

Random Walks on the Symmetric Group: Cutoff for One-sided Transposition Shuffles

Oliver Thomas Matheau-Raven

PHD

UNIVERSITY OF YORK
MATHEMATICS

AUGUST 2020.

Abstract

In this thesis we introduce a new type of card shuffle called the *one-sided transposition shuffle*. At each step a card is chosen uniformly from the pack and then transposed with another card chosen uniformly from *below* it. This defines a random walk on the symmetric group generated by a distribution which is non-constant on the conjugacy class of transpositions. Nevertheless, we provide an explicit formula for all eigenvalues of the shuffle by demonstrating a useful correspondence between eigenvalues and standard Young tableaux. This allows us to prove the existence of a total-variation cutoff for the one-sided transposition shuffle at time $n \log n$. We also study weighted generalisations of the one-sided transposition shuffle called *biased one-sided transposition shuffles*. We compute the full spectrum for every biased one-sided transposition shuffle, and prove the existence of a total variation cutoff for certain choices of weighted distribution. In particular, we recover the eigenvalues and well known mixing time of the classical random transposition shuffle. We study the hyperoctahedral group as an extension of the symmetric group, and formulate the one-sided transposition shuffle and random transposition shuffle as random walks on this new group. We determine the spectrum of each hyperoctahedral shuffle by developing a correspondence between their eigenvalues and standard Young bi-tableaux. We prove that the one-sided transposition shuffle on the hyperoctahedral group exhibits a cutoff at $n \log n$, the same time as its symmetric group counterpart. We conjecture that this results extends to the biased one-sided transposition shuffles and the random transposition shuffle on the hyperoctahedral group.

Contents

Abstract	i
List of Tables	iii
List of Figures	iv
Introduction	v
Acknowledgements	viii
Declarations	ix
1 Preliminaries	1
1.1 Random Walks on Finite Groups	1
1.1.1 Markov Chains from Random Walks	1
1.1.2 Convergence to a Stationary Distribution	5
1.1.3 The Cutoff Phenomenon	11
1.2 Representations, Modules, and Applications to Random Walks	14
1.2.1 Representations, Characters and Modules	14
1.2.2 Discrete Fourier Transforms	21
1.2.3 The Upper Bound Lemma	23
2 The Symmetric Group and its Modules	25
2.1 The Symmetric Group	25
2.2 The Structure of Modules for The Symmetric Group	26
2.2.1 Young Diagrams	27
2.2.2 Young Tableaux	28
2.2.3 Permutation Modules	29
2.2.4 Specht Modules	31
2.2.5 Branching Rules for Specht Modules	34
2.2.6 Switching to Words	35
2.3 Random Walks on The Symmetric Group	37
2.3.1 The Random Transposition Shuffle	38
2.3.2 The Top-to-Random Shuffle	43
2.3.3 The Random-to-Random Shuffle	47
3 The One-sided Transposition Shuffle	54
3.1 Main Results	54
3.2 Lifting Eigenvectors for Transposition Shuffles	56
3.2.1 Lifting Eigenvectors for the One-sided Transposition Shuffle	56
3.2.2 Lifting Eigenvectors for the Random Transposition Shuffle	66
3.3 Eigenvalues Analysis	68

3.4	Upper Bound for the One-sided Transposition Shuffle	73
3.4.1	Large Partitions	75
3.4.2	Small Partitions	76
3.5	Lower Bound for the One-sided Transposition Shuffle	79
3.6	Biased One-sided Transposition shuffles	82
3.6.1	Cutoff for Biased One-sided Transposition Shuffles with $\alpha \leq 1$	85
3.6.2	Cutoff for Biased One-sided Transposition Shuffles with $\alpha \geq 1$	90
3.6.3	Cutoff for Biased One-sided Transposition Shuffles with General Weight Functions	96
3.7	Separation Distance for the One-sided Transposition Shuffle	98
3.7.1	Cutoff for the Unbiased One-sided Transposition Shuffle	98
3.7.2	Generalising for Biased One-sided Transposition Shuffles	101
4	The Hyperoctahedral Group and Random Walks	104
4.1	The Hyperoctahedral Group	104
4.2	The Structure of Modules for the Hyperoctahedral Group.	106
4.2.1	Bi-partitions	106
4.2.2	Bi-tableaux	108
4.2.3	The Module Structure of the Hyperoctahedral Group	109
4.2.4	Switching to Words	114
4.3	The Random Transposition Shuffle on The Hyperoctahedral Group	117
4.3.1	Upper Bound for the Random Transposition Shuffle	118
4.3.2	Lifting Eigenvectors for the Random Transposition Shuffle	121
4.4	One-sided Transposition Shuffles on The Hyperoctahedral Group	130
4.4.1	Lifting Eigenvectors for One-sided Transposition Shuffles	131
4.4.2	A Strong Stationary Time for the Unbiased One-sided Transposition Shuffle on The Hyperoctahedral Group.	132
4.5	Further Work	136
	Bibliography	140

List of Tables

1.1	The values of separation and total variation distance for the simple random walk on \mathbb{Z}_5	8
1.2	The action of two group homomorphisms of the symmetric group S_3	15
3.1	Eigenvalues corresponding to $T \in \text{SYT}((3, 2))$	69
3.2	A list of tableaux formed from T in Example 3.3.4 by transposing the value 9 with a value from $\{1, \dots, 8\}$	70
3.3	The eigenvalues $T_\lambda^{\rightarrow}, T_\lambda^{\downarrow}$ for every partition $\lambda \vdash 4$	72
3.4	Asymptotics of $t_{n,\alpha} \log n$ as $n \rightarrow \infty$	85
3.5	Asymptotics of binomial and factorial terms using Stirling's approximation	88

List of Figures

1.1	Total variation distance of Ehrenfest's urn from two different starting distributions	11
2.1	Young's lattice for partitions of size $n \in \{0, 1, 2, 3, 4\}$	28
4.1	Young's lattice for bi-partitions of size $n \in \{0, 1, 2, 3\}$	107

Introduction

Consider a stacked deck of n distinct cards, whose positions are labelled by elements of the set $[n] := \{1, \dots, n\}$ from bottom to top. Any shuffle which involves choosing two positions and switching the cards found there (if the two positions coincide then no cards are moved) is called a *transposition shuffle*, and may be viewed as a random walk on the symmetric group S_n . We refer to the two positions as being picked by our right and left hands, associating with them random variables, R^i and L^i respectively, which take values in $[n]$. Diaconis and Shahshahani [18] were the first to study random walks on groups using representation theory; they famously showed that the *random transposition shuffle*, in which the two positions are chosen independently and uniformly on $[n]$, takes $(n/2) \log n$ steps to randomise the order of the deck. The time taken to randomise the order of the deck is known as the *mixing time* of the shuffle.

We may expect random walks to converge to their stationary distribution smoothly as time increases but there are many classes of random walks for which convergence happens only once a certain threshold time has been met. This behaviour was first observed in the random transposition shuffle by Diaconis and Shahshahani and has been given the name *the cutoff phenomenon*. Suppose the mixing time of a succession of random walks on groups $\{G_n\}_{n \in \mathbb{N}}$ may be tightly bounded above and below by a time of the form $t_n \pm cw_n$, where c is of constant order and $w_n = o(t_n)$, then the sequence of random walks is said to exhibit a *cutoff* at time t_n . The mixing time of the random transposition shuffle may be shown to be upper and lower bounded around at time $(n/2) \log n \pm cn$, thus it shows a cutoff at time $(n/2) \log n$. The cutoff phenomenon is prevalent throughout the study of random walks on groups with many examples shown to exhibit a cutoff. However, there is currently no known sufficient criterion for the existence of a cutoff, and proving a cutoff for any particular random walk is generally difficult.

Since the analysis of the random transposition shuffle a variety of algebraic and probabilistic techniques have been employed to study the mixing time of different transposition shuffles. Notable examples include the *top-to-random transposition* shuffle [39], the *adjacent transposition* shuffle [29], and the *random transposition shuffle on a cycle* where our positions are constrained to be within some fixed distance [8]. All of these shuffles have the property that, at each step, the transposition to be applied is chosen uniformly from a subset of transpositions which generate the entire group S_n .

An interesting class of transposition shuffles which are not chosen uniformly from a set are the *semi-random transposition shuffles* (see e.g. [34, 36, 37]). Any semi-random transposition shuffle is driven by the following procedure: pick one position with our right hand uniformly at random, and then independently choose a second position with our left hand following some stochastic (possibly time-inhomogeneous) process on the set $[n]$, then transpose the cards in the chosen positions. This class includes the random transposition shuffle by letting the stochastic process for the left hand be repeated uniform draws from $[n]$. Mossel, Peres, and Sinclair [34] were able to establish a universal upper bound of $O(n \log n)$ on the mixing time of any semi-random transposition shuffle. The key idea behind this proof relies on the positions chosen being independent. The card picked by our stochastic process is uniformly transposed with a card in the deck, therefore, once our

stochastic process has picked every possible card our deck will be at a uniformly chosen permutation of S_n . This reduces the analysis of the mixing time to a coupon collector's problem, and thus we find an upper bound of order $n \log n$.

In this thesis we introduce a new class of shuffles called the *one-sided transposition shuffles*: these have the defining property that at step i the right hand's position (R^i) is chosen according to some arbitrary distribution on $[n]$, and given the value of R^i the distribution of the left hand's position (L^i) is supported on the set $\{1, \dots, R^i\}$. When our right hand is chosen uniformly from $[n]$ the one-sided transposition shuffle follows a similar description to that of a semi-random transposition shuffle however with the key difference that our choice for L^i now depends on our first uniform choice. This change of dependence greatly influences the behaviour of the one-sided transposition shuffles. There does not exist a universal upper bound on the mixing time of the one-sided transposition shuffles without imposing further constraints on the distribution of the left hand, since a arbitrarily slow shuffle can be created by increasing the probability that the two hands choose the same position. For the majority of Chapter 5 we focus on the case when the right and left hands are both chosen uniformly from their possible ranges: we call this the *(unbiased) one-sided transposition shuffle*. Although the support for the one-sided transposition shuffle is the entire conjugacy class of transpositions, our probability distribution on this set is in general far from uniform. (E.g. when right and left hands are both uniformly chosen from their permitted ranges, the probabilities attached to different transpositions range from $1/n^2$ to $1/2n$)

The main results of this thesis are as follows: we recover the eigenvalues for the one-sided transposition shuffle using the technique of lifting eigenvectors which was recently constructed by Dieker and Saliola [19]. Analysis of the spectrum for the one-sided transposition shuffle allows us to find tight bounds on the mixing time of the shuffle and prove that it exhibits a cutoff at time $n \log n$. Afterwards we modify our analysis to study one-sided transposition shuffles where the right hand R^i is chosen via a weighted distribution: we call these the *biased one-sided transposition shuffles*. We extend the technique of lifting eigenvectors to compute the full spectrum of all biased one-sided transposition shuffles. Furthermore, for particular choices of weighted distribution we are able to prove the existence of a cutoff for the biased one-sided transposition shuffles. We go on to study random walks on the hyperoctahedral group B_n as an extension of the symmetric group where each card now has a distinguishable back and front. We create transposition shuffles on B_n from existing shuffles on S_n by allowing ourselves a chance of turning cards over. We show that the method of lifting eigenvectors may be modified to the setting of the hyperoctahedral group, and compute the spectrum of the random transposition shuffle and one-sided transposition shuffle on B_n . We end by showing that the unbiased one-sided transposition shuffle on the hyperoctahedral group exhibits a cutoff at time $n \log n$, the same time as its symmetric group counterpart.

The structure of this thesis is as follows: Chapter 1 gives an introduction to the topic of random walks on groups, including all the necessary probabilistic and algebraic background. We restrict our attention to random walks which converge to a unique stationary

distribution and introduce total variation distance to measure the convergence of random walks to their stationary distributions. We define the mixing time of a random walk, and what it means for a random walk to exhibit a cutoff in total variation distance. We recall the definitions of representations, modules and characters for a group G , and state classical results including Schur's Lemma and Maschke's Theorem. We also explore the properties of discrete Fourier transforms on a group G . Finally we recall the *upper bound lemma* which allows us to analyse the mixing time of a random walk on a group G using the irreducible representations of G .

In Chapter 2 we specialise the results of Chapter 1 to the symmetric group S_n . We build up a picture of the structure of the symmetric group including a construction of its simple modules. To demonstrate a variety of techniques used to analyse random walks on groups we present three longer examples: the random transposition shuffle [18], the top-to-random shuffle [2], and the random-to-random shuffle [19]. The most recent of these is the random-to-random shuffle where Dieker and Saliola first used the technique of lifting eigenvectors. We explore their work showing that the eigenvalues of the random-to-random shuffle may be computed by transforming eigenvalues of the random-to-random shuffle on n cards to those of the shuffle on $n + 1$ cards.

The main body of Chapter 3 is comprised of a paper jointly written with Michael E. Bate and Stephen B. Connor [3]. In this chapter we explore the one-sided transposition shuffle which stems from a modification of the classical random transposition shuffle. Our main result is proving that the one-sided transposition shuffle exhibits a cutoff in total variation distance at time $n \log n$. To prove this result we analyse the spectrum of the shuffle which we compute by inductively lifting the eigenvectors for the shuffle. We present the details of our lifting in full as it differs in interesting ways from that of the random-to-random shuffle, in particular we are able to recover all eigenvectors and eigenvalues for the one-sided transposition shuffle. We generalise our lifting to recover eigenvalues for biased one-sided transposition shuffles and for particular choices of weighted distribution show them to exhibit a cutoff in total variation distance. Finally we show the one-sided transposition shuffle to exhibit a cutoff in separation distance at time $n \log n$ via the use of a strong stationary time.

Chapter 4 establishes the hyperoctahedral group B_n as an extension of the symmetric group. We recall facts about the hyperoctahedral group, relating them to the definitions and results of the symmetric group in Chapter 2. We examine the module structure of B_n and construct its permutation and simple modules which are now indexed by bi-partitions of n . We extend the random transposition shuffle and one-sided transposition shuffle to random walks on the group B_n by adding in a chance to flip cards over during our shuffling procedure. We generalise the technique of lifting eigenvectors to the hyperoctahedral group allowing us to recover the entire spectrum of the random transposition shuffle and one-sided transposition shuffle. We conjecture that the two shuffles on the hyperoctahedral group exhibit a cutoff in total variation distance at the same time as their symmetric group counterparts.

Acknowledgements

First and Foremost, I would like to thank Michael Bate and Stephen Connor for introducing me to this wonderful area of mathematics. Thanks to your attentive supervision and patience I have grown as a mathematician and as a communicator in general. This thesis would not have been possible without your help.

Thank you to my examiners, Harry Geranios and Evita Nestoridi, for making my viva an enjoyable experience and for your comments on this thesis. Thank you to EPSRC and the department of mathematics for their the financial support which made my PhD possible. I also want to thank all the staff at York for their support, particularly Vicky Gould and Alet Roux for being on my TAP.

I want to give a big thank you to the rats of Millfield Road: Allan, Chris, Christoph, and Scott, for being the best housemates, board game players and quarantine companions I could ask for. Thank you to the whole mathematics PhD community at York, including those that have left and those that have just joined, for all the Wednesday talks, pub trips and tea breaks, which provided needed respite from mathematics. Thank you to the archery society for fostering a hopefully life long hobby in me. I also extend my gratitude to everyone who has been a part of my PhD journey over the last three years.

Finally, I want to thank my family for all the love and support they have always given me.

Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has been carried out under the supervisor of Dr. Michael E. Bate, and Dr. Stephen B. Connor, and has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References. Chapters 1 and 2, form a review of background material needed for this thesis, and the appropriate literature is referenced before each chapter. Chapter 3 is based on a paper written in collaboration with Michael E. Bate and Stephen B. Connor, which has been accepted for publication in the *Annals of Applied Probability* [3]. Chapter 4 consists of original work by the author.

This thesis is dedicated to the memory of little gran.

Chapter 1

Preliminaries

1.1 Random Walks on Finite Groups

There are many different ways we may formulate a random walk on a group G . We could designate one step transition probabilities for every two states $g \rightarrow h$, or even have the probabilities depend on our past like a self-avoiding random walk. However, in the first case we have not used any of the symmetry a group has, and in the second case our walk is not Markovian. We want to exploit the structure of groups to make our random walk inherit sensible properties and we also want the walk to form a Markov chain. In this thesis we focus on one particular Markovian description of a random walk where the walk is driven by a single probability distribution $P : G \rightarrow [0, 1]$. All groups we consider are finite and have identity element e , unless otherwise stated. A detailed account of the theory of Markov chains presented in this section can be found in *Markov Chains and Mixing Times* by Levin, Peres, and Wilmer [32]. We begin this section with the definition of a random walk on a group before moving on to recalling important definition and facts about Markov chains. The section ends with the introduction of the cutoff phenomenon for random walks on groups.

1.1.1 Markov Chains from Random Walks

Definition 1.1.1. Let G be a finite group, suppose $P : G \rightarrow [0, 1]$, and $\mu : G \rightarrow [0, 1]$, are probability distributions on G . Define each individual step of our random walk as sequence of i.i.d. (independent and identically distributed) random variables $\{\xi^t\}_{t \geq 1}$ on G distributed according to P . Create a sequence of random variables $\{X^t\}_{t \geq 0}$, by setting $X^0 \sim_d \mu$ and inductively defining

$$X^{t+1} = \xi^{t+1} X^t$$

for all $t > 0$. Then $\{X^t\}_{t \geq 0}$ defines a Markov chain with one step transition probabilities $P(g, h) := P(hg^{-1})$. We call this the *random walk on G driven by P* .

Given a random walk on G , its driving probability P gives rise to the transition matrix of our random walk defined using one step transition probabilities $P(g, h)$. We distinguish between these two interpretations of the symbol P by the number of arguments each one takes. Note that from the transition matrix we may recover the driving probability by setting $P(g) = P(e, g)$.

Definition 1.1.2. Let P, Q be probabilities on a finite group G , define the *convolution* of P with Q , denoted $P \star Q : G \rightarrow [0, 1]$ as

$$(P \star Q)(g) := \sum_{h \in G} P(gh^{-1})Q(h).$$

This is the probability of a random walk ending at element g if we start at e and take the first step according to Q and the second step to P . We denote $P \star P = P^{\star 2}$, and the t^{th} convolution of P with itself as $P^{\star t}$.

Lemma 1.1.3. Let P^t be the t^{th} power of our transition matrix P . Then $P^t(g, h) = P^{\star t}(hg^{-1})$. Hence, we drop the \star in all subsequent work and just write P^t .

Proof. We proceed by induction, by definition the statement holds for $t = 1$. Now

$$\begin{aligned} P^{t+1}(g, h) &= \sum_{a \in G} P(g, a)P^t(a, h) = \sum_{a \in G} P(ag^{-1})P^{\star t}(ha^{-1}) \\ &= \sum_{b \in G} P(b)P^{\star t}(hg^{-1}b^{-1}) = P^{\star(t+1)}(hg^{-1}). \quad \square \end{aligned}$$

Lemma 3 tells us that convolution and matrix multiplication amount to the same thing. We use P^t to stand for both the t -step transition matrix and the t -fold convolution of our probability P , with the assumption that if a starting state is not specified we assume it to be the identity.

Our starting state X^0 is picked via a probability distribution μ , i.e. $\mathbb{P}(X^0 = g) = \mu(g)$. Given the starting distribution μ the probability of being in state g at time t is $(P^t \mu)(g)$. In practice we often fix the starting distribution to be a single element of our state space, usually the identity of our finite group. We note, however, that the choice of starting distribution can have a large impact on the behaviour of a random walk, as the following example shows.

Example 1.1.4. The simple random walk on \mathbb{Z}_n , is the random walk generated by probability $P : \mathbb{Z}_n \rightarrow [0, 1]$ with $P(1) = P(-1) = 1/2$.

Consider P to be the simple random walk on \mathbb{Z}_6 . Suppose we start at state 0 then after 3 steps we have $P^3(0, 0) = 0$ and $P^3(0, 1) = 6/16$. This is because we always take one step at each time meaning at odd times we are on odd elements on \mathbb{Z}_6 and at even times we are on even elements. Instead now suppose we start at element 1, then $P^3(1, 0) = 6/16$, and $P^3(1, 1) = 0$ for the same reason as above. Combining the two starting distributions let μ now have $\mu(0) = \mu(1) = 1/2$ then $(P^3 \mu)(0) = \frac{1}{2}(P^3(0, 0) + P^3(1, 0)) = 3/16 = (P^3 \mu)(1)$.

Understanding how the distribution P^t evolves in time is key to the study of Markov chains. In theory we could always compute P^t given the probability P but this becomes impractical as we consider larger state spaces and behaviour at times with t large. Instead, we often look to bound these probabilities to understand their behaviour without exact computation. Development of tools to help understand how $P^t \mu$ behaves are an active area of research. Often we can use characteristics of our random walk: the starting distribution μ , the transition matrix P , and the group G , in order to form bounds on the probability distribution $P^t \mu$. In particular the eigenvalues of the transition matrix P play a big part in its behaviour. Furthermore, random walks on groups allow the use of algebraic tools using

the representations of G , which we explore in section 1.2. Next we recall the definition of stationary distributions for a Markov chain.

Definition 1.1.5. Let $\{X^t\}_{t \geq 0}$ be a Markov chain with transition matrix P . A *stationary distribution* for the Markov chain is a probability distribution π such that $P\pi = \pi$. Note that for any stationary distribution π we have $P^t\pi = \pi$ for all $t \geq 1$.

Generic Markov chains do not necessarily have stationary distributions nor do they have to be unique.

Example 1.1.6. Let $G = \mathbb{Z}$, consider a Markov chain started at 0 and driven by probability $P(1) = 1$. This random walk has no stationary distribution because for any distribution μ on \mathbb{Z} we have $(P\mu)(i) = \mu(i + 1)$.

Let $G = \mathbb{Z}_2^2$, consider a random walk with probability $P((0,0)) = P((0,1)) = 1/2$. Suppose we have μ_0, μ_1 defined by $\mu_0((i,j)) = 1/2$ if $i = 0$, and $\mu_1((i,j)) = 1/2$ if $i = 1$. Then both μ_0, μ_1 are stationary distributions for P . In fact let $a \in [0, 1]$ and μ_a be defined by

$$\mu_a((i,j)) = \begin{cases} \frac{1-a}{2} & \text{if } i = 0 \\ \frac{a}{2} & \text{if } i = 1 \end{cases}.$$

Then μ_a is a stationary distribution for P , so we have infinitely many stationary distributions for this random walk.

Lemma 1.1.7. Let π denote the uniform distribution on a finite group G , that is for all $g \in G$, $\pi(g) = 1/|G|$. Then π is a stationary distribution for any random walk on G .

Proof. To prove this all we need to show is that $P\pi = P \star \pi = \pi$ for any probability P defined on G . Following from the definition

$$(P \star \pi)(g) = \sum_{h \in G} P(gh^{-1})\pi(h) = \frac{1}{|G|} \sum_{h \in G} P(gh^{-1}) = \pi(g).$$

□

Under mild assumptions we can prove that a Markov chain has a unique stationary distribution which we denote by π . Below we list some of the properties a Markov chain may exhibit.

Definition 1.1.8. Let $\{X^t\}$ be a Markov chain on a space \mathcal{X} with transition matrix P . The Markov chain may have the following properties.

Irreducible: $\{X^t\}$ is called *irreducible* if for all x, y , there exists $t > 0$ such that $P^t(x, y) > 0$.

Aperiodic: Let $r(x) = \{t \geq 1 : P^t_{x,x} > 0\}$ these are the *return times* of the state x , the *period* of x is $\gcd(r(x))$. The Markov chain is called *aperiodic* if all states have period 1.

Transitive: A Markov chain is called *transitive* if for all pairs $(x_1, x_2) \in \mathcal{X} \times \mathcal{X}$ there exists a bijection ϕ , such that $\phi(x_1) = x_2$ and ϕ preserves all one step transition probabilities, i.e. for all $y, z \in \mathcal{X}$ we have $P(y, z) = P(\phi(y), \phi(z))$.

Reversible: A Markov chain is called *reversible* if there exists a probability distribution π on \mathcal{X} such that

$$\pi(x)P(x, y) = \pi(y)P(y, x) \quad (1.1)$$

for all $x, y \in \mathcal{X}$. In this case π is a stationary distribution. The set of equations defined by (1.1) are called the *detailed balance equations*.

All random walks we study in detail in this thesis will be irreducible, aperiodic, and transitive. We shall see one by one why these conditions are necessary for the study of mixing times of random walks on groups. First we prove that all random walks on finite groups are transitive.

Lemma 1.1.9 (Section 2.6.2 [32]). *Random walks on finite groups define transitive Markov chains.*

Proof. Let $\{X^t\}_{t \geq 0}$ be a random walk on a finite group with driving probability P . Take a pair $(x_1, x_2) \in G \times G$, and define a bijection $\phi_{x_1, x_2} : G \rightarrow G$ by $\phi_{x_1, x_2}(g) = gx_1^{-1}x_2$. Then $\phi_{x_1, x_2}(x_1) = x_2$, and for any pair $g, h \in G$ we get

$$P(g, h) = P(hg^{-1}) = P((hx_1^{-1}x_2)(x_2^{-1}x_1g^{-1})) = P(\phi_{x_1, x_2}(g), \phi_{x_1, x_2}(h)) \quad \square.$$

Irreducibility implies that the Markov chain has a unique stationary distribution [32, Corollary 1.17]. Therefore, by Lemma 1.1.7 if a random walk on a finite group G is irreducible its unique stationary distribution must in fact be the uniform distribution.

Lemma 1.1.10. *Let P be the transition matrix of an aperiodic irreducible Markov chain, there exists a unique stationary distribution π for P .*

Corollary 1.1.11. *Let P be a probability on a finite group G . Define the support of P as the set $\text{Supp}(P) = \{g : P(x) > 0\}$. Suppose $\text{Supp}(P)$ is a generating set for G , then the random walk on G driven by P is irreducible, and hence has the uniform distribution as its unique stationary distribution.*

Suppose a random walk P on a group G is not irreducible, then if we restrict our state space to the subgroup $\langle \text{Supp}(P) \rangle = H \triangleleft G$ we recover an irreducible random walk on group H . When choosing random walks on groups we like to fix the support of a probability to be a set of generators $S \subseteq G$ and vary our probability to see what different behaviour can arise. One important case we focus on is transposition shuffles of the symmetric group. We shall see that the one-sided transposition shuffle has very different behaviour to the random transposition shuffle despite being generated by the same conjugacy class. If our support S is a union of conjugacy classes of our group G then we gain access to extra algebraic tools to help our analysis, we explore these in section 1.2. From here onwards we assume all the random walks on groups are irreducible with stationary distribution $\pi : G \rightarrow G$ being uniform.

Once we know a unique stationary distribution exists our next question is whether our probability P^t will ever reach this equilibrium? Convergence of P^t to a stationary distribution generally depends on the starting distribution of our Markov chain.

Example 1.1.12. Consider the simple random walk on \mathbb{Z}_n started at 0, this Markov chain is irreducible therefore has stationary distribution $\pi_n(i) = 1/n$. Suppose $n > 2$ is even, then the walk has a period of 2: at all even times X^{2t} we must be at an even integer, and at all odd times X^{2t+1} we are at an odd integer. Therefore, our probability P^t will always be distinguishable from π . However, if we start at the distribution $\mu(0) = \mu(1) = 0.5$, then we have $(P^t\mu)(i) \rightarrow 1/n$ for all i as $t \rightarrow \infty$.

Consider instead $n > 2$ being odd, the simple random walk on \mathbb{Z}_n has period 1. Now for any single starting point i we have $P^t(i, j) \rightarrow 1/n$ for all j as $t \rightarrow \infty$.

The key difference between the two random walks in Example 1.1.12 is their period. In Section 1.2.2. we establish that aperiodic random walks always converge to their stationary distributions. To get rid of periodicity concerns for a random walk, we introduce the notion of a *lazy* random walk.

Lemma 1.1.13. *Let $\{X^t\}$ be an periodic Markov chain driven by P . Create a new Markov chain $\{Y^t\}$ with $Y^0 = X^0$, and new transition probabilities Q created in the following way: flip a fair coin, if heads do nothing, otherwise proceed according to P . Therefore, $Q(e) = \frac{1}{2} + \frac{1}{2}P(e)$ and $Q(g) = \frac{1}{2}P(g)$ for $g \neq e$. The Markov chain $\{Y^t\}$ is called the lazy version of $\{X^t\}$, and is aperiodic.*

Proof. The lazy walk has $Q(e) > 1/2$, and so for any $g \in G$ we find $r(g) = \{t \geq 1 : Q^t(g, g) > 0\} = \mathbb{N}$. Hence, all states have period 1. \square

Example 1.1.14. The lazy simple random walk on \mathbb{Z}_n has probability $P(0) = 1/2$ and $P(-1) = P(1) = 1/4$. This walk is now aperiodic for all $n \geq 2$.

The proof of Lemma 1.1.13 exploits the fact that so long as we have $P(e) > 0$ our random walk is aperiodic. All random walks we study we detail will be aperiodic because they will have a non-zero probability of remaining still. In the next section we define pair of measures on the space of probability distributions, and define a notion of convergence with respect to these measures.

1.1.2 Convergence to a Stationary Distribution

Consider a Markov chain $\{X^t\}$ with transition matrix P and stationary distribution π . We know that if our chain starts at $X^0 \sim_d \pi$ then we remain at the distribution π after every step. We show that for any starting distribution an irreducible, aperiodic Markov chain always converges to its unique stationary distributions. To prove this we first establish a measure on our space of probability distributions, and then use this to define convergence. To this end we introduce two notions of distance on the space of probability distributions over a finite group G .

Total Variation Distance

The first metric we introduce is ubiquitous within the study of Markov chains, it is called total variation distance.

Definition 1.1.15. Let μ, ν be two probabilities on a finite group G . Define the *total variation distance* between μ, ν as follows:

$$\|\mu - \nu\|_{\text{TV}} = \sup_{A \subseteq G} |\mu(A) - \nu(A)|. \quad (1.2)$$

The total variation distance between any two probabilities always lies in the range $[0, 1]$. The formulation of total variation distance provided in Definition 1.1.15 is frequently too cumbersome for use. We may reformulate definition (1.2) to involve summing over the group G instead of its subsets.

Lemma 1.1.16 (Section 4.1 [32]). *Let μ, ν be two probabilities on G , then we have*

$$\|\mu - \nu\|_{\text{TV}} = \frac{1}{2} \sum_{g \in G} |\mu(g) - \nu(g)|.$$

Proof. Let $A = \{g : \mu(g) \geq \nu(g)\}$, and so $A^c = \{g : \nu(g) > \mu(g)\}$. First notice that the supremum in Definition 1.1.15 is reached by set A , and secondly that $\mu(A) - \nu(A) = \nu(A^c) - \mu(A^c)$. Putting these facts together we find,

$$\|\mu - \nu\|_{\text{TV}} = \mu(A) - \nu(A) = \frac{1}{2} (\mu(A) - \nu(A) + \nu(A^c) - \mu(A^c)) = \frac{1}{2} \sum_{g \in G} |\mu(g) - \nu(g)|. \quad \square$$

Following from our definitions it is now easy to see that total variation distance forms a metric on the space of probability distributions of G (the triangle inequality follows from 1.1.16).

We are interested in measuring the convergence of a Markov chain P^t to its stationary distribution. To understand this we need to analyse the total variation distance $\|P^t(g, \cdot) - \pi\|_{\text{TV}}$ as a function of t . The following lemma demonstrates that $\|P^t(g, \cdot) - \pi\|_{\text{TV}}$ is a non-increasing function; this follows the logic that performing another step of our random walk should never bring us further away from our equilibrium.

Lemma 1.1.17. *Let P be a transition matrix and μ, ν be probabilities on G . Then*

$$\|P\mu - P\nu\|_{\text{TV}} \leq \|\mu - \nu\|_{\text{TV}}.$$

Proof. Going from left to right we have:

$$\begin{aligned} \|P\mu - P\nu\|_{\text{TV}} &= \frac{1}{2} \sum_{g \in G} |P\mu(g) - P\nu(g)| \\ &= \frac{1}{2} \sum_{g \in G} \left| \sum_{h \in G} P(gh^{-1})(\mu(h) - \nu(h)) \right| \\ &\leq \frac{1}{2} \sum_{g \in G} \sum_{h \in G} |P(gh^{-1})| |\mu(h) - \nu(h)| = \|\mu - \nu\|_{\text{TV}}. \quad \square \end{aligned}$$

Corollary 1.1.18. *Let P be a transition matrix for a random walk on a finite group G with stationary distribution π , then for all $t \in \mathbb{N}$ and $g \in G$ we have:*

$$\|P^{t+1}(g, \cdot) - \pi\|_{\text{TV}} \leq \|P^t(g, \cdot) - \pi\|_{\text{TV}}$$

Proof. This is a consequence of Lemma 1.1.17 with $\mu = P^t(g, \cdot)$, $\nu = \pi$. \square

Our definition of total variation distance so far has depended on our starting state g . When bounding the speed at which a random walk converges to its stationary distribution we look at the time it takes from its worst possible starting state. For irreducible random walks on finite groups every starting state gives the same total variation distance.

Lemma 1.1.19. *Let P define an irreducible random walk on G . Then for any $g, h \in G$ we have*

$$\|P^t(g, \cdot) - \pi\|_{\text{TV}} = \|P^t(h, \cdot) - \pi\|_{\text{TV}}.$$

Proof. Let P be the transition probability for our Markov chain, and ϕ be our probability preserving bijection with $\phi(g) = h$, from the definition of transitivity. It follows that from the proof of Lemma 1.1.9 that $P^t(a, b) = P^t(\phi(a), \phi(b))$ for all $a, b \in G$. Therefore,

$$\|P^t(g, \cdot) - \pi\|_{\text{TV}} = \frac{1}{2} \sum_{a \in G} |P^t(g, a) - \pi(a)| = \frac{1}{2} \sum_{a \in G} |P^t(\phi(g), \phi(a)) - \pi(\phi(a))| = \|P^t(h, \cdot) - \pi\|_{\text{TV}}. \quad \square$$

Following Lemma 1.1.19 we start all our random walks on groups at the identity of the group. We introduce the function $d_{\text{TV}}(t)$ as condensed notation for the (worst) total variation distance of a Markov chain.

Definition 1.1.20. Let P be a probability which defines an irreducible random walk on G and $t \geq 0$. Define the function $d_{\text{TV}}(t)$ as follows

$$d_{\text{TV}}(t) := \max_{g \in G} \|P^t(g, \cdot) - \pi\|_{\text{TV}} = \|P^t - \pi\|_{\text{TV}}.$$

Separation Distance

We now provide a second notion of distance on probabilities as an alternative to total variation distance. Separation distance is a commonly used measure of distance between a Markov chain driven by P and its stationary distribution. Unlike total variation distance to define separation distance we require the stationary distribution to exist and be unique.

Definition 1.1.21 (Section 6.4 [32]). Let P be a transition matrix for a Markov chain with unique stationary distribution π . Define *separation distance* as follows:

$$\|P^t(g, \cdot) - \pi\|_{\text{Sep}} = \max_{h \in G} \left(1 - \frac{P^t(g, h)}{\pi(h)} \right). \quad (1.3)$$

Separation distance takes values in $[0, 1]$ and tells us the maximum ratio of the probability being in single state h against the uniform distribution.

Similarly to total variation distance for a random walk on a group G separation distance does not depend on the starting state g , that is for any $g, h \in G$, we have $\|P^t(g, \cdot) - \pi\|_{\text{Sep}} = \|P^t(h, \cdot) - \pi\|_{\text{Sep}}$.

Definition 1.1.22. Let P define by an irreducible random walk on a group G and $t \geq 0$. Define the function $d_{\text{Sep}}(t)$ as follows

$$d_{\text{Sep}}(t) := \max_{g \in G} \|P^t(g, \cdot) - \pi\|_{\text{Sep}} = \|P^t - \pi\|_{\text{Sep}}.$$

The separation distance of a random walk on a group is always decreasing in time, $d_{\text{sep}}(t+1) \leq d_{\text{sep}}(t)$, for any $t \geq 0$ [32, Section 6.4]. From our definitions it is unclear whether separation distance is statistically different from total variation distance.

Example 1.1.23. Consider the simple random walk on \mathbb{Z}_5 started at point 0. The table below gives the separation and total variation distance for $0 \leq t \leq 6$:

t	0	1	2	3	4	5	6
$d_{\text{sep}}(t)$	1	1	1	1	0.6875	0.6875	0.453125
$d_{\text{TV}}(t)$	0.8	0.6	0.4	0.35	0.275	0.225	0.18125

Table 1.1: The values of separation and total variation distance for the simple random walk on \mathbb{Z}_5 .

In Table 1.1 we can see that total variation distance is always smaller than separation distance. This is not specific to our example and holds for all Markov chains for which separation distance is well defined.

Lemma 1.1.24. *Let P define an irreducible random walk on a finite group G with stationary distribution π . Separation distance forms an upper bound on total variation distance, that is for all $t \geq 0$,*

$$d_{\text{TV}}(t) \leq d_{\text{sep}}(t). \quad (1.4)$$

Proof. Following from the definition of total variation distance we find:

$$\|P^t - \pi\|_{\text{TV}} = \sum_{\substack{h \in G \\ P(h) < \pi(h)}} \pi(h) - P^t(h) = \sum_{\substack{h \in G \\ P(h) < \pi(h)}} \pi(h) \left(1 - \frac{P^t(h)}{\pi(h)}\right) \leq \max_{h \in G} \left(1 - \frac{P^t(h)}{\pi(h)}\right) \quad \square$$

Bounds on Total Variation and Separation Distance

With our new definitions in hand we may state one of the cornerstone theorems in Markov chain theory, the convergence of irreducible aperiodic Markov chains to their stationary distributions. A proof of the following result may be found in [32, Theorem 4.9].

Theorem 1.1.25. *Suppose that P is an irreducible, aperiodic random walk on a finite group G , with stationary distribution π . Then there exists some constants $c \in [0, 1)$ and $A > 0$ such that*

$$\|P^t - \pi\|_{\text{TV}} \leq Ac^t \quad (1.5)$$

Therefore, $d_{\text{TV}}(t) \rightarrow 0$ as $t \rightarrow \infty$. Similarly $d_{\text{sep}}(t) \rightarrow 0$ as $t \rightarrow \infty$.

This theorem tells us that any irreducible aperiodic random walk necessarily converges to its stationary distribution and moreover we can bound the rate of this convergence. The time at which our random walk is close to uniform is called the *mixing time* of the Markov chain.

Definition 1.1.26 (Section 4.5 [32]). Let $\{X^t\}$ be a Markov chain with transition matrix P . Define the ε -total variation mixing time $t_{\text{TV}}(\varepsilon)$ of our Markov chain as:

$$t_{\text{TV}}(\varepsilon) := \min\{t : d_{\text{TV}}(t) \leq \varepsilon\} \quad (1.6)$$

Now define the *total variation mixing time* t_{TV} of a Markov chain to be $t_{\text{TV}} := t_{\text{TV}}(1/4)$ (the choice of $1/4$ is semi-arbitrary, see [32, Section 4.5]). We may define the *separation mixing time* $t_{\text{sep}}(\varepsilon)$ and t_{sep} , in an analogous way replacing d_{TV} with d_{sep} .

From now on all random walks on groups we study will be irreducible, and aperiodic, meaning that they always have convergence to their uniform distribution. It is then of interest to ask what this convergence looks like for a given random walk. Theorem 1.1.25 allows us to bound the mixing time of a random walk from above. Without any more assumptions about our Markov chain there is not much more we can say about the rate of the convergence. The techniques and tools used to bound mixing times have long been developed by probabilists. However, it is often the case that ad hoc techniques need to be developed specific to the random walk being studied. We introduce the key probabilistic and algebraic tools in this chapter. In Chapter 2 we shall see how we have to use a combination of tools to prove precise bounds on the mixing time for the random transposition shuffle and the top-to-random shuffle. To begin with we state a classical upper bound on total variation distance for reversible random walks.

Theorem 1.1.27 (Classical ℓ^2 bound - see Lemma 12.16 [32]). *Let P be the transition matrix for a reversible, transitive, irreducible, aperiodic, Markov chain on a finite group G , with stationary distribution π . Then we may label the eigenvalues so that $1 = \beta_1 > \beta_2 \geq \dots \geq \beta_{|G|} > -1$. Furthermore, we have the following bound,*

$$4\|P^t - \pi\|_{\text{TV}}^2 \leq \sum_{i \neq 1} \beta_i^{2t}. \quad (1.7)$$

For irreducible random walks on groups the detailed balanced equations (1.1) reduce to $P(g, h) = P(h, g)$, therefore a random walk on a group is reversible if and only if $P(g^{-1}) = P(g)$. We make use of Theorem 1.1.27 in the analysis of the random-to-random shuffle in Chapter 2, and the one-sided transposition shuffle in Chapter 3. Another popular method to establish an upper bound on separation distance (and therefore total variation distance) is the use of strong stationary times. To define a strong stationary time, we first need to define the notion of a stopping time.

Definition 1.1.28. Let τ be a random variable taking values in \mathbb{N}^0 . We call τ a *stopping time* for a Markov chain $\{X^t\}$, if we can decide the event $\{\tau = t\}$ with the knowledge of our Markov chain up to time t , that is the states $\{X^0, \dots, X^t\}$.

Example 1.1.29. Let $G = S_5$ be the symmetric group on 5 elements, and define a random walk on G given by the uniform distribution π . We may view S_5 as the permutations of a deck of 5 cards. We start at the identity permutation, and at every step of our random walk choose a permutation to apply uniformly at random. Let τ be the first time card 5 is moved. Then τ is a stopping time for our random walk because if we observe our walk up to time t , we can tell whether card 5 has been moved and if so when was the first time this

happened in our random walk. Suppose instead that τ is the last time card 5 is moved up until a fixed future time T . To even decide the event $\{\tau = 1\}$ we would need to know the entire future of our Markov chain up to time T , and so τ is no longer a stopping time.

Strong stationary times are a special kind of stopping times. We impose the extra conditions that once a strong stationary time is met the distribution of X^t must be stationary and independent of τ .

Definition 1.1.30. Let τ be a stopping time for an irreducible, aperiodic Markov chain $\{X^t\}$ which is a random walk on a group. We say τ is a *strong stationary time* if

$$\mathbb{P}(X^t = g | \tau = t) = \pi(g). \quad (1.8)$$

From the definition we swiftly prove how strong stationary times may be used to bound separation distance and thus total variation distance for random walks.

Lemma 1.1.31. *Let τ be a strong stationary time for an irreducible, aperiodic random walk on G with driving probability P . The following holds for all $t \geq 1$:*

$$\|P^t - \pi\|_{\text{TV}} \leq \|P^t - \pi\|_{\text{Sep}} \leq \mathbb{P}(\tau > t).$$

Proof. The first inequality comes directly from Lemma 1.1.24. For the second inequality we start at the definition of separation distance,

$$1 - \frac{P^t(g)}{\pi(g)} \leq 1 - \frac{\mathbb{P}(X^t = g, \tau \leq t)}{\pi(g)} = 1 - \frac{\pi(g)\mathbb{P}(\tau \leq t)}{\pi(g)} = \mathbb{P}(\tau > t) \quad \square.$$

Strong stationary times are useful because they reduce the analysis of separation distance based on the probability P^t , to the analysis of a single random variable τ . In practice strong stationary times can be tricky to find but we will see two examples of them in this thesis: a classical argument for the top-to-random shuffle in Chapter 2, and an original argument for the one-sided transposition shuffle in Chapter 3.

All the methods we have mentioned so far give upper bounds on the mixing time of a random walk. When analysing the rate of convergence of a Markov chain it is also useful to bound the mixing time from below. To formulate a lower bound we frequently make use of the following simple lemma.

Lemma 1.1.32. *Let P be a random walk on a finite group G with stationary distribution π . Suppose $A \subseteq G$ then*

$$|P^t(A) - \pi(A)| \leq \|P^t - \pi\|_{\text{TV}}.$$

Proof. The inequality follows from Definition 1.1.15. □

This bound may seem naïve but throughout this thesis we shall see that it is a very adaptable tool. The key idea behind using Lemma 1.1.32 for a lower bound on the mixing time of a random walk P is to find a set A which tells apart P^t from π . Usually we look for A that has small probability under the stationary distribution, and that has high probability under the probability P^t for time $t < t_{\text{TV}}$. The set A is often chosen based on information about our group G . For random walks on the symmetric group a commonly

used idea is to take A as a set of permutations with certain fixed points with the intention of reducing the calculation of $P^t(A)$ to an estimation of how long it takes for each fixed point to be eliminated. This allows us to use a modified coupon collectors argument to estimate our probability. We use this idea to formulate lower bounds for both the random transposition shuffle and the one-sided transposition shuffle.

1.1.3 The Cutoff Phenomenon

In practice, for many Markov chains, the total variation distance to uniform does not decrease steadily as suggested by Theorem 1.1.25. Rather, a phenomenon is often seen where the distance decreases sharply past a critical time as show in Figure 1.1: The blue line shows an exponential convergence from the start, whereas the red line only shows convergence to 0 after a certain time threshold has been met.

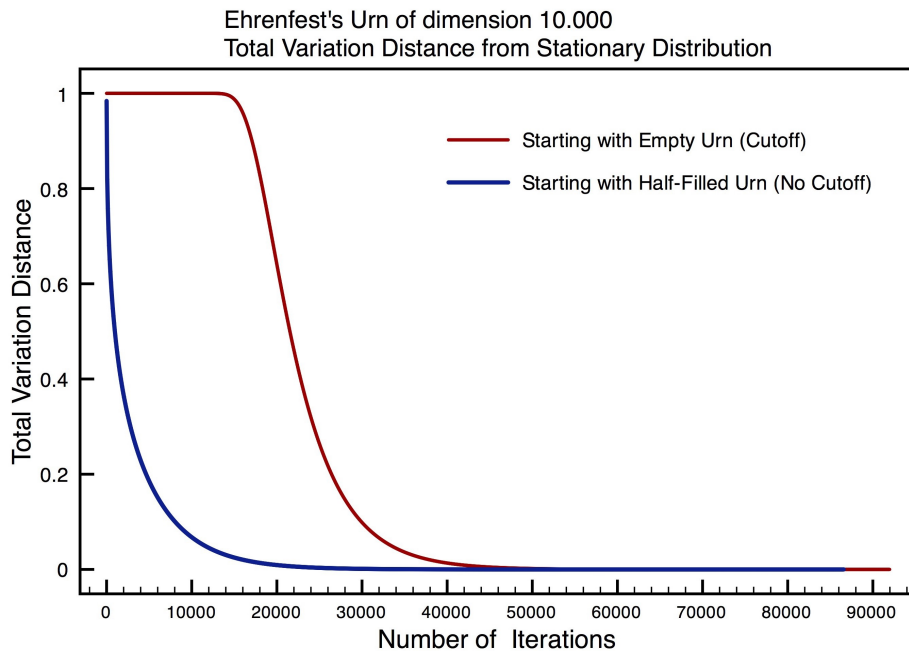


Figure 1.1: Ehrenfest’s urn is a random process starting with n balls distributed between two urns, and at each step a ball is chosen uniformly at random and moved to the opposite urn. The lines shows the total variation distance of the process from its uniform distribution from different starting states. The blue process starts with both urns containing $1/2$ the total amount of balls, the red process starts with one urn full and the other empty. The figure is taken from [31].

Furthermore, there exist many families of Markov chains $\{X_n^t\}$ with transition probabilities $\{P_n\}$ that have the transition from 1 to 0 becoming sharper as $n \rightarrow \infty$, Thus when such a phase transition is present we may think about the mixing time of a random walk being the time of this phase transition. We now switch our view to now consider a family of Markov chains $\{X_n^t\}$ with transition matrices $\{P_n\}$ and stationary distributions $\{\pi_n\}$ where $n \in \mathbb{N}$ indexes the family. The phase transition which has been observed in many families of random walks has been given the name *the cutoff phenomenon*.

Definition 1.1.33 (Section 18.1 [32]). A sequence of random walks $\{X_n^t\}$ on finite groups $\{G_n\}$ with driving probabilities $\{P_n\}$ and stationary distributions $\{\pi_n\}$ exhibits a (total

variation/separation) *cutoff* at time t_n with a *window* of size w_n if $w_n = o(t_n)$ and the following limits hold:

$$\lim_{c \rightarrow \infty} \limsup_{n \rightarrow \infty} \|P_n^{t_n + cw_n} - \pi_n\|_{\text{TV/Sep}} = 0 \quad (1.9)$$

$$\lim_{c \rightarrow \infty} \liminf_{n \rightarrow \infty} \|P_n^{t_n - cw_n} - \pi_n\|_{\text{TV/Sep}} = 1. \quad (1.10)$$

The limits (1.10) and (1.9), respectively define upper and lower bounds on the mixing time of the random walk .

The heuristic behind this definition is the phase transition for the family of chains happens at time t_n and this transition becomes sharper as $n \rightarrow \infty$. If a family shows a cutoff at time t_n then we must have that $t_{\text{TV}} \sim t_n$. An equivalent definition (see [32, Lemma 18.1]) of a cutoff is the function $d_{\text{TV}}(\cdot)$ converges to a step function around the mixing time of the random walk. Precisely, a family of random walks P_n exhibits a total variation cutoff if and only if the following limit holds:

$$\lim_{n \rightarrow \infty} d_{\text{TV}}(c t_{\text{TV}}) = \begin{cases} 1 & \text{if } c > 1 \\ 0 & \text{if } c < 1 \end{cases}.$$

Once we have established a cutoff for a family of Markov chains we have complete asymptotic information about the families convergence to their stationary distributions. Definition 1.1.33 is given in terms of both total variation and separation distance. We may hope that a cutoff in one would imply a cutoff in the other, however this is not the case [24]. This does not mean that a cutoff in total variation cannot help us prove a cutoff in separation or vice versa. By Lemma 1.1.24 if we establish an upper bound on separation distance we get an upper bound on total variation distance, similarly a lower bound on total variation distance gives a lower bound on separation distance. For the one-sided transposition shuffle we first prove a cutoff in total variation distance then a cutoff in separation distance with the lower bound following from our first proof.

To establish a cutoff it is not enough to just find the order of t_n , although this is a useful first step. Computing the exact constant factor of the correct order for a cutoff is often where the difficulty in the analysis of mixing times stems from. Frequently finding the correct cutoff time for a random walk involves an initial ansatz, and then refinement to make sure both the upper and lower bounds hold. We prove the existence of a cutoff for several random walks on the symmetric group in Chapter 2, for the moment we state some simple cutoff results.

Example 1.1.34. Consider the family of lazy simple random walks on \mathbb{Z}_n , which are driven by probabilities

$$P_n(i) = \begin{cases} \frac{1}{2} & \text{if } i = 0 \\ \frac{1}{4} & \text{if } i = \pm 1 \\ 0 & \text{otherwise} \end{cases}.$$

This random walk is irreducible, aperiodic and transitive. We may ask whether $d_{\text{TV}}(t)$ shows a cutoff as $n \rightarrow \infty$. The answer is no, this family of random walks fails to show a cutoff [39, Section 3.3]. The lazy simple random walk on a circle steadily spreads out from

the starting point 0 meaning there is no phase transition and thus no cutoff.

Example 1.1.35. Consider the hypercube \mathbb{Z}_2^n , define the identity of the group as e_0 , the zero vector. For $i \in \{1, \dots, n\}$, define e_i as the element which has 1 in position i , and 0 in every other position. The nearest neighbour walk on the hypercube is driven by probability P_n defined by

$$P_n(g) = \begin{cases} \frac{1}{n+1} & \text{if } g = e_i \text{ for } i \in \{0, 1, \dots, n\} \\ 0 & \text{otherwise} \end{cases}.$$

The probabilities $\{P_n\}$ define a family of irreducible, aperiodic, and transitive random walks, which show a cutoff in total variation distance at time $t_n = \frac{n+1}{4} \log n$, with a window of $w_n = \frac{n+1}{4}$. Therefore we may say the mixing time of this random walk is exactly $t_{\text{TV}} = \frac{n+1}{4} \log n$. Details of this random walk may be found at [16, Chapter 3C].

Since its formulation the cutoff phenomenon has been a great area of interest in the study of random walks. Establishing a cutoff for any particular example requires detailed knowledge and understanding of the inner workings of the Markov chain. The upper bound (1.9) and the lower bound (1.10) have to be proved separately and require the use of a wide variety of techniques. We would instead prefer to find a simple criterion for a random walk exhibiting a cutoff; however, no such result currently exists and in fact there are very few global results on total variation cutoff of random walks. Contrary to this much is known about cutoffs in ℓ_p -distance with $1 < p < \infty$ [14], which further highlights why the lack of results for $p = 1$ (total variation) is so surprising. Why any family of Markov chains shows a total variation cutoff is a mystery with many explanations proposed. One of the most popular theories is that the existence of a cutoff is linked to the multiplicity and value of the second biggest eigenvalue.

Conjecture 1.1.36 (Peres's Conjecture [15]). *Let $\{P_n\}$ be a family of transition matrices Markov chains, with second biggest (in absolute value) eigenvalue β_n . The family shows a total variation cutoff at time t_n if and only if*

$$t_n(1 - \beta_n) \rightarrow \infty \text{ as } n \rightarrow \infty. \quad (1.11)$$

If we measure cutoff instead in ℓ_p -distance for $1 < p < \infty$, Chen and Saloff-Coste [14] have managed to prove a criterion for a cutoff in the spirit of Peres' conjecture. If a family of Markov chains shows a total variation cutoff, condition (1.11) is known to be satisfied. On the other hand Aldous managed to find a counterexample for Peres' conjecture using random walks on graphs, see [13, Section 4.2] for details. However, if we specialise to random walks on finite groups then there is currently no known counterexample to Peres's conjecture. Random walks on finite groups are an active area of research, the algebraic setting allowing for specialised techniques in order to prove cutoff results. In the next section we cover the algebraic tools available to find the mixing time of a random walk on G , these require background about the representations, modules and characters of groups.

1.2 Representations, Modules, and Applications to Random Walks

Group representation theory aims to study groups using linear algebra. The goal is to reduce the abstract elements of groups into the well understood elements of linear algebra. To do this we represent the elements of our group as matrices in such a way that the structure of the group is respected. For the applications in this thesis it is enough to develop the representation theory of groups over the complex numbers, so all vector spaces, matrices, etc., are over the field \mathbb{C} unless otherwise stated. This assumption helps to simplify the representation theory presented and allows the use of several theorems which do not hold over general fields. A detailed account of all the representation theory presented in this section can be found in *Representations and Characters of Groups* by James and Liebeck [27]. We begin this section by introducing the representations and characters of a finite group G . We focus on the irreducible representations of G from which any representation may be built. The following section is devoted to the properties of the discrete Fourier transform, which is used to analyse probabilities on G using the group's irreducible characters. Our last section covers the upper bound lemma, which establishes a close link between random walks on G and the representation theory of G .

1.2.1 Representations, Characters and Modules

Representations and Modules

Definition 1.2.1. Let G be a finite group. A *representation* of G is a vector space V together with a *group homomorphism* from G to the general linear group $GL(V)$, denoted $\rho: G \rightarrow GL(V)$, which has the property $\rho(gh) = \rho(g)\rho(h)$ for all $g, h \in G$. We say the pair (ρ, V) is a representation. The dimension of V is denoted d_ρ and is called the *dimension* of the representation (ρ, V) .

The vector space V our matrices $\rho(\cdot)$ act on is critical to our knowledge of the representation. Thus, we write (ρ, V) when we want to emphasise the space V . If the space V is clear from context we will just talk about the group homomorphism ρ being a representation by itself. We now introduce the notion of G modules, by equipping our group with a vector space structure.

Definition 1.2.2. Let G be a finite group. The *group algebra* of G is the vector space $\mathbb{C}[G] = \langle g \mid g \in G \rangle$, spanned by elements of G with coefficients in \mathbb{C} . This vector space has a natural multiplication given by multiplying complex coefficients and the elements of G separately, e.g. for $g, h, f \in G$ we have $(4g)(ih - \frac{1}{2}f) = 4i(gh) - 2(gf)$.

Definition 1.2.3. Let G be a finite group. A $\mathbb{C}[G]$ -*module* is a vector space V with a bilinear multiplication $\mu: \mathbb{C}[G] \times V \rightarrow V$ such that: $\mu(1, v) = v$ for all $v \in V$, and for all $g, h \in G$, $\mu(h, \mu(g, v)) = \mu(hg, v)$. To simplify notation we frequently write the multiplication $\mu(g, v)$ as just gv , and say $\mathbb{C}[G]$ acts on V .

We can think of a $\mathbb{C}[G]$ -module as a vector space V which our group G acts on. For any group G , and vector space V , we may always define a trivial action by setting, $gv = v$ for all $g \in G, v \in V$. The *trivial module* of a group G is formed from a one-dimensional

vector space with the trivial action. The *trivial representation* of a group G is a one-dimensional vector space with group homomorphism $\rho(g) = \text{Id}$ for all $g \in G$. The trivial module and trivial representation are two different ways to view the same algebraic object. The next result established that the representations of G (over field \mathbb{C}) and $\mathbb{C}[G]$ -modules are in one-to-one correspondence.

Theorem 1.2.4 (Theorem 4.4 [27]). *Let G be a finite group. The representations of G and $\mathbb{C}[G]$ -modules are in one-to-one correspondence. If (ρ, V) is a representation for G then V is a $\mathbb{C}[G]$ -module with multiplication defined by $\mu(g, v) = \rho(g)v$ for all $g \in G$. On the other hand if V is a $\mathbb{C}[G]$ -module with multiplication μ , define $\rho(g) = \mu(g, \cdot)$, then (ρ, V) is a representation for G .*

Proof. The constructions in the theorem may be verified by checking the definitions of modules and representations. To prove this is a one-to-one correspondence take a representation (ρ, V) and form the module V with action μ via the process described in the theorem. Now we turn the module V into a representation (φ, V) with a new mapping $\varphi(g) = \mu(g, \cdot) = \rho(g)$, therefore we recover our original representation (ρ, V) . We may perform a similar check starting with a module V and transforming it to a representation and back to a module recovering V . Thus, representations and modules are in one-to-one correspondence. \square

From now on out we may talk about representations and modules interchangeably, and every result we state is valid for either view point via the correspondence described in Theorem 1.2.4. We now give a brief example of some natural representations.

Example 1.2.5. Let S_3 denote the symmetric group on 3 elements. Consider the space \mathbb{C}^3 , we may form an action of S_3 on \mathbb{C}^3 , by permuting the basis elements. Realising this action in terms of matrices gives us a group homomorphism ρ detailed in Table 1.2, therefore the pair (ρ, \mathbb{C}^3) defines a representation of S_3 . For another module of S_3 take the 2-dimensional subspace $W = \{v = (v_1, v_2, v_3)^T \in \mathbb{C}^3 \mid v_1 + v_2 + v_3 = 0\}$, spanned by vectors $w_1 = (1, -1, 0)^T$, and $w_2 = (0, 1, -1)^T$. Again consider an action of S_3 on W by permuting the coordinates of \mathbb{C}^3 , e.g., $(12)w_1 = -w_1$ and $(12)w_2 = w_1 + w_2$. Realising this action as matrices on $W = \langle w_1, w_2 \rangle$ we find the group homomorphism φ shown in Table 1.2, hence (φ, W) is a representation of S_3 .

S_3	e	(12)	(13)	(23)	(123)	(132)
$\rho(\cdot)$	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}$
$\varphi(\cdot)$	$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$	$\begin{pmatrix} -1 & 0 \\ 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 1 \\ 0 & -1 \end{pmatrix}$	$\begin{pmatrix} 0 & 1 \\ -1 & -1 \end{pmatrix}$	$\begin{pmatrix} -1 & -1 \\ 1 & 0 \end{pmatrix}$

Table 1.2: The table shows the group homomorphisms ρ, φ on the vectors spaces \mathbb{C}^3 and W respectively. Note that the elements $\rho(\cdot)$ are the 3-dimensional permutation matrices.

Definition 1.2.6. Let V be a $\mathbb{C}[G]$ -module. We call a subspace $W \subseteq V$ *G-stable* if $\{gw \mid g \in \mathbb{C}[G], w \in W\} = W$. A *submodule* of V is a stable subspace $W \subseteq V$. A module

V is called *simple* if it has no non-trivial ($W \neq V, \{0\}$) submodules. Equivalently a representation (ρ, V) is called *irreducible* if it has no non-trivial subrepresentations.

Take any one-dimensional vector space V , we know that $\{0\}$ and V are the only subspaces of V . Therefore, any one dimension module is necessarily simple. The simple modules of a group G are the building blocks of its representation theory, in the sense that every module may be decomposed into its simple submodules, see Theorem 1.2.14. We now introducing mappings between modules which preserve their structure, these are called homomorphisms.

Definition 1.2.7. Let V, W be $\mathbb{C}[G]$ -modules. A *homomorphism of modules* is a linear map $\psi : V \rightarrow W$ such that $\psi(gv) = g\psi(v)$. An *isomorphism of modules* is a bijective homomorphism $\psi : V \rightarrow W$, in this case we say V is isomorphic to W as $\mathbb{C}[G]$ -modules, denoted $V \cong W$. Define the *kernel* and *image* of a homomorphism $\psi : V \rightarrow W$ as follows: $\text{Ker}(\psi) = \{v : \psi(v) = 0\}$, $\text{Im}(\psi) = \{w : \psi(v) = w \text{ for some } v\}$.

Lemma 1.2.8. Let V, W be $\mathbb{C}[G]$ -modules, and $\psi : V \rightarrow W$ a homomorphism. Then $\text{Ker}(\psi)$ is a submodule of V , and $\text{Im}(\psi)$ is a submodule of W . If $\text{Ker}(\psi) = \{0\}$ then ψ is an injective map, and if $\text{Im}(\psi) = W$ then ψ is a surjective map.

Proof. Take $v \in \text{Ker}(\psi)$. For any element $g \in \mathbb{C}[G]$ we have $gv \in \text{Ker}(\psi)$ because $\psi(gv) = g\psi(v) = 0$. Hence, the kernel of ψ is stable under the action of $\mathbb{C}[G]$ and so is a $\mathbb{C}[G]$ -module contained in V . Similarly the image of ψ is stable under the action of $\mathbb{C}[G]$ and so is a $\mathbb{C}[G]$ -module contained in W . The last statements are standard facts of linear maps. \square

We are now in a position to state Schur's Lemma which tells us that homomorphisms between simple modules are trivial.

Lemma 1.2.9 (Schur's Lemma). Let V, W be simple $\mathbb{C}[G]$ -modules. If $\psi : V \rightarrow W$ is a homomorphism then either ψ is the zero morphism or an isomorphism of representations. Furthermore, the only homomorphisms $\psi : V \rightarrow V$ are scalar multiples of the identity map, i.e., $\psi(v) = \beta v$ for some $\beta \in \mathbb{C}$.

Proof. By V being a simple module we have two choices for the kernel of ψ : V or $\{0\}$. If $\text{Ker}(\psi) = V$ then ψ is clearly the zero morphism. If $\text{Ker}(\psi) = \{0\}$ then our map is injective, moreover $\text{Im}(\psi) \neq \{0\}$ and hence must have $\text{Im}(\psi) = W$ because W is a simple. Therefore the map ψ is a bijective homomorphism.

Now assume $\psi : V \rightarrow V$ is a homomorphism. Since we are working over the field \mathbb{C} the map ψ must have an eigenvalue β with eigenvector v . Form a new homomorphism $\psi' = \psi - \beta \text{Id}$, with Id being the identity homomorphism. We can see that $\psi'(v) = 0$ thus $\text{Ker}(\psi') = V$, and so our new map ψ' must be the zero morphism. Hence, we have $\psi = \beta \text{Id}$. \square

Isomorphism forms an equivalence relation on the set of $\mathbb{C}[G]$ -modules. For any group G we would like to form a complete collection of its simple modules (equivalently irreducible representations) up to isomorphism.

Lemma 1.2.10. *Let G be a finite group. Then there are only finitely many simple $\mathbb{C}[G]$ -modules up to isomorphism (see Definition 1.2.18).*

Definition 1.2.11. Let G be a finite group. Define the *complete collection of simple $\mathbb{C}[G]$ -modules* as a set, denoted $\text{Irr}(G)$, such that if V is an simple $\mathbb{C}[G]$ -module then $V \cong M$ for exactly one $M \in \text{Irr}(G)$.

We now look to decompose any reducible module into a direct sum of smaller submodules.

Definition 1.2.12. Let V_1, V_2 be $\mathbb{C}[G]$ -modules. The *direct sum* of vector spaces $V_1 \oplus V_2$ is a $\mathbb{C}[G]$ -module under the action $g(v_1, v_2) = (g v_1, g v_2)$ for $v_1 \in V_1, v_2 \in V_2$ and $g \in G$. Moreover, we have $\dim(V) = \dim(V_1) + \dim(V_2)$. Conversely, suppose W is a $\mathbb{C}[G]$ -module, with submodules W_1, W_2 , such that $W_1 \oplus W_2 = W$, and $W_1 \cap W_2 = \{0\}$. Then our module W may be decomposed into a direct sum $W = W_1 \oplus W_2$.

