



Non-Uniform Sampling of Fixed Margin Binary Matrices

Alex Fout

fout@colostate.edu
Department of Statistics, Colorado
State University
Fort Collins, Colorado

Bailey K. Fosdick

bailey.fosdick@colostate.edu
Department of Statistics, Colorado
State University
Fort Collins, Colorado

Matthew P. Hitt

matthew.hitt@colostate.edu
Department of Political Science,
Colorado State University
Fort Collins, Colorado

ABSTRACT

Data sets in the form of binary matrices are ubiquitous across scientific domains, and researchers are often interested in identifying and quantifying noteworthy structure. One approach is to compare the observed data to that which might be obtained under a null model. Here we consider sampling from the space of binary matrices which satisfy a set of marginal row and column sums. Whereas existing sampling methods have focused on uniform sampling from this space, we introduce modified versions of two elementwise swapping algorithms which sample according to a non-uniform probability distribution defined by a weight matrix, which gives the relative probability of a one for each entry. We demonstrate that values of zero in the weight matrix, i.e. *structural zeros*, are generally problematic for swapping algorithms, except when they have special monotonic structure. We explore the properties of our algorithms through simulation studies, and illustrate the potential impact of employing a non-uniform null model using a classic bird habitation dataset.

CCS CONCEPTS

• **Mathematics of computing** → **Markov-chain Monte Carlo methods**; *Graph algorithms*; • **Theory of computation** → **Random walks and Markov chains**.

KEYWORDS

binary matrix, MCMC sampling, checkerboard swap, curveball

ACM Reference Format:

Alex Fout, Bailey K. Fosdick, and Matthew P. Hitt. 2020. Non-Uniform Sampling of Fixed Margin Binary Matrices. In *Proceedings of the 2020 ACM-IMS Foundations of Data Science Conference (FODS '20)*, October 19–20, 2020, Virtual Event, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3412815.3416887>

1 INTRODUCTION

In 1975, Diamond [7] 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 [6]

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 [1, 8, 11, 12, 19, 24, 25], Connor & Simberloff sparked interest in a novel form of “null model analysis” [13, 15], 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 1 if the corresponding species inhabits the corresponding island, and 0 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 1. 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 1’s. Connor & Simberloff propose two methods, one which begins with a 0-filled matrix and sequentially *fills* in each row randomly, being careful to adhere to the constraints, and another which begins with the observed matrix, and repeatedly *swaps* 1’s such that the constraints are maintained. Subsequent authors have proposed additional methods for generating the desired matrices, which are generally either a “fill” [12, 24, 25, 27] or “swap” [1, 19, 23, 25, 28] 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, 23]. However, focus eventually shifted from whether to use this approach to which method of producing matrices is best [24]. This accompanied theoretical work in binary matrices with fixed marginal sums [2] and Markov Chain Monte Carlo (MCMC) sampling [1], paving the way for swap and fill methods with proven statistical properties [4, 18].

Generally speaking, sampling from a null model can be useful for quantifying the degree of “non-randomness” in an observed matrix

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FODS 2020, October 18–20, 2020, Seattle, WA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8103-1/20/10...\$15.00

<https://doi.org/10.1145/3412815.3416887>



(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 [5], evaluating goodness of fit for the Rasch item-response model [1, 4], and in the analysis of contingency tables [4].

In this paper, 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.

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 0, 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.

1.1.1 Fill Methods. Fill methods begin with a matrix of all 0's, and randomly *fill* in 1's in such a way that adheres to the margin (row and column) constraints.

Connor & Simberloff's [6] 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 1's, until the desired row sum is reached. If the column is already 1, or it has met its column sum, another random selection is made. If the procedure reaches a point where no 1 can be filled 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 [25] adapt the fill method of Connor & Simberloff to check whether the proposed fill will lead to a dead end.

Sanderson et al [24] 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 row 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 [4] 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 [20] 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 [14] adopt the sequential importance sampling approach of Chen et al for non-uniform sampling. This is, to our knowledge, the only paper 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 0.

