sensitive to differences in matrix determinant, while the statistic using  $d_{Euc}$  may be better equipped to detect differences between matrices of similar determinants.

Comparing the covariance matrices between groups implicitly compares the metric space to metric space dependence of the random object vectors across groups. Another way to compare dependence between groups is to compare their centered Fréchet correlations. Define the  $(s, s')$

element of the centered sample correlation matrix  $\hat{P}_j$  of group  $j$  be defined as

$$\hat{P}_j(s, s') = \begin{cases} 1 & s = s' \\ \hat{\rho}'_{jss'} & s \neq s' \end{cases}, \quad (5.45)$$

where  $\hat{\rho}'_{jss'}$  is the sample centered covariance between metric spaces  $s$  and  $s'$  in group  $j$ . Since the centered Fréchet correlation is just the classical Euclidean correlation applied to the distances of the random objects from their means,  $\hat{P}_j$  is necessarily positive semi-definite, so we can use the same approach to compare correlation matrices as we used to compare covariance matrices above:

$$R_{P,Euc} = \sum_{j < j'} \gamma_j \gamma_{j'} d_{Euc}(\hat{P}_j, \hat{P}_{j'}) = \sum_{j < j'} \gamma_j \gamma_{j'} \sqrt{\sum_i \lambda_i^2 (\hat{P}_j - \hat{P}_{j'})} \quad (5.46)$$

$$R_{P,AIRM} = \sum_{j < j'} \gamma_j \gamma_{j'} d_{AIRM}(\hat{P}_j, \hat{P}_{j'}) = \sum_{j < j'} \gamma_j \gamma_{j'} \sqrt{\sum_i \log^2 [\lambda_i (\hat{P}_j^{-1} \hat{P}_{j'})]} \quad (5.47)$$

$$R_{P,LERM} = \sum_{j < j'} \gamma_j \gamma_{j'} d_{LERM}(\hat{P}_j, \hat{P}_{j'}) = \sum_{j < j'} \gamma_j \gamma_{j'} \sqrt{\sum_i \lambda_i^2 (\text{Log}(\hat{P}_j) - \text{Log}(\hat{P}_{j'}))} \quad (5.48)$$

It is possible that comparing correlations directly will be more sensitive to changes in dependence than comparing covariance matrices, so the correlation based statistics are worth consideration.

### 5.5.3 Statistics Based on the Fréchet ANOVA

Just as Dubey and Müller defined the  $T_s$  statistic to compare group means and variances, we may analogously define the  $T_{ss'}$  statistic as

$$T_{ss'} = \frac{nU_{ss'}}{\sum_{j=1}^J \frac{\gamma_j}{\widehat{\text{Var}}(\hat{\sigma}_{jss'})}} + \frac{nF_{ss'}^2}{\sum_{j=1}^J \gamma_j^2 \widehat{\text{Var}}(\hat{\sigma}_{jss'})}, \quad (5.49)$$

where

$$F_{ss'} = \hat{\sigma}_{pss'} - \sum_{j=1}^J \gamma_j \hat{\sigma}_{jss'} \quad (5.50)$$

$$U_{ss'} = \sum_{j < j'} \frac{\gamma_j \gamma_{j'}}{\widehat{\text{Var}}(\widehat{\sigma}_{jss'}) \widehat{\text{Var}}(\widehat{\sigma}_{j'ss'})} (\widehat{\sigma}_{jss'} - \widehat{\sigma}_{j'ss'})^2, \quad (5.51)$$

Unlike  $T_s$ , we do not have asymptotic results for  $T_{ss'}$ , though if the Fréchet covariance is indeed asymptotically normal, then under the null hypothesis  $T_{ss'}$  would be asymptotically  $\chi_{J-1}^2$  as  $n \rightarrow \infty$ , the proof being identical to the proof of Theorem 2 in [50]. Based on simulation studies, we suspect asymptotic normality of the Fréchet covariance may hold at least in some circumstances, suggesting we may be justified in comparing  $T_{ss'}$  to a  $\chi_{J-1}^2$  distribution to test the null hypothesis. We explore the validity of this approach in a few example scenarios below.

In the preceding three sections we introduced several statistics, two inspired by the classical MANOVA test, three comparing pooled and group mean covariance matrices, three comparing group covariance matrices, three comparing group centered correlation matrices, and one analogous to the statistic involved in the Fréchet ANOVA. In the next section we will employ these statistics in the construction of tests for the null hypothesis in (5.32).

#### 5.5.4 Testing the Null Hypothesis

Broadly speaking, we will consider three main approaches to testing the null hypothesis in (5.32). The first approach will combine the Riemannian mean, covariance matrix, and correlation matrix based statistics from the previous two sections into a global permutation-based test, achieving the desired Type I error rate using a Bonferroni correction. Doing this for each Riemannian metric gives three global tests. The second approach will combine several Fréchet ANOVA test statistics, each of which applies to just one or two metric spaces, into a single test, also using a Bonferroni correction to account for multiple tests. Using the  $T_s$  and  $T_{ss'}$  statistics, we will consider a test which compares against the  $\chi_{J-1}^2$  distribution, as well as a permutation test. The third approach will apply both extensions of the Pillai-Bartlett trace statistic discussed above, resulting in two tests, one based on permutations and one based on an approximation using the F distribution. Each test will be discussed in more detail below.

In the absence of any distributional assumptions on the underlying metric space probability measures, we lack distributions for the Riemannian statistics under the null hypothesis. The null hypothesis only assumes equality of the first and second moments between the groups, however if we make the stronger assumption that all moments are equal, then observations are exchangeable between groups, and we may generate an approximate sampling distribution for the statistics using either label permutation or bootstrap. Permutation tests have recently been employed to compare groups of persistence diagrams [91], another non-euclidean setting, and Dubey & Müller [50] found in simulation studies that a bootstrapped version of their  $T_s$  statistic showed improved power over the version comparing against a  $\chi_{J-1}^2$  distribution.

We create three global tests, one for each Riemannian-based metric, each comprising of a mean test, a covariance matrix test, and a correlation matrix test, then apply a Bonferroni correction by rejecting if any statistic is greater than the  $1 - \alpha/3$  upper-quantile of its permutation distribution. Explicitly, the three tests are:

1.  $\mathbf{R}_{Euc}$ : reject if the permutation test for any of  $R_{\mu,Euc}$ ,  $R_{\Sigma,Euc}$ ,  $R_{P,Euc}$  rejects at the  $\alpha/3$  level
2.  $\mathbf{R}_{AIRM}$ : reject if the permutation test for any of  $R_{\mu,AIRM}$ ,  $R_{\Sigma,AIRM}$ ,  $R_{P,AIRM}$  rejects at the  $\alpha/3$  level
3.  $\mathbf{R}_{LERM}$ : reject if the permutation test for any of  $R_{\mu,LERM}$ ,  $R_{\Sigma,LERM}$ ,  $R_{P,LERM}$  rejects at the  $\alpha/3$  level

Another way to test the null hypothesis is by applying the Fréchet ANOVA in (5.15) to each of the  $S$  metric spaces separately, and also apply the analogous ANOVA for covariances in (5.49) for each distinct pairing of metric spaces, performing  $S(S + 1)/2$  tests in total. Again we can use a Bonferroni correction to construct a single global test,  $\mathbf{T}_{FA}$ , which we reject if any of the  $S$  ANOVA tests or any of the  $S(S - 1)/2$  ANOVA for covariance tests reject at the  $2\alpha/(S(S + 1))$  significance level. We compare each statistic against the  $1 - 2\alpha/(S(S + 1))$  upper quantile of a  $\chi_{J-1}^2$  distribution. We also consider permutation based tests, denoted  $\mathbf{T}_{FA,perm}$ , where each

statistic is compared against the  $1 - 2\alpha/(S(S + 1))$  upper quantile of the permutation distribution, to investigate if it produces higher power at lower sample sizes as was observed in [50].

Finally, we construct tests using each version of the Pillai-Bartlett trace,  $\Lambda_{pillai}$  and  $\Lambda_{pillai,d}$  to examine a more straightforward application of existing MANOVA methods. Again, we lack distributional results for  $\Lambda_{pillai}$ , so we resort to a permutation test comparing  $\Lambda_{pillai}$  to the upper  $1 - \alpha$  quantile of its permutation distribution. On the other hand,  $\Lambda_{pillai,d}$  is just the Euclidean Pillai-Bartlett trace, which typically assumes normality and homogeneity of variance. The statistic is usually transformed and compared to an  $F$  distribution, and tends to be robust against departure from normality assumptions, see [92] or [93] for more details. We'll denote these tests as  $\Lambda_{pillai}$  and  $\Lambda_{pillai,d}$  respectively.

In total, we have seven tests reflecting three broad approaches to comparing two or more groups in multiple metric spaces:  $R_{Euc}$ , Riemannian permutation test based on the Euclidean Riemannian metric;  $R_{AIRM}$ , Riemannian permutation test based on the affine invariant Riemannian metric;  $R_{LERM}$ , Riemannian permutation test based on the log-Euclidean Riemannian metric;  $T_{FA}$ , asymptotic test based on the Fréchet ANOVA (recall that asymptotic results for the Fréchet covariance remain to be proven);  $T_{FA,perm}$ , permutation test based on the Fréchet ANOVA;  $\Lambda_{pillai}$ , permutation test adapted from the Pillai-Bartlett trace;  $\Lambda_{pillai,d}$ , classical Pillai-Bartlett trace applied to the distances of observations to their respective group means. In the next section, we will evaluate these statistics under various departures from the null hypothesis in two separate scenarios

## 5.6 Simulation Studies

To compare the power and Type I error rate of the tests described in the previous section, we simulate data under the null hypothesis and under various departures from the null hypothesis for two different scenarios, performing each test on each sample. All permutation-based tests permuted group labels 500 times to generate a null distribution, and 6,000 simulated data sets were used to estimate the power and Type I error rate. In both scenarios, we use a desired Type I error rate of  $\alpha = 0.05$ .

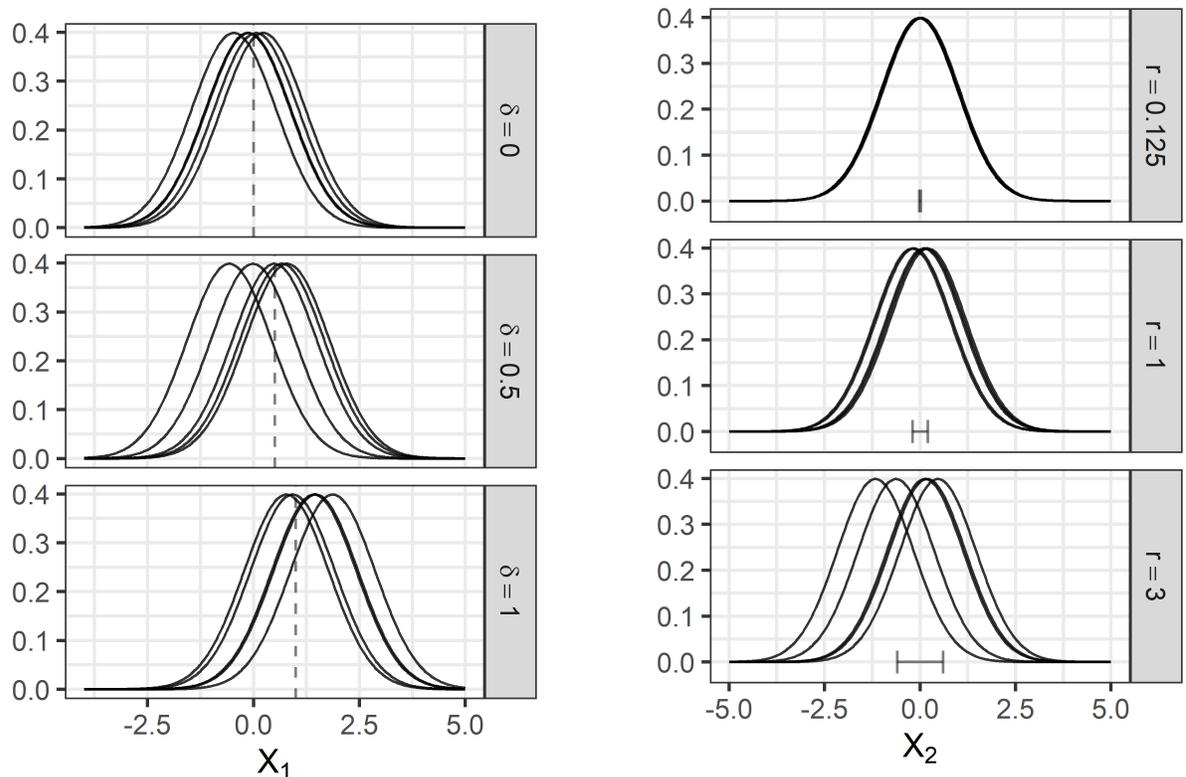
### 5.6.1 Scenario 1: Normal Distributions in Two Metric Spaces, Two Groups

