PROJECTIONS OF THE ALDOUS CHAIN ON BINARY TREES: INTERTWINING AND CONSISTENCY

NOAH FORMAN, SOUMIK PAL, DOUGLAS RIZZOLO, AND MATTHIAS WINKEL

ABSTRACT. Consider the Aldous Markov chain on the space of rooted binary trees with n labeled leaves in which at each transition a uniform random leaf is deleted and reattached to a uniform random edge. Now, fix $1 \le k < n$ and project the leaf mass onto the subtree spanned by the first k leaves. This yields a binary tree with edge weights that we call a "decorated k-tree with total mass n." We introduce label swapping dynamics for the Aldous chain so that, when it runs in stationarity, the decorated k-trees evolve as Markov chains themselves, and are projectively consistent over $k \le n$. The construction of projectively consistent chains is a crucial step in the construction of the Aldous diffusion on continuum trees by the present authors, which is the $n \to \infty$ continuum analogue of the Aldous chain and will be taken up elsewhere. Some of our results have been generalized to Ford's alpha model trees.

1. INTRODUCTION

The Aldous chain [3, 31] is a Markov chain on the space of (rooted, for our purposes) binary trees with n labeled leaves. Each transition of the chain, called a down-up move, has two stages. First, a uniform random leaf is deleted and its parent branch point is contracted away. Next, a uniform random edge is selected, we insert a new branch point into the middle of that edge, and we extend a new leaf-edge out from that branch point. This is illustrated in Figure 1 where $n = 6$ and the leaf labeled 3 is deleted and re-inserted. The stationary distribution of this chain is the uniform distribution on rooted binary trees with n labeled leaves.

Binary trees with leaf labels are called "cladograms." Markov chains on cladograms are used in MCMC algorithms for phylogenetic inference. The Aldous chain is used in this role [4], alongside other chains such as the subtree prune and re-graft chain [9, 30], which makes larger changes to the tree structure at each step.

Despite its simplicity, it is difficult to understand how the structure of a tree evolves over time under the Aldous chain. In his original work, Aldous gave bounds on the relaxation time for the chain. Soon thereafter, Schweinsberg [31] showed that the relaxation time is $\Theta(n^2)$. More recently, Pal [23] has studied connections, initially proposed by Aldous [2], between the Aldous chain and an extension of Wright–Fisher diffusions to negative mutation rates. Specifically, suppose $(T_n(j))_{j \ge 0}$ is a stationary version of the Aldous chain on trees with n leaves. Let L_1 and L_2 be leaves sampled uniformly from $T_n(0)$ without replacement and let L_0 be the root of
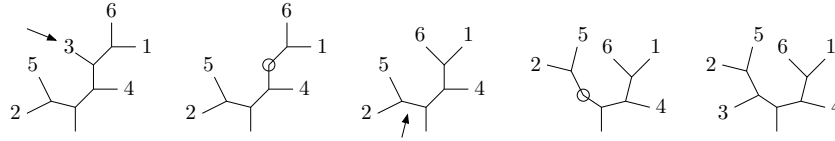
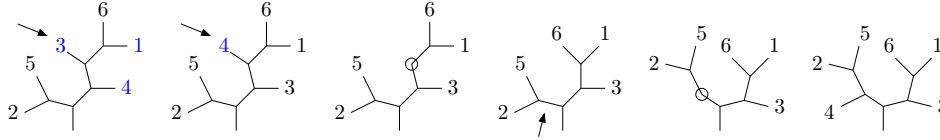FIGURE 1. From left to right, one Aldous down-up move.

FIGURE 2. The same down-up move as in Figure 1, but with label swapping as in Definition 1. Here, $i = 3, a = 1, b = 4$, so $\tilde{\imath} = 4$. We swap labels $i = 3$ and $\tilde{\imath} = 4$ before deleting the chosen leaf. Label 4 then regrows in the up-move.

$T_n(0)$. The subtree of $T_n(0)$ spanned by these leaves and the root consists of a root, a single branch point, and two leaves. This branch point naturally partitions the tree $T_n(0)$ into three components. As the Aldous chain runs, leaves move among components until the branch point disappears, i.e. a component becomes empty. If we let $m_i(j)$, $i \in \{0, 1, 2\}$, be the proportion of leaves of $T_n(j)$ in the component that initially contained L_i (stopped when the branch point disappears), then

$$(1) \qquad \left(m_0^{(n)}(\lfloor n^2 t \rfloor), m_1^{(n)}(\lfloor n^2 t \rfloor), m_2^{(n)}(\lfloor n^2 t \rfloor) \right)_{t \geq 0} \xrightarrow[n \to \infty]{d} (X_0(t), X_1(t), X_2(t))_{t \geq 0},$$

where $(X_0(t), X_1(t), X_2(t))_{t \geq 0}$ is a Wright–Fisher diffusion with mutation rate parameters $(1/2, -1/2, -1/2)$, stopped when one of the latter two coordinates vanishes. Similar results hold if we initially sample more leaves and look at more branch points, but the limiting description always only holds until the first time a branch point disappears. As the number of branch points goes to infinity the time for which the limiting description is valid tends to zero. Moreover, zero is an exit boundary for the coordinates of a Wright–Fisher diffusion that have negative mutation rates, so the limiting process does not shed light on how to continue beyond the disappearance of a branch point. In this paper, we give a mechanism through which one can select a new branch point when the original one disappears and, crucially, the process tracking the proportion of leaves in each component remains Markovian. In fact, we can do this while keeping track of any fixed number of branch points in a way that is projectively consistent over different numbers of branch points being tracked. We are currently working to understand the consequences of this construction for extending limits as in (1) beyond the first time a branch point disappears.

Our strategy is to modify the dynamics of leaf labels in the Aldous chain. Importantly, this does not affect the evolution of the unlabeled tree; it only affects the labels. These new label dynamics give rise to a natural mechanism for selecting new branch points about which to partition the tree when old branch points die.
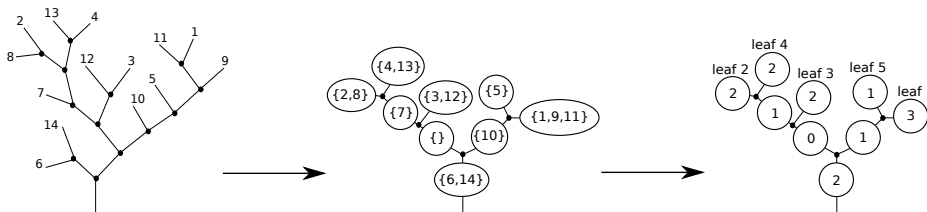
FIGURE 3. Left: $\mathbf{t} \in \mathbb{T}_{[14]}$. Middle: $\rho_5^{\star(14)}(\mathbf{t})$ – k-tree ($k = 5$) with edges marked by labels of collapsed leaves. Right: $\rho_5^{\bullet(14)}(\mathbf{t})$ – k-tree with edges decorated by leaf mass. See Section 3.

Given a finite set A, let $\mathbb{T}_A$ be the space of rooted binary trees with $\#A$ leaves, with the leaves labeled by the set A. See Definition 2 below. Similar to the Aldous chain, we introduce a new down-up chain on $\mathbb{T}_{[n]}$, where $[n] := \{1, 2, 3, \ldots, n\}$.

Definition 1 (Uniform chain). Fix $n \geq 3$. We define a Markov chain $(T(j))_{j \geq 0}$ on $\mathbb{T}_{[n]}$ in which each transition comprises a down-move followed by an up-move. Given $T(j) = \mathbf{t}_n \in \mathbb{T}_{[n]}$, $j \geq 0$, we randomly construct $T(j + 1)$ as follows.

  (i) Down-move: select the leaf labeled $i \in [n]$ for removal with probability $1/n$. Then compare i with the smallest leaf labels a and b in the first two subtrees on the ancestral path from leaf i to the root (with the convention $b = 0$ if the path has only one subtree). Let $\tilde{\imath} = \max\{i, a, b\}$. If $\tilde{\imath} \neq i$, swap labels i and $\tilde{\imath}$. Remove the leaf now labeled $\tilde{\imath}$ (which had been labeled i) and contract away its parent branch point. Call the resulting tree $\mathbf{t}_{n-1} \in \mathbb{T}_{[n]\setminus\{\tilde{\imath}\}}$.
  (ii) Up-move: insert a new leaf $\tilde{\imath}$ at an edge in $\mathbf{t}_{n-1}$ chosen uniformly at random.

Figure 2 shows one down-up move of this chain. Note that, if we erase the leaf labels, the above chain is identical to the Aldous chain. Somewhat surprisingly, although the dynamics of labels are not exchangeable in this down-up chain as they are in the Aldous chain, the two chains have the same stationary distribution.

Theorem 1. *The unique invariant distribution for the uniform chain is the uniform distribution on $\mathbb{T}_{[n]}$.*

To state our main result, we require an appropriate notion of projecting a tree $\mathbf{t} \in \mathbb{T}_{[n]}$ onto a subtree spanned by k leaves. We give an informal definition here; for a formal construction see Section 3. For $\mathbf{t} \in \mathbb{T}_{[n]}$ and $k \leq n$, we define $\rho_k^{\bullet(n)}(\mathbf{t}) := (\mathbf{s}, f)$, where $\mathbf{s}$ is the subtree of $\mathbf{t}$ spanned by $[k]$, reduced by contracting vertices of degree 2, and f is a weight function on the edge set of $\mathbf{s}$ that assigns weight to each edge equal to the number of leaves in the corresponding subtree of $\mathbf{t}$. See Figure 3. The following is our main result.

Theorem 2. *Let $(T(j))_{j \geq 0}$ be a uniform chain running in stationarity. Then, for $k \leq n$, $(\rho_k^{\bullet(n)}(T(j)))_{j \geq 0}$ is a Markov chain running in stationarity.*

In Definition 13 we explicitly describe the transition probabilities of the chains $(\rho_k^{\bullet(n)}(T(j)))_{j \geq 0}$. We remark that this theorem is not true for $\rho_k^{\bullet(n)}$-projections of the Aldous chain: relatively small examples show that those projected processes are not Markovian. It is also important that the uniform chain starts in stationarity.

There is a broad literature on when a function of a Markov chain is again Markov [8, 11, 29], but to our knowledge, Theorem 2 does not directly fit any known framework. In particular, the chains studied here resemble those in the literature on down-up chains on branching graphs [7, 14, 15, 20, 25]: (i) each transition comprises a down-move and an up-move, and (ii) the set of rooted, leaf-labeled binary trees forms a branching graph.

However, as Fulman notes [16], despite these similarities, the Aldous chain is not naturally amenable to these down-up methods. One of the impediments when working with either the Aldous chain or the uniform chain of Definition 1 is that the same label that is removed in the down-move is reinserted in the up-move. This dependence between down-move and up-move differs from the standard down-up framework. Our approach combines two well-known criteria for functions of Markov chains to be Markov: intertwining, which arises elsewhere in the down-up literature [6], and the Kemeny–Snell criterion [19], also known as Dynkin's criterion. See also [24]. The uniform n-tree down-up chain $(T(j))_{j\geq0}$ is intertwined on top of an intermediate Markov chain $(\rho_k^{\star(n)}(T(j)))_{j\geq0}$ (Figure 3, middle panel), which then relates to $(\rho_k^{\bullet(n)}(T(j)))_{j\geq0}$ via the Kemeny–Snell criterion.

1.1. **Connection to a conjecture of Aldous.** This paper is a part of a project by the current authors to resolve a long-standing conjecture by Aldous [2] that the Wright-Fisher limits of (1), taken consistently for any number of branch points, record aspects of an underlying "diffusion on continuum trees." This would be a continuum analogue of the Aldous Markov chain on finite trees. Since the stationary law of the Aldous chain, noted in Theorem 1, converges in a scaling limit to the Brownian continuum random tree (BCRT) [1], the latter should be the stationary distribution of the conjectured continuum process.

One way to study continuum tree structures is to randomly sample k leaves and consider the subtree generated by these k leaves and the root. In [1], Aldous constructs the BCRT by taking a limit of such projections that are consistent over k. One naturally wonders whether projections of the Aldous chain onto the subtree generated by the first k leaves of the n-tree give rise to projectively consistent Markov processes, thus enabling us to take a projective limit. However, this approach quickly runs into non-trivial road blocks. It has been argued elsewhere [23] that, in order to take a continuum limit, one must speed up time by a factor of n^2 for the Aldous chain on n-leaf binary trees. Consequently, in the limit, leaves become perfectly ephemeral: the entire leaf set dies at every instant. So how can we project onto the "same" leaves over time?

To get around this problem, Aldous proposed that one consider branch points rather than leaves. This leads to the challenges raised after (1) above, which are resolved by Theorem 2. In light of this theorem, our program for constructing the conjectured Aldous diffusion in future work involves constructing a continuum analogue to (an extension, discussed in the Appendix, of) the chains described in Theorem 2, for each k, and then using projective consistency in k to construct the full continuum-tree-valued process.

1.2. **Applications to statistics of partially classified hierarchical data.** Hierarchically arranged data sets are a common feature in areas such as topic modeling [5, 21] used in machine learning and natural language processing, as well as in phylogenetics [10]. Both types of data can be displayed as a leaf-labeled tree. The

leaves of a phylogenetic tree are labeled by a set of taxon (i.e. organism) names. The edges have lengths, corresponding to "time" between "mutations" in the tree. The goal of topic modeling is to classify a given set of documents according to a set of topics and to organize the topics in a hierarchy with more abstract topics near the root. This is the standard way to arrange books in libraries or assign AMS subject classifications (MSC2010) to mathematical research papers. The goal in natural language processing is to *discover* these topics from the words in the given documents. The leaves of the tree may be labeled by the documents themselves.

The decorated k-trees that we study in this article can be used as a model of hierarchically arranged clusters of data or as partially classified hierarchical data. As an example, consider again the AMS subject classifications. One such classification is *Probability theory and stochastic processes* (60XX). This classification has seventeen sub-classifications, including *Limit theorems* (60FXX) which has seven further sub-classifications, including *Large deviations* (60F10) and *None of the above but in this section* (60F99). Multiple papers are grouped into each classification. Each classification corresponds to a vertex of a tree. The leaves (say 60F10) can be thought of as having an associated mass given by the number (or proportion) of papers that are classified by that edge (i.e. papers on large deviation). On the other hand, papers tagged 60F99 contain possible further classifications which will all share the same three digits '60F' but are not fully known yet. The number of papers in this last category appear as a mass on the internal edge 60F, representing the number of leaves that would grow out of this edge if each paper received its own, individual classification. A similar data structure emerges when we consider a complete phylogenetic tree of k organisms and classify an additional m organisms according to the branch of the tree from which each has diverged.

To the best of our knowledge, no thorough statistical analyses have been done regarding these decorated trees, although models in non-parametric Bayesian analysis on nested Chinese restaurant processes [5] and nested hierarchical Dirichlet processes [21] are somewhat related. Our contribution to this research is the introduction of a natural model for decorated binary trees as a projection of large random binary trees (Section 3) and analyzing a sequence of natural Markov chains on the space of such trees that have explicit stationary distributions and are consistent over finer and coarser partial classifications (Section 4.1).

2. Down-up chains on binary trees

2.1. **Preliminaries.** For $k \in \mathbb{N}$, let $[k] := \{1, 2, \ldots, k\}$. We denote by $\mathbb{T}_{[k]}$ the set of rooted binary trees with k leaves labeled $1, \ldots, k$, and more generally for finite $A \subset \mathbb{N}$ by $\mathbb{T}_A$ the set of rooted binary trees with $\#A$ leaves labeled by the elements of the label set A. Specifically, we have in mind

$$\mathbb{T}_{[1]} = \left\{ \begin{array}{c} \text{①} \end{array} \right\}, \qquad \mathbb{T}_{[2]} = \left\{ \begin{array}{c} \text{①②} \end{array} \right\},$$

$$\mathbb{T}_{[3]} = \left\{ \begin{array}{c} \text{①②③} \end{array}, \quad \begin{array}{c} \text{①③②} \end{array}, \quad \begin{array}{c} \text{①②③} \end{array} \right\}, \quad \text{etc.,}$$

where we have ordered subtrees by their least labels to uniquely choose the plane tree representatives. It will be convenient to label each *edge* by the set of leaf labels in the subtree above the edge. This inspires the following definition.

Definition 2. Given a finite subset A of $\mathbb{N}$, a collection $\mathbf{t}$ of non-empty subsets of A is called a *rooted binary tree with leaves labeled by A* if the following hold.

