

UNCERTAINTY QUANTIFICATION FOR BAYESIAN CART

BY ISMAËL CASTILLO^{*,‡} AND VERONIKA ROČKOVÁ^{†,§}

Sorbonne Université[‡]
University of Chicago[§]

Sorbonne Université & Institut Universitaire de France
Laboratoire de Probabilités, Statistique et Modélisation
4, Place Jussieu, 75005 Paris, France
ismael.castillo@upmc.fr

University of Chicago
Booth School of Business
5807 S. Woodlawn Avenue
Chicago, IL, 60637, USA
veronika.rockova@chicagobooth.edu

This work affords new insights into Bayesian CART in the context of structured wavelet shrinkage. The main thrust is to develop a formal inferential framework for Bayesian tree-based regression. We reframe Bayesian CART as a *g-type* prior which departs from the typical wavelet product priors by harnessing correlation induced by the tree topology. The practically used Bayesian CART priors are shown to attain adaptive near rate-minimax posterior concentration in the *supremum norm* in regression models. For the fundamental goal of uncertainty quantification, we construct *adaptive* confidence bands for the regression function with uniform coverage under self-similarity. In addition, we show that tree-posteriors enable optimal inference in the form of efficient confidence sets for smooth functionals of the regression function.

1. Introduction. The widespread popularity of Bayesian tree-based regression has raised considerable interest in theoretical understanding of their empirical success. However, theoretical literature on methods such as Bayesian CART and BART is still in its infancy. In particular, statistical *inferential* theory for regression trees and forests (both frequentist and Bayesian) has been severely under-developed.

This work sheds light on Bayesian CART [20, 23] which is a popular learning tool based on ideas of recursive partitioning and which forms an integral constituent of BART [19]. Bayesian Additive Regression Trees (also known as BART) have emerged as one of today’s most effective general approaches to predictive modeling under minimal assumptions. Their empirical success has been amply illustrated in the context of non-parametric regression [19], classification [43], variable selection [8, 42, 40], shape constrained inference [18], causal inference [36, 35], to name a few. The BART model deploys an additive aggregate of individual trees using Bayesian CART as its building block. While theory for random forests, the frequentist counterpart, has seen numerous recent developments [57, 6, 50, 39, 56], theory for Bayesian CART and BART has not kept pace with its application. With the first theoretical

^{*}The author gratefully acknowledges support from the Institut Universitaire de France and from the ANR grant ANR-17-CE40-0001 (BASICS).

[†] The author gratefully acknowledges support from the James S. Kemper Foundation Faculty Research Fund at the University of Chicago Booth School of Business and the National Science Foundation (grant DMS-1944740).
MSC2020 subject classifications: Primary 62G20, 62G15.

Keywords and phrases: Bayesian CART, Posterior Concentration, Recursive Partitioning, Regression Trees, Nonparametric Bernstein–von Mises theorem.

results (Hellinger convergence rates) emerging very recently [49, 41, 48], many fundamental questions pertaining to, e.g., convergence in stronger losses such as the supremum norm, as well as *uncertainty quantification* (UQ), have remained to be addressed. This work takes a leap forward in this important direction by developing a formal frequentist statistical framework for uncertainty quantification with confidence bands for Bayesian CART.

We first show that Bayesian CART reaches a (near-)optimal posterior convergence rate under the *supremum-norm* loss, a natural loss for UQ of regression functions. Many methods that are adaptive for the L^2 -loss actually fail to be adaptive in an L^∞ -sense, as we illustrate below. We are actually not aware of any supremum-norm convergence rate result for related machine learning methods in the literature, including CART, random forests and deep learning. Regarding inference, we provide a construction of an *adaptive* credible band for the unknown regression function with (nearly, up to a logarithmic term) optimal uniform coverage under self-similarity. In addition, we provide efficient confidence sets and bands for a family of smooth functionals. Uncertainty quantification for related random forests or deep learning has been an open problem, with distributional results available only for point-wise prediction using bootstrap techniques [39]. Our results make a needed contribution to the literature on the widely sought-after UQ for (tree-based) machine learning methods.

Regarding supremum-norm (and its associated discrete ℓ_∞ version) posterior contraction rates, their derivation is typically more delicate compared to the more familiar testing distances (e.g. L^2 or Hellinger) for which general theory has been available since the seminal work [29]. Despite the lack of unifying theory, however, advances have been made in the last few years [31, 12, 38] including specific models [52, 63, 45, 46]. However, Bayesian *adaptation* for the supremum loss has been obtained, to the best of our knowledge, *only* through spike-and-slab priors (the work [62] uses Gaussian process priors, but adaptation is obtained via Lepski’s method). In particular, [38] show that spike-and-slab priors on wavelet coefficients yield the *exact* adaptive minimax rate in the white noise model and [60] considers the anisotropic case in a regression framework. For density estimation, [13, 14] derive optimal $\|\cdot\|_\infty$ -rates for Pólya tree priors, while [44] considers adaptation for log-density spike and slab priors. In this work, we consider Gaussian white noise and non-parametric regression with Bayesian CART which is widely used in practice.

Bayesian CART is a method of function estimation based on ideas of recursive partitioning of the predictor space. The work [25] highlighted the link between dyadic CART and best ortho-basis selection using Haar wavelets in two dimensions; [27] furthered this connection by considering unbalanced Haar wavelets of [33]. CART methods have been also studied in the machine learning literature, see e.g. [7, 51, 58] and references therein. Unlike plain wavelet shrinkage methods and standard spike-and-slab priors, general Bayesian CART priors have extra flexibility by allowing for (some) *basis selection*. First results in this direction are derived in Section 4. This aspect is particularly useful in higher-dimensional data, where CART methods have been regarded as an attractive alternative to other methods [26].

By taking the Bayesian point of view, we relate Bayesian CART to structured wavelet shrinkage using libraries of *weakly* balanced Haar bases. Each tree provides an underlying skeleton or a ‘sparsity structure’ which supervises the sparsity pattern (see e.g. [2]). We show that Bayesian CART borrows strength between coefficients in the tree ancestry by giving rise to a variant of the *g-prior* [64]. Similarly as independent product priors, we show that these dependent priors *also* lead to adaptive supremum norm concentration rates (up to a logarithmic factor). To illustrate that local (internal) sparsity is a key driver of adaptivity, we show that dense trees are incapable of adaptation.

To convey the main ideas, the mathematical development will be performed through the lense of a Gaussian white noise model. Our techniques, however, also apply in non-parametric regression. Results in this setting are briefly presented in Section 3.4 with details

postponed until the Supplement (Section 7.1). The white noise model is defined through the following stochastic differential equation, for an integer $n \geq 1$,

$$(1) \quad dX(t) = f_0(t)dt + \frac{1}{\sqrt{n}}dW(t), \quad t \in [0, 1],$$

where $X(t)$ is an observation process, $W(t)$ is the standard Wiener process on $[0, 1]$ and f_0 is unknown and belongs to $L^2[0, 1]$, set of squared-integrable functions on $[0, 1]$. The model (1) is observationally equivalent to a Gaussian sequence space model after projecting the observation process onto a wavelet basis $\{\psi_{lk} : l \geq 0, 0 \leq k \leq 2^l - 1\}$ of $L^2[0, 1]$. This sequence model writes as

$$(2) \quad X_{lk} = \beta_{lk}^0 + \frac{\varepsilon_{lk}}{\sqrt{n}}, \quad \varepsilon_{lk} \stackrel{iid}{\sim} \mathcal{N}(0, 1),$$

where the wavelet coefficients $\beta_{lk}^0 = \langle f_0, \psi_{lk} \rangle = \int_0^1 f_0(t)\psi_{lk}(t)dt$ of f_0 are indexed by a scale index $l \geq -1$ and a location index $k \in \{0, \dots, (2^l - 1)_+\}$. A paradigmatic example is the standard Haar wavelet basis

$$(3) \quad \psi_{-10}(x) = \mathbb{I}_{[0,1]}(x) \quad \text{and} \quad \psi_{lk}(x) = 2^{l/2}\psi(2^l x - k) \quad (l \geq 0),$$

obtained with orthonormal dilation-translations of $\psi = \mathbb{I}_{(0,1/2]} - \mathbb{I}_{(1/2,1]}$, where \mathbb{I}_A denotes the indicator of a set A . Later in the text, we also consider weakly balanced Haar wavelet relaxations (Section 4), as well as smooth wavelet bases (Section 10.2).

One of the key motivations behind the Bayesian approach is the mere fact that the posterior is an actual distribution, whose limiting shape can be analyzed towards obtaining *uncertainty quantification* and inference. Our results in this direction can be grouped in two subsets. First, for uncertainty quantification for f_0 itself, we construct adaptive and honest confidence bands under self-similarity (with coverage converging to one). Exact asymptotic coverage is achieved through intersections with a multiscale credible band (along the lines of [47]). Confidence bands construction for regression surfaces is a fundamental task in nonparametric regression and can indicate whether there is empirical evidence to support conjectured features such as multi-modality or exceedance of a level. Results of this type are, to date, unavailable for classical CART, random forests and/or deep learning. Second, we consider inference for smooth functionals of f_0 , including linear ones and the primitive functional $\int_0^\cdot f_0$, for which exact optimal confidence sets are derived from posterior quantiles. While these results for functionals are stated in the main paper (Theorem 3 below), their derivation is most naturally obtained through a general limiting shape result, stated and proved in the Supplement (Theorem 7). Such an adaptive Bernstein-von Mises theorem for Bayesian CART is obtained following the approach of [15, 16]; it is only the second result of this kind (providing *adaptation*) after the recent result of Ray [47].

The paper is structured as follows. Section 2 introduces regression tree-priors, as well as the notion of tree-shaped sparsity and the g -prior for trees. In Section 3, we state supremum-norm inference properties of Bayesian dyadic CART (estimation and confidence bands). Section 4 considers flexible partitionings allowing for basis choice. A brief discussion can be found in Section 5. The proof of our master Theorem 1 can be found in Section 6. The supplementary file [17] gathers the proofs of the remaining results. The sections and equations of this supplement are referred to with an additional symbol “S-” in the numbering.