1.1.2 Swap Methods. Swap methods have developed in parallel to fill methods, beginning with the original paper of Connor & Simberloff [6]. These methods involve identifying “checkerboard” patterns of 0's and 1's in an existing valid matrix, and swapping the 1's for the 0's to produce a new matrix which still adheres to the constraints (see Figure 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 [2] showed this to be true using elementary circuits. This paved the way for more development in swap methods [19, 25].

As with fill methods, early swap methods gave no proof that matrices are sampled uniformly, until Kannan et al [18] 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 [12] 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 [28], 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 [26] 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 [3] proves that curveball samples uniformly and demonstrate that multiple swaps can lead to faster mixing compared to Kannan et al's algorithm.

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 [6] for not accounting for such “significant 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 probability 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 [14], we assume the existence of a weight matrix indicating the relative likelihood of a 1 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 0, which further restrict the set of matrices with positive probability, and affect the validity of the proposed sampling methods.

In the rest of this paper, 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 curveball MCMC algorithm, which in most circumstances is a faster alternative to the weighted checkerboard swap algorithm. We investigate the correctness, efficiency under different weighting and structural zero schemes via simulation. Finally, we revisit the bird species-island dataset first discussed by Diamond [7] and Connor & Simberloff [6] and consider incorporating a weight matrix in the null distribution to show the impact that non-uniform sampling can have on scientific conclusions.

2 NON-UNIFORM SAMPLING FROM BINARY MATRICES

2.1 Probability Model

Let $\mathcal{A}(\mathbf{r}, \mathbf{c})$ be the set of all $m \times n$ binary matrices satisfying row sums $\mathbf{r} = (r_1, r_2, \dots, r_m)$ and column sums $\mathbf{c} = (c_1, c_2, \dots, c_n)$. Furthermore let $W = (w_{ij}) \in [0, \infty)^{m \times n}$ be a non-negative weight matrix representing the relative likelihood of a 1 in element (i, j) . Following Harrison & Miller [14], define the probability of observed matrix $A = (a_{ij}) \in \mathcal{A}(\mathbf{r}, \mathbf{c})$ as

$$\mathbb{P}(A) = \frac{1}{\kappa} \prod_{ij} w_{ij}^{a_{ij}}, \quad \kappa = \sum_{A \in \mathcal{A}(\mathbf{r}, \mathbf{c})} \prod_{ij} w_{ij}^{a_{ij}}. \quad (1)$$

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 $\mathcal{A}'(\mathbf{r}, \mathbf{c}) = \{A \in \mathcal{A}(\mathbf{r}, \mathbf{c}) : \mathbb{P}(A) > 0\}$ be the set of matrices with positive probability. We interpret the weights by considering the




Figure 1: A 2×2 submatrix of A from rows i and i' , and columns j and j' , which are not necessarily consecutive. Swapping the positions of 1's and 0's preserves the row and column sums of A .

relative probability between two matrices $A, B \in \mathcal{A}'(\mathbf{r}, \mathbf{c})$:

$$\frac{\mathbb{P}(A)}{\mathbb{P}(B)} = \frac{\prod_{\substack{a_{ij}=1 \\ i,j: b_{ij}=0}} w_{ij}}{\prod_{\substack{b_{ij}=1 \\ i,j: a_{ij}=0}} w_{ij}}, \quad (2)$$

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

$$\mathbb{P}(B|\{A\} \cup \{B\}) = \frac{\prod_{\substack{b_{ij}=1 \\ i,j: a_{ij}=0}} w_{ij}}{\prod_{\substack{a_{ij}=1 \\ i,j: b_{ij}=0}} w_{ij} + \prod_{\substack{b_{ij}=1 \\ i,j: a_{ij}=0}} w_{ij}}. \quad (3)$$

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

2.2 Weighted Checkerboard Swap Algorithm

We propose a swapping algorithm similar to that of Brualdi [2], which identifies 2×2 submatrices with diagonal 1's and 0's, as in Figure 1, where the rows and columns are not necessarily consecutive. Let (i, j) and (i', j') be the indices of the 1's in the submatrix, so that (i, j') and (i', j) are 0's. The positions of the 1's and 0's can be swapped without altering the row or column sums. This swap operation produces a matrix $B \in \mathcal{A}(\mathbf{r}, \mathbf{c})$ which differs from A (see Figure 1). 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

$$\mathbb{P}(B|\{A\} \cup \{B\}) = \frac{w_{i'j}w_{ij'}}{w_{ij}w_{i'j'} + w_{i'j}w_{ij'}} := p_{ij,i'j'}. \quad (4)$$

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.