In the first scenario, we compare  $J = 2$  groups in  $S = 2$  metric spaces,  $\Omega_1$  and  $\Omega_2$ . Random objects  $X_1 \in \Omega_1$  are normal distributions with random mean  $a$  and unit variance, and random objects  $X_2 \in \Omega_2$  are normal distributions with random mean  $b$  and unit variance. Both  $\Omega_1$  and  $\Omega_2$  are equipped with the Wasserstein-2 metric which reduces to the  $L_2$  Euclidean distance between the means of the probability distributions. The Fréchet mean of  $X_1$  is a  $N(E(a), 1)$  distribution, and the mean of  $X_2$  is a  $N(E(b), 1)$  distribution. The Fréchet variance of  $X_1$  is simply  $var(a)$ , and the Fréchet variance of  $X_2$  is  $var(b)$ . For this scenario, we performed five simulation studies to examine different ways the data might depart from the null hypothesis, summarized in Table 5.1. In each study, both groups have sample size  $n = 100$ .

The first study investigates a change in the Fréchet mean of  $X_1$  for group 2. That is,  $a \sim N(0, 0.5^2)$  and  $b \sim N(0, 0.2^2)$  for group 1, and  $a \sim N(\delta, 0.5^2)$  and  $b \sim N(0, 0.2^2)$  for group 2, where we let  $\delta$  range between  $-1$  and  $1$ . Note that  $0$  corresponds to the null hypothesis, and  $a$  and  $b$  are independent. Figure 5.2a shows five realizations of  $X_1$  for group 2, under different values of  $\delta$ .

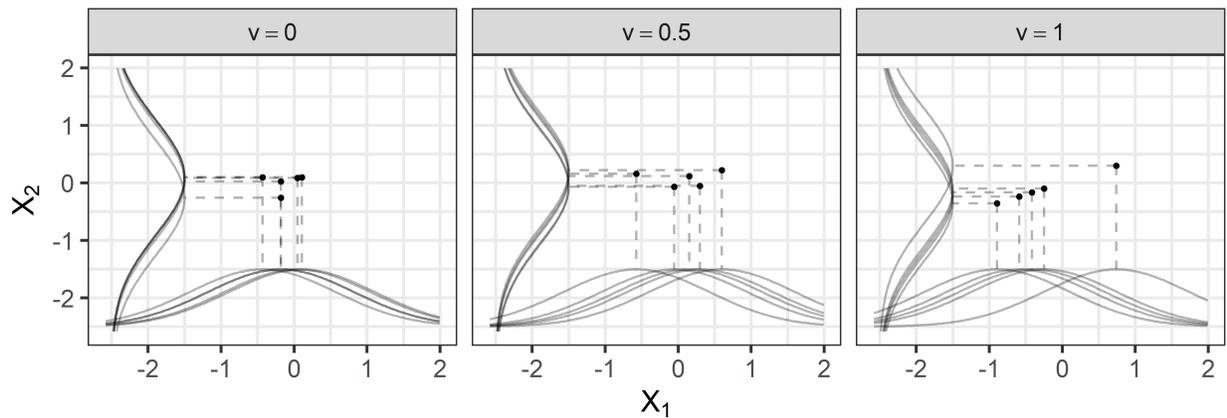
The second study investigates a change in the variance of  $X_2$  for group 2, such that  $b \sim N(0, (0.2r)^2)$  with  $r$  ranging between  $0.125$  and  $3$ , with  $r = 1$  corresponding to the null hypothesis. Figure 5.2b shows five realizations of  $X_2$  for group 2, under different values of  $r$ . The distribution of group 1 is the same as in study 1, and again  $a$  and  $b$  are independent.

The third and fourth studies investigate changes in the dependence between  $X_1$  and  $X_2$ , where dependence is induced by giving  $a$  and  $b$  a non-zero covariance. For study 3,  $a$  and  $b$  for group 1 remain independent, while  $cor(a, b) = v$  for group 2, with  $v$  ranging between  $0$  to  $1$ , and  $v = 0$  corresponding to the null hypothesis. Note that  $a$  and  $b$  are real numbers, so  $v$  is the classic (Euclidean) correlation. For study 4, we let  $a$  and  $b$  for group 1 have some dependence (i.e.  $cor(a, b) = \sqrt{0.5}$ ), while we let  $cor(a, b) = v$  range between  $0$  and  $1$  for group 2. Figure 5.2c shows five realizations of  $(X_1, X_2)$  in group 2 under different values of  $v$ , where  $X_1$  and  $X_2$  are illustrated in the margins, and the points correspond to  $(a, b)$  for each realization.



(a) Five realizations of  $X_1$  for group 2 for simulation scenario 1, study 1, for three values of  $\delta$ , which is the mean of  $a$ .

(b) Five realizations of  $X_1$  for group 2 in simulation scenario 1, study 2, for three values of  $r$ . The length of the whiskers in each plot shows the standard deviation of  $b$ , which is  $0.2r$ .



(c) Five realizations of  $(X_1, X_2)$  for group 2 in simulation scenario 1, study 3, for three values of  $v = \text{cor}(a, b)$ .  $X_1$  and  $X_2$  are illustrated in the margins, and the points correspond to  $(a, b)$ . Dashed lines connect each instance of  $(X_1, X_2)$  to its corresponding point  $(a, b)$ .

**Figure 5.2:** Data examples for the first four studies in scenario 1

The fifth study investigates multiple simultaneous departures from the null hypothesis. Here  $a \sim N(\delta, 0.5^2)$  and  $b \sim N(0, (0.2r)^2)$ , and  $\text{cor}(a, b) = v$  for group 2, whereas in group 1,  $a$  and  $b$

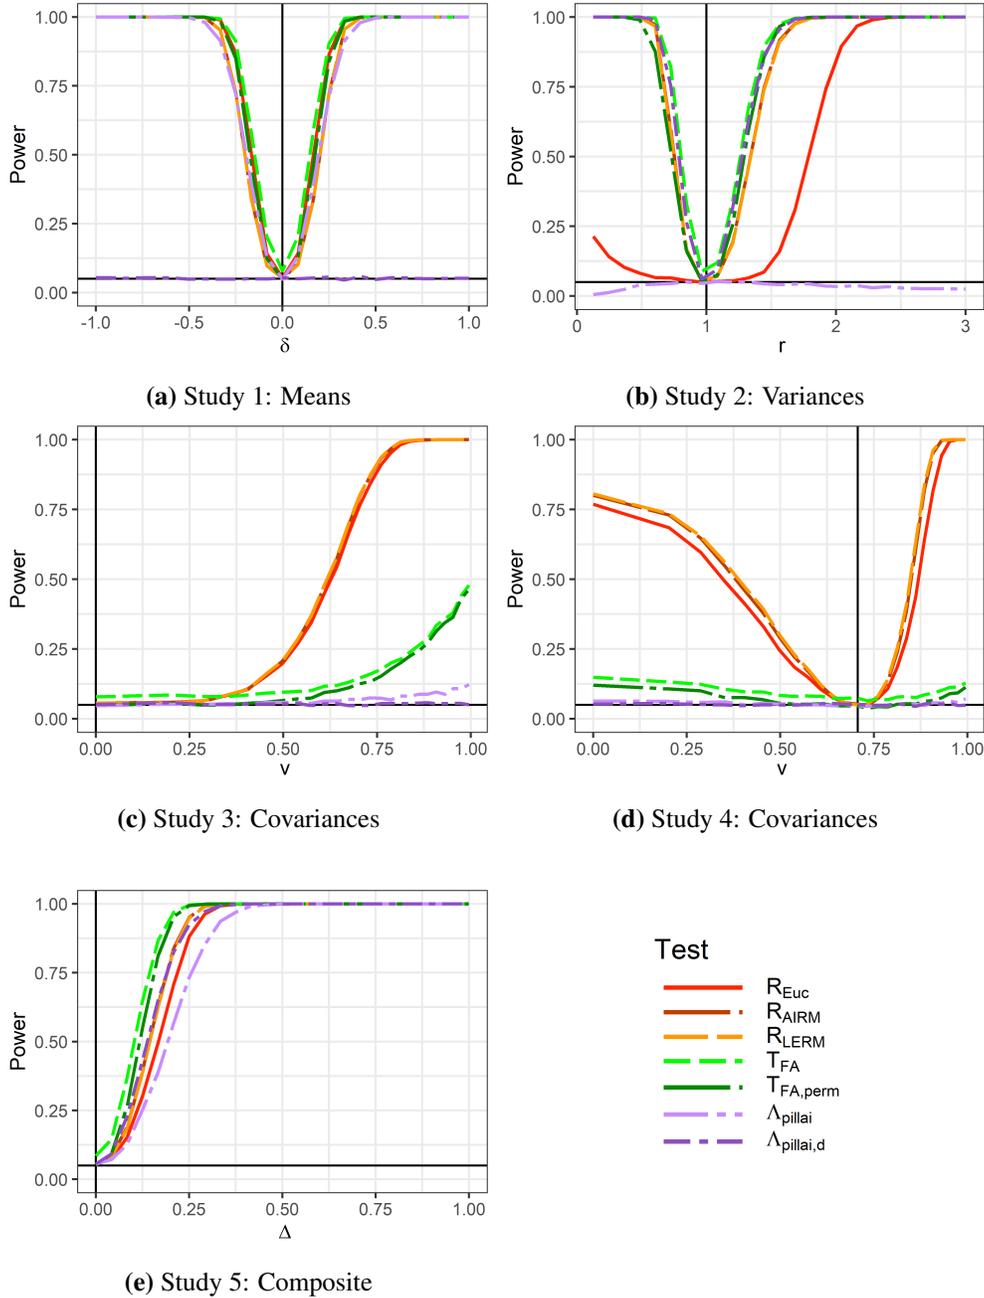
have the same distributions as in study 1. We allow  $\delta$  to range between 0 to 1,  $r$  range from 1 to 3, and  $v$  range from 0 to 1 all simultaneously. Specifically, we let  $(\delta, r, v) = (1 - \Delta) \times (0, 1, 0) + \Delta \times (1, 3, 1)$ , with  $\Delta \in [0, 1]$  indicating the overall “effect size”.

