

Local limit of the fixed point forest*

Tobias Johnson[†] Anne Schilling[‡] Erik Slivken[§]

Abstract

Consider the following partial “sorting algorithm” on permutations: take the first entry of the permutation in one-line notation and insert it into the position of its own value. Continue until the first entry is 1. This process imposes a forest structure on the set of all permutations of size n , where the roots are the permutations starting with 1 and the leaves are derangements. Viewing the process in the opposite direction towards the leaves, one picks a fixed point and moves it to the beginning. Despite its simplicity, this “fixed point forest” exhibits a rich structure. In this paper, we consider the fixed point forest in the limit $n \rightarrow \infty$ and show using Stein’s method that at a random permutation the local structure weakly converges to a tree defined in terms of independent Poisson point processes. We also show that the distribution of the length of the longest path from a random permutation to a leaf converges to the geometric distribution with mean $e - 1$, and the length of the shortest path converges to the Poisson distribution with mean 1. In addition, the higher moments are bounded and hence the expectations converge as well.

Keywords: permutations; trees; fixed points; sorting algorithm; Poisson point process; Stein’s method; weak convergence.

AMS MSC 2010: Primary 60J80, Secondary 05C05; 05A05; 60C05; 60F99.

Submitted to EJP on July 20, 2016, final version accepted on February 4, 2017.

Supersedes arXiv:1605.09777.

1 Introduction

1.1 Fixed point forest

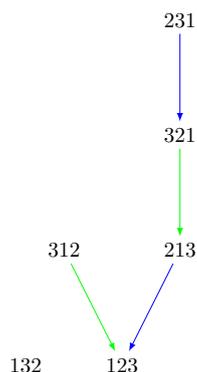
Consider a deck of n cards labeled $1, 2, \dots, n$ given in an arbitrary order. Take the top card and reinsert it into the pile at the position of its value. This gives rise to a partial sorting algorithm, where the algorithm stops when card 1 is on the top. This can be formulated in terms of the set \mathfrak{S}_n of permutations of size n . Each permutation

*TJ was partially supported by NSF grant DMS-1401479. AS was partially supported by NSF grants OCI-1147247 and DMS-1500050.

[†]New York University. E-mail: tobias.johnson@nyu.edu <http://cims.nyu.edu/~tjohnson>

[‡]University of California, Davis. E-mail: anne@math.ucdavis.edu <http://www.math.ucdavis.edu/~anne>

[§]University of California, Davis. E-mail: erikslivken@math.ucdavis.edu <http://www.math.ucdavis.edu/~erikslivken>

Figure 1: Fixed point forest F_3 .

π represents an ordered deck of cards and can be written in one-line notation as $\pi(1)\pi(2)\dots\pi(n)$, where $\pi(i)$ is the value of the i -th card in the deck with $\pi(1)$ being the topmost card. Then $\pi(1)$ is removed and reinserted into the position $\pi(1)$. For example, the permutation $3142 \in \mathfrak{S}_4$ goes to 1432 . This defines a graph with vertices being all permutations in \mathfrak{S}_n and vertex π being connected by an edge to vertex π' , if π' is obtained from π by moving the first entry $\pi(1)$ to position $\pi(1)$. It was shown in [13] that the graph is in fact a rooted forest, which we call F_n . A *rooted forest* is a union of *rooted trees*, where a tree is a graph that does not contain any closed loops involving distinct vertices. Since later on we will consider trees with random permutations as roots, we call each root of F_n a *base*. The bases of the forest are the permutations with $\pi(1) = 1$, and the leaves are the derangements, that is, permutations without fixed points. The partial sorting algorithm describes a path from a permutation to a base. Viewing the process in the opposite direction towards the leaves (i.e. vertices without incoming edge), one picks a fixed point $\pi(i) = i$ of the permutation and moves it to the beginning, which we call *bumping* the fixed point i . For this reason, we call this forest the *fixed point forest*. Examples for $n = 3$ and $n = 4$ are given in Figures 1 and 2, respectively.

The goal of this paper is to understand the local structure of this forest as $n \rightarrow \infty$. Put differently, we would like to describe the neighborhood of a “typical” permutation in the forest for large n . We carry this out in Theorem 3.6, where we find the limit of the forest as $n \rightarrow \infty$ in the sense of *local weak convergence*, as defined in Section 2.1 (see [2] for more background). The limit is a random tree that we construct from an infinite collection of independent Poisson point processes. The proof relies on Stein’s method for Poisson approximation. The limit tree seems to have interesting properties; see Section 1.4 for comparisons to other processes. As a corollary of this local weak convergence, we show that the distributions of the distance from a random permutation to the farthest and nearest leaves descending from it converge, respectively, to a geometric distribution with mean $e - 1$ and a Poisson distribution with mean 1. With some additional work, we prove that the higher moments are bounded and hence the expectations converge as well (see Theorem 5.1).

We were first made aware of the fixed point forest by Gwen McKinley who studied it in her undergraduate thesis [13], which she wrote under the guidance of the second author. McKinley was introduced to the problem at the Missouri State University summer REU program by Les Reid to whom the process was suggested by Gerhardt Hinkle. In her thesis, McKinley investigated several global properties of the forest, finding for instance that the longest path to the base is of length $2^{n-1} - 1$. Despite the simplicity of the description of the fixed point forest, many basic questions about it seem difficult and

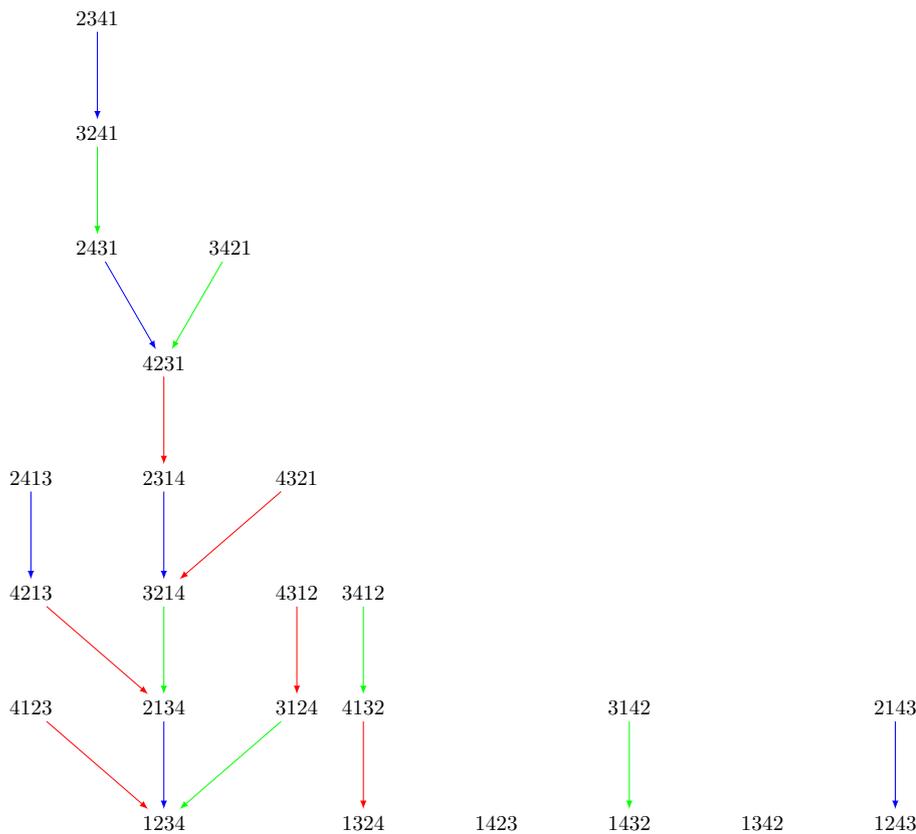


Figure 2: Fixed point forest F_4 .

remain unanswered. We discuss some of them in more detail in Section 6. Both locally and globally, the structure of the fixed point forest seems quite rich.

In the remainder of the introduction, we provide a non-technical discussion of the local structure of the forest F_n in the limit $n \rightarrow \infty$. Let π_n be a uniformly random permutation in F_n . The essential information for determining the neighborhood of π_n is the location of its fixed and near-fixed points. Our idea for understanding the limiting structure of F_n is to construct a sort of limit of the fixed and near-fixed points of π_n . By rescaling by a factor of $1/n$, these are represented as Poisson point processes on $[0, 1]$. Then, we define a tree from these point processes, which will turn out to be the local limit of F_n (see Section 3).

1.2 Moving towards leaves in the tree

Suppose that π is a permutation and that we would like to enumerate its descendants up to three levels in the forest; that is, we want to determine all permutations obtained from π by bumping fixed points to the beginning no more than three times. What information about π do we need?

The answer is that we must know all i such that $\pi(i) = i$, $\pi(i) = i + 1$, or $\pi(i) = i + 2$. This is best seen by example. To unify our terminology, say that the letter $\pi(i)$ or the position i is *k-separated* if $\pi(i) = i + k$. A 0-separated letter is simply a fixed point.

Example 1.1. Suppose that π has 0-separated letters at positions 7 and 27, that it has a 1-separated letter at 18, and that it has a 2-separated letter at 13. Then π has two children in the forest, given by bumping the letters in either position 7 or 27.

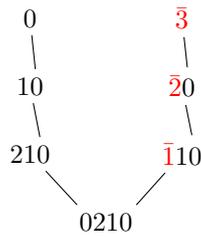


Figure 3: The descendants of π from Example 1.1, encoded by their words. Only three levels are given here, since we are using the information only from the 0-, 1-, and 2-separated letters. To construct a tree like this given only the encoding of π by its word, when we bump the 0-separated letter at a given position, we remove the 0 from this position and subtract 1 from all letters previous to the bumped position. A value of $-k$ is indicated here by \bar{k} , that is, barred and colored red.

If the letter in position 7 is bumped, then the resulting permutation still has a 0-separated letter at 27, a 1-separated letter at 18, and a 2-separated letter at 13. From here, one child is given by bumping the letter in position 27. Then, this turns the 1-separated letter at position 18 into a 0-separated letter at 19, and it turns the 2-separated letter at 13 into a 1-separated letter at 14. From here, there is a child given by bumping the letter in position 19 (as well as another child after that).

If the letter in position 27 is bumped in the first step, then the 0-separated letter at 7 is destroyed, the 1-separated letter at 18 becomes a 0-separated letter at 19, and the 2-separated letter becomes a 1-separated letter at 14. Now there is a child given by bumping the letter in position 19, which turns the 1-separated letter at 14 into a 0-separated letter at 15, which can then be bumped.

This example suggests a way to encode a permutation according to its k -separated letters for $k = 0, \dots, K$. We create a word containing a k for each k -separated position in the permutation, in the order that they appear. For example, with $K = 2$, the word corresponding to π in Example 1.1 is 0210. When we bump the 0-separated letter at a given position, we remove the 0 from this position, subtract 1 from all letters previous to the bumped position, and we leave alone all letters after the bumped position. We write a $-k$ letter as \bar{k} , that is, barred and colored red. These negative letters are irrelevant for determining the descendants of a permutation, but we leave them in the word for consistency with the next section. Thus bumping the first 0 from 0210 yields 210, and bumping the second 0 yields $\bar{1}10$. This gives us the compact depiction of the descendants of π in Figure 3.

We warn the reader that this picture is incomplete in one way. When a 0-separated letter at position i is bumped, it creates a new $(i - 1)$ -separated letter at position 1. In Example 1.1, this makes no difference, but in the following example it does.

Example 1.2. Suppose that π is the permutation 42135, in one-line notation. Then π has 0-separated letters at locations 2 and 5, and it has no 1- or 2-separated letters. If 5 is bumped, the resulting permutation has no 0-separated letters and hence no children. If 2 is bumped, then we still have a 0-separated letter at 5. We also have a new 1-separated letter at position 1, namely the 2 that was just bumped. Thus, if 5 is bumped, we can then bump the 2 again.