We now prove that every stable subspace of a module V has a stable complement. Thus, any non-simple module may be split into a direct sum of smaller modules. This allows us to prove the powerful Maschke's Theorem.

Lemma 1.2.13 (See Theorem 8.1 [27]). *Let V be a $\mathbb{C}[G]$ -module. Suppose V contains a stable subspace W . Then V contains a second stable subspace W^0 such that $W \cap W^0 = \{0\}$ and $W \oplus W^0 = V$. The vector space W^0 is called the complement of W .*

Proof. Suppose v_1, \dots, v_n form a basis of V with v_1, \dots, v_l simultaneously being a basis for W . Let $\phi : V \rightarrow V$ be the projection onto the subspace W , defined for basis elements $\phi(v_i) = v_i$ if $1 \leq i \leq l$ and 0 otherwise. Now define a new linear map $\psi : V \rightarrow V$ as follows:

$$\psi(v) = \frac{1}{|G|} \sum_{g \in G} g^{-1} \phi(gv).$$

The map ψ is a homomorphism of G -modules. Indeed for any $h \in G$,

$$h\psi(v) = \frac{1}{|G|} \sum_{g \in G} hg^{-1} \phi(gv) = \frac{1}{|G|} \sum_{g \in G} g^{-1} \phi(ghv) = \psi(hv).$$

The image of ψ must be contained in the submodule W because of the presence of the projection ϕ . In fact it is equal to W , taking $w \in W$ we have,

$$\psi(w) = \frac{1}{|G|} \sum_{g \in G} g^{-1} \phi(gw) = \frac{1}{|G|} \sum_{g \in G} g^{-1} gw = w.$$

Set $W^0 = \text{Ker}(\psi)$, by Lemma 1.2.8 this forms a stable submodule of V . To summarise, ψ is a module homomorphism, with $\text{Im}(\psi) = W$ and $\text{Ker}(\psi) = W^0$ therefore we must have $W \oplus W^0 = V$, and $W \cap W^0 = \{0\}$. \square

Theorem 1.2.14 (Maschke's Theorem). *Suppose V is a module of a finite group G . Then we may decompose V into a direct sum of simple modules. This means the module V may be decomposed as follows,*

$$V \cong \bigoplus_{M_i \in \text{Irr}(G)} k_i M_i \tag{1.12}$$

where the direct sum is over all simple modules $M_i \in \text{Irr}(G)$ and k_i denotes the multiplicity of M_i in the direct sum.

Proof. We proceed by induction on dimension. The statement is clearly true for one-dimensional modules. If $\dim(V) > 1$ then either V is simple, in which case we are done, or V has a proper submodule W . By 1.2.13 we may write $V = W \oplus W^0$, for modules W, W^0 with $\dim(W), \dim(W^0) < \dim(V)$. Hence, W and W^0 can be written as direct sums of simple submodules, which implies the same for V . \square

Theorem 1.2.14 is essential to the study of modules (and representations) because it tells us that to understand the modules of G it is enough to understand all the simple modules. However, to completely describe any module V we also need to know the value of the constants k_i , before finding these multiplicities we should be sure that the decomposition (1.12) is unique up to isomorphism. This follows from an application of Schur's Lemma.

Lemma 1.2.15. *Let G be a finite group. Then for constants $k_i, l_i \in \mathbb{N}^0$ we have:*

$$\bigoplus_{M_i \in \text{Irr}(G)} k_i M_i \cong \bigoplus_{M_i \in \text{Irr}(G)} l_i M_i \Leftrightarrow k_i = l_i \text{ for all } i.$$

Proof. If $k_i = l_i$ for all i then the conclusion is immediate. Suppose that our direct sums are isomorphic with isomorphism ψ . If we restrict to a summand M_i on the left hand side then we must get an isomorphic copy of M_i on the right hand side, by Schur's lemma. Hence, the summand $k_i M_i$ must land in the summand $l_i M_i$ on the right. Counting the dimensions on each side now gives $l_i = k_i$ for all i . \square

Corollary 1.2.16. *The decomposition shown in Theorem 1.2.14 is unique up to isomorphism.*

Definition 1.2.17. Let G be a finite group, M_i a simple module, and V a module with decomposition $V \cong \bigoplus k_i M_i$. The *isotypic component* of M_i in V is the unique submodule $k_i M_i$.

Corollary 1.2.16 tells us we can decompose any module into its simple parts but it does not tell us the values of the constants k_i . Finding efficient ways to know or compute the constants in a decomposition is an important problem in representation theory. To end this section we introduce the group algebra as a $\mathbb{C}[G]$ -module itself, called the *regular module* of G .

Definition 1.2.18 (Definition 6.5 [27]). The *regular module* $\mathbb{C}[G]$ has the same vector space and action given by the multiplication of the group algebra and it has dimension $|G|$.

Lemma 1.2.19. *Let G be a finite group. The regular module $\mathbb{C}[G]$ has canonical decomposition*

$$\mathbb{C}[G] \cong \bigoplus_{M_i \in \text{Irr}(G)} d_i M_i \tag{1.13}$$

where d_i is the dimension of the simple module M_i . We prove this decomposition in Lemma 1.2.27, after the introduction of characters.

Knowledge of the simple modules of the symmetric group S_n will be key to our analysis of the one-sided transposition shuffle in Chapter 3. In particular we study the action of the shuffle P_n on the regular module of S_n . This allows us to use the decomposition in Definition 1.2.18 to reduce our goal of finding eigenvalues of the regular module to finding eigenvalues of the action on the simple modules of S_n . Before we perform our analysis we shall present a detailed construction of the simple modules of the symmetric group in Chapter 2.

Character Theory

In the beginning of this section we assumed the field we are working over to be \mathbb{C} . This assumption lets us condense the information of any representation (ρ, V) into a single function $\chi_\rho : G \rightarrow \mathbb{C}$, called the *character* of our representation. The character gives a simple way to view the information given by any representation.

Definition 1.2.20. Let (ρ, V) be a representation of a group G . Define the *character* $\chi_\rho : G \rightarrow \mathbb{C}$ of a representation ρ by $\chi_\rho(g) = \text{Tr} \rho(g)$, where $\text{Tr}(\cdot)$ denotes the trace of a matrix. We say the *dimension* of the character, denoted d_χ , is the same as its corresponding representation. If (ρ, V) is an irreducible representation then χ_ρ is called an *irreducible character*.

Lemma 1.2.21. Let χ be a character for the group G . We find:

1. $\chi(e) = d_\chi$
2. For any $g \in G$ we have $\chi(g^{-1}) = \overline{\chi(g)}$, where \bar{c} denotes the complex conjugate of $c \in \mathbb{C}$
3. For any $g, h \in G$, we have $\chi(hgh^{-1}) = \chi(g)$.

Proof. 1. We clearly have $\chi(e) = \text{Tr}(\text{Id}) = d_\chi$

2. We know that $\rho(g)^n = \rho(g^n) = \text{Id}$ for some n , so the eigenvalues of $\rho(g)$ must be roots of unity, denoted ξ_i . Then we have

$$\overline{\chi(g)} = \overline{(\text{Tr} \rho(g))} = \sum_i \overline{\xi_i} = \sum_i \xi_i^{-1} = \text{Tr}(\rho(g)^{-1}) = \text{Tr} \rho(g^{-1}) = \chi(g^{-1})$$

3. Follows from the property $\text{Tr}(AB) = \text{Tr}(BA)$ of the trace .

□

The last of these properties tells us that characters are constant on the conjugacy classes of any group G . Representations and characters have a close connection to the conjugacy classes of the group they are defined over. In fact the number of irreducible representations of G is exactly the number of conjugacy classes of G . To establish this we introduce an inner product on the space of class functions of G .

Definition 1.2.22. Let G be a finite group. A *class function* for G is a function $\phi : G \rightarrow \mathbb{C}$ that is constant on the conjugacy classes of G , i.e., for all $g, h \in G$ we have $\phi(g) = \phi(h^{-1}gh)$. Let χ, ϕ be class functions on G . Define the inner product $\langle \chi | \phi \rangle$ as follows:

$$\langle \chi | \phi \rangle = \frac{1}{|G|} \sum_{g \in G} \chi(g) \overline{\phi(g)}. \quad (1.14)$$

Theorem 1.2.23 (See Chapter 15 [27]). *Let χ_i denote the character of the irreducible representation $\rho_i \in \text{Irr}(G)$. The irreducible characters are orthonormal, that is:*

$$\langle \chi_i | \chi_j \rangle = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}.$$

Furthermore, the irreducible characters form a basis for the class functions of G .

Corollary 1.2.24. *Let G be a finite group. The number of irreducible representations of G is exactly the number of conjugacy classes of G .*

Proof. Let n be the number of conjugacy classes of G . The space of class functions on G is spanned by exactly n functions which take value 1 on a single conjugacy class and zero otherwise. Therefore, our basis of irreducible characters must be formed from n characters. \square

Using the inner product on class functions and the irreducible characters χ_i we now give a second proof of the uniqueness of decomposition (1.12).

Lemma 1.2.25. *Let $(\rho, V), (\varphi, W)$ be representations of a finite group G . Then we have $\chi_{\rho \oplus \varphi} = \chi_\rho + \chi_\varphi$.*

Proof. We may choose a basis of the space $V \oplus W$ such that the group homomorphism has form $(\rho \oplus \varphi)(g) = \begin{pmatrix} \rho(g) & 0 \\ 0 & \varphi(g) \end{pmatrix}$, then take traces. \square

Corollary 1.2.26. *Let (ρ, V) be a representation of a finite group G , with decomposition $\oplus_i k_i \rho_i$ into irreducible representations. Then $\langle \chi_\rho | \chi_i \rangle = k_i$.*

Proof. We split the character χ_ρ into a sum of its irreducible characters via Lemma 1.2.25. Taking the inner product, all non χ_i terms disappear by the orthogonality of irreducible characters. Note that the character χ_ρ does not depend on our decomposition, thus the constants k_i must be unique. \square

Finally we may prove the decomposition of the regular module stated in Definition 1.2.18. As a consequence we find that $\sum_{\rho \in \text{Irr}(G)} d_\rho^2 = |G|$.

Lemma 1.2.27. *Let $(\rho, \mathbb{C}[G])$ be the regular representation for a finite group G . Let χ_i be an irreducible character with dimension d_i . Then $\langle \chi_\rho | \chi_i \rangle = d_i$, and $|G| = \chi_\rho(e) = \sum_i d_i \chi_i(e) = \sum_{\rho \in \text{Irr}(G)} d_\rho^2$.*

Proof. By definition of the regular module for every $g, h \in G$ we have $\rho(g)(h) = gh$, therefore, $\rho(g)(h) \neq h$ unless $g = e$. This implies that if $g \neq e$ then all diagonals of $\rho(g)$

must be 0, and thus $\chi_\rho(g) = 0$. Computing the inner product gives us,

$$\langle \chi_\rho | \chi_i \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_i(g) \overline{\chi_\rho(g)} = \frac{1}{|G|} \chi_i(e) \chi_\rho(e) = d_i \quad \square.$$

1.2.2 Discrete Fourier Transforms

The discrete Fourier transform links a \mathbb{C} -valued function P on a group G and a representation ρ of G into a single algebraic object. Understanding Fourier transforms and how they interact with convolutions is a key step in linking the worlds of probability and algebra. Our goal in this section is to demonstrate the key properties of the transform, which are used in the next section to bound total variation distance of a random walk. The results below and their proofs come from [16, Chapter 2]

Definition 1.2.28. Let (ρ, V) be a representation of G , and P a function on G . The *Fourier transform* of P at ρ is:

$$\hat{P}(\rho) = \sum_{g \in G} P(g) \rho(g). \quad (1.15)$$

Note that $\hat{P}(\rho)$ may be viewed as a mapping from $V \rightarrow V$.

It may seem opaque at the moment why we want to consider taking Fourier transforms of functions P . It turns out that Fourier transforms have properties which help simplify the analysis of the convolution P^t . One of the most immediate results is that the discrete Fourier transform allows us to split up convolutions.

Lemma 1.2.29. Let P, Q be functions on a finite group G , and ρ a representation. Then

$$\widehat{P \star Q}(\rho) = \hat{P}(\rho) \hat{Q}(\rho). \quad (1.16)$$

Proof. Proceeding from left to right we have:

$$\begin{aligned} \widehat{P \star Q}(\rho) &= \sum_{g \in G} (P \star Q)(g) \rho(g) \\ &= \sum_{g \in G} \sum_{h \in G} P(gh^{-1}) Q(h) \rho(g) \\ &= \sum_{h \in G} P(gh^{-1}) \rho(gh^{-1}) \sum_{g \in G} Q(h) \rho(h) = \hat{P}(\rho) \hat{Q}(\rho). \quad \square \end{aligned}$$

This simple lemma is key to the usefulness of the discrete Fourier transform, applying it repeatedly we can see that $\widehat{P^t} = \hat{P}^t$. The second property of the Fourier transform which we exploit is the Fourier Inversion Theorem, which allows us to recover P from the information provided by $\hat{P}(\rho)$ on all the irreducible representations of G .

Lemma 1.2.30 (Fourier Inversion Theorem - See Chapter 2C [16]). Let P be a function on a finite group G . Then

$$P(g) = \frac{1}{|G|} \sum_{\rho \in \text{Irr}(G)} d_\rho \text{Tr}(\rho(g^{-1}) \hat{P}(\rho)). \quad (1.17)$$

The Fourier inversion Theorem is important because it allows us to go back and forth between the \mathbb{C} -valued function P and the Fourier transform \hat{P} . However, to be able to use this technique effectively we also need the details of the irreducible representations of our group G . Using discrete Fourier transforms we may reduce the study of P^t to P . This connection was one of the crucial insights which permitted the study of random walks on groups via their representation theory. Let us now state some simple results following from the Fourier inversion Theorem.

Corollary 1.2.31 (Plancherel Theorem). *Let P, Q be functions on a finite group G . Then*

$$\sum_{g \in G} P(g)Q(g^{-1}) = \frac{1}{|G|} \sum_{\rho \in \text{Irr}(G)} d_\rho \text{Tr}(\hat{P}(\rho)\hat{Q}(\rho)). \quad (1.18)$$

Proof. Both sides of (1.18) are linear in Q , so taking $Q(h) = \delta_{g,h}$ we just have to show

$$P(g^{-1}) = \frac{1}{|G|} \sum_{\rho \in \text{Irr}(G)} d_\rho \text{Tr}(\rho(g)\hat{P}(\rho)).$$

This is just an application of the Fourier Inversion Theorem. Further Details may be found in [16, Chapter 2C] \square

Corollary 1.2.32. *Let P be a function on a finite group G . Then*

$$\sum_{g \in G} P(g)P(g) = \frac{1}{|G|} \sum_{\rho \in \text{Irr}(G)} d_\rho \text{Tr}(\hat{P}(\rho)\hat{P}(\rho)^*). \quad (1.19)$$

Proof. Apply Plancherel's Theorem with Q defined as $Q(g) = P(g^{-1})$. \square

Corollary 1.2.33. *Let P be a probability distribution on a finite group G , and π the uniform distribution on G . Then*

$$\sum_{g \in G} P(g)\pi(g) = \frac{1}{|G|}. \quad (1.20)$$

Proof. Following from Plancherel's Theorem we have

$$\begin{aligned} \sum_{g \in G} P(g)\pi(g) &= \frac{1}{|G|} \sum_{\rho \in \text{Irr}(G)} d_\rho \text{Tr}(\hat{P}(\rho)\hat{\pi}(\rho)) \\ &= \frac{1}{|G|} \sum_{\rho \in \text{Irr}(G)} d_\rho \text{Tr}(\hat{\pi}(\rho)) = \frac{1}{|G|} \sum_{\rho \in \text{Irr}(G)} d_\rho \frac{d_\rho}{|G|} = \frac{1}{|G|}. \end{aligned}$$

The second equality above comes from merging the Fourier transforms into one and applying Lemma 1.1.7. \square

The trace of Fourier transforms present in the Fourier inversion and Plancherel's Theorem turn the representation present in $\hat{P}(\rho), \hat{Q}(\rho)$ into the associated character χ_ρ . We have seen that characters are constant on conjugacy classes of G , so if also we impose this condition on P we may simplify the trace of $\hat{P}(\rho)$.

Lemma 1.2.34. *Let C_i denote the conjugacy classes of a finite group G , with representatives g_i . Let P be a probability on G which is constant on conjugacy classes of G . Then*

we have $\hat{P}(\rho) = \beta \text{Id}$ with

$$\beta = \sum_i |C_i| P(g_i) \frac{\chi_\rho(g_i)}{d_\rho} \quad (1.21)$$

Proof. It is enough to show that $\hat{P} : V \rightarrow V$ is a module homomorphism. This follows from the following calculation: for any $g \in G$ we have $\rho(g)\hat{P}(\rho)\rho(g^{-1}) = \sum_h P(h)\rho(ghg^{-1}) = \sum_h P(h)\rho(h) = \hat{P}(\rho)$. Therefore, applying Schur's Lemma gives $\hat{P}(\rho) = \beta \text{Id}$, and taking traces gives the required value of β . \square

1.2.3 The Upper Bound Lemma

We are now in a position to state one of the most important results on the mixing times of random walks on groups. The upper bound lemma links the worlds of probability and algebra in a surprising way, stating that the total variation distance of a random walk on a group G may be upper bounded via Fourier analysis on the driving probability of the walk. If the Fourier transforms on irreducible representation are well understood this reduces the issue of analysing $\|P^t - \pi\|_{\text{TV}}$ to bounding a summation of numbers in \mathbb{C} (often just \mathbb{R}). The results and proof presented below are taken from [16, Chapter 3B]

Lemma 1.2.35 (Upper Bound Lemma [18]). *Let P be a probability on a finite group G , π the uniform distribution on G . Then*

$$\|P - \pi\|_{\text{TV}}^2 \leq \frac{1}{4} \sum_{\substack{\rho \in \text{Irr}(G) \\ \rho \neq \text{Triv}}} d_\rho \text{Tr} \left(\hat{P}(\rho) \overline{\hat{P}(\rho)} \right) \quad (1.22)$$

Proof. Following from the definition of total variation distance we find:

$$\begin{aligned} 4\|P - \pi\|_{\text{TV}}^2 &= \left(\sum_{g \in G} |P(g) - \pi(g)| \right)^2 \\ &\leq |G| \sum_{g \in G} |P(g) - \pi(g)|^2 \end{aligned} \quad (1.23)$$

$$\begin{aligned} &= |G| \sum_{g \in G} P(g)P(g) - 2|G| \sum_{g \in G} P(g)\pi(g) + |G| \sum_{g \in G} \pi(g)\pi(g) \\ &= \sum_{\substack{\rho \in \text{Irr}(G) \\ \rho \neq \text{Triv}}} d_\rho \text{Tr} \left(\hat{P}(\rho) \overline{\hat{P}(\rho)} \right) \end{aligned} \quad (1.24)$$

Step (1.23) is due to the Cauchy-Schwarz inequality, and the final equality is Plancherel's Theorem and its Corollaries. \square

Lemma 1.2.36 (Lower Bound Lemma [16]). *Let P be a probability on a finite group G , π the uniform distribution on G . Then*

$$\frac{1}{4|G|} \sum_{\substack{\rho \in \text{Irr}(G) \\ \rho \neq \text{Triv}}} d_\rho \text{Tr} \left(\hat{P}(\rho) \overline{\hat{P}(\rho)} \right) \leq \|P - \pi\|_{\text{TV}}^2 \quad (1.25)$$

Proof. Following from the definition of total variation distance we find:

$$4\|P - \pi\|_{\text{TV}}^2 = \left(\sum_{g \in G} |P(g) - \pi(g)| \right)^2 \geq \sum_{g \in G} |P(g) - \pi(g)|^2 = \frac{1}{|G|} \sum_{\substack{\rho \in \text{Irr}(G) \\ \rho \neq \text{Triv}}} d_\rho \text{Tr} \left(\hat{P}(\rho) \overline{\hat{P}(\rho)} \right)$$

The first inequality is a general one, and the final equality is Plancherel's Theorem and its Corollaries. \square

Corollary 1.2.37. *Let P be a probability on a finite group G , and π the uniform distribution on G . Then*

$$\frac{1}{4|G|} \sum_{\substack{\rho \in \text{Irr}(G) \\ \rho \neq \text{Triv}}} d_\rho \text{Tr} \left(\hat{P}(\rho)^t \overline{\hat{P}(\rho)^t} \right) \leq \|P^t - \pi\|_{\text{TV}}^2 \leq \frac{1}{4} \sum_{\substack{\rho \in \text{Irr}(G) \\ \rho \neq \text{Triv}}} d_\rho \text{Tr} \left(\hat{P}(\rho)^t \overline{\hat{P}(\rho)^t} \right) \quad (1.26)$$

The first use of the Upper Bound Lemma appeared in the seminal paper of Diaconis and Shahshahani, where it was used to analyse the mixing time of the random transposition shuffle [18, Lemma 14]. We review this classic argument in detail in Chapter 2 to demonstrate the power of the upper bound lemma. Since its first appearance the upper bound lemma has seen frequent use in proving cutoff for many random walks on groups. It gives an explicit link between the areas of mixing time of Markov chains and representation theory, reducing the study of $\|P^t - \pi\|_{\text{TV}}$ to understanding $\hat{P}(\rho)$. The use of the Cauchy-Schwarz inequality in the Upper Bound Lemma may seem like a naïve way to bound this complicated summation, however we shall see that it frequently gives sharp bounds of the correct mixing time. The upper and lower bounds presented in Corollary 1.2.37 are of different orders, this means that we can not get the exact mixing time for a random walk just using representation theory.

In this thesis the main results will be focused on random walks on the symmetric group S_n . To get a better understanding of the techniques introduced in this chapter we dedicated the next chapter to exploring the symmetric group fully. We concentrate the tools previously seen to specific examples which have guided the analysis of the one-sided transposition shuffle, these are the random transposition shuffle, top-to-random shuffle, and random-to-random shuffle.

Chapter 2

The Symmetric Group and its Modules

In this chapter we review all the information needed about the symmetric group in order to analyse the mixing times of random walks on S_n . The first section is dedicated to the structure of the group. The second section presents a detailed construction of the simple modules of the symmetric group, by way of permutation modules. We then see how to view these modules as vector spaces spanned by words of length n . We end this chapter with detailed examples of shuffles which have guided the analysis of the one-sided transposition shuffle. These include cutoff results for the random transposition shuffle and the top-to-random shuffle, as well as an exploration into lifting eigenvectors for the random-to-random shuffle. A detailed account of the symmetric group and its representations may be found in *The Symmetric Group* by Sagan [38] or *The Representation Theory of the Symmetric Groups* by James [28].

2.1 The Symmetric Group

The symmetric group on n objects, denoted S_n , is defined as the group of all bijections $\sigma : [n] \rightarrow [n]$, where $[n] = \{1, \dots, n\}$. There are exactly $n!$ of these bijections. Given two bijections $\sigma, \eta \in S_n$ define the product $\sigma\eta$ as the composition of functions from right to left. Another common way to think about the symmetric group S_n is the arrangements of a deck of n cards labelled, $1, \dots, n$, from bottom to top. If we have n positions in a deck of cards labelled 1 to n , and n cards labelled 1 to n then we may view the bijection σ as telling us what position each card is in, i.e., card i is in position $\sigma(i)$. The identity element e has every card in its labelled position, i.e., for all $i \in [n]$ we have $e(i) = i$. Viewing the symmetric group as a deck of cards is often useful for random walks on S_n because it allows us to formulate our random walks in the expressive terms of shuffling a deck of cards.

We express elements of the symmetric group using *cycle notation*, for example,

$$\sigma = (125)(46)(3) = (125)(46).$$

We read the cycle (125) as 1 maps to 2 (i.e. $\sigma(1) = 2$), 2 maps to 5, and 5 maps to 1. The *cycle structure* of a permutation is the tuple of the lengths of its cycles arranged

in non-increasing order, example the cycle structure of σ defined above is $(3, 2, 1)$. Note that the sum of cycle lengths for $\sigma \in S_n$ is always n , therefore the cycle structure for any permutation is a *partition* of n .

Definition 2.1.1. A *partition* of n is a tuple of positive integers $\lambda = (\lambda_1, \dots, \lambda_r)$ such that, $\sum_{i=1}^r \lambda_i = n$, and $\lambda_1 \geq \dots \geq \lambda_r$, we denote this by $\lambda \vdash n$. We call n the *size* of the partition and r the *length* of the partition, denoted $|\lambda|$ and $l(\lambda)$ respectively. If a partition contains repeated digits we may write them as a power for brevity, for example we denote the partition $(1, \dots, 1) = (1^n)$, similarly $(3, 3, 2, 2, 2, 1) = (3^2, 2^3, 1)$.

Cycle notation can help us identify the conjugacy classes of the symmetric group. To know if two elements are conjugate in S_n we need only compare their cycle structures.

Lemma 2.1.2 (See Chapter 1 [38]). *Two elements in S_n are conjugate if and only if they have the same cycle structure. Hence, the conjugacy classes of S_n are labelled by partitions of n .*

The identity has cycle structure (1^n) , as every element belongs to its own trivial cycle. One particularly important conjugacy class of S_n is the class of transpositions, which is formed from all two cycles (ij) for $i < j$, these elements have cycle type $(2, 1^{n-2})$. We call the transposition (ij) an *adjacent transposition* if $j = i + 1$. The conjugacy class of transpositions generates the whole symmetric group, in fact we only need the $n-1$ adjacent transpositions for this, $\langle (i, i+1) \mid 1 \leq i \leq n-1 \rangle = S_n$. We may decompose any element in S_n into a product of transpositions, and in any decomposition the number of transpositions required remains constant modulo 2. This leads us to make the following definition.

Definition 2.1.3. Let $\sigma \in S_n$, and decompose our permutation as $\sigma = \tau_k \tau_{k-1} \dots \tau_1$, where τ_i is a transposition. Define the *sign* function for S_n , denoted $\text{sgn} : S_n \rightarrow \{1, -1\}$, as follows

$$\text{sgn}(\sigma) = (-1)^k.$$

This is a well defined function, i.e. independent of the decomposition of σ into transpositions. An element of $\sigma \in S_n$ is called *odd* if $\text{sgn}(\sigma) = -1$, and called *even* if $\text{sgn}(\sigma) = 1$. The sign function is also multiplicative, i.e., for $\sigma, \eta \in S_n$ we have $\text{sgn}(\sigma\eta) = \text{sgn}(\sigma)\text{sgn}(\eta)$. Therefore, the set of even permutations defines a subgroup of S_n called the *alternating group* and denoted A_n .

Random walks supported on the conjugacy class of transpositions have been well studied since the random transposition shuffle, variants include the semi-random transposition shuffles [34], adjacent transposition shuffle [29] and biased transposition shuffle [9]. In Chapter 3 we present a novel modification to general transposition shuffles called the one-sided transposition shuffle where the probability of applying transposition (ij) depends only on the position j . We now move on to describe the module structure of the symmetric group.

2.2 The Structure of Modules for The Symmetric Group

Denote the group algebra of the symmetric group as $\mathfrak{S}_n := \mathbb{C}[S_n]$. All the simple modules of \mathfrak{S}_n are indexed by partitions of $\lambda \vdash n$, we denote the simple module indexed by λ as

S^λ , these are also called *Specht* modules. Understanding the simple modules of \mathfrak{S}_n will be important to our analysis of the one-sided transposition shuffle presented in Chapter 3. Before we construct modules of the symmetric group, we need to recall some facts about partitions, Young diagrams and Young tableaux.

2.2.1 Young Diagrams

Every partition λ has an associated *Young diagram*, made by forming a left adjusted stack of boxes with rows labelled downwards and row i having λ_i boxes. We often blur the distinction between a partition and its Young diagram, e.g. $(3, 2) = \begin{array}{|c|c|} \hline \square & \square \\ \hline \square & \square \\ \hline \end{array}$. We may refer to the boxes of a diagram λ by using coordinates (i, j) to mean the box in the i^{th} row and j^{th} column. Define the *diagonal index* of the box (i, j) in partition λ to be the value $j - i$. Define the *diagonal sum* of λ to be $\text{Diag}(\lambda) = \sum_{(i,j) \in \lambda} (j - i)$, this is the sum over all the diagonal indexes of λ .

Given a partition $\lambda \vdash n$, we may form the *transpose* of λ , denoted λ' , by swapping rows and columns in the Young diagram, e.g. $(3, 2)' = (2, 2, 1)$. We have $\lambda \vdash n$ if and only if $\lambda' \vdash n$. Define a partial order on partitions of n called the *dominance order*: in terms of Young diagrams, for two partitions $\mu, \lambda \vdash n$, we say λ *dominates* μ if we can form μ by moving boxes of λ down and to the left, we denote this by $\lambda \trianglerighteq \mu$. Equivalently $\lambda \trianglerighteq \mu$ if and only if $\sum_{i=1}^j \lambda_i \geq \sum_{i=1}^j \mu_i$ for all choices of j . Furthermore, we have $\lambda \trianglerighteq \mu$ if and only if $\mu' \trianglerighteq \lambda'$, see [26, Lemma 1.4.11].

Example 2.2.1. Let $n = 8$ and consider the partitions $(3, 2, 2, 1)$ and (2^4) . We may see that $(3, 2, 2, 1) \trianglerighteq (2^4)$, and $(2^4)' = (4, 4) \trianglerighteq (4, 3, 1) = (3, 2, 1, 1)'$. Now consider partitions $(5, 1^3)$ and (4^2) , neither partition here dominates the other. Hence, the dominance order on partitions is not necessarily a total ordering.

Given two partitions μ, λ of different sizes, we write $\mu \subseteq \lambda$ if μ is fully contained in λ when we align the Young diagrams of μ and λ at the top left corners; equivalently, if we write $\lambda = (\lambda_1, \dots, \lambda_r)$ and $\mu = (\mu_1, \dots, \mu_s)$, this means that $s \leq r$ and $\mu_i \leq \lambda_i$ for each $1 \leq i \leq s$. For example $(3, 2) \subseteq (4, 3)$ – this is simpler to see from the corresponding Young diagrams, $\begin{array}{|c|c|} \hline \square & \square \\ \hline \square & \square \\ \hline \end{array} \subseteq \begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \square & \square & \square \\ \hline \end{array}$. If we have two partitions μ, λ such that $\mu \subseteq \lambda$ we may define the *skew diagram* λ/μ as the diagram containing all boxes which are in λ but not in μ . For example the skew diagram of $(4, 3)/(3, 2)$ is $\begin{array}{|c|c|} \hline \square & \square \\ \hline \square & \square \\ \hline \end{array}$. A skew diagram λ/μ is called a horizontal strip if it has at most one box per column, e.g. $(4, 3)/(3, 2)$ is a horizontal strip but $(4, 3)/(2, 2) = \begin{array}{|c|c|} \hline \square & \square \\ \hline \square & \square \\ \hline \end{array}$ is not.

Given a partition of n we may turn it into a partition of $n + 1$ by adding a box to the corresponding Young diagram. In order to formalise this we view partitions as n -tuples, filling the end with zeros if necessary. Given an n -tuple $\lambda = (\lambda_1, \dots, \lambda_n)$ of non-negative integers summing to n and an element $i \in [n + 1]$, we form an $(n + 1)$ -tuple denoted $\lambda + e_i$ by first adding a zero to the end of λ and then adding 1 to this $(n + 1)$ -tuple in position i . Then $\lambda + e_i$ is an $(n + 1)$ -tuple of non-negative integers summing to $n + 1$, e.g. $(1, 1, 1) + e_3 = (1, 1, 1, 0) + (0, 0, 1, 0) = (1, 1, 2, 0)$. In terms of Young diagrams $\lambda + e_i$ represents adding a box to λ on row i . Note that if $\mu \subseteq \lambda$ we may add boxes to μ to form λ , the boxes we have to add are exactly those contained in the skew diagram λ/μ , e.g. for

$(3, 2) \subseteq (4, 3)$ we may see that $(3, 2) + e_1 + e_2 = (4, 3)$. If we restrict our attention to choices of box e_i that result in another partition we uncover a structure on Young diagrams called *Young's lattice*, seen in Figure 2.1. We allow $n = 0$ as a special case with empty partition (0) and corresponding Young diagram \emptyset .




Figure 2.1: Young's lattice for partitions of size $n \in \{0, 1, 2, 3, 4\}$.

Each path in Figure 2.1 represents the placement or removal of a box to form a new partition. Young's lattice is an important structure as it allows us to link the partitions of n to those of $n + 1$. This structure play an important role in the linking the modules of S_n and S_{n+1} in Section 2.2.5.

2.2.2 Young Tableaux

Given a partition $\lambda \vdash n$, we may form a *Young tableau* (alternatively a λ -*tableau*) T by putting numbers $1, \dots, n$ into the boxes of (the Young diagram of) λ , such that each number only appears once. The set of Young tableaux of shape λ is denoted $\text{YT}(\lambda)$. We says a value m *occurs* in the Young tableau T if m is present in a box of T . A *standard Young tableau* T is a Young tableau where the values in the boxes of T are increasing across rows and down columns. The set of standard Young tableaux of shape λ is denoted by $\text{SYT}(\lambda)$. The size of the set $\text{SYT}(\lambda)$ is called the *dimension* of λ , denoted d_λ . For a tableau T , form the transpose of T , denoted T' , by swapping rows and columns while preserving the value in each box. A tableau T has $T \in \text{SYT}(\lambda)$ if and only if $T' \in \text{SYT}(\lambda')$. If a tableau T has a box in position (i, j) , we let $T(i, j)$ denote the value in that box (i, j) ; otherwise $T(i, j)$ is undefined. The partition $(3, 2)$ has 5 standard Young tableaux therefore $d_{(3,2)} = 5$, these are given below:

$$\begin{array}{ccc}
 \begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & 5 & \\ \hline \end{array} &
 \begin{array}{|c|c|c|} \hline 1 & 2 & 4 \\ \hline 3 & 5 & \\ \hline \end{array} &
 \begin{array}{|c|c|c|} \hline 1 & 2 & 5 \\ \hline 3 & 4 & \\ \hline \end{array} \\
 \\
 \begin{array}{|c|c|c|} \hline 1 & 3 & 4 \\ \hline 2 & 5 & \\ \hline \end{array} &
 \begin{array}{|c|c|c|} \hline 1 & 3 & 5 \\ \hline 2 & 4 & \\ \hline \end{array} &
 \end{array}$$

Each standard Young tableau corresponds to one path up Young's lattice, starting at \emptyset and ending at λ . To form this correspondence take a standard Young tableau $T \in \text{SYT}(\lambda)$ and form a path up Young's lattice from \emptyset to λ by adding the boxes to \emptyset in the order given by the entries in T . A standard Young tableau $T \in \text{SYT}(\lambda)$ is called a *desarrangement* tableau if there is no box $(1, 2) \in \lambda$ and $|\lambda|$ is even, or the value $T(1, 2)$ is odd. We denote the number of desarrangement tableaux of a shape λ as d^λ . For example, in the collection of standard Young tableaux of shape $(3, 2)$ displayed above, those in the bottom row are desarrangement tableaux.

Finally we end with a result that links the tableaux of different shapes to the dominance order of partitions.

Lemma 2.2.2 (Lemma 2.2.4 [38]). *Let $\lambda, \mu \vdash n$. Take T_λ to be a λ -tableau and T_μ to be a μ -tableau. Suppose that for each index i , the elements in row i of T_μ are all in different columns of T_λ . Then $\lambda \succeq \mu$.*

Proof. By our hypothesis the elements in row 1 of T_μ are in different columns of T_λ . We may rearrange the columns of T_λ so that these elements all appear in row 1 of T_λ . Continuing this we may sort the columns of T_λ so that the elements of rows $1, 2, \dots, j$ of T_μ all occur in the first j rows of T_λ . Then,

$$\begin{aligned} \sum_{i=1}^j \lambda_i &= \text{number of elements in the first } j \text{ rows of } T_\lambda \\ &\geq \text{number of elements in the first } j \text{ rows of } T_\mu = \sum_{i=1}^j \mu_i. \quad \square \end{aligned}$$

2.2.3 Permutation Modules

On our way to defining the simple modules S^λ for \mathfrak{S}_n we first need to define the *permutation modules*, denoted M^λ , which contain the simple modules as submodules. To define the permutation module M^λ we use the natural action of the symmetric group on the Young tableaux of shape λ .

Definition 2.2.3. Let T be a Young tableau of n , and $\sigma \in S_n$. Define an action of σ on T , by applying σ to the values in T box wise. Thus σT defines a new Young tableau of the same shape as T . For example, given $T = \begin{array}{|c|c|c|} \hline 5 & 4 & 1 \\ \hline 2 & 3 & \\ \hline \end{array}$ and $\sigma \in S_n$, the product

$$\sigma T = \begin{array}{|c|c|c|} \hline \sigma(5) & \sigma(4) & \sigma(1) \\ \hline \sigma(2) & \sigma(3) & \\ \hline \end{array}.$$

To form the permutation module M^λ we first establish an important equivalence relation on the tableaux of shape λ . Using the action in Definition 2.2.3 we may define the groups of row and column permutations of a tableau T .

Definition 2.2.4. Let T be a tableau of size n . Define the subgroup of *row permutations* of T , denoted R_T , as the set of elements of S_n which preserves each row of T . Define the subgroup of *column permutations* of T , denoted C_T , as the set of elements of S_n which preserves each columns of T . For example, if T is the tableau given in Definition 2.2.3

then $R_T \cong S_3 \times S_2$ as we are allowed to swap the numbers 5, 4, 1 freely and 2, 3 freely but not swap between them, similarly $C_T \cong S_2 \times S_2 \times S_1$.

Definition 2.2.5. Let T_1, T_2 be two tableaux of the same Young diagram λ . We say T_1 is *row equivalent* to T_2 , denoted $T_1 \sim_R T_2$, if there exists $\sigma \in R_{T_1}$ such that $\sigma T_1 = T_2$. Similarly we say T_1 is *column equivalent* to T_2 , denoted $T_1 \sim_C T_2$, if there exists $\sigma \in C_{T_1}$ such that $\sigma T_1 = T_2$.

Example 2.2.6. Let $T = \begin{array}{|c|c|c|} \hline 5 & 4 & 1 \\ \hline 2 & 3 & \\ \hline \end{array}$, we know that $|R_T| = 12$, therefore there are 12 row equivalent tableaux. 6 of these are found by permuting the first row of T , we list them below:

$$\begin{array}{|c|c|c|} \hline 1 & 4 & 5 \\ \hline 2 & 3 & \\ \hline \end{array} \quad \begin{array}{|c|c|c|} \hline 1 & 5 & 4 \\ \hline 2 & 3 & \\ \hline \end{array} \quad \begin{array}{|c|c|c|} \hline 4 & 1 & 5 \\ \hline 2 & 3 & \\ \hline \end{array} \quad \begin{array}{|c|c|c|} \hline 4 & 5 & 1 \\ \hline 2 & 3 & \\ \hline \end{array} \quad \begin{array}{|c|c|c|} \hline 5 & 1 & 4 \\ \hline 2 & 3 & \\ \hline \end{array} \quad \begin{array}{|c|c|c|} \hline 5 & 4 & 1 \\ \hline 2 & 3 & \\ \hline \end{array} .$$

The remaining 6 row equivalent tableaux can be found by swapping 2 and 3 in the second row.

Lemma 2.2.7. *The relation \sim_R defines an equivalence relation on the tableaux of shape λ .*

Proof. This follows from R_T forming a subgroup of S_n . □

Definition 2.2.8. The row equivalence class of a tableau T is denoted by $\{T\}$ and called a *tabloid*.

Our action of S_n on tableaux extends naturally to an action on tabloids given by $\sigma\{T\} = \{\sigma T\}$. We are now in a position to define our permutation modules using tabloids of a given partition.

Definition 2.2.9. Let $\lambda \vdash n$. The *permutation module* for \mathfrak{S}_n , denoted M^λ , is the following vector space

$$M^\lambda = \langle \{T\} \mid T \text{ is a tableau of shape } \lambda \rangle$$

with action of \mathfrak{S}_n given by extending the action of the symmetric group on tabloids (Definition 2.2.3) linearly.

There are $5!$ Young tableaux of shape $(3, 2)$ (in fact of any shape $\lambda \vdash 5$), but the module $M^{(3,2)}$ has a basis of 10 tabloids, because there are only two rows to split the five numbers $\{1, 2, 3, 4, 5\}$ across, giving $\binom{5}{3}$ choices for the first row, with the second row containing the leftover numbers. Note that for the permutation module we may have basis elements that are not in the equivalence class of a standard Young tableau. E.g., the module $M^{(3,2)}$ is spanned by the following 10 tabloids:

$$\begin{array}{ccccc} \left\{ \begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & 5 & \\ \hline \end{array} \right\} & \left\{ \begin{array}{|c|c|c|} \hline 1 & 2 & 4 \\ \hline 3 & 5 & \\ \hline \end{array} \right\} & \left\{ \begin{array}{|c|c|c|} \hline 1 & 2 & 5 \\ \hline 3 & 4 & \\ \hline \end{array} \right\} & \left\{ \begin{array}{|c|c|c|} \hline 1 & 3 & 4 \\ \hline 2 & 5 & \\ \hline \end{array} \right\} & \left\{ \begin{array}{|c|c|c|} \hline 1 & 3 & 5 \\ \hline 2 & 4 & \\ \hline \end{array} \right\} \\ \left\{ \begin{array}{|c|c|c|} \hline 1 & 4 & 5 \\ \hline 2 & 3 & \\ \hline \end{array} \right\} & \left\{ \begin{array}{|c|c|c|} \hline 2 & 3 & 4 \\ \hline 1 & 5 & \\ \hline \end{array} \right\} & \left\{ \begin{array}{|c|c|c|} \hline 2 & 3 & 5 \\ \hline 1 & 4 & \\ \hline \end{array} \right\} & \left\{ \begin{array}{|c|c|c|} \hline 2 & 4 & 5 \\ \hline 1 & 3 & \\ \hline \end{array} \right\} & \left\{ \begin{array}{|c|c|c|} \hline 3 & 4 & 5 \\ \hline 1 & 2 & \\ \hline \end{array} \right\} \end{array}$$

Only the first row above are the tabloid classes of standard Young tableaux.

2.2.4 Specht Modules

Every M^λ contains a unique copy of the simple module S^λ . We construct the Specht module S^λ as a stable subspace of M^λ before going on to show it is simple. We form a basis for the Specht module S^λ by taking taking linear combinations of tabloids, we call these new elements *polytabloids*.

Definition 2.2.10. Given a Young tableau T form the element $k_T \in \mathfrak{S}_n$ as the following sum:

$$k_T = \sum_{\sigma \in C_T} \text{sgn}(\sigma) \sigma \quad (2.1)$$

The *polytabloid* generated by tableau T , denoted e_T , is given by the element $e_T = k_T \{T\}$. Every polytabloid is an element of the permutation module M^λ . We say the polytabloid e_{T_1} contains the tabloid $\{T_2\}$ if it appears with a non-zero coefficient in e_{T_1}

Lemma 2.2.11 (Lemma 2.3.2 [38]). *Let T be a tableau of size n , and η is an element in S_n . Then $\eta e_T = e_{\eta T}$.*

Proof.

$$\eta e_T = \eta \left(\sum_{\sigma \in C_T} \text{sgn}(\sigma) \sigma \right) \eta^{-1} \eta \{T\} = \sum_{\sigma \in C_T} \text{sgn}(\sigma) \eta \sigma \eta^{-1} \{\eta T\} = \sum_{\sigma \in C_{\eta T}} \text{sgn}(\sigma) \sigma \{\eta T\}. \quad \square$$

Definition 2.2.12. Let $\lambda \vdash n$, the *Specht module* S^λ is defined to be the following vector space

$$S^\lambda = \langle e_T \mid T \in \text{YT}(\lambda) \rangle$$

with the natural action of \mathfrak{S}_n on tabloids. In fact we may restrict this spanning set to a basis by considering only polytabloids formed from standard Young tableaux,

$$S^\lambda = \langle e_T \mid T \in \text{SYT}(\lambda) \rangle$$

see [38, Theorem 2.6.2] for a proof of this. Thus, the Specht module S^λ has dimension d_λ . The Specht modules are also cyclic, that is S^λ is generated as an \mathfrak{S}_n -module by any one polytabloid.

Example 2.2.13. Let $\lambda = (3, 1)$, up to row equivalence we have 4 Young tableaux of this shape, we label them by which number appears in the second row

$$T_4 = \begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & & \\ \hline \end{array} \quad T_3 = \begin{array}{|c|c|c|} \hline 1 & 2 & 4 \\ \hline 3 & & \\ \hline \end{array} \quad T_2 = \begin{array}{|c|c|c|} \hline 1 & 3 & 4 \\ \hline 2 & & \\ \hline \end{array} \quad T_1 = \begin{array}{|c|c|c|} \hline 2 & 3 & 4 \\ \hline 1 & & \\ \hline \end{array}.$$

The permutation module M^λ is four dimensional and spanned by the four tabloids above. We can see that T_4, T_3, T_2 are all standard Young tableau of shape λ , and so the dimension of $S^{(3,1)}$ is 3. To form the polytabloids that make up S^λ we need to use the group of column permutations of each tableau. In this example $C_T \cong S_2$ with each tableau having a single transposition which may be applied to the first column. Therefore,

$$e_{T_4} = \left\{ \begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & & \\ \hline \end{array} \right\} - (14) \left\{ \begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & & \\ \hline \end{array} \right\} = \left\{ \begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & & \\ \hline \end{array} \right\} - \left\{ \begin{array}{|c|c|c|} \hline 4 & 2 & 3 \\ \hline 1 & & \\ \hline \end{array} \right\} = \{T_4\} - \{T_1\}.$$

Performing the same calculation for the tableaux T_3 and T_2 , we may write a basis for S^λ as follows:

$$S^{(3,1)} = \langle e_{T_4}, e_{T_3}, e_{T_2} \rangle = \{ \{T_4\} - \{T_1\}, \{T_3\} - \{T_1\}, \{T_2\} - \{T_1\} \}.$$

Furthermore, the Specht module $S^{(3,1)}$ may be generated by a single polytabloid e_{T_4} ,

$$\begin{aligned} (34)e_{T_4} &= (34)\{T_4\} - (34)\{T_1\} = \{T_3\} - \{T_1\} = e_{T_3} \\ (24)e_{T_4} &= (24)\{T_4\} - (24)\{T_1\} = \{T_2\} - \{T_1\} = e_{T_2} \end{aligned}$$

as Lemma 2.2.11 asserts.

We know the conjugacy classes of S_n are labelled by partitions of n , and we have found exactly one Specht module for each partition. To be sure $\{S^\lambda \mid \lambda \vdash n\}$ forms a complete set of simple modules we have to check that they are in fact simple, and all distinct from one another. To do this we study how k_T acts on tableaux of different shapes. The following arguments were presented by James in [25], and further details of them can be found in Sections 2.4 and 2.5 in [38]

Lemma 2.2.14. *Let $\lambda, \mu \vdash n$. Take T_λ to be a λ -tableau and T_μ to be a μ -tableau. Suppose there exists a, b which occur in the same row of T_μ and same column of T_λ . Then $k_{T_\lambda}\{T_\mu\} = 0$.*

Proof. From our hypothesis we have $(e - (ab))\{T_\mu\} = 0$. Take signed coset representatives $\{\sigma_1, \dots, \sigma_k\}$ for the subgroup $\{e, (ab)\} \subseteq C_{T_\lambda}$. Then $K_{T_\lambda}\{T_\mu\} = \left(\sum_{\sigma \in C_{T_\lambda}} \text{sgn}(\sigma) \sigma \right) \{T_\mu\} = \left(\sum_{i=1}^k \sigma_i (e - (ab)) \right) \{T_\mu\} = 0$. \square

Applying Lemma 2.2.14 in combination with Lemma 2.2.15, we are able to use the value of $k_{T_\lambda}\{T_\mu\}$ to tell us information about λ and μ .

Lemma 2.2.15. *Let $\lambda, \mu \vdash n$. Take T_λ to be a λ -tableau and T_μ to be a μ -tableau. If $k_{T_\lambda}\{T_\mu\} \neq 0$, then $\lambda \supseteq \mu$. Furthermore if $\lambda = \mu$ then $k_{T_\lambda}\{T_\mu\} = \pm e_{T_\lambda}$.*

Proof. If $k_{T_\lambda}\{T_\mu\} \neq 0$ then for all a, b in the same row of μ they are in different columns of λ . Therefore, by Lemma 2.2.2 we know $\lambda \supseteq \mu$. If $\lambda = \mu$, then we must have our tabloids being column equivalent otherwise we violate $k_{T_\lambda}\{T_\mu\} \neq 0$ (see [38, Corollary 2.4.2]). Hence, there must exist $\eta \in C_{T_\lambda}$ such that $\{T_\mu\} = \eta\{T_\lambda\}$, therefore

$$k_{T_\lambda}\{T_\mu\} = \left(\sum_{\sigma \in C_T} \text{sgn}(\sigma) \sigma \eta \right) \{T_\lambda\} = \left(\sum_{\sigma \in C_T} \text{sgn}(\eta^{-1}) \text{sgn}(\sigma) \sigma \right) \{T_\lambda\} = \text{sgn}(\eta) e_{T_\lambda}. \quad \square$$

Corollary 2.2.16. *Let $v \in M^\mu$, and take T_μ a μ -tableau. Then $k_{T_\mu} v$ is a multiple of e_{T_μ} .*

Proof. Write $v = \sum_i c_i \{T_i\}$ where T_i are μ -tableaux. By Lemma 2.2.15 each summand $k_{T_\mu}\{T_i\}$ is either a multiple of e_{T_μ} or 0. \square

We are now in a position to prove that the Specht module S^λ is simple, and that $S^\lambda \cong S^\mu$ if and only if $\lambda = \mu$. These facts together tell us that the Specht modules form a complete set of simple modules for \mathfrak{S}_n .

Theorem 2.2.17 (Submodule Theorem). *Let $V \subseteq M^\lambda$ be a submodule. Then $V \supseteq S^\lambda$ or $V \subseteq (S^\lambda)^0$. Therefore, the Specht module S^λ is simple.*

Proof. Take $v \in V$, and a λ -tableau T . By Corollary 2.2.16 we know that $k_T v = c e_T$ for some $c \in \mathbb{C}$. We consider two cases: $c \neq 0$ for some T , and $c = 0$ always. Suppose that there exists v and T with $k_T v = c e_T \neq 0$. Then we have $c^{-1} k_T v = e_T \in V$. Hence, by Lemma 2.2.11 we may generate S^λ inside of V .

Now suppose we always have $k_T v = 0$. Consider the inner product on M^λ defined on tabloids by $\langle \{T_1\}, \{T_2\} \rangle = \delta_{\{T_1\}, \{T_2\}}$, this inner product is S_n -invariant. We now find,

$$\langle v, e_T \rangle = \langle v, \sum_{\sigma \in C_T} \text{sgn}(\sigma) \sigma \{T\} \rangle = \langle \sum_{\sigma \in C_T} \text{sgn}(\sigma) \sigma^{-1} v, \{T\} \rangle = \langle k_T v, \{T\} \rangle = \langle 0, \{T\} \rangle = 0.$$

A single polytabloid e_T spans S^λ , therefore $v \notin S^\lambda \Rightarrow v \in (S^\lambda)^0$. \square

Theorem 2.2.18. *Let $\lambda, \mu \vdash n$. Suppose we have a non-zero homomorphism $\psi : S^\lambda \rightarrow M^\mu$. Then $\lambda \supseteq \mu$ and if $\lambda = \mu$ then ψ is multiplication by a scalar.*

Proof. Take a basis vector $e_T \in S^\lambda$ such that $\psi(e_T) \neq 0$. Extend the homomorphism ψ to a homomorphism $\psi : M^\lambda \rightarrow M^\mu$ by setting it to be zero on the complement of S^λ . Then,

$$0 \neq \psi(e_T) = k_T \psi(\{T\}) = k_T \left(\sum_i c_i \{T_i\} \right)$$

where the T_i are μ -tableaux, and we must have at least one c_i being non-zero. Hence, by Lemma 2.2.15 we have $\lambda \supseteq \mu$. If $\lambda = \mu$, then we know $\psi(e_T) = c e_T$ for some constant $c \in \mathbb{C}$, and for any $\sigma \in S_n$ we have

$$\psi(e_{\sigma T}) = \psi(\sigma e_T) = \sigma \psi(e_T) = c \cdot \sigma e_T = c \cdot e_{\sigma T} \quad \square.$$

Corollary 2.2.19. *Let $\lambda, \mu \vdash n$. Then $S^\lambda \cong S^\mu$ if and only if $\mu = \lambda$.*

Proof. If $\mu = \lambda$ then the conclusion is immediate. Suppose $S^\lambda \cong S^\mu$, then there exist non-zero homomorphisms $\psi : S^\lambda \rightarrow M^\mu$, and $\varphi : S^\mu \rightarrow M^\lambda$. Therefore, by Theorem 2.2.18, $\lambda \supseteq \mu$ and $\mu \supseteq \lambda$, which implies $\lambda = \mu$. \square

Corollary 2.2.20 (Theorem 2.4.6 [38]). *The Specht modules S^λ for $\lambda \vdash n$ form a complete set of non-isomorphic simple modules for S_n .*

Proof. We have found the correct number of non-isomorphic simple modules. \square

Now we have a complete set of simple modules for the symmetric groups we proceed to decompose each permutation module into its simple parts. The following result is sometimes known as *Young's Rule*.

Lemma 2.2.21 (Young's Rule). *For $\mu \vdash n$ we have,*

$$M^\mu \cong \bigoplus_{\lambda \supseteq \mu} K_{\lambda, \mu} S^\lambda,$$

where $K_{\lambda, \mu} S^\mu$ denotes a direct sum of $K_{\lambda, \mu}$ copies of S^μ . The coefficients $K_{\lambda, \mu} \in \mathbb{N}$ are called Kostka numbers, and for all $\lambda \vdash n$ we know $K_{\lambda, \lambda} = 1$.

Proof. If S^λ appears with a non-zero coefficient in the decomposition of M^μ then we clearly have a non-zero homomorphism $\psi : S^\lambda \rightarrow M^\mu$, therefore $\lambda \succeq \mu$. To establish that all $\lambda \succeq \mu$ appear in our decomposition requires information about semi-standard Young tableaux which will not feature elsewhere in this thesis, so we leave the details which may be found in Section 2.10 [38]. If $\lambda = \mu$, we know any morphism $\psi : S^\lambda \rightarrow M^\lambda$ is multiplication by a scalar hence there is only one copy in of S^λ in M^λ . \square

Lemma 2.2.21 shows a instance of Maschke's Theorem (Theorem 1.2.14). Notably this theorem does not tell us the value of the Kostka numbers but importantly it tells us that S^μ appears as a submodule of M^λ if and only if $\mu \succeq \lambda$. Consider the permutation module $M^{(3,1)}$ from Example 2.2.13, we know that $S^{(3,1)}$ appears as one composition factor. This leaves a 1 dimensional submodule left to find, and Young's rule tells us it must be the Specht module $S^{(4)}$, as the only partition of 4 which dominates $(3,1)$ is (4) .

We have previously seen that the regular module for any group has a canonical decomposition into the simple modules for that group. For the symmetric group the regular module \mathfrak{S}_n may be seen to be the permutation module $M^{(1^n)}$.

Lemma 2.2.22. *The permutation module $M^{(1^n)} \cong \mathfrak{S}_n$ as modules. Therefore, $M^{(1^n)}$ has canonical decomposition*

$$M^{(1^n)} \cong \bigoplus_{\lambda \vdash n} d_\lambda S^\lambda \text{ as } \mathfrak{S}_n\text{-modules.}$$

This decomposition satisfies Lemma 2.2.21, as every partition of n dominates (1^n) .

Proof. The module $M^{(1^n)}$ is spanned by $n!$ tabloids. To define a linear map $\psi : \mathfrak{S}_n \rightarrow M^{(1^n)}$ it is enough to define it on each permutation $\sigma \in S_n$. To do this set,

$$\psi(\sigma) = \left\{ \begin{array}{c} \sigma(1) \\ \dots \\ \sigma(n) \end{array} \right\}.$$

The map ψ respects the action of both modules and is an isomorphism of vector spaces, therefore $\mathfrak{S}_n \cong M^{(1^n)}$ as \mathfrak{S}_n -modules. \square

The decomposition present in Lemma 2.2.22 is important for the analysis of the random-to-random shuffle and the one-sided transposition shuffle. It allows us to focus on the Specht modules of S_n as opposed to the larger space \mathfrak{S}_n .

2.2.5 Branching Rules for Specht Modules

Every symmetric group has a natural embedding into symmetric groups of a greater size. To embed $S_m \hookrightarrow S_n$ for $m < n$, we extend every permutation $\sigma \in S_m$ to a permutation in S_n by choosing it to fix all elements of $[m] \setminus [n]$. The group algebras of the symmetric groups and subsequently \mathfrak{S}_n -modules inherit this recursive structure. A module S^λ for $\lambda \vdash n$ may be viewed as an \mathfrak{S}_{n-1} module by the action of S_{n-1} inside of S_n . Conversely we may also turn \mathfrak{S}_n -modules into \mathfrak{S}_{n+1} -modules by a process known as induction. We now introduce the definition of restricted and induced modules.

Definition 2.2.23. Let $H \subseteq G$ be groups. Suppose W is a H -module and V is a G -module. The *restriction* $\text{Res}_H^G V$ of V consists of the vector space V viewed as a module for H by restricting the action of G to the subgroup H .

Let $\{g_1, \dots, g_n\}$ be a set of coset representatives for H in G . Then the *induced module* of W to a G -module, denoted $\text{Ind}_H^G W$, has vector space $\oplus_i (\{g_i\} \otimes W)$ and action given by

$$g \left(\sum_{i=1}^n (g_i, w_i) \right) = \sum_{i=1}^n (g_{r(i)}, h_i w_i)$$

where r_i, h_i are the unique values satisfying $g \cdot g_i = g_{r(i)} \cdot h_i$. This definition may be extended to include restricting/inducing representations and characters as well, see [27, Chapter 21].

There is a close link between the restriction and induction of modules, this is shown off best in a theorem known as *Frobenius reciprocity* (see [27, Theorem 21.16]).

Theorem 2.2.24 (Frobenius Reciprocity). *Let $H \subseteq G$ be groups, and suppose φ and χ are characters of H and G respectively. Then*

$$\langle \text{Ind}_H^G \varphi, \chi \rangle_G = \langle \varphi, \text{Res}_H^G \chi \rangle_H,$$

where the inner product for characters is computed over G and H respectively.

The decomposition of an induced module $\text{Ind}_H^G V$ into the simple modules of G , or the restricted module $\text{Res}_H^G W$ into simple modules of H , is an important question in representation theory. Results which allow us to relate the of the modules of G in terms of those for H and vice versa are called *branching rules*. The natural recursive structure of the symmetric group allows us an answer to this question for $S_{n-1} \subseteq S_n$. A proof of the following result can be found in [38, Theorem 2.8.3]

Theorem 2.2.25 (Branching rules for S_n). *Let $n \geq 1$, and $\lambda \vdash n$. The branching rules for the simple module of the symmetric group are as follows:*

$$\text{Res}_{S_{n-1}}^{S_n} S^\lambda \cong \bigoplus_{\substack{\mu \vdash n-1 \\ \mu \subseteq \lambda}} S^\mu \text{ as } \mathfrak{S}_{n-1}\text{-modules} \quad (2.2)$$

$$\text{Ind}_{S_n}^{S_{n+1}} S^\lambda \cong \bigoplus_{\substack{\mu \vdash n+1 \\ \lambda \subseteq \mu}} S^\mu \text{ as } \mathfrak{S}_{n+1}\text{-modules} \quad (2.3)$$

The branching rules for the symmetric group are closely related to Young's lattice. The direct sum of (2.2) could be rephrased as take a direct sum of all Specht modules found by removing a box of λ , similarly the direct sum (2.3) may be thought as taking a sum of all Specht modules found by adding a box to λ . The branching rules allow us to describe the restriction or induction of any \mathfrak{S}_n -module, all we must do is decompose it into its simple summands and then apply Theorem 2.2.25 to each part individually. The recursive structure of \mathfrak{S}_n -modules is key to the study of the random-to-random shuffle and one-sided transposition shuffle.

2.2.6 Switching to Words

The notation of tabloids is cumbersome, we therefore introduce a one-to-one correspondence between certain words of length n and tabloids of size n , allowing us to describe our

permutation and Specht modules in more succinct notation. This notation is particularly useful when studying the action of \mathfrak{S}_n on these modules.

Given $n \in \mathbb{N}^0$ we denote by W^n the set of words of length n with letters in $[n]$, where by a word of length n we simply mean a string $w = w_1 w_2 \dots w_n$ with $w_i \in [n]$ for all i . For $n = 0$ we allow a special case where W^0 is comprised solely of the empty word denoted ω . Note that in forming words we regard the elements of $[n]$ as distinct symbols. Later on it will be notationally convenient to have our words comprised of positive integers, for example $W^2 = \{11, 12, 21, 22\}$. The size of the set W^n is n^n .

There is a natural action of the symmetric group S_n on W^n . For a word $w = w_1 w_2 \dots w_n \in W^n$ and an element $\sigma \in S_n$, we let $\sigma w := w_{\sigma^{-1}(1)} w_{\sigma^{-1}(2)} \dots w_{\sigma^{-1}(n)} \in W^n$. We emphasise that this is the action of S_n on words by place permutations, it is not the action of S_n acting on the individual letters that comprise a word, e.g. if $\sigma = (123)$ then $\sigma(232) = 223 \neq 313$. Let M^n be the vector space over the field \mathbb{C} with basis of words in W^n . The action of S_n on words in W^n extends linearly to an action of \mathfrak{S}_n on the vector space M^n . Thus, M^n is an n^n -dimensional module for the group algebra \mathfrak{S}_n . For elements of M^n we use the notation \cdot to separate the complex coefficients from the words in W^n , e.g. $2 \cdot 232 + 4i \cdot 213$.

To each word $w \in W^n$ we can associate an n -tuple of non-negative integers, called its *evaluation*, denoted $\text{eval}(w)$, as follows. For $1 \leq i \leq n$, let $\text{eval}_i(w)$ count the number of occurrences of the symbol i in the word w , and then let $\text{eval}(w) := (\text{eval}_1(w), \dots, \text{eval}_n(w))$. Note that $\sum_{i=1}^n \text{eval}_i(w) = n$ for any word w in W^n . For example taking $232 \in W^3$, its evaluation is $\text{eval}(232) = (0, 2, 1)$. If in addition $\text{eval}(w)$ is a non-increasing sequence of integers, then we identify $\text{eval}(w)$ with the corresponding partition of n , ignoring possible ending zeros at the end of the partition. For example $w = 1231 \in W^4$ has evaluation $\text{eval}(w) = (2, 1, 1, 0)$ and we associate it to the partition $(2, 1, 1)$. Note that the evaluation of a word is unchanged by the action of S_n , i.e. for any $\sigma \in S^n$ and $w \in W^n$, we have $\text{eval}(w) = \text{eval}(\sigma w)$. Thus we may find a stable subspace of M^n corresponding to the words with a given evaluation.

Definition 2.2.26. Let ν be a n -tuple of non-negative integers. Define the module M^ν as the following vector space

$$M^\nu = \{w \in W^n \mid \text{eval}(w) = \nu\} \subseteq M^n.$$

For a partition $\lambda \vdash n$ the module M^λ from Definition 2.2.26 is exactly the permutation module we defined in Definition 2.2.9. To establish this we create a correspondence between words of evaluation λ and Young tableau of shape λ . For a partition $\lambda \vdash n$, a Young tableaux T of shape λ naturally corresponds to a word in W^n .

Definition 2.2.27. Let $\lambda \vdash n$. Define a map $w : \text{YT}(\lambda) \rightarrow W^n$ as follows: for each tableau T of shape λ , let $w(T) = w_1 \dots w_n$ be the word with $w_{T(i,j)} = i$ for each box (i, j) in T . Equivalently, the numerical entries in the i^{th} row of T tell us in which positions to put the symbol i in the word $w(T)$. Importantly the words formed by tableau of shape λ have evaluation $w(T) = \lambda$ (possibly ignoring some zeroes).

The map w respects the action of S_n , i.e. for any $\sigma \in S_n$, we have $\sigma w(T) = w(\sigma T)$. Also for any $w \in M^n$ with $\text{eval}(w) = \lambda$, there exists a tableau $T \in \text{YT}(\lambda)$ such that $w(T) = w$. Thus, the map is surjective on words of the correct evaluation.

Example 2.2.28. Let $\lambda = (3, 2)$, and $T = \begin{array}{|c|c|c|} \hline 5 & 4 & 1 \\ \hline 2 & 3 & \\ \hline \end{array}$. The word in W^n corresponding to T is $w(T) = 12211$. Now take the element $(123) \in S_n$, then the action on tableaux (equivalently words) gives us: $(123)w(T) = 21211 = w((123)T)$. Note that the map w is not injective, for example

$$w\left(\begin{array}{|c|c|c|} \hline 5 & 4 & 1 \\ \hline 2 & 3 & \\ \hline \end{array}\right) = 12211 = w\left(\begin{array}{|c|c|c|} \hline 1 & 4 & 5 \\ \hline 3 & 2 & \\ \hline \end{array}\right).$$

To link Definitions 2.2.9 and 2.2.26 we need to find a bijection from tabloids of shape λ to words of evaluation λ , to do this we observe that two tableaux have the same corresponding word if and only if they belong to the same class of tabloids.

