

SMIRNOV TREES

MATJAŽ KONVALINKA AND VASU TEWARI

ABSTRACT. We introduce a generalization of Smirnov words in the context of labeled binary trees, which we call Smirnov trees. We study the generating function for ascent-descent statistics on Smirnov trees and establish that it is e -positive, which is akin to the classical case of Smirnov words. Our proof relies on an intricate weight-preserving bijection.

CONTENTS

1. Introduction	1
2. Background	3
2.1. Words	4
2.2. Smirnov trees	4
3. Bijection and applications	6
3.1. Outline of the bijection	6
3.2. e -positivity and bleeding trees	7
3.3. Another functional equation and exponential specialization	11
4. The gory details	13
5. Final remarks	21
Acknowledgements	22
References	22

1. INTRODUCTION

The study of permutation statistics has been an active area of research since the seminal work of MacMahon [9]. A particular permutation statistic that plays a prominent role in algebraic combinatorics is the descent statistic. The associated integer sequence obtained by counting permutations according to their number of descents is the well-known sequence of Eulerian numbers. Rather than giving an exhaustive list of areas in mathematics where this sequence makes an appearance, we refer the reader to [11] for a beautiful book exposition.

2010 *Mathematics Subject Classification*. Primary 05E05, 20C08; Secondary 05A05, 05E10, 05E15, 06A07, 16T05, 20C30.

Key words and phrases. binary trees, Smirnov words, ascent-descent, e -positivity.

The first author acknowledges the financial support from the Slovenian Research Agency (research core funding No. P1-0294). The second author was supported by an AMS-Simons travel grant.

Given that permutations may be considered to be linear trees, it is natural to consider generalizations of classical permutation statistics in the context of labeled plane binary trees.

Gessel [6] was the first to study the analogue of the descent statistic for labeled binary trees, and he further pointed out intriguing connections to the enumerative theory of Coxeter arrangements. There has been a flurry of activity towards understanding these connections better, and the reader is referred to [2, 3, 4, 7, 19, 20] for more details. Gessel-Griffin-Tewari [7] investigated these connections from the perspective of symmetric functions; they attach a multivariate generating function tracking ascent-descents over all labeled binary trees and subsequently prove that this generating function expands positively in terms of ribbon Schur functions. This result is part of the motivation for our article. The rest of it stems from recent work of Shareshian-Wachs [13] on chromatic quasisymmetric functions, which we discuss next.

Stanley [16] introduced chromatic symmetric functions of graphs as a way to generalize chromatic polynomials of graphs. In the case where the graph has n nodes and no edges, the chromatic polynomial is the sum of \mathbf{x}_w where w ranges over all words of length n . Here, and throughout the rest of this article, set $\mathbf{x}_w := x_{w_1} \cdots x_{w_n}$ if $w = w_1 \dots w_n$, where $\mathbf{x} = \{x_1, x_2, \dots\}$ is an alphabet with commuting indeterminates. It is worth remarking that the refined version tracking the distribution of descents over the set of all words of length n is ribbon Schur-positive. This motivates studying an analogue of the descent statistic for general graphs in tandem with the chromatic symmetric function, and this was done by Shareshian and Wachs in [13], wherein they introduced chromatic quasisymmetric functions. In the special case where the underlying graph is a path on n nodes, the chromatic quasisymmetric function tracks descents in Smirnov words of length n , i.e. words where two adjacent letters are distinct. More importantly, in contrast to the ribbon-positivity in the case where the graph was completely disconnected, the chromatic quasisymmetric function of the path graph is e -positive, i.e. it can be expressed as a non-negative integer linear combination of elementary symmetric functions. This raises the following natural questions:

- (1) Is there an analogue of Smirnov words and the descent statistic in the context of labeled binary trees?
- (2) If yes, is the generating function tracking the distribution of descents e -positive?

In this article, we provide answers to both these questions in the positive by introducing Smirnov trees — labeled rooted plane binary trees with the property that if the parent has the same label as one of its children, then the left child must have a larger label than the right child. Denote the set of all Smirnov trees by \mathcal{T} . For any labeled binary tree T , let $\mathbf{lasc}(T)$, $\mathbf{ldes}(T)$, $\mathbf{rasc}(T)$ and $\mathbf{rdes}(T)$ denote the number of ascents and descents in the labeling to the left and right, such that \mathbf{rasc} and \mathbf{lasc} are determined by weak inequalities, whereas \mathbf{rdes} and \mathbf{ldes} are strict. See Subsection 2.2 for details. We associate a monomial \mathbf{x}_T with T as follows. For a node $v \in T$ labeled i , let x_v be x_i . Then

$$\mathbf{x}_T = \prod_{v \in T} x_v.$$

Consider the formal power series in \mathbf{x} with coefficients in $\mathbb{Q}[\bar{\rho}, \rho, \bar{\lambda}, \lambda]$,

$$G := G(\bar{\rho}, \rho, \bar{\lambda}, \lambda) = \sum_{T \in \mathcal{T}} \bar{\rho}^{\text{rasc}(T)} \rho^{\text{rdes}(T)} \bar{\lambda}^{\text{lasc}(T)} \lambda^{\text{lides}(T)} \mathbf{x}_T.$$

It is not immediate that G is a symmetric function in \mathbf{x} with coefficients in $\mathbb{Q}[\bar{\rho}, \rho, \bar{\lambda}, \lambda]$. We establish the preceding fact via the following functional equation satisfied by G , which is also our first main result.

Theorem 1.1. *Let \mathcal{W}_n be the set of Smirnov words of length n . Then we have*

$$G(\bar{\rho}, \rho, \bar{\lambda}, \lambda) = \sum_{n \geq 1} \sum_{w \in \mathcal{W}_n} (\bar{\rho} \bar{\lambda} G + \bar{\rho} + \bar{\lambda})^{\text{asc}(w)} (\rho \lambda G + \rho + \lambda)^{\text{des}(w)} \mathbf{x}_w.$$

Theorem 1.1 follows from an intricate weight-preserving bijection between the sets $\bigcup_{n \geq 1} \mathcal{W}_n \times (\mathcal{T} \cup \{D, U\})^{n-1}$ and \mathcal{T} , and this forms the technical crux of our article.

Let $e_n := e_n(\mathbf{x})$ denote the n -th elementary symmetric function, and let $E(z) := \sum_{n \geq 0} e_n z^n$. It is well known [13, Theorem C.4] that the generating function tracking ascents and descents over all Smirnov words is e -positive. This fact in conjunction with Theorem 1.1 implies the following theorem.

Theorem 1.2. *$G(\bar{\rho}, \rho, \bar{\lambda}, \lambda)$ is e -positive.*

In fact, we provide a combinatorially explicit description for the coefficients in the basis of elementary symmetric functions. Additionally, we establish the following functional equation satisfied by G .

Theorem 1.3.

$$\frac{(1 + \bar{\rho}G)(1 + \bar{\lambda}G)}{(1 + \rho G)(1 + \lambda G)} = \frac{E(\bar{\rho} \bar{\lambda} G + \bar{\rho} + \bar{\lambda})}{E(\rho \lambda G + \rho + \lambda)}.$$

Outline of the article: In Section 2, we provide the necessary definitions for Smirnov words and Smirnov trees. In Section 3, we outline how we prove Theorem 1.1, how e -positivity follows from the functional equation, and how to compute the coefficients in the e -expansion. We give a proof of Theorem 1.3 as well. Finally, in Section 4, we construct an intermediate bijection. The bijection that proves the main result is then the result of successive applications of this intermediate bijection. We provide many examples and a technical proof.

2. BACKGROUND

We denote the set of positive integers by \mathbb{N} . For $n \in \mathbb{N}$, set $[n] := \{1, \dots, n\}$. In particular, $[0] = \emptyset$. For any undefined terms, we refer the reader to [17].

2.1. **Words.** Let \mathbb{N}^* denote the set of all words in the alphabet \mathbb{N} . Given $w = w_1 \dots w_n \in \mathbb{N}^*$, we call n the *length* of w . An index $i \in [n - 1]$ is called a *descent* of w if $w_i > w_{i+1}$ and an *ascent* otherwise. We denote the set of descents (respectively ascents) of w by $\text{Des}(w)$ (respectively $\text{Asc}(w)$) and denote $|\text{Des}(w)|$ (respectively $|\text{Asc}(w)|$) by $\text{des}(w)$ (respectively $\text{asc}(w)$).