Notation. Let $\mathcal{C}([0, 1])$ denote the set of continuous functions on $[0, 1]$ and let ϕ_σ denote the normal density with zero mean and variance σ^2 . Let $\mathbb{N} = \{0, 1, 2, \dots\}$ be the set of natural integers and $\mathbb{N}^* = \mathbb{N} \setminus \{0\}$. We denote by I_K the $K \times K$ identity matrix. Also, B^c denotes the complement of a set B . For an interval $I = (a, b] \subset (0, 1]$, let $|I| = b - a$ be its diameter and $a \vee b = \max(a, b)$. The notation $x \lesssim y$ means $x \leq Cy$ for C a large enough universal constant, and $:=$ (or $=:$) means “the left-hand side is defined as”.

2. Trees and Wavelets. In this section, we discuss multiscale prior assignments on functions $f \in L^2[0, 1]$ (i.e. priors on the sequence of wavelet coefficients $\beta_{lk} = \langle f, \psi_{lk} \rangle$) inspired by (and including) Bayesian CART. Such methods recursively subdivide the predictor space into cells where f can be estimated locally. The partitioning process can be captured with a tree object (a hierarchical collection of nodes) and a set of splitting rules attached to each node. Section 2.1 discusses priors on the tree object. The splitting rules are ultimately tied to a chosen basis, where the traditional Haar wavelet basis yields deterministic dyadic splits (as we explain in Section 2.1.2). Later in Section 4, we extend our framework to random unbalanced Haar bases which allow for more flexible splits. Beyond random partitioning, an integral component of CART methods are histogram heights assigned to each partitioning cell. We flesh out connections between Bayesian histograms and wavelets in Section 2.2. Finally, we discuss Bayesian CART priors over histogram heights in Section 2.3.

2.1. Priors on Trees $\Pi_{\mathbb{T}}(\cdot)$. First, we need to make precise our definition of a tree object which will form a skeleton of our prior on (β_{lk}) for each given basis $\{\psi_{lk}\}$. Throughout this paper, we will largely work with the Haar basis.

DEFINITION 1 (Tree terminology). We define a binary tree \mathcal{T} as a collection of nodes (l, k) , where $l \in \mathbb{N}$, $k \in \{0, \dots, 2^l - 1\}$, that satisfies

$$(l, k) \in \mathcal{T}, l \geq 1 \Rightarrow (l-1, \lfloor k/2 \rfloor) \in \mathcal{T}.$$

In the last display, the node (l, k) is a child of its parent node $(l-1, \lfloor k/2 \rfloor)$. A full binary tree consists of nodes with exactly 0 or 2 children. For a node (l, k) , we refer to l as the layer index (or also depth) and k as the position in the l^{th} layer (from left to right). The cardinality $|\mathcal{T}|$ of a tree \mathcal{T} is its total number of nodes and the depth is defined as $d(\mathcal{T}) = \max_{(l,k) \in \mathcal{T}} l$.

A node $(l, k) \in \mathcal{T}$ belongs to the set \mathcal{T}_{ext} of *external* nodes (also called *leaves*) of \mathcal{T} if it has no children and to the set \mathcal{T}_{int} of *internal* nodes, otherwise. By definition $|\mathcal{T}| = |\mathcal{T}_{int}| + |\mathcal{T}_{ext}|$, where, for full binary trees, we have $|\mathcal{T}| = 2|\mathcal{T}_{int}| + 1$. An example of a full binary tree is depicted in Figure 1(a). In the sequel, \mathbb{T} denotes the set of full binary trees of depth no larger than $L = L_{max} = \lfloor \log_2 n \rfloor$, a typical cut-off in wavelet analysis. Indeed, trees can be associated with certain wavelet decompositions, as will be seen in Section 2.2.2.

Before defining tree-structured priors over the entire functions f 's, we first discuss various ways of assigning a prior distribution over \mathbb{T} , that is over trees themselves. We focus on the Bayesian CART prior [20], which became an integral component of many Bayesian tree regression methods including BART [19].

2.1.1. Bayesian CART Priors. The Bayesian CART construction of [20] assigns a prior over \mathbb{T} via the heterogeneous Galton-Watson (GW) process. The prior description utilizes the following top-down left-to-right exploration metaphor (see also [48]). Denote with Q a queue of nodes waiting to be explored. Each node (l, k) is assigned a random binary indicator $\gamma_{lk} \in \{0, 1\}$ for whether or not it is split. Starting with $\mathcal{T} = \emptyset$, one initializes the exploration process by putting the root node $(0, 0)$ tentatively in the queue, i.e. $Q = \{(0, 0)\}$. One then repeats the following three steps until $Q = \emptyset$:

- (a) Pick a node $(l, k) \in Q$ with the highest priority (i.e. the smallest index $2^l + k$) and if $l < L_{max}$, split it with probability

$$(4) \quad p_{lk} = \mathbb{P}(\gamma_{lk} = 1).$$

If $l = L_{max}$, set $\gamma_{lk} = 0$.

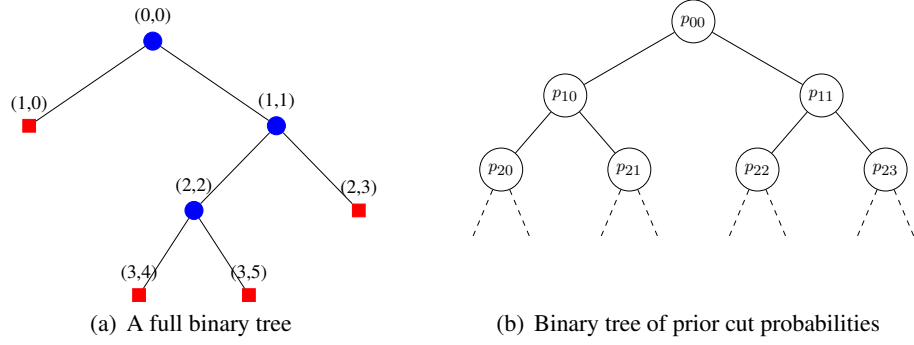


FIG 1. (Left) A full binary tree $\mathcal{T} = \mathcal{T}_{int} \cup \mathcal{T}_{ext}$. Red nodes are external nodes \mathcal{T}_{ext} and blue nodes are internal nodes \mathcal{T}_{int} . (Right) A binary tree of cut probabilities p_{lk} in (4).

(b) If $\gamma_{lk} = 0$, remove (l, k) from Q .

(c) If $\gamma_{lk} = 1$, then

(i) add (l, k) to the tree, i.e. $\mathcal{T}_{int} \leftarrow \mathcal{T}_{int} \cup \{(l, k)\}$,

(ii) remove (l, k) from Q and if $l < L_{max}$ add its children to Q , i.e.

$$Q \leftarrow Q \setminus \{(l, k)\} \cup \{(l+1, 2k), (l+1, 2k+1)\}.$$

The tree skeleton is probabilistically underpinned by the cut probabilities (p_{lk}) which are typically assumed to decay with the depth l as a way to penalise too complex trees. While [20] suggest $p_{lk} = \alpha/(1+l)^\gamma$ for some $\alpha \in (0, 1)$ and $\gamma > 0$, [48] point out that this decay may not be fast enough and suggest instead $p_{lk} = \Gamma^{-l}$ for some $2 < \Gamma < n$, which leads to a (near) optimal empirical L^2 -convergence rate. We use a similar assumption in our analysis, and also assume that the split probability depends only on l , and simply denote $p_l = p_{lk}$.

Independently of [20], [23] proposed another variant of Bayesian CART, which first draws the number of leaves (i.e. external nodes) $K = |\mathcal{T}_{ext}|$ at random from a certain prior on integers, e.g. a Poisson distribution (say, conditioned to be non-zero). Then, a tree \mathcal{T} is sampled uniformly at random from all full binary trees with K leaves. Noting that there are \mathbb{C}_{K-1} such trees, with \mathbb{C}_K the K -th Catalan number (see Lemma 6), this leads to $\Pi(\mathcal{T}) = (\lambda^K/[K!(e^\lambda - 1)]) \cdot \mathbb{C}_{K-1}^{-1}$. As we restrict to trees in \mathbb{T} , i.e. with depth at most $L = L_{max}$, we slightly update the previous prior choice by setting, for some $\lambda > 0$, with $K = |\mathcal{T}_{ext}|$,

$$(5) \quad \Pi_{\mathbb{T}}(\mathcal{T}) \propto \frac{\lambda^K}{(e^\lambda - 1)K!} \frac{1}{\mathbb{C}_{K-1}} \mathbb{I}_{\mathcal{T} \in \mathbb{T}},$$

where \propto means ‘proportional to’. We call the resulting prior $\Pi_{\mathbb{T}}$ the ‘conditionally uniform prior’ with a parameter λ .

2.1.2. Trees and Random Partitions. Trees provide a structured framework for generating random partitions of the predictor space (here we choose $(0, 1]$ for simplicity of exposition). In CART methodology, each node $(l, k) \in \mathcal{T}$ is associated with a partitioning interval $I_{lk} \subseteq (0, 1]$. Starting from the trivial partition $I_{00} = (0, 1]$, the simplest way to obtain a partition is by successively dividing each I_{lk} into $I_{lk} = I_{l+1, 2k} \cup I_{l+1, 2k+1}$. One central example is *dyadic* intervals I_{lk} which correspond to the domain of the balanced Haar wavelets ψ_{lk} in (3), i.e.

$$(6) \quad I_{00} = (0, 1], \quad I_{lk} = (k2^{-l}, (k+1)2^{-l}] \quad \text{for } l \geq 0 \text{ and } 0 \leq k < 2^l.$$

For any fixed depth $l \in \mathbb{N}$, the intervals $\cup_{0 \leq k < 2^l} I_{lk}$ form a deterministic regular (equispaced) partition of $(0, 1]$. Trees, however, generate *more flexible* partitions $\cup_{(l,k) \in \mathcal{T}_{ext}} I_{lk}$ by keeping

only those intervals I_{lk} attached to the leaves of the tree. Since \mathcal{T} is treated as random with a prior $\Pi_{\mathbb{T}}$ (as defined in Section 2.1), the resulting partition will also be random.