(i) $\{j\} \in \mathbf{t}$ for all $j \in A$. These singleton sets are called the *external edges* of $\mathbf{t}$. Any other set B in $\mathbf{t}$ with $\#B \geq 2$ will be called an *internal edge*.

(ii) For all $B_1, B_2 \in \mathbf{t}$, we have $B_1 \cap B_2 = \varnothing$, or $B_1 \subseteq B_2$, or $B_2 \subseteq B_1$.

(iii) For all $B \in \mathbf{t} \setminus \{A\}$, there is a unique $B' \in \mathbf{t}$ such that $B' \cap B = \varnothing$ with $B \cup B' \in \mathbf{t}$; then B' is called the *sibling* edge of B, while $\overleftarrow{B} := B \cup B'$ is called the *parent* edge of B and B', and B and B' the *children* of $\overleftarrow{B}$. The sibling of the parent of B will be called the uncle of B.

For short, elements of $\mathbb{T}_{[n]}$ will be called *n-trees* (and elements of $\mathbb{T}_{[k]}$ *k-trees*).

Conditions (i) and (iii) imply that $A \in \mathbf{t}$ for all $\mathbf{t} \in \mathbb{T}_A$. This notion of a tree $\mathbf{t}$ is consistent with the graph-theoretic notion of a rooted tree with $\mathbf{t}$ as the set of edges between vertices. In that setting, there is a one-to-one correspondence between non-root vertices and their parental edges, so we can use edge labels to denote non-root vertices. For example, the edge labeled $\{j\}$ connects leaf j to its parent. Moreover, we denote the *root vertex* by the empty set $\varnothing$ and note that the edge A is the unique edge incident on it.

With this definition of a tree, we have

$$\mathbb{T}_{[1]} = \Big\{\{\{1\}\}\Big\}, \quad \mathbb{T}_{[2]} = \Big\{\{\{1\}, \{2\}, [2]\}\Big\},$$

and $\mathbb{T}_{[3]}$ consists of the three trees

$$\big\{\{1\}, \{2\}, \{3\}, \{2,3\}, [3]\big\}, \big\{\{1\}, \{2\}, \{3\}, \{1,3\}, [3]\big\}, \text{ and } \big\{\{1\}, \{2\}, \{3\}, \{1,2\}, [3]\big\}.$$

Given a tree $\mathbf{t}$ labeled by A, the *ancestral line* or *ancestral path* of an edge B_0 is the sequence of edges, starting with $B_1 = \overleftarrow{B_0}$, and, inductively, $B_{k+1} = \overleftarrow{B_k}$, until we reach $B_m = A$, for some m. This will be referred to as the *path from B_0 to the root*. Terms such as the *most recent ancestor* of B_0 satisfying some property will be naturally defined by this notion of ancestry. If $B_0 = \{j\}$ for some leaf labeled $j \in A$, there is a natural *spinal decomposition* of the tree $\mathbf{t}$ into the sequence of subtrees rooted along this ancestral line. That is, consider the vertices labeled by $B_1, B_2, \ldots, B_m$. At each B_k, there is a binary tree, rooted at B_k, whose leaves are labeled by B'_{k-1}. Since $\{j\}, B'_{k-1}, k = 1, 2, \ldots, m$, are disjoint and exhaust A, this gives a decomposition of $\mathbf{t}$. We will refer to the *kth subtree on this path* to refer to the subtree rooted at B_k.

If $C \subset A$ and $\mathbf{t}$ is a tree labeled by A, we define $\mathbf{t} \cap C$ to be the collection of non-empty elements of the set $\{B \cap C, B \in \mathbf{t}\}$. For $\mathbf{t} \in \mathbb{T}_A$ and $j \in A$, define

$$\mathbf{s} := \mathbf{t} - j := \mathbf{t} \cap (A \setminus \{j\}) = \{B \setminus \{j\}, B \in \mathbf{t}\} \setminus \{\varnothing\}$$

to be the tree $\mathbf{t}$ with leaf j removed. Note that the parent $\overleftarrow{\{j\}}$ and the sibling $\{j\}'$ in $\mathbf{t}$ have been identified as one edge in $\mathbf{s}$, and that all edges in $\mathbf{t}$ on the ancestral line of leaf j have lost j from their label sets in $\mathbf{s}$.

If we are given an edge $S \in \mathbf{s}$, we can add the leaf j as a sibling of S by splitting S into two edges S and $S \cup \{j\}$, and adding j to the label set of all edges on the ancestral line of S in $\mathbf{s}$. The resulting tree, denoted by $\mathbf{s} + (S, j)$, is therefore

$$\{B \in \mathbf{s} \colon B \cap S = \varnothing\} \cup \{B \in \mathbf{s} \colon B \subseteq S\} \cup \{B \cup \{j\} \colon B \in \mathbf{s} \text{ and } S \subseteq B\} \cup \{\{j\}\}.$$

We refer to the operation $(\mathbf{t}, j) \mapsto \mathbf{t} - j$ as the *deletion of leaf j from $\mathbf{t}$*, and to the operation $(\mathbf{s}, S, j) \mapsto \mathbf{s} + (S, j)$ as the *insertion of leaf j into $\mathbf{s}$ on edge S*.

It is now easy to see that for each fixed $j \in A$, we have

$$\mathbb{T}_A = \{\mathbf{s} + (S, j) \colon \mathbf{s} \in \mathbb{T}_{A \setminus \{j\}}, S \in \mathbf{s}\},$$

and that the representation of an element of $\mathbb{T}_A$ as $\mathbf{s} + (S, j)$ is unique so that, inductively, we obtain the well-known enumeration formula for numbers

$$(2) \qquad \#\mathbb{T}_A = \prod_{1 \leq i \leq \#A - 1} (2i - 1)$$

of rooted binary trees with $\#A$ labeled leaves.

2.2. Proof of Theorem 1 and generalization to Ford's alpha model trees.

The proof of Theorem 1 requires the concept of a *binary tree growth process* (see [26]). This is a sequence $(T_n, n \geq 1)$ of trees $T_n \in \mathbb{T}_{[n]}$ with the consistency property that $T_n - n = T_{n-1}$ for all $n \geq 2$. As this entails that $T_n = T_{n-1} + (S_{n-1}, n)$ for some $S_{n-1} \in T_{n-1}$, the conditional distribution of T_n given T_{n-1} is specified by the conditional distribution of S_{n-1} given T_{n-1}.

Definition 3 (Rémy's uniform growth process). Let T_1 be the unique element in $\mathbb{T}_{[1]}$. For every $n \geq 2$, let $T_n := T_{n-1} + (S_{n-1}, n)$, where conditionally given T_{n-1}, the random edge S_{n-1} is uniformly distributed on the set of edges of T_{n-1}. Then $(T_n, n \geq 1)$ is called Rémy's uniform binary tree growth process or just *uniform growth process* [28].

In particular, T_n is uniformly distributed on $\mathbb{T}_{[n]}$. We will denote the uniform distribution on $\mathbb{T}_{[n]}$ by $q_{n,1/2}$. The parameter $1/2$ alludes to a one-parameter family of tree growth process, due to Ford [12]. See also Pitman and Winkel [26].

Definition 4 (Ford's alpha growth process). Fix $\alpha \in [0, 1]$. Let T_1 and T_2 be the unique elements in $\mathbb{T}_{[1]}$ and $\mathbb{T}_{[2]}$, respectively. Now, inductively, given T_{n-1} for some $n \geq 3$, assign a weight $1 - \alpha$ to each external edge and a weight α to each internal edge of T_{n-1}, let S_{n-1} be a random edge of T_{n-1} whose distribution is proportional to these weights, and define $T_n = T_{n-1} + (S_{n-1}, n)$. Then $(T_n, n \geq 1)$ is called Ford's alpha model growth process or just *alpha growth process*.

For $\alpha = 1/2$, this is the uniform growth process. We denote the distribution of T_n under the alpha growth process by $q_{n,\alpha}$. The case for $\alpha = 0$ is also known as the Yule model while $\alpha = 1$ is called the Comb. Crucially, the leaf labels of T_n are only exchangeable in the uniform $\alpha = 1/2$ case. Naturally, one can ask if Theorem 1 generalizes to other values $\alpha \in [0, 1]$. The answer is yes and no. A slightly different Markov chain given below generalizes to all $\alpha \in [0, 1]$. However, for $\alpha = 1/2$, this does not give us the uniform chain of Definition 1, which is our main interest.

Definition 5 (Alpha chain). Fix $n \geq 3$. We define a Markov chain $(T(j))_{j \geq 0}$ on $\mathbb{T}_{[n]}$ in which each transition comprises a down-move followed by an up-move. Given $T(j) = \mathbf{t}_n \in \mathbb{T}_{[n]}$, $j \geq 0$, we randomly construct $T(j + 1)$, as follows.

  (i) Down-move: select the leaf labeled $i \in \{1, \ldots, n\}$ for removal with probability $1/n$. Compare i with the smallest leaf labels a and b in the first two subtrees on the ancestral path from leaf i to the root ($b = 0$ if the path has only one subtree). Let $\tilde{\imath} = \max\{i, a, b\}$. Swap labels i and $\tilde{\imath}$ and remove the leaf now labeled $\tilde{\imath}$ (which had been labeled i). Relabel leaves $\tilde{\imath}+1, \ldots, n$ by $\tilde{\imath}, \ldots, n-1$, respectively. This produces an element $\mathbf{t}_{n-1}$ in $\mathbb{T}_{[n-1]}$.
  (ii) Up-move: insert a new leaf labeled n according to the alpha growth process.
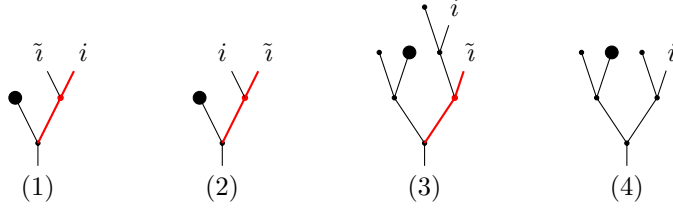
FIGURE 4. From left to right, we swap i and $\tilde{\imath}$ in (2) and insert two new leaves in (3). No leaves attach to the red edges. The final tree in (4) is obtained by dropping $\tilde{\imath}$ and the red edges.

Theorem 3. *For each $n \geq 3$ and each $\alpha \in [0, 1]$, the unique invariant distribution for the alpha chain on n-trees is $q_{n,\alpha}$.*

Key to the proof of Theorems 1 and 3 is a consistency property of the trees T_n, $n \geq 1$, in the alpha growth process under the down-move part of the alpha chain:

Lemma 4. *Consider the tree T_n in the alpha growth process. Fix $1 \leq i \leq \tilde{\imath} \leq n$, with $\tilde{\imath} \geq 2$. Let $E_{i,\tilde{\imath}}$ be the event that $\tilde{\imath} = \max(i, a, b)$ where a is the smallest label in the first subtree on the ancestral path from leaf i to the root in T_n, and b is the smallest label in the second subtree on that same path. Consider the conditional probability distribution $\mathbf{P}\left(\cdot \mid E_{i,\tilde{\imath}}\right)$. Swap leaves i and $\tilde{\imath}$, drop leaf $\tilde{\imath}$, and relabel leaves $\tilde{\imath}+1, \ldots, n$ by $\tilde{\imath}, \ldots, n-1$, respectively. Then the resulting tree, $\widetilde{T}_{n-1}$, under $\mathbf{P}\left(\cdot \mid E_{i,\tilde{\imath}}\right)$, is distributed as $q_{n-1,\alpha}$.*

Proof. Consider an alpha growth process $(T_k, k \geq 1)$. First we claim that the event $E_{i,\tilde{\imath}}$ is independent of $T_{\tilde{\imath}-1}$. This statement is trivial for $\tilde{\imath} = 2$, when $T_{\tilde{\imath}-1}$ is deterministic. Henceforth, assume $\tilde{\imath} \geq 3$. Now $E_{i,\tilde{\imath}}$ can be expressed as follows.

  (i) If $i < \tilde{\imath}$, consider the ancestral path of leaf i in $T_{\tilde{\imath}-1}$. There are at least two edges on this path (say, $\{i\}$ and $[\tilde{\imath}-1]$), and $E_{i,\tilde{\imath}}$ is equivalent to the event that $\tilde{\imath}$ gets attached to one of the first two edges on the ancestral path from leaf i to the root, and none of the leaves with label in $\{\tilde{\imath}+1, \ldots, n\}$ gets attached to the first two edges on the ancestral path of leaf i in $T_{\tilde{\imath}}$.
 (ii) Otherwise $i = \tilde{\imath}$. Then $E_{\tilde{\imath},\tilde{\imath}}$ is the event that none of the leaves with label in $\{\tilde{\imath}+1, \ldots, n\}$ gets attached to the first two edges on the ancestral path from $\tilde{\imath}$ to the root in $T_{\tilde{\imath}}$.

In case (i) the leaf $\tilde{\imath}$ has to get attached to a given pair of internal and external edges, and, in both cases (i) and (ii), all higher labeled leaves are conditioned not to get attached to a given pair of internal and external edges. Hence, in both cases, $E_{i,\tilde{\imath}}$ is independent of $T_{\tilde{\imath}-1}$, and $T_{\tilde{\imath}-1}$ remains distributed as $q_{\tilde{\imath}-1,\alpha}$ under the conditional law $\mathbf{P}\left(\cdot \mid E_{i,\tilde{\imath}}\right)$,.

Now, consider T_n under $\mathbf{P}\left(\cdot \mid E_{i,\tilde{\imath}}\right)$. Swapping i and $\tilde{\imath}$, and dropping the new leaf labeled $\tilde{\imath}$ and all the higher-labeled ones, we still obtain $T_{\tilde{\imath}-1}$, which is distributed as $q_{\tilde{\imath}-1,\alpha}$. See Figure 4. Under $\mathbf{P}\left(\cdot \mid E_{i,\tilde{\imath}}\right)$, given $T_{\tilde{\imath}-1}$, consider how new leaves get attached to form $\widetilde{T}_{n-1}$. The old leaf $\tilde{\imath}$ never gets attached, therefore the two red edges in Figure 4 (2) do not exist anymore. Leaves $(\tilde{\imath}+1, \ldots, n)$ have been relabeled $(\tilde{\imath}, \ldots, n-1)$. Consider the distribution of the edges $(S_{\tilde{\imath}}, \ldots, S_n)$ where these successive leaves get attached. Each new leaf still gets attached to every internal

edge with probability proportional to α and to an external edge with probability proportional to $(1 - \alpha)$. Hence, we still have an alpha growth process starting from $T_{\tilde{\imath}-1}$. Since the law of $T_{\tilde{\imath}-1}$ is $q_{\tilde{\imath}-1,\alpha}$, this completes the proof. $\qquad\square$

Proof of Theorems 1 *and* 3. The uniqueness of the invariant distribution in both theorems is immediate since the Markov chain is finite and either is irreducible (for $\alpha \in (0,1)$) or has a unique recurrent communicating class (for $\alpha \in \{0,1\}$).

Consider Theorem 3 and one step of the Markov chain. Suppose $T_n(0) \in \mathbb{T}_{[n]}$ is distributed as $q_{n,\alpha}$ and let I be uniformly chosen from $[n]$, independent of $T_n(0)$. For a given pair of $1 \le i \le \tilde{\imath} \le n$, $\tilde{\imath} \ge 2$, let $F_{i,\tilde{\imath}}$ be the event that $I = i$ is selected for the down-move and swapped with label $\tilde{\imath}$. Condition on $F_{i,\tilde{\imath}}$. Under this conditional distribution, apply Lemma 4 to get a tree $\widetilde{T}_{n-1}(0)$, which is distributed as $q_{n-1,\alpha}$. Add a leaf labeled n according to the alpha growth process to get a tree $T_n(1)$ that is distributed according to $q_{n,\alpha}$. Since the conditional distribution of $T_n(1)$ is $q_{n,\alpha}$, given $F_{i,\tilde{\imath}}$, irrespective of the pair $(i, \tilde{\imath})$, its unconditional distribution is also $q_{n,\alpha}$. This completes the proof of Theorem 3.

