



Inference for conditioned Galton-Watson trees from their Harris path

Romain Azaïs, Alexandre Genadot, Benoît Henry

► To cite this version:

Romain Azaïs, Alexandre Genadot, Benoît Henry. Inference for conditioned Galton-Watson trees from their Harris path. *ALEA : Latin American Journal of Probability and Mathematical Statistics*, 2019, 16 (1), pp.1-45. 10.30757/ALEA.v16-21 . hal-01360650v3

HAL Id: hal-01360650

<https://hal.science/hal-01360650v3>

Submitted on 1 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Inference for conditioned Galton-Watson trees from their Harris path

Romain Azaïs^a, Alexandre Genadot^b, and Benoît Henry^c

^a Laboratoire Reproduction et Développement des Plantes, Univ Lyon, ENS de Lyon, UCB Lyon 1, CNRS, INRA, Inria,
F-69342, Lyon, France.

^b Institut de Mathématiques de Bordeaux, Univ Bordeaux, CNRS, UMR 5251 and INRIA Bordeaux-Sud Ouest, Team CQFD,
F-33400 Talence, France.

^c Institut Mines-Télécom Lille Douai, Villeneuve d’Ascq, France.

Abstract

Tree-structured data naturally appear in various fields, particularly in biology where plants and blood vessels may be described by trees, but also in computer science because XML documents form a tree structure. This paper is devoted to the estimation of the relative scale parameter of conditioned Galton-Watson trees. New estimators are introduced and their consistency is stated. A comparison is made with an existing approach of the literature. A simulation study shows the good behavior of our procedure on finite-sample sizes and from missing or noisy data. An application to the analysis of revisions of Wikipedia articles is also considered through real data.

1 Introduction

Many data are naturally modeled by an ordered tree structure: from blood vessels in biology to XML files in computer science through the secondary structure of RNA in biochemistry. The statistical analysis of a dataset of hierarchical records is thus of great interest. In this paper, our aim is to propose new methods to estimate the scale parameter arising in Galton-Watson trees conditioned on their number of nodes from various statistical settings.

A Galton-Watson tree is the genealogical tree of a population starting from one initial ancestor (the root) in which each individual gives birth to a random number of children according to the same probability distribution, independently of each other. In this article, we focus on Galton-Watson trees conditional on their number of nodes. Several main classes of random trees can be seen as conditioned Galton-Watson trees [9, 16]. For instance, an ordered tree picked uniformly at random in the set of all ordered trees of a given size is a conditioned Galton-Watson tree with offspring distribution the geometric law with parameter $1/2$. In addition, an ordered tree picked uniformly at random in the set of d -ary trees, i.e., trees in which each node has no more than d children, is a conditioned Galton-Watson tree with offspring distribution the binomial law with parameter d and $1/d$. In particular, binary trees but also full binary trees (taking the uniform law on the set $\{0, 2\}$ as offspring distribution) are thus encoded by a conditioned Galton-Watson model. Binary trees are widely used in computer science, through binary search trees [20] and Huffman coding [21] commonly used for data compression. They also appear in biology in the approximation of phylogenetic trees [3] for example. One also refers the reader to [16, 10. Examples of simply generated random trees] for other examples of conditioned Galton-Watson trees arising from random trees (that can even be unordered and labelled). To sum up, conditioned Galton-Watson trees model a large variety of random hierarchical structures. Developing specific statistical methods for this stochastic model is thus of first importance.

Any ordered tree may be encoded by its Harris path which returns height of nodes in depth-first order (see Subsection 2.2, Algorithm 1 and Figure 1). Aldous [2, Theorem 23] stated the following asymptotic property of the Harris path $\mathcal{H}[\tau_n]$ of a Galton-Watson tree τ_n conditioned on having n nodes,

$$\left(\frac{\mathcal{H}[\tau_n](2nt)}{\sqrt{n}}, t \in [0, 1] \right) \xrightarrow{(d)} \left(\frac{2}{\sigma} \mathbf{e}_t, t \in [0, 1] \right), \quad (1)$$

in the uniform topology of $\mathcal{C}([0, 1], \mathbf{R})$, when n goes to infinity whenever the offspring distribution is 1 on average with standard deviation σ and \mathbf{e} denotes the normalized Brownian excursion. This means that conditioned Galton-Watson trees asymptotically share a common form (the so-called *continuum random tree*) given by the Brownian excursion, and can be differentiated only by the scale parameter of interest σ^{-1} . This unknown quantity is to be estimated from only one tree or from a forest of independent trees generated from the same birth distribution.

Estimating (functions of) σ from a forest of independent conditioned Galton-Watson trees has only been considered in a recent paper [4]. The authors of [4] exploit a corollary of the weak convergence (1) providing the asymptotic distribution of the height of a uniformly sampled node in the tree (see [4, Proposition 4]) to construct estimators of the variance σ^2 and develop asymptotic tests. It should be already noticed that estimation strategies based on the convergence in distribution (1) can only lead to weak convergence results for estimators computed from a unique tree. The aim of the present paper is twofold. First, we establish in Theorem 6 that the empirical variance of the numbers of children is a consistent estimator of σ^2 (in particular even from the observation of only one tree), whereas, even if a Galton-Watson tree is generated from a sequence of i.i.d. random variables, this is not the case for the conditioned structure. This new result based on a corollary of Bartlett's formula shows that the empirical variance provides a better estimate of σ^2 than any other statistical method based on Aldous' theorem (1). Secondly, we propose two new estimation strategies for σ^{-1} from a forest of independent conditioned Galton-Watson trees based on the weak convergence established by Aldous and we compare them with the procedure developed in [4].

These estimation strategies rely on the idea motivated by the weak convergence (1) that, on average, the normalized Harris paths of the forest should look like the expected process $(2\sigma^{-1}E(t), t \in [0, 1])$ at least asymptotically, where $E(t) = \mathbf{E}[\mathbf{e}(t)]$. The parameter σ^{-1} can thus be expressed as the solution of a least square problem. Our first method consists in computing the least square estimator of σ^{-1} from the concatenation of the Harris paths of the forest. We establish two results of convergence in Subsection 4.2. For only one Galton-Watson tree τ_n conditioned on having n nodes, this estimator of σ^{-1} is given (see Subsection 3.1) by

$$\hat{\lambda}[\tau_n] = \frac{\langle \mathcal{H}[\tau_n](2n\cdot), E \rangle}{2\sqrt{n}\|E\|_2^2},$$

where $\langle \cdot, \cdot \rangle$ is the scalar product of $\mathbf{L}^2([0, 1], \mathbf{R})$. By virtue of the weak convergence (1), one may remark (see Corollary 7) that

$$\hat{\lambda}[\tau_n] \xrightarrow{(d)} \sigma^{-1} \Lambda_\infty,$$

where $\Lambda_\infty = \frac{\langle \mathbf{e}, E \rangle}{\|E\|_2^2}$. Actually, the aforementioned least square estimator only exploits the average behavior of Λ_∞ (in other words, the average asymptotic behavior of Harris paths) and not its complete distribution. Our second strategy takes into account the shape of the distribution of Λ_∞ : we estimate σ^{-1} by the parameter x that aligns the theoretical distribution of $x\Lambda_\infty$ and the empirical measure of the $\hat{\lambda}[\tau_{n_i}^i]$'s in terms of Wasserstein distance, the considered forest being composed of N trees $\tau_{n_i}^i$. Convergence results are stated in Subsection 4.3. We point out that the theoretical properties of Λ_∞ are far from obvious. In particular, we establish by Malliavin calculus that Λ_∞ is absolutely continuous w.r.t. the Lebesgue measure in Proposition 10, which is required in some proofs.

The authors of [4] do not focus on the problem of estimating σ^{-1} but, for the sake of comparison, we rely on their approach to provide another estimator of this quantity. We compare these alternative strategies from both theoretical and numerical points of view. In particular, we show in Subsection 3.1 that the variances

of our estimators are approximately 4 times lower than the one of the estimator based on this competitive approach of the literature. Our results are better in terms of dispersion because the estimators take into account all the behavior of the tree and not only the behavior of a randomly chosen node. We also point out that the theoretical setting of [4] is slightly different because investigations are directly based on infinite trees (i.e., *continuum random trees*, unobservable in practice) and not on large but finite trees.

At this step, one may wonder whether an approach that only yields weak convergence results is relevant considering the empirical variance is a consistent estimator of σ^2 . Our idea is to explore statistical inference for trees from coding processes, i.e., via functional data analysis. This connection has been first established in the recent paper [31]. In the present article we aim at investigating this strategy when the data have been generated from the stochastic model of conditioned Galton-Watson trees. In Subsection 5.3, we prove from simulations that our estimators based on the weak convergence of Harris paths provide good results even from missing or noisy data, in particular when the empirical variance presents a large bias or can not be computed, showing the great interest of this approach.

The application of our estimators on simulated and real data in Section 5 appears to be a non trivial task, in particular because it requires important preliminary computations. For this reason and to provide a turnkey solution, we have developed a **Matlab** toolbox that enables users to quickly and easily apply our methods to data. This toolbox as well as a detailed user documentation can be found at the webpage: <http://agh.gforge.inria.fr>. The numerical experiments presented in Subsection 5.2 show that both our estimators and the approach developed in [4] are intrinsically biased for binary trees because of the approximation of the Harris paths of finite trees by the average Brownian excursion. Indeed, we empirically observe on simulation examples that Harris paths of binary trees weakly converge to the Brownian excursion from below (see Figure 7). As a consequence, we introduce a numerical correction of this negative bias, also implemented in the toolbox. The simulation study illustrates the good behavior of the corrected estimates on finite-sample sizes.

Visualizing the evolution of historical hierarchical data is a difficult issue in particular because such objects have no representation in a Euclidean space. This problem occurs in the study of the sequence of revisions of a given Wikipedia article. Indeed, the famous free Internet encyclopedia allows its users (the *Wikipedians*) to edit almost any articles. Starting from the creation of a given article, the history of revisions is accessible and can be investigated to understand how the contributors agree on its structure, or to automatically detect vandalism¹ [1, 25]. *IBM's History Flow* is a visualization tool for documents in various stages of their development which has been applied to Wikipedia articles [32, 33]. We think that our method may be a complementary tool to this famous technique. Indeed the structure of HTML documents, such as Wikipedia articles, may be encoded by an ordered tree structure (see Figure 19). Furthermore, all the Wikipedia webpages share the same template, i.e., standardized HTML/CSS files, and thus can be differentiated by their relative scale. In Section 6, we apply our estimators to the analysis of two Wikipedia articles. We highlight that Wikipedia articles undergo “running in” period before reaching some kind of steady state in which the contributors had agreed on the structure of the article. In addition, we show that our techniques may be used to detect improper editions of an article.

The organization of the paper is as follows. Section 2 is devoted to the formulation of the problem at hand: definition of conditioned Galton-Watson trees in Subsection 2.1, definition of Harris paths in Subsection 2.2, asymptotic behavior of Harris paths of conditioned Galton-Watson trees in Subsection 2.3. In addition, we state in Subsection 2.4 the consistency of the empirical variance of the numbers of children. The two estimation procedures from Harris paths are presented in Section 3, while Section 4 focuses on the results of convergence. Simulation techniques for conditioned Galton-Watson trees, numerical experiments and application to real data are presented in Sections 5 and 6. In particular, Subsection 5.3 is dedicated to the difficult context of missing or noisy data in which the empirical variance performs less well than our Harris paths-based estimators or even can not be computed. In this preliminary version, all the figures obtained from simulated or real data have been deferred to the end of the article.

¹It frequently happens that malicious people willingly disrupt the content of an article, for instance, for political or ideological reasons.

2 Conditioned Galton-Watson trees

2.1 Definition

Trees are connected graphs with no cycles. A rooted tree τ is a tree in which one node has been distinguished as the root, denoted by $r(\tau)$ (always drawn at the bottom of the tree in this paper). In this case, the edges are assigned a natural orientation, away from the root towards the leaves. One obtains a directed rooted tree in which there exists a parent-child relationship: the parent of a node v is the first vertex met on the path to the root starting from v . The length of this path (in number of nodes) is called the height $h(v)$ of v . The set $c(v)$ of children of a vertex v is the set of nodes that have v as parent. An ordered or plane tree is a rooted tree in which an ordering has been specified for the set of children of each node, conventionally drawn from left to right. In this paper we consider ordered rooted trees simply referred to as trees. In addition, for any node v , $\tau[v]$ denotes the subtree of τ composed of v and all of its descendants in τ .

Intuitively, a Galton-Watson tree can be seen as a tree encoding the dynamic of a population generated from some offspring distribution μ on \mathbf{N} . A Galton-Watson tree τ with offspring distribution μ is a random ordered rooted tree constructed recursively as follows.

- ◇ The number of children $\#c(r(\tau))$ emanating from the root is a random variable with law μ . The first generation consists thus in $\#c(r(\tau))$ vertices.
- ◇ Assume that the n^{th} generation of children has been constructed and consists in a list of vertices \mathcal{V}_n . Then, the generation $n + 1$ is constructed such that $\{\#c(v) : v \in \mathcal{V}_n\}$ is a collection of independent random variables with law μ .

The asymptotic behavior of Galton-Watson trees may exhibit different regimes depending on the average number of children per capita,

$$\bar{\mu} = \sum_{k \geq 0} k\mu(k),$$

where $\mu(k)$ is the measure of the singleton $\{k\}$ by μ .

- ◇ The subcritical case: $\bar{\mu} < 1$. In this case, the average number of nodes is finite. This means that the population goes extinct almost surely.
- ◇ The critical case: $\bar{\mu} = 1$. The fact that the offspring distribution μ is critical also ensures the almost sure finiteness of the tree, except when $\mu(1) = 1$ where the number of nodes is almost surely infinite. When $\mu(1) < 1$, in contrary to the sub-critical case, the expected number of nodes is infinite.
- ◇ The supercritical case: $\bar{\mu} > 1$. In this case, the number of vertices is infinite with positive probability.

We use the notation $\text{GW}_n(\mu)$ for the distribution of Galton-Watson trees with offspring distribution μ conditioned on having n nodes.

Remark 1 *In this paper, we will always state our results in terms of critical Galton-Watson trees. However, this is not really a restriction since, as noted in [28, 6.3 Brownian asymptotics for conditioned Galton-Watson trees], for any offspring distribution μ , there exists a critical law μ' such that*

$$\text{GW}_n(\mu) \stackrel{(d)}{=} \text{GW}_n(\mu').$$

In particular, this means that the average number of children $\bar{\mu}$ is not identifiable from conditioned Galton-Watson trees without some additional assumptions on μ .

2.2 From ordered trees to Harris paths

The Harris walk $\mathcal{H}[\tau]$ of an ordered rooted tree τ is defined from the depth-first search algorithm and the notion of height of nodes already presented in Subsection 2.1. Depth-first search is an algorithm for traversing a tree which one explores as far as possible along each branch before backtracking. The version of the algorithm used to define the Harris walk of a tree is presented in Algorithm 1.

Function DFS($\tau, l = \emptyset$):
Data: an ordered tree τ
Result: vertices of τ in depth-first order
 add $r(\tau)$ to l
for v **in** $c(r(\tau))$ **do**
 if $r(t/v)$ **is not in** l **then**
 call DFS($t[v], l$)
 add again $r(\tau)$ to l
return l

Algorithm 1: Recursive depth-first search.




Figure 1: Construction of the Harris path (right) from 0 to $2n = 14$ as the contour of an ordered tree (left) with $n = 7$ nodes.

Remark 2 In Algorithm 1, each node v appears $\#c(v) + 1$ times. Starting from the root of a tree τ , the result is thus a sequence of length

$$\sum_{v \in \tau} (\#c(v) + 1) = \#\tau + \sum_{v \in \tau} \#c(v) = 2\#\tau - 1,$$

because the root is the only vertex not to be counted.

The Harris walk $\mathcal{H}[\tau]$ of τ is defined as a sequence of integers indexed by the set $\{0, \dots, 2\#\tau\}$ as follows:

- ◇ $\mathcal{H}[\tau](0) = \mathcal{H}[\tau](2\#\tau) = 0$,
- ◇ for $1 \leq k < 2\#\tau$, $\mathcal{H}[\tau](k) = h(v) + 1$ where v is the k^{th} node in depth-first traversal of τ .

The Harris process is then defined as the linear interpolation of the Harris walk (see example in Figure 1). Note that, as displayed in Figure 2, the tree can be recovered from its Harris process such that the correspondence is one to one.




Figure 2: The ordered tree of Figure 1 in its Harris path (left): each vertical axis represents a node of the original structure (right). A common picture helping to see how to recover the tree from the contour is to imagine putting glue under the contour and then squeezing the contour together horizontally such that the inner parts of the contour which face each other are glued.