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

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

Proof of Theorem 1 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.

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  $A_{ij}^{(k)} = 1$  or  $A_{i'j'}^{(k)} = 1$  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  $A_{ij}^{(k)} = A_{ij}^{(k)} = 0$  and  $A_{i'j'}^{(k)} = A_{i'j'}^{(k)} = 1$  (perform
      | swap);
11  end
12 end

```

2.3 Structural Zeros

As mentioned, structural zeros have received little attention since originally introduced by Connor & Simberloff [6]. 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 [2] upper bound on the diameter of the Markov space no longer holds with structural zeros.

Given matrices $A, B \in \mathcal{A}(\mathbf{r}, \mathbf{c})$ 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 following matrices, both with unit row and column sums where structural zeros are represented with grey fill:

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad B = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}$$

The Hamming distance 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 1 per column: $((1, 2), (1, 3), (1, 4), (1, 5), (1, 2), (2, 3))$.

Additionally, the following example from Rao et al [22] illustrates how structural zeros can render it impossible to transition between two matrices via checkerboard swaps:

$$A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \quad B = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

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 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 1 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 1's 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:

Theorem 2. 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 (1).

The only modification of the proof for Theorem 1 needed to prove Theorem 2 is to show irreducibility of the chain given structural zeros (see Appendix C).

2.4 Weighted Curveball Algorithm

Strona et al. [26] proposed a faster version of checkerboard swapping called the curveball, which imagines two rows trading their 1's like baseball cards. Two rows are randomly selected, and the 1's in their symmetric difference become candidates for the 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 [3] proves that the unweighted curveball algorithm can be used to sample uniformly, and we contribute the following theorem for the weighted case:

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  $A_{i \setminus i'} = \{j : A_{ij} = 1, A_{i'j} = 0, w_{ij} > 0\}$ ;
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  $j \in \mathbf{j}_i$ ;
19    Set  $A_{ij}^{(k)} = 0$  and  $A_{i'j}^{(k)} = 1$  for  $j \in \mathbf{j}_{i'}$ ;
20  end
21 end

```

Theorem 3. 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 (1).

In Appendix B, we provide a proof of this theorem when there are no structural zeros, and in Appendix C this is extended to the case of monotonic structural zeros. R code implementing Algorithms 1 and 2 is provided on Github at (<https://github.com/fouticus/nsfbm>).

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 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 1's of a square ($n \times n$) matrix A are from its diagonal:

$$T(A) = \frac{1}{|A|n} \sum_{ij} |i - j| \mathbb{I}(A_{ij} = 1) \quad (5)$$

$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$	$B = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$	$C = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$
$D = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$	$E = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$	$F = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$

Figure 2: All 3×3 matrices with row and column sums equal to 1.

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

Table 1: Theoretical vs. empirical probabilities for matrices in Figure 2 with binary matrix and weight matrix as defined in (6). Empirical probabilities also show Tukey-Hanning [9] MCMC standard errors.

where $|A|$ is the number of 1's in A and $\mathbb{I}(\cdot)$ is the indicator function.

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 following binary matrix A , and weight matrix W :

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad W = \begin{bmatrix} 1 & 2 & 1 \\ 2 & 1 & 2 \\ 1 & 2 & 1 \end{bmatrix}. \quad (6)$$

Including A , there are six possible states with row and column sums equal to 1 (see Figure 2). Under a uniform distribution, all matrices would have probability $\frac{1}{6}$. W induces a non-uniform probability distribution among the matrices (see Table 1), where matrices A and D have lower probability because their non-zero elements correspond to 1's 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 1 shows agreement between the true and empirical probability distributions. The KL-Divergence between the empirical and true distributions is 9.0×10^{-4} .