Lemma 2.2.29. *Let T_1, T_2 be two tableaux of shape λ . Then*

$$w(T_1) = w(T_2) \Leftrightarrow T_1 \sim_R T_2 \Leftrightarrow \{T_1\} = \{T_2\}.$$

Proof. This is clear from the description of $w(T)$. □

Following from this we can see that the map $w(T)$ on all tableaux induces a bijection between tabloids and words. Thus we can quickly establish our permutation modules as vector spaces over words.

Lemma 2.2.30. *Let $\lambda \vdash n$, the permutation module M^λ for \mathfrak{S}_n may be seen as a vector space over the following bases:*

$$\begin{aligned} M^\lambda &= \langle \{T\} \mid T \text{ is a tableau of shape } \lambda \rangle \\ &\cong \langle w(T) \mid T \text{ is a tableau of shape } \lambda \rangle \\ &\cong \langle w \in M^n \mid \text{eval}(w) = \lambda \rangle \end{aligned}$$

Proof. The map $w(T)$ forms a bijection between tabloids of shape λ and words of evaluation λ , which respects the action of \mathfrak{S}_n . □

Thus to find elements of the module M^λ we now only have to think of words with a fixed evaluation. The regular module $M^{(1^n)}$ may now be viewed as the module of words of length n which contain every letter of $[n]$, e.g.

$$M^{(1^3)} \cong \langle 123, 132, 213, 231, 312, 321 \rangle. \tag{2.4}$$

We may also establish a basis for the Specht module S^λ using words, by replacing each tabloid by its corresponding word. E.g. rewriting $S^{(3,1)}$ from Example 2.2.13,

$$S^{(3,1)} \cong \langle 1112 - 2111, 1121 - 2111, 1211 - 2111 \rangle. \tag{2.5}$$

2.3 Random Walks on The Symmetric Group

We now use the techniques introduced in Chapter 1 to study random walks on the symmetric group, otherwise known as *shuffles*. We explore three different shuffles which have

inspired the analysis of the one-sided transposition shuffle. These illustrate both the algebraic and probabilistic techniques needed to prove cutoff results on mixing time. We begin with the classic random transposition shuffle, where we present Diaconis and Shahshahani’s argument using the upper bound lemma (Lemma 1.2.36). Next we study the top-to-random shuffle where we use strong stationary times to get an effective bound on our mixing time. Finally we end the section with a discussion of the random-to-random shuffle, for which the existence of a cutoff was recently proven. The random-to-random shuffle is our first look at the technique of lifting eigenvectors. We give explicit examples of this procedure but leave the full algebraic details for our analysis of the one-sided transposition shuffle.

2.3.1 The Random Transposition Shuffle

Our first significant example is the random transposition shuffle. This was first analysed by Diaconis and Shahshahani in 1981 [18]. In this seminal paper they proved the upper bound lemma and used it to find tight bounds on the mixing time of the random transposition shuffle. Since this important work there has been much research into the behaviour of the random transposition shuffle. Berestycki has studied the hyperbolic geometry of the random transposition shuffle when formulated as a random walk on the Cayley graph of S_n generated by transpositions [4]. Furthermore, Berestycki and Durrett showed that under a continuous time random transposition shuffle the number of transposition required to return to the identity undergoes a phase transition around time $n/2$ [6]. A natural extension of the random transposition shuffle is the random k -cycles shuffle where at each step we of the shuffle we apply a uniformly chosen k -cycle. In 2011, Berestycki, Schramm and Zeitouni proved that the random k -cycles shuffle has mixing time $(n/k) \log n$ [7].

Lately there has been an effort to study the cutoff (or limit) profile of random walks on S_n , that is the exact behaviour of $\|P_n^{t_n+cn} - \pi_n\|_{\text{TV}}$ as $n \rightarrow \infty$ where P_n exhibits a cutoff at time t_n . Naturally the first cutoff profile we want to understand is that of the random transposition shuffle and in 2016 Berestycki posed this exact question at an AIM workshop [5]. This conjecture was recently settled in a breakthrough paper by Teyssier where he proved that

$$\lim_{n \rightarrow \infty} \|\text{RT}_n^{(n/2) \log n + cn} - \pi\|_{\text{TV}} = \|\text{Po}(1 + e^{-2c}) - \text{Po}(1)\|_{\text{TV}}$$

where $\text{Po}(r)$ represents the Poisson distribution with rate r [42]. In order to prove the limit profile of the random transposition shuffle Teyssier derived a improvement of Diaconis’ upper bound lemma (Lemma 1.2.35). This improvement of the upper bound lemma has opened the path to study limit profiles of many other random walks, including the previously mentioned random k -cycles shuffle [21, 35]. This collection of work demonstrates how important the random transposition shuffle is to the study of random walks on groups and how it is still influencing the field even today.

Throughout this section we use information about the irreducible representations of S_n which we proved in Section 2.2. The argument we present here is given by Diaconis in his book *Group Representations in Probability and Statistics* [16, Chapter 3D, Theorem 5].

The random transposition shuffle is a random walk on S_n given by the following procedure: at time t , choose position i uniformly at random with your left hand, and independently choose position j uniformly at random with your right hand, then swap the cards at the positions. Swapping the cards at positions i and j amounts to applying the transposition (ij) .

Definition 2.3.1. The *random transposition shuffle* is the random walk on S_n generated by the following probability distribution:

$$\text{RT}_n(\sigma) = \begin{cases} 1/n & \text{if } \sigma = e \\ 2/n^2 & \text{if } \sigma = (ij) \text{ with } i < j \\ 0 & \text{otherwise} \end{cases}$$

We now state the cutoff result for the random transposition shuffle which we work towards throughout this section.

Theorem 2.3.2. Let $t_n = (n/2) \log n$, and $w_n = n$. The random transposition shuffle RT_n satisfies the following bounds for any $c > 0$:

$$\limsup_{n \rightarrow \infty} \|\text{RT}_n^{t_n + cw_n} - \pi_n\|_{\text{TV}} \leq Ae^{-2c} \text{ for a universal constant } A \quad (2.6)$$

$$\liminf_{n \rightarrow \infty} \|\text{RT}_n^{t_n - cw_n} - \pi_n\|_{\text{TV}} \geq \frac{1}{e} - e^{-e^{2c}} \quad (2.7)$$

Thus, the random transposition shuffle exhibits a cutoff in total variation distance at time $(n/2) \log n$ with a window of size n .

Upper Bound

We begin by establishing the upper bound present in Theorem 2.3.2. The random transposition shuffle is constant on the conjugacy classes of S_n , thus making it amenable to analysis via discrete Fourier transforms. We have seen that the simple modules of S_n are labelled by partitions of n . Let $\lambda \vdash n$ be a partition, and let χ_λ denote the character of the corresponding irreducible representation.

The random transposition shuffle only takes non-zero values of the conjugacy class of the identity and of transpositions. The conjugacy class of the identity has size 1, and the class of transpositions has size $\binom{n}{2} = (n(n-1))/2$. Applying Lemma 1.2.34 we find the Fourier transform of the random transposition shuffle at the irreducible representation λ to be,

$$\widehat{\text{RT}_n}(\lambda) = \left(\sum_i |C_i| \text{RT}_n(g_i) \frac{\chi_\lambda(g_i)}{d_\lambda} \right) \cdot \text{Id} = \left(\frac{1}{n} + \frac{n-1}{n} \frac{\chi_\lambda(\tau)}{d_\lambda} \right) \cdot \text{Id}$$

where the first sum is over the conjugacy classes of S_n with g_i a class representative, and τ in the second equation is any transposition. A straightforward application of the Upper

Bound Lemma (Lemma 1.2.35) now gives us

$$4\|\mathrm{RT}_n^t - \pi_n\|_{\mathrm{TV}}^2 \leq \sum_{\substack{\lambda \vdash n \\ \lambda \neq (n)}} d_\lambda \mathrm{Tr} \left(\widehat{\mathrm{RT}}_n(\lambda)^t \overline{\widehat{\mathrm{RT}}_n(\lambda)^t} \right) = \sum_{\substack{\lambda \vdash n \\ \lambda \neq (n)}} d_\lambda^2 \left(\frac{1}{n} + \frac{n-1}{n} \frac{\chi_\lambda(\tau)}{d_\lambda} \right)^{2t}. \quad (2.8)$$

To bound this sum we need to understand the irreducible characters of S_n , particularly their values on the conjugacy class of transpositions. Below we give a combinatorial formula for the value of $\chi_\lambda(\tau)/d_\lambda$ in terms of the Young diagram λ , see [16, Chapter 3D, Fact 2] for details.

Lemma 2.3.3. *Let $\lambda \vdash n$, and τ be any transposition. Then*

$$\frac{n(n-1)}{2} \frac{\chi_\lambda(\tau)}{d_\lambda} = \mathrm{Diag}(\lambda)$$

where $\mathrm{Diag}(\lambda) = \sum_{(i,j) \in \lambda} (j-i)$.

Using this equality we may rewrite the sum (2.8) as

$$\sum_{\substack{\lambda \vdash n \\ \lambda \neq (n)}} d_\lambda^2 \left(\frac{1}{n} + \frac{n-1}{n} \frac{\chi_\lambda(\tau)}{d_\lambda} \right)^{2t} = \sum_{\substack{\lambda \vdash n \\ \lambda \neq (n)}} d_\lambda^2 \left(\frac{n + 2\mathrm{Diag}(\lambda)}{n^2} \right)^{2t}. \quad (2.9)$$

The combinatorics of Young diagrams allows us to find and prove bounds on the terms present in (2.9). We now state several facts which have intuitive proofs using our knowledge of Young diagrams.

Lemma 2.3.4. *Let λ, μ be partitions of n , such that $\lambda \supseteq \mu$. Then we have:*

1. $\mathrm{Diag}(\lambda) \geq \mathrm{Diag}(\mu)$.
2. $\mathrm{Diag}(\lambda') = -\mathrm{Diag}(\lambda)$.
3. *The following bound holds,*

$$2\mathrm{Diag}(\lambda) \leq \begin{cases} (n-1)n - 2(n-\lambda_1)(\lambda_1+1) & \text{if } \lambda_1 \geq \frac{n}{2} \\ (\lambda_1-1)n & \text{for all } \lambda \end{cases} \quad (2.10)$$

Proof. For our first assertion suppose that μ is one step below λ , that is we may form μ by moving one box of λ down and to the left. When we move this box its diagonal index (column - row) must decrease, therefore $\mathrm{Diag}(\lambda) \geq \mathrm{Diag}(\mu)$. If μ is more than one step down from λ we may apply this result inductively to find $\mathrm{Diag}(\lambda) \geq \mathrm{Diag}(\mu)$. The second assertion follows from the construction of the transpose $(i,j) \in \lambda' \Leftrightarrow (j,i) \in \lambda$. For the last property take any $\lambda \vdash n$, in general we know that the first row of λ gives the biggest contribution to its diagonal index sum. Therefore, by taking this value and multiplying by the number of rows in λ we may form the following bound

$$\mathrm{Diag}(\lambda) = \sum_{(i,j) \in \lambda} (j-i) \leq \frac{n}{\lambda_1} \sum_{j=1}^{\lambda_1} (j-1) = \frac{(\lambda_1-1)n}{2}.$$

Now suppose $\lambda_1 \geq n/2$, then we may see that our partition is dominated by $(\lambda_1, n - \lambda_1) \triangleright \lambda$, by fixing the first row and move all other boxes up and to the right to join the second row. Therefore, by property 2 we have:

$$\text{Diag}(\lambda) \leq \text{Diag}(\lambda_1, n - \lambda_1) = \sum_{j=1}^{\lambda_1} (j-1) + \sum_{j=1}^{n-\lambda_1} (j-2) = \frac{n(n-1)}{2} - (n - \lambda_1)(\lambda_1 + 1). \quad \square$$

To simplify our upper bound we also need a bound on the dimension sum of partitions with fixed first row λ_1 . Details of the following result may be found in [16, Chapter 3D].

Lemma 2.3.5. *Let $\lambda \vdash n$, then*

$$\sum_{\substack{\lambda \vdash n \\ \lambda_1 = n-k}} d_{\lambda}^2 \leq \binom{n}{k}^2 \sum_{\substack{\lambda \vdash n \\ \lambda_1 = n-k}} d_{\lambda/\lambda_1}^2 \leq \binom{n}{k}^2 k! \quad (2.11)$$

Using the results we have just proved we may simplify the summation (2.9). We first double up our sum by only concentrating on those partitions which have positive diagonal sum. Note that we have to deal with the partition (1^n) as a special case because its transpose $(1^n)' = (n)$ is not presented in the sum (2.9). We then parametrise our sum based on the value of λ_1 splitting our sum around the value of $\lambda_1 \leq 3n/4$. Lemma 2.3.4 helps to bound the terms with given λ_1 as follows:

$$\frac{n + 2\text{Diag}(\lambda)}{n^2} \leq \begin{cases} 1 - \frac{2(\lambda_1+1)(n-\lambda_1)}{n^2} & \text{if } \lambda_1 \geq 3n/4 \\ \frac{\lambda_1}{n} & \text{for all } \lambda \end{cases}. \quad (2.12)$$

Subsequently Lemma 2.3.5 may be used to bound the multiplicities of the new terms. Performing these steps gives us,

$$\begin{aligned} (2.9) &= \sum_{\substack{\lambda \vdash n \\ \lambda \neq (n) \\ \text{Diag}(\lambda) \geq 0}} d_{\lambda}^2 \left(\frac{n + 2\text{Diag}(\lambda)}{n^2} \right)^{2t} + \sum_{\substack{\lambda \vdash n \\ \lambda \neq (1^n) \\ \text{Diag}(\lambda) < 0}} d_{\lambda}^2 \left(\frac{n + 2\text{Diag}(\lambda)}{n^2} \right)^{2t} + \left(\frac{n + 2\text{Diag}((1^n))}{n^2} \right)^{2t} \\ &\leq \sum_{\substack{\lambda \vdash n \\ \lambda \neq (n) \\ \text{Diag}(\lambda) \geq 0}} d_{\lambda}^2 \left(\frac{n + 2\text{Diag}(\lambda)}{n^2} \right)^{2t} + \sum_{\substack{\lambda \vdash n \\ \lambda \neq (1^n) \\ \text{Diag}(\lambda) < 0}} d_{\lambda'}^2 \left(\frac{n + 2\text{Diag}(\lambda')}{n^2} \right)^{2t} + \left(\frac{n - n(n-1)}{n^2} \right)^{2t} \\ &= 2 \sum_{\substack{\lambda \vdash n \\ \lambda \neq (n) \\ \text{Diag}(\lambda) \geq 0}} d_{\lambda}^2 \left(\frac{n + 2\text{Diag}(\lambda)}{n^2} \right)^{2t} + \left(1 - \frac{2}{n} \right)^{2t} \quad (2.13) \end{aligned}$$

$$\begin{aligned} &\leq 2 \sum_{k=1}^{n/4} \left(1 - \frac{2(n-k+1)k}{n^2} \right)^{2t} \sum_{\substack{\lambda \vdash n \\ \lambda_1 = n-k}} d_{\lambda}^2 + 2 \sum_{k > n/4}^{n-2} \left(1 - \frac{k}{n} \right)^{2t} \sum_{\substack{\lambda \vdash n \\ \lambda_1 = n-k}} d_{\lambda}^2 + \left(1 - \frac{2}{n} \right)^{2t} \\ &\leq 2 \sum_{k=1}^{n/4} \left(1 - \frac{2(n-k+1)k}{n^2} \right)^{2t} \binom{n}{k}^2 k! + 2 \sum_{k > n/4}^{n-2} \left(1 - \frac{k}{n} \right)^{2t} \binom{n}{k}^2 k! + \left(1 - \frac{2}{n} \right)^{2t} \quad (2.14) \end{aligned}$$

After our simplifications we are left with summations in n, k , and t . Finally we are in a position to analyse the behaviour of these sums around time $t = (n/2) \log n + cn$. The

singular term given by the partition (1^n) may be easily dealt with at this time,

$$\limsup_{n \rightarrow \infty} \left(\frac{n - n(n-1)}{n^2} \right)^{2t} = \limsup_{n \rightarrow \infty} \left(1 - \frac{2}{n} \right)^{n \log n + 2cn} = 0.$$

The first summation of (2.14) contains the largest values of the Fourier transforms $\widehat{\text{RT}}_n(\lambda)$. These values control the mixing time of the random transposition shuffle. The first term, $k = 1$, in the summation is given by

$$n^2 \left(1 - \frac{2}{n} \right)^{2t} \tag{2.15}$$

this term is tightly bounded by 1 at time $t = (n/2) \log n$ as $n \rightarrow \infty$. If we add an additional window of cn to give time $t = (n/2) \log n + cn$ we find the first term bounded by an exponential decay,

$$n^2 \left(1 - \frac{2}{n} \right)^{2((n/2) \log n + cn)} \leq n^2 e^{-2 \log n - 4c} \leq e^{-4c}.$$

We may show that the ratio of subsequent terms in the first sum is less than 1 for $n \geq 17$. This allows us to bound the first summation via a geometric series and conclude that it may be bounded by Ae^{-4c} for a universal constant A at the time $(n/2) \log n + cn$ for n sufficiently large. Now we look to bound the second summation in (2.14) at time $(n/2) \log n + cn$. The terms in the second summation can be shown to be decreasing in k for n sufficiently large and thus we may bound the sum by $3n/4$ times the first term. Using Stirling's formula the first term can be seen to be tending to 0 at time $(n/2) \log n + cn$ as $n \rightarrow \infty$. Further details of this analysis can be found in [16]. Putting these bounds together we are now in a position to prove our upper bound on mixing time for the random transposition shuffle.

Proof of the upper bound in Theorem 2.3.2. We reduced the analysis of total variation distance to the bound of summation (2.14). Taking the limit as $n \rightarrow \infty$ we may establish the following bound on total variation distance at time $(n/2) \log n + cn$,

$$\begin{aligned} \limsup_{n \rightarrow \infty} 4 \|\text{RT}_n^{t_n + cw_n} - \pi_n\|_{\text{TV}}^2 &\leq \limsup_{n \rightarrow \infty} \sum_{\substack{\lambda \vdash n \\ \lambda \neq (n)}} d_\lambda^2 \left(\frac{1}{n} + \frac{n-1}{n} \frac{\chi_\lambda(\tau)}{d_\lambda} \right)^{2(t_n + cw_n)} \\ &\leq Ae^{-4c} \text{ for a universal constant } A. \quad \square \end{aligned}$$

Lower Bound

We now prove the lower bound of Theorem 2.3.2. Let $F_n = \{\sigma \in S_n \mid \sigma \text{ has at least one fixed point}\}$. Using the definition of total variation distance we may find

$$\|\text{RT}_n - \pi_n\|_{\text{TV}} \geq |\text{RT}_n(F_n) - \pi_n(F_n)|.$$

It is a classical result [20] that as $n \rightarrow \infty$ the chance that a permutation picked uniformly at random has no fixed points tends to $1/e$. Therefore,

$$\lim_{n \rightarrow \infty} \pi_n(F_n) = 1 - \frac{1}{e}.$$

Next we show that under the random transposition measure we have a high probability of being in set F_n at time $(n/2)\log n - cn$. Consider the set of random variables $\{R^i, L^i\}_{i \in \mathbb{N}}$ which track which positions our hands have picked at time i in the random transposition shuffle. Clearly if the random variables R^i, L^i , have not chosen every position by time t then the permutation we find ourselves at after t steps of random transposition shuffle must have at least one fixed point. Let J_n^t be the event $\{\cup_{i=1}^t \{R^i, L^i\} \not\subseteq [n]\}$, then by the reasoning above we have $J_n^t \subseteq F_n$, and hence $\mathbb{P}(J_n^t) \leq \text{RT}_n^t(F_n)$. The probability $\mathbb{P}(J_n^t)$ may be computed using a coupon collector's argument. In the uniform coupon collector's problem, the expected time to collect all coupons is $n \log n$. However, we are performing two independent trials at each step, leading our expected time to choose every card to quicken to $(n/2)\log n$. More careful arguments in Feller [20, Section IV.2] show that,

$$\mathbb{P}(J_n^t) = 1 - e^{-ne^{-2t/n}} + o(1), \text{ as } n \rightarrow \infty.$$

Proof of the lower bound in Theorem 2.3.2. Let $t_n = (n/2)\log n$ and $w_n = n$. Putting together our bounds on the probability of F_n under our different distributions we may complete our lower bound on total variation distance as follows,

$$\liminf_{n \rightarrow \infty} \|\text{RT}_n^{t_n - cw_n} - \pi_n\|_{\text{TV}} \geq \liminf_{n \rightarrow \infty} \text{RT}_n^{t_n - cw_n}(F_n) - \pi(F_n) \geq \frac{1}{e} - e^{-e^{2c}}.$$

This completes the proof of a total variation cutoff for the random transposition shuffle. \square

The analysis of the random transposition shuffle is exemplary in a few important ways. We always look to reduce the upper bound on total variation distance to a collection of summations like those in (2.14). At this point we may perform analysis of the summations to find a value of t for which they are bounded. The mixing time of random walks on groups are thought to be controlled by their second biggest eigenvalue. The first sum (2.14) contains the largest eigenvalues of the random transposition shuffle, including the second biggest eigenvalue (2.15), and is closely bounded at time $(n/2)\log n$, whereas the second summation disappears at time $(n/2)\log n$ as $n \rightarrow \infty$. Therefore, the first summation in (2.14) can be thought of as determining the mixing time for the random transposition shuffle.

The lower bound argument presented for the random transposition shuffle shows how versatile the simple bound in Lemma 1.1.32 can be when applied to a careful choice of set. Often we like to reduce a lower bound in total variation distance to other probabilistic problems, in this instance a coupon collector's problem. This gives us another insight into why a cutoff occurs, hard to reach sets stop the total variation decaying before they are hit, keeping it close to 1 until we pass the critical time of $(n/2)\log n$.

2.3.2 The Top-to-Random Shuffle

The top-to-random shuffle is defined by the following procedure: choose a position of the deck uniformly at random, then insert the top card into this position. The top-to-random shuffle is an example of a random walk which is not constant on the conjugacy classes of S_n . This means the technique of discrete Fourier transforms does not work as effectively

for the top-to-random shuffle, which leads us to use the method of strong stationary times instead.

Definition 2.3.6. The *top-to-random shuffle* has driving probability TTR_n defined by:

$$\text{TTR}_n(\sigma) = \begin{cases} 1/n & \text{if } \sigma = (k \ k-1 \ \dots \ 1) \text{ for some } k \in [n] \\ 0 & \text{otherwise} \end{cases}.$$

The elements $(k \ k-1 \ \dots \ 1)$ represent cycling the top k cards of the deck one space up the deck with the card in position k moving to position $k-1$, and finally the top card of the deck moving to position k .

We may clearly see that the top-to-random shuffle is not constant on the conjugacy classes of S_n , for a quick counter example notice that $\text{TTR}_n((23)) \neq \text{TTR}_n((12))$ for $n \geq 3$. In this section we prove the following total variation cutoff for the top-to-random shuffle, the arguments we present are taken from [16, Chapter 4A, Theorem 1] and [32, Section 6.5.3].

Theorem 2.3.7. *Let $t_n = n \log n$ and $w_n = n$. The top-to-random shuffle TTR_n satisfies the following bounds for any $c > 0$:*

$$\limsup_{n \rightarrow \infty} \|\text{TTR}_n^{t_n + cw_n} - \pi_n\|_{\text{TV}} \leq e^{-c} \quad (2.16)$$

$$\liminf_{n \rightarrow \infty} \|\text{TTR}_n^{t_n - cw_n} - \pi_n\|_{\text{TV}} \geq 1 - \frac{2}{e^{c-1}} \quad (2.17)$$

Thus, the top-to-random shuffle exhibits a cutoff in total variation distance at time $n \log n$ with a window of size n .

Upper Bound

We start by proving the upper bound (2.16), but for separation distance – we can then appeal to Lemma 1.1.24 to give a bound on total variation distance. Recall that our deck of cards is labelled $\{1, \dots, n\}$ from bottom to top. As mentioned above to find a bound on separation distance we use a strong stationary time for the top-to-random shuffle. To this end define T to be the first time card 1 is randomly inserted into the deck when following the top-to-random shuffle. We prove that T is a strong stationary time and then reduce the analysis of $\mathbb{P}(T > n \log n + cn)$ to the uniform coupon collector’s problem.

Lemma 2.3.8. *Let T be the first time card 1 is inserted into the deck during the top-to-random shuffle. Then T is a strong stationary time for the top-to-random shuffle.*

Proof. We start our shuffle from the identity permutation with card 1 at the bottom of the deck and card n at the top. Define a sequence of stopping times T_i , as the first time i cards have been placed under the bottom card of the deck, i.e. card 1. Notice that T_{n-1} is the first time card 1 has reached the top of our deck, define the stopping time T_n as the time when card 1 is inserted into the deck uniformly at random. Our proposed strong stationary time is $T = T_n = T_{n-1} + 1$. By induction we prove that at time T_i the bottom i cards of the deck are distributed according to a uniformly chosen permutation of

S_i . Clearly at time T_1 the first card to be inserted below card 1 is in a uniformly chosen permutation of S_1 . Assume our hypothesis to be true for time T_i , so the bottom i cards of our deck are arranged in a uniformly chosen permutation of S_i . At time T_{i+1} we have moved a new card into one of $i+1$ positions below card 1, thus forming a permutation of S_{i+1} amongst the bottom $i+1$ cards. To see that this permutation is uniformly chosen from S_{i+1} we must recognise that the new card has equal probability of being inserted into of any of the $i+1$ positions below card 1. Pairing this fact with our inductive hypothesis means that all permutations of the bottom $i+1$ cards are equiprobable. Furthermore, the arrangement of the bottom $i+1$ cards is independent of the value of T_{i+1} . At time $T = T_n$ the entire deck is distributed according to a uniformly chosen permutation of S_n , and so T is a strong stationary time. \square

Lemma 2.3.9. *The time T satisfies the following bound:*

$$\mathbb{P}(T > n \log n + cn) \leq e^{-c} \quad (2.18)$$

Proof. Define a sequence of stopping times T_i as in Lemma 2.3.8. Our first observation is that

$$T = T_n = (T_n - T_{n-1}) + (T_{n-1} - T_{n-2}) + \cdots + (T_2 - T_1) + T_1.$$

The increments $T_i - T_{i-1}$ are independent and geometrically distributed with parameter i/n . The random variable T shares the same distribution as the uniform coupon collector's problem which may be described as follows: Suppose at each time step we choose a card from our deck of n cards uniformly at random, if this is the first time we have seen this card we say we *collect* this card. The coupon collector's problem asks how long does it takes to collect every card of the deck?

Let \mathcal{T}_i be the first time we have collected i cards from the deck, clearly we have

$$\mathcal{T}_n = (\mathcal{T}_n - \mathcal{T}_{n-1}) + \cdots + (\mathcal{T}_2 - \mathcal{T}_1) + \mathcal{T}_1.$$

The increments $\mathcal{T}_i - \mathcal{T}_{i-1}$ are independent and are geometrically distributed with parameter $(n-i)/n$. Therefore, $\mathcal{T}_{n-i} - \mathcal{T}_{n-i-1} \sim_D \mathcal{T}_i - \mathcal{T}_{i-1}$ and so the random variables $\mathcal{T}_n \sim_D \mathcal{T}_n$ follow the same law. In the uniform coupon collector's problem let C_i^t be the event we have not collected the card i by time t . Using the events C_i^t we may find a simple bound on \mathcal{T}_n (and respectively T) at the time $t = n \log n + cn$,

$$\mathbb{P}(T_n > t) = \mathbb{P}(\mathcal{T}_n > t) \leq \mathbb{P}\left(\bigcup_{i=1}^n C_i^t\right) \leq \sum_{i=1}^n \mathbb{P}(C_i^t) \leq \sum_{i=1}^n \left(1 - \frac{i}{n}\right)^{n \log n + cn} \leq e^{-c}. \quad (2.19)$$

\square

Proof of Upper Bound in Theorem 2.3.7. Using the strong stationary time T , together with Lemma 2.3.9, we proven the upper bound (2.16) as follows,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \|\text{TTR}_n^{n \log n + cn} - \pi\|_{\text{TV}} &\leq \limsup_{n \rightarrow \infty} \|\text{TTR}_n^{n \log n + cn} - \pi_n\|_{\text{Sep}} \\ &\leq \limsup_{n \rightarrow \infty} \mathbb{P}(T > n \log n + cn) \leq e^{-c} \quad \square \end{aligned}$$

Lower Bound

To prove the lower bound present in Theorem 2.3.7 we again appeal to Lemma 1.1.32. For $k \geq 2$, let I_k be the set of permutations our deck such that the bottom k cards remain in their original relative order, in other words $\sigma \in I_k$ if and only if $\sigma(1) < \sigma(2) < \dots < \sigma(k)$. Under the uniform distribution we have $\binom{n}{k}$ ways to arrange the bottom k cards while retaining their relative order, after which we may place the remaining $n - k$ cards in any order giving us a factor of $(n - k)!$. Therefore,

$$\pi_n(I_k) = \frac{(n - k)!}{n!} \binom{n}{k} = \frac{1}{k!} \leq \frac{1}{k - 1}. \quad (2.20)$$

To bound the probability of being in set I_k under the top-to-random shuffle we make use of the random variables T_i introduced in the proof of Lemma 2.3.9. The time $T - T_{k-1}$ is the first time the card k has come to the top of our deck and been reinserted into the deck. Before this time the bottom k cards always remain in the same relative order, therefore $\text{TTR}_n^t(I_k) \geq \mathbb{P}(T - T_{k-1} > t)$. Using the decomposition of T in Lemma 2.3.9 we may rewrite the random variable into independent geometrically distributed increments $T - T_{k-1} = \sum_{i=k}^n (T_i - T_{i-1})$. The expectation and variance of the random variable $T - T_{k-1}$ is as follows:

$$\mathbb{E}[T - T_{k-1}] = \sum_{i=k}^n \frac{n}{i} \geq n(\log n - \log k) \quad (2.21)$$

$$\text{Var}[T - T_{i-1}] \leq \sum_{i=k}^{\infty} \frac{n^2}{i(i-1)} \leq \frac{n^2}{k-1}. \quad (2.22)$$

Applying Chebyshev's inequality to the random variable $T - T_{k-1}$ we find,

$$\mathbb{P}(T - T_{k-1} \leq n \log n - cn) \leq \frac{1}{k-1} \quad (2.23)$$

provided that $c \geq \log k + 1$. We are now in a position to prove the lower bound on the mixing time of the top-to-random shuffle.

Proof of Lower Bound in Theorem 2.3.7. Let I_k be the set defined above and set $k - 1 = \lceil e^{c-1} \rceil$. Provided $n \geq e^{c-1} + 1$ we may use our bounds (2.20), (2.23), to form the following lower bound on total variation distance,

$$\|\text{TTR}_n^{t_n - cw_n} - \pi_n\|_{\text{TV}} \geq \text{TTR}_n^{t_n - cw_n}(I_k) - \pi_n(I_k) \geq 1 - \frac{2}{e^{c-1}}.$$

Taking the \liminf as $n \rightarrow \infty$ gives the desired result. \square

The top-to-random shuffle showcases the advantages of strong stationary times. In this case analysis via discrete Fourier transforms would be difficult owing to its support being a mix of elements from different conjugacy classes. Instead a strong stationary time leads to a simple and effective upper bound on its mixing time, reducing the analysis to the well studied uniform coupon collector's problem. The strong stationary time T was the limit of a sequence of stopping times, which build up a uniform set from S_1 to S_n . In Chapter 3 we present a novel strong stationary time for the one-sided transposition shuffle which

does not rely on building up a sequence of uniform subgroups. The lower bound for the top-to-random shuffle further showcases the usefulness of Lemma 1.1.32 in reducing the analysis of total variation distance to discrete random variables.

2.3.3 The Random-to-Random Shuffle

The last shuffle we study in this chapter is the random-to-random shuffle. The random-to-random shuffle is described by the following procedure: at each step pick a card uniformly at random and take it from the deck, then pick a position uniformly at random and insert the card back into the deck at this position. This shuffle has a long history in the literature of mixing times of random walks. First introduced by Diaconis and Saloff-Coste [17] in 1993, together they proved its mixing time was of order $n \log n$, and conjectured that a total variation cutoff occurs at time $(3/4)n \log n$. This bound was improved by Uyemura-Reyes [43] who showed the mixing time to be in the range $[(n/2) \log n, 4n \log n]$. In 2012, Subag [41] was able to prove a lower bound of $(3/4)n \log n - (1/4)n \log \log n$ on the mixing time and in 2017, Bernstein and Nestoridi [10] proved a matching upper bound of $(3/4)n \log n + (1/4)n \log \log n$ on the mixing time. Together these results prove the existence of a total variation cutoff for the random-to-random shuffle.

The recent work of Bernstein and Nestoridi [10] relied on knowledge of the eigenvalues for the random-to-random shuffle. These were computed by Dieker and Saliola [19], in a breakthrough publication where they invented the technique of *lifting eigenvectors and eigenvalues*. We dedicate this subsection to an exploration of the technique of lifting eigenvectors for the the random-to-random shuffle. This technique uses the branching structure of the symmetric group to turn eigenvalues of the random-to-random shuffle on n cards into eigenvalues of the random-to-random shuffle on $n + 1$ cards.

To describe the random-to-random shuffle we introduce the concept of *symmetrizing* a random walk. Consider a Markov chain $\{X^t\}_{t \geq 0}$ on S_n driven by a probability P . By viewing our Markov chain backwards in time we may create a new a Markov chain $\{Y^t\}_{t \geq 0}$ driven by probability $P^{-1}(\sigma) := P(\sigma^{-1})$, this is called the *time reversal process*. Note that for a reversible random walk on a group driven by P we necessarily have that $P^{-1} = P$, this follows straight from the detailed balance equations (1.1). From any shuffle P we may create a reversible random walk called the *symmetrization* of P , by taking the convolution of P with its time reversal process, i.e., $P^{-1} \star P$.

Consider the top-to-random shuffle, its time reversal process does the following: it picks a card uniformly at random and places it back on top of the deck, the shuffle described by this process is called the *random-to-top shuffle*, denoted RTT_n . Imagine that we perform one random-to-top shuffle and then one top-to-random shuffle, the two steps together define a process whereby we pick a card uniformly at random and insert it back into the deck uniformly at random, we know this as the random-to-random shuffle.

Definition 2.3.10. The *random-to-random shuffle* on S_n picks a card uniformly at random and places it back into the deck uniformly at random. It is defined by the following

probability distribution:

$$\text{RTR}_n(\sigma) = \begin{cases} 1/n & \text{if } \sigma = e \\ 1/n^2 & \text{if } \sigma = (i \ i+1 \ \dots \ j-1 \ j) \text{ for } 1 \leq i < j \leq n \\ 1/n^2 & \text{if } \sigma = (i \ i-1 \ \dots \ j+1 \ j) \text{ for } 1 \leq j < i \leq n \\ 0 & \text{otherwise} \end{cases}. \quad (2.24)$$

The first non-identity permutations above correspond to taking a card from position j and moving it down the deck to position i . The second non-identity permutation correspond to taking a card from position j and moving it up the deck to position i . We may also define the random-to-random shuffle as the *symmetrization* of the top-to-random (or equivalently random-to-top) shuffle, $\text{RTR}_n = \text{TTR}_n \star \text{RTT}_n$. The random-to-random shuffle defines an aperiodic, transitive, reversible Markov chain.

The random-to-random shuffle is not constant on the conjugacy classes of S_n , again this may see by considering two transpositions, $\text{RTR}_n((12)) \neq \text{RTR}_n((1n))$ for all $n \geq 3$. This makes the discrete Fourier transforms of probability RTR_n difficult to compute, leading us to require a different method to study this shuffle. The random-to-random shuffle is reversible, which means if we can compute its eigenvalues we may apply the classical ℓ^2 bound given in Theorem 1.1.27 to find an upper bound on the total variation distance of the random-to-random shuffle and uniform distribution. Dieker and Saliola provide an interesting combinatorial description (involving Young diagrams) of all the eigenvalues for the random-to-random shuffle which follows directly from the procedure of lifting eigenvectors.

Theorem 2.3.11 (Theorem 5 [19]). *The eigenvalues of the random-to-random shuffle on n cards are indexed by pairs of partitions (λ, μ) with $\lambda \vdash n$ and λ/μ a horizontal strip. The eigenvalue $\text{eig}(\lambda, \mu)$ corresponding to (λ, μ) occurs with multiplicity $d_\lambda d^\mu$ and has value,*

$$\text{eig}(\lambda, \mu) = \frac{1}{n^2} \left(\binom{|\lambda|+1}{2} - \binom{|\mu|+1}{2} + \text{Diag}(\lambda) - \text{Diag}(\mu) \right).$$

The combinatorial description of the eigenvalues of the random-to-random shuffle hides away the algebraic techniques used to find them. The key idea behind Deiker and Saliola's work is to exploit the module structure of S_n in order to create new eigenvectors from known ones. The first step is to change from studying the random-to-random shuffle as a probability to instead studying it as an element of the group algebra \mathfrak{S}_n .

Definition 2.3.12. The random-to-random shuffle on n cards can be viewed as the following element of \mathfrak{S}_n :

$$\begin{aligned} \text{ARTR}_n &:= n^2 \sum_{\sigma \in S_n} \text{RTR}_n(\sigma) \sigma \\ &= ne + \sum_{1 \leq i < j \leq n} (i \ i+1 \ \dots \ j-1 \ j) + \sum_{1 \leq j < i \leq n} (i \ i-1 \ \dots \ j+1 \ j) \end{aligned} \quad (2.25)$$

We call this element the *algebraic random-to-random shuffle*. Note that to create ARTR_n we have scaled our probability by a factor of n^2 .

This new viewpoint allows the random-to-random shuffle to act on modules of S_n via the element ARTR_n . In particular we are interested in the action of ARTR_n on the regular module of \mathfrak{S}_n , because the eigenvectors and eigenvalues of this action are exactly the eigenvectors and eigenvalues for the probability RTR_n .

Lemma 2.3.13. *A probability distribution $\nu : S_n \rightarrow [0, 1]$ is an eigenvector for RTR_n with eigenvalue ϵ , if and only if $v = \sum_{\eta \in S_n} \nu(\eta) \eta \in \mathfrak{S}_n$ is an eigenvector for ARTR_n with eigenvalue $n^2 \epsilon$.*

Proof. Suppose ν is an eigenvector for RTR_n with eigenvalue ϵ , this means that $P \star \nu = \epsilon \nu$. Therefore, $\epsilon \nu(\sigma) = \sum_{\eta} \text{RTR}_n(\sigma \eta^{-1}) \nu(\eta)$. Now consider the action

$$\begin{aligned} \text{ARTR}_n \left(\sum_{\eta \in S_n} \nu(\eta) \eta \right) &= n^2 \sum_{\sigma \in S_n} \sum_{\eta \in S_n} \text{RTR}_n(\sigma) \nu(\eta) \sigma \eta \\ &= n^2 \sum_{\tau \eta^{-1} \in S_n} \sum_{\eta \in S_n} \text{RTR}_n(\tau \eta^{-1}) \nu(\eta) \tau \\ &= n^2 \sum_{\tau \eta^{-1} \in S_n} \epsilon \nu(\tau) \tau = n^2 \epsilon \left(\sum_{\eta \in S_n} \nu(\eta) \eta \right). \end{aligned}$$

In the second equality we have relabelled $\sigma \eta = \tau$, and the last equality we relabel the sum to match our original vector. \square

Following Lemma 2.3.13 we may focus on the action of ARTR_n on the regular module with the understanding that all of the eigenvectors and eigenvalues of the random-to-random shuffle may be recovered. Recall that the regular module has the following decomposition

$$\mathfrak{S}_n \cong \bigoplus_{\lambda \vdash n} d_\lambda S^\lambda \text{ as } \mathfrak{S}_n\text{-modules.} \quad (2.26)$$

Note that because ARTR_n is an element of our group algebra it acts on the module \mathfrak{S}_n and stabilises the submodules S^λ . Therefore, the decomposition (2.26) means we can reduce the problem of finding eigenvectors for the random-to-random shuffle on \mathfrak{S}_n to the problem of finding eigenvectors belonging to the individual Specht modules S^λ . Moreover, since the shuffle is acting as an element of the group algebra, we are then free to study its action on the single copy of S^λ inside M^λ to solve this problem rather than having to stick to the copies of S^λ which appear inside \mathfrak{S}_n . In work that follows we present the eigenvectors of ARTR_n as elements of the Specht modules comprised of words as described in Section 2.2.6.

The Specht modules of \mathfrak{S}_n and \mathfrak{S}_{n+1} are closely related by the branching structure in Theorem 2.2.25, and by Young's lattice shown in Figure 2.1. The random-to-random shuffle also exhibits a recursive structure which can be seen by taking the difference of the shuffle on $n+1$ cards and the shuffle on n -cards,

$$\text{ARTR}_{n+1} - \text{ARTR}_n = e + \sum_{\substack{1 \leq i \leq n \\ j = n+1}} (i \ i+1 \ \dots \ n \ n+1) + \sum_{\substack{1 \leq j \leq n \\ i = n+1}} (n+1 \ n \ \dots \ j+1 \ j).$$

The permutations that do not involve the new card in position $n + 1$ have disappeared in our comparison of the two elements. The key idea is that moving the card in position j to position $i < j$ does not affect any of the cards above j , thus it does not matter how many cards above j our deck contains. This recursive structure is at the heart of why lifting eigenvectors works and we recover a similar recursive structure for the one-sided transposition shuffle (see equation (3.6)). The technique of lifting eigenvectors uses the branching structure for Specht modules and the recursive structure of the random-to-random shuffle, to turn eigenvectors of ARTR_n into eigenvectors of ARTR_{n+1} . In particular we lift eigenvectors belonging to the Specht modules S^μ into eigenvectors belonging to Specht modules S^λ where $\mu \subset \lambda$, i.e., for λ we can create from μ by adding boxes.

The following theorem summarises the main result of lifting eigenvectors for the random-to-random shuffle.

Theorem 2.3.14 (Theorem 21 [19]). *Let $\lambda \vdash n$, and $1 \leq i \leq l(\lambda) + 1$ be such that $\lambda + e_i \vdash n + 1$. Then there exists a linear map $\mathcal{L}_i^\lambda : S^\lambda \rightarrow S^{\lambda+e_i}$, such that, \mathcal{L}_i^λ maps eigenvectors of the random-to-random shuffle on n cards belonging to S^λ , to eigenvectors of the random-to-random shuffle on $n + 1$ cards belonging to $S^{\lambda+e_i}$. In particular, if $v \in S^\lambda$ is an eigenvector of ARTR_n with eigenvalue ϵ , then $\mathcal{L}_i^\lambda(v)$, if **non-zero**, is an eigenvector of ARTR_{n+1} with eigenvalue $\epsilon + (n + 1) + (\lambda_i + 1) - i$.*

Using Theorem 2.3.14 we may find eigenvectors of a Specht module S^λ from by lifting eigenvectors from below, however we must be careful that the map \mathcal{L}_i^λ we apply does not kill the eigenvectors by mapping them to 0. The non-injectivity of the lifting operators stops all the eigenvectors belonging to a Specht module S^λ being described by lifting. However, Deiker and Saliola managed to recover all the eigenvalues belonging to S^λ by proving that the eigenvectors we fail to find by lifting belong to the kernel of ARTR_{n+1} and thus have eigenvalue 0. This was taken into account when formulating Theorem 2.3.11. The lifting operators \mathcal{L}_i^λ may be fully described in terms of linear operators on the space of words M^n (see Section 2.2.6), which we now introduce.

Definition 2.3.15. Define two linear operators on the space M^n spanned by words, to do so, it is enough to define the effect on any given word. Let $w = w_1 \dots w_n$ in W^n .

1. Let $a \in [n + 1]$. Define the *shuffling operator*, denoted $\text{sh}_i : M^n \rightarrow M^{n+1}$ as the following linear map

$$\text{sh}_a(w) = \sum_{j=0}^{n+1} w_1 \dots w_j a w_{j+1} \dots w_n.$$

The operator sh_a is a linear combination of all words that are formed by inserting a into one position. The shuffling operator is a linear map from M^λ (or S^λ) to $M^{\lambda+e_a}$.

2. Let $a, b \in [n]$. Define the *switching operator*, denoted $\Theta_{b,a} : M^n \rightarrow M^n$ as follows:

$$\Theta_{b,a}(w) := \sum_{\substack{1 \leq k \leq n \\ w_k = b}} w_1 \dots w_{k-1} \cdot a \cdot w_{k+1} \dots w_n.$$

The operator $\Theta_{b,a}$ forms a linear sum of all words created from w by replacing one

occurrence of b with a . If we restrict to domain M^λ (or S^λ) then we end in the module M^μ with $\lambda + e_a = \mu + e_b$.

Remark 2.3.16. We should define the shuffling and switching operators for all possible domains M^λ , however to do so would burden us with more notation, when it will always be clear from context which domain and codomain we are considering.

Example 2.3.17. Let $w = 1231 \in M^4$. Then $\text{sh}_1(w) = 11231 + 11231 + 12131 + 12311 + 12311 = 2 \cdot 11231 + 12131 + 2 \cdot 12311$, this is an element belongs to the vector space M^5 . We may see that $1231 \in M^{(2,1,1)} \subseteq M^4$ and $\text{sh}_1(1231) \in M^{(3,1,1)} \subseteq M^5$. An example of our switching operator is $\Theta_{3,1}(w) = 1211$, and $\Theta_{1,3}(w) = 3231 + 1233$. Our images for the switching operator are $\Theta_{3,1}(w) \in M^{(3,1)}$ and $\Theta_{1,3}(w) \in M^{(1,1,2)}$. Another characteristic of the switching operator is that for any $w \in M^n$ and $i \in [n]$ we have $\Theta_{i,i}(w) = \text{eval}_i(w) w$.

We may construct the lifting operators \mathcal{L}_i as linear combinations of the shuffling and switching operators defined in Definition 2.3.15. Moreover, we may tell exactly when there is no possible lifting from module S^μ to S^λ by analysis of the partitions μ and λ .

Lemma 2.3.18 (Theorem 21 [19]). *Let $\lambda \vdash n$. The lifting operators as defined in Theorem 2.3.14 are a linear combination of the shuffling and switching operators, explicitly they are:*

$$\mathcal{L}_i^\lambda = \sum_{1 \leq b_1 < \dots < b_m < b_{m+1} = i} \left(\prod_{j=1}^m \frac{1}{(\lambda_i - i) - (\lambda_{b_j} - b_j)} \Theta_{b_j, b_{j+1}} \right) \circ \text{sh}_{b_1}. \quad (2.27)$$

For $i = 1, 2$ the above equation simplifies to:

$$\mathcal{L}_1^\lambda = \text{sh}_1 \quad (2.28)$$

$$\mathcal{L}_2^\lambda = \text{sh}_2 + \frac{1}{(\lambda_2 - 2) - (\lambda_1 - 1)} \Theta_{1,2} \circ \text{sh}_1 \quad (2.29)$$

Furthermore, let $\mu \subset \lambda$, such that $\mu + \sum_{j=1}^k e_{i_j} = \lambda$ where $i_1 \leq \dots \leq i_k$. Define the lifting operator from μ to λ , denoted $\mathcal{L}^{\lambda/\mu}$, as the map $\mathcal{L}^{\lambda/\mu} = \prod_{i=1}^k \mathcal{L}_{i_j}$, that is lift μ to λ by forming each row in order.

Theorem 2.3.19 (Theorem 26 [19]). *Suppose $\lambda \vdash n$. Every non-kernel eigenvector of ARTR_n belonging to S^λ may be found by lifting a eigenvector in the kernel of $\text{ARTR}_{|\mu|}$ belonging S^μ for some $\mu \subset \lambda$ with λ/μ a horizontal strip. In other words if $v \in \mu$ is a eigenvector in the kernel of $\text{ARTR}_{|\mu|}$, then $\mathcal{L}^{\lambda/\mu}(v) \neq 0$ if and only if λ/μ is a horizontal strip. Thus, all eigenvalues for the module S^λ may be recovered as they are the eigenvalue of a lifted eigenvector, or 0 otherwise.*

We end this chapter with explicit examples of lifting eigenvectors of the random-to-random shuffle to Specht modules of \mathfrak{S}_1 , \mathfrak{S}_2 , and \mathfrak{S}_3 , using the lifting operators $\mathcal{L}_1^\lambda, \mathcal{L}_2^\lambda$. We start with the only Specht module of \mathfrak{S}_1 which is $S^{(1)} = \langle 1 \rangle$. Define the trivial partition (0) to have a trivial eigenvector of the empty word, denoted ω , with eigenvalue 0. The only partition strictly contained in the partition (1) is the empty partition (0) and so we the only possible lifting is

$$\mathcal{L}^{(1)/(0)}(\omega) = \mathcal{L}_1^{(0)}(\omega) = 1.$$

We easily see that $\mathcal{L}^{(1)/(0)}(\omega)$ is the only eigenvector of ARTR_1 belonging to $S^{(1)}$ with eigenvalue 1.

We now compute the eigenvectors for ARTR_2 belonging to the Specht modules of \mathfrak{S}_2 . For the simple modules $S^{(2)} = \langle 11 \rangle$ and $S^{(1,1)} = \langle 12 - 21 \rangle$ we must start our lifting in the module $S^{(0)}$ because the we saw above that $S^{(1)}$ has no eigenvector in the kernel of ARTR_1 thus it can not be used for lifting. The skew diagram of $(2)/(0)$ is a horizontal strip, therefore

$$\mathcal{L}^{(2)/(0)}(\omega) = \mathcal{L}_1^{(1)}\mathcal{L}_1^{(0)}(\omega) = \mathcal{L}_1^{(1)}(1) = 2 \cdot 11$$

is an eigenvector of ARTR_2 . The skew diagram of $(1,1)/(0)$ is not a horizontal strip therefore when we applying the lifting operator we find

$$\mathcal{L}^{(1,1)/(0)}(\omega) = \mathcal{L}_2^{(1)}\mathcal{L}_1^{(0)}(\omega) = (12 + 21) - (12 + 21) = 0$$

as Theorem 2.3.19 asserts. The eigenvector of ARTR_2 belonging to $S^{(1,1)}$ is actually the basis element $12 - 21$ which has corresponding eigenvalue 0.

Our final example is lifting eigenvectors to the 2 dimensional Specht module $S^{(2,1)} = \langle 112 - 211, 121 - 211 \rangle$. Consider $\mu \subset (2,1)$, from our above work we know that the only modules S^μ which contain eigenvectors in the kernel of $\text{ARTR}_{|\mu|}$ are $\mu = (0)$ or $\mu = (1,1)$. The skew diagram of $(2,1)/(0)$ is not a horizontal strip so lifting ω does not recover an eigenvector. The skew diagram $(2,1)/(1,1)$ is a horizontal strip and so lifting the eigenvector $12 - 21$ we find

$$\mathcal{L}_1^{(1,1)}(12 - 21) = 2 \cdot (112 - 211)$$

which belongs to $S^{(2,1)}$ and is an eigenvector of ARTR_3 with eigenvalue 4. To find the remaining eigenvector for $S^{(2,1)}$ it is enough to look at the orthogonal complement of the subspace $\langle 112 - 211 \rangle \subset S^{(2,1)}$, this gives an eigenvector $112 - 2 \cdot 121 + 211$ with eigenvalue 0.

This ends our discussion of the random-to-random shuffle. Since Dieker and Saliola's work Lafrenière has shown that similar techniques can be applied to more general symmetrized shuffling operators [30]. In the next chapter we show that the technique of lifting eigenvectors can be used to recover the spectrum of a variety of transposition shuffles, including the one-sided transposition shuffle. The lifting we present in Chapter 3 differs in key ways from Deiker and Saliola's [19] because we are analysing very different random walks. The random-to-random shuffle is a uniform measure on its non-identity support which covers many conjugacy class of S_n , whereas the one-sided transposition shuffle is not uniform on its single conjugacy class of transpositions. These differences lead to changes in the underlying algebra of lifting eigenvectors and subsequently a change in the required lifting operators. The biggest difference for the one-sided transposition shuffle is the corresponding lifting operators are injective and thus never kill our eigenvectors (see Theorem 3.2.18). This means that *all* the eigenvectors as well as the eigenvalues of the one-sided transposition shuffle belonging to S^λ may be recovered by lifting eigenvectors from modules S^μ with $\mu \subset \lambda$. In Chapter 3 we give the full details of this technique for the one-sided transposition shuffle. Later in Chapter 4 we extend the technique of lifting eigenvectors to the random transposition shuffle and one-sided transposition shuffle on the

hyperoctohedral group.

Chapter 3

The One-sided Transposition Shuffle

In this chapter we introduce a new class of shuffles called *one-sided transposition shuffles*: these have the defining property that at step i the right hand's position (R^i) is chosen according to a distribution supported on $[n]$, and given the value of R^i the distribution of the left hand's position (L^i) is supported on the set $\{1, \dots, R^i\}$. In this chapter we restrict ourselves to the case where the left hand is chosen uniformly from $\{1, \dots, R^i\}$. We begin by focusing on the situation where our right hand is chosen uniformly from its possible range, we call this the *(unbiased) one-sided transposition shuffle*. Afterwards we generalise our techniques to allow our right hand to be driven by a non-uniform distribution, we call these the *biased one-sided transposition shuffles*.

3.1 Main Results

In order to state our main results we introduce the (unbiased) one-sided transposition shuffle formally as follows.

Definition 3.1.1. The *(unbiased) one-sided transposition shuffle* is the random walk on S_n generated by the following distribution on the conjugacy class of transpositions:

$$\text{OST}_n(\tau) = \begin{cases} \frac{1}{n} \cdot \frac{1}{j} & \text{if } \tau = (ij) \text{ for some } 1 \leq i < j \leq n \\ 0 & \text{otherwise.} \end{cases}$$

We use the convention that all permutations (ii) are equal to the identity element e , and therefore $\text{OST}_n(e) = \frac{1}{n}(1 + \frac{1}{2} + \dots + \frac{1}{n}) = H_n/n$, where H_k denotes the k^{th} harmonic number.

This shuffle is clearly reversible, transitive, and has stationary distribution equal to the uniform distribution on S_n , denoted π_n . However, the shuffle is not constant on the conjugacy class of transpositions, unlike the previously seen random transposition shuffle. We look to study the total variation mixing time of the family of one-sided transposition shuffles $\{\text{OST}_n\}_{n \in \mathbb{N}}$. Recall that mixing time is defined as follow:

$$t_{\text{TV}}(\varepsilon) = \min\{t : \|\text{OST}_n^t - \pi_n\|_{\text{TV}} < \varepsilon\}.$$

Existence of a cutoff at time t_n implies that $t_n^{\text{mix}}(\varepsilon) \sim t_n$ for all $\varepsilon \in (0, 1)$ (see Definition 1.1.33). The main conclusion of our work is that the one-sided transposition shuffle exhibits a cutoff at time $t_n = n \log n$.

Theorem 3.1.2. *The one-sided transposition shuffle OST_n satisfies the following bounds for any $c_1 > 0$ and $c_2 > 2$:*

$$\limsup_{n \rightarrow \infty} \|\text{OST}_n^{n \log n + c_1 n} - \pi_n\|_{\text{TV}} \leq \sqrt{2}e^{-c_1}, \quad (3.1)$$

$$\text{and } \liminf_{n \rightarrow \infty} \|\text{OST}_n^{n \log n - n \log \log n - c_2 n} - \pi_n\|_{\text{TV}} \geq 1 - \frac{\pi^2}{6(c_2 - 2)^2}. \quad (3.2)$$

Thus, the one-sided transposition shuffle exhibits a cutoff at time $n \log n$ with a window of order $n \log \log n$.

The lower bound on the total variation distance in (3.2) will be obtained via a coupling argument which allows us to compare the one-sided transposition shuffle to a variation of a coupon collector's problem. To establish the upper bound on total variation distance we make use of the classical ℓ^2 bound given in Theorem 1.1.27. In order to use this result we compute the eigenvalues of the one-sided transposition shuffle.

To analyse the spectrum of the one-sided transposition shuffle we make use of the technique of lifting eigenvectors. We make several non-trivial changes to the technique presented in Section 2.3.3 in order to employ it in the analysis of transposition shuffles: we believe that this is the first time such a technique has been shown to be applicable to non-symmetrized shuffles or to a transposition shuffle. In Section 3.2 we describe an explicit method for obtaining the eigenvectors of OST_{n+1} from those of OST_n . The key to our method is to show that each eigenvalue of OST_n corresponds to a standard Young tableau, and may be computed explicitly from the entries in the tableau. We state here the main result which we aim towards with our analysis.

Theorem 3.1.3. *The eigenvalues of OST_n are labelled by standard Young tableaux of size n , and the eigenvalue represented by a tableau of shape λ appears d_λ times, where d_λ is the dimension of λ . For a standard Young tableau T of shape λ the eigenvalue corresponding to T is given by*

$$\text{eig}(T) = \frac{1}{n} \sum_{\substack{\text{boxes} \\ (i,j)}} \frac{j - i + 1}{T(i, j)}, \quad (3.3)$$

where the sum is performed over all boxes (i, j) in T .

The organisation of this chapter is as follows. Section 3.2 will be dedicated to the proof of Theorem 3.1.3. We give full details of the lifting procedure, highlighting original contributions to the method which allow the recovery of the eigenvalues for the one-sided transposition shuffle. In Section 3.3 we first explore some important properties of the eigenvalues for OST_n , and then use these to prove the upper bound on the mixing time given by Theorem 3.1.2. The corresponding lower bound will be proved in Section 3.5, using entirely probabilistic arguments. In Section 3.6 we consider a generalisation of the one-sided transposition shuffle, in which R^i is chosen according to a non-uniform

distribution: we show that the algebraic technique developed for OST_n holds in this more general setting, and that the well-known mixing time for the random transposition shuffle may be recovered in this way. Finally in Section 3.7 we show that the unbiased one-sided transposition shuffle exhibits a cutoff in separation distance at $n \log n$, the same time as our cutoff in total variation distance.

3.2 Lifting Eigenvectors for Transposition Shuffles

In this section we explore the technique of lifting eigenvalues for transposition shuffles. Our analysis follows a similar path to Dieker and Saliola's but with several novel changes which allow us to describe every eigenvector of the one-sided transposition shuffle by lifting. We end this section showing how lifting may be used to recover all the eigenvalues for the random transposition shuffle which were previously stated in Section 2.3.1.

3.2.1 Lifting Eigenvectors for the One-sided Transposition Shuffle

Recall that the permutation modules M^λ of \mathfrak{S}_n are formed of a basis of words in the alphabet $[n]$ with length n and evaluation $\lambda \vdash n$. The simple modules S^λ of \mathfrak{S}_n are called Specht modules and labelled by λ a partition of n . Within each permutation module M^λ we find one copy of the Specht module S^λ . The permutation module $M^{(1^n)}$ is isomorphic to the regular module for \mathfrak{S}_n , and has decomposition,

$$M^{(1^n)} \cong \bigoplus_{\lambda \vdash n} d_\lambda S^\lambda \text{ as } \mathfrak{S}_n\text{-modules.} \quad (3.4)$$

To model our shuffle OST_n acting on the space \mathfrak{S}_n we need to turn it into a linear operator. In fact we turn it into an element of our group algebra \mathfrak{S}_n .

Definition 3.2.1. Let $n \in \mathbb{N}$. The one-sided transposition shuffle on n cards may be viewed as the following element of the group algebra \mathfrak{S}_n .

$$\sum_{1 \leq i < j \leq n} \text{OST}_n((ij))(ij) = \sum_{1 \leq i < j \leq n} \frac{1}{nj}(ij).$$

To simplify our calculations it is convenient to scale this operator by n , so we introduce a new element called the *algebraic one-sided transposition shuffle*:

$$\text{AOST}_n := \sum_{1 \leq i < j \leq n} \frac{1}{j}(ij). \quad (3.5)$$

Lemma 3.2.2. Let ν be a distribution of S_n , and define $v = \sum_{\eta \in S_n} \nu(\eta) \eta$ an element of \mathfrak{S}_n . Then ν is an eigenvector for OST_n with eigenvalue ε if and only if v is an eigenvalue for AOST_n with eigenvalue $n\varepsilon$.

Proof. The proof follows from the same argument as Lemma 2.3.13. \square

By realising the one-sided transposition shuffle as an element of the group algebra we can concentrate on finding the eigenvalues of AOST_n acting on \mathfrak{S}_n or equivalently $M^{(1^n)}$. Furthermore, applying equation (3.4), we can reduce the problem of finding eigenvalues for the shuffle on $M^{(1^n)}$ to the problem of finding eigenvalues belonging to the individual

Specht modules S^λ comprised of the words formed by polytabloids. To lift our eigenvectors from \mathfrak{S}_n -modules S^λ to \mathfrak{S}_{n+1} modules $S^{\lambda+e_i}$ we introduce linear operators on the vector space of words M^n . The one-sided transposition shuffle admits a recursive structure which is seen when we focus on the difference of AOST_{n+1} and AOST_n ,

$$\text{AOST}_{n+1} - \text{AOST}_n = \frac{1}{(n+1)} \sum_{1 \leq i \leq (n+1)} (i \ n+1). \quad (3.6)$$

This signifies that the only difference between shuffle $n+1$ and n cards is the movement of the new card in position $n+1$. This relation is key to lifting eigenvectors as it allows us to model the action of AOST_{n+1} using AOST_n and special linear operators which we now define.

Definition 3.2.3. We define two linear operators on the spaces spanned by words. To do so, it is enough to define the effect on any given word. Let $w = w_1 \dots w_n \in W^n$.

1. Let $a \in [n+1]$. Define the *adding operator* $\Phi_a : M^n \rightarrow M^{n+1}$ as follows:

$$\Phi_a(w) := w a \quad (3.7)$$

The adding operator appends the symbol a to the end of the word. If we think about the adding operator acting on tabloids $\{T\}$ instead of words, then $\Phi_a(\{T\})$ is the process of adding a box labelled $n+1$ onto row a of the tabloid $\{T\}$.

2. Let $a, b \in [n]$ Define the *switching operator* $\Theta_{b,a} : M^n \rightarrow M^n$ as follows:

$$\Theta_{b,a}(w) := \sum_{\substack{1 \leq k \leq n \\ w_k = b}} w_1 \dots w_{k-1} a w_{k+1} \dots w_n. \quad (3.8)$$

The operator $\Theta_{b,a}$ forms a linear sum of all words created from w by replacing one occurrence of b with a .

We previously defined the switching operator in Definition 2.3.15. Again we should define the operators for all possible domains separately but to do so would burden us with more notation: it will always be clear from context which domain and codomain we are considering. The adding operator is our analogue of the shuffling operator (Section 2.3.3) and it allow us to lift eigenvectors of the one-sided transposition shuffle.

Recall that given a partition λ we may add boxes onto certain rows to form a new partition. Since we are allowing ourselves to blur the distinction between n -tuples and partitions of n , if we add a box on row i the new tuple/partition formed is $\lambda + e_i$. By the end of our analysis we only need the cases where $\lambda + e_i$ is a partition.

Lemma 3.2.4. *Given $a \in [n+1]$ and an n -tuple λ of non-negative integers summing to n , we have*

$$\Phi_a : M^\lambda \rightarrow M^{\lambda+e_a}.$$

In other words the restriction of Φ_a to M^λ (or S^λ) has image in $M^{\lambda+e_a}$.

Given $a, b \in [n]$ and n -tuples λ, μ of non-negative integers summing to n with $\lambda + e_a = \mu + e_b$, we have

$$\Theta_{b,a} : M^\lambda \rightarrow M^\mu,$$

i.e., the restriction of $\Theta_{b,a}$ to M^λ (or S^λ) has image in M^μ .

Our next result establishes the crucial equation upon which all the subsequent results in this section rely. It relates the shuffle on n cards to that on $n+1$ cards, and gives us the basis of lifting eigenvectors. The following theorem is an analogue of [19, Theorem 38].