If we start with the word 00 corresponding to 42135 and then follow the rules laid out before this example for manipulating these words, we miss this last descendant (see Figure 4).

Local limit of the fixed point forest

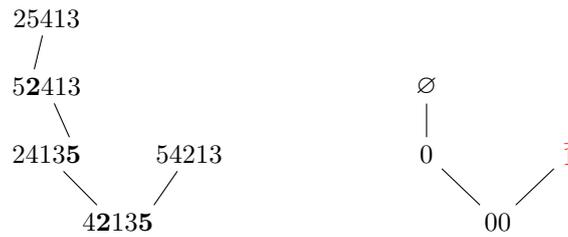


Figure 4: On the left, we depict the descendants of the permutation 42135 in the forest. Fixed points are in bold. On the right, we show the apparent forest computed only using the words giving the order of the k -separated letters. This misses a descendant created by the “reentry” of 2 as a fixed point.

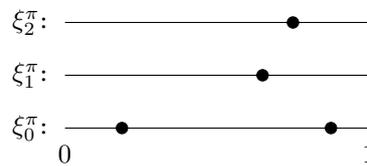


Figure 5: The point processes ξ_0^π , ξ_1^π , and ξ_2^π on the interval $[0, 1]$, representing the 0-, 1-, and 2-separated letters in the abstracted permutation π . The word associated with this permutation for $K = 2$ is 0120, and the tree of its descendants up to three levels is the same as in Figure 3.

This possible “reentry” of a bumped letter as a fixed point complicates the picture. When we determine the descendants of a permutation π up to level K , this reentry can only occur if there is a 0-separated letter at one of positions $1, \dots, K - 1$. For fixed K , this is vanishingly unlikely as $n \rightarrow \infty$, and so in constructing the limit tree, we ignore it. The essential idea to this construction is to take a random word on the alphabet $0, \dots, K$ and make a tree by the procedure described in Figure 3. In constructing the limit tree, one complication is that we want K to be infinite. To address this, we represent the locations of k -separated letters for each k in an abstracted permutation π as a point process ξ_k^π , representing a set of locations on the interval $[0, 1]$. These will be *independent Poisson point processes* with intensity one in the limit. For any fixed K , we can then obtain a string by writing a k for each point of ξ_k^π for $0 \leq k < K$, sorted by the positions of the points in $[0, 1]$. This string then determines a tree up to level K following the procedure sketched out after Example 1.1.

Example 1.3. Suppose ξ_0^π , ξ_1^π , and ξ_2^π contain points as depicted in Figure 5. Then the word associated with the abstracted permutation π for $K = 2$ is 0120, and the tree generated by it up to three levels is the same one as in Example 1.1 and Figure 3.

One can also construct the tree directly from the point processes, without the intermediate step of converting to a word. The abstracted permutation π with associated point processes $\xi_0^\pi, \xi_1^\pi, \dots$ has $|\xi_0^\pi|$ children, with $|\xi_0^\pi|$ denoting the number of points in ξ_0^π . The point processes obtained by bumping a point $x \in [0, 1]$ of ξ_0^π are given by removing the point x from ξ_0^π , and by “shifting down” each point process on $[0, x)$ as in Figure 6.

Local limit of the fixed point forest

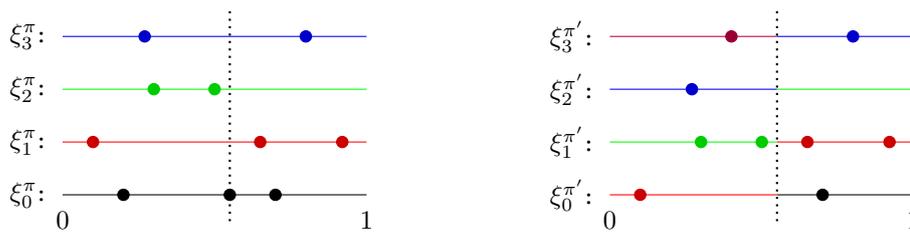


Figure 6: If π is an abstracted permutation and π' is its child given by bumping the middle point x in ξ_0^π , then the point processes $\xi_k^{\pi'}$ is equal to ξ_{k+1}^π on $[0, x)$ and is equal to ξ_k^π on $(x, 1]$, as depicted above.

1.3 Moving towards the base in the tree

In the previous section, we gave a loose account of how to define the descendants of an abstracted permutation π in the limit tree. We need to define the entire tree, however, which includes the ancestors of π and their descendants.

Suppose π is a (non-abstracted) permutation, and we would like to determine both how many children and how many siblings it has in the forest. Again, we ask the question of what information about π we need to find this out.

As before, we need to know the locations of 0-separated letters in π . We also need to know the locations of -1 -separated letters, which can become 0-separated letters in the parent of π . Finally, we need to know the value of $\pi(1)$, as this determines the ancestor of π in the forest.

Example 1.4. Let π be the permutation from Example 1.1, which has 0-separated letters at positions 7 and 27. Suppose that it has -1 -separated letters at 15 and 36, and suppose that $\pi(1) = 20$.

As before, π has two children in the forest, given by bumping positions 7 and 27. When we move towards the base in the tree to the parent of π , the 0-separated letter at position 7 becomes a 1-separated letter at position 6, and the -1 -separated letter at position 15 becomes a 0-separated letter at position 14, while the separated letters after position 20 remain the same. The permutation also has a new 0-separated letter at position 20, which if bumped leads to π . Thus the parent of π has three children total, and π has two siblings.

Again, we can view this in terms of words encoding the k -separated letters. We can view the permutation π of Example 1.4 as the word $02\bar{1}1|0\bar{1}$, with the $|$ symbol specifying the value of $\pi(1)$. When moving towards the base in the tree, a 0 is inserted in the position of the $|$ symbol, and all values to the left of the $|$ are incremented. See Figure 7 for a depiction of Example 1.4 in these terms. As before, this picture is incomplete: if the first character in the word corresponds to a separated letter at position 1, this character is deleted rather than incremented when moving towards the base in the tree. This will be irrelevant in the limit, since for a fixed K and r , a random permutation is vanishingly unlikely to have a k -separated letter with $|k| \leq K$ occurring in the first r positions.

The extra ingredient in moving towards the base in the tree rather than towards the leaves is knowledge of $\pi(1)$. Looking back at Figure 6, to go backward from π' to π in the limit tree, we need the location of the dotted line, which cannot be determined from $\xi_k^{\pi'}$. In the limit case, these locations will be uniform over $[0, 1]$ and independent of the point processes.

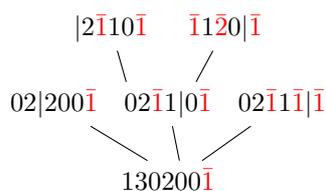


Figure 7: The children, parent, and siblings of the permutation π given in Example 1.4.

1.4 Comparison to other known processes

It is a natural question to compare the fixed point forest or its subtrees to other known random trees, such as the Galton–Watson tree [18]. The tree component of the fixed point forest containing the identity permutation has approximately $(n - 1)!$ vertices. As we discussed earlier, the height of this component is $2^{n-1} - 1$ as shown in [13]. The Galton–Watson tree on the other hand has height \sqrt{N} if the tree has N vertices [1], which is much larger than 2^{n-1} when $N = (n - 1)!$. This is consistent with the fact that the fixed point tree has offspring sizes that are correlated across generations, whereas the Galton–Watson tree has independent siblings. The fixed point forest locally also is quite different from a Galton–Watson tree. For example, no leaf in the fixed point forest has a sibling that is also a leaf.

The process of picking a fixed point at random and moving it to the front (which corresponds to a walk to a leaf) has some resemblance to the Tsetlin library [17, 8, 9], which is a model for the evolution of an arrangement of books in a library shelf over time. It is a Markov chain on permutations, where the entry in the i -th position is moved to the front with probability p_i . However, in the fixed point forest this process eventually stops when a derangement is reached, whereas in the Tsetlin library the process can go on arbitrarily long.

2 The limiting objects

In this section, we provide the precise definition of the limiting tree.

2.1 Local weak convergence in general

The main result of this paper is the *local weak convergence* of the fixed point forest to a certain limiting tree. This mode of convergence is sometimes called *Benjamini–Schramm convergence* after the paper [4]. We will give a short introduction to local weak convergence now, but see [2, Section 2] for a more in depth discussion.

Let G, G_1, G_2, \dots be a sequence of random rooted graphs. For any rooted graph H , $H(r)$ denotes the r -ball around the root of H ; that is, $H(r)$ is the subgraph of H induced by all vertices at distance r or less from the root. We write $H \cong H'$ to signify that H and H' are isomorphic as rooted graphs. We say that G is the *local weak limit* of G_n if for every $r \geq 0$ and every finite graph H ,

$$\mathbf{P}[G_n(r) \cong H] \rightarrow \mathbf{P}[G(r) \cong H]$$

as $n \rightarrow \infty$. Roughly speaking, this says that the view from the root of G_n resembles the view from the root of G in distribution more and more as $n \rightarrow \infty$. Frequently, G_n is a finite, deterministic graph with its root chosen uniformly at random, as will be the case in this paper.

2.2 Construction of the limit tree

The ingredients of our construction are a collection of Poisson point processes $(\xi_k^\rho)_{k \in \mathbb{Z}}$ of unit intensity on $[0, 1]$ and a sequence U_1, U_2, \dots of independent $\text{Unif}[0, 1]$ random variables. Formally, a *point process* is an integer-valued random measure on the Borel sets of \mathbb{R} . One should think of it as a random collection of points, represented as the atoms of the measure. A *Poisson point process* ξ of unit intensity on $[0, 1]$ is characterized by two properties: First, for any interval of length x , the number of points of ξ in the interval is distributed as $\text{Poi}(x)$. Second, the numbers of points of ξ in disjoint intervals are independent. We use the terminology *point process configuration* to mean a deterministic collection of points, also represented formally as a measure.

The point process ξ_k^ρ represents the k -separated letters in the abstracted permutation ρ . Let ρ_1 be the parent of $\rho_0 = \rho$, let ρ_2 be the parent of ρ_1 , and so on. The random variable U_i represents the 0-separated letter in ρ_i that was bumped to create ρ_{i-1} .

To construct the tree, we first define maps corresponding to moving forwards and backwards from a given vertex.

Definition 2.1 (The forward map f). *Let $\xi = (\xi_k)_{k \in \mathbb{Z}}$ be a collection of point process configurations on $[0, 1]$. For any atom x of ξ_0 , let $f(\xi, x) = (\xi'_k)_{k \in \mathbb{Z}}$, where*

$$\xi'_k = \xi_{k+1}|_{[0,x]} + \xi_k|_{(x,1]}.$$

This is the down-shift operation depicted in Figure 6, corresponding to moving forwards towards a leaf in the tree from a permutation to one of its children by bumping the abstracted fixed point x .

Definition 2.2 (The backward map b). *Let $\xi = (\xi_k)_{k \in \mathbb{Z}}$ be a collection of point process configurations on $[0, 1]$. For any $u \in [0, 1]$, let $b(\xi, u) = (\xi'_k)_{k \in \mathbb{Z}}$, where*

$$\begin{aligned} \xi'_0 &= \xi_{-1}|_{[0,u)} + \xi_0|_{(u,1]} + \delta_u, \\ \xi'_k &= \xi_{k-1}|_{[0,u)} + \xi_k|_{(u,1]} \quad \text{for } k \neq 0. \end{aligned}$$

This is the reverse of the forward map, in the following sense: if x is an atom of ξ_0 and $f(\xi, x) = \xi'$, then $b(\xi', x) = \xi$.

Next, we define a tree by applying f and b to map out the abstracted permutations.