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 [16]. 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 [21].

To assess the impact of weights, we constructed 20×20 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 1's, 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 0 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 20×20 binary starting matrix by setting each element to 1 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 shows ESS for the different conditions, where each line is the average over 10 individual chains with different random seeds. For illustration, Figure 4 shows ESS for all chains for the condition in the bottom-right panel of Figure 3.

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 1's 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 1's" weighting schemes. This is because the former weight schemes contain weights near 0, making 1 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 1 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 0 (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 0. 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 0, 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




Figure 3: Average effective sample size for the simulations described in Section 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 10 separate chains using the same starting matrix.

weights away from 0, 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 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 C).

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 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 50×50 , and define the base weight matrix as:

$$W_{ij} = \frac{100 - |i - j|}{100}, \quad (7)$$

which has value 1 on the diagonal and value near 0.5 in the corners (see Figure 5). Sampling proceeded with a burn-in period of 5,000




Figure 4: Detailed view of effective sample size for the simulations described in Section 3.2, using a weight matrix of all 1's and with 50% structural zeros. Average ESS is shown in bold, with each iteration shown partially transparent.




Figure 5: Initial binary matrix A and base weight matrix W for the simulation described in Section 3.3. Both use a shading scale where white denotes 0 and black denotes 1. 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 1's corresponding to lower weighted regions of W .




Figure 6: Sampling distribution of the diagonal divergence statistic, under different weighting matrices, defined by raising the base weight matrix to the power p .

proposed swaps, before retaining every 1,000th matrix, to generate a sample of 5,000 retained matrices.

Figure 6 shows the approximate sampling distributions of the diagonal divergence T for selected powers of the weight matrix. Since W favors matrices with 1's 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 1's 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 $\mathcal{A}'(\mathbf{r}, \mathbf{c})$ is as there is virtually no overlap in the sampling distributions under different weighting schemes.

4 EXAMPLE: NEW HEBRIDES BIRD SPECIES

We revisit the original Vanuatu (formerly New Hebrides) bird species data provided by Diamond in 1975 [7], 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 7a). The elements of the matrix are 1 if the corresponding species is found on the corresponding island, and 0 otherwise. If a checkerboard pattern exists between two species on two islands, this can indicate competitive exclusion (see [23]). Accordingly, Stone & Roberts [25] propose the C -score to measure competitive exclusion by counting how many checkerboards exist between species pairs in the matrix:

$$C(A) = \frac{2}{m(m-1)} \sum_{j=2}^m \sum_{i<j} (\mathbf{r}_i - S_{ij})(\mathbf{r}_j - S_{ij}). \quad (8)$$

Here, S_{ij} is the number of shared islands between species i and j , $\mathbf{r}_i$ 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 7b). This could result, for example, from several environmental and biological factors. We then proceed, as did Connor & Simberloff [6], to generate a sample of matrices under the null distribution, assuming equal weights (all 1's) 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 8, 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 7: Observed New Hebrides bird species combination A [7], and non-uniform weight matrix W for the data analysis described in Section 4. For A , white denotes 0 and grey denotes 1. For W , light grey denotes 1 and grey denotes 15. With higher values indicating species (rows) preference for certain island (columns) habitats.

5 DISCUSSION

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,




Figure 8: 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.

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 [11]: “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 0, 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.

Structural zeros have also received attention from Rao et al [22] who noted that when structural zeros are on the matrix diagonal, the Markov space may be disconnected when limited to “alternating rectangles”(checkerboard swaps). Their solution was to also consider “alternating hexagons”, which they prove connects the space. Alternative swapping mechanisms like this can be incorporated into the non-uniform setting straightforwardly, by defining the swap probability as a ratio of the relevant weights.

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). 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 [10]. However, we know of no implementation of graph sampling where edges are weighted, as in this paper.