Study / Group	$a$	$b$	$a$ vs. $b$	Parameter Range
Study 1: Means				
- Group 1	$N(\mathbf{0}, \mathbf{0.5}^2)$	$N(0, 0.2^2)$	Independent	$-1 \leq \delta \leq 1$
- Group 2	$N(\delta, \mathbf{0.5}^2)$	$N(0, 0.2^2)$	Independent	
Study 2: Variances				
- Group 1	$N(0, 0.5^2)$	$N(\mathbf{0}, \mathbf{0.2}^2)$	Independent	$0.125 \leq r \leq 3$
- Group 2	$N(0, 0.5^2)$	$N(\mathbf{0}, (\mathbf{0.2}r)^2)$	Independent	
Study 3: Covariances				
- Group 1	$N(0, 0.5^2)$	$N(0, 0.2^2)$	<b>Independent</b>	$0 \leq v \leq 0.9$
- Group 2	$N(0, 0.5^2)$	$N(0, 0.2^2)$	$\text{cor}(\mathbf{a}, \mathbf{b}) = v$	
Study 4: Covariances				
- Group 1	$N(0, 0.5^2)$	$N(0, 0.2^2)$	$\text{cor}(\mathbf{a}, \mathbf{b}) = \sqrt{\mathbf{0.5}}$	$0 \leq v \leq 0.9$
- Group 2	$N(0, 0.5^2)$	$N(0, 0.2^2)$	$\text{cor}(\mathbf{a}, \mathbf{b}) = v$	
Study 5: Composite				
- Group 1	$N(\mathbf{0}, \mathbf{0.5}^2)$	$N(\mathbf{0}, \mathbf{0.2}^2)$	<b>Independent</b>	$0 \leq \delta \leq 1$
- Group 2	$N(\delta, \mathbf{0.5}^2)$	$N(\mathbf{0}, (\mathbf{0.2}r)^2)$	$\text{cor}(\mathbf{a}, \mathbf{b}) = v$	$1 \leq r \leq 3$
				$0 \leq v \leq 0.9$

**Table 5.1:** Summary of simulation parameters for groups 1 and 2 in each study of scenario 1. Bold sections highlight the differences between groups in each study.

The simulation parameters for scenario 1 are summarized in Table 5.1 and the results of the simulation studies are shown in Figure 5.3. When the group means differ in study 1, all tests except  $\Lambda_{pillai,d}$  show similar ability to reject  $H_0$  at varying effect sizes. Since  $\Lambda_{pillai,d}$  only works with distances, it cannot pick up on a shift in the mean of  $X_2$ , so it rejects with rate  $\alpha = 0.05$  at all effect sizes.

When changing group variances in study 2, however,  $\Lambda_{pillai,d}$  is able to detect the difference, while  $\Lambda_{pillai}$  is not. Recall the latter test, being reconfigured from the Euclidean Pillai-Bartlett MANOVA test, is only designed to detect differences in means. Both the Riemannian ( $R_{Euc}$ ,  $R_{AIRM}$ ,  $R_{LERM}$ ) and Fréchet ANOVA based tests ( $T_{FA}$  and  $T_{FA,perm}$ ) reject under departure of the null as well, though the Riemannian test using the underlying Euclidean Metric requires a



**Figure 5.3:** Results from each simulation study in scenario 1. In each plot, a vertical black line indicates the null hypothesis, and a horizontal black line indicates the desired Type I error rate,  $\alpha = 0.05$ .

much more drastic group difference to detect the difference. This shows the utility of the LERM and AIRM Riemannian geometries when comparing PSD matrices in certain scenarios.

For studies 3 and 4, the Riemannian tests show the greatest sensitivity to changes in dependence. Through additional simulations (not shown), we identified that the statistics which compare

correlation matrices ( $R_{P,Euc}$ ,  $R_{P,AIRM}$ ,  $R_{P,LERM}$ ) are the the most sensitive to the differences in dependence in these studies. Omitting the correlation statistics from the Riemannian metric based tests drastically decreases their power against this alternative. The Fréchet ANOVA based tests ( $T_{FA}$ ,  $T_{FA,perm}$ ) include a statistic testing for changes in covariance,  $T_{12}$ , resulting in power against the alternatives in these studies, although substantially lower than that of the Riemannian tests. The Pillai-Bartlett tests ( $\Lambda_{pillai,d}$ ,  $\Lambda_{pillai}$ ) show that adapting classical MANOVA procedures may not prove useful in some instances, with  $\Lambda_{pillai}$  showing low power even at the largest effect size of  $v = 1$  in scenario 3, and  $\Lambda_{pillai,d}$  rejecting close to  $\alpha = 0.05$ .

Though studies 1-4 show that the tests considered can behave very differently under different types of departures from the null hypothesis, they perform similarly when groups 1 and 2 differ in several ways simultaneously as in study 5. Presumably each test is leaning on its own strengths in this case, and the similar power curves are a function of the relative magnitude of change in mean, variance, and covariance affected by the chosen ranges of  $\delta$ ,  $r$ , and  $v$  respectively.

Across all studies, the non-Euclidean Riemannian metric based tests ( $R_{LERM}$ ,  $R_{AIRM}$ ) perform the most consistently well, being on par or superior to the other tests examined.

### 5.6.2 Scenario 2: Random Networks with Node Covariates, Two Groups

In the second scenario, we consider random networks with undirected, simple edges (no loops or multi-edges) with ten nodes each and a real valued covariate on each node. Our interest is in detecting a difference between the groups with respect to network topology and/or node covariates, or some change in the relationship between topology and covariates. Therefore we identify two metric spaces, one which captures network topology and one which captures the node covariates. Let  $\Omega_1$  consist of graph Laplacians equipped with the Frobenius metric, and  $\Omega_2$  consist of vectors in  $\mathbb{R}^{10}$  equipped with the standard Euclidean metric, so that each network can be represented by a random object vector  $(X_1, X_2) \in \Omega_1 \times \Omega_2$ .

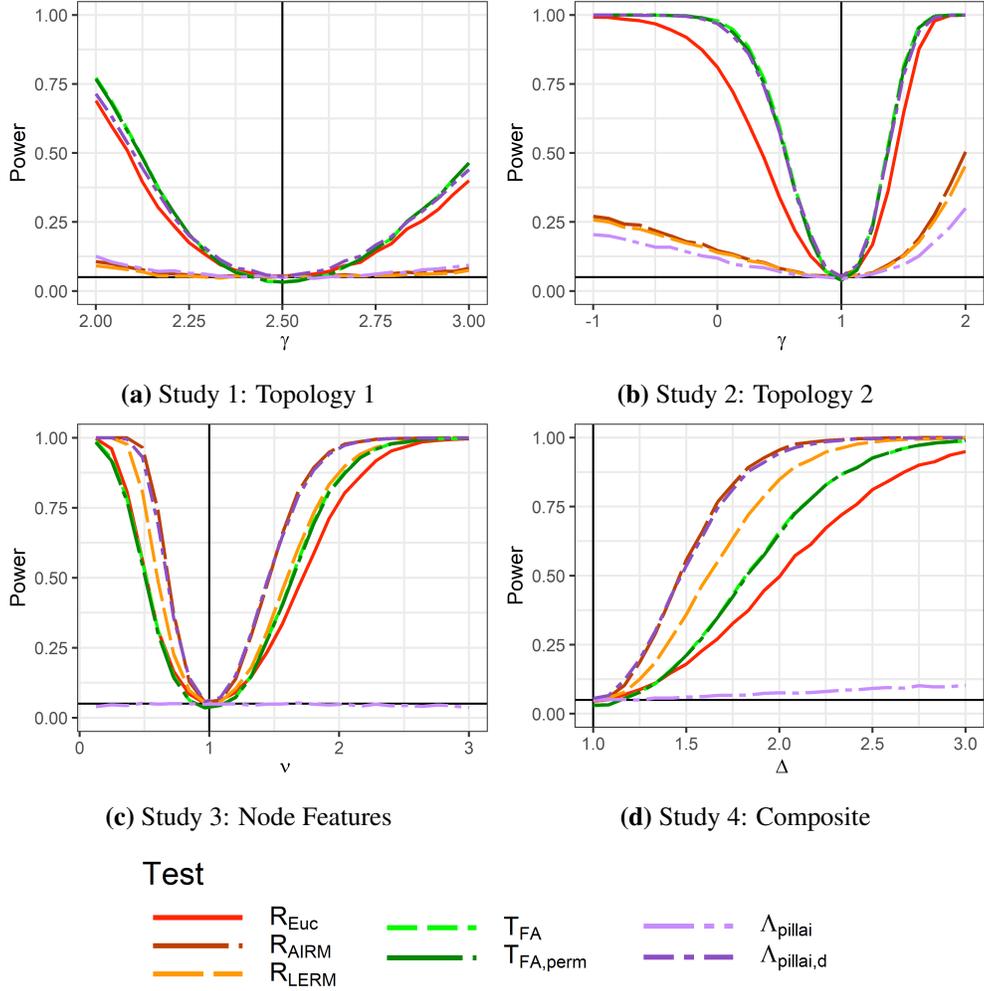
We generate observations of  $(X_1, X_2)$  by first sampling  $X_1$ , the network topology, then sampling  $X_2$ , the node covariates, conditional on  $X_1$ . We sample random networks using the Barabasi-

<b>Study / Group</b>	$\gamma$	$\nu$
Study 1: Power Law		
- Group 1	$\in (2, 3)$	1
- Group 2	<b>2.5</b>	1
Study 2: Power Law		
- Group 1	$\in (-1, 2)$	1
- Group 2	<b>1</b>	1
Study 3: Dependence		
- Group 1	2.5	<b>1</b>
- Group 2	2.5	$\in (0.125, 3)$
Study 4: Composite		
- Group 1	$\in (2.5, 3)$	1
- Group 2	<b>2.5</b>	$\in (1, 3)$

**Table 5.2:** Summary of simulation parameters for Groups 1 and 2 in each study of scenario 2. Bold sections highlight the differences between groups in each study.

Albert growth-plus-preferential-attachment mechanism [94] in which each network starts with a single node, then nodes are added one at a time, each connecting ("attaching") to a single existing node at random. The probability of attaching to an existing node is proportional to  $k^\gamma$  with  $\gamma \in \mathbb{R}$ , where  $k$  is the current number of edges currently incident to that node, its *degree*. Note that a node's degree  $k$  can change as the network is constructed, so we let  $k_f$  denote a node's final degree after the network is fully sampled. Values of  $\gamma$  closer to 0 approach more uniform degree distributions while values further from 0 tend to produce more skewed degree distributions. Each node covariate is generated according to a Gamma distribution with mean  $k_f$ , and variance  $\nu$ . Because network topology and node covariates both depend on node degree  $k_f$ , there is an inherent dependence between  $X_1$  and  $X_2$ , and the strength of that dependence depends on the magnitudes of  $\gamma$  and  $\nu$ : larger magnitudes of  $\gamma$  and smaller values of  $\nu$  correspond with stronger dependence between  $X_1$  and  $X_2$ . We consider four studies, summarized in Table 5.2, and for each, both groups have sample size  $n = 100$ .

For the first study,  $\gamma$  is fixed at 2.5 for group 2 while it varies between 2 and 3 for group 1. Both groups have  $\nu = 1$ . The second study is identical to study 1, except that  $\gamma$  is fixed at 1 for group 2 while it varies between  $-1$  and 2 for group 1. Changing  $\gamma$  will affect the mean and variance



**Figure 5.4:** Results from each simulation study in scenario 2. In each plot, a vertical black line indicates the null hypothesis, and a horizontal black line indicates the desired Type I error rate,  $\alpha = 0.05$

of  $X_1$  (network topology) as well as the mean of  $X_2$  (node covariates), leading to a multifaceted departure from the null hypothesis.

In the third study,  $\gamma = 2.5$  for both groups, but  $\nu$  is fixed at 1 for group 1 and varies between 0.125 and 3 for group 2. Higher values of  $\nu$  lead to weaker dependence between  $X_1$  and  $X_2$ . In the fourth study both groups vary at the same time, but in different ways. For group 1,  $\gamma$  varies between 2.5 and 3 with  $\nu$  fixed at 1, while for group 2,  $\nu$  varies between 1 and 3 simultaneously with  $\gamma$  fixed at 2.5. Specifically,  $(\gamma_{\text{group 1}}, \nu_{\text{group 2}}) = (1 - \Delta) \times (2.5, 1) + \Delta(3, 3)$ , with  $\Delta \in [0, 1]$  indicating the overall “effect size”.

The results of the simulation studies are shown in Figure 5.4. In study 1 the changing topology is best detected by the Fréchet ANOVA based tests, the Riemannian test using the underlying Euclidean metric, and the Pillai-Bartlett trace using distances. In this case, the non-Euclidean geometry imposed by the AIRM and LERM metrics works against those tests. It turns out that those tests are most sensitive to changes in determinants, owing to the behavior of the underlying  $d_{LERM}$  and  $d_{AIRM}$  distances. For instance, the *shortest* path between two PSD matrices  $A$  and  $B$  with the same determinant will always follow a curve of constant determinant. Empirically, changing  $\gamma$  results in group covariance matrices which are different but have similar determinant, similarly for the pooled and mean group covariance matrices, meaning this study is particularly challenging for the  $R_{LERM}$  and  $R_{AIRM}$  tests.