Definition 2.3. *Given point process configurations $\xi^\rho = (\xi_k^\rho)_{k \in \mathbb{Z}}$ and a sequence $u = (u_1, u_2, \dots)$ of elements of $[0, 1]$, we construct a rooted tree $\varphi(u, \xi^\rho)$ as follows. We think of each vertex v of this tree as an abstracted permutation, represented by a collection of point processes $\xi^v = (\xi_k^v)_{k \in \mathbb{Z}}$. Let $\rho = \rho_0$ be the root of the tree. First, we give ρ_0 an infinite chain of ancestors ρ_1, ρ_2, \dots . Starting with ξ^{ρ_0} , which is given to us, we inductively define*

$$\xi^{\rho_{i+1}} = b(\xi^{\rho_i}, u_{i+1}).$$

Next, we construct descendants of each ρ_i . For every atom x in $\xi_0^{\rho_i}$ other than u_i , give ρ_i a child $\rho_i(x)$ and define

$$\xi^{\rho_i(x)} = f(\xi^{\rho_i}, x).$$

for all k . (We avoid doing this with $x = u_i$ since this would just recreate ρ_{i-1} .) From here on, we proceed inductively, continuing to extend the tree forwards. Suppose that ξ^v has already been defined. For each atom x in ξ_0^v , extend the tree by creating a child $v(x)$ of v , and define

$$\xi^{v(x)} = f(\xi^v, x).$$

We define $\varphi(u, \xi^\rho)$ as the resulting tree. Also, observing that the r -neighborhood of the root of the tree depends only on $(\xi_{-r+1}^\rho, \dots, \xi_{r-1}^\rho)$ and on u_1, \dots, u_r , define the map φ_r by setting $\varphi_r(u_1, \dots, u_r; \xi_{-r+1}^\rho, \dots, \xi_{r-1}^\rho)$ as the r -neighborhood of the root of $\varphi(u, \xi^\rho)$.

Finally, we construct the limit tree T .

Definition 2.4. Define $T = \varphi(U, \xi^\rho)$, where $U = (U_1, U_2, \dots)$ consists of independent $\text{Unif}[0, 1]$ random variables, and $\xi^\rho = (\xi_k^\rho)_{k \in \mathbb{Z}}$ consists of independent Poisson point processes on $[0, 1]$ with unit intensity, and ξ^ρ and U are independent of each other.

In the following section, we will prove that T is the local limit of F_n as $n \rightarrow \infty$.

3 Local weak convergence of the fixed point tree

Recall that i is a k -separated position in a permutation π if $\pi(i) = i + k$. The main idea in this section is that the neighborhood of a permutation π in F_n up to distance r can typically be reassembled from two pieces of information: the k -separated positions in π for $-r < k < r$, and the values of $\pi(1), \dots, \pi(r)$. The first piece lets us work out the tree forwards (towards leaves) from π . When we rescale by $1/n$, these locations converge to independent Poisson point processes on $[0, 1]$. The second piece of information lets us move backwards in the tree. With the same rescaling, these random variables converge to independent points sampled uniformly from $[0, 1]$. The two pieces of information converge jointly, as shown in Proposition 3.5, and the weak local convergence of the fixed point tree F_n follows easily from this in Theorem 3.6.

While this convergence is what one would expect from well known Poisson approximations of fixed points of random permutations (see [6, Theorem 11], for example), it will take some technical work to prove our precise statement. We will use Stein’s method via size-bias couplings using the framework from [3], which we introduce now. See also [14, Section 4.3] for a more detailed introduction to size-bias couplings. The general idea for our purposes is that we have a collection of 0-1 random variables, and we would like to show that they are well approximated by independent Poisson random variables. If there exist certain couplings described below, Stein’s method gives a quantitative version of this approximation. The bound is given in terms of the covariances of the random variables and does not depend on the couplings, once they are shown to exist.

Condition 3.1. Let $\mathbf{I} = (I_\alpha)_{\alpha \in \mathcal{I}}$ be a collection of 0-1 random variables. For each $\alpha \in \mathcal{I}$, there is a random vector $J_{\bullet, \alpha} = (J_{\beta, \alpha})_{\beta \in \mathcal{I}}$ coupled with \mathbf{I} such that

- $J_{\bullet, \alpha}$ is distributed as \mathbf{I} conditioned on $I_\alpha = 1$;
- we can partition \mathcal{I} into disjoint sets

$$\mathcal{I} = \mathcal{I}_\alpha^+ \cup \mathcal{I}_\alpha^- \cup \{\alpha\}$$

such that with probability one,

$$J_{\beta, \alpha} \leq I_\beta \quad \text{if } \beta \in \mathcal{I}_\alpha^-, \tag{3.1}$$

$$J_{\beta, \alpha} \geq I_\beta \quad \text{if } \beta \in \mathcal{I}_\alpha^+. \tag{3.2}$$

When this condition holds, one can estimate the distance between \mathbf{I} and a vector of independent Poisson random variables in terms of the covariances of the components of \mathbf{I} , with no mention of the coupling. Here, as usual, the covariance between two random variables X and Y is

$$\text{Cov}(X, Y) = \mathbf{E}[XY] - \mathbf{E}[X]\mathbf{E}[Y].$$

The bound is on the *total variation distance* between the laws of the random vectors. For random variables X and Y taking values in some space \mathcal{S} , this distance is defined as

$$d_{TV}(X, Y) = \sup_{S \subseteq \mathcal{S}} |\mathbf{P}[X \in S] - \mathbf{P}[Y \in S]|.$$

Another characterization of the total variation distance between the laws of X and Y is as the minimum of $\mathbf{P}[X \neq Y]$ over all couplings of X and Y ; see [3, Section A.1].

Proposition 3.2 (Corollary 10.J.1 in [3]). *Assume the coupling condition. Let $\mathbf{Y} = (Y_\alpha)_{\alpha \in \mathcal{I}}$ be a vector of independent Poisson random variables with $\mathbf{E}Y_\alpha = \mathbf{E}I_\alpha$. Then*

$$d_{TV}(\mathbf{I}, \mathbf{Y}) \leq \sum_{\alpha \in \mathcal{I}} (\mathbf{E}I_\alpha)^2 + \sum_{\alpha \in \mathcal{I}} \sum_{\beta \in \mathcal{I}_\alpha^-} |\mathbf{Cov}(I_\alpha, I_\beta)| + \sum_{\alpha \in \mathcal{I}} \sum_{\beta \in \mathcal{I}_\alpha^+} \mathbf{Cov}(I_\alpha, I_\beta). \quad (3.3)$$

We will also need the following technical lemma.

Lemma 3.3. *Let $\mathbf{Y} = (Y_\alpha)_{\alpha \in \mathcal{I}}$ and let $\mathbf{Z} = (Z_\alpha)_{\alpha \in \mathcal{I}}$ be vectors of independent Poisson random variables. Then*

$$d_{TV}(\mathbf{Y}, \mathbf{Z}) \leq \sum_{\alpha \in \mathcal{I}} |\mathbf{E}Y_\alpha - \mathbf{E}Z_\alpha|.$$

Proof. Suppose that U and V are Poisson with means $a \leq b$. Then U and V can be coupled by setting $V = U + W$ where $W \sim \text{Poi}(b - a)$ and is independent of U . By Markov's inequality,

$$\mathbf{P}[U \neq V] = \mathbf{P}[W \geq 1] \leq b - a.$$

Applying this coupling to Y_α and Z_α for each α , we obtain a coupling of \mathbf{Y} and \mathbf{Z} in which they differ with probability at most $\sum_{\alpha \in \mathcal{I}} |\mathbf{E}Y_\alpha - \mathbf{E}Z_\alpha|$. \square

Next, we apply Proposition 3.2 to some indicators derived from a random permutation π on $[n] := \{1, \dots, n\}$ with distribution to be specified. Let $I(i, k)$ be an indicator on position i being k -separated in π . Fix r and n , and let

$$\mathcal{I} = \left\{ (i, k) \mid k \in \{-r + 1, \dots, r - 1\}, i \in \{r + 1, \dots, n\}, \text{ and } i + k \in \{1, \dots, n\} \right\},$$

which is the set of (i, k) such that $i > r$ and i might possibly be k -separated in π , for $|k| \leq r - 1$.

Proposition 3.4. *Let π be a uniformly random permutation on $[n]$ conditioned on*

$$\pi(1) = a_1, \dots, \pi(r) = a_r \quad (3.4)$$

for some set of r distinct values $A = \{a_1, \dots, a_r\} \subseteq [n]$. Let $\mathbf{I} = (I(i, k))_{(i, k) \in \mathcal{I}}$ be the vector of indicators defined above. Let $\mathbf{Z} = (Z(i, k))_{(i, k) \in \mathcal{I}}$ where the components of \mathbf{Z} are drawn independently from $\text{Poi}(1/n)$. Then

$$d_{TV}(\mathbf{I}, \mathbf{Z}) \leq \frac{16r^2 + 2r}{n - r - 1}.$$

Proof. The proof proceeds in three steps: First, we construct a coupling satisfying the coupling condition. Next, we bound the expression on the right hand side of (3.3) to obtain a total variation bound between \mathbf{I} and a random vector $\mathbf{Y} = (Y(i, k))_{(i, k) \in \mathcal{I}}$ whose components are independent Poisson random variables with $\mathbf{E}Y(i, k) = \mathbf{E}I(i, k)$. Last, we bound the total variation distance between \mathbf{Y} and \mathbf{Z} .

One thing to observe before we start is that if $i + k \in A$, then $I(i, k) = 0$ deterministically for $(i, k) \in \mathcal{I}$. The pairs (i, k) where this holds are irrelevant when we bound the distance between \mathbf{I} and \mathbf{Y} , as the corresponding term in \mathbf{Y} is also deterministically zero. Thus we can ignore these terms in Steps 1 and 2 by removing them from \mathbf{I} and \mathbf{Y} . Let $\mathcal{I}' = \{(i, k) \in \mathcal{I} \mid i + k \notin A\}$. In a slight abuse of notation, we take \mathbf{I} and \mathbf{Y} to be indexed by \mathcal{I}' rather than by \mathcal{I} in Steps 1 and 2 only.

Step 1. Constructing the coupling.

Fix some $(i_0, k_0) \in \mathcal{I}'$. Our goal is to construct $(J(i, k))_{(i, k) \in \mathcal{I}'}$ distributed as \mathbf{I} conditioned on $I(i_0, k_0) = 1$ and to partition \mathcal{I} so that (3.1)–(3.2) hold. (In the notation used in the coupling condition, $J(i, k)$ would be written $J_{(i, k)(i_0, k_0)}$. We omit mention of (i_0, k_0) to simplify notation.)

Let τ be the random swap $(\pi(i_0), i_0 + k_0)$, and let $\pi' = \tau \circ \pi$. This forces π' to map i_0 to $i_0 + k_0$, making i_0 a k_0 -separated point for π' . As $\pi(i_0)$ cannot be an element of A (because $i_0 > r$) and $i_0 + k_0$ is not in A by definition of \mathcal{I}' , the permutation π' also satisfies (3.4).

We show now that π' is distributed as π conditioned on mapping i_0 to $i_0 + k_0$. Let Π be the set of permutations on n elements satisfying the conditions specified in (3.4), and let $\Pi' \subseteq \Pi$ be the set of permutations that also map i_0 to $i_0 + k_0$. One can easily check that for any $\sigma' \in \Pi'$, there are exactly $n - r$ permutations $\sigma \in \Pi$ such that swapping $\sigma(i_0)$ and $i_0 + k_0$ yields σ' . As π is distributed uniformly over Π , this implies that π' is distributed uniformly over Π' , which shows that π' is distributed as π conditioned on mapping i_0 to $i_0 + k_0$.

Now, for $(i, k) \in \mathcal{I}'$ we define

$$J(i, k) = \mathbf{1}\{i \text{ is } k\text{-separated in } \pi'\},$$

and the distribution of $(J(i, k))_{(i, k) \in \mathcal{I}'}$ is as we wanted. To partition \mathcal{I}' , we define \mathcal{I}_{i_0, k_0}^- as all pairs $(i, k) \in \mathcal{I}'$ where either

- (i) $i = i_0$ and $k \neq k_0$; or
- (ii) $i = i_0 + k_0 - k$ and $k \neq k_0$.

