

Combinatorics of Tableau Inversions

Jonathan E. Beagley Paul Drube

Department of Mathematics and Statistics

Valparaiso University

Valparaiso, IN 46383-6493, U.S.A.

{jon.beagley, paul.drube}@valpo.edu

Submitted: Dec 21, 2014; Accepted: May 23, 2015; Published: Jun 3, 2015

Mathematics Subject Classifications: 05A19, 05A05

Abstract

A tableau inversion is a pair of entries in row-standard tableau T that lie in the same column of T yet lack the appropriate relative ordering to make T column-standard. An i -inverted Young tableau is a row-standard tableau along with precisely i inversion pairs. Tableau inversions were originally introduced by Fresse to calculate the Betti numbers of Springer fibers in Type A, with the number of i -inverted tableaux that standardize to a fixed standard Young tableau corresponding to a specific Betti number of the associated fiber. In this paper we approach the topic of tableau inversions from a completely combinatorial perspective. We develop formulas enumerating the number of i -inverted Young tableaux for a variety of tableaux shapes, not restricting ourselves to inverted tableau that standardize a specific standard Young tableau, and construct bijections between i -inverted Young tableaux of a certain shape with j -inverted Young tableaux of different shapes. Finally, we share some the results of a computer program developed to calculate tableaux inversions.

Keywords: Young tableaux, inversions of Young tableaux

1 Introduction

Let $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_m)$ be a non-increasing sequence of positive integers that partition N . A **Young diagram** Y of shape λ is a left-justified array of $N = \lambda_1 + \dots + \lambda_m$ boxes such that there are λ_i boxes in the i^{th} row of Y . A standard filling of Y is a bijective assignment of the integers $1, 2, \dots, N$ to the boxes of Y , producing what is known as a tableau of shape λ . The resulting tableau is said to be **row-standard** if its entries are increasing from left-to-right along each row, and **column-standard** if its entries are increasing from top-to-bottom in each column. A **standard Young tableau** of shape λ

is a tableau of shape λ that is both row-standard and column-standard. For more basic information on Young Tableaux, see [4].

One established result about standard Young tableaux that we will make repeated use of is the Hook-Length formula. Let Y be a Young diagram of shape λ . To each box in Y one may assign a **hook-length** that equals one plus the number of boxes directly below or directly to the right of the given box (i.e.- the number of boxes in the “hook”-shaped sub-tableaux whose corner lies at the chosen box). As stated below, the Hook-Length formula uses these hook-lengths to determine the number of standard Young tableaux of shape λ . For more discussion of this result, see Chapter 7 of [6].

Theorem 1.1 (The Hook-Length Formula). *Let $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_m)$ be a non-increasing partition of N , and let h_{ij} denote the hook-length of the (i, j) -entry in a Young diagram of shape λ . Then the number of standard Young tableaux of shape λ is:*

$$|S(\lambda)| = \frac{N!}{\prod h_{ij}}$$

Now take a permutation $\sigma \in S_n$. An inversion of σ is a pair of integers $1 \leq i, j \leq n$ such that $i < j$ yet $\sigma(i) > \sigma(j)$. We write $(i, j)_\sigma$ or simply (i, j) if i and j form an inversion pair of σ with $i < j$. We denote the total number of inversions of σ by $n_{inv}(\sigma)$.

In [3], Fresse generalized this notion of inversion to multi-column tableaux. Let τ be a row-standard tableau of shape λ . A pair of entries $1 \leq i, j \leq N$ from the same column of τ are an **inversion** of τ if $i < j$ and one of the following holds:

1. Either i or j lacks an entry to its right and i is below j .
2. i is bordered immediately on its right by i' , j is bordered immediately on its right by j' , and $i' > j'$.

As above, we write $(i, j)_\tau$ or simply (i, j) if i and j constitute an inversion pair of τ with $i < j$, while we denote the total number of inversion pairs of τ by $n_{inv}(\tau)$. Notice that a row-standard tableau τ is a standard Young tableau if and only if $n_{inv}(\tau) = 0$, and that this definition specializes to the notion of permutation inversion if you represent $\sigma \in S_n$ by the single-column tableau whose entries appear in the order $\sigma(1), \dots, \sigma(n)$.

Given any row-standard tableau τ of shape λ , one may independently reorder each of the columns of τ to produce a tableau that is column-standard. As argued in [3], the resulting tableau is also guaranteed to be row-standard, and hence is a standard Young tableau of shape λ . This standard Young tableau is clearly unique, and is referred to as the **standardization** of τ , denoted $st(\tau)$; we equivalently say that τ is “based” on the standard Young tableau $st(\tau)$ ¹. As any row-standard tableau may be “standarized” by removing its inversions, we henceforth refer to the collection of all row-standard tableaux of shape λ as **inverted (standard) Young tableaux** of shape λ . We denote the set of all inverted Young tableaux of shape λ by $I(\lambda)$.

¹We note that this is a distinct definition from standardization of a semi-standard Young tableau, as found in the literature.

Figure 1 shows an inverted Young tableau of shape $\lambda = (3, 3, 3)$ alongside its standardization. Throughout this paper, we will highlight the location of inversions with shaded boxes, although do notice that this convention becomes slightly ambiguous when there are multiple inversions per column.

1	2	8		1	2	6
4	5	6	→	3	5	8
3	7	9		4	7	9

Figure 1: An inverted tableau τ with inversions $(3, 4), (2, 5), (6, 8)$ and its standardization $st(\tau)$

In [3], Fresse utilizes inverted Young tableaux to determine the Betti numbers of Springer fibers in type A. In particular, he fixes a specific standard Young tableau T of shape λ and considers the corresponding set of inverted tableaux $inv(T) = \{\tau \mid st(\tau) = T\}$. He then argues that the component of the Springer variety F_λ corresponding to T has m^{th} Betti number equaling the number of inverted Young tableaux in $inv(T)$ with precisely $d - m$ inversions, where d is the dimension of the entire Springer variety. Ranging over all standard Young tableaux T then gives:

Theorem 1.2 (Fresse). *Let λ be a partition of N and consider the Springer variety F_λ . If $\dim(F_\lambda) = d$, then the m^{th} Betti number $b_m = \dim(H^{2m}(F_\lambda, \mathbb{Q}))$ equals the number of inverted Young tableaux τ of shape λ with $n_{inv}(\tau) = d - m$.*

For a given shape $\lambda = (\lambda_1, \dots, \lambda_m)$, let $S_i(\lambda) = S_i(\lambda_1, \dots, \lambda_m)$ denote the set of row-standard tableaux τ of shape λ with $n_{inv}(\tau) = i$. Theorem 1.2 then states that $b_m = |S_{d-m}(\lambda)|$. We henceforth refer to elements of $S_i(\lambda)$ as **i-inverted (standard) Young tableau** of shape λ . Note that standard Young tableaux correspond to $S_0(\lambda) = S(\lambda)$, and that the size of the set $S_0(\lambda)$ is explicitly determined by the Hook-Length formula.

This paper won't assume detailed knowledge of Springer varieties. We present Fresse's results primarily to motivate our study of tableaux inversions. Our approach is purely combinatorial, and is concerned with developing tractable methods for enumerating elements in the sets $S_i(\lambda)$ for arbitrary λ and any $i \geq 0$.

The primary difficulty in adapting Fresse's approach to these goals is that his enumerative results require knowledge of specific standard Young tableau as a starting point. Even if one can easily produce a generating function that gives the sizes of the sets $S_i(\lambda) \cap inv(T)$ for a fixed standard tableau T , there is no workable way for determining all possible generating functions when one ranges over all underlying T . We avoid this difficulty by presenting techniques for directly calculating the sizes of the entire sets $S_i(\lambda)$.