To prove Theorem 1, given $F_{i,\tilde{\imath}}$, get $T_n(1) \sim q_{n,1/2}$ as above. Relabel the new leaf n by $\tilde{\imath}$ and leaves labeled $(\tilde{\imath}, \ldots, n-1)$ by $(\tilde{\imath}+1, \ldots, n)$. By exchangeability of labels under $q_{n,1/2}$, this new tree is again uniformly distributed. But, of course, this is the same as simply inserting a new leaf labeled $\tilde{\imath}$ while leaving the higher-labeled leaves unchanged. Integrate over $F_{i,\tilde{\imath}}$, as before, to obtain the result. $\qquad\square$

Consider again the uniform chain from Definition 1. Let $\widetilde{I}$ denote the random label that is inserted in the first up-move. We say that the leaf $\widetilde{I}$ is resampled. As an application of our proof technique, we show that the distribution of the resampled label is not uniform over $[n]$. This is a special case of the alpha chain (Definition 5) where $\widetilde{I}$ now refers to the label removed in the down-move.

Corollary 5. *Suppose T_n is distributed according to $q_{n,\alpha}$ on $\mathbb{T}_{[n]}$. Consider one step of the alpha chain starting at T_n and let $\widetilde{I}$ denote the random leaf label removed in the down-move. Then*

$$(3) \qquad \mathbf{P}\left(\widetilde{I} = \tilde{\imath}\right) = \begin{cases} \dfrac{1}{n(n-1-\alpha)}, & \text{if } \tilde{\imath} = 2, \\ \dfrac{2\tilde{\imath} - 2 - \alpha}{n(n-1-\alpha)}, & \text{if } 3 \le \tilde{\imath} \le n. \end{cases}$$

In particular, for the stationary uniform chain, the case $\alpha = 1/2$ of the above gives the distribution of the resampled leaf label.

Proof. Consider the event $E_{i,\tilde{\imath}}$ from Lemma 4. First consider the special case of $\alpha = 1/2$ which is simpler. Assume $\tilde{\imath} \ge 3$ and $i < \tilde{\imath}$. Then, it follows from alternative expression (i) of $E_{i,\tilde{\imath}}$ that

$$\mathbf{P}\left(E_{i,\tilde{\imath}}\right) = \frac{2}{2\tilde{\imath} - 3} \prod_{j=\tilde{\imath}}^{n-1} \left(1 - \frac{2}{2j-1}\right) = \frac{2}{2\tilde{\imath} - 3} \prod_{j=\tilde{\imath}}^{n-1} \left(\frac{2j-3}{2j-1}\right) = \frac{2}{2n-3}.$$

Similarly, for $i = \tilde{\imath}$, $\mathbf{P}\left(E_{\tilde{\imath},\tilde{\imath}}\right) = \prod_{j=\tilde{\imath}}^{n-1} \left(1 - \frac{2}{2j-1}\right) = \prod_{j=\tilde{\imath}}^{n-1} \left(\frac{2j-3}{2j-1}\right) = \frac{2\tilde{\imath} - 3}{2n-3}.$

Let I be the leaf chosen uniformly in the down-move, independent of $T_n(0)$. Denote by $\widetilde{I}$ the random leaf label that gets removed in the down-step when I is selected. Integrating over the choice of $\{I = i\}$, $i \in [n]$, we get

$$\mathbf{P}\left(\widetilde{I} = \widetilde{\imath}\right) = \frac{1}{n} \sum_{i=1}^{\widetilde{\imath}} \mathbf{P}\left(E_{i,\widetilde{\imath}}\right) = \frac{1}{n} \left[\frac{2(\widetilde{\imath} - 1)}{2n - 3} + \frac{2\widetilde{\imath} - 3}{2n - 3} \right] = \frac{4\widetilde{\imath} - 5}{n(2n - 3)}, \quad \widetilde{\imath} \geq 3.$$

Then, when $\widetilde{\imath} = 2$, $\mathbf{P}\left(\widetilde{I} = 2\right) = 1 - \sum_{j=3}^{n} \mathbf{P}\left(\widetilde{I} = j\right) = \frac{2}{n(2n - 3)}$.

For any other $\alpha \in (0, 1)$, the proof is very similar. On the event $E_{i,\widetilde{\imath}}$, new leaves are forbidden to get attached to one internal and one external edge. The probability that leaf $j + 1$ gets attached to any given internal edge is $\alpha/(j - \alpha)$ and the probability that it gets attached to a given external edge is $(1 - \alpha)/(j - \alpha)$. Thus, for $\widetilde{\imath} \geq 3$,

$$\mathbf{P}\left(E_{i,\widetilde{\imath}}\right) = \begin{cases} \dfrac{1}{\widetilde{\imath} - 1 - \alpha} \displaystyle\prod_{j=\widetilde{\imath}}^{n-1} \left(1 - \dfrac{1}{j - \alpha}\right) = \dfrac{1}{n - 1 - \alpha}, & \text{for } 1 \leq i < \widetilde{\imath}, \\[4mm] \displaystyle\prod_{j=\widetilde{\imath}}^{n-1} \left(1 - \dfrac{1}{j - \alpha}\right) = \dfrac{\widetilde{\imath} - 1 - \alpha}{n - 1 - \alpha}, & \text{for } \widetilde{\imath} = i. \end{cases}$$

Summing up,

$$\mathbf{P}(\widetilde{I} = \widetilde{\imath}) = \frac{1}{n} \left[\frac{\widetilde{\imath} - 1}{n - 1 - \alpha} + \frac{\widetilde{\imath} - 1 - \alpha}{n - 1 - \alpha} \right] = \frac{2\widetilde{\imath} - 2 - \alpha}{n(n - 1 - \alpha)}, \quad \widetilde{\imath} \geq 3,$$

and, as above, $\mathbf{P}(\widetilde{I} = 2) = 1/n(n - 1 - \alpha)$. This completes the proof. $\qquad\square$

3. Collapsing a tree and decorating edges with masses

Definition 6 (Decorated k-trees of mass n)**.** Fix integers $n \geq k \geq 1$. For $\mathbf{s} \in \mathbb{T}_{[k]}$, denote by $\mathcal{N}_n^{\mathbf{s}}$ the set of functions $f \colon \mathbf{s} \to (\mathbb{N} \cup \{0\})$ that satisfy two conditions

(i) $\sum_{B \in \mathbf{s}} f(B) = n$ and
(ii) $f(\{j\}) \geq 1$ for each $j \in [k]$.

Then the set of *k-trees decorated with total edge mass n*, or *decorated k-trees* for short, is given by

$$\mathbb{T}_{[k]}^{\bullet(n)} := \bigcup_{\mathbf{s} \in \mathbb{T}_{[k]}} \{\mathbf{s}\} \times \mathcal{N}_n^{\mathbf{s}}.$$

We refer to an element $\mathbf{t}^{\bullet} = (\mathbf{s}, (x_1, \ldots, x_k, (y_B)_{B \in \mathbf{s}:\ \#B \geq 2})) \in \{\mathbf{s}\} \times \mathcal{N}_n^{\mathbf{s}}$ as a tree $\mathbf{s}$ with integer masses attached to every edge. The labeled external edges have positive masses $x_1, \ldots, x_k$. Each internal edge B has a nonnegative mass y_B. We say that $\mathbf{t}^{\bullet}$ has *tree shape* $\mathbf{s}$.

A natural way to obtain decorated k-trees is to project the leaf mass from an n-tree to the subtree spanned by the first k leaves. We describe this concept formally.

Let $\mathbf{t} \in \mathbb{T}_{[n]}$ and $k \leq n$. We will associate a decorated tree $\mathbf{t}^{\bullet} \in \mathbb{T}_{[k]}^{\bullet(n)}$, as follows. Define $\mathbf{s} := \mathbf{t} \cap [k]$. For every $j \in [n]$, consider the ancestral line of the leaf j, $\overleftarrow{\{j\}} = B_1^{(j)} \subseteq B_2^{(j)} \subseteq \cdots \subseteq B_{m_j}^{(j)} = [n]$. Additionally, let $B_0^{(j)} = \{j\}$. In this ordering let B_j be the most recent ancestor $B_i^{(j)}$ of $\{j\}$, including $\{j\}$ itself, such that $B_i^{(j)} \cap [k] \neq \varnothing$. Let μ be the measure on $\mathbf{t}$, of total mass n, that assigns
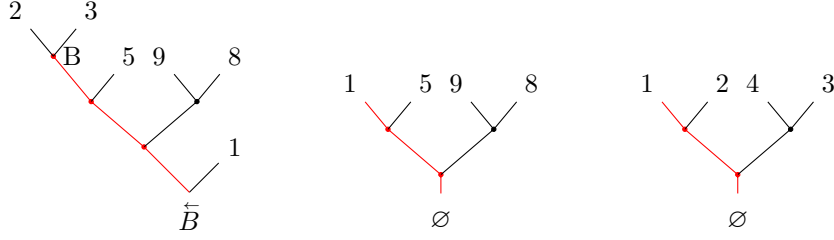
FIGURE 5. On the left, consider the subtree generated by leaves [3]. The three red edges contract to an internal edge B in the 3-tree. The unranked internal structure $\mathbf{v}_B$ is shown in the middle. The vertex $\overleftarrow{B}$ is replaced by the root, while B itself is replaced by a leaf labeled 1. The internal structure $\mathrm{int}(\mathbf{v}_B)$ is on the right.

mass 1 to each of the n leaves. Consider the projection maps $\pi \colon [n] \to \mathbf{s}$ given by $j \mapsto B_j \cap [k]$. The push-forward of μ by this projection assigns a nonnegative integer value to every edge of $\mathbf{s}$ producing a decorated k-tree $\mathbf{t}^\bullet$ as in Definition 6. Since $B_j = \{j\}$ for each $j \in [k]$, the projected mass on each leaf of $\mathbf{s}$ is at least one. The total mass is n. We denote the map that assigns the decorated k-tree $\mathbf{t}^\bullet \in \mathbb{T}_{[k]}^{\bullet(n)}$ to the tree $\mathbf{t} \in \mathbb{T}_{[n]}$ by $\rho_k^{\bullet(n)} \colon \mathbb{T}_{[n]} \to \mathbb{T}_{[k]}^{\bullet(n)}$, $\mathbf{t} \mapsto \rho_k^{\bullet(n)}(\mathbf{t}) = \mathbf{t}^\bullet$.

For later use we also introduce an intermediate stage between the full n-tree and the decorated k-tree with mass n obtained from it.

Definition 7 (Collapsed n-trees with k leaves)**.** Consider an n-tree $\mathbf{t}$ and a decorated k-tree $\mathbf{t}^\bullet$ obtained by the projection map $\pi \colon [n] \to \mathbf{s} = \mathbf{t} \cap [k]$. Let π^{-1} be the pre-image function of π. That is, for every edge B of $\mathbf{s}$, we associate a subset of $[n]$ given by $\pi^{-1}(B)$. Thus $\mu(B) = \#\pi^{-1}(B)$ is the number of elements in $\pi^{-1}(B)$. The resulting object

$$\mathbf{t}^\star := \left(\mathbf{s}, \left(\pi^{-1}(B) \right)_{B \in \mathbf{s}} \right)$$

will be called a *collapsed n-tree with k leaves* or just a *collapsed n-tree*. Denote the set of all collapsed n-trees with k leaves (as $\mathbf{t}$ varies over $\mathbb{T}_{[n]}$) by $\mathbb{T}_{[k]}^{\star(n)}$. The map from $\mathbb{T}_{[n]}$ to $\mathbb{T}_{[k]}^{\star(n)}$ that takes $\mathbf{t}$ to $\mathbf{t}^\star$ will be denoted by $\rho_k^{\star(n)}$. That is, $\mathbf{t}^\star = \rho_k^{\star(n)}(\mathbf{t})$.

See Figure 3 for examples of decorated and collapsed trees obtained by projecting a tree with $n = 14$ leaves on the subtree spanned by the first $k = 5$ leaves.

Fix $1 \le k < n$ and let $\mathbf{t}^\star = \left(\mathbf{s}, \left(\pi^{-1}(B) \right)_{B \in \mathbf{s}} \right) \in \mathbb{T}_{[k]}^{\star(n)}$ be obtained by collapsing $\mathbf{t} \in \mathbb{T}_{[n]}$. The collection of subsets $\left(\pi^{-1}(B) \right)_{B \in \mathbf{s}}$ is an ordered partition of $[n]$, ordered by the edges of $\mathbf{s}$. Fix an edge B of $\mathbf{s}$. Let $\mathbf{r}$ be the subtree of $\mathbf{t}$ spanned by the leaves $1, 2, \ldots, k$, including any degree-two vertices. Label the edges of $\mathbf{r}$ by their corresponding labels in $\mathbf{t}$. Consider the set of edges S of $\mathbf{r}$ such that $S \cap [k] = B$ and consider the vertices corresponding to those edges. Then, there is a sequence of subtrees of $\mathbf{t}$ with leaf labels in $\pi^{-1}(B)$ that are rooted at those vertices. Along with the edges S of $\mathbf{r}$ such that $S \cap [k] = B$, this gives us a binary tree $\mathbf{v}_B$ rooted at the vertex corresponding to $\overleftarrow{B}$, the parent of B in $\mathbf{s}$. Note that, if B is an internal edge of $\mathbf{s}$, the vertex corresponding to B becomes a new leaf of $\mathbf{v}_B$. We will always label this leaf by 1 in $\mathbf{v}_B$. Other leaves of $\mathbf{v}_B$ retain their labels from $\mathbf{t}$, including

the case when B is itself an external edge in $\mathbf{s}$. In particular, the leaf set of $\mathbf{v}_B$ is $\pi^{-1}(B)$ if B is external in $\mathbf{s}$ and $\pi^{-1}(B) \cup \{1\}$, if B is internal. These labelings are local in the sense that leaves labeled 1 appear in $\mathbf{v}_B$ for all internal edges B in $\mathbf{s}$. See Figure 5 for an example where B is internal.

Definition 8 (Internal structure). Let $A = \{a_1, \ldots, a_k\} \subseteq [n]$, $k \leq n$, where the elements are arranged in increasing order $a_1 < a_2 < \cdots < a_k$. Then we associate with any tree in $\mathbb{T}_A$ the tree in $\mathbb{T}_{[k]}$ obtained by relabeling each leaf a_i by its rank i. In the case where we apply this relabeling to $\mathbf{v}_B$ for $\mathbf{t} \in \mathbb{T}_n$, $B \in \mathbf{t} \cap [k]$, we use the function notation $\mathrm{int}(\mathbf{v}_B) \in \mathbb{T}_{[\#\pi^{-1}(B)]} \cup \mathbb{T}_{[\#\pi^{-1}(B)+1]}$. We will refer to this tree as the *internal structure* of $\rho_k^{\bullet(n)}(\mathbf{t})$ supported on B, while $\mathbf{v}_B$ itself will be called the *unranked internal structure* of $\rho_k^{\bullet(n)}(\mathbf{t})$ supported on B.

We write $\widetilde{q}_{n,\alpha}$ for the law of the n-leaf tree formed by Ford growth, as in Definition 4, with the exception that new leaves select to grow on edge $\{1\}$, specifically, with probability proportional to α, rather than $1 - \alpha$. Note that $\widetilde{q}_{n,\alpha} = q_{n,\alpha}$ for $\alpha = \frac{1}{2}$.

Lemma 6 (Spatial Markov property of $q_{n,\alpha}$). *Fix $\alpha \in [0,1]$. If $T_n \sim q_{n,\alpha}$ conditioned on $\rho_k^{\bullet(n)}(T_n) = \big(\mathbf{s}, (\mu(B))_{B \in \mathbf{s}}\big)$, then the random internal structures $(\mathrm{int}(V_B), B \in \mathbf{s})$ of $\rho_k^{\bullet(n)}(T_n)$ are mutually independent, jointly independent of $\big(\pi^{-1}(B)\big)_{B \in \mathbf{s}}$, and each $\mathrm{int}(V_B)$ is distributed according to $q_{\mu(B),\alpha}$ if B is an external edge, or $\widetilde{q}_{\mu(B)+1,\alpha}$ otherwise. In particular, these structures remain independent when conditioned on $\rho_k^{\star(n)}(T_n) = \mathbf{t}^{\star} = \big(\mathbf{s}, \big(\pi^{-1}(B)\big)_{B \in \mathbf{s}}\big)$.*

Proof. Consider Ford's alpha growth process starting from $\mathbf{s}$. It is clear from the definition that, given that a label i is in $\pi^{-1}(B)$ for some $B \in \mathbf{s}$, the location where it grows in $V_B \cap [i-1]$ is independent of the internal structures of all other subtrees growing on other edges. Moreover, the probability that i will grow out of a given edge in $V_B \cap [i-1]$ is as in Definition 4 of Ford growth, with the exception that if B is internal, then the external edge $\{1\}$ in V_B, which corresponds to an internal edge of $T_n \cap [i-1]$, is selected with probability proportional to α, instead of $1 - \alpha$. Therefore, given $\mathbf{t}^{\star}$, each $\mathrm{int}(V_B)$ is distributed as $q_{\mu(\{i\}),\alpha}$ if $B = \{i\}$ is external or as $\widetilde{q}_{\mu(B)+1,\alpha}$ if B is internal, and these internal structures are jointly independent of one another. Since, conditionally given $(\mu(B), B \in \mathbf{s})$, the joint distribution of the internal structures is independent of the leaf labels $\big(\pi^{-1}(B), B \in \mathbf{s}\big)$, the statement of the lemma follows. $\square$

Lemma 7 (Regenerative spinal compositions). *Fix $\alpha \in (0,1)$ and consider $T_n \sim \widetilde{q}_{n,\alpha}$. Let U denote the first subtree along the spine from $\{1\}$ to the root and M_1 the number of leaves in U. Then*

$$(4) \qquad \mathbf{P}(M_1 = m) = \delta_\alpha(n : m) := \alpha \binom{n}{m} \frac{\Gamma(m - \alpha)\Gamma(n - m + \alpha)}{\Gamma(1 - \alpha)\Gamma(n + \alpha)}, \quad m \in [n].$$

Moreover, if V is T_n with U removed but leaf 1 retained, then $\mathrm{int}(U)$ is conditionally independent of $\mathrm{int}(V)$ given M_1, with $\mathrm{int}(U) \sim q_{M_1,\alpha}$ and $\mathrm{int}(V) \sim \widetilde{q}_{n-M_1,\alpha}$.