2.3 Asymptotic behavior of Harris paths

Let $\tau_n \sim \text{GW}_n(\mu)$ with $\bar{\mu} = 1$. The variance of the offspring distribution μ is denoted by σ^2 ,

$$\sigma^2 = \sum_{k \geq 1} (k-1)^2 \mu(k).$$

We focus on the asymptotic behavior of the Harris process $\mathcal{H}[\tau_n](2n\cdot)$ when n tends to infinity. The convergence in distribution has been stated by Aldous [2, Theorem 23].

Theorem 3 *When n goes to infinity, we have*

$$\left(\frac{\mathcal{H}[\tau_n](2nt)}{\sqrt{n}}, t \in [0, 1] \right) \xrightarrow{(d)} \left(\frac{2}{\sigma} \mathbf{e}_t, t \in [0, 1] \right),$$

where \mathbf{e} is a standard Brownian excursion, the convergence holding in law in the space $\mathcal{C}([0, 1], \mathbf{R})$. An illustration of this convergence in distribution may be found in Figure 5.

[Figure 5 approximately here]

Let us simply recall that a standard Brownian excursion is a Brownian motion conditioned on being positive and on taking the value 0 at time 1. The density of \mathbf{e}_t , for $0 \leq t \leq 1$, is given in [30, XI. 3. Bessel Bridges] and writes

$$\forall x \in \mathbf{R}, f_{\mathbf{e}_t}(x) = \sqrt{\frac{2}{\pi}} \frac{x^2}{\sqrt{t(1-t)}^3} \exp\left(-\frac{x^2}{2t(1-t)}\right) \mathbf{1}_{\mathbf{R}_+}(x).$$

From this, we can compute some simple functionals of the excursion. For instance, we have,

$$\forall 0 \leq t \leq 1, \quad E_t = \mathbf{E}[\mathbf{e}_t] = 4\sqrt{\frac{t(1-t)}{2\pi}} \quad \text{and} \quad \mathbf{E}[\mathbf{e}_t^2] = 3t(1-t). \quad (2)$$

The easiest way to simulate a Brownian excursion is certainly from its identity in law with a three-dimensional Bessel bridge [30, Theorem XII.4.2], which is simply the Euclidean norm of a three-dimensional Brownian bridge,

$$(\mathbf{e}_t, t \in [0, 1]) \stackrel{(d)}{=} \left(\sqrt{\sum_{i=1}^3 (B_t^i - tB_1^i)^2}, t \in [0, 1] \right), \quad (3)$$

where the B^i 's are three independent Brownian motions. The convergence presented in Theorem 3 also holds in expectation [10, Theorem 1].

Theorem 4 *When n goes to infinity, we have,*

$$\forall 0 \leq t \leq 1, \quad \mathbf{E} \left[\frac{\mathcal{H}[\tau_n](2nt)}{\sqrt{n}} \right] \longrightarrow \frac{2}{\sigma} E_t,$$

where the function $(E_t, 0 \leq t \leq 1)$ has been defined in (2).

Remark 5 *Theorem 3 establishes that, in the asymptotic regime, the shape of a conditioned Galton-Watson tree is given by the normalized Brownian excursion, regardless of the offspring distribution μ . However, there is one scale parameter given by the inverse of the standard deviation of μ . As a consequence, when μ is unknown, the only quantity of interest that one may access by asymptotic inference from Theorem 3 is σ^{-1} . From Section 3, we shall focus on the estimation of σ^{-1} .*

2.4 Empirical estimators

The purpose of this section is to study the behavior of empirical estimators of the mean $\bar{\mu}$ and of the variance σ^2 . These estimators are privileged candidates in the case of standard, i.e., unconditioned, Galton-Watson trees. Unfortunately, the lack of independency and homogeneity in the numbers of children of the nodes in conditioned Galton-Watson trees suggests such methods should not work in this more complex framework.

2.4.1 Empirical mean

We begin our study with the empirical mean. Let τ_n be a Galton-Watson tree with birth distribution μ conditioned on having n nodes. Denote, for any i in $\{1, \dots, n\}$, X_i the number of children of the i^{th} individual indexed in depth-first order in τ_n . Set

$$M[\tau_n] = \frac{1}{n} \sum_{i=1}^n X_i$$

the empirical mean of the number of children of the individuals in τ_n . However, it is easily seen that, whatever the underlying stochastic model,

$$M[\tau_n] = \frac{\#\tau_n - 1}{n} = 1 - \frac{1}{n}.$$

As a consequence, this estimator is deterministic and always estimates 1 asymptotically whatever the real mean of the birth distribution.

2.4.2 Empirical variance

This section is devoted to the study of the empirical variance,

$$V[\tau_n] = \frac{1}{n} \sum_{i=1}^n (X_i - M[\tau_n])^2.$$

Our main result of convergence is given below.

Theorem 6 *When n goes to infinity,*

$$\mathbf{E} [|V[\tau_n] - \sigma^2|] \rightarrow 0.$$

In particular, $V[\tau_n]$ converges to σ^2 in probability.

Proof. Let $(\xi_i)_{1 \leq i \leq n}$ be a sequence of i.i.d. random variables distributed according to μ . Using the work of Devroye [9, Turning to random walks], it is known that there exists a random permutation Σ such that the random vector $(\xi_{\Sigma(1)}, \dots, \xi_{\Sigma(n)})$ conditioned to $\sum_{i=1}^n \xi_i = n - 1$ is equal in distribution to (X_1, \dots, X_n) . However, since the empirical estimators are invariant up to permutation, it follows that we can work directly with the vector (ξ_1, \dots, ξ_n) conditioned to $\sum_{i=1}^n \xi_i = n - 1$. Consequently, our main goal is simply to prove, keeping in mind that $M[\tau_n] = 1 - 1/n$, that

$$R(n) = \mathbf{E} \left[\left| \frac{1}{n} \sum_{i=1}^n \left(\xi_i - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right| \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right]$$

converges to 0 as n goes to infinity. The proof lies on the asymptotic behavior of conditional probabilities which were obtained in [14] and allows to get the expected convergence. Let F be a measurable and bounded real valued function. Since $\xi_1/(n-1)$ goes to 0 in probability as n goes to infinity, it follows, according to [14, Theorem 1] (using the alternative hypothesis of Remark 2.9 of this paper), that the conditional expectation

$$\mathbf{E} \left[F(\xi_1) \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right]$$

converges to $\mathbf{E}[F(\xi_1)]$, as n goes to infinity. Since, μ does not have necessarily third order moments, we need to use a truncation method. Hence, let us consider, for any positive integer k ,

$$R(n) = \mathbf{E} \left[\left| \frac{1}{n} \sum_{i=1}^n \left(\xi_i \mathbf{1}_{\xi_i > k} + \xi_i \mathbf{1}_{\xi_i \leq k} - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right| \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right].$$

Now, since

$$\left(\xi_i \mathbf{1}_{\xi_i > k} + \xi_i \mathbf{1}_{\xi_i \leq k} - 1 + \frac{1}{n} \right)^2 = \left(\xi_i \mathbf{1}_{\xi_i \leq k} - 1 + \frac{1}{n} \right)^2 + (\xi_i \mathbf{1}_{\xi_i > k})^2 + 2\xi_i \mathbf{1}_{\xi_i > k} \left(\frac{1}{n} - 1 \right),$$

we get

$$R(n) \leq T_1(n, k) + T_2(n, k) + 2 \left(1 - \frac{1}{n} \right) T_3(n, k),$$

where

$$\begin{aligned} T_1(n, k) &= \mathbf{E} \left[\left| \frac{1}{n} \sum_{i=1}^n \left(\xi_i \mathbf{1}_{\xi_i \leq k} - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right| \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right], \\ T_2(n, k) &= \mathbf{E} \left[\left| \frac{1}{n} \sum_{i=1}^n (\xi_i \mathbf{1}_{\xi_i > k})^2 \right| \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right], \\ T_3(n, k) &= \mathbf{E} \left[\left| \frac{1}{n} \sum_{i=1}^n \xi_i \mathbf{1}_{\xi_i > k} \right| \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right]. \end{aligned}$$

In order to treat $T_1(n, k)$, we consider

$$Q(n, k) = \mathbf{E} \left[\left(\frac{1}{n} \sum_{i=1}^n \left(\xi_i \mathbf{1}_{\xi_i \leq k} - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right)^2 \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right].$$

We have

$$Q(n, k) = \frac{1}{n^2} \sum_{1 \leq i, j \leq n} \mathbf{E} \left[\left(\left(\xi_i \mathbf{1}_{\xi_i \leq k} - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right) \left(\left(\xi_j \mathbf{1}_{\xi_j \leq k} - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right) \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right].$$

Using the exchangeability of the vector (ξ_1, \dots, ξ_n) under $\mathbf{P}(\cdot \mid \sum_{i=1}^n \xi_i = n-1)$, we get

$$\begin{aligned} Q(n, k) &= \frac{1}{n} \mathbf{E} \left[\left(\left(\xi_1 \mathbf{1}_{\xi_1 \leq k} - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right)^2 \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right] \\ &\quad + \frac{n(n-1)}{n^2} \mathbf{E} \left[\left(\left(\xi_1 \mathbf{1}_{\xi_1 \leq k} - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right) \left(\left(\xi_2 \mathbf{1}_{\xi_2 \leq k} - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right) \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right]. \end{aligned}$$

Then, [14, Theorem 1] allows to understand the asymptotic behavior of both terms in the above sum,

$$\lim_{n \rightarrow \infty} \mathbf{E} \left[\left(\left(\xi_1 \mathbf{1}_{\xi_1 \leq k} - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right)^2 \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right] = \mathbf{E} \left[\left((\xi_1 \mathbf{1}_{\xi_1 \leq k} - 1)^2 - \sigma^2 \right)^2 \right] > 0,$$

and

$$\mathbf{E} \left[\left(\left(\xi_1 \mathbf{1}_{\xi_1 \leq k} - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right) \left(\left(\xi_2 \mathbf{1}_{\xi_2 \leq k} - 1 + \frac{1}{n} \right)^2 - \sigma^2 \right) \middle| \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 \right]$$

goes, as n tends to infinity, to

$$\mathbf{E} \left[\left((\xi_1 \mathbf{1}_{\xi_1 \leq k} - 1)^2 - \sigma^2 \right) \left((\xi_2 \mathbf{1}_{\xi_2 \leq k} - 1)^2 - \sigma^2 \right) \right] = \left(\mathbf{E} \left[(\xi_1 \mathbf{1}_{\xi_1 \leq k} - 1)^2 \right] - \sigma^2 \right)^2 \rightarrow 0,$$

when k goes to infinity. Hence, we get, for any positive integer k ,

$$\lim_{n \rightarrow \infty} Q(n, k) = \left(\mathbf{E} \left[(\xi_1 \mathbf{1}_{\xi_1 \leq k} - 1)^2 \right] - \sigma^2 \right)^2, \quad (4)$$

without any assumption on the moments of μ because of the truncation $\xi_1 \mathbf{1}_{\xi_1 \leq k}$. First, the Cauchy-Schwarz inequality entails $T_1(n, k) \leq \sqrt{Q(n, k)}$, which gives, according to (4),

$$\limsup_{n \rightarrow \infty} T_1(n, k) \leq \left| \mathbf{E} \left[(\xi_1 \mathbf{1}_{\xi_1 \leq k} - 1)^2 \right] - \sigma^2 \right|.$$

In addition, we have according to [14, Theorem 1],

$$\limsup_{n \rightarrow \infty} T_2(n, k) \leq \limsup_{n \rightarrow \infty} \mathbf{E} \left[(\xi_i \mathbf{1}_{\xi_i > k})^2 \right] \frac{1}{n-1} \sum_{i=1}^n \xi_i = 1 = \mathbf{E} \left[(\xi_i \mathbf{1}_{\xi_i > k})^2 \right].$$

The case of $T_3(n, k)$ can be treated similarly and leads to

$$\limsup_{n \rightarrow \infty} R(n) \leq \left| \mathbf{E} \left[(\xi_1 \mathbf{1}_{\xi_1 \leq k} - 1)^2 \right] - \sigma^2 \right| + \mathbf{E} \left[(\xi_1 \mathbf{1}_{\xi_1 > k})^2 \right] + 2\mathbf{E} \left[|\xi_1 \mathbf{1}_{\xi_1 > k}| \right].$$

Now, letting k going to infinity leads to the result. \square

3 Estimation procedure

In the previous part, we have shown that the empirical variance is a consistent estimator of σ^2 in critical conditioned Galton-Watson trees despite the lack of independency in the numbers of children. In this section, we aim at developing estimation procedures for σ^2 that exploit the weak convergence towards the Brownian excursion stated by Aldous and presented in Theorem 3.

3.1 Adequacy of the Harris path with the expected contour

Let $\tau_n \sim \text{GW}_n(\mu)$ with $\bar{\mu} = 1$. We assume that the offspring distribution μ is unknown. By virtue of Theorem 4, the asymptotic average behavior of the normalized Harris process $(n^{-1/2} \mathcal{H}[\tau_n](2nt), 0 \leq t \leq 1)$ is given by $(2\sigma^{-1} E_t, 0 \leq t \leq 1)$, where σ^{-1} is obviously also unknown. We propose to estimate σ^{-1} by minimizing the \mathbf{L}^2 -error defined by

$$\lambda \mapsto \left\| \frac{\mathcal{H}[\tau_n](2n\cdot)}{\sqrt{n}} - 2\lambda E \right\|_2^2,$$

where, and in all the sequel, $\mathbf{L}^2 = \mathbf{L}^2([0, 1], \mathbf{R})$ and its usual norm is denoted $\|\cdot\|_2$ for the sake of readability. The solution of this least square problem is well-known and is given by

$$\hat{\lambda}[\tau_n] = \frac{\langle \mathcal{H}[\tau_n](2n\cdot), E \rangle}{2\sqrt{n} \|E\|_2^2}, \quad (5)$$

where $\langle \cdot, \cdot \rangle$ is the scalar product of \mathbf{L}^2 .

Corollary 7 *When n goes to infinity, we have*

$$\hat{\lambda}[\tau_n] \xrightarrow{(d)} \sigma^{-1} \Lambda_\infty,$$

where the random variable Λ_∞ is defined by

$$\Lambda_\infty = \frac{\langle \mathbf{e}, E \rangle}{\|E\|_2^2}. \quad (6)$$

Proof. The result directly follows from Theorem 3 because the functional $x \mapsto \langle x, E \rangle$ is continuous on $\mathcal{C}([0, 1], \mathbf{R})$. \square

Remark 8 *The convergence in distribution stated in Corollary 7 seems quite unsatisfactory because this means that $\hat{\lambda}[\tau_n]$ is not a consistent estimator of σ^{-1} and the least square strategy thus seems like inadequate in regards to the consistency of $V[\tau_n]$. In the sequel, we shall focus on the estimation of the parameter of interest σ^{-1} from a forest of conditioned Galton-Watson trees as in [4], only chance to get consistent estimates from Aldous' theorem. As mentioned in the introduction, our goal in this paper is to explore statistical inference for trees via functional data analysis of their Harris paths.*

Computing $\hat{\lambda}[\tau_n]$ is a first step in the estimation of the inverse standard deviation from a large number of conditioned Galton-Watson trees. As a consequence, the distribution of the limit variable Λ_∞ is of first importance.