Define \mathcal{I}_{i_0, k_0}^+ to be the rest of \mathcal{I}' except for (i_0, k_0) . For any $(i, k) \in \mathcal{I}_{i_0, k_0}^-$, it is impossible that $\pi'(i) = i + k$: If (i) holds, then we already know that $\pi'(i) = i + k_0$, since we have conditioned π' to make this so. If (ii) holds, then we already know that $\pi'(i_0) = i_0 + k_0$, and so it cannot be that $\pi'(i) = i + k = i_0 + k_0$. Thus, for $(i, k) \in \mathcal{I}_{i_0, k_0}^-$, we have $J(i, k) = 0$, and the coupling satisfies (3.1).

To see that (3.2) is satisfied, suppose that $J(i, k) = 0$ and $I(i, k) = 1$ for some $(i, k) \in \mathcal{I}'$. We will show that $(i, k) \in \mathcal{I}_{i_0, k_0}^-$. By our assumption, $\pi(i) = i + k$, and $\pi'(i) \neq i + k$. Thus τ swaps $i + k$ with some other value. By the definition of τ , either $i + k = \pi(i_0)$, or $i + k = i_0 + k_0$. In the first case, we have $\pi(i_0) = \pi(i)$, implying that $i_0 = i$; thus (i, k) satisfies (i). In the second case, we have $i = i_0 + k_0 - k$, satisfying (ii). Thus our coupling satisfies (3.2).

Step 2. Bounding $d_{TV}(\mathbf{I}, \mathbf{Y})$.

The conditions of Proposition 3.2 are now satisfied, and we just need to bound the three terms on the right hand side of (3.3). We start with the observation that π can be thought of as a uniformly random bijection from $\{r + 1, \dots, n\}$ to $\{1, \dots, n\} \setminus A$. Thus, for any $(i, k) \in \mathcal{I}'$, the probability that π maps i to $i + k$ is $1/(n - r)$. Equivalently, $\mathbf{E}I(i, k) = 1/(n - r)$.

Now, we bound the first term in (3.3). As $|\mathcal{I}'| \leq 2r(n-r)$, we have

$$\sum_{(i,k) \in \mathcal{I}'} (\mathbf{E}I(i,k))^2 \leq \frac{2r}{n-r}. \tag{3.5}$$

For the next term, observe that if $(i,k) \in \mathcal{I}_{i_0,k_0}^-$, then $I(i,k)$ and $I(i_0,k_0)$ cannot simultaneously be 1. Thus

$$\mathbf{Cov}(I(i,k), I(i_0,k_0)) = -\mathbf{E}I(i,k)\mathbf{E}I(i_0,k_0) = -1/(n-r)^2.$$

For any (i_0,k_0) , the number of pairs (i,k) satisfying (i) is at most $2r$, and the number satisfying (ii) is at most $2r$. Thus $|\mathcal{I}_{i_0,k_0}^-| \leq 4r$, and

$$\sum_{(i_0,k_0) \in \mathcal{I}'} \sum_{(i,k) \in \mathcal{I}_{i_0,k_0}^-} |\mathbf{Cov}(I(i,k), I(i_0,k_0))| \leq \frac{2r(n-r)(4r)}{(n-r)^2} = \frac{8r^2}{n-r}. \tag{3.6}$$

For the final term, suppose that $(i,k) \in \mathcal{I}_{i_0,k_0}^+$. As π conditioned on $I_{i_0,k_0} = 1$ is a uniformly random element of the set Π' from Step 1, we have $\mathbf{E}[I_{i,k} \mid I_{i_0,k_0} = 1] = 1/(n-r-1)$. Thus

$$\begin{aligned} \mathbf{Cov}(I(i,k), I(i_0,k_0)) &= \mathbf{E}[I(i,k)I(i_0,k_0)] - \frac{1}{(n-r)^2} \\ &= \frac{1}{n-r} \mathbf{E}[I_{i,k} \mid I_{i_0,k_0} = 1] - \frac{1}{(n-r)^2} \\ &= \frac{1}{(n-r)(n-r-1)} - \frac{1}{(n-r)^2} = \frac{1}{(n-r)^2(n-r-1)}. \end{aligned}$$

Using the bound $|\mathcal{I}_{i_0,k_0}^+| \leq |\mathcal{I}'| \leq 2r(n-r)$,

$$\sum_{(i_0,k_0) \in \mathcal{I}'} \sum_{(i,k) \in \mathcal{I}_{i_0,k_0}^+} \mathbf{Cov}(I(i,k), I(i_0,k_0)) \leq \frac{(2r(n-r))^2}{(n-r)^2(n-r-1)} = \frac{4r^2}{n-r-1}. \tag{3.7}$$

Summing (3.5)–(3.7) and applying Proposition 3.2 shows that

$$d_{TV}(\mathbf{I}, \mathbf{Y}) \leq \frac{12r^2 + 2r}{n-r-1}. \tag{3.8}$$

Step 3. Bounding $d_{TV}(\mathbf{Y}, \mathbf{Z})$.

If $(i,k) \in \mathcal{I} \setminus \mathcal{I}'$, then $\mathbf{E}Y(i,k) = 0$. If $(i,k) \in \mathcal{I}'$, then $\mathbf{E}Y(i,k) = 1/(n-r)$. Applying Lemma 3.3 and using the bounds $|\mathcal{I} \setminus \mathcal{I}'| \leq 2r^2$ and $|\mathcal{I}'| \leq 2r(n-r)$,

$$d_{TV}(\mathbf{Y}, \mathbf{Z}) \leq 2r^2 \left(\frac{1}{n} - 0 \right) + 2r(n-r) \left(\frac{1}{n-r} - \frac{1}{n} \right) = \frac{4r^2}{n}.$$

Summing this bound and the one in (3.8) proves the theorem. □

Proposition 3.5. *Let π_n be a uniform permutation on n elements. Let $\xi_k^{(n)}$ be the point process on $[0, 1]$ with atoms at i/n for every position i that is k -separated in π_n , and let $(\xi_i)_{i \in \mathbb{Z}}$ be independent Poisson point processes on $[0, 1]$ with unit intensity. Let U_1, U_2, \dots be distributed uniformly on $[0, 1]$, and let them be independent of each other and of $(\xi_i)_{i \in \mathbb{Z}}$. For any fixed r ,*

$$\left(\frac{\pi_n(1)}{n}, \dots, \frac{\pi_n(r)}{n}, \xi_{-r+1}^{(n)}, \dots, \xi_{r-1}^{(n)} \right) \xrightarrow{d} (U_1, \dots, U_r, \xi_{-r+1}, \dots, \xi_{r-1}) \tag{3.9}$$

as $n \rightarrow \infty$.

Proof. As a first step, let $\bar{\xi}_k^{(n)}$ be the point process obtained by restricting $\xi_k^{(n)}$ to the interval $((r+1)/n, 1]$. For any i ,

$$\mathbf{P}[\pi_n(i) \in \{i-r+1, \dots, i+r-1\}] \leq \frac{2r}{n}.$$

By a union bound, the probability that any of $1, \dots, r$ is k -separated in π_n for some $-r < k < r$ is at most $2r^2/n$. Thus

$$\left(\frac{\pi_n(1)}{n}, \dots, \frac{\pi_n(r)}{n}, \xi_{-r+1}^{(n)}, \dots, \xi_{r-1}^{(n)} \right) = \left(\frac{\pi_n(1)}{n}, \dots, \frac{\pi_n(r)}{n}, \bar{\xi}_{-r+1}^{(n)}, \dots, \bar{\xi}_{r-1}^{(n)} \right)$$

except with probability at most $2r^2/n$. Since this probability vanishes as $n \rightarrow \infty$, it suffices to show the convergence in distribution of the right hand side of the above equation to the limit in (3.9).

Now, observe that

$$\left(\frac{\pi_n(1)}{n}, \dots, \frac{\pi_n(r)}{n} \right) \xrightarrow{d} (U_1, \dots, U_r) \tag{3.10}$$

by directly calculating

$$\lim_{n \rightarrow \infty} \mathbf{P} \left[\left(\frac{\pi_n(1)}{n}, \dots, \frac{\pi_n(r)}{n} \right) \in E_1 \times \dots \times E_r \right]$$

for any intervals E_1, \dots, E_r . To finish off the theorem, we will show that the law of $(\bar{\xi}_{-r+1}^{(n)}, \dots, \bar{\xi}_{r-1}^{(n)})$ conditional on $\pi_n(1), \dots, \pi_n(r)$ converges weakly to the law of $(\xi_{-r+1}, \dots, \xi_{r-1})$.

Observe that $(\bar{\xi}_{-r+1}^{(n)}, \dots, \bar{\xi}_{r-1}^{(n)})$ is a function of \mathbf{I} , with its entries given by

$$\bar{\xi}_k^{(n)} = \sum_{i=r+1}^n I(i, k) \delta_{i/n}.$$

Let $\zeta_k^{(n)} = \sum_{i=r+1}^n Z(i, k) \delta_{i/n}$, where $(Z(i, k))_{(i,k) \in \mathcal{I}}$ has entries independent and distributed as $\text{Poi}(1/n)$. Since $d_{TV}(f(X), f(Y)) \leq d_{TV}(X, Y)$ for any random variables X and Y and measurable map f , Proposition 3.4 implies that the total variation distance between $(\bar{\xi}_{-r+1}^{(n)}, \dots, \bar{\xi}_{r-1}^{(n)})$ conditional on $\pi_n(1), \dots, \pi_n(r)$ and $(\zeta_{-r+1}^{(n)}, \dots, \zeta_{r-1}^{(n)})$ vanishes as $n \rightarrow \infty$.

Thus we only need to show that $(\zeta_{-r+1}^{(n)}, \dots, \zeta_{r-1}^{(n)})$, a collection of discretized Poisson point processes, converges in distribution to the continuous Poisson point processes $(\xi_{-r+1}, \dots, \xi_{r-1})$. By the independence of entries of both of these vectors, we only need to show that $\zeta_k^{(n)} \xrightarrow{d} \xi_k$ as $n \rightarrow \infty$. To prove this, by Theorem 16.16 and Chapter 16, Exercise 11 in [10], it suffices to show that $\zeta_k^{(n)}(B) \xrightarrow{d} \xi_k(B)$ for any finite union of intervals B . Both $\zeta_k^{(n)}(B)$ and $\xi_k(B)$ are Poisson, and so the convergence holds as the portion of points of the form i/n with $i \geq r+1$ contained in B approaches the length of B . \square

Let us consider the shift operation from the perspective of the point processes $(\xi_k^{(n)})_{k \in \mathbb{Z}}$ of Proposition 3.5. Suppose that π_n has a fixed point at position i and we bump it to the beginning of the permutation, moving forward in the tree. At times after i/n , all point processes $\xi_k^{(n)}$ are unchanged. At times before i/n , every point in $\xi_k^{(n)}$ becomes a point in $\xi_{k-1}^{(n)}$ shifted to the right by $1/n$. Additionally, $\xi_0^{(n)}$ loses its point at i/n , and $\xi_{i-1}^{(n)}$ gains a point at $1/n$. This completely describes how the point processes change when this fixed point is shifted.

Theorem 3.6. *As $n \rightarrow \infty$, the randomly rooted graph F_n converges to T in the local weak sense.*

Proof. Fix some integer $r \geq 1$, and assume throughout that $n > r$. We need to show that the r -ball around π_n in F_n converges in distribution to the r -ball around the root in T , which is

$$\varphi_r(U_1, \dots, U_r; \xi_{-r+1}, \dots, \xi_{r-1}).$$

Let

$$T_n = \varphi_r\left(\frac{\pi_n(1)}{n}, \dots, \frac{\pi_n(r)}{n}; \xi_{-r+1}^{(n)}, \dots, \xi_{r-1}^{(n)}\right).$$

The idea will be to show that this tree T_n is almost the same thing as the r -neighborhood of π_n in F_n , and then to apply Proposition 3.5 to show that T_n converges in distribution to the r -ball around the root in T .