We remark that, when $\alpha = \frac{1}{2}$, (4) is also the probability that the first return to the origin of a simple random walk bridge of length $2n$ occurs at time $2m$:

$$(5) \qquad \delta(n : m) := \delta_{1/2}(n : m) = \frac{1}{2m - 1} \binom{2m}{m} \binom{2n - 2m}{n - m} \bigg/ \binom{2n}{n}, \quad m \in [n].$$

Proof. Think of T_n as growing according to Ford's alpha growth, with the modification noted above Lemma 6, and consider how subtrees grow along the spine to leaf 1. At each step, prior to adding the new leaf, suppose the spinal subtrees have masses $m_1, \ldots, m_k$. The new leaf attaches to one of the j^{th} subtree with probability proportional to $m_j - \alpha$, or begins a new spinal subtree with probability proportional to α at any of the $k+1$ possible sites between adjacent spinal subtrees or at the top or bottom of the spine. This is the growth rule for an (α, α) ordered Chinese restaurant process [26]. As noted in [26], this means that after $\mu(B)$ growth steps, the sequence of subtree masses going down the spine is an (α, α) regenerative composition structure, as in [17, Section 8.4]. Therefore, the mass of the first such subtree is distributed according to the *decrement matrix*, which we have quoted from [17] in (4). This proves the first assertion.

By the same argument as in the proof of Lemma 6, the spinal subtrees, including U, are conditionally independent of each other given the sequence of their leaf counts $(M_i)_{i \in [K]}$, each with law $q_{M_i, \alpha}$. This gives the claimed conditional law for U. Moreover, by definition of regenerative composition structures, the subtree masses down the spine from leaf 1 in V also form an (α, α) regenerative composition structure. And again, by the argument in Lemma 6, this sequence of masses is conditionally independent of U given M_1. This is enough information to specify the conditional law of $\text{int}(V)$, given U, as $\widetilde{q}_{n - M_1, \alpha}$. □

4. DOWN-UP CHAINS ON DECORATED TREES

4.1. **Definitions.** We now show that both the alpha chain and the uniform chain induce Markov chains on decorated k-trees with total mass n. The idea is very general. Suppose we have a time-homogeneous Markov chain $(\mathcal{X}(j))_{j \geq 0}$, with a finite state space $\mathbb{X}$. Let P denote the transition matrix for this Markov chain. Let $\mathbb{Y}$ be another finite space and let $g \colon \mathbb{X} \to \mathbb{Y}$ be a surjective function. For each $y \in \mathbb{Y}$, suppose we have a probability function $\Lambda(y, \cdot)$ on $\mathbb{X}$. Then, one can define a Markov operator S acting on any function $h \colon \mathbb{Y} \to \mathbb{R}$ by

$$(6) \qquad S(h)(y) := \sum_{x \in \mathbb{X}} \Lambda(y, x) \, \mathbf{E}\left[h \circ g(\mathcal{X}(1)) \mid \mathcal{X}(0) = x\right].$$

We will denote the corresponding transition matrix on $\mathbb{Y}$ by $Q = \Lambda P g$, where we abuse notation and write g for the transition matrix $(\mathbf{1}_{(g(x) = y)})_{x \in \mathbb{X}, y \in \mathbb{Y}}$ that corresponds to deterministic transitions from each $x \in \mathbb{X}$ to $g(x) \in \mathbb{Y}$. If $(\mathcal{X}(j))_{j \geq 0}$ has a stationary distribution q, a natural choice for $\Lambda(y, \cdot)$ is the conditional distribution $q(\,\cdot \mid g = y)$. In particular, this holds when $(\mathcal{X}(j))_{j \geq 0}$ is the alpha chain (or, the uniform chain), q is $q_{n, \alpha}$ (or, the uniform distribution) on $\mathbb{T}_{[n]}$ and $g = \rho_k^{\bullet(n)}$. In this section we show that the resulting Markov transition kernel Q on $\mathbb{T}_{[k]}^{\bullet(n)}$ has an autonomous description that does not refer to the chains on n-trees. This is useful if we wish to send n to infinity while keeping k fixed. Throughout, we will consider the parameter α restricted to the open interval $(0, 1)$.

Recall the Dirichlet-multinomial (DM) distributions. This is a family of distributions that takes two parameters, a positive integer m and a vector v in $(0, \infty)^d$, for some $d \geq 2$. Given such m and v, this is a discrete probability distribution on the set of nonnegative integers $(j_1, \ldots, j_d)$ that add up to m that can be described as follows. Pick $(p_1, \ldots, p_d)$ from a Dirichlet distribution with parameters v. Let $(J_1, \ldots, J_d)$ be conditionally distributed according to a multinomial distribution of

counts when sampling with replacement m balls of d colors, where at each step the ball of color i has a probability p_i of getting sampled. Then the marginal distribution of $(J_1, \ldots, J_d)$, when integrated over $(p_1, \ldots, p_d)$, is DM with parameters m and v. For our purpose we take $d = 2k - 1$, for some $1 \leq k < n$, and, for some $\alpha \in (0, 1)$, the first k coordinates of v are equal to $1 - \alpha$ and the rest are equal to α. Denote this distribution by $\mathrm{DM}_{2k-1}^\alpha(m)$. We also define $\mathrm{DM}_{2k-1}^\alpha(0)$ to be the probability distribution that puts mass one on the zero-vector of dimension $2k - 1$.

Just as in the case of n-trees, the induced Markov chains on decorated k-trees involve a down-move followed by an up-move. However, the down- and up-moves can be of different kinds depending on the decoration. The essential point is that an external edge in the k-tree is not at risk of being dropped until its mass is one and the mass of its parent is zero. We now define these special cases first. Fix $1 \leq k < n$ and $\alpha \in (0, 1)$ throughout.

Definition 9. (Dropping a label) Let $\mathbf{t}^\bullet = (\mathbf{s}, (x_1, \ldots, x_k, (y_B)_{B \in \mathbf{s}, \#B \geq 2})) \in \mathbb{T}_{[k]}^{\bullet (n)}$. Suppose label i is such that $x_i = 1$ and $y_{\overleftarrow{\{i\}}} = 0$. Then, by *dropping label i from $\mathbf{t}^\bullet$* we mean performing the following sequentially.

(i) Drop leaf i from $\mathbf{s}$ and relabel leaves $i + 1, \ldots, k$ by $i, \ldots, k - 1$, respectively, to get a tree shape $\mathbf{s}^{(i)} \in \mathbb{T}_{[k-1]}$.

(ii) Drop x_i and $y_{\overleftarrow{\{i\}}}$ from the list of weights in $\mathbf{t}^\bullet$ and relabel the remaining weights by the edges of $\mathbf{s}^{(i)}$ to obtain an element of $\mathbb{T}_{[k-1]}^{\bullet (n-1)}$.

Definition 10 (Inserting a label)**.** Suppose $\mathbf{t}^\bullet = (\mathbf{s}, (x_1, \ldots, x_k, (y_B)_{B \in \mathbf{s}, \#B \geq 2})) \in \mathbb{T}_{[k-1]}^{\bullet (n-1)}$. By *inserting the label k in $\mathbf{t}^\bullet$* we mean performing the following actions on $\mathbf{t}^\bullet$ sequentially to obtain a $\mathbb{T}_{[k]}^{\bullet (n-1)}$-valued random variable.

(i) Choose an edge B with probability

$$
\begin{cases}
\dfrac{x_i - 1}{n - k}, & \text{if } B = \{i\} \text{ is external,} \\[2mm]
\dfrac{y_B}{n - k}, & \text{if } B \text{ is internal.}
\end{cases}
$$

(ii) If an external edge $\{i\}$ has been chosen, sample three integers (j_1^*, j_2^*, j_3^*) using the kernel $\mathrm{DM}_3^\alpha(x_i - 2)$. Update the tree shape $\mathbf{s}$ to $\mathbf{s} + (\{i\}, k)$, by inserting a new leaf labeled k on edge $\{i\}$. Split the corresponding x_i into three random integers $(x_i, x_k, y_{\{i,k\}})$ where

$$
x_i = j_1^* + 1, \quad x_k = j_2^* + 1, \quad y_{\{i,k\}} = j_3^*.
$$

(iii) If an internal edge has been chosen, sample three integers (j_1^*, j_2^*, j_3^*) using the kernel $\mathrm{DM}_3^{1-\alpha}(y_B - 1)$. Update the tree shape $\mathbf{s}$ to $\mathbf{s} + (B, k)$, by inserting a new leaf labeled k on edge B. Split the original y_B in three random integers $(y_{B \cup \{k\}}, y_B, x_k)$ where

$$
y_{B \cup \{k\}} = j_1^*, \quad y_B = j_2^*, \quad x_k = j_3^* + 1.
$$

Definition 11 (Up-move on decorated trees).
Let $\mathbf{t}^\bullet = (\mathbf{s}, (x_1, \ldots, x_k, (y_B)_{B \in \mathbf{s}, \, \#B \geq 2})) \in \mathbb{T}_{[k]}^{\bullet(n-1)}$. Choose an edge B with probability

$$\begin{cases} \dfrac{x_i - \alpha}{n - 1 - \alpha}, & \text{if } B = \{i\} \text{ is external,} \\[2ex] \dfrac{y_B + \alpha}{n - 1 - \alpha}, & \text{if } B \text{ is internal,} \end{cases}$$

and increase the mass of the chosen edge by one to obtain an element of $\mathbb{T}_{[k]}^{\bullet(n)}$.

For the uniform decorated chain we combine variations of dropping a label (without relabeling leaves), up-move and inserting a label (using the edge chosen for the up-move also for the label insertion), in a single resampling step.

Definition 12 (Resampling a label). Let $\mathbf{t}^\bullet = (\mathbf{s}, (x_1, \ldots, x_k, (y_B)_{B \in \mathbf{s}, \, \#B \geq 2})) \in \mathbb{T}_{[k]}^{\bullet(n)}$. Suppose label i is such that $x_i = 1$ and $y_{\overleftarrow{\{i\}}} = 0$. Then, by *resampling label i* from $\mathbf{t}^\bullet$ we mean the following list of actions to get a $\mathbb{T}_{[k]}^{\bullet(n)}$-valued random variable.

(i) Drop leaf i from $\mathbf{s}$ to get a tree shape $\mathbf{s}^{(i)} \in \mathbb{T}_{[k] \setminus \{i\}}$.
(ii) Drop x_i and $y_{\overleftarrow{\{i\}}}$ from the list of weights in $\mathbf{t}^\bullet$.
(iii) Choose an edge B of $\mathbf{s}^{(i)}$ with probability

$$\begin{cases} \dfrac{x_j - 1/2}{n - 3/2}, & \text{if } B = \{j\} \text{ is external,} \\[2ex] \dfrac{y_B + 1/2}{n - 3/2}, & \text{if } B \text{ is internal.} \end{cases}$$

(iv) If an external edge $B = \{j\}$ has been chosen, sample three integers (j_1^*, j_2^*, j_3^*) using the kernel $\mathrm{DM}_3^{1/2}(x_j - 1)$. Increase the current x_j by one and split into three new random integers $(x_j, x_i, y_{\{i,j\}})$ where

$$x_j = j_1^* + 1, \quad x_i = j_2^* + 1, \quad y_{\{i,j\}} = j_3^*.$$

Update the tree shape $\mathbf{s}$ to $\mathbf{s}^{(i)} + (\{j\}, i)$, inserting a new leaf i below leaf j.
(v) If an internal edge B has been chosen, sample three integers (j_1^*, j_2^*, j_3^*) using the kernel $\mathrm{DM}_3^{1/2}(y_B)$. Increase the original y_B by one and split into three new random integers $(y_{B \cup \{i\}}, y_B, x_i)$ where

$$y_{B \cup \{i\}} = j_1^*, \quad y_B = j_2^*, \quad x_i = j_3^* + 1.$$

Update the tree shape $\mathbf{s}$ to $\mathbf{s}^{(i)} + (B, i)$, inserting a new leaf i below vertex B.

We will now define a Markov chain on the space of decorated k-trees $\mathbb{T}_{[k]}^{\bullet(n)}$ of total edge mass n.

Definition 13 (Uniform decorated chain). Fix $n \geq 3$ and $k \in [n]$. We define a Markov chain $(T_k^{\bullet(n)}(j))_{j \geq 0}$ on the state space $\mathbb{T}_{[k]}^{\bullet(n)}$ with the following transition rules. Given $T_k^{\bullet(n)}(j) = \mathbf{t}^\bullet = (\mathbf{s}, (x_1, \ldots, x_k, (y_B)_{B \in \mathbf{s}, \, \#B \geq 2})) \in \mathbb{T}_{[k]}^{\bullet(n)}$, $j \geq 0$, randomly construct $T_k^{\bullet(n)}(j+1)$ by first making the following random selection:

(i) select the internal edge B for any $B \in \mathbf{s} : \#B \geq 2$ with probability y_B/n, or
(ii) select the external edge to leaf i for any $i \in [k]$ with probability x_i/n.

Then, perform the following random move in the case determined by this selection.

(A) If some internal B or an external i with $x_i \geq 2$ is chosen, reduce the chosen weight by one and perform an up-move as in Definition 11 with $\alpha = 1/2$.

(B) If some external i with $x_i = 1$ is chosen and $y_{\overleftarrow{\{i\}}} > 0$, sample a random m from the kernel $\delta\left(y_{\overleftarrow{\{i\}}} : \cdot\right)$ in (5). Set $x_i = m$ and reduce $y_{\overleftarrow{\{i\}}}$ by m. Perform an up-move as in Definition 11.

(C) If some external i with $x_i = 1$ is chosen and $y_{\overleftarrow{\{i\}}} = 0$, consider the k-tree shape $\mathbf{s}$ and compare i with the smallest leaf labels a and b in the first two subtrees on the ancestral path from leaf i to the root (with the convention $b = 0$ if the path has only one subtree). Let $\tilde{\imath} = \max\{i, a, b\}$. Swap i and $\tilde{\imath}$, if different. Resample $\tilde{\imath}$ as in Definition 12.

The above chain can be generalized for all $\alpha \in (0,1)$ by using the alpha chain given in Definition 5.

Definition 14 (Alpha decorated chain). Fix $\alpha \in (0,1)$, $n \geq 3$, $k \in [n]$. We define $(T_k^{\bullet(n)}(j))_{j \geq 0}$. Given $T_k^{\bullet(n)}(j) = (\mathbf{s}, (x_1, \ldots, x_k, (y_B)_{B \in \mathbf{s}, \#B \geq 2})) \in \mathbb{T}_{[k]}^{\bullet(n)}$, $j \geq 0$, randomly construct $T_k^{\bullet(n)}(j+1)$ by first making the following random selection:

(i) select the internal edge B for any $B \in \mathbf{s} : \#B \geq 2$ with probability y_B/n, or

(ii) select the external edge to leaf i for any $i \in [k]$ with probability x_i/n.

Then, perform the following random move in the case determined by this selection.

(A) If some internal B or an external i with $x_i \geq 2$ is chosen, reduce the chosen weight by one and perform an up-move as in Definition 11.