A word $w = w_1 \dots w_n$ is a *Smirnov word* if adjacent letters are different, i.e. if $w_i \neq w_{i+1}$ for $i \in [n - 1]$. The set of all Smirnov words is denoted by \mathcal{W} and the set of all Smirnov words of length n by \mathcal{W}_n . Consider the multivariate generating function tracking the distribution of ascents and descents over \mathcal{W}_n

$$Q_n(s, t) = \sum_{w \in \mathcal{W}_n} s^{\text{asc}(w)} t^{\text{des}(w)} \mathbf{x}_w,$$

and let $Q(z; s, t) := \sum_{n \geq 0} Q_n(s, t) z^n$.

2.2. **Smirnov trees.** A *plane binary tree* is a rooted tree in which every node has at most two children, of which at most one is called a *left child* and at most one is called a *right child*. Henceforth, we simply say binary tree instead of plane binary tree. A *labeled plane binary tree* (or simply a *labeled binary tree*) is a binary tree whose nodes have labels drawn from \mathbb{N} . We assign a weight to labeled binary trees as follows. An edge between a parent (resp. a left child) with label a and its right child (resp. its parent) with label b is weighted $\bar{\rho}$ (resp. $\bar{\lambda}$) if $a \leq b$ and ρ (resp. λ) if $a > b$. In other words, if the edge is between a parent and the right child, we use rho (for right), and if it is between a parent and its left child, we use lambda (for left). We add the bar if the nodes form a weak ascent when reading from left to right (either diagonally up or down). A node with label a has weight x_a . The weight of a binary tree T , denoted by $\text{wt}(T)$, is the product of the weights of all its edges and nodes.

A labeled binary tree is called *Smirnov* if the following holds:

- if the left (resp. right) child has the same label i as its parent, then the parent must also have a right (resp. left) child with label $< i$ (resp. $> i$).

In other words, if the parent has the same label as its child, the left child must have a larger label than the right child. Smirnov trees inherit a weight function from the one defined for ordinary labeled binary trees.

Denote the set of all Smirnov trees by \mathcal{T} , and the set of all Smirnov trees whose root has label c by \mathcal{T}^c . If $T \in \mathcal{T}$, define its *principal path* $\text{P}(T)$ as follows: it starts at the root; if the current node has no right child, stop; if the current node has a left child with the same label (and therefore a right child with a smaller label), move down left; otherwise, move down right. The last node on the principal path is called the *principal node* and is denoted $\alpha(T)$; its label is $a(T)$. The maximal label on $\text{P}(T)$ is $\text{M}(T)$, and the minimal label is $\text{m}(T)$.

Example 2.1. *As an example, let T be the binary tree in Figure 1. There are two nodes with a child with the same label; in both cases, the parent has both children, and the left label is larger than the right label. Therefore T is Smirnov. In particular, $T \in \mathcal{T}^3$. The edges*




FIGURE 1. A Smirnov tree with principal path in bold.

lying on the principal path $\mathbf{P}(T)$ are thickened, and the labels on $\mathbf{P}(T)$ are 3, 3, 4, 1, 3. So $a(T) = 3$, $\mathbf{M}(T) = 4$, and $\mathbf{m}(T) = 1$. The principal node $\alpha(T)$ is gray. Furthermore,

$$\mathbf{wt}(T) = \bar{\rho}^4 \rho^3 \bar{\lambda}^2 \lambda^3 x_1^3 x_2^2 x_3^6 x_4^2.$$

As mentioned in the introduction, we prove Theorem 1.1 by establishing a weight-preserving bijection between $\bigcup_{n \geq 1} \mathcal{W}_n \times (\mathcal{T} \cup \{D, U\})^{n-1}$ and \mathcal{T} . We have already described how to assign weights to elements of the latter, so we proceed to describe weights assigned to elements of the former. For the sake of brevity, set

$$\mathcal{X} := \bigcup_{n \geq 1} \mathcal{W}_n \times (\mathcal{T} \cup \{D, U\})^{n-1}.$$

Given a statement P , let $[P]$ equal 1 if P is true and 0 otherwise. If a, b are distinct positive integers and $Y \in \mathcal{T} \cup \{D, U\}$, define

$$\mathbf{f}(a, b, Y) := \begin{cases} \bar{\rho} \bar{\lambda} \mathbf{wt}(T) [Y \in \mathcal{T}] + \bar{\rho} [Y = D] + \bar{\lambda} [Y = U] & \text{if } a < b, \\ \rho \lambda \mathbf{wt}(T) [Y \in \mathcal{T}] + \rho [Y = D] + \lambda [Y = U] & \text{if } a > b. \end{cases}$$

This given, consider $(w, S) \in \mathcal{X}$, and let $w := w_1 \dots w_n$ and $S := S_1 \dots S_{n-1}$. The weight of (w, S) , denoted by $\mathbf{wt}(w, S)$, is defined as follows.

$$\mathbf{wt}(w, S) := \mathbf{x}_w \prod_{1 \leq i \leq n-1} \mathbf{f}(w_i, w_{i+1}, S_i).$$

Example 2.2. Consider the Smirnov word $w = 42534242 \in \mathcal{W}_8$ and the sequence

$$S = \left(\circ, D, \circ, D, D, U, \overset{\circ}{\circ} \right) \in (\mathcal{T} \cup \{D, U\})^7$$

Then we have

$$\begin{aligned}
f(w_1, w_2, S_1) &= \rho\lambda x_1 \\
f(w_2, w_3, S_2) &= \bar{\rho} \\
f(w_3, w_4, S_3) &= \rho\lambda x_2 \\
f(w_4, w_5, S_4) &= \bar{\rho} \\
f(w_5, w_6, S_5) &= \rho \\
f(w_6, w_7, S_6) &= \bar{\lambda} \\
f(w_7, w_8, S_7) &= \rho^2\bar{\lambda}\lambda x_1x_2^2.
\end{aligned}$$

It follows that

$$\text{wt}(w, S) = \bar{\rho}^2\rho^5\bar{\lambda}^2\lambda^3 x_1^2x_2^6x_3x_4^3x_5.$$

It is straightforward to see that

$$(2.1) \quad \sum_{(w,S) \in \mathcal{X}} \text{wt}(w, S) = \sum_{n \geq 1} \sum_{w \in \mathcal{W}_n} (\bar{\rho}\bar{\lambda}G + \bar{\rho} + \bar{\lambda})^{\text{asc}(w)} (\rho\lambda G + \rho + \lambda)^{\text{des}(w)} \mathbf{x}_w.$$

Note that the right hand side in equation (2.1) is exactly the right hand side of the equality in Theorem 1.1. This explains why it suffices to exhibit a weight-preserving bijection between \mathcal{T} and \mathcal{X} in order to prove Theorem 1.1.

3. BIJECTION AND APPLICATIONS

3.1. Outline of the bijection. Our weight-preserving bijection between \mathcal{T} and \mathcal{X} uses an intermediate bijection that we call Φ . Consider a triple (T, S, b) , where $S \in \mathcal{T} \cup \{D, U\}$ and $b \in \mathbb{N}$ is such that $a(T) \neq b$. Recall that $a(T)$ is the label of the principal node of T . We assign a weight to (T, S, b) as follows:

$$\text{wt}(T, S, b) = x_b \text{wt}(T) f(a(T), b, S).$$

Our map Φ , whose precise description we postpone to Section 4, satisfies the following.

Theorem 3.1. *The map Φ is a well-defined weight-preserving bijection from the set of triples (T, S, b) in $\mathcal{T} \times (\mathcal{T} \cup \{D, U\}) \times \mathbb{N}$ satisfying $a(T) \neq b$ to the set of Smirnov trees with at least 2 nodes, and the principal node of the image has label equal to the third argument. In other words, $\text{wt}(\Phi(T, S, b)) = \text{wt}(T, S, b)$ and $a(\Phi(T, S, b)) = b$.*

The proof of Theorem 3.1, which is fairly involved, is also presented in Section 4. For the moment, we describe how Theorem 1.1 follows from Theorem 3.1, and subsequently we discuss various applications.