Claim 3.7. *If $\xi_{-r+1}^{(n)}, \dots, \xi_{r-1}^{(n)}$ contain no points in the interval $[0, r/n]$, then T_n is identical to the r -ball around π_n in F_n .*

Proof. The algorithm defining T_n differs from the true r -ball around π_n in F_n in two ways. First, when moving forward in the tree by bumping a fixed point at position i to position 1, no point at $1/n$ is inserted into the point process $\xi_{i-1}^{(n)}$. For this to cause T_n to lack a vertex in the r -ball around π_n , it must occur that after $s \geq 0$ steps backward in the tree from π_n , there is a fixed point at i , and then after bumping it one can move forward in the tree another $i - 1$ times to return the fixed point to i , and then one can bump the fixed point again, all while remaining within the r -ball around π_n . Thus, it is necessary that $s + i < r$. Under the conditions of this claim, $\pi_n(1), \dots, \pi_n(r) > r$, since otherwise π_n would have a k -separated point at position i for some $-r + 1 \leq k \leq r - 1$ and $1 \leq i \leq r$. Hence, one must move backwards from π_n at least $r - i + 1$ times to create a fixed point at i , showing that this circumstance can never occur.

The second difference in the algorithm is that points in the point processes are not shifted by $1/n$ at each step. One could equally well define the algorithm giving T_n only in terms of the order of the points in $\xi_{-r+1}^{(n)}, \dots, \xi_{r-1}^{(n)}$, and doing this one sees that the shifting of points by $1/n$ would not change T_n . \square

As $n \rightarrow \infty$, the condition in this claim holds with probability approaching 1, as we showed with a union bound in the proof of Proposition 3.5. Thus we only need to show that T_n converges in distribution to the r -ball around the root in T , that is, that

$$\varphi_r\left(\frac{\pi_n(1)}{n}, \dots, \frac{\pi_n(r)}{n}; \xi_{-r+1}^{(n)}, \dots, \xi_{r-1}^{(n)}\right) \xrightarrow{d} \varphi_r(U_1, \dots, U_r; \xi_{-r+1}, \dots, \xi_{r-1}).$$

Once we show that φ_r is continuous on a set that almost surely contains

$$(U_1, \dots, U_r; \xi_{-r+1}, \dots, \xi_{r-1}),$$

Proposition 3.5 and the continuous mapping theorem [5, Theorem 2.7] immediately give us this result.

We claim that φ is continuous at any point $(u_1, \dots, u_r; \zeta_{-r+1}, \dots, \zeta_{r-1})$ where all atoms of the point process configurations $\zeta_{-r+1}, \dots, \zeta_{r-1}$ are distinct, all $u_1, \dots, u_r \in [0, 1]$ are distinct from each other and any points in the configurations. In fact, φ_r is constant on a neighborhood of $(u_1, \dots, u_r; \zeta_{-r+1}, \dots, \zeta_{r-1})$, as a slight perturbation that does not change the order of any of the points u_1, \dots, u_r or the points in $\zeta_{-r+1}, \dots, \zeta_{r-1}$ does not change the resulting tree. As $(U_1, \dots, U_r; \xi_{-r+1}, \dots, \xi_{r-1})$ has these properties almost surely, this proves the continuity property of φ_r that we needed. \square

As a consequence of Theorem 3.6, any statistic of π_n determined by its r -neighborhood in F_n converges in distribution to the corresponding statistic of the limit tree T . This excludes many statistics like the distance from π_n to the nearest leaf of F_n , which is nearly local but cannot be deduced from the r -neighborhood of π_n for any fixed value of r . The following corollary addresses this by a truncation argument.

Corollary 3.8. *Let f be a function defined on rooted trees, and suppose that $\min(f(T_0), r)$ is determined by the r -neighborhood of the root of T_0 , for any rooted tree T_0 and any $r \geq 0$. Then $f(F_n) \xrightarrow{d} f(T)$.*

Proof. By assumption, we have $\min(f(F_n), r) = \gamma(F_n(r))$, where $F_n(r)$ denotes the r -ball around the root of F_n and γ is some deterministic function on rooted trees. By Theorem 3.6, as $n \rightarrow \infty$ we have $F_n(r) \xrightarrow{d} T(r)$ with respect to the discrete topology on rooted trees. Hence, by the continuous mapping theorem, $\gamma(F_n(r)) \xrightarrow{d} \gamma(T(r))$, and in our original notation, $\min(f(F_n), r) \xrightarrow{d} \min(f(T), r)$. As $\mathbf{P}[X \leq x] = \mathbf{P}[\min(X, r) \leq x]$ for $x < r$, it holds for any x where the distribution function $\mathbf{P}[f(T) \leq x]$ is continuous that

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbf{P}[f(F_n) \leq x] &= \lim_{n \rightarrow \infty} \mathbf{P}[\min(f(F_n), x + 1) \leq x] \\ &= \mathbf{P}[\min(f(T), x + 1) \leq x] = \mathbf{P}[f(T) \leq x]. \end{aligned} \quad \square$$

4 Combinatorics of paths to leaves

In the next section, we will be interested in the limiting distributions of the shortest and longest distance of a random permutation to a leaf. In preparation for this, we consider the combinatorics related to the shortest and longest path to a leaf in this section.

For a permutation $\pi \in \mathfrak{S}_n$, let $T(\pi)$ be the fixed point tree with π as root. We call i a **true fixed point** of π if $\pi(i) = i$ and $i \neq 1$.

Proposition 4.1. *Given $\pi \in \mathfrak{S}_n$, a shortest path from π to a leaf in $T(\pi)$ is obtained by bumping the rightmost true fixed point at each step.*

Remark 4.2. Note that the procedure of picking the rightmost true fixed point in $\pi \in \mathfrak{S}_n$ and then bumping the letter can be reformulated in the following way. Replace each letter in the one-line notation for π by its k -separation. Then scan the word from right to left and successively pick the first $k = 0, 1, 2, \dots$ until no further subsequent bigger k can be picked to the left. Picking 0 in position 1 is not allowed. The letters in these positions in the original π are the letters that are bumped.

For example, take $\pi = 32415$. The k -separation word is given by 20130 (where as in Figure 3 a negative separation is indicated in red). The selected positions are underlined 20130, meaning that the letters 5, 4, 3 are bumped in this order:

$$32415 \rightarrow 53241 \rightarrow 45321 \rightarrow 34521. \tag{4.1}$$

Proof of Proposition 4.1. Let $v = (v_0, v_1, \dots, v_\ell)$ with $v_0 = \pi$ and v_ℓ a leaf be a sequence of vertices in which at some step, the rightmost true fixed point is not bumped. Let v_k be the last vertex (meaning k is maximal) where anything other than the rightmost true fixed point is bumped in the step $v_k \rightarrow v_{k+1}$. Denote by $v' = (v_k = v'_k, v'_{k+1}, \dots, v'_{\ell'})$ the path in which the rightmost true fixed point is always bumped. We will show that $\ell' \leq \ell$. Then moving from right to left replacing all non-rightmost bumps produces a (weakly) shorter path than v , which shows that the path where the rightmost true fixed point is always bumped is a shortest path.

Let f_1, \dots, f_a be the true fixed points in $v_k = v'_k$ ordered from left to right and suppose that f_i with $1 \leq i < a$ is the fixed point bumped to reach v_{k+1} . Since k is maximal

such that $v_k \rightarrow v_{k+1}$ does not bump the rightmost true fixed point, both (v_{k+1}, \dots, v_ℓ) and $(v'_k, \dots, v'_{\ell'})$ follow the algorithm outlined in Remark 4.2 starting at v_{k+1} and v'_k , respectively. Note that by Remark 4.2, the length of the path formed by always choosing the rightmost true fixed point can be read off directly from the first vertex (in our case v_{k+1} and v'_k , respectively) by determining the longest consecutive sequence $0, 1, 2, \dots$ from right to left in the corresponding separation word. Now the separation word of v_{k+1} does not differ from that of v'_k in the positions to the right of the position of f_i . In the positions to the left of the position of f_i , the separations of v_{k+1} are one less than those of v'_k and shifted one position to the right. In addition there is a new separation of value $i - 1$ in position 1. Hence the chosen $0, 1, 2, \dots$ sequence in v_{k+1} is at most one shorter than that in v'_k . This proves that $\ell' \leq \ell$. \square

Proposition 4.3. *Given $\pi \in \mathfrak{S}_n$, the longest path from π to a leaf in $T(\pi)$ is obtained by bumping the leftmost true fixed point at each step.*

Proof. Let $v = (v_0, v_1, \dots, v_\ell)$ with $v_0 = \pi$ and v_ℓ a leaf be a sequence of vertices in which at some step, the leftmost true fixed point is not bumped. We will show the existence of a strictly longer path to a leaf. Let v_k be the last vertex such that from $v_k \rightarrow v_{k+1}$ anything but the leftmost true fixed point is bumped. Denote by f_1, \dots, f_a the true fixed points in v_k ordered from left to right and suppose that f_i with $i > 1$ is the fixed point bumped in the transition $v_k \rightarrow v_{k+1}$.

We are going to construct a new path to a leaf $v' = (v'_0, \dots, v'_{\ell'})$ as follows. First, $v'_i = v_i$ for all $0 \leq i \leq k$. In the step $v'_k \rightarrow v'_{k+1}$ the true fixed point f_{i-1} is bumped. Except for the letters f_{i-1} versus f_i , v'_{k+1} and v_{k+1} are the same on the first f_{i-1} positions. Hence we can bump the same letters in v and v' until a letter beyond position f_{i-1} is bumped in $v_h \rightarrow v_{h+1}$ (with $h > k$). At this point, the letter f_i is bumped in $v'_h \rightarrow v'_{h+1}$, which is still a true fixed point (no letters to the right of position f_{i-1} have changed position up to this point in the bumping process). The letters in positions weakly between $f_{i-1} + 2$ and f_i agree in v_h and v'_{h+1} . The letters different from f_{i-1} and f_i weakly before position $f_{i-1} + 1$ are shifted one position to the right in v'_{h+1} compared to v_h . Denote the sequence of bumped letters in $v_h \rightarrow v_{h+1} \rightarrow \dots \rightarrow v_\ell$ by $R_1 L_1 R_2 L_2 \dots$, where L_i are letters in positions weakly before position $f_{i-1} + 1$ and R_i are letters after position $f_{i-1} + 1$. Write $R_i = R'_i r_i$, where r_i is a single letter. Then the sequence $R'_1 L_1 r_1 R'_2 L_2 r_2 \dots$ is a bumpable sequence on v'_{h+1} due to the relative shift of letter in the left $f_{i-1} + 1$ positions. This indeed shows that v' is strictly longer than v .

Hence the path given by always bumping the leftmost true fixed point is the unique longest path to a leaf. \square

Note that unlike for the longest path to a leaf, the shortest path is not unique.

Example 4.4. Take $\pi = 32415$. Then the longest path to a leaf is of length 9:

$$32415 \rightarrow 23415 \rightarrow 52341 \rightarrow 25341 \rightarrow 32541 \rightarrow 23541 \rightarrow 42351 \\ \rightarrow 24351 \rightarrow 32451 \rightarrow 23451.$$

The shortest path to a leaf given by always bumping the rightmost fixed point is given in (4.1). Here is another shortest path:

$$32415 \rightarrow 23415 \rightarrow 52341 \rightarrow 45231.$$

Denote by $\mathcal{B} := \mathcal{B}(\pi)$ the set of *bumped values* in the longest path from π to a leaf. For example, $\mathcal{B}(32415) = \{2, 3, 4, 5\}$ as can be seen from Example 4.4. We now provide a characterization of the set \mathcal{B} that enables us to determine if $\pi(i) \in \mathcal{B}$ given only knowledge of $\pi(i + 1), \dots, \pi(n)$.