(B) If some external i with $x_i = 1$ is chosen and $y_{\overleftarrow{\{i\}}} > 0$, sample a random m from the kernel $\delta_\alpha\left(y_{\overleftarrow{\{i\}}} : \cdot\right)$ in (5). Set $x_i = m$ and reduce $y_{\overleftarrow{\{i\}}}$ by m. Perform an up-move as in Definition 11.

(C) If some external i with $x_i = 1$ is chosen and $y_{\overleftarrow{\{i\}}} = 0$, consider the k-tree shape $\mathbf{s}$ and compare i with the smallest leaf labels a and b in the first two subtrees on the ancestral path from leaf i to the root ($b = 0$ if the path has only one subtree). Let $\tilde{\imath} = \max\{i, a, b\}$. Swap i and $\tilde{\imath}$, if different. Drop $\tilde{\imath}$ from $\mathbf{t}^\bullet$ as in Definition 9 to get an element in $\mathbb{T}_{[k-1]}^{\bullet(n-1)}$. Insert label k as in Definition 10 and then perform an up-move as in Definition 11.

4.2. Representation of uniform and alpha decorated transition matrices.
Recall the discussion following (6).

Proposition 8. *Let P denote the Markov transition kernel for the uniform chain of Definition 1 and let $\Lambda(\mathbf{t}^\bullet, \cdot)$ denote the conditional law $q_{n,1/2}(\cdot \,|\, \rho_k^{\bullet(n)} = \mathbf{t}^\bullet)$. Then, the Markov transition kernel for the uniform decorated chain of Definition 13 is $\Lambda P \rho_k^{\bullet(n)}$. Similarly, let $P^{(\alpha)}$ denote the transition kernel of the alpha chain of Definition 5. Then the Markov transition kernel for the alpha decorated chain in Definition 14 is $\Lambda^{(\alpha)} P^{(\alpha)} \rho_k^{\bullet(n)}$, where $\Lambda^{(\alpha)}(\mathbf{t}^\bullet, \cdot) = q_{n,\alpha}(\cdot \,|\, \rho_k^{\bullet(n)} = \mathbf{t}^\bullet)$.*

To prove this proposition we first describe the decorated k-trees as a process when projected from an alpha growth process $(T_n, n \geq 1)$. Let $k \geq 2$. Condition on $T_k = \mathbf{s}$ and consider the conditional law of the decorated k-tree process

$$T_k^{\bullet(k+r)} := \rho_k^{\bullet(k+r)}(T_{k+r}), \ r \geq 0.$$

Clearly every $T_k^{\bullet(k+r)}$ has the same k-tree shape given by $\mathbf{s}$ and nondecreasing edge weights. Let $x_j(r)$ and $y_B(r)$, for $j \in [k]$ and $B \in \mathbf{s}$, $\#B \geq 2$, denote the masses assigned to edges of $\mathbf{s}$ in $T_k^{\bullet(k+r)}$.

Lemma 9. *The joint distribution of the multi-dimensional process*

$$\left(x_1(r) - 1, \ldots, x_k(r) - 1, \ (y_B(r))_{B \in \mathbf{s}: \#B \geq 2} \right), \ r \geq 0,$$

is given by a generalized Pólya urn model $H(r) = (H_B(r), B \in \mathbf{s}), r \geq 0$, with weights

$$\begin{cases} \alpha + H_B(\cdot), & \text{when } B \text{ is internal, and} \\ 1 - \alpha + H_B(\cdot), & \text{when } B \text{ is external.} \end{cases}$$

The initial condition is $H_B(0) = 0$ for all $B \in \mathbf{s}$.

In this Pólya urn, the edges $B \in \mathbf{s}$ are the colors in the urn, and the model follows the standard update rule, i.e. selecting colors with probabilities proportional to weights and setting $H_B(r + 1) = H_B(r) + 1$ if color B is selected at step $r + 1$. See [18, Proposition 18], [26, Proposition 14].

It follows from Lemma 9 that $(T_k^{\bullet(k+r)})_{r \geq 0}$ is a Markov chain with a limiting Dirichlet distribution for the weights. The following corollaries are easy consequences of Lemma 9.

Corollary 10. *The up-move in Definition 11 describes the conditional distribution of $T_k^{\bullet(n)}$, given $T_k^{\bullet(n-1)} = \mathbf{t}^\bullet$.*

Corollary 11. *Let $T_k \sim q_{k,\alpha}$. Given $T_k = \mathbf{s}$, use the edges as an indexing set where the external edges are labeled by $1, 2, \ldots, k$ corresponding to the leaf labels, and the internal edges are labeled (e.g. in lexicographical order) by $k+1, \ldots, 2k-1$. Now, generate a vector $(J_1, \ldots, J_{2k-1}) \sim \mathrm{DM}_{2k-1}^\alpha(n - k)$. Assign J_i to the edge labeled i for each $i \in [2k - 1]$. For $i \in [k]$ define $x_i = J_i + 1$, and for $B \in \mathbf{s}$, $\#B \geq 2$, set $y_B = J_l$, where $l \in [k + 1, 2k - 1]$ is the label of edge B. Then $(\mathbf{s}, (x_1, \ldots, x_k, (y_B)_{B \in \mathbf{s}: \#B \geq 2}))$ has the same distribution as $\rho_k^{\bullet(n)}$ under $q_{n,\alpha}$.*

Proof. By definition, T_k is distributed as $q_{k,\alpha}$. The $\mathrm{DM}_{2k-1}^\alpha$ distribution is due to the well-known exchangeability of the Pólya urn model. If we let $r \to \infty$ in Lemma 9, the proportion of counts of each color has a limit $\mathbf{p}$ that is distributed according to the Dirichlet distribution with a parameter vector v that is α in the first k coordinates and $(1 - \alpha)$ in the rest. By de Finetti's theorem, given a realization of this limiting proportion $\mathbf{p}$, the counts of balls is conditionally multinomial with probability vector $\mathbf{p}$. This, by definition, is $\mathrm{DM}_{2k-1}^\alpha(n - k)$. $\qquad\square$

We now vary k in Lemma 9. Fix $n \in \mathbb{N}$, and consider the stochastic process $\left(T_k^{\bullet(n)}, \ k = 1, 2, \ldots, n \right)$. It is not hard to see that this is also a Markov chain. To describe its transition probabilities, notice that the alpha growth process gives the following relation between $T_{k-1}^{\bullet(n)}$ and $T_k^{\bullet(n)}$. When we attach the kth leaf to T_{k-1}, it introduces two new edges on which the future leaves can get attached. If we imagine the count of leaves growing on the edges of T_{k-1} as an urn model, the kth leaf chooses one of the colors in T_{k-1} at random and splits it into three new colors. The transition kernel from $T_{k-1}^{\bullet(n)}$ to $T_k^{\bullet(n)}$ is given by the conditional probability of (i) which of the original colors got split and (ii) what the proportions of the new colors are after all n leaves have been attached.

Lemma 12. *Let $T_{n-1} \sim q_{n-1,\alpha}$, set $T_k^{\bullet(n-1)} = \rho_k^{\bullet(n-1)}(T_{n-1})$ and $T_{k-1}^{\bullet(n-1)} = \rho_{k-1}^{\bullet(n-1)}(T_{n-1})$. Then inserting a label, as in Definition 10, describes the conditional distribution of $T_k^{\bullet(n-1)}$ given $T_{k-1}^{\bullet(n-1)} = (\mathbf{s}, (x_1, \ldots, x_k, (y_B)_{B \in \mathbf{s}, \#B \geq 2})) \in \mathbb{T}_{[k-1]}^{\bullet(n-1)}$.*

Proof. Consider the Pólya urn set-up from Lemma 9 and Corollary 11. Let $r = n - k$. Fix $B^* \in \mathbf{s}$. Given $H_B(r) = j_B$, $B \in \mathbf{s}$, the conditional probability that the kth leaf got attached to B^* is j_{B^*}/r by exchangeability of the urn scheme. Once the kth leaf gets attached, the three new colors evolve as a Pólya urn model themselves for time $r - 1$. Let B_1^*, B_2^*, B_3^* be the three new colors. Then $H_{B_1^*}(0) = H_{B_2^*}(0) = H_{B_3^*}(0) = 0$. If j_1^*, j_2^*, j_3^* form the state of the three colors at time $r - 1$ (when $n - 1$ leaves have been attached), then, we must have $j_1^* + j_2^* + j_3^* = j_{B^*} - 1$.

Suppose B^* is external. Then the introduction of the new leaf creates two external edges (say B_1^* and B_2^*) and one internal (B_3^*). Thus the distribution of $\left(H_{B_1^*}(r - 1), H_{B_2^*}(r - 1), H_{B_3^*}(r - 1) \right)$, given $H_{B^*}(r) = j_{B^*} = x_{B^*} - 1$, is DM with parameters $x_{B^*} - 2$ and $(1 - \alpha, 1 - \alpha, \alpha)$. On the other hand, if B^* is internal, the introduction of the new leaf creates two internal edges and one external. This gives us a DM distribution with parameters $y_{B^*} - 1$ and $(\alpha, \alpha, 1 - \alpha)$, where the first two coordinates stand for the internal edges and the last one for the external edge.

Hence, the edge selection and splitting probabilities are as in Definition 10. $\square$

Lemma 13. *Fix $i \in [k]$. Let $T_n \sim q_{n,1/2}$, $T_{n-1} = T_n \cap [n-1]$, $T_k^{\bullet(n)} = \rho_k^{\bullet(n)}(T_n)$ and $T_{[k] \setminus \{i\}}^{\bullet(n-1)} := \rho_{[k] \setminus \{i\}}^{\bullet(n-1)}(T_{n-1})$. Then performing steps (iii), (iv) and (v) of Definition 12, describes the conditional distribution of $T_k^{\bullet(n)}$ given $T_{[k] \setminus \{i\}}^{\bullet(n-1)} = \mathbf{t}^\bullet \in \mathbb{T}_{[k] \setminus \{i\}}^{\bullet(n-1)}$.*

Proof. By exchangeability of leaf labels in $q_{n,1/2}$, it is enough to consider the case when $i = k$. Thus, we compute the conditional distribution of $T_k^{\bullet(n)}$, given $T_{k-1}^{\bullet(n-1)}$. The remainder of this proof is very similar to the previous proof, and hence we give an outline of the argument and leave the details for the reader.

Let $\mathbf{s} \in \mathbb{T}_{k-1}$. We condition on $\{T_{k-1} = \mathbf{s}\}$. Consider an up-move that adds the mass of a new leaf n to $T_{k-1}^{\bullet(n-1)}$ at a randomly chosen edge $B \in \mathbf{s}$ according to weights as in Definition 11 for $\alpha = 1/2$. By Corollary 10, this is a step by the conditional distribution of $T_k^{\bullet(n)}$ given $T_{k-1}^{\bullet(n-1)}$. By exchangeability, again, this leaf n could be relabeled leaf k, and the total leaf mass of T_n can now be projected onto the tree spanned by leaves labeled by $[k]$. This will split the weight of edge B according to a DM distribution as in the previous proof. Hence, the edge selection, mass insertion and splitting probabilities are as in Definition 12 (iii)-(v). $\square$

Proof of Proposition 8. We only prove the case of the alpha chain. The case of the uniform chain is similar and is left to the reader.

Fix $\alpha \in (0, 1)$ and $\mathbf{t}^\bullet = (\mathbf{s}, (x_1, x_2, \ldots, x_k, (y_B)_{B \in \mathbf{s}, \#B \geq 2})) \in \mathbb{T}_{[k]}^{\bullet(n)}$. Proposition 8 claims the equality of two distributions on $\mathbb{T}_{[k]}^{\bullet(n)}$, one obtained by a transition of the alpha decorated chain as defined in Definition 14, the other by transitions via the alpha chain. Specifically, let us consider random variables that exhibit the latter transition, with the alpha chain transition further split into the down-move and the up-move, as follows:

- an initial n-tree T_n, which is distributed according to $\Lambda(\mathbf{t}^\bullet, \cdot)$;

- a leaf label I, uniformly distributed over $[n]$ and independent of T_n, that is initially selected for removal in the down-move; then $T'_{n-1} \in \mathbb{T}_{[n-1]}$ is obtained from T_n and I by dropping $\widetilde{I}$ and relabeling as in Definition 5;
- an edge S of T'_{n-1} distributed as in the up-move of Definition 5; then the tree $T'_n = T'_{n-1} + (S, n) \in \mathbb{T}_{[n]}$ realises the up-move and we want to study the distribution of $\rho_k^{\bullet(n)}(T'_n)$.

To compare this distribution with the alpha decorated chain, we further define on the same probability space the common tree shape $T'_k = T'_{n-1} \cap [k]$ of $\rho_k^{\bullet(n-1)}(T'_{n-1})$ and $\rho_k^{\bullet(n)}(T'_n)$, and the edges $\beta \in \mathbf{s}$ and $\beta' \in T'_k$ to which I and S project, respectively. Let us introduce β and β' more formally.

Let π be the projection map (see above Definition 7) that takes leaves of T_n to edges in $\mathbf{s}$. Let $\beta = \pi(I)$ denote the edge of $\mathbf{s}$ on which I projects. The conditional probability $\mathbf{P}(\beta = B \mid T_n)$, for any edge $B \in \mathbf{s}$, is clearly $\mu(B)/n$. Note that this probability is independent of T_n, thus β and T_n are independent and $\mathbf{P}(\beta = B) = \mu(B)/n$ has the same distribution as the edge selected in the transition of the alpha decorated chain. It therefore suffices to identify the conditional distribution of $\rho_k^{\bullet(n)}(T'_n)$ given $\beta = B$ with the distribution of the random move specified in (A), (B) or (C) of Definition 14, as the case may be for each $B \in \mathbf{s}$.

Let π' be the projection from T'_{n-1} to $\rho_k^{\bullet(n-1)}(T'_{n-1})$. By construction of $T'_n = T'_{n-1} + (S, n)$, there is a unique edge $\beta' \in T'_k$ whose weight changes when passing from $\rho_k^{\bullet(n-1)}(T'_{n-1})$ to $\rho_k^{\bullet(n)}(T'_n)$, and this weight increases by one. Then, for any edge $B' \in T'_k$, the conditional probability $\mathbf{P}(\beta' = B' \mid T_n, I)$ is proportional to $\mu'(B') - \alpha$ for external edges B', or $\mu'(B') + \alpha$ for internal edges, where $\mu'(B')$ is the weight of the edge B' in $\rho_k^{\bullet(n-1)}(T'_{n-1})$. In particular, it is a function only of $\rho_k^{\bullet(n-1)}(T'_{n-1})$, as is required as the last step in all three cases (A), (B) and (C) of Definition 14.

Case (A). Consider an edge $B \in \mathbf{s}$ such that either (i) B is internal with $\mu(B) > 0$ or (ii) B is external and $\mu(B) > 1$. Let us study the effect of the alpha chain down-move on the decorated tree $\rho_k^{\bullet(n-1)}(T'_{n-1})$ when conditioning on $\{\beta = B\}$.

(i) First suppose B is internal. If $I = i$, obviously $i > k$. Every leaf label j such that $\pi(j) = B$ must also be larger than k. In particular, $\tilde{\imath} \geq i > k$. Hence, once we swap and drop $\tilde{\imath}$, we reduce y_B by one without changing any other weights on $\mathbf{t}^\bullet$.

(ii) Now suppose B is an external edge $\{i\}$ such that $x_i > 1$. Consider the set of leaves j in T_n such that $\pi(j) = \{i\}$. Since $x_i > 1$, there are at least two of them, including i itself. All such j, $j \neq i$, must be at least $k + 1$. Among them, during the n-tree down-move, either some $j \neq i$ is removed, which is similar to the paragraph above, or the leaf i itself is chosen for removal. For the latter, let $\tilde{\imath}$ be the leaf that gets swapped with i and is dropped in the down-move. Consider the first two subtrees on the ancestral path from i to the root in T_n. Since $x_i > 1$, the leaf labels in at least one of the subtrees is contained in $\pi^{-1}(\{i\})$, and therefore $\tilde{\imath} > k$. In fact, $\pi(\tilde{\imath})$ must be $\{i\}$. To see this, note that, either both subtrees are contained in $\pi^{-1}(\{i\})$, in which case the conclusion is obvious, or, the first subtree is in $\pi^{-1}(\{i\})$, while the second subtree is not. In the latter case, the second subtree must have a leaf $\leq k$ (consider a leaf label in the sibling of $\{i\}$ in $\mathbf{s}$) while in the first subtree all leaves must be labeled $\geq k + 1$. Thus $\tilde{\imath}$, being the maximum of the two minimum labels in two subtrees, must belong to the first subtree and hence

$\tilde{\imath} \geq k + 1$ and $\pi(\tilde{\imath}) = \{i\}$. Therefore, after swapping and removing $\tilde{\imath}$ and relabeling higher labeled leaves simply reduces x_i by one and has no other effects on $\mathbf{t}^\bullet$.