EXAMPLE 1. Figure 1(a) shows a full binary tree $\mathcal{T} = \mathcal{T}_{int} \cup \mathcal{T}_{ext}$, where $\mathcal{T}_{int} = \{(0,0), (1,1), (2,2)\}$ and $\mathcal{T}_{ext} = \{(1,0), (2,3), (3,4), (3,5)\}$, resulting in the partition of $(0,1]$ given by

$$(7) \quad (I_{lk})_{(l,k) \in \mathcal{T}_{ext}} = \{(0, 1/2], (1/2, 5/8], (5/8, 3/4], (3/4, 1]\}.$$

The set of possible *split points* obtained with (6) is confined to dyadic rationals. One can interpret the resulting partition as the result of recursive splitting where, at each level l , intervals I_{lk} for each internal node $(l,k) \in \mathcal{T}_{int}$ are cut in half and intervals I_{lk} for each external node $(l,k) \in \mathcal{T}_{ext}$ are left alone. We will refer to such a recursive splitting process as *dyadic CART*. There are several ways to generalize this construction, for instance by considering arbitrary splitting rules that iteratively dissect the intervals at values other than the midpoint. We explore such extensions in Section 4.

2.2. *Tree-shaped Priors on f .* This section outlines two strategies for assigning a tree-shaped prior distribution on f underpinned by a tree skeleton $\mathcal{T} \in \mathbb{T}$. Each tree $\mathcal{T} = \mathcal{T}_{int} \cup \mathcal{T}_{ext}$ can be associated with two sets of coefficients: (a) *internal* coefficients β_{lk} attached to wavelets ψ_{lk} for $(l,k) \in \mathcal{T}_{int}$ and (b) *external* coefficients $\tilde{\beta}_{lk}$ attached to partitioning intervals I_{lk} for $(l,k) \in \mathcal{T}_{ext}$ (see Section 2.1.2). While wavelet priors (Section 2.2.1) assign the prior distribution internally on β_{lk} , Bayesian CART priors [20, 23] (Section 2.2.2) assign the prior externally on $\tilde{\beta}_{lk}$. We discuss and relate these two strategies in more detail below.

2.2.1. *Tree-shaped Wavelet Priors.* Traditional (linear) Haar wavelet reconstructions for f deploy *all* wavelet coefficients β_{lk} with resolutions l smaller than some $d > 0$. This strategy amounts to fitting a *flat tree* with d layers (i.e. a tree that contains all nodes up to a level d , see Figure 2) or, equivalently, a regular dyadic regression histogram with 2^d bins. This construction can be made more flexible by selecting coefficients prescribed by trees that are not necessarily flat. Given a full binary tree $\mathcal{T} \in \mathbb{T}$, one can build the following wavelet reconstruction of f using *only* active wavelet coefficients that are *inside* a tree \mathcal{T}

$$(8) \quad f_{\mathcal{T}, \beta}(x) = \beta_{-10} \psi_{-10}(x) + \sum_{(l,k) \in \mathcal{T}_{int}} \beta_{lk} \psi_{lk}(x) = \sum_{(l,k) \in \mathcal{T}'_{int}} \beta_{lk} \psi_{lk}(x),$$

where $\beta = (\beta_{-10}, (\beta_{lk})_{0 \leq l \leq L-1, 0 \leq k < 2^l})'$ is a vector of wavelet coefficients and where $\mathcal{T}'_{int} = \mathcal{T}_{int} \cup \{(-1,0)\}$ is the ‘rooted’ tree with the index $(-1,0)$ added to \mathcal{T}_{int} . Note that $|\mathcal{T}'_{int}| = |\mathcal{T}_{ext}|$.

Define a *tree-shaped wavelet prior* on $f_{\mathcal{T}, \beta}$ as the prior induced by the hierarchical model

$$(9) \quad \begin{aligned} \mathcal{T} &\sim \Pi_{\mathbb{T}} \\ (\beta_{lk})_{lk} | \mathcal{T} &\sim \bigotimes_{(l,k) \in \mathcal{T}'_{int}} \pi(\beta_{lk}) \otimes \bigotimes_{(l,k) \notin \mathcal{T}'_{int}} \delta_0(\beta_{lk}), \end{aligned}$$

where $\Pi_{\mathbb{T}}$ is a prior on trees as described in Section 2.1.1 and where the active wavelet coefficients β_{lk} for $(l,k) \in \mathcal{T}_{int}$ follow a distribution with a bounded and positive density $\pi(\beta_{lk})$ on \mathbb{R} . The prior (9) is seen as a distribution on \mathbb{R}^{2^L} , where all remaining coefficients, i.e. β_{lk} ’s for $l \geq L$, are set to 0.

The prior (9) contains the so-called *sieve priors* [16] (i.e. flat trees) as a special case, where the sieve is with respect to the approximating spaces $\text{Vect}\{\psi_{lk}, l < d\}$ for some $d \geq 0$. For nonparametric estimation of f_0 , it is well-known that sieve priors can achieve (nearly)

adaptive rates in the L^2 -sense (see e.g. [30]). It turns out, however, that sieve priors (and therefore flat tree priors) are too rigid to enable adaptive results for stronger losses such as the supremum norm, as we demonstrate in Theorem 9 in Section 7.5 (Supplement). This theorem illustrates that supremum norm adaptation using Bayesian (or other likelihood-based) methods is a delicate phenomenon that is not attainable by many typical priors.

By definition, the prior (9) weeds out all wavelet coefficients β_{lk} that are not supported by the tree skeleton (i.e. are not *internal* nodes in \mathcal{T}). This has two shrinkage implications: global and local. First, the global level of truncation (i.e. the depth of the tree) in (9) is not fixed but random. Second, unlike in sieve priors, only some low resolution coefficients are active depending on whether or not the tree splits the node (l, k) . These two shrinkage aspects create hope that tree-shaped wavelet priors (9) attain adaptive adaptive supremum norm rates (up to log factors) and enable construction of adaptive confidence bands. We see later in Section 3 that this optimism is indeed warranted.

For adaptive wavelet shrinkage, [21] propose a Gaussian mixture spike-and-slab prior on the wavelet coefficients. The point mass spike-and-slab incarnation of this prior was studied by [38] and [47]. Independently for each wavelet coefficient β_{lk} at resolutions larger than some $l_0(n)$ (strictly increasing sequence), the prior in [47] can be written in the standard spike-and-slab form

$$(10) \quad \pi(\beta_{lk} | \gamma_{lk}) = \gamma_{lk} \pi(\beta_{lk}) + (1 - \gamma_{lk}) \delta_0(\beta_{lk}),$$

where $\gamma_{lk} \in \{0, 1\}$ for whether or not the coefficient is active with $\mathbb{P}(\gamma_{lk} = 1 | \theta_l) = \theta_l$. Moreover, the prior on all coefficients at resolutions no larger than $l_0(n)$ is dense, i.e. $\theta_l = 1$ for $l \leq l_0(n)$. The value θ_l can be viewed as the probability that a given wavelet coefficient β_{lk} at resolution l will contain ‘signal’.

There are undeniable similarities between (9) and (10), in the sense that the binary inclusion indicator γ_{lk} in (10) can be regarded as the node splitting indicator γ_{lk} in (4). While the indicators γ_{lk} in (10) are *independent* under the spike-and-slab prior, they are hierarchically constrained under the CART prior, where the pattern of non-zeroes encodes the tree oligarchy. The seeming resemblance of the CART-type prior (9) to the spike-and-slab prior (10) makes one naturally wonder whether, unlike sieve-type priors, CART posteriors attain adaptive supremum-norm inference.

2.2.2. Bayesian CART Priors. A perhaps more transparent approach to assigning a tree-shaped prior on f is through histograms (as opposed to wavelet reconstructions from Section 2.2.1). Each tree $\mathcal{T} \in \mathbb{T}$ generates a random partition via intervals I_{lk} (see Section 2.1.2) and gives rise to the following histogram representation

$$(11) \quad \tilde{f}_{\mathcal{T}, \tilde{\beta}}(x) = \sum_{(l,k) \in \mathcal{T}_{ext}} \tilde{\beta}_{lk} \mathbb{I}_{I_{lk}}(x),$$

where $\tilde{\beta} = (\tilde{\beta}_{lk} : (l, k) \in \mathcal{T}_{ext})'$ is a vector of reals interpreted as step heights and where I_{lk} ’s are obtained from the tree \mathcal{T} as in Section 2.1.2 (and as illustrated in Example 1). We now define the (*Dyadic*) *Bayesian CART prior* on f using the following hierarchical model on the *external* coefficients rather than *internal* coefficients (compare with (9))

$$(12) \quad \begin{aligned} \mathcal{T} &\sim \Pi_{\mathbb{T}} \\ (\tilde{\beta}_{lk})_{(l,k) \in \mathcal{T}_{ext}} | \mathcal{T} &\sim \bigotimes_{(l,k) \in \mathcal{T}_{ext}} \tilde{\pi}(\tilde{\beta}_{lk}), \end{aligned}$$

where $\Pi_{\mathbb{T}}$ is as in Section 2.1, and where the height $\tilde{\beta}_{lk}$ at a specific $(l, k) \in \mathcal{T}_{ext}$ has a bounded and positive density $\tilde{\pi}(\tilde{\beta}_{lk})$ on \mathbb{R} . This model *coincides* with the widely used




FIG 2. Flat tree with edges weighted by the amplitude of the Haar wavelets.

Bayesian CART priors using a midpoint dyadic splitting rule (as we explained in Section 2.1.2). In practice, the density $\tilde{\pi}$ is often chosen as centered Gaussian with some variance $\sigma^2 > 0$ [20, 23].