Proof. (of Theorem 1.1) By applying Φ successively, we construct a weight-preserving bijection

$$\Psi: \mathcal{X} \longrightarrow \mathcal{T},$$

which proves our main theorem. If $w \in \mathcal{W}_1$, let $\Psi(w, \emptyset)$ be the binary tree with a single node labeled w_1 . Here \emptyset denotes the empty sequence. If $n \geq 2$, $w = w_1 \cdots w_n \in \mathcal{W}_n$, $S = (S_1, \dots, S_{n-1}) \in (\mathcal{T} \cup \{D, U\})^{n-1}$, define $w' = w_1 \cdots w_{n-1}$, $S' = (S_1, \dots, S_{n-2})$, and

$$\Psi(w, S) = \Phi(\Psi(w', S'), S_{n-1}, w_n).$$

We prove by induction on n that for $w \in \mathcal{W}_n$ and $S \in (\mathcal{T} \cup \{D, U\})^{n-1}$, $\Psi(w, S)$ is well defined and $a(\Psi(w, S)) = w_n$. For $n = 1$, this is obvious, and if it holds for $n - 1$, then $a(\Psi(w', S')) = w_{n-1} \neq w_n$, so $\Phi(\Psi(w', S'), S_{n-1}, w_n)$ is well defined, and by Theorem 3.1, $a(\Phi(\Psi(w', S'), S_{n-1}, w_n)) = w_n$.

As far as the weight-preserving aspect of Ψ is concerned, note that

$$(3.1) \quad \mathbf{wt}(w, S) = x_{w_n} \cdot \mathbf{wt}(w', S') \cdot \mathbf{f}(w_{n-1}, w_n, S_{n-1})$$

Furthermore,

$$(3.2) \quad \begin{aligned} \mathbf{wt}(\Psi(w, S)) &= \mathbf{wt}(\Phi(\Psi(w', S'), S_{n-1}, w_n)) \\ &= x_{w_n} \cdot \mathbf{wt}(\Psi(w', S')) \cdot \mathbf{f}(w_{n-1}, w_n, S_{n-1}). \end{aligned}$$

In going from the first equality to the second, we have utilized the fact that Φ is a weight-preserving bijection. Induction along with a comparison of equations (3.1) and (3.2) implies that Ψ is weight preserving. Bijectivity of Ψ also follows from the bijectivity of Φ . \square

3.2. e -positivity and bleeding trees. To write an explicit expression for $Q_n(s, t)$, we invoke [13, Theorem C.4] which states that the generating function for Smirnov words of length n is e -positive. More precisely, Shareshian and Wachs prove that

$$(3.3) \quad \sum_{w \in \mathcal{W}_n} t^{\mathbf{des}(w)} \mathbf{x}_w = \sum_{m=1}^{\lfloor \frac{n+1}{2} \rfloor} \sum_{\substack{k_1, \dots, k_m \geq 2 \\ \sum k_i = n+1}} e_{(k_1-1, k_2, \dots, k_m)} t^{m-1} \prod_{i=1}^m [k_i - 1]_t,$$

where $[a]_t = 1 + t + \cdots + t^{a-1}$ for $a \in \mathbb{N}$. For example, it is easy to check that

$$\sum_{w \in \mathcal{W}_3} t^{\mathbf{des}(w)} \mathbf{x}_w = (1 + t + t^2)e_3 + te_{21},$$

which agrees with the formula.

It immediately follows from equation (3.3) that

$$\begin{aligned}
(3.4) \quad Q_n(s, t) &= s^{n-1} \sum_{m=1}^{\lfloor \frac{n+1}{2} \rfloor} \sum_{\substack{k_1, \dots, k_m \geq 2 \\ \sum k_i = n+1}} e_{(k_1-1, k_2, \dots, k_m)}(t/s)^{m-1} \prod_{i=1}^m [k_i - 1]_{t/s} \\
&= \sum_{m=1}^{\lfloor \frac{n+1}{2} \rfloor} \sum_{\substack{k_1, \dots, k_m \geq 2 \\ \sum k_i = n+1}} e_{(k_1-1, k_2, \dots, k_m)} s^{m-1} t^{m-1} \prod_{i=1}^m (s^{k_i-2} + s^{k_i-3}t + \dots + t^{k_i-2}) \\
&= \sum_{m=1}^{\lfloor \frac{n+1}{2} \rfloor} \sum_{\substack{k_1, \dots, k_m \geq 2 \\ \sum k_i = n+1}} e_{(k_1-1, k_2, \dots, k_m)} (s^{k_1-2} + s^{k_1-3}t + \dots + t^{k_1-2}) \prod_{i=2}^m (s^{k_i-1}t + s^{k_i-2}t^2 + \dots + st^{k_i-1})
\end{aligned}$$

Our main result Theorem 1.1 states that the generating function for Smirnov trees satisfies the functional equation

$$(3.5) \quad G = \sum_{n \geq 1} \sum_{w \in \mathcal{W}_n} s^{\text{asc}(w)} t^{\text{des}(w)} \mathbf{x}_w = Q(1; s, t) - 1,$$

where

$$s = \bar{\rho} \bar{\lambda} G + \bar{\rho} + \bar{\lambda}, \quad t = \rho \lambda G + \rho + \lambda.$$

From equations (3.4) and (3.5), it follows that G is e -positive. Our goal is to describe $c_\pi = [e_\pi]G$, the coefficient of e_π in the e -expansion of G , for an arbitrary partition π . We will describe c_π in terms of certain trees. For motivation, let us extract the coefficient at e_{321} .

Among the terms in (3.5) (with s and t as above) there are

$$e_{132} \cdot 1 \cdot (s^2t + st^2) \cdot (st) \quad \text{and} \quad e_{123} \cdot 1 \cdot (st) \cdot (s^2t + st^2)$$

corresponding to compositions $(k_1, k_2, k_3) = (2, 3, 2)$, $(k_1, k_2, k_3) = (2, 2, 3)$. For every occurrence of s (resp. t), we can select either $\bar{\rho} \bar{\lambda} G$ (resp. $\rho \lambda G$) or $\bar{\rho} + \bar{\lambda}$ (resp. $\rho + \lambda$). Of course, selecting a term with G means that the resulting symmetric function will have degree > 6 , so the two terms contribute

$$2((\bar{\rho} + \bar{\lambda})^2(\rho + \lambda) + (\bar{\rho} + \bar{\lambda})(\rho + \lambda)^2)((\bar{\rho} + \bar{\lambda})(\rho + \lambda))$$

to c_{321} .

There are other terms in c_{321} , however. For example, among the terms in (3.4) there is also

$$e_{32} \cdot (s^2 + st + t^2)(st),$$

corresponding to the composition $(k_1, k_2) = (4, 2)$. Since $e_{321} = e_{32}e_1$, the contribution of this term to c_{321} is equal to the coefficient of e_1 in $(s^2 + st + t^2)(st)$. This is equal to

$$([e_1](s^2 + st + t^2))([1](st)) + ([1](s^2 + st + t^2))([e_1](st)).$$

The function e_1 in $s \cdot s + s \cdot t + t \cdot t$ comes from selecting one of the three summands, and then selecting e_1 from the first or the second term in the product. Therefore

$$[e_1](s^2 + st + t^2) = 2(\bar{\rho}\bar{\lambda}c_1)(\bar{\rho} + \bar{\lambda}) + (\bar{\rho}\bar{\lambda}c_1)(\rho + \lambda) + (\bar{\rho} + \bar{\lambda})(\rho\lambda c_1) + 2(\rho\lambda c_1)(\rho + \lambda).$$

It is clear that c_1 , the coefficient of e_1 in G , is 1. We also have

$$\begin{aligned} [1](st) &= (\bar{\rho} + \bar{\lambda})(\rho + \lambda), \\ [1](s^2 + st + t^2) &= (\bar{\rho} + \bar{\lambda})^2 + (\bar{\rho} + \bar{\lambda})(\rho + \lambda) + (\rho + \lambda)^2, \\ [e_1](st) &= \bar{\rho}\bar{\lambda}(\rho + \lambda) + (\bar{\rho} + \bar{\lambda})\rho\lambda, \end{aligned}$$

which gives further terms of c_{321} .

We can distill the above reasoning into the following definition. Note that it follows from (3.4) that when extracting a coefficient of e_π from some expression, we must pick a part of π that will play the role of k_1 . That is the meaning of red edges below.

