PERMUTATION TABLEAUX AND PERMUTATION PATTERNS

EINAR STEINGRÍMSSON AND LAUREN K. WILLIAMS

Abstract. In this paper we introduce and study a class of tableaux which we call permutation tableaux; these tableaux are naturally in bijection with permutations, and they are a distinguished subset of the J-diagrams of Alex Postnikov [7, 9]. The structure of these tableaux is in some ways more transparent than the structure of permutations; therefore we believe that permutation tableaux will be useful in furthering the understanding of permutations. We give two bijections from permutation tableaux to permutations. The first bijection carries tableaux statistics to permutation statistics based on relative sizes of pairs of letters in a permutation and their places. We call these statistics *weak excedance statistics* because of their close relation to weak excedances. The second bijection carries tableaux statistics (via the weak excedance statistics) to statistics based on generalized permutation patterns. We then give enumerative applications of these bijections. For example, we give generating functions enumerating permutation tableaux according to their content, which therefore enumerate permutations according to the number of occurrences of various generalized patterns. We conclude our paper with a list of open problems.

1. Introduction

The aim of this article is to advertise a new class of tableaux together with two curious bijections for the study of permutations. We call these tableaux *permutation tableaux*; they are naturally in bijection with permutations, and are a distinguished subset of Alex Postnikov's J-diagrams [7], which were enumerated by the second author [9] because of their connection with the totally nonnegative part of the Grassmannian.

Recall that a *partition* $\lambda = (\lambda_1, \ldots, \lambda_k)$ is a weakly decreasing sequence of nonnegative integers. For a partition λ, where $\sum \lambda_i = m$, the *Young diagram* Y_λ of shape λ is a left-justified diagram of m boxes, with λ_i boxes in the i-th row.

We define a *permutation tableau* $\mathcal{T}_n^k$ to be a partition λ such that Y_λ is contained in a $k \times (n-k)$ rectangle, together with a filling of the boxes of Y_λ with 0's and 1's such that the following properties hold:

(1) Each column of the rectangle contains at least one 1.
(2) There is no 0 which has a 1 above it in the same column *and* a 1 to its left in the same row.

We call such a filling a *valid* filling of Y_λ. Observe that the requirement in (1) implies that the Young diagram must have $n-k$ columns, whereas the number of rows may be smaller than k.

Date: July 9, 2005.

Key words and phrases. Le-tableau, permutation patterns, permutation tableaux, q-analogs.

The first author was partially supported by the European Commission's IHRP Programme, grant HPRN-CT-2001-00272, "Algebraic Combinatorics in Europe".

If we forget the requirement (1) above we recover the definition of a J-diagram [7]. Figure 1 gives an example of a permutation tableau.

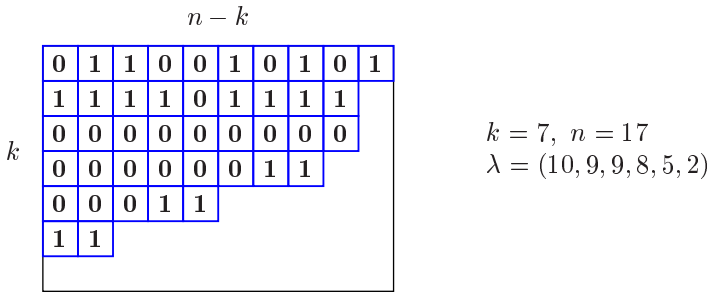

FIGURE 1. A permutation tableau

We will also think of a permutation tableau $\mathcal{T}_n^k$ as a $k \times (n-k)$ array of 0's, 1's, and 2's, by simply taking the previous description of a permutation tableaux and putting a 2 in every box of the rectangle which is not in Y_λ, as in Figure 2. We position the partition shape so that its top row lies at the top of the rectangle; therefore the 2's cut out a (rotated) Young diagram in the southeast corner of the rectangle.

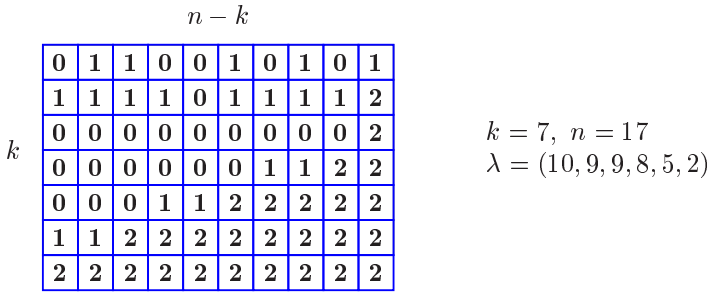

FIGURE 2. Another representation of a permutation tableau

Postnikov [8] has described a map that takes permutation tableaux contained in a $k \times (n-k)$ rectangle to permutations in $\mathcal{S}_n$ with k weak excedances. In this paper we give a simpler description of this map, prove that it is a bijection, and show that this map in fact preserves many more statistics. Namely, the numbers of 0's, 1's and 2's, respectively, in a permutation tableau equal certain linear combinations of certain statistics defined on the corresponding permutation. Each of these statistics counts pairs of indices (i, j) in a permutation, according to the relative sizes of the letters in those places and the place numbers themselves. These statistics were defined by Corteel in [5].

We then define another bijection, taking permutations to permutations and translating the statistics mentioned above into certain linear combinations of generalized permutation patterns. These combinations between them contain precisely half the generalized patterns of length 3 with one dash (see Section 4).

We conclude our paper by giving various enumerative applications of our bijections. The structure of permutation tableaux is in many ways more transparent than the structure of permutations, and hence lends itself more easily to enumeration. For example, by using our bijections together with results of the second author [9], we are able to give the entire distribution of permutations according to the number of occurrences of the generalized pattern $(2-31)$. This is the first such result for any pattern of length 3 (or more). However, this particular result, although first conjectured by the present authors, was first proved by Corteel [5], whose work provided us with a crucial piece of the puzzle solved here.

Another interesting consequence of the results presented here is that the statistic counting permutation tableaux according to the number of rows and number of 0's is an Euler-Mahonian statistic, that is, has the same distribution as the bistatistic on permutations consisting of the number of descents and the major index. It follows that if we define $D_{k,n}(p,q,r)$ to be the polynomial enumerating permutation tableaux $\mathcal{T}_n^k$ according to the number of 0's, 1's, and 2's, then $D_{k,n}(p,q,r)$ specializes to Carlitz' classical q-analog of the Eulerian numbers [4], as well as to the more recent q-analog of the Eulerian numbers that was studied in [9].

2. Bijection from Permutation Tableaux to Permutations

In this section we describe a bijection Φ from permutation tableaux to permutations. More precisely, Φ is a bijection from permutation tableaux $\mathcal{T}_n^k$ to permutations in the symmetric group $\mathcal{S}_n$ with k weak excedances. Here, a *weak excedance* of a permutation π is a value $\pi(i)$ such that $\pi(i) \geq i$. In this situation we say that i is a *weak excedance bottom* of π and that $\pi(i)$ is a *weak excedance top* of π. To make the notation less cumbersome, we abbreviate these as *wexbottoms* and *wextops*, respectively. *Non-weak excedance bottoms* and *non-weak excedance tops* are defined in the obvious way, in terms of i and $\pi(i)$ such that $\pi(i) < i$, and are abbreviated *non-wexbots* and *non-wextops*, respectively. The *number* of weak excedances in π will be denoted $\operatorname{wex}\pi$. Also, we let WEXBOTSUM be the sum of all the wexbottoms in π.

We remark that Postnikov [8] defined a map that is equal to Φ but his description was much more complicated and went through the intermediate step of *web diagrams*.

Before giving the bijection Φ, we must define the *diagram* $D(\mathcal{T}_n^k)$ associated with $\mathcal{T}_n^k$ as follows. Regard the south-east border of the partition Y_λ contained in the $k \times (n-k)$ rectangle as giving a path (the *partition path*) $P = \{P_i\}_{i=1}^n$ of length n from the northeast corner of the rectangle to the southwest corner of the rectangle: label each of the (unit) steps in this path with a number from 1 to n according to the order in which the step was taken. Then, remove the 0's from $\mathcal{T}_n^k$ and replace each 1 in $\mathcal{T}_n^k$ with a vertex. We will call the top vertex in each column a white vertex and all other vertices black vertices. Finally, from each vertex v, draw an edge to the east and an edge to the south; each such edge should connect v to either a closest vertex in the same row or column, or to one of the labels from 1 to n. The resulting picture is the *diagram* $D(\mathcal{T}_n^k)$. See Figure 3.