Lemma 9 *For any $0 \leq t < u \leq 1$, we have*

$$\mathbf{E}[\mathbf{e}_t \mathbf{e}_u] = \frac{2}{\pi} \left[3\sqrt{t(u-t)(1-u)} + (2t(1-u) + u(1-t)) \arcsin \left(\sqrt{\frac{t(1-u)}{u(1-t)}} \right) \right].$$

Proof. This identity is derived from the joint density of $(\mathbf{e}_t, \mathbf{e}_u)$ given in [30, XI. 3. Bessel Bridges]. The density of $(\mathbf{e}_t, \mathbf{e}_u)$ for $0 \leq t < u \leq 1$ is given, for any positive numbers x and y , by

$$f_{t,u}(x, y) = \frac{2}{\pi} \frac{xy}{\sqrt{(u-t)t^3(1-u)^3}} \sinh \left(\frac{xy}{u-t} \right) \exp \left(-\frac{x^2 u}{2t(u-t)} \right) \exp \left(-\frac{y^2(1-t)}{2(1-u)(u-t)} \right).$$

Thus,

$$\mathbf{E}[\mathbf{e}_t \mathbf{e}_u] = \int_0^\infty \int_0^\infty \frac{2}{\pi} \frac{x^2 y^2}{\sqrt{(u-t)t^3(1-u)^3}} \sinh \left(\frac{xy}{u-t} \right) \exp \left(-\frac{x^2 u}{2t(u-t)} \right) \exp \left(-\frac{y^2(1-t)}{2(1-u)(u-t)} \right) dx dy.$$

At this point, one can use the power series of \sinh to separate the variables x and y and obtain

$$\mathbf{E}[\mathbf{e}_t \mathbf{e}_u] = \frac{2}{\pi} \frac{\sqrt{(u-t)^5}}{\sqrt{u^3(1-t)^3}} \sum_{k=0}^{\infty} \frac{4^{k+1}((k+1)!)^2}{(2k+1)!} \left(\sqrt{\frac{t(1-u)}{u(1-t)}} \right)^{2k+1}.$$

Note that when $0 \leq t < u \leq 1$ we have indeed $\sqrt{\frac{t(1-u)}{u(1-t)}} < 1$. Then, using

$$\sum_{k=0}^{\infty} \frac{4^{k+1}((k+1)!)^2}{(2k+1)!} x^{2k+1} = \frac{3x\sqrt{1-x^2} + (2x^2+1)\arcsin(x)}{\sqrt{(1-x^2)^5}},$$

we obtain the desired expression for $\mathbf{E}[\mathbf{e}_t \mathbf{e}_u]$. \square

Proposition 10 *The random variable Λ_∞ admits a density f_{Λ_∞} w.r.t. the Lebesgue measure. Furthermore,*

$$\mathbf{E}[\Lambda_\infty] = 1 \quad \text{and} \quad \mathbf{Var}(\Lambda_\infty) = \frac{1}{\|E\|_2^4} \int_0^1 \int_0^1 g(s, u) E_s E_u ds du - 1, \quad (7)$$

where the mapping $g : [0, 1]^2 \rightarrow \mathbf{R}_+$ is defined from

$$g(t, u) = \begin{cases} \frac{2}{\pi} \left[3\sqrt{t(u-t)(1-u)} + (2t(1-u) + u(1-t)) \arcsin \left(\sqrt{\frac{t(1-u)}{u(1-t)}} \right) \right] & \text{if } 0 \leq t \leq u \leq 1, \\ g(u, t) & \text{otherwise.} \end{cases}$$

Proof. We consider the probability space $(\mathcal{C}([0, 1], \mathbf{R}^3), \mathcal{F}, \mathbf{W})$, where $\mathcal{C}([0, 1], \mathbf{R}^3)$ is endowed with the uniform topology, \mathcal{F} is the corresponding Borel σ -field and \mathbf{W} is the Wiener measure. Let T be the continuous linear operator defined by

$$\begin{aligned} T : \mathcal{C}([0, 1], \mathbf{R}^3) &\rightarrow \mathcal{C}([0, 1], \mathbf{R}^3), \\ \varphi &\mapsto (T\varphi(s) = \varphi(s) - s\varphi(1)). \end{aligned}$$

Let also Γ be the following function,

$$\Gamma : \varphi \mapsto \int_0^1 |\varphi(s)|_2 \frac{E_s}{\|E\|_2^2} ds.$$

where $|x|_2$ denotes the Euclidian norm on \mathbf{R}^3 . With these notations and (3), we have that the pushforward measure of \mathbf{W} through the application

$$F : \varphi \mapsto \Gamma(T\varphi),$$

is the law of Λ_∞ . In other words, the random variable F is equal in distribution to Λ_∞ . Now for every φ in $\mathcal{C}([0, 1], \mathbf{R}^3)$ such that $\text{Leb}(\{t \in \mathbf{R}_+ : \varphi(t) = 0\}) = 0$, we have that Γ is Fréchet differentiable at the point φ . Moreover, the derivative at such point φ is given by

$$\begin{aligned} D_\varphi \Gamma : \mathcal{C}([0, 1], \mathbf{R}^3) &\rightarrow \mathbf{R}, \\ h &\mapsto \int_0^1 \frac{(\varphi(s), h(s))}{|\varphi(s)|_2} \frac{E_s}{\|E\|_2^2} ds, \end{aligned}$$

where (\cdot, \cdot) denotes the Euclidean scalar product on \mathbf{R}^3 . Indeed, some straightforward manipulations give

$$\int_0^1 \left[|\varphi(s) + h(s)|_2 - |\varphi(s)|_2 - \frac{(\varphi(s), h(s))}{|\varphi(s)|_2} \right] \frac{E_s}{\|E\|_2^2} ds = \int_0^1 \left[\frac{|h(s)|_2^2 + (\varphi(s), h(s)) \left(1 - \frac{|\varphi(s) + h(s)|_2}{|\varphi(s)|_2} \right)}{|\varphi(s) + h(s)|_2 + |\varphi(s)|_2} \right] \frac{E_s}{\|E\|_2^2} ds.$$

Now, since $\frac{E_s}{\|E\|_2^2} \leq \frac{3\sqrt{\pi}}{2\sqrt{2}}$ and using the Cauchy-Schwarz inequality, we obtain

$$\begin{aligned} &\left| \int_0^1 \left[|\varphi(s) + h(s)|_2 - |\varphi(s)|_2 - \frac{(\varphi(s), h(s))}{|\varphi(s)|_2} \right] \frac{E_s}{\|E\|_2^2} ds \right| \\ &\leq \frac{3\sqrt{\pi}}{2\sqrt{2}} \int_0^1 \left[\frac{|h(s)|_2^2 + |h(s)|_2 \left| |\varphi(s)|_2 - |\varphi(s) + h(s)|_2 \right|}{|\varphi(s) + h(s)|_2 + |\varphi(s)|_2} \right] ds \\ &\leq \frac{3\sqrt{\pi}}{2\sqrt{2}} \|h\|_\infty \int_0^1 \left[\frac{|h(s)|_2 + \left| |\varphi(s)|_2 - |\varphi(s) + h(s)|_2 \right|}{|\varphi(s) + h(s)|_2 + |\varphi(s)|_2} \right] ds, \end{aligned}$$

with $\|h\|_\infty = \sup_{s \in [0, 1]} |h(s)|_2$. Since

$$\int_0^1 \left[\frac{|h(s)|_2 + \left| |\varphi(s)|_2 - |\varphi(s) + h(s)|_2 \right|}{|\varphi(s) + h(s)|_2 + |\varphi(s)|_2} \right] ds$$

is well-defined (because the integrand is bounded by 2) and goes to zero as $\|h\|_\infty$ goes to zero, this proves that $D_\varphi \Gamma$ is the Fréchet derivative of Γ at point φ . The functional T being linear, F is also Fréchet differentiable

with Fréchet derivative given by

$$\begin{aligned} D_\varphi F : \quad (\mathcal{C}([0, 1], \mathbf{R}^3) &\rightarrow \mathbf{R}, \\ h &\mapsto \int_0^1 \frac{(T\varphi(s), Th(s))}{|T\varphi(s)|_2} \frac{E_s}{\|E\|_2^2} ds. \end{aligned}$$

Moreover, let h be an element of $\mathbf{L}^2([0, 1], \mathbf{R}^3)$, we have, since $\|E\|_2^2 = \frac{4}{3\pi}$,

$$\left| F\left(\omega + \int_0^\cdot h(s)ds\right) - F(\omega) \right| \leq \frac{3\pi}{4} \int_0^1 \left\{ \left| \int_0^t h(s)ds \right|_2 + t \left| \int_0^1 h(s)ds \right|_2 \right\} E_t dt.$$

But in the right hand side of the last inequality, we have, using Jensen's inequality,

$$\begin{aligned} \int_0^1 \left\{ \sqrt{\sum_{i=1}^3 \left(\int_0^t h^i(s)ds \right)^2} + t \sqrt{\sum_{i=1}^3 \left(\int_0^1 h^i(s)ds \right)^2} \right\} E_t dt &\leq \int_0^1 \sqrt{\sum_{i=1}^3 \left(\int_0^1 h^i(s)^2 ds \right)} (1+t) E_t dt \\ &= \int_0^1 \|h\|_{\mathbf{L}^2([0,1], \mathbf{R}^3)} (1+t) E_t dt. \end{aligned}$$

From this, using the results of [26, p. 35], we have that F belongs to the space $\mathbf{D}^{1,2}$, which is the domain of the Malliavin operator D in $\mathbf{L}^2([0, 1], \mathbf{R}^3)$ (see [26, pp. 25–27] for more details). Before going further let us recall some facts on Malliavin derivative. When working with the probability space $(\mathcal{C}([0, 1], \mathbf{R}^3), \mathcal{F}, \mathbf{W})$, it is known [26, 1.2.1 The derivative operator in the white noise case] that there exist strong connections between Malliavin derivative and Fréchet derivative for a random variable G of $\mathbf{D}^{1,2}$ defined from $(\mathcal{C}([0, 1], \mathbf{R}^3), \mathcal{F}, \mathbf{W})$ to \mathbf{R} . Since, the Fréchet derivative $D_\omega G$ at point ω of G is a continuous linear form from $\mathcal{C}([0, 1], \mathbf{R}^3)$ into \mathbf{R} , it can be identified to a triple $(\mu_1^\omega, \mu_2^\omega, \mu_3^\omega)$ of σ -finite measures on \mathbf{R} such that,

$$\forall h \in \mathcal{C}([0, 1], \mathbf{R}^3), \quad D_\varphi G h = \sum_{i=1}^3 \int_{[0,1]} h^i(s) \mu_i^\omega(ds).$$

In such a case, the Malliavin derivative of G is the random process belonging to $\mathbf{L}^2([0, 1], \mathbf{R}^3)$ given by

$$\{ (\mu_1^\omega(u, 1], \mu_2^\omega(u, 1], \mu_3^\omega(u, 1]) : u \in [0, 1] \}.$$

In our case, it follows that the Malliavin derivative of F is given by

$$DF(\omega) = \left(\int_0^1 \frac{(\omega_s - s\omega_1)E_s}{|\omega_s - s\omega_1|_2 \|E\|_2^2} (\mathbf{1}_{s>u} - s) ds, \quad u \in [0, 1] \right) \in \mathbf{L}^2([0, 1], \mathbf{R}^3).$$

Now, since DF is not zero in $\mathbf{L}^2([0, 1], \mathbf{R}^3)$ for \mathbf{W} -almost every ω , we get, together with [26, Theorem 2.1.2], the existence of a density for F w.r.t. the Lebesgue measure. The calculation of the variance is derived from the expectation of $\mathbf{e}_t \mathbf{e}_s$, $(s, t) \in [0, 1]^2$, stated in Lemma 9. \square

Remark 11 *The existence of a density was already known [23, 24] for the random variable $\int_0^1 \mathbf{e}_s ds$ but to the best of our knowledge no paper investigates the existence of a density for Λ_∞ . In these papers [23, 24] the study is performed thanks to the analysis of the double Laplace transform*

$$\lambda \mapsto \int_0^\infty \exp(-\lambda t) \mathbf{E} \left[\exp \left(-t \int_0^1 \mathbf{e}_s ds \right) \right] dt.$$

Thanks to the Feynman-Kac formula, the authors express this quantity in terms of Airy functions. Then, they take the inverse of the Laplace transform via analytical methods. Unfortunately, their method does not extend to our case. Indeed, in their case, an expression of the double Laplace transform given above is derived from the Feynman-Kac formula for standard Brownian motion which tells us that the function

$$u(t, x) = \mathbf{E}_x \left[f(B_t) \exp \left(\int_0^t B_s ds \right) \right], \quad \forall (t, x) \in \mathbf{R}_+ \times \mathbf{R},$$

is solution of the partial differential equation

$$\begin{cases} \partial_t u(t, x) = \frac{1}{2} \Delta u(t, x) + xu(t, x) & \forall x \in \mathbf{R}, t \in \mathbf{R}_+, \\ u(0, x) = f(x) & \forall x \in \mathbf{R}. \end{cases}$$

In this case, taking the Laplace transform in time of u leads to an ordinary differential equation whose solution can be expressed in terms of Airy functions [15]. In our problem, this partial differential equation becomes inhomogeneous in time which prevents us to use this Laplace transform. As a consequence, we think that one can not obtain information by this method. This is why we have established that Λ_∞ admits a density using Malliavin calculus and the representation of the Brownian excursion as a three-dimensional Bessel bridge (3).

Of course, $\widehat{\lambda}[\tau_n]$ is not a consistent estimator of σ^{-1} but it should be noted that its weak limit is unbiased by (7) and Corollary 7. The expression (7) of the variance of Λ_∞ is an explicit but quite intractable formula. Nevertheless, it may be at least evaluated numerically to compute the variance of Λ_∞ . Otherwise, we can also use Monte Carlo simulations to produce a sample with the same law as Λ_∞ to achieve this task. Both methods lead to

$$\mathbf{Var}(\Lambda_\infty) \simeq 0.0690785.$$

At this point, it is quite interesting to compare our approach to the one developed in [4]. The authors of [4] construct estimators for the variance of the offspring distribution of a forest of conditioned critical Galton-Watson trees. Their strategy relies on the distance to the root of a uniformly sampled node v of the considered tree $\tau_n \sim \text{GW}_n(\mu)$,

$$\delta[\tau_n] = \frac{h(v)}{\sqrt{n}}. \quad (8)$$

Using Theorem 3, it has been shown that $\delta[\tau_n]$ converges in law, when the number of nodes n goes to infinity, towards $\sigma^{-1} \Delta_\infty$ where the random variable Δ_∞ follows the Rayleigh distribution with scale 1 [4, Proposition 4] with density,

$$\forall x \in \mathbf{R}_+, f_{\Delta_\infty}(x) = x \exp\left(-\frac{1}{2}x^2\right).$$

We emphasize that $\delta[\tau_n]$ is somehow biased because $\mathbf{E}[\Delta_\infty] = \sqrt{\frac{\pi}{2}} \neq 1$. Nevertheless, one may avoid this issue by considering the quantity

$$\widehat{\delta}[\tau_n] = \sqrt{\frac{2}{\pi}} \delta[\tau_n] \quad (9)$$

that converges to $\sigma^{-1} \sqrt{\frac{2}{\pi}} \Delta_\infty$ which is σ^{-1} on average. As a consequence, $\widehat{\lambda}[\tau_n]$ and $\widehat{\delta}[\tau_n]$ are two quantities directly computable from the tree τ_n and that may be used to construct an estimator of the inverse standard deviation of interest. We propose to compare them from their respective asymptotic dispersion which should be as small as possible in order to get an accurate estimator. A first comparison may be done by computing the variances of Λ_∞ and $\sqrt{\frac{2}{\pi}} \Delta_\infty$. One has

$$\mathbf{Var}\left(\sqrt{\frac{2}{\pi}} \Delta_\infty\right) \simeq 0.2732395 \quad \text{and} \quad \mathbf{Var}(\Lambda_\infty) \simeq 0.0690785.$$

This difference in the dispersions is quite apparent in Figure 3 where the densities of $\sqrt{\frac{2}{\pi}} \Delta_\infty$ and Λ_∞ have been displayed. Consequently, one may expect better results in terms of dispersion from our approach.




Figure 3: Densities of $\sqrt{\frac{2}{\pi}}\Delta_\infty$ (full line) where Δ_∞ follows the Rayleigh distribution given by $f(x) = \frac{\pi}{2}x \exp\left(-\frac{\pi x^2}{4}\right)$ for $x \in \mathbf{R}_+$ and of Λ_∞ (dashed line) estimated from 1 000 000 simulated Brownian excursions.

3.2 Interpretation in functional principal component analysis

On the suggestion of a reviewer, we performed a functional principal component analysis (FPCA) from normalized Harris paths of large conditioned Galton-Watson trees with different values of σ . We refer the reader to [17] and the references therein for explanations on this statistical tool and to Subsection 5.1 for simulation methods. The FPCA has been carried out with function `FPCA` from the R package `fdapace`. The results are presented in Figure 4.

These numerical experiments show that the first eigenfunction is very close to the average Brownian excursion (see Figure 4 (bottom right)). As a consequence, $\hat{\lambda}[\tau]$ can be interpreted as the first eigenvalue in FPCA of the normalized Harris path of τ . The first eigenspace expresses approximately 70% of the total dataset inertia. For the sake of comparison, the second eigenspace gets only 7% of the inertia. In addition, one can see on Figure 4 (top left, top right and bottom left) that only the first dimension of FPCA captures information on the value of σ . These results highlight that $\hat{\lambda}[\tau]$ is a relevant quantity in our estimation problem as well as one can not expect more significant information from the projection on the other eigenspaces.

Alternative functionals of the Brownian excursion can be investigated in order to develop statistical methods for conditioned Galton-Watson trees. For instance, the vector of peaks and valleys at random [8] or fixed [27] times could be considered. If the distribution of such objects is simple enough, this could enable the use of maximum likelihood methods. However, as shown above, the functional $\hat{\lambda}[\tau]$ considered in our work seems to be one of the best in our setting in terms of quantity of information.

3.3 Estimation strategies

In this section, we give details on two ideas in order to estimate σ^{-1} from a forest of conditioned Galton-Watson trees. A forest is defined as a tuple of trees. Let N be a positive integer. In this section, we consider a forest \mathcal{F} made of N independent trees τ^1, \dots, τ^N with respective sizes n_1, \dots, n_N and respective laws $\text{GW}_{n_1}(\mu), \dots, \text{GW}_{n_N}(\mu)$.