1.1 Outline of Results

In this paper we begin with a series of basic results about the number of inverted Young Tableaux. We give a closed formula (Proposition 2.1) for the total number of inverted

tableaux for a fixed shape λ in terms of its row lengths. Then we show that, for any shape λ , there is a maximal inversion number M_λ that is determined by sums of triangular numbers, and that there exists a single tableau of shape λ realizing that maximal number of inversions (Proposition 2.2). After directly enumerating all one-column and two-row tableaux with a fixed number of inversions (Theorem 2.3), we proceed to our primary combinatorial results. Our first major result, introduced as Theorem 3.1 (and generalized to non-rectangular tableaux in Theorem 3.2), relates the number of rectangular tableaux with precisely one inversion to the number of standard Young tableaux of a related “stair-step” shape:

Theorem 1.3. *Let $n, m \geq 1$, and consider the m -row shapes $\lambda = (n, \dots, n)$, $\tilde{\lambda} = (n + 1, n, \dots, n, n - 1)$. Then $|S_1(\lambda)| = |S_0(\tilde{\lambda})|$.*

We then turn to tableaux with more than one inversion. Working downward from the maximum inversion number M_λ , Theorem 4.1 and 4.2 combine to prove the following closed formulas enumerating the number of $m \times n$ rectangular tableaux with precisely $M_\lambda - 1$ and $M_\lambda - 2$ inversions:

Theorem 1.4. *Take $n \geq 1, m \geq 3$ and consider the m -row rectangular shape $\lambda = (n, \dots, n)$. Then $|S_{M_\lambda-1}(\lambda)| = mn - 1$ and $|S_{M_\lambda-2}(\lambda)| = \frac{(mn-2)(mn+1)}{2}$.*

Closing the paper are a series of more specific results and conjectures enumerating inverted tableaux whose inversion number is “sufficiently large” (a realm that we refer to as the “tail end” of a shape’s inversion distribution). All results and conjectures in this paper were developed with the help of a computer program developed by the authors. Details about this program, as well as several tables of generated results, may be found in Appendix A.

2 Basic Enumerative Results About Tableau Inversions

In this section, we present a number of foundational enumerative results involving tableau inversions that are independent from the more involved methods of Sections 3 and 4.

2.1 Total Number of Inverted Young Tableaux of Shape λ

Let $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_m)$ be a non-increasing partition of the positive integer N , and consider the set $I(\lambda) = \bigcup_i S_i(\lambda)$ of inverted Young tableaux of shape λ (with any number of inversions $i \geq 0$).

$I(\lambda)$ is the collection of all tableaux of shape λ that are row-standard. As there is no restriction on the columns of these tableaux, and since there is a unique way to order each row of a tableau so that it is increasing, elements of $I(\lambda)$ are in bijection with ordered partitions of $\{1, \dots, N\}$ into m subsets of respective sizes $\lambda_1, \dots, \lambda_m$. This immediately proves the following.

Proposition 2.1. *Let $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_m)$ be a non-increasing sequence of positive integers. Then the total number of inverted Young tableaux of shape λ is:*

$$|I(\lambda)| = \sum_{i=0}^{\infty} |S_i(\lambda)| = \binom{\lambda_1 + \lambda_2 + \cdots + \lambda_m}{\lambda_m} \binom{\lambda_1 + \lambda_2 + \cdots + \lambda_{m-1}}{\lambda_{m-1}} \cdots \binom{\lambda_1 + \lambda_2}{\lambda_2} \binom{\lambda_1}{\lambda_1}$$

In terms of Fresse's results in [3], Proposition 2.1 states that the sum of all Betti numbers of F_λ is given by $|I(\lambda)|$. As one application of this straightforward result, for $\lambda = (n, n)$ the quantity $|I(\lambda)|$ gives the rank of the sl_2 skein module of surfaces over the solid torus with $2n$ boundary points [5]. In an upcoming paper [2], one of the authors proves that $|I(\lambda)|$ with $\lambda = (n, n, n)$ gives the rank of the sl_3 skein module of surfaces over the solid torus with $2n$ boundary points.

2.2 Maximum Number of Inversions for Shape λ

A natural question to arise from consideration of Proposition 2.1 is how many nonzero $|S_i(\lambda)|$ appear in the summation for $|I(\lambda)|$. Clearly, a row-standard tableau with a finite number of entries can only admit a finite number of inversions, as each entry can only be involved in a maximum of one inversion pair with every other entry in its column. To determine “maximum inversion number”, we consider column heights in inverted Young tableaux of shape λ .

First consider the case of single-column tableau with N total entries, so that $\lambda = (1, 1, \dots, 1)$. The maximum number of inversions for a tableau of this shape occurs when the entries appear in descending order $N, \dots, 2, 1$. In this situation we have $N - 1$ inversion pairs $(1, 2), (1, 3), \dots, (1, N)$ where 1 is the smaller entry, $N - 2$ inversion pairs $(2, 3), \dots, (2, N)$ where 2 is the smaller entry, etc. It follows that the maximum number of inversions for a tableau of shape λ is $T_{N-1} = 1 + 2 + \cdots + (N - 1) = \binom{N}{2}$, which is the $(N - 1)^{\text{st}}$ triangular number. Notice that the “reverse order” tableau described above is the unique inverted Young tableau of this shape realizing the maximum possible inversion number T_{N-1} .

Generalizing to tableaux of arbitrary shape λ is essentially a repeated application of the procedure above:

Proposition 2.2. *Let $\lambda = (\lambda_1, \dots, \lambda_m)$, and define $h_j = |\{\lambda_i \mid \lambda_i \geq j\}|$ to be the height of the j^{th} column for any tableau of shape λ . Then the maximum number of inversions for any inverted Young tableau of shape λ is:*

$$M_\lambda = \sum_j T_{h_j-1} = \sum_j \binom{h_j}{2}$$

Moreover, this maximum inversion number is realized by precisely one inverted Young tableau of shape λ , so that $|S_{M_\lambda}(\lambda)| = 1$.

Proof. We have already shown that the maximum number of inversions within a single column of height h_j is T_{h_j-1} , so when we range over all columns of our tableau we clearly can't obtain more than $\sum_j T_{h_j-1}$ total inversions. Thus we merely need to construct a tableau that exhibits this maximum number of inversions.

Let $N = \lambda_1 + \cdots + \lambda_m$. We work from right-to-left through the columns of the tableau. For the rightmost (n^{th}) column, place $N, N-1, \dots, N-h_n$ in decreasing order from top-to-bottom as in the single column case. There are then T_{h_n-1} inversions involving elements of this column. Moving one column to the left, we place the next h_{n-1} largest remaining entries in the unique order that guarantees $T_{h_{n-1}-1}$ inversions in its column. This is accomplished by “flipping” the order of the column relative to the column on its right, with the important convention that any empty boxes on the right are filled with arbitrary large numbers that increase from top-to-bottom, so that the smallest entry in the leftward column is directly to the right of the largest entry in the rightward column, etc. Continuing to work leftward through the columns, we repeat this procedure by placing the h_j largest remaining entries in the unique order that guarantees T_{h_j-1} inversions in the j^{th} column. An example of this procedure for a non-rectangular shape is shown in Figure 2.

Notice that this resulting tableau is guaranteed to be row-standard, as every entry in a given column is smaller than every entry in columns to its right. Also notice that this is the only possible tableau with the maximum of $\sum_j T_{h_j-1}$ inversions: changing the ordering within any column reduces the number of inversions in that column, and switching entries between columns is guaranteed to reduce the number of inversions possible in the leftward column (as there is now an entry in that leftward column that is larger than some entry in the rightward column). \square

2	7	10
1	8	9
3	6	
4	5	

Figure 2: The unique inverted Young tableau of shape $\lambda = (3, 3, 2, 2)$ with the maximum number of $T_3 + T_3 + T_1 = 13$ inversion pairs. When giving the middle column the “reverse” ordering relative to the rightmost column, we treat the empty slots in the lower-right as 11 and 12.

For a quick specialization to rectangular tableaux of size $m \times n$, note that Proposition 2.2 gives a maximum inversion number of $M_{n,\dots,n} = n \binom{m}{2}$. In particular, two-row rectangular tableaux have maximal inversion number $M_{n,n} = n$ equal to their number of columns.

2.3 Enumerating i -Inverted Young Tableaux of Shape $\lambda = (1, 1, \dots, 1)$

In general, it is extremely difficult to determine the number of inverted tableaux of shape λ with precisely i inversions, but there are several basic shapes λ for which this is computable without the complicated methods of Sections 3 and 4. The first of these shapes are single-column tableaux, so that $\lambda = (1, 1, \dots, 1)$. In this case, tableau inversions are identical to ordinary permutation inversions, which are well-studied in the literature.