We now define the permutation $\pi = \Phi(\mathcal{T}_n^k)$ via the following procedure. For each $i \in \{1, \ldots, n\}$, find the corresponding position on $D(\mathcal{T}_n^k)$ which is labeled by i. If the label i is on a vertical step of P, start from this position and travel straight west as far as possible on edges of $D(\mathcal{T}_n^k)$. Then, take a "zig-zag" path southeast,

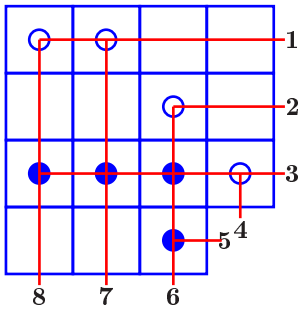

FIGURE 3. The diagram of a tableau. The topmost 1 in each column becomes a white vertex, and the other 1's become black vertices.

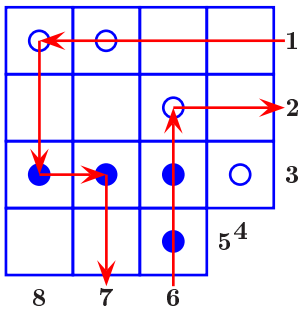

FIGURE 4. The paths taken by 1 and 6: $\pi(1) = 7$, $\pi(6) = 2$.

by traveling on edges of $D(\mathcal{T}_n^k)$ south and east and turning at each opportunity (i.e. at each new vertex). This path will terminate at some label $j \geq i$, and we let $\pi(i) = j$. If i is not connected to any edge (equivalently, if there are no vertices in the row of i) then we set $\pi(i) = i$. Similarly, if the label i is on a horizontal step of P, start from this position and travel north as far as possible on edges of $D(\mathcal{T}_n^k)$. Then, as before, take a zig-zag path south-east, by traveling on edges of $D(\mathcal{T}_n^k)$ east and south, and turning at each opportunity. This path will terminate at some label $j < i$, and we let $\pi(i) = j$.

See Figure 4 for a picture of the path taken by i.

The following three lemmas are clear from the construction above.

Lemma 1. *In $\Phi(\mathcal{T}_n^k)$, the letter i is a fixed point if and only if there is an entire row in $\mathcal{T}_n^k$ that has no 1's and whose right hand edge is labeled by i. In particular, $n, n-1, \ldots, n-m+1$ are fixed points in π if and only if the bottom m rows of $\mathcal{T}_n^k$ (in the $k \times (n-k)$ rectangle) consist entirely of 2's.*

Lemma 2. *Any directed step in a path on $D(\mathcal{T}_n^k)$ determines the path completely.*

It follows that $\Phi(\mathcal{T}_n^k)$ is a permutation.

As an example, if $\mathcal{T}_n^k$ is the permutation tableau whose diagram is given in Figures 3 and 4, then $\Phi(\mathcal{T}_n^k) = 74836215$.

Lemma 3. *The weak excedance bottoms of $\pi = \Phi(\mathcal{T}_n^k)$ are precisely the labels on the vertical edges of P. The non-weak excedance bottoms of π are precisely the labels on the horizontal edges of P. In particular, $\Phi(\mathcal{T}_n^k)$ is a permutation in $\mathcal{S}_n$ with precisely k weak excedances.*

Theorem 4. *The map Φ is a bijection from permutation tableaux to permutations.*

Proof. To prove that Φ is a bijection, we will give an explicit description of its inverse, again via an algorithmic procedure. First we need to define a *relative fixed point* in any *biword* $\begin{pmatrix} a_1 \dots a_n \\ b_1 \dots b_n \end{pmatrix}$, where the a_i's are distinct integers and the b_i's also. A relative fixed point is a pair (b_p, a_p) such that if a_p is the j-th smallest letter among the a_i then b_p is also the j-th smallest letter among the b_i.

For example, in $\begin{pmatrix} 5\,3\,1\,4 \\ 6\,2\,1\,3 \end{pmatrix}$, the pair $(2,3)$ is a relative fixed point, since each is the second smallest letter in its row.

Note that we will use the biword notation $\begin{pmatrix} a_1 \dots a_n \\ 1 \dots n \end{pmatrix}$ as an alternative representation of the permutation $a_1 \dots a_n$.

Let π be the permutation $\begin{pmatrix} a_1 \dots a_n \\ 1 \dots n \end{pmatrix}$. We now give the procedure for computing $\Phi^{-1}(\pi)$.

0. Compute the weak excedance bottoms of π to get the shape of the partition in $\mathcal{T}_n^k$ (see Lemma 3). Let $\tilde{\pi} = \pi$.
1. Check for relative fixed points in $\tilde{\pi}$. If (i, j) is a relative fixed point then fill with 0's the as yet undetermined part of the row corresponding to the weak excedance bottom i. Recompute $\tilde{\pi}$ by removing the relative fixed points.
2. Suppose we have determined the content of the k rightmost columns. Then look at the next column to the left, which is indexed by a non-excedance bottom r (that is, by the label on the horizontal step at the bottom of that column (Lemma 3)). Knowing that $r \to a_r$ in π uniquely determines the position p of the highest 1 in the column corresponding to r, since there is a unique zig-zag path going backwards (north-west) from a_r to a box in the column above i. Insert a 1 at that position and 0's in all boxes above it which are in the same column. Also, insert 1's into all undetermined boxes below p. (Note that we know that all nonzero boxes below position p must also be 1's because otherwise, if there were some 0 below the 1 then everything to its left would have to be a 0 also, implying that we had a relative fixed point in $\tilde{\pi}$.) Reduce $\tilde{\pi}$ by removing the column $\begin{pmatrix} a_r \\ r \end{pmatrix}$ from the biword for $\tilde{\pi}$. Go to step 1.

It is clear from this construction that our resulting tableau $\mathcal{T}_n^k$ will be a permutation tableau. And it is easy to check that the map described above is indeed the inverse of Φ. □

Example 5. Let $\pi = 514263$. Since π is in S_6 and has three weak excedances and $a_6 \neq 6$, our permutation tableau $\mathcal{T}_n^k$ will be contained in a 3×3 rectangle (the resulting tableau is shown in Figure 5). As in step 0, we want to first compute the shape of the associated partition. Since $1, 3, 5$ are the wexbottoms, and $2, 4, 6$ are the non-wexbottoms, this uniquely determines a path (the *partition path*) from

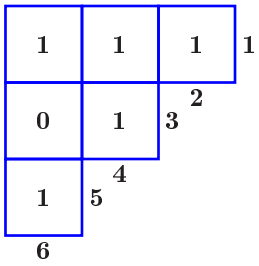

FIGURE 5. The permutation tableau for $\pi = 514263$

the northeast corner of the rectangle to the southwest corner of the rectangle with vertical steps in positions $1, 3, 5$ and horizontal steps in positions $2, 4, 6$. That is, our partition has the shape $(3, 2, 1)$. We now draw this partition, labeling the edges of its southeast border accordingly with the numbers $1, \ldots, 6$, and set $\tilde{\pi} = (5, 1, 4, 2, 6, 3)$.

Going to step 1, we see that $\tilde{\pi}$ has no relative fixed points.

Going to step 2, the fact that $2 \to 1$ in π implies that the rightmost column (which consists of a single box) contains a 1 in the top row. We now reduce the permutation $\tilde{\pi} = \begin{pmatrix} 5\,1\,4\,2\,6\,3 \\ 1\,2\,3\,4\,5\,6 \end{pmatrix}$ by removing $(2, 1)$, obtaining $\tilde{\pi} = \begin{pmatrix} 5\,4\,2\,6\,3 \\ 1\,3\,4\,5\,6 \end{pmatrix}$.

Going back to step 1, we see that there are no relative fixed points in $\tilde{\pi}$.

Going to step 2, since $4 \to 2$ in $\tilde{\pi}$ it is clear that the highest box in the column indexed by 4 must contain a 1. All undetermined boxes below this 1 must contain 1's also. We now reduce the permutation $\tilde{\pi} = \begin{pmatrix} 5\,4\,2\,6\,3 \\ 1\,3\,4\,5\,6 \end{pmatrix}$ by removing $4 \to 2$, obtaining $\tilde{\pi} = \begin{pmatrix} 5\,4\,6\,3 \\ 1\,3\,5\,6 \end{pmatrix}$.