Lemma 4.5. For $\pi \in \mathfrak{S}_n$ and $1 \leq i \leq n$,

$$\pi(i) \in \mathcal{B} \iff \pi(i) \neq 1 \text{ and } 0 \leq \pi(i) - i \leq \#\{j > i \mid \pi(j) \in \mathcal{B}\}. \quad (4.2)$$

Proof. Let $\pi(i) \in \mathcal{B}$. First we show that then the conditions on the right of (4.2) hold. Since only true fixed points can be bumped, we certainly have $\pi(i) \neq 1$. Note that $\pi(i) - i$ is the i -separation of $\pi(i)$ at position i . In the longest path to a leaf, a letter in one-line notation either stays in its position or moves to the right unless it is bumped. Hence the separation of a letter either remains the same or becomes smaller (unless it is bumped). This implies that a letter with a negative separation can never be bumped, which shows that $0 \leq \pi(i) - i$. Now suppose that $\pi(i) - i = k$, meaning that $\pi(i)$ is k -separated. Since $\pi(i) \in \mathcal{B}$ and hence must be bumped in the path to the leaf, there must be at least k letters to the right of $\pi(i)$ that are bumped to move $\pi(i)$ k positions to the right and make it 0-separated. This is precisely the condition $\pi(i) - i \leq \#\{j > i \mid \pi(j) \in \mathcal{B}\}$.

Conversely, suppose that $\pi(i)$ satisfies the conditions on the right of (4.2) and set $\pi(i) - i = k$. By Proposition 4.3, the longest path is obtained by always bumping the leftmost true fixed point, and hence any point that becomes 0-separated will eventually be bumped. Since by assumption $0 \leq k \leq \#\{j > i \mid \pi(j) \in \mathcal{B}\}$, there are at least k points to the right of $\pi(i)$ that get bumped, which makes $\pi(i)$ eventually 0-separated, and hence $\pi(i) \in \mathcal{B}$. \square

Using the above ideas, we can give a bound on the length of the longest path to a leaf, which will be useful in Section 5. For $\pi \in \mathfrak{S}_n$, we denote by $\ell(\pi)$ the length of the longest path from π to a descendent leaf.

Lemma 4.6. Fix $\pi \in \mathfrak{S}_n$ and let $\mathcal{B} := \mathcal{B}(\pi) = \{b_1, b_2, \dots, b_k\}$, where $k = |\mathcal{B}|$ and $2 \leq b_1 < b_2 < \dots < b_k \leq n$. Then

$$\ell(\pi) \leq 1 + \sum_{m=1}^{k-1} \prod_{i=1}^m \left(1 + \frac{1}{b_i - i}\right) \leq k \prod_{i=1}^{k-1} \left(1 + \frac{1}{b_i - i}\right). \quad (4.3)$$

Proof. Let N_i denote the number of times b_i is bumped in the longest path in $T(\pi)$ and let $M_i = \sum_{j \geq i} N_j$. With this notation $\ell(\pi) = M_1$.

Every time the letter b_j is bumped to the beginning, it needs to be moved to the right one position at a time by $b_j - 1$ bumps of letters to the right of b_j . The position of b_j can only be increased once by bumps of smaller letters $b_i < b_j$. At least $b_j - j$ letters greater than b_j must be bumped for the value b_j to return to position b_j and be eligible to be bumped again. Therefore $N_j \leq 1 + \frac{1}{b_j - j} M_{j+1}$, and hence

$$M_j \leq 1 + \left(1 + \frac{1}{b_j - j}\right) M_{j+1}.$$

By inductively applying this inequality starting from $M_k \leq 1$, we obtain

$$M_j \leq 1 + \sum_{m=j}^{k-1} \prod_{i=j}^m \left(1 + \frac{1}{b_i - i}\right)$$

for $j = 1, \dots, k$. Thus

$$\begin{aligned} \ell(\pi) = M_1 &\leq 1 + \sum_{m=1}^{k-1} \prod_{i=1}^m \left(1 + \frac{1}{b_i - i}\right) \\ &\leq 1 + (k-1) \prod_{i=1}^{k-1} \left(1 + \frac{1}{b_i - i}\right) \leq k \prod_{i=1}^{k-1} \left(1 + \frac{1}{b_i - i}\right). \end{aligned} \quad \square$$

The first bound in (4.3) is sharp, giving the correct length $2^{n-1} - 1$ for the identity permutation of length n .

5 Limiting distributions and higher moments

In this section we compute the limiting distributions of two statistics of the fixed point tree F_n using the local weak limit we have constructed in Section 3. As usual, π_n is a uniformly random permutation of length n . We study two statistics related to leaves:

1. the distance M_n from π_n to the nearest leaf descending from it;
2. the distance L_n from π_n to the farthest leaf descending from it.

We use $\text{Geo}(q)$ to refer to the geometric distribution on $\{0, 1, \dots\}$ with parameter q , the number of failures before the first success in independent trials that are successful with probability q . This distribution places probability $(1 - q)^k q$ on k and has mean $(1 - q)/q$.

Theorem 5.1.

- (i) Let $M \sim \text{Poi}(1)$. As $n \rightarrow \infty$, we have $M_n \xrightarrow{d} M$ and $\mathbf{E}M_n^p \rightarrow \mathbf{E}M^p$ for all $p > 0$.
- (ii) Let $L \sim \text{Geo}(e^{-1})$. As $n \rightarrow \infty$, we have $L_n \xrightarrow{d} L$ and $\mathbf{E}L_n^p \rightarrow \mathbf{E}L^p$ for all $p > 0$.

Proof. We will prove that $M_n \xrightarrow{d} M$ and $L_n \xrightarrow{d} L$ in Propositions 5.3 and 5.5, respectively. In Proposition 5.11, we will show that $\sup_n \mathbf{E}M_n^p < \infty$ and $\sup_n \mathbf{E}L_n^p < \infty$ for any $p > 0$. It is a standard result that this proves the convergence of all moments (see [7, Theorem 4.5.2], for instance). \square

An interesting open problem is to determine the limiting distribution of the number of steps from π_n to a leaf when moving randomly towards a leaf.

Sections 5.1 and 5.2 are devoted to proving the convergence of M_n and L_n to their limiting distributions. Both M_n and L_n are functionals of F_n satisfying the criteria of Corollary 3.8. This corollary then proves that M_n and L_n converge in distribution to the corresponding functionals of the limit tree. Thus, our task in these sections is to work out the distributions of the distances in T from the root to the nearest and farthest leaves descending from it.

Section 5.3 gives our proof that M_n and L_n are bounded in L^p , establishing the convergence of moments. We emphasize that this result is not just a technicality. The reentry phenomenon mentioned on Page 4 can cause L_n to be extremely large. For example, if π_n begins with the string $1 \cdots k$, then $L_n \geq 2^{k-1}$, with the same small letters bumped repeatedly. The probability of such reentry is vanishingly small, making it irrelevant to the distributional convergence of L_n . But these unlikely events nonetheless contribute to the moments of L_n , and a priori it is plausible that even the expectation of L_n tends to infinity. To prove that this is not the case, we are forced to develop several bounds that should be useful in future work on the fixed point forest.

5.1 Shortest distance to a leaf

Our first lemma is the limiting analogue of Proposition 4.1.

Lemma 5.2. *A shortest path from the root of T to a descendant leaf is obtained by bumping the rightmost abstracted fixed point at each step.*

Proof. The statement can be proved in an analogous fashion to Proposition 4.1. \square

Proposition 5.3. *The distance $M_n \xrightarrow{d} \text{Poi}(1)$ as $n \rightarrow \infty$.*

Proof. Let M be the distance in the random tree T from the root to its nearest descendent leaf. As we mentioned in our summary of Section 5, Corollary 3.8 shows that $M_n \rightarrow M$ in distribution, and thus we need only show that $M \sim \text{Poi}(1)$. By Lemma 5.2, M is the

number of steps taken if we start at the root of T and bump the rightmost fixed point until no fixed points remain.

Walking towards a leaf in the tree in this way can be viewed as follows. Recall that ρ is the root of T . If ξ_0^ρ has no points, then the walk is over, and $M = 0$. Otherwise, let X_1 be the rightmost point of ξ_0^ρ , and let v_1 be the child of ρ corresponding to bumping it. Now, if $\xi_0^{v_1}$ has no points, then the walk is over and $M = 1$. Otherwise, let X_2 be its rightmost point of $\xi_0^{v_1}$, which is necessarily to the left of X_1 , since $\xi_0^{v_1} = \xi_1^\rho|_{[0, X_1)} + \xi_0^\rho|_{(X_1, 0]}$. Continue in this way, making a sequence of vertices ρ, v_1, \dots, v_M and a sequence of points corresponding to them, X_1, \dots, X_M .

In this procedure, X_1 is the rightmost point of ξ_0^ρ , then X_2 is the rightmost point of $\xi_1^\rho|_{[0, X_1)}$, then X_3 is the rightmost point of $\xi_2^\rho|_{[0, X_2)}$, and so on. We can interpret this as follows: We start at 1, moving backwards in a Poisson process until we encounter a point, which takes $\text{Exp}(1)$ time to arrive. Then, looking at a different Poisson process, we move backwards until we encounter a point, which again will arrive in $\text{Exp}(1)$ time, independent of the first arrival. We then continue on in this way, and M is the total number of points encountered before we make it backwards to time 0. This is the same as counting the number of arrivals in a *single* Poisson process between time 0 and 1, which is distributed as $\text{Poi}(1)$. \square

5.2 Farthest distance to a leaf

The next lemma is the limiting analogue of Proposition 4.3.

Lemma 5.4. *The unique longest path from the root of T to a descendent leaf is obtained by bumping the leftmost abstracted fixed point at each step.*

Proof. Suppose we have some sequence of points bumped in a path from the root of T to a descendent leaf in which some point y was not the leftmost when it was bumped. We will show the existence of a strictly longer path to a leaf. Let $x < y$ be a point that could have been bumped instead of y . Decompose the original sequence of bumped points as

$$PyL_1a_1L_2a_2 \cdots a_{r-1}L_r. \tag{5.1}$$

Capital letters in (5.1) denote (possibly empty) strings of points and lowercase letters denote single points. The string P consists of all bumped points prior to y . Next, L_1 is made up of points smaller than x , and a_1 is the first point larger than x . Then L_2 consists of points smaller than x , and a_2 is the next subsequent point larger than x , and so on. Note that almost surely, no point of $\xi_0^\rho, \xi_1^\rho, \dots$ occurs more than once, so we never need to worry about repeated values in these sequences.

We claim that we can instead bump the following sequence of points:

$$PxL_1yL_2a_1L_3a_2 \cdots a_{r-2}L_r a_{r-1}. \tag{5.2}$$

As this sequence is one longer than (5.1), this claim completes the proof. Thus, we just need to show that this sequence is in fact *bumpable*, by which we mean that when each point is bumped in turn, the next one is an abstracted fixed point. Before we do this, we make a key observation: Suppose that $z_1 \cdots z_k$ and $z'_1 \cdots z'_k$ are two bumpable sequences, each of which contains the same number of points larger than z . Then $z_1 \cdots z_k z$ is bumpable if and only if $z'_1 \cdots z'_k z$ is bumpable.

Clearly, Px is bumpable. Since all points in L_1 are smaller than x and $x < y$, as each point of L_1 is encountered in (5.1) and (5.2) the same number of larger points has already been encountered in each sequence. By our observation, the bumpability of PyL_1 implies the bumpability of PxL_1 . Next, since Py is bumpable and all points in xL_1 are smaller than y , the sequence PxL_1y is also bumpable. Repeating this reasoning,

bumpability of $PyL_1a_1L_2$ implies bumpability of PxL_1yL_2 , and bumpability of PyL_1a_1 implies bumpability of $PxL_1yL_2a_1$. Continuing in this way, we arrive at the bumpability of (5.2). \square

The following proof involves a continuous-time Markov chain known as a *Yule process*. At state k , it jumps to $k + 1$ at rate k , with no other transitions allowed. It is well known that if $(Y_t)_{t \geq 0}$ is a Yule process starting at 1, then $Y_t - 1 \sim \text{Geo}(e^{-t})$; see [11, Section 4.1.D], for example.