The histogram prior (11) can be rephrased in terms of wavelets. Indeed, the histogram representation (11) can be rewritten in terms of the *internal* coefficients, i.e. $\tilde{f}_{\mathcal{T},\tilde{\beta}}(x) = f_{\mathcal{T},\beta}(x)$ as in (8), with β_{lk} 's and $\tilde{\beta}_{lk}$'s linked via

$$(13) \quad \tilde{\beta}_{lk} = \beta_{-10} + \sum_{j=0}^{l-1} s_{\lfloor k/2^{l-j-1} \rfloor} 2^{j/2} \beta_{j \lfloor k/2^{l-j} \rfloor},$$

where $s_k = (-1)^{k+1}$. The identity (13) follows the fact that for $x \in I_{lk}$ we obtain $\tilde{\beta}_{lk} = \sum_{(l',k') \in P_{lk}} \beta_{l'k'} \psi_{l'k'}$ from (11), where $P_{lk} \equiv \{(j, \lfloor k/2^{l-j} \rfloor) : j = 0, \dots, l-1\}$ are the ancestors of the bottom node (l, k) . Note that $\psi_{j \lfloor k/2^{l-j} \rfloor} = 2^{j/2} s_{\lfloor k/2^{l-j-1} \rfloor}$ where $s = (-1)^{k+1}$ for whether x belongs to the left (positive sign) or right (negative sign) of the wavelet piece. There is a pinball game metaphor behind (13). A ball is dropped through a series of dyadically arranged pins of which the ball can bounce off to the right (when $s_k = +1$) or to the left (when $s_k = -1$). The ball ultimately lands in one of the histogram bins I_{lk} whose coefficient $\tilde{\beta}_{lk}$ is obtained by aggregating β_{lk} 's of those pins (l, k) that the ball encountered on its way down. The pinball aggregation process can be understood from Figure 3. The duality between the equivalent representations (11) and (8) through (13) provides various avenues for constructing prior distributions, and enables an interesting interpretation of Bayesian CART [20, 23] as a correlated wavelet prior, as we now see.

2.3. The g -prior for Trees. We now discuss various ways of assigning a prior distribution on the bottom node histogram heights $\tilde{\beta}_{lk}$ and, equivalently, the internal Haar wavelet coefficients β_{lk} . This section also describes an interesting connection between the widely used Bayesian CART prior [20, 23] and a g -prior [64] on wavelet coefficients. For a given tree \mathcal{T} , let $\beta_{\mathcal{T}} = (\beta_{lk} : (l, k) \in \mathcal{T}'_{int})'$ denote the vector of *ordered internal* node coefficients β_{lk} including the extra root node $(-1, 0)$ (and with ascending ordering according to $2^l + k$). Similarly, $\tilde{\beta}_{\mathcal{T}} = (\beta_{lk} : (l, k) \in \mathcal{T}_{ext})'$ is the vector of *ordered external* node coefficients $\tilde{\beta}_{lk}$. The duality between $\beta_{\mathcal{T}}$ and $\tilde{\beta}_{\mathcal{T}}$ is apparent from the pinball equation (13) written in matrix form

$$(14) \quad \tilde{\beta}_{\mathcal{T}} = A_{\mathcal{T}} \beta_{\mathcal{T}},$$

where $A_{\mathcal{T}}$ is a square $|\mathcal{T}_{ext}| \times |\mathcal{T}'_{int}|$ matrix (noting $|\mathcal{T}_{ext}| = |\mathcal{T}'_{int}|$), further referred to as the *pinball matrix*. Each row of $A_{\mathcal{T}}$ encodes the ancestors of the external node, where the nonzero entries correspond to the internal nodes in the family pedigree. The entries are rescaled, where younger ancestors are assigned more weight. For example, the tree \mathcal{T} in Figure 3(a) induces




FIG 3. (a) Example of a full binary tree, edges weighted by the amplitude of the Haar wavelets. (b) Pinball matrix of the tree in (a).

a pinball matrix $A_{\mathcal{T}}$ in Figure 3(b). The pinball matrix $A_{\mathcal{T}}$ can be easily expressed in terms of a diagonal matrix and an orthogonal matrix as

$$(15) \quad A_{\mathcal{T}} A'_{\mathcal{T}} = D_{\mathcal{T}}, \quad \text{where} \quad D_{\mathcal{T}} = \text{diag}\{\tilde{d}_{lk, lk}\}_{(l,k) \in \mathcal{T}_{ext}}, \quad \tilde{d}_{lk, lk} = 2^l.$$

This results from the fact that the collection $(2^{l/2} \mathbb{I}_{lk}, (l, k) \in \mathcal{T}_{ext})$ is an orthonormal system spanning the same space as $(\psi_{jk}, (j, k) \in \mathcal{T}'_{int})$, so $D_{\mathcal{T}}^{-1/2} A_{\mathcal{T}}$ is an orthonormal change-of-basis matrix. We now exhibit precise connections between the theoretical wavelet prior (9) which draws $\beta_{lk} \sim \pi$ and the practical Bayesian CART histogram prior which draws $\tilde{\beta}_{lk} \sim \tilde{\pi}$.

Recall that the wavelet prior (9) assumes independent wavelet coefficients, e.g. through the standard Gaussian prior $\beta_{\mathcal{T}} \sim \mathcal{N}(0, I_{|\mathcal{T}_{ext}|})$. Starting *from within* the tree, this translates into the following independent product prior on the bottom coefficients $\tilde{\beta}_{lk}$ through (14)

$$(16) \quad \tilde{\beta}_{\mathcal{T}} \sim \mathcal{N}(0, D_{\mathcal{T}}), \quad \text{where} \quad D_{\mathcal{T}} \quad \text{was defined in (15),}$$

i.e. $\text{var } \tilde{\beta}_{lk} = 2^l$ where the variances increase with the resolution l .

The Bayesian CART prior [20, 23], on the other hand, starts *from outside* the tree by assigning $\tilde{\beta}_{\mathcal{T}} \sim \mathcal{N}(0, g_n I_{|\mathcal{T}_{ext}|})$ for some $g_n > 0$, ultimately setting the bottom node variances equal. This translates into the following ‘ g -prior’ on the *internal* wavelet coefficients through the duality (14).

DEFINITION 2. Let $\mathcal{T} \in \mathbb{T}$ with a pinball matrix $A_{\mathcal{T}}$ and denote with $\beta_{\mathcal{T}}$ the internal wavelet coefficients. We define the g -prior for trees as

$$(17) \quad \beta_{\mathcal{T}} \sim \mathcal{N}(0, g_n (A'_{\mathcal{T}} A_{\mathcal{T}})^{-1}) \quad \text{for some } g_n > 0.$$

Note that, except for very special cases (e.g. flat trees) $A'_{\mathcal{T}} A_{\mathcal{T}}$ is in general not diagonal, unlike $A_{\mathcal{T}} A'_{\mathcal{T}}$. This means that the correlation structure induced by the Bayesian CART prior on internal wavelet coefficients is non-trivial, although $A'_{\mathcal{T}} A_{\mathcal{T}}$ admits some partial sparsity. We characterize basic properties of the pinball matrix in Section 8.1 in the Supplement. For example, Proposition 3 shows that matrices $A'_{\mathcal{T}} A_{\mathcal{T}}$ and $A_{\mathcal{T}} A'_{\mathcal{T}}$ have the same eigenspectrum consisting of values 2^l where l corresponds to the depth of the bottom nodes. This means that the g -prior variances (diagonal elements of $g_n (A'_{\mathcal{T}} A_{\mathcal{T}})^{-1}$) are lower-bounded by the minimal eigenvalue of $g_n (A'_{\mathcal{T}} A_{\mathcal{T}})^{-1}$ which equals $g_n 2^{-l}$ (where l is the depth of the deepest external node) which is lower-bounded by g_n/n . Since the traditional wavelet prior assumes variance 1, the choice $g_n = n$ matches the lower bound 1 by undersmoothing all possible variance combinations. While other choices could be potentially used (see [?, ?, ?] in the context of linear regression), we will consider $g_n = n$ in our results below.

We regard (17) as the ‘ g -prior for trees’ due to its apparent similarity to g -priors for linear regression coefficients [64]. The g -prior has been shown to have many favorable properties in terms of invariance or predictive matching [5, 4]. Here, we explore the benefits of the g -type correlation structure in the context of structured wavelet shrinkage where each ‘model’ is defined by a tree topology. The correlation structure (17) makes this prior very different from any other prior studied in the context of wavelet shrinkage.

3. Inference with (Dyadic) Bayesian CART. In this section we investigate the inference properties of tree-based posteriors, showing that (a) they attain the minimax rate of posterior concentration in the supremum-norm sense (up to a log factor), and (b) enable uncertainty quantification: for f in the form of adaptive confidence bands, and for smooth functionals thereof, in terms of Bernstein-von Mises type results. For clarity of exposition, we focus now on the one-dimensional case, but the results readily extend to the multi-dimensional setting with \mathbb{R}^d , $d \geq 1$ fixed, as predictor space; see Section 7.6 for more details.

3.1. Posterior supremum-norm convergence. Let us recall the standard inequality (see e.g. (57) below), for f_0 a continuous function and f a Haar histogram (8), with coefficients β_{lk}^0 and β_{lk} ,

$$(18) \quad \|f - f_0\|_\infty \leq |\beta_{-10} - \beta_{-10}^0| + \sum_{l \geq -1} 2^{l/2} \max_{0 \leq k < 2^l} |\beta_{lk} - \beta_{lk}^0| =: \ell_\infty(f, f_0).$$

As ℓ_∞ dominates $\|\cdot\|_\infty$, it is enough to derive results for the ℓ_∞ -loss.