Study 2 shows that  $R_{LERM}$  and  $R_{AIRM}$  have *some* ability to detect changes in  $\gamma$ , but are much less powerful than other tests. Study 3 reflects a more isolated difference between groups, where only the variance of  $X_2$  and covariance between  $X_1$  and  $X_2$  change. In this case, the  $R_{AIRM}$  and  $R_{LERM}$  are comparable to or are better than the Fréchet ANOVA based tests.

When combining both a changing  $\gamma$  and a changing  $\nu$  in study 4, we see that the  $R_{AIRM}$  and  $\Lambda_{pillai,d}$  tests have the greatest sensitivity to change, but again this is influenced by the relative magnitude of change of  $\gamma$  and  $\nu$ .

In all four networks studies, the distance based Pillai-Bartlett test is among the best performing of all tests considered, which makes sense given that all studies involve changing variances and hence changing distances between the groups. On the other hand, the adapted Euclidean Pillai-Bartlett test shows some power against changing  $\gamma$ , which affects the mean of  $X_1$ , but is not sensitive to changing  $\nu$  at all, which affects only variances and covariances.

### 5.6.3 Type I Error Rate

Table 5.3 summarizes the Type I error for scenarios 1 and 2. All tests are below the desired Type I error rate of 0.05 in both scenarios, with the exception of the Fréchet ANOVA based  $T_{FA}$  in Scenario 1, at 0.074. This statistic relies on asymptotic results for the Fréchet variance, and

Scenario	$R_{Euc}$	$R_{AIRM}$	$R_{LERM}$	$T_{FA}$	$T_{FA,perm}$	$\Lambda_{pillai}$	$\Lambda_{pillai,d}$
1	0.046	0.048	0.047	0.074	0.041	0.049	0.050
2	0.045	0.046	0.044	0.033	0.030	0.051	0.051

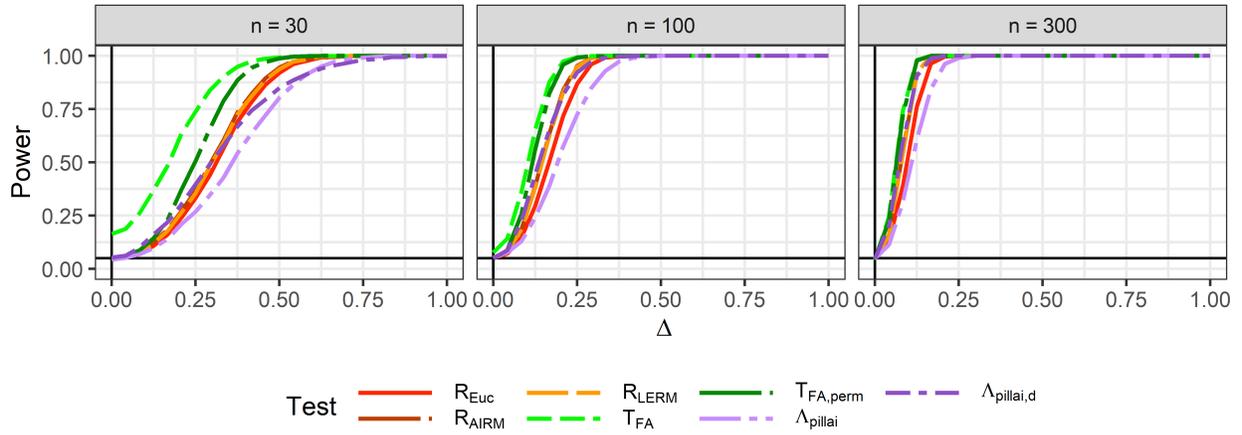
**Table 5.3:** Type I error rates for each test in the studies of scenarios 1 and 2. The Type I error is not a function of the alternative hypothesis so it is the same across all studies within each scenario. We simulated 6,000 data sets under each scenario, giving a standard error of approximately 0.003 for each estimate, and used 1,000 permutations for each permutation test. .

presumptive (not yet proven) asymptotic results for the Fréchet covariance, so the over-rejection is likely explained by a lack of convergence to the asymptotic distribution in the tail of the Fréchet ANOVA statistics. Indeed, the same over-rejection is not seen in the permutation version of the test.

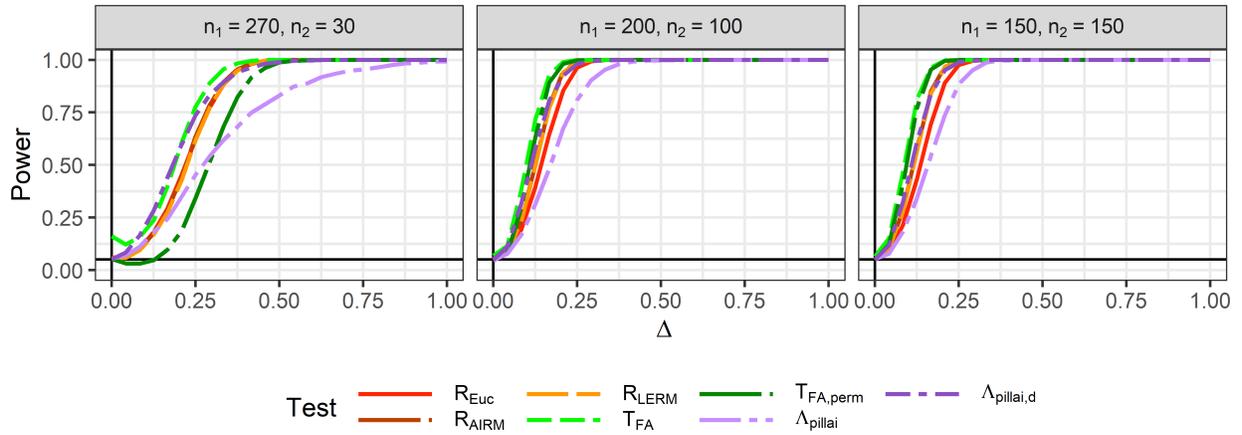
The conservative nature of the Bonferroni correction can be seen in both the Riemannian based tests,  $R_{Euc}$ ,  $R_{AIRM}$ ,  $R_{LERM}$ , and the Fréchet ANOVA based tests,  $T_{FA}$  and  $T_{FA,perm}$ , which reject below the nominal level. As the number of metric spaces  $S$  increases, the Fréchet ANOVA based tests must correct for a greater number of underlying tests, while the number of underlying test statistics for the Riemannian based tests remain at three. The lowest Type I error rate is seen for  $T_{FA}$  and  $T_{FA,perm}$  in scenario 2, resulting from correlation between the underlying test statistics. We expect correlation to be common among the Fréchet ANOVA based statistics due to the correlation between the estimated elements of the Fréchet covariance matrices. The Pillai tests,  $\Lambda_{pillai}$  and  $\Lambda_{pillai,d}$ , don't use a Bonferroni correction, and appear to reject at the nominal level.

#### 5.6.4 Effect of Sample Size and Group Balance

Finally, we investigate how sample size and balanced group sizes affect the performance of each test in study 5 of scenario 1. Figure 5.5 shows power curves under three different sample sizes:  $n_1 = n_2 = 30$ ,  $n_1 = n_2 = 100$ , and  $n_1 = n_2 = 300$ . As expected, larger sample sizes lead to increased power at all effect sizes across all tests examined. However we see the inflated Type I error rate (shown at  $\Delta = 0$  in Figure 5.5) for the asymptotic Fréchet ANOVA based test,  $T_{FA}$ , unsurprising given its reliance on large sample asymptotics. However, the permutation version of



**Figure 5.5:** Results of study 5 in scenario 1 at three different sample sizes (equal for each group). In each plot, a vertical black line indicates the null hypothesis, and a horizontal black line indicates the desired Type I error rate,  $\alpha = 0.05$ .



**Figure 5.6:** Results of study 5 in scenario 1 with varying balance (unequal samples sizes). In each plot, a vertical black line indicates the null hypothesis, and a horizontal black line indicates the desired Type I error rate,  $\alpha = 0.05$ .

the Fréchet ANOVA based test,  $T_{FA,perm}$ , shows a more appropriate Type I error and the best power among the tests (excluding  $T_{FA}$ ) at the lower sample sizes.

Figure 5.6 shows the results of varying the relative sample size of each group, keeping the total sample size the same ( $n_1 + n_2 = 300$ ). We consider highly imbalanced ( $n_1 = 270, n_2 = 30$ ), moderately imbalanced ( $n_1 = 100, n_2 = 200$ ), and balanced ( $n_1 = 150, n_2 = 150$ ) situations. The inflated Type I error rate ( $\Delta = 0$ ) of  $T_{FA}$  is again present in the highly imbalanced case, for the same reason as in the balanced design with  $n_1 = n_2 = 30$ , since group 2 has few observations. Among the other tests, the Riemannian tests and distance based Pillai-Bartlett tests are best in the highly imbalanced case.

## 5.7 Conclusions

In expanding the setting of random objects to multiple metric spaces, we defined the Fréchet analog to covariance. We motivated two possible definitions and explored their interpretations, as well as established consistency for the sample estimator for one of them, which allowed the construction of a full metric space covariance matrix. Using this matrix, we considered various tests for differences in mean and covariance structure in multiple metric spaces between two or more groups. Finally, we examined the performance of these tests under two simulation scenarios and studied their power characteristics under a variety of departures from the null distribution.

It is clear from our simulations that no single proposed test is universally preferred, a result reminiscent of prior work comparing the performance of the various classical MANOVA tests [95, 96]. However, the adaptations of the Pillai-Bartlett test appeared to fall short of the performance of the other tests:  $\Lambda_{pillai,d}$  because it only sees the within group distances,  $\Lambda_{pillai}$  because it is not designed to test for variances. A possible augmentation to either  $\Lambda_{pillai,d}$  or  $\Lambda_{pillai}$  would be to incorporate some adaptation of Box's M test [97], though performance may be poor when normality assumptions do not hold, as noted in [98].

Permutation tests provided a path to testing for group differences in the absence of asymptotic results, which is particularly useful when distributional properties are unknown. However, the

desire remains to establish asymptotic results for the Fréchet covariance, and theoretically justify the Riemannian approaches, if possible showing asymptotically favorable power as sample size grows.

Despite recent progress on working with random objects in metric spaces, many existing challenges remain. Among these are the selection of an appropriate metric, the mechanics of its computation, and efficient estimation of sample Fréchet mean. Regardless, future prospects of the metric space framing of data sets appear promising.

# Chapter 6

## Conclusion

The work in this dissertation explores two approaches to the analysis of network data. In one approach we wish to know how special is an observed network, and in the other approach we wish to compare several groups of networks (or other random objects) represented in multiple metric spaces. These approaches differ both in their purpose and the generality of their scope, but they are both intriguing avenues worthy of further research.

To address the question of how special is an observed network we considered the popular method of comparing an observed network against random networks with the same degree distribution. Sampling such networks can be reduced to sampling binary matrices with fixed marginal sums, for which several sampling methods exist. We extended two MCMC techniques based on edge swaps to allow for non-uniform sampling of the space through specification of a weight matrix. This gave rise to structural zeros, which in some cases prohibit Markov transitions from one state to another, resulting in a reducible chain. However when structural zeros are monotonic, we prove the sample space is still connected. In principle, this type of non-uniform extension may be applied to other property-preserving sampling algorithms, like the recently proposed method of sampling chordal graphs [99]. It may also be possible to apply the non-uniform extension to future methods which preserve properties like modularity, assortativity, or community structure, which we believe would also be of interest to practitioners. We observe that non-uniform weights and the presence of structural zeros increases autocorrelation between samples and affects mixing time, which can be addressed using the methods in [100] for example, but further analysis is needed to better quantify this impact theoretically. It would also be valuable to identify other structural zero configurations which retain irreducibility in the Markov chain. The estimation or specification of weights remains an open problem.