The number of length- m permutations with precisely i inversions, and hence the number of tableaux of size $m \times 1$ with precisely i inversions, is given by the Mahonian number $M(m-1, i)$ [7]. For a fixed $m \geq 1$, in [7] it is also shown the Mahonian numbers have generating function:

$$\sum_{i=0}^{\infty} M(m-1, i)x^i = \prod_{j=0}^{m-1} \sum_{k=0}^j x^k = (x+1)(x^2+x+1)\cdots(x^{m-1}+\cdots+x+1) \quad (1)$$

Notice that the generating polynomial of Equation 1 is of degree T_{m-1} and that it is always monic, corroborating our results from Proposition 2.2.

2.4 Enumerating i-Inverted Young Tableaux of Shape $\lambda = (n, n)$

Although we can no longer rely upon previously-established results about permutation inversions, the other situation where it is still tractable to directly compute the number of i-inverted tableaux is with two-row rectangular tableaux. We immediately jump to the two-row case because one-row tableaux are entirely trivial: if $\lambda = (n)$ then $|S_0(\lambda)| = 1$ and $|S_i(\lambda)| = 0$ for all $i \geq 1$, as inversions require at least two elements in a column.

Before proceeding to the two-row case, we establish some new terminology. An $m \times n$ rectangular tableau τ is said to be a **(vertically) split tableau** if the first j columns of τ (where $j < n$) contain the first jm entries of τ . Hence, a split tableau is merely two disjoint tableaux sitting side-by-side, with the entries of the rightward tableau re-indexed to continue where the leftward tableau stopped. Notice that the number of split (standard Young) tableaux of size $m \times n$ that split after the j^{th} column equals the number of (standard Young) tableaux of size $m \times j$ times the number of (standard Young) tableaux of size $m \times (n-j)$. Also notice that an inverted Young tableau τ splits after its j^{th} column iff its standardization $st(\tau)$ splits after its j^{th} column.

It is well-known that the number of rectangular standard Young tableaux of size $2 \times n$ equals the n^{th} Catalan number C_n , a fact that can be directly calculated from the Hook-Length Formula of Theorem 1.1 or by setting up a bijection with noncrossing matchings. This means that $|S_0(n, n)| = C_n$.

To find $|S_i(n, n)|$ for $i > 0$, we begin by noting that each column in a two-row inverted Young tableau has either 0 or 1 inversion pair. If a column admits an inversion pair, this means that the larger entry in that column containing the inversion must be smaller than the smallest of the two entries in the column directly to the right. It follows that a two-row inverted Young tableau must split after any column in which it has an inversion. For an example of this phenomenon see Figure 3, where larger subscripts correspond to larger entries. In that example, $b_2 < c_1$ ensures that the tableau must split.

Theorem 2.3. *Let $\lambda = (n, n)$ for some $n \geq 1$. Then the number of inverted Young tableaux of shape λ with precisely i inversions is:*

$$|S_i(n, n)| = \left(\sum_{k_1 + \cdots + k_i = n} C_{k_1} C_{k_2} \cdots C_{k_i} \right) + \left(\sum_{l_1 + \cdots + l_{i+1} = n} C_{l_1} C_{l_2} \cdots C_{l_{i+1}} \right)$$

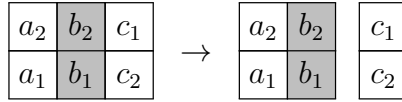


Figure 3: An Inversion Guarantees a Split Two-Row Tableau

Where C_j is the k^{th} Catalan number, and the summations run over all ordered partitions of n of length i and of length $i + 1$, respectively.

Proof. First observe that, if $i = 0$, then the first summation is empty and the second summation contains the single summand C_n corresponding to the unique length-one partition (n) of n . Thus $|S_0(n, n)| = C_n$, as detailed above.

For $i \geq 1$, consider the subset of i -inverted two-row tableaux τ where the i inversions are located at columns j_1, \dots, j_i . As mentioned above, any such tableau must vertically split after each of these columns. If $j_i = n$, the tableau splits into i sub-tableaux $\tau_0, \dots, \tau_{i-1}$ such that τ_q ends after column j_{q+1} . If $j_i < n$, the tableau splits into $i + 1$ sub-tableaux τ_0, \dots, τ_i such that τ_q similarly ends after column j_{q+1} for $q < i$ and τ_i ends after column n . If $j_i = n$, we may conclude that the number of such tableaux is $C_{j_1} C_{j_2 - j_1} \cdots C_{j_i - j_{i-1}}$, which is a summand from the first summation in the Theorem. If $j_i < n$, we may conclude that the number of such tableaux is $C_{j_1} C_{j_2 - j_1} \cdots C_{n - j_i}$, which is a summand from the second summation in the Theorem. Noting that any ordered partition of n of length i or $i + 1$ uniquely corresponds to a selection of i columns from our tableau (including the n^{th} column or not including the n^{th} column, respectively), ranging over all possible partitions gives the full summation for $|S_i(n, n)|$ in the Theorem. \square

Example 2.4. For $n = 3$, Theorem 2.3 gives:

i	$\lambda \vdash 3$	$ S_i(3, 3) $
0	(3)	5
1	(3), (2, 1), (1, 2)	$5 + 2 * 1 + 1 * 2 = 9$
2	(2, 1), (1, 2), (1, 1, 1)	$2 * 1 + 1 * 2 + 1 * 1 * 1 = 5$
3	(1, 1, 1)	$1 * 1 * 1 = 1$

Readers in search of additional combinatorial identities should note that, when successively ranging over all $n > 0$, the sequence formed by the $|S_i(n, n)|$ is A039599 on OEIS [1].

In light of Theorem 1.2, Theorem 2.3 applies to give a closed formula for all Betti numbers of the Springer variety F_λ when $\lambda = (n, n)$. Also note that, as $(1, \dots, 1)$ is the only partition of n of length at least n , then $M_{n,n} = n$ and $|S_n(n, n)| = 1$ as predicted by Proposition 2.2.

3 Enumerating 1-Inverted Young Tableaux

For the remainder of this paper, we introduce a number of techniques that allow us to calculate the number of inverted Young tableaux $|S_i(\lambda)|$ with precisely i inversions for a

variety of shapes λ . One of our primary techniques will be to develop bijections between i -inverted tableaux of shape λ and j -inverted tableaux of some other shape $\tilde{\lambda}$. In some cases, this second set of inverted Young tableaux will then be directly enumerable via the Hook-length formula or the methods of Section 2. In this particular section we begin by restricting our attention to tableaux with precisely one inversion. We present an algorithm for directly calculating $|S_1(\lambda)|$ for rectangular tableaux of shape $\lambda = (n, \dots, n)$. We then offer a direct generalization to arbitrary (non-rectangular) shapes λ .

3.1 Calculating $|S_1(\lambda)|$, Rectangular Case

Let $\lambda = (n, \dots, n)$, with m consecutive n 's, so that we are dealing with rectangular tableaux of size $m \times n$. The goal of this subsection is to develop a bijection that proves the following.

Theorem 3.1. *Let $n, m \geq 1$, and consider the m -row shapes $\lambda = (n, \dots, n)$, $\tilde{\lambda} = (n + 1, n, \dots, n, n - 1)$. Then $|S_1(\lambda)| = |S_0(\tilde{\lambda})|$.*

Theorem 3.1 relates the number of 1-inverted rectangular Young tableaux of size $m \times n$ to the number of standard Young tableaux of a certain “stair-step” shape $\tilde{\lambda}$. $\tilde{\lambda}$ is obtained from λ by removing the box at the lower-right corner of our $m \times n$ rectangular diagram and moving it to the top of a new, $(n + 1)^{\text{st}}$ column. See Figure 4 for an example of this shape change. Note that $|S_0(\tilde{\lambda})|$ is calculable using the Hook-Length Formula.

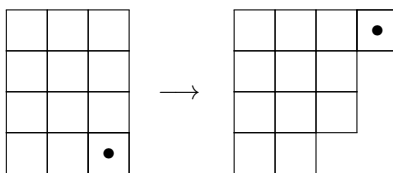


Figure 4: Shape Change in the Bijection of Theorem 3.1 when $\lambda = (3, 3, 3, 3)$

Proof of Theorem 3.1. We outline two procedures that give well-defined maps $\phi_1 : S_1(\lambda) \rightarrow S_0(\tilde{\lambda})$, $\phi_2 : S_0(\tilde{\lambda}) \rightarrow S_1(\lambda)$, and then argue that those maps are inverses of one another.