Definition 3.2. A **bleeding tree** is an unordered rooted tree with the following properties:

- nodes at even depth (including the root) are red and unlabeled, nodes at odd depth are black and labeled with positive integers
- every red node has children, and it is connected with a red edge with exactly one of them; all other edges in the tree are black
- if a black node is connected to its parent with a black edge, the number of children it has is less than or equal to its label, which cannot be 1; otherwise, it is strictly smaller than its label (in particular, it has no children if its label is 1).

For a partition π , we denote by \mathcal{B}_π the (clearly finite) set of bleeding trees whose labels (with repetitions) are precisely the parts of π .

We give black nodes weights. Assume that a black node has label r and k children.

- If the edge to its parent is black, its weight is


$$p(r, k) = \sum_{i=1}^{r-1} \sum_{j=0}^i \binom{i}{j} \binom{r-i}{r-k-j} (\bar{\rho}\bar{\lambda})^{i-j} (\bar{\rho} + \bar{\lambda})^j (\rho\lambda)^{k-i+j} (\rho + \lambda)^{r-k-j}$$

- If the edge to its parent is red, its weight is

$$\bar{p}(r, k) = \sum_{i=0}^{r-1} \sum_{j=0}^i \binom{i}{j} \binom{r-1-i}{r-1-k-j} (\bar{\rho}\bar{\lambda})^{i-j} (\bar{\rho} + \bar{\lambda})^j (\rho\lambda)^{k-i+j} (\rho + \lambda)^{r-1-k-j}$$

The weight $\mathbf{wt}(U)$ of a bleeding tree U is defined as the product of the weights of all its black nodes and the number of ways to draw a bleeding tree in the plane so that red edges always come first.

For example, the bleeding tree with one red node, three black nodes with labels 1, 2 and 3, and a red edge between the red node and the node with label 1 has weight $2\bar{p}(1, 0)p(2, 0)p(3, 0)$, with the factor 2 coming from drawings



Note that the requirement that a black node with a black edge to its parent cannot have label 1 is superfluous, as the weight of such a node would be the empty sum, i.e. 0. The requirement about the number of children being smaller than (or equal to) the label is also not needed, as the second binomial coefficient in the sums would be 0 in that case. In other words, given the weighting of trees, only the first two assumptions are necessary for a tree to be bleeding.

Example 3.3. *Figure 2 depicts a bleeding tree in $\mathcal{B}_{4332222111}$ with weight*

$$2\bar{p}(3, 2)p(2, 1)p(2, 1)\bar{p}(2, 0)p(4, 2)\bar{p}(1, 0)\bar{p}(3, 0)\bar{p}(2, 0)\bar{p}(1, 0)\bar{p}(1, 0).$$




FIGURE 2. A bleeding tree in $\mathcal{B}_{4332222111}$.

All trees in \mathcal{B}_{321} are shown in Figure 3.


Given the straightforward yet cumbersome details of the induction, we omit the proof of the following.

Theorem 3.4. *The generating function for Smirnov trees is e -positive, with coefficients in $\mathbb{N}[\bar{\rho}, \rho, \bar{\lambda}, \lambda]$. More precisely, we can write*

$$G = \sum_{\pi} \left(\sum_{U \in \mathcal{B}_{\pi}} w(U) \right) e_{\pi},$$

where the sum is over all partitions π , $|\pi| \geq 1$.

In fact, the coefficient of e_{π} in G is a polynomial with positive integer coefficients in the “variables” $\bar{\rho}, \bar{\lambda}, \bar{\rho} + \bar{\lambda}, \rho, \lambda, \rho + \lambda$. Thus, our Theorem 3.4 may be interpreted as an instance

FIGURE 3. All trees in \mathcal{B}_{321} .

of Schur- γ -nonnegativity defined in [14]. See also [1] for more on this theme. The first few terms of the expansion are

$$e_1 + (\bar{\rho} + \rho + \bar{\lambda} + \lambda)e_2 + (\bar{\rho}^2 + \bar{\rho}\rho + \rho^2 + 2\bar{\rho}\bar{\lambda} + \rho\bar{\lambda} + \bar{\lambda}^2 + \bar{\rho}\lambda + 2\rho\lambda + \bar{\lambda}\lambda + \lambda^2)e_3 \\ + (\bar{\rho}\rho + \bar{\rho}\bar{\lambda} + \rho\bar{\lambda} + \bar{\rho}\lambda + \rho\lambda + \bar{\lambda}\lambda)e_{21} + \dots$$

Figure 3 allows us to quickly compute the coefficient at e_{321} :

$$2\bar{p}(1,0)p(3,0)p(2,0) + \bar{p}(3,1)p(2,0)\bar{p}(1,0) + \bar{p}(3,0)p(2,1)\bar{p}(1,0) + \bar{p}(2,1)p(3,0)\bar{p}(1,0) \\ + \bar{p}(2,0)p(3,1)\bar{p}(1,0) + \bar{p}(1,0)p(3,1)\bar{p}(2,0) + \bar{p}(1,0)p(2,1)\bar{p}(3,0) + \bar{p}(3,1)p(2,0)\bar{p}(1,0) \\ + 2\bar{p}(3,2)\bar{p}(2,0)\bar{p}(1,0) + \bar{p}(3,1)\bar{p}(2,1)\bar{p}(1,0) + \bar{p}(2,1)\bar{p}(1,0)p(3,0) + \bar{p}(2,1)\bar{p}(3,1)\bar{p}(1,0) \\ = 2\bar{\rho}^3\rho^2 + 2\bar{\rho}^2\rho^3 + 12\bar{\rho}^3\rho\bar{\lambda} + 18\bar{\rho}^2\rho^2\bar{\lambda} + 8\bar{\rho}\rho^3\bar{\lambda} + 6\bar{\rho}^3\bar{\lambda}^2 + 28\bar{\rho}^2\rho\bar{\lambda}^2 + 18\bar{\rho}\rho^2\bar{\lambda}^2 + 2\rho^3\bar{\lambda}^2 + 6\bar{\rho}^2\bar{\lambda}^3 \\ + 12\bar{\rho}\rho\bar{\lambda}^3 + 2\rho^2\bar{\lambda}^3 + 8\bar{\rho}^3\rho\lambda + 18\bar{\rho}^2\rho^2\lambda + 12\bar{\rho}\rho^3\lambda + 12\bar{\rho}^3\bar{\lambda}\lambda + 56\bar{\rho}^2\rho\bar{\lambda}\lambda + 56\bar{\rho}\rho^2\bar{\lambda}\lambda + 12\rho^3\bar{\lambda}\lambda + 28\bar{\rho}^2\bar{\lambda}^2\lambda \\ + 56\bar{\rho}\rho\bar{\lambda}^2\lambda + 18\rho^2\bar{\lambda}^2\lambda + 12\bar{\rho}\bar{\lambda}^3\lambda + 8\rho\bar{\lambda}^3\lambda + 2\rho^3\lambda^2 + 18\bar{\rho}^2\rho\lambda^2 + 28\bar{\rho}\rho^2\lambda^2 + 6\rho^3\lambda^2 + 18\bar{\rho}^2\bar{\lambda}\lambda^2 + 56\bar{\rho}\rho\bar{\lambda}\lambda^2 \\ + 28\rho^2\bar{\lambda}\lambda^2 + 18\bar{\rho}\bar{\lambda}^2\lambda^2 + 18\rho\bar{\lambda}^2\lambda^2 + 2\bar{\lambda}^3\lambda^2 + 2\bar{\rho}^2\lambda^3 + 12\bar{\rho}\rho\lambda^3 + 6\rho^2\lambda^3 + 8\bar{\rho}\bar{\lambda}\lambda^3 + 12\rho\bar{\lambda}\lambda^3 + 2\bar{\lambda}^2\lambda^3$$

3.3. Another functional equation and exponential specialization. We now give a proof of the functional equation satisfied by G stated in Theorem 1.3, and then relate it to earlier work of Gessel.