We also generalized a divide and conquer approach to the uniform sampling of binary matrices with fixed margins. In exploring this generalization, we identified emergent properties of the space

of margin fills and spreads which satisfy the Gale-Ryser conditions. Importantly, this includes the convexity of the integer polytope describing the set of valid fills, enabling full enumeration with local search. In simulation studies, we investigated factors which contribute to algorithm performance metrics like lookup table size and running time, which are partially implementation dependent, but also depend on characteristics of the margins, the size of the matrix, etc. Many open questions exist for this generalized method. A vital contribution would be a bound on the worst case running time, which may be achievable using a similar approach to [33], which bounds the number of nodes in the recursive computation tree. It would be tempting to explore a non-uniform extension to the divide and conquer approach as well. However, we foresee challenges in the bookkeeping, which is already quite complicated, required to incorporate the non-uniform weights of all the matrices represented by a certain choice of margin, into the probability of choosing that margin. It is possible this would require full enumeration of the sample space, defeating one purpose of divide and conquer. Though we focus on sampling binary matrices, it is possible a similar approach could be taken for sampling contingency tables. Though our generalized algorithm is not simple, its success may indicate that related methods are awaiting discovery, perhaps a probabilistic divide and conquer approach similar to [101].

To compare several groups of networks, we framed the statistics problem in terms of random objects in metric spaces, where the lack of a coordinate system renders most classical statistical methods unusable. Inspired by recent progress extending the ANOVA test to metric space random objects, we considered a multivariate extension, presuming that multiple metric spaces may more adequately capture the various facets of random objects better than a single metric space can. We developed the notion of Fréchet covariance, and proved its consistency. We then proposed and evaluated several tests for differences in means and covariance matrices between two or more groups, with some tests being inspired by classical MANOVA tests, some drawing inspiration from the Riemannian geometry of symmetric positive definite matrices, and some relying on our theoretical results. Most of these tests relied on permutation testing, but asymptotic results under the null distribution may be achievable. At a minimum, asymptotic power results for one or more of

the proposed tests would help validate them as possible choices for the metric space version of MANOVA. Finally, extending to multiple metric spaces not only allows more expressive representations of random objects, but also opens the door to the development of analogs of classical multivariate approaches, for example, metric space on metric space regression and metric space principal component analysis, both promising paths for future research.

# Bibliography

- [1] James M. Flegal and Galin L. Jones. Batch means and spectral variance estimators in Markov chain Monte Carlo. *The Annals of Statistics*, 38(2):1034–1070, April 2010.
- [2] Eric Goles, Michel Morvan, and Ha Duong Phan. Sandpiles and order structure of integer partitions. *Discrete Applied Mathematics*, 117(1-3):51–64, 2002. Publisher: Elsevier.
- [3] Richard A. Brualdi. Matrices of zeros and ones with fixed row and column sum vectors. *Linear Algebra and its Applications*, 33:159–231, 1980.
- [4] Jared M. Diamond. Assembly of species communities. In *Ecology and Evolution of Communities*, pages 342–443. Harvard University Press, 1975.
- [5] D. R. Fulkerson and H. J. Ryser. Multiplicities and minimal widths for  $(0, 1)$ -matrices. *Canadian Journal of Mathematics*, 14:498–508, 1962.
- [6] Adery C. A. Hope. A simplified Monte Carlo significance test procedure. *Journal of the Royal Statistical Society: Series B (Methodological)*, 30(3):582–598, 1968.
- [7] Edward F. Connor and Daniel Simberloff. The assembly of species communities: chance or competition? *Ecology*, 60(6):1132, December 1979.
- [8] Jared M. Diamond and Michael E. Gilpin. Examination of the “null” model of Connor and Simberloff for species co-occurrences on islands. *Oecologia*, 52(1):64–74, 1982.
- [9] Julian Besag and Peter Clifford. Generalized Monte Carlo significance tests. *Biometrika*, 76(4):633–642, December 1989.
- [10] Lewi Stone and Alan Roberts. The checkerboard score and species distributions. *Oecologia*, 85(1):74–79, 1990.
- [11] Bryan F. J. Manly. A note on the analysis of species co-occurrences. *Ecology*, 76(4):1109–1115, 1995.

- [12] James G. Sanderson, Michael P. Moulton, and Ralph G. Selfridge. Null matrices and the analysis of species co-occurrences. *Oecologia*, 116(1-2):275–283, 1998.
- [13] Nicholas J. Gotelli. Null model analysis of species co-occurrence patterns. *Ecology*, 81(9):2606–2621, 2000.
- [14] Nicholas J. Gotelli and Gary L. Entsminger. Swap and fill algorithms in null model analysis: rethinking the knight’s tour. *Oecologia*, 129(2):281–291, October 2001.
- [15] P H Harvey, R K Colwell, J W Silvertown, and R M May. Null models in ecology. *Annual Review of Ecology and Systematics*, 14(1):189–211, 1983.
- [16] Nicholas J. Gotelli and Gary R. Graves. *Null models in ecology*. Smithsonian, April 1996.
- [17] Georg Rasch. An individualistic approach to item analysis. *Readings in mathematical social science*, pages 89–108, 1966.
- [18] Jan-Eric Gustafsson. Testing and obtaining fit of data to the Rasch model. *British Journal of Mathematical and Statistical Psychology*, 33(2):205–233, 1980.
- [19] Hendrikus Kelderman. Loglinear Rasch model tests. *Psychometrika*, 49(2):223–245, 1984. Publisher: Springer.
- [20] Ivo Ponocny. Nonparametric goodness-of-fit tests for the Rasch model. *Psychometrika*, 66(3):437–459, 2001. Publisher: Springer.
- [21] Bailey K. Fosdick, Daniel B. Larremore, Joel Nishimura, and Johan Ugander. Configuring random graph models with fixed degree sequences. *SIAM Review*, 60(2):315–355, January 2018. Publisher: Society for Industrial and Applied Mathematics.
- [22] M. Harrison and Jeffrey W. Miller. Importance sampling for weighted binary random matrices with specified margins. *arXiv*, 2013.

- [23] Martin Dyer, Ravi Kannan, and John Mount. Sampling contingency tables. *Random Structures & Algorithms*, 10(4):487–506, 1997. Publisher: Wiley Online Library.
- [24] Yuguo Chen. Conditional inference on tables with structural zeros. *Journal of Computational and Graphical Statistics*, 16(2):445–467, 2007. Publisher: Taylor & Francis.
- [25] J. Bastow Wilson. Methods for detecting non-randomness in species co-occurrences: a contribution. *Oecologia*, 73(4):579–582, 1987.
- [26] Alan Roberts and Lewi Stone. Island-sharing by archipelago species. *Oecologia*, 83(4):560–567, 1990.
- [27] Arif Zaman and Daniel Simberloff. Random binary matrices in biogeographical ecology—instituting a good neighbor policy. *Environmental and Ecological Statistics*, 9(4):405–421, 2002.
- [28] Ravi Kannan, Prasad Tetali, and Santosh Vempala. Simple Markov-chain algorithms for generating bipartite graphs and tournaments. *Random Structures & Algorithms*, 13(4):16, 1999.
- [29] Yuguo Chen, Persi Diaconis, Susan P. Holmes, and Jun S. Liu. Sequential Monte Carlo methods for statistical analysis of tables. *Journal of the American Statistical Association*, 100(469):109–120, 2005.
- [30] C. J. Carstens. Proof of uniform sampling of binary matrices with fixed row sums and column sums for the fast Curveball algorithm. *Physical Review E*, 91(4):042812, April 2015.
- [31] Guanyang Wang. A fast MCMC algorithm for the uniform sampling of binary matrices with fixed margins. *Electronic Journal of Statistics*, 14(1):1690–1706, January 2020. Publisher: Institute of Mathematical Statistics and Bernoulli Society.

- [32] Yuguo Chen and Dylan Small. Exact tests for the rasch model via sequential importance sampling. *Psychometrika*, 70(1):11–30, March 2005.
- [33] Jeffrey W. Miller and Matthew T. Harrison. Exact sampling and counting for fixed-margin matrices. *The Annals of Statistics*, 41(3):1569–1592, June 2013.
- [34] Karl Pearson. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50(302):157–175, July 1900.
- [35] A. W. Van Der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998. Google-Books-ID: udhfQgAACAAJ.
- [36] Alan Agresti. *Analysis of ordinal categorical data*, volume 656. John Wiley & Sons, 2010.
- [37] Alan Agresti. *Categorical data analysis*. John Wiley & Sons, 2003.
- [38] Jörg Blasius and Michael Greenacre. *Visualization of categorical data*. Academic Press, 1998.
- [39] Maurice George Kendall. *The advanced theory of statistics*. Charles Griffin and Co., Ltd., London, 2nd edition, 1946.
- [40] Clara Happ and Sonja Greven. Multivariate functional principal component analysis for data observed on different (dimensional) domains. *Journal of the American Statistical Association*, 113(522):649–659, April 2018.
- [41] P. J. Brockwell. Time series analysis. *Encyclopedia of Statistics in Behavioral Science*, 2005. Publisher: Wiley Online Library.
- [42] Mark EJ Newman. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003. Publisher: SIAM.

- [43] Brian D. Ripley. *Spatial statistics*. John Wiley & Sons, 2005.
- [44] Alan E. Gelfand, Peter Diggle, Peter Guttorp, andMontserrat Fuentes. *Handbook of spatial statistics*. CRC press, 2010.
- [45] Haonan Wang and J. S. Marron. Object oriented data analysis: Sets of trees. *The Annals of Statistics*, 35(5):1849–1873, 2007. Publisher: Institute of Mathematical Statistics.
- [46] Cedric E. Ginestet. Strong Consistency of Frechet Sample Mean Sets for Graph-Valued Random Variables. *arXiv:1204.3183 [math, q-bio, stat]*, April 2012. arXiv: 1204.3183 version: 1.
- [47] J. Steve Marron and Andrés M. Alonso. Overview of object oriented data analysis. *Biometrical Journal*, 56(5):732–753, 2014. Publisher: Wiley Online Library.
- [48] Rabi Bhattacharya and Vic Patrangenaru. Large sample theory of intrinsic and extrinsic sample means on manifolds. *Annals of Statistics*, 31(1):1–29, February 2003. Publisher: Institute of Mathematical Statistics.
- [49] Abhishek Bhattacharya and Rabi Bhattacharya. *Nonparametric inference on manifolds: with applications to shape spaces*. Cambridge University Press, Cambridge, 2012.
- [50] Paromita Dubey and Hans-Georg Müller. Fréchet analysis of variance for random objects. *Biometrika*, page 19, 2019.
- [51] Raoul Müller, Dominic Schuhmacher, and Jorge Mateu. ANOVA for data in metric spaces, with applications to spatial point patterns. *arXiv:2201.08664 [stat]*, January 2022. arXiv: 2201.08664.
- [52] Eric D. Kolaczyk, Lizhen Lin, Steven Rosenberg, Jackson Walters, and Jie Xu. Averages of unlabeled networks: geometric characterization and asymptotic behavior. *Annals of Statistics*, 48(1):514–538, February 2020. Publisher: Institute of Mathematical Statistics.

- [53] K. R. Parthasarathy. *Probability measures on metric spaces*. American Mathematical Soc., 2005. Google-Books-ID: 640bAgAAQBAJ.
- [54] Paromita Dubey and Hans-Georg Müller. Modeling time-varying random objects and dynamic networks. *Journal of the American Statistical Association*, pages 1–16, June 2021.
- [55] Paromita Dubey and Hans-Georg Müller. Functional models for time-varying random objects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(2):275–327, 2020.
- [56] Alfred Young. On quantitative substitutional analysis. *Proceedings of the London Mathematical Society*, 1(1):97–145, 1900. Publisher: Wiley Online Library.
- [57] David Gale. A theorem on flows in networks. *Pacific Journal of Math*, 7(2):1073–1082, 1957.
- [58] Herbert J. Ryser. Combinatorial properties of matrices of zeros and ones. *Canadian Journal of Mathematics*, 9:371–377, 1957. Publisher: Cambridge University Press.
- [59] Thomas Brylawski. The lattice of integer partitions. *Discrete Mathematics*, 6(3):201–219, 1973. Publisher: Elsevier.
- [60] Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A Riemannian Framework for Tensor Computing. *International Journal of Computer Vision*, 66(1):41–66, January 2006.
- [61] Vincent Arsigny, Pierre Fillard, Xavier Pennec, and Nicholas Ayache. Geometric Means in a Novel Vector Space Structure on Symmetric Positive-Definite Matrices. *SIAM Journal on Matrix Analysis and Applications*, 29(1):328–347, January 2007.
- [62] Kisung You and Hae-Jeong Park. Re-visiting Riemannian geometry of symmetric positive definite matrices for the analysis of functional connectivity. *NeuroImage*, 225:117464, January 2021.