Theorem 3.2.5. *Given $n \in \mathbb{N}$, acting on words in M^n we have*

$$\text{AOST}_{n+1} \circ \Phi_a - \Phi_a \circ \text{AOST}_n = \frac{1}{n+1} \Phi_a + \frac{1}{n+1} \sum_{1 \leq b \leq n} \Phi_b \circ \Theta_{b,a}. \quad (3.9)$$

Proof. It suffices to prove the result on a generic word in M^n . Let $w = w_1 \dots w_n$ be a word of length n and let $a \in [n+1]$. Consider the two terms on the left hand side applied to w :

$$(\text{AOST}_{n+1} \circ \Phi_a)(w) = \frac{1}{n+1} \sum_{\substack{j=n+1 \\ 1 \leq i \leq n+1}} (i j)(w a) + \sum_{1 \leq i \leq j \leq n} \frac{1}{j} (i j)(w a) \quad (3.10)$$

$$(\Phi_a \circ \text{AOST}_n)(w) = \left(\sum_{1 \leq i \leq j \leq n} \frac{1}{j} (i j)(w) \right) a. \quad (3.11)$$

The second summation in (3.10) cancels with (3.11) because the adjoined a is in the $(n+1)$ -th place, therefore it never moves and may be brought outside the sum. This leaves us with the following:

$$(\text{AOST}_{n+1} \circ \Phi_a - \Phi_a \circ \text{AOST}_n)(w) = \frac{1}{n+1} \sum_{1 \leq i \leq n+1} (i \ n+1)(w a). \quad (3.12)$$

If $i = n+1$ we move nothing, giving the term $w a = \Phi_a(w)$. Otherwise we apply the transposition $(i \ n+1)$ to $w a$. This has the same effect as replacing the i^{th} symbol w_i in w with a and then appending w_i on the end of the new word. Since we do this for all symbols in w , the net effect is the same as $\sum_{1 \leq b \leq n} \Phi_b \circ \Theta_{b,a}$ applied to w . The operator $\Theta_{b,a}$ systematically finds all occurrences of the letter b in w and replaces with an a , and then Φ_b puts the b back on the end. Since $w \in W^n$, all possibilities are exhausted by letting b range over every possible letter $1 \leq b \leq n$. This completes the proof. \square

In terms of shuffling cards, we can interpret (3.9) as taking into account the difference between shuffling a deck and then adding a card versus adding a card and then shuffling. If we can understand how the operators Φ_a and $\Theta_{b,a}$ behave, then this inductively tells us how the shuffle on $n+1$ cards behaves using information about the shuffle on n cards, vastly simplifying our original problem. We now record a key property of the linear maps $\Theta_{a,b}$.

Lemma 3.2.6 (See Section 2.9 of [38]). *The switching operators $\Theta_{b,a}$ are \mathfrak{S}_n -module morphisms.*

Proof. This is clear from the definitions, since S_n is acting by place permutations, the operator $\Theta_{b,a}$ commutes with the action of \mathfrak{S}_n . It amounts to the same thing to replace

an occurrence of the symbol b with a symbol a and then permute the word as to first permute the word and then replace the same symbol b in its new position with an a . \square

The above result is helpful in understand how our adding operators behave when restricted to Specht modules.

Lemma 3.2.7 (Lemma 44 of [19]). *Let $\lambda \vdash n$ be such that $\lambda + e_a = \mu + e_b$ for some $a, b \in [n]$. Then $\Theta_{b,a}$ is non-zero on S^λ if and only if λ dominates the non-increasing rearrangement of μ . In particular, if $b > a$ then $\Theta_{b,a}(S^\lambda) = 0$.*

Proof. Since S^λ is simple and $\Theta_{b,a}$ is a module homomorphism, the image $\Theta_{b,a}(S^\lambda)$ is 0 or isomorphic to S^λ , by Schur's lemma. But $\Theta_{b,a}(S^\lambda)$ lies in M^μ because of the relationship $\lambda + e_a = \mu + e_b$. Let ν be the non-increasing rearrangement of μ , then $\nu \vdash n$, and $M^\mu \cong M^\nu$, so M^μ has a submodule isomorphic to S^λ if and only if λ dominates ν , see Lemma 2.2.21. Therefore, if λ does not dominate ν we have $\Theta_{b,a}(S^\lambda) = 0$. Conversely if λ dominates ν then $\Theta_{b,a}(S^\lambda) \cong S^\lambda$ (see [38, Section 2.10] for further details).

To finish, note that in terms of diagrams the fact that $\lambda + e_a = \mu + e_b$ corresponds to the fact that we can get from the diagram for λ to that for μ by moving a box from row b to row a . Hence, under the given hypothesis, we have that λ dominates the non-increasing rearrangement of μ if and only if $b \leq a$. \square

Example 3.2.8. Consider the Specht module $S^{(3,1)}$. We have seen previously that this is spanned by the following polytabloids represented as words

$$S^{(3,1)} = \langle 1112 - 2111, 1121 - 2111, 1211 - 2111 \rangle.$$

If we apply $\Theta_{2,1}$ we go from partition $(3,1)$ to the partition $(2,2)$. Lemma 3.2.7 tells us that all the elements of $S^{(3,1)}$ belong to the kernel of $\Theta_{2,1}$. Verifying this result we find,

$$\begin{aligned} \Theta_{2,1}(1112 - 2111) &= 1111 - 1111 = 0 \\ \Theta_{2,1}(1211 - 2111) &= 0 \\ \Theta_{2,1}(1121 - 2111) &= 0 \end{aligned}$$

Instead if we apply $\Theta_{1,2}$ we find a non-zero elements belonging to the module $M^{(2,2)}$.

$$\begin{aligned} \Theta_{1,2}(1112 - 2111) &= 1212 + 1122 - 2211 - 2121 \\ \Theta_{1,2}(1211 - 2111) &= 1212 + 1221 - 2121 - 2112 \\ \Theta_{1,2}(1121 - 2111) &= 1221 + 1122 - 2211 - 2121. \end{aligned}$$

The preceding result shows that when we restrict equation (3.9) to a Specht module S^λ we can change the index of the summation in the final term on the right hand side, as follows.

Corollary 3.2.9 (Corollary 45 of [19]).

$$(\text{AOST}_{n+1} \circ \Phi_a - \Phi_a \circ \text{AOST}_n)|_{S^\lambda} = \frac{1}{n+1} \Phi_a|_{S^\lambda} + \frac{1}{n+1} \sum_{1 \leq b \leq a} \Phi_b \circ \Theta_{b,a}|_{S^\lambda} \quad (3.13)$$

Having restricted equation (3.9) to the Specht module S^λ , we analyse the image in the module $M^{\lambda+e_a}$: note that it is clear from the left hand side of (3.9) that we land in $M^{\lambda+e_a}$. The following lemma and its proof follow Lemma 41 of Deiker and Saliola [19].

Lemma 3.2.10 (Lemma 41 of [19]). *Suppose $\lambda \vdash n$. Then the subspace $\Phi_a(S^\lambda)$ is contained in an \mathfrak{S}_{n+1} submodule of $M^{\lambda+e_a}$ that is isomorphic to $\bigoplus_{\mu} S^\mu$, where the sum ranges over the partitions μ obtained from λ by adding a box in row i for $i \leq a$.*

Proof. Let w be a word of length n , so that $\Phi_a(w) = wa$. If the symbol b does not occur in w then

$$\Phi_a(w) = \Theta_{b,a}(\Phi_b(w)).$$

Let $b = l(\lambda) + 1$, so b does not appear in any $w \in M^\lambda$, and consider the \mathfrak{S}_{n+1} -submodule N of $M^{\lambda+e_b}$ generated by the elements $x b$ with $x \in S^\lambda$,

$$N = \langle x b : x \in S^\lambda \rangle.$$

The submodule N is isomorphic to $\text{Ind}_{\mathfrak{S}_n \times \mathfrak{S}_1}^{\mathfrak{S}_{n+1}}(S^\lambda \otimes S^1)$ (this is essentially the definition of how to induce), and using the branching rules on S_n this decomposes as a multiplicity free direct sum of Specht modules S^μ , where $\mu \vdash n+1$ and $\lambda \subset \mu$ (see [22, Theorem I.7]). Using the observation at the start of the proof, we obtain

$$\Phi_a(S^\lambda) = \Theta_{b,a}(\Phi_b(S^\lambda)) \subseteq \Theta_{b,a}(\langle \Phi_b(S^\lambda) \rangle) = \Theta_{b,a}(N) \cong \bigoplus_{\substack{\mu \vdash n+1 \\ \lambda \subset \mu}} \Theta_{b,a}(S^\mu).$$

Now note that $\Theta_{b,a}$ sends any word with evaluation $\lambda + e_b$ to a word with evaluation $\lambda + e_a$, and hence $\Theta_{b,a}(M^{\lambda+e_b}) \subseteq M^{\lambda+e_a}$. It follows that all nonzero summands S^μ appearing on the right hand side occur for $\mu \vdash n+1$ dominating the non-decreasing rearrangement of $\lambda + e_a$, and then by Lemma 3.2.7 we can conclude that μ is obtained from λ by adding a cell in row i with $i \leq a$, as required. \square

Lemma 3.2.7 tells us a great deal about the image of $\Phi_a(S^\lambda)$. Specifically we know that $\Phi_a(S^\lambda)$ contains exactly one submodule isomorphic to $S^{\lambda+e_a}$. This means that after applying our adding operator we may project down onto the required simple module to get a linear operator from $S^\lambda \rightarrow S^{\lambda+e_a}$: these will become our lifting operators. To project onto the required simple module we introduce the isotypic projections.

Definition 3.2.11. Let G be a finite group, V be a simple G -module with corresponding character χ_V . Let M be a module for G , we may define the *isotypic projection* π^V from M onto its unique defined V -isotypic component (see Definition 1.2.17) in the following way:

$$\pi^V : M \rightarrow M, \quad \pi^V(m) = \left(\frac{\dim(V)}{|G|} \sum_{g \in G} \overline{\chi_V(g)} g \right) m.$$

Note that the isotypic projection is given by the action of the element $\frac{\dim(V)}{|G|} \sum_{g \in G} \overline{\chi_V(g)} g$ which belongs to the centre of the group algebra $\mathbb{C}[G]$. Thus, π^V is a morphism of G -modules and commutes with any G -module endomorphism.

Using these projections, we can now define our lifting operators, which will be proven to map eigenvectors of AOST_n to those of AOST_{n+1} .

Definition 3.2.12. Suppose $\lambda \vdash n$ and $\lambda + e_a = \mu \vdash n + 1$ are two partitions. Define the *lifting operator*

$$\kappa_a^{\lambda, \mu} := \pi^\mu \circ \Phi_a : S^\lambda \rightarrow S^\mu \subseteq M^\mu.$$

Note that since $\Phi_a(S^\lambda) \subseteq M^\mu$ and M^μ contains a unique copy of S^μ , therefore the S^μ isotypic component in M^μ is just S^μ and $\kappa_a^{\lambda, \mu}(S^\lambda)$ is actually contained in S^μ .

We next prove that our lifting operators $\kappa_a^{\lambda, \mu}$ are injective \mathfrak{S}_n -module morphisms. This differs from the lifting operators \mathcal{L}_i^λ for the random to random shuffle which we have seen are not injective. The difference in lifting operators depends in an essential way on our adding operator Φ_a . Using the lifting operators $\kappa_a^{\lambda, \mu}$ we are able to find *all* the eigenvectors for a module S^μ by lifted eigenvectors from partitions $\lambda \subset \mu$.

Corollary 3.2.13. *For any $\lambda \vdash n$ and $\lambda + e_a \vdash n + 1$, there exists some $v \in S^\lambda$ with*

$$\kappa_a^{\lambda, \lambda + e_a}(v) \neq 0.$$

Proof. If $\kappa_a^{\lambda, \lambda + e_a}(S^\lambda) = \{0\}$, then the image $\Phi_a(S^\lambda)$ lies in the kernel of the projection $\pi^{\lambda + e_a} : M^{\lambda + e_a} \rightarrow S^{\lambda + e_a}$, which is an \mathfrak{S}_{n+1} -submodule with no component equal to $S^{\lambda + e_a}$. Hence the submodule generated by $\Phi_a(S^\lambda)$ has no component equal to $S^{\lambda + e_a}$. But we previously observed that (with notation as in the proof of Lemma 3.2.10)

$$\langle \Phi_a(S^\lambda) \rangle = \langle \Theta_{b,a}(\Phi_b(S^\lambda)) \rangle = \Theta_{b,a}(\langle \Phi_b(S^\lambda) \rangle) \cong \Theta_{b,a}(N) \cong \bigoplus_{1 \leq i \leq a} S^{\lambda + e_i}.$$

Since the right hand side contains $S^{\lambda + e_a}$ as a summand, we have a contradiction. \square

We already know the map π^μ is an \mathfrak{S}_{n+1} module morphism. Let us realise \mathfrak{S}_n inside \mathfrak{S}_{n+1} as the stabilizer of the $(n + 1)^{\text{th}}$ position. Then any \mathfrak{S}_{n+1} -module gives rise to an \mathfrak{S}_n -module by restriction (see Definition 2.2.23). Hence, π^μ is a \mathfrak{S}_n -module morphism.

Lemma 3.2.14. *The linear operator $\kappa_a^{\lambda, \lambda + e_a}$ is a \mathfrak{S}_n -module morphism with trivial kernel. Therefore, the lifting operators are injective.*

Proof. Our key observation is that our adding operator commutes with elements of \mathfrak{S}_n inside of \mathfrak{S}_{n+1} , i.e., $\Phi_a(\sigma(v)) = \sigma(\Phi_a(v))$ for all $v \in S^\lambda$ and $\sigma \in \mathfrak{S}_n \subset \mathfrak{S}_{n+1}$. This is obvious, since Φ_a adds an element in the final position which is not affected by σ . Hence $\kappa_a^{\lambda, \lambda + e_a}$ is the composition of two \mathfrak{S}_n -module morphisms. The final observation follows from Corollary 3.2.13 – since $\kappa_a^{\lambda, \lambda + e_a}$ is a nonzero module morphism with a simple module as its domain, it must be injective by Schur’s lemma (Lemma 1.2.9). \square

The lifting operators $\kappa_a^{\lambda, \lambda + e_a}$ being injective is a key point which simplifies our analysis compared to that of the random to random shuffle - in [19] the lifting operators can kill eigenvectors. The next results show that $\kappa_a^{\lambda, \lambda + e_a}$ lifts eigenvectors of AOST_n into those of AOST_{n+1} . To establish this we apply our projection $\pi^{\lambda + e_a}$ to equation (3.13). Before we do this we state an identity between our adding and switching operators which is critical to the proofs which follow.

Lemma 3.2.15. *Take $a, b \in [n]$. Then our adding and switching operators satisfy the following equality*

$$\Phi_b \circ \Theta_{b,a} = \Theta_{b,a} \circ \Phi_b - \Phi_a. \quad (3.14)$$

Proof. Let w be a word of length n . Let $v := \Phi_b(w)$ so that $v_{n+1} = b$. Consider the affect of switching operator $\Theta_{b,a}$ on the word v of length $n+1$. We find,

$$\begin{aligned} \Theta_{b,a} \circ \Phi_b(w) &= \sum_{\substack{1 \leq k \leq n+1 \\ v_k = b}} v_1 \dots v_{k-1} a v_{k+1} \dots v_{n+1} \\ &= \Phi_a(w) + \sum_{\substack{1 \leq k \leq n \\ v_k = b}} v_1 \dots v_{k-1} a v_{k+1} \dots v_n b = \Phi_a(w) + \Phi_b \circ \Theta_{b,a}(w) \end{aligned}$$

The second equality follows from taking the $k = n+1$ out of the sum. The last equality follows from the fact that $v_i = w_i$ for all $i \in [n]$. Rearranging this final expression gives us the desired equality. \square

We can now state our versions of [19, Lemma 48, Theorem 49]; the proofs follow *mutatis mutandis* from the ones given there (the changes needed are to the constants in equation (3.9)).

Lemma 3.2.16 (Lemma 48 of [19]). *Let $\lambda \vdash n$, and $a \in \{1, 2, \dots, l(\lambda) + 1\}$. Take $i \in [n]$ such that $1 \leq i \leq a$ and set $\mu = \lambda + e_i$. Then,*

$$\text{AOST}_{n+1} \circ \kappa_a^{\lambda, \mu} - \kappa_a^{\lambda, \mu} \circ \text{AOST}_n = \frac{2 + \lambda_a - a}{n+1} \kappa_a^{\lambda, \mu} + \frac{1}{n+1} \sum_{i \leq b < a} \Theta_{b,a} \circ \kappa_b^{\lambda, \mu}.$$

Proof. This follows from the work in [19] because we have not changed the switching operators $\Theta_{b,a}$. The values on the right hand side change to reflect our adding operators and our new equation (3.9). We present the modified proof in full for completeness.

Continuing from Corollary 3.2.9 we know

$$\text{AOST}_{n+1} \circ \Phi_a - \Phi_a \circ \text{AOST}_n|_{S^\lambda} = \frac{1}{n+1} \Phi_a|_{S^\lambda} + \frac{1}{n+1} \sum_{1 \leq b \leq a} \Phi_b \circ \Theta_{b,a}|_{S^\lambda}.$$

Apply the isotypic projection π^μ to both sides of the equation. Since AOST_{n+1} is given by the action of an element of the group algebra of \mathfrak{S}_{n+1} and π^μ is an \mathfrak{S}_{n+1} -module morphism, these operators commute and so we have

$$\text{AOST}_{n+1} \circ \kappa_a^{\lambda, \mu} - \kappa_a^{\lambda, \mu} \circ \text{AOST}_n = \frac{1}{n+1} \kappa_a^{\lambda, \mu} + \frac{1}{n+1} \sum_{1 \leq b \leq a} (\pi^\mu \circ \Phi_b \circ \Theta_{b,a})|_{S^\lambda}. \quad (3.15)$$

Our adding operator Φ_a satisfies the equation (3.14) and so

$$\pi^\mu \circ (\Phi_b \circ \Theta_{b,a}) = \pi^\mu \circ (\Theta_{b,a} \circ \Phi_b - \Phi_a) = \Theta_{b,a} \circ \pi^\mu \circ \Phi_b - \pi^\mu \circ \Phi_a.$$

The right side side of the equation (3.15) now becomes

$$\frac{1-a}{n+1} \kappa_a^{\lambda, \mu} + \frac{1}{n+1} \sum_{1 \leq b \leq a} \Theta_{b,a} \circ \kappa_b^{\lambda, \mu}. \quad (3.16)$$

Notice that if $b = a$ then $\Theta_{a,a}(w)$ acts as a scalar by the number of occurrences of the symbol a in w . In our case all words in $\kappa_a^{\lambda,\mu}(S^\lambda)$ contain $\lambda_a + 1$ occurrences of a . Finally if $b < i$ we know that $\Phi_b(S^\lambda)$ does not contain the module $S^{\lambda+e_i}$ (by Lemma 3.2.10), so $\kappa_b^{\lambda,\mu} = 0$. Thus, equation (3.16) is equal to

$$\frac{2 + \lambda_a - a}{n + 1} \kappa_a^{\lambda,\mu} + \frac{1}{n + 1} \sum_{i \leq b < a} \Theta_{b,a} \circ \kappa_b^{\lambda,\mu}. \quad \square$$

Theorem 3.2.17 (Theorem 49 of [19]). *Let $\lambda \vdash n$, and $a \in \{1, 2, \dots, l(\lambda) + 1\}$. Take $i \in [n]$ such that $1 \leq i \leq a$ and set $\mu = \lambda + e_i$. Then,*

$$\text{AOST}_{n+1} \circ \kappa_a^{\lambda,\mu} - \kappa_a^{\lambda,\mu} \circ \text{AOST}_n = \frac{(2 + \lambda_i - i)}{n + 1} \kappa_a^{\lambda,\mu}. \quad (3.17)$$

In particular if $v \in S^\lambda$ is an eigenvector of AOST_n with eigenvalue ε , then $\kappa_a^{\lambda,\mu}(v)$ is an eigenvector of AOST_{n+1} belonging to S^μ with eigenvalue

$$\varepsilon + \frac{2 + \lambda_i - i}{n + 1}. \quad (3.18)$$

Proof. This proof follows from the work in [19] with minor changes to reflect the one-sided transposition shuffle. For $i = a$ this follows from Lemma 3.2.16, this will be the key to showing it holds in all cases. Let $\mu = \lambda + e_i$, we know from Lemma 3.2.16 that lifting via $\kappa_i^{\lambda,\mu}$ gives us

$$\text{AOST}_{n+1} \circ \kappa_i^{\lambda,\mu} - \kappa_i^{\lambda,\mu} \circ \text{AOST}_n = \frac{2 + \lambda_i - i}{n + 1} \kappa_i^{\lambda,\mu}. \quad (3.19)$$

Applying the linear operator $\Theta_{i,a}$ to the above equation,

$$\text{AOST}_{n+1} \circ \Theta_{i,a} \circ \kappa_i^{\lambda,\mu} - \Theta_{i,a} \circ \kappa_i^{\lambda,\mu} \circ \text{AOST}_n = \frac{2 + \lambda_i - i}{n + 1} \Theta_{i,a} \circ \kappa_i^{\lambda,\mu}. \quad (3.20)$$

Consider the left hand side of (3.20) break up the lifting operator into $\kappa_i^{\lambda,\mu} = \pi^\mu \circ \Phi_i$. The projection π^μ commutes with the \mathfrak{S}_{n+1} module morphism $\Theta_{i,a}$. Performing this we obtain the equation below, restricted to S^λ ;

$$\begin{aligned} & \text{AOST}_{n+1} \circ \Theta_{i,a} \circ \kappa_i^{\lambda,\mu} - \Theta_{i,a} \circ \kappa_i^{\lambda,\mu} \circ \text{AOST}_n \\ &= \text{AOST}_{n+1} \circ \pi^\mu \circ \Theta_{i,a} \circ \Phi_i - \pi^\mu \circ \Theta_{i,a} \circ \Phi_i \circ \text{AOST}_n \end{aligned} \quad (3.21)$$

From the identity $\Theta_{i,a} \circ \Phi_i = \Phi_a + \Phi_i \circ \Theta_{i,a}$ (which is a rearrangement of equation (3.14)), the equation (3.21) becomes

$$(\text{AOST}_{n+1} \circ \pi^\mu \circ \Phi_a + \text{AOST}_{n+1} \circ \pi^\mu \circ \Phi_i \circ \Theta_{i,a}) - (\pi^\mu \circ \Phi_a \circ \text{AOST}_n + \pi^\mu \circ \Phi_i \circ \Theta_{i,a} \circ \text{AOST}_n)$$

Our one-sided transposition shuffle AOST_n commutes with $\Theta_{i,a}$, swapping these around

and using our observation (3.19), we reduce the above equation to:

$$\begin{aligned}
& (\text{AOST}_{n+1} \circ \pi^\mu \circ \Phi_a - \pi^\mu \circ \Phi_a \circ \text{AOST}_n) + (\text{AOST}_{n+1} \circ \pi^\mu \circ \Phi_i - \pi^\mu \circ \Phi_i \circ \text{AOST}_n) \circ \Theta_{i,a}. \\
&= (\text{AOST}_{n+1} \circ \kappa_a^{\lambda,\mu} - \kappa_a^{\lambda,\mu} \circ \text{AOST}_n) + (\text{AOST}_{n+1} \circ \kappa_i^{\lambda,\mu} - \kappa_i^{\lambda,\mu} \circ \text{AOST}_n) \circ \Theta_{i,a}. \\
&= (\text{AOST}_{n+1} \circ \kappa_a^{\lambda,\mu} - \kappa_a^{\lambda,\mu} \circ \text{AOST}_n) + \frac{2 + \lambda_i - i}{n + 1} \kappa_i^{\lambda,\mu} \circ \Theta_{i,a}. \tag{3.22}
\end{aligned}$$

We are finished manipulating the left hand side of equation (3.20), and now focus our attention on the right hand side. Taking the right hand side we split the lifting operator into $\kappa_i^{\lambda,\mu} = \pi^\mu \circ \Phi_i$, and use the properties of $\Theta_{i,a}$ and equation (3.14) in order to find the following:

$$\begin{aligned}
& \frac{2 + \lambda_i - i}{n + 1} \Theta_{i,a} \circ \kappa_i^{\lambda,\mu} \\
&= \frac{2 + \lambda_i - i}{n + 1} \pi^\mu \circ \Theta_{i,a} \circ \Phi_i|_{S^\lambda} = \frac{2 + \lambda_i - i}{n + 1} \pi^\mu \circ \Phi_a|_{S^\lambda} + \frac{2 + \lambda_i - i}{n + 1} \pi^\mu \circ \Phi_i \circ \Theta_{i,a}|_{S^\lambda} \\
&= \frac{2 + \lambda_i - i}{n + 1} \kappa_a^{\lambda,\mu} + \frac{2 + \lambda_i - i}{n + 1} \kappa_i^{\lambda,\mu} \circ \Theta_{i,a}. \tag{3.23}
\end{aligned}$$

Combining equations (3.22) and (3.23) gives the desired result. \square

The last theorem tells us exactly how to turn eigenvectors of AOST_n into those of AOST_{n+1} and critically it shows how the eigenvalues change in value. An important observation here is the numerator of the change $\frac{2+\lambda_i-i}{n+1}$ depends on what box (or letter) we are adding and the denominator depends on the step of our lifting, thus they are independent of one another. The final part of our analysis rests on showing that *all* of the eigenvectors in a Specht module S^μ can be retrieved by lifting from Specht modules S^λ with $\mu = \lambda + e_a$. In fact, we show that these lifted eigenvectors form a basis of S^μ .

Theorem 3.2.18. *For any $\mu \vdash n + 1$ we may find a basis of eigenvectors of AOST_{n+1} for the module S^μ , by lifting the eigenvectors of AOST_n belonging to the modules S^λ with $\lambda \vdash n$ and $\lambda \subset \mu$.*

Proof. We proceed by induction. For $n = 1$ we know that the simple modules $S^{(2)}, S^{(1,1)}$ of \mathfrak{S}_2 are both one dimensional. Therefore, the eigenvector $1 \in S^{(1)}$ when lifted indeed forms a basis for each simple module, this was demonstrated in Example 3.2.20.

Consider the simple module of S^μ with $\mu \vdash n + 1$. We know from the classic branching rules of S_n (Theorem 2.2.25) that the restriction of this module to \mathfrak{S}_n is given by

$$\text{Res}_{\mathfrak{S}_n}^{\mathfrak{S}_{n+1}}(S^\mu) \cong \bigoplus_{\substack{\lambda \vdash n \\ \lambda \subset \mu}} S^\lambda.$$

Importantly S^μ is isomorphic as a vector space to $\bigoplus_{\substack{\lambda \vdash n \\ \lambda \subset \mu}} S^\lambda$. Now suppose we have a basis of eigenvectors for every S^λ . By Lemma 3.2.14 the map $\kappa^{\lambda,\mu}(S^\lambda)$ gives a basis for the submodule S^λ inside of the vector space of $\text{Res}_{\mathfrak{S}_n}^{\mathfrak{S}_{n+1}}(S^\mu)$. Hence, considering all of the lifted eigenvectors from every S^λ together we find a basis for S^μ . By Theorem 3.2.17 the lifted eigenvectors form a basis of eigenvectors for S^μ . \square

Inductively, for any $\lambda \vdash n$, Theorem 3.2.18 gives us the way to find all the eigenvectors for AOST_n belonging to the Specht module S^λ : starting at S^\emptyset and recursively applying

lifting operators until we reach S^λ gives us an eigenvector, and all eigenvectors arise in this way. Note that S^\emptyset has no eigenvectors attached to it, but we allow the empty word ω to be an eigenvector with eigenvalue 0, and $\Phi_a(\omega) = a$. This agrees with the formula in Theorem 3.2.17 because a is the only eigenvector of AOST_1 with eigenvalue $1 = 0 + (2 + 0 - 1)/(1)$. The inductive process of lifting naturally forms one path up Young's lattice which starts at \emptyset and ends at λ . Furthermore, by Theorem 3.2.18 each unique path we take $\emptyset \rightarrow \lambda$ results in a unique eigenvector for S^λ , and all these eigenvectors together form a basis. We now are in a position to prove Theorem 3.1.3.

Proof of Theorem 3.1.3. Every eigenvector in our constructed basis gives a distinct eigenvalue of S^λ , hence there are d_λ distinct eigenvalues. These are eigenvalues for the shuffle AOST_n , and each one appears d_λ times due to the isomorphism in equation (3.4). Overall we have found $\sum_{\lambda \vdash n} d_\lambda^2 = n!$ eigenvalues and thus have a complete set. Given a standard tableau $T \in \text{SYT}(\lambda)$, we build up the tableau following its labelling and keeping track of the changes in eigenvalue given by Theorem 3.2.17. When box (i, j) is added to T we get a change in eigenvalue of $\frac{2+\lambda_i-i}{n+1} = \frac{2+(j-1)-i}{T(i,j)} = \frac{j-i+1}{T(i,j)}$. After summing these changes for all boxes (i, j) in T we divide by n to normalise the eigenvalue, recovering equation 3.3. \square

We have given an explicit description of how to compute the eigenvalues of the shuffle AOST_n . We now state a description of maps $\kappa_a^{\lambda, \lambda+e_a}$ in terms of adding and switching operators, these allow explicit computation of any eigenvector if required. The proof of Theorem 3.2.19 follows directly from the work of [19, Section 5.6].

Lemma 3.2.19 (Theorem 21 [19]). *Let $\lambda \vdash n$. The lifting maps as defined in Definition 3.2.12 are a linear combination of the shuffling and switching operators, explicitly they are:*

$$\kappa_i^{\lambda, \lambda+e_i} = \sum_{1 \leq b_1 < \dots < b_m < b_{m+1} = i} \left(\prod_{j=1}^m \frac{1}{(\lambda_i - i) - (\lambda_{b_j} - b_j)} \Theta_{b_j, b_{j+1}} \right) \Phi_{b_1}. \quad (3.24)$$

For example, the first three lifting operators are:

$$\begin{aligned} \kappa_1^{\lambda, \lambda+e_1} &= \Phi_1 \\ \kappa_2^{\lambda, \lambda+e_2} &= \Phi_2 + \frac{1}{(\lambda_2 - 2) - (\lambda_1 - 1)} \Theta_{1,2} \Phi_1 \\ \kappa_3^{\lambda, \lambda+e_3} &= \Phi_3 + \frac{1}{(\lambda_3 - 3) - (\lambda_2 - 2)} \Theta_{2,3} \circ \Phi_2 + \frac{1}{(\lambda_3 - 3) - (\lambda_1 - 1)} \Theta_{1,3} \circ \Phi_1 \\ &\quad + \frac{1}{((\lambda_3 - 3) - (\lambda_2 - 2))((\lambda_3 - 3) - (\lambda_1 - 1))} \Theta_{2,3} \circ \Theta_{1,2} \circ \Phi_1 \end{aligned}$$

Example 3.2.20. Consider the Specht module $S^{(1)} = \langle 1 \rangle$. The eigenvector of AOST_1 belonging to this module is 1 with eigenvalue 1. Using our lifting operators $\kappa_1^{(1),(2)}$ and $\kappa_2^{(1,1)}$ we find:

$$\begin{aligned} \kappa_1^{(1),(2)}(1) &= \Phi_1(1) = 11 \\ \kappa_2^{(1,1)}(1) &= \left(\Phi_2 - \frac{1}{2} \Theta_{1,2} \Phi_1 \right)(1) = \frac{1}{2} \cdot (12 - 21) \end{aligned}$$

The above elements belong to the modules $S^{(2)}$ and $S^{(1,1)}$ respectively. Furthermore 11 is

an eigenvector of AOST_2 with eigenvalue 2, and $12 - 21$ is an eigenvector of AOST_2 with eigenvalue 1.

For the analysis of the mixing time of the one-sided transposition shuffle we only use the eigenvalues computed in Theorem 3.1.3. To end this section we give an explicit example of computing an eigenvalue using the process described in the proof Theorem 3.1.3.

Example 3.2.21. Let $\lambda = (4, 2, 1)$ and choose a standard Young tableau

$$T_1 = \begin{array}{|c|c|c|c|} \hline 1 & 3 & 6 & 7 \\ \hline 2 & 4 & & \\ \hline 5 & & & \\ \hline \end{array}.$$

We build up T_1 using our lifting maps, at each step Theorem 3.2.17 tells us the change in our eigenvalue. We start at $n = 0$ with eigenvector $\omega \in S^\emptyset$ with eigenvalue 0. Applying the lifting map $\kappa_1^{\emptyset, \emptyset + e_1}$ gives us:

$$\kappa_1^{\emptyset, (1)}(\omega) = 1 \in S^{(1)} \text{ with the eigenvalue } 0 + \frac{2 + 0 - 1}{1} = 1.$$

On the second row, this corresponds to applying $\kappa_2^{(1), (1) + e_2}$, which gives:

$$\kappa_2^{(1), (1,1)} \kappa_1^{\emptyset, (1)}(\omega) = \kappa_2^{(1), (1,1)}(1) = \frac{1}{2}(12 - 21) \text{ with eigenvalue } 1 + \frac{2 + 0 - 2}{2} = 1.$$

Continuing the procedure in the proof of Lemma 3.1.3 we build up T_1 keeping track of the change in eigenvalue at each step.

$$\emptyset \xrightarrow{\frac{1}{1}} \begin{array}{|c|} \hline 1 \\ \hline \end{array} \xrightarrow{\frac{0}{2}} \begin{array}{|c|} \hline 1 \\ \hline 2 \\ \hline \end{array} \xrightarrow{\frac{2}{3}} \begin{array}{|c|c|} \hline 1 & 3 \\ \hline 2 & \\ \hline \end{array} \xrightarrow{\frac{1}{4}} \begin{array}{|c|c|} \hline 1 & 3 \\ \hline 2 & 4 \\ \hline \end{array} \xrightarrow{\frac{-1}{5}} \begin{array}{|c|c|} \hline 1 & 3 \\ \hline 2 & 4 \\ \hline 5 & \\ \hline \end{array} \xrightarrow{\frac{3}{6}} \begin{array}{|c|c|c|} \hline 1 & 3 & 6 \\ \hline 2 & 4 & \\ \hline 5 & & \\ \hline \end{array} \xrightarrow{\frac{4}{7}} \begin{array}{|c|c|c|c|} \hline 1 & 3 & 6 & 7 \\ \hline 2 & 4 & & \\ \hline 5 & & & \\ \hline \end{array}$$

The arrows represent the lifting operators and the value above each arrow is the change in eigenvalue at each step. From here computing the eigenvalue for T_1 is a matter of summing all the changes then dividing by n (because we scaled OST_n in Definition 3.2.1). We find the eigenvalue for T_1 being

$$\text{eig}(T_1) = \frac{1}{7} \left(\frac{1}{1} + \frac{0}{2} + \frac{2}{3} + \frac{1}{4} - \frac{1}{5} + \frac{3}{6} + \frac{4}{7} \right) = \frac{1}{n} \sum_{\substack{\text{boxes} \\ (i,j)}} \frac{j - i + 1}{T_1(i, j)} = \frac{1171}{2940}.$$

3.2.2 Lifting Eigenvectors for the Random Transposition Shuffle

The eigenvalues of the random transposition shuffle were stated in Section 2.3.1 and shown to correspond to Fourier transforms of the random transposition shuffle at the irreducible representations of S_n . Given a partition λ of n we computed these to be

$$\frac{1}{n} + \frac{n-1}{n} \frac{\chi_\lambda(\tau)}{d_\lambda} = \frac{n + 2\text{Diag}(\lambda)}{n^2} \text{ with multiplicity } d_\lambda^2.$$

We show our lifting operators $\kappa_a^{\lambda, \lambda + e_i}$ may be used to lift eigenvectors of the random transposition shuffle. Thus, we recover the a description of the eigenvalues of the random transposition shuffle in terms of partitions of n . The analysis follows from the work

of Section 3.2.1, but we need to account for changes in constants needed to represent a different shuffle. Recall that in Definition 2.3.1 the random transposition shuffle was defined with driving probability RT_n . We begin our analysis by turning the random transposition shuffle into an element of the group algebra \mathfrak{S}_n .

Definition 3.2.22. The random transposition shuffle on n cards can be viewed as the following element of \mathfrak{S}_n :

$$\text{ART}_n := n^2 \sum_{\sigma \in \mathfrak{S}_n} \text{RT}_n(\sigma) \sigma = n \cdot e + 2 \sum_{1 \leq i < j \leq n} (ij). \quad (3.25)$$

We call this element the *algebraic random transposition shuffle*. Note that here we have scaled by n^2 .

We remark that the eigenvalues of RT_n may be recovered from those of ART_n , as we have done for the the random-to-random and the one-sided transposition shuffles (see Lemma 3.2.2). The random transposition shuffle also exhibits a recursive structure shown in the equation below

$$\text{ART}_{n+1} - \text{ART}_n = e + 2 \sum_{1 \leq i \leq n} (i \ n+1). \quad (3.26)$$

This allows us to lift the eigenvectors of ART_n to those of ART_{n+1} . The next result replicates Theorem 3.2.5, which relates the shuffle ART_{n+1} to a combination of ART_n and our adding and switching operators.

Theorem 3.2.23. *Given $n \in \mathbb{N}$, we have*

$$\text{ART}_{n+1} \circ \Phi_a - \Phi_a \circ \text{ART}_n = \Phi_a + 2 \sum_{1 \leq b \leq n} \Phi_a \circ \Theta_{b,a}. \quad (3.27)$$

Proof. This follows from the same proof as Theorem 3.2.5 allowing for changes in constants in equation (3.25) \square

From here the rest of the analysis in Section 3.2.1 can be followed without fundamental changes to the algebra. Keeping track of the coefficients in Theorem 3.2.23 carefully we can swiftly arrive at a modified Theorem 3.2.17 for the random transposition shuffle.

Theorem 3.2.24 (Theorem 49 of [19]). *Let $\lambda \vdash n$, and $a \in \{1, 2, \dots, l(\lambda) + 1\}$. Take $i \in [n]$ such that $1 \leq i \leq a$ and set $\mu = \lambda + e_i$. Then,*

$$\text{ART}_{n+1} \circ \kappa_a^{\lambda, \mu} - \kappa_a^{\lambda, \mu} \circ \text{ART}_n = (1 + 2(\lambda_i - i)) \kappa_a^{\lambda, \mu}. \quad (3.28)$$

In particular if $v \in S^\lambda$ is an eigenvector of ART_n with eigenvalue ε , then $\kappa_a^{\lambda, \mu}(v)$ is an eigenvector of ART_{n+1} belonging to μ with eigenvalue

$$\varepsilon + 1 + 2(\lambda_i + 1 - i) \quad (3.29)$$

The value $\lambda_i + 1 - i$ is the diagonal index of the added box $(i, \lambda_i + 1)$.

Notice that the change in eigenvalue $(1 + 2(\lambda_i + 1 - i))$ has no dependence on n . Thus, the eigenvalues we recover are not dependent on the order in which we lift our eigenvectors

(path we take up Young's lattice), in other words, we can label the eigenvalues by partitions $\lambda \vdash n$, rather than standard Young tableaux of λ . Using Theorem 3.2.24 and the proof of Theorem 3.2.18 we can recover exactly the eigenvalues we found in Section 2.3.1.

Lemma 3.2.25. *The eigenvalues of the random transposition shuffle RT_n are indexed by partitions λ . For a partition λ the corresponding eigenvalue is:*

$$\frac{n + 2\text{Diag}(\lambda)}{n^2} \text{ with multiplicity } d_\lambda^2$$

Proof. This follows from following the procedure of Theorem 3.1.3 keeping track of the new changes in eigenvalue from Theorem 3.2.24. \square

Comparison of the lifting for the one-sided transposition and random transposition shuffles allows us to see what parts of the method are coloured by our choice of shuffle. The equations in Theorems 3.2.5 and 3.2.23 are crafted to fit with our chosen shuffle, but the lifting operators $\kappa_a^{\lambda, \mu}$ remain the same in both cases. The key here is how our shuffles act on $n + 1$ versus n cards, they both show a similar structure (compare equations (3.6) and (3.26)) which allows the recovery of their eigenvectors using our lifting operators. For lifting eigenvectors to work on a generic family of transposition shuffles $\{P_n\}$ we need the difference of P_{n+1} and P_n to only involve the movement of the new card $n + 1$. This allows the adding operator Φ_a to mimic the addition of this card, and for the recovery of the eigenvalues of P_{n+1} using those known for P_n . In Section 3.6 we show this technique to be applicable to an entire class of transposition shuffles, called the biased one-sided transposition shuffles.

3.3 Eigenvalues Analysis

In this section we establish important results about the eigenvalues of the one-sided transposition shuffle. Throughout this section we will use standard facts and definitions about Young diagrams and tableaux, these were laid out in Sections 2.2.1, 2.2.2. We also need some specialised notation to deal with the eigenvalues of OST_n which we introduce now.

Definition 3.3.1. For any $\lambda \vdash n$, define the tableau T_λ^\rightarrow by inserting the numbers $1, \dots, n$ from left to right. Define the tableau T_λ^\downarrow by inserting the numbers $1, \dots, n$ from top to bottom.

Following from Section 3.2.1 we know the eigenvalues for OST_n are labelled by Young tableaux of size n , and Theorem 3.1.3 gives an explicit formula for the eigenvalue associated to any given tableau. Before applying the classical ℓ^2 bound on total variation distance, we first investigate relationships between the eigenvalues. We show that the eigenvalue corresponding to $T \in \text{SYT}(\lambda)$ is bounded by the eigenvalues for T_λ^\rightarrow and T_λ^\downarrow . To simplify our upper bound calculation, we prove that we only need to consider the partitions for which T_λ^\rightarrow gives a positive eigenvalue. Lastly, we prove that the eigenvalues corresponding to $T_\lambda^\rightarrow, T_\lambda^\downarrow$ decrease as one moves down the dominance order of partitions. We first illustrate the preceding definitions and discussion with an example.

Example 3.3.2. Let $\lambda = (3, 2) \vdash 5$. There are 5 standard Young tableaux of shape λ , these tableaux together with the associated eigenvalues are given in Table 3.1 below.

In this table, T_λ^{\rightarrow} is the first tableau listed and T_λ^{\downarrow} is the last one; we can see that the corresponding eigenvalues bound all the others.

$T \in \text{SYT}((3, 2))$	$\begin{array}{ c c c } \hline 1 & 2 & 3 \\ \hline 4 & 5 & \\ \hline \end{array}$	$\begin{array}{ c c c } \hline 1 & 2 & 4 \\ \hline 3 & 5 & \\ \hline \end{array}$	$\begin{array}{ c c c } \hline 1 & 2 & 5 \\ \hline 3 & 4 & \\ \hline \end{array}$	$\begin{array}{ c c c } \hline 1 & 3 & 4 \\ \hline 2 & 5 & \\ \hline \end{array}$	$\begin{array}{ c c c } \hline 1 & 3 & 5 \\ \hline 2 & 4 & \\ \hline \end{array}$
$\text{eig}(T)$	0.64	0.59	0.57	$0.52\bar{3}$	$0.50\bar{3}$

Table 3.1: Eigenvalues corresponding to $T \in \text{SYT}((3, 2))$.

For a Young tableau T which is not necessarily standard, define $\text{eig}(T)$ to be the value given by the formula in Theorem 3.1.3 (if T is not standard this value has no relation to the eigenvalues of OST_n). We begin our analysis by showing how swapping numbers in a tableau affects the corresponding eigenvalue. Throughout this section we accompany results with explicit examples to aid understanding of the results and their proofs.

Lemma 3.3.3. *Let T be a Young tableau. Suppose we form a new tableau S by swapping two values in T which have coordinates $(i_1, j_1), (i_2, j_2)$ in T . WLOG assume $T(i_1, j_1) < T(i_2, j_2)$. Then the change in corresponding eigenvalues satisfies the following inequality:*

$$\text{eig}(S) - \text{eig}(T) \begin{cases} \geq 0 & \text{if } (i_1 - i_2) + (j_2 - j_1) \geq 0 \\ < 0 & \text{if } (i_1 - i_2) + (j_2 - j_1) < 0. \end{cases}$$

Importantly, if we move the larger entry down and to the left the change in eigenvalue is non-negative; if it moves up and to the right then the change is negative.

Proof. Since S and T agree in all but two entries the difference in eigenvalues is given by

$$\begin{aligned} \text{eig}(S) - \text{eig}(T) &= \frac{1}{n} \left(\frac{j_1 - i_1 + 1}{T(i_2, j_2)} + \frac{j_2 - i_2 + 1}{T(i_1, j_1)} \right) - \frac{1}{n} \left(\frac{j_1 - i_1 + 1}{T(i_1, j_1)} + \frac{j_2 - i_2 + 1}{T(i_2, j_2)} \right) \\ &= \frac{(i_1 - i_2) + (j_2 - j_1)}{n} \left(\frac{1}{T(i_1, j_1)} - \frac{1}{T(i_2, j_2)} \right). \quad \square \end{aligned}$$

Example 3.3.4. Let $\lambda = (3, 3, 3)$ and take T to be the Young tableau of shape λ ,

$$T = \begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 8 & 9 & 4 \\ \hline 7 & 6 & 5 \\ \hline \end{array}.$$

To demonstrate Lemma 3.3.3 we transpose the value 9 with every other value in the tableau. Table 3.2 gives the new tableaux and the corresponding change in eigenvalue. We see that when 9 is moved down and to the left (swapping with values 6, 7, 8) the change in eigenvalue is non-negative, and when it is moved up and to the right (swapping with values 2, 3, 4) the change in eigenvalue is negative.

$S \in \text{YT}((3, 2))$	<table border="1"><tr><td>9</td><td>2</td><td>3</td></tr><tr><td>8</td><td>1</td><td>4</td></tr><tr><td>7</td><td>6</td><td>5</td></tr></table>	9	2	3	8	1	4	7	6	5	<table border="1"><tr><td>1</td><td>9</td><td>3</td></tr><tr><td>8</td><td>2</td><td>4</td></tr><tr><td>7</td><td>6</td><td>5</td></tr></table>	1	9	3	8	2	4	7	6	5	<table border="1"><tr><td>1</td><td>2</td><td>9</td></tr><tr><td>8</td><td>3</td><td>4</td></tr><tr><td>7</td><td>6</td><td>5</td></tr></table>	1	2	9	8	3	4	7	6	5	<table border="1"><tr><td>1</td><td>2</td><td>3</td></tr><tr><td>8</td><td>4</td><td>9</td></tr><tr><td>7</td><td>6</td><td>5</td></tr></table>	1	2	3	8	4	9	7	6	5	<table border="1"><tr><td>1</td><td>2</td><td>3</td></tr><tr><td>8</td><td>5</td><td>4</td></tr><tr><td>7</td><td>6</td><td>9</td></tr></table>	1	2	3	8	5	4	7	6	9	<table border="1"><tr><td>1</td><td>2</td><td>3</td></tr><tr><td>8</td><td>6</td><td>4</td></tr><tr><td>7</td><td>9</td><td>5</td></tr></table>	1	2	3	8	6	4	7	9	5	<table border="1"><tr><td>1</td><td>2</td><td>3</td></tr><tr><td>8</td><td>7</td><td>4</td></tr><tr><td>9</td><td>6</td><td>5</td></tr></table>	1	2	3	8	7	4	9	6	5	<table border="1"><tr><td>1</td><td>2</td><td>3</td></tr><tr><td>9</td><td>8</td><td>4</td></tr><tr><td>7</td><td>6</td><td>5</td></tr></table>	1	2	3	9	8	4	7	6	5
9	2	3																																																																														
8	1	4																																																																														
7	6	5																																																																														
1	9	3																																																																														
8	2	4																																																																														
7	6	5																																																																														
1	2	9																																																																														
8	3	4																																																																														
7	6	5																																																																														
1	2	3																																																																														
8	4	9																																																																														
7	6	5																																																																														
1	2	3																																																																														
8	5	4																																																																														
7	6	9																																																																														
1	2	3																																																																														
8	6	4																																																																														
7	9	5																																																																														
1	2	3																																																																														
8	7	4																																																																														
9	6	5																																																																														
1	2	3																																																																														
9	8	4																																																																														
7	6	5																																																																														
$5670 \cdot (\text{eig}(S) - \text{eig}(T))$	0	-245	-280	-87.5	0	35	40	8.75																																																																								

Table 3.2: A list of tableaux formed from T in Example 3.3.4 by transposing the value 9 with a value from $\{1, \dots, 8\}$. To aid comparison of the eigenvalues we have found the difference with $\text{eig}(T)$ and scaled by 5670.

Lemma 3.3.3 allows us to prove that the eigenvalue for any $T \in \text{SYT}(\lambda)$ is bounded between those for T_λ^\downarrow and T_λ^\rightarrow .

Lemma 3.3.5. *Let $\lambda \vdash n$. For any $T \in \text{SYT}(\lambda)$ we have the following inequality:*

$$\text{eig}(T_\lambda^\downarrow) \leq \text{eig}(T) \leq \text{eig}(T_\lambda^\rightarrow). \quad (3.30)$$

Proof. Reading across the rows of T , beginning with the first row, identify the first box in which T and T_λ^\rightarrow have different entries; write (i, j) for the coordinates of this box. Due to the way in which T_λ^\rightarrow is constructed, $T(i, j) > T_\lambda^\rightarrow(i, j)$. Furthermore, the number $T(i, j) - 1$ must occur strictly below and to the left of $T(i, j)$, since T is a standard Young tableau. Swapping entries $T(i, j) - 1$ and $T(i, j)$ in tableau T produces a new element of $\text{SYT}(\lambda)$ whose corresponding eigenvalue is no smaller than $\text{eig}(T)$, thanks to Lemma 3.3.3.

We iterate this procedure, swapping $T(i, j) - 1$ with $T(i, j) - 2$ etc, until $T(i, j) = T_\lambda^\rightarrow(i, j)$. Note that at this point the entries in the first $T(i, j)$ boxes of T and T_λ^\rightarrow must agree, moreover these entries are now fixed in place. We now proceed to the next box in which T and T_λ^\rightarrow differ, and repeat: this results in a sequence of swaps which make the entries of T agree with those in T_λ^\rightarrow , and which can only ever cause the corresponding eigenvalue to increase. This proves the second inequality in Lemma 3.3.5, and the first one follows via an analogous argument on the columns of T this time with our eigenvalue decreasing after each iteration. □

Example 3.3.6. Let $\lambda = (3, 2, 1, 1)$, to illustrate the procedure defined in Lemma 3.3.5, we show the first two iterations of the algorithm for turning T_λ^\downarrow into T_λ^\rightarrow . The first two iterations involve transforming the boxes $(1, 2)$ and $(1, 3)$ in T_λ^\downarrow to match the boxes of $(1, 2)$ and $(1, 3)$ in T_λ^\rightarrow (box $(1, 1)$ is always matched for any pair of standard Young tableaux). We keep track of the eigenvalue of the tableaux at each step of the algorithm to show it is always increasing. Applying transpositions in order to fix box $(1, 2)$ we find,

<table border="1" style="display: inline-table; border-collapse: collapse;"><tr><td>1</td><td>5</td><td>7</td><td>8</td></tr><tr><td>2</td><td>6</td><td></td><td></td></tr><tr><td>3</td><td></td><td></td><td></td></tr><tr><td>4</td><td></td><td></td><td></td></tr></table>	1	5	7	8	2	6			3				4				→	<table border="1" style="display: inline-table; border-collapse: collapse;"><tr><td>1</td><td>4</td><td>7</td><td>8</td></tr><tr><td>2</td><td>6</td><td></td><td></td></tr><tr><td>3</td><td></td><td></td><td></td></tr><tr><td>5</td><td></td><td></td><td></td></tr></table>	1	4	7	8	2	6			3				5				→	<table border="1" style="display: inline-table; border-collapse: collapse;"><tr><td>1</td><td>3</td><td>7</td><td>8</td></tr><tr><td>2</td><td>6</td><td></td><td></td></tr><tr><td>4</td><td></td><td></td><td></td></tr><tr><td>5</td><td></td><td></td><td></td></tr></table>	1	3	7	8	2	6			4				5				→	<table border="1" style="display: inline-table; border-collapse: collapse;"><tr><td>1</td><td>2</td><td>7</td><td>8</td></tr><tr><td>3</td><td>6</td><td></td><td></td></tr><tr><td>4</td><td></td><td></td><td></td></tr><tr><td>5</td><td></td><td></td><td></td></tr></table>	1	2	7	8	3	6			4				5			
1	5	7	8																																																																			
2	6																																																																					
3																																																																						
4																																																																						
1	4	7	8																																																																			
2	6																																																																					
3																																																																						
5																																																																						
1	3	7	8																																																																			
2	6																																																																					
4																																																																						
5																																																																						
1	2	7	8																																																																			
3	6																																																																					
4																																																																						
5																																																																						
0.207738		0.232738		0.263988		0.305655																																																																

Notice that after each swap we still remain at a standard Young tableau. Once box $(1, 2)$ is fixed we proceed to fix box $(1, 3)$ in the same way

1	2	7	8	→	1	2	6	8	→	1	2	5	8	→	1	2	4	8	→	1	2	3	8
3	6	3	7		3	7	4			3	7	5			4	7	5			4	7	5	
4	5	4	5		4	6	3	6		5	6	2	6		5	6	7	8		5	6	7	8
0.305655					0.311607					0.33244					0.35744					0.38869			

If we repeat this process until all boxes (i, j) match T_λ^\rightarrow , we only ever increase the eigenvalue and thus find an upper bound. If instead we wanted to turn T_λ^\rightarrow into T_λ^\downarrow we would first fix box $(2, 1)$ like so:

1	2	3	4	→	1	2	3	5	→	1	2	4	5	→	1	3	4	5
5	6	4	6		3	6	7			3	6	2	6		2	6	7	8
7	8	7	8		7	8	7	8		7	8	7	8		7	8	7	8
0.471726					0.446726					0.415476					0.37381			

Now the eigenvalue is decreasing after each step, if we fix every column in turn we eventually end at T_λ^\downarrow .

The next result and its corollary establish that when bounding eigenvalues, we only need to consider those given by T_λ^\rightarrow .

Lemma 3.3.7. *Let $\lambda \vdash n$. For any $T \in \text{SYT}(\lambda)$ we have*

$$\text{eig}(T) + \text{eig}(T') = \frac{2H_n}{n}.$$

Proof. Let $T \in \text{SYT}(\lambda)$. Then

$$\begin{aligned} \text{eig}(T) + \text{eig}(T') &= \frac{1}{n} \sum_{\substack{\text{boxes} \\ (i,j) \in T}} \frac{j-i+1}{T(i,j)} + \frac{1}{n} \sum_{\substack{\text{boxes} \\ (j,i) \in T'}} \frac{i-j+1}{T'(j,i)} \\ &= \frac{1}{n} \sum_{\substack{\text{boxes} \\ (i,j) \in T}} \frac{j-i+1}{T(i,j)} + \frac{1}{n} \sum_{\substack{\text{boxes} \\ (i,j) \in T}} \frac{-(j-i)+1}{T(i,j)} = \frac{2H_n}{n}. \quad \square \end{aligned}$$

Corollary 3.3.8. *Let $\lambda \vdash n$, and suppose we have $\text{eig}(T_\lambda^\downarrow) \leq 0$, then we have*

$$\text{eig}(T_\lambda^\rightarrow) \geq |\text{eig}(T_\lambda^\downarrow)| \geq 0. \quad (3.31)$$

Proof. It follows from Lemma 3.3.7 that $\text{eig}(T_\lambda^\rightarrow) + \text{eig}(T_\lambda^\downarrow) = 2H_n/n$. Thus if $\text{eig}(T_\lambda^\downarrow) \leq 0$ then

$$\text{eig}(T_\lambda^\rightarrow) = \frac{2H_n}{n} - \text{eig}(T_\lambda^\downarrow) \geq -\text{eig}(T_\lambda^\downarrow) = |\text{eig}(T_\lambda^\downarrow)| \geq 0. \quad \square \quad (3.32)$$

We end this section by establishing a relationship between eigenvalues and the dominance ordering on partitions.

Lemma 3.3.9. *Let $\lambda, \mu \vdash n$. If $\lambda \triangleright \mu$ then*

$$\text{eig}(T_\lambda^\rightarrow) \geq \text{eig}(T_\mu^\rightarrow) \quad (3.33)$$

$$\text{and } \text{eig}(T_\lambda^\downarrow) \geq \text{eig}(T_\mu^\downarrow). \quad (3.34)$$

Proof. If we can show the statements hold for any partition μ which is formed from λ by moving only one box then inductively it will hold for all $\lambda \succeq \mu$. Suppose μ is formed from λ by moving a box from row a to row b , with $a < b \leq l(\lambda) + 1$ (if $b = l(\lambda) + 1$ then a new row is created by placing the removed box on the very bottom of the diagram). The box we move goes from coordinates (a, λ_a) of λ to $(b, \lambda_b + 1)$ in μ .

We begin by proving $\text{eig}(T_\lambda^\rightarrow) \geq \text{eig}(T_\mu^\rightarrow)$. Since T_λ^\rightarrow and T_μ^\rightarrow are both numbered from left to right, the effect of moving a box from row a to row b is that $T_\mu^\rightarrow(i, j) = T_\lambda^\rightarrow(i, j) - 1$ for any box $(i, j) \in T_\lambda^\rightarrow \cap T_\mu^\rightarrow$ with $a < i \leq b$; boxes in all other rows contain the same values in both tableaux. Using equation (3.3), and remembering to include a term to account for the box being moved, we find that:

$$\begin{aligned} n(\text{eig}(T_\lambda^\rightarrow) - \text{eig}(T_\mu^\rightarrow)) &= \left(\frac{\lambda_a - a + 1}{T_\lambda^\rightarrow(a, \lambda_a)} - \frac{(\lambda_b + 1) - b + 1}{T_\mu^\rightarrow(b, \lambda_b + 1)} \right) + \sum_{\substack{(i,j) \in T_\lambda^\rightarrow \cap T_\mu^\rightarrow \\ \text{with } a < i \leq b}} \left[\frac{1}{T_\lambda^\rightarrow(i, j)} - \frac{1}{T_\mu^\rightarrow(i, j)} \right] (j - i + 1) \\ &\geq \left(\frac{\lambda_a - a + 1}{T_\lambda^\rightarrow(a, \lambda_a)} - \frac{(\lambda_b + 1) - b + 1}{T_\mu^\rightarrow(b, \lambda_b + 1)} \right) + (\lambda_a - a + 1) \left(\frac{1}{T_\mu^\rightarrow(b, \lambda_b + 1)} - \frac{1}{T_\lambda^\rightarrow(a, \lambda_a)} \right) \\ &= \frac{(\lambda_a - \lambda_b) + (b - a) - 1}{T_\mu^\rightarrow(b, \lambda_b + 1)} \geq 0. \end{aligned}$$

The first inequality holds because all the square-bracketed terms in the sum are negative; we upper bound $j - i + 1 \leq \lambda_a - a + 1$, and the resulting sum telescopes. The final inequality holds because $(\lambda_a - \lambda_b) \geq 1$ and $(b - a) \geq 1$.

For the inequality (3.34), recall that $\lambda \succeq \mu$ if and only if $\mu' \succeq \lambda'$. Therefore, using inequality (3.33) we find that $\text{eig}(T_{\mu'}^\rightarrow) \geq \text{eig}(T_{\lambda'}^\rightarrow)$. Now Lemma 3.3.7 gives $-\text{eig}(T_\mu^\downarrow) \geq -\text{eig}(T_\lambda^\downarrow)$ and thus we recover the desired inequality. \square

Example 3.3.10. The dominance ordering on partitions of size 4 is a linear order. Table 3.3 lists T_λ^\rightarrow and T_λ^\downarrow for every partition $\lambda \vdash 4$. Reading from the table we can see the eigenvalues of T_λ^\rightarrow and T_λ^\downarrow decrease down the dominance ordering. Also we may use the eigenvalues in the table to verify that $\text{eig}(T) + \text{eig}(T') = 2H_4/4 = 25/24$.

T_λ^\rightarrow	$\begin{array}{ c c c c } \hline 1 & 2 & 3 & 4 \\ \hline \end{array}$	$\begin{array}{ c c c } \hline 1 & 2 & 3 \\ \hline 4 \\ \hline \end{array}$	$\begin{array}{ c c } \hline 1 & 2 \\ \hline 3 & 4 \\ \hline \end{array}$	$\begin{array}{ c c } \hline 1 & 2 \\ \hline 3 \\ \hline 4 \\ \hline \end{array}$	$\begin{array}{ c } \hline 1 \\ \hline 2 \\ \hline 3 \\ \hline 4 \\ \hline \end{array}$
$24 \cdot \text{eig}(T_\lambda^\rightarrow)$	24	18	13.5	10.5	1
T_λ^\downarrow	$\begin{array}{ c c c c } \hline 1 & 2 & 3 & 4 \\ \hline \end{array}$	$\begin{array}{ c c c } \hline 1 & 3 & 4 \\ \hline 2 \\ \hline \end{array}$	$\begin{array}{ c c } \hline 1 & 3 \\ \hline 2 & 4 \\ \hline \end{array}$	$\begin{array}{ c c } \hline 1 & 4 \\ \hline 2 \\ \hline 3 \\ \hline \end{array}$	$\begin{array}{ c } \hline 1 \\ \hline 2 \\ \hline 3 \\ \hline 4 \\ \hline \end{array}$
$24 \cdot \text{eig}(T_\lambda^\downarrow)$	24	14.5	11.5	7	1

Table 3.3: The eigenvalues $T_\lambda^\rightarrow, T_\lambda^\downarrow$ for every partition $\lambda \vdash 4$. The dominance order follows from left to right. We have scaled the values by 24 for ease of comparison

3.4 Upper Bound for the One-sided Transposition Shuffle

In this section we complete the proof of the upper bound present in Theorem 3.1.2, making use of the results of Section 3.3. The analysis splits into two parts, dealing separately with those partitions λ having either *large* or *small* first row.

Theorem 1.1.27 allows us to upper bound the total variation distance in terms of the non-trivial eigenvalues of the transition matrix. Using Theorem 3.1.3 we see that the trivial eigenvalue corresponds to the one-dimensional partition $\lambda = (n)$, and so Theorem 3.1.3 implies that

$$4\|\text{OST}_n^t - \pi_n\|_{\text{TV}}^2 \leq \sum_{\substack{\lambda \vdash n \\ \lambda \neq (n)}} d_\lambda \sum_{T \in \text{SYT}(\lambda)} \text{eig}(T)^{2t}. \quad (3.35)$$

Recall from Lemma 3.3.5 that for any $T \in \text{SYT}(\lambda)$ the eigenvalue corresponding to T may be bounded by those corresponding to T_λ^\downarrow and T_λ^\rightarrow . With this in mind, we let $\Lambda_n^\rightarrow = \{\lambda \vdash n : |\text{eig}(T_\lambda^\downarrow)| \leq |\text{eig}(T_\lambda^\rightarrow)|\}$ and $\Lambda_n^\downarrow = \{\lambda \vdash n : |\text{eig}(T_\lambda^\downarrow)| > |\text{eig}(T_\lambda^\rightarrow)|\}$; note that these are disjoint sets, with $\Lambda_n^\rightarrow \subseteq \{\lambda \vdash n : \text{eig}(T_\lambda^\rightarrow) \geq 0\}$ and $\Lambda_n^\downarrow \subseteq \{\lambda \vdash n : \text{eig}(T_\lambda^\downarrow) < 0\}$. Using Lemma 3.3.5 and then Corollary 3.3.8 we relax the upper bound as follows:

$$\begin{aligned} 4\|\text{OST}_n^t - \pi_n\|_{\text{TV}}^2 &\leq \text{eig}(T_{(1^n)})^{2t} + \sum_{\substack{\lambda \in \Lambda_n^\rightarrow \\ \lambda \neq (n)}} d_\lambda \sum_{T \in \text{SYT}(\lambda)} \text{eig}(T)^{2t} + \sum_{\substack{\lambda \in \Lambda_n^\downarrow \\ \lambda \neq (1^n)}} d_\lambda \sum_{T \in \text{SYT}(\lambda)} \text{eig}(T)^{2t} \\ &\leq \text{eig}(T_{(1^n)})^{2t} + \sum_{\substack{\lambda \in \Lambda_n^\rightarrow \\ \lambda \neq (n)}} d_\lambda^2 \text{eig}(T_\lambda^\rightarrow)^{2t} + \sum_{\substack{\lambda \in \Lambda_n^\downarrow \\ \lambda \neq (1^n)}} d_\lambda^2 \text{eig}(T_\lambda^\downarrow)^{2t} \\ &\leq \text{eig}(T_{(1^n)})^{2t} + \sum_{\substack{\lambda : \text{eig}(T_\lambda^\rightarrow) \geq 0 \\ \lambda \neq (n)}} d_\lambda^2 \text{eig}(T_\lambda^\rightarrow)^{2t} + \sum_{\substack{\lambda : \text{eig}(T_\lambda^\downarrow) < 0 \\ \lambda \neq (1^n)}} d_\lambda^2 \text{eig}(T_\lambda^\downarrow)^{2t} \\ &\leq \text{eig}(T_{(1^n)})^{2t} + \sum_{\substack{\lambda : \text{eig}(T_\lambda^\rightarrow) \geq 0 \\ \lambda \neq (n)}} d_\lambda^2 \text{eig}(T_\lambda^\rightarrow)^{2t} + \sum_{\substack{\lambda : \text{eig}(T_\lambda^\downarrow) < 0 \\ \lambda \neq (1^n)}} d_{\lambda'}^2 \text{eig}(T_{\lambda'}^\rightarrow)^{2t} \\ &\leq \text{eig}(T_{(1^n)})^{2t} + 2 \sum_{\substack{\lambda : \text{eig}(T_\lambda^\rightarrow) \geq 0 \\ \lambda \neq (n)}} d_\lambda^2 \text{eig}(T_\lambda^\rightarrow)^{2t}. \end{aligned} \quad (3.36)$$

In the penultimate line we have used Corollary 3.3.8 and the fact that $d_{\lambda'} = d_\lambda$. The final inequality follows by a second application of Corollary 3.3.8: if λ satisfies $\text{eig}(T_\lambda^\downarrow) < 0$ then $\text{eig}(T_{\lambda'}^\rightarrow)$ must be non-negative.