Proof. (of Theorem 1.3) Recall that $Q(z; s, t)$ tracks the distribution of ascents and descents over the set of all Smirnov words, and we have

$$(3.6) \quad Q(z; s, t) = 1 + s^{-1} (Q(sz; 1, ts^{-1}) - 1).$$

From [13, Theorem C.3], it follows that

$$(3.7) \quad Q(z; 1, t) - 1 = \frac{E(z) - E(zt)}{E(zt) - tE(z)}.$$

From (3.7) and (3.6), we obtain

$$Q(1; s, t) - 1 = \frac{E(s) - E(t)}{sE(t) - tE(s)}.$$

Set $s = \bar{\rho}\bar{\lambda}G + \bar{\rho} + \bar{\lambda}$ and $t = \rho\lambda G + \rho + \lambda$ henceforth. By Theorem 1.1, we have

$$G = Q(1; s, t) - 1,$$

which in turn implies $(1 + \bar{\rho}G)(1 + \bar{\lambda}G) = 1 + Gs = s(Q(1; s, t) - 1) + 1$ and $(1 + \rho G)(1 + \lambda G) = 1 + Gt = t(Q(1; s, t) - 1) + 1$. Thus we obtain

$$(3.8) \quad \frac{(1 + \bar{\rho}G)(1 + \bar{\lambda}G)}{(1 + \rho G)(1 + \lambda G)} = \frac{s \frac{E(s) - E(t)}{sE(t) - tE(s)} + 1}{t \frac{E(s) - E(t)}{sE(t) - tE(s)} + 1} = \frac{E(s)}{E(t)}.$$

This establishes the claim. \square

While the equality in Theorem 1.3 is less transparent than Theorem 1.1, it immediately allows us to establish a result present in unpublished work of Gessel, and then proved in [8] and [4]. We call a labeled binary tree on n nodes *standard* if the labels are all distinct and drawn from $[n]$. Note that a standard labeled binary trees is necessarily Smirnov. Gessel considered the following generating function that tracks the distributions of ascents and descents over standard labeled binary trees:

$$(3.9) \quad B := B(\bar{\rho}, \rho, \bar{\lambda}, \lambda) = \sum_{n \geq 1} \sum_{T \text{ standard}} \bar{\rho}^{\text{rasc}(T)} \rho^{\text{rdes}(T)} \bar{\lambda}^{\text{lasc}(T)} \lambda^{\text{lides}(T)} \frac{x^n}{n!}.$$

Consider the homomorphism ex from the ring of symmetric functions to $\mathbb{Q}[[x]]$ defined by setting $\text{ex}(e_n) = \frac{x^n}{n!}$ for $n \in \mathbb{N}$. A key property of ex is the following. Given a symmetric function f , we have $[x_1 \cdots x_n]f = [\frac{x^n}{n!}] \text{ex}(f)$. It follows that $\text{ex}(G) = B$, and we obtain the following corollary.

Corollary 3.5.

$$\frac{(1 + \bar{\rho}B)(1 + \bar{\lambda}B)}{(1 + \rho B)(1 + \lambda B)} = e^{((\bar{\lambda}\bar{\rho} - \lambda\rho)B + \bar{\rho} + \bar{\lambda} - \rho - \lambda)x}.$$

4. THE GORY DETAILS

The crucial ingredient of the proof of Theorem 1.1 is the map

$$\Phi: \{(T, S, b) \in \mathcal{T} \times (\mathcal{T} \cup \{D, U\}) \times \mathbb{N}: a(T) \neq b\} \longrightarrow \mathcal{T},$$

which we define as follows. Take a Smirnov tree T ; S , which is either a Smirnov tree, a down step D or an up step U ; and an integer b that is different from the label of the principal node of T . Write $\alpha = \alpha(T)$, $a = a(T)$, $P = \mathbf{P}(T)$, $M = \mathbf{M}(T)$, and $m = \mathbf{m}(T)$. By definition, $a \neq b$, and α has no right child. If $S \in \mathcal{T}$, let c be the label of its root. If $a < b$ (resp. $a > b$), let δ be the last node on P with label $\geq b$ (resp. $\leq b$), and let d be its label. If there are no such nodes on P , δ and d are undefined. For example, if T is the tree from the previous example and $b = 2$, then $d = 1$, and if $b = 5$, d is undefined. Note that saying, for example, $M < b$ is equivalent to saying that all the labels on the principal path of the tree are smaller than b , and $M \geq b$ means that at least one label on the principal path is at least b .




FIGURE 4. The map Φ .

The map Φ is illustrated in Figure 4, and the formal definition is as follows.

- (1) If
 - (1a) $S = D$ or
 - (1b) $S \in \mathcal{T}^c$ & $a, c < b$ or
 - (1c) $S \in \mathcal{T}^c$ & $a, c > b$,

then $\Phi(T, S, b)$ is the tree we obtain if we add a right child β with label b to α . If $S = D$, β has no children, and if $S \in \mathcal{T}^c$, β has S as its left subtree.

(2) If

- (2a) $S = U$ & $a, M < b$ or
- (2b) $S = U$ & $a, m > b$ or
- (2c) $S \in \mathcal{T}^c$ & $a, M < b \leq c$ or
- (2d) $S \in \mathcal{T}^c$ & $a, m > b \geq c$,

then $\Phi(T, S, b)$ is the tree we obtain if we add a new root β with label b , and make T its left subtree. If $S \in \mathcal{T}^c$, we make S the right subtree of α .

(3) If

- (3a) $S = U$ & $a < b \leq M$ & $d > b$ or
- (3b) $S = U$ & $a > b \geq m$ & $d < b$ or
- (3c) $S \in \mathcal{T}^c$ & $a < b \leq c, M$ & $d > b$ or
- (3d) $S \in \mathcal{T}^c$ & $a > b \geq c, m$ & $d < b$,

then $\Phi(T, S, b)$ is the tree we obtain if we replace the right subtree of δ with a new node β with label b ; β has no right child, and its left subtree is the previous right subtree of δ . Furthermore, if $S \in \mathcal{T}^c$, we make S the right subtree of α .

(4) If

- (4a) $S = U$ & $a < b \leq M$ & $d = b$ or
- (4b) $S \in \mathcal{T}^c$ & $a < b \leq c, M$ & $d = b$,

then $\Phi(T, S, b)$ is the tree we obtain if we replace the left subtree of δ with a new node β with label b ; β has no right child, and its left subtree is the previous left subtree of δ . Furthermore, if $S \in \mathcal{T}^c$, we make S the right subtree of α .

(5) If

- (5a) $S = U$ & $a > b \geq m$ & $d = b$ or
- (5b) $S \in \mathcal{T}^c$ & $a > b \geq c, m$ & $d = b$,

then $\Phi(T, S, b)$ is the tree we obtain if we replace the left subtree of δ with its right subtree, and its right subtree with a new node β with label b ; β has no right child, and its left subtree is the previous left subtree of δ . Furthermore, if $S \in \mathcal{T}^c$, we make S the right subtree of α .

Remark 4.1. *Let us emphasize that while the map is complicated, there is not much freedom if we want to preserve weights and Smirnovness, and if we want the new principal path to end in the node with the newly added label.*

Let us illustrate the simplest rule, rule 1. See Figure 5. We take the same tree T for all four examples, with the principal node $\alpha = \alpha(T)$ being the left child of the right child of the root, and $a = a(T) = 3$ (again, the principal node is gray). By definition, b cannot equal 3. If $S = D$, we simply add a right child to α , with label b . The weight of T is multiplied by x_b and, if $a < b$ (resp. $a > b$), by $\bar{\lambda}$ (resp. λ); the edge that contributes this weight is red. The situation is similar if S is a tree with root label c , and b is either larger or smaller than both a and c . In that case, add a right child with label b to α , and this right child has S as

its left subtree (light gray). The weight of T is multiplied by x_b , $\text{wt}(S)$, and $\bar{\lambda}\bar{\rho}$ (resp. $\lambda\rho$) if $a < b$ (resp. $a > b$); the edges that contribute these weights are red. Note that in all three cases, the principal path is extended by one right step, and $a(\Phi(T, S, b)) = b$.

T	S	a	b	c	d	M	m	rule	$\Phi(T, S, b)$	T	S	a	b	c	d	M	m	rule	$\Phi(T, S, b)$
	D	3	4					1a			D	3	2					1a	
		3	4	2				1b				3	2	3				1c	

FIGURE 5. An illustration of rule 1.