- [63] Maher Moakher. A differential geometric approach to the geometric mean of symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications*, 26(3):735–747, January 2005.
- [64] Zhenhua Lin. Riemannian Geometry of Symmetric Positive Definite Matrices via Cholesky Decomposition. *SIAM Journal on Matrix Analysis and Applications*, 40(4):1353–1370, January 2019. arXiv: 1908.09326.
- [65] Edward F. Connor, Michael D. Collins, and Daniel Simberloff. The checkered history of checkerboard distributions. *Ecology*, 94(11):2403–2414, November 2013.
- [66] Giovanni Strona, Domenico Nappo, Francesco Boccacci, Simone Fattorini, and Jesus San-Miguel-Ayanz. A fast and unbiased procedure to randomize ecological binary matrices with fixed row and column totals. *Nature Communications*, 5(1):1–9, June 2014.
- [67] Peter D. Hoff. *A first course in Bayesian statistical methods*. Springer texts in statistics. Springer, London ; New York, 2009. OCLC: ocn310401109.
- [68] Martyn Plummer, Nicky Best, Kate Cowles, and Karen Vines. CODA: convergence diagnosis and output analysis for MCMC. *R News*, 6(1):7–11, 2006.
- [69] Alexander Barvinok and James E. Pommersheim. An algorithmic theory of lattice points in polyhedra. *New perspectives in algebraic combinatorics*, 38:91–147, 1999.
- [70] Alexander Barvinok. 7: Lattice points and lattice polytopes. In *Handbook of discrete and computational geometry*, pages 185–210. Chapman and Hall/CRC, 2017.
- [71] Godfrey H. Hardy and Srinivasa Ramanujan. Asymptotic formulæ in combinatory analysis. *Proceedings of the London Mathematical Society*, 2(1):75–115, 1918. Publisher: Wiley Online Library.
- [72] Hans Rademacher. On the partition function  $p(n)$ . *Proceedings of the London Mathematical Society*, 2(1):241–254, 1938. Publisher: Wiley Online Library.

- [73] Jeffrey W. Miller and Matthew T. Harrison. Exact Enumeration and Sampling of Matrices with Specified Margins. *arXiv:1104.0323 [cs, stat]*, April 2011. arXiv: 1104.0323.
- [74] Eric D Kolaczyk and Gábor Csárdi. *Statistical analysis of network data with R*. Springer, 2014. OCLC: 880842217.
- [75] Frank W. Marrs and Bailey K. Fosdick. Regression of binary network data with exchangeable latent errors. *arXiv preprint arXiv:1911.00082*, 2019.
- [76] Frank W. Marrs, Bailey K. Fosdick, and Tyler H. McCormick. Standard errors for regression on relational data with exchangeable errors. *arXiv e-prints*, pages arXiv–1701, 2017.
- [77] Paul Erdos and Alfréd Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60, 1960. Publisher: Citeseer.
- [78] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47, 2002. Publisher: APS.
- [79] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, 1998.
- [80] Cedric E. Ginestet, Jun Li, Prakash Balachandran, Steven Rosenberg, and Eric D. Kolaczyk. Hypothesis testing for network data in functional neuroimaging. *The Annals of Applied Statistics*, 11(2):725–750, June 2017. Publisher: Institute of Mathematical Statistics.
- [81] Aasa Feragen and Tom Nye. 8 - Statistics on stratified spaces. In Xavier Pennec, Stefan Sommer, and Tom Fletcher, editors, *Riemannian geometric statistics in medical image analysis*, pages 299–342. Academic Press, January 2020.
- [82] Maurice Fréchet. Les éléments aléatoires de nature quelconque dans un espace distancié. In *Annales de l’institut Henri Poincaré*, volume 10, pages 215–310, 1948. Issue: 4.

- [83] Alexander Petersen and Hans-Georg Müller. Fréchet integration and adaptive metric selection for interpretable covariances of multivariate functional data. *Biometrika*, 103(1):103–120, March 2016.
- [84] Alexander Petersen and Hans-Georg Müller. Fréchet regression for random objects with Euclidean predictors. *The Annals of Statistics*, 47(2):691–719, 2019. Publisher: Institute of Mathematical Statistics.
- [85] Wassily Hoeffding. A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 19(3):293–325, September 1948. Publisher: Institute of Mathematical Statistics.
- [86] Howard Levene. Robust tests for equality of variances. *Contributions to probability and statistics. Essays in honor of Harold Hotelling*, pages 279–292, 1960. Publisher: Stanford University Press.
- [87] Samuel S. Wilks. Certain generalizations in the analysis of variance. *Biometrika*, pages 471–494, 1932. Publisher: JSTOR.
- [88] KC Sreedharan Pillai. Some new test criteria in multivariate analysis. *The Annals of Mathematical Statistics*, pages 117–121, 1955. Publisher: JSTOR.
- [89] Harold Hotelling. A generalized T test and measure of multivariate dispersion. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, pages 23–41. University of California Press, 1951.
- [90] Samarendra Nath Roy. *Some aspects of multivariate analysis*. Asia Publishing House, 1958.
- [91] Eric Berry, Yen-Chi Chen, Jessi Cisewski-Kehe, and Brittany Terese Fasy. Functional summaries of persistence diagrams. *Journal of Applied and Computational Topology*, 4(2):211–262, June 2020.

- [92] T. Anderson. *An introduction to multivariate statistical analysis*. Wiley, 1984. Publisher: John Wiley & sons, New York. US.
- [93] W. J. Krzanowski and F. H. C. Marriott. Multivariate analysis, part I: distributions. *Ordination and Inference*, Edward Arnold, London, 1994.
- [94] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999. Publisher: American Association for the Advancement of Science.
- [95] Holmes Finch. Comparison of the performance of nonparametric and parametric MANOVA test statistics when assumptions are violated. *Methodology: European Journal of Research Methods for the Behavioral and Social Sciences*, 1:27–38, January 2005.
- [96] Can Ateş, Özlem Kaymaz, H. Emre Kale, and Mustafa Agah Tekindal. Comparison of test statistics of nonnormal and unbalanced samples for multivariate analysis of variance in terms of type-I error rates. *Computational and Mathematical Methods in Medicine*, 2019:1–8, July 2019.
- [97] Joan Fisher Box. RA Fisher and the design of experiments, 1922–1926. *The American Statistician*, 34(1):1–7, 1980. Publisher: Taylor & Francis.
- [98] Bryan FJ Manly and Jorge A. Navarro Alberto. *Multivariate statistical methods: a primer*. Chapman and Hall/CRC, 2016.
- [99] Wenbo Sun and Ivona Bezáková. Sampling random chordal graphs by MCMC (student abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13929–13930, 2020. Issue: 10.
- [100] Upasana Dutta and Aaron Clauset. Convergence criteria for sampling random graphs with specified degree sequences. *arXiv:2105.12120 [physics, stat]*, May 2021. arXiv: 2105.12120.

- [101] Stephen DeSalvo and James Y. Zhao. Random sampling of contingency tables via probabilistic divide-and-conquer. *arXiv preprint arXiv:1507.00070*, 2015.
- [102] Mark Jerrum and Alistair Sinclair. The Markov chain Monte Carlo method: an approach to approximate counting and integration. In *Approximation algorithms for NP-hard problems*, pages 482–520. PWS Publishing Co., Boston, 1996. Publisher: PWS Boston.
- [103] Michael R. Garey and David S. Johnson. *Computers and intractability*, volume 174. free-man San Francisco, 1979.

# Appendix A

## Proofs

### A.1 Nonuniform Sampling of Fixed Margin Binary Matrices

*Proof of Theorem 2.* Here we prove that the weighted checkerboard swap algorithm samples from  $\mathbb{P}(\cdot)$  defined in (3.1) when there are no structural zeros, by showing that the Markov chain implied by the algorithm is ergodic (irreducible and aperiodic) with stationary distribution  $\mathbb{P}(\cdot)$  [28, 102].

**Irreducibility.** Since  $w_{ij} > 0$  we have  $0 < p_{ij:i'j'} < 1$ , so irreducibility follows directly from the result of Brualdi [3], which showed that a series of swaps could change any matrix  $A$  into any other matrix  $B$  when row and column sums are preserved.

**Aperiodicity.** Since every checkerboard swap has probability less than 1, it is possible to remain in the same state after a step. This is enough to show aperiodicity.

**Detailed Balance.** Lastly, we show that  $P(\cdot)$  is the stationary distribution by showing that it satisfies the detailed balance condition:

$$\mathbb{P}(A) p_{A \rightarrow B} = \mathbb{P}(B) p_{B \rightarrow A}. \quad (\text{A.1})$$

for any matrices  $A, B \in \Omega'(p, q)$ , where  $p_{A \rightarrow B}$  and  $p_{B \rightarrow A}$  are transition probabilities. If  $A$  and  $B$  do not differ by a checkerboard swap, then both transition probabilities are 0, and detailed balance holds trivially. If  $A$  and  $B$  differ by a checkerboard swap, then let  $u$  and  $u'$  be its rows and let  $v$  and  $v'$  be its columns. Also, let  $k$  be the sum of the row counts, and define the set  $\mathcal{S} = \{(u, v), (u', v), (u, v'), (u', v')\}$ , so that  $a_{ij} = b_{ij}$  for  $(i, j) \notin \mathcal{S}$ . The transition probability  $p_{A \rightarrow B}$  is the probability of first selecting the two 1's in the checkerboard,  $\binom{k}{2}^{-1}$ , times the probability of performing the swap  $p_{uv;u'v'}$ . Starting from the left side of (A.1),

$$\mathbb{P}(A) p_{A \rightarrow B} = \quad (\text{A.2})$$

$$= \left[ \frac{1}{\kappa} \prod_{ij} w_{ij}^{a_{ij}} \right] \left[ \binom{k}{2}^{-1} \frac{w_{u'v} w_{uv'}}{w_{uv} w_{u'v'} + w_{u'v} w_{uv'}} \right] \quad (\text{A.3})$$

$$= \frac{1}{\kappa} \binom{k}{2}^{-1} \left[ \prod_{ij \notin \mathcal{S}} w_{ij}^{a_{ij}} \right] w_{uv} w_{u'v'} \frac{w_{u'v} w_{uv'}}{w_{uv} w_{u'v'} + w_{u'v} w_{uv'}} \quad (\text{A.4})$$

$$= \frac{1}{\kappa} \binom{k}{2}^{-1} \left[ \prod_{ij \notin \mathcal{S}} w_{ij}^{b_{ij}} \right] w_{u'v} w_{uv'} \frac{w_{uv} w_{u'v'}}{w_{uv} w_{u'v'} + w_{u'v} w_{uv'}} \quad (\text{A.5})$$

$$= \left[ \frac{1}{\kappa} \prod_{ij} w_{ij}^{b_{ij}} \right] \left[ \binom{k}{2}^{-1} \frac{w_{u'v} w_{uv'}}{w_{uv} w_{u'v'} + w_{u'v} w_{uv'}} \right] \quad (\text{A.6})$$

$$= \mathbb{P}(B) p_{B \rightarrow A}, \quad (\text{A.7})$$

so detailed balance holds. □

*Proof of Theorem 3.* When structural zeros are present, that is, at least one of  $w_{ij} = 0$ , the Markov chain implied by checkerboard swapping may become reducible. Here we prove that in the special case of monotonic structural zeros, irreducibility still holds. Specifically we show that between any two matrices  $A$  and  $B$ , there exists a checkerboard swap which reduces the Hamming distance between them, so that repeated application of such swaps eventually reduces the distance to 0.