3.3.1 Least square estimation

This first strategy lies on the goodness of fit between the Harris path of the forest with the expected limiting contour. This adequacy is measured thanks to an $\mathbf{L}^2([0, N], \mathbf{R})$ -norm. More precisely, we denote

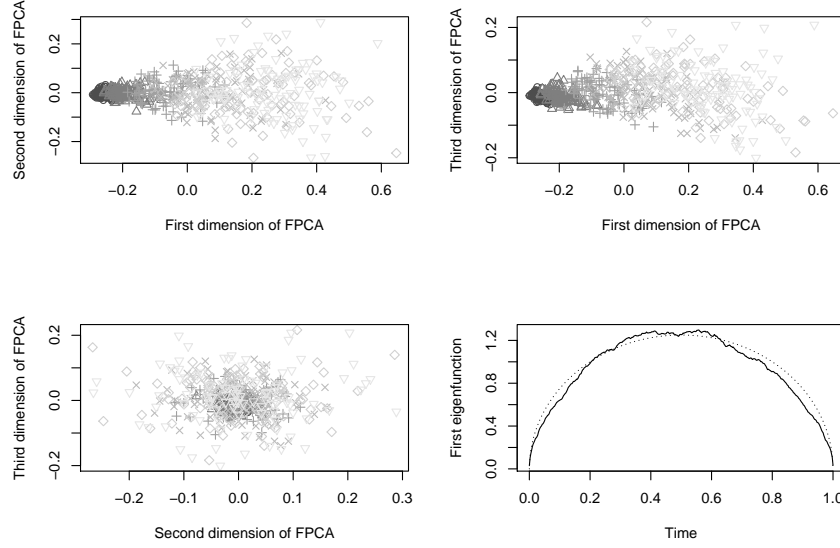


Figure 4: FPCA performed from 600 large conditioned Galton-Watson trees with different standard deviations: from dark to bright gray, $\sigma \in \{0.1, 0.2, 0.4, 0.6, 0.8, 0.95\}$. Coefficients of projection on the three first eigenfunctions (top left, top right and bottom left) and first eigenfunction (full line, bottom right) compared to the average Brownian excursion (dashed line).

$(\mathcal{H}[\mathcal{F}](t), t \in [0, N])$ the Harris path of the forest \mathcal{F} . This process is defined by

$$\forall 0 \leq t \leq N, \mathcal{H}[\mathcal{F}](t) = \sum_{i=1}^N \frac{1}{\sqrt{n_i}} \mathcal{H}[\tau^i](2n_i(t - i + 1)) \mathbf{1}_{[i-1, i)}(t).$$

The Harris path of a forest is simply the concatenation of the Harris paths of each tree, in the natural order. We propose to estimate σ^{-1} by $\hat{\lambda}_{ls}[\mathcal{F}]$ that minimizes the $\mathbf{L}^2([0, N], \mathbf{R})$ -error

$$\lambda \mapsto \|\mathcal{H}[\mathcal{F}](\cdot) - \lambda H(\cdot - \lfloor \cdot \rfloor)\|_{\mathbf{L}^2([0, N], \mathbf{R})}^2,$$

the function $H(\cdot - \lfloor \cdot \rfloor)$ mapping $x \in [0, N]$ to $H(x - \lfloor x \rfloor)$. As aforementioned in (5), $\hat{\lambda}_{ls}[\mathcal{F}]$ can be explicitly computed. Indeed, one can check that

$$\hat{\lambda}_{ls}[\mathcal{F}] = \frac{\langle \mathcal{H}[\mathcal{F}](\cdot), H(\cdot - \lfloor \cdot \rfloor) \rangle}{\|H(\cdot - \lfloor \cdot \rfloor)\|_2^2}.$$

We remark that $\hat{\lambda}_{ls}[\mathcal{F}]$ is only the average of the quantities $\hat{\lambda}[\tau^i]$ (defined in (5)),

$$\hat{\lambda}_{ls}[\mathcal{F}] = \frac{1}{N} \sum_{i=1}^N \hat{\lambda}[\tau^i].$$

Thus, according to Theorems 4 and 7, one can expect that $\hat{\lambda}_{ls}[\mathcal{F}]$ tends to σ^{-1} in some sense, when both N and n_i go to infinity, by virtue of the law of large numbers.

3.3.2 Estimation by minimal Wasserstein distance

In the preceding method, we did not use our knowledge of the limiting distribution of the random variable of type $\lambda[\tau^n]$. In order to take this into account, one may want to test the goodness of fit between the empirical

measure $\widehat{\mathcal{P}}$ defined by

$$\widehat{\mathcal{P}} = \frac{1}{N} \sum_{i=1}^N \delta_{\widehat{\lambda}[\tau^i]} \quad (10)$$

and the law of Λ_∞ . Using Wasserstein metrics to align distributions is rather natural since it corresponds to the transportation cost between two probability laws. In particular, this feature appears to be useful in a statistical framework [6, 12]. In our case, $\widehat{\mathcal{P}}$ is expected to be close in terms of Wasserstein distance to $\sigma^{-1}\Lambda_\infty$ in the asymptotic regime of an infinite forest of infinite trees. That is why, we propose to estimate σ^{-1} with the real number λ which minimizes the distance between $\widehat{\mathcal{P}}$ and $\lambda\Lambda_\infty$. More precisely, our estimator $\widehat{\lambda}_W[\mathcal{F}]$ is defined by

$$\widehat{\lambda}_W[\mathcal{F}] = \arg \min_{\lambda > 0} d_W \left(\widehat{\mathcal{P}}, \mathbf{P}_{\lambda\Lambda_\infty} \right),$$

where d_W denotes the Wasserstein distance of order 2 and $\mathbf{P}_{\lambda\Lambda_\infty}$ denotes the law of $\lambda\Lambda_\infty$.

The Wasserstein distance of order 2, denoted $d_W(\nu_1, \nu_2)$, between two probability measures ν_1 and ν_2 on \mathbf{R} can be defined from their cumulative distribution functions F_1 and F_2 as follows,

$$d_W(\nu_1, \nu_2) = \|F_1^{-1} - F_2^{-1}\|_2. \quad (11)$$

Let \widehat{F} be the cumulative function of the empirical measure $\widehat{\mathcal{P}}$, while $F_{\lambda\Lambda_\infty}$ stands for the cumulative function of the random variable $\lambda\Lambda_\infty$. As a consequence of (11), one has

$$\begin{aligned} d_W \left(\frac{1}{N} \sum_{i=1}^N \delta_{\widehat{\lambda}[\tau^i]}, \mathbf{P}_{\lambda\Lambda_\infty} \right)^2 &= \int_0^1 \left(\widehat{F}^{-1}(s) - F_{\lambda\Lambda_\infty}^{-1}(s) \right)^2 ds \\ &= \int_0^1 \left(\widehat{F}^{-1}(s) - \lambda F_{\Lambda_\infty}^{-1}(s) \right)^2 ds, \end{aligned}$$

thanks to the fact that $F_{\lambda\Lambda_\infty}^{-1} = \lambda F_{\Lambda_\infty}^{-1}$. It follows that minimizing the Wasserstein distance boils down to solving a least square minimization problem. Hence, it comes that

$$\begin{aligned} \widehat{\lambda}_W[\mathcal{F}] &= \frac{\langle \widehat{F}^{-1}, F_{\Lambda_\infty}^{-1} \rangle}{\|F_{\Lambda_\infty}^{-1}\|_2^2} \\ &= \frac{1}{\|F_{\Lambda_\infty}^{-1}\|_2^2} \sum_{i=1}^N \widehat{\lambda}[\tau^{(i)}] \int_{\frac{i-1}{N}}^{\frac{i}{N}} F_{\Lambda_\infty}^{-1}(s) ds, \end{aligned} \quad (12)$$

where $(\widehat{\lambda}[\tau^{(i)}])_{1 \leq i \leq N}$ denotes the order statistic associated with the family $(\widehat{\lambda}[\tau^i])_{1 \leq i \leq N}$.

Remark 12 We point out the fact that there is no problem of definition in the above quantities because both \widehat{F}^{-1} and $F_{\Lambda_\infty}^{-1}$ belong to \mathbf{L}^2 . In the first case, this follows from the fact that \widehat{F}^{-1} is bounded (because $\widehat{\mathcal{P}}$ has compact support). For $F_{\Lambda_\infty}^{-1}$, this comes from the uniform sampling principle which entails that

$$\int_0^1 F_{\Lambda_\infty}^{-1}(u)^2 du = \mathbf{E}[\Lambda_\infty^2].$$

Remark 13 The proposed methodology consists in identifying the best parameter λ that allows to align the distributions $\widehat{\mathcal{P}}$ and $\mathbf{P}_{\lambda\Lambda_\infty}$. The Wasserstein distance is well-adapted to this problem because it is computed from the inverse cumulative distribution functions together with the fact that $F_{\lambda\Lambda_\infty}^{-1} = \lambda F_{\Lambda_\infty}^{-1}$. As a consequence, one may get the optimal parameter $\widehat{\lambda}_W[\mathcal{F}]$ from only a numerical estimate of $F_{\Lambda_\infty}^{-1}$. The same trick does not hold for the maximum likelihood method: one can not express the likelihood of $\lambda\Lambda_\infty$ as a function of the two variables λ and f_{Λ_∞} . Thus this alternative method is not adequate without an explicit formula for f_{Λ_∞} , which seems to be out of our reach.

4 Main results

4.1 Increasing sequences of random forests

Before going further, the statistical framework needs to be precisely formulated. In the sequel, the set of integer sequences is denoted by \mathbf{S} . For any positive real number A , we denote by \mathbf{S}_A the subset of \mathbf{S} defined by

$$\mathbf{S}_A = \left\{ u \in \mathbf{S} : \min_{i \geq 1} u_i \geq A \right\}.$$

In addition, for any sequence u in \mathbf{S} and any positive integer N , \vec{u}_N is the multi-integer made of the N first components of u , that is

$$\vec{u}_N = (u_1, \dots, u_N).$$

Now, let us introduce our probabilistic framework. Let $(\tau_n^k)_{n,k \geq 1}$ be a family of independent conditioned Galton-Watson trees such that, for a given n , the family $(\tau_n^k)_{k \geq 1}$ is i.i.d. $\text{GW}_n(\mu)$. From this family, we define, for any multi-integer $\vec{u}_N = (u_1, \dots, u_N)$, the random forest $\mathcal{F}_{\vec{u}_N}$ made of the trees $(\tau_{u_1}^1, \dots, \tau_{u_N}^N)$.

The idea of this construction is to consider increasing (in the sense of inclusion) sequences of random forests. Indeed, assume we are given a sequence $(u_n)_{n \geq 1}$ of integers (corresponding with the sizes of our trees), then the N first trees of the forest $\mathcal{F}_{\vec{u}_{N+1}}$ are the same as the trees of the forest $\mathcal{F}_{\vec{u}_N}$.

We point out that the hypothesis of independence may be thought to be too strong in some applications. Exchangeability is a weaker assumption that could be considered. In such a statistical setting, the reference [13] is particularly relevant.

4.2 Least square estimation

This first result focuses on the large trees regime and gives the asymptotic unbiasedness of the least square estimator in this regime.

Proposition 14 *The least square estimator is asymptotically unbiased in the large trees regime, that is*

$$\forall \epsilon > 0, \exists A \in \mathbf{N}, \forall u \in \mathbf{S}_A, \forall N \in \mathbf{N} \quad \left| \mathbf{E} \left[\hat{\lambda}_{ls}[\mathcal{F}_{\vec{u}_N}] \right] - \sigma^{-1} \right| < \epsilon.$$

This means that the expectation of the least square estimator converges to σ^{-1} as the sizes of the trees increase.

Proof. Since the family $(\tau_{u_i}^i)_{1 \leq i \leq N}$ is made of independent random variables, it follows from Theorem 4 and the definition (5) of $\hat{\lambda}[\tau_{n_i}]$ that the proof of this last statement boils down to proving that, when n goes to infinity,

$$\int_0^1 \mathbf{E} \left[\frac{\mathcal{H}[\tau_n](2ns)}{\sqrt{n}} \right] E_s \, ds \longrightarrow \frac{2}{\sigma} \int_0^1 E_s^2 \, ds,$$

where τ_n is some tree with law $\text{GW}_n(\mu)$. It is known from [10, Lemma 4] that, for any positive integer n and real number $0 < t < 1$,

$$\forall x \in \mathbf{R}_+, \mathbf{P} \left(\frac{\mathcal{H}[\tau_n](2nt)}{\sqrt{n}} > x \right) \leq \frac{C}{t} \exp \left(\frac{-Dx}{\sqrt{t}} \right). \quad (13)$$

From this last estimate, one can easily show that $\mathbf{E} \left[\frac{\mathcal{H}[\tau_n](2n \cdot)}{\sqrt{n}} \right]$ is uniformly bounded (w.r.t. n) by an integrable function. Finally, the result follows from Theorem 4 and the dominated convergence theorem. \square

The spirit of the following result is that, given an increasing sequence of random forests, the least square estimator can not be too far from σ^{-1} as soon as the sizes of the trees are large enough. In particular, due to the weakness of the convergence of conditioned Galton-Watson trees given in Theorem 4, one can not expect a stronger result of convergence.

Proposition 15 *We have,*

$$\forall \epsilon > 0, \exists A \in \mathbf{N}, \forall u \in \mathbf{S}_A, \mathbf{P} \left(\limsup_{N \rightarrow \infty} \left| \widehat{\lambda}_{ls}[\mathcal{F}_{\bar{u}_N}] - \sigma^{-1} \right| < \epsilon \right) = 1.$$

Proof. We begin the proof by showing that the family $(\widehat{\lambda}[\tau_n^k])_{n,k \geq 1}$ has uniformly bounded fourth moments. By Jensen's inequality, there exists a positive constant c such that

$$\begin{aligned} \mathbf{E} \left[\left(\widehat{\lambda}[\tau_n^i] \right)^4 \right] &\leq c \int_0^1 \mathbf{E} \left[\left(\frac{\mathcal{H}[\tau_n^i](2ns)}{\sqrt{n}} \right)^4 \right] ds \\ &= 4c \int_0^1 \int_{\mathbf{R}_+} x^3 \mathbf{P} \left(\frac{\mathcal{H}[\tau_n^i](2ns)}{\sqrt{n}} > x \right) dx ds. \end{aligned} \quad (14)$$

Finally, using again equation (13) gives the desired bound,

$$\mathbf{E} \left[\widehat{\lambda}[\tau_n^i]^4 \right] \leq \frac{12cC}{D^4}. \quad (15)$$

From this point we consider a sequence u of integers. This sequence corresponds to the sizes of the trees in our increasing sequence of random forests $(\mathcal{F}_{\bar{u}_N})_{N \geq 1}$. We recall according to the definitions given in the beginning of this section that the random forest $\mathcal{F}_{\bar{u}_N}$ is composed of the trees $(\tau_{u_1}^1, \dots, \tau_{u_N}^N)$.

Let $m_{u_i}^i$ be the expectation of $\widehat{\lambda}[\tau_{u_i}^i]$. It is worth noting that this expectation depends only on the integer u_i . Now, using the uniform bound on the fourth moment (15), Markov's inequality applied to the fourth power of

$$\frac{1}{N} \sum_{i=1}^N \left(\widehat{\lambda}[\tau_{u_i}^i] - m_{u_i}^i \right)$$

gives the convergence in probability of the above sum to zero at rate N^{-2} which implies, in light of the Borel-Cantelli lemma, that

$$\frac{1}{N} \sum_{i=1}^N \left(\widehat{\lambda}[\tau_{u_i}^i] - m_{u_i}^i \right) \xrightarrow{a.s.} 0, \quad (16)$$

when N goes to infinity. Moreover, using Theorem 4, we have that $m_{u_i}^i$ converges to σ^{-1} as u_i goes to infinity, from which it follows that for any $\epsilon > 0$, there exists an integer A such that

$$|m_{u_i}^i - \sigma^{-1}| < \epsilon, \quad (17)$$

whenever $u_i > A$. Finally, letting all the u_i 's be greater than A , we have that there exists a measurable set Ω_u , with mass 1, such that, using (16) and (17), for all ω in this set,

$$\limsup_{N \rightarrow \infty} \left| \frac{1}{N} \sum_{i=1}^N \widehat{\lambda}[\tau_{u_i}^i](\omega) - \sigma^{-1} \right| \leq \limsup_{N \rightarrow \infty} \frac{1}{N} \left| \sum_{i=1}^N \widehat{\lambda}[\tau_{u_i}^i](\omega) - m_{u_i}^i \right| + \limsup_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N |m_{u_i}^i - \sigma^{-1}| \leq \epsilon,$$

which establishes the expected convergence. \square

Remark 16 *According to the proof of the preceding theorem, it would be very interesting to control the rate of convergence in Theorem 4. Indeed, this would enable us to get a control of the error in the convergence stated in Proposition 15 given in terms of the smallest tree in the increasing sequence of random forests.*

Remark 17 *Let us point out that equation (13) gives the exponential decay of the tail distribution of $n^{-1/2}\mathcal{H}[\tau_n](2nt)$ uniformly w.r.t. n . In particular, one can apply the method used in equation (14) to obtain uniform (w.r.t. n) bounds like*

$$\mathbf{E} \left[\widehat{\lambda}[\tau_n]^k \right] \leq \frac{(k-1)!cC}{D^k},$$

for any positive integer k and some positive constant c .