Let us present examples of rule 2, which is the only case when we produce a new root. See Figure 6. Again, we can take T to be the same tree throughout. If $S = U$, and either $a < b$ and b is larger than all labels on the principal path (i.e. $M < b$), or $a > b$ and b is smaller than all the labels on the principal path (i.e. $m > b$), create a new root with label b , with no right child, and with the entire T as its left subtree. The subcases 2c and 2d are similar. We have a tree S with root label c , and either $a, M < b \leq c$ or $a, m > b \geq c$. In addition to the new root, we attach the tree S as the right subtree of α . The new root multiplies the weight of the tree by x_b , and if S is a tree, the weight is multiplied by $\text{wt}(S)$. Furthermore, the edge to the new root adds $\bar{\lambda}$ if $a > b$ and λ if $a < b$; if S is a tree, the edge between α and its new right child adds the weight $\bar{\rho}$ if $a > b$ and ρ if $a < b$.

T	S	a	b	c	d	M	m	rule	$\Phi(T, S, b)$	T	S	a	b	c	d	M	m	rule	$\Phi(T, S, b)$
	U	3	4			3		2a			U	3	1				2	2b	
		3	4	4		3		2c				3	1	1			2	2d	

FIGURE 6. An illustration of rule 2.

For rule 3, let us take a slightly larger example; note that the T 's Figure 7 differ only in the root label. Unlike in rule 2, there is a label on the principal path that is $\geq b$ if $a < b$ and $\leq b$ if $a > b$. Recall that δ is the last node on the principal path with that property (the root in all our four examples), with label d . The crucial assumption now is that $d \neq b$. If S is a tree, its root label must be $\geq b$ if $a < b$ and $\leq b$ if $a > b$ (since otherwise we would use rule 1). Now break up the tree: add a new node β as the right child of δ . The new node has no right child, and its left subtree is the old right subtree of δ . Again, if S is a tree, add it as the right subtree of α .

T	S	a	b	c	d	M	m	rule	$\Phi(T, S, b)$	T	S	a	b	c	d	M	m	rule	$\Phi(T, S, b)$
	U	4	5		6	6		3a			U	4	2		1		1	3b	
		4	5	5	6	6		3c				4	2	2	1		1	3d	

FIGURE 7. An illustration of rule 3.

Rules 4 and 5 deal with the case when $d = b$. Now δ becomes a node with the same label as one of its children: left if $a < b$ and right if $a > b$. If $a > b$, the subtrees of δ are flipped (and we also insert β and S).

T	S	a	b	c	d	M	m	rule	$\Phi(T, S, b)$	T	S	a	b	c	d	M	m	rule	$\Phi(T, S, b)$
	U	4	5		5	5		4a			U	4	2		2		2	5a	
		4	5	5	5	5		4b				4	2	2	2		2	5b	

FIGURE 8. An illustration of rules 4 and 5.

As an example of the use of Ψ , consider the Smirnov word $w = 42534242 \in \mathcal{W}_8$ and the sequence

$$(4.1) \quad \mathcal{S} = \left(\circlearrowleft, D, \circlearrowleft, D, D, U, \circlearrowright \circlearrowleft \right) \in (\mathcal{T} \cup \{D, U\})^7.$$

Then successive applications of Φ are depicted in Figure 9. The last tree obtained is T from the example for rule (5b), and $S = S_7$ and $b = w_8$ also match that example, so $\Psi(w, \mathcal{S})$ is the Smirnov tree in Figure 10.




FIGURE 9. Φ applied to a certain element of \mathcal{X} .

Proof of Theorem 3.1. We are given a Smirnov tree T ; write $\alpha = \alpha(T)$, $a = a(T)$, $M = M(T)$, and $m = m(T)$. We also have $S \in \mathcal{T} \cup \{D, U\}$ and $b \neq a$. We prove that for every appropriate choice of T, S, b , exactly one of the rules applies and the result is a Smirnov tree T' with $a(T') = b$ and $\text{wt}(T') = \text{wt}(T, S, b)$.

The first case is when $S = D$. In that case, rule (1a) applies. Since $a \neq b$ and the principal node α has no right child by definition, we can give it a right child β with label b . This creates a new right edge in the tree; if $a < b$, then its weight is $\bar{\rho}$, and if $a > b$, then its weight is ρ . In both cases, the weight of the new tree, which is clearly Smirnov, equals the weight of (T, S, b) , which is $\bar{\rho} \text{wt}(T)x_b$ or $\rho \text{wt}(T)x_b$, depending on whether $a < b$ or $a > b$. It is also clear that the principal path is the same as before with an added right step at the end. So the new principal node has label b , as claimed. *Note that in case (1a), the new principal node has no children, and its left parent has a different label. (Statements in italics will be crucial in the construction of the inverse of Φ .)*




FIGURE 10. The Smirnov tree corresponding to $w = 42534242$ and \mathcal{S} as in (4.1).

Now assume that $S = U$ and $a < b$. The principal path contains a node with label $< b$, and either M , the maximum of all labels on the principal path, is $< b$, or it is $\geq b$. The first case gives rule (2a). Indeed, since $M < b$, the root label is also $< b$. When we add a new root β with label b and make T its left subtree, we create a new edge with weight $\bar{\lambda}$, so the weight changes as claimed. The new tree is also obviously Smirnov. The principal path contains only the root, and the new principal node again has label b . *Note that in case (2a), the root of the new tree is the principal node with a left child, and all the labels on the principal path of its left subtree are smaller than the root label.*

The next case is when $S = U$, $a < b$, and $M \geq b$. On the principal path, there is at least one node with label at least b . Let δ be the last node on the principal path with that property, and let $d \geq b$ be its label. The children of δ must have a label different from d (otherwise, the node after δ on the principal path is its child with the same label, and therefore δ is not the last node on the principal path with label $\geq b$). That means that the principal path goes right from δ . Again, we have two options, $d > b$ and $d = b$. The first case gives rule (3a). Since $d \neq b$, we can replace the right child of δ with a new node β with label b , and make the old right subtree of δ the left subtree of the new node (which has no right child). Since $d > b$, the edge between δ and its new right child has label ρ ; since all the labels on the principal path of the right subtree of δ in T are $< b$ by the selection of δ , the old edge between δ and its right child had weight ρ , and the new edge between β and its left child has label $\bar{\lambda}$. In other words, the weight of the new tree is the previous weight times $\bar{\lambda}x_b$, as required. Also, the new principal path again ends at β , which has label b . *Note that in case (3a), the new principal node is a right child whose parent has a larger label, and all the labels on the principal path of its (nonempty) left subtree are smaller.*

The other option is that $S = U$, $a < b$, $M \geq b$, and $d = b$. This gives rule (4a): we insert a new node β with label b on the edge between δ and its left child (if δ has no left child, we just add a left child with label b). Since the right child of δ has label $< b$, and the left child of δ has label $\neq d$, the new tree is Smirnov. We added an edge between β and δ , which has, since both have labels b , weight $\bar{\lambda}$. The new principal path now goes to the left after δ , and stops there. The label of the new principal node is b . *Note that in case (4a), the new principal node is a left child, and its label is larger than all the labels on the principal path of its parent's (nonempty) right subtree.*

This completes the analysis of the case $S = U$, $a < b$. The analysis for $S = U$, $a > b$ is similar (the major difference being in the case $d = b$) and we will provide fewer details. The principal path contains a node with label $> b$, and m , the minimum of all labels on the principal path, can be $> b$ or $\leq b$. The first case gives rule (2b): add a new root, and make T its left subtree. The new edge has label λ , so the weight changes as required. *Note that in case (2b), the root is the principal node, and all the labels on the principal path of its (nonempty) left subtree are larger than the root label.*

If $S = U$, $a > b$, and $m \leq b$. On the principal path, there is at least one node with label at most b . Let δ be the last node on the principal path with that property, and let $d \leq b$ be its label. Again, the children of δ must have a label different from d , and that means that

the principal path goes right from δ . Again, we have two options, $d < b$ and $d = b$. The first case gives rule (3b). Replace the right child of δ with a new node β with label b , and make the old right subtree of δ the left subtree of the new node (which has no right child). The new tree is Smirnov and its weight is the previous weight times λx_b . The new principal path again ends at β , which has label b . *Note that in case (3b), the new principal node is a right child whose parent has a smaller label, and all the labels on the principal path of its (nonempty) left subtree are larger.*

On the other hand, if $S = U$, $a > b$, $m \leq b$, and $d = b$, we have rule (5a). We insert a new node β with label b as a right child of δ (not left child as in case (4a)!), make its left subtree the old left subtree of δ , and make the old right subtree of δ its left subtree. Since the root of the old right subtree of δ has label $> b$, this is still a Smirnov tree. The difference in weights comes from the edge between δ and its new left child, which has weight λ , and, as always, from the new node β . The new principal path ends at β . *Note that in case (5a), the new principal node is a right child, and its label is the same as its parent's label and smaller than all the labels on the principal path of its parent's (nonempty) left subtree. The new principal node might have no children (if δ had no left child in T).*

This leaves us with the case when S is a tree with, say, root label c . The analysis is similar and we will be brief.