The first term in (3.36) is simple to deal with at time $t = n \log(n) + cn$. We have already observed that $\text{eig}(T_{(n)}) = 1$, and so Lemma 3.3.8 implies that $\text{eig}(T_{(1^n)}) = 2H_n/n - 1$. This means that

$$\text{eig}(T_{(1^n)})^{2t} = \left(1 - \frac{2H_n}{n}\right)^{2(n \log n + cn)} \leq e^{-4H_n(\log + c)} \text{ for } n \geq 5 \quad (3.37)$$

here we have used the bound $1 - x \leq e^{-x}$, we see that this tends to zero for any positive c as $n \rightarrow \infty$.

It therefore remains to bound the sum in (3.36). The partitions with the biggest eigenvalues are those with large first rows λ_1 , and so we split the analysis into two

parts according to this value; by *large* partitions we mean those with $\lambda_1 \geq 3n/4$, and *small* partitions are those with $\lambda_1 < 3n/4$. Large partitions give the biggest eigenvalues for OST_n and must be dealt with carefully; it is these which will determine the mixing time of the shuffle. Small partitions have correspondingly larger dimensions, but eigenvalues which are small enough to give control around time of order $n \log(n)$. We begin by identifying the partition at the top of the dominance ordering for any fixed value of λ_1 , which allows us to employ Lemma 3.3.9.

Definition 3.4.1. Let $k \in [n]$, define the partition $(n-k, \star)$ to have as many rows of $n-k$ as possible, with the last row being formed of $n - (n-k)\lfloor n/(n-k) \rfloor$ boxes. For example,

$$(n-k, \star) = \begin{cases} (n-k, k) & \text{if } k \in [1, n/2] \\ (n-k, n-k, 2k-n) & \text{if } k \in (n/2, n/3] \end{cases}$$

Lemma 3.4.2. Suppose $\lambda \vdash n$ has first row equal to $\lambda_1 = n-k$, then by moving boxes up and to the right it follows trivially that $\lambda \triangleleft (n-k, \star)$.

The notation of $(n-k, \star)$ is driven by our analysis as it will transpire that only the size of the first two rows are important for our bounds. For each k we also need a bound on sum of the squared dimensions of all partitions with $\lambda_1 = n-k$, and for this recall that Lemma 2.3.5 tells us

$$\sum_{\substack{\lambda \vdash n \\ \lambda_1 = n-k}} d_\lambda^2 \leq \binom{n}{k}^2 k! \leq \frac{n^{2k}}{k!}.$$

The Eigenvalues of Partition $(n-1, 1)$

Before we proceed with the analysis of large and small partitions let us look at the partition $(n-1, 1)$ in more detail. This partition gives the first and largest term in the sum in (3.36), and controls the mixing time of the one-sided transposition shuffle.

Label the different standard Young tableau of shape $(n-1, 1)$ by T_i , where i denotes the value in the second row of $(n-1, 1)$. The eigenvalue of T_i is given by,

$$\text{eig}(T_i) = 1 - \frac{1}{n} \left(1 + \sum_{j=i+1}^n \frac{1}{j} \right)$$

We have $\text{eig}(T_n) = \text{eig}(T_{(n-1,1)}^\rightarrow)$, and a clear linear order on the eigenvalues given by

$$\text{eig}(T_n) > \text{eig}(T_{n-1}) > \dots > \text{eig}(T_3) > \text{eig}(T_2).$$

The tableau T_n gives us the second biggest (in absolute value) eigenvalue of the one-sided transposition shuffle. The contribution of this tableau to the sum (3.35) is,

$$(n-1) \left(1 - \frac{1}{n} \right)^{2t} \tag{3.38}$$

and we have seen previously that this eigenvalue is killed at time $t = (n/2) \log n$. We have said before that the time needed to kill the second biggest eigenvalue often tells us the time expected for a random walk to converge to its stationary distribution. For the one-sided

transposition shuffle the eigenvalue T_n is tightly grouped around the other eigenvalues for partition $(n-1, 1)$, with the biggest difference being $T_n - T_2 = (H_n - 1)/n$. This means that in order to estimate the mixing time it is not good enough to consider the single eigenvalue (3.38); rather, we must consider the sum,

$$\sum_{i=2}^n (n-1) \text{eig}(T_i)^{2t} = (n-1) \sum_{i=2}^n \left(1 - \frac{1}{n} \left(1 + \sum_{j=i+1}^n \frac{1}{j} \right) \right)^{2t}. \quad (3.39)$$

This sum is bounded at time close to $n \log n$ as $n \rightarrow \infty$. Therefore, we expect the mixing time of the one-sided transposition shuffle to be around $n \log n$, as stated in Theorem 3.1.2. Note that in sum (3.36) we reduce the sum (3.39) of eigenvalues for $(n-1, 1)$, to

$$(n-1)^2 \left(1 - \frac{1}{n} \right)^{2t}$$

which is bounded in n exactly at time $t = n \log n$.

3.4.1 Large Partitions

In this subsection we prove that the sum of large partitions is bounded with a decay of e^{-2c} at time $n \log n + cn$. Let λ be a partition satisfying $\text{eig}(T_\lambda^\rightarrow) \geq 0$, and for which $\lambda_1 = n - k$ for some $k \leq n/4$. We have observed that $\lambda \triangleleft (n-k, k)$, and so Lemma 3.3.9 suggests that we look at the eigenvalue of $T_{(n-k, k)}^\rightarrow$. Using our eigenvalue formula from Theorem 3.1.3 we calculate this as follows, with the first and second sum corresponding to the first and second row of $T_{(n-k, k)}^\rightarrow$ respectively:

$$\begin{aligned} \text{eig}(T_{(n-k, k)}^\rightarrow) &= \frac{1}{n} \sum_{j=1}^{n-k} \frac{j}{T_{(n-k, k)}^\rightarrow(1, j)} + \frac{1}{n} \sum_{j=1}^k \frac{j-1}{T_{(n-k, k)}^\rightarrow(2, j)} \\ &= \frac{n-k}{n} + \frac{1}{n} \sum_{j=1}^k \frac{j-1}{n-k+j} \end{aligned} \quad (3.40)$$

$$= 1 - \frac{(n-k+1)}{n} (H_n - H_{n-k+1}) - \frac{1}{n}. \quad (3.41)$$

We now use this, along with the inequality $1 - x \leq e^{-x}$, to bound the contribution of large partitions to the sum in (3.36):

$$\begin{aligned} \sum_{k=1}^{n/4} \sum_{\substack{\lambda: \text{eig}(T_\lambda^\rightarrow) \geq 0 \\ \lambda_1 = n-k}} d_\lambda^2 \text{eig}(T_\lambda^\rightarrow)^{2t} &\leq \sum_{k=1}^{n/4} \text{eig}(T_{(n-k, k)}^\rightarrow)^{2t} \sum_{\substack{\lambda: \text{eig}(T_\lambda^\rightarrow) \geq 0 \\ \lambda_1 = n-k}} d_\lambda^2 \\ &\leq \sum_{k=1}^{n/4} \frac{n^{2k}}{k!} \text{eig}(T_{(n-k, k)}^\rightarrow)^{2t} \quad (\text{by Lemma 2.3.5}) \\ &\leq \sum_{k=1}^{n/4} \frac{n^{2k}}{k!} \left(1 - \frac{(n-k+1)}{n} (H_n - H_{n-k+1}) - \frac{1}{n} \right)^{2t} \\ &\leq \sum_{k=1}^{n/4} \frac{n^{2k}}{k!} e^{-2t \left(\frac{(n-k+1)}{n} (H_n - H_{n-k+1}) + \frac{1}{n} \right)} \\ &\leq e^{-2c} \sum_{k=1}^{n/4} \frac{n^{2k-2(n-k+1)(H_n - H_{n-k+1})-2}}{k!}, \end{aligned} \quad (3.42)$$

in the last step we have substituted $t = n \log n + cn$. When $k = 1$ we get the following term

$$\left(n^{2-2(n)(H_n-H_n)-2} \right) / 1! = 1.$$

Thus if the ratio between consecutive terms is less than 1 for n suitably large we may bound the sum via a geometric series. The ratio of the $(k+1)^{\text{th}}$ term to the k^{th} term in (3.42) is given by

$$\frac{n^{2(H_n-H_{n-k})}}{k+1}. \quad (3.43)$$

For large n this ratio is approximated by $(n^{2 \log(n/n-k)}) / (k+1)$. To see that this ratio is less than 1 for all $k \in \{1, \dots, n/4\}$ as $n \rightarrow \infty$ we consider the two cases, $k = O(1)$ and $k = O(n)$. In both cases $\lim_{n \rightarrow \infty} n^{2(H_n-H_{n-k})} = \lim_{n \rightarrow \infty} n^{2 \log n / (n-k)}$. Now if k is constant then we have $\lim_{n \rightarrow \infty} n^{2 \log n / (n-k)} = 1$ and so

$$\lim_{n \rightarrow \infty} (n^{2 \log(n/n-k)}) / (k+1) = \frac{1}{k+1}.$$

If $k = an$ with $a \in (0, 1/4]$ then $n^{2 \log(n/(n-an))} = n^{2 \log 1 / (1-a)} \leq n^{2 \log(4/3)} = n^{0.58}$. Taking the limit in n we have,

$$\lim_{n \rightarrow \infty} (n^{2 \log(n/n-k)}) / (k+1) \leq \lim_{n \rightarrow \infty} n^{0.58-1} / a = 0.$$

Therefore, for large enough n the ratio (3.43) is less than 1 for all $k \in \{1, \dots, n/4\}$. Indeed for $n \rightarrow \infty$ the largest value of the ratio over this range of k is achieved when $k = 1$, at which point it equals $n^{2/n} / 2$. For sufficiently large n this ratio is thus bounded above by $3/4$, say, which permits us to bound the sum in (3.42) by a geometric series with initial term 1:

$$e^{-2c} \sum_{k=1}^{n/4} \frac{n^{2k-2(n-k+1)(H_n-H_{n-k+1})-2}}{k!} \leq e^{-2c} \sum_{k=1}^{n/4} (3/4)^{k-1} \leq 4e^{-2c}. \quad (3.44)$$

3.4.2 Small Partitions

Now consider a partition λ satisfying $\text{eig}(T_\lambda^{\rightarrow}) \geq 0$ and for which $\lambda_1 = n - k$ with $n/4 < k \leq n - 2$. Suppose first of all that $n/4 < k \leq n/2$; as in the large partition case, any such partition is dominated by $(n - k, k)$, and the same calculation as in equation (3.40) shows that

$$\text{eig}\left(T_{(n-k,k)}^{\rightarrow}\right) = \frac{n-k}{n} + \frac{1}{n} \sum_{j=1}^k \frac{j-1}{n-k+j}. \quad (3.45)$$

Now consider the case when $k > n/2$. We have already identified that $\lambda \triangleleft (n - k, \star)$, and so we proceed by calculating the eigenvalue of $T_{(n-k,\star)}^{\rightarrow}$. Note first that for any box on row three or below – that is (i, j) with $i \geq 3$ – its contribution to equation (3.3) may be bounded by:

$$\frac{j-i+1}{T_{(n-k,\star)}^{\rightarrow}(i,j)} = \frac{j-i+1}{(i-1)(n-k)+j} \leq \frac{(n-k)}{(i-1)(n-k)+(n-k)} \leq \frac{1}{3}.$$

Using this inequality in conjunction with Theorem 3.1.3 we bound $\text{eig}(T_{(n-k,*)}^{\rightarrow})$ as follows:

$$\begin{aligned} \text{eig}(T_{(n-k,*)}^{\rightarrow}) &= \frac{1}{n} \sum_{j=1}^{n-k} \frac{j}{T_{(n-k,*)}^{\rightarrow}(1, j)} + \frac{1}{n} \sum_{j=1}^{n-k} \frac{j-1}{T_{(n-k,*)}^{\rightarrow}(2, j)} + \frac{1}{n} \sum_{\substack{(i,j) \\ i \geq 3}} \frac{j-i+1}{T_{(n-k,*)}^{\rightarrow}(i, j)} \\ &\leq \frac{n-k}{n} + \frac{1}{n} \sum_{j=1}^{n-k} \frac{j-1}{n-k+j} + \frac{n-2(n-k)}{3n}. \end{aligned} \quad (3.46)$$

Observe that if we substitute $k \in (n/4, n/2]$ in (3.46) it provides an upper bound for the expression in (3.45). Indeed, for $n/4 < k \leq n/2$ we may write

$$\begin{aligned} \frac{n-k}{n} + \frac{1}{n} \sum_{j=1}^{n-k} \frac{j-1}{n-k+j} + \frac{n-2(n-k)}{3n} - \text{eig}(T_{(n-k,k)}^{\rightarrow}) &= \frac{1}{n} \sum_{j=k+1}^{n-k} \left(\frac{j-1}{n-k+j} - \frac{1}{3} \right) \\ &= \frac{2(n-2k)}{3n} - \frac{(n-k+1)}{n} (H_{2(n-k)} - H_n). \end{aligned}$$

Substituting $k = \gamma n$, the final expression is bounded below by

$$\frac{2(1-2\gamma)}{3} - \left(1 - \gamma + \frac{1}{n} \right) \log(2(1-\gamma)). \quad (3.47)$$

For $n \geq 15$ the expression (3.47) is non-negative for all $\gamma \in [1/4, 1/2]$ thus completing our claim. We have just shown that (3.46) provides a bound on $\text{eig}(T_{(n-k,*)}^{\rightarrow})$ for all $k \in (n/4, n-2]$, therefore, it provides a bound for all $\text{eig}(T_{\lambda}^{\rightarrow})$ with $\lambda_1 = n-k$. Working with our new bound we rearrange it to:

$$\begin{aligned} \text{eig}(T_{\lambda}^{\rightarrow}) &\leq \frac{n-k}{n} + \frac{1}{n} \sum_{j=1}^{n-k} \frac{j-1}{n-k+j} + \frac{n-2(n-k)}{3n} \\ &= \frac{n-k}{n} + \frac{n-k-1 - (n-k+1)(H_{2(n-k)} - H_{n-k+1})}{n} + \frac{2k-n}{3n} \\ &= 1 - \frac{(4k-2n+3)}{3n} - \frac{(n-k+1)}{n} (H_{2(n-k)} - H_{n-k+1}). \end{aligned} \quad (3.48)$$

Using the inequalities $1-x \leq e^{-x}$ for all x , and, our eigenvalue bound (3.48), we are able to bound the contributions of small partitions in the sum (3.36) at time $t = n \log n + cn$ as follows:

$$\begin{aligned} \sum_{k=n/4}^{n-2} \sum_{\substack{\lambda: \text{eig}(T_{\lambda}^{\rightarrow}) \geq 0 \\ \lambda_1 = n-k}} d_{\lambda}^2 \text{eig}(T_{\lambda}^{\rightarrow})^{2t} &\leq \sum_{k=n/4}^{n-2} \frac{n^{2k}}{k!} e^{-\frac{2t}{n} \left(\frac{4k-2n+3}{3} + (n-k+1)(H_{2(n-k)} - H_{n-k+1}) \right)} \\ &\leq e^{-2c} \sum_{k=n/4}^{n-2} \frac{n^{\frac{4n-2k-6}{3} - 2(n-k+1)(H_{2(n-k)} - H_{n-k+1})}}{k!}. \end{aligned} \quad (3.49)$$

To analyse this bound further we require the following inequality.

Lemma 3.4.3. $(n+1)(H_{2n} - H_{n+1}) > (n-1) \log 2$ for all integers $n \geq 1$

Proof. Rearranging our inequality we have to show that $(n+1)(H_{2n} - H_{n+1} - \log 2) > -2 \log 2$. The n^{th} harmonic number may be bounded by $\log n + \gamma + \frac{1}{2n+1} \leq H_n \leq \log n + \gamma + \frac{1}{2n-1}$ where γ is the Euler–Mascheroni constant. Using the lower bound for H_{2n} and upper

bound for H_{n+1} , we have:

$$\begin{aligned}
(n+1)(H_{2n} - H_{n+1} - \log 2) &\geq (n+1) \left(\frac{1}{4n+1} - \frac{1}{2n+1} + \log \left(\frac{n}{n+1} \right) \right) \\
&= -\frac{2n(n+1)}{(4n+1)(2n+1)} - (n+1) \log \left(\frac{n+1}{n} \right) \\
&\geq -\frac{2n(n+1)}{(4n+1)(2n+1)} - \frac{(n+1)}{n} \\
&\geq -2 \log 2 \text{ for all } n \geq 8.
\end{aligned}$$

The last inequality comes from $\frac{2n(n+1)}{(4n+1)(2n+1)} + \frac{(n+1)}{n}$ being a decreasing function, and $n = 8$ is the first time it passes $2 \log 2$. The original inequality may be verified for the remaining integer values $1 \leq n \leq 7$ \square

Using the new bound provided by Lemma 3.4.3 we may bound (3.49) via the following:

$$e^{-2c} \sum_{k=n/4}^{n-2} \frac{n^{\frac{4n-2k-6}{3} - 2(n-k-1) \log 2}}{k!}. \quad (3.50)$$

Once again writing $k = \gamma n$, now for $\gamma \in [1/4, 1]$, the terms in the summand of (3.50) may be rewritten as

$$n^{\frac{2n(2-\gamma)}{3} - 2n(1-\gamma) \log 2 - 0.5 + 2 \log 2} / (\gamma n!) \quad (3.51)$$

An application of Stirling's formula to $\gamma n!$ tells us that $\gamma n! > \sqrt{2\pi} (\gamma n)^{\gamma n + 0.5} e^{-\gamma n}$. Combining this with equation (3.51) we get

$$\frac{e^{\gamma n} n^{2 \log 2 - 1}}{\gamma^{\gamma n} \sqrt{2\pi}} n^{\frac{n}{3}(4 - 5\gamma - 6(1-\gamma) \log 2)} \quad (3.52)$$

Thus the dominant term of (3.52) takes the form $n^{\frac{n}{3}g(\gamma)}$, where $g(\gamma) = 4 - 5\gamma - 6(1 - \gamma) \log 2 < 0$ for all $\gamma \in [1/4, 1]$. It follows that, for any positive c ,

$$\lim_{n \rightarrow \infty} e^{-2c} \sum_{k=n/4}^{n-2} \frac{n^{\frac{4n-2k-6}{3} - 2(n-k-1) \log 2}}{k!} = 0 \quad (3.53)$$

This completes the analysis of small partitions.

Proof of the Upper Limit in Theorem 3.1.2

Combining the results and bounds of (3.36), (3.37), (3.44) and (3.53) we find at time $t = n \log n + cn$,

$$\begin{aligned}
4 \limsup_{n \rightarrow \infty} \|\text{OST}_n^t - \pi_n\|_{\text{TV}}^2 &\leq \limsup_{n \rightarrow \infty} \text{eig}(T_{(1^n)})^{2t} + 2 \limsup_{n \rightarrow \infty} \sum_{\substack{\lambda: \text{eig}(T_\lambda^\rightarrow) \geq 0 \\ \lambda \neq (n)}} d_\lambda^2 \text{eig}(T_\lambda^\rightarrow)^{2t} \\
&\leq 2 \limsup_{n \rightarrow \infty} \sum_{k=1}^{n/4} d_\lambda^2 \text{eig}(T_{(n-k, k)}^\rightarrow)^{2t} + 2 \limsup_{n \rightarrow \infty} \sum_{k=n/4}^{n-2} d_\lambda^2 \text{eig}(T_{(n-k, \star)}^\rightarrow)^{2t} \\
&\leq 8e^{-2c}
\end{aligned}$$

This completes the proof that $\limsup_{n \rightarrow \infty} \|\text{OST}_n^{\log n + cn} - \pi_n\|_{\text{TV}} \leq \sqrt{2}e^{-c}$. Thus, we have an upper bound on the mixing time of the one-sided transposition shuffle of time $n \log n$.

3.5 Lower Bound for the One-sided Transposition Shuffle

To complete Theorem 3.1.2 we need to prove our lower limit on total variation distance. To do this we employ the usual trick of finding a set of permutations $F_n \subseteq S_n$ which has significantly different probability under the uniform distribution π_n and the one-sided transposition measure OST_n^t before a certain time t . The definition of total variation distance then immediately yields a simple lower bound:

$$\|\text{OST}_n^t - \pi_n\|_{\text{TV}} \geq \text{OST}_n^t(F_n) - \pi_n(F_n).$$

In particular, we follow in the steps of the random transposition shuffle from Section 2.3.1 and find a suitable set F_n by considering the number of fixed points within - a certain part of - the deck. Estimation of $\text{OST}_n^t(F_n)$ then reduces to a novel variant of the classical coupon collector's problem.

Recall that one step of the one-sided transposition shuffle may be modelled by firstly choosing a position $R^i \sim_d U\{1, \dots, n\}$ with our right hand, and then choosing a position $L^i \sim_d U\{1, \dots, R^i\}$ with our left hand and transposing the cards in the chosen positions. Since the left hand always chooses a position below that of the right hand, it is intuitively clear that our shuffle is relatively unlikely to transpose two cards near to the top of the deck. For example, taking $n > 3$, $\text{OST}_n((12)) = 1/2n$, whereas $\text{OST}_n((n-1n)) = 1/n^2$, therefore we are $(n/2)$ times as likely to apply transposition (12) , than $(n-1n)$.

This leads us to focus the attention of our analysis on a set of positions at the top of the deck: write V_n for the top part of the deck,

$$V_n = \{n - n/m + 1, \dots, n - 1, n\},$$

where $m = m(n)$ is growing in n and to be chosen later. We want to keep track of fixed points within this part of the deck, let $F_n = \{\sigma \in S_n \mid \sigma \text{ has at least 1 fixed point in } V_n\}$. Note that V_n contains n/m positions, and so we find an easy upper bound on the uniform distribution $\pi_n(F_n) \leq 1/m \rightarrow 0$ as $n \rightarrow \infty$.

To bound the value of $\text{OST}_n^t(F_n)$ we reduce the problem to studying a simpler Markov chain linked to coupon collecting. When either of our hands (R^i, L^i) picks a new (previously untouched) card we say that this card gets *collected*. The uncollected cards in V_n at time t are those which have not yet been picked by either hand, and thus the size of this set gives us a lower bound on the number of fixed points in V_n . Writing U_n^t for the set of uncollected cards in V_n after t steps of the one-sided transposition shuffle, it follows that

$$\text{OST}_n^t(F_n) \geq \mathbb{P}(|U_n^t| \geq 1). \tag{3.54}$$

We wish to show that at time $t = n \log n - n \log \log n$ the probability on the right hand side of (3.54) is large, thus recovering a lower bound on OST_n^t . For the one-sided

transposition shuffle the value of L^i is clearly not independent of the value R^i . Importantly if $R^i \notin V_n$ then $\mathbb{P}(L^i \in V_n) = 0$. This means that a standard coupon-collecting argument for the time taken to collect all of the cards/positions in V_n cannot be applied to our shuffle, and a little more work is therefore required.

Note that at each step there are four possibilities: both hands collect new cards, only one hand does (left or right) or neither does. This permits us to bound the change in the number of *collected* cards as follows:

$$\begin{aligned} |V_n \setminus U_n^{t+1}| &= |V_n \setminus U_n^t| + |\{L^{t+1}, R^{t+1}\} \cap U_n^t| \\ &\leq |V_n \setminus U_n^t| + 2 \cdot \mathbb{1}[L^{t+1} \in U_n^t] + \mathbb{1}[L^{t+1} \notin U_n^t, R^{t+1} \in U_n^t], \end{aligned} \quad (3.55)$$

where $\mathbb{1}[\cdot]$ is an indicator function. We may find a useable upper bound for $\mathbb{P}(L^{t+1} \in U_n^t)$ by conditioning only on the size of U_n^t . Furthermore, we shall show that $\mathbb{P}(L^{t+1} \in U_n^t)$ is relatively small compared with $\mathbb{P}(L^{t+1} \notin U_n^t, R^{t+1} \in U_n^t)$, thus our approximation does not stop us finding the correct mixing time.

The probability $\mathbb{P}(L^{t+1} \in U_n^t)$, naturally depends on what positions the uncollected cards are in at time t . However, our left hand is more likely to choose positions towards the bottom of the pack so, letting $\hat{U}_n^t = \{n - n/m + 1, \dots, n - n/m + |U_n^t|\}$, i.e. the $|U_n^t|$ lowest numbered positions in V_n , we may form an upper bound on our probabilities as follows:

$$\mathbb{P}(L^{t+1} \in U_n^t) \leq \mathbb{P}(L^{t+1} \in \hat{U}_n^t).$$

Given the number of uncollected cards, the set \hat{U}_n^t is completely determined, so we may compute a bound in terms of $|U_n^t|$,

$$\begin{aligned} \mathbb{P}(L^{t+1} \in \hat{U}_n^t) &= \frac{1}{n} \sum_{k \in \hat{U}_n^t} \mathbb{P}(L^{t+1} \in \hat{U}_n^t | R^{t+1} = k) + \frac{1}{n} \sum_{k \in V_n \setminus \hat{U}_n^t} \mathbb{P}(L^{t+1} \in \hat{U}_n^t | R^{t+1} = k) \\ &= \frac{1}{n} \sum_{k \in \hat{U}_n^t} \frac{k - (n - n/m)}{k} + \frac{1}{n} \sum_{k \in V_n \setminus \hat{U}_n^t} \frac{|U_n^t|}{k} \\ &\leq \frac{\sum_{k=1}^{|U_n^t|} k + (n/m - |U_n^t|)|U_n^t|}{n(n - n/m)} \leq \frac{|U_n^t|}{(m-1)n}. \end{aligned} \quad (3.56)$$

The probability of the event $\{L^{t+1} \notin U_n^t, R^{t+1} \in U_n^t\}$ in (3.55) is simple to bound:

$$\mathbb{P}(L^{t+1} \notin U_n^t, R^{t+1} \in U_n^t) \leq \mathbb{P}(R^{t+1} \in U_n^t) \leq \frac{|U_n^t|}{n}. \quad (3.57)$$

Using (3.55), (3.56) and (3.57) together, we now define a counting process M_n^t which stochastically dominates the number of collected cards $|V_n \setminus U_n^t|$ at all times:

$$\begin{aligned} M_n^0 &= 0; \\ \mathbb{P}(M_n^{t+1} = M_n^t + k) &= \begin{cases} \frac{1}{(m-1)n} \left(\frac{n}{m} - M_n^t \right) & \text{if } k = 2 \\ \frac{1}{n} \left(\frac{n}{m} - M_n^t \right) & \text{if } k = 1 \\ 1 - \frac{m}{(m-1)n} \left(\frac{n}{m} - M_n^t \right) & \text{if } k = 0. \end{cases} \end{aligned} \quad (3.58)$$

Combining this with (3.54) we obtain the following bound on $\text{OST}_n^t(F_n)$:

$$\text{OST}_n^t(F_n) \geq \mathbb{P}(M_n^t < n/m). \quad (3.59)$$

The idea behind this counting process is that it increases by 1 whenever R^i collects a card (i.e. the event $\{R^{t+1} \in U_n^{t+1}\}$), and increases by 2 whenever L^i collects a card (i.e. the event $\{L^{t+1} \in U_n^{t+1}\}$). We know that for R^i to collect every card in V_n would take $O(n \log n)$ time by standard coupon collecting. In our case L^i is also helping to collect cards speeding up this time but we mitigate this additional help by focusing on collecting the cards in V_n .

From this point on we are interested in the time at which the process M_n^t first reaches level n/m , we now take $m = m(n) = \log n$ for the remainder of this section.

Lemma 3.5.1. *Let $\mathcal{T} = \min\{t : M_n^t \geq n/\log n\}$. Then for any $c > 2$,*

$$\mathbb{P}(\mathcal{T} \leq n \log n - n \log \log n - cn) \leq \frac{\pi^2}{6(c-2)^2}.$$

Proof. Let \mathcal{T}_i be the time spent by the process M_n^t in each state $i \geq 0$. We have $\mathcal{T} = \mathcal{T}_0 + \mathcal{T}_1 + \dots + \mathcal{T}_{(n/m)-1}$. Define p_i to be the probability that we leave state i after one step of M_n^t , from (3.58) we see that

$$p_i := \mathbb{P}(M_n^{t+1} > M_n^t | M_n^t = i) = \frac{m}{(m-1)n} \left(\frac{n}{m} - i \right). \quad (3.60)$$

In the standard coupon collector's problem each of the random variables \mathcal{T}_i has a geometric distribution with success probability p_i (see Section 2.3.2). Here, however, we have to take into account the chance that our counting process M_n increments by two, leading it to spend zero time at some state. Note first that

$$\mathbb{P}(M_n^{t+1} = M_n^t + 2 | M_n^{t+1} > M_n^t) = \frac{1}{m},$$

independently of the value of M_n^t , taking $m = \log n$ the probability of us skipping a state is tending to 0 as $n \rightarrow \infty$.

Prior to spending any time in state i , the process M_n^t must visit (at least) one of the states $i-1$ or $i-2$. A simple argument shows that

$$\mathbb{P}(\mathcal{T}_i > 0 | \mathcal{T}_{i-1} > 0) = 1 - \frac{1}{m}, \quad \text{and} \quad \mathbb{P}(\mathcal{T}_i > 0 | \mathcal{T}_{i-2} > 0) = 1 - \frac{1}{m} \left(1 - \frac{1}{m} \right) \geq 1 - \frac{1}{m}.$$

Therefore $\mathbb{P}(\mathcal{T}_i > 0) \geq 1 - \frac{1}{m}$ for all states i , and so \mathcal{T}_i stochastically dominates the random variable \mathcal{T}'_i with mass function

$$\mathbb{P}(\mathcal{T}'_i = k) = \begin{cases} 1/m & k = 0 \\ (1 - 1/m)p_i(1 - p_i)^{k-1} & k \geq 1. \end{cases} \quad (3.61)$$

Define $\mathcal{T}' = \mathcal{T}'_0 + \mathcal{T}'_1 + \dots + \mathcal{T}'_{(n/m)-1}$, it follows that $\mathbb{P}(\mathcal{T} < t) \leq \mathbb{P}(\mathcal{T}' < t)$ for any t .

Substituting $m = m(n) = \log n$ we may bound the expectation and variance of \mathcal{T}' using

the variables \mathcal{T}'_i :

$$\begin{aligned} \mathbb{E}[\mathcal{T}'] &= \sum_{i=0}^{n/m-1} \frac{m-1}{mp_i} = \left(\frac{m-1}{m}\right)^2 n \log(n/m) \geq n \log n - n \log \log n - 2n; \\ \text{Var}[\mathcal{T}'] &\leq \sum_{i=0}^{n/m-1} \frac{1}{p_i^2} \leq \sum_{i=1}^{n/m} \frac{n^2}{i^2} \leq \frac{\pi^2}{6} n^2. \end{aligned}$$

Finally, applying Chebyshev's inequality yields the following for any $c > 2$:

$$\mathbb{P}(\mathcal{T}' \leq n \log(n) - n \log \log n - cn) \leq \mathbb{P}(|\mathcal{T}' - \mathbb{E}[\mathcal{T}']| \geq (c-2)n) \leq \frac{\pi^2}{6(c-2)^2}. \quad \square$$

Proof of the Lower Limit in Theorem 3.1.2

Lemma 3.5.1 quickly leads to a proof of the lower bound in Theorem 3.1.2. Setting $t = n \log n - n \log \log n$, and $c > 2$ we obtain

$$\begin{aligned} \|\text{OST}_n^{t-cn} - \pi_n\|_{\text{TV}} &\geq \text{OST}_n^{t-cn}(F_n) - \pi_n(F_n) \geq \mathbb{P}(M_n^{t-cn} < n/m(n)) - 1/m(n) \\ &= \mathbb{P}(\mathcal{T} > t - cn) - \frac{1}{\log n} \geq 1 - \frac{\pi^2}{6(c-2)^2} - \frac{1}{\log n}. \end{aligned}$$

Therefore taking a limit in n , we recover $\liminf_{n \rightarrow \infty} \|\text{OST}_n^{t-cn} - \pi_n\|_{\text{TV}} \geq 1 - \frac{\pi^2}{6(c-2)^2}$ as required.

This finishes our proof of Theorem 3.1.2, showing the unbiased one-sided transposition shuffle exhibits a total variation cutoff at time $n \log n$. The mixing time for the unbiased one-sided transposition shuffle is therefore half as slow as the mixing time of the random transposition shuffle. This fact fits our intuition that restricting the choice of our left hand should slow down the randomisation of our deck. We now explore the behaviour of a new class of one-sided transposition shuffles where our left hand retains its behaviour but the right hand may be biased on $[n]$.

3.6 Biased One-sided Transposition shuffles

The unbiased one-sided transposition shuffle had our right hand choosing uniformly from the support $[n]$. In this section we allow the right hand to choose a position according to a weighted distribution on $[n]$, and we call these *biased one-sided transposition shuffles*. We generalise the work of Section 3.2 to compute the eigenvalues for our biased shuffles and show that within this class certain natural choices for the weights lead to shuffles which exhibit a cutoff in total variation distance.

Definition 3.6.1. Given a *weight function* $w : \mathbb{N} \rightarrow (0, \infty)$, let $N_w(n) = \sum_{i=1}^n w(k)$ denote the cumulative weight up to n . Then the *biased one-sided transposition shuffle* $\text{OST}_{n,w}$ is the random walk on S_n generated by the following distribution on transpositions:

$$\text{OST}_{n,w}(\tau) = \begin{cases} \frac{w(j)}{N_w(n)} \cdot \frac{1}{j} & \text{if } \tau = (ij) \text{ for some } 1 \leq i \leq j \leq n \\ 0 & \text{otherwise.} \end{cases} \quad (3.62)$$

Note that if $w(j) = 1$ for all j , we recover the unbiased one-sided transposition shuffle.

According to a biased one-sided transposition shuffle with weight function w the random variables R^t, L^t defined by the choices of our right and left hands follow the distributions,

$$\mathbb{P}(R^t = j) = \frac{w(j)}{N_w(n)}, \text{ for } 1 \leq j \leq n \quad (3.63)$$

$$\mathbb{P}(L^t = i | R^t = j) = \begin{cases} \frac{1}{j} & \text{if } i \leq j \\ 0 & \text{otherwise} \end{cases}. \quad (3.64)$$

Notice the left hand L^t still chooses a position uniformly on the set $\{1, \dots, R^t\}$. Importantly in Definition 3.6.1 the weight of each position $w(j)$ may only depend on j and not the size of the deck n . This setup preserves the recursive algebraic structure identified in Section 3.2, and allows us to recover the eigenvalues of the biased shuffles by minor modifications to our analysis.

Lifting Eigenvectors for the Biased One-sided Transposition Shuffle

Our first step is to turn the shuffle $\text{OST}_{n,w}$ into an element of our group algebra.

Definition 3.6.2. Let $n \in \mathbb{N}$. The biased one-sided transposition shuffle on n cards, with bias $w(j)$, may be viewed as the following element of the group algebra \mathfrak{S}_n :

$$\text{AOST}_{n,w} := \sum_{1 \leq i \leq j \leq n} \frac{w(j)}{j} (i j). \quad (3.65)$$

Note that to form $\text{AOST}_{n,w}$ we have scaled $\text{OST}_{n,w}$ by a factor of $N_w(n)$.

The eigenvalues of the biased one-sided transposition shuffle $\text{OST}_{n,w}$ may be recovered from the eigenvalues of the element $\text{AOST}_{n,w}$ acting on the simple Specht modules S^λ . Thus, as before we may focus our attention on the action of our shuffle as an element of the group algebra. The biased one-sided transposition shuffle inherits the recursive structure present in the unbiased case. Compare

$$\text{AOST}_{n+1,w} - \text{AOST}_{n,w} = \frac{w(n+1)}{n+1} \sum_{1 \leq i \leq n+1} (i \ n+1) \quad (3.66)$$

to equation (3.6), we can see that the only change is a new factor of $w(n+1)$. If our weight function is allowed to depend on things other than the position j then we run into issues with our lifting because $w(n+1)$ may no longer solely be a function of $(n+1)$. Now we establish an analogue of Theorem 3.2.5 for our biased shuffles.

Theorem 3.6.3. *Given $n \in \mathbb{N}$, we have*

$$\text{AOST}_{n+1,w} \circ \Phi_a - \Phi_a \circ \text{AOST}_{n,w} = \frac{w(n+1)}{n+1} \Phi_a + \frac{w(n+1)}{n+1} \sum_{1 \leq b \leq n} \Phi_b \circ \Theta_{b,a}. \quad (3.67)$$

Proof. This follows the same proof as Theorem 3.2.5 up to changes in constants. \square

If our bias is uniform, i.e. $w(j) = 1$, then we recover exactly the equation used for the unbiased one-sided transposition shuffle. Overall we can see that our bias affects Theorem 3.2.5 minimally. From here we may use the previously established lifting operators $\kappa_a^{\lambda,\mu}$ to skip to the conclusion of lifting eigenvectors. We restate Theorem 3.2.17 for our biased one-sided transposition shuffle.

Theorem 3.6.4 (Theorem 49 of [19]). *Let $\lambda \vdash n$, and $a \in \{1, 2, \dots, l(\lambda) + 1\}$. Take $i \in [n]$ such that $1 \leq i \leq a$ and set $\mu = \lambda + e_i$. Then,*

$$\text{AOST}_{n+1,w} \circ \kappa_a^{\lambda,\mu} - \kappa_a^{\lambda,\mu} \circ \text{AOST}_{n,w} = \frac{w(n+1)(2 + \lambda_i - i)}{n+1} \kappa_a^{\lambda,\mu}. \quad (3.68)$$

In particular if $v \in S^\lambda$ is an eigenvector of $\text{AOST}_{n,w}$ with eigenvalue ε , then $\kappa_a^{\lambda,\mu}(v)$ is an eigenvector of $\text{AOST}_{n+1,w}$ with eigenvalue

$$\varepsilon + \frac{w(n+1)(2 + \lambda_i - i)}{n+1}. \quad (3.69)$$

Proof. This follows from Theorem 3.2.17 with changes in constants for our biased shuffles. \square

Lemma 3.6.5. *The eigenvalues for the biased one-sided transposition shuffle $P_{n,w}$ on n cards are indexed by standard Young tableaux of size n . Moreover, the eigenvalue corresponding to a tableau T is given by*

$$\text{eig}(T) = \frac{1}{N_w(n)} \sum_{\substack{\text{boxes} \\ (i,j) \in T}} \frac{j-i+1}{T(i,j)} \cdot w(T(i,j))$$

Mixing Times of Biased One-sided Transposition Shuffles

We have recovered the eigenvalues for the biased one-sided transposition shuffles. To find the mixing time of the shuffle $\text{OST}_{n,w}$, we need some extra conditions on the weight function w . We focus on a class of weight functions with form $w(j) = j^\alpha$ for $\alpha \in \mathbb{R}$; we denote the resulting shuffle as $\text{OST}_{n,\alpha}$, and write $N_\alpha(n)$ in place of $N_w(n)$. For $\alpha = 0$ we recover our original one-sided transposition shuffle OST_n , while if $\alpha > 0$ ($\alpha < 0$) the right hand is biased towards the top (respectively, bottom) of the deck. Restating Lemma 3.6.5 for this class of weight functions we find the eigenvalue for $\text{OST}_{n,\alpha}$ given by tableau T to be,

$$\text{eig}(T) = \frac{1}{N_\alpha(n)} \sum_{\substack{\text{boxes} \\ (i,j) \in T}} (j-i+1) \cdot T(i,j)^{\alpha-1} = \frac{1}{N_\alpha(n)} \sum_{m=1}^n T(m) m^{\alpha-1}, \quad (3.70)$$

where $T(m)$ is the index $j-i+1$ of the box containing value m in T .

In the following sections we analyse the the mixing time of biased one-sided transposition shuffles generalising work from Sections 3.4 and 3.5. The conclusion of this section is the biased one-sided transposition shuffles exhibit a cutoff for all real choices of α .

Theorem 3.6.6. Define the time $t_{n,\alpha}$ as,

$$t_{n,\alpha} = \begin{cases} N_\alpha(n)/n^\alpha & \text{if } \alpha \leq 1 \\ N_\alpha(n)/N_{\alpha-1}(n) & \text{if } \alpha \geq 1 \end{cases}.$$

The biased one-sided transposition shuffle $\text{OST}_{n,\alpha}$ satisfies the following bounds for any $c_1 > 5/2, c_2 > \max(2, 3 - \alpha)$:

$$\limsup_{n \rightarrow \infty} \|\text{OST}_{n,\alpha}^{t_{n,\alpha}(\log n + c_1)} - \pi_n\| \leq Ae^{-2c_1} \text{ for a universal constant } A, \text{ for all } \alpha$$

and

$$\liminf_{n \rightarrow \infty} \|\text{OST}_{n,\alpha}^{t_{n,\alpha}(\log n - \log \log n - c_2)} - \pi_n\| \geq \begin{cases} 1 - \frac{\pi^2}{6(c_2 - 3 + \alpha)^2} & \text{if } \alpha \leq 1 \\ 1 - \frac{\pi^2}{6(c_2 - 2)^2} & \text{if } \alpha \geq 1 \end{cases}$$

Thus, the biased one-sided transposition shuffle exhibits a total variation cutoff at time $t_{n,\alpha} \log n$ for all α with a window of size $t_{n,\alpha} \log \log n$.

The asymptotics of the times $t_{n,\alpha}$ as $n \rightarrow \infty$ for the one-sided transposition shuffle are summarised in Table 3.4

	$\alpha \in (-\infty, -1)$	$\alpha = -1$	$\alpha \in (-1, 1]$	$\alpha \in (1, \infty)$
$t_{n,\alpha} \log n$	$\zeta(-\alpha)n^{-\alpha} \log n$	$n(\log n)^2$	$\frac{1}{1+\alpha}n \log n$	$\frac{\alpha}{1+\alpha}n \log n$

Table 3.4: Asymptotics of $t_{n,\alpha} \log n$ as $n \rightarrow \infty$.

The fastest mixing time of a biased one-sided transposition shuffle is obtained when $\alpha = 1$; using this weight function the shuffle is constant on the conjugacy class of transpositions, with probability close to that of the random transposition shuffle, $P_{n,1}((ij)) = 2/(n(n+1))$. We obtain a mixing time of $t_{n,1} \sim (n/2) \log n$ which agrees with that of the random transposition shuffle. The mixing time increases as α moves away from 1 in either direction but as $\alpha \rightarrow \infty$ the time stays bounded above by $n \log n$ whereas if $\alpha \rightarrow -\infty$ then the mixing time is unbounded. As $\alpha \rightarrow -\infty$ our right hand is choosing positions near to the bottom of the deck frequently, and since the left hand is restricted to the range $\{1, \dots, R^i\}$, this leads to more and more mass being placed on the identity element as $\alpha \rightarrow \infty$, thus slowing the mixing time.

3.6.1 Cutoff for Biased One-sided Transposition Shuffles with $\alpha \leq 1$

The proof of a total variation cutoff for the biased one-sided transposition shuffle with $\alpha \leq 1$ follows from generalisations of Sections 3.4 and 3.5. For the upper bound we establish bounds on the eigenvalues of $\text{OST}_{n,\alpha}$ for large and small partitions, and reduce the analysis to a previously bounded summation. The lower bound follows from the same argument as Section 3.5 with careful attention paid to how the weighting affects the probabilities of L^i, R^i .

Upper Bound

First of all we once again use Lemma 1.1.27 to form an upper bound on the total variation distance of $\text{OST}_{n,\alpha}$ from π_n . Furthermore, for $\alpha \leq 1$ every result of Section 3.3 holds after

suitable adjustments to account for the new eigenvalues. In particular, Lemma 3.3.7 now becomes: for any $T \in \text{SYT}(\lambda)$ we have $\text{eig}(T) + \text{eig}(T') = 2N_{\alpha-1}(n)/N_{\alpha}(n)$. Applying the same analysis as equation (3.36) we reduce our bound to:

$$\begin{aligned} 4\|\text{OST}_{n,\alpha}^t - \pi_n\|_{\text{TV}}^2 &\leq \sum_{\substack{\lambda \vdash n \\ \lambda \neq (n)}} \sum_{T \in \text{SYT}(\lambda)} d_{\lambda} (\text{eig}(T))^{2t} \\ &\leq (\text{eig}(T_{(1^n)}))^{2t} + 2 \sum_{\substack{\lambda: \text{eig}(T_{\lambda}^{\rightarrow}) \geq 0 \\ \lambda \neq (n)}} d_{\lambda}^2 \text{eig}(T_{\lambda}^{\rightarrow})^{2t}. \end{aligned} \quad (3.71)$$

We bound the eigenvalues of large and small partitions separately.

Lemma 3.6.7. *Let $\lambda \vdash n$ with $\lambda_1 = n - k$. Then the eigenvalue $\text{eig}(T_{\lambda}^{\rightarrow})$ for the shuffle $\text{OST}_{n,\alpha}$ with $\alpha \leq 1$ may be bounded as follows:*

$$\text{eig}(T_{\lambda}^{\rightarrow}) \leq \begin{cases} 1 - \frac{(n-k+1)kn^{\alpha}}{nN_{\alpha}(n)} & \text{if } k \leq n/4 \\ 1 - \frac{kn^{\alpha}}{2N_{\alpha}(n)} & \text{if } k > n/4. \end{cases}$$

Proof. For $k \leq n/2$, the maximum partition in the dominance order of partitions with $\lambda_1 = n - k$ is $(n - k, k)$, and so $\text{eig}(T_{\lambda}^{\rightarrow}) \leq \text{eig}(T_{(n-k,k)}^{\rightarrow})$. The eigenvalue of $T_{(n-k,k)}^{\rightarrow}$ may be calculated by summing over the rows of the partition $(n - k, k)$ and using equation (3.70), as follows:

$$\begin{aligned} N_{\alpha}(n) \text{eig}(T_{(n-k,k)}^{\rightarrow}) &= \sum_{m=1}^{n-k} m^{\alpha} + \sum_{m=n-k+1}^n (m - n + k - 1)m^{\alpha-1} \\ &= \sum_{m=1}^n m^{\alpha} - (n - k + 1) \sum_{m=n-k+1}^n m^{\alpha-1} \\ &\leq N_{\alpha}(n) - \frac{(n - k + 1)kn^{\alpha}}{n}. \end{aligned}$$

This immediately proves the desired inequality for $k \leq n/4$. The above bound also holds for $k \in (n/4, n/2]$, and in this case

$$N_{\alpha}(n) - \frac{(n - k + 1)kn^{\alpha}}{n} \leq N_{\alpha}(n) - \frac{kn^{\alpha}}{2}$$

giving us our required bound.

For $k > n/2$ we once again need to bound $\text{eig}(T_{(n-k,\star)}^{\rightarrow})$. Letting $\nu = (n - k, \star)$ for ease of notation we calculate as follows:

$$\begin{aligned} N_{\alpha}(n) \text{eig}(T_{\nu}^{\rightarrow}) &= \sum_{j=1}^{n-k} (j - 1 + 1)^{\alpha} + \sum_{i=2}^{l(\nu)} \sum_{j=1}^{\nu_i} (j - i + 1) \frac{((i - 1)(n - k) + j)^{\alpha}}{(i - 1)(n - k) + j} \\ &= N_{\alpha}(n) - \sum_{i=2}^{l(\nu)} \sum_{j=1}^{\nu_i} (i - 1)(n - k + 1) \frac{((i - 1)(n - k) + j)^{\alpha}}{(i - 1)(n - k) + j} \\ &\leq N_{\alpha}(n) - \frac{n^{\alpha}}{n} \sum_{i=2}^{l(\nu)} \sum_{j=1}^{\nu_i} ((i - 1) \cdot (n - k + 1)) \\ &\leq N_{\alpha}(n) - \frac{(n - k + 1)n^{\alpha}}{n} \sum_{i=2}^{l(\nu)} (i - 1)\nu_i. \end{aligned} \quad (3.72)$$

By definition of the partition ν , each row but the last has size $n - k$, and the final row has size $\nu_{l(\nu)} = n - (l(\nu) - 1)(n - k)$. In addition, since $l(\nu)$ is equal to the ceiling of $n/(n - k)$ we may write $l(\nu) = n/(n - k) + x$ for some $0 \leq x < 1$. Substituting these values into equation (3.72) we obtain:

$$\begin{aligned}
N_\alpha(n) \operatorname{eig}(T_\nu^{\rightarrow}) &\leq N_\alpha(n) - \frac{(n - k)n^\alpha (l(\nu) - 1)(2n - (n - k)l(\nu))}{n \cdot 2} \\
&= N_\alpha(n) - \frac{n^\alpha}{2n} (n - (1 - x)(n - k))(n - x(n - k)) \\
&= N_\alpha(n) - \frac{n^\alpha}{2n} (nk + x(1 - x)(n - k)^2) \\
&\leq N_\alpha(n) - \frac{kn^\alpha}{2}.
\end{aligned}$$

□

Using these eigenvalue bounds we complete the proof of the upper bound for the biased one-sided transposition shuffles with $\alpha \leq 1$. Splitting the sum in (3.71) around the big and small partitions we find,

$$\begin{aligned}
4\|\operatorname{OST}_{n,\alpha}^t - \pi_n\|_{\text{TV}}^2 &\leq (\operatorname{eig}(T_{1^n}))^{2t} + 2 \sum_{k=1}^{n/4} \binom{n}{k}^2 k! \left(1 - \frac{2(n - k + 1)kn^\alpha}{2nN_\alpha(n)}\right)^{2t} \\
&\quad + 2 \sum_{k>n/4}^{n-1} \binom{n}{k}^2 k! \left(1 - \frac{kn^\alpha}{2N_\alpha(n)}\right)^{2t}.
\end{aligned}$$

The singular term $\operatorname{eig}(T_{1^n})$ may be seen to be tending to 0 as $n \rightarrow \infty$ at time $t = t_{n,\alpha}(\log n + c)$,

$$\lim_{n \rightarrow \infty} \operatorname{eig}(T_{1^n})^{2t} = \lim_{n \rightarrow \infty} \left(\frac{1}{N_\alpha(n)} \sum_{i=1}^n (2 - i)^{\alpha-1} \right)^{2t} \leq \lim_{n \rightarrow \infty} \left(\frac{n}{N_\alpha(n)} \right)^{2t_{n,\alpha}(\log n + c)} = 0$$

We are left with two sums to control, taking $t = t_{n,\alpha}(\log n + c)$ these sums may be reduced too:

$$e^{-2c} \sum_{k=1}^{n/4} \frac{n^{\frac{2(k(k-1))}{n}}}{k!} + \sum_{k>n/4}^{n-1} \binom{n}{k}^2 k! n^{-k} e^{-kc} \tag{3.73}$$

The first summation in (3.73) was shown to be bounded by a universal constant by Diaconis [16, Chapter 3D Theorem 5]. The ratio between consecutive terms in the second summation is decreasing in k and less than 1 at its start ($k = n/4$) if $c > 2.5$. Thus we may bound the summation by bounding it by n times its first term; this reduces to bounding

$$\sum_{k>n/4}^{n-1} \binom{n}{k}^2 k! n^{-k} e^{-kc} \leq n \binom{n}{n/4}^2 (n/4)! n^{-n/4} e^{-(n/4)c}. \tag{3.74}$$

We now compute the asymptotics of the binomial and factorial terms using Stirling's approximation, these results are given in Table 3.5.

Term	$\binom{n}{\delta n}^2$	$(\delta n)!$
Stirling's approximation	$(\delta^{-2\delta n}(1-\delta)^{-2(1-\delta)n})/(2\pi\delta(1-\delta)n)$	$\sqrt{2\pi\delta n}(\delta n/e)^{\delta n}$
$\log(\text{Stirling's approximation})$	$-\log(2\pi\delta) - 2\delta n \log(\delta) - 2(1-\delta)n \log(1-\delta)$	$\frac{1}{2} \log(2\pi\delta n) + \delta n \log(\delta n) - \delta n$

Table 3.5: Asymptotics of binomial and factorial terms computed using Stirling's approximation.

Using the results in Table 3.5 we find the asymptotics of the logarithm of (3.74) to be:

$$-(n/4)(c+1+2\log(0.25)+6\log(0.75))+o(n). \quad (3.75)$$

The coefficient of the leading order term

$$c+1+2\log(0.25)+6\log(0.75)$$

is positive if $c > 5/2$. Therefore for $c > 5/2$, the term (3.75) is tending to $-\infty$ as $n \rightarrow \infty$. Hence, the second summation in (3.73) is tending to 0 as $n \rightarrow \infty$. Putting together all of the bounds above we find,

$$\limsup_{n \rightarrow \infty} \|\text{OST}_{n,\alpha}^{t_{n,\alpha}(\log n^{++c})} - \pi_n\|_{\text{TV}} \leq Ae^{-c} \text{ for some universal constant } A.$$

This completes the upper bound for $\alpha \leq 1$ present in Theorem 3.6.6.

Lower Bound

We use a coupon-collecting argument as in Section 3.5, once again letting $V_n = \{n - n/m + 1, \dots, n - 1, n\}$ with $m = \log n$, and considering the set $F_n = \{\sigma \in S_n \mid \sigma \text{ has at least 1 fixed point in } V_n\}$. We have seen previously that under the uniform distribution we have $\pi_n(F_n) \leq 1/m = 1/\log n$. For the biased one-sided transposition shuffle we modify the bounds (3.56) and (3.57) as follows, using the inequality $k^\alpha \leq (\frac{m}{m-1})^{1-\alpha} n^\alpha$ for all $k \in V_n$ (which holds for all $\alpha \leq 1$):

$$\begin{aligned} \mathbb{P}(L^{t+1} \in \hat{U}_n^t) &= \sum_{k \in \hat{U}_n^t} \frac{w(k)}{N_\alpha(n)} \frac{k - (n - n/m)}{k} + \sum_{k \in V_n \setminus \hat{U}_n^t} \frac{w(k)}{N_\alpha(n)} \frac{|U_n^t|}{k} \\ &\leq \frac{(\frac{m}{m-1})^{1-\alpha} n^\alpha \sum_{k=1}^{|U_n^t|} k + (n/m - |U_n^t|)|U_n^t|}{N_\alpha(n)(n - n/m)} \leq \frac{(\frac{m}{m-1})^{1-\alpha} n^\alpha |U_n^t|}{N_\alpha(n)(m-1)}; \end{aligned} \quad (3.76)$$

$$\mathbb{P}(R^{t+1} \in U_n^t) \leq \frac{(\frac{m}{m-1})^{1-\alpha} n^\alpha |U_n^t|}{N_\alpha(n)}. \quad (3.77)$$

Using these as before we construct a counting process $M_{n,\alpha}^t$ which stochastically dominates the number of collected cards $|V_n \setminus U_n^t|$ at all times:

$$M_{n,\alpha}^0 = 0;$$

$$\mathbb{P}(M_{n,\alpha}^{t+1} = M_{n,\alpha}^t + k) = \begin{cases} \frac{1}{(m-1)} \left(\frac{\left(\frac{m-1}{m}\right)^{1-\alpha} n^\alpha |U_n^t|}{N_\alpha(n)} \right) & \text{if } k = 2 \\ \left(\frac{\left(\frac{m-1}{m}\right)^{1-\alpha} n^\alpha |U_n^t|}{N_\alpha(n)} \right) & \text{if } k = 1 \\ 1 - \frac{m}{(m-1)} \left(\frac{\left(\frac{m-1}{m}\right)^{1-\alpha} n^\alpha |U_n^t|}{N_\alpha(n)} \right) & \text{if } k = 0. \end{cases} \quad (3.78)$$

Our counting process $M_{n,\alpha}^t$ gives a lower bound on the probability of being in F_n as follows,

$$\text{OST}_{n,\alpha}^t(F_n) \geq \mathbb{P}(M_{n,\alpha}^t < n/m).$$

We are now interested in the time it takes $M_{n,\alpha}^t$ to pass value n/m .

Lemma 3.6.8. *Let $\mathcal{T} = \min\{t : M_{n,\alpha}^t \geq n/\log n\}$. Then for any $c > 3 - \alpha$,*

$$\mathbb{P}(\mathcal{T} \leq t_{n,\alpha}(\log n - \log \log n - c)) \geq 1 - \frac{\pi^2}{6(c - (3 - \alpha))^2}.$$

Proof. Construct the time $\mathcal{T} = \mathcal{T}_0 + \mathcal{T}_1 + \dots + \mathcal{T}_{n/m-1}$ with T_i being the time spent in state i . Denote the escape probability of state i as p_i , the expression for p_i modified from (3.60) becomes

$$p_i = \left(\frac{m}{m-1} \right)^{2-\alpha} \frac{n^\alpha}{N_\alpha(n)} \left(\frac{n}{m} - i \right),$$

and this is easily checked to be strictly less than one for all values of $\alpha \leq 1$ if n is sufficiently large. The remainder of the analysis mirrors the unbiased case: using the new expression for p_i the distribution of the random variable \mathcal{T}'_i is exactly as given in (3.61), that is

$$\mathbb{P}(\mathcal{T}'_i = k) = \begin{cases} 1/m & k = 0 \\ (1 - 1/m)p_i(1 - p_i)^{k-1} & k \geq 1. \end{cases}$$

Setting $\mathcal{T}' = \mathcal{T}'_0 + \dots + \mathcal{T}'_{n/m-1}$ then \mathcal{T} stochastically dominates \mathcal{T}' . Analysing the expectation and variance of \mathcal{T}' we arrive at:

$$\begin{aligned} \mathbb{E}[\mathcal{T}'] &= \sum_{i=0}^{n/m-1} \frac{m-1}{mp_i} = \frac{N_\alpha(n)}{n^\alpha} \left(1 - \frac{1}{m}\right)^{3-\alpha} \sum_{i=1}^{n/m} \frac{1}{i} \geq \frac{N_\alpha(n)}{n^\alpha} (\log n - \log \log n - (3 - \alpha)); \\ \text{Var}[\mathcal{T}'] &\leq \sum_{i=0}^{n/m-1} \frac{1}{p_i^2} = \frac{N_\alpha(n)^2}{n^{2\alpha}} \left(1 - \frac{1}{m}\right)^{2(2-\alpha)} \sum_{i=1}^{n/m} \frac{1}{i^2} \leq \frac{N_\alpha(n)^2 \pi^2}{n^{2\alpha} 6}. \end{aligned}$$

An application of Chebyshev's inequality give us:

$$\begin{aligned} \mathbb{P}(\mathcal{T}' \leq t_{n,\alpha}(\log n - \log \log n - (3 - \alpha))) &\leq \mathbb{P}(|\mathcal{T}' - \mathbb{E}[\mathcal{T}']| \geq (c - (3 - \alpha))n) \\ &\leq \frac{\pi^2}{6(c - (3 - \alpha))^2}. \end{aligned}$$

□

Lemma 3.6.8 quickly leads to a lower bound, taking $c > 3 - \alpha$ we have

$$\begin{aligned} \|\text{OST}_{n,\alpha}^{t_{n,\alpha}(\log n - \log \log n - c)} - \pi_n\|_{\text{TV}} &\geq \mathbb{P}(\mathcal{T} > t_{n,\alpha}(\log n - \log \log n - c)) - \pi_n(F_n) \\ &\geq 1 - \frac{\pi^2}{6(c - (3 - \alpha))^2} - \frac{1}{\log n}. \end{aligned}$$

Taking limits in n recovers the limit present in Theorem 3.6.6.

3.6.2 Cutoff for Biased One-sided Transposition Shuffles with $\alpha \geq 1$

The proof of a total variation cutoff for biased one-sided transposition shuffle with $\alpha \geq 1$ requires more work than the $\alpha \leq 1$ case. To prove the upper limit we construct a new special tableau, T_λ^\searrow , which bounds all the eigenvalues associated to partition λ . This new tableau is then used to reduce our analysis to summations which we have previous seen to be bounded at the correct mixing time. The lower bound follows from the same argument as Section 3.5 up to the replacement of a normalising factor.

Upper Bound

Unlike the $\alpha \leq 1$ previous case not all results of Section 3.3 generalise to the bias $\alpha \geq 1$. The results which do generalise include Lemma 3.3.5 and 3.3.8, we restate these below with all eigenvalues calculated for a biased one-sided transposition shuffle with $\alpha \geq 1$.

Lemma 3.6.9. *Let $\lambda \vdash n$. For any $T \in \text{SYT}(\lambda)$ we have the following inequality :*

$$\text{eig}(T_\lambda^{\rightarrow}) \leq \text{eig}(T) \leq \text{eig}(T_\lambda^{\downarrow}). \quad (3.79)$$

Proof. This follows from the proof of Lemma 3.3.5 with the roles of T_λ^{\rightarrow} and T_λ^{\downarrow} reversed. \square

Corollary 3.6.10. *Let $\lambda \vdash n$, and suppose we have $\text{eig}(T_\lambda^{\rightarrow}) \leq 0$, then we have*

$$\text{eig}(T_{\lambda'}^{\downarrow}) \geq |\text{eig}(T_\lambda^{\rightarrow})| \geq 0. \quad (3.80)$$

Proof. This follows from the proof of Lemma 3.3.8 with the roles of T_λ^{\rightarrow} and T_λ^{\downarrow} reversed. \square

Using the above lemmas allows us to reduce our analysis to looking at T_λ^{\downarrow} but we may not solely focus on partitions at the top of the dominance order like before. Applying the same analysis as equation (3.36) we reduce our bound to:

$$\begin{aligned} 4\|\text{OST}_{n,\alpha}^t - \pi_n\|_{\text{TV}}^2 &\leq \sum_{\substack{\lambda \vdash n \\ \lambda \neq (n)}} \sum_{T \in \text{SYT}(\lambda)} d_\lambda (\text{eig}(T))^{2t} \\ &\leq (\text{eig}(T_{1^n}))^{2t} + 2 \sum_{\substack{\lambda: \text{eig}(T_\lambda^{\downarrow}) \geq 0 \\ \lambda \neq (n)}} d_\lambda^2 \text{eig}(T_\lambda^{\downarrow})^{2t}. \end{aligned} \quad (3.81)$$

To bound the eigenvalue of T_λ^{\downarrow} when $\lambda_1 = n - k$ we introduce a new special tableau for a Young diagram λ .

Definition 3.6.11. Let $\lambda \vdash n$. Define the λ -tableau T_λ^\searrow by filling in the diagonals of λ from left to right, with each diagonal filled in top to bottom. For example, the tableaux $T_{(4^4)}^\searrow, T_{(4,3,1^2)}^\searrow, T_{(6,4)}^\searrow$ are given below,

7	11	14	16
4	8	12	15
2	5	9	13
1	3	6	10

4	6	8	9
3	5	7	
2			
1			

2	4	6	8	9	10
1	3	5	7		

Lemma 3.6.12. For $\alpha \geq 1$, and $\lambda \vdash n$ with $\lambda_1 = n - k$ we have,

$$\text{eig}(T_\lambda^\downarrow) \leq \text{eig}(T_{(n-k,*)}^\searrow).$$

Proof. The eigenvalue associated to the tableau T_λ^\downarrow is given by,

$$N_\alpha(n) \cdot \text{eig}(T) = \sum_{m=1}^n T(m) m^{\alpha-1}.$$

Now consider all indexes $T_\lambda^\downarrow(m)$ for $m \in [n]$ including repeats and order them from smallest to largest as c_i for $i \in [n]$. The fact that $\alpha \geq 1$ allows us to upper bound our eigenvalue in the following way:

$$N_\alpha(n) \text{eig}(T_\lambda^\downarrow) = \sum_{m=1}^n T_\lambda^\downarrow(m) m^{\alpha-1} \leq \sum_{m=1}^n c_m m^{\alpha-1} = \sum_{m=1}^n T_\lambda^\searrow(m) m^{\alpha-1}.$$

The first inequality holds because we have matched up the pairs of values c_i and $m^{\alpha-1}$ in such a way as to maximise the summations value. The last equality comes from the fact that every diagonal in λ has the same index $(j - i + 1)$ and this increases from left to right matching the ordering of T_λ^\searrow .

To complete our proof we need to show that moving from T_λ^\searrow to $T_{(n-k,*)}^\searrow$ increases the eigenvalue. To get from λ to $(n - k, *)$ we move boxes up and to the right. Every box that moves increases its index $(j - i + 1)$ and thus its value in our eigenvalue summation. Hence when we put the values of the boxes of $(n - k, *)$ in via \searrow , we find $T_\lambda^\searrow(m) \leq T_{(n-k,*)}^\searrow(m)$ for all m , establishing

$$\sum_{m=1}^n T_\lambda^\searrow(m) \cdot m^{\alpha-1} \leq \sum_{m=1}^n T_{(n-k,*)}^\searrow(m) \cdot m^{\alpha-1} = N_\alpha(n) \cdot \text{eig}(T_{(n-k,*)}^\searrow). \quad \square$$

Example 3.6.13. As a demonstration of Lemma 3.6.12 consider the partition $\lambda = (4, 2)$. The tableaux $T_{(4,2)}^\downarrow$, and $T_{(4,3)}^\searrow$ are given below:

1	3	5	6
2	4		

2	4	5	6
1	3		

For any α we may compute the eigenvalues of both tableaux using equation (3.70):

$$\begin{aligned} N_\alpha(8) \text{eig}(T_{(5,3)}^\downarrow) &= 1(1^{\alpha-1}) + 0(2^{\alpha-1}) + 2(3^{\alpha-1}) + 1(4^{\alpha-1}) + 3(5^{\alpha-1}) + 4(6^{\alpha-1}) \\ N_\alpha(8) \text{eig}(T_{(5,3)}^\searrow) &= 0(1^{\alpha-1}) + 1(2^{\alpha-1}) + 1(3^{\alpha-1}) + 2(4^{\alpha-1}) + 3(5^{\alpha-1}) + 4(6^{\alpha-1}). \end{aligned}$$

For $\alpha \geq 1$ the values $i^{\alpha-1}$ are increasing in i . Therefore rearranging the coefficients given by the diagonal value of each box (in our case the values $\{0, 1, 1, 2, 3, 4\}$) to be in increasing order only increases our eigenvalue of the tableau.