4.3 Estimation by minimal Wasserstein distance

As in the preceding section, we begin by looking at the asymptotic bias of the considered estimator.

Proposition 18 *The Wasserstein estimator is asymptotically unbiased in the large trees and large forests regime. That is,*

$$\forall \epsilon > 0, \exists (\mathfrak{N}, A) \in \mathbf{N}^2, \forall u \in \mathbf{S}_A, \forall N \geq \mathfrak{N}, \left| \mathbf{E} \left[\hat{\lambda}_W[\mathcal{F}_{\vec{u}_N}] \right] - \sigma^{-1} \right| < \epsilon.$$

Proof. Let u in \mathbf{S} and N in \mathbf{N} . Let $(\Lambda_{\infty, i})_{i \geq 1}$ be an i.i.d. sequence of random variables with the same distribution as Λ_∞ . The first step of the proof is to show that

$$\eta_{\vec{u}_N} = \left| \mathbf{E} \left[\hat{\lambda}_W[\mathcal{F}_{\vec{u}_N}] - \frac{1}{\sigma \|F_{\Lambda_\infty}^{-1}\|_2^2} \sum_{i=1}^N \mathbf{E}[\Lambda_{\infty, (i)}] \int_{\frac{i-1}{N}}^{\frac{i}{N}} F_{\Lambda_\infty}^{-1}(s) ds \right] \right|$$

converges to 0 as $\min(\vec{u}_N)$ goes to infinity uniformly w.r.t. N , where the order statistic in the above formula has to be understood w.r.t. the random vector $(\Lambda_{\infty, 1}, \dots, \Lambda_{\infty, N})$. The convergence easily follows from the results of [10]. However, getting it uniformly w.r.t. N requires to give some new insights. To this end let $(\mathcal{U}_i)_{i \geq 1}$ be a sequence of i.i.d. random variables with uniform distribution on $[0, 1]$. Denote, for any positive integer n , F_n^{-1} the right inverse of the cumulative distribution function associated to the random variable $\hat{\lambda}[\tau_n^i]$. According to the right inverse principle, the random vector $(F_{u_i}^{-1}(\mathcal{U}_i))_{1 \leq i \leq N}$ is equal in distribution to $(\hat{\lambda}[\tau_{u_i}^i])_{1 \leq i \leq N}$ and the vector $(F_{\Lambda_\infty}^{-1}(\mathcal{U}_i))_{1 \leq i \leq N}$ is a vector of i.i.d. random variables with the same distribution as Λ_∞ . At this point, let us highlight that the sequences $(\hat{\lambda}[\tau_n^i])_{n \geq 1}$ and $(F_n^{-1}(\mathcal{U}_i))_{n \geq 1}$ do not have the same distribution since the first one is made of independent random variables whereas this is clearly not the case for the second one. However, this feature is not important since, in the above formula, we only look at averaged behaviors. First note that, by definition of the Wasserstein estimator and the right inverse sampling principle, we have

$$\eta_{\vec{u}_N} = \frac{1}{\sigma \|F_{\Lambda_\infty}^{-1}\|_2^2} \left| \int_0^1 F_{\Lambda_\infty}^{-1}(s) \mathbf{E} \left[\hat{G}_{\vec{u}_N}^{-1}(s) - \hat{H}_N^{-1}(s) \right] ds \right|, \quad (18)$$

where $\hat{G}_{\vec{u}_N}^{-1}$ and \hat{H}_N^{-1} denote the inverse distribution functions of the empirical measures respectively associated to the vectors $(F_{u_i}^{-1}(\mathcal{U}_i))_{1 \leq i \leq N}$ and $(F_{\Lambda_\infty}^{-1}(\mathcal{U}_i))_{1 \leq i \leq N}$. Now, the Cauchy-Schwartz inequality entails that

$$\left| \int_0^1 F_{\Lambda_\infty}^{-1}(s) \mathbf{E} \left[\hat{G}_{\vec{u}_N}^{-1}(s) - \hat{H}_N^{-1}(s) \right] ds \right| \leq \|F_{\Lambda_\infty}^{-1}\|_2 \sqrt{\int_0^1 \mathbf{E} \left[\hat{G}_{\vec{u}_N}^{-1}(s) - \hat{H}_N^{-1}(s) \right]^2 ds}.$$

By the definition of the inverse distribution function, we get

$$\int_0^1 \mathbf{E} \left[\hat{G}_{\vec{u}_N}^{-1}(s) - \hat{H}_N^{-1}(s) \right]^2 ds = \int_0^1 \left(\sum_{i=1}^N \mathbf{1}_{[\frac{i-1}{N}, \frac{i}{N})}(s) \mathbf{E} \left[F_{u_{(i)}}^{-1}(\mathcal{U}_{(i)}) - F_{\Lambda_\infty}^{-1}(\mathcal{U}_{(i)}) \right] \right)^2 ds,$$

where $(F_{n_{(i)}}^{-1}(\mathcal{U}_{(i)}))_{1 \leq i \leq N}$ and $(F_{\Lambda_\infty}^{-1}(\mathcal{U}_{(i)}))_{1 \leq i \leq N}$ denote the order statistics respectively associated with the vectors $(F_{n_i}^{-1}(\mathcal{U}_i))_{1 \leq i \leq N}$ and $(F_{\Lambda_\infty}^{-1}(\mathcal{U}_i))_{1 \leq i \leq N}$. Using two times Jensen's inequality leads to

$$\int_0^1 \mathbf{E} \left[\hat{G}_{\vec{u}_N}^{-1}(s) - \hat{H}_N^{-1}(s) \right]^2 ds \leq \frac{1}{N} \sum_{i=1}^N \mathbf{E} \left[\left| F_{n_{(i)}}^{-1}(\mathcal{U}_{(i)}) - F_{\Lambda_\infty}^{-1}(\mathcal{U}_{(i)}) \right|^2 \right].$$

Finally, using that the order function $(x_1, \dots, x_n) \mapsto (x_{(1)}, \dots, x_{(n)})$ is 1-Lipschitz w.r.t. the Euclidean norm (as a consequence of the rearrangement inequality), we have

$$\left| \int_0^1 F_{\Lambda_\infty}^{-1}(s) \mathbf{E} \left[\hat{G}_{\vec{u}_N}^{-1}(s) - \hat{H}_N^{-1}(s) \right] ds \right| \leq \|F_{\Lambda_\infty}^{-1}\|_2 \sqrt{\frac{1}{N} \sum_{i=1}^N \mathbf{E} \left[\left| F_{n_i}^{-1}(\mathcal{U}_i) - F_{\Lambda_\infty}^{-1}(\mathcal{U}_i) \right|^2 \right]}. \quad (19)$$

Now by construction, for any i , $F_n^{-1}(\mathcal{U}_i)$ converges almost surely to $F_{\Lambda_\infty}^{-1}(\mathcal{U}_i)$. Moreover, the uniform square-integrability of the laws of the $\widehat{\lambda}[\tau_{n_i}^i]$'s provided by Remark 17 gives that

$$\lim_{n \rightarrow \infty} \mathbf{E} \left[\left| F_n^{-1}(\mathcal{U}_i) - F_{\Lambda_\infty}^{-1}(\mathcal{U}_i) \right|^2 \right] = 0.$$

Now, let $0 < \epsilon < 1$ be some positive real number and let A such that for any $n \geq A$, we have

$$\mathbf{E} \left[\left| F_n^{-1}(\mathcal{U}_i) - F_{\Lambda_\infty}^{-1}(\mathcal{U}_i) \right|^2 \right] \leq \sigma^2 \|F_{\Lambda_\infty}^{-1}\|_2^2 \epsilon^2.$$

Hence, as soon as $\min(u) \geq A$, we have, together with (19),

$$\frac{1}{\sigma \|F_{\Lambda_\infty}^{-1}\|_2^2} \left| \int_0^1 F_{\Lambda_\infty}^{-1}(s) \mathbf{E} \left[\widehat{G}_{\bar{u}_N}^{-1}(s) - \widehat{H}_N^{-1}(s) \right] ds \right| \leq \epsilon.$$

Finally, (18) gives the desired uniform convergence. It remains to prove that

$$\frac{1}{\sigma \|F_{\Lambda_\infty}^{-1}\|_2^2} \sum_{i=1}^N \mathbf{E} [\Lambda_{\infty, (i)}] \int_{\frac{i-1}{N}}^{\frac{i}{N}} F_{\Lambda_\infty}^{-1}(s) ds$$

converges to σ^{-1} as N goes to infinity. It is well known, since Λ_∞ has a density, that, for any $1 \leq i \leq N$, one has (see for instance [7])

$$\mathbf{E} [\Lambda_{\infty, (i)}] = N \binom{N-1}{i-1} \int_0^\infty x F_{\Lambda_\infty}(x)^{i-1} (1 - F_{\Lambda_\infty}(x))^{N-i} f_{\Lambda_\infty}(x) dx.$$

Hence,

$$\begin{aligned} \mathbf{E} \left[\sum_{i=1}^N \Lambda_{\infty, (i)} \int_{\frac{i-1}{N}}^{\frac{i}{N}} F_{\Lambda_\infty}^{-1}(s) ds \right] \\ = N \int_0^\infty x f_{\Lambda_\infty}(x) \sum_{i=1}^N \binom{N-1}{i-1} F_{\Lambda_\infty}(x)^{i-1} (1 - F_{\Lambda_\infty}(x))^{N-i} \int_0^{\frac{1}{N}} F_{\Lambda_\infty}^{-1} \left(s + \frac{i-1}{N} \right) ds dx. \end{aligned}$$

This rewrites thanks to the right inverse sampling principle as

$$\mathbf{E} \left[\sum_{i=1}^N \Lambda_{\infty, (i)} \int_{\frac{i-1}{N}}^{\frac{i}{N}} F_{\Lambda_\infty}^{-1}(s) ds \right] = \int_0^1 F_{\Lambda_\infty}^{-1}(y) K_N(F_{\Lambda_\infty}^{-1})(y) dy,$$

where K_N is defined for all function φ in \mathbf{L}^2 by

$$K_N(\varphi)(y) = N \sum_{i=1}^N \binom{N-1}{i-1} y^{i-1} (1-y)^{N-i} \int_0^{\frac{1}{N}} \varphi \left(s + \frac{i-1}{N} \right) ds, \quad \forall y \in [0, 1].$$

The operators K_N are known as Bernstein-Kantorovich operators which were introduced in the 30's by Kantorovich in order to extend the properties of Bernstein polynomials to non-continuous functions [18]. In particular, it is known that, for all φ in \mathbf{L}^2 , $K_N(\varphi)$ converges to φ in \mathbf{L}^2 [22, Theorem 2.1.2 and p. 33]. Now, according to the Cauchy-Schwarz inequality we have that

$$\left| \int_0^1 F_{\Lambda_\infty}^{-1}(y) K_N(F_{\Lambda_\infty}^{-1})(y) dy - \int_0^1 F_{\Lambda_\infty}^{-1}(y)^2 dy \right| \leq \|F_{\Lambda_\infty}^{-1}\|_2 \|K_N(F_{\Lambda_\infty}^{-1}) - F_{\Lambda_\infty}^{-1}\|_2.$$

But since $K_N(F_{\Lambda_\infty}^{-1})$ converges to $F_{\Lambda_\infty}^{-1}$ in \mathbf{L}^2 , we finally obtain

$$\mathbf{E} \left[\sum_{i=1}^N \Lambda_{\infty, (i)} \int_{\frac{i-1}{N}}^{\frac{i}{N}} F_{\Lambda_\infty}^{-1}(s) ds \right] \longrightarrow \|F_{\Lambda_\infty}^{-1}\|_2^2,$$

when N goes to infinity. This gives the result. \square

We also have a stronger convergence result for this estimator. It relies on the fact that the empirical measure $\widehat{\mathcal{P}}$ defined in (10) must be close (in Wasserstein distance) to the law of $\sigma^{-1}\Lambda_\infty$ as soon as the trees are large enough. More precisely, we have the following result of consistency.

Proposition 19 *Let \mathcal{P} be the law of $\sigma^{-1}\Lambda_\infty$. Let also $\widehat{\mathcal{P}}_{\vec{u}_N}$ be the empirical distribution defined for any multi-integer \vec{u}_N by*

$$\widehat{\mathcal{P}}_{\vec{u}_N} = \frac{1}{N} \sum_{i=1}^N \delta_{\widehat{\lambda}[\tau_{u_i}^i]}.$$

Then, the following statement holds,

$$\forall \epsilon > 0, \exists A \in \mathbf{N}, \forall u \in \mathbf{S}_A, \quad \mathbf{P} \left(\limsup_{N \rightarrow \infty} d_W \left(\widehat{\mathcal{P}}_{\vec{u}_N}, \mathcal{P} \right) < \epsilon \right) = 1.$$

Proof. Let Π_δ be the canonical projection of \mathbf{R} on $[-\delta, \delta]$, for a positive real number δ . We have

$$d_W \left(\widehat{\mathcal{P}}_{\vec{u}_N}(\omega), \mathcal{P} \right) \leq d_W \left(\widehat{\mathcal{P}}_{\vec{u}_N}(\omega), \Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega) \right) + d_W \left(\Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega), \Pi_\delta \mathcal{P} \right) + d_W(\mathcal{P}, \Pi_\delta \mathcal{P}), \quad (20)$$

where $\Pi_\delta \mu$ denotes the image measure of μ by Π_δ . To obtain the desired result, we need to control each of the three terms in the right hand side of (20).

Third term. First, it is clear, for any probability measure μ , that Π_δ is a transport of μ on $\Pi_\delta \mu$ which need not be optimal [5, 2. Generalities on Kantorovich transport distances]. Hence,

$$d_W(\mu, \Pi_\delta \mu) \leq \sqrt{\int_{\mathbf{R}} |x - \Pi_\delta(x)|^2 \mu(dx)}.$$

It follows, since $x \mapsto x^2$ is integrable w.r.t. \mathcal{P} , that δ can be chosen in order to have

$$d_W(\mathcal{P}, \Pi_\delta \mathcal{P}) \leq \sqrt{\mathbf{E} \left[(\sigma^{-1}\Lambda_\infty)^2 \mathbf{1}_{|\sigma^{-1}\Lambda_\infty| > \delta} \right]} < \frac{\epsilon}{3}. \quad (21)$$

First term. On the other hand, following the same lines as in the proof of Proposition 15 (and using Remark 17), one can show that, for any $\epsilon > 0$,

$$\exists A \in \mathbf{N}, \forall u \in \mathbf{S}_A, \quad \mathbf{P} \left(\limsup_{N \rightarrow \infty} \left| \frac{1}{N} \sum_{i=1}^N \widehat{\lambda}[\tau_{u_i}^i]^2 \mathbf{1}_{|\widehat{\lambda}[\tau_{u_i}^i]| > \delta} - \mathbf{E} \left[\sigma^{-2} \Lambda_\infty^2 \mathbf{1}_{|\sigma^{-1}\Lambda_\infty| > \delta} \right] \right| < \epsilon \right) = 1. \quad (22)$$

This bound allows us to control the first term in the right hand side of (20) since

$$\begin{aligned} d_W \left(\widehat{\mathcal{P}}_{\vec{u}_N}(\omega), \Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega) \right) &\leq \sqrt{\int_{\mathbf{R}} |x - \Pi_\delta(x)|^2 \widehat{\mathcal{P}}_{\vec{u}_N}(\omega)(dx)} \\ &\leq \sqrt{\frac{1}{N} \sum_{i=1}^N \widehat{\lambda}[\tau_{u_i}^i](\omega)^2 \mathbf{1}_{|\widehat{\lambda}[\tau_{u_i}^i](\omega)| > \delta}}. \end{aligned}$$

Hence, it remains to control the second term.