Given a tree $\mathcal{T} \in \mathbb{T}$, and recalling that trees in \mathcal{T} have depth at most $L := L_{\max} = \lfloor \log_2 n \rfloor$, we consider a generalized tree-shaped prior Π on the *internal* wavelet coefficients, recalling the notation \mathcal{T}'_{int} from Section 2.2,

$$(19) \quad \begin{aligned} \mathcal{T} &\sim \Pi_{\mathbb{T}} \\ (\beta_{lk})_{l \leq L, k < 2^l} \mid \mathcal{T} &\sim \pi(\beta_{\mathcal{T}}) \otimes \bigotimes_{(l,k) \notin \mathcal{T}'_{int}} \delta_0(\beta_{lk}), \end{aligned}$$

where $\pi(\beta_{\mathcal{T}})$ is a law to be chosen on $\mathbb{R}^{|\mathcal{T}'_{int}|}$, not necessarily of a product form. This is a generalization of (9), which allows for *correlated* wavelet coefficients (e.g. the g -prior). Let $\mathbf{X}_{\mathcal{T}}$ denote the vector of ordered responses X_{lk} in (2) for $(l, k) \in \mathcal{T}'_{int}$. From the white noise model, we have

$$\mathbf{X}_{\mathcal{T}} = \beta_{\mathcal{T}} + \frac{1}{\sqrt{n}} \varepsilon_{\mathcal{T}}, \quad \text{with} \quad \varepsilon_{\mathcal{T}} \sim \mathcal{N}(0, I_{|\mathcal{T}_{ext}|}) \quad (\text{given } \mathcal{T}).$$

By Bayes’ formula, the posterior distribution $\Pi[\cdot \mid X]$ of the variables $(\beta_{lk})_{l \leq L, k}$ has density

$$(20) \quad \sum_{\mathcal{T} \in \mathbb{T}} \Pi[\mathcal{T} \mid X] \cdot \pi(\beta_{\mathcal{T}} \mid X) \cdot \prod_{(l,k) \notin \mathcal{T}'_{int}} \mathbb{I}_0(\beta_{lk}),$$

where, denoting as shorthand $N_X(\mathcal{T}) = \int e^{-\frac{n}{2} \|\beta_{\mathcal{T}}\|_2^2 + n \mathbf{X}'_{\mathcal{T}} \beta_{\mathcal{T}}} \pi(\beta_{\mathcal{T}}) d\beta_{\mathcal{T}}$,

$$(21) \quad \pi(\beta_{\mathcal{T}} \mid X) = \frac{e^{-\frac{n}{2} \|\beta_{\mathcal{T}}\|_2^2 + n \mathbf{X}'_{\mathcal{T}} \beta_{\mathcal{T}}} \pi(\beta_{\mathcal{T}})}{N_X(\mathcal{T})},$$

$$(22) \quad \Pi[\mathcal{T} \mid X] = \frac{W_X(\mathcal{T})}{\sum_{\mathcal{T} \in \mathbb{T}} W_X(\mathcal{T})}, \quad \text{with} \quad W_X(\mathcal{T}) = \Pi_{\mathbb{T}}(\mathcal{T}) N_X(\mathcal{T}).$$

Let us note that the sum in the last display is finite, as we restrict to trees of depth at most $L = L_{max}$. Note that the classes of priors $\Pi_{\mathbb{T}}$ from Section 2 are non-conjugate, i.e. the posterior on trees is given by the somewhat intricate expression (22) and does not belong to one of the classes of $\Pi_{\mathbb{T}}$ priors. While the posterior expression (21) allows for general priors $\pi(\beta_{\mathcal{T}})$, we will focus on conditionally conjugate Gaussian priors for simplicity. This assumption is not essential and can be relaxed. For instance, in case $\pi(\beta_{\mathcal{T}})$ is of a product form, one could use a product of e.g. Laplace distributions, using similar ideas as in [15], Theorem 5.

Our first result exemplifies the potential of tree-shaped priors by showing that Dyadic Bayesian CART achieves the minimax rate of posterior concentration over Hölder balls in the sup-norm sense, i.e. $\varepsilon_n = (n/\log n)^{-\alpha/(2\alpha+1)}$, up to a logarithmic term. Define Hölder-type spaces of functions on $[0, 1]$ as

$$(23) \quad \mathcal{H}(\alpha, M) := \left\{ f \in \mathcal{C}[0, 1] : \max_{l \geq 0, 0 \leq k < 2^l} 2^{l(\frac{1}{2} + \alpha)} |\langle f, \psi_{lk} \rangle| \vee |\langle f, \psi_{-10} \rangle| \leq M \right\}.$$

For balanced Haar wavelets ψ_{lk} as in (3), $\mathcal{H}(\alpha, M)$ contains the standard space of α -Hölder (resp. Lipschitz when $\alpha = 1$) functions for any $\alpha \in (0, 1]$, defined as

$$(24) \quad \mathcal{H}_M^\alpha := \left\{ f : \|f\|_\infty \leq M, \frac{|f(x) - f(y)|}{|x - y|^\alpha} \leq M \quad \forall x, y \in [0, 1] \right\}.$$

Our master rate-theorem, whose proof can be found in Section 6, is stated below. It will be extended in various directions in the sequel.

THEOREM 1. *Let $\Pi_{\mathbb{T}}$ be the Galton-Watson process prior from Section 2.1 with $p_{lk} = \Gamma^{-l}$ and $\Gamma > 2e^3$. Consider the tree-shaped wavelet prior (19) with $\pi(\beta_{\mathcal{T}}) \sim \mathcal{N}(0, \Sigma_{\mathcal{T}})$, where $\Sigma_{\mathcal{T}}$ is either $I_{|\mathcal{T}'_{int}|}$ or $g_n(A'_{\mathcal{T}} A_{\mathcal{T}})^{-1}$ with $g_n = n$. Define*

$$(25) \quad \varepsilon_n = \left(\frac{\log^2 n}{n} \right)^{\frac{\alpha}{2\alpha+1}} \quad \text{for } \alpha > 0.$$

Then for any $\alpha \in (0, 1]$, $M > 0$, any sequence $M_n \rightarrow \infty$ we have for $n \rightarrow \infty$

$$(26) \quad \sup_{f_0 \in \mathcal{H}(\alpha, M)} E_{f_0} \Pi[f_{\mathcal{T}, \beta} : \ell_\infty(f_{\mathcal{T}, \beta}, f_0) > M_n \varepsilon_n \mid X] \rightarrow 0.$$

By (18), the statement (26) also holds for the supremum loss $\|\cdot\|_\infty$.

Theorem 1 encompasses both original Bayesian CART proposals for priors on bottom coefficients $\tilde{\beta}_{\mathcal{T}} \sim \mathcal{N}(0, I_{|\mathcal{T}_{ext}|})$ (the case $\Sigma_{\mathcal{T}} = g_n(A_{\mathcal{T}} A'_{\mathcal{T}})^{-1}$ discussed in Section 2.3) as well as the mathematically slightly simpler wavelet priors $\Sigma_{\mathcal{T}} = I_{|\mathcal{T}_{ext}|}$ (discussed in Section 2.2.1). We did not fully optimize the constants in the statement; for instance, one can check that $\Gamma > 2$ for the g -prior works.

EXTENSION 1. *While Theorem 1 is formulated for Bayesian CART obtained with Haar wavelets, the concept of tree-shaped sparsity extends to general wavelets that give rise to smoother objects than just step functions. With $\{\psi_{lk}\}$ an S -regular wavelet basis on $[0, 1]$, e.g. the boundary-corrected wavelet basis of [22] (see also [32], Chapter 4), and with $f_0 \in \mathcal{H}(\alpha, M)$ defined in (23) for some $M > 0$ and arbitrary $0 < \alpha \leq S$, one indeed obtains the statement (26) by choosing $\Gamma \geq \Gamma_0(S) > 0$ or $c \geq c_0 > 0$ large enough, see Section 10.2.*

EXTENSION 2. *Theorem 1 holds for a variety of other tree priors. This includes the conditionally uniform prior mentioned in Section 2.1.1 with $\lambda = 1/n^c$ in (5), or an exponential-type prior $\Pi_{\mathbb{T}}(\mathcal{T}) \propto e^{-c|\mathcal{T}_{ext}|\log n} \mathbb{I}_{\mathcal{T} \in \mathbb{T}}$ for some $c > 0$. One can also assume a general Gaussian prior on active wavelet coefficients with an unstructured covariance matrix $\Sigma_{\mathcal{T}}$ which satisfies $\lambda_{\min}(\Sigma_{\mathcal{T}}) \gtrsim 1/\sqrt{\log n}$ and $\lambda_{\max}(\Sigma_{\mathcal{T}}) \lesssim n^a$ for some $a > 0$. Detailed proofs can be found in the Supplement (Section 10.1).*

Only very few priors (actually *only* point mass spike-and-slab based priors, as discussed in the Introduction) were shown to attain adaptive posterior sup-norm concentration rates. Theorem 1 now certifies Dyadic Bayesian CART as one of them. The rate ε_n in (25) coincides with the minimax rate for the supremum norm in the white noise model up to a logarithmic factor, which equals $(\log n)^{\frac{1}{2\alpha+1}}$. This means that there may be a slight logarithmic term to pay for tree priors (no more than $\sqrt{\log n}$ in the limit $\alpha \rightarrow \infty$ and no price in the limit $\alpha \rightarrow 0$). This log-factor is in fact real, i.e. it is *not* an artifact from the upper-bound proof. This is proved in Section 7.4, where a lower-bound result, implying that the rate ε_n is sharp, is established. The logarithmic penalty reflects that Bayesian CART priors occupy the middle ground between flat trees (with only a depth cutoff) and spike-and-slab priors (with general sparsity patterns). As mentioned earlier, flat trees are incapable of supremum-norm adaptation, as we formally prove in Section 7.5. The fact that the more flexible Bayesian CART priors still achieves supremum-norm adaptation in a near-optimal way is a rather notable feature. From a more general perspective, we note that while general tools are available to derive adaptive L^2 - or Hellinger-rate results in broad settings (e.g. model selection techniques, or the theory of posterior rates in [29]), deriving adaptive L^∞ -results is often obtained in a case-by-case basis; two possible techniques are wavelet thresholding (when empirical estimates of wavelet coefficients are available) and Lepski's method (which requires some 'ordered' set of estimators, typically in terms of variance; for tree-estimators for instance it would not readily be applicable). The fact that tree methods enable for supremum-norm adaptation in nonparametric settings is one of the main take-away messages of this work.

3.2. *Adaptive Honest Confidence Bands for f_0 .* We now turn to the ultimate landing point of this paper, uncertainty quantification for f_0 and its functionals. The existence of adaptive confidence sets in general is an interesting and delicate question (see Chapter 8 of [32]). In the present context of regression function estimation under the supremum norm loss, it is in fact impossible to build adaptive confidence bands without further restricting the parameter space. We do so by imposing some classical self-similarity conditions (see [32], [47] for more details).

DEFINITION 3. (*Self-similarity*) *Given an integer $j_0 > 0$, we say that $f \in \mathcal{H}(\alpha, M)$ is self-similar if, for some constant $\varepsilon > 0$,*

$$(27) \quad \|K_j(f) - f\|_\infty \geq \varepsilon 2^{-j\alpha} \quad \text{for all } j \geq j_0,$$

where $K_j(f) = \sum_{l \leq j-1} \sum_k \langle \psi_{lk}, f \rangle \psi_{lk}$ is the wavelet projection at level j . The class of all such self-similar functions will be denoted by $\mathcal{H}_{SS}(\alpha, M, \varepsilon)$.