Going back to step 1, we now see that $\tilde{\pi}$ has the relative fixed point $(3, 4)$. Therefore the undetermined part of the row corresponding to 3 consists of zeros. Now we reduce the permutation $\tilde{\pi} = \begin{pmatrix} 5\,4\,6\,3 \\ 1\,3\,5\,6 \end{pmatrix}$, obtaining $\tilde{\pi} = \begin{pmatrix} 5\,6\,3 \\ 1\,5\,6 \end{pmatrix}$.

Going to step 2, since $6 \to 3$ in $\tilde{\pi}$ the top box in the column corresponding to 6 has a 1. All undetermined boxes below that contain 1's. We have now filled in all columns of the tableau – obtaining the permutation tableau in Figure 5 – so we are done.

3. How Φ translates statistics

The six permutation statistics in the following definition will be related to the statistics recording the numbers of 0's, 1's and 2's in permutation tableaux. The first four of these refine Postnikov's definition of alignment [7] (see [9]); all of these statistics were defined by Corteel [5];

Definition 6. Given a permutation $\pi = a_1 a_2 \dots a_n$, let

$$\begin{aligned}
\mathrm{A_{EE}}(i) &= \{j \mid j < i \le a_i < a_j\},\\
\mathrm{A_{NN}}(i) &= \{j \mid a_j < a_i < i < j\},\\
\mathrm{A_{EN}}(i) &= \{j \mid j \le a_j < a_i < i\},\\
\mathrm{A_{NE}}(i) &= \{j \mid a_i < i < j \le a_j\},\\
\mathrm{C_{EE}}(i) &= \{j \mid j < i \le a_j < a_i\},\\
\mathrm{C_{NN}}(i) &= \{j \mid a_i < a_j < i < j\}.
\end{aligned}$$

We then set

$$\mathrm{A_{EE}}(\pi) = \sum_i |\mathrm{A_{EE}}(i)|,$$

and likewise for the other five statistics.

Observe that if we draw the permutation as a chord diagram on a circle, as in Figure 6, then $j \in A_{**}(i)$ means that the chords starting at i and j do not intersect and roughly "point in the same direction"; we will say that this is an *alignment of type* $\mathrm{A_{**}}$. And if $j \in \mathrm{C_{**}}(i)$ then the chords starting at i and j cross each other; we will say that this is a *crossing of type* $\mathrm{C_{**}}$. Note that the subscripts in our notation refer to whether the positions i and j are wexbottoms or non-wexbottoms of the permutation. For example, in Figure 6, the chords beginning at 3 and 5 form an alignment of type $\mathrm{A_{NE}}$, and the chords beginning at 2 and 4 form a crossing of type $\mathrm{C_{NN}}$.

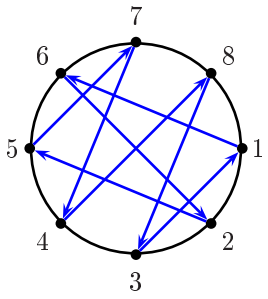

FIGURE 6. A chord diagram for the permutation 65187243

Theorem 7. *Let $T(d, a, b, c)$ be the set of permutation tableaux with d rows, $(n-d)$ columns, a 0's, b 1's and c 2's. Let $M(d, a, b, c)$ be the set of all permutations $\pi \in \mathcal{S}_n$ with*

- $d = \mathrm{wex}(\pi)$,
- $a = \mathrm{A_{EE}}(\pi) + \mathrm{A_{NN}}(\pi) + \mathrm{A_{EN}}(\pi)$,
- $b = \mathrm{C_{EE}}(\pi) + \mathrm{C_{NN}}(\pi)$,
- $c = \mathrm{A_{NE}}(\pi)$.

Then $|T(d, a, b, c)| = |M(d, a, b, c)|$.

To prove the above result, we will show that Φ is a bijection from $T(d, a, b, c)$ to $M(d, a, b, c)$. The structure of the argument is as follows:

- First we prove, in Proposition 8, that under Φ, the c $2's$ in $\mathcal{T}_n^k$ correspond to exactly c alignments of type $\mathrm{A_{NE}}$ in π.
- Second, in Proposition 9, we prove that $\mathrm{C_{EE}}(\pi)+\mathrm{C_{NN}}(\pi)$ equals the number of black vertices in the diagram D. This implies that $\mathrm{C_{EE}}(\pi) + \mathrm{C_{NN}}(\pi) = b - (n-k)$.
- We complete the proof with the following argument. Let $m = \mathrm{A_{EE}}(\pi) + \mathrm{A_{NN}}(\pi)+\mathrm{A_{EN}}(\pi)$. We know that $a+b+c = k(n-k)$. Corteel [5] proved that $\mathrm{A_{EE}}(\pi)+\mathrm{A_{NN}}(\pi)+\mathrm{A_{EN}}(\pi)+\mathrm{A_{NE}}(\pi)+\mathrm{C_{EE}}(\pi)+\mathrm{C_{NN}}(\pi) = (k-1)(n-k)$. Therefore $m+c+b-(n-k) = (k-1)(n-k)$, which implies that $m+c+b = k(n-k) = a+b+c$, and hence $m = a$.

Proposition 8. *If $\Phi(\mathcal{T}_n^k) = \pi$ then the number of 2's in $\mathcal{T}_n^k$ is equal to the number of alignments of type $\mathrm{A_{NE}}$ in π.*

Proof. Recall that if $\pi = \Phi(\mathcal{T}_n^k)$ then the wexbottoms and the non-wexbottoms of π correspond to the labels of the vertical and horizontal steps, respectively, in the south east border of the partition underlying $\mathcal{T}_n^k$. Note that the position of every 2 in $\mathcal{T}_n^k$ can be given by specifying the label of the edge above it and the edge to its left. The label i of the edge above it will be a non-wexbottom, and the label $j > i$ of the edge to its left will be a weak excedance bottom. Since $j > i$, and j is a wexbottom, and i is a non-wexbottom, the pair (i, j) is precisely an alignment of type $\mathrm{A_{NE}}$. Conversely, any alignment of type $\mathrm{A_{NE}}$ is a pair (i, j) where $i < j$, and i is a non-wexbottom, and j is a wexbottom. This implies that i is the label of a horizontal step, and j is the label of a vertical step. The fact that $i < j$ implies that the box of the tableau indexed as above by i and j contains a 2. □

Proposition 9. *Under the bijection Φ, there is a one-to-one correspondence between black vertices in the diagram of the permutation tableau, and crossings of types $\mathrm{C_{EE}}$ and $\mathrm{C_{NN}}$ in the permutation.*

Proof. Recall that black vertices correspond to those 1's in a tableau that are not topmost in their columns. Let D be the diagram of $\mathcal{T}_n^k$, and let $\pi = (a_1, \ldots, a_n)$ be $\Phi(\mathcal{T}_n^k)$. We will construct a map ϕ (induced by Φ) which takes each crossing (i, j) (where $i < j$) of type $\mathrm{C_{EE}}$ or $\mathrm{C_{NN}}$ in π to a black vertex d in D, and show that this is a bijection. The map ϕ is defined as follows. Let (i, j) be a crossing of type $\mathrm{C_{EE}}$ or $\mathrm{C_{NN}}$. We claim that the paths $(i \to a_i)$ and $(j \to a_j)$ intersect in a unique edge. If that edge is horizontal, then let d be the left vertex of the edge. If that edge is vertical, then let d be the bottom vertex of the edge.

First we need to show that the paths $(i \to a_i)$ and $(j \to a_j)$ intersect in an edge. We will prove this when (i, j) is a crossing of type $\mathrm{C_{EE}}$; the proof for $\mathrm{C_{NN}}$ is similar. Since $i < j$ and $a_i < a_j$, it is clear that the paths must cross each other at least once.

Consider the first point x at which the path $(j \to a_j)$ intersects the path $(i \to a_i)$. We will show that the intersection here will contain an edge. Clearly this intersection must be in the zig-zag portion of the path $(i \to a_i)$. If we let $d_1, d_2, \ldots, d_t$ be the sequence of vertices encountered by the path $(i \to a_i)$ in its zig-zag portion, then, by construction of that path, there are no vertices in the diagram D between any d_r and d_{r+1}. Note that if the path $(j \to a_j)$ intersects the path $(i \to a_i)$ in only the point x (rather than an edge containing x), then it is easy to see—using condition (2) in the definition of permutation tableaux—that x must actually be a vertex in D, located between some d_r and d_{r+1}. This is a contradiction.