Thus, in both (i) and (ii), $\rho_k^{\bullet(n-1)}(T'_{n-1})$ has the same tree shape as $\mathbf{t}^\bullet$ and all the same weights except on edge B, whose weight is reduced by one, as needed for the down-move in Definition 14. The up-move probabilities have already been discussed.

Case (B). Suppose $B = \{i\} \in \mathbf{s}$ for some $i \in [k]$ such that $x_i = 1$ and $y_{\overleftarrow{\{i\}}} > 0$. We condition on $\{\beta = B\}$. Let $\{i\}'$ and $\overleftarrow{\{i\}}$ be the labels of the sibling edge and the parent edge of $\{i\}$ in $\mathbf{s}$, the tree shape in $\mathbf{t}^\bullet$. Consider the spinal decomposition (recall Section 2.1) from leaf $\{i\}$ to the root in T_n. Let

$$U := \bigcup_{B \in \mathbf{s}:\, B \subseteq \{i\}'} \pi^{-1}(B) \quad \text{and} \quad V := \pi^{-1}\left(\overleftarrow{\{i\}}\right).$$

Since $x_i = 1$, U is the set of labels in the first subtree on the ancestral line from $\{i\}$ to the root, while the labels in the second subtree comprise a non-empty subset $W \subseteq V$. Since U contains some leaf label at most k, whereas every leaf in V (and thus W) has label at least $k + 1$, we conclude that $\tilde{\imath}$ comes from W.

In fact, the above U and W are edges in T_n: the sibling and uncle of $\{i\}$. Let U' and W' denote the resulting label sets after replacing $\tilde{\imath}$ with i in W and reducing all labels greater than $\tilde{\imath}$ by one. The step of dropping $\tilde{\imath}$ in the down-move makes U' and W' siblings in T'_{n-1}. Set $M_i := \#W \leq \#V = y_{\overleftarrow{\{i\}}}$ in $\mathbf{t}^\bullet$. All weights and labels are the same in $\rho_k^{\bullet(n-1)}(T'_{n-1})$ as in $\mathbf{t}^\bullet$, with two exceptions: weight $x_i = 1$ is replaced by M_i and weight $y_{\overleftarrow{\{i\}}}$ is reduced by M_i. Therefore, the law of $\rho_k^{\bullet(n-1)}(T'_{n-1})$ under $\mathbf{P}(\cdot \,|\, \beta = B)$ is specified by the conditional law of M_i. Since T_n has law $q_{n,\alpha}(\cdot \,|\, \rho_k^{\bullet(n)} = \mathbf{t}^\bullet)$ and is independent of β, it follows from Lemmas 6 and 7 that the law of M_i is given by the stochastic kernel $\delta_\alpha\left(y_{\overleftarrow{\{i\}}} : \cdot\right)$ of (4).

Again, the up-move probabilities have been discussed, which completes this case.

Case (C). Suppose $B = \{i\}$ with $x_i = 1$ and $y_{\overleftarrow{\{i\}}} = 0$ in $\mathbf{t}^\bullet$. We condition on $\{\beta = B\}$. Since $x_i = 1$ and $y_{\overleftarrow{\{i\}}} = 0$, each of the first two subtrees on the ancestral line of $\{i\}$ in T_n must have a leaf label at most k. Hence, the n-tree down-move drops $\tilde{\imath} \leq k$ and shifts labels $(\tilde{\imath} + 1, \ldots, n)$ down to by one. This gives us T'_{n-1}. However, unlike the previous cases, the tree shape of $\rho_k^{\bullet(n-1)}(T'_{n-1})$ is not deterministic under $\mathbf{P}(\cdot \,|\, \beta = B)$. This is because leaf $k + 1$ was relabeled as k, so the weight on $\{k\}$ in $\rho_k^{\bullet(n-1)}(T'_{n-1})$ corresponds to part of some other, larger weight in $\mathbf{t}^\bullet$, according to where $k + 1$ was placed when generating T_n.

Now note that relabeling $(k + 1, \ldots, n)$ does not affect the weights on the decorated trees. Thus $\mathbf{v}^\bullet = \rho_{k-1}^{\bullet(n-1)}\left(T'_{n-1}\right)$ is deterministic under $\mathbf{P}(\cdot \,|\, \beta = B)$. To be consistent with Definition 14, it is enough to argue (due to Lemma 12) that the distribution of T'_{n-1} under $\mathbf{P}(\cdot \,|\, \beta = B)$ is $q_{n-1,\alpha}(\cdot \,|\, \rho_{k-1}^{\bullet(n-1)} = \mathbf{v}^\bullet)$. However, this is very similar to the proof of Lemma 4 since T_n is distributed as $q_{n,\alpha}(\cdot \,|\, \rho_k^{\bullet(n)} = \mathbf{t}^\bullet)$. In fact, if one repeats the proof of Lemma 4 for decorated k-trees, instead of n-trees, we reach our conclusion. The up-move is as usual. This completes the proof. $\square$

Before we end this section, let us extend the result in Corollary 5.

Corollary 14. *Consider an alpha decorated chain (Definition* 14*) running in sta-tionarity. Let $\widetilde{I}$ be the first leaf label in $[k]$ that gets dropped in case (C). Then, the distribution of $\widetilde{I}$ is given by*

(7)
$$
\mathbf{P}\left(\widetilde{I} = \tilde{\imath}\right) = \begin{cases} \dfrac{1}{k(k-1-\alpha)}, & \text{if } \tilde{\imath} = 2. \\ \dfrac{2\tilde{\imath} - 2 - \alpha}{k(k-1-\alpha)}, & \text{if } 3 \leq \tilde{\imath} \leq k. \end{cases}
$$

For $\alpha = 1/2$ this is also the distribution of the first resampled leaf label in the uniform decorated chain (Definition 13*).*

Proof. Let $(T(r))_{r \geq 0}$ be an alpha chain on $\mathbb{T}_{[n]}$ running in stationarity. As before, let $(T_k^{\bullet (n)}(r))_{r \geq 0}$ denote the projected decorated chain. Let τ denote the first time r such that a label in $[k]$ is dropped. Note that τ is a stopping time with respect to the natural filtration of $(T_k^{\bullet (n)}(r))_{r \geq 0}$.

Let $\mathbf{s}$ denote the k-tree shape of $T_k^{\bullet (n)}(0)$. By our assumption of stationarity, $\mathbf{s}$ is distributed according to $q_{k,\alpha}$ over $\mathbb{T}_{[k]}$. By definition of τ, each of the decorated trees $T_k^{\bullet (n)}(r)$, $0 \leq r \leq \tau - 1$, has the same trees shape $\mathbf{s}$. At time τ, let $I \in [k]$ denote the random leaf label that is selected for removal. That is, $I = i$ is the label satisfying case (C) in Definition 13. The statement of the corollary will follow by the exact same proof as for Corollary 5 if we show that I is uniformly distributed over $[k]$, independent of $\mathbf{s}$.

However, this follows from symmetry of the conditional law $T_k^{\bullet (n)}(0)$, given the tree shape $\mathbf{s}$, as described in Corollary 11. It shows the following exchangeability property: suppose we relabel the leaves of $\mathbf{s}$ by a permutation σ of $[k]$. For simplicity assume that σ transposes leaf labels 1 and 2. Switch the labels of the leaves 1 and 2 in $T_k^{\bullet (n)}(0)$ to get a new decorated k-tree. The tree-shape of this new tree is identical to $\mathbf{s}$, except for the switching of leaf labels 1 and 2. However, the conditional law of the both the decorated k-trees, given their tree shapes, is identical and given by Corollary 11. This follows from the exchangeability of the first k coordinates of $\mathrm{DM}_{2k-1}^\alpha(n-k)$, given the last $(k-1)$ coordinates. That is, if we generate a vector from $\mathrm{DM}_{2k-1}^\alpha(n-k)$, switch coordinates 1 and 2, while leaving the rest of the coordinates unchanged, the distribution remains unchanged.

Now, this initial coupling of the two decorated trees extends to a natural coupling of the alpha decorated chains where we perform the same moves on the unlabeled trees but keep the swapped leaf labels 1 and 2. Thus, we have a measure-preserving bijection between the sample paths in the events $\{I = 1\}$ and $\{I = 2\}$, given $\mathbf{s}$. Hence the probability of the two events, given $\mathbf{s}$, must be the same. Replacing 2 by any other label $\tilde{\imath} \in [k]$ shows the conditional uniform distribution of I. Hence, I and $\mathbf{s}$ are independent and the proof is complete. $\square$

5. Consistency in stationarity for the uniform down-up chains

For the rest of the paper we consider the uniform chain of Definition 1 and the uniform decorated chain of Definition 13. Our main result is Theorem 2 of the Introduction. With the preparations done so far, we restate it here, as follows.

Theorem 15 (Consistency in stationarity). *Fix $n \geq 2$. Let $(T(j))_{j\geq 0}$ be a uniform chain running in stationarity. For $1 \leq k \leq n-1$, let*

$$T_k^{\bullet(n)}(j) = \rho_k^{\bullet(n)}(T(j)), \quad j \geq 0,$$

denote the corresponding sequence of decorated k-trees obtained by projections. Then each $(T_k^{\bullet(n)}(j))_{j\geq 0}$, is a uniform decorated chain as described in Definition 13, also running in stationarity.

Recall the concept of collapsed trees from Definition 7. Define the process

$$T_k^{\star(n)}(j) := \rho_k^{\star(n)}(T(j)), \quad j \geq 0.$$

Thus, on the same filtered probability space that supports $(T(j))_{j\geq 0}$ we have three stochastic processes $T(\cdot) := (T(j))_{j\geq 0}$, $T_k^{\star(n)}(\cdot) := (T_k^{\star(n)}(j))_{j\geq 0}$ and $T_k^{\bullet(n)}(\cdot) := (T_k^{\bullet(n)}(j))_{j\geq 0}$. We will show that, when the first is running in stationarity, all three processes are Markov chains by themselves, also running in stationarity.

Recall the set-up described at the beginning of Section 4.1. Suppose we have a Markov chain $\mathcal{X}(\cdot) := (\mathcal{X}(j))_{j\geq 0}$, with a finite state space $\mathbb{X}$, a transition kernel $K(x, \cdot)$, and an initial probability distribution μ. Let $\mathbb{Y}$ be another finite state space and let $g \colon \mathbb{X} \to \mathbb{Y}$ be a given surjective function. We ask the question: when is $\mathcal{Y}(j) := g(\mathcal{X}(j))$, $j \geq 0$, a Markov chain?

There have been a number sufficient conditions proposed for this problem. We will require two of them: (i) the Kemeny–Snell criterion and (ii) the intertwining criterion due to Rogers and Pitman. We describe both these conditions below. The proof of Theorem 15 will consist of two key steps. The first step is to show that the pair $\mathcal{X}(\cdot) = T_k^{\star(n)}(\cdot)$, $\mathcal{Y}(\cdot) := T_k^{\bullet(n)}(\cdot)$ satisfies the Kemeny–Snell criterion, where $T_k^{\star(n)}(\cdot)$ is a Markov chain with a transition kernel induced from $T(\cdot)$ as described in (6). The next step is to show that the pair $\mathcal{X}(\cdot) := T(\cdot)$, $\mathcal{Y}(\cdot) := T_k^{\star(n)}(\cdot)$ satisfies the Rogers–Pitman intertwining criterion. Combining the two gives us Theorem 15. We are forced to take this two-step approach since it can be seen from explicit calculation and examples that neither intertwining nor Kemeny–Snell hold between the chains $\mathcal{X}(\cdot) := T(\cdot)$ and $\mathcal{Y}(\cdot) := T_k^{\bullet(n)}(\cdot)$.

Let us now describe the Kemeny–Snell [19] criterion. Consider the Markov transition kernel from $\mathbb{X}$ to $\mathbb{Y}$ given by

$$Q_g(x, \{y\}) := K\left(x, g^{-1}(\{y\})\right), \quad y \in \mathbb{Y}.$$

Suppose that the kernel $Q_g(x, \cdot)$ is a function of $g(x)$, i.e., for any $y \in \mathbb{Y}$,

$$Q_g(x_1, y) = Q_g(x_2, y), \quad \text{if } g(x_1) = g(x_2) \quad \text{(Kemeny–Snell criterion)}.$$

For any pair of $y_1, y_2 \in \mathbb{Y}$, denote the common value of $Q_g(x_1, y_2)$, for any $x_1 \in g^{-1}(\{y_1\})$, by $Q(y_1, y_2)$. The following result is in [19, Theorem 6.3.2].

Lemma 16. *If the Kemeny–Snell criterion holds, then $\mathcal{Y}(\cdot)$ is a Markov chain with transition kernel $Q(\cdot, \cdot)$ and an initial distribution given by the push-forward of the measure μ by the function g. In particular, $\mathcal{Y}(\cdot)$ is running in stationarity if $\mathcal{X}(\cdot)$ is running in stationarity.*

Consider again the set-up at the beginning of Section 4.1. Take the Markov chain $\mathcal{X}$ to be the down-up chain $T(\cdot)$ and the function $g = \rho_k^{\star(n)}$. Then, by (6), there is a natural transition probability kernel on the set of collapsed trees $\mathbb{T}_{[k]}^{\star(n)}$.

$$T_k^{\star(n)}(i) \xrightarrow{\quad Q \quad} T_k^{\star(n)}(i+1)$$
$$\downarrow \Lambda \qquad\qquad\qquad \downarrow \Lambda$$
$$T(i) \xrightarrow{\quad P \quad} T(i+1)$$

FIGURE 6. Commutative diagram of intertwining of Markov chains on collapsed trees with that on n-trees.

We now take $\mathbb{X} = \mathbb{T}_{[k]}^{\star(n)}$ and $K(\cdot, \cdot)$ to be this induced probability kernel. Redefine g to denote the projection map that takes a collapsed tree $\mathbf{t}^\star$ and replaces each $\pi^{-1}(B)$ by the mass $\mu(B) = \#\pi^{-1}(B)$ to obtain the decorated tree with masses $\mathbf{t}^\bullet$.

Proposition 17. *If we take $\mathcal{X}(j) = T_k^{\star(n)}(j)$, $j \geq 0$, with transition kernel K and $\mathcal{Y}(j) = T_k^{\bullet(n)}(j)$, then the Kemeny–Snell criterion holds.*

Proof. Suppose we are given $\mathbf{t}^\bullet(0), \mathbf{t}^\bullet(1) \in \mathbb{T}_k^{\bullet(n)}$ and $T_k^{\star(n)}(0) = \mathbf{t}^\star(0)$ for some $\mathbf{t}^\star(0) \in g^{-1}(\{\mathbf{t}^\bullet(0)\})$ aiming to study $K(\mathbf{t}^\star(0), g^{-1}(\{\mathbf{t}^\bullet(1)\}))$. The induced kernel K on $\mathbb{T}_k^{\star(n)}$ can be described in the following way. Generate $T_n(0)$ from the conditional distribution $q_{n,1/2}(\cdot \mid \rho_k^{\star(n)} = \mathbf{t}^\star(0))$ specified in Lemma 6. Let $T_n(1)$ be one step in the uniform chain, starting at $T_n(0)$, and let $T_k^{\star(n)}(1) = \rho_k^{\star(n)}(T_n(1))$.

Consider the unranked internal structures defined above Definition 8. Specifically, let $\mathbf{s}(0)$ and $T_k(1)$ denote the k-tree shapes at times zero and one, respectively. Then the internal structures $(V_B(0), B \in \mathbf{s}(0))$ corresponding to $T_n(0)$ and $(V_B(1), B \in T_k(1))$ corresponding to $T_n(1)$ are closely related. Indeed, we see by conditioning as in the different cases in the proof of Proposition 8, that $\mathbf{P}(T_k^{\star(n)}(1) \in g^{-1}(\{\mathbf{t}^\bullet(1)\}) \mid T_n(0))$ depends only on the ranked internal structures $(\text{int}\,(V_B(0)),\ B \in \mathbf{s}(0))$ and not on the labels in $V_B(0)$, $B \in \mathbf{s}(0)$. Specifically, only the splitting via the regenerative composition in Case (B) and via resampling in Case (C) depends on $T_n(0)$. By Lemma 6, the internal structures are jointly independent of the labels in $V_B(0)$, $B \in \mathbf{s}(0)$. Hence, $K(\mathbf{t}^\star(0), g^{-1}(\{\mathbf{t}^\bullet(1)\})) = \mathbf{P}(T_k^{\star(n)}(1) \in g^{-1}(\{\mathbf{t}^\bullet(1)\}))$ does not depend on the labels in $V_B(0)$, $B \in \mathbf{s}(0)$, i.e. does not depend on $\mathbf{t}^\star(0) \in g^{-1}(\{\mathbf{t}^\bullet(0)\})$. This independence property is the Kemeny–Snell criterion. $\square$

Note that the Kemeny–Snell criterion fails between $\mathcal{X}(\cdot) = T(\cdot)$ and $\mathcal{Y}(\cdot) = T_k^{\bullet(n)}(\cdot)$ for instance because, for each edge $B \in \mathbf{s}(0)$, the internal structure $\text{int}(V_B)$ determines the size of the first spinal subtree, and hence influences the transition probabilities $K(\mathbf{t}(0), g^{-1}(\{\mathbf{t}^\bullet(1)\}))$.