Second term. Since $\Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega)$ and $\Pi_\delta \mathcal{P}$ are compactly supported measures, we have

$$d_W \left(\Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega), \Pi_\delta \mathcal{P} \right) \leq C \sqrt{d_W^{(1)} \left(\Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega), \Pi_\delta \mathcal{P} \right)},$$

where $d_W^{(1)}$ is the first order Wasserstein metric. As a consequence, if one gets the result for $d_W^{(1)}$, it gives the result for d_W . First of all, we have

$$d_W^{(1)} \left(\Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega), \Pi_\delta \mathcal{P} \right) \leq d_W^{(1)} \left(\Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega), \overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}} \right) + d_W^{(1)} \left(\overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}}, \Pi_\delta \mathcal{P} \right), \quad (23)$$

with

$$\overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}} = \frac{1}{N} \sum_{i=1}^N \mathbf{P}_{\Pi_\delta(\widehat{\lambda}[\tau_{n_i}^i])},$$

where $\mathbf{P}_{\Pi_\delta(\widehat{\lambda}[\tau_{n_i}^i])}$ denotes the law of $\Pi_\delta(\widehat{\lambda}[\tau_{n_i}^i])$. Since the space $\mathcal{C}([-\delta, \delta], \mathbf{R})$ of continuous functions on $[-\delta, \delta]$, endowed with the uniform topology is separable, there exists a countable dense subset $(f_k)_{k \geq 1}$ of $\mathcal{C}([-\delta, \delta], \mathbf{R})$. Once again, using the method developed in Proposition 15, it is easy to get that, for any positive integer k and all ω in a set Ω_k of mass 1,

$$\lim_{N \rightarrow \infty} \left| \Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega) f_k - \overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}} f_k \right| = 0,$$

where μf denotes $\int f(x) \mu(dx)$, for any a measure μ and any μ -integrable function f . Now, take ω in $\bigcap_{k \geq 1} \Omega_k$ and f in $\mathcal{C}([-\delta, \delta], \mathbf{R})$. Since $(f_k)_{k \geq 1}$ is dense in $\mathcal{C}([-\delta, \delta], \mathbf{R})$, there exists for any $\epsilon > 0$ an integer k such that $\|f - f_k\| < \epsilon/2$. This implies that

$$\begin{aligned} \limsup_{N \rightarrow \infty} \left| \Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega) f - \overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}} f \right| &\leq \limsup_{N \rightarrow \infty} \left| \Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega) f - \Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega) f_k \right| \\ &+ \limsup_{N \rightarrow \infty} \left| \Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega) f_k - \overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}} f_k \right| \\ &+ \limsup_{N \rightarrow \infty} \left| \overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}} f_k - \overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}} f \right| < \epsilon. \end{aligned}$$

Since this last inequality holds for any $\epsilon > 0$, $\Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N} - \overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}}$ converges weakly to 0 in the space $\mathcal{M}_s([-\delta, \delta])$ of signed measures on $[-\delta, \delta]$ with probability 1. Consequently, since $d_W^{(1)}$ metricizes the subspace of probability measures, we get that, almost surely,

$$\lim_{N \rightarrow \infty} d_W^{(1)} \left(\Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega), \overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}} \right) = 0. \quad (24)$$

In order to control the second term and, hence, end the proof, it remains to show that

$$\forall \epsilon > 0, \exists A \in \mathbf{N}, \forall u \in \mathbf{S}_A, d_W^{(1)} \left(\overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}}, \Pi_\delta \mathcal{P} \right) < \epsilon.$$

To get this, we use the duality formula for the first order Wasserstein distance,

$$d_W^{(1)} \left(\overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}}, \Pi_\delta \mathcal{P} \right) = \sup_{\phi \in Lip_1([-\delta, \delta], \mathbf{R})} \left| \overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}} \phi - \Pi_\delta \mathcal{P} \phi \right|, \quad (25)$$

where $Lip_1([-\delta, \delta], \mathbf{R})$ denotes the set of 1-Lipschitz continuous functions on $[-\delta, \delta]$. Now, let ϕ be an element of $Lip_1([-\delta, \delta], \mathbf{R})$, we have that

$$\left| \overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}} \phi - \Pi_\delta \mathcal{P} \phi \right| \leq \frac{1}{N} \sum_{i=1}^N \left| \mathbf{E} \left[\phi \left(\Pi_\delta(\widehat{\lambda}[\tau_{n_i}^i]) \right) \right] - \mathbf{E}[\phi(\Pi_\delta(\Lambda_\infty))] \right|.$$

To prove that the supremum taken in the above inequality will be small as soon as the trees are large enough, we use the following lemma that gives the uniform convergence of the expectation of Lipschitz functionals of the $\widehat{\lambda}[\tau_{n_i}^i]$'s.

Lemma 20 *We have,*

$$\forall \epsilon > 0, \exists A \in \mathbf{N}, \forall n > A, \forall f \in Lip_1([-\delta, \delta], \mathbf{R}), \left| \mathbf{E} \left[f \left(\Pi_\delta(\widehat{\lambda}[\tau_{n_i}^i]) \right) \right] - \mathbf{E} [f(\Pi_\delta(\Lambda_\infty))] \right| < \epsilon.$$

The proof of the lemma has been postponed to the end of the section. In light of this lemma, we get

$$\forall \epsilon > 0, \exists A \in \mathbf{N}, \forall n > A, d_W^{(1)}(\overline{\Pi_\delta \mathcal{P}_{\vec{u}_N}}, \Pi_\delta \mathcal{P}) < \epsilon.$$

Consequently, using the above result in conjunction with (23), (24) and (25), we have

$$\forall \epsilon > 0, \exists A \in \mathbf{N}, \forall n > A, \limsup_{N \rightarrow \infty} d_W^{(1)}(\Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega), \Pi_\delta \mathcal{P}) < \epsilon,$$

for all ω in $\bigcup_{k \geq 1} \Omega_k$, where the sets Ω_k have been defined above. Hence,

$$\forall \epsilon > 0, \exists A \in \mathbf{N}, \forall u \in \mathbf{S}_A, \mathbf{P} \left(\limsup_{N \rightarrow \infty} d_W(\Pi_\delta \widehat{\mathcal{P}}_{\vec{u}_N}(\omega), \Pi_\delta \mathcal{P}) < \epsilon \right) = 1. \quad (26)$$

To end, using (21), (22) and (26) in (20) leads to the result. \square

Finally, we get the following consistency result.

Proposition 21 *We have,*

$$\forall \epsilon > 0, \exists A \in \mathbf{N}, \forall u \in \mathbf{S}_A, \mathbf{P} \left(\limsup_{N \rightarrow \infty} \left| \widehat{\lambda}_W[\mathcal{F}_{\vec{u}_N}] - \frac{1}{\sigma} \right| < \epsilon \right) = 1.$$

Proof. By the Cauchy-Schwarz inequality, the convergence of this estimator is conditioned to the convergence of the Wasserstein distance in the following manner,

$$\begin{aligned} \left| \widehat{\lambda}_W[\mathcal{F}_{\vec{u}_N}] - \frac{1}{\sigma} \right| &= \frac{|\langle \widehat{F}_{\vec{u}_N}^{-1} - \sigma^{-1} F_{\Lambda_\infty}^{-1}, F_{\Lambda_\infty}^{-1} \rangle|}{\|F_{\Lambda_\infty}^{-1}\|_2^2} \\ &\leq \frac{\|\widehat{F}_{\vec{u}_N}^{-1} - \sigma^{-1} F_{\Lambda_\infty}^{-1}\|_2 \|F_{\Lambda_\infty}^{-1}\|_2}{\|F_{\Lambda_\infty}^{-1}\|_2^2} \\ &= \frac{d_W(\widehat{\mathcal{P}}_{\vec{u}_N}, \mathcal{P})}{\|F_{\Lambda_\infty}^{-1}\|_2}. \end{aligned}$$

The result finally arises from Proposition 19 concerning the convergence of the empirical measure in the sense of the Wasserstein distance. \square

Proof of Lemma 20. Let $(n_i)_{i \geq 1}$ be a strictly increasing sequence of integers and, for any $i \geq 1$, let $\tau_{n_i}^i$ be a conditioned Galton-Watson tree with size n_i . By virtue of Skorokhod's representation theorem, there exists a probability space on which is defined a sequence $(X_{n_i})_{i \geq 1}$ of random variables and a random variable X_∞ such that:

- ◊ the sequence $(X_{n_i})_{i \geq 1}$ has the same distribution as the sequence $(\Pi_\delta(\widehat{\lambda}[\tau_{n_i}^i]))_{i \geq 1}$;
- ◊ X_∞ has the same distribution as $\Pi_\delta(\Lambda_\infty)$;
- ◊ X_{n_i} converges to X_∞ in probability as i goes to infinity.

As a consequence of the uniform integrability properties showed in the proof of Proposition 18, we get

$$\lim_{i \rightarrow \infty} \mathbf{E} [|X_\infty - X_{n_i}|] = 0. \quad (27)$$

Now, choose f in $Lip_1([-\delta, \delta], \mathbf{R})$, we have

$$|\mathbf{E} [f(X_{n_i})] - \mathbf{E} [f(X_\infty)]| \leq \mathbf{E} [|X_{n_i} - X_\infty|].$$

Together with (27), we obtain the expected convergence. \square

5 Simulation study

5.1 Simulation of conditioned Galton-Watson trees

In order to illustrate our estimation techniques on Galton-Watson forests, we need to make some numerical experiments. However, simulation of conditioned Galton-Watson trees is a difficult problem of independent importance. In this section, we briefly present an algorithm due to Devroye [9] allowing to achieve this aim. Given an integer n and a distribution μ on the set $\{0, \dots, K\}$, this algorithm provides, in two steps, the simulation of the Łukasiewicz walk $\mathcal{L}[\tau_n]$ of a tree τ_n with distribution $\text{GW}_n(\mu)$. Three more steps are required to obtain the corresponding Harris path $\mathcal{H}[\tau_n]$ through other coding processes (see for example [11]).

- ♦ **Simulation of numbers of children.** The multinomial distribution of parameters $(\mu(k))_{0 \leq k \leq K}$ and n may be defined by its probability mass function,

$$\mathbf{P}(N_0 = n_0, \dots, N_K = n_K) = \begin{cases} \frac{n!}{n_0! \dots n_K!} \mu(0)^{n_0} \dots \mu(K)^{n_K} & \text{if } \sum_{k=0}^K n_k = n, \\ 0 & \text{else.} \end{cases}$$

Simulation of the multinomial distribution presents no difficulty. By rejection sampling, we simulate multinomial random variables until obtaining a sequence $(N_k)_{0 \leq k \leq K}$ satisfying

$$\sum_{k=0}^K k N_k = n - 1.$$

We define the sequence $(\zeta_i)_{1 \leq i \leq n}$ from

$$(\zeta_i)_{1 \leq i \leq n} = (\underbrace{0, \dots, 0}_{N_0}, \underbrace{1, \dots, 1}_{N_1}, \dots, \underbrace{K, \dots, K}_{N_K}).$$

Let $(\xi_i)_{1 \leq i \leq n}$ be a sequence obtained as a random permutation of $(\zeta_i)_{1 \leq i \leq n}$. A suitable technique for random shuffling is presented in [19, Algorithm P (p. 139)]. The sequence $(\xi_i)_{1 \leq i \leq n}$ represents the vertices' numbers of children in the depth-first search order.

- ♦ **Computation of the Łukasiewicz walk.** Let L be the process defined by $L(0) = 0$ and,

$$\forall 0 \leq k \leq n - 2, \quad L(k + 1) = L(k) + \xi_{k+1} - 1.$$

Set $l = 1 + \arg \min \{L(k) : 0 \leq k \leq n - 1\}$. Then there exists a tree τ_n with n nodes whose Łukasiewicz walk is defined by

$$\mathcal{L}[\tau_n](k) = \begin{cases} L(l + k) + \min L - 1 & \text{if } 0 \leq k \leq n - 1 - l, \\ L(k - n + l) + \min L - 1 & \text{if } n - l \leq k \leq n - 1. \end{cases}$$

- ♦ **From the Łukasiewicz walk to the height process.** Now, we compute the corresponding height process [11, eq.(2)],

$$\forall 0 \leq k \leq n - 1, \quad \mathfrak{H}[\tau_n](k) = \# \left\{ 0 \leq j \leq k - 1 : \mathcal{L}[\tau_n](j) = \min_{j \leq l \leq n} \mathcal{L}[\tau_n](l) \right\}.$$

- ♦ **From the height process to the contour process.** Let $(b_k)_{0 \leq k \leq n-1}$ be the sequence defined from $b_k = 2k - \mathfrak{H}[\tau_n](k)$ if $0 \leq k \leq n - 1$ and $b_n = 2(n - 1)$. Then the b_i 's are sorted in increasing order. The contour process $\mathcal{C}[\tau_n](k)$ is defined for any $0 \leq k \leq 2n - 2$ by [11, eq.(1)]

$$\mathcal{C}[\tau_n](k) = \begin{cases} \mathfrak{H}[\tau_n](i) - (k - b_i) & \text{if } \exists 0 \leq i \leq n - 2, \quad b_i \leq k < b_{i+1} - 1, \\ k - b_{i+1} + \mathfrak{H}[\tau_n](i + 1) & \text{if } \exists 0 \leq i \leq n - 2, \quad b_{i+1} - 1 \leq k < b_{i+1}, \\ \mathfrak{H}[\tau_n](b_{n-1}) - (k - b_{n-1}) & \text{if } b_{n-1} \leq k \leq b_n. \end{cases}$$

◊ **From the contour process to the Harris path.** The Harris path is only a small modification of the contour process, defined by $\mathcal{H}[\tau_n](0) = \mathcal{H}[\tau_n](2n) = 0$ and

$$\forall 1 \leq k \leq 2n-1, \quad \mathcal{H}[\tau_n](k) = \mathcal{C}[\tau_n](k-1) + 1.$$

5.2 Inference for a forest of binary conditioned Galton-Watson trees

The aim of this section is to analyze the finite-sample behavior of both estimators introduced in this paper by means of numerical experiments. The theoretical study achieved in Section 4 shows that we can expect to obtain good numerical results, at least for large trees. To this goal, we consider a forest of independent conditioned Galton-Watson trees with common critical birth distribution μ such that $\mu(k) = 0$ for $k \geq 3$. Such a distribution satisfies the following linear system of equations,

$$\begin{cases} \mu(0) + \mu(1) + \mu(2) &= 1 \\ \mu(1) + 2\mu(2) &= 1 \\ \mu(1) + 4\mu(2) - 1 &= \sigma^2 \end{cases}$$

which is equivalent to

$$\mu(0) = \mu(2) = \frac{\sigma^2}{2} \quad \text{and} \quad \mu(1) = 1 - \sigma^2.$$

In other words, μ is entirely characterized by its variance σ^2 . Simulations of Galton-Watson trees $\text{GW}_n(\mu)$ are performed with the method provided in Subsection 5.1.

Let $\mathcal{F} = (\tau^i)_{1 \leq i \leq N}$ be a forest of N independent trees such that, for any $1 \leq i \leq N$, $\tau^i \sim \text{GW}_{n_i}(\mu)$ for some integer n_i . From the Harris process of each tree τ^i , one first computes the quantity

$$\hat{\lambda}[\tau^i] = \frac{\langle \mathcal{H}[\tau^i](2n_i \cdot), E \rangle}{2\sqrt{n_i} \|E\|_2^2},$$

where E is known and defined in (2). Then, we propose to estimate σ^{-1} in the two following ways, where $(\hat{\lambda}[\tau^{(i)}])_{1 \leq i \leq N}$ denotes the order statistic associated to the family $(\hat{\lambda}[\tau^i])_{1 \leq i \leq N}$.

Least Squares	Wasserstein
$\hat{\lambda}_{ls}[\mathcal{F}] = \frac{1}{N} \sum_{i=1}^N \hat{\lambda}[\tau^i]$	$\hat{\lambda}_W[\mathcal{F}] = \frac{1}{\ F_{\Lambda_\infty}^{-1}\ _2^2} \sum_{i=1}^N \hat{\lambda}[\tau^{(i)}] \int_{\frac{i-1}{N}}^{\frac{i}{N}} F_{\Lambda_\infty}^{-1}(s) ds$

Remark 22 *In order to compute $\hat{\lambda}_W[\mathcal{F}]$, we need to be able to perform computations using the function $F_{\Lambda_\infty}^{-1}$. Unfortunately, in view of the theoretical study of Λ_∞ made in Proposition 3.1, one can not expect to have an explicit expression for this function. In the following of this section, we use a numerical estimation of $F_{\Lambda_\infty}^{-1}$ by Monte Carlo simulations. To achieve this goal, we perform simulations of Λ_∞ thanks to formula (6) by simulating Brownian excursions thanks to (3). In order to ensure that the error made on $F_{\Lambda_\infty}^{-1}$ does not propagate too much in our results, $F_{\Lambda_\infty}^{-1}$ is estimated from one million simulations of Λ_∞ .*

The theoretical investigations of Section 4 establish that our estimators are asymptotically unbiased. Nevertheless, the problem is not as simple when working with finite trees. A clear illustration of this comes from the numerical evaluations of the average Harris processes of finite trees. Indeed, the numerical study of Figure 6 shows that the average Harris processes of small trees seem to be lower than the limiting Harris process. Hence, the quantities $\hat{\lambda}[\tau^i]$ are expected to underestimate the target σ^{-1} . But any estimator based on the asymptotic behavior of conditioned Galton-Watson trees is expected to present such a bias. In particular, we state in our numerical experiments that the estimator proposed in [4] presents the same bias.