For our map $\phi_1 : S_1(\lambda) \rightarrow S_0(\tilde{\lambda})$, take a 1-inverted rectangular tableau $\tau \in S_1(\lambda)$. Assume that the sole inversion pair of τ is (a, b) , where $a < b$, and that this inversion appears in the k^{th} column of τ . Notice that, since τ contains only one inversion, b must appear directly above a in the k^{th} column. We begin our first procedure with the larger element b of the inversion pair (a, b) . We recursively “bump” an increasing sequence distinguished elements $\{b = c_k, c_{k+1}, \dots, c_n\}$, one from each column of τ from the k^{th} column rightward, to the right by one column each. Our procedure is as follows:

1. If $j < n$, c_j isn't in the rightmost column of τ . Define c_{j+1} to be the smallest entry d in the $(j + 1)^{\text{st}}$ column of τ such that $d > c_j$. Then “bump” c_j into the box occupied by c_{j+1} , temporarily producing a box in the $(j + 1)^{\text{st}}$ column that is shared by two

elements. There will then be an empty box in the j^{th} column: repeatedly fill any open boxes in the j^{th} column by sliding the smaller of the elements lying directly to the right or directly below that empty box into the box. Do this until the open box has been moved into the $(j+1)^{\text{st}}$ column. Then repeat this entire procedure with c_{j+1} until you reach c_n .

2. If $j = n$, c_n lies in the rightmost column of τ . Move c_n to the top of a new $(n+1)^{\text{st}}$ column of τ . There will be an empty box in the n^{th} column: slide all entries that are below that empty box up one slot, moving the empty box to the bottom of the column.

An example of the procedure for $S_1(3, 3, 3) \hookrightarrow S_0(4, 3, 2)$ is shown in Figure 5. This algorithm always transforms τ into a tableau $\tilde{\tau}$ of shape $\tilde{\lambda}$. As there is a unique allowable move at each step, the resulting $\tilde{\tau}$ is unique. If τ begins with only a single inversion, the “back-filling” component of the procedure from case #1 ensures that $\tilde{\tau}$ is a standard Young tableau. Thus this algorithm describes a well-defined map $\phi : S_1(\lambda) \rightarrow S_0(\tilde{\lambda})$.

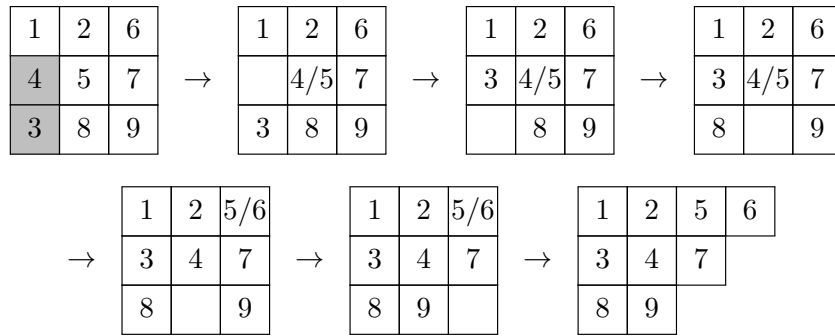


Figure 5: $S_1(3, 3, 3) \hookrightarrow S_0(4, 3, 2)$ Example

For our second map $\phi_2 : S_0(\tilde{\lambda}) \rightarrow S_1(\lambda)$, take a standard tableau $T \in S_0(\tilde{\lambda})$ and let c_n be the sole entry in the $(n+1)^{\text{st}}$ column of T . Here our goal is to define a decreasing sequence of distinguished entries $\{c_n, c_{n-1}, \dots\}$, one from each column of T beginning with the $(n+1)^{\text{st}}$ column, and “reverse bump” those elements one column to the left until one of the distinguished elements becomes the larger member of an inversion pair. Here our procedure is as follows:

1. Consider c_j , which originally lies in the $(j+1)^{\text{st}}$ column of T . There will always be an empty box in the j^{th} column of T . Repeatedly fill that empty box with the largest of c_j , the entry directly above the box, and the entry directly to the left of the box, stopping when the empty box is moved leftward into the $(j-1)^{\text{st}}$ column of T or when c_j directly fills the empty box.
2. If the empty box is moved into the $(j-1)^{\text{st}}$ column, define c_{j-1} to be the largest entry d in the j^{th} column such that $d < c_j$. Then move c_j into the box occupied by c_{j-1} , so that the tableau temporarily has a box containing two entries. Now repeat step #1 with c_{j-1} .

3. If c_j directly fills an empty box in the j^{th} column and a is the entry that lies directly above c_j after this insertion, create a single inversion pair (a, c_j) by flipping the rows containing a and c_j from the j^{th} column leftward. The most important thing to note here is that the final tableau always admits the inversion (a, c_j) at this step: i.e.- c_j is guaranteed to be smaller than the entry directly to the right of a . This is because, when c_j fills an empty box, it always must move down by at least one row. See Figure 6 for additional explanation of this fact.

a_1	b_1	d_1
a_2	b_2	$\frac{c_j}{c_{j+1}}$
a_3		d_3
a_4	b_4	d_4

Figure 6: The $S_0(n+1, n, \dots, n, n-1) \hookrightarrow S_1(n, \dots, n)$ procedure at the j^{th} column. As $b_2 < c_j$, at this stage either c_j fills the empty box and an inversion is possible or a_3 is slid right and we return to step #1 for the next column to the left. It is not possible for b_2 to be slid down, which is the only way for c_j not to slide down at least one row when filling an empty box.

An example of this procedure for $S_0(4, 3, 2) \hookrightarrow S_1(3, 3, 3)$ is shown in Figure 7. Notice that the “front sliding” move of step #1 ensures that the tableau remains both row and column standard, so that the final tableau only has the single inversion that is introduced in #3. Also note that #3 eventually has to apply at some point in the procedure, as $c_j > 1$ for all j , and thus that an inversion is always eventually added. As there is a unique allowable move at each step, the resulting 1-inversion tableau \tilde{T} is unique and the algorithm represents a well-defined map $\phi_2 : S_0(\tilde{\lambda}) \hookrightarrow S_1(\lambda)$.

ϕ_2 is constructed so that it is clearly the inverse of ϕ_1 . In particular, it is straightforward to check that $\phi_2 \circ \phi_1(\tau) = \tau$ and $\phi_1 \circ \phi_2(T) = T$. This demonstrates the bijectivity of both maps and allows us to conclude that $|S_1(\lambda)| = |S_0(\tilde{\lambda})|$.

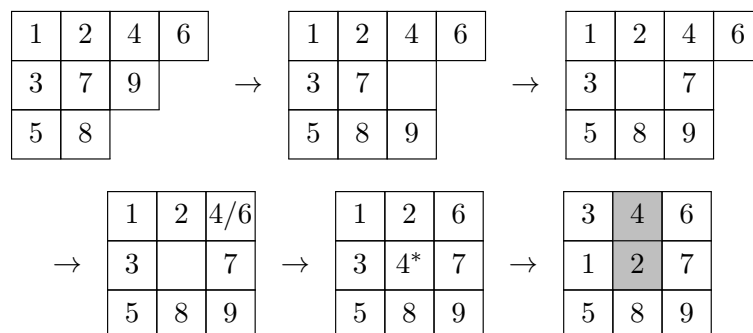


Figure 7: $S_0(4, 3, 2) \hookrightarrow S_1(3, 3, 3)$ Example

□

3.2 Calculating $|S_1(\lambda)|$, Non-Rectangular Case

The goal of this subsection is to prove a generalization of Theorem 3.1 that enumerates 1-inverted Young tableaux for an arbitrary (non-rectangular) shape λ . Before explicitly giving our result in Theorem 3.2, we need to establish some new notation describing non-rectangular tableaux shapes.

For a tableau of shape $\lambda = (\lambda_1, \dots, \lambda_m)$, define $d_i = \lambda_i - \lambda_{i+1}$ for $1 \leq i < m$ and $d_m = \lambda_m$. Notice that $d_i \geq 0$ for all i and that $d_i > 0$ iff the rightmost entry in the i^{th} row of our tableau is a “lower-right corner” for our tableau. Similarly define $\tilde{d}_i = \lambda_{i-1} - \lambda_i$ for $1 < i \leq m$ and $\tilde{d}_1 = \infty$. Notice that $\tilde{d}_i \geq 0$ for all i and that $\tilde{d}_i > 0$ iff an additional box may be added to the end of the i^{th} row of a tableau of shape λ to produce a valid tableau shape.