**Finding a Swap.** Consider  $A, B \in \Omega'(p, q)$ , and let  $d_H$  be the Hamming distance between them. The structural zeros are monotonic, so assume the rows and columns have been arranged such that  $w_{ij} = 0 \implies w_{kj} = w_{il} = 0$  for  $k > i, l > j$ . This places all the structural zeros in the bottom-right corner (See Figure A.1).

Let  $(i, j)$  denote the first element where  $A$  and  $B$  differ, where “first” means no rows above  $j$  contain a difference, and in row  $j$ , no column to the right of column  $i$  contains a difference. Specifically,  $(i, j)$  satisfies:

1.  $a_{ij} \neq b_{ij}$
2.  $a_{kl} = b_{kl}$  for  $k < i$  and  $1 \leq l \leq n$  ( $A$  and  $B$  agree above row  $i$ )
3.  $a_{il} = b_{il}$  for  $j < l$  ( $A$  and  $B$  agree in row  $j$  to the right of column  $j$ )

In Figure A.1, we show agreement between  $A$  and  $B$  with cross hatching.

Without loss of generality, assume  $a_{ij} = 0$  and  $b_{ij} = 1$ .  $A$  and  $B$  must have the same sum in row  $i$ , but they differ at  $(i, j)$ , so there must be some column  $j'$  such that  $a_{ij'} = 1$  and  $b_{ij'} = 0$ . We also know  $j' < j$  because  $j$  is the last column where  $A$  and  $B$  differ in row  $i$ . Similarly, there must be some row  $i'$ , such that  $i' > i$ , where  $a_{i'j} = 1$  and  $b_{i'j} = 0$ , since the column sums agree.

Element  $(i', j')$ , filled orange in Figure A.1, cannot be a structural zero due to the selection criterion so one of the following is true:

1. Checkerboard in  $A$  and  $B$ :  $a_{i'j'} = 0$  and  $b_{i'j'} = 1$  so both  $A$  and  $B$  contain a checkerboard with the elements in rows  $\{i, i'\}$  and columns  $\{j, j'\}$ . Performing the swap in either  $A$  or  $B$  reduces  $d_H$  by 4.
2. Checkerboard in  $A$  only:  $a_{i'j'} = b_{i'j'} = 0$ . Performing the swap in  $A$  reduces  $d_H$  by 2.
3. Checkerboard in  $B$  only:  $a_{i'j'} = b_{i'j'} = 1$ . Performing the swap in  $B$  reduces  $d_H$  by 2.
4. No checkerboard:  $a_{i'j'} = 1$  and  $b_{i'j'} = 0$ .

Cases (1-3) allow checkerboard swaps which reduce the distance between  $A$  and  $B$  by at least 2.

Case (4) does not permit a swap, but if this is the case, a swap can be found as outlined below.

Let  $c$  be the last column in row  $i'$  which is not a structural zero. Since row  $i'$  in  $A$  and row  $i'$  in  $B$  have the same sum, the *partial* sums up to column  $c$  are also equal (see region with blue border in Figure A.1).

Note that  $j, j' \leq c$ . In row  $i$ ,  $A$  and  $B$  agree after column  $j$  (due to the selection procedure), and since  $j < c$ , they agree after column  $c$  as well. Therefore the partial sums up to column  $c$  must agree in row  $i$  as well (see region with red border in Figure A.1).

We will use the agreement of these partial sums to show that there exists some column  $j''$  which permits a checkerboard swap in either  $A$  or  $B$ . Consider the difference matrix  $D = A - B$  with elements  $d_{kl} \in \{-1, 0, 1\}$ . Because the partial sums described above agree, we have:

$$0 = \sum_{l=1}^c d_{il} = \sum_{l \in \{1, \dots, c\} \setminus \{j, j'\}} d_{il}, \quad (\text{A.8})$$

since  $d_{ij} = -1$ ,  $d_{ij'} = 1$  (recall we are in Case 4 from above). Furthermore,

$$0 = \sum_{l=1}^c d_{i'l} \implies -2 = \sum_{l \in \{1, \dots, c\} \setminus \{j, j'\}} d_{i'l}, \quad (\text{A.9})$$

since  $d_{i'j} = d_{i'j'} = 1$ . Therefore there must be some  $j'' \in \{1, \dots, c\} \setminus \{j, j'\}$  such that  $d_{ij''} > d_{i'j''}$ , and one of the following is true:

1. Checkerboard in B,  $d_{ij''} = 1, d_{i'j''} = 0$ :  $a_{ij''} = 1, b_{ij''} = 0$ , and  $a_{i'j''} = b_{i'j''} = 1$ , a checkerboard swap exists in  $B$  in rows  $\{i, i'\}$  and columns  $\{j, j''\}$ .
2. Checkerboard in A,  $d_{ij''} = 1, d_{i'j''} = 0$ :  $a_{ij''} = 1, b_{ij''} = 0$ , and  $a_{i'j''} = b_{i'j''} = 0$ , a checkerboard swap exists in  $A$  in rows  $\{i, i'\}$  and columns  $\{j, j''\}$ .
3. Checkerboard in A,  $d_{ij''} = 0, d_{i'j''} = -1$ :  $a_{ij''} = b_{ij''} = 1$ , and  $a_{i'j''} = 0, b_{i'j''} = 1$ , a checkerboard swap exists in  $A$  in rows  $\{i, i'\}$  and columns  $\{j, j''\}$ .
4. Checkerboard in B,  $d_{ij''} = 0, d_{i'j''} = -1$ :  $a_{ij''} = b_{ij''} = 0$ , and  $a_{i'j''} = 0, b_{i'j''} = 1$ , a checkerboard swap exists in  $B$  in rows  $\{i, i'\}$  and columns  $\{j, j''\}$ .

Thus there exists a checkerboard swap which decreases  $d_H$  by 2. Weighted curveball trades include checkerboard swaps, so the proof applies to irreducibility of the weighted curveball algorithm as well. This also implies that the number of swaps between any two matrices is bounded by  $\frac{d_H}{2}$ , and hence the diameter of the Markov state space is bounded by half the maximum possible Hamming distance between two matrices.

□

*Proof of Theorem 4.* Here we prove that the weighted curveball algorithm samples from  $\mathbb{P}(\cdot)$  defined in (3.1) when there are no structural zeros.

**Irreducibility.** Any checkerboard swap can be viewed as a curveball trade, so the set of all weighted curveball trades contains the set of all weighted checkerboard swaps. Therefore any state transition in the checkerboard swap Markov chain is also a state transition in the curveball trade chain, hence irreducibility holds for curveball as well.

**Aperiodicity.** Again, since  $w_{ij} > 0$ , then  $0 < p_{ii'} < 1$  for any proposed trade, so the probability of remaining in the same state is non-zero, and the chain is aperiodic.

**Detailed Balance.** For detailed balance, consider matrices  $A$  and  $B$  which differ by a curveball trade in rows  $u$  and  $u'$ . The curveball algorithm involves permuting the elements of the trade candidates in  $A_{u \setminus u'}$  and  $A_{u' \setminus u}$  to create a proposed trade, captured in  $(B_{u \setminus u'}, B_{u' \setminus u})$ . The probability of obtaining  $(B_{u \setminus u'}, B_{u' \setminus u})$  is

$$p_c = |A_{u \setminus u'}|! |A_{u' \setminus u}|! / (|A_{u \setminus u'}| + |A_{u' \setminus u}|)! \quad (\text{A.10})$$

$$= |B_{u \setminus u'}|! |B_{u' \setminus u}|! / (|B_{u \setminus u'}| + |B_{u' \setminus u}|)!, \quad (\text{A.11})$$

where  $|\cdot|$  is the cardinality operator. Note that  $p_c$  is the same for all such partitions, that is, all partitions are equally likely. Let  $\mathcal{S} = \{ij : i \in \{u, u'\}, j \in \{\mathbf{j}_i \cup \mathbf{j}_{i'}\}\}$ , and for a generic index  $t$  and index set  $\mathbf{s}$ , define  $w_{t\mathbf{s}} = \prod_{j \in \mathbf{s}} w_{tj}$ . Again starting from the left side of (A.1):

$$\mathbb{P}(A)p_{A \rightarrow B} = \mathbb{P}(A)p_c p_{uu'} \quad (\text{A.12})$$

$$= \frac{p_c}{\kappa} \prod_{ij} w_{ij}^{a_{ij}} \left[ \frac{w_{u\mathbf{j}_i} w_{u'\mathbf{j}_{i'}}}{w_{u\mathbf{j}_i} w_{u'\mathbf{j}_{i'}} + w_{u'\mathbf{j}_i} w_{u\mathbf{j}_{i'}}} \right] \quad (\text{A.13})$$

$$= \frac{p_c}{\kappa} \left[ \prod_{ij \notin \mathcal{S}} w_{ij}^{a_{ij}} \right] w_{u'\mathbf{j}_i} w_{u\mathbf{j}_{i'}} \left[ \frac{w_{u\mathbf{j}_i} w_{u'\mathbf{j}_{i'}}}{w_{u\mathbf{j}_i} w_{u'\mathbf{j}_{i'}} + w_{u'\mathbf{j}_i} w_{u\mathbf{j}_{i'}}} \right] \quad (\text{A.14})$$

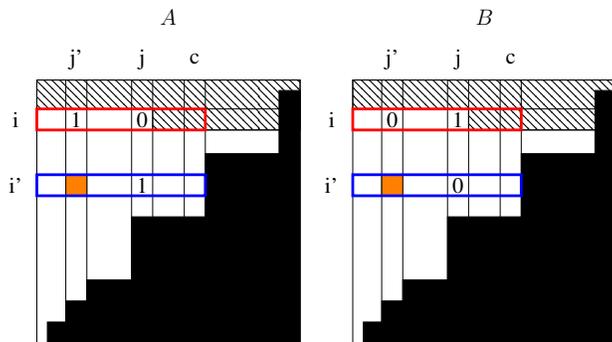
$$= \frac{p_c}{\kappa} \left[ \prod_{ij \notin \mathcal{S}} w_{ij}^{b_{ij}} \right] w_{u\mathbf{j}_i} w_{u'\mathbf{j}_{i'}} \left[ \frac{w_{u'\mathbf{j}_i} w_{u\mathbf{j}_{i'}}}{w_{u\mathbf{j}_i} w_{u'\mathbf{j}_{i'}} + w_{u'\mathbf{j}_i} w_{u\mathbf{j}_{i'}}} \right] \quad (\text{A.15})$$

$$= \frac{p_c}{\kappa} \prod_{ij} w_{ij}^{b_{ij}} \left[ \frac{w_{u'\mathbf{j}_i} w_{u\mathbf{j}_{i'}}}{w_{u\mathbf{j}_i} w_{u'\mathbf{j}_{i'}} + w_{u'\mathbf{j}_i} w_{u\mathbf{j}_{i'}}} \right] \quad (\text{A.16})$$

$$= \mathbb{P}(B)p_c p_{u'u} \quad (\text{A.17})$$

$$= \mathbb{P}(B)p_{B \rightarrow A} \quad (\text{A.18})$$

Hence the detailed balance condition holds, where the stationary distribution is again given by (3.1). □



**Figure A.1:** Illustration of irreducibility proof Theorem 3.  $A$  and  $B$  have structural zeros (in black), and agree for elements in the hashed region. Their partial sums must agree in the red regions, and their partial sums must agree in the blue regions as well. Four cases are possible for the orange element.