[Figure 6 approximately here]

The natural question arising from the preceding comments is: how is the bias of a conditioned Galton-Watson tree related to its size and/or the unknown parameter σ ? The numerical study presented in Figure 7 shows that the quantity $\eta(n) = \sigma^{-1} \mathbf{E}[\hat{\lambda}[\tau_n]]^{-1}$, where $\tau_n \sim \text{GW}_n(\mu)$, seems close to uncorrelated to σ at least when σ is large enough. This allows us to construct a bias corrector which is independent of the unknown standard deviation σ . In addition, the dependency on n may be modeled by the relation $\eta(n) = 1 - (a\sqrt{n} + b)^{-1}$. The coefficients appearing in η may be estimated from simulated data,

$$\hat{\eta}(n) = 1 - (0.504273\sqrt{n} + 0.9754839)^{-1}$$

(see Figure 7 again). The correction is obviously expected to be better for large values of σ . Finally, we construct the following corrected versions of the estimators $\hat{\lambda}_{ls}[\mathcal{F}]$ and $\hat{\lambda}_W[\mathcal{F}]$.

Corrected Least Squares	Corrected Wasserstein
$\hat{\lambda}_{ls}^c[\mathcal{F}] = \frac{1}{N} \sum_{i=1}^N \hat{\eta}(\#\tau^i) \hat{\lambda}[\tau^i]$	$\hat{\lambda}_W^c[\mathcal{F}] = \frac{1}{\ F_{\Lambda_\infty}^{-1}\ _2^2} \sum_{i=1}^N \hat{\eta}(\#\tau^{(i)}) \hat{\lambda}[\tau^{(i)}] \int_{\frac{i-1}{N}}^{\frac{i}{N}} F_{\Lambda_\infty}^{-1}(s) ds$

[Figure 7 approximately here]

In light of the previous comments, computing the estimators proposed in this paper is not an easy task. According to Remark 22, one needs to perform a significant number of simulations of Λ_∞ in order to get an accurate approximation of $F_{\Lambda_\infty}^{-1}$. Moreover, to be able to correct the aforementioned bias, one needs to perform many simulations of finite trees. Together with this work, we propose a `Matlab` toolbox which already includes these preliminary computations and allows to directly and quickly compute our estimators for forests. This toolbox as well as its documentation and the scripts used in this paper are available at the webpage: <http://agh.gforge.inria.fr>.

For improved comparison, we also compute the estimator $\hat{\lambda}_{un}[\mathcal{F}]$ of σ^{-1} based on the work [4] (see Subsection 3.1) given by

$$\hat{\lambda}_{un}[\mathcal{F}] = \frac{1}{N} \sum_{i=1}^N \hat{\delta}[\tau^i],$$

where $\hat{\delta}[\tau^i]$ is defined (see equations (8) and (9)) from a node v randomly chosen in τ^i by

$$\hat{\delta}[\tau^i] = \frac{\sqrt{2} h(v)}{\sqrt{\pi \#\tau^i}}.$$

The estimator $\hat{\lambda}_{un}[\mathcal{F}]$ is expected to present the bias due to the approximation of Harris paths by their expected limit. We correct it by the aforementioned method,

$$\hat{\lambda}_{un}^c[\mathcal{F}] = \frac{1}{N} \sum_{i=1}^N \hat{\eta}(\#\tau^i) \hat{\delta}[\tau^i].$$

In Figures 8, 9 and 10, estimators $\hat{\lambda}_{ls}[\mathcal{F}]$, $\hat{\lambda}_W[\mathcal{F}]$ and $\hat{\lambda}_{un}[\mathcal{F}]$ are denoted by “LSE”, “Wasserstein” and “Uniform node” (or “UN” in short), respectively.

The study of Figure 8 shows that for values of σ greater than 0.5, the bias correction works properly. Moreover, it also shows that the approach developed in [4] presents the same kind of bias as ours. In the case of small parameter σ , the bias correction is not as accurate. This was expected because the bias corrector does not fit as well to the bias curve for small values of sigma as it does for greater values of σ .

[Figure 8 approximately here]

Since we have an estimation procedure which seems to work, the natural further study is to see how the quality of our estimators varies as the characteristics of the forest change. We begin by looking at the variations when the sizes of the trees increase. A priori, the sizes of the trees in the considered forest should not have influence on the variability of the estimators. Indeed, our estimation strategy is based on the approximation of the Harris path of a finite tree by its limit. As a consequence, the size parameter only governs the quality of this approximation. Whatever the sizes of the trees, the variability will be given by the variance of the limit distribution Λ_∞ . As expected Figure 9 shows that the variability of the estimators does not change as the sizes of the trees change when σ takes great values. Similarly, as shown in Figure 10, for small values of σ , the sizes of the trees do not influence the dispersion of the estimator. However, Figure 10 also shows that the sizes of the trees have a positive influence on the bias of the estimators.

[Figures 9 and 10 approximately here]

Finally, Figure 11 shows the variation of the dispersion of the least square estimator as the size of the forest changes. It appears to be consistent with the theoretical tolerance intervals given by the central limit theorem. Similar results have been obtained from the Wasserstein method (see Figure 12).

[Figures 11 and 12 approximately here]

5.3 Missing or noisy data

We focus here on the application of our statistical methods in the framework of missing or noisy data. This section is key for this paper because we exhibit difficult contexts in which Harris paths-based estimators perform well while the empirical variance is biased or can not be computed.

5.3.1 Estimation with outliers

We assume that the forest $\mathcal{F} = (\tau^i)_{1 \leq i \leq N}$ is mainly composed of binary conditioned Galton-Watson trees with the same variance σ^2 . However the forest also contains trees τ^i that have not been generated from the model and such that $\hat{\lambda}[\tau^i]$ is significantly lesser or greater than the true parameter σ^{-1} . As a consequence, the estimators $\hat{\lambda}_{ls}[\mathcal{F}]$ and $\hat{\lambda}_W[\mathcal{F}]$ should be disturbed by these outliers. In the following simulation experiments, the forest contains 500 conditioned Galton-Watson trees and 50 outliers (addition of 10% outliers).

First, we consider the presence of outliers $\hat{\lambda}[\tau^i] \simeq 0.03$ smaller than the expected parameter $\sigma^{-1} = 2$ and we compare our estimation strategies (see Figure 13). One may observe that the Wasserstein method is less sensitive to these outliers than the least square estimator. This may be explained by the small weights $\int_{\frac{i-1}{N}}^{\frac{i}{N}} F_{\Lambda_\infty}^{-1}(s)ds$ given to the smallest values of $\hat{\lambda}[\tau^{(i)}]$, that is to say, to the outliers, in formula (12) of $\hat{\lambda}_W[\mathcal{F}]$. For the same reason, the Wasserstein estimator is more sensitive to large outliers $\hat{\lambda}[\tau^i] \simeq 2.9$ (see Figure 14) than the least square strategy.

[Figures 13 and 14 approximately here]

The numerical results of Subsection 5.2 show that the two strategies developed in this paper perform in a similar way on a dataset without outliers, which is clearly a benefit for the least square estimator $\hat{\lambda}_{ls}[\mathcal{F}]$ easier to compute than $\hat{\lambda}_W[\mathcal{F}]$ (see Remark 22). Nevertheless, if one suspects the presence of small (large, respectively) outliers, these numerical experiments yield that the Wasserstein (least square, respectively) estimator should be privileged. In addition, it seems that the two estimators behave differently only under the presence of outliers. This observation could be used to detect suspicious data.

5.3.2 Missing leaves

We assume that we observe a conditioned critical Galton-Watson tree τ through a noise hiding its leaves. It means that, in the Harris path $\mathcal{H}[\tau]$, all the sections $[i-1, i+1]$ with

$$\mathcal{H}[\tau](i-1) = \mathcal{H}[\tau](i+1) = \mathcal{H}[\tau](i) - 1,$$

i.e., sections corresponding to leaves, are also hidden. We refer the reader to Figure 15 for an example of conditioned Galton-Watson tree observed through this deleting noise. From a tree τ with n nodes, we propose to estimate $\hat{\lambda}[\tau]$ from the partially observed Harris path as follows. Let us denote $\zeta \subset [0, 2n]$ the union of the unobserved sections of the Harris path and $\eta = \frac{\zeta}{2n} \subset [0, 1]$. Even if $\hat{\lambda}[\tau]$ is uncomputable, we can approximate it by the solution $\tilde{\lambda}[\tau]$ of the least square problem on η^c ,

$$\lambda \mapsto \left\| \left(\frac{\mathcal{H}[\tau](2n \cdot)}{\sqrt{n}} - 2\lambda E \right) \mathbf{1}_{\eta^c} \right\|_2^2.$$

Mimicking the expression (5) of $\hat{\lambda}[\tau]$, $\tilde{\lambda}[\tau]$ is given by

$$\tilde{\lambda}[\tau] = \frac{\int_{[0,1] \cap \eta^c} \mathcal{H}[\tau](2nt) E_t dt}{2\sqrt{n} \int_{[0,1] \cap \eta^c} E_t^2 dt}.$$

From a forest of conditioned Galton-Watson trees τ , we compute the values of $\hat{\lambda}[\tau]$ and $\tilde{\lambda}[\tau]$ in order to evaluate the influence of the noise on the quality of the estimation. We also compute the empirical variance $V[\tau]$ of the numbers of children appearing in τ from the complete tree and from its noisy version, σ^{-1} being estimated by $V[\tau]^{-1/2}$. Some numerical results are presented in Figure 16.

[Figures 15 and 16 approximately here]

First, we remark that the distributions of $\hat{\lambda}[\tau]$ and $\tilde{\lambda}[\tau]$ are quite close, whereas the behavior of the empirical variance is highly disturbed by the absence of zeros in the set of numbers of children. As a consequence, statistical estimators computed from the Harris path seem to be more robust than empirical estimators of the birth distribution, even when the Harris path is largely hidden (see the example of Figure 15).

5.3.3 Partial observation of the Harris path

Here we assume that the Harris path is partially observed in such a way that the underlying tree can not be reconstructed. This kind of disturbance may appear in data transmission where unwanted electromagnetic energy can degrade the quality of the signal. We consider two types of partial observation: (i) large sections of the Harris path are hidden (see Figure 17) and (ii) the Harris path is observed through an additive Gaussian noise (see Figure 18).

[Figures 17 and 18 approximately here]

Since the tree can not be deduced from these noisy observations, empirical estimators of the birth distribution can not be computed, while statistical methods based on the Harris path are still feasible. The numerical results of Figures 17 and 18 show that the distribution of $\hat{\lambda}[\tau]$ is only slightly disturbed by the noise proving again the robustness of statistical estimators computed on Harris paths.

6 Real data analysis: history of Wikipedia webpages

The aim of this section is to show that the methodology developed in this paper can be used to analyze the history of some real hierarchical data. More precisely, we focus on the evolution over time of a given

webpage on the World Wide Web. HTML is the standard markup language for creating webpages. Documents encoded in a markup language naturally presents a tree structure: the area delimited by opening and closing tags represents a node of the tree; the children of this node are given by the tags directly found in this area in the order they appear (see Figure 19 for an example of HTML document and the corresponding ordered tree structure). It should be noted that the ordered tree representing an HTML document does not take into account the text between tags but only the hierarchical structure.

[Figure 19 approximately here]

Nowadays, Wikipedia is probably the most famous free Internet encyclopedia. It allows its users to create and edit almost any article. All past changes are listed in reverse-chronological order and are accessible from the current version of the Wikipedia webpage. Consequently, each article forms a time series composed of hundreds of revisions. The analysis of this chronological dataset is difficult because of the complex structure of the data which has no representation in an Euclidean state space. We propose to apply the strategy presented in this paper to investigate this question and obtain informations on the history of articles. Wikipedia webpages do not look like conditioned Galton-Watson trees (see Figure 20 for a typical Wikipedia webpage and its Harris path which should be compared with the conditioned Galton-Watson tree with a comparable number of nodes of Figure 5) but they share the same structure with a typical layout that consists in standardized HTML/CSS files on which articles are based. Thus webpages at hand might not be differentiated by considering their shape but some scale parameter, as it is the case for conditioned Galton-Watson trees. We claim that the quantity $\hat{\lambda}[\tau]$, where τ is the underlying tree of a given webpage, is a good estimate of its relative scale and may be used to represent the revision history.

[Figure 20 approximately here]

We begin with the English version of the Wikipedia article *Gravitational wave*². This article has been edited 2001 times by 810 Wikipedians since its creation on September 3rd 2001 (information acquired on August 11 2016). For each month since January 2005, we compute $\hat{\lambda}_{ls}$ from the forest of the versions revised during this month. If no revision has been found during this period, $\hat{\lambda}_{ls}$ is equal to the estimate of the previous month, and recursively. Figure 21 displays the evolution of $\hat{\lambda}_{ls}$ over time. First, we remark two spikes (a), negative in May 2007, and (b), positive in May 2016. Both these spikes correspond to massive vandalism of the article on May 7 2007 (addition of 720 pointless sections with random text) and May 23 2016 (complete deletion of the article) by malicious people. Indeed, if we do not consider these two vandalized webpages in our estimation, we obtain the graph of Figure 22 (left) that has no spikes. In addition, we observe in Figure 22 (left) that the time series of $\hat{\lambda}_{ls}$ has roughly two regimes (c) and (d). The first period (c) corresponds to the “running in” required to find the adequate structure of the article. In this period, the webpage is subject to major changes that are most often additions of new sections or paragraphs but may be deletions of inappropriate content. When a good structure arises, the webpage is then slowly broadened during the second regime (d). It should be remarked in Figure 22 (right) that two important modifications occur in the period (d): (e) between July 2013 and April 2014 and (f) in February 2016. The (e) period is related to major changes in the webpage (mainly addition of references and reorganization of some sections) especially following advances in this field. The second event (f) corresponds to extensive adding following the announce of the first observation of gravitational waves using the Advance LIGO detectors.

[Figures 21 and 22 approximately here]

We perform the same methodology on the history of the Wikipedia article *Chocolate*³ (see Figure 23). This article has been edited 6332 times by 3105 Wikipedians since its creation on November 13 2001 (information acquired on August 11 2016). All the spikes observed on the graph of Figure 23 correspond to acts of vandalism (deletion of substantial content). For the sake of example, we highlight two major events (a)

²Wikipedia article *Gravitational wave*: https://en.wikipedia.org/wiki/Gravitational_wave

³Wikipedia article *Chocolate*: <https://en.wikipedia.org/wiki/Chocolate>

(in May 2008) and (b) (in June 2010) occurring during the “running in” period (c): (a) corresponds to important additions in the article (sections *Etymology*, *Holidays* and *Manufacturers* have been added), while (b) is related to the creation of the parallel article *Health effects of chocolate* leading to deletion of the corresponding sections in the main article.

[Figure 23 approximately here]

For both examples, we empirically observe that, when $\hat{\lambda}_{ls}$ decreases, some content has been added to the webpage, and conversely, when $\hat{\lambda}_{ls}$ increases, some parts of the article have been removed. Our analysis shows that, starting from their creation, these Wikipedia articles are broadened over time after a long “running in” period used to unconsciously find the adequate structure. One may also detect vandalism on Wikipedia articles by identifying spikes a posteriori. Vandalism is usually removed by dedicated individuals who patrol Wikipedia webpages, but this is an onerous task with a rate of 10 edits per second⁴ and around 7% of edits have been estimated to be vandalism [29]. Vandalism detection is often based on a combination of various indicating features [1, 25]. Our algorithm might be used as a new feature for identifying acts of vandalism on the structure of the article.

Acknowledgment

The authors would like to show their gratitude to an anonymous reviewer who provided many relevant comments on the manuscript.

References

- [1] ADLER, B. T., DE ALFARO, L., MOLA-VELASCO, S. M., ROSSO, P., AND WEST, A. G. *Wikipedia Vandalism Detection: Combining Natural Language, Metadata, and Reputation Features*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 277–288.
- [2] ALDOUS, D. The Continuum Random Tree III. *Ann. Probab.* 21, 1 (01 1993), 248–289.
- [3] ALDOUS, D. Probability distributions on cladograms. In *Random discrete structures*. Springer, 1996, pp. 1–18.
- [4] BHARATH, K., KAMBADUR, P., DEY, D., ARVIN, R., AND BALADANDAYUTHAPANI, V. Statistical tests for large tree-structured data. *Journal of the American Statistical Association* (2016).
- [5] BOBKOV, S., AND LEDOUX, M. One-dimensional empirical measures, order statistics and Kantorovich transport distances. *Preprint* (2014).
- [6] CZADO, C., AND MUNK, A. Nonparametric validation of similar distributions and assessment of goodness of fit. *Journal of the Royal Statistical Society: Series B* 60, 1 (1998), 223–241.
- [7] DAVID, H. A., AND NAGARAJA, H. N. *Order statistics*, third ed. Wiley Series in Probability and Statistics. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, 2003.
- [8] DERRIDA, B., ENAUD, C., AND LEBOWITZ, J. L. The asymmetric exclusion process and brownian excursions. *Journal of Statistical Physics* 115, 1 (Apr 2004), 365–382.
- [9] DEVROYE, L. Simulating size-constrained Galton-Watson trees. *SIAM Journal on Computing* 41, 1 (2012), 1–11.
- [10] DRMOTA, M., AND MARCKERT, J.-F. Reinforced weak convergence of stochastic processes. *Statist. Probab. Lett.* 71, 3 (2005), 283–294.