Theorem 3.2. *Let $\lambda = (\lambda_1, \dots, \lambda_m)$ be a non-decreasing sequence of positive integers with $m \geq 1$. Then:*

$$|S_1(\lambda_1, \dots, \lambda_m)| = \sum_E |S_0(\lambda_1 + \epsilon_1, \dots, \lambda_m + \epsilon_m)|$$

Where the summation runs over all tuples $E = (\epsilon_1, \dots, \epsilon_m)$ such that $\epsilon_i = 0, \pm 1$ for all i , $\epsilon_i = -1$ for precisely one $i = i_1$ such that $i_1 > 1$ and $d_{i_1} > 0$, and $\epsilon_i = 1$ for precisely one $i = i_2$ such that $i_2 < i_1$ and $\tilde{d}_{i_2} > 0$.

Observe that valid tuples $E = (\epsilon_1, \dots, \epsilon_m)$ include all zeroes apart from a single $+1$ and a single -1 , with the -1 appearing after the $+1$. In terms of actual tableaux, the summation of Theorem 3.2 runs over all shapes that are obtained from λ by moving a single lower-right corner (that isn’t in the first row) to create a new lower-right corner that is at least one row higher than the box’s original location. See Figure 8 for an example of the relevant tableaux shapes in the summation.

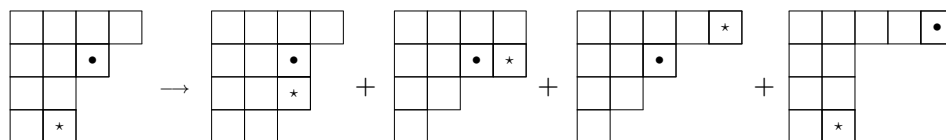


Figure 8: Shape Change in the Bijection of Theorem 3.2 when $\lambda = (4, 3, 2, 2)$

Also notice that when $\lambda = (n, \dots, n)$ there is a single valid tuple $E = (1, 0, \dots, 0, -1)$, and thus that the summation of Theorem 3.2 specializes to the equality of Theorem 3.1 in the case of rectangular tableaux.

Proof of Theorem 3.2. (\Leftarrow) Take $\tau \in S_1(\lambda_1, \dots, \lambda_m)$, and assume that the sole inversion pair (a, b) lies in the k^{th} column of τ . We repeatedly “bump” a sequence of distinguished elements rightward, as in the first map ϕ_1 from the proof of Theorem 3.1, following the exact same rules at each intermediate step. The difference here is that we terminate the procedure (possibly before the rightmost column) as soon as a distinguished element can be placed in a valid tableau position that does not fit inside the original shape $(\lambda_1, \dots, \lambda_m)$.

The resulting tableau shape will differ from λ in that it will have a new lower-right corner added at the final step, and that it will have a single vacated box at the bottom of some column (k^{th} or leftward) that is created by the “back-filling” component of the procedure. The specific location of these boxes depends on the inverted tableau, but the procedure ensures that it will always be one of the shapes from the summation in the Theorem. As in the proof of Theorem 3.1 this is a well-defined map that always yields a standard Young tableau.

(\leftarrow) Now take $\tau \in S_0(\tilde{\lambda})$, where $\tilde{\lambda}$ is any one of the shapes from the summation in the Theorem, and consider the entry in the sole box (always one of the lower-right corners of $\tilde{\lambda}$) that does not fit inside the shape λ . As in the second map ϕ_2 from the proof of Theorem 3.1, we recursively “back-bump” a sequence of distinguished elements one column leftward, following the exact same rules as in that earlier procedure. Now the difference is that we are only allowed to slide entries into boxes that are contained in λ , even if that spot is a valid tableau position. For the same reasons as described in the proof of Theorem 3.1, this map is well-defined and produces a 1-inversion tableau of the desired shape. This map is also clearly the inverse of the map defined above, so if you define a piecewise map by ranging over all valid shapes $\tilde{\lambda}$ the result is a bijection. \square

4 Enumerating i -Inverted Young Tableaux, $i > 1$

In Section 3 we developed bijections between 1-inverted rectangular tableaux of an arbitrary shape λ and 0-inverted tableaux of some related collection of shapes $\{\tilde{\lambda}_i\}$. One may ask whether the methods of that section readily extend to higher numbers? More specifically, does there exist a clear relationship between the number of i -inverted tableaux of shape λ and the numbers of $(i - j)$ -inverted tableaux of shapes $\{\tilde{\lambda}_i\}$ for some integer $j > 0$?

In Appendix A, we present a full comparison of $|S_i|$ for several choices of three-row rectangular shapes $\lambda = (n, n, n)$ and the associated “stair-step” shapes $\tilde{\lambda} = (n+1, n, n-1)$. As predicted by Theorem 3.1, we have $|S_1(\lambda)| = |S_0(\tilde{\lambda})|$. Sadly, even in our specialization to rectangular λ , Theorem 3.1 does not extend to a bijection between $|S_i(\lambda)|$ and $|S_{i-1}(\tilde{\lambda})|$ for all $i > 1$.

What those tables do reveal are discernible patterns in $|S_i(\lambda)|$ for sufficiently high choices of i . In addition to $|S_{M_\lambda}(\lambda)| = 1$ for the maximum inversion number M_λ (see Proposition 2.2), notice that $|S_{M_\lambda-1}(\lambda)| = 3n - 1$ and $|S_{M_\lambda-2}(\lambda)| = \frac{(3n-2)(3n+1)}{2}$ for all $n \geq 1$. Also notice that $|S_i(\lambda)| = |S_{i-2}(\tilde{\lambda})|$ for all values of i beyond a certain point. We henceforth refer to the region where this latter equality holds as the “tail-end” of the inversion table. For our three-row rectangular tableaux with $\lambda = (n, n, n)$, this tail-end consistently begins at $i = n + 1$.

In this section we prove formulas for $|S_{M_\lambda-1}(\lambda)|$ and $|S_{M_\lambda-2}(\lambda)|$ in the rectangular case, generalizing our observations about the three-row tableaux in Appendix A. We then prove a special case of the “tail-end” bijection and conjecture as to how this result extends to the general rectangular case. For this entire section, we restrict our attention

to rectangular tableaux, although the authors suspect that there may exist generalizations to non-rectangular tableaux that mirror Subsection 3.2.

4.1 Calculating $|S_{M_\lambda-1}(\lambda)|$ and $|S_{M_\lambda-2}(\lambda)|$, Rectangular Case

Directly enumerating inverted tableaux becomes increasingly tractable as one approaches the maximum inversion number M_λ . This is because having large numbers of inversions in a fixed column increases the likelihood that the tableau splits after that column. In the $(M_\lambda - 1)$ - and $(M_\lambda - 2)$ -inversion cases, non-split tableaux are so rare that they may be easily enumerated. As split tableaux are easily counted, this allows for relatively straightforward closed formulas for $|S_{M_\lambda-1}(\lambda)|$ and $|S_{M_\lambda-2}(\lambda)|$.

Theorem 4.1. *Take $n \geq 1$ and $m \geq 2$, and consider the m -row shape $\lambda = (n, \dots, n)$. If $M_\lambda = nT_{m-1}$ is the maximum inversion number, then $|S_{M_\lambda-1}(\lambda)| = mn - 1$.*

Proof. Fix any $m \geq 2$. With our choice of m implicit, temporarily denote $S_i(n) = S_i(n, \dots, n)$ and $M_n = M_{(n, \dots, n)} = nT_{m-1}$. Thus we are looking to show that $|S_{M_n-1}(n)| = mn - 1$. We proceed by induction on $n \geq 1$, with the base case of $|S_{M_1-1}(1)| = m - 1$ following from Subsection 2.3 and the Mahonian number identity $M(m - 1, T_{m-1} - 1) = m - 1$.

For the inductive case, we think about isolating the leftmost column in an $(M_n - 1)$ -inverted tableau of size $m \times n$. To obtain the required number of inversions, one of the following must hold:

1. The leftmost column has precisely M_1 inversions and the remaining $(n - 1)$ columns have a total of precisely $M_{(n-1)} - 1$ inversions.
2. The leftmost column has precisely $M_1 - 1$ inversions and the remaining $(n - 1)$ columns have a total of precisely $M_{(n-1)}$ inversions.