Section 8.3.3 in [32] describes self-similar functions as typical representatives of the Hölder class. As shown in Proposition 8.3.21 of [32], self-dissimilar functions are nowhere dense in the sense that they cannot approximate any open set in $\mathcal{H}(\alpha, M)$. In addition, Bayesian non-parametric priors for Hölder functions charge self-similar functions with probability 1. Finally, self-similarity does not affect the difficulty of the statistical estimation

problem, where the (ℓ_∞) minimax rate is not changed after adding this assumption. A variant of the self-similarity condition was shown to be *necessary* for adaptive inference, in that such condition cannot essentially be weakened for uniform coverage with an optimal rate to hold [?].

Following [47], we construct adaptive honest credible sets by first defining a pivot centering estimator, and then determining a data-driven radius.

DEFINITION 4. (*The Median Tree*) Given a posterior $\Pi_{\mathbb{T}}[\cdot | X]$ over trees, we define the median tree $\mathcal{T}_X^* = \mathcal{T}^*(\Pi_{\mathbb{T}}[\cdot | X])$ as the set of nodes

$$(28) \quad \mathcal{T}_X^* = \{(l, k), l \leq L_{max}, \Pi[(l, k) \in \mathcal{T}_{int} | X] \geq 1/2\}.$$

Similarly as in the median probability model [4, 3], a node belongs to \mathcal{T}_X^* if its (marginal) posterior probability to be selected by a tree estimator exceeds 1/2. Interestingly, as the terminology suggests, \mathcal{T}_X^* is an *actual tree*, i.e. the nodes follow hereditary constraints (see Lemma 13 in the Supplement). We define the resulting median tree estimator as

$$(29) \quad \hat{f}_T(x) = \sum_{(l,k) \in \mathcal{T}_X^*} X_{lk} \psi_{lk}(x).$$

Moreover, we define a *radius*, for some $v_n \rightarrow \infty$ to be chosen, as

$$(30) \quad \sigma_n = \sigma_n(X) = \sup_{x \in [0,1]} \sum_{l=0}^{L_{max}} v_n \sqrt{\frac{\log n}{n}} \sum_{k=0}^{2^l-1} \mathbb{I}_{(l,k) \in \mathcal{T}_X^*} |\psi_{lk}(x)|.$$

A credible band with a radius $\sigma_n(X)$ as in (30) and a center \hat{f}_T as in (29) is

$$(31) \quad \mathcal{C}_n = \left\{ f : \|f - \hat{f}_T\|_\infty \leq \sigma_n(X) \right\}.$$

Theorem 2, proved in Section 10.3, shows that valid frequentist uncertainty quantification with Bayesian CART is attainable (up to log factors). Indeed, the confidence band (31) has a near-optimal diameter and a uniform frequentist coverage under self-similarity.

THEOREM 2. Let $0 < \alpha_1 \leq \alpha_2 \leq 1$, $M \geq 1$ and $\varepsilon > 0$. Let Π be any prior as in the statement of Theorem 1. Let σ_n be as in (30) with v_n such that $(\log n)^{1/2} = o(v_n)$ and let \hat{f}_T denote the median tree estimator (29). Then for \mathcal{C}_n defined in (31), uniformly over $\alpha \in [\alpha_1, \alpha_2]$, as $n \rightarrow \infty$,

$$\inf_{f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)} P_{f_0}(f_0 \in \mathcal{C}_n) \rightarrow 1.$$

For every $\alpha \in [\alpha_1, \alpha_2]$ and uniformly over $f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)$, the diameter $|\mathcal{C}_n|_\infty = \sup_{f, g \in \mathcal{C}_n} \|f - g\|_\infty$ and the credibility of the band verify, as $n \rightarrow \infty$,

$$(32) \quad |\mathcal{C}_n|_\infty = O_{P_{f_0}}((n/\log n)^{-\alpha/(2\alpha+1)} v_n),$$

$$(33) \quad \Pi[\mathcal{C}_n | X] = 1 + o_{P_{f_0}}(1).$$

Similarly as for Theorem 1, the results of Theorem 2 carry over to wavelet priors over a smooth wavelet basis, leading to the construction of confidence sets with arbitrary regularities $0 < \alpha_1 \leq \alpha_2 < \infty$. The undersmoothing factor v_n is commonplace in the context of confidence bands, with the condition $v_n \gg (\log n)^{1/2}$ reflecting the slight logarithmic price to pay for trees noted earlier in terms of L^∞ -estimation accuracy. In the previous statement both confidence and credibility of \mathcal{C}_n tend to 1. It is possible to achieve exact coverage by

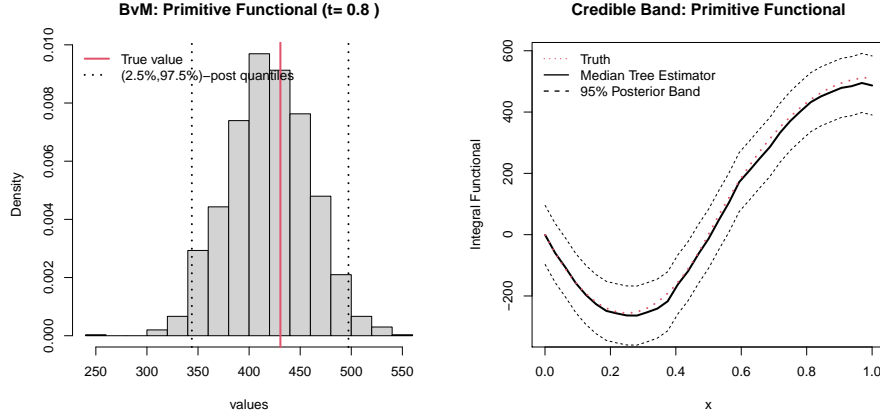


FIG 4. (Left) 0.95% credible interval for the primitive functional $F(x)$ with $x = 0.8$; (Right) the confidence band (34) obtained for $f_0(x) = (4x - 1)\mathbb{I}(x \leq 1/2) + (-2x + 2)\mathbb{I}(x > 1/2)$.

intersecting \mathcal{C}_n further with another ball. A natural way to do so (from the ‘estimating many functionals’ perspective, see [16]) is to intersect with a multiscale ball (we refer to Section 7.3 in the Supplement for details). For stability reasons, this intersection-band seems also preferable in practice and we present in Figure 6 on the right an illustration of coverage of such a band in nonparametric regression.

3.3. Inference for Functionals of f_0 : Bernstein-von Mises Theorems. By slightly modifying the Bayesian CART prior on the coarsest scales, it is possible to obtain asymptotic normality results, in the form of Bernstein-von Mises theorems, that imply that posterior quantile-credible sets are optimal-size confidence sets. In the next result, β_S denotes the bounded-Lipschitz metric on the metric space S (see also the Supplement Section 7.3).

THEOREM 3. Assume the Bayesian CART priors $\Pi_{\mathbb{T}}$ from Theorem 1 constrained to trees that fit $j_0(n)$ layers, i.e. $\gamma_{lk} = 1$ for $l \leq j_0(n)$, for $j_0(n) \asymp \sqrt{\log n}$.

1. BvM for smooth functionals $\psi_b(f) := \langle f, b \rangle$. Let $b \in L^\infty[0, 1]$ with coefficients $(b_{lk} = \langle b, \psi_{lk} \rangle)$. Assume $\sum_k |b_{lk}| \leq c_l$ for all $l \geq 1$ with $\sum_l l^2 c_l < \infty$. Then, in P_{f_0} -probability,

$$\beta_{\mathbb{R}} \left(\mathcal{L}(\sqrt{n}(\psi_b(f) - \hat{\psi}_b) | X), \mathcal{L}(\mathcal{N}(0, \|b\|_2^2)) \right) \rightarrow 0.$$

2. Functional BvM for the primitive $F(\cdot) = \int_0^\cdot f$. Let $(G(t) : t \in [0, 1])$ be a Brownian motion. Then, in P_{f_0} -probability,

$$\beta_{C([0,1])} \left(\mathcal{L} \left(\sqrt{n} \left(F(\cdot) - \int_0^\cdot dX^{(n)} \right) | X \right), \mathcal{L}(G) \right) \rightarrow 0$$

As a consequence of this result, quantile credible sets for the considered functionals are optimal confidence sets. For $\alpha \in (0, 1)$, let $q_{\alpha/2}^{\psi_b}(X)$ and $q_{1-\alpha/2}^{\psi_b}(X)$ be the $\alpha/2$ and $1 - \alpha/2$ quantiles of the induced posterior distribution on the functional $\psi_b = \int_0^1 f(u)b(u)du$ and set $I_b(X) := [q_{\alpha/2}^{\psi_b}(X), q_{1-\alpha/2}^{\psi_b}(X)]$. Theorem 3 (part 1) then implies (see [16] for a proof) that

$$P_{f_0} [\psi_b(f_0) \in I_b(X)] \rightarrow 1 - \alpha.$$

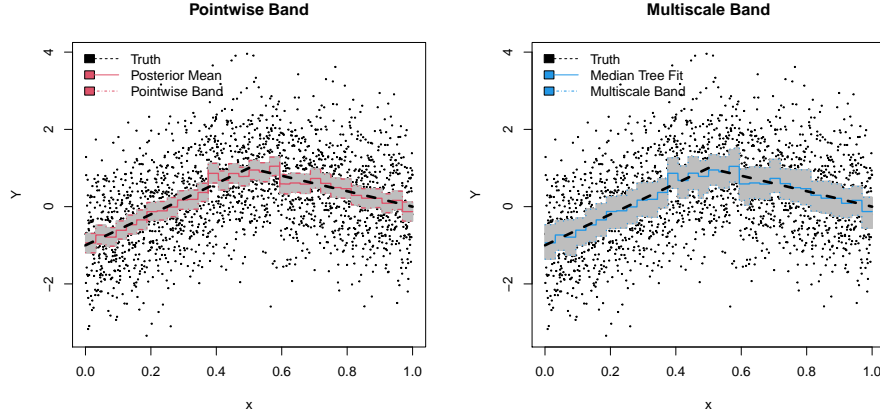


FIG 5. (Left) Pointwise 0.95% credible intervals. (Right) Multiscale 0.95% credible band (75) (Supplement, Section 7.3.5) obtained with $v_n = 1$ and $w_l = l^{1/2+0.01}$. The true function is $f_0(x) = (4x - 1)\mathbb{I}(x \leq 1/2) + (-2x + 2)\mathbb{I}(x > 1/2)$.