Proposition 5.5. *As $n \rightarrow \infty$, we have $L_n \xrightarrow{d} \text{Geo}(e^{-1})$.*

Proof. Let L be the length of the longest path in the limit tree T from the root to a descendent leaf. By Corollary 3.8, we have $L_n \xrightarrow{d} L$. By Lemma 5.4 the random variable L is the length of the path down T given by bumping the leftmost fixed point until none remain.

First, we sketch one way to compute the distribution of L . Suppose we start at the root ρ of T . With probability e^{-1} , there are no points in ξ_0^ρ , and $L = 0$. If not, we bump the leftmost point X of ξ_0^ρ to move to a vertex v . The next set of abstracted fixed points is ξ_0^v , which is equal to ξ_1^ρ on $[0, X)$ and to ξ_0^ρ on $(X, 1]$. This is a new Poisson point process independent of the past, and so again we stop with probability e^{-1} and $L = 1$. Continuing on in this way, L is the number of failures before the first success in independent trials that succeed with probability e^{-1} .

The above argument works and is the most direct, but we give a different proof now whose ideas will be useful in Section 5.3. Let $\mathcal{B} \subseteq [0, 1]$ be the set of points bumped in the longest path in T from the root to a descendent leaf. Lemma 4.5 can be restated in this setting: A point x of ξ_k^ρ is an element of \mathcal{B} if and only if $0 \leq k \leq |\mathcal{B} \cap (x, 1]|$. Thus, we can progressively build our set \mathcal{B} by the following procedure. Start with \mathcal{B} empty. Scan ξ_0^ρ from right to left starting at time 1 until a point is encountered, and add it to \mathcal{B} . Now, scan ξ_0^ρ and ξ_1^ρ from right to left from this point until a point is encountered, and add it to \mathcal{B} . Now scan ξ_0^ρ , ξ_1^ρ , and ξ_2^ρ , and so on, stopping when we reach time 0.

In this procedure, the first point arrives at rate 1, since it is the first arrival time of a unit intensity Poisson process. The next point arrives at rate 2, since it is the first arrival time of two independent unit intensity Poisson processes, and the next at rate 3, and so on. Thus, the size of \mathcal{B} is the number of increases of a Yule process in time 1, which is distributed as $\text{Geo}(e^{-1})$. \square

It follows from this proposition that the root of T has finitely many descendants almost surely. In forthcoming work, this will be investigated further, and it will be shown that the expected number of descendants is infinite.

5.3 Higher moments

To bound the moments of L_n and M_n , we must leave behind the limit tree T and work directly with the finite fixed point forest. Since $M_n \leq L_n$ and we are only looking for upper bounds, we will deal exclusively with L_n . Our first result gives an upper bound on L_n in terms of $\mathcal{B} = \mathcal{B}(\pi_n) = \{b_1, b_2, \dots, b_k\}$, the set of letters bumped in the longest path from π_n to a descendent leaf. Since $b_1 < b_2 < \dots < b_k$ is increasing, $b_i - i$ is weakly increasing. Fix some $x > 0$ and denote by \mathcal{B}_x the subset of \mathcal{B} such that $b_i - i < x$ for $b_i \in \mathcal{B}_x$.

Lemma 5.6. *Fix $x > 0$. Then*

$$L_n \leq 2^{|\mathcal{B}_x|} |\mathcal{B}| \left(1 + \frac{1}{x}\right)^{|\mathcal{B}|}. \tag{5.3}$$

Proof. For $b_i \in \mathcal{B}_x$, we have $1 + \frac{1}{b_i - i} \leq 2$ and for $b_i \in \mathcal{B} \setminus \mathcal{B}_x$, we have $1 + \frac{1}{b_i - i} < 1 + \frac{1}{x}$. With this, the bound in (5.3) follows directly from Lemma 4.6. \square

We will bound EL_n^p using (5.3) by showing that $|\mathcal{B}_x|$ and $|\mathcal{B}|$ are unlikely to be large. Our first result in this direction, interesting in its own right, is an exponential tail bound on $|\mathcal{B}|$. Proposition 5.5 exactly computes the distribution of the corresponding quantity in the limit case as $\text{Geo}(e^{-1})$. Our proof here follows the same intuition, but it will be considerably more difficult.

Proposition 5.7. *For some constant $C < 1$, it holds for all $k, n \geq 0$ that*

$$\mathbf{P}[|\mathcal{B}| \geq k] \leq C^k.$$

We will need several preliminary lemmas first. Recall from Lemma 4.5 that \mathcal{B} can be constructed by moving leftward through the random permutation π_n , successively revealing $\pi_n(n), \pi_n(n-1), \dots, \pi_n(1)$ and tracking the set \mathcal{B} as we go. Reversing the indexing so that we can count up instead of down, define

$$X_i = \#\{j > n - i \mid \pi_n(j) \in \mathcal{B}\}.$$

This yields a *pure birth process* (that is, one that either stays the same or increases by one at each step) starting at $X_0 = 0$ and ending at $X_n = |\mathcal{B}|$. Our goal is to prove an exponential tail bound for X_n that does not depend on n . In Theorem 5.5, we show that the analogous process for the limit tree is a Yule process. The process $(X_i)_{0 \leq i \leq n}$ is not as straightforward, but the following lemma shows that it approximates a Yule process in the sense that it increases with probability approximately proportional to its current value.

Lemma 5.8. *For $k < n/2$,*

$$\mathbf{P}[X_{i+1} = X_i + 1 \mid X_i = k] \leq \frac{1+k}{n-2k}. \tag{5.4}$$

Proof. For a given permutation $\sigma \in \mathfrak{S}_n$, let $x_i(\sigma) = \#\{j > n - i \mid \sigma(j) \in \mathcal{B}\}$, so that $X_i = x_i(\pi_n)$. Taking i and k to be fixed, define

$$\begin{aligned} \mathfrak{T}_\ell &= \{\sigma \in \mathfrak{S}_n \mid x_i(\sigma) = k, \sigma(n-i) = n-i+\ell\}, \\ \mathfrak{U} &= \{\sigma \in \mathfrak{S}_n \mid x_i(\sigma) = k\}. \end{aligned}$$

By Lemma 4.5, the process $x_i(\sigma)$ increases in its next step if and only if $0 \leq \sigma(n-i) - (n-i) \leq x_i(\sigma)$ and $\sigma(n-i) \neq 1$, showing that

$$\mathbf{P}[X_{i+1} = X_i + 1 \mid X_i = k] \leq \frac{\sum_{\ell=0}^k |\mathfrak{T}_\ell|}{|\mathfrak{U}|}. \tag{5.5}$$

Now, we compare the sizes of \mathfrak{T}_ℓ and \mathfrak{U} by a combinatorial switching argument. Suppose that $\sigma \in \mathfrak{T}_\ell$. For any $j \in [n]$, let $\sigma^j = (\sigma(n-i), \sigma(j)) \circ \sigma$, the permutation given by swapping the values at positions $n-i$ and j in σ . Given σ^j , i , ℓ , and n but without σ or j , we can recover σ by the formula $\sigma = (n-i+\ell, \sigma^j(n-i)) \circ \sigma^j$ and j by $j = (\sigma^j)^{-1}(n-i+\ell)$. This shows that the map from $\mathfrak{T}_\ell \times [n] \rightarrow \mathfrak{S}_n$ given by $(\sigma, j) \mapsto \sigma^j$ is injective.

We claim that for any $\sigma \in \mathfrak{T}_\ell$, the permutation σ^j falls in \mathfrak{U} if either of the following holds:

- (i) $j \leq n - i$;
- (ii) $j > n - i + \ell$ and $\sigma(j) \notin \mathcal{B}(\sigma)$.

Indeed, in the first case, σ and σ^j differ only at locations $n - i$ and smaller. Since $x_i(\sigma)$ is determined by $\sigma(n - i + 1), \dots, \sigma(n)$, we have $x_i(\sigma^j) = x_i(\sigma) = k$, and hence $\sigma^j \in \mathfrak{U}$. In the second case, the only way for $x_i(\sigma^j)$ to be different from $x_i(\sigma)$ is if $\sigma^j(j) \in \mathcal{B}$. But this cannot occur, since $\sigma^j(j) = \sigma(n - i) \leq n - i + \ell$ and $j > n - i + \ell$.

For $\sigma \in \mathfrak{T}_\ell$, there are $n - i$ choices of j satisfying (i). There are another $i - \ell$ choices of j satisfying $j > n - i + \ell$; at most $x_{i-\ell}(\sigma) \leq x_i(\sigma) = k$ of these have $\sigma(j) \in \mathcal{B}(\sigma)$, giving us at least $i - \ell - k$ choices of j satisfying (ii). Thus, for each $\sigma \in \mathfrak{T}_\ell$, there are at least $n - \ell - k$ choices of j for which $\sigma^j \in \mathfrak{U}$. By injectivity of the map $(\sigma, j) \mapsto \sigma^j$,

$$|\mathfrak{U}| \geq (n - \ell - k)|\mathfrak{T}_\ell| \geq (n - 2k)|\mathfrak{T}_\ell|$$

for $\ell \leq k$. Substituting this into (5.5) gives (5.4). □

Now, we couple X_i with a true Yule process, whose marginal distributions we know to be exactly geometrically distributed. The only complication is that we lose control of (X_i) if it becomes too large, which we deal with by considering it only up to a stopping time.

Lemma 5.9. *Let Y_t be a Yule process starting from 1. Let $S = \min\{i \mid X_i \geq n/4\}$, taking $S = n$ if the minimum is over the empty set. The processes $(X_i)_{0 \leq i \leq n}$ and $(Y_t)_{0 \leq t \leq 4}$ can be coupled so that*

$$1 + X_i \leq Y_{4i/n}$$

for all $i \leq S$.

Proof. Essentially, we just need to confirm that X_i is less likely to increase from time i to $i + 1$ than Y is from time $4i/n$ to $4(i + 1)/n$. Formally, we start with $(X_i, i \in [n])$ and then build up the Yule process inductively on the same probability space. Since a Yule process is Markov, we can construct (Y_t) up to some time t_0 and then extend it by attaching to its end a new Yule process starting at Y_{t_0} .

Assume for some $j \in \{0, \dots, n\}$ that we have already constructed a Yule process $(Y_t)_{0 \leq t \leq 4j/n}$ so that $1 + X_i \leq Y_{4i/n}$ for $i \leq \min(j, S)$. Note that this is trivial in the starting case $j = 0$. We want to show that if $S > j$, then we can extend (Y_t) up to time $4(j + 1)/n$ so that $1 + X_{j+1} \leq Y_{4(j+1)/n}$. (If $S \leq j$, then we can just extend (Y_t) up to time $4(j + 1)/n$ independently from X_{j+1} , since the relationship between the two is irrelevant.)

So, we assume $S > j$, which implies that $X_j < n/4$. According to Lemma 5.8,

$$\mathbf{P}[X_{j+1} = X_j + 1 \mid X_j] \leq \frac{1 + X_j}{n - 2X_j} \leq \frac{2(1 + X_j)}{n}$$

under this assumption. The next increase of a Yule process at $Y_{4j/n}$ arrives at rate $Y_{4j/n}$. Thus, conditional on $Y_{4j/n}$, the probability that a Yule process increases from $Y_{4j/n}$ to $Y_{4j/n} + 1$ within time $4/n$ is

$$1 - \exp\left(-\frac{4Y_{4j/n}}{n}\right) \geq 1 - \exp\left(-\frac{4(1 + X_j)}{n}\right) \geq \frac{2(1 + X_j)}{n}.$$

The first inequality uses the inductive hypothesis that $1 + X_j \leq Y_{4j/n}$, and the second uses the inequality $1 - e^{-x} \geq x/2$, which holds for $0 \leq x \leq 1$. Thus, conditional on X_j and $Y_{4j/n}$, a Yule process starting at $Y_{4j/n}$ is more likely to increase in time $4/n$ than is the process (X_i) from time j to $j + 1$. Thus, conditionally on X_j and $Y_{4j/n}$, we can couple a Yule process starting at $Y_{4j/n}$ of duration $4/n$ with (X_i) so that it increases only if (X_i) does from j to $j + 1$. Tacking this Yule process onto the end of (Y_t) , we have extended our coupling as desired. □

Proof of Proposition 5.7. Recall that $X_n = |\mathcal{B}|$. Couple (X_i) and (Y_t) according to Lemma 5.9. Fix $k \leq n/4$, and let $S' = \min\{i \mid X_i \geq k\}$, taking $S' = n$ if the minimum is over the empty set. Since $S' \leq S$, we have $1 + X_{S'} \leq Y_{4S'/n}$.