In case #1, there is a unique ordering of the elements in the leftmost column (relative to the column immediately on its right) and the tableau necessarily splits after the first column. It follows that the number of tableaux satisfying this case equals $|S_{M_{(n-1)}-1}(n - 1)|$.

In case #2, there are two sub-cases depending upon whether the largest entry a_m or the second largest entry a_{m-1} in the leftmost column is directly to the left of the smallest entry b_1 in the column immediately on its right. For an illustration of these two sub-cases when $m = 4$, see Figure 9 below.

In the first sub-case, the tableau must split after the first column, so that the first column must contain the entries $1, 2, \dots, m$. Since there is a unique way to produce $M_{(n-1)}$ inversions across the latter columns, the placement of $m + 1, m + 2, \dots, mn$ is predetermined. In the first column, there are $m - 1$ inversions involving a_m and we must have $T_{m-1} - 1 - (m - 1) = T_{m-2} - 1$ inversions among the remaining $m - 1$ entries in that column. There are $M(m - 2, T_{m-2} - 1) = m - 2$ distinct arrangements that produce these $T_{m-2} - 1$ inversions. This means that there are precisely $m - 2$ inverted tableaux satisfying this sub-case.

In the second sub-case, the tableau need not split after the first column, as the largest entry a_m in the leftmost column may or may not be greater than the smallest entry b_1 in the column directly to its right. No matter the relationship between these two entries, there is a unique arrangement that produces the requisite $M_{(n-1)}$ inversions across the latter columns. In the first column, there are always $m - 2$ inversions involving the top entry a_{m-1} , leaving T_{m-2} inversions among the remaining $m - 1$ entries in the column. There is only $M(m - 2, T_{m-2}) = 1$ arrangement of the remaining entries in the leftmost column that produces these T_{m-2} inversions. Also notice that the required arrangement in the leftmost column ensures that a_m is directly to the left of the second-smallest entry b_2 of the next column, and that the tableau always splits after the second column. This leaves the relationship between a_m and b_1 as the only open question about such a tableau. It follows that there are precisely 2 inverted tableaux in this sub-case (one for $a_m < b_1$, one for $a_m > b_1$).

Collecting results from all sub-cases and then applying the inductive assumption gives the desired result: $|S_{M_n}(n)| = |S_{M_{n-1}}(n-1)| + (m-2) + 2 = (m(n-1)-1) + (m-2) + 2 = mn - 1$. \square

a_4	b_1	a_4	b_1	a_3	b_1
a_3	b_2		b_2	a_4	b_2
a_2	b_3	(2)	b_3	a_2	b_3
a_1	b_4		b_4	a_1	b_4

Figure 9: The leftmost columns of an $(M_\lambda - 1)$ -inverted 4-row tableau, with subscripts denoting the relative ordering within each column (we've reordered the second column so it appears increasing). Case #1 from the proof of Theorem 4.1 is on the left, whereas the two sub-cases of case #2 from the proof are on the right. Parentheses denote number of inversions required in shaded sub-tableau.

Theorem 4.2. Take $n \geq 1$ and $m \geq 3$, and consider the m -row shape $\lambda = (n, \dots, n)$. If $M_\lambda = nT_{m-1}$ is the maximum inversion number, then $|S_{M_\lambda-2}(\lambda)| = \frac{(mn-2)(mn+1)}{2}$.

Proof. This requires a more complicated variation on the method from the proof of Theorem 4.2. Once again we fix $m \geq 3$ and temporarily denote $S_i(n) = S_i(n, \dots, n)$, $M_n = M_{(n, \dots, n)} = nT_{m-1}$. In this case we are looking to show that $|S_{M_n-2}(n)| = \frac{(mn-2)(mn+1)}{2}$. We proceed by induction on $n \geq 1$.

For the base case $|S_{M_n-2}(1)|$, one may directly verify the Mahonian number $M(m-1, T_{m-1}-2) = \frac{(m-2)(m+1)}{2}$. For the inductive case we once again isolate the leftmost column in an $(M_n - 2)$ -inverted tableau of size $m \times n$. We now have three ways to obtain the required number of inversions:

1. The leftmost column has precisely $T_m = M_1$ inversions and the remaining $(n-1)$ columns have precisely $M_{(n-1)} - 2$ inversions.

2. The leftmost column has precisely $T_m - 1 = M_1 - 1$ inversions and the remaining $(n - 1)$ columns have precisely $M_{(n-1)} - 1$ inversions.
3. The leftmost column has precisely $T_m = M_1 - 2$ inversions and the remaining $(n - 1)$ columns have precisely $M_{(n-1)}$ inversions.

In case #1, there is a unique ordering of the leftmost column to obtain the maximal number of M_1 inversions and the tableau necessarily splits after the first column. It follows that the number of inverted tableaux in this case equals $|S_{M_{(n-1)}-2}(n-1)|$.

In case #2, there are two sub-cases depending upon whether the largest entry a_m or the second largest entry a_{m-1} in the leftmost column is directly left of the smallest entry b_1 in the second column. For an illustration of these sub-cases when $m = 4$, see Figure 10.

In the first sub-case of case #2, the tableau must split after the first column. In the first column, a_m participates in $m - 1$ inversions and thus the remaining elements must account for $T_{m-1} - 1 - (m - 1) = T_{m-2} - 1$ total inversions. By Theorem 4.1, there are $m - 2$ ways to achieve these $T_{m-2} - 1$ inversions. It follows that there are precisely $(m - 2)|S_{M_{(n-1)}-1}(n - 1)|$ tableaux in this sub-case.

In the second sub-case of case #2, the tableau needn't split after the first column. In the first column, a_{m-1} accounts for $m - 2$ inversions and hence the remaining elements in that column must account for T_{m-2} inversions. This implies that the rest of the first column must be "fully inverted", and thus that a_m is directly to the left of the second-smallest entry b_2 in the second column. We then have $a_{m-1} < b_1$ and $a_m < b_2$, leaving the relationship between a_m and b_1 as the only question about the first two columns. Also noticing that b_1 is directly to the left of either c_m or c_{m-1} in the third column, as $m \geq 3$ we know that $a_m < b_2 < c_2 \leq c_{m-1}/c_m$. This means we needn't worry about the relationship between a_m and any elements of the third column. For each possible arrangement of the tableau's final $n - 1$ columns, it follows that there are precisely two inverted tableaux fitting this sub-case (one for $a_m < b_1$, one for $a_m > b_1$). Thus there are precisely $2|S_{M_{(n-1)}-1}(n - 1)|$ tableaux in this sub-case.

In case #3 there are four sub-cases, each depending upon which elements in the first column are directly to the left of the smallest two entries b_1, b_2 in the second column. Once again, reference Figure 10 for an illustration of these sub-cases in the situation where $m = 4$. Notice that in all of these sub-cases, the $M_{(n-1)}$ inversions needed over the final $n - 1$ columns ensures that there is precisely one arrangement of those columns for each arrangement of the first column. It also guarantees that any such tableau must split after the second column.

In the first sub-case of case #3, a_m is directly to the left of b_1 and the tableau must split after the first column. In the first column, a_m participates in $m - 1$ inversions and thus the remaining elements account for $T_{m-1} - 2 - (m - 1) = T_{m-2} - 2$ total inversions. By the base case of this theorem, there are $\frac{(m-1-2)(m-1+1)}{2} = \frac{(m-3)m}{2}$ ways to achieve these inversions. Thus there are exactly $\frac{(m-3)m}{2}$ tableaux in this sub-case.

In the second sub-case of case #3, a_{m-1} is directly to the left of b_1 and a_m is directly to the left of b_2 . As a_{m-1} is involved in $m - 2$ inversions and a_m is involved in $m - 2$

inversions, this leaves $T_{m-1} - 2 - (m-2) - (m-2) = T_{m-3} - 1$ inversions required across the remaining $m-2$ elements of that column. By Theorem 4.1, there are precisely $(m-2) - 1$ distinct ways to accomplish these inversions. For each of these $m-3$ arrangements, there is a single open choice as to the relationship between a_m and b_1 , seeing as a_m is to the left of b_2 and a_{m-1} is to the left of b_1 . Thus there are a total of $2(m-3)$ tableaux in this sub-case.