Similarly, let $R_n(X)$ be the data-dependent radius chosen from the induced posterior distribution on $F(\cdot) = \int_0^\cdot f$ as follows, for $\hat{F}(\cdot) = \int_0^\cdot dX^{(n)}$,

$$(34) \quad \Pi[\|F - \hat{F}\|_\infty \leq R_n(X) | X] = 1 - \alpha.$$

Consider the band $\mathcal{C}^F(X) := \{F : \|F - \hat{F}\|_\infty \leq R_n(X)\}$. Then Theorem 3 (part 2) implies (see [16], Corollary 2 for a related statement and proof), for $F_0(\cdot) = \int_0^\cdot f_0$,

$$P_{f_0}[F_0 \in \mathcal{C}^F(X)] \rightarrow 1 - \alpha.$$

In other words, the band (34) has exact asymptotic coverage. It can also be checked that it is optimal efficient in semiparametric terms (that is, its width is optimal asymptotically). We derive Theorem 3 as a consequence of an adaptive nonparametric BvM (Theorem 7 in the Supplement; see Section 10.4 for a proof, where other possible choices for $j_0(n)$ are discussed), only obtained so far for *adaptive* priors in the work of Ray [47], which considered (conjugate) spike and slab priors. Derivation of the band (34) in practice is easily obtained once posterior samples are available. Theorem 3 is illustrated, in the regression framework studied in Section 7.1, on a numerical example with a piece-wise linear regression function (details on the implementation are in Section 7.2) in Figure ?? . The left panel presents a histogram of posterior samples (together with 2.5% and 97.5% quantiles) of the primitive functional $F(x) = \sum_{t_i \leq x} f(t_i)$ for $x = 0.8$ with true value is marked with a red solid line. The right panel portrays the confidence band (34) which uniformly captures the true functional (dotted line).

3.4. Nonparametric Regression: Overview of Results. Our results obtained under the white noise model can be transported to the more practical nonparametric regression model. While these two models are asymptotically equivalent [?] (under uniform smoothness assumptions satisfied, e.g., by α -Hölderian functions with $\alpha > 1/2$), it is not automatic that the knowledge of a (wavelet shrinkage/non-linear) minimax procedure in one model implies the optimality in the other. It turns out, however, that our results *can* be carried over to fixed-design regression without necessarily assuming $\alpha > 1/2$. We assume outcomes $Y = (Y_1, \dots, Y_n)'$ arising from

$$(35) \quad Y_i = f_0(t_i) + \varepsilon_i, \quad \varepsilon_i \stackrel{iid}{\sim} \mathcal{N}(0, 1), \quad i = 1, \dots, n = 2^{L_{max}+1}$$

where f_0 is an unknown regression function and $\{t_i \in [0, 1] : 1 \leq i \leq n\}$ are fixed design points. For simplicity, we consider a regular grid, i.e. $t_i = i/n$ for $1 \leq i \leq n$ and assume n is a power of 2. In Section 7.1, we show that most results for Bayesian CART obtained earlier in white noise carry over to the model (61) with a few minor changes. One minor modification concerns the loss function. We mainly consider the ‘canonical’ supremum-norm loss for the fixed design setting, that is, the ‘max-norm’ defined for given functions f, g by

$$\|f - g\|_{\infty, n} = \max_{1 \leq i \leq n} |f(t_i) - g(t_i)|,$$

but it is also possible to consider the whole supremum-norm loss $\|\cdot\|_{\infty}$. We postpone statements and proofs to the Supplement, Sections 7.1 and 12.1. In a numerical study (Section 7.2), we illustrate that the implementation of Bayesian CART [20, 23] and the construction of our confidence bands is rather straightforward. For example, Figure 7.2 shows how inference can be carried out with Bayesian CART posteriors in non-parametric regression with a piece-wise linear regression function using the intersecting band construction (detailed in Section 7.3.5. Contrary to point-wise credible intervals (on the left) that are easy to produce but do not cover, our multiscale confidence band (on the right) uniformly captures the true regression function. More details on this example are presented in Section 7.2.

4. Non-dyadic Bayesian CART. A limitation of midpoint splits in dyadic trees is that they treat the basis as fixed, allowing the jumps to occur *only* at pre-specified dyadic locations even when not justified by data. General CART regression methodology [9, 28] avoids this restriction by treating the basis as *unknown*, where the partitioning cells shrink and stretch with data. In this section, we leave behind ‘static’ dyadic trees to focus on the analysis of Bayesian (non-dyadic) CART [20, 23] and its connection to Unbalanced Haar (UH) wavelet basis selection.

4.1. Unbalanced Haar Wavelets. UH wavelet basis functions [33] are *not* necessarily translates/dilates of any mother wavelet function and, as such, allow for different support lengths and design-adapted split locations. Here, we particularize the constructive definition of UH wavelets given by [27]. Assume that possible values for splits are chosen from a set of $n = 2^{L_{max}}$ breakpoints $\mathcal{X} = \{x_i : x_i = i/n, 1 \leq i \leq n\}$. Using the scale/location index enumeration, pairs (l, k) in the tree are now equipped with (a) a *breakpoint* $b_{lk} \in \mathcal{X}$ and (b) *left and right brackets* $(l_{lk}, r_{lk}) \in \mathcal{X} \cup \{0, 1\}$. Unlike balanced Haar wavelets (3), where $b_{lk} = (2k + 1)/2^{l+1}$, the breakpoints b_{lk} are *not required* to be regularly dyadically constrained and are chosen from \mathcal{X} in a hierarchical fashion as follows. One starts by setting $l_{00} = 0, r_{00} = 1$. Then

(a) The first breakpoint b_{00} is selected from $\mathcal{X} \cap (0, 1)$.

(b) For each $1 \leq l \leq L_{max}$ and $0 \leq k < 2^l$, set

$$(36) \quad \begin{aligned} l_{lk} &= l_{(l-1)\lfloor k/2 \rfloor}, & r_{lk} &= b_{(l-1)\lfloor k/2 \rfloor}, & \text{if } k \text{ is even,} \\ l_{lk} &= b_{(l-1)\lfloor k/2 \rfloor}, & r_{lk} &= r_{(l-1)\lfloor k/2 \rfloor}, & \text{if } k \text{ is odd.} \end{aligned}$$

If $\mathcal{X} \cap (l_{lk}, r_{lk}] \neq \emptyset$, choose b_{lk} from $\mathcal{X} \cap (l_{lk}, r_{lk}]$.

Let A denote the set of *admissible* nodes (l, k) , in that (l, k) is such that $\mathcal{X} \cap (l_{lk}, r_{lk}] \neq \emptyset$, obtained through an instance of the sampling process described above and let

$$B = (b_{lk})_{(l,k) \in A}$$

be the corresponding set of breakpoints. Each collection of split locations B gives rise to nested intervals

$$L_{lk} = (l_{lk}, b_{lk}] \quad \text{and} \quad R_{lk} = (b_{lk}, r_{lk}].$$

Starting with the mother wavelet $\psi_{-10}^B = \psi_{-10} = \mathbb{I}_{(0,1)}$, one then recursively constructs wavelet functions ψ_{lk}^B with a support $I_{lk}^B = L_{lk} \cup R_{lk}$ as

$$(37) \quad \psi_{lk}^B(x) = \frac{1}{\sqrt{|L_{lk}|^{-1} + |R_{lk}|^{-1}}} \left(\frac{\mathbb{I}_{L_{lk}}(x)}{|L_{lk}|} - \frac{\mathbb{I}_{R_{lk}}(x)}{|R_{lk}|} \right).$$

By construction, the system $\Psi_A^B = \{\psi_{-10}^B, \psi_{lk}^B : (l, k) \in A\}$ is orthonormal in $L^2[0, 1]$. With UH wavelets, the decay of wavelet coefficients $\beta_{lk} = \langle f, \psi_{lk}^B \rangle$ for a α -Hölder function f verifies $|\beta_{lk}^B| \lesssim \max\{|L_{lk}|, |R_{lk}|\}^{\alpha+1/2}$, see Lemma 9. [27] points out that the computational complexity of the discrete UH transform could be unnecessarily large and imposes the balancing requirement $\max\{|L_{lk}|, |R_{lk}|\} \leq E(|L_{lk}| + |R_{lk}|) \forall (l, k) \in A$, for some $1/2 \leq E < 1$. Similarly, in order to control the combinatorial complexity of the basis system, we require that the UH wavelets are *weakly balanced* in the following sense.

DEFINITION 5. A system $\Psi_A^B = \{\psi_{-10}^B, \psi_{lk}^B : (l, k) \in A\}$ of UH wavelets is weakly balanced with balancing constants $E, D \in \mathbb{N}^*$ if, for any $(l, k) \in A$,

$$(38) \quad \max(|L_{lk}|, |R_{lk}|) = \frac{M_{lk}}{2^{l+D}} \text{ for some } M_{lk} \in \{1, \dots, E + l\}.$$

Examples of weakly balanced systems are given in Section 9. These include quantile-based splits, where the split locations are distributed along the quantiles of a given distribution. Note that in the actual BART implementation, the splits are chosen from sample quantiles to ensure balancedness (similar to our condition (38)). Quantile splits are a natural way to generate many weakly balanced systems, providing a much increased flexibility compared to dyadic splits, which correspond to uniform quantiles.