Now we recall the concept of intertwining from Rogers and Pitman [29]. However, since we work in discrete time, unlike the continuous time setting in that paper, some of their statements get simplified. Once more consider the set-up described at the beginning of Section 4.1. The Markov chain $\mathcal{X}$ has transition kernel P and a stationary distribution q. Let $\mathcal{Y}(j) = g(\mathcal{X}(j))$. For $y \in \mathbb{Y}$, set $\Lambda(y, \cdot) = q(\cdot \mid g = y)$. Let Q be a transition kernel on $\mathbb{Y}$ given by $\Lambda P g$, which corresponds to the operator (6). The following is a simplification of [29, Theorem 2] adapted to our case.

Lemma 18. *Given the Markov chain* $(\mathcal{X}(j))_{j\geq 0}$*, suppose*
(a) the distribution of $\mathcal{X}(0)$*, given* $\mathcal{Y}(0) = g(\mathcal{X}(0)) = y$*, is* $\Lambda(y, \cdot)$*.*
(b) And, the triplet (Λ, P, Q) *satisfies the equality*

(8) $$\Lambda P = Q\Lambda, \qquad \text{(Intertwining criterion)}$$

as an equality of two stochastic kernels.
Then $\mathcal{Y}(\cdot)$ *is a Markov chain with transition kernel* Q*. Moreover, the conditional*
distribution of any $\mathcal{X}(j + 1)$*, given* $\mathcal{Y}(0), \ldots, \mathcal{Y}(j)$*, is* $\Lambda(\mathcal{Y}(j), \cdot)$

Proof. The function Φ in [29] is our function g. Hence condition (a) in [29, Theorem 2] that $\Lambda g = I$, the identity kernel, obviously follows from our definition of Λ as a conditional distribution, given g.

To verify condition (b) in [29, Theorem 2], note that our time is discrete. Hence P^j, for $j \geq 0$, is the transition semigroup (the semigroup is denoted by P_t, $t \geq 0$ in [29]). Let $Q_j = \Lambda P^j g$, for $j \geq 0$. We claim that $Q_j = Q^j$, where the latter is Q raised to the power j. This is true for $j = 0$ by the previous paragraph, and for $j = 1$ by (8). Now, take $j = 2$. From usual matrix multiplication, it follows that

$$Q_2 = \Lambda P^2 g = (\Lambda P) P g = (Q\Lambda) P g = Q(\Lambda P) g = Q^2 \Lambda g = Q^2.$$

An induction over j establishes our claim. Condition (b) in [29] requires $\Lambda P^j = Q_j \Lambda$. As $Q_j = Q^j$, we need to show $\Lambda P^j = Q^j \Lambda$. But this is obvious from (8). $\square$

Proposition 19. *The Intertwining criterion* (8) *holds for the pair* $\mathcal{Y}(\cdot) = T_k^{\star(n)}(\cdot)$ *and* $\mathcal{X}(\cdot) = T(\cdot)$ *with the kernel* Λ *given by the conditional distribution under* $q_{n,1/2}$*.*

Intuitively the intertwining criterion refers to the commutation of the diagram in Figure 6. Let P and Q be the Markov transition kernels for the uniform chain on n-trees and for collapsed trees, respectively. Let $\Lambda(\mathbf{t}^\star, \mathbf{t}) = q_{n,1/2}(\mathbf{t} \,|\, \rho_k^{\star(n)} = \mathbf{t}^\star)$. The function g is $\rho_k^{\star(n)}$. Suppose we are given the realization of $T_k^{\star(n)}(j)$, for some $j \geq 0$. Intertwining means that there are two equivalent ways to sample $T(j + 1)$.

(i) Generate $T(j)$ from $\Lambda\left(T_k^{\star(n)}(j), \cdot\right)$ and run one step of the uniform n-tree chain to get $T(j + 1)$.

(ii) Run one step of the Markov chain with transition Q, starting from $T_k^{\star(n)}(j)$, to get $T_k^{\star(n)}(j + 1)$. Then, generate $T(j + 1)$ from $\Lambda\left(T_k^{\star(n)}(j + 1), \cdot\right)$.

Before arguing Proposition 19, let us complete the proof of Theorem 15.

Proof of Theorem 15. The proof follows from the commutativity of the diagram in Figure 7. In fact, one can make a stronger statement. Suppose a realization of $T_k^{\bullet(n)}(0)$ is given. Generate $T_k^{\star(n)}(0)$, given $T_k^{\bullet(n)}(0)$, from the conditional distribution under $q_{n,1/2}$. This amounts to throwing, uniformly at random, balls labeled $k + 1, \ldots, n$ to boxes labeled $[2k - 1]$, one box for each edge of the tree shape of $T_k^{\star(n)}(0)$, such that the number of balls in each box is consistent with the corresponding edge mass in $T_k^{\bullet(n)}(0)$. Also generate $T(0)$ according to the kernel $\Lambda\left(T_k^{\star(n)}(0), \cdot\right)$. Now run the uniform n-tree chain starting from $T(0)$ to get a sequence $(T(j))_{j\geq 0}$. Set $T_k^{\star(n)}(j) := \rho_k^{\star(n)}(T(j))$ and $T_k^{\bullet(n)}(j) := \rho_k^{\bullet(n)}(T(j))$, $j \geq 0$. By Lemma 18 and Proposition 19, $T_k^{\star(n)}(\cdot)$ is a Markov chain with transition kernel Q and initial state $T_k^{\star(n)}(0)$. Hence, by Lemma 17, $T_k^{\bullet(n)}(\cdot)$ is a Markov chain on

$$T_k^{\bullet(n)}(i) \xrightarrow{\quad R \quad} T_k^{\bullet(n)}(i+1)$$

$$T_k^{\star(n)}(i) \xrightarrow{\quad Q \quad} T_k^{\star(n)}(i+1)$$

$$T(i) \xrightarrow{\quad P \quad} T(i+1)$$

FIGURE 7. Combining Kemeny–Snell with intertwining. R is the transition kernel for the down-up chain on decorated k-trees. The up-arrows represent deterministic projection maps.

decorated k-trees. When $T_k^{\bullet(n)}(0)$ is random and distributed according to the stationary distribution (which is the push-forward of the uniform distribution $q_{n,1/2}$ by the map $\rho_k^{\bullet(n)}$), we get the desired result. $\qquad \square$

To prove Proposition 19 we need the following lemma.

Lemma 20. *Let $T_m \sim q_{m,1/2}$ for some $m \geq 2$. Then, given $T_2^{\star(m)} := \rho_2^{\star(m)}(T_m) = (T_m \cap [2], (L_B)_{B \in T_m \cap [2]})$, the internal structures $\mathrm{int}(V_{\{1\}})$, $\mathrm{int}(V_{\{2\}})$ and $\mathrm{int}(V_{\{1,2\}})$ of $T_2^{\star(m)}$ are independent and uniformly distributed on the set of binary trees with numbers of leaves given by $\#L_{\{1\}}, \#L_{\{2\}}$ and $\#L_{\{1,2\}} + 1$, respectively.*

Proof. The proof is a direct application of Lemma 6. $\qquad \square$

Proof of Proposition 19. By Remark (ii) in [29, page 575], the Intertwining criterion (8) is equivalent to showing that the conditional distribution of $\mathcal{X}(1)$, given $\mathcal{Y}(0) = y_0$ and $\mathcal{Y}(1) = y_1$, is $\Lambda(y_1, \cdot)$. We verify this reformulated condition in our case.

Fix $\mathbf{t}^\star = \left(\mathbf{s}, (L_B)_{B \in \mathbf{s}}\right) \in \mathbb{T}_{[k]}^{\star(n)}$ throughout this proof. Consider the following uniform down- and up-move random variables, as in the proof of Proposition 8:

- an initial n-tree T_n, which is distributed according to $\Lambda(\mathbf{t}^\star, \cdot)$;
- a leaf label I, uniformly distributed over $[n]$ and independent of T_n; then the swapped label $\widetilde{I}$ and the tree T'_{n-1} after the down-move (leaving the other labels unchanged as in Definition 1) are determined, as is the random edge β of $\mathbf{s}$ to which I projects;
- an edge S of T'_{n-1} distributed as in the up-move of Definition 1; this determines $T'_n = T'_{n-1} + (S, \widetilde{I})$ after the up-move, as well as the edge β' of $T'_n \cap ([k] \setminus \{\widetilde{I}\})$ to which S projects (either to insert $\widetilde{I}$ into the collapsed tree after the down-move, if $\widetilde{I} > k$, or to split when $\widetilde{I}$ is resampled, if $\widetilde{I} \leq k$).

We claim that the conditional distribution of T'_n given $\rho_k^{\star(n)}(T'_n)$ is $\Lambda\left(\rho_k^{\star(n)}(T'_n), \cdot\right)$.
We write $\rho_k^{\star(n)}(T'_n) = (T'_k, (L'_B)_{B \in T'_k})$, where $T'_k = T'_n \cap [k]$.

By the spatial Markov property Lemma 6, the above claim amounts to showing that, given $\rho_k^{\star(n)}(T'_n) = \mathbf{v}^\star$, the internal structures of $\mathbf{v}^\star$ in T'_n are conditionally independent trees that are conditionally uniformly distributed with sizes given by the label sets of $\mathbf{v}^\star$.

We will prove this claim on a case-by-case basis that mirrors the proof of Proposition 8 showing that, conditioned on each case, the claim holds.

Case (A). Consider any $i, \tilde{\imath} \in L_B$ for an edge $B \in \mathbf{s}$ such that either (i) B is internal or (ii) B is external and $\#L_B > 1$. Also consider any $B' \in \mathbf{s}$. We will condition on the event that $I = i$, $\widetilde{I} = \tilde{\imath}$ and $\beta' = B'$. Note that on this event the collapsed tree $\rho_k^{\star(n)}(T_n')$ after the down-move and up-move is determined as a deterministic function of $\mathbf{t}^\star$, $\tilde{\imath}$ and B'.

Consider a uniform growth process up to some step $m \geq 1$. Suppose we remove leaf m and the edge below it, it follows from the uniform growth rules that the remaining subtree with leaves labeled by $[m-1]$ is uniformly distributed. By exchangeability of leaf labels, removing any other leaf $\tilde{\imath} \in [m-1]$, the internal structure of the remaining tree is uniformly distributed. Similarly, if we insert a new label, the new tree with the additional leaf is uniformly distributed with the given set of labels.

In this case, as in the proof of Proposition 8, the dropped label $\tilde{\imath}$ is at least $k+1$. There is no resampling: the edge B loses label $\tilde{\imath}$, the label set of the edge B' gains the label. Hence, the tree shape remains $\mathbf{s}$ after the down-move and also after the up-move. By Lemma 6, the internal structures of $\mathbf{t}^\star$ in T_n are mutually independent uniform binary trees under $\mathbf{P}$, and since I is independent of T_n, this also holds under $\mathbf{P}(\,\cdot\,|\,I=i)$. Consider the internal structure $\text{int}(V_B)$ of $\mathbf{t}^\star$ in T_n supported by edge B. Applying Lemma 4 to $\text{int}(V_B)$ and further conditioning on $\widetilde{I} = \tilde{\imath}$, the internal structures after the down-move are independent under $\mathbf{P}(\,\cdot\,|\,I=i, \widetilde{I} = \tilde{\imath})$. By the previous paragraph, the same holds for the internal structures after the up-move, under $\mathbf{P}(\,\cdot\,|\,I=i, \widetilde{I} = \tilde{\imath}, \beta' = B')$, as required.

Case (B). Consider any $i \in [k]$ with $x_i = 1$ and $y_{\overleftarrow{\{i\}}} > 0$ and $\tilde{\imath} \in L_{\overleftarrow{\{i\}}}$. Also consider any $B' \in \mathbf{s}$. We will condition on the event that $I = i$, $\widetilde{I} = \tilde{\imath}$ and $\beta' = B'$. Note that on this event the collapsed tree $\rho_k^{\star(n)}(T_n')$ is not usually a deterministic function of $\mathbf{t}^\star$, i, $\tilde{\imath}$ and B', since the down-move splits a new label set for $\{i\}$ off the label set $L_{\overleftarrow{\{i\}}}$. The tree shape is unaffected by this.

Let $B_1 = \{i\}$ and consider the associated unranked internal structure V_{B_1} of $\mathbf{t}^\star$ in T_n. The swapped label $\tilde{\imath}$ is in V_{B_1} as for the corresponding case in the proof of Proposition 8. Since B_1 is an internal edge of $\mathbf{s}$, V_{B_1} has a leaf labeled 1 that corresponds to B_1 (see Figure 5). Consider the first subtree U on the spinal decomposition from leaf 1 to the root in V_{B_1}. Then, during the down-move, leaf i in T_n swaps with the smallest label $\tilde{\imath}$ in this subtree to give us the new unranked internal structure of edge $\{i\} \in \mathbf{s}$ in T_{n-1}'. Remove this subtree U from V_{B_1}, but leave leaf 1. The remaining subtree (say, V) in V_{B_1} is precisely the new unranked internal structure of edge $B_1 \in \mathbf{s}$ in T_{n-1}'. We note that their label sets, denoted by L_{B_1}'' and $L_{\{i\}}''$, determine $\rho_k^{\star(n-1)}(T_{n-1}')$ and are recorded in $\rho_k^{\star(n-1)}(T_{n-1}')$. The rest of the internal structures is not affected. Applying Lemma 7 to $\text{int}(V_{B_1})$, we obtain that the internal structures after the down-move are independent under $\mathbf{P}(\,\cdot\,|\,I=i, \widetilde{I} = \tilde{\imath}, L_{B_1}'', L_{\{i\}}'')$. The up-move is similar to the previous case.

Case (C). Consider any $i \in [k]$ with $x_i = 1$ and $y_{\overleftarrow{\{i\}}} = 0$. On $\{I = i\}$, this determines $\widetilde{I} = \tilde{\imath} \in [k]$ and the $(k-1)$-tree shape $T_n' \cap ([k] \setminus \{\widetilde{I}\}) = \tilde{\mathbf{v}}$ after the

down-move. Also consider $B' \in \widetilde{\mathbf{v}}$. As usual, condition on the event that $I = i$, $\widetilde{I} = \tilde{\imath}$ and $\beta' = B'$.

In this case some label in $[k]$ is dropped and resampled. Hence the fact that the internal structures after the down-move are mutually independent uniform binary trees under $\mathbf{P}(\,\cdot\,|\,I = i, \widetilde{I} = i)$ is obvious. We focus on the resampling step.

When the label $\tilde{\imath}$ resamples it affects the collapsed tree after the down-move in the following way. The label $\tilde{\imath}$ is added to the label set of the edge $\beta' = B'$ increasing its weight by one, and the label set is then split according to a three-color Pólya urn, as follows, into three sets that we denote by $L'_{B'}$, $L'_{\{\tilde{\imath}\}}$ and $L'_{B' \cup \{\tilde{\imath}\}}$. If $B' = \{j\}$ is external, initially the two lowest labels $\tilde{\imath}, j \in [k]$ have the first two colors. If B' is internal, the lowest label $\tilde{\imath} \in [k]$ is initially allocated the last color. In both cases, the other labels of B' are added in increasing order according to the standard update rules with initial weights $(1/2, 1/2, 1/2)$. This is so, because the internal structure supported by B' after the down-move is uniform and stays uniform after insertion of $\tilde{\imath}$, so that this split can be read from the proof of Corollary 11. In particular, this achieves the split of the edge weight of edge B' according to the Dirichlet multinomial kernel given in Definition 12. This split determines $\rho_k^{\star(n)}(T'_n)$ and is recorded in $\rho_k^{\star(n)}(T'_n)$. Resampling further inserts the label corresponding to $\tilde{\imath}$ into $\mathrm{int}(V_{B'})$ as in the up-move of Definition 1 as either the smallest (external edge with $\tilde{\imath} < j$) or second smallest (internal edge or external with $\tilde{\imath} > j$) label. This results in a uniform tree, by uniform growth and exchangeability. By Lemma 20, the three new internal structures are independent. The rest of the internal structures being unaffected, we conclude that the internal structures of $\rho_k^{\star(n)}(T'_n)$ in T'_n are independent under $\mathbf{P}(\,\cdot\,|\,I = i, \widetilde{I} = \tilde{\imath}, \beta' = B', L'_{B'}, L'_{\{\tilde{\imath}\}}, L'_{B' \cup \{\tilde{\imath}\}})$.