In the third sub-case of case #3, a_{m-1} still lies directly to the left of b_1 but now a_{m-2} lies directly to the left of b_2 . In this situation, a_{m-1} is involved in $m-2$ inversions and a_m is involved in $m-3$ inversions, which leaves $T_{m-1} - 2 - (m-2) - (m-3) = T_{m-3}$ inversions to be accounted for over the remaining $m-2$ entries of the first column. There is precisely one way to account for these remaining inversions, with a_m necessarily appearing directly to the left of b_3 . With a_{m-1} to the left of b_1 and a_m to the left of b_3 , the relationship of a_m to both b_1 and b_2 is an open question. Here the possible orderings are $a_m < b_1 < b_2$, $b_1 < a_m < b_2$, and $b_1 < b_2 < a_m$. As this is the only undetermined aspect of these tableaux, there are precisely 3 tableaux in this sub-case.

In the final sub-case of case #3, a_{m-2} lies directly to the left of b_1 . In the first column, a_{m-2} is involved in $m-3$ inversions, leaving $T_{m-1} - 2 - (m-3) = T_{m-2}$ inversions for the remaining $m-1$ entries of that column. There is precisely one arrangement of those entries that gives T_{m-2} inversions, but in this situation we still need to determine the relationship of both a_m and a_{m-1} to b_1 . Here there are possible orderings $a_{m-1} < a_m < b_1$, $a_{m-1} < b_1 < a_m$, and $b_1 < a_{m-1} < a_m$. As this is the only undetermined aspect of these tableaux, there are precisely 3 tableaux in this sub-case.

Collecting all sub-cases and citing Theorem 4.1 along with the inductive assumption gives:

$$\begin{aligned} |S_{M_\lambda-2}(\lambda)| &= |S_{M_{(n-1)}-2}(n-1)| + (m-2)|S_{M_{(n-1)}-1}(n-1)| + 2|S_{M_{(n-1)}-1}(n-1)| \\ &\quad + \frac{1}{2}(m-3)m + 2(m-3) + 6 \\ &= \frac{(m(n-1)-2)(m(n-1)+1)}{2} + m(m(n-1)-1) + \frac{m(m+1)}{2} \\ &= \frac{(mn-2)(mn+1)}{2}. \end{aligned}$$

□

Notice that the $m \geq 3$ condition of Theorem 4.2 (as opposed to the $m \geq 2$ condition of Theorem 4.1) sidesteps the semantic difficulties arising from the fact that a column in a two-row tableau can have at most 1 inversions, which implies that $M_1 - 2$ would need to be negative. The fact that $m \geq 3$ was also required in the appeal to which entry c_i lay directly to the right of b_1 in case #2. This precluded the situation where a_m may be smaller than elements of the third columns, which would have necessitated an even more sophisticated consideration of sub-cases.

In theory, the technique of Theorems 4.1 and 4.2 could be extended to the $M_\lambda - i$ inversion case is we require rectangular tableaux with at least $m = i+1$ columns. However,

a_4	b_1	a_4	b_1	a_3	b_1	a_4	b_1	a_3	b_1	a_3	b_1	a_2	b_1
a_3	b_2		b_2	a_4	b_2		b_2	a_4	b_2	a_2	b_2	a_4	b_2
a_2	b_3	(2)	b_3	a_2	b_3	(1)	b_3	a_1	b_3	a_4	b_3	a_3	b_3
a_1	b_4		b_4	a_1	b_4		b_4	a_2	b_4	a_1	b_4	a_1	b_4

Figure 10: The leftmost columns of an $(M_\lambda - 2)$ -inverted 4-row tableau, with subscripts denoting relative ordering within each column. Case #1 from the proof of Theorem 4.2 is on the left, the two sub-cases of case #2 are in the middle, and the four sub-cases of case #3 are on the right. Parentheses denote number of inversions required in shaded sub-tableaux.

the massive number of sub-cases required would make the proof of even the $M_\lambda - 3$ case extremely daunting.

4.2 Calculating $|S_i(\lambda)|$, “Tail-End”

Consider the three-row rectangular shape $\lambda = (n, n, n)$ and let $\tilde{\lambda} = (n + 1, n, n - 1)$ be the associated stair-step shape, as in the proofs of Section 3. Appendix A reveals that $|S_i(\lambda)| = |S_{i-2}(\tilde{\lambda})|$ when $i > n$. Since Proposition 2.2 gives $M_{\tilde{\lambda}} = M_\lambda - 2$, it follows that the inversion tables for λ and $\tilde{\lambda}$ can be matched “from the bottom up” until we reach $i = n$. Once again, we refer to this range as the “tail-end” of the inversion table for λ . In this Subsection, we take steps towards generalized proof of this “tail-end” phenomenon.

So consider the m -row rectangular shape $\lambda = (n, \dots, n)$ and the associated stair-step shape $\tilde{\lambda} = (n + 1, n, \dots, n, n - 1)$. By Proposition 2.2 we have $M_{\tilde{\lambda}} = M_\lambda - (m - 1)$. Thus one would expect an equality $|S_i(\lambda)| = |S_{i-m+1}(\tilde{\lambda})|$ for “sufficiently large” i . Predicting what constitutes “sufficiently large” i in this generalized tail-end is far less obvious. To find appropriate bounds on i , as well as to develop a general method of proof, we begin with a consideration of one-column tableaux:

Lemma 4.3. *Let $m \geq 2$, and consider the m -row column tableaux shape $\lambda = (1, \dots, 1)$. For the $(m-1)$ -row “hook” shape $\tilde{\lambda} = (2, 1, \dots, 1)$, we have $|S_i(\lambda)| = |S_{i-m+1}(\tilde{\lambda})|$ for all $i > T_{m-2}$, where T_{m-2} is the triangular number. Moreover, there exists a bijection $\phi : S_i(\lambda) \rightarrow S_{i-m+1}(\tilde{\lambda})$ such that, if the top entry of $\tau \in S_i(\lambda)$ is k , then the sole entry in the second column of $\phi(\tau)$ is also k .*

Proof. For each $1 \leq k \leq m$, denote the set of inverted tableaux in $S_i(\lambda)$ whose first entry is k by $S_i^k(\lambda)$. Similarly, denote the set of inverted tableaux in $S_i(\tilde{\lambda})$ whose sole second column entry is k by $S_i^k(\tilde{\lambda})$. To expedite notation, for the p -row column tableaux we refer to our $p \times 1$ shape as $(1, \dots, 1) = 1^p$, so that our original shape is $\lambda = 1^m$.

So fix any $1 \leq k \leq m$ and enforce our condition that $i > T_{m-2}$. This is equivalent to saying that $i - m + 2 > T_{m-3}$. As the maximum number of inversions for a column tableau of size $(m - 2) \times 1$ is $M_{m-2} = T_{m-3}$, it follows that $|S_{i-k+1}(1^{m-2})| = \dots =$

$|S_{i-m+2}(1^{m-2})| = 0$ no matter our choice of k . This ensures:

$$\begin{aligned} & |S_{i-k+1}(1^{m-2})| + |S_{i-k}(1^{m-2})| + \cdots + |S_{i-k-m+3}(1^{m-2})| \\ &= |S_{i-m+1}(1^{m-2})| + \cdots + |S_{i-k-m+3}(1^{m-2})| \end{aligned} \quad (2)$$

Consider any inverted tableau $\tau \in S_{i-k+1}(1^{m-1})$, and assume that the top entry of τ is a . τ contains precisely $a - 1$ inversion pairs that involve a . Removing a and re-indexing the remaining entries of τ so that a isn't skipped then produces an element of $S_{i-k+1-(a-1)}(1^{m-2})$. Ranging over all possible top entries for elements of $S_{i-k+1}(1^{m-1})$ gives:

$$|S_{i-k+1}(1^{m-1})| = |S_{i-k+1}(1^{m-2})| + |S_{i-k}(1^{m-2})| + \cdots + |S_{i-k+1-(m-2)}(1^{m-2})| \quad (3)$$

Applying similar reasoning to $S_{i-m+1}(1^{m-1})$ then allows us to rewrite Equation 2 as:

$$\begin{aligned} & |S_{i-k+1}(1^{m-1})| \\ &= |S_{i-m+1}(1^{m-1})| - |S_{i-m-k+2}(1^{m-2})| - |S_{i-m-k+1}(1^{m-2})| - \cdots - |S_{i-2m+3}(1^{m-2})| \end{aligned} \quad (4)$$