⁴Wikipedia statistics (last consulted on August 11 2016): <https://en.wikipedia.org/wiki/Wikipedia:Statistics>

- [11] DUQUESNE, T. A limit theorem for the contour process of conditioned Galton-Watson trees. *Ann. Probab.* 31, 2 (04 2003), 996–1027.
- [12] GALLÓN, S., LOUBES, J.-M., AND MAZA, E. Statistical properties of the quantile normalization method for density curve alignment. *Mathematical Biosciences vol. 242*, 2 (Apr. 2013), pp. 129–142.
- [13] HAULK, C. K. *Aspects of Exchangeable Partitions and Trees*. PhD thesis, UC Berkeley, 2011.
- [14] JANSON, S. Asymptotic distribution for the cost of linear probing hashing. *Random Structures & Algorithms* 19, 3-4 (2001), 438–471.
- [15] JANSON, S. Brownian excursion area, Wright’s constants in graph enumeration, and other Brownian areas. *Probability Surveys* 4 (2007), 80–145.
- [16] JANSON, S. Simply generated trees, conditioned Galton-Watson trees, random allocations and condensation. *Probab. Surveys* 9 (2012), 103–252.
- [17] JONES, M., AND RICE, J. A. Displaying the important features of large collections of similar curves. *The American Statistician* 46, 2 (1992), 140–145.
- [18] KANTOROVITCH, L. V. Sur certains développements suivant les polynômes de la forme de S. Bernstein. *C. R. Acad. Sci. URSS* (1930).
- [19] KNUTH, D. Seminumerical algorithms, 2nd edn, vol. 2 of the art of computer programming, 1981.
- [20] KNUTH, D. Sorting and searching, 2nd edn, vol. 3 of the art of computer programming, 1981.
- [21] KNUTH, D. E. Dynamic Huffman coding. *Journal of algorithms* 6, 2 (1985), 163–180.
- [22] LORENTZ, G. G. G. *Bernstein polynomials*. Mathematical Expositions, no. 8. University of Toronto Press, Toronto, 1953.
- [23] LOUCHARD, G. Kac’s formula, Levy’s local time and Brownian excursion. *J. Appl. Probab.* 21, 3 (1984), 479–499.
- [24] LOUCHARD, G., AND JANSON, S. Tail estimates for the Brownian excursion area and other Brownian areas. *Electron. J. Probab.* 12 (2007), no. 58, 1600–1632.
- [25] MOLA-VELASCO, S. Wikipedia vandalism detection through machine learning: Feature review and new proposals. In *CLEF 2010 LABs and Workshops, Notebook papers* (September 2010).
- [26] NUALART, D. *The Malliavin calculus and related topics*, second ed. Probability and its Applications (New York). Springer-Verlag, Berlin, 2006.
- [27] PITMAN, J. Brownian motion, bridge, excursion, and meander characterized by sampling at independent uniform times. *Electron. J. Probab.* 4 (1999), 33 pp.
- [28] PITMAN, J. *Combinatorial stochastic processes*, vol. 32. Springer Science & Business Media, 2006.
- [29] POTTHAST, M. Crowdsourcing a Wikipedia vandalism corpus. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (New York, NY, USA, 2010), SIGIR ’10, ACM, pp. 789–790.
- [30] REVUZ, D., AND YOR, M. *Continuous Martingales and Brownian Motion*. Grundlehren der mathematischen Wissenschaften A series of comprehensive studies in mathematics. Springer, 1999.
- [31] SHEN, D., SHEN, H., BHAMIDI, S., MALDONADO, Y. M., KIM, Y., AND MARRON, J. S. Functional Data Analysis of Tree Data objects. *Journal of Computational and Graphical Statistics* 23, 2 (2014), 418–438.

- [32] VIÉGAS, F. B., WATTENBERG, M., AND DAVE, K. Studying cooperation and conflict between authors with history flow visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2004), CHI '04, ACM, pp. 575–582.
- [33] VIEGAS, F. B., WATTENBERG, M., KRISS, J., AND VAN HAM, F. Talk before you type: Coordination in Wikipedia. In *Proceedings of the 40th Annual Hawaii International Conference on System Sciences* (Washington, DC, USA, 2007), HICSS '07, IEEE Computer Society, pp. 78–.

Figures obtained from simulated data




Figure 5: A Galton-Watson tree conditional on having 1000 nodes generated from the geometric birth distribution with variance $\sigma^2 = 2$ (top) and its Harris path (bottom).




Figure 6: Estimated mean Harris processes of binary conditioned Galton-Watson trees with size n and $\sigma = 0.7$ calculated from 2000 trees for each value of n .




Figure 7: Estimation of the quantity $\eta(n) = \sigma^{-1}\mathbf{E}[\hat{\lambda}[\tau_n]]^{-1}$, where $\tau_n \sim \text{GW}_n(\mu)$, for different values of σ and different numbers of nodes n , together with the fitted bias corrector function $\hat{\eta}$. Estimations have been made by Monte Carlo method with samples of 2000 trees for each couple (n, σ) .

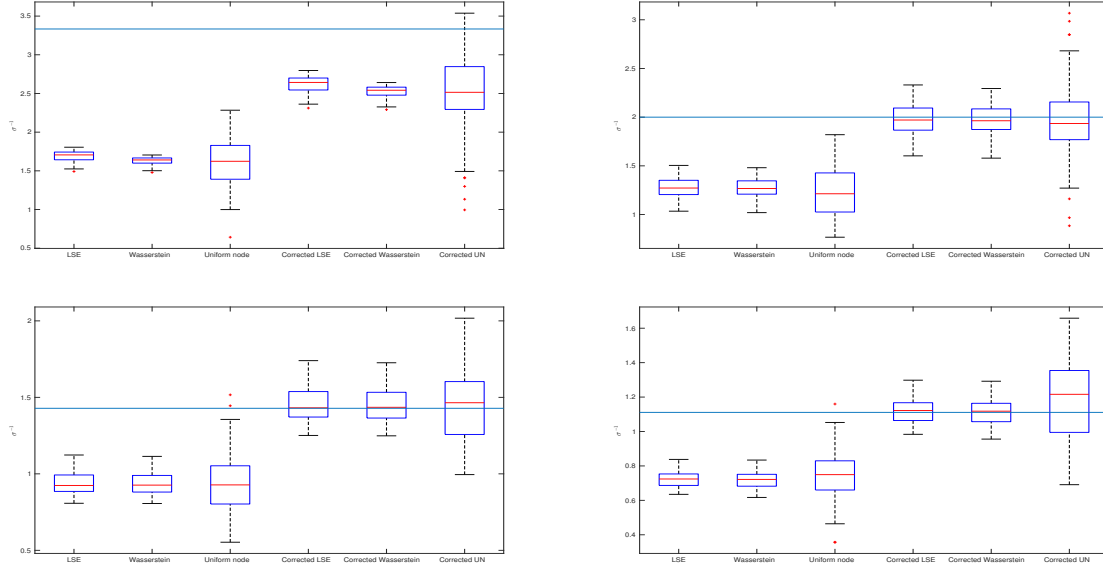


Figure 8: Estimation and bias correction for forests of 10 trees with 20 nodes for σ equals to 0.3 (top, left) 0.5 (top, right), 0.7 (bottom, left) and 0.9 (bottom right). Boxplots have been drawn from 100 replicates each.




Figure 9: Influence of the size of the trees for σ equals to 0.9: tree sizes varying from 20 nodes (left), 50 nodes (center), to 100 nodes (right). Forests of 50 trees. Boxplots have been drawn from 100 replicates each.




Figure 10: Influence of the size of the trees for σ equals to 0.3: tree sizes varying from 20 nodes (left), 50 nodes (center), to 100 nodes (right). Forests of 50 trees. Boxplots have been drawn from 100 replicates each.




Figure 11: Least square estimation of σ^{-1} for different sizes of forests ($\sigma = 0.5$ with trees of size 20). Boxplots have been drawn from 100 replicates each.




Figure 12: Wasserstein estimation of σ^{-1} for different sizes of forests ($\sigma = 0.5$ with trees of size 20). Boxplots have been drawn from 100 replicates each.




Figure 13: Boxplots of $\hat{\lambda}_{ls}[\mathcal{F}]$ and $\hat{\lambda}_W[\mathcal{F}]$ from forests containing (right) or not (left) small outliers impacting the quality of the estimation.




Figure 14: Boxplots of $\hat{\lambda}_{ls}[\mathcal{F}]$ and $\hat{\lambda}_W[\mathcal{F}]$ from forests containing (right) or not (left) large outliers impacting the quality of the estimation.

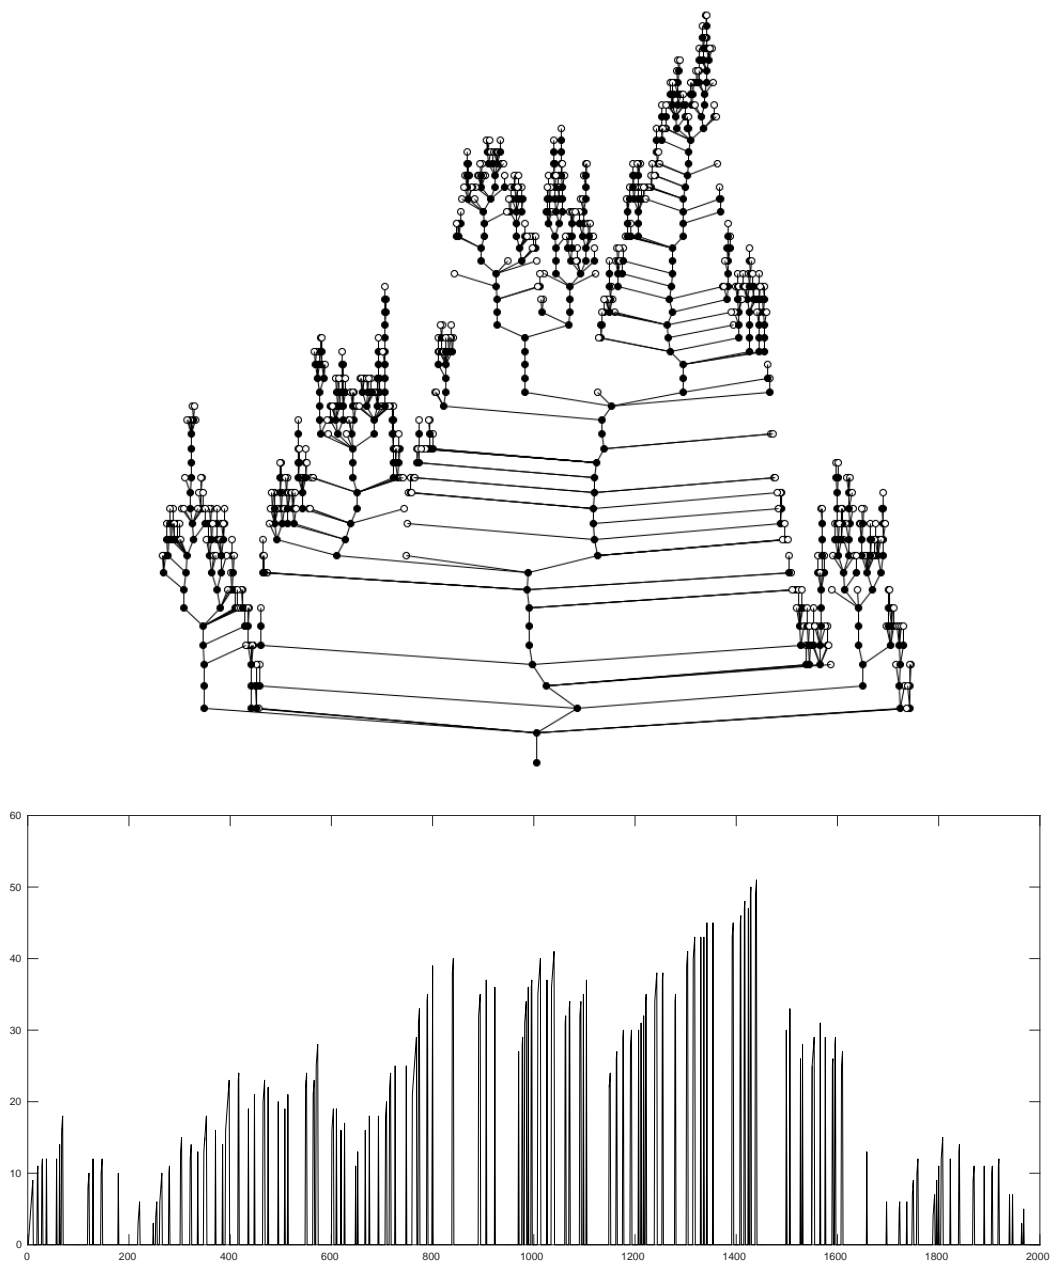


Figure 15: The conditioned Galton-Watson tree of Figure 5 (top) and its Harris path (bottom) are partially observed: leaves (in white) are missing.




Figure 16: Least square and empirical estimators of σ^{-1} from complete trees and trees with missing leaves.




Figure 17: The Harris path of the conditioned Galton-Watson tree of Figure 5 observed only on the intervals $[0, 500]$ and $[1000, 1500]$ (left) and boxplots of $\hat{\lambda}[\tau]$ from complete and partially observed Harris paths (right).




Figure 18: The Harris path of the conditioned Galton-Watson tree of Figure 5 observed through a Gaussian noise with standard deviation 5 (left) and boxplots of $\hat{\lambda}[\tau]$ from complete and partially observed Harris paths (right).

Figures obtained from real data

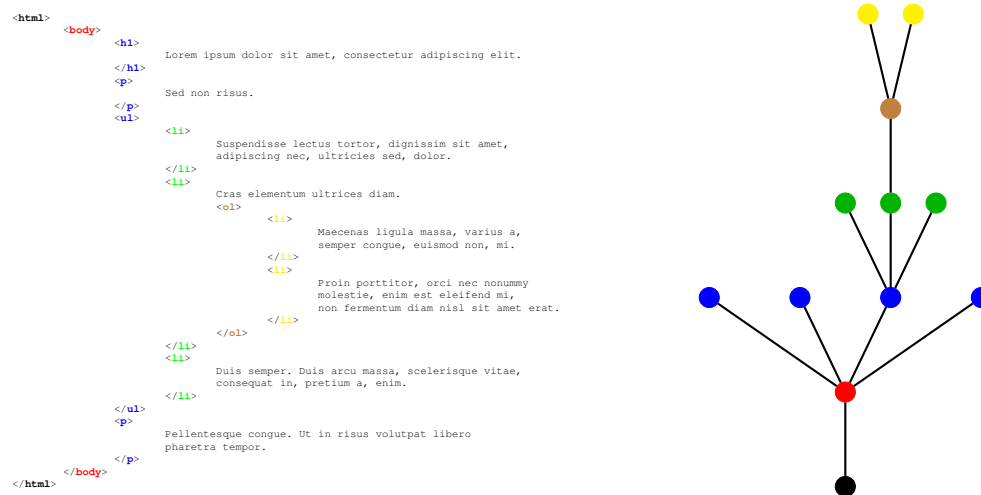


Figure 19: Underlying ordered tree structure (right) present in an HTML document (left). Each level in the tree is colored in the same way as the corresponding tags in the document. Natural order from top to bottom in the HTML document corresponds to left-to-right order in the tree.




Figure 20: Underlying tree of the main page of Wikipedia accessed on April 12 2017 with 906 nodes (top) and its Harris path (bottom).




Figure 21: History of $\hat{\lambda}_{ls}$ between January 2005 and June 2016 for the Wikipedia article *Gravitational wave*. Events (a) and (b) are related to vandalism.




Figure 22: History of $\hat{\lambda}_{ls}$ for the Wikipedia article *Gravitational wave* without taking into account the two vandalism pages related to (a) and (b) between 2005 and 2016 (left) and 2010 and 2016 (right).




Figure 23: History of $\hat{\lambda}_{l,s}$ between January 2005 and June 2016 for the Wikipedia article *Chocolate*. All spikes are related to vandalism.