## A.2 Divide and Conquer Sampling of Fixed Margin Binary Matrices

*Proof of Lemma 1.* Denote by  $h$  and  $v$  the horizontal and vertical transitions described in [2]. We will show there is a minimal element for each block of  $p$ , and by combining the minimal elements in each block we prove the existence of the minimal element of  $\mathcal{T}(f^l)$ . Our definition of majorization by orbits assumes the vectors have been sorted non-increasing, so we may treat them as integer partitions. Since  $f^l$  is fixed, every  $\tau(p^l, p^r) \in \mathcal{T}(f^l)$  has the same number of elements in each level set (block) of  $p$ . This means the  $i$ th block of every  $\tau(p^l)$  is a partition of  $f_i^l$  into at most  $|s_i|$  parts. Similarly, the  $i$ th block of every  $\tau(p^r)$  is a partition of  $f_i^r$  into at most  $|s_i|$  parts. Denote these partitions with  $\lambda^{(i)}$  and  $\mu^{(i)}$  for the  $i$ th block of the left and right sides, respectively. Because  $p_j^l + p_j^r = i$  for each  $j \in s_i$ , every  $h$  transition in  $\lambda^{(i)}$  corresponds to an  $h$  transition in  $\mu^{(i)}$ , and every  $v$  transition in  $\lambda^{(i)}$  corresponds to a  $v$  transition in  $\mu^{(i)}$ . Repeated application of  $h$  and  $v$  transitions will produce the minimal partitions  $\lambda_*^{(i)}$  and  $\mu_*^{(i)}$ . Define  $\tau(p_*^l)$  such that block  $i$  is  $\lambda_*^{(i)}$  and define  $\tau(p_*^r)$  such that block  $i$  is  $\mu_*^{(i)}$ . Given some  $\tau(p^l, p^r) \in \mathcal{T}(f^l)$  where the  $i$ th block of  $\tau(p^l)$  is  $\lambda^{(i)}$  and the  $i$ th block of  $\tau(p^r)$  is  $\mu^{(i)}$ . Since  $\lambda_*^{(i)} \preceq \lambda^{(i)}$  for each  $i$ , then  $\tau(p_*^l) \preceq \tau(p^l)$ , and similarly since  $\mu_*^{(i)} \preceq \mu^{(i)}$ , then  $\tau(p_*^r) \preceq \tau(p^r)$ .

□

*Proof of Theorem 5.* We show that  $F_1$  and  $F_2$  are individually convex, from which the result immediately follows. Consider mapping from the fill  $f^l$  to the conjugate partition of the minimal element  $\tau(p_*^l) \in \mathcal{T}(f^l)$ , specifically define  $h_k : \mathbb{Z}_{\geq 0}^b \rightarrow \mathbb{Z}_{\geq 0}$  such that  $h_k(f^l) = \sum_{i=1}^m \mathbb{I}(\tau(p_*^l)_i \geq k)$  for  $k \in (1, \dots, n)$ . Then  $h_k(f^l)$  is given by

$$h_k(f^l) = \sum_{i=1}^b \text{rem} \left( \frac{f_i^l}{|s_i|} \right) \mathbb{I} \left( k - 1 < \frac{f_i^l}{|s_i|} < k \right) + |s_i| \mathbb{I} \left( k \leq \frac{f_i^l}{|s_i|} \right), \quad (\text{A.19})$$

where  $\mathbb{I}$  is the indicator function, and  $\text{rem}(a/b) = a - b\lfloor a/b \rfloor$  is the remainder function.

Each  $h_k(f^l)$  is monotone in each  $f_i^l$  and the domain  $F$  is convex, so the image  $H_k = h_k(F)$  is convex for each  $k$ . The Gale-Ryser condition implies constraints to each  $h_k(F)$ , namely

$$\sum_{k=1}^t h_k \geq \sum_{k=1}^t q_k^l \quad (\text{A.20})$$

for  $t \in (1, \dots, n)$ . Let  $H'_i \subseteq H_i$  be the subsets of each image which satisfy (A.20) for all  $t$ . Because these are half space constraints, each  $H'_i$  is still convex. Taking the preimage  $h_i^{-1}(H'_i)$  of each subsetted range preserves convexity because the  $h_i$  are monotone. Satisfying all constraints simultaneously implies taking the intersection of all preimages:  $F_1 = \bigcap_i h_i^{-1}(H'_i)$ , which must be convex as well. Applying this reasoning to the right fill analogously shows that  $F_2$  is also convex.  $\square$

*Proof of Lemma 2.* As in the proof of Lemma 1, we represent the spread orbits as lists of integer partitions  $\lambda^{(i)}$  and  $\mu^{(i)}$ , for the  $i$ th block of the left and right sides, respectively. We will show that any partition lists  $(\lambda^{(1)}, \dots, \lambda^{(b)})$  and  $(\mu^{(1)}, \dots, \mu^{(b)})$  can be reached from the minimum lists  $(\lambda_*^{(1)}, \dots, \lambda_*^{(b)})$  and  $(\mu_*^{(1)}, \dots, \mu_*^{(b)})$ , implied by  $\tau(p_*^l, p_*^r)$  via reverse  $h$  and  $v$  transitions. Brylawski [59] showed that for any partitions  $a$  and  $b$  of  $k$  with  $a \succeq b$ , there exist a sequence of  $h$  and  $v$  transitions which transition from  $a$  to  $b$ . We apply that process to each level set in turn, which “reduces”  $\lambda^{(i)}$  to  $\lambda_*^{(i)}$  and  $\mu^{(i)}$  to  $\mu_*^{(i)}$  using  $h$  and  $v$  transitions. As shown in the proof for Lemma 1, a  $v$  or  $h$  transition in  $\lambda^{(i)}$  corresponds to a  $v$  or  $h$  transition in  $\mu^{(i)}$  as well, so reducing the left

partition also reduces the right partition. Assuming the starting partitions satisfy the Gale-Ryser conditions, each reduction maintains satisfaction of the conditions by definition, and eventually the minimum spread orbit is reached. Applying this chain of  $h$  and  $v$  transitions in reverse would therefore change the minimal orbit into the desired valid spread orbit. Therefore an exhaustive search over  $h$  and  $v$  transitions can produce all valid spread orbits within  $\mathcal{T}(f^l)$ .  $\square$

### A.3 Fréchet Covariance and MANOVA Tests for Random Objects in Multiple Metric Spaces

*Proof of Theorem 6.* Note that  $d(\hat{\mu}_s, \mu_s) \xrightarrow{p} 0$  by the continuous mapping theorem and consistency of  $\hat{\mu}_s$ , similarly for  $X_{s'}$ . Then

$$|\hat{\sigma}_{ss'} - \sigma_{ss'}| = \left| \frac{1}{n} \sum_{i=1}^n d(\hat{\mu}_s, X_{si})d(\hat{\mu}_{s'}, X_{s'i}) - E[d(\mu_s, X_s)d(\mu_{s'}, X_{s'})] \right| \quad (\text{A.21})$$

$$= |A + B + C| \quad (\text{A.22})$$

$$\leq |A| + |B| + |C| \quad (\text{A.23})$$

where

$$A = \frac{1}{n} \sum_{i=1}^n [d(X_{si}, \hat{\mu}_s) - d(X_{si}, \mu_s)] [d(X_{s'i}, \hat{\mu}_{s'}) - d(X_{s'i}, \mu_{s'})] \quad (\text{A.24})$$

$$B = \frac{1}{n} \sum_{i=1}^n d(X_{si}, \hat{\mu}_s)d(X_{s'i}, \mu_{s'}) + d(X_{si}, \mu_s)d(X_{s'i}, \hat{\mu}_{s'}) - 2d(X_{si}, \mu_s)d(X_{s'i}, \mu_{s'}) \quad (\text{A.25})$$

$$C = \frac{1}{n} \sum_{i=1}^n d(X_{si}, \mu_s)d(X_{s'i}, \mu_{s'}) - E[d(X_{si}, \mu_s)d(X_{s'i}, \mu_{s'})] \quad (\text{A.26})$$

Each of  $|A|$ ,  $|B|$ , and  $|C|$  converges to zero in probability as  $n \rightarrow \infty$ , as follows:

$$|A| \leq \frac{1}{n} \sum_{i=1}^n |d(X_{si}, \hat{\mu}_s) - d(X_{si}, \mu_s)| |d(X_{s'i}, \hat{\mu}_{s'}) - d(X_{s'i}, \mu_{s'})| \quad (\text{A.27})$$

$$\leq d(\hat{\mu}_s, \mu_s)d(\hat{\mu}_{s'}, \mu_{s'}). \quad (\text{A.28})$$

Since  $d(\hat{\mu}_s, \mu_s)$  and  $d(\hat{\mu}_{s'}, \mu_{s'})$  converge to zero in probability, so does  $|A|$ . For  $B$ ,

$$|B| = \left| \frac{1}{n} \sum_{i=1}^n [d(X_{si}, \hat{\mu}_s) - d(X_{si}, \mu_s)] d(X_{s'i}, \mu_{s'}) + d(X_{si}, \mu_s) [d(X_{s'i}, \hat{\mu}_{s'}) - d(X_{s'i}, \mu_{s'})] \right| \quad (\text{A.29})$$

$$\leq \frac{1}{n} \sum_{i=1}^n |d(X_{si}, \hat{\mu}_s) - d(X_{si}, \mu_s)| d(X_{s'i}, \mu_{s'}) + \frac{1}{n} \sum_{i=1}^n d(X_{si}, \mu_s) |d(X_{s'i}, \hat{\mu}_{s'}) - d(X_{s'i}, \mu_{s'})| \quad (\text{A.30})$$

$$\leq \frac{1}{n} \sum_{i=1}^n d(\hat{\mu}_s, \mu_s) d(X_{s'i}, \mu_{s'}) + d(X_{si}, \mu_s) d(\hat{\mu}_{s'}, \mu_{s'}), \quad (\text{A.31})$$

which converges to zero in probability from the consistency of  $\hat{\mu}_s$  and  $\hat{\mu}_{s'}$ , as well as the boundedness of both metric spaces. Finally,  $|C|$  converges to zero in probability by the weak law of large numbers and the boundedness of the metric spaces. Consistency follows from Slutsky's theorem.  $\square$

## Appendix B

### Enumerating Valid Spreads: The Naïve Approach

The set  $C_1$  is defined by multiple constraints. In this section we discuss how to obey those constraints while enumerating spreads. Theorem 1 limits the set of valid partition pairs to those where  $\lambda^l$  and  $\lambda^r$  are majorized by  $q^{l'}$  and  $q^{r'}$  respectively. Thankfully the Young diagram framework provides an intuitive way to enumerate all partitions majorized by some starting partition. Starting from  $q^{l'}$ , we may repeatedly apply  $v$  and  $h$  shifts to enumerate all candidates for  $p^l$  for which the conditions of Theorem 1 hold, similarly for  $q^{r'}$  and  $p^r$ .

Given candidates  $\{c_{(i)}^l\}$  for  $p^l$ , and candidates  $\{c_{(j)}^r\}$  for  $p^r$ , we must identify all possible pairs  $(c_{(i)}^l, c_{(j)}^r)$  which can add together to produce  $p$ . It is important to note that for integer partitions, order is unimportant, so the elements are ordered non-increasing by convention. However,  $c_{(i)}^l$  and  $c_{(j)}^r$  may require permuting in order to add to  $p$ . This gives rise to the 3D matching problem:

$$\sigma^l c_{(i)}^l + \sigma^r c_{(j)}^r = p \quad (\text{B.1})$$

where  $\sigma^l$  and  $\sigma^r$  are unknown permutation matrices. This problem is famously NP-Hard [103], and ostensibly must be solved for every possible  $(i, j)$  pair of candidates. Furthermore, our desire is not simply to discover *some*  $(\sigma^l, \sigma^r)$  which satisfy (B.1), rather *all* valid  $(\sigma^l, \sigma^r)$  which give a valid matching. As an example, consider the simple case  $p = (5, 4, 3)$  with candidates  $c^l = (3, 2, 1)$  and  $c^r = (3, 2, 1)$ . The choice of permutations is not unique, since  $\{\sigma^l = (2, 1, 3), \sigma^r = (1, 3, 2)\}$  and  $\{\sigma^l = (1, 3, 2), \sigma^r = (2, 1, 3)\}$  both give rise to a valid matching. This example shows that a 3D matching algorithm which produces *a* valid match may not be sufficient, and exhaustive search may be the only way to ensure all matches are found. This approach is therefore ill-advised, and below we outline a method of enumerating valid spreads that avoids the prohibitive computational complexity of the naïve method.