Pause to consider $S_i^k(\lambda) = S_i^k(1^m)$. Similarly to above, any $\tau \in S_i^k(1^m)$ has $k - 1$ inversion pairs involving k , so that if we simply remove k (and re-index so that we don't skip k) we have an element of $S_{i-k+1}(1^{m-1})$. It follows that:

$$|S_i^k(1^m)| = |S_{i-k+1}(1^{m-1})| \quad (5)$$

Now consider $S_{i-m+1}^k(\tilde{\lambda})$, and compare this set to $S_{i-m+1}(1^{m-1})$ by removing the sole entry in the second column (and then re-indexing so k isn't skipped). Note that removing this entry does not effect the number of inversions in the tableau. In general, $|S_{i-m+1}(1^{m-1})| \geq |S_{i-m+1}^k(\tilde{\lambda})|$, as simply adding k to the right of the first entry in a element from $S_{i-m+1}(m-1)$ may result in a tableau that isn't row-standard. Eliminating elements of $S_{i-m+1}(1^{m-1})$ that fail to be row-standard after inserting k gives $|S_{i-m+1}^k(\tilde{\lambda})| = |S_{i-m+1}(1^{m-1})| - |S_{i-m+1}^{k+1}(1^{m-1})| - \cdots - |S_{i-m+1}^m(1^{m-1})|$. This gives:

$$\begin{aligned} & |S_{i-m+1}^k(\tilde{\lambda})| \\ &= |S_{i-m+1}(1^{m-1})| - |S_{i-m-k+2}(1^{m-2})| - |S_{i-m-k+1}(1^{m-2})| - \cdots - |S_{i-2m+3}(1^{m-2})| \end{aligned} \quad (6)$$

Inserting Equations 5 and 6 into Equation 4 yields:

$$|S_i^k(1^m)| = |S_{i-m+1}^k(\tilde{\lambda})| \quad (7)$$

As Equation 7 holds for all $1 \leq k \leq m$, it is possible to construct a bijection $\phi : S_i(\lambda) \rightarrow S_{i-m+1}(\tilde{\lambda})$ that takes the top entry of each element of $S_i(\lambda)$ to the sole entry in the second column of an element from $S_{i-m+1}(\tilde{\lambda})$ \square

Lemma 4.3 suggests that any bijection $|S_i(\lambda)| = |S_{i-m+1}(\tilde{\lambda})|$ requires the existence of a column in the λ -shaped tableau with greater than T_{m-2} inversions. Notice that this is in

agreement with what Appendix A suggests about the tail-end in the $m = 3$ row case: the condition that $i > n$ is precisely the number of inversions necessary to guarantee that an i -inverted tableau of shape λ possesses at least one column with greater than $T_{m-2} = 1$ inversions. For an m -row rectangular tableau, the number of inversions necessary to guarantee that every i -inverted tableau has a column with greater than T_{m-2} inversions is $i > nT_{m-2}$. These observations lead us to postulate the following:

Conjecture 4.4. *Let $n \geq 1$, $m \geq 2$, and take the m -row shapes $\lambda = (n, \dots, n)$, $\tilde{\lambda}(n + 1, n, \dots, n, n - 1)$. Then $|S_i(\lambda)| = |S_{i-m+1}(\tilde{\lambda})|$ for all $i > nT_{m-2}$, where T_{m-2} is the triangular number.*

It is the authors' belief that it may be possible to prove Conjecture 4.4 via a "recursive bumping" algorithm superficially similar to the method of proof in Theorem 3.1. In this modified algorithm, the $\phi_1 : S_i(\lambda) \hookrightarrow S_{i-m+1}(\tilde{\lambda})$ map would begin in the rightmost column that contains greater than nT_{m-2} inversions. The first rightward "bump" would entail a reordering of the initial column in a manner that is uniquely determined by the bijection of Lemma 4.3. Subsequent rightward "bumps" would then be required to conserve inversion number in a manner that can be uniquely undone by some inverse algorithm $\phi_2 : S_{i-m+1}(\tilde{\lambda}) \hookrightarrow S_i(\lambda)$.

The difficulty in proving Conjecture 4.4 derives from these later "bumps". In particular, if a tableau possesses additional inversions to the right of the starting column, those inversions make the preservation of inversion number very complicated when attempting to bump a distinguished element past their column. Resolving the difficulties in this proof, as well as developing tractable methods for dealing with $|S_i(\lambda)|$ when i is greater than 1 but too small to be in the "tail-end", are the most significant remaining questions about enumerating tableaux inversions.

References

- [1] The On-Line Encyclopedia of Integer Sequences. <http://oeis.org/>
- [2] Paul Drube, *sl₃ skein modules of surfaces*, in preparation.
- [3] Lucas Fresse, *Betti numbers of Springer fibers in type A*, J. Algebra **322** (2009), 2566–2579.
- [4] William Fulton, *Young tableaux, with applications to representations theory and geometry*, Cambridge University Press, 1996.
- [5] Heather M. Russell, *The Bar-Natan skein module of the solid torus and the homology of (n, n) Springer varieties*, Geometriae Dedicata **142** (2009), 71–89.
- [6] Richard Stanley, *Enumerative Combinatorics Vol. 2*, Cambridge University Press, 1999.
- [7] Richard Stanley, *Enumerative Combinatorics Vol. 1*, Cambridge University Press, 2011.

A Inversion Tables

In this section we summarize some results from a computer program written by the authors to generate inverted Young tableaux. This is done by first generating all standard Young tableaux of a given shape. Each of the columns of each tableau have their elements permuted, and if the resulting Young tableau is row-standard, the number of inversions are counted and aggregated across every generated standard Young tableaux. We highlight some entries to draw attention to a potential bijection, given in Conjecture 4.4. In its current form, the program can handle Young tableaux of at most three rows. A version of the program is available [here](https://sites.google.com/a/valpo.edu/beagley/research/YoungTableauxInversions.java)², and questions about the program should be directed to the first author.

	(2, 2, 2)	(3, 2, 1)	
$m = 0$	5		
$m = 1$	16	16	$m = 0$
$m = 2$	25		
$m = 3$	24	24	$m = 1$
$m = 4$	14	14	$m = 2$
$m = 5$	5	5	$m = 3$
$m = 6$	1	1	$m = 4$
TOTAL	90	60	TOTAL

	(3, 3, 3)	(4, 3, 2)	
$m = 0$	42		
$m = 1$	168	168	$m = 0$
$m = 2$	330	366	$m = 1$
$m = 3$	414		
$m = 4$	357	357	$m = 2$
$m = 5$	222	222	$m = 3$
$m = 6$	103	103	$m = 4$
$m = 7$	35	35	$m = 5$
$m = 8$	8	8	$m = 6$
$m = 9$	1	1	$m = 7$
TOTAL	1680	1260	TOTAL

	(4, 4, 4)	(5, 4, 3)	
$m = 0$	462		
$m = 1$	2112	2112	$m = 0$
$m = 2$	4785	5643	$m = 1$
$m = 3$	7051	7161	$m = 2$
$m = 4$	7436		
$m = 5$	5951	5951	$m = 3$
$m = 6$	3773	3773	$m = 4$
$m = 7$	1937	1937	$m = 5$
$m = 8$	803	803	$m = 6$
$m = 9$	263	263	$m = 7$
$m = 10$	65	65	$m = 8$
$m = 11$	11	11	$m = 9$
$m = 12$	1	1	$m = 10$
TOTAL	34650	27720	TOTAL

	(5, 5, 5)	(6, 5, 4)	
$m = 0$	6006		
$m = 1$	30030	30030	$m = 0$
$m = 2$	75075	91520	$m = 1$
$m = 3$	123552	137137	$m = 2$
$m = 4$	148512	137163	$m = 3$
$m = 5$	138801		
$m = 6$	105495	105495	$m = 4$
$m = 7$	67158	67158	$m = 5$
$m = 8$	36297	36297	$m = 6$
$m = 9$	16667	16667	$m = 7$
$m = 10$	6448	6448	$m = 8$
$m = 11$	2065	2065	$m = 9$
$m = 12$	531	531	$m = 10$
$m = 13$	104	104	$m = 11$
$m = 14$	14	14	$m = 12$
$m = 15$	1	1	$m = 13$
TOTAL	756756	630630	TOTAL

²<https://sites.google.com/a/valpo.edu/beagley/research/YoungTableauxInversions.java>