Since the joint conditional distributions of internal structures given the events that appeared in the three (exhaustive) cases, only depend on the collapsed tree $\rho_k^{\star(n)}(T'_n)$, this joint distribution is also the conditional distribution given just the collapsed tree. This completes the proof. $\qquad\square$

APPENDIX A. TREES DECORATED WITH STRINGS OF BEADS

Although a decorated k-tree $\rho_k^{\bullet(n)}(\mathbf{t})$ contains information on the leaf masses of collapsed subtrees of $\mathbf{t}$, it does not contain much information on lengths in these subtrees. For an illustration of this point, see the left-most tree in Figure 5. The three red edges in $\mathbf{t}$ contract to a single edge in $\rho_k^{\bullet(n)}(\mathbf{t})$. This number, 3, is lost in the projection. To construct the Aldous diffusion as discussed in Section 1.1, we must retain more of this length information. But if we simply enrich our state space by decorating edges with lengths, in addition to weights, then the resulting projected processes would be non-Markovian. To retain the Markov property while incorporating length information, we decorate each edge by a composition, which we call a *(discrete) string of beads*.

For $m \in \mathbb{N}$, denote by $\mathcal{C}_m = \bigcup_{\ell \geq 0}\{(y_1, \ldots, y_\ell) \in \mathbb{N}^\ell \colon y_1 + \cdots + y_\ell = m\}$ the set of compositions of m, with the convention that $\mathcal{C}_0$ is a singleton set containing the empty composition, denoted by '()'. We think of compositions as sequences of bead sizes, strung along a length of string. Continuum strings of beads were introduced to study weighted continuum random trees [26, 27]. In that setting, a string of beads is a purely atomic measure on an interval. Here, the composition $(y_1, \ldots, y_\ell)$ could be thought of as describing a measure $y_1 \delta_1 + \cdots + y_\ell \delta_\ell$ on an interval $[0, \ell + 1]$.
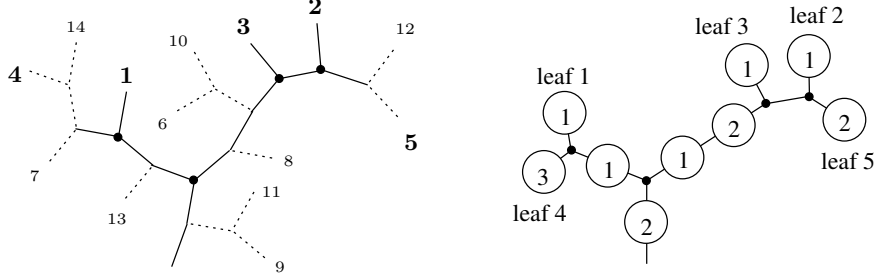
FIGURE 8. $\mathbb{T}_{[14]} \ni \mathbf{t} \mapsto \rho_5^{\circ(14)}(\mathbf{t}) \in \mathbb{T}_5^{\circ(14)}$.

For $\mathbf{s} \in \mathbb{T}_{[k]}$, recall the set of edge weight functions $\mathcal{N}_m^{\mathbf{s}}$ from Definition 6. The set of *k-trees decorated with strings of beads* is given by

$$\mathbb{T}_{[k]}^{\circ(n)} = \bigcup_{\mathbf{s} \in \mathbb{T}_{[k]}} \{\mathbf{s}\} \times \left(\bigcup_{f \in \mathcal{N}_n^{\mathbf{s}}} \{(f(\{1\}), \ldots, f(\{k\}))\} \times \prod_{B \in \mathbf{s}: \#B \geq 2} \mathcal{C}_{f(B)} \right).$$

We interpret an element $\mathbf{t}^\circ = (\mathbf{s}, x_1, \ldots, x_k, (y_1^B, \ldots, y_{\ell_B}^B)_{B \in \mathbf{s}: \#B \geq 2}) \in \mathbb{T}_{[k]}^{\circ(n)}$ as a k-tree $\mathbf{s} \in \mathbb{T}_{[k]}$ with leaf masses $x_1, \ldots, x_k$, each at least one, and with each internal edge B decorated by a (possibly empty) string of beads $(y_1^B, \ldots, y_{\ell_B}^B)$, with bead masses listed in order of increasing distance from the root. For example, in the right panel of Figure 8, the edges have decorations

$$x_1 = 1, \quad x_2 = 1, \quad x_3 = 1, \quad x_4 = 3, \quad x_5 = 2,$$
$$\left(y_\bullet^{\{1,4\}}\right) = (1), \quad \left(y_\bullet^{\{2,5\}}\right) = (), \quad \left(y_\bullet^{\{2,3,5\}}\right) = (1,2), \quad \left(y_\bullet^{[5]}\right) = (2).$$

Now, $\ell_B + 1$ is the *length of edge* B in $\mathbf{t}^\circ$, and $y_1^B + \cdots + y_{\ell_B}^B$ is its *mass*.

Analogously to Section 3, we define a projection $\rho_k^{\circ(n)} \colon \mathbb{T}_{[n]} \to \mathbb{T}_{[k]}^{\circ(n)}$. Fix $\mathbf{t} \in \mathbb{T}_{[n]}$. We define

$$\mathbf{r} = \left\{ B \in \mathbf{t} \colon \#(B \cap [k]) \geq 2 \right\} \cup \left\{ \bigcup_{B \in \mathbf{t}: \, B \cap [k] = \{j\}} B \ \middle| \ j \in [k] \right\}.$$

We view $\mathbf{r}$ as a subgraph of $\mathbf{t}$ formed by pruning away fringe subtrees that contain no labels from $[k]$, and if there is a path of degree-2 branch points terminating in a leaf, contracting that down to a single edge. See the left panel of Figure 8, in which the subgraph $\mathbf{r}$ is shown in solid lines, with $\mathbf{t} \setminus \mathbf{r}$ shown in dashed lines.

Denote by μ the measure on $\mathbf{t}$ that assigns to each of the leaves $\{1\}, \ldots, \{n\}$ mass 1. Consider the projection map $\pi_1 \colon \mathbf{t} \to \mathbf{r}$ given by $B \mapsto \bigcap_{C \in \mathbf{r}: \, C \supseteq B} C$. Denote by $\widetilde{\mu}$ the π_1-pushforward of μ. Note that, for $B \in \mathbf{r}$, $\widetilde{\mu}\{B\} = 0$ if and only if B is a degree-3 branch point in $\mathbf{r}$. Indeed, if B is a leaf in $\mathbf{r}$ then its π_1-pre-image is a subtree containing one leaf with low label ($\leq k$), and possibly more high-labeled leaves. If B is a degree-2 branch point in $\mathbf{r}$ then its π_1-pre-image is a fringe subtree that branches off of $\mathbf{r}$ at B, which therefore contains at least one high-labeled leaf. But if B has degree 3 in $\mathbf{r}$ then its π_1-pre-image is only $\{B\}$.

Now define $\mathbf{s} := \mathbf{t} \cap [k] = \mathbf{r} \cap [k]$, and let $\pi_2 \colon \mathbf{r} \to \mathbf{s}$ denote the map $B \mapsto B \cap [k]$. This map bijects the leaves and degree-3 branch points of $\mathbf{r}$ with the vertices of $\mathbf{s}$, but the branch points of $\mathbf{s}$ may also have degree-2 vertices of $\mathbf{r}$ in their π_2-pre-image.

For $j \in [k]$ we define x_j to be the $\widetilde{\mu}$-mass of the leaf of $\mathbf{r}$ that maps to $\{j\}$ via π_2. For $B \in \mathbf{s}$ with $\#B \geq 2$, let $\ell_B := \#\pi_2^{-1}(B) - 1$. For the purpose of the following, denote the elements of $\pi_2^{-1}(B)$ by $B_1, \ldots, B_{\ell_B+1}$, in order of increasing distance from the root, so that $B_1, \ldots, B_{\ell_B}$ all have degree 2 and B_{ℓ_B+1} has degree 3. We define $\left(y_j^B\right)_{j \in [\ell_B]} := (\widetilde{\mu}(B_j))_{j \in [\ell_B]}$. Finally,

$$\rho_k^{\circ(n)}(\mathbf{t}) := \left((x_j)_{j \in [k]}, \left(y_1^B, \ldots, y_{\ell_B}^B\right)_{B \in \mathbf{s}:\ \#B \geq 2} \right).$$

Theorem 21. *Let $(T(j))_{j \geq 0}$ be a uniform chain running in stationarity. Then, for $k \leq n$, $(\rho_k^{\circ(n)}(T(j)))_{j \geq 0}$ is a Markov chain running in stationarity.*

The proof of this extension of Theorem 2 is very similar to that already given, so we omit it. We note this because we will give a completely different proof of the analogous result in the continuum, and the discrete proofs presented in this paper will shed more light on the intuition behind those results.

We also note that the discrete strings of beads considered here could also be taken as representations of finite *interval partitions*. In the continuum, however, there is a crucial advantage working with interval partitions rather than strings of beads (purely atomic measures). This is the approach taken in [13], which the reader can consult for a discussion of interval partitions in this setting.

As a final remark, one may wonder whether the down-up chain in Definition 1 is the only modification of label dynamics for the Aldous chain that would make Theorem 2 possible. We found our down-up chain as a natural way to project the stationary Aldous chain down to a stationary evolving 2-tree with leaf masses and a string of beads. This idea goes back to [22]. By working our way upwards, from 2 up to n leaves, while keeping the consistency in stationarity, we arrived at Definition 1. The proof presented here is reversed from the order in which it was discovered, beginning with label dynamics for n leaves and then projecting down, and only considering edge masses rather than strings of beads for simplicity.

References

[1] D. Aldous. The continuum random tree. I. *Ann. Probab.*, 19(1):1–28, 1991.

[2] D. Aldous. Wright-Fisher diffusions with negative mutation rate! `http://www.stat.berkeley.edu/~aldous/Research/OP/fw.html`, 1999.

[3] D. J. Aldous. Mixing time for a Markov chain on cladograms. *Combin. Probab. Comput.*, 9(3):191–204, 2000.

[4] D. Barker. *LVB 1.0: reconstructing evolution with parsimony and simulated annealing.* Daniel Barker, 1997.

[5] D. M. Blei and M. I. Jordan. The nested Chinese restaurant process and Bayesian nonparametric inference of topic hierarchies. *J. ACM*, 57(2), 2010.

[6] A. Borodin and P. L. Ferrari. Anisotropic growth of random surfaces in 2+ 1 dimensions. *Communications in Mathematical Physics*, 325(2):603–684, 2014.

[7] A. Borodin and G. Olshanski. Infinite-dimensional diffusions as limits of random walks on partitions. *Probab. Theory Related Fields*, 144(1-2):281–318, 2009.

[8] P. Diaconis and J. A. Fill. Strong stationary times via a new form of duality. *Ann. Probab.*, 18(4):1483–1522, 1990.

[9] S. N. Evans and A. Winter. Subtree prune and regraft: a reversible real tree-valued Markov process. *Ann. Probab.*, 34(3):918–961, 2006.

[10] J. Felsenstein. *Inferring Phylogenies.* Sinauer, 2003.

[11] J. A. Fill. Strong stationary duality for continuous-time Markov chains. I. Theory. *J. Theoret. Probab.*, 5(1):45–70, 1992.

[12] D. J. Ford. *Probabilities on cladograms: introduction to the alpha model.* Ph. D. thesis, Stanford University, 241 p., article version also available at arXiv:math.PR/0511246, 2006.

[13] N. Forman, S. Pal, D. Rizzolo, and M. Winkel. Diffusions on a space of interval partitions with Poisson-Dirichlet stationary distributions. arXiv:1609.06706, 2016.

[14] J. Fulman. Stein's method and Plancherel measure of the symmetric group. *Trans. Amer. Math. Soc.*, 357(2):555–570, 2005.

[15] J. Fulman. Commutation relations and Markov chains. *Probab. Theory Related Fields*, 144(1-2):99–136, 2009.

[16] J. Fulman. Mixing time for a random walk on rooted trees. *Electron. J. Combin.*, 16(1):Research Paper 139, 13, 2009.

[17] A. Gnedin and J. Pitman. Regenerative composition structures. *Ann. Probab.*, 33(2):445–479, 2005.

[18] B. Haas, G. Miermont, J. Pitman, and M. Winkel. Continuum tree asymptotics of discrete fragmentations and applications to phylogenetic models. *Ann. Probab.*, 36(5):1790–1837, 2008.

[19] J. Kemeny and J. Snell. *Finite Markov chains*. University series in undergraduate mathematics. Van Nostrand, 1960.

[20] S. Kerov. The boundary of Young lattice and random Young tableaux. In *Formal power series and algebraic combinatorics (New Brunswick, NJ, 1994)*, volume 24 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 133–158. Amer. Math. Soc., Providence, RI, 1996.

[21] J. Paisley, C. Wang, D. M. Blei, and M. I. Jordan. Nested hierarchical Dirichlet processes. *IEEE Transactions on pattern analysis and machine intelligence*, 37(2), 2015.

[22] S. Pal. On the Aldous diffusion on Continuum Trees. I. *arXiv:1104.4186 [math.PR]*, 2011.

[23] S. Pal. Wright-Fisher diffusion with negative mutation rates. *Ann. Probab.*, 41(2):503–526, 2013.

[24] C. Y. A. Pang. Lumpings of algebraic Markov chains arise from subquotients. arXiv: 1508.01570v2 [math.CO].

[25] L. A. Petrov. A two-parameter family of infinite-dimensional diffusions on the Kingman simplex. *Funktsional. Anal. i Prilozhen.*, 43(4):45–66, 2009.

[26] J. Pitman and M. Winkel. Regenerative tree growth: binary self-similar continuum random trees and Poisson-Dirichlet compositions. *Ann. Probab.*, 37(5):1999–2042, 2009.

[27] F. Rembart and M. Winkel. Recursive construction of continuum random trees. To appear in *Ann. Probab.*. Preprint at arXiv:1607.05323 [math.PR], 2016.

[28] J.-L. Rémy. Un procédé itératif de dénombrement d'arbres binaires et son application à leur génération aléatoire. *RAIRO Inform. Théor.*, 19(2):179–195, 1985.

[29] L. C. G. Rogers and J. W. Pitman. Markov functions. *The Annals of Probability*, 9(4):573–582, 1981.

[30] F. Ronquist, M. Teslenko, P. Van Der Mark, D. L. Ayres, A. Darling, S. Höhna, B. Larget, L. Liu, M. A. Suchard, and J. P. Huelsenbeck. MrBayes 3.2: efficient Bayesian phylogenetic inference and model choice across a large model space. *Systematic biology*, 61(3):539–542, 2012.

[31] J. Schweinsberg. An $O(n^2)$ bound for the relaxation time of a Markov chain on cladograms. *Random Structures Algorithms*, 20(1):59–70, 2002.

NOAH FORMAN, DEPARTMENT OF STATISTICS, UNIVERSITY OF OXFORD, 24–29 ST GILES', OXFORD OX1 3LB, UK, EMAIL: NOAH.FORMAN@GMAIL.COM

SOUMIK PAL, DEPARTMENT OF MATHEMATICS, UNIVERSITY OF WASHINGTON, SEATTLE WA 98195, USA, EMAIL: SOUMIKPAL@GMAIL.COM

DOUGLAS RIZZOLO, DEPARTMENT OF MATHEMATICAL SCIENCES, UNIVERSITY OF DELAWARE, NEWARK DE 19716, USA, EMAIL: DRIZZOLO@UDEL.EDU

MATTHIAS WINKEL, DEPARTMENT OF STATISTICS, UNIVERSITY OF OXFORD, 24–29 ST GILES', OXFORD OX1 3LB, UK, EMAIL: WINKEL@STATS.OX.AC.UK