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' :

$$\frac{\mathbb{P}(A)}{\mathbb{P}(B)} = \frac{w_{ij}w_{i'j'}}{w_{i'j}w_{ij}}. \quad (9)$$

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.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] Julian Besag and Peter Clifford. 1989. Generalized Monte Carlo significance tests. *Biometrika* 76, 4 (Dec. 1989), 633–642. <https://doi.org/10.1093/biomet/76.4.633>
- [2] Richard A. Brualdi. 1980. Matrices of zeros and ones with fixed row and column sum vectors. *Linear algebra and its applications* 33 (1980), 159–231.
- [3] C. J. Carstens. 2015. Proof of uniform sampling of binary matrices with fixed row sums and column sums for the fast Curveball algorithm. *Physical Review E* 91, 4 (April 2015), 042812. <https://doi.org/10.1103/PhysRevE.91.042812>
- [4] Yuguo Chen, Persi Diaconis, Susan P. Holmes, and Jun S. Liu. 2005. Sequential Monte Carlo methods for statistical analysis of tables. *J. Amer. Statist. Assoc.* 100, 469 (2005), 109–120.
- [5] Edward F. Connor, Michael D. Collins, and Daniel Simberloff. 2013. The checkered history of checkerboard distributions. *Ecology* 94, 11 (Nov. 2013), 2403–2414. <https://doi.org/10.1890/12-1471.1>
- [6] Edward F. Connor and Daniel Simberloff. 1979. The Assembly of Species Communities: Chance or Competition? *Ecology* 60, 6 (Dec. 1979), 1132. <https://doi.org/10.2307/1936961>
- [7] Jared M. Diamond. 1975. Assembly of Species Communities. In *Ecology and Evolution of Communities*. Harvard University Press, 342–443.
- [8] Jared M. Diamond and Michael E. Gilpin. 1982. Examination of the “null” model of Connor and Simberloff for species co-occurrences on islands. *Oecologia* 52, 1 (1982), 64–74.
- [9] James M. Flegal and Galin L. Jones. 2010. Batch means and spectral variance estimators in Markov chain Monte Carlo. *The Annals of Statistics* 38, 2 (April 2010), 1034–1070. <https://doi.org/10.1214/09-AOS735>
- [10] Bailey K. Fosdick, Daniel B. Larremore, Joel Nishimura, and Johan Ugander. 2018. Configuring Random Graph Models with Fixed Degree Sequences. *SIAM Rev.* 60, 2 (Jan. 2018), 315–355. <https://doi.org/10.1137/16M1087175> Publisher: Society for Industrial and Applied Mathematics.
- [11] Nicholas J. Gotelli. 2000. Null Model Analysis of Species Co-Occurrence Patterns. *Ecology* 81, 9 (2000), 2606–2621. [https://doi.org/10.1890/0012-9658\(2000\)081\[2606:NMAOSC\]2.0.CO;2](https://doi.org/10.1890/0012-9658(2000)081[2606:NMAOSC]2.0.CO;2)
- [12] Nicholas J. Gotelli and Gary L. Entsminger. 2001. Swap and fill algorithms in null model analysis: rethinking the knight’s tour. *Oecologia* 129, 2 (Oct. 2001), 281–291. <https://doi.org/10.1007/s004420100717>
- [13] Nicholas J. Gotelli and Gary R. Graves. 1996. *Null models in ecology*. Smithsonian.
- [14] Matthew T. Harrison and Jeffrey W. Miller. 2013. Importance sampling for weighted binary random matrices with specified margins. *arXiv:1301.3928 [math, stat]* (Jan. 2013). <http://arxiv.org/abs/1301.3928> arXiv: 1301.3928.
- [15] P H Harvey, R K Colwell, J W Silvertown, and R M May. 1983. Null Models in Ecology. *Annual Review of Ecology and Systematics* 14, 1 (1983), 189–211. <https://doi.org/10.1146/annurev.es.14.110183.001201>
- [16] Peter D. Hoff. 2009. *A first course in Bayesian statistical methods*. Springer, London; New York. OCLC: ocn310401109.
- [17] Mark Jerrum and Alistair Sinclair. 1996. The Markov chain Monte Carlo method: an approach to approximate counting and integration. In *Approximation algorithms for NP-hard problems*. PWS Publishing Co., Boston. Publisher: PWS Boston.
- [18] Ravi Kannan, Prasad Tetali, and Santosh Vempala. 1999. Simple Markov-chain algorithms for generating bipartite graphs and tournaments. *Random Structures & Algorithms* 13, 4 (1999), 16.
- [19] Bryan F. J. Manly. 1995. A Note on the Analysis of Species Co-Occurrences. *Ecology* 76, 4 (1995), 1109–1115. <https://doi.org/10.2307/1940919>
- [20] Jeffrey W. Miller and Matthew T. Harrison. 2013. Exact sampling and counting for fixed-margin matrices. *The Annals of Statistics* 41, 3 (June 2013), 1569–1592. <https://doi.org/10.1214/13-AOS1131>
- [21] Martyn Plummer, Nicky Best, Kate Cowles, and Karen Vines. 2006. CODA: Convergence Diagnosis and Output Analysis for MCMC. *R News* 6, 1 (2006), 7–11. <https://journal.r-project.org/archive/>
- [22] A. Ramachandra Rao, Rabindranath Jana, and Suraj Bandyopadhyay. 1996. A Markov chain Monte Carlo method for generating random (0, 1)-matrices with given marginals. *Sankhyā: The Indian Journal of Statistics, Series A* (1996), 225–242.
- [23] Alan Roberts and Lewis Stone. 1990. Island-sharing by archipelago species. *Oecologia* 83, 4 (1990), 560–567.
- [24] James G. Sanderson, Michael P. Moulton, and Ralph G. Selfridge. 1998. Null matrices and the analysis of species co-occurrences. *Oecologia* 116, 1-2 (1998), 275–283.
- [25] Lewi Stone and Alan Roberts. 1990. The checkerboard score and species distributions. *Oecologia* 85, 1 (1990), 74–79.
- [26] Giovanni Strona, Domenico Nappo, Francesco Boccacci, Simone Fattorini, and Jesus San-Miguel-Ayanz. 2014. A fast and unbiased procedure to randomize ecological binary matrices with fixed row and column totals. *Nature Communications* 5, 1 (June 2014), 1–9. <https://doi.org/10.1038/ncomms5114>
- [27] J. Bastow Wilson. 1987. Methods for detecting non-randomness in species co-occurrences: a contribution. *Oecologia* 73, 4 (1987), 579–582.
- [28] Arif Zaman and Daniel Simberloff. 2002. Random binary matrices in biogeographical ecology—instituting a good neighbor policy. *Environmental and Ecological Statistics* 9, 4 (2002), 405–421.