Next, we show that the paths $(i \to a_i)$ and $(j \to a_j)$ intersect in a *unique* edge. If the two paths were to intersect a second time (and they may indeed intersect again in a *vertex*), then this intersection must take place in the zig-zag portion of *both* paths. Such a point e of intersection must be approached via a south step by $(i \to a_i)$ and must be approached via an east step by $(j \to a_j)$. But then, according to the procedure defining Φ, the path $(i \to a_i)$ will immediately turn east, and the path $(j \to a_j)$ will immediately turn south. Therefore this intersection is not an edge intersection.

We have thus shown that Φ induces a well-defined map from crossings to black vertices. We will now show that this map is a bijection by constructing its inverse. Namely, to each black vertex in D we need to produce a crossing of type $\mathrm{C_{EE}}$ or $\mathrm{C_{NN}}$. We do this as follows. Given a black vertex d, there is a path $(i \to a_i)$ on D which enters d by going south, and then leaves d going east. (It is easy to see that such a path exists by tracing backwards through the algorithm that defined the map Φ.)

If the path $(i \to a_i)$ is an excedance, then consider the unique path $(j \to a_j)$ which enters d traveling west. This path must be a weak excedance, as it is only the paths of weak excedances which contain steps to the west. Moreover, (i, j) must form a crossing of type $\mathrm{C_{EE}}$, since the two paths intersect in an edge (and we have seen that two paths which are both weak excedances may not intersect in an edge more than once).

On the other hand, if the path $(i \to a_i)$ is a non-excedance, then consider the unique path $(j \to a_j)$ which enters d traveling north. Clearly this path must be a non-excedance, as it is only the paths of non-excedances which contain steps north. Moreover, (i, j) must form a crossing of type $\mathrm{C_{NN}}$, since the two paths must intersect in a unique edge.

Therefore ϕ is a bijection between the set of $\mathrm{C_{EE}}$- and $\mathrm{C_{NN}}$-crossings in π, and the set of black vertices in D. □

This completes the proof of Theorem 7.

4. Permutation patterns

In this section we introduce necessary terminology and definitions that will be used in the next section, where we construct a bijection $\Psi : \mathcal{S}_n \longrightarrow \mathcal{S}_n$. This bijection proves the equidistribution of certain linear combinations of the statistics in Definition 6 (alignments and crossings) with certain linear combinations of generalized permutation patterns, which we define below. The composition of Ψ and the bijection Φ from Section 2 then proves the equidistribution of our tableaux statistics (numbers of 0's, 1's and 2's) with the pattern statistics to be defined here.

A *classical permutation pattern* $p = p_1 p_2 \dots p_k$ is simply a permutation, and an *occurrence* of p in a permutation $\pi = a_1 a_2 \dots a_n$ is a subsequence $a_{i_1}, a_{i_2}, \dots, a_{i_k}$ of π (where $i_1 < i_2 < \cdots < i_k$) whose letters are in the same relative order as in p. For example, the permutation 416235 has two occurrences of the pattern $(2-3-1)$, namely the subsequences 462 and 463.

In the literature, the pattern $(2-3-1)$ is usually denoted simply by 231. We write it here with dashes between consecutive letters in order to emphasize that there are no restrictions on the distance between the letters in a permutation that form an occurrence of the pattern. A *generalized pattern* is a pattern where some pairs of adjacent letters may lack a dash between them. Such an absence indicates

that the corresponding letters must be adjacent in an occurrence of the pattern in a permutation. For example, the pattern $(2-31)$ occurs only once in 416235, namely as 462. In the subsequence 463, whose letters are in the same relative order as those of $(2-31)$, the last two letters are not adjacent in 416235 as required for an occurrence of $(2-31)$.

The reason for writing patterns in parentheses is that we will consider them as *functions* from the set of permutations to the natural numbers, where the value of a pattern p on a permutation π is the number of different occurrences of p in π. For example, if $\pi = 416235$, as above, then $(2-3-1)\pi = 2$ and $(2-31)\pi = 1$.

It is easy to see that there are exactly twelve different patterns of length 3 with one dash. Six of these will be considered here, namely $(1-32)$, $(2-31)$, $(3-21)$, $(21-3)$, $(31-2)$ and $(32-1)$. These are all the patterns of length 3 with one dash whose two letters *not* separated by a dash are in decreasing order.

A *descent* in a permutation $\pi = a_1 a_2 \dots a_n$ is an i such that $a_i > a_{i+1}$. We say that a_i is a *descent top* and a_{i+1} a *descent bottom*. The set of descent tops is denoted DESTOPSET, and the set of descent bottoms DESBOTSET. Moreover, we let DESTOPSUM be the sum of the elements of DESTOPSET, and likewise for DESBOTSUM.

We now define the linear combinations of patterns whose joint distribution on permutations matches the distribution of 0's, 1's and 2's on permutation tableaux.

Definition 10. Given a permutation π, let

$$\begin{aligned} a(\pi) &= (21-3)\pi + (3-21)\pi + (31-2)\pi - \binom{\operatorname{des}\pi}{2}, \\ b(\pi) &= (2-31)\pi + n - 1 - \operatorname{des}\pi, \\ c(\pi) &= (1-32)\pi + (32-1)\pi - \binom{\operatorname{des}\pi}{2}. \end{aligned}$$

It is important to note that since we will be considering the quadruple statistic consisting of a, b, c and the number of descents, the terms $\binom{\operatorname{des}\pi}{2}$ and $n-1-\operatorname{des}\pi$ in the above definition only effect a *shift*, by a constant, of the statistics involved, but not an essential modification.

It is also important to note that the sum $(1-32)\pi + (32-1)\pi$ in c is equal to $\text{DESBOTSUM}\,\pi - \operatorname{des}\pi$. Namely, for each descent $\dots yx \dots$ in π, the pattern $(1-32)$ counts the letters to the left of the descent that are smaller than its descent bottom (x). The letters to the right of the descent, and smaller than x, are counted by $(32-1)$, so clearly we are counting *all* letters in π that are smaller than x. Analogously, the sum $(21-3)\pi + (3-21)\pi$ in a equals the sum of $n-t$ over all descent tops t in π.

This leaves $(31-2)$ in a, which sums the *left embracing numbers* in π, so called because $(31-2)$ counts, for each letter x in π, the descents to the left of x that embrace x, that is, where the letters of the descent are larger and smaller, respectively, than x. Analogously, the pattern $(2-31)$ in b sums the *right embracing numbers* in π.

To be more precise, we define the right embracing number of each letter ℓ in π, denoted $\text{REMBR}(\ell)$, as the number of descents $\dots yx \dots$ to the right of ℓ in π such that $x < \ell < y$.

Lemma 11. *Let* des *be the number of descents in a permutation π, and let $a(\pi)$, $b(\pi)$ and $c(\pi)$ be as above. Then*

$$a(\pi) + b(\pi) + c(\pi) = (\mathrm{des}+1)(n - \mathrm{des} - 1).$$

Proof. Each of the patterns involved in $a + b + c$ counts certain letters to the left or to the right of each descent in π. Together they count, for each descent in π, all the letters in π not belonging to the descent itself. There are, of course, $n-2$ such letters for each descent. Thus, the sum of all the patterns in $a(\pi) + b(\pi) + c(\pi)$ is $\mathrm{des} \cdot (n-2)$. Completing the proof now only requires a routine calculation. □

5. Another bijection

We now describe the construction of a bijection $\Psi : \mathcal{S}_n \longrightarrow \mathcal{S}_n$ that takes a permutation π to a permutation τ such that the set of descent tops in π determines the set of weak excedance tops in τ and the set of descent bottoms in π determines the set of weak excedance bottoms in τ. Moreover, the right embracing number of i in π becomes $\mathrm{C}_{\mathrm{EE}}(i)$ in τ if i is a wexbottom in τ and becomes $\mathrm{C}_{\mathrm{NN}}(i)$ in τ otherwise.