The first case is when c is “on the same side” of b as a , i.e. if $a, c < b$ (use (1b)) or $a, c > b$ (use (1c)). Apart from adding the new node β as the right child of α , we also add the tree S as the left subtree of β . If $a < b$, the new edges have weights $\bar{\rho}$ and $\bar{\lambda}$, and if $a > b$, the weights are ρ and λ . Of course, we also added the tree S , which contributes $\text{wt}(S)$ to the weight of the new tree. Thus in both cases, the weight of the new tree equals $\text{wt}(T, S, b)$, and the new principal node is β . *Note that in case (1b), the new principal node is a right child and has a left child, and its label is larger than the labels of its parent and its left child. In case 1c, the new principal node is a right child and has a left child, and its label is smaller than the labels of its parent and its left child.*

Otherwise, c is “on the other side” of b as a : we have either $a < b \leq c$ or $a > b \geq c$. Let us assume that $a < b \leq c$. Again, we look at the labels on the principal path of T . If they are all smaller than b , we have case (2c). Apart from adding the new root and making T its left subtree, we also add S as the right subtree of α . The new edges have weights $\bar{\lambda}$ (since the principal label is $< b$) and $\bar{\rho}$ (since $a < c$), which ensures that the weight of the new tree is $\text{wt}(T, S, b)$. *Note that in case (2c), the root is the principal node, its label is larger than the label of its left child, but there is a node on the principal path of its left subtree (the root of S !) that has a label at least as large.*

We can also have nodes on the principal path that have labels $\geq b$. Again, let δ be the last such node, and let $d \geq b$ be its label (different from the nodes of δ 's children). If $d > b$, we have rule (3c): do the same as in case (3a), but also add the tree S as the right subtree of α . *In case (3c), the new principal node is a right child whose parent has a larger label, the left child has a smaller label, and there is a node on the principal path of its left subtree with a label that is at least as large.*

If, however, $d = b$, we are in case (4b), which is similar to (4a). *In case (4b), the new principal node is a left child and its label is larger than the label of its parent's right child, but there is a label on the principal path of its parent's right subtree with a label that is at least as large.*

This leaves us with the analysis of the cases when $S \in \mathcal{T}^c$ and $a > b \geq c$. If $m > b$, we have case (2d), which has the same construction as (2c), but new edge weights λ and ρ instead of $\bar{\lambda}$ and $\bar{\rho}$. Of course, this is consistent with the requirement that $a > b$, $S \in \mathcal{T}$ adds weight $\lambda \rho w(S)x_b$. *In case (2d), the root is the principal node, its label is smaller than the label of its left child, but there is a node on the principal path of its left subtree that has a label at most as large.*

If $m \leq b$ and $d < b$, we have case (3d). Again, do the same thing as in case (3c). *Note that in case (3d), the new principal node is a right child whose parent has a smaller label, the left child has a larger label, and there is a node on the principal path of its left subtree with a label that is at most as large.*

Finally, consider the case $m \leq b$, $d = b$, which is rule (5b). The situation is similar as in (5a), except that we add the tree S as the right child of α . *In case (5b), the new principal node is a right child, its label is the same as its parent's label and smaller than the label of its parent's left child, but there is a label on the principal path of its parent's left subtree with a label that is at most as large. The new principal node might not have children.*

This concludes the proof that Φ is a well-defined weight-preserving map. The proof is finished if we construct the inverse. The detailed analysis of the previous paragraphs makes that easy. We are given a tree T' with at least two nodes, and we want to see that there is a unique triple (T, S, b) with $\Phi(T, S, b) = T'$. Clearly, b must be $a(T')$; denote $\alpha(T')$ by β .

As the principal node, β cannot have a right child. Let us first assume that it has no parent, i.e. that it is the root of T' . Since T' is not a single node, it must have a nonempty left subtree T . We obtain T' by applying one of the two rules to T or to T with a subtree removed. More precisely, if all the labels on the principal path of the left subtree of the root are smaller (resp. larger) than b , then $T' = \Phi(T, U, b)$ is the result of rule (2a) (resp. (2b)). If, however, the left child of the root has a label that is smaller (resp. larger) than b , but there is a node on the principal path of T whose label is $\geq b$ (resp. $\leq b$), denote the first such node γ . Then we obtain T by applying rule (2c) (resp. (2d)) to $(T \setminus S, S, b)$, where S is the subtree of T with root γ .

Now assume that β is not a root. The first case is that it is a left child. By construction of the principal path and the fact that T' is Smirnov, that must mean that its parent's label is also b , and that the parent has a right child with a smaller label. If all the labels on the principal path of β 's parent's right subtree are smaller than b , we obtain T' by applying rule (4a) to (T, U, b) , where T is obtained from T by deleting the node β (if β has a left subtree, simply make this subtree β 's parent's left subtree). If, on the other hand, γ is the first node on the principal path of β 's parent's right subtree with label $\geq b$ and S is the subtree of T' with root γ , then $T' = \Phi(T \setminus S, S, b)$.

The last (and most common, as the principal path prefers the right direction) case is when

β is the right child. The first case is when its parent's label is also b . That must necessarily mean that β 's parent has a left child with label $> b$. If all the labels on the principal path of β 's parent's left subtree are $> b$, we obtain T' as an application of rule (5a) to (T, U, b) , where T is the tree we obtain if we remove β and swap the resulting left and right trees of β 's parent. Otherwise, T' results by using rule (5b).

Now we can assume that β is the right child, and its parent has a different label. If β has no children, we can produce it by using rule (1a) on $(T' \setminus \{\beta\}, D, b)$. If β has a left child, it will necessarily have a different label. If the labels of β 's parent and child are both smaller than b , T' is obtained by rule (1b), and if they are both larger than b , by rule (1c).

That leaves us with the case when one of the labels is larger and the other one is smaller. If the parent's label is larger (resp. smaller) and all the labels on the principal path of β 's left subtree are smaller (resp. larger), we applied rule (3a) (resp. (3b)) to get T' . If, however, the parent's label is larger (resp. smaller), the left child's label is smaller (resp. larger), but there is a node on the principal path of β 's child whose label is at least b (resp. at most b), we applied rule (3c) (resp. (3d)). This completes the construction of the inverse, and the proof of the theorem. \square

A suggested exercise for the reader is to take the tree drawn just before the proof and reproduce the word w and the sequence \mathcal{S} .

5. FINAL REMARKS

- (1) As mentioned in the introduction, Smirnov words can be interpreted as proper colorings of path graphs. Smirnov trees, on the other hand, are a labeled tree-analogue of Smirnov words. This raises the natural question of constructing labeled (binary) tree analogues for graphs other than path graphs, and see if one can define ascent-descent statistics on these tree analogues that relate to the Shareshian-Wachs chromatic quasisymmetric function. Another potential avenue is to consider the directed or cyclic analogue of the aforementioned question given the recent work of Ellzey-Wachs [5] and Panova-Alexandersson [10].
- (2) As established in Subsection 3.3, the exponential specialization of G tracks ascents-descents in standard labeled binary trees. In particular, this implies the equalto preserveities

$$\begin{aligned} [x_1 \dots x_n]G(1, 1, 1, 1) &= n! \times \text{Cat}_n, \\ [x_1 \dots x_n]G(1, 1, 1, 0) &= (n + 1)^{n-1}, \end{aligned}$$

where Cat_n denotes the n th Catalan number. It is unclear to us how one could derive these equalities starting from our e -positive expansion for G in terms of weights of bleeding trees.