A PROOF OF WEIGHTED CHECKERBOARD ALGORITHM

Here we prove that the weighted checkerboard swap algorithm samples from $\mathbb{P}(\cdot)$ defined in (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 $\mathbb{P}(\cdot)$ [17, 18].

Irreducibility. Since $w_{ij} > 0$ we have $0 < p_{ij:i'j'} < 1$, so irreducibility follows directly from the result of Brualdi [2], 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 equation:

$$\mathbb{P}(A) p_{A \rightarrow B} = \mathbb{P}(B) p_{B \rightarrow A}. \quad (10)$$

for any matrices $A, B \in \mathcal{A}'(\mathbf{r}, \mathbf{c})$, 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 (10),

$$\mathbb{P}(A) p_{A \rightarrow B} = \quad (11)$$

$$= \left[\frac{1}{\kappa} \prod_{ij} w_{ij}^{a_{ij}} \right] \left[\binom{k}{2}^{-1} \frac{w_{u'v} w_{uv'}}{w_{uv} w_{u'v'} + w_{u'v} w_{uv'}} \right] \quad (12)$$

$$= \frac{1}{\kappa} \binom{k}{2}^{-1} \left[\prod_{ij \notin \mathcal{S}} w_{ij}^{a_{ij}} \right] w_{uv} w_{u'v'} \frac{w_{u'v} w_{uv'}}{w_{uv} w_{u'v'} + w_{u'v} w_{uv'}} \quad (13)$$