Recall the biword notation for permutations. For example, we write the permutation 31524 as

$$\begin{pmatrix} 3\,1\,5\,2\,4 \\ 1\,2\,3\,4\,5 \end{pmatrix}.$$

In order to construct $\Psi(\pi)$ (where $\pi \in \mathcal{S}_n$), we first construct two biwords, $\begin{pmatrix} f' \\ f \end{pmatrix}$ and $\begin{pmatrix} g' \\ g \end{pmatrix}$, and then form the biword $\tau' = \begin{pmatrix} f'\,g' \\ f\,g \end{pmatrix}$ by concatenating f and g, and f' and g', respectively. The words f, f', g, g' are defined as follows (we will prove later, in Theorem 12, that this is possible):

- The letters of f consist of the set obtained by adding 1 to each of the descent bottoms in π and then adjoining the letter 1. The letters of f are ordered increasingly.
- The letters of g consist of the set obtained from the non-descent bottoms in π by removing the letter n and adding 1 to the remaining letters. The letters of g are ordered increasingly.
- The letters of f' consist of the set obtained by subtracting 1 from each of the descent tops of π and then adjoining the letter n. The letters of f' are ordered so that, for a in f', $\mathrm{C}_{\mathrm{EE}}(a)$ in τ is the right embracing number of a in π. (Observe that $\mathrm{C}_{\mathrm{EE}}(a)$ only depends on the relative order of the wextops in τ, together with their corresponding wexbottoms.)
- The letters of g' consist of the set obtained by removing 1 from the set of non-descent tops in π and then subtracting 1 from the remaining letters. The letters of g' are ordered so that, for a in g', $\mathrm{C}_{\mathrm{NN}}(a)$ is the right embracing number of a in π.

Rearranging the columns of τ', so that the bottom row is in increasing order, we obtain the permutation $\tau = \Psi(\pi)$ as the top row of the rearranged biword. Before we prove that this can always be done in the way described, we give an example.

Let $\pi = 215896374$. Then π has

```
Descent bottoms:  1 3 4 6     Non-descent bottoms:  2 5 7 8 9

Descent tops:     2 6 7 9     Non-Descent tops:     1 3 4 5 8
```

The right embracing numbers are 2 for 5, 1 for 6 and 8, and 0 for all others:

```
21-5-8-963-74
  2 1  1
```

We construct a permutation with the corresponding wexbottoms and wextops, and the corresponding nonzero values for C_{EE} and C_{NN}, that is, with $C_{EE}(5) = 2$ and $C_{NN}(6) = C_{NN}(8) = 1$. First, the wexbottoms are obtained by adding 1 to each descent bottom, and adjoining 1, which is always a wexbottom. The wextops are obtained by subtracting 1 from the descent tops, and adjoining n, which is always a wextop. Thus, we get

```
Wexbottoms:  (1) 2 4 5 7       Non-wexbots: 3 6 8 9

Wextops:      1 5 6 8 (9)      Non-wextops: 2 3 4 7
```

We construct the permutation in two parts, one for the weak excedances, the other for the non-weak excedances.

Now, the definitions of C_{EE} and C_{NN} are such that C_{EE} only applies to pairs of weak excedances, and C_{NN} only to pairs of non-weak excedances. We first construct the weak excedance part of the permutation, by deciding where to place each of the wextops, in the places given by the wexbots:

```
_ _ _ _ _
1 2 4 5 7
```

We start from the right, in place 7, which has a 0 associated to it (since $\text{REMBR}(7) = 0$ in π). We need to put there the smallest number among the wextops that is at least as large as 7 (otherwise, $C_{EE}(7)$ would exceed 0 in the resulting permutation). This is the number 8:

```
        8
_ _ _ _ _
1 2 4 5 7
```

This leaves the wextops 1, 5, 6, 9. The next place, 5, has a 2 associated to it (since $\text{REMBR}(5) = 2$ in π), so we have to put there a wextop that it is bigger than exactly two of the remaining wextops that are at least as big as 5. This forces us to make this 9 (and the two remaining wextops between 5 and 9 in size are 5 and 6).

We continue in this way until we have placed all the wextops, in such a way that the values of C_{EE} for the remaining places are 0, since 5 is the only letter among the wexbottoms here with a nonzero right embracing number in π:

```
1 6 5 9 8
_ _ _ _ _
1 2 4 5 7
```

The non-wex part is done in a similar way, but starting from the left, and we get:

```
2 3 4 7
```

- - - -
3 6 8 9

Observe that $C_{NN}(6) = C_{NN}(8) = 1$ and $C_{NN}(3) = C_{NN}(9) = 0$, as required. Concatenating these two biwords, and sorting the columns to get the bottom row in increasing order, the permutation we obtain is $\Psi(215896374) = 162593847$.

We now prove that the above procedure can always be carried out in the way described.

Theorem 12. *Let* $\mathrm{DB}'(\pi)$ *be the set obtained from* $\mathrm{DESBOTSET}(\pi)$ *by adding 1 to each of its elements, and adjoining the letter 1.*

Let $\mathrm{DT}'(\pi)$ *be the set obtained from* $\mathrm{DESTOPSET}(\pi)$ *by subtracting 1 from each of its elements, and adjoining the letter* n*.*

For a permutation τ *let* $\mathrm{WB}(\tau)$ *be the set of weak excedance bottoms of* τ *and let* $\mathrm{WT}(\tau)$ *be the set of weak excedance tops of* τ*.*

The map Ψ *described above is well defined, and has the following properties, where* $\tau = \Psi(\pi)$*:*

(i) $\mathrm{WB}(\tau) = \mathrm{DB}'(\pi)$,
(ii) $\mathrm{WT}(\tau) = \mathrm{DT}'(\pi)$,
(iii) $C_{EE}(\tau) + C_{NN}(\tau) = \mathrm{REMBR}(\pi)$.

Moreover, Ψ *is a bijection.*

Proof. Recall that $C_{EE}(i) = 0$ unless i is an excedance bottom, and that $C_{NN}(i) = 0$ unless i is a non-excedance bottom.

Let the letters of $\mathrm{WB}(\tau)$ be $b_1, b_2, \ldots, b_\ell$, ordered so that $b_\ell < \cdots < b_2 < b_1$. Look at the largest letter in $\mathrm{WB}(\tau)$, that is, b_1. Suppose the embracing number of b_1 in π is e_1. Then there are at least e_1 descent tops in π that are larger than b_1. Thus, by the construction of $\mathrm{WT}(\tau)$ from the descent top set of π, there are at least e_1+1 elements x in $\mathrm{WT}(\tau)$ such that $b_1 \le x$. So, we can find an element t_1 in $\mathrm{WT}(\tau)$ such that $\mathrm{WT}(\tau)$ contains precisely e_1 elements x satisfying $b_1 \le x \le t_1$. Setting $\tau(b_1) = t_1$ guarantees that $C_{EE}(b_1) = e_1$ in τ.

Look next at b_2, the second largest element in $\mathrm{WB}(\tau)$. Suppose its embracing number in π is e_2. There are then at least $e_2 + 1$ elements x in $\mathrm{WT}(\tau)$ such that $b_2 \le x$. However, one of these elements is t_1, which has already been placed to the right of place b_2 in τ, and so t_1 cannot contribute to $C_{EE}(b)$ in τ. But, $b_1 + 1$ is a descent bottom in π and so its corresponding descent top, d, must be larger than $b_1 + 1$ and hence larger than b_2. Thus, b_2 cannot be embraced by the descent $\ldots d(b_1+1)\ldots$ in π. Hence, the embracing number of b_2 in π can be at most one less than the number of elements x in $\mathrm{WT}(\tau)$ satisfying $b_2 \le x$. We can therefore find an element $t_2 \neq t_1$ in $\mathrm{WT}(\tau)$ such that precisely e_2 of the elements x in $\mathrm{WT}(\tau)$ apart from t_1 satisfy $b_2 \le x \le t_2$.

An analogous argument shows that the embracing number of b_i in π can be at most $N + 1 - i$, where N is the number of elements x in $\mathrm{WT}(\tau)$ with $b_i \le x$. We can thus place each of the elements t_i of $\mathrm{WT}(\tau)$ in τ so that $C_{EE}(b_i)$ in τ equals $\mathrm{REMBR}(b_i)$ in π.

In particular, each placement according to the above algorithm will result in the creation of a weak excedance. Namely, clearly the k-th largest wexbottom is smaller than or equal to the k-th largest wextop. Thus, by induction, since we consider

the wexbots in decreasing order, the largest wextop unused at each stage of the algorithm is greater than or equal to the wexbottom being considered.

To construct the subword of τ consisting of non-wextops, we proceed in a similar way, except that we start from the smallest non-wexbottom. At each stage, for the non-wexbottom b_i we find a non-wextop d that satisfies $d < x < b_i$ for precisely e elements x among the remaining non-wextops, where $e = \textsc{rembr}(b_i)$ in π. The argument showing that this is always possible, and that each placement results in a non-weak excedance, is analogous to the case of the weak excedance subword, and is omitted.