Lemma 3.6.14. *Let $n \in \mathbb{N}$, $k \in \{1, \dots, n-2\}$. Suppose $\lambda = (n-k, \star)$, then for every box $(i, j) \in \lambda$ we have*

$$j - i + 1 \leq \frac{n-k}{n} T_{\lambda}^{\succ}(i, j) \quad (3.82)$$

Proof. Let us write $l(n-k, \star) = l^+ = \lceil \frac{n}{n-k} \rceil$ and $l^- = \lfloor \frac{n}{n-k} \rfloor$. If $(i_1, j_1), (i_2, j_2)$ belong to the same diagonal of λ then $j_1 - i_1 + 1 = j_2 - i_2 + 1$, and hence if (3.82) holds for the smallest value $T_{\lambda}^{\succ}(i, j)$ on a diagonal it holds for every entry of that diagonal. Furthermore, the bound trivially holds for any box (i, j) which satisfies $j - i + 1 \leq 0$ (for which the left hand side of (3.82) is non-positive). Combining these two observations, we see that it suffices to prove the bound for boxes $(1, j)$, which appear on the first row of $(n-k, \star)$ as they must contain the smallest value on their given diagonal.

Note that no diagonal can contain more than l^+ boxes: call diagonals with l^- or fewer boxes *short* diagonals, and all others *long* diagonals. Note that long diagonals can only exist when $l^+ = l^- + 1$. Any long diagonals clearly occur strictly before the short ones, when working left to right along the first row. If the box $(1, j)$ lies on a long diagonal, then the numbering pattern for T^{\succ} implies that this box will contain the integer $T_{(n-k, \star)}^{\succ}(i, j) = \binom{l^+}{2} + 1 + (j-1)l^+$. For this value of m , the left hand side of (3.82) becomes

$$\frac{n-k}{n} \left(\frac{l^+(l^+-1)}{2} + 1 + (j-1)l^+ \right) = \frac{(n-k)l^+}{n} \left(\frac{l^+-1}{2} + \frac{1}{l^+} - 1 + j \right) \geq j$$

thanks to the definition of l^+ and the fact that $(x-1)/2 + 1/x \geq 1$ if $x \geq 2$.

It remains to deal with the short diagonals which contain a box on the first row. For these diagonals we now work from right ($j = n-k$) to left ($j = 1$). For rightmost box on the first row $(1, n-k)$ we know that $T_{(n-k, \star)}^{\succ}((1, n-k)) = n$, and it is clear that (3.82) holds for box $(1, n-k)$. Suppose (3.82) holds for a box $(1, j)$, then for box $(1, j-1)$ on a small diagonal it is straightforward to see that

$$j-1 \leq \frac{n-k}{n} T_{\lambda}^{\succ}(1, j) - 1 \leq \frac{n-k}{n} T_{\lambda}^{\succ}(1, j-1) \quad (3.83)$$

because $T_{\lambda}^{\succ}(1, j) - n/(n-k) \leq T_{\lambda}^{\succ}(1, j-1)$. The result for all short diagonals now follows quickly by induction. \square

Example 3.6.15. To illustrate Lemma 3.6.14 consider the partition $\lambda = (10, 10, 9)$, which may be represented as $\lambda = (n-k, \star)$ for $n = 29, k = 19$. For $\lambda = (n-k, \star)$ we have $(n-k/n) = 10/29$. Define T a λ -tableau, such that $T(i, j) = \lfloor (10/29)T_{\lambda}^{\succ}(i, j) \rfloor$ for all $(i, j) \in \lambda$. The two tableaux T_{λ}^{\succ}, T are given below,

$$T_{\lambda}^{\succ} = \begin{array}{|c|c|c|c|c|c|c|c|c|c|} \hline 4 & 7 & 10 & 13 & 16 & 19 & 22 & 25 & 27 & 29 \\ \hline 2 & 5 & 8 & 11 & 14 & 17 & 20 & 23 & 26 & 28 \\ \hline 1 & 3 & 6 & 9 & 12 & 15 & 18 & 21 & 24 & \\ \hline \end{array} \quad T = \begin{array}{|c|c|c|c|c|c|c|c|c|c|} \hline 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 \\ \hline 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ \hline 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & \\ \hline \end{array}.$$

We may clearly see that for any box $(i, j) \in \lambda$ we have $j - i + 1 \leq T(i, j)$.

Lemma 3.6.16. *Let $\lambda \vdash n$ with $\lambda_1 = n - k$. Then the eigenvalue $\text{eig}(T_\lambda^\downarrow)$ for the shuffle $\text{OST}_{n,\alpha}$ with $\alpha \geq 1$ may be bounded as follows:*

$$\text{eig}(T_\lambda^\downarrow) \leq \begin{cases} 1 - \frac{k(n-k)N_{\alpha-1}(n)}{nN_\alpha(n)} & \text{if } k \leq n/4 \\ 1 - \frac{k}{n} & \text{for all } k \end{cases}.$$

Proof. We begin by quickly proving the second bound for all k using Lemmas 3.6.12 and 3.6.14:

$$\begin{aligned} \text{eig}(T_\lambda^\downarrow) \leq \text{eig}\left(T_{(n-k,*)}^\downarrow\right) &= \frac{1}{N_\alpha(n)} \sum_{(i,j) \in T_{(n-k,*)}^\downarrow} (j-i+1) T_{(n-k,*)}^\downarrow(i,j)^{\alpha-1} \\ &\leq \frac{1}{N_\alpha(n)} \frac{n-k}{n} \sum_{(i,j) \in T_{(n-k,*)}^\downarrow} T_{(n-k,*)}^\downarrow(i,j)^\alpha = 1 - \frac{k}{n}. \end{aligned}$$

Now we prove the bound for $k \leq n/4$, by finding a tighter bound on $T_{(n-k,*)}^\downarrow$ and once again using Lemma 3.6.12. For $k \leq n/4$, we know $(n-k, *) = (n-k, k)$, and thanks to the order in which the boxes are filled, we see that $T_{(n-k,*)}^\downarrow(m) \leq m/2$ for $1 \leq m \leq 2k$, and $T_{(n-k,*)}^\downarrow(m) = m - k$ for $2k + 1 \leq m \leq n$. This gives us the following simple bound:

$$\begin{aligned} N_\alpha(n) \cdot \text{eig}(T_{(n-k,*)}^\downarrow) &\leq \sum_{m=1}^{2k} \frac{m}{2} m^{\alpha-1} + \sum_{m=2k+1}^n (m-k) m^{\alpha-1} \\ &= N_\alpha(n) - kN_{\alpha-1}(n) + \left(kN_{\alpha-1}(2k) - \frac{1}{2}N_\alpha(2k) \right). \quad (3.84) \end{aligned}$$

Now,

$$\begin{aligned} kN_{\alpha-1}(2k) - \frac{1}{2}N_\alpha(2k) &= k \sum_{m=1}^{2k} m^{\alpha-1} \left(1 - \frac{m}{2k}\right) \leq k \int_0^{2k} x^{\alpha-1} \left(1 - \frac{x}{2k}\right) dx \\ &= \frac{k(2k)^\alpha}{\alpha(1+\alpha)} \\ &\leq \frac{k^2 n^{\alpha-1}}{\alpha} \frac{2^{2-\alpha}}{(1+\alpha)} \quad (\text{since } k \leq n/4) \\ &\leq \frac{k^2 N_{\alpha-1}(n)}{n} \quad (3.85) \end{aligned}$$

in the last step we have used the inequality $n^\alpha/\alpha = \int_0^n j^{\alpha-1} \leq \sum_0^n j^{\alpha-1} = N_{\alpha-1}(n)$ which holds for $\alpha \geq 1$. Combining (3.85) and (3.84) yields the desired result. \square

Note that the bound for $k \leq n/4$ in Lemma 3.6.16 is not the same as our previous bound (Lemma 3.6.7) when $\alpha = 1$. The accuracy lost here helps us establish one bound for all $\alpha \geq 1$. Using Lemma 3.6.16 above we complete our upper bound argument for $P_{n,\alpha}$

with $\alpha \geq 1$. The total variation distance is bounded as follows,

$$\begin{aligned} 4\|\text{OST}_{n,\alpha}^t - \pi_n\|_{\text{TV}}^2 &\leq (\text{eig}(1^n))^{2t} + 2 \sum_{k=1}^{\binom{n/4}{2}} \binom{n}{k}^2 k! \left(\text{eig}(T_{(n-k,k)}^\lambda)\right)^{2t} \\ &\quad + 2 \sum_{k>\binom{n/4}{2}}^{\binom{n}{2}} \binom{n}{k}^2 k! \left(\text{eig}(T_{(n-k,*)}^\lambda)\right)^{2t} \end{aligned}$$

Substituting $t = t_{n,\alpha}(\log n + c)$, the first term $\text{eig}(T_1^n)$ disappears as $n \rightarrow \infty$,

$$\lim_{n \rightarrow \infty} \text{eig}(T_1^n)^{2t} = \lim_{n \rightarrow \infty} \left(\frac{1}{N_\alpha(n)} \sum_{i=1}^n (2-i)^{\alpha-1} \right)^{2t} \leq \lim_{n \rightarrow \infty} \left(\frac{n^\alpha}{N_\alpha(n)} \right)^{2t_{n,\alpha}(\log n + c)} = 0.$$

We are left with the following two sums to control:

$$2 \sum_{k=1}^{\binom{n/4}{2}} \binom{n}{k}^2 k! \left(1 - \frac{(n-k)kN_{\alpha-1}(n)}{nN_\alpha(n)} \right)^{2t} + 2 \sum_{k>\binom{n/4}{2}}^{\binom{n}{2}} \binom{n}{k}^2 k! \left(1 - \frac{k}{n} \right)^{2t}.$$

The first sum at time $t = t_{n,\alpha}(\log n + c)$ may be reduced to:

$$\sum_{k=1}^{\binom{n/4}{2}} \binom{n}{k}^2 k! \left(1 - \frac{(n-k)kN_{\alpha-1}(n)}{nN_\alpha(n)} \right)^{2t_{n,\alpha}(\log n + c)} \leq e^{-2c} \sum_{k=1}^{n/4} \frac{n^{2k^2/n}}{k!}.$$

Following the work of Diaconis presented in Section 2.3.1 this may be shown to be bounded by Ae^{-2c} for sufficiently large n . For the second sum note that $N_\alpha(n)/N_{\alpha-1}(n)$ is increasing in α , and at $\alpha = 1$ we have $N_\alpha(n)/N_{\alpha-1}(n) = (n+1)/2$. The second summation has been shown to be tending to 0 at time $(n/2) \log n + cn$ as $n \rightarrow \infty$ by Diaconis [16, Chapter 3D Theorem 5]. Therefore, the sum tending to 0 at the larger time $t = (N_\alpha(n)/N_{\alpha-1}(n))(\log n + c)$ as $n \rightarrow \infty$. Putting these bounds together we get find,

$$\limsup_{n \rightarrow \infty} \|\text{OST}_{n,\alpha}^{t_{n,\alpha}(\log n + c)} - \pi_n\|_{\text{TV}} \leq Ae^{-c} \text{ for some universal constant } A.$$

This completes the upper bound for $\alpha \geq 1$ present in Theorem 3.6.6.

Lower Bound

For the case of $\alpha \geq 1$ a straightforward generalisation of Section 3.5 does not work. Instead, we are going to present a symmetric argument, swapping the roles of the left and right hands. For $\alpha \geq 1$, let $V_n = \{1, \dots, n/m\}$ (i.e. the bottom n/m cards) and

$$F_n = \{\sigma \in S_n \mid \sigma \text{ has at least 1 fixed point in } V_n\}.$$

We have already seen that $\pi_n(F_n) \leq 1/m$. We again want to bound the probability of $\text{OST}_{n,\alpha}(F_n)$ via a coupon collector's argument. Writing U_n^t for the set of uncollected cards in V_n after t steps of the biased one-sided transposition shuffle, it follows that

$$\text{OST}_{n,\alpha}^t(F_n) \geq \mathbb{P}(|U_n^t| \geq 1). \quad (3.86)$$

Given $\alpha \geq 1$ our right hand is now choosing the top of the deck with higher probability meaning it is not likely to touch any card in U_n^t . Whereas, our left hand is comparatively more likely to touch cards in U_n^t . In one step we still only have four choices: both hands touch uncollected cards, only one does (left or right), or neither does. This permits us to bound the change in the number of *collected* cards as follows:

$$\begin{aligned} |V_n \setminus U_n^{t+1}| &= |V_n \setminus U_n^t| + |\{L^{t+1}, R^{t+1}\} \cap U_n^t| \\ &\leq |V_n \setminus U_n^t| + 2 \cdot \mathbf{1}[R^{t+1} \in U_n^t] + \mathbf{1}[R^{t+1} \notin U_n^t, L^{t+1} \in U_n^t], \end{aligned} \quad (3.87)$$

This is the ‘reverse’ of equation (3.55), this time we increase by two if our right hand collects a card. For sufficiently large n we may bound the probability of the events in (3.87) as follows:

$$\begin{aligned} \mathbb{P}(R^{t+1} \in U_n^t) &\leq \frac{1}{N_\alpha(n)} \sum_{i=(n/m)-|U_n^t|}^{n/m} i^\alpha \leq \frac{|U_n^t|}{N_\alpha(n)} (n/m)^\alpha \leq \frac{|U_n^t|}{N_\alpha(n)} \frac{n^\alpha}{\alpha(m-1)} \\ &\leq \frac{|U_n^t| N_{\alpha-1}(n)}{N_\alpha(n)(m-1)} \end{aligned} \quad (3.88)$$

$$\mathbb{P}(R^{t+1} \notin U_n^t, L^{t+1} \in U_n^t) \leq \frac{|U_n^t|}{N_\alpha(n)} \sum_{i=|U_n^t|+1}^n \frac{i^\alpha}{i} \leq \frac{|U_n^t| N_{\alpha-1}(n)}{N_\alpha(n)}. \quad (3.89)$$

In the last inequality of (3.88) we have used the fact that $n^\alpha/\alpha \leq N_{\alpha-1}(n)$ for $\alpha \geq 1$. Using (3.87), (3.89) and (3.88) together, we now define a counting process $M_{n,\alpha}^t$ which stochastically dominates the number of collected cards $|V_n \setminus U_n^t|$ at all times:

$$\begin{aligned} M_{n,\alpha}^0 &= 0; \\ \mathbb{P}(M_{n,\alpha}^{t+1} = M_{n,\alpha}^t + k) &= \begin{cases} \frac{N_{\alpha-1}(n)}{N_\alpha(n)(m-1)} \left(\frac{n}{m} - M_{n,\alpha}^t\right) & \text{if } k = 2 \\ \frac{N_{\alpha-1}(n)}{N_\alpha(n)} \left(\frac{n}{m} - M_{n,\alpha}^t\right) & \text{if } k = 1 \\ 1 - \frac{m}{m-1} \frac{N_{\alpha-1}(n)}{N_\alpha(n)} \left(\frac{n}{m} - M_{n,\alpha}^t\right) & \text{if } k = 0. \end{cases} \end{aligned} \quad (3.90)$$

Setting $m = \log n$, this forms a valid probability distribution for n sufficiently large. The counting process $M_{n,\alpha}$ is analogous to the process M_n defined in (3.58) with $1/n$ replaced with $N_{\alpha-1}(n)/N_\alpha(n)$. We recover the bound,

$$\text{OST}_{n,\alpha}^t(F_n) \geq \mathbb{P}(M_{n,\alpha}^t < n/m).$$

Lemma 3.6.17. *Let $\mathcal{T} = \min\{t : M_n^t \geq n/\log n\}$. Then for any $c > 2$,*

$$\mathbb{P}(\mathcal{T} \leq t_{n,\alpha}(\log n - n \log \log n - cn)) \leq \frac{\pi^2}{6(c-2)^2}.$$

Proof. This follows from the proof of Lemma 3.5.1 accounting for the change of $1/n$ to $N_{\alpha-1}(n)/N_\alpha(n)$. \square

Following from Lemma 3.6.17 and taking $c > 2$ we find,

$$\begin{aligned} \|\text{OST}_{n,\alpha}^{t_{n,\alpha}(\log n - \log \log n - c)} - \pi_n\|_{\text{TV}} &\geq \mathbb{P}(\mathcal{T} > t_{n,\alpha}(\log n - \log \log n - c)) - \pi_n(F_n) \\ &\geq 1 - \frac{\pi^2}{6(c-2)^2} - \frac{1}{\log n}. \end{aligned}$$

Taking limits in n recovers the limit present in Theorem 3.6.6 for biased shuffles with $\alpha \geq 1$.

3.6.3 Cutoff for Biased One-sided Transposition Shuffles with General Weight Functions

We have established a total variation cutoff for the biased one-sided transposition shuffles where the weight function takes the form $w(j) = j^\alpha$ for some α . In this section we look to extend our previous results to a more general class of weight functions. Notice that j^α/j is monotonically increasing if and only if $\alpha \geq 1$ and monotonically decreasing if and only if $\alpha \leq 1$. We now consider the class of weight functions $w : \mathbb{N} \rightarrow (0, \infty)$ with the property $w(j)/j$ is monotonic. We first consider the case where $w(j)/j$ is monotonically decreasing (corresponding to $\alpha \leq 1$). In this situation we can prove the following extension of Lemma 3.6.7.

Lemma 3.6.18. *Let $\lambda \vdash n$ with $\lambda_1 = n - k$. Then the eigenvalue $\text{eig}(T_\lambda^{\rightarrow})$ for the shuffle $\text{OST}_{n,w}$ with $w(j)/j$ monotonically decreasing may be bounded as follows:*

$$\text{eig}(T_\lambda^{\rightarrow}) \leq \begin{cases} 1 - \frac{(n-k+1)kw(n)}{nN_w(n)} & \text{if } k \leq n/4 \\ 1 - \frac{kw(n)}{2N_w(n)} & \text{if } k > n/4. \end{cases}$$

Proof. This follows the same proof as Lemma 3.6.7 with m^α/m (or $m^{\alpha-1}$) replaced with $w(m)/m$. □

Corollary 3.6.19. *The mixing time of the biased one-sided transposition shuffle $\text{OST}_{n,w}$ with $w(j)/j$ monotonically decreasing is at most $(N_w(n)/w(n)) \log n$.*

Proof. Applying the same argument as Section 3.6.1 Upper Bound with the new bounds from Lemma 3.6.18 immediately yields the desired result. □

To complete the existence of a total variation cutoff we need to find a matching lower bound of the same time. However, the argument presented in Section 3.6.1 is not easily generalised for our new choices of weight function. The main issue being the relationship between $w(k)$ and $w(n)$ for $k \in V_n = \{n - n/m, \dots, n\}$. In the previous case with $w(j) = j^\alpha$ we can easily see that $w(k) = w(k/n)w(n)$. In the current situation we only know $w(n)/n \leq w(k)/k$; this is not enough information to establish a relationship between $w(k)$ and $w(n)$ for use in equations (3.76), and (3.77). A new approach is likely to be needed to prove a lower bound on the mixing time of the biased one-sided transposition shuffle for this class of weight functions.

Conjecture 3.6.20. *The biased one-sided transposition shuffle $\text{OST}_{n,w}$ with $w(j)/j$ monotonically decreasing exhibits a total variation cutoff at time $(N_w(n)/w(n)) \log n$.*

We now consider the class of weight functions where $w(j)/j$ is monotonically increasing (corresponding to $\alpha \geq 1$). Previously we proved the mixing time for the biased one-sided transposition shuffle with $\alpha \geq 1$ to be $(N_\alpha(n)/N_{\alpha-1}(n)) \log n$. For a general function w we replace the factor of $N_{\alpha-1}(n)$ by the summation $N'_w(n) := \sum_{i=1}^n w(i)/i$. Therefore, we now look to prove a cutoff at time $(N_w(n)/N'_w(n)) \log n$. The lower bound argument presented in Section 3.6.2 Lower Bound, can easily be modified for this new class of weight functions.

Lemma 3.6.21. *The mixing time of the biased one-sided transposition shuffle $\text{OST}_{n,w}$ with $w(j)/j$ monotonically decreasing is at least $(N_w(n)/N'_w(n)) \log n$.*

Proof. The result follows from the work of Section 3.6.2 Lower Bound with $N_{\alpha-1}(n)/N_\alpha(n)$ replaced by $N'_w(n)/N_w(n)$, and the inequalities (3.88) and (3.89) replaced with:

$$\begin{aligned} \mathbb{P}(R^{t+1} \in U_n^t) &\leq \frac{1}{N_w(n)} \sum_{i=(n/m)-|U_n^t|}^{n/m} w(i) \leq \frac{|U_n^t|}{(m-1)N_w(n)} (m-1)w(n/m) \\ &\leq \frac{|U_n^t|}{(m-1)N_w(n)} \sum_{i=1}^n \frac{w(i)}{i} = \frac{|U_n^t|N'_w(n)}{(m-1)N_w(n)} \end{aligned} \quad (3.91)$$

$$\mathbb{P}(R^{t+1} \notin U_n^t, L^{t+1} \in U_n^t) \leq \frac{|U_n^t|}{N_w(n)} \sum_{i=|U_n^t|+1}^n \frac{w(i)}{i} \leq \frac{|U_n^t|N'_w(n)}{N_w(n)}. \quad (3.92)$$

□

To find a matching upper bound on the mixing time we look to generalise the results of Section 3.6.2 Upper Bound. For a general weight function with $w(j)/j$ monotonically increasing Lemma 3.6.12 still holds and so we can bound the eigenvalues associated to a partition λ using the value $\text{eig}(T_\lambda^\lambda)$. For small partitions we are able to use the bound constructed in Lemma 3.6.16 without any modifications. For large partitions we encounter an issue because the proof of the bound in Lemma 3.6.16 relies on exact knowledge of the weight function w . In particular, we use the integral of w in order to obtain an upper bound on the value of $\text{eig}(T_\lambda^\lambda)$, this argument can not be generalised to a generic weight function w . However, we conjecture that an upper bound of the correct mixing time is possible.

Conjecture 3.6.22. *The biased one-sided transposition shuffle $\text{OST}_{n,w}$ with $w(j)/j$ monotonically decreasing exhibits a total variation cutoff at time $(N_w(n)/N'_w(n)) n \log n$.*

The class of functions $\{w : \mathbb{N} \rightarrow (0, \infty) \mid w(j)/j \text{ is monotonic}\}$ is still quite restrictive. We suspect that cutoffs in total variation distance may be proved for biased one-sided transposition shuffles with more general weight function w (for example, slowly and regularly varying functions), but any proof of this would certainly need different techniques to those developed in this section.

3.7 Separation Distance for the One-sided Transposition Shuffle

3.7.1 Cutoff for the Unbiased One-sided Transposition Shuffle

The unbiased one-sided transposition shuffle exhibits a cutoff in total variation distance at time $n \log n$. To prove that we also have a cutoff in separation distance at time $n \log n$ we identify a strong stationary time for the unbiased one-sided transposition shuffle. The aim of this section is to prove the following result.

Theorem 3.7.1. *There exists a strong stationary time T for the one-sided transposition shuffle, with $\mathbb{P}(T > n \log n + cn) \leq e^{-c}$.*

From Theorem 3.7.1 we may quickly establish that the one-sided transposition shuffle exhibits a cutoff in separation distance at time $n \log n$.

Theorem 3.7.2. *The one-sided transposition shuffle exhibits a cutoff in separation distance at time $n \log n$.*

Proof of Theorem 3.7.2. The lower limit on separation distance mixing time follows from Lemma 1.1.24 and Theorem 3.1.2. The upper limit on separation distance mixing time follows from Theorem 3.7.1. □

A Strong Stationary Time Argument

Recall that we may view the elements of S_n as the possible permutations of deck of cards which is made up of cards and positions both indexed by $[n]$. Any permutation $\sigma \in S_n$ is a bijection from cards to positions and σ^{-1} is a bijection from positions to cards, i.e. $\sigma(i)$ tells us the position of card i whereas $\sigma^{-1}(i)$ tells us what card is in position i . We begin all our random walks at the identity permutation with positions and labels fully matched. Throughout the rest of this section let $(X^t)_{t \in \mathbb{N}}$ denote a Markov chain driven by the unbiased one-sided transposition shuffle, and let $(Y^t)_{t \in \mathbb{N}}$ be a Markov chain defined by setting $Y^t = (X^t)^{-1}$ for all times t . For any $\sigma \in S_n$ and time t we have $\mathbb{P}(X^t = \sigma) = \mathbb{P}(Y^t = \sigma^{-1})$. The Markov chains (X^t) and (Y^t) represent two different ways to view the one-sided transposition shuffle.

Let τ^t be the transposition chosen at step t of the unbiased one-sided transposition shuffle. To construct our strong stationary time we need to condition on the exact permutation of cards in positions above position j at time t , that is the random variables $Y^t(i)$ for $j < i \leq n$. Given this information we also know which cards can be in positions 1 to j at time t , define this set as,

$$A_j^t = [n] \setminus \{Y^t(i) \mid j < i \leq n\}.$$

Definition 3.7.3. We say the random walk $(Y^t)_{t \in \mathbb{N}}$ satisfies *property \mathcal{P}_j at time t* if we have:

$$\mathbb{P}(Y^t(j) = l \mid Y^t(i) \text{ for all } j < i \leq n) = \begin{cases} 1/j & \text{if } l \in A_j^t \\ 0 & \text{otherwise} \end{cases} \quad (3.93)$$

This property tells us that given total information about the deck strictly above position j , the card in position j is equally likely to be any of the remaining cards.

Lemma 3.7.4. *Let T_j be the first time our right hand chooses the position j when performing the unbiased one-sided transposition shuffle. If $T_j \leq t$ then the Markov chain $(Y^t)_{t \in \mathbb{N}}$ satisfies property \mathcal{P}_j at time t .*

Proof. We prove this by induction: once property \mathcal{P}_j holds for some t , it holds for all times after t .

Consider the time T_j , at this step of our Markov chain we must have applied a transposition (ij) with $i \leq j$. The probability of picking any one of the transpositions (ij) at time T_j is $\mathbb{P}(\tau^{T_j} = (ij)) = 1/j$ for all $i \leq j$. Therefore, the card in position j at time T_j has a uniform chance of being any of the cards in $A_j^{T_j-1} = A_j^{T_j}$. Thus, we may clearly see that,

$$\mathbb{P}(Y^{T_j}(j) = l | Y^{T_j}(i) \text{ for all } j < i \leq n) = \begin{cases} 1/j & \text{if } l \in A_j^{T_j} \\ 0 & \text{otherwise} \end{cases} \quad (3.94)$$

so \mathcal{P}_j holds at time T_j .

Now suppose property \mathcal{P}_j holds at time t . We study the time $t+1$ and split the analysis into cases based on which transposition (ab) (with $a \leq b$) was applied at time $t+1$,

$$\begin{aligned} & \mathbb{P}(Y^{t+1}(j) = l | Y^{t+1}(i) \text{ for all } j < i \leq n) \\ &= \sum_{1 \leq a \leq b \leq n} \mathbb{P}(\tau^{t+1} = (ab)) \mathbb{P}(Y^{t+1}(j) = l | Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = (ab)). \end{aligned} \quad (3.95)$$

Using knowledge of the transposition τ^{t+1} we evolve our deck backwards in time ($X^t = \tau^{t+1} X^{t+1}$) to recover the random variables $Y^t(i)$ from $Y^{t+1}(i)$ for $j < i \leq n$, and relate $Y^{t+1}(j)$ to $Y^t(j)$. This allows us to use our inductive hypothesis.

If $b = j$ then our random walk satisfies property \mathcal{P}_j at time $t+1$ for the same reasoning as time T_j . Suppose $a \leq b < j$ then we know that $Y^t(i) = Y^{t+1}(i)$ for all $j < i \leq n$. Suppose instead that $a, b > j$, then we have $Y^t(b) = Y^{t+1}(a)$ and $Y^t(a) = Y^{t+1}(b)$, with $Y^t(i) = Y^{t+1}(i)$ for all other $j < i \leq n$. In either case we know $A_j^{t+1} = A_j^t$, and the card in position j has not moved from time t to $t+1$. Therefore, we have

$$\begin{aligned} & \mathbb{P}(Y^{t+1}(j) = l | Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = (ab)) \\ &= \mathbb{P}(Y^t(j) = l | Y^t(i) \text{ for all } j < i \leq n) = \begin{cases} 1/j & \text{if } l \in A_j^t = A_j^{t+1} \\ 0 & \text{otherwise} \end{cases}. \end{aligned} \quad (3.96)$$

We now study in detail the effects of the remaining transpositions (ab) with $a \leq j < b$. In this case we can not fully recover the random variables $Y^t(i)$ with $i < j \leq n$ without extra assumptions. To this end fix $b > j$, a card $C \in A_j^{t+1}$, and suppose that $\tau^{t+1} = (X^{t+1}(C) b)$, i.e., card C is moved from position b into a position below j by τ^{t+1} . Letting C range over all choices in A_j^{t+1} will recover every transposition (ab) with $b > j$ fixed and $a \leq j$. In the case that $\tau^{t+1} = (X^{t+1}(C) b)$, we know that $Y^t(b) = C$ and the other positions above j have $Y^t(i) = Y^{t+1}(i)$ for $j < i \leq n$ and $i \neq b$. Therefore, for this choice of τ^{t+1} we know

that $A_j^t = (A_j^{t+1} \sqcup \{Y^{t+1}(b)\}) \setminus \{C\}$. Now consider the probability:

$$\mathbb{P}(Y^{t+1}(j) = l | Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = (X^{t+1}(C) \ b)).$$

If $l = C$ then the event in question can only occur if the card currently in position b , i.e. $Y^{t+1}(b)$, was in position j at time t . Noting that $Y^{t+1}(b) \in A_j^t$, and using our inductive hypothesis we find,

$$\begin{aligned} & \mathbb{P}(Y^{t+1}(j) = C | Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = (X^{t+1}(C) \ b)) \\ &= \mathbb{P}(Y^t(j) = Y^{t+1}(b) | Y^t(i) \text{ for all } j < i \leq n) = 1/j \end{aligned} \quad (3.97)$$

Alternatively, suppose $l \in A_j^{t+1} \setminus \{C\}$, we know the card l does not move from its position at time t to time $t+1$, and we know $l \in A_j^t$. Therefore, we find

$$\begin{aligned} & \mathbb{P}(Y^{t+1}(j) = l | Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = (X^{t+1}(C) \ b)) \\ &= \mathbb{P}(Y^t(j) = l | Y^t(i) \text{ for all } j < i \leq n) = 1/j \end{aligned} \quad (3.98)$$

Putting the equations (4.48) and (4.50) together gives us,

$$\mathbb{P}(Y^{t+1}(j) = l | Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = (X^{t+1}(C) \ b)) = \begin{cases} 1/j & \text{if } l \in A_j^{t+1} \setminus \{C\} \\ 1/j & \text{if } l = C \\ 0 & \text{otherwise} \end{cases}.$$

Letting C range over all possible choices of card in A_j^{t+1} while keeping $b > j$ fixed, we cover the desired probability for every transposition (ab) with $a \leq j < b$. Finally applying every separate case to (3.95) we have established that

$$\mathbb{P}(Y^{t+1}(j) = l | Y^{t+1}(i) \text{ for all } j < i \leq n) = \begin{cases} 1/j & \text{if } l \in A_j^{t+1} \\ 0 & \text{otherwise} \end{cases}$$

as required, thus by induction our hypothesis holds for all $t \geq T_j$. □

Lemma 3.7.5. *Let T be the first time our right hand has chosen every position. Then T is a strong uniform time for $(X^t)_{t \in \mathbb{N}}$ and $(Y^t)_{t \in \mathbb{N}}$.*

Proof. Note that X^t is uniformly distributed if and only if Y^t is uniformly distributed. Lemma 3.7.4 implies that by time $T \geq T_j$ the Markov chain $(Y^t)_{t \in \mathbb{N}}$ satisfies all properties \mathcal{P}_j . Hence, for any $\sigma \in S_n$,

$$\begin{aligned} \mathbb{P}(Y^t = \sigma^{-1} | T \leq t) &= \mathbb{P}(\cap_{j=1}^n \{Y^t(j) = \sigma^{-1}(j)\} | T \leq t) \\ &= \prod_{j=1}^n \mathbb{P}(Y^t(j) = \sigma^{-1}(j) | \cap_{i=j+1}^n \{Y^t(i) = \sigma^{-1}(i)\}, T \leq t) \\ &= \prod_{j=1}^n \frac{1}{j} = \pi_n(\sigma). \end{aligned}$$

□

We have found a strong stationary time for the unbiased one-sided transposition shuffle. Following quickly from this we may prove Theorem 3.7.1, and thus establish a cutoff in separation distance for the one-sided transposition shuffle.

Proof of Theorem 3.7.1. Let T be the first time our right hand has chosen every position j . Our right hand is choosing positions via a uniform probability on $[n]$. Thus T is modelled by the standard coupon collector's problem with n coupons. To complete our argument recall (Section 2.3.2, equation (2.19)) that for the standard coupon collector's problem on n cards we have $\mathbb{P}(T_n > n \log n + cn) \leq e^{-c}$. \square

The strong stationary time T is novel in a few ways. Often strong stationary times rely on a sequence of stopping times $(T_i)_{i \in \mathbb{N}}$ where at time T_i we know there exists a subset of i cards which are uniformly distributed with respect to S_i . This is exactly the technique employed to find a strong stationary time for the top-to-random shuffle which we covered in Section 2.3.2. The strong stationary time T can not be split up in this way. Suppose we require that we touch positions in order from 1 to n , then after we choose position i the cards in the subset of positions $[i] \subseteq [n]$ are uniformly distributed on S_i , similarly to the top-to-random shuffle. However, if instead we require that positions are touched in reverse order, i.e. from n to 1, then we do not find any uniform subsets of S_n before becoming completely uniform when we touch card 1. Therefore, because we do not impose a condition on the order in which we must touch positions to get to the time T , at any particular time before T there is no certainty that we have a uniformly distributed subset of the cards $[n]$. This sudden uniformity is the kind of behaviour which drives the existence of a cutoff in separation distance.

3.7.2 Generalising for Biased One-sided Transposition Shuffles

In Section 3.6 we were able to prove cutoff in total variation for all biased one-sided transposition shuffles. Given that the unbiased one-sided transposition shuffle $\text{OST}_{n,0}$ also shows a cutoff in separation at the same time as for total variation we may hope the same can be proven for all biased shuffles $\text{OST}_{n,\alpha}$. However, in this section we show that simply extending the method used for $\text{OST}_{n,0}$ is not enough to prove the existence of a cutoff in separation distance for all α . Recall that for a biased one-sided transposition shuffle $\text{OST}_{n,\alpha}$ our right hand chooses a position j with probability $j^\alpha/N_\alpha(n)$.

Lemma 3.7.6. *Let T_α be the time for the right hand of the biased shuffle $\text{OST}_{n,\alpha}$ to choose every position in $[n]$ at least once. Then T_α is a strong stationary time for the shuffle $\text{OST}_{n,\alpha}$, for all α .*

Proof. For the biased shuffle $\text{OST}_{n,\alpha}$ our left hand is still choosing uniformly below our right hand, therefore this result follows from the proofs of Lemmas 3.7.4 and 3.7.5. \square

Given T_α is still a strong stationary time all we have left to do is analyse the probability $\mathbb{P}(T_\alpha > t)$. This reduces to a biased coupon collector's problem. The next lemma establishes a bound on the time T_α which will allow us to prove a cutoff in separation distance for biased one-sided transposition shuffles with $\alpha \leq 0$.

Lemma 3.7.7. Let $t_{n,\alpha}$ be the time defined in Theorem 3.6.6, that is

$$t_{n,\alpha} = \begin{cases} N_\alpha(n)/n^\alpha & \text{if } \alpha \leq 1 \\ N_\alpha(n)/N_{\alpha-1}(n) & \text{if } \alpha \geq 1 \end{cases}.$$

For the biased one-sided transposition shuffle with $\alpha \leq 0$ we have that $\mathbb{P}(T_\alpha > t_{n,\alpha}(\log n + c)) \leq e^{-c}$.

Proof. We say position j is *collected* the first time it is chosen by our right hand. Its collection probability at each step of our random walk is $j^\alpha/N_\alpha(n)$. We bound the time T_α by a coupon collector's argument with inspiration taken from the top to random shuffle (Lemma 2.3.9). Let C_j^t be the event that we have not collected position j by time t . We may find a simple upper bound as follows:

$$\mathbb{P}(T_\alpha > t) = \mathbb{P}(\cup_{j=1}^n C_j^t) \leq \sum_{j=1}^n \mathbb{P}(C_j^t) = \sum_{j=1}^n \left(1 - \frac{j^\alpha}{N_\alpha(n)}\right)^t. \quad (3.99)$$

Now substituting $t = t_{n,\alpha}(\log n + c) = (N_\alpha(n)/n^\alpha)(\log n + c)$ for $\alpha \leq 0$, and applying the bound $(1-x)^t \leq e^{-tx}$, we find,

$$\mathbb{P}(T_\alpha > t) \leq \sum_{j=1}^n \left(1 - \frac{j^\alpha}{N_\alpha(n)}\right)^{t_{n,\alpha}(\log n + c)} \leq e^{-c} \sum_{j=1}^n n^{-(j/n)^\alpha} \leq e^{-c} n^{1-(1)^\alpha} \leq e^{-c},$$

with the second to last inequality following from $(j/n)^\alpha = (n/j)^{-\alpha} \geq 1^{-\alpha}$ for $\alpha \leq 0$. □

Theorem 3.7.8. The biased one-sided transposition shuffle $\text{OST}_{n,\alpha}$ with $\alpha \leq 0$ exhibits a cutoff in separation distance at time $t_{n,\alpha} \log n$.

Proof. The lower bound on separation distance mixing time follows from Lemma 1.1.24 and Theorem 3.6.6. The upper bound on separation distance mixing time follows from Theorem 3.7.1 □

The problem stopping us from copying Lemma 3.7.7 for shuffles with $\alpha > 0$ is the approximation in equation (3.99). For $\alpha \in (0, 1]$ we still find that

$$\mathbb{P}(T_\alpha > t) \leq \sum_{j=1}^n \left(1 - \frac{j^\alpha}{N_\alpha(n)}\right)^{t_{n,\alpha}(\log n + c)} \leq e^{-c} \sum_{j=1}^n n^{-(j/n)^\alpha}.$$

The summation present above is unbounded as $n \rightarrow \infty$ for $\alpha \in (0, 1]$. Similarly for $\alpha \in [1, \infty)$ we may find the bound (3.99) to be,

$$\mathbb{P}(T_\alpha > t) \leq \sum_{j=1}^n \left(1 - \frac{j^\alpha}{N_\alpha(n)}\right)^{t_{n,\alpha}(\log n + c)} \leq \sum_{j=1}^n n^{-j^\alpha/N_{\alpha-1}(n)} e^{-c(j^\alpha/N_{\alpha-1}(n))}$$

and this summation is still unbounded as $n \rightarrow \infty$. Therefore, to use bound (3.99) for $\alpha > 0$ we would need to increase the time we consider from $t_{n,\alpha} \log n$ until the respective

summation is bounded in n . For example if $\alpha = 1$ the summation (3.99) becomes

$$e^{-c} \sum_{j=1}^n \left(1 - \frac{2j}{n(n+1)} \right)^t$$

which is bounded at time $t = O(n^2)$, which is of a greater order than our lower bound of $(n+1)/2 \log n$ from Theorem 3.6.6.

Overall the bound presented in the proof of Lemma 3.7.7 is not good enough to establish an upper bound on our separation distance mixing time that matches our lower bound of $t_{n,\alpha} \log n$ for the biased one-sided transposition shuffles with $\alpha > 0$.

The shuffle $\text{OST}_{n,1}$ is closely related to the random transposition shuffle. Strong stationary times for the random transposition shuffle have historically been difficult to find. In [11] Broder found a strong stationary time for random transpositions which gave an upper bound on the mixing time of $2n \log n$. Matthews claimed to have improved this to $n \log n$ and then to $(n/2) \log n$, thus establishing a cutoff in separation distance [33]. However, the argument Matthew presented had a subtle mistake which was only recently discovered by White and explored in his thesis [44]. White has constructed a strong stationary time for random transpositions which gives the correct upper bound of $(n/2) \log n$ [45]. The proof of this involves the creation of a sophisticated stationary time involving keeping track of the possible cycle structure the deck could be in at certain times. In light of this it is not a surprise that a simple generalisation does not yield a satisfactory upper bound on separation distance for the shuffle $\text{OST}_{n,1}$. Despite not finding a strong stationary time for the biased one-sided transposition shuffles with $\alpha > 0$, we conjecture that they show a cutoff in separation distance at the same time as their respective total variation cutoff.

Conjecture 3.7.9. *The biased one-sided transposition shuffle $\text{OST}_{n,\alpha}$ with $\alpha \in (0, \infty)$ exhibits a cutoff in separation distance at $t_{n,\alpha} \log n$.*

Chapter 4

The Hyperoctahedral Group and Random Walks

The hyperoctahedral group B_n is a natural extension of the symmetric group. We may view the hyperoctahedral group as the arrangement of a deck of n cards with the extra information of knowing whether a given card is face up or face down. The group has a lot of structural similarity with the symmetric group; its elements may still be decomposed into a product of different cycles, and the conjugacy classes are determined by cycle type. This allows us to extend shuffles on S_n into those on B_n , in particular we focus on extensions of the random transposition shuffle and one-sided transposition shuffle. The module structure of the hyperoctahedral group also resembles that of S_n with its simple modules being generalised Specht modules $S^{\bar{\lambda}}$ labelled now by bi-partitions of n . This enables us to modify the technique of lifting eigenvectors for the hyperoctahedral group and extended shuffles.

In the first few sections we cover the basics of the hyperoctahedral group and its module structure, good references for the material we cover are *The Representations of the Weyl Groups of Type B_n* by Aamili, Morris and Peel [1], and *Representations of the hyperoctahedral Groups* by Geissinger and Kinch [22]. The last two sections detail lifting eigenvectors for the hyperoctahedral group, we recover the full spectrum of the random transposition shuffle and one-sided transposition shuffle.

4.1 The Hyperoctahedral Group

Define the set $[\pm n] := \{1, \dots, n\} \cup \{-1, \dots, -n\}$. The hyperoctahedral group on n elements, denoted B_n , may be defined as the group of all bijections $\sigma : [\pm n] \rightarrow [\pm n]$ such that $\sigma(-i) = -\sigma(i)$, there are $2^n n!$ such bijections. Any map $\sigma \in B_n$ is completely determined by its values on the positive elements $\{1, \dots, n\}$. Another common way to think about the group B_n is the arrangements of a deck of n cards where we may now distinguish between cards that are face up or face down. The bijection σ tells us what position card i is in and whether it is face up ($\sigma(i)$ positive), or face down ($\sigma(i)$ negative). The last way to view the hyperoctahedral group is as the wreath product $\mathbb{Z}_2 \wr S_n$. We may identify important subgroups $S_n = \{\sigma \mid \sigma(i) > 0 \text{ if } i > 0\}$ and $\mathbb{Z}_2^n = \{\sigma \mid \sigma(i) = \pm i \text{ for all } i\}$ inside the hyperoctahedral group.

Define multiplication performed in the group B_n to be the composition of functions from right to left. We may decompose elements of B_n into a product of negative transpositions, and a permutation from S_n . This allows us to use cycle notation from S_n (Section 2.1) to represent the elements of B_n .

Definition 4.1.1. The conjugacy class of *negative transpositions* in B_n is composed of the elements ξ_i for $1 \leq i \leq n$, with:

$$\xi_i(j) = \begin{cases} -i & \text{if } j = i \\ +i & \text{if } j = -i \\ j & \text{otherwise} \end{cases} .$$

The elements ξ_i can be thought of as flipping card i over.

Lemma 4.1.2. Let $\sigma \in B_n$. We may decompose σ into a permutation from S_n and a product of negative transpositions. I.e., there exists $x_i \in \{0, 1\}$ and $\eta \in S_n \subset B_n$ such that,

$$\sigma = \left(\prod_{i=1}^n \xi_i^{x_i} \right) \eta.$$

Proof. By applying negative transpositions to σ we may recover an element of S_n inside of B_n , i.e., there exists a collection of $y_i \in \{0, 1\}$ such that, $(\prod_{i=1}^n \xi_i^{y_i}) \sigma = \eta$ for some $\eta \in S_n$. From here taking inverses of the elements ξ_i gives the desired form. \square

The conjugacy class of *positive transpositions* for B_n is formed from the elements (ij) and $\xi_i \xi_j (ij)$ with $i < j$. The set of *transpositions* for the hyperoctahedral group is the union of the sets of positive and negative transpositions. Lemma 4.1.2 may be restated as; every element of B_n may be decomposed into a product of transpositions. Previously we described the conjugacy classes of S_n in terms of their cycle type. Any element of B_n may be split into positive and negative cycles which define the cycle type of that element. To label the possible cycles types for B_n we introduce the concept of bi-partitions.

Definition 4.1.3. Let $n \in \mathbb{N}^0$, a *bi-partition* of n , denoted $\bar{\lambda}$, is a pair of partitions $\bar{\lambda} = (\lambda^1, \lambda^2)$ such that $\lambda^1 \vdash m$ and $\lambda^2 \vdash n - m$ for some $0 \leq m \leq n$. If $\bar{\lambda}$ is a bi-partition of n we write $\bar{\lambda} \vdash n$.

Definition 4.1.4. Let $\sigma \in B_n$ with decomposition $\sigma = (\prod_{i=1}^n \xi_i^{x_i}) \eta$. We call a cycle in η a *positive cycle* if $\sum_j x_j \equiv 0 \pmod{2}$ for elements j in the cycle. We call a cycle in η a *negative cycle* if $\sum_j x_j \equiv 1 \pmod{2}$ for elements j in the cycle. Form two non-increasing tuples λ^1, λ^2 of the lengths of positive and negative cycles respectively. The *cycle type* of σ is defined as the bi-partition (λ^1, λ^2) .

Example 4.1.5. Let $n = 6$ and consider $\sigma \in B_6$ as defined by,

$$\sigma = (\xi_1 \xi_2 \xi_4 \xi_5) (1)(234)(56).$$

The element σ only has one positive cycle of length 3, being $\xi_2 \xi_4 (234)$. We also have two negative cycles, the element ξ_1 forms a cycle of length 1, and the elements $\xi_5 (56)$ giving a cycle of length 2. Therefore, σ has cycle type $((3), (2, 1))$.

Lemma 4.1.6. *Elements of B_n belong to the same conjugacy class if and only if they share the same cycle type. Hence, bi-partitions label the conjugacy classes of B_n .*

The conjugacy classes of negative and positive transpositions have cycle types $((1^{n-1}), (1))$ and $((2, 1^{n-2}), (0))$ respectively. Following from the decomposition of Lemma 4.1.2, we may introduce the generalised sign function for the hyperoctahedral group, which counts the number of transpositions (positive and negative) needed to form an element.

Definition 4.1.7. Let $\sigma \in B_n$ with decomposition $\sigma = (\prod_{i=1}^n \xi_i^{x_i}) \eta$. Define the *sign function* for the hyperoctahedral group, denoted $\text{sgn} : B_n \rightarrow \{-1, 1\}$, as follows,

$$\text{sgn}(\sigma) = (-1)^{\sum_{i=1}^n x_i} \text{sgn}(\eta).$$

An element in B_n is called *even* if it has positive sign, and is called *odd* if it has negative sign. The sign function on B_n is multiplicative, that is $\text{sgn}(\sigma\tau) = \text{sgn}(\sigma)\text{sgn}(\tau)$. Thus, the set of positive elements of B_n forms a normal subgroup of index 2.

4.2 The Structure of Modules for the Hyperoctahedral Group.

Define the group algebra of the hyperoctahedral group as $\mathfrak{B}_n := \mathbb{C}[B_n]$. To construct the modules of the hyperoctahedral group we extend the notions of partitions and tableaux to bi-partitions and bi-tableaux. This new notation allows for the formulation of permutation and Specht modules for B_n following a combinatoric method similar to that of S_n .

4.2.1 Bi-partitions

Every bi-partition has an associated *Young diagram*, which is formed from the Young diagrams for partitions λ^1, λ^2 separately. For example, the bi-partition $((3, 1), (2^2, 1))$ has Young diagram

$$\left(\begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \square & & \\ \hline \end{array}, \begin{array}{|c|c|} \hline \square & \square \\ \hline \square & \square \\ \hline \square & \\ \hline \end{array} \right).$$

In a Young diagram we denote the empty partition (0) as \emptyset . We refer to the boxes of $\bar{\lambda}$ by triples (i, j, k) , which mean box (i, j) in λ^k . For any bi-partition $\bar{\lambda}$ we may form the *transpose* of $\bar{\lambda}$, denoted $\bar{\lambda}'$, by swapping the partitions λ^1, λ^2 and taking their transposes separately, i.e., $\bar{\lambda}' = (\lambda^{2'}, \lambda^{1'})$. For example the transpose of bi-partition $((3, 1), (2^2, 1))$ is $((3, 2), (2, 1, 1))$, which has corresponding Young diagram,

$$\left(\begin{array}{|c|c|} \hline \square & \square \\ \hline \square & \square \\ \hline \square & \\ \hline \end{array}, \begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \square & & \\ \hline \square & & \\ \hline \end{array} \right).$$

We may extend the dominance ordering on partitions to a partial ordering on bi-partitions in the following way.

Definition 4.2.1. Let $\bar{\lambda}, \bar{\mu}$ be bi-partitions. Define the *dominance ordering* on bi-partitions as follows:

$$\bar{\lambda} \succeq \bar{\mu} \Leftrightarrow \begin{cases} \text{we have} & |\lambda^1| > |\mu^1|, \\ \text{or} & |\lambda^1| = |\mu^1| \text{ and } \lambda^1 \succeq \mu^1, \lambda^2 \succeq \mu^2. \end{cases}$$

In the case $\bar{\lambda} \succeq \bar{\mu}$ we say $\bar{\lambda}$ dominates $\bar{\mu}$.

Consider the bi-partitions $\bar{\lambda} = ((2, 1), (2, 2))$, $\bar{\mu} = ((3), (2, 1, 1))$, we do not have $\bar{\lambda}$ dominating $\bar{\mu}$ nor vice versa, because $\mu^1 \succeq \lambda^1$ and $\lambda^2 \succeq \mu^2$. Therefore, the dominance order of bi-partitions is not necessarily a total ordering (however, it may be for small n). The dominance ordering does have maximal and minimal elements, which are given by $((n), (0))$ and $((0), (1^n))$ respectively. Note that if $\bar{\lambda} \succeq \bar{\mu}$ we may form $\bar{\mu}$ from $\bar{\lambda}$ by moving boxes from λ^1 to λ^2 , or moving boxes down and to the left within each separate partition.

Lemma 4.2.2. *Let $\bar{\lambda}, \mu$ be bi-partitions of n . Then $\bar{\lambda} \succeq \bar{\mu}$ if and only if $\bar{\mu}' \succeq \bar{\lambda}'$.*

Proof. Suppose $\bar{\lambda} \succeq \bar{\mu}$, if $|\lambda^1| > |\mu^1|$ then the lemma follows immediately from the definitions of transpose and dominance ordering. If $|\lambda^1| = |\mu^1|$ then also $|\lambda^2| = |\mu^2|$ and we know $\lambda^k \succeq \mu^k$ if and only if $\mu^{k'} \succeq \lambda^{k'}$. Hence, $\bar{\mu}' \succeq \bar{\lambda}'$. \square

Given two bi-partitions $\bar{\lambda}, \bar{\mu}$ of possibly different sizes, we say $\bar{\lambda} \subseteq \bar{\mu}$ if $\lambda^1 \subseteq \mu^1$ and $\lambda^2 \subseteq \mu^2$. Given a bi-partition of n we may turn it into a bi-partition of $n + 1$ by adding a single box to one of the partitions λ^1, λ^2 . Let the element e_i^1 denote the tuple $(e_i, (0))$ and $e_i^2 = ((0), e_i)$, with e_i as defined in Section 2.2.1. When we add a box to $\bar{\lambda}$ on the i^{th} row of partition k we form a new bi-tuple $\bar{\lambda} + e_i^k$. We may extend Young's lattice (Figure 2.1) to bi-partitions by restricting our attention to choices of e_i^k that result in a new bi-partition. This structure may be seen in Figure 4.1.

Each path upwards (respectively downwards) in Figure 4.1 represents the placement (respectively removal) of a box to form a Young diagram. This is an important structure as it connects the bi-partitions of n to bi-partitions of $n + 1$, this in turn allows us to link the modules of \mathfrak{B}_n and \mathfrak{B}_{n+1} .




Figure 4.1: Young's lattice for bi-partitions of size $n \in \{0, 1, 2, 3\}$.

4.2.2 Bi-tableaux

Given a bi-partition $\bar{\lambda}$ we may form a *Young bi-tableau* (alternatively a $\bar{\lambda}$ -*tableau*), denoted $\bar{T} = (T^1, T^2)$, by putting the numbers $\pm 1, \dots, \pm n$ into the boxes of (the Young diagram of) $\bar{\lambda}$ such that the value a (positive or negative) only appears once. We say the value a *occurs* in a bi-tableau \bar{T} if $\pm a$ appears in a box of \bar{T} . In this chapter we refer to bi-tableaux as just tableaux with the understanding that for the hyperoctahedral group we are working with bi-partitions, and bi-tableaux. The set of Young tableaux of shape $\bar{\lambda}$ is denoted $\text{YT}(\bar{\lambda})$. A *standard Young tableau* \bar{T} is a Young tableau where the values in the boxes of \bar{T} are all positive, and increasing across rows and down columns within T^1 and T^2 . The set of standard Young tableaux of shape $\bar{\lambda}$ is denoted $\text{SYT}(\bar{\lambda})$. The size of the set $\text{SYT}(\bar{\lambda})$ is denoted $d_{\bar{\lambda}}$ and called the dimension of $\bar{\lambda}$.

Lemma 4.2.3. *The dimension $d_{\bar{\lambda}}$ is equal to the product $\binom{n}{|\lambda^1|} d_{\lambda^1} d_{\lambda^2}$ where d_{λ^k} is the dimension of λ^k as a partition of $|\lambda^k|$.*

Proof. To form a standard Young tableau of shape $\bar{\lambda}$ first split the values of $[n]$ across the two Young diagrams, there are $\binom{n}{|\lambda^1|}$ ways to do this. Then form standard Young tableaux of shape λ^1, λ^2 with their assigned values, there are $d_{\lambda^1} d_{\lambda^2}$ way to do this. \square

The transpose of a tableau \bar{T} , denoted \bar{T}' , is defined by swapping the tableaux T^1, T^2 and taking each transpose separately, i.e., $\bar{T}' = (T^{2'}, T^{1'})$. A tableau \bar{T} is a standard Young tableau if and only if its transpose \bar{T}' is a standard Young tableau. If a tableau \bar{T} has a box in position (i, j, k) , we let $\bar{T}(i, j, k)$ denote the value in that box; otherwise $\bar{T}(i, j, k)$ is undefined.

Example 4.2.4. The bi-partition $\bar{\lambda} = ((3, 1), (2, 1))$ has $\binom{7}{4} d_{(3,1)} d_{(2,1)} = 35 * 6 = 210$ standard Young tableaux, 6 of these are given below:

$$\begin{array}{ccc} \left(\begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 5 & 6 \\ \hline 7 & \\ \hline \end{array} \right) & \left(\begin{array}{|c|c|c|} \hline 1 & 2 & 4 \\ \hline 3 & & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 5 & 6 \\ \hline 7 & \\ \hline \end{array} \right) & \left(\begin{array}{|c|c|c|} \hline 1 & 3 & 4 \\ \hline 2 & & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 5 & 6 \\ \hline 7 & \\ \hline \end{array} \right) \\ \left(\begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 5 & 7 \\ \hline 6 & \\ \hline \end{array} \right) & \left(\begin{array}{|c|c|c|} \hline 1 & 2 & 4 \\ \hline 3 & & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 5 & 7 \\ \hline 6 & \\ \hline \end{array} \right) & \left(\begin{array}{|c|c|c|} \hline 1 & 3 & 4 \\ \hline 2 & & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 5 & 7 \\ \hline 6 & \\ \hline \end{array} \right). \end{array}$$

The remaining standard Young tableaux may be found by choosing different divisions of the numbers $\{1, 2, 3, 4, 5, 6, 7\}$ between the partitions $(3, 1)$ and $(2, 1)$.

Each standard Young tableau of shape $\bar{\lambda}$ corresponds to one path up Young's lattice (Figure 4.1) starting at \emptyset and ending at $\bar{\lambda}$. To form this correspondence take any $\bar{T} \in \text{SYT}(\bar{\lambda})$ and create a path up Young's lattice by adding boxes to \emptyset in the order given by the entries in \bar{T} .

The following result is an extension of Lemma 2.2.2 which establishes a link between Young tableaux and their respective shapes. In particular by comparing two young tableaux of possibly different shapes we may learn whether one bi-partition dominates the other.

Definition 4.2.5. Let $T_{\bar{\lambda}}$ be a $\bar{\lambda}$ -tableau and $T_{\bar{\mu}}$ be a $\bar{\mu}$ -tableau. We say that $T_{\bar{\lambda}}$ *agrees* with $T_{\bar{\mu}}$, if for all a that occur in $T_{\bar{\lambda}}^k$ we have $\pm a$ occurring in $T_{\bar{\mu}}^k$ for $k \in \{1, 2\}$. If the tableaux $T_{\bar{\lambda}}, T_{\bar{\mu}}$ agree then we must have $|\lambda^k| = |\mu^k|$ for $k \in \{1, 2\}$.

Lemma 4.2.6. Let $T_{\bar{\lambda}}$ be a $\bar{\lambda}$ -tableau and $T_{\bar{\mu}}$ be a $\bar{\mu}$ -tableau that agree. Suppose for each index i the elements in row i of $T_{\bar{\mu}}^k$ appear in different columns of $T_{\bar{\lambda}}^k$. Then $\bar{\lambda} \succeq \bar{\mu}$.

Proof. We know $|\lambda^k| = |\mu^k|$ for $k \in \{1, 2\}$ and thus using the same argument as Lemma 2.2.2 on the first and second tableaux separately we have $\lambda^1 \succeq \mu^1$, $\lambda^2 \succeq \mu^2$. \square

4.2.3 The Module Structure of the Hyperoctahedral Group

We are now in a position to define the permutation and Specht modules for the hyperoctahedral group. We find one permutation module for every bi-partition $\bar{\lambda}$ which contains the simple Specht modules as special submodules. The following proofs are adapted from several sources, including Geissinger and Kinch [22], Aamily, Morris, and Peel [1], and Can [12]. We have adapted the constructions presented in the literature to allow us to re-establish Young's rule (Lemma 2.2.21) for the permutation modules of the hyperoctahedral group.

Permutation Modules

To define the permutation modules of B_n we reintroduce the idea of a row permutation for a tableau \bar{T} , this defines an equivalence relation on the Young tableaux of a fixed bi-partition.

Definition 4.2.7. Let \bar{T} be a Young tableau of size n , and $\sigma \in B_n$. Define an action of σ on \bar{T} by applying σ to the values of \bar{T} box wise.

Definition 4.2.8. Let \bar{T} be a bi-tableau of size n . Define the subgroup of *row permutations* of \bar{T} , denoted $R_{\bar{T}}$, as the set of elements in B_n which fix rows of T^1 up to changes in sign, and fix rows in T^2 completely. Define the subgroup of *column permutations* of \bar{T} , denoted $C_{\bar{T}}$, as the set of elements in B_n which fix columns in T^2 up to changes in sign, and fix columns in T^1 completely.

Definition 4.2.9. Define an equivalence relation on tableaux of shape $\bar{\lambda}$ called *row equivalence*, denoted \sim_R , in the following way:

$$\bar{T}_1 \sim_R \bar{T}_2 \Leftrightarrow \exists \sigma \in R_{\bar{T}_1} \text{ such that } \sigma \bar{T}_1 = \bar{T}_2.$$

The row equivalence class of a tableau \bar{T} is denoted by $\{\bar{T}\}$ and called a *(bi-)tabloid*.

Example 4.2.10. Let $\bar{\lambda} = ((2, 1), (2))$. Consider the following three $\bar{\lambda}$ -tableaux,

$$\bar{T}_1 = \left(\begin{array}{|c|c|} \hline 1 & 5 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 4 & 2 \\ \hline \end{array} \right) \quad \bar{T}_2 = \left(\begin{array}{|c|c|} \hline 5 & -1 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline -4 & 2 \\ \hline \end{array} \right) \quad \bar{T}_3 = \left(\begin{array}{|c|c|} \hline -1 & -5 \\ \hline -3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 2 & 4 \\ \hline \end{array} \right).$$

The tableaux \bar{T}_1, \bar{T}_2 are not row equivalent because we have changed the sign of 4 in the second tableau. The tableaux \bar{T}_1, \bar{T}_3 are row equivalent for the permutation $(\xi_1 \xi_3 \xi_5(24)) \in R_{\bar{T}_1}$, therefore, $\{\bar{T}_1\} = \{\bar{T}_3\}$.

The action of B_n on tableaux, naturally extends to an action on tabloids, given by $\sigma \{\bar{T}\} = \{\sigma \bar{T}\}$. We now construct the permutation modules for \mathfrak{B}_n .

Definition 4.2.11. Let $\bar{\lambda} \vdash n$. The *permutation module* for \mathfrak{B}_n , denoted $M^{\bar{\lambda}}$, is the following vector space

$$M^{\bar{\lambda}} = \langle \{\bar{T}\} \mid T \text{ is a tableau of shape } \bar{\lambda} \rangle$$

with action of \mathfrak{B}_n given by extending the action of B_n on tabloids (Definition 4.2.7) linearly.

Note that just simply extending the previous definition of row permutations (Definition 2.2.4) for single tableaux T to both tableaux T^1, T^2 is not strong enough to distinguish all the permutation modules of B_n from one another. For example, if we study row permutations of bi-tableaux which fix rows completely in both separate tableaux, then for any partition $\lambda \vdash n$ the modules $M^{(\lambda, (0))}$ and $M^{((0), \lambda)}$ would be isomorphic to one another.

Lemma 4.2.12. *The permutation module $M^{\bar{\lambda}}$ is cyclic and has dimension*

$$2^{|\lambda^2|} \frac{n!}{\prod_{i=1}^2 \left(\prod_{j=1}^{l(\lambda^i)} \lambda_j^i! \right)}.$$

Proof. The main fraction above is a multinomial coefficient formed from n and the size of every row in $\bar{\lambda}$. The extra factor $2^{|\lambda^2|}$ represents the fact that all entries of λ^2 have fixed signs. The permutation module being cyclic follows from the row equivalence relation. \square

Specht Modules

We construct the Specht modules $S^{\bar{\lambda}}$ for the hyperoctahedral group as stable subspaces of the corresponding permutation modules $M^{\bar{\lambda}}$. We form a basis for the Specht module $S^{\bar{\lambda}}$ by taking linear combinations of tabloids, we call these new elements *polytabloids*.

Definition 4.2.13. Given a Young tableau \bar{T} , form the element $k_{\bar{T}} \in \mathfrak{B}_n$ as the following sum:

$$k_{\bar{T}} = \sum_{\sigma \in C_{\bar{T}}} \text{sgn}(\sigma) \sigma. \quad (4.1)$$

The *polytabloid* associated to tabloid \bar{T} , denoted $e_{\bar{T}}$, is the element $e_{\bar{T}} = k_{\bar{T}} \{\bar{T}\}$. We say that the tabloid $\{\bar{T}'\}$ is contained in the polytabloid $e_{\bar{T}}$ if it appears with a non-zero coefficient.

Lemma 4.2.14. *Let \bar{T} be a Young tableau of size n , and $\eta \in B_n$. Then $\eta e_{\bar{T}} = e_{\eta \bar{T}}$.*

Proof.

$$\eta e_{\bar{T}} = \sum_{\sigma \in C_{\bar{T}}} \text{sgn}(\sigma) \eta \sigma \eta^{-1} \{\eta \bar{T}\} = \sum_{\sigma \in C_{\eta \bar{T}}} \text{sgn}(\sigma) \sigma \{\eta \bar{T}\} = e_{\eta \bar{T}}. \quad \square$$

Definition 4.2.15. Let $\bar{\lambda} \vdash n$. The *Specht module* for \mathfrak{B}_n , denoted $S^{\bar{\lambda}}$, is defined as the following vector space,

$$S^{\bar{\lambda}} = \langle e_{\bar{T}} \mid \bar{T} \in \text{YT}(\bar{\lambda}) \rangle$$

with the natural action of \mathfrak{B}_n on tabloids. In fact we may restrict the spanning set $\{e_{\bar{T}} | \bar{T} \in \text{YT}(\bar{\lambda})\}$ to a basis of polytabloids formed from standard Young tableaux of $\bar{\lambda}$ (see [12, Sections 4 and 5]), that is

$$S^{\bar{\lambda}} = \langle e_{\bar{T}} | \bar{T} \in \text{SYT}(\bar{\lambda}) \rangle. \quad (4.2)$$

Thus, the Specht module $S^{\bar{\lambda}}$ has dimension $d_{\bar{\lambda}}$, and is cyclic by applying Lemma 4.2.14.