- (3) The h -positivity of ωG (here ω is the standard involution on the algebra of symmetric functions) raises the question of constructing a natural permutation representation that realizes ωG as its Frobenius characteristic. Stanley [15, Proposition 7.7], using

a recurrence due to Procesi [12], established that $\omega G(\bar{\rho}, \rho, 0, 0)$ can be realized as the generating function of the Frobenius characteristic of the representation of S_n on the cohomology of the toric variety associated with the Coxeter complex of type A_{n-1} . Stembridge [18] also constructed a symmetric group representation which realizes $\omega G(\bar{\rho}, \rho, 0, 0)$ as the Frobenius characteristic via a bijection from permutations to what he calls codes. It would be interesting to generalize these results to Smirnov trees.

The following are the characters that such permutation representations would have. For example, there are 288 standard trees on five nodes that have weight $\bar{\rho}^2 \bar{\lambda} \lambda$ (i.e., that have two right ascents, one left ascent and one left descent). According to line 7 of the last table, 94 of them should be fixed under the action of a transposition, 8 under the action of a 4-cycle etc.

	111	21	3		1111	211	22	31	4		
$\bar{\rho}^2, \rho^2, \bar{\lambda}^2, \lambda^2$	1	1	1	$\bar{\rho}^3, \rho^3, \bar{\lambda}^3, \lambda^3$	1	1	1	1	1		
$\bar{\rho}\rho, \bar{\rho}\lambda, \rho\lambda, \bar{\lambda}\lambda$	4	2	1	$\bar{\rho}^2\rho, \bar{\rho}^2\lambda, \bar{\rho}\rho^2, \bar{\rho}\lambda^2, \rho^2\bar{\lambda}, \rho\bar{\lambda}^2, \bar{\lambda}^2\lambda, \bar{\lambda}\lambda^2$	11	5	3	2	1		
$\bar{\rho}\bar{\lambda}, \rho\lambda$	5	3	2	$\bar{\rho}^2\bar{\lambda}, \bar{\rho}\bar{\lambda}^2, \rho^2\lambda, \rho\lambda^2$	17	9	5	5	3		
				$\bar{\rho}\rho\bar{\lambda}, \bar{\rho}\rho\lambda, \bar{\rho}\bar{\lambda}\lambda, \rho\bar{\lambda}\lambda$	44	16	8	5	2		
					11111	2111	221	311	32	41	5
				$\bar{\rho}^4, \rho^4, \bar{\lambda}^4, \lambda^4$	1	1	1	1	1	1	1
				$\bar{\rho}^3\rho, \bar{\rho}^3\lambda, \bar{\rho}\rho^3, \bar{\rho}\lambda^3, \rho^3\bar{\lambda}, \rho\bar{\lambda}^3, \bar{\lambda}^3\lambda, \bar{\lambda}\lambda^3$	26	12	6	5	3	2	1
				$\bar{\rho}^3\bar{\lambda}, \bar{\rho}\bar{\lambda}^3, \rho^3\lambda, \rho\lambda^3$	49	25	13	13	7	7	4
				$\bar{\rho}^2\rho^2, \bar{\rho}^2\lambda^2, \rho^2\bar{\lambda}^2, \bar{\lambda}^2\lambda^2$	66	22	10	6	4	2	1
				$\bar{\rho}^2\bar{\lambda}^2, \rho^2\lambda^2$	146	60	26	26	12	12	6
				$\bar{\rho}^2\rho\lambda, \bar{\rho}\rho^2\bar{\lambda}, \bar{\rho}\bar{\lambda}\lambda^2, \rho\bar{\lambda}^2\lambda$	237	73	29	18	10	5	2
				$\bar{\rho}^2\rho\bar{\lambda}, \bar{\rho}^2\bar{\lambda}\lambda, \bar{\rho}\rho^2\lambda, \bar{\rho}\rho\bar{\lambda}^2, \bar{\rho}\rho\lambda^2, \bar{\rho}\bar{\lambda}^2\lambda, \rho^2\bar{\lambda}\lambda, \rho\bar{\lambda}\lambda^2$	288	94	36	27	13	8	3
				$\bar{\rho}\rho\bar{\lambda}\lambda$	824	228	80	50	24	12	4

ACKNOWLEDGEMENTS

The authors would like to thank Guillaume Chapuy, Ira Gessel, Ian Goulden, Sean Griffin, David Jackson, Andrea Sportiello, Stephan Wagner, and Philip Zhang for interesting conversations.

REFERENCES

- [1] C. ATHANASIADIS, *Gamma-positivity in combinatorics and geometry*, Sém. Lothar. Combin. 77 ([2016-2018]), Art. B77i, 64 pp. 11
- [2] O. BERNARDI, *Deformations of the braid arrangement and trees*, Adv. Math. 335 (2018), 466–518. 2
- [3] S. CORTEEL, D. FORGE AND V. VENTOS, *Bijections between affine hyperplane arrangements and valued graphs*, European J. Combin. 50 (2015), 30–37. 2
- [4] B. DRAKE, *An inversion theorem for labeled trees and some limits of areas under lattice paths*, Ph.D. Dissertation, Brandeis University (2008). 2, 12

- [5] B. ELLZEY AND M. WACHS, *On enumerators of Smirnov words by descents and cyclic descents*, arxiv: <https://arxiv.org/abs/1901.01591>. 21
- [6] I. GESSEL, *Oberwolfach Reports (Enumerative Combinatorics)*, https://www.mfo.de/document/1410/OWR_2014_12.pdf, 2014, Page 709. 2
- [7] I. GESSEL, S. GRIFFIN AND V. TEWARI, *Labeled plane binary trees and Schur-positivity*, arxiv: <https://arxiv.org/abs/1706.03055>. 2
- [8] L. KALIKOW, *Symmetries in trees and parking functions*, Discrete Math. 256 (2002), 719–741. 12
- [9] P.A. MACMAHON, *Combinatory Analysis*, 2 volumes, Cambridge University Press, London, 1915–1916. 1
- [10] G. PANOVA AND P. ALEXANDERSSON, *LLT polynomials, chromatic quasisymmetric functions and graphs with cycles*, Disc. Math. 341 (2018), 3453–3482. 21
- [11] T. KYLE PETERSEN, *Eulerian numbers*, Birkhäuser Advanced Texts: Basler Lehrbücher, Birkhäuser/Springer, New York, 2015. 1
- [12] C. PROCESI, *The toric variety associated to Weyl chambers*, Mots, Lang. Raison. Calc., Hermès, Paris, 1990, 153–161. 22
- [13] J. SHARESHIAN AND M. WACHS, *Chromatic quasisymmetric functions*, Adv. Math. 295 (2016), 497–551. 2, 3, 7, 12
- [14] J. SHARESHIAN AND M. WACHS, *Gamma-positivity of variations of Eulerian polynomials*, arxiv: <https://arxiv.org/abs/1702.06666> 11
- [15] R. STANLEY, *Log-concave and unimodal sequences in algebra, combinatorics, and geometry*. Ann. New York Acad. Sci. 576 (1989), 500–535. 21
- [16] R. STANLEY, *A symmetric function generalization of the chromatic polynomial of a graph*, Adv. Math. 111 (1995), 166–1994. 2
- [17] R. STANLEY, *Enumerative Combinatorics vol. 2*, Cambridge University Press, 1999. 3
- [18] J. STEMBRIDGE, *Eulerian numbers, tableaux, and the Betti numbers of a toric variety*, Disc. Math. 99 (1992), 307–320. 22
- [19] V. TEWARI, *Gessel polynomials, rooks, and extended Linnial arrangements*, J. Combin. Theory Ser. A 163 (2019), 98–117. 2
- [20] V. TEWARI AND S. VAN WILLIGENBURG, *Permuted composition tableaux, 0-Hecke algebra and labeled binary trees*, J. Combin. Theory Ser. A 161 (2019), 420–452. 2

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF LJUBLJANA & INSTITUTE OF MATHEMATICS, PHYSICS AND MECHANICS, LJUBLJANA, SLOVENIA

E-mail address: matjaz.konvalinka@fmf.uni-lj.si

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF PENNSYLVANIA, PHILADELPHIA, PA 19104, USA

E-mail address: vvtewari@math.upenn.edu