To prove that Ψ is a bijection, it suffices to show that it is injective, since it is a map from $\mathcal{S}_n$ to itself. Let σ_1 and σ_2 be two permutations with $\Psi(\sigma_1) = \Psi(\sigma_2)$. From the definition of Ψ it is clear that σ_1 and σ_2 must have the same descent tops and descent bottoms and also the same right embracing numbers for each letter. It follows from the proof of Theorem 4 in [2, p. 249] that a permutation is uniquely determined by its sets of descent bottoms and tops, respectively, together with the right embracing numbers of its letters. Thus, we must have $\sigma_1 = \sigma_2$.

In fact, the proof of Theorem 4 in [2] can be applied directly to our situation with trivial modifications, and yields a description of the inverse of Ψ. □

Recall that $\textsc{wexbotsum}(\pi)$ is the sum of all the wexbottoms in π. The following corollary of Theorem 12 requires only straightforward calculations.

Corollary 13.

$$\begin{aligned} \textsc{wextopsum}\,\Psi(\pi) &= \textsc{destopsum}\,\pi + n - \operatorname{des}\pi, \\ \textsc{wexbotsum}\,\Psi(\pi) &= \textsc{desbotsum}\,\pi + \operatorname{des}\pi + 1. \end{aligned}$$

We will use the following two lemmas, the first one proved by Corteel [5], in our proofs of the equidistribution results between our tableaux statistics and permutation statistics.

Lemma 14 (Corteel [5])**.** *Let* wex, $\mathrm{A_{EE}}$, $\mathrm{A_{NN}}$, $\mathrm{A_{EN}}$, $\mathrm{A_{NE}}$, $\mathrm{C_{EE}}$, $\mathrm{C_{NN}}$ *be as above. Then*

$$\mathrm{A_{EE}} + \mathrm{A_{NN}} + \mathrm{A_{EN}} + \mathrm{A_{NE}} + \mathrm{C_{EE}} + \mathrm{C_{NN}} = (\mathrm{wex} - 1)(n - \mathrm{wex}).$$

Lemma 15.

$$\mathrm{A_{EN}} = \binom{n}{2} - \binom{n - \mathrm{wex}}{2} + \mathrm{wex} - \textsc{wextopsum}, \tag{1}$$

$$\mathrm{A_{NE}} = \textsc{wexbotsum} - \binom{\mathrm{wex}}{2}. \tag{2}$$

Proof. Equation (1) is equivalent to

$$\mathrm{A_{EN}} + \textsc{wextopsum} - \mathrm{wex} = \binom{n}{2} - \binom{n - \mathrm{wex}}{2}.$$

We will show that the sum in the left-hand-side above counts all pairs (i, j), with $1 \leq i < j \leq n$, except those for which neither of i and j is a weak excedance.

Recall that $\mathrm{A_{EN}}$ counts the pairs (i, j) such that $j \leq a_j < a_i < i$. Each such pair can be described as consisting of a wextop w in the permutation, and a non-wextop that is larger than w and to the right of w.

We can interpret $(\textsc{wextopsum} - \mathrm{wex})$ as the sum, over all wextops, of the size of the wextop, minus 1. Counting this for each wextop w can be done by counting

all the letters in the permutation that are strictly smaller than w. Doing this for all wextops is equivalent to counting all pairs of letters in the permutation that either consist of two wextops, or a wextop and a non-wextop, where the wextop is the larger of the two.

Therefore A_{EN} and $(\textsc{wextopsum} - \mathrm{wex})$ together count all pairs of letters in the permutation, *except* those consisting of two non-wextops. (Observe that it is impossible to have a non-wextop z and a wextop w such that z is left of w and $z > w$.) The total number of pairs of letters in a permutation in $\mathcal{S}_n$ is of course $\binom{n}{2}$, and the number of pairs of non-wexbots is $\binom{n-\mathrm{wex}}{2}$, which completes the proof.

Equation 2 can be proved in a similar manner. □

We can now prove the main results about the equidistribution implied by the bijection Ψ.

Theorem 16. *Let $\sigma = \Psi(\pi)$. We have*

$$\mathrm{des}\,\pi = \mathrm{wex}\,\sigma - 1, \tag{3}$$

$$(31\text{-}2)\pi = \mathrm{A}_{\mathrm{EE}}\sigma + \mathrm{A}_{\mathrm{NN}}\sigma, \tag{4}$$

$$(21\text{-}3)\pi + (3\text{-}21)\pi - \binom{\mathrm{des}\,\pi}{2} = \mathrm{A}_{\mathrm{EN}}\sigma, \tag{5}$$

$$(2\text{-}31)\pi = \mathrm{C}_{\mathrm{EE}}\sigma + \mathrm{C}_{\mathrm{NN}}\sigma, \tag{6}$$

$$(1\text{-}32)\pi + (32\text{-}1)\pi - \binom{\mathrm{des}\,\pi}{2} = \mathrm{A}_{\mathrm{NE}}\sigma. \tag{7}$$

Proof. Equations (3) and (6) follow directly from Theorem 12, since $(2\text{-}31)\pi$ is the sum of the right embracing numbers for all the letters in π. We will prove (5) here; the proof of (7) is analogous and is omitted. Having done this, Equation (4) follows from the other four identites in the present theorem, together with Lemmas 14 and 11 and routine calculations.

To prove Equation (5), observe that

$$(21\text{-}3)\pi + (3\text{-}21)\pi = n \cdot \mathrm{des}\,\pi - \textsc{destopsum}\,\pi.$$

This is because $(21\text{-}3)\pi + (3\text{-}21)\pi$ counts the letters in π larger than the descent top b for each descent $\ldots ba \ldots$ in π. According to Corollary 13, the right-hand-side in the equation above can be rewritten as follows:

$$n \cdot \mathrm{des}\,\pi - \textsc{destopsum} = n \cdot \mathrm{des}\,\pi - \textsc{wextopsum}\,\sigma + n - \mathrm{des}\,\pi.$$

By Lemma 15, this is equal to

$$n \cdot \mathrm{des}\,\pi + \left(\mathrm{A}_{\mathrm{EN}}\sigma - \mathrm{wex}\,\sigma - \binom{n}{2} + \binom{n - \mathrm{wex}\,\sigma}{2}\right) + n - \mathrm{des}\,\pi,$$

which, in turn, is equal to

$$\mathrm{A}_{\mathrm{EN}}\sigma + n \cdot \mathrm{des}\,\pi - (\mathrm{des}\,\pi + 1) - \binom{n}{2} + \binom{n - (\mathrm{des}\,\pi + 1)}{2} + n - \mathrm{des}\,\pi.$$

To show that this last expression is equal to $\mathrm{A}_{\mathrm{EN}}\sigma + \binom{\mathrm{des}\,\pi}{2}$ is straightforward. □

Note that Equations (4) and (5) together imply that

$$(31\text{-}2)\pi + (21\text{-}3)\pi + (3\text{-}21)\pi - \binom{\mathrm{des}}{2} = \mathrm{A}_{\mathrm{EE}}\sigma + \mathrm{A}_{\mathrm{NN}}\sigma + \mathrm{A}_{\mathrm{EN}}\sigma.$$

This, together with Theorem 7, leads to the following corollary.

Corollary 17. *Let $T(d,a,b,c)$ be the set of permutation tableaux with $d+1$ rows and $(n-d-1)$ columns, which are filled with precisely a 0's, b 1's and c 2's. Let $P(d,a,b,c)$ be the set of all permutations $\pi \in \mathcal{S}_n$, such that*

- $d = \mathrm{des}(\pi)$,
- $a = [(31\text{-}2) + (21\text{-}3) + (3\text{-}21)]\,\pi - \binom{\mathrm{des}\,\pi}{2}$,
- $b = (2\text{-}31)\pi + n - 1 - \mathrm{des}\,\pi$,
- $c = [(1\text{-}32) + (32\text{-}1)]\,\pi - \binom{\mathrm{des}\,\pi}{2}$.

Then $|T(d,a,b,c)| = |P(d,a,b,c)|$.

6. Enumeration results

One nice application of permutation tableaux is that they facilitate enumeration of permutations according to various statistics. This is because permutation tableaux satisfy a rather simple recurrence, which we now explain.

Fix a partition $\lambda = (\lambda_1, \ldots, \lambda_k)$. Let $F_\lambda(p,q)$ be the polynomial in p and q such that the coefficient of $p^s q^t$ is the number of valid fillings of the Young diagram Y_λ which contain s 0's and t 1's. As Figure 7 illustrates, there is a simple recurrence for $F_\lambda(p,q)$.