$$= \frac{1}{\kappa} \binom{k}{2}^{-1} \left[\prod_{ij \notin \mathcal{S}} w_{ij}^{b_{ij}} \right] w_{u'v} w_{uv'} \frac{w_{uv} w_{u'v'}}{w_{uv} w_{u'v'} + w_{u'v} w_{uv'}} \quad (14)$$

$$= \left[\frac{1}{\kappa} \prod_{ij} w_{ij}^{b_{ij}} \right] \left[\binom{k}{2}^{-1} \frac{w_{uv} w_{u'v'}}{w_{uv} w_{u'v'} + w_{u'v} w_{uv'}} \right] \quad (15)$$

$$= \mathbb{P}(B) p_{B \rightarrow A}, \quad (16)$$

so detailed balance holds.

B PROOF OF WEIGHTED CURVEBALL ALGORITHM

Here we prove that the weighted curveball algorithm samples from $\mathbb{P}(\cdot)$ defined in (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

$$p_c = |A_{u \setminus u'}|! |A_{u' \setminus u}|! / (|A_{u \setminus u'}| + |A_{u' \setminus u}|)! \quad (17)$$

$$= |B_{u \setminus u'}|! |B_{u' \setminus u}|! / (|B_{u \setminus u'}| + |B_{u' \setminus u}|)!, \quad (18)$$

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 $\mathcal{S} = \{ij : i \in \{u, u'\}, j \in \{j_i \cup j_{i'}\}\}$, 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 (10):

$$\mathbb{P}(A) p_{A \rightarrow B} = \mathbb{P}(A) p_c p_{uu'} \quad (19)$$

$$= \frac{p_c}{\kappa} \prod_{ij} w_{ij}^{a_{ij}} \left[\frac{w_{uj_i} w_{u'j_{i'}}}{w_{uj_i} w_{u'j_{i'}} + w_{u'j_i} w_{uj_{i'}}} \right] \quad (20)$$

$$= \frac{p_c}{\kappa} \left[\prod_{ij \notin \mathcal{S}} w_{ij}^{a_{ij}} \right] w_{u'j_i} w_{uj_{i'}} \left[\frac{w_{uj_i} w_{u'j_{i'}}}{w_{uj_i} w_{u'j_{i'}} + w_{u'j_i} w_{uj_{i'}}} \right] \quad (21)$$

$$= \frac{p_c}{\kappa} \left[\prod_{ij \notin \mathcal{S}} w_{ij}^{b_{ij}} \right] w_{uj_i} w_{u'j_{i'}} \left[\frac{w_{u'j_i} w_{uj_{i'}}}{w_{uj_i} w_{u'j_{i'}} + w_{u'j_i} w_{uj_{i'}}} \right] \quad (22)$$

$$= \frac{p_c}{\kappa} \prod_{ij} w_{ij}^{b_{ij}} \left[\frac{w_{u'j_i} w_{uj_{i'}}}{w_{uj_i} w_{u'j_{i'}} + w_{u'j_i} w_{uj_{i'}}} \right] \quad (23)$$

$$= \mathbb{P}(B) p_c p_{u'u} \quad (24)$$

$$= \mathbb{P}(B) p_{B \rightarrow A} \quad (25)$$

Hence the detailed balance equation holds, where the stationary distribution is again given by (1).

C IRREDUCIBILITY WITH MONOTONIC STRUCTURAL ZEROS

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 $A, B \in \mathcal{A}'(\mathbf{r}, \mathbf{c})$, 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 9).

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




Figure 9: Illustration of irreducibility proof in Appendix C. 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.

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 9, 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 9, 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 9).

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 9).

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:

$$0 = \sum_{l=1}^c d_{il} = \sum_{l \in \{1, \dots, c\} \setminus \{j, j'\}} d_{il}, \quad (26)$$

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

$$0 = \sum_{l=1}^c d_{i'l} \implies -2 = \sum_{l \in \{1, \dots, c\} \setminus \{j, j'\}} d_{i'l}, \quad (27)$$

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.