We claim that if $X_n \geq k$, then $Y_4 \geq 1 + k$. To see this, observe that if $X_n \geq k$, then $X_{S'} \geq k$. Thus

$$Y_4 \geq Y_{4S'/n} \geq 1 + X_{S'} \geq 1 + k.$$

This demonstrates that

$$\mathbf{P}[X_n \geq k] \leq \mathbf{P}[Y_4 - 1 \geq k] = (1 - e^{-4})^k \leq \exp(-e^{-4}k) \tag{5.6}$$

since $Y_t - 1$ is distributed as $\text{Geo}(e^{-t})$ (see [11, Section 4.1.D], for example).

For $n/4 < k \leq n$,

$$\mathbf{P}[X_n \geq k] \leq \mathbf{P}[X_n \geq n/4] \leq \exp\left(-\frac{e^{-4}n}{4}\right) \leq \exp\left(-\frac{e^{-4}k}{4}\right). \tag{5.7}$$

The right hand side of (5.7) is larger than that of (5.6). As X_n does not take values larger than n , this shows that (5.7) holds for all $k \geq 0$. \square

Next, we prove a subexponential tail bound on the size of \mathcal{B}_x for any fixed x .

Lemma 5.10. For $x > 0$ and $t \leq n - x$,

$$\mathbf{P}[|\mathcal{B}_x| \geq t] \leq \binom{t+x}{t} \prod_{i=1}^t \left(\frac{x}{n-i}\right) \leq \left(\frac{ex}{t}\right)^t. \tag{5.8}$$

Proof. In order for each of the first t values in \mathcal{B} to satisfy $b_i - i < x$, we must have that all of the values are less than $t + x$. For a specific choice of letters $1 < c_1 < \dots < c_t$, let $E(c_1, \dots, c_t)$ denote the event that the first t values of \mathcal{B} are c_1, \dots, c_t . If $|\mathcal{B}_x| \geq t$ then $E(c_1, \dots, c_t)$ holds for one of at most $\binom{t+x}{t}$ possible choices of (c_1, \dots, c_t) .

For a particular choice of letters $c_1 < \dots < c_t < t + x$,

$$\mathbf{P}[E(c_1, \dots, c_t)] \leq \mathbf{P}\left[\bigcap_{i=1}^t \{\pi_n^{-1}(c_i) < x + i\}\right] = \prod_{i=1}^t \frac{x}{n - i + 1}.$$

To prove the final bound in (5.8), we apply the union bound and evaluate

$$\binom{t+x}{t} \prod_{i=1}^t \left(\frac{x}{n-i+1}\right) = \frac{x^t(x+t)\cdots(x+1)}{t!n\cdots(n-t+1)} \leq \frac{x^t}{t!},$$

since $x + t \leq n$ under the assumptions of the lemma, and then we apply the bound $t! \geq (t/e)^t$. \square

Proposition 5.7 and Lemma 5.10 now combine to bound $\mathbf{E}L_n^p$.

Proposition 5.11. For any $p > 0$, it holds that $\sup_n \mathbf{E}L_n^p < \infty$.

Proof. Fix $p > 0$, and choose x large enough that $(1 + \frac{1}{x})^{2p} < C^{-1}$ for the constant C from Proposition 5.7. By (5.3) and the Cauchy-Schwarz inequality,

$$\mathbf{E}L_n^p \leq \mathbf{E}\left[2^{2p|\mathcal{B}_x|}\right]^{1/2} \mathbf{E}\left[|\mathcal{B}|^{2p}\left(1 + \frac{1}{x}\right)^{2p|\mathcal{B}|}\right]^{1/2}. \tag{5.9}$$

By Proposition 5.7,

$$\mathbf{E} \left[|\mathcal{B}|^{2p} \left(1 + \frac{1}{x} \right)^{2p|\mathcal{B}|} \right] \leq \sum_{t=0}^{\infty} C^t t^{2p} \left(1 + \frac{1}{x} \right)^{2pt}, \quad (5.10)$$

which is summable. Similarly, applying Lemma 5.10 and using $|\mathcal{B}_x| \leq n$,

$$\begin{aligned} \mathbf{E} \left[2^{2p|\mathcal{B}_x|} \right] &\leq \sum_{t=0}^{n-x} \mathbf{P}[|\mathcal{B}_x| = t] 2^{2pt} + \mathbf{P}[|\mathcal{B}_x| \geq n-x] 2^{2pn} \\ &\leq \sum_{t=0}^{\infty} \left(\frac{ex}{t} \right)^t 2^{2pt} + \left(\frac{ex}{n-x} \right)^{n-x} 2^{2pn}. \end{aligned}$$

The first term is finite, and the last term vanishes as $n \rightarrow \infty$. Together with (5.10), this yields an upper bound on (5.9) with no dependence on n . \square

6 Further Questions

Many of the open questions in [13] about the global properties of the fixed point forest F_n remain unanswered. Recall that each base of F_n is a permutation with 1 as a fixed point. Let T_n denote the tree in F_n with the identity permutation as its base. In fact, T_n is the largest tree in F_n , which can be seen as follows. Let π be any base with $\pi(1) = 1$. Let i be the largest index of such that $\pi(i) \neq i$. Then $\pi(i)$ is never bumped, so switching $\pi(i)$ and $\pi(1)$ creates a new permutation π' with subtree at least as big as the tree with base π . Note that $\pi' \in T_n$, since all letters after 1 are fixed points. Hence an isomorphic copy of the tree starting from π is also in T_n .

In [13, Propositions 24 and 25], it is shown that for a uniformly random permutation π_n ,

$$\frac{1}{n} \leq \mathbf{P}[\pi_n \in T_n] \leq \frac{e}{n},$$

and it is conjectured that $\mathbf{P}[\pi_n \in T_n] \sim \frac{1}{n}$. We highlight this question here and tack on a few more of our own:

Question 6.1. Prove that $n\mathbf{P}[\pi_n \in T_n]$ converges as $n \rightarrow \infty$ and determine its limit. Characterize all permutations in T_n . How do the next largest components compare in size to it?

Another question posed in [13] is:

Question 6.2. Let R_n be the distance from π_n to the base of its tree in the fixed point forest. What are the limiting asymptotics of $\mathbf{E}R_n$?

One could also ask about the limiting fluctuations of R_n from its mean.

Though some of our work in Section 4 could be helpful in addressing these questions, our limit tree has nothing to say about them. For example, the limiting tree has no base at all, reflecting that questions about the distance from π_n to a base are not local. One could instead look for a different limit of the fixed point forest where edges are scaled so that the diameter of each component of the forest stays bounded, along the lines of the continuum random tree (see [1]).

The present work has created other avenues to explore. Continuing in the same theme as Section 5, what can we say about paths from root to leaf in the limiting tree besides the longest and shortest ones? In particular,

Question 6.3. Walk from the root towards the leaves in the limiting tree by choosing randomly among all children at each step. What is the distribution of the number of steps before reaching a leaf?

Other properties of the tree are interesting as well:

Question 6.4. Is simple random walk on the limiting tree transient or recurrent? What is the branching number (see [12, Section 1.2]) of the tree?

A random walk on the nonlimit fixed point forest can be interpreted as a stochastic version of the bumping process where we randomly bump and unbump letters of a permutation. A solution to the above problem would likely give information on how quickly this process moves from the starting permutation.

Another problem is to determine what happens when the root of the fixed point forest is chosen nonuniformly:

Question 6.5. Determine the local limit of the fixed point forest when the root is sampled from the Ewens or Mallows distributions.

The fixed point of Mallows distributions are not distributed evenly over the permutation, so we would expect convergence to a limiting tree defined by Poisson point processes of nonuniform intensity.

References

- [1] David Aldous, *The continuum random tree. II. An overview*, Stochastic analysis (Durham, 1990), London Math. Soc. Lecture Note Ser., vol. 167, Cambridge Univ. Press, Cambridge, 1991, pp. 23–70. MR-1166406
- [2] David Aldous and J. Michael Steele, *The objective method: probabilistic combinatorial optimization and local weak convergence*, Probability on discrete structures, Encyclopaedia Math. Sci., vol. 110, Springer, Berlin, 2004, pp. 1–72. MR-2023650
- [3] A. D. Barbour, Lars Holst, and Svante Janson, *Poisson approximation*, Oxford Studies in Probability, vol. 2, The Clarendon Press, Oxford University Press, New York, 1992, Oxford Science Publications. MR-1163825
- [4] Itai Benjamini and Oded Schramm, *Recurrence of distributional limits of finite planar graphs*, Electron. J. Probab. **6** (2001), no. 23, 13 pp. (electronic). MR-1873300
- [5] Patrick Billingsley, *Convergence of probability measures*, second ed., Wiley Series in Probability and Statistics: Probability and Statistics, John Wiley & Sons, Inc., New York, 1999, A Wiley-Interscience Publication. MR-1700749
- [6] Sourav Chatterjee, Persi Diaconis, and Elizabeth Meckes, *Exchangeable pairs and Poisson approximation*, Probab. Surv. **2** (2005), 64–106. MR-2121796
- [7] Kai Lai Chung, *A course in probability theory*, third ed., Academic Press, Inc., San Diego, CA, 2001. MR-1796326
- [8] W. J. Hendricks, *The stationary distribution of an interesting Markov chain*, J. Appl. Probability **9** (1972), 231–233. MR-0292178
- [9] W. J. Hendricks, *An extension of a theorem concerning an interesting Markov chain*, J. Appl. Probability **10** (1973), 886–890. MR-0358998
- [10] Olav Kallenberg, *Foundations of modern probability*, second ed., Probability and its Applications (New York), Springer-Verlag, New York, 2002. MR-1876169
- [11] Samuel Karlin and Howard M. Taylor, *A first course in stochastic processes*, second ed., Academic Press [A subsidiary of Harcourt Brace Jovanovich, Publishers], New York-London, 1975. MR-0356197
- [12] Russell Lyons and Yuval Peres, *Probability on trees and networks*, Cambridge University Press, 2016, Available at <http://pages.iu.edu/~rdlyons/>.
- [13] Gwen McKinley, *A problem in card shuffling*, UC Davis Undergraduate Thesis, 2015, https://www.math.ucdavis.edu/files/1114/3950/6599/McKinley_UG_Thesis_SP15.pdf, pp. 1–29.
- [14] Nathan Ross, *Fundamentals of Stein’s method*, Probab. Surv. **8** (2011), 210–293. MR-2861132
- [15] The Sage-Combinat community, *Sage-Combinat: enhancing Sage as a toolbox for computer exploration in algebraic combinatorics*, 2008, <http://combinat.sagemath.org>.

- [16] The Sage Developers, *Sage Mathematics Software (Version 7.2)*, 2016, <http://www.sagemath.org>.
- [17] M L Tsetlin, *Finite automata and models of simple forms of behaviour*, Russian Mathematical Surveys **18** (1963), no. 4, 1.
- [18] H. W. Watson and Francis Galton, *On the probability of the extinction of families*, J. Anthr. Inst. Great Britain **4** (1875), 138–144.

Acknowledgments. We would like to thank David Aldous, Persi Diaconis, Gwen McKinley, and Graham White for discussions. This work benefited from computations in SageMath [15, 16].