Explicitly, any valid filling of λ is obtained in one of the following ways:

- inserting a column whose bottom entry is 1 and whose other entries are 0 after the $(\lambda_k - 1)$st column of a valid filling of $(\lambda_1 - 1, \lambda_2 - 1, \ldots, \lambda_k - 1)$;
- adding a 1 to the end of the bottom row of a valid filling of the shape $(\lambda_1, \lambda_2, \ldots, \lambda_{k-1}, \lambda_k - 1)$;
- adding an all-zero row of length λ_{k-1} to a valid filling of $(\lambda_1, \ldots, \lambda_{k-1})$.

Thus, we have the following recurrence.

Proposition 18.

$$F_\lambda(p,q) = p^{k-1} q F_{(\lambda_1-1,\lambda_2-1,\ldots,\lambda_k-1)}(p,q) + q F_{(\lambda_1,\lambda_2,\ldots,\lambda_{k-1},\lambda_k-1)}(p,q) + p^{\lambda_k} F_{(\lambda_1,\ldots,\lambda_{k-1})}(p,q).$$

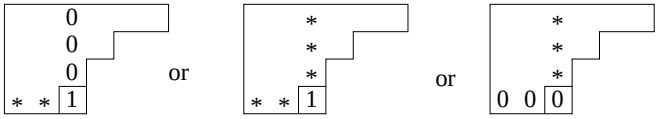

FIGURE 7. Recurrence for $F_\lambda(p,q)$

It is straightforward to compute $F_\lambda(p,q)$ when k (the number of rows of λ) is small. Here are the first two formulas.

Proposition 19.

$$F_{(\lambda_1)}(p,q) = q^{\lambda_1}.$$
$$F_{(\lambda_1,\lambda_2)}(p,q) = -q^{\lambda_1-1}p^{\lambda_2+1} + q^{\lambda_1-1}[2]_{p,q}^{\lambda_2+1}.$$

In the above expression, $[2]_{p,q}$ is the p,q-analog of 2. Recall that the p,q-analog of the number n is $p^{n-1} + p^{n-2}q + p^{n-3}q^2 + \cdots + q^{n-1}$, denoted $[n]_{p,q}$.

Let $D_{k,n}(p,q,r) := \sum_\lambda F_\lambda(p,q) r^{k(n-k)-|\lambda|}$, where λ ranges over all partitions contained in a $k \times (n-k)$ rectangle. By Theorem 7, $D_{k,n}(p,q,r)$ enumerates

permutations according to the number of weak excedances, several kinds of alignments, and crossings. And by Corollary 17, $D_{k,n}(p,q,r)$ enumerates permutations according to the number of descents and occurrences of various generalized patterns. Therefore it would be nice to get an explicit expression for $D_{k,n}(p,q,r)$, for example by solving explicitly for $F_\lambda(p,q)$ and then by summing over partitions λ contained in a $k \times (n-k)$ rectangle.

For fixed small k, it is not too difficult to compute the generating function $D_k(p,q,r,x) := \sum_n D_{k,n}(p,q,r)x^n$. Here are the first few formulas. Note that it is easy to determine what the denominator should be for $D_k(p,q,r,x)$, but the numerator is significantly more complicated.

Proposition 20.

$$D_1(p,q,r,x) = \frac{x}{1-qx}$$

$$D_2(p,q,r,x) = \frac{x^2}{(1-pqx)(1-qrx)(1-q[2]_{p,q}x)}$$

$$D_3(p,q,r,x) = \frac{x^3(1+pq^2x-p^3q^2rx^2-2p^2q^3rx^2-pq^4rx^2)}{(1-p^2qx)(1-pqrx)(1-qr^2x)(1-pq[2]_{p,q}x)(1-qr[2]_{p,q}x)(1-q[3]_{p,q}x)}$$

One can derive these formulas by either using the methodology outlined above (i.e. by summing $F_\lambda(p,q)$), or else by translating the problem of enumerating permutation tableaux into a problem about enumerating certain weighted lattice paths, and then by enumerating these lattice paths. In order to sketch the latter method, let us define a *bad* zero in a permutation tableau to be a 0 which lies directly underneath some 1. Note that if some column C in a permutation tableau $\mathcal{T}_n^k$ contains a bad zero in the rth row, then every column to the left of C must also contain a zero in the rth row.

In the lattice path method for enumeration of permutation tableaux, we associate to each permutation tableau $\mathcal{T}_n^k$ a weighted lattice path $L = \{L_i\}_{i=1}^n$ consisting of n steps in the plane, which must be of the following types: $(1,1)$ (a northeast step), $(1,0)$ (an east step), and $(1,-j)$, where $1 \le j \le k-1$ (a southeast step). Each step L_i in the lattice path represents the step P_i in the partition path $\{P_i\}_{i=1}^n$. (Recall that the partition path follows the shape of the partition Y_λ and travels from the northeast corner to the southwest corner of the $k \times (n-k)$ rectangle containing $\mathcal{T}_n^k$.) The steps $(1,1)$ in L correspond to vertical steps in the partition path, and have weight x. A step $(1,0)$ in L corresponds to a horizontal step in the partition path such that the corresponding column C of $\mathcal{T}_n^k$ does not introduce any bad zeros (except those that were forced by bad zeros to the right of C). Such a step has weight $p^a q^b r^c x$, where a, b, and c are the numbers of 0's, 1's, and 2's, respectively, in column C. (Note that $a+b+c=k$.) Finally, a step $(1,-j)$ in L corresponds to a horizontal step in the partition path such that the corresponding column C of $\mathcal{T}_n^k$ introduces exactly j new bad zeros (that were not forced by bad zeros in columns to the right of C). As before, such a step has weight $p^a q^b r^c x$, where a, b, and c are the numbers of 0's, 1's, and 2's, respectively, in column C. Observe that the height of any point in the lattice path L is equal to the number of boxes of the corresponding column of $\mathcal{T}_n^k$ which can be filled with either a 0 or a 1. By associating weighted

lattice paths to permutation tableaux in this way, we can facilitate computation of the generating functions $D_k(p,q,r,x)$ for small k.

Now we will give complete results about a certain specialization of $D_{k,n}(p,q,r)$. Let $E_{k,n}(q) := D_{k,n}(1,q,1)$. An explicit formula for $E_{k,n}(q)$ was found in [9]; the proof utilized a recurrence similar to that in Proposition 18.

Theorem 21 (Williams [9])**.**

$$E_{k,n}(q) = q^{n-k^2} \sum_{i=0}^{k-1} (-1)^i [k-i]^n q^{ki-k} \left(\binom{n}{i} q^{k-i} + \binom{n}{i-1} \right).$$

The polynomials above have many nice properties. It was observed in [9] that if one renormalizes $E_{k,n}(q)$ by defining $\hat{E}_{k,n}(q) := q^{k-n} E_{k,n}(q)$, then $\hat{E}_{k,n}(q)$ is a new q-analog of the Eulerian numbers (distinct from Carlitz' classical q-analog of the Eulerian numbers [4]). Furthermore, $\hat{E}_{k,n}(q)$ specializes at $q = -1, 0, 1$ to the binomial coefficients, the Narayana numbers, and the Eulerian numbers. Additionally, $\hat{E}_{k,n}(q) = \hat{E}_{n+1-k,n}(q)$. It was shown more recently by Corteel [5] that the polynomials $\hat{E}_{k,n}(q)$ naturally relate to the ASEP model in statistical physics.

Table 1 lists $\hat{E}_{k,n}(q)$ for $n = 4, 5, 6, 7$.

Theorem 21 together with Corollary 17 implies the following result.

Corollary 22. *The number of permutations in $\mathcal{S}_n$ with $k-1$ descents and m occurrences of the pattern $(2-31)$ is equal to the coefficient of q^m in*

$$\hat{E}_{k,n}(q) = q^{-k^2} \sum_{i=0}^{k-1} (-1)^i [k-i]^n q^{ki} \left(\binom{n}{i} q^{k-i} + \binom{n}{i-1} \right).$$

This result was first conjectured by the authors of this paper, and first proved by Corteel [5]. The formula $\hat{E}_{k,n}(q)$ is the first known polynomial expression which gives the *complete* distribution of a permutation pattern of length greater than 2 (the two cases of length 2 correspond to the Eulerian numbers and the coefficients of $[n]!$, respectively).