Example 4.2.16. Let $\bar{\lambda} = ((2, 1), (2))$, the Specht module $S^{\bar{\lambda}}$ has dimension $d_{\bar{\lambda}} = \binom{5}{3} d_{(2,1)} d_2 = 20$. Take \bar{T} to be the following standard Young tableau of shape $\bar{\lambda}$,

$$\bar{T} = \left(\begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 4 & 5 \\ \hline \end{array} \right).$$

The group of column permutations for any tableau of shape $\bar{\lambda}$ has size 8, as we may swap elements in the first column of \bar{T}^1 , and swap the sign of every value in \bar{T}^2 . Starting from the tabloid \bar{T} we can compute its associated polytabloid:

$$\begin{aligned} e_{\bar{T}} = k_{\bar{T}} \{\bar{T}\} &= \left\{ \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 4 & 5 \\ \hline \end{array} \right\} - \left\{ \begin{array}{|c|c|} \hline 3 & 2 \\ \hline 1 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 4 & 5 \\ \hline \end{array} \right\} - \left\{ \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline -4 & 5 \\ \hline \end{array} \right\} + \left\{ \begin{array}{|c|c|} \hline 3 & 2 \\ \hline 1 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline -4 & 5 \\ \hline \end{array} \right\} \\ &- \left\{ \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 4 & -5 \\ \hline \end{array} \right\} + \left\{ \begin{array}{|c|c|} \hline 3 & 2 \\ \hline 1 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 4 & -5 \\ \hline \end{array} \right\} + \left\{ \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline -4 & -5 \\ \hline \end{array} \right\} - \left\{ \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline -4 & -5 \\ \hline \end{array} \right\}. \end{aligned}$$

We have seen that the conjugacy classes of B_n are labelled by bi-partitions, and we have found exactly one Specht module for each bi-partition $\bar{\lambda}$. We now prove that the Specht modules form a complete set of simple modules for the hyperoctahedral group. To do this we study the properties of the groups action on polytabloids. We first show every Specht module to be simple and then prove that $S^{\bar{\lambda}} \cong S^{\bar{\mu}}$ if and only if $\bar{\lambda} = \bar{\mu}$. The following arguments take inspiration from the construction of Specht modules for the symmetric group (Section 2.2.4), and work by Can [12].

Lemma 4.2.17. *Let $\bar{\lambda}, \bar{\mu} \vdash n$. Let $T_{\bar{\lambda}}$ be a $\bar{\lambda}$ -tableau and $T_{\bar{\mu}}$ be a $\bar{\mu}$ -tableau. Suppose that at least one of the following conditions holds: 1) there exists $a \in T_{\bar{\mu}}^1$ such that $\pm a \in T_{\bar{\lambda}}^2$, 2) there exists a, b in the same row of $T_{\bar{\mu}}^k$, with $c = \pm a, d = \pm b$ in the same column of $T_{\bar{\lambda}}^k$. Then $k_{T_{\bar{\lambda}}} \{T_{\bar{\mu}}\} = 0$.*

Proof. First we establish that $C_{T_{\bar{\lambda}}} \cap R_{T_{\bar{\mu}}} \neq \{e\}$. If the first condition is met then clearly $\xi_a \in C_{T_{\bar{\lambda}}} \cap R_{T_{\bar{\mu}}}$. If the second condition is met with $k = 1$ then $(ab) \in C_{T_{\bar{\lambda}}} \cap R_{T_{\bar{\mu}}}$, if $k = 2$ then $(cd) \in C_{T_{\bar{\lambda}}} \cap R_{T_{\bar{\mu}}}$. Let τ be a transposition in $C_{T_{\bar{\lambda}}} \cap R_{T_{\bar{\mu}}}$, then $(e - \tau) \{T_{\bar{\mu}}\} = 0$. All choices of τ generate a subgroup of order 2 in $C_{T_{\bar{\lambda}}}$, taking signed coset representatives $\sigma_1, \dots, \sigma_m$ for this subgroup. Then,

$$k_{T_{\bar{\lambda}}} \{T_{\bar{\mu}}\} = \left(\sum_{\sigma \in C_{T_{\bar{\lambda}}}} \text{sgn}(\sigma) \sigma \right) \{T_{\bar{\mu}}\} = \left(\sum_{i=1}^m \sigma_i (e - \tau) \right) \{T_{\bar{\mu}}\} = 0. \quad \square$$

Lemma 4.2.18. *Let $\bar{\lambda}, \bar{\mu} \vdash n$. Let $T_{\bar{\lambda}}$ be a $\bar{\lambda}$ -tableau and $T_{\bar{\mu}}$ be a $\bar{\mu}$ -tableau. Suppose $\bar{\lambda} \not\subseteq \bar{\mu}$, then $k_{T_{\bar{\lambda}}} \{T_{\bar{\mu}}\} = 0$.*

Proof. Suppose $T_{\bar{\lambda}}$ does not agree with $T_{\bar{\mu}}$ then by our hypothesis there must exist $a \in T_{\bar{\mu}}^1$ such that $\pm a \in T_{\bar{\lambda}}^2$ thus the conclusion holds by Lemma 4.2.17. However, if $T_{\bar{\lambda}}$ agrees with $T_{\bar{\mu}}$ the conclusion holds by Lemma 4.2.17 and Lemma 4.2.6. \square

Corollary 4.2.19. *Let $\bar{\lambda}, \bar{\mu} \vdash n$. Let $T_{\bar{\lambda}}$ be a $\bar{\lambda}$ -tableau and $T_{\bar{\mu}}$ be a $\bar{\mu}$ -tableau. Suppose $k_{T_{\bar{\lambda}}}\{T_{\bar{\mu}}\} \neq 0$, then $\bar{\lambda} \triangleright \bar{\mu}$. Furthermore, if $\bar{\lambda} = \bar{\mu}$ then $k_{T_{\bar{\lambda}}}\{T_{\bar{\mu}}\} = \pm e_{T_{\bar{\lambda}}}$.*

Proof. The first statement is the contrapositive of Lemma 4.2.18. Let $\bar{\lambda} = \bar{\mu}$ and $k_{T_{\bar{\lambda}}}\{T_{\bar{\mu}}\} \neq 0$. The tableaux $T_{\bar{\lambda}}, T_{\bar{\mu}}$ must agree and furthermore, if a, b belong to the same row of $T_{\bar{\lambda}}^k$ then $c = \pm a, d = \pm b$ belong to different columns of $T_{\bar{\mu}}$. This implies that there exists a column permutation $\eta \in C_{T_{\bar{\lambda}}}$ such that $\{T_{\bar{\mu}}\} = \eta\{T_{\bar{\lambda}}\}$ for more details see [12, Lemma 3.6]. Therefore,

$$k_{T_{\bar{\lambda}}}\{T_{\bar{\mu}}\} = \left(\sum_{\sigma \in C_{T_{\bar{\lambda}}}} \text{sgn}(\sigma) \sigma \eta \right) \{T_{\bar{\lambda}}\} = \left(\sum_{\sigma \in C_{T_{\bar{\lambda}}}} \text{sgn}(\eta^{-1}) \text{sgn}(\sigma) \sigma \right) \{T_{\bar{\lambda}}\} = \text{sgn}(\eta) e_{T_{\bar{\lambda}}}. \quad \square$$

Corollary 4.2.20. *Let $T_{\bar{\lambda}}$ be a $\bar{\lambda}$ -tableau and $v \in M^{\bar{\lambda}}$. Then $k_{T_{\bar{\lambda}}}v$ is a multiple (possibly zero) of $e_{T_{\bar{\lambda}}}$.*

Proof. Write $v = \sum_i c_i \{\bar{T}_i\}$, where \bar{T}_i is a $T_{\bar{\lambda}}$ -tableau. Each summand $k_{T_{\bar{\lambda}}}\{\bar{T}_i\}$ is either zero or a multiple of $e_{T_{\bar{\lambda}}}$ by Corollary 4.2.19. \square

We are now in a position to prove that every Specht module is simple, and that they are pairwise non-isomorphic to each other. The proofs of the next theorems follow closely from their symmetric group counterparts (Theorems 2.2.17 and 2.2.18).

Theorem 4.2.21 (Submodule Theorem). *Let $V \subseteq M^{\bar{\lambda}}$ be a submodule. Then $V \supseteq S^{\bar{\lambda}}$ or $V \subseteq (S^{\bar{\lambda}})^0$. Therefore, the Specht module $S^{\bar{\lambda}}$ is simple.*

Proof. Take $v \in V$, and $T_{\bar{\lambda}}$ a $\bar{\lambda}$ -tableau. By Corollary 4.2.20 we have $k_{T_{\bar{\lambda}}}v = ce_{T_{\bar{\lambda}}}$ for some $c \in \mathbb{C}$. Suppose we have $c \neq 0$ for some choice of $T_{\bar{\lambda}}$, then $c^{-1}k_{T_{\bar{\lambda}}}v = e_{T_{\bar{\lambda}}} \in V$ hence we may generate $S^{\bar{\lambda}}$ inside V .

Now suppose that $c = 0$ for all choices of tableau $T_{\bar{\lambda}}$, then consider the inner product on $M^{\bar{\lambda}}$ defined on its basis by $\langle \{\bar{T}_1\}, \{\bar{T}_2\} \rangle = \delta_{\{\bar{T}_1\}, \{\bar{T}_2\}}$. This inner product is invariant under the action of B_n , so we find

$$\langle v, e_{T_{\bar{\lambda}}} \rangle = \sum_{\sigma \in C_{T_{\bar{\lambda}}}} \langle v, \text{sgn}(\sigma) \sigma \{T_{\bar{\lambda}}\} \rangle = \sum_{\sigma \in C_{T_{\bar{\lambda}}}} \langle \text{sgn}(\sigma) \sigma^{-1} v, \{T_{\bar{\lambda}}\} \rangle = \langle k_{T_{\bar{\lambda}}}v, \{T_{\bar{\lambda}}\} \rangle = \langle 0, \{T_{\bar{\lambda}}\} \rangle = 0.$$

Therefore no polytabloid appears in $M^{\bar{\lambda}}$, and because any single polytabloid would span $S^{\bar{\lambda}}$, we have $v \notin S^{\bar{\lambda}} \Rightarrow v \in (S^{\bar{\lambda}})^0$. \square

Theorem 4.2.22. *Let $\bar{\lambda}, \bar{\mu} \vdash n$. Suppose there is a non-zero homomorphism $\psi : S^{\bar{\lambda}} \rightarrow M^{\bar{\mu}}$. Then $\bar{\lambda} \triangleright \bar{\mu}$, and if $\bar{\lambda} = \bar{\mu}$ then ψ is multiplication by a scalar.*

Proof. Take a basis vector $e_{T_{\bar{\lambda}}} \in S^{\bar{\lambda}}$ such that $\psi(e_{T_{\bar{\lambda}}}) \neq 0$. Extend the homomorphism ψ to the whole module $M^{\bar{\lambda}}$ by setting it to be 0 on the complement of $S^{\bar{\lambda}}$. Then

$$0 \neq \psi(e_{T_{\bar{\lambda}}}) = k_{T_{\bar{\lambda}}}\psi(\{T_{\bar{\lambda}}\}) = k_{T_{\bar{\lambda}}}\left(\sum_i c_i \{\bar{T}_i\}\right)$$

where the \bar{T}_i are $\bar{\mu}$ -tableau. We must have at least one c_i and product $k_{T_{\bar{\lambda}}} c_i \{\bar{T}_i\}$ being non-zero. Therefore, by Corollary 4.2.19 we have $\bar{\lambda} \succeq \bar{\mu}$.

If $\bar{\lambda} = \bar{\mu}$ we have that $\psi(e_{T_{\bar{\lambda}}}) = c \cdot e_{T_{\bar{\lambda}}}$ for some constant $c \in \mathbb{C}$. Hence, for any $\sigma \in B_n$, using Lemma 4.2.14 we have

$$\psi(e_{\sigma T_{\bar{\lambda}}}) = \psi(\sigma e_{T_{\bar{\lambda}}}) = \sigma \psi(e_{T_{\bar{\lambda}}}) = c \cdot \sigma e_{T_{\bar{\lambda}}} = c \cdot e_{\sigma T_{\bar{\lambda}}}. \quad \square$$

Corollary 4.2.23. *The Specht modules $S^{\bar{\lambda}}$ for $\bar{\lambda} \vdash n$ form a complete set of pairwise non-isomorphic simple modules for \mathfrak{B}_n .*

Proof. The Submodule Theorem tells us the Specht modules are simple. Now if $\bar{\lambda} = \bar{\mu}$ we may easily see that $S^{\bar{\lambda}} \cong S^{\bar{\mu}}$. Conversely if $S^{\bar{\lambda}} \cong S^{\bar{\mu}}$ then we have non-zero homomorphisms $\psi : S^{\bar{\lambda}} \rightarrow M^{\bar{\mu}}$ and $\phi : S^{\bar{\mu}} \rightarrow M^{\bar{\lambda}}$. Therefore, by Theorem 4.2.22 we must have $\bar{\lambda} \succeq \bar{\mu}$ and $\bar{\mu} \succeq \bar{\lambda}$ which implies $\bar{\lambda} = \bar{\mu}$. \square

Given a permutation module $M^{\bar{\lambda}}$ we would like to decompose it into its simple submodules. An application of Theorem 4.2.22 allows us to re-establish Young's rule for the the permutation modules of the hyperoctahedral group.

Lemma 4.2.24 (Young's Rule). *For $\bar{\mu} \vdash n$ we have,*

$$M^{\bar{\mu}} \cong \bigoplus_{\substack{\bar{\lambda} \vdash n \\ \bar{\lambda} \succeq \bar{\mu}}} K_{\bar{\lambda}, \bar{\mu}} S^{\bar{\lambda}}, \quad (4.3)$$

where $K_{\bar{\lambda}, \bar{\mu}} S^{\bar{\lambda}}$ denotes a direct sum of $K_{\bar{\lambda}, \bar{\mu}}$ copies of $S^{\bar{\lambda}}$. We call the coefficients $K_{\bar{\lambda}, \bar{\mu}} \in \mathbb{N}^0$ generalised Kostka numbers. We know that $K_{\bar{\lambda}, \bar{\lambda}} = 1$ for all $\bar{\lambda} \vdash n$. Furthermore, if $|\lambda^1| = |\mu^1|$ then $K_{\bar{\lambda}, \bar{\mu}} = K_{\lambda^1, \mu^1} K_{\lambda^2, \mu^2}$.

Proof. Theorem 4.2.22 tells us that if $S^{\bar{\lambda}}$ appears as a summand of $M^{\bar{\mu}}$ we must have $\bar{\lambda} \succeq \bar{\mu}$. If $\bar{\lambda} = \bar{\mu}$ we know that any homomorphism $S^{\bar{\lambda}} \rightarrow M^{\bar{\lambda}}$ is multiplication by a scalar hence there must only be one copy of $S^{\bar{\lambda}}$ in $M^{\bar{\lambda}}$. The last assertion follows from work by Geissinger and Kinch [22, Corollary II.3]. \square

Similarly to the symmetric group case the regular module \mathfrak{B}_n is isomorphic to the permutation module associated to the bi-partition at the bottom of the dominance ordering, that is the partition $((0), (1^n))$.

Lemma 4.2.25. *The permutation module $M^{((0), (1^n))} \cong \mathfrak{B}_n$ as \mathfrak{B}_n -modules. Thus, it has canonical decomposition,*

$$M^{((0), (1^n))} \cong \bigoplus_{\bar{\lambda} \vdash n} d_{\bar{\lambda}} S^{\bar{\lambda}} \text{ as } \mathfrak{B}_n\text{-modules.}$$

Proof. Define a isomorphism of \mathfrak{B}_n -modules $\psi : \mathfrak{B}_n \rightarrow M^{((0), (1^n))}$ by its action on a single permutation $\sigma \in B_n$:

$$\psi(\sigma) = \left\{ \emptyset, \boxed{\sigma(1)} \mid \dots \mid \boxed{\sigma(n)} \right\}. \quad \square$$

The version of Young's rule we stated for the hyperoctahedral group is weaker than the result for the symmetric group (Lemma 2.2.21). For the purpose of lifting eigenvectors we do not require that every Specht module $S^{\bar{\lambda}}$ with $\bar{\lambda} \succeq \bar{\mu}$ appears in the decomposition of $M^{\bar{\mu}}$.

The hyperoctahedral group has a natural inclusion structure $B_m \hookrightarrow B_n$ for $m < n$, by extending every permutation of B_m to a permutation of B_n by choosing it to fix all elements of $[\pm m] \setminus [\pm n]$. The group algebras of the hyperoctahedral group and subsequently \mathfrak{B}_n -modules inherit this recursive structure. To end this section we state the branching rules for the Specht modules of the hyperoctahedral group, a proof and further details of this result may be found in [22, Section III].

Theorem 4.2.26 (Branching rules for B_n). *Let $n \geq 2$, and $\bar{\lambda} \vdash n$. The branching rules for the simple modules of the hyperoctahedral group are as follows:*

$$\text{Res}_{B_{n-1}}^{B_n} S^{\bar{\lambda}} \cong \bigoplus_{\substack{\bar{\mu} \vdash n-1 \\ \bar{\mu} \subseteq \bar{\lambda}}} S^{\bar{\mu}} \text{ as } \mathfrak{B}_{n-1}\text{-modules} \quad (4.4)$$

$$\text{Ind}_{B_n}^{B_{n+1}} S^{\bar{\lambda}} \cong \bigoplus_{\substack{\bar{\mu} \vdash n+1 \\ \bar{\lambda} \subseteq \bar{\mu}}} S^{\bar{\mu}} \text{ as } \mathfrak{B}_{n+1}\text{-modules} \quad (4.5)$$

The branching rules for Specht modules are closely related to the edges of Young's lattice. When we induce a Specht module $S^{\bar{\lambda}}$ one step we create a direct sum of the Specht modules associated to bi-partitions directly above $\bar{\lambda}$ in Young's lattice, conversely when we restrict one step we find a direct sum of Specht modules associated to bi-partitions directly below $\bar{\lambda}$ in Young's lattice.

4.2.4 Switching to Words

Let $\bar{\lambda} \vdash n$, and take a tableau \bar{T} of shape $\bar{\lambda}$. If we add a box to \bar{T} containing the value $\pm(n+1)$ we transform \bar{T} to a tableau of size $n+1$. Similarly we can transform a tabloid $\{\bar{T}\}$ if size n into a tabloid of size $n+1$ by adding a single box. This is exactly how we lift vectors from permutation module $M^{\bar{\lambda}}$ to permutation modules of \mathfrak{B}_{n+1} . Notationally this process is cumbersome, therefore we instead form a correspondence between the study of tabloids to the study of words over a finite alphabet. This allows the algebra of lifting eigenvectors in Section 4.2.4 to be presented in a concise way. The construction of the correspondence follows closely to that of the symmetric group, however the alphabet and words now contain signed and unsigned letters.

Define the set $[n^\pm] = \{1^+, \dots, n^+\} \cup \{1^-, \dots, n^-\}$, viewed as $2n$ distinct symbols. We call the letters in $[n^\pm]$ *signed*, and letters in $[n]$ *unsigned*. Given $n \in \mathbb{N}^0$ we denote by $W^{\bar{n}}$ the set of words of length n with letters from the set $[\bar{n}] = [n] \cup [n^\pm]$. For $n = 0$ we allow the empty word ω as the unique word of length zero. The size of $W^{\bar{n}}$ is $(3n)^n$.

The hyperoctahedral group has a natural action on $W^{\bar{n}}$. Take $\sigma \in B_n$ with decomposition $\sigma = \left(\prod_{i=1}^n \xi_i^{x_i}\right) \eta$, then σ acts on the word $w \in W^{\bar{n}}$ in the following way: permute the letters in w by η (as defined in Section 2.2.6), then for all $x_i = 1$ flip the sign of w_i if possible. For example, take $w = 1^+24^-3^+3$ and $\sigma = (\xi_1\xi_4\xi_5)(123)$, then $\sigma w = 4^+1^+23^-3$, note that ξ_5 has no affect on $w_5 = 3$ because it is an unsigned letter. Let $M^{\bar{n}}$ be the vector

space over \mathbb{C} spanned by words in $W^{\bar{n}}$. The action of B_n on $W^{\bar{n}}$ extends linearly to an action of \mathfrak{B}_n on $M^{\bar{n}}$, thus the vector space $M^{\bar{n}}$ is a \mathfrak{B}_n -module.

To each word $w \in W^{\bar{n}}$ we may associate a bi-tuple of size n of non-negative integers called its *evaluation*, denoted \overline{eval} , as follows. Define $eval_i(w)$ to count the number of occurrences of unsigned symbol i in w , and let $eval(w) := (eval_1, \dots, eval_n)$. Define $eval_i^\pm(w)$ to count the number of occurrences of signed symbol i^\pm in w and let $eval^\pm(w) := (eval_1^\pm, \dots, eval_n^\pm)$. Then for a word $w \in W^{\bar{n}}$ define its evaluation as the tuple $\overline{eval}(w) := (eval(w), eval^\pm(w))$. Note that the evaluation of any word in $W^{\bar{n}}$ sums to n . For example, the word $1^+24^-3^+3$ has evaluation $((0, 1, 1, 0, 0), (1, 0, 1, 1, 0))$. If the bi-tuple $\overline{eval}(w)$ is non-decreasing we associate it with the corresponding bi-partition, e.g. 1^+112^-2 has evaluation $((2, 1, 0^2), (1, 1, 0^3))$, and we associate this bi-tuple with the bi-partition $((2, 1), (1, 1))$. The evaluation of any word is stable under the action of B_n because we only permute symbols and change the sign of signed symbols, both of which do not affect $eval^\pm(w)$. Therefore, we may form a submodule of $M^{\bar{n}}$ by restricting to words with a given evaluation.

Definition 4.2.27. Let $\bar{\nu}$ be a bi-tuple of non-negative integers which sum to n . Define the \mathfrak{B}_n -module $M^{\bar{\nu}}$ as the following stable vector space

$$M^{\bar{\nu}} = \{w \in W^{\bar{n}} \mid \overline{eval}(w) = \bar{\nu}\} \subseteq M^{\bar{n}}.$$

Let $\bar{\nu} \vdash n$, then the module $M^{\bar{\nu}}$ is isomorphic to the permutation module defined in Definition 4.2.11. To establish the equivalence of these two definitions we construct a bijection from tabloids of shape $\bar{\lambda}$ to words with evaluation $\bar{\lambda}$.

Definition 4.2.28. Let $\bar{\lambda} \vdash n$. Define a map $w : \text{YT}(\bar{\lambda}) \rightarrow W^{\bar{n}}$ as follows: for each tableau \bar{T} of shape $\bar{\lambda}$, let $w(\bar{T}) = w_1 \dots w_n$ be the word with $w_{\bar{T}(i,j,1)} = i$ for each box $(i, j, 1) \in \bar{T}$, and $w_{\bar{T}(i,j,2)} = i^\pm$ for each box $(i, j, 2) \in \bar{T}$ with the sign of i given by the sign present in box $(i, j, 2)$ of \bar{T} . The word $w(\bar{T})$ formed by this process has evaluation $\bar{\lambda}$. The map w respects the action of B_n on tableaux and words, thus forming a homomorphism between the \mathfrak{B}_n -modules of $\bar{\lambda}$ -tableaux and words in $M^{\bar{\lambda}}$.

Example 4.2.29. Let $\bar{\lambda} = ((2, 1), (1, 1))$, and take \bar{T} a $\bar{\lambda}$ -tableau,

$$\bar{T} = \left(\begin{array}{|c|c|} \hline 2 & 1 \\ \hline -4 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline -3 & 5 \\ \hline \end{array} \right) \text{ which has corresponding word } w(\bar{T}) = 111^-21^+.$$

Now take $\sigma \in B_n$, the action of σ commutes with the linear map w . For example take $\sigma = (\xi_1 \xi_2)(13)(245)$, then

$$\sigma \bar{T} = \left(\begin{array}{|c|c|} \hline 4 & 3 \\ \hline -5 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 1 & -2 \\ \hline \end{array} \right) \text{ which has corresponding word } w(\sigma \bar{T}) = 1^+1^-112 = \sigma w(\bar{T}).$$

The homomorphism $w : \text{YT}(\bar{\lambda}) \rightarrow M^{\bar{\lambda}}$ is surjective but not injective. To form an isomorphism between the two definitions of the permutation modules $M^{\bar{\lambda}}$ we restrict the domain of w from tableaux of shape $\bar{\lambda}$, to tabloids of shape $\bar{\lambda}$ which are in one-to-one correspondence with words of evaluation $\bar{\lambda}$.

Lemma 4.2.30. *Let $\bar{\lambda} \vdash n$ and \bar{T}_1, \bar{T}_2 be two $\bar{\lambda}$ -tableaux. Then*

$$\{\bar{T}_1\} = \{\bar{T}_2\} \Leftrightarrow \bar{T}_1 \sim_R \bar{T}_2 \Leftrightarrow w(\bar{T}_1) = w(\bar{T}_2).$$

Proof. This follows from the definition of w . Any row permutation fixes the occurrences of letters in our word, but may change the sign of those in the first tableau. However, the letters represented by the first tableau are unsigned, therefore changing the sign does not create a new word. \square

Lemma 4.2.31. *Let $\bar{\lambda} \vdash n$, the permutation module $M^{\bar{\lambda}}$ may be seen as a vector space over the following bases:*

$$\begin{aligned} M^{\bar{\lambda}} &= \langle \{\bar{T}\} \mid \bar{T} \text{ is a tableau of shape } \bar{\lambda} \rangle \\ &\cong \langle w(\bar{T}) \mid \bar{T} \text{ is a tableau of shape } \bar{\lambda} \rangle \\ &\cong \langle w \in M^{\bar{n}} \mid \text{eval}(w) = \bar{\lambda} \rangle \end{aligned}$$

Proof. The map $w : M^{\bar{\lambda}} \rightarrow M^{\bar{\lambda}}$ forms a bijection between tabloids of shape $\bar{\lambda}$ and words of evaluation $\bar{\lambda}$, which respects the action of \mathfrak{B}_n . \square

Using the one-to-one correspondence between tabloids and words we may also use words to describe the Specht modules for \mathfrak{B}_n . In example 4.2.16 we saw an explicit description of a polytabloid $e_{\bar{T}}$ belonging to the module $S^{((2,1),(2))}$. We can restate this element replacing every tabloid by its corresponding word:

$$\begin{aligned} e_{\bar{T}} &= \left\{ \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 4 & 5 \\ \hline \end{array} \right\} - \left\{ \begin{array}{|c|c|} \hline 3 & 2 \\ \hline 1 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 4 & 5 \\ \hline \end{array} \right\} - \left\{ \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline -4 & 5 \\ \hline \end{array} \right\} + \left\{ \begin{array}{|c|c|} \hline 3 & 2 \\ \hline 1 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline -4 & 5 \\ \hline \end{array} \right\} \\ &- \left\{ \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 4 & -5 \\ \hline \end{array} \right\} + \left\{ \begin{array}{|c|c|} \hline 3 & 2 \\ \hline 1 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 4 & -5 \\ \hline \end{array} \right\} + \left\{ \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline -4 & -5 \\ \hline \end{array} \right\} - \left\{ \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline -4 & -5 \\ \hline \end{array} \right\} \\ &= 1121^+1^+ - 2111^+1^+ - 1121^-1^+ + 2111^-1^+ \\ &- 1121^+1^- + 2111^+1^- + 1121^-1^- - 2111^-1^-. \end{aligned} \tag{4.6}$$

The structure of words allow us to describe the lifting from modules of \mathfrak{B}_n to \mathfrak{B}_{n+1} by adding letters to words instead of adding boxes to tabloids. The lifting operators for the Specht modules of \mathfrak{S}_n were described by appending letters to the end of a word. We perform a similar analysis for the Specht modules of \mathfrak{B}_n but we now we have a choice of adding a signed or unsigned letter. Overall the correspondence with words helps to simplify the presentation of the lifting operators for \mathfrak{B}_n and the results in Section 4.3.2.

Random Walks on The Hyperoctahedral Group

We may extend any shuffle on the symmetric group to a shuffle on B_n by adding an additional action by negative transpositions to the end of the shuffle. In the following sections we study extended versions of the random transposition shuffle and one-sided transposition shuffle. We first extend the random transposition shuffle to B_n and show how to recover its eigenvalues by lifting eigenvectors. Then we extend the one-sided transposition shuffle to B_n and recover its eigenvalues using the same technique.

4.3 The Random Transposition Shuffle on The Hyperoctahedral Group

The random transposition shuffle for the hyperoctahedral group is described by the following procedure: pick two positions in $[n]$ uniformly at random and switch the cards in these positions, then flip a fair coin; if heads do nothing, if tails flip the moved cards over to their opposite sides. There are other ways we could choose to extend the random transposition shuffle to the hyperoctahedral group – Schoolfield [40] studied an alternative description where each card’s flip are done independently at the end of the shuffle, e.g., we could flip one card and not the other. Schoolfield went on to analyse the mixing time of this shuffle using Fourier transforms and found it to be tightly bounded at time $(n/2) \log n$, i.e., the same time as the random transposition shuffle on S_n . Instead of using Fourier transformations we analyse the random transposition shuffle using the technique of lifting eigenvectors.

Definition 4.3.1. The *random transposition shuffle for B_n* , denoted $\overline{\text{RT}}_n$, is driven by the following probability distribution:

$$\overline{\text{RT}}_n(\sigma) = \begin{cases} 1/2n & \text{if } \sigma = e \\ 1/2n^2 & \text{if } \sigma = \xi_i \text{ for } i \in [n] \\ 1/n^2 & \text{if } \sigma = (ij) \text{ for } i, j \in [n] \text{ with } i < j \\ 1/n^2 & \text{if } \sigma = \xi_i \xi_j (ij) \text{ for } i, j \in [n] \text{ with } i < j \\ 0 & \text{otherwise} \end{cases} . \quad (4.7)$$

This shuffle is defined on the conjugacy classes of positive and negative transpositions. Lifting eigenvectors for the random transposition shuffle on B_n requires modification from the symmetric group case in order to account for the new Specht modules associated to bi-partitions. In particular we require two different lifting operators to take eigenvectors of \mathfrak{B}_n to those of \mathfrak{B}_{n+1} . Throughout this section we will be working towards the following result.

Theorem 4.3.2. *The eigenvalues for the random transposition shuffle $\overline{\text{RT}}_n$ are indexed by bi-partitions $\bar{\lambda} \vdash n$. The eigenvalue corresponding to partition $\bar{\lambda}$ occurs with multiplicity $d_{\bar{\lambda}}^2$, and is given by*

$$\text{eig}(\bar{\lambda}) = \frac{1}{2n^2} (2|\lambda^1| + 4 \text{Diag}(\lambda^1) + 4 \text{Diag}(\lambda^2)). \quad (4.8)$$

The random transposition shuffle for B_n is closely related to that for S_n , and in fact by considering a group homomorphism from B_n to S_n we recover an easy lower bound on the total variation distance between $\overline{\text{RT}}_n^t$ and $\pi_n := 1/(2^n n!)$.

Lemma 4.3.3. *The random transposition shuffle $\overline{\text{RT}}_n$ satisfies the following bound for any $c > 0$:*

$$\liminf_{n \rightarrow \infty} \|\overline{\text{RT}}_n^{(n/2) \log n - cn} - \pi_n\|_{\text{TV}} \geq \frac{1}{e} - e^{-e^{2c}}. \quad (4.9)$$

Proof. The proof of this lemma follows from the fact that total variation distance can only decrease under projections, see [32, Lemma 7.10]. Form a surjective homomorphism $\psi : B_n \rightarrow S_n$ by ignoring the signs on permutations in B_n , i.e., for $\sigma \in B_n$ with decomposition $\sigma = \left(\sum_{i=1}^n \xi_i^{x_i}\right) \eta$ we have $\psi(\sigma) = \eta$. Consider the preimage of ψ at a permutation in $\eta \in S_n$, we find that

$$\sum_{\sigma \in \psi^{-1}(\eta)} \overline{\text{RT}}_n(\sigma) = \text{RT}_n(\eta).$$

Extending this equality forward in time we have,

$$\sum_{\sigma \in \psi^{-1}(\eta)} \overline{\text{RT}}_n^t(\sigma) = \text{RT}_n^t(\eta) \text{ for all } t \geq 0 \text{ and } \eta \in S_n.$$

Thus, we may reduce the total variation distance of random transposition shuffle on B_n to random transposition shuffle on S_n :

$$\begin{aligned} \|\overline{\text{RT}}_n^t - \pi_n\|_{\text{TV}} &= \frac{1}{2} \sum_{\sigma \in B_n} |\overline{\text{RT}}_n^t(\sigma) - \pi_n(\sigma)| \\ &\geq \frac{1}{2} \sum_{\eta \in S_n} \left| \sum_{\substack{\sigma \in B_n \\ \sigma \in \psi^{-1}(\eta)}} \overline{\text{RT}}_n^t(\sigma) - \pi_n(\sigma) \right| = \|\text{RT}_n^t - \pi_{S_n}\|_{\text{TV}}. \end{aligned}$$

Using the above inequality the lower bound is then established from Theorem 2.3.2. \square

4.3.1 Upper Bound for the Random Transposition Shuffle

To complete a cutoff argument for the random transposition shuffle on B_n we need to show that its mixing time is at most $(n/2) \log n$. The random transposition shuffle is reversible so we can use the eigenvalues given by Theorem 4.3.2 to upper bound the total variation distance between $\overline{\text{RT}}_n^t$ and π_n (see Theorem 1.1.27).

$$\begin{aligned} 4\|\overline{\text{RT}}_n^t - \pi_n\|_{\text{TV}}^2 &\leq \sum_{\substack{\bar{\lambda} \vdash n \\ \bar{\lambda} \neq ((n), (0))}} d_{\bar{\lambda}}^2 \text{eig}(\bar{\lambda})^{2t} \tag{4.10} \\ &= \sum_{k=0}^n \sum_{\substack{\bar{\lambda} \vdash n \\ \bar{\lambda} \neq ((n), (0)) \\ |\lambda^1|=k}} \binom{n}{k}^2 d_{\lambda^1}^2 d_{\lambda^2}^2 \left(\frac{k + 2\text{Diag}(\lambda^1) + 2\text{Diag}(\lambda^2)}{n^2} \right)^{2t}. \tag{4.11} \end{aligned}$$

The first step in bounding the summation (4.11) is to decrease the number of bi-partitions $\bar{\lambda}$ we have to consider. For any bi-partition $\bar{\lambda} = (\lambda^1, \lambda^2)$, we can create at most eight other bi-partitions by taking transposes of λ^1, λ^2 , as well as swapping the positions of λ^1 and λ^2 . For example two of the bi-partitions we can form following these rules are $(\lambda^{2'}, \lambda^{1'})$, and $(\lambda^{1'}, \lambda^2)$. Furthermore, all the bi-partitions we can form have the same dimension because $d_{\lambda} = d_{\lambda'}$ and $\binom{n}{|\lambda^1|} = \binom{n}{|\lambda^2|}$. If we consider the eigenvalues associated to these eight possible bi-partitions, the largest eigenvalue in magnitude must come from the bi-partition $\bar{\lambda}$ with $|\lambda^1| \geq |\lambda^2|$ and $\text{Diag}(\lambda^1) \geq 0, \text{Diag}(\lambda^2) \geq 0$. Therefore, the contribution of any bi-partition to (4.11) may be upper bounded by a bi-partition $\bar{\lambda}$ with $|\lambda^1| \geq |\lambda^2|, \text{Diag}(\lambda^1) \geq 0, \text{Diag}(\lambda^2) \geq 0$. This allows us to reduce the number of bi-partitions we

have to consider by roughly a factor of 8. However, we need to make a special exception for the bi-partitions $((1^n), (0)), ((0), (n)), ((0), (1^n))$, because $((n), (0))$ does not belong to our sum. The eigenvalues associated to these bi-partitions are: $\text{eig}((1^n), (0)) = -1 + (2/n)$ and $\text{eig}((0), (n)) = -\text{eig}((0), (1^n)) = 1 - (1/n)$. These eigenvalues all have dimension 1 and so can easily be seen to be bounded at time $(n/2) \log n$,

$$\limsup_{n \rightarrow \infty} \left(1 - \frac{2}{n}\right)^{n \log n} + 2 \left(1 - \frac{1}{n}\right)^{n \log n} = 0.$$

Altogether applying our reduction we may see that the summation (4.11) is upper bounded by

$$3 \left(1 - \frac{1}{n}\right)^{2t} + 8 \sum_{k \geq n/2} \sum_{\substack{\bar{\lambda} \vdash n \\ \bar{\lambda} \neq ((n), (0)) \\ |\lambda^1| = k \\ \text{Diag}(\lambda^1) \geq 0 \\ \text{Diag}(\lambda^2) \geq 0}} \binom{n}{k}^2 d_{\lambda^1}^2 d_{\lambda^2}^2 \left(\frac{k + 2\text{Diag}(\lambda^1) + 2\text{Diag}(\lambda^2)}{n^2} \right)^{2t}. \quad (4.12)$$

The values $\text{Diag}(\lambda^1)$ and $\text{Diag}(\lambda^2)$ are closely related to the eigenvalues $\text{eig}(\lambda^1)$ and $\text{eig}(\lambda^2)$ for the random transposition shuffle on S_k and S_{n-k} respectively. In fact we can upper bound $\text{eig}(\bar{\lambda})$ by

$$\text{eig}(\bar{\lambda}) \leq \frac{k + 2\text{Diag}(\lambda^1) + (n-k) + 2\text{Diag}(\lambda^2)}{n^2} = \frac{k^2}{n^2} \text{eig}(\lambda^1) + \frac{(n-k)^2}{n^2} \text{eig}(\lambda^2).$$

We have previously bounded the value of $\text{Diag}(\lambda)$ for λ a partition of n using information about λ_1 . The following bound was given in Lemma 2.3.4 for $\lambda \vdash n$:

$$2\text{Diag}(\lambda) \leq \begin{cases} (n-1)n - 2(n-\lambda_1)(\lambda_1+1) & \text{if } \lambda_1 \geq \frac{n}{2} \\ (\lambda_1-1)n & \text{for all } \lambda \end{cases}. \quad (4.13)$$

Using this result we may establish bounds on the value $\text{eig}(\bar{\lambda})$ for particular choices of λ^1 and λ^2 . We would like to analyse equation (4.12) by reducing it to previously studied bounds for the random transposition shuffle on the symmetric group. This technique was used by Schoolfield to analyse the mixing time of the random transposition shuffle on B_n with independent card flips [40], and recently by Ghosh to analyse the mixing time of the flip-transpose top with random shuffle [23]. However, for the shuffle $\overline{\text{RT}}_n$ this technique leads to problems which stem from how our eigenvalues $\text{eig}(\bar{\lambda})$ are composed from the eigenvalues $\text{eig}(\lambda^1)$ and $\text{eig}(\lambda^2)$. To see where this approach fails let us focus on an example. Suppose that $\lambda^1 = k - i$ with $i \leq k/4$, and $\lambda^2 = (n-k) - j$ with $j < (n-k)/4$, in other words both λ^1, λ^2 are large partitions of k and $n-k$ respectively (see Section 3.4). In this case the eigenvalue $\text{eig}(\bar{\lambda})$ may be bounded as follows:

$$\frac{k + 2\text{Diag}(\lambda^1) + 2\text{Diag}(\lambda^2)}{n^2} \leq 1 - \frac{2k(n-k)}{n^2} - \frac{2i(k-i+1)}{n^2} - \frac{2j(n-k-j+1)}{n^2}. \quad (4.14)$$

In the above bound there are three different negative terms, we want to use the first to bound the binomial coefficient $\binom{n}{k}^2$, the second to bound $d_{\lambda^1}^2$, and the third to bound $d_{\lambda^2}^2$. Making use of Lemma 2.3.5 and setting $t = (n/2) \log n + cn$ we can reduce the contributions

of large partitions in (4.12) to:

$$\begin{aligned}
& \sum_{k \geq n/2}^n \sum_{i=0}^{k/4} \sum_{j=0}^{(n-k)/4} \sum_{\substack{\bar{\lambda} \vdash (n), (0) \\ |\lambda^1|=k \\ \lambda_1^1=k-i \\ \lambda_1^2=n-k-j}} \binom{n}{k}^2 d_{\lambda^1}^2 d_{\lambda^2}^2 \left(1 - \frac{2k(n-k)}{n^2} - \frac{2i(k-i+1)}{n^2} - \frac{2j(n-k-j+1)}{n^2} \right)^{2t} \\
&= \sum_{k \geq n/2}^n \sum_{i=0}^{k/4} \sum_{j=0}^{(n-k)/4} \binom{n}{k}^2 \left(1 - \frac{2k(n-k)}{n^2} - \frac{2i(k-i+1)}{n^2} - \frac{2j(n-k-j+1)}{n^2} \right)^{2t} \sum_{\lambda_1^1=k-i} d_{\lambda^1}^2 \sum_{\lambda_1^2=n-k-j} d_{\lambda^2}^2 \\
&\leq \sum_{k \geq n/2}^n \sum_{i=0}^{k/4} \sum_{j=0}^{(n-k)/4} \binom{n}{k}^2 \binom{k}{i}^2 i! \binom{n-k}{j}^2 j! \left(1 - \frac{2k(n-k)}{n^2} - \frac{2i(k-i+1)}{n^2} - \frac{2j(n-k-j+1)}{n^2} \right)^{2t} \\
&\leq \sum_{k \geq n/2}^n \sum_{i=0}^{k/4} \sum_{j=0}^{(n-k)/4} \binom{n}{k}^2 \binom{k}{i}^2 i! \binom{n-k}{j}^2 j! e^{-(n \log n + 2cn) \left(\frac{2k(n-k)}{n^2} + \frac{2i(k-i+1)}{n^2} + \frac{2j(n-k-j+1)}{n^2} \right)} \\
&\leq e^{-2c} \sum_{k \geq n/2}^n \sum_{i=0}^{k/4} \sum_{j=0}^{(n-k)/4} \binom{n}{k}^2 \binom{k}{i}^2 i! \binom{n-k}{j}^2 j! n^{-\frac{2k(n-k)}{n} - \frac{2i(k-i+1)}{n} - \frac{2j(n-k-j+1)}{n}} \tag{4.15}
\end{aligned}$$

From here we split equation (4.15) into three separate summations, one for the binomial coefficients and one each for the dimensions of λ^1 and λ^2 ,

$$e^{-2c} \sum_{k \geq n/2}^n \binom{n}{k}^2 n^{-\frac{2k(n-k)}{n}} \sum_{i=0}^{k/4} \binom{k}{i}^2 i! n^{-\frac{2i(k-i+1)}{n}} \sum_{j=0}^{(n-k)/4} \binom{n-k}{j}^2 j! n^{-\frac{2j(n-k-j+1)}{n}}.$$

We can now try to show that equation (4.15) is bounded in n by analysing each of these summations separately. The sums corresponding to partitions λ^1 and λ^2 given by

$$\sum_{i=0}^{k/4} \binom{k}{i}^2 i! n^{-\frac{2i(k-i+1)}{n}} \quad \text{and} \quad \sum_{j=0}^{(n-k)/4} \binom{n-k}{j}^2 j! n^{-\frac{2j(n-k-j+1)}{n}} \tag{4.16}$$

are closely related to the bounds analysed by Diaconis [16, Chapter 3D] for the random transposition shuffle on S_k and S_{n-k} , these are respectively given by,

$$\sum_{i=0}^{k/4} \binom{k}{i}^2 i! k^{-\frac{2i(k-i+1)}{k}} \quad \text{and} \quad \sum_{j=0}^{(n-k)/4} \binom{n-k}{j}^2 j! (n-k)^{-\frac{2j(n-k-j+1)}{(n-k)}}. \tag{4.17}$$

Diaconis showed the summations in (4.17) are bounded by a universal constant as $k \rightarrow \infty$ and $n-k \rightarrow \infty$ respectively. The difference between equations (4.16) and (4.17) comes from the eigenvalue $\text{eig}(\bar{\lambda})$ being of a different dimension compared with those of $\text{eig}(\lambda^1)$, $\text{eig}(\lambda^2)$. These differences make the sums in (4.16) challenging to bound. For example, if $n-k = O(n)$ we find that the summation given by λ^2 is unbounded as $n \rightarrow \infty$. A similar issue is found when we consider λ^1 and λ^2 to be any other combination of large or small partitions. Hence, trying to reduce the analysis of (4.12) to separate bounds related to the symmetric group case would appear to fail.

There are several ways we could try to circumvent this issue. The first way is to consider a more holistic approach to bounding equation (4.15). The binomial term given by $\binom{n}{k}^2 n^{-\frac{2k(n-k)}{n}}$ is not only bounded in n but is decaying as $k \rightarrow n/2$, whereas the terms

$\binom{k}{i}^2 i! n^{-\frac{2i(k-i+1)}{n}}$ and $\binom{n-k}{j}^2 j! n^{-\frac{2j(n-k-j+1)}{n}}$ corresponding to λ^1 and λ^2 respectively are growing as $k \rightarrow n/2$. By considering all these terms together instead of as three separate sums we may be able to use the decay of the binomial term to counteract the growth of the λ^1 and λ^2 terms. Alongside this it may be useful to consider a more fine-grain analysis of the eigenvalues, either by further restricting the choices of λ^1 and λ^2 , or by splitting the analysis into cases depending on whether $n - k = O(1)$ or $n - k = O(n)$.

To date we have been unable to successfully use the eigenvalues to establish an upper bound on the mixing time of the random transposition shuffle on B_n , but we conjecture that an upper bound of time $(n/2) \log n$ may be found. Thus, we expect the random transposition shuffle on B_n to exhibit a total variation cutoff at the same time as the random transposition shuffle on S_n .

Conjecture 4.3.4. *The random transposition shuffle $\overline{\text{RT}}_n$ satisfies the following bound:*

$$\lim_{c \rightarrow \infty} \liminf_{n \rightarrow \infty} \|\overline{\text{RT}}_n^{(n/2) \log n + cn} - \pi_n\|_{\text{TV}} = 0 \quad (4.18)$$

Thus, the random transposition shuffle on the hyperoctahedral group exhibits a cutoff in total variation distance at time $(n/2) \log n$.

4.3.2 Lifting Eigenvectors for the Random Transposition Shuffle

We lift the eigenvectors of $\overline{\text{RT}}_n$ by reducing the lifting of bi-partitions $\bar{\lambda}$ to the lifting of the individual partitions λ^1 and λ^2 . This allows us to extend the technique of lifting eigenvectors that we developed for the random transposition shuffle on S_n to the random transposition shuffle on B_n . The first step as before is to turn the probability $\overline{\text{RT}}_n$ into an element of the group algebra \mathfrak{B}_n .

Definition 4.3.5. The random transposition shuffle on B_n may be viewed as the following element of the group algebra \mathfrak{B}_n .

$$\overline{\text{ART}}_n = n \cdot e + \sum_{1 \leq i \leq n} \xi_i + 2 \sum_{1 \leq i < j \leq n} (ij) + 2 \sum_{1 \leq i < j \leq n} \xi_i \xi_j (ij). \quad (4.19)$$

Note that here we have scaled by $2n^2$.

The eigenvectors and eigenvalues of $\overline{\text{RT}}_n$ are in one-to-one correspondence with those given by $\overline{\text{ART}}_n$ acting on the regular module \mathfrak{B}_n . Using the canonical decomposition of the regular module we now focus our attention on finding the eigenvectors of $\overline{\text{ART}}_n$ belonging to the Specht modules $S^{\bar{\lambda}}$. If we compare the elements $\overline{\text{ART}}_{n+1}$ and $\overline{\text{ART}}_n$ we can see they have a close relationship to one another,

$$\overline{\text{ART}}_{n+1} - \overline{\text{ART}}_n = e + \xi_{n+1} + 2 \sum_{1 \leq i \leq n} (i \ n+1) + 2 \sum_{1 \leq i \leq n} \xi_i \xi_{n+1} (i \ n+1). \quad (4.20)$$

The equation (4.20) only depends on transpositions involving $n + 1$, this is similar to equation (4.20), and gives us an insight into why lifting eigenvectors works for the random transposition shuffle on B_n . To turn eigenvectors for the module $S^{\bar{\lambda}}$ into those for $S^{\bar{\lambda} + e_i^k}$ we introduce new adding and switching operators.

Definition 4.3.6. Let w be a word in $M^{\bar{\lambda}}$, we may add a letter to w to form a word of size $n + 1$. Define the linear *adding operators*, denoted Φ_a , for letter $a \in [n]$ as follows,

$$\Phi_a(w) = wa, \quad \Phi_a^+(w) = wa^+, \quad \Phi_a^-(w) = wa^-.$$

Note that the first operator takes us from $M^{\bar{\lambda}}$ to space $M^{\bar{\lambda}+e_a^1}$ whereas the second and third operators takes us to $M^{\bar{\lambda}+e_a^2}$. Using the above basic adding operators define two new adding operators:

$$\Phi_a^1 = \Phi_a, \quad \Phi_a^2 = \Phi_a^+ - \Phi_a^-.$$

These new adding operators are critical to our analysis of the random transposition shuffle.

Definition 4.3.7. Let w be a word in $M^{\bar{\lambda}}$ and $a, b \in [\bar{n}]$. We define the linear *switching operators* as follows:

$$\Theta_{b,a}(w) = \sum_{\substack{1 \leq i \leq n \\ w_i = b}} w_1 \dots w_{i-1} a w_{i+1} \dots w_n \quad (4.21)$$

Using this single operator we define three other switching operators for $a, b \in [n]$:

$$\Theta_{b^\pm, a}(w) = \Theta_{b^-, a} + \Theta_{b^+, a} \quad (4.22)$$

$$\Theta_{b^\pm, a^\pm}^+(w) = \Theta_{b^+, a^+} + \Theta_{b^-, a^-} \quad (4.23)$$

$$\Theta_{b^\pm, a^\pm}^-(w) = \Theta_{b^+, a^-} + \Theta_{b^-, a^+} \quad (4.24)$$

The operator $\Theta_{b^\pm, a}$ takes all signed occurrences of b and replaces them by the unsigned a . The operators $\Theta_{b^\pm, a^\pm}^\pm$ swap signed occurrences of b for signed occurrences of a , with Θ^+ fixing the signs, and Θ^- swapping them. Define two new operators $\Theta_{b,a}^1, \Theta_{b,a}^2$ with domain $M^{\bar{\lambda}}$ for the unsigned symbols $a, b \in [n]$ as follows:

$$\Theta_{b,a}^1 = \Theta_{b,a} \quad (4.25)$$

$$\Theta_{b,a}^2 = \Theta_{b^\pm, a^\pm}^+ - \Theta_{b^\pm, a^\pm}^- \quad (4.26)$$

The operator $\Theta_{b,a}^1$ acts only on unsigned letters, and $\Theta_{b,a}^2$ acts only on signed letters. The image of $\Theta_{b,a}^1$ is $\bar{\lambda} - e_b^1 + e_a^1$, and the image of $\Theta_{b,a}^2$ is $\bar{\lambda} - e_b^2 + e_a^2$.

Lemma 4.3.8. *The switching operators $\Theta_{b^\pm, a}, \Theta_{b,a}^1$ and $\Theta_{b,a}^2$ are \mathfrak{B}_n -module morphisms.*

Proof. The operators $\Theta_{b,a}^1$ and $\Theta_{b^\pm, a}$, swap a letter b (signed or unsigned) for an unsigned letter a , this is unaffected by movement of the letter and changes in sign (as it always ends up unsigned), therefore it commutes with elements B_n . The operator $\Theta_{b^\pm, a^\pm}^\pm$ swaps a signed letter b^\pm for a^\pm while preserving the current sign of b^\pm so it does not matter if the letter is moved or the sign of the letter is changed before or after applying the operator. Similar logic holds for the operator $\Theta_{b^\pm, a^\pm}^\mp$. Hence, the signed lifting operator $\Theta_{b,a}^2$ is a \mathfrak{B}_n -module morphism. □

Notice that we have not defined a switching operator which takes unsigned letters and turns them into a signed letters, the reason behind this is there is no way to give a letter a

sign and respect the action of B_n . The adding operators are not module morphisms, however, they can be seen to commute with the switching operators with an extra adjustment term.

Lemma 4.3.9. *The adding and switching operators satisfy the following equalities:*

$$\Phi_b^1 \circ \Theta_{b,a}^1 = \Theta_{b,a}^1 \circ \Phi_b^1 - \Phi_a^1 \quad (4.27)$$

$$\Phi_b^2 \circ \Theta_{a,b}^2 = \Theta_{a,b}^2 \circ \Phi_b^2 - 2 \cdot \Phi_a^2. \quad (4.28)$$

We can already see that lifting eigenvectors for B_n requires more careful tools than for S_n ; we have gone from requiring one adding and switching operator to two of them each. We now establish a relationship between the modules of \mathfrak{B}_n and \mathfrak{B}_{n+1} using the newly defined operators. This is similar to Theorem 3.2.5 however we now consider two cases, based on if we lift the partition λ^1 or λ^2 .

Theorem 4.3.10. *Let $n \in \mathbb{N}$ and $\bar{\lambda} \vdash n$. For words in $M^{\bar{\lambda}}$ we have the following equalities:*

$$\overline{\text{ART}}_{n+1} \circ \Phi_a^1 - \Phi_a^1 \circ \overline{\text{ART}}_n = 2\Phi_a^1 + 2 \sum_{1 \leq b \leq n} 2 \cdot \Phi_b^1 \circ \Theta_{b,a}^1 + (\Phi_b^+ + \Phi_b^-) \circ \Theta_{b^\pm, a}. \quad (4.29)$$

$$\overline{\text{ART}}_{n+1} \circ \Phi_a^2 - \Phi_a^2 \circ \overline{\text{ART}}_n = 2 \sum_{1 \leq b \leq n} \Phi_b^2 \circ \Theta_{b,a}^2. \quad (4.30)$$

Proof. Take $w \in M^{\bar{\lambda}}$ a generic element to verify these equations with. In both equations all non-identity elements of $\overline{\text{ART}}_{n+1}$ which do not involve $n+1$ commute with the adding operators Φ_a^1, Φ_a^2 , and therefore cancel with $\overline{\text{ART}}_n$. This just leaves us the elements present in equation (4.20) to work with. First we prove equation (4.29), using the above observation we reduce it to,

$$\overline{\text{ART}}_{n+1} \circ \Phi_a^1 - \Phi_a^1 \circ \overline{\text{ART}}_n = \left(e + \xi_{n+1} + 2 \sum_{1 \leq i \leq n} (i \ n+1) + 2 \sum_{1 \leq i \leq n} \xi_i \xi_{n+1} (i \ n+1) \right) \Phi_a^1(w).$$

Clearly $(e + \xi_{n+1})\Phi_a^1(w) = 2\Phi_a^1(w)$. The letter we add to w is unsigned, define the word $v_i = w_1 \dots w_{i-1} a w_{i+1} \dots w_n$. Then the positive transpositions act on $w a$ in the following way:

$$((i \ n+1) + \xi_i \xi_{n+1} (i \ n+1)) \Phi_a^1(w) = \begin{cases} v_i b + v_i b = 2v_i b & \text{if } w_i = b \\ v_i b^+ + v_i b^- & \text{if } w_i = b^+ \\ v_i b^- + v_i b^+ & \text{if } w_i = b^- \end{cases}$$

Thus, summing over all i or equivalently all signed and unsigned letters we recover (4.29). Now we prove equation (4.30), again it may be reduced it to

$$\overline{\text{ART}}_{n+1} \circ \Phi_a^2 - \Phi_a^2 \circ \overline{\text{ART}}_n = \left(e + \xi_{n+1} + 2 \sum_{1 \leq i \leq n} (i \ n+1) + 2 \sum_{1 \leq i \leq n} \xi_i \xi_{n+1} (i \ n+1) \right) \Phi_a^2(w).$$

Looking at the negative transposition we find $(e + \xi_{n+1})\Phi_a^2(w) = (e + \xi_{n+1})(w a^+ - w a^-) = 0$. For the positive transpositions we must be careful because the added letter a is now signed. Define $v_i^+ = w_1 \dots w_{i-1} a^+ w_{i+1} \dots w_n$, and $v_i^- = w_1 \dots w_{i-1} a^- w_{i+1} \dots w_n$. The posi-

tive transpositions act on $(wa^+ - wa^-)$ in the following way:

$$((i \ n + 1) + \xi_i \xi_{n+1}(i \ n + 1)) \Phi_a^2(w) = \begin{cases} (v_i^+ b - v_i^- b) + (v_i^- b - v_i^+ b) = 0 & \text{if } w_i = b \\ (v_i^+ b^+ - v_i^- b^+) + (v_i^- b^- - v_i^+ b^-) = \Phi_b^2(v_i^+ - v_i^-) & \text{if } w_i = b^+ \\ (v_i^+ b^- - v_i^- b^-) + (v_i^- b^+ - v_i^+ b^+) = \Phi_b^2(v_i^- - v_i^+) & \text{if } w_i = b^- \end{cases}$$

Thus summing over all $1 \leq i \leq n$ or equivalently all signed letters b^\pm in w completes equation (4.30). \square

Once refined, equations (4.29), (4.30) will enable us to turn eigenvectors of $\overline{\text{ART}}_n$ into those for $\overline{\text{ART}}_{n+1}$. The next step is to restrict their domains to a Specht module $S^{\bar{\lambda}}$.

Lemma 4.3.11. *Let $\bar{\lambda} \vdash n$ be such that $\bar{\mu} + e_b^k = \bar{\lambda} + e_a^l$ for some $k, l \in \{1, 2\}$ and $a, b \in [n]$. If $k > l$ then $\Theta_{b^\pm, a}(S^{\bar{\lambda}}) = 0$. If $k = l$ then $\Theta_{b, a}^k$ is non-zero on $S^{\bar{\lambda}}$ if and only if λ^k dominates the non-increasing rearrangement of μ^k . In particular, if $a < b$, then $\Theta_{b, a}^k(S^{\bar{\lambda}}) = 0$.*

Proof. We know that for both cases the respective switching operators $\Theta(S^{\bar{\lambda}})$ must belong in $M^{\bar{\mu}}$, and so for the map to be non-zero we must have $\bar{\lambda}$ dominating the non-increasing rearrangement of $\bar{\mu}$ by Lemma 4.2.24. If $k > l$ then $\bar{\mu} \not\geq \bar{\lambda}$ therefore $\Theta_{b^\pm, a}(S^{\bar{\lambda}}) = 0$. If $k = l$ the assertion holds by the same reasoning as Lemma 3.2.7. \square

For an illustration of Lemma 4.3.11 take the polytabloid $e_{\bar{T}}$ belonging to module $S^{((2,1), (2))}$ from Example 4.2.16 (also equation (4.6)), the switching operators $\Theta_{1^\pm, 1}$, $\Theta_{2, 1}^1$, and $\Theta_{1, 3}^1$ applied to $e_{\bar{T}}$ gives us:

$$\begin{aligned} \Theta_{1^\pm, 1}(e_{\bar{T}}) &= (11211^+ + 1121^+1) - (21111^+ + 2111^+1) - (11211^+ + 1121^-1) + (21111^+ + 2111^-1) \\ &\quad - (11211^- + 1121^+1) + (21111^- + 2111^+1) + (11211^- + 1121^-1) - (21111^- + 2111^-1) = 0 \\ \Theta_{2, 1}^1(e_{\bar{T}}) &= 1111^+1^+ - 1111^+1^- - 1111^-1^+ + 1111^-1^- \\ &\quad - 1111^+1^- + 1111^+1^- + 1111^-1^- - 1111^-1^- = 0 \\ \Theta_{1, 3}^1(e_{\bar{T}}) &= (1321^+1^+ + 3121^+1^+) - (2131^+1^+ + 2311^+1^+) \\ &\quad - (1321^-1^+ + 3121^-1^+) + (2131^-1^+ + 2311^-1^+) \\ &\quad - (1321^+1^- + 3121^+1^-) + (2131^+1^- + 2311^+1^-) \\ &\quad + (1321^-1^- + 3121^-1^-) - (2131^-1^- + 2311^-1^-) \neq 0. \end{aligned}$$

Applying Lemma 4.3.11 we now restrict equations (4.29) and (4.30) to the domain of a Specht module $S^{\bar{\lambda}}$.

Corollary 4.3.12. *Let $\bar{\lambda} \vdash n$, restricting equations (4.29) and (4.30) to have domain $S^{\bar{\lambda}}$ we find the following:*

$$\overline{\text{ART}}_{n+1} \circ \Phi_a^1 - \Phi_a^1 \circ \overline{\text{ART}}_n|_{S^{\bar{\lambda}}} = 2\Phi_a^1|_{S^{\bar{\lambda}}} + 4 \sum_{1 \leq b \leq a} \Phi_b^1 \circ \Theta_{b, a}^1|_{S^{\bar{\lambda}}} \quad (4.31)$$

$$\overline{\text{ART}}_{n+1} \circ \Phi_a^2 - \Phi_a^2 \circ \overline{\text{ART}}_n|_{S^{\bar{\lambda}}} = 2 \sum_{1 \leq b \leq a} \Phi_b^2 \circ \Theta_{b, a}^2|_{S^{\bar{\lambda}}}. \quad (4.32)$$

Notice that equations (4.31) and (4.32) now only depend on unsigned or signed letters, respectively, being added or moved. This allows us to stop focusing on the bi-partition

$\bar{\lambda}$ and instead focus on the two single partitions λ^1, λ^2 . Equations (4.31), (4.32) are used to lift eigenvectors by adding a box to λ^1 or λ^2 respectively. Before we project onto the Specht modules $S^{\bar{\lambda}+e_a^1}$ or $S^{\bar{\lambda}+e_a^2}$ we must be sure they are contained as a submodule of the image. The next lemma is an analogue of Lemma 3.2.10 for the hyperoctahedral group.

Lemma 4.3.13. *The subspace $\Phi_a^k(S^{\bar{\lambda}})$ is contained in a \mathfrak{B}_{n+1} submodule of $M^{\bar{\lambda}+e_a^k}$, isomorphic to $\bigoplus S^{\bar{\mu}}$ where $\bar{\mu}$ ranges over the partitions obtained from $\bar{\lambda}$ by adding a box e_i^k with $i \leq a$.*

Proof. We begin by proving our hypothesis for $k = 1$, that is adding a box to λ^1 . Let w be a word in $M^{\bar{\lambda}}$, let $a \in [n+1]$, if letter b does not appear in w then

$$\Phi_a^1(w) = \Theta_{b,a}^1(\Phi_b^1(w)) = \Theta_{b,a}^1(wb).$$

Let $b = l(\lambda^1) + 1$, and consider the \mathfrak{B}_{n+1} -module

$$N_1 = \langle xb : x \in S^{\bar{\lambda}} \rangle \cong \text{Ind}_{\mathfrak{B}_n \times \mathfrak{B}_1}^{\mathfrak{B}_{n+1}} (S^{\bar{\lambda}} \otimes S^{((1), \emptyset)}) \cong \bigoplus_{\substack{\bar{\mu} \vdash n+1 \\ \lambda^1 \subseteq \mu^1 \\ \lambda^2 = \mu^2}} S^{\bar{\mu}}$$

where the last isomorphism follows from the branching rules for B_n (see [22, Theorem III.2]). Using the observation at the start of the proof, we obtain

$$\Phi_a^1(S^{\bar{\lambda}}) = \Theta_{b,a}^1(\Phi_b^1(S^{\bar{\lambda}})) \subseteq \Theta_{b,a}^1(\langle \Phi_b^1(S^{\bar{\lambda}}) \rangle) = \Theta_{b,a}^1(N_1) \cong \bigoplus_{\substack{\bar{\mu} \vdash n+1 \\ \lambda^1 \subseteq \mu^1 \\ \lambda^2 = \mu^2}} \Theta_{b,a}^1(S^{\bar{\mu}}).$$

Now note that $\Theta_{b,a}^1$ sends any word with evaluation $\bar{\lambda} + e_b^1$ to a word with evaluation $\bar{\lambda} + e_a^1$, and hence $\Theta_{b,a}^1(N_1) \subseteq \Theta_{b,a}^1(M^{\bar{\lambda}+e_b^1}) \subseteq M^{\bar{\lambda}+e_a^1}$. It follows that all nonzero summands $S^{\bar{\mu}}$ appearing on the right hand side occur for $\bar{\mu} \vdash n+1$ dominating the non-decreasing rearrangement of $\bar{\lambda} + e_a^1$, and then by Lemma 4.3.11 we can conclude that $\bar{\mu}$ is obtained from $\bar{\lambda}$ by adding a box e_i^1 with $i \leq a$, as required.