For the polynomials $\hat{E}_{k,n}(q)$, both ordinary and continued fraction generating functions are known. It can be shown [9] that $E(q,x,y) := \sum_{n,k} q^{n-k} \hat{E}_{k,n}(q) y^k x^n$ is equal to

$$\sum_{i=0}^{\infty} \frac{y^i (q^{2i+1} - y)}{q^{i^2+i+1} (q^i - q^{i+1}[i]x + [i]xy)}.$$

Additionally, Corteel [5] used results of Clark, Steingrímsson, and Zeng [2] to show the following:

Theorem 23 (Corteel [5])**.**

$$\hat{E}(q,x,y) = \cfrac{1}{1 - b_0 x - \cfrac{\lambda_1 x^2}{1 - b_1 x - \cfrac{\lambda_2 x^2}{1 - b_2 x - \cfrac{\lambda_3 x^2}{\ddots}}}},$$

$\hat{E}_{1,4}(q)$	1
$\hat{E}_{2,4}(q)$	$6+4q+q^2$
$\hat{E}_{3,4}(q)$	$6+4q+q^2$
$\hat{E}_{4,4}(q)$	1
$\hat{E}_{1,5}(q)$	1
$\hat{E}_{2,5}(q)$	$10+10q+5q^2+q^3$
$\hat{E}_{3,5}(q)$	$20+25q+15q^2+5q^3+q^4$
$\hat{E}_{4,5}(q)$	$10+10q+5q^2+q^3$
$\hat{E}_{5,5}(q)$	1
$\hat{E}_{1,6}(q)$	1
$\hat{E}_{2,6}(q)$	$15+20q+15q^2+6q^3+q^4$
$\hat{E}_{3,6}(q)$	$50+90q+84q^2+50q^3+21q^4+6q^5+q^6$
$\hat{E}_{4,6}(q)$	$50+90q+84q^2+50q^3+21q^4+6q^5+q^6$
$\hat{E}_{5,6}(q)$	$15+20q+15q^2+6q^3+q^4$
$\hat{E}_{6,6}(q)$	1
$\hat{E}_{1,7}(q)$	1
$\hat{E}_{2,7}(q)$	$21+35q+35q^2+21q^3+7q^4+q^5$
$\hat{E}_{3,7}(q)$	$105+245q+308q^2+259q^3+161q^4+77q^5+28q^6+7q^7+q^8$
$\hat{E}_{4,7}(q)$	$175+441q+588q^2+532q^3+364q^4+196q^5+84q^6+28q^7+7q^8+q^9$
$\hat{E}_{5,7}(q)$	$105+245q+308q^2+259q^3+161q^4+77q^5+28q^6+7q^7+q^8$
$\hat{E}_{6,7}(q)$	$21+35q+35q^2+21q^3+7q^4+q^5$
$\hat{E}_{7,7}(q)$	1

TABLE 1. $\hat{E}_{k,n}(q)$

where $b_n = y[n+1]_q + [n]_q$, $\lambda_n = y[n]_q^2$, *and* $[n]_q = 1+q+\cdots+q^{n-1}$.

Finally, we remark that, up to a shift by a constant, the bistatistic counting the number of rows and number of 0's in permutation tableaux has the Euler-Mahonian distribution, that is, it has the same distribution as the number of descents and the *major index* for permutations (the major index is the sum of the places of the descents in a permutation). Note first that the statistic a in Corollary 17, when stripped of $\binom{\mathrm{des}}{2}$ (which is a constant when considering a fixed number of descents), has the same distribution as the statistic

$$(1-32)+(32-1)+(2-31). \tag{8}$$

This is because the statistic in (8) is obtained by taking the reverse complement of the statistic $(31-2)+(21-3)+(3-21)$, that is, by reversing each of the patterns and then replacing each letter i by $4-i$. Doing the same with each permutation in $\mathcal{S}_n$ (with 4 replaced by $n+1$) is a bijection from $\mathcal{S}_n$ to itself, and this bijection clearly proves the equidistribution of $(31-2)+(21-3)+(3-21)$ with $(1-32)+(32-1)+(2-31)$, even when each statistic is taken jointly with the number of descents (which is invariant under reverse complement). The statistic $(1-32)+(32-1)+(2-31)$ is

equal to the statistic MAK, as pointed out in [1], and it was shown by Foata and Zeilberger [6] that (des, MAK) has the Euler-Mahonian distribution.[1]

Because the bistatistic counting number of rows and number of 0's has the Euler-Mahonian distribution, the specialization $D_{k,n}(p,1,1)$ is essentially Carlitz' classical q-analog of the Eulerian numbers [4]. Therefore the polynomials $D_{k,n}(p,q,r)$ generalize both the classical q-analog of the Eulerian numbers and the new q-analog of the Eulerian numbers found in [9].

7. Open problems

The bijections taking permutation tableaux to permutations seem to have several properties that merit further investigation. We list a few of these, along with some other open problems.

(1) Find an explicit expression for $D_{k,n}(p,q,r)$.

(2) We say that a 1 in a permutation tableau is *essential* if it is the topmost one in its column or the leftmost 1 in its row. A tableau is determined by its essential 1's: all the other 1's are determined by these, because of condition (2) in the definition of a permutation tableau. What do the essential 1's correspond to in the corresponding permutation?

We conjecture that the distribution of permutation tableau according to the number of essential 1's is equal to that for number of cycles in permutations. This distribution is the same as that for Left-to-Right minima. Moreover, we conjecture that the joint distribution of tableaux according to the number of rows and the number of essential 1's equals that of permutations according to descents and Left-to-Right-minima. That bistatistic, in turn, has the same distribution as the number of weak excedances and the number of cycles of a permutation, when written in standard cycle form.

(3) The number of 0's in a tableau corresponds to the total number of occurrences of the patterns $(3-21)$, $(21-3)$ and $(31-2)$. It is easy to see that these patterns have the same distributions as $(1-32)$, $(32-1)$ and $(2-31)$, respectively. To prove this, simply reverse each permutation in $\mathcal{S}_n$. Can we partition the 0's in a tableau into two sets, one corresponding to occurrences of $(3-21)+(21-3)$ and the other to occurrences of $(31-2)$? Observe that the first one of these sets would correspond to descent tops and the second one to left embracings. Thus, these sets would be symmetric counterparts of 2's and 1's respectively, although this symmetry is not transparent in the tableaux.

(4) The reflection of a permutation tableau $\mathcal{T}_n^k$ in its north-east/south-west diagonal yields a permutation tableau if and only if $\mathcal{T}_n^k$ has a 1 in each row. That is equivalent to the associated permutation being fixed point free. Which permutation is associated to the reflected tableau (that tableau is also fixed point free because it has a 1 in each row)?

(5) A permutation tableau T must have at least one 1 in each column. If it has only this minimum number of 1's, then the corresponding permutation, that is, $\Psi^{-1}(\Phi(T))$, has no occurrences of the pattern $(2-31)$. It has been

[1] Actually, the statistic $(1-32)+(32-1)+(2-31)$ is a slight variation on MAK, as defined by Foata and Zeilberger, but is easily seen to have the same distribution when taken together with the number of descents. Foata and Zeilberger's MAK is actually equal to the statistic called MAKL in [1, Table 1]

shown (see [3]), that permutations avoiding this pattern are enumerated by the Catalan numbers. Is there a bijection from these tableaux to any of the well known objects enumerated by Catalan numbers, such as Dyck paths?

(6) Find a better description of the bijection Ψ.

References

[1] E. Babson and E. Steingrímsson: Generalized permutation patterns and a classification of the Mahonian statistics, Sém. Lothar. Combin., **B44b** (2000), 18 pp.

[2] R. J. Clarke, E. Steingrímsson and J. Zeng: New Euler-Mahonian statistics on permutations and words, Adv. in Appl. Math. **18** (1997), 237–270.

[3] A. Claesson: Generalized pattern avoidance, Europ. J. Combin. **22** (2001), 961–971.

[4] L. Carlitz: q-Bernoulli and Eulerian numbers, Trans. Amer. Math. Soc. **76** (1954), 332–350.

[5] S. Corteel: Crossings and alignments of permutations, preprint.

[6] D. Foata and D. Zeilberger: Denert's permutation statistic is indeed Euler-Mahonian, Studies in Appl. Math. **83** (1990), 31–59.

[7] A. Postnikov: Webs in totally positive Grassmann cells, in preparation.

[8] A. Postnikov: personal communication.

[9] L. Williams: Enumeration of totally positive Grassmann cells, Advances in Math, **190** (2005), 319–342.

Reykjavik University, Reykjavik, Iceland, and Dept. of Mathematics, Chalmers Univ. of Technology, Göteborg, Sweden

E-mail address: einar@math.chalmers.se

Department of Mathematics, MIT, Cambridge, MA 02139

E-mail address: lauren@math.mit.edu