Suppose $k = 2$, let w be a word in $\bar{\lambda}$, let $a \in [n+1]$, if the letter b^\pm does not appear in w then

$$\Phi_a^2(w) = \frac{1}{2} \Theta_{b,a}^2 \circ \Phi_b^2(w) = \frac{1}{2} \Theta_{b,a}^2 (wb^+ - wb^-).$$

Let $b = l(\lambda^2) + 1$, consider the set

$$N_2 = \langle xb^+ - xb^- : x \in S^{\bar{\lambda}} \rangle \cong \text{Ind}_{\mathfrak{B}_n \times \mathfrak{B}_1}^{\mathfrak{B}_{n+1}} (S^{\bar{\lambda}} \otimes S^{(\emptyset, (1))}) \cong \bigoplus_{\substack{\bar{\mu} \vdash n+1 \\ \lambda^1 = \mu^1 \\ \lambda^2 \subseteq \mu^2}} S^{\bar{\mu}}$$

where the last isomorphism follows from the branching rules for B_n (see [22, Theorem III.2]) Using the observation at the start of the proof, we obtain as modules over \mathfrak{B}_{n+1} :

$$\Phi_a^2(S^{\bar{\lambda}}) = \Theta_{b,a}^2(\Phi_b^2(S^{\bar{\lambda}})) \subseteq \Theta_{b,a}^2(\langle \Phi_b^2(S^{\bar{\lambda}}) \rangle) = \Theta_{b,a}^2(N_2) \cong \bigoplus_{\substack{\bar{\mu} \vdash n+1 \\ \lambda^1 = \mu^1 \\ \lambda^2 \subseteq \mu^2}} \Theta_{b,a}^2(S^{\bar{\mu}}).$$

By the same reasoning as before the final direct sum must be over all $\bar{\mu}$ with μ^2 formed

by adding a box e_i^2 to $\bar{\lambda}$ with $i \leq a$. \square

Definition 4.3.14. Let $\pi^{\bar{\mu}}$ be the isotypic projector onto the simple module $S^{\bar{\mu}}$. Define the *lifting operator* for $k \in \{1, 2\}$ as:

$$\kappa_a^{\bar{\lambda}, \bar{\mu}, k} = (\pi^{\bar{\mu}} \circ \Phi_a^k)|_{S^{\bar{\lambda}}} : S^{\bar{\lambda}} \rightarrow M^{\bar{\lambda} + e_a^k}. \quad (4.33)$$

Lemma 4.3.13 tells us that the image of Φ_a^k may contain at most one copy of any Specht module $S^{\bar{\mu}}$. Therefore, if there exists a copy in the image the element $\kappa_a^{\bar{\lambda}, \bar{\mu}, k}(w)$ belongs exactly to the Specht module $S^{\bar{\mu}}$ within $M^{\bar{\lambda} + e_a^k}$.

Corollary 4.3.15. For any $\bar{\lambda} \vdash n$, $a \in \{1, 2, \dots, l(\lambda^k) + 1\}$ and $1 \leq i \leq a$, there exists $v \in S^{\bar{\lambda}}$ such that

$$\kappa_a^{\bar{\lambda}, \bar{\lambda} + e_i^k, k}(v) \neq 0.$$

Proof. Suppose $\kappa_a^{\bar{\lambda}, \bar{\lambda} + e_i^k, k}(v) = 0$, then the image $\Phi_a^k(S^{\bar{\lambda}})$ lies in the kernel of the projection $\pi^{\bar{\lambda} + e_i^k} : M^{\bar{\lambda} + e_i^k} \rightarrow S^{\bar{\lambda} + e_i^k}$ which is an \mathfrak{B}_{n+1} -submodule with no component equal to $S^{\bar{\lambda} + e_i^k}$. Hence, the submodule generated by $\Phi_a^k(S^{\bar{\lambda}})$ has no component equal to $S^{\bar{\lambda} + e_i^k}$. But we previously observed that (with notation from Lemma 4.3.13 and the two cases presented together)

$$\langle \Phi_a^k(S^{\bar{\lambda}}) \rangle = \langle \Theta_{b,a}^k(\Phi_b^k(S^{\bar{\lambda}})) \rangle = \Theta_{b,a}^k(\langle \Phi_b^k(S^{\bar{\lambda}}) \rangle) = \Theta_{b,a}^k(N_k) \cong \bigoplus_{1 \leq i \leq a} S^{\bar{\lambda} + e_i^k} \quad (4.34)$$

Since the corresponding right hand side contains $S^{\bar{\lambda} + e_i^k}$ as a summand, we have a contradiction. \square

Lemma 4.3.16. The linear operators $\kappa_a^{\bar{\lambda}, \bar{\lambda} + e_i^k, k}$ for $1 \leq i \leq a$ are \mathfrak{B}_n -module morphisms with trivial kernels. Therefore, these lifting operators are injective.

Proof. Our key observation is that $\sigma(\Phi_a^k(w)) = \Phi_a^k(\sigma w)$ for any $k \in \{1, 2\}$ and $\sigma \in \mathfrak{B}_n \subset \mathfrak{B}_{n+1}$, which fixes the coordinate $n+1$. Thus our lifting operators $\kappa_a^{\bar{\lambda}, \bar{\lambda} + e_i^k, k}$ are a composition of two \mathfrak{B}_n module morphisms. Finally we know from Corollary 4.3.15 that the map is non-zero and therefore by Schur's Lemma must be injective. \square

We are now ready to lift eigenvectors for the random transposition shuffle. We state two theorems, one for lifting λ^1 , and another for lifting λ^2 . Afterwards we combine the theorems to recover all the eigenvalues for the random transposition shuffle on B_n . The results below follow from work presented in [19] and Section 3.2.1 (Lemma 3.2.16 and Theorem 3.2.17) with suitable modifications to the new setting of the hyperoctahedral group.

Lemma 4.3.17. Let $\bar{\lambda} \vdash n$, and $a \in \{1, 2, \dots, l(\lambda^1) + 1\}$. Take $i \in [n]$ such that $1 \leq i \leq a$ and set $\bar{\mu} = \bar{\lambda} + e_i^1$. Then,

$$\overline{\text{ART}}_{n+1} \circ \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} - \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} \circ \overline{\text{ART}}_n = (2 + 4(\lambda_a^1 + 1 - a))\kappa_a^{\bar{\lambda}, \bar{\mu}, 1} + 4 \sum_{i \leq b < a} \Theta_{b,a}^1 \circ \kappa_b^{\bar{\lambda}, \bar{\mu}, 1}.$$

Proof. Continuing from Corollary 4.3.12 we know

$$\overline{\text{ART}}_{n+1} \circ \Phi_a^1 - \Phi_a^1 \circ \overline{\text{ART}}_n|_{S^{\bar{\lambda}}} = 2\Phi_a^1|_{S^{\bar{\lambda}}} + 4 \sum_{1 \leq b \leq a} \Phi_b^1 \circ \Theta_{b,a}^1|_{S^{\bar{\lambda}}}.$$

Apply the isotypic projection $\pi^{\bar{\mu}}$ to both sides of the equation. Since $\overline{\text{ART}}_{n+1}$ is given by the action of an element of the group algebra \mathfrak{B}_n and $\pi^{\bar{\mu}}$ is an \mathfrak{B}_{n+1} -module morphism, these operators commute and so we have

$$\overline{\text{ART}}_{n+1} \circ \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} - \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} \circ \overline{\text{ART}}_n = 2\kappa_a^{\bar{\lambda}, \bar{\mu}, 1} + 4 \sum_{1 \leq b \leq a} (\pi^{\bar{\mu}} \circ \Phi_b^1 \circ \Theta_{b,a}^1)|_{S^{\bar{\lambda}}}. \quad (4.35)$$

Applying equation (4.27) we may see that

$$(\pi^{\bar{\mu}} \circ \Phi_b^1 \circ \Theta_{b,a}^1) = \Theta_{b,a}^1 \circ \kappa_b^{\bar{\lambda}, \bar{\mu}, 1} - \kappa_a^{\bar{\lambda}, \bar{\mu}, 1}.$$

The right hand side of equation (4.35) now becomes,

$$2\kappa_a^{\bar{\lambda}, \bar{\mu}, 1} + 4 \sum_{1 \leq b \leq a} \left(\Theta_{b,a}^1 \circ \kappa_b^{\bar{\lambda}, \bar{\mu}, 1} - \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} \right) \quad (4.36)$$

Notice that if $b = a$ then $\Theta_{a,a}^1(w)$ acts as a scalar by the number of occurrences of the symbol a in w . All words in $\kappa_a^{\bar{\lambda}, \bar{\mu}, 1}(S^{\bar{\lambda}})$ contain $\lambda_a^1 + 1$ occurrences of a . Finally if $b < i$ we know that $\Phi_b^1(S^{\bar{\lambda}})$ does not contain $S^{\bar{\lambda} + e_i^1}$ as a submodule (by Lemma 4.3.13), so $\kappa_b^{\bar{\lambda}, \bar{\mu}, 1} = 0$. Thus, equation (4.36) is equal to

$$(2 + 4(\lambda_a^1 + 1 - a))\kappa_a^{\bar{\lambda}, \bar{\mu}, 1} + 4 \sum_{i \leq b < a} \Theta_{b,a}^1 \circ \kappa_b^{\bar{\lambda}, \bar{\mu}, 1}.$$

□

Lemma 4.3.18. *Let $\bar{\lambda} \vdash n$, and $a \in \{1, 2, \dots, l(\lambda^2) + 1\}$. Take $i \in [n]$ such that $1 \leq i \leq a$ and set $\bar{\mu} = \bar{\lambda} + e_i^2$. Then,*

$$\overline{\text{ART}}_{n+1} \circ \kappa_a^{\bar{\lambda}, \bar{\mu}, 2} - \kappa_a^{\bar{\lambda}, \bar{\mu}, 2} \circ \overline{\text{ART}}_n = 4(\lambda_a^2 + 1 - a)\kappa_a^{\bar{\lambda}, \bar{\mu}, 2} + 2 \sum_{i \leq b < a} \Theta_{b,a}^2 \circ \kappa_b^{\bar{\lambda}, \bar{\mu}, 2}.$$

Proof. This follows from the same proof as Lemma 4.3.17 replacing the use of the first equality in Lemma 4.3.11 with the second equality. □

Theorem 4.3.19 (Lifting for λ^1). *Let $\bar{\lambda} \vdash n$, and $a \in \{1, 2, \dots, l(\lambda^1) + 1\}$. Take $i \in [n]$ such that $1 \leq i \leq a$ and set $\bar{\mu} = \bar{\lambda} + e_i^1$. Then,*

$$\overline{\text{ART}}_{n+1} \circ \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} - \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} \circ \overline{\text{ART}}_n = (2 + 4(\lambda_i^1 + 1 - i))\kappa_a^{\bar{\lambda}, \bar{\mu}, 1}$$

Thus if we have an eigenvector $v \in S^{\bar{\lambda}}$ with eigenvalue ϵ , we find $\kappa_a^{\bar{\lambda}, \bar{\mu}, 1}(v)$ to be an eigenvector of $S^{\bar{\mu}}$ with eigenvalue $\epsilon + (2 + 4(\lambda_i^1 + 1 - i))$.

Proof. For $i = a$ the result follows from Lemma 4.3.17. Let $\mu = \bar{\lambda} + e_i^1$, again from Lemma 4.3.17 we know that

$$\overline{\text{ART}}_{n+1} \circ \kappa_i^{\bar{\lambda}, \bar{\mu}, 1} - \kappa_i^{\bar{\lambda}, \bar{\mu}, 1} \circ \overline{\text{ART}}_n = (2 + 4(\lambda_i^1 + 1 - i))\kappa_i^{\bar{\lambda}, \bar{\mu}, 1}.$$

Applying the linear operator $\Theta_{i,a}^1$ to the above equation,

$$\overline{\text{ART}}_{n+1} \circ \Theta_{i,a}^1 \circ \kappa_i^{\bar{\lambda}, \bar{\mu}, 1} - \Theta_{i,a}^1 \circ \kappa_i^{\bar{\lambda}, \bar{\mu}, 1} \circ \overline{\text{ART}}_n = (2 + 4(\lambda_i^1 + 1 - i))\Theta_{i,a}^1 \circ \kappa_i^{\bar{\lambda}, \bar{\mu}, 1}. \quad (4.37)$$

Consider the left hand side of (4.37), break up the lifting operator into $\kappa_i^{\bar{\lambda}, \bar{\mu}, 1} = \pi^{\bar{\mu}} \circ \Phi_i^1|_{S^{\bar{\lambda}}}$. The projection $\pi^{\bar{\mu}}$ commutes with the \mathfrak{B}_{n+1} -module morphism $\Theta_{i,a}^1$. Performing this and using equation (4.27) we obtain the equation below, restricted to $S^{\bar{\lambda}}$;

$$\begin{aligned}
& \overline{\text{ART}}_{n+1} \circ \pi^{\bar{\mu}} \circ \Theta_{i,a}^1 \circ \Phi_i^1 - \pi^{\bar{\mu}} \circ \Theta_{i,a}^1 \circ \Phi_i^1 \circ \overline{\text{ART}}_n \\
= & \overline{\text{ART}}_{n+1} \circ \pi^{\bar{\mu}} \circ (\Phi_i^1 \circ \Theta_{i,a}^1 + \Phi_i^1) - \pi^{\bar{\mu}} \circ (\Phi_i^1 \circ \Theta_{i,a}^1 + \Phi_i^1) \circ \overline{\text{ART}}_n \\
= & (\overline{\text{ART}}_{n+1} \circ \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} - \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} \circ \overline{\text{ART}}_n) + (\overline{\text{ART}}_{n+1} \circ \kappa_i^{\bar{\lambda}, \bar{\mu}, 1} - \kappa_i^{\bar{\lambda}, \bar{\mu}, 1} \circ \overline{\text{ART}}_n) \Theta_{i,a}^1 \\
= & (\overline{\text{ART}}_{n+1} \circ \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} - \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} \circ \overline{\text{ART}}_n) + (2 + 4(\lambda_i^1 + 1 - i)) \kappa_i^{\bar{\lambda}, \bar{\mu}, 1} \circ \Theta_{i,a}^1, \tag{4.38}
\end{aligned}$$

Manipulating the right hand side of (4.35) by again splitting our lifting operator $\kappa_i^{\bar{\lambda}, \bar{\mu}, 1}$ and applying (4.27) we recover,

$$\begin{aligned}
(4.35) &= (2 + 4(\lambda_i^1 + 1 - i)) \pi^{\bar{\mu}} \circ \Theta_{i,a}^1 \circ \Phi_i^1 \\
&= (2 + 4(\lambda_i^1 + 1 - i)) \pi^{\bar{\mu}} \circ (\Phi_i^1 \circ \Theta_{i,a}^1 + \Phi_i^1) \\
&= (2 + 4(\lambda_i^1 + 1 - i)) \kappa_i^{\bar{\lambda}, \bar{\mu}, 1} \circ \Theta_{i,a}^1 + (2 + 4(\lambda_i^1 + 1 - i)) \kappa_a^{\bar{\lambda}, \bar{\mu}, 1}. \tag{4.39}
\end{aligned}$$

Combining equations (4.38) and (4.39) completes the proof. \square

Theorem 4.3.20 (Lifting for λ^2). *Let $\bar{\lambda} \vdash n$, and $a \in \{1, 2, \dots, l(\lambda^2) + 1\}$. Take $i \in [n]$ such that $1 \leq i \leq a$ and set $\bar{\mu} = \bar{\lambda} + e_i^2$. Then,*

$$\overline{\text{ART}}_{n+1} \circ \kappa_a^{\bar{\lambda}, \bar{\mu}, 2} - \kappa_a^{\bar{\lambda}, \bar{\mu}, 2} \circ \overline{\text{ART}}_n = 4(\lambda_i^2 + 1 - i) \kappa_a^{\bar{\lambda}, \bar{\mu}, 2}$$

Thus if we have an eigenvector $v \in S^{\bar{\lambda}}$ with eigenvector ϵ , we find $\kappa_a^{\bar{\lambda}, \bar{\mu}, 2}(v)$ to be an eigenvector of $S^{\bar{\mu}}$ with eigenvalue $\epsilon + 4(\lambda_i^2 + 1 - i)$.

Proof. This follows the same proof as Theorem 4.3.19. \square

Notice that there is a quantifiable difference in the change in eigenvalue depending on whether we lift λ^1 or λ^2 . The value $\lambda_i^k + 1 - i$ refers to the diagonal index of the new box added to λ^k , this value does not depend on current size of $\bar{\lambda}$. Combining Theorems 4.3.19 and 4.3.20 we recover all the eigenvectors for the random transposition shuffle belonging to each Specht module, this allows us to give a swift proof of Theorem 4.3.2

Lemma 4.3.21. *For any $\bar{\mu} \vdash n + 1$ we may find a basis of eigenvectors of $\overline{\text{ART}}_{n+1}$ for the module $S^{\bar{\mu}}$ by lifting the eigenvectors of $\overline{\text{ART}}_n$ belonging in the modules $S^{\bar{\lambda}}$ with $\bar{\lambda} \vdash n$ and $\bar{\lambda} \subset \bar{\mu}$.*

Proof. We proceed by induction. For $n = 1$ we know the simple modules $\langle a \rangle = S^{((1), (0))}$, $\langle a^+ - a^- \rangle = S^{((0), (1))}$ of \mathfrak{B}_1 are both one dimensional, and their elements are eigenvectors for $\overline{\text{ART}}_1$.

Consider the simple module $S^{\bar{\mu}}$ with $\bar{\mu} \vdash n + 1$. We know classically from the branching rules of the hyperoctahedral group (Theorem 4.2.26) that the restriction of this module to \mathfrak{B}_n is given by

$$\text{Res}_{\mathfrak{B}_n}^{\mathfrak{B}_{n+1}}(S^{\bar{\mu}}) \cong \bigoplus_{\substack{\bar{\lambda} \vdash n \\ \bar{\lambda} \subset \bar{\mu}}} S^{\bar{\lambda}}.$$

Now suppose we have a basis of eigenvectors for every $S^{\bar{\lambda}}$. By Lemma 4.3.16 the map $\kappa^{\bar{\lambda}, \bar{\mu}}(S^{\bar{\lambda}})$ gives a basis for the submodule $S^{\bar{\lambda}}$ inside of the vector space of $\text{Res}_{\mathfrak{S}_n}^{\mathfrak{S}_{n+1}}(S^{\bar{\mu}})$ which vector space is isomorphic to the vector space $S^{\bar{\mu}}$. Hence, considering all of the lifted eigenvectors from every $S^{\bar{\lambda}}$ together we find a basis for $S^{\bar{\mu}}$. By Theorems 4.3.19 and 4.3.20 the lifted eigenvectors form a basis of eigenvectors for $S^{\bar{\mu}}$. \square

Proof of Theorem 4.3.2. We have found a basis of eigenvectors of the random transposition shuffle for every Specht module $S^{\bar{\lambda}}$. Each eigenvector in this basis may be constructed by repeated applications of Theorems 4.3.19 and 4.3.20, starting at the empty word ω belonging to $S^{((0),(0))}$, and ending at the eigenvector belonging to $S^{\bar{\lambda}}$. When we add box e_i^1 in the process of forming $\bar{\lambda}$ the change in eigenvalue is $2 + 4(\lambda_i^1 + 1 - i)$, the second coefficient here is the diagonal index of the box $(i, \lambda_i^1 + 1)$ added to λ^1 . Similarly when we add box e_i^2 the change in eigenvalue is the change in eigenvalue is $4(\lambda_i^2 + 1 - i)$, which is the diagonal index of the box $(i, \lambda_i^2 + 1)$ added to λ^2 . Both of these changes are independent of when in the lifting we choose to add each box. Therefore, building up $\bar{\lambda}$ in any order produces an eigenvalue $2|\lambda^1| + 4\text{Diag}(\lambda^1) + 4\text{Diag}(\lambda^2)$ with multiplicity $d_{\bar{\lambda}}^2$. Normalising the eigenvalue by $2n^2$ gives the expression (4.8). \square

If we choose to ignore the second partition λ^2 when we lift eigenvalues we recover exactly the eigenvalues for the random transposition shuffle on S_n . That is, for a partition $\lambda \vdash n$, the bi-partition $(\lambda, (0))$ has corresponding eigenvalue,

$$\text{eig}((\lambda, (0))) = \frac{1}{n^2} (n + 2\text{Diag}(\lambda)).$$

as seen in Lemma 3.2.25. We end the lifting of the random transposition shuffle by providing an explicit description of the lifting operators $\kappa_i^{\bar{\lambda}, \bar{\lambda} + e_i^k, k}$, this result follows from work of Dieker and Saliola [19].

Lemma 4.3.22 (Theorem 21 [19]). *Let $\lambda \vdash n$. The lifting maps as defined in Definition 4.3.14 are a linear combination of the shuffling and switching operators, explicitly they are:*

$$\kappa_i^{\bar{\lambda}, \bar{\lambda} + e_i^k, k} = \sum_{1 \leq b_1 < \dots < b_m < b_{m+1} = i} \left(\prod_{j=1}^m \frac{1}{k \left((\lambda_i^k - i) - (\lambda_{b_j}^k - b_j) \right)} \Theta_{b_j, b_{j+1}}^k \right) \Phi_{b_1}^k. \quad (4.40)$$

Example 4.3.23. Let $\bar{\lambda} = ((1), (1))$, its associated Specht module is 2-dimensional and spanned by $S^{\bar{\lambda}} = \langle 11^+ - 11^-, 1^+1 - 1^-1 \rangle$. Both these basis vectors of $S^{\bar{\lambda}}$ are eigenvectors for $\overline{\text{ART}}_2$. Take $w = 11^+ - 11^-$ using the formulae given in Lemma 4.3.22 we lift w to an eigenvector of $\overline{\text{ART}}_3$ belonging to the Specht modules $S^{((1,1),(1))}$ and $S^{((1),(1,1))}$ by adding the box e_2^1, e_2^2 respectively.

$$\begin{aligned} \kappa_2^{\bar{\lambda}, \bar{\lambda} + e_2^1, 1}(w) &= \left(\Phi_2^1 - \frac{1}{2} \Theta_{1,2}^1 \Phi_1^1 \right) (w) = \frac{1}{2} [(11^+2 - 21^+1) - (11^-2 - 21^-1)] \\ \kappa_2^{\bar{\lambda}, \bar{\lambda} + e_2^2, 2}(w) &= \left(\Phi_2^2 - \frac{1}{4} \Theta_{1,2}^2 \Phi_1^2 \right) (w) \\ &= \frac{1}{2} [(11^+2^+ - 12^+1^+) + (11^-2^- - 12^-1^-) - (11^+2^- - 12^-1^+) - (11^-2^+ - 12^+1^-)]. \end{aligned}$$

4.4 One-sided Transposition Shuffles on The Hyperoctahedral Group

The one-sided transposition shuffle for the hyperoctahedral group is described by the following procedure: apply a transposition chosen according to the one-sided transposition shuffle for S_n , i.e. $\text{OST}_{n,w}$, then flip a fair coin; if heads do nothing, if tails flips the cards that were moved this step to their opposite sides.

Definition 4.4.1. The *biased one-sided transposition shuffle* for B_n with bias $w(j)$, denoted $\overline{\text{OST}}_{n,w}$, is driven by the following probability distribution:

$$\overline{\text{OST}}_{n,w}(\sigma) = \begin{cases} \frac{\sum_{j=1}^n w(j)/j}{2N_w} & \text{if } \sigma = e \\ \frac{w(j)}{N_w} \frac{1}{2j} & \text{if } \sigma = \xi_j \text{ for } j \in [n] \\ \frac{w(j)}{N_w} \frac{1}{2j} & \text{if } \sigma = (ij) \text{ for } i, j \in [n] \text{ with } i < j \\ \frac{w(j)}{N_w} \frac{1}{2j} & \text{if } \sigma = \xi_i \xi_j (ij) \text{ for } i, j \in [n] \text{ with } i < j \\ 0 & \text{otherwise} \end{cases} . \quad (4.41)$$

The one-sided transposition shuffle is a transitive, aperiodic, and reversible random walk on B_n , and is not constant on the conjugacy classes of B_n . The method of lifting eigenvectors for the random transposition shuffle on B_n may be modified to be applicable for the biased one-sided transposition shuffle on B_n . In Section 4.4.1 we prove the following result.

Theorem 4.4.2. *The eigenvalues for the biased one-sided transposition shuffle $\overline{\text{OST}}_{n,w}$ are indexed by standard Young tableaux of shape $\bar{\lambda} \vdash n$, and the eigenvalue represented by a tableau of shape $\bar{\lambda}$ has multiplicity $d_{\bar{\lambda}}$. For a standard Young tableau \bar{T} of shape $\bar{\lambda}$ the eigenvalue corresponding to \bar{T} is given by*

$$\text{eig}(\bar{T}) = \frac{1}{N_w} \left(\sum_{(i,j,1) \in \bar{T}} \frac{j-i+1}{T^1(i,j)} w(T^1(i,j)) + \sum_{(i,j,2) \in \bar{T}} \frac{j-i}{T^2(i,j)} w(T^2(i,j)) \right).$$

If we focus on bi-tableau $\bar{T} = (T^1, T^2)$ with $T^2 = \emptyset$ (those we can form by lifting only the first partition) then we recover the eigenvalues for the biased one-sided transposition shuffle on S_n (Lemma 3.6.5). We once again restrict our attention to weight functions of the form $w(j) = j^\alpha$, and define the biased one-sided transposition shuffle with this weight function as $\overline{\text{OST}}_{n,\alpha}$. Recall that the time $t_{n,\alpha}$ was defined as,

$$t_{n,\alpha} = \begin{cases} N_\alpha(n)/n^\alpha & \text{if } \alpha \leq 1 \\ N_\alpha(n)/N_{\alpha-1}(n) & \text{if } \alpha \geq 1 \end{cases} .$$

Applying the same projection argument used in Lemma 4.3.3 we may establish a lower bound of $t_{n,\alpha}(\log n)$ on the mixing time of the biased one-sided transposition shuffle $\overline{\text{OST}}_{n,\alpha}$. Using the eigenvalues established in Theorem 4.4.2 we can compute an upper bound on the total variation distance between $\overline{\text{OST}}_{n,\alpha}$ and π_n . We suspect that any analysis of this bound is likely to encounter similar problems to those described in Section

4.3.1 for the random transposition shuffle. We conjecture that it is possible to find a matching upper bound on the mixing time of $\overline{\text{OST}}_{n,\alpha}$ for all α .

Lemma 4.4.3. *The biased one-sided transposition shuffle on the hyperoctahedral group $\overline{\text{OST}}_{n,\alpha}$ satisfies the following bound for any $c > \max(2, 3 - \alpha)$:*

$$\liminf_{n \rightarrow \infty} \|\overline{\text{OST}}_{n,\alpha}^{t_{n,\alpha}(\log n - \log \log n - c)} - \pi_n\|_{\text{TV}} \geq \begin{cases} 1 - \frac{\pi^2}{6(c-3+\alpha)^2} & \text{if } \alpha \leq 1 \\ 1 - \frac{\pi^2}{6(c-2)^2} & \text{if } \alpha \geq 1 \end{cases}$$

Conjecture 4.4.4. *The biased one-sided transposition shuffle on the hyperoctahedral group $\overline{\text{OST}}_{n,\alpha}$ satisfies the following bound:*

$$\lim_{c \rightarrow \infty} \limsup_{n \rightarrow \infty} \|\overline{\text{OST}}_{n,\alpha}^{t_{n,\alpha}(\log n + c)} - \pi_n\|_{\text{TV}} = 0$$

Thus, the biased one-sided transposition shuffle exhibits a cutoff in total variation distance at time $t_{n,\alpha} \log n$ for all α .

4.4.1 Lifting Eigenvectors for One-sided Transposition Shuffles

The eigenvectors of the one-sided transposition shuffle $\overline{\text{OST}}_{n,w}$ may be computed using the same lifting operators $\kappa_a^{\bar{\lambda}, \bar{\mu}, k}$ as the random transposition shuffle. The difference in the lifting is reflected by changes to Theorems 4.3.10 and 4.3.19. We begin by transforming $\overline{\text{OST}}_{n,w}$ into an element of the group algebra \mathfrak{B}_n .

Definition 4.4.5. The one-sided transposition shuffle on B_n may be viewed as the following element of the group algebra \mathfrak{B}_n ,

$$\overline{\text{AOST}}_{n,w} = \sum_{1 \leq j \leq n} \frac{w(j)}{2j} e + \sum_{1 \leq j \leq n} \frac{w(j)}{2j} \xi_i + \sum_{1 \leq i < j \leq n} \frac{w(j)}{2j} (ij) + \sum_{1 \leq i < j \leq n} \frac{w(j)}{2j} \xi_i \xi_j (ij). \quad (4.42)$$

Note that above we have scaled our probability by $N_w(n)$.

The one-sided transposition shuffle on B_n has a similar recursive structure to the shuffle on S_n . Taking the difference of $\overline{\text{AOST}}_{n+1,w}$ and $\overline{\text{AOST}}_{n,w}$ we find

$$\overline{\text{AOST}}_{n+1,w} - \overline{\text{AOST}}_{n,w} = \frac{w(n+1)}{2(n+1)} (e + \xi_{n+1}) + \sum_{1 \leq i < n} \frac{w(n+1)}{2(n+1)} ((in+1) + \xi_i \xi_{n+1} (in+1))$$

which only depends the movement of the new card $n+1$ (compare this with equation (3.66)). Using this relationship we recover a new version of Theorem 4.3.10 for the one-sided transposition shuffle.

Theorem 4.4.6. *Let $\bar{\lambda} \vdash n$ for $n \in \mathbb{N}$, for words in $M^{\bar{\lambda}}$ we have the following equalities:*

$$\begin{aligned} \overline{\text{AOST}}_{n+1,w} \circ \Phi_a^1 - \Phi_a^1 \circ \overline{\text{AOST}}_{n,w} &= \frac{w(n+1)}{n+1} \left(\Phi_a^1 + \sum_{1 \leq b \leq n} \Phi_b^1 \circ \Theta_{b,a}^1 + \frac{1}{2} \sum_{1 \leq b \leq n} (\Phi_b^+ + \Phi_b^-) \circ \Theta_{b^\pm, a} \right). \\ \overline{\text{AOST}}_{n+1,w} \circ \Phi_a^2 - \Phi_a^2 \circ \overline{\text{AOST}}_{n,w} &= \frac{w(n+1)}{2(n+1)} \sum_{1 \leq b \leq n} \Phi_b^2 \circ \Theta_{b,a}^2. \end{aligned}$$

Proof. This follows the same proof as Theorem 4.3.10 with changes in constants to reflect equation (4.42). \square

Following from Theorem 4.4.6 the work of Section 4.3.2 may be replicated for the one-sided transposition shuffle. We summarise the results of the lifting in the following theorems whose proofs follow from those of the random transposition shuffle with changes in coefficients from Theorem 4.4.6.

Theorem 4.4.7 (Lifting for λ^1). *Let $\bar{\lambda} \vdash n$, and $a \in \{1, 2, \dots, l(\lambda^1) + 1\}$. Take $i \in [n]$ such that $1 \leq i \leq a$ and set $\bar{\mu} = \bar{\lambda} + e_i^1$. Then,*

$$\overline{\text{AOST}}_{n+1,w} \circ \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} - \kappa_a^{\bar{\lambda}, \bar{\mu}, 1} \circ \overline{\text{AOST}}_{n,w} = \frac{w(n+1)(1 + (\lambda_i^1 + 1 - i))}{n+1} \kappa_a^{\bar{\lambda}, \bar{\mu}, 1}$$

Thus if we have an eigenvector $v \in S^{\bar{\lambda}}$ with eigenvalue ϵ , we find $\kappa_a^{\bar{\lambda}, \bar{\mu}, 1}(v)$ to be an eigenvector of $S^{\bar{\mu}}$ with eigenvalue $\epsilon + \frac{w(n+1)(1 + (\lambda_i^1 + 1 - i))}{n+1}$.

Theorem 4.4.8 (Lifting for λ^2). *Let $\bar{\lambda} \vdash n$, and $a \in \{1, 2, \dots, l(\lambda^2) + 1\}$. Take $i \in [n]$ such that $1 \leq i \leq a$ and set $\bar{\mu} = \bar{\lambda} + e_i^2$. Then,*

$$\overline{\text{AOST}}_{n+1,w} \circ \kappa_a^{\bar{\lambda}, \bar{\mu}, 2} - \kappa_a^{\bar{\lambda}, \bar{\mu}, 2} \circ \overline{\text{AOST}}_{n,w} = \frac{w(n+1)(\lambda_i^2 + 1 - i)}{n+1} \kappa_a^{\bar{\lambda}, \bar{\mu}, 2}$$

Thus if we have an eigenvector $v \in S^{\bar{\lambda}}$ with eigenvalue ϵ , we find $\kappa_a^{\bar{\lambda}, \bar{\mu}, 2}(v)$ to be an eigenvector of $S^{\bar{\mu}}$ with eigenvalue $\epsilon + \frac{w(n+1)(\lambda_i^2 + 1 - i)}{n+1}$.

Note that Lemma 4.3.21 still holds for the one-sided transposition shuffle because the lifting operators have not changed. This leads us to a proof of Theorem 4.4.2

Proof of Theorem 4.4.2. For any bi-tableau \bar{T} of shape $\bar{\lambda}$ we build up a distinct eigenvector belonging to $S^{\bar{\lambda}}$ by lifting from partition $((0), (0))$ to $\bar{\lambda}$ in the order specified by the tableau \bar{T} . Summing the changes in eigenvalue given by Theorems 4.4.7 and 4.4.8 and normalising by $N_w(n)$ we recover the expression for the eigenvalue corresponding to \bar{T} given in the theorem. \square

4.4.2 A Strong Stationary Time for the Unbiased One-sided Transposition Shuffle on The Hyperoctahedral Group.

In this section we prove that the unbiased one-sided transposition shuffle on B_n exhibits a cutoff in separation distance at time $n \log n$. To establish a cutoff in separation distance we prove that the strong stationary time introduced in Section 3.7 is also a strong stationary time for the unbiased one-sided transposition shuffle on B_n . We record our main results below before constructing the strong stationary time in detail.

Theorem 4.4.9. *There exists a strong stationary time T for the unbiased one-sided transposition shuffle for B_n . Furthermore, for $c > 0$ we have $\mathbb{P}(T > n \log n + cn) \leq e^{-c}$.*

From Theorem 4.4.9 we may quickly establish that the one-sided transposition shuffle exhibits a cutoff in total variation distance and separation distance at time $n \log n$. This proves Conjecture 4.4.4 for the case $\alpha = 0$.

Theorem 4.4.10. *The unbiased one-sided transposition shuffle $\overline{\text{OST}}_{n,0}$ satisfies the following bounds, for $c_1 > 0, c_2 > 2$:*

$$\limsup_{n \rightarrow \infty} \|\overline{\text{OST}}_{n,0}^{n \log n + c_1 n} - \pi_n\|_{\text{sep}} \leq e^{-c_1} \quad (4.43)$$

$$\liminf_{n \rightarrow \infty} \|\overline{\text{OST}}_{n,0}^{n \log n - n \log \log n - c_2 n} - \pi_n\|_{\text{TV}} \geq 1 - \frac{\pi^2}{6(c_2 - 2)^2} \quad (4.44)$$

Thus, unbiased one-sided transposition shuffle exhibits a cutoff in total variation distance and separation distance at time $n \log n$.

Proof of Theorem 4.4.10. The lower limit was given in Lemma 4.4.3. The upper limit follows from Theorem 4.4.9, using the strong stationary time T , we see that,

$$\limsup_{n \rightarrow \infty} \|\overline{\text{OST}}_{n,0}^{n \log n + c_1 n} - \pi_n\|_{\text{sep}} \leq \limsup_{n \rightarrow \infty} \mathbb{P}(T > n \log n + c_1 n) \leq e^{-c_1}.$$

Separation distance gives an upper bound on total variation distance (Lemma 1.1.24), therefore combining the two bounds establishes a cutoff in both total variation distance and separation distance. \square

A Strong Stationary Time Argument

Recall that we may view the elements of B_n as a subset of permutations of a deck of cards made up of cards and positions both indexed by $[\pm n]$. Any permutation $\sigma \in B_n$ is a bijection from cards to positions and σ^{-1} is a bijection from positions to cards, i.e., $\sigma(i)$ tells us the position of card i whereas $\sigma^{-1}(i)$ tells us what card is in position i . We begin all our random walks at the identity permutation with positions and labels fully matched. Throughout the rest of this section let $(X^t)_{t \in \mathbb{N}}$ denote a Markov chain on B_n driven by the unbiased one-sided transposition shuffle, and let $(Y^t)_{t \in \mathbb{N}}$ be a Markov chain on B_n defined by setting $Y^t = (X^t)^{-1}$ for all t . The Markov chains (X^t) and (Y^t) represent two different ways to view the unbiased one-sided transposition shuffle.

Let τ^t be the transposition chosen at step t of the unbiased one-sided transposition shuffle. To construct our strong stationary time we need to condition on the exact permutation of cards in positions (both positive and negative) above position j at time t , that is the random variables $Y^t(\pm i)$ for $j < i \leq n$. Note that if we know the random variable $Y^t(i)$ we may compute $Y^t(-i) = -Y^t(i)$. Given this information we also know which cards can be in positions $[\pm j]$ at time t , define this set as,

$$A_j^t = [\pm n] \setminus \{Y^t(\pm i) \mid j < i \leq n\}.$$

Definition 4.4.11. We say the random walk $(Y^t)_{t \in \mathbb{N}}$ satisfies *property \mathcal{P}_j at time t* if we have:

$$\mathbb{P}(Y^t(j) = l \mid Y^t(i) \text{ for all } j < i \leq n) = \begin{cases} 1/2j & \text{if } l \in A_j^t \\ 0 & \text{otherwise} \end{cases} \quad (4.45)$$

This property tells us that given total information about the deck strictly above position j , the card in position j is equally likely to be any of the remaining cards.

Lemma 4.4.12. *Let T_j be the first time our right hand chooses the position j when performing the unbiased one-sided transposition shuffle. If $T_j \leq t$ then the Markov chain $(Y^t)_{t \in \mathbb{N}}$ satisfies property \mathcal{P}_j at time t .*

Proof. We prove this by induction: once property \mathcal{P}_j holds for some t , it holds for all times after t .

Consider the time T_j , at this step of our Markov chain we must have applied a transposition (ij) or $\xi_i \xi_j(ij)$ with $i \leq j$. The probability of picking any one of the transpositions (ij) at time T_j is $\mathbb{P}(\tau^{T_j} = (ij)) = \mathbb{P}(\tau^{T_j} = \xi_i \xi_j(ij)) = 1/2j$ for all $i \leq j$. Therefore, the card in position j at time T_j has a uniform chance of being any of the cards in $A_j^{T_j-1} = A_j^{T_j}$ (which contains both positive and negative cards). Thus, we may clearly see that,

$$\mathbb{P}(Y^{T_j}(j) = l | Y^{T_j}(i) \text{ for all } j < i \leq n) = \begin{cases} 1/2j & \text{if } l \in A_j^{T_j} \\ 0 & \text{otherwise} \end{cases} \quad (4.46)$$

so \mathcal{P}_j holds at time T_j .

Now suppose property \mathcal{P}_j holds at time t . We study the time $t+1$ and split the analysis into cases based on which permutation was applied at time $t+1$,

$$\begin{aligned} & \mathbb{P}(Y^{t+1}(j) = l | Y^{t+1}(i) \text{ for all } j < i \leq n) \\ &= \sum_{\sigma \in B_n} \mathbb{P}(\tau^{t+1} = \sigma) \mathbb{P}(Y^{t+1}(j) = l | Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = \sigma). \end{aligned} \quad (4.47)$$

Using knowledge of the transposition τ^{t+1} we evolve our deck backwards in time ($X^t = \tau^{t+1} X^{t+1}$) to recover the random variables $Y^t(i)$ from $Y^{t+1}(i)$ for $j < i \leq n$, and relate $Y^{t+1}(j)$ to $Y^t(j)$. This allows us to use our inductive hypothesis.

If $\tau^{t+1} \in \{(aj), \xi_a \xi_j(aj) | a \leq j\}$ then our random walk satisfies property \mathcal{P}_j at time $t+1$ for the same reasoning as time T_j . Suppose that $\tau^{t+1} \in \{(ab), \xi_a \xi_b(ab) | a, b < j\}$. Then we know that $Y^t(i) = Y^{t+1}(i)$ for all $j < i \leq n$. Suppose instead that $\tau^{t+1} \in \{(ab), \xi_a \xi_b(ab) | j < a, b\}$, then we have $Y^t(\pm b) = Y^{t+1}(\pm a)$ and $Y^t(\pm a) = Y^{t+1}(\pm b)$ (with the signs swapping in the second case), with $Y^t(i) = Y^{t+1}(i)$ for all other $j < i \leq n$. In either case we know $A_j^{t+1} = A_j^t$, and the card in position j has not moved from time t to $t+1$. Therefore, we have

$$\begin{aligned} & \mathbb{P}(Y^{t+1}(j) = l | Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} \in \{(ab), \xi_a \xi_b(ab) | a, b < j \text{ or both } b, a > j\}) \\ &= \mathbb{P}(Y^t(j) = l | Y^t(i) \text{ for all } j < i \leq n) = \begin{cases} 1/2j & \text{if } l \in A_j^t = A_j^{t+1} \\ 0 & \text{otherwise} \end{cases}. \end{aligned}$$

We now study in detail the effects of the remaining transpositions (ab) and $\xi_a, \xi_b(ab)$ with $a \leq j < b$. In this case we can not fully recover the random variables $Y^t(i)$ with $i < j \leq n$ without extra assumptions. To this end fix $b > j$, a card $C \in A_j^{t+1}$, and suppose that $\tau^{t+1} = (X^{t+1}(C) b)$, i.e., card C is moved from position b into a position below j by τ^{t+1} . Letting C range over all choices in A_j^{t+1} will recover every transposition (ab) and $\xi_a \xi_b(ab)$ with $b > j$ fixed and $a \leq j$. In the case that $\tau^{t+1} = (X^{t+1}(C) b)$, we know that $Y^t(b) = C$ and the other positions above j have $Y^t(i) = Y^{t+1}(i)$ for $j < i \leq n$ and $i \neq b$. Therefore, for this choice of τ^{t+1} we know that $A_j^t = (A_j^{t+1} \sqcup \{Y^{t+1}(\pm b)\}) \setminus \{\pm C\}$. Now

consider the probability:

$$\mathbb{P}\left(Y^{t+1}(j) = l \mid Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = (X^{t+1}(C) \ b)\right).$$

If $l = C$ then the event in question can only occur if the card currently in position b , i.e. $Y^{t+1}(b)$, was in position j at time t . Noting that $Y^{t+1}(b) \in A_j^t$, and using our inductive hypothesis we find,

$$\begin{aligned} & \mathbb{P}\left(Y^{t+1}(j) = C \mid Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = (X^{t+1}(C) \ b)\right) \\ &= \mathbb{P}\left(Y^t(j) = Y^{t+1}(b) \mid Y^t(i) \text{ for all } j < i \leq n\right) = 1/2j \end{aligned} \quad (4.48)$$

If $l = -C$ then the event in question can only occur if the card currently in position $-b$, i.e. $Y^{t+1}(-b)$, was in position j at time t . Noting that $Y^{t+1}(-b) \in A_j^t$, and using our inductive hypothesis we find,

$$\begin{aligned} & \mathbb{P}\left(Y^{t+1}(j) = -C \mid Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = (X^{t+1}(C) \ b)\right) \\ &= \mathbb{P}\left(Y^t(j) = Y^{t+1}(-b) \mid Y^t(i) \text{ for all } j < i \leq n\right) = 1/2j \end{aligned} \quad (4.49)$$

Alternatively, suppose $l \in A_j^{t+1} \setminus \{\pm C\}$, we know the card l does not move from its position at time t to time $t+1$, and we know $l \in A_j^t$. Therefore, we find

$$\begin{aligned} & \mathbb{P}\left(Y^{t+1}(j) = l \mid Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = (X^{t+1}(C) \ b)\right) \\ &= \mathbb{P}\left(Y^t(j) = l \mid Y^t(i) \text{ for all } j < i \leq n\right) = 1/2j \end{aligned} \quad (4.50)$$

Putting the equations above together gives us,

$$\mathbb{P}\left(Y^{t+1}(j) = l \mid Y^{t+1}(i) \text{ for all } j < i \leq n, \tau^{t+1} = (X^{t+1}(C) \ b)\right) = \begin{cases} 1/2j & \text{if } l \in A_j^{t+1} \setminus \{\pm C\} \\ 1/2j & \text{if } l \in \{\pm C\} \\ 0 & \text{otherwise} \end{cases}.$$

Letting C range over all possible choices of card in A_j^{t+1} while keeping $b > j$ fixed, we cover the desired probability for all transpositions (ab) , $\xi_a \xi_b(ab)$ with $a \leq j < b$. Finally applying every separate case to (4.47) we have established that

$$\mathbb{P}\left(Y^{t+1}(j) = l \mid Y^{t+1}(i) \text{ for all } j < i \leq n\right) = \begin{cases} 1/2j & \text{if } l \in A_j^{t+1} \\ 0 & \text{otherwise} \end{cases}$$

as required, thus by induction our hypothesis holds for all $t \geq T_j$. □

Lemma 4.4.13. *Let $T = \min\{t \geq 0 \mid t \geq T_1, \dots, T_n\}$ be the first time our right hand has chosen every position j . Then T is a strong uniform time for $(X^t)_{t \in \mathbb{N}}$ and $(Y^t)_{t \in \mathbb{N}}$.*

Proof. Note that X^t is uniformly distributed if and only if Y^t is uniformly distributed. Lemma 4.4.12 implies that by time T our Markov chain $(Y^t)_{t \in \mathbb{N}}$ satisfies all properties \mathcal{P}_j .

Hence, we have

$$\begin{aligned}
\mathbb{P}(Y^t = \sigma^{-1} | T \leq t) &= \mathbb{P}(\cap_{j=1}^n \{Y^t(j) = \sigma^{-1}(j)\} | T \leq t) \\
&= \prod_{j=1}^n \mathbb{P}(Y^t(j) = \sigma^{-1}(j) | \cap_{i=j+1}^n \{Y^t(i) = \sigma^{-1}(i)\}, T \leq t) \\
&= \prod_{j=1}^n \frac{1}{2j} = \frac{1}{2^n n!} = \pi_n(\sigma).
\end{aligned}$$

□

We have found a strong stationary time for the unbiased one-sided transposition shuffle. Following quickly from this we may prove Theorem 4.4.9, and thus establish a cutoff in separation distance for the one-sided transposition shuffle.

Proof of Theorem 4.4.9. Let T be the first time our right hand has chosen every position j . Our right hand is choosing positions via a uniform probability on $[n]$. Thus T is modelled by the uniform coupon collectors problem with n coupons. To complete our argument recall (Section 2.3.2, equation (2.19)) that for the uniform coupon collectors problem on n cards we have $\mathbb{P}(T > n \log n + cn) \leq e^{-c}$. □

The strong stationary time we have constructed in this section is actually a strong stationary time for any biased one-sided transposition shuffle. Let T_α be the first time our right hand has chosen every position j following the biased one-sided transposition shuffle $\overline{\text{OST}}_{n,\alpha}$. In Lemma 3.7.7 we showed that for $\alpha \leq 0$ we have the bound $\mathbb{P}(T_\alpha > t_{n,\alpha}(\log n + c)) \leq e^{-c}$. Therefore, extending the proof of Theorem 4.4.10 we are able to prove the following.

Corollary 4.4.14. *The biased one-sided transposition shuffle $\overline{\text{OST}}_{n,\alpha}$ with $\alpha \leq 0$ exhibits a cutoff in separation distance and total variation distance at time $t_{n,\alpha} \log n$. Thus, conjecture 4.4.4 holds for all $\alpha \leq 0$.*

4.5 Further Work

In this chapter we have explored the hyperoctahedral group as an extension of the symmetric group. The hyperoctahedral group is just one instance of the generalised symmetric group.

Definition 4.5.1. Let $r, n \in \mathbb{N}$, and let ξ be a r^{th} root of unity. The *generalised symmetric group*, denoted $G_{r,n}$, is defined as the group of all bijections σ on elements $\{\xi^k i | i \in [n], k \in [r]\}$ such that $\sigma(\xi^k i) = \xi^k \sigma(i)$. The generalised symmetric group $G_{r,n}$ is isomorphic to the wreath product $\mathbb{Z}_r \wr S_n$.

The symmetric and hyperoctahedral groups are isomorphic to the generalised symmetric groups $G_{1,n}$ and $G_{2,n}$ respectively. We were able to describe the module structure of the symmetric group and the hyperoctahedral group using a similar set of techniques in both cases. The arguments presented in Sections 2.2 and 4.2 can be modified to work for the generalised symmetric group $G_{r,n}$ for any $r \in \mathbb{N}$. Below we describe the module structure of $G_{r,n}$ over the field \mathbb{C} , a full construction of the permutation and Specht modules for any choice of $r \in \mathbb{N}$ may be found in [12].

Definition 4.5.2. Let $r, n \in \mathbb{N}$. An r -partition of n , denoted $\bar{\lambda}$, is a tuple of partitions $\bar{\lambda} = (\lambda^1, \dots, \lambda^r)$ such that $\sum_{k=1}^r |\lambda^k| = n$. Let $\bar{\lambda}, \bar{\mu}$ be r -partitions. Define the *dominance ordering* on r -partitions as follows:

$$\bar{\lambda} \succeq \bar{\mu} \Leftrightarrow \begin{cases} |\lambda^1| > |\mu^1| \\ \text{or } |\lambda^k| = |\mu^k| \text{ for } 1 \leq k \leq i, \text{ and } |\lambda^{i+1}| > |\mu^{i+1}| \text{ for some } i \in \{1, \dots, r-2\}. \\ \text{or } |\lambda^k| = |\mu^k| \text{ and } \lambda^k \succeq \mu^k \text{ for all } k \in [r] \end{cases}$$

Lemma 4.5.3. *The permutation and simple modules for the generalised symmetric group $G_{r,n}$ are indexed by r -partitions of n . The permutation module and simple module corresponding to a r -partition $\bar{\lambda}$ are denoted $M^{\bar{\lambda}}$ and $S^{\bar{\lambda}}$ respectively. Furthermore, the permutation and simple modules respect Young's rule, that is*

$$M^{\bar{\mu}} \cong \bigoplus_{\bar{\lambda} \succeq \bar{\mu}} K_{\bar{\lambda}, \bar{\mu}} S^{\bar{\lambda}},$$

for constants $K_{\bar{\lambda}, \bar{\mu}} \in \mathbb{N}^0$.

Conjecture 4.5.4. *Let $n \geq 2$, and $\bar{\lambda} \vdash n$. The branching rules for the simple modules of $G_{r,n}$ are as follows:*

$$\text{Res}_{G_{r,n-1}}^{G_{r,n}} S^{\bar{\lambda}} \cong \bigoplus_{\substack{\bar{\mu} \vdash n-1 \\ \bar{\mu} \subseteq \bar{\lambda}}} S^{\bar{\mu}} \text{ as } G_{r,n-1}\text{-modules} \quad (4.51)$$

$$\text{Ind}_{G_{r,n}}^{G_{r,n+1}} S^{\bar{\lambda}} \cong \bigoplus_{\substack{\bar{\mu} \vdash n+1 \\ \bar{\lambda} \subseteq \bar{\mu}}} S^{\bar{\mu}} \text{ as } G_{r,n+1}\text{-modules.} \quad (4.52)$$

The structure of $G_{r,n}$ allows us to generalise the random transposition shuffle and the one-sided transposition shuffle for any choice of r . To do this we first need to define the negative transpositions ξ_i for the group $G_{r,n}$. Let ξ be an r^{th} -root of unity, and define the negative transpositions $\xi_i \in G_{r,n}$ for $i \in [n]$ as follows:

$$\xi_i(j) = \begin{cases} \xi i & \text{if } j = i \\ j & \text{otherwise} \end{cases}.$$

The random transposition shuffle on the group $G_{r,n}$ is defined by the following procedure: pick a transposition (ij) of $S_n \subseteq G_{r,n}$ via the random transposition shuffle and apply it, then uniformly at random choose an element $k \in [r]$ and apply the negative transpositions ξ_i^k, ξ_j^k to the positions i, j . Overall we have applied an element of $G_{n,r}$ of the form $\xi_i^k \xi_j^k (ij)$ or $\xi_j^k (jj)$. If $r = 1, 2$ we recover the random transposition shuffle on the symmetric group and hyperoctahedral group respectively which were described in Sections 2.3.1 and 4.3. Similarly we may form the biased one-sided transposition shuffles on the group $G_{r,n}$ by the following procedure: pick a transposition (ij) of S_n via the biased one-sided transposition shuffle and apply it, then uniformly at random choose an element $k \in [r]$ and apply the negative transpositions ξ_i^k, ξ_j^k to the positions i, j . It is natural to think how we might extend the technique of lifting eigenvectors to the random transposition and the one-sided transposition shuffles on the group $G_{r,n}$ for $r > 2$. In line with the $r = 2$ case we expect

there to be r different lifting operators, one for adding a box to each separate partition λ^k which together constitute an r -partition. Furthermore, we expect the eigenvectors for each shuffle to correspond to standard Young tableaux of r -partitions.

Conjecture 4.5.5. *The eigenvalues for the random transposition shuffle on the generalised symmetric group $G_{r,n}$ are labelled by r -partitions of n , and may be described by the technique of lifting eigenvectors.*

Conjecture 4.5.6. *The eigenvalues for the biased one-sided transposition shuffle on the generalised symmetric group $G_{r,n}$ are labelled by standard Young tableaux of r -partitions of n , and may be described by the technique of lifting eigenvectors.*

Once the eigenvalues for each shuffle have been found they could be used to analyse the mixing time of each shuffle on the generalised symmetric group $G_{r,n}$. We conjecture that the random transposition shuffle and the one-sided transposition shuffles on the group $G_{r,n}$ exhibit a cutoff in total variation distance at the same time as their symmetric group counterparts.

Conjecture 4.5.7. *The random transposition shuffle on the generalised symmetric group $G_{r,n}$ exhibits a cutoff in total variation distance at time $(n/2) \log n$.*

Conjecture 4.5.8. *The biased one-sided transposition shuffle on the generalised symmetric group $G_{r,n}$ with weight function w such that $w(j) = j^\alpha$ exhibits a cutoff in total variation distance at time $t_{n,\alpha} \log n$.*

In fact by using the strong stationary time introduced in Section 3.7 we may prove the above conjecture in the case of the unbiased one-sided transposition shuffle ($\alpha = 0$) on the group $G_{r,n}$.

Theorem 4.5.9. *The unbiased one-sided transposition shuffle on the generalised symmetric group $G_{r,n}$ exhibits a cutoff in total variation distance and separation distance at time $n \log n$.*

Sketch Proof. The proof of this Theorem follows from extending the techniques of Sections 3.7.1 and 4.4.2. Let ξ be an r^{th} root of unity. Let $(X^t)_{t \in \mathbb{N}}$ denote a Markov chain on $G_{r,n}$ driven by the unbiased one-sided transposition shuffle, and let $(Y^t)_{t \in \mathbb{N}}$ be a Markov chain on $G_{r,n}$ defined by setting $Y^t = (X^t)^{-1}$ for all t . To construct a strong stationary time we need to condition on exact knowledge of the positions above position j at time t , that is the random variables $Y^t(\xi^k i)$ for $j < i \leq n$ and $k \in [r]$. Note that if we know the random variable $Y^t(i)$ we may compute $Y^t(\xi^k i) = \xi^k Y^t(i)$ for any $k \in [r]$. Given this information we also know the value of positions $\xi^k 1$ to $\xi^k j$ at time t , define this set as,

$$A_j^t = \{\xi^k i \mid i \in [n], k \in [r]\} \setminus \{Y^t(\xi^k i) \mid j < i \leq n, k \in [r]\}.$$

We say our random walk satisfies *property \mathcal{P}_j at time t* if we have,

$$\mathbb{P}(Y^t(j) = l \mid Y^t(i) \text{ for all } j < i \leq n) = \begin{cases} 1/rj & \text{if } l \in A_j^t \\ 0 & \text{otherwise} \end{cases}. \quad (4.53)$$

Let T_j be the first time we apply a transposition $\xi_i^k \xi_j^k(ij)$ with $i \leq j$ and $k \in [r]$ following the unbiased one-sided transposition shuffle on $G_{r,n}$. Extending the arguments of Lemmas 3.7.4 and 4.4.12, we may see that at any time $t \geq T_j$ the Markov chain $(Y^t)_{t \in \mathbb{N}}$ satisfies property \mathcal{P}_j .

Let $T = \min\{t \geq 0 : t \geq T_j \text{ for all } j\}$, this is a stopping time for our random walk. At any time after T the random walk $(Y^t)_{t \in \mathbb{N}}$, satisfies all properties \mathcal{P}_j , therefore for any $\sigma \in G_{r,n}$ we have

$$\begin{aligned} \mathbb{P}(Y^t = \sigma^{-1} | T \leq t) &= \mathbb{P}(\cap_{j=1}^n \{Y^t(j) = \sigma^{-1}(j)\} | T \leq t) \\ &= \prod_{j=1}^n \mathbb{P}(Y^t(j) = \sigma^{-1}(j) | \cap_{i=j+1}^n \{Y^t(i) = \sigma^{-1}(i)\}, T \leq t) \\ &= \prod_{j=1}^n \frac{1}{rj} = \frac{1}{r^n n!} = \pi_n(\sigma). \end{aligned}$$

Hence, the time T is a strong stationary time for $(Y^t)_{t \in \mathbb{N}}$ and so $(X^t)_{t \in \mathbb{N}}$. Using the same coupon collector's argument as before we know that $\mathbb{P}(T > n \log n + cn) \leq e^{-c}$, hence we have an upper bound of $n \log n$ on the on separation distance mixing time of the unbiased one-sided transposition shuffle. To find a matching lower bound we reduce the analysis of the unbiased one-sided transposition shuffle on $G_{r,n}$ to the $r = 1$ (symmetric group case) using the same argument as Lemma 4.3.3. This gives a lower bound of $n \log n$ on the total variation distance mixing time of the unbiased one-sided transposition shuffle. Combining the two bounds above proves the existence of a cutoff in total variation distance and separation distance at time $n \log n$. \square

Bibliography

- [1] E Al-Aamily, AO Morris, and MH Peel. The representations of the weyl groups of type bn. *Journal of Algebra*, 68(2):298–305, 1981.
- [2] David Aldous and Persi Diaconis. Shuffling cards and stopping times. *The American Mathematical Monthly*, 93(5):333–348, 1986.
- [3] Michael E Bate, Stephen B Connor, and Oliver Matheau-Raven. Cutoff for a one-sided transposition shuffle. *Preprint, arXiv:1907.12074*, 2019.
- [4] Nathanaël Berestycki. The hyperbolic geometry of random transpositions. *The Annals of Probability*, 34(2):429–467, 2006.
- [5] Nathanaël Berestycki. Markov chain mixing times. <http://aimpl.org/markovmixing/5/>, 2016. AIM Workshop.
- [6] Nathanaël Berestycki and Rick Durrett. A phase transition in the random transposition random walk. *Probability theory and related fields*, 136(2):203–233, 2006.
- [7] Nathanaël Berestycki, Oded Schramm, and Ofer Zeitouni. Mixing Times for Random k-Cycles and Coalescence-Fragmentation Chains. *The Annals of Probability*, 39(5):1815–1843, 2011.
- [8] Nathanaël Berestycki and Richard Pymar. Effect of scale on long-range random graphs and chromosomal inversions. *The Annals of Applied Probability*, 22(4):1328–1361, 2012.
- [9] Megan Bernstein, Nayantara Bhatnagar, and Igor Pak. Cutoff for biased transpositions. *Preprint, arXiv: 1709.03477*, 2017.
- [10] Megan Bernstein and Evita Nestoridi. Cutoff for random to random card shuffle. *The Annals of Probability*, 47(5):3303–3320, 2019.
- [11] Andrei Z Broder. Weighted random mappings; properties and applications. PhD thesis, Stanford University, CA, Department of Computer Science, 1985.
- [12] Himmet Can. Representations of the generalized symmetric groups. *Beitrage zur Algebra und Geometrie*, 37(2):289–308, 1996.
- [13] Guan-Yu Chen. The cutoff phenomenon for finite Markov chains. PhD thesis, Cornell University, 2006.

- [14] Guan-Yu Chen and Laurent Saloff-Coste. The cutoff phenomenon for ergodic Markov processes. *Electronic Journal of Probability*, 13:26–78, 2008.
- [15] Amir Dembo, Yuval Peres, and David Revelle. Sharp thresholds for mixing times. <https://aimath.org/pastworkshops/mixingtimes.html>, 2004. ARCC Workshop Report.
- [16] P Diaconis. Group representations in probability and statistics. *Lecture Notes-Monograph Series*, 11:1–192, 1988.
- [17] Persi Diaconis and Laurent Saloff-Coste. Comparison techniques for random walk on finite groups. *The Annals of Probability*, pages 2131–2156, 1993.
- [18] Persi Diaconis and Mehrdad Shahshahani. Generating a random permutation with random transpositions. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 57(2):159–179, 1981.
- [19] AB Dieker and FV Saliola. Spectral analysis of random-to-random Markov chains. *Advances in Mathematics*, 323:427–485, 2018.
- [20] William Feller. An introduction to probability theory and its applications, volume 1. John Wiley & Sons, 2008.
- [21] Amaury Freslon, Lucas Teyssier, and Simeng Wang. Cutoff profiles for quantum \mathfrak{e} vy processes and quantum random transpositions. *Preprint, arXiv:2010.03273*, 2020.
- [22] Ladnor Geissinger and D Kinch. Representations of the hyperoctahedral groups. *J. Algebra*, 53(1):1–20, 1978.
- [23] Subhajit Ghosh. Total variation cutoff for the flip-transpose top with random shuffle. *Preprint, arXiv:1906.11544*, 2019.
- [24] Jonathan Hermon, Hubert Lacoïn, and Yuval Peres. Total variation and separation cutoffs are not equivalent and neither one implies the other. *Electronic Journal of Probability*, 21, 2016.
- [25] G. D. James. The irreducible representations of the symmetric groups. *Bulletin of the London Mathematical Society*, 8(3):229–232, 11 1976.
- [26] Gordon James and Adalbert Kerber. The representation theory of the symmetric group. *Encyclopedia of Mathematics and its Applications. Cambridge University Press*, 1984.
- [27] Gordon James, Martin W. Liebeck, and Martin Liebeck. Representations and characters of groups. *Cambridge University Press*, 2001.
- [28] Gordon D. James. The representation theory of the symmetric groups, volume 682. *Springer*, 1978.
- [29] Hubert Lacoïn. Mixing time and cutoff for the adjacent transposition shuffle and the simple exclusion. *The Annals of Probability*, 44(2):1426–1487, 2016.

- [30] Nadia Lafrenière. Eigenvalues of symmetrized shuffling operators. *Preprint, arXiv:1811.07196*, 2018.
- [31] Carlo Lancia. The cutoff phenomenon in diffusion processes. <https://www.slideshare.net/pilmotz/presentazione-1522809>, 2009. *Seminar given at University of Rome Tor Vergata*.
- [32] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. Markov chains and mixing times: second edition. *American Mathematical Society*, Providence, RI, 2008.
- [33] Peter Matthews. A strong uniform time for random transpositions. *Journal of Theoretical Probability*, 1(4):411–423, 1988.
- [34] Elchanan Mossel, Yuval Peres, and Alistair Sinclair. Shuffling by semi-random transpositions. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 572–581. IEEE, 2004.
- [35] Evita Nestoridi and Sam Thomas. Limit profiles for markov chains. *Preprint, arXiv:2005.13437*, 2020.
- [36] Ross G. Pinsky. Cyclic to random transposition shuffles. *Preprint, arXiv:1204.2081*, 2012.
- [37] Richard Pymar. Partial mixing of semi-random transposition shuffles. *Preprint, arXiv:1302.2601*, 2013.
- [38] Bruce E Sagan. The symmetric group: representations, combinatorial algorithms, and symmetric functions, volume 203. Springer Science & Business Media, 2013.
- [39] Laurent Saloff-Coste. Random walks on finite groups. *Probability on discrete structures*, pages 263–346, 2004.
- [40] Clyde H Schoolfield. Random walks on wreath products of groups. *Journal of Theoretical Probability*, 15(3):667–693, 2002.
- [41] Eliran Subag. A lower bound for the mixing time of the random-to-random insertions shuffle. *Electronic Journal of Probability*, 18, 2013.
- [42] Lucas Teyssier. Limit profile for random transpositions. *Annals of Probability*, 48(5):2323–2343, 2020.
- [43] Jay-Calvin Uyemura-Reyes. Random walk, semi-direct products, and card shuffling. PhD thesis, Stanford University, 2002.
- [44] Graham White. Combinatorial methods in Markov chain mixing. PhD thesis, Stanford University, 2017.
- [45] Graham White. A strong stationary time for random transpositions. *Preprint, arXiv:1910.00770*, 2019.