The *non-dyadic Bayesian CART* prior is then defined as follows:

- *Step 1. (Basis Generation)* Sample $B = (b_{lk})_{0 \leq k < 2^l - 1, l \leq L}$ from $\Pi_{\mathbb{B}}$ by following the steps a)–b) around (36) subject to satisfying the *balancing condition* (38).
- *Step 2. (Tree Generation)* Independently of B , sample a binary tree \mathcal{T} from one of the priors $\Pi_{\mathbb{T}}$ described in Section 2.1.
- *Step 3. (Step Heights Generation)* Given \mathcal{T} , we obtain the coefficients (β_{lk}^B) from the tree-shaped prior (19). Using the UH wavelets, the prior on the internal coefficients β_{lk}^B can be translated into a model on the histogram heights $\tilde{\beta}_{lk}^B$ through (8).

An example of such a prior is obtained by first randomly drawing quantiles (e.g. by drawing a density at random verifying conditions as in Example 3) to generate the breakpoints for Step 1 and then following the construction from Section 2 for Steps 2–3. The following theorem is proved in Section 11.

THEOREM 4. Let $\Pi_{\mathbb{B}}$ be any prior on breakpoint collections that satisfy weak balancedness according to Definition 5. Let $\Pi_{\mathbb{T}}$ be the Galton-Watson process prior from Section 2.1 with $p_{lk} = \Gamma^{-l^4}$. Consider the tree-shaped wavelet prior (19) with $\pi(\beta_{\mathcal{T}}) \sim \mathcal{N}(0, I_{|\mathcal{T}_{ext}|})$. Let $f_0 \in \mathcal{H}_M^\alpha$ as in (24) for some $M > 0$ and $0 < \alpha \leq 1$ and define

$$(39) \quad \varepsilon_n = (\log n)^{1+\frac{3}{2}} \left(\frac{\log n}{n} \right)^{\frac{\alpha}{2\alpha+1}}.$$

Then, there exist $\Gamma_0, c_0 > 0$ depending only on the constants E, D in the weak balancedness condition such that, for any $\Gamma \geq \Gamma_0$ and $c \geq c_0$, for any $M_n \rightarrow \infty$, we have, for $n \rightarrow \infty$

$$(40) \quad E_{f_0} \Pi[\ell_\infty(f_{\mathcal{T}, \beta}, f_0) \geq \|f_{\mathcal{T}, \beta} - f_0\|_\infty > M_n \varepsilon_n \mid X] \rightarrow 0.$$

In the context of piecewise constant priors, Theorem 4 allows further flexibility in the choice of the prior as compared to Theorem 1 in that the location of the breakpoints, on the top of their structure given by the tree prior, can vary in their location according to its own specific prior. Whether one can further weaken the balancing condition to still get optimal multiscale results is an interesting open question that goes beyond the scope of this paper. In addition, the log-factor in (39) could be further optimized, similarly as in Theorem 1.

5. Discussion. In this paper we explored connections between Bayesian tree-based regression methods and structured wavelet shrinkage. We demonstrated that Bayesian tree-based methods attain (almost) optimal convergence rates in the supremum norm and obtain limiting results for functionals, that follow from a non-parametric and adaptive Bernstein–von Mises theorem. The developed framework also allows us to construct adaptive credible bands around f_0 under self-similarity. To allow for non-dyadically organized splits, we introduced weakly balanced Haar wavelets (an elaboration on unbalanced Haar wavelets of [33]) and showed that Bayesian CART performs basis selection from this library and attains a near-minimax rate of posterior concentration under the sup-norm loss.

Although for clarity of exposition we focused on the white noise model, our results can be extended to the more practical regression model for fixed regular designs (Section 7.1 in the Supplement) or possibly more general designs under some conditions. We note that the techniques of proof are non-conjugate in their key tree aspect, which opens the door to applications in many other statistical settings. A version of Bayesian CART for density estimation following the ideas of the present work is currently investigated by T. Randrianarisoa as part of his PhD thesis. More precisely, using the present techniques, it is possible to develop multiscale rate results for Pólya trees with ‘optional stopping’ along a tree, in the spirit of [59]. Our confidence set construction can be also shown to have local adaptation properties. The ability of Bayesian CART to spatially adapt in this way will be investigated in a followup work. Further natural extensions include high-dimensional versions of the model, extending the multi-dimensional version briefly presented here, as well as forest priors. These will be considered elsewhere.

6. Proof of Theorem 1. The proof proceeds in three steps. In Section 6.1 we first show that the posterior concentrates on not too deep trees. In Section 6.2, we then show that the posterior probability of missing signal vanishes and, finally, in Section 6.3 we show that the posterior distribution concentrates around signals. To better convey main ideas, we present the proof for the independent prior $\beta_{\mathcal{T}} \sim \mathcal{N}(0, \Sigma_{\mathcal{T}})$ with $\Sigma_{\mathcal{T}} = I_K$ for $K = |\mathcal{T}_{ext}|$ and the Galton-Watson (GW) tree prior from Section 2.1.1 with a split probability p_l . The proof for the g -prior $\Sigma_{\mathcal{T}} = g_n(A'_{\mathcal{T}}A_{\mathcal{T}})^{-1}$ is more technically involved and is presented in Section 10.1 in the Supplement.

We will be working conditionally on the event

$$(41) \quad \mathcal{A} = \left\{ \max_{-1 \leq l \leq L, 0 \leq k < 2^l} \varepsilon_{lk}^2 \leq 2 \log(2^{L+1}) \right\},$$

where $L = L_{max} = \lfloor \log_2 n \rfloor$. Since $\varepsilon_{lk} \sim \mathcal{N}(0, 1)$, this event has a large probability in the sense that $P(\mathcal{A}^c) \lesssim (\log n)^{-1}$, which follows from $P\left[\max_{1 \leq i \leq N} |Z_i| > \sqrt{2 \log N}\right] \leq c_0 / \sqrt{\log N}$ for some $c_0 > 0$ when $Z_i \sim \mathcal{N}(0, 1)$ for $1 \leq i \leq N$.

6.1. Posterior Probability of Deep Trees. The first step is to show that, on the event \mathcal{A} , the posterior concentrates on reasonably small trees, i.e. trees whose depth $d(\mathcal{T})$ is no larger

than an ‘optimal’ depth which depends on the unknown smoothness α . Let us define such a depth $\mathcal{L}_c = \mathcal{L}_c(\alpha, M)$ as

$$(42) \quad \mathcal{L}_c = \left\lceil \log_2 \left((8M)^{\frac{1}{\alpha+1/2}} \left(\frac{n}{\log n} \right)^{\frac{1}{2\alpha+1}} \right) \right\rceil.$$

LEMMA 1. *Under the assumptions of Theorem 1, on the event \mathcal{A} ,*

$$(43) \quad \Pi[d(\mathcal{T}) > \mathcal{L}_c \mid X] \rightarrow 0 \quad (n \rightarrow \infty).$$

PROOF. Consider one tree $\mathcal{T} \in \mathbb{T}$ such that $d(\mathcal{T}) \geq 1$ and denote with \mathcal{T}^- a pruned subtree obtained from \mathcal{T} by turning its deepest rightmost internal node, say (l_1, k_1) , into a terminal node. Then $\mathcal{T}^- = \mathcal{T}_{int}^- \cup \mathcal{T}_{ext}^-$, where

$$\mathcal{T}_{int}^- = \mathcal{T}_{int} \setminus \{(l_1, k_1)\}, \quad \mathcal{T}_{ext}^- = \mathcal{T}_{ext} \setminus \{(l_1 + 1, 2k_1), (l_1 + 1, 2k_1 + 1)\} \cup \{(l_1, k_1)\}.$$

Note that \mathcal{T}^- is a full binary tree and that the mapping $\mathcal{T} \rightarrow \mathcal{T}^-$ is not necessarily injective. Indeed, there are up to $2^{d(\mathcal{T}^-)}$ trees \mathcal{T} that give rise to the same pruned tree \mathcal{T}^- . Let $\mathbb{T}_d = \{\mathcal{T} \in \mathbb{T} : d(\mathcal{T}) = d\}$ denote the set of all full binary trees of depth *exactly* $d \geq 1$. Then, using the notation (22),

$$(44) \quad \Pi[\mathbb{T}_d \mid X] = \frac{\sum_{\mathcal{T} \in \mathbb{T}_d} W_X(\mathcal{T})}{\sum_{\mathcal{T} \in \mathbb{T}} W_X(\mathcal{T})} = \frac{\sum_{\mathcal{T} \in \mathbb{T}_d} \frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^-)} W_X(\mathcal{T}^-)}{\sum_{\mathcal{T} \in \mathbb{T}} W_X(\mathcal{T})},$$

$$\text{where} \quad \frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^-)} = \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^-)} \frac{\int \prod_{(l,k) \in \mathcal{T}'_{int}} e^{nX_{lk}\beta_{lk} - n\beta_{lk}^2/2} d\pi(\beta_{\mathcal{T}})}{\int \prod_{(l,k) \in \mathcal{T}'_{int}} e^{nX_{lk}\beta_{lk} - n\beta_{lk}^2/2} d\pi(\beta_{\mathcal{T}^-})}.$$

Let $\mathbf{X}_{\mathcal{T}} = (X_{lk} : (l, k) \in \mathcal{T}'_{int})'$ and $\beta_{\mathcal{T}} = (\beta_{lk} : (l, k) \in \mathcal{T}'_{int})'$ be the top-down left-to-right ordered sequences (recall that we order nodes according to the index $2^l + k$). Assuming $\beta_{\mathcal{T}} \sim \mathcal{N}(0, \Sigma_{\mathcal{T}})$, and denoting $K = |\mathcal{T}_{ext}| = |\mathcal{T}_{int}| + 1$,

$$(45) \quad \begin{aligned} \frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^-)} &= \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^-)} \sqrt{\frac{|\Sigma_{\mathcal{T}^-}|}{2\pi|\Sigma_{\mathcal{T}}|}} \frac{\int e^{n\mathbf{X}'_{\mathcal{T}}\beta_{\mathcal{T}} - \beta'_{\mathcal{T}}[nI_K + \Sigma_{\mathcal{T}}^{-1}]\beta_{\mathcal{T}}/2} d\beta_{\mathcal{T}}}{\int e^{n\mathbf{X}'_{\mathcal{T}^-}\beta_{\mathcal{T}^-} - \beta'_{\mathcal{T}^-}[nI_{K-1} + \Sigma_{\mathcal{T}^-}^{-1}]\beta_{\mathcal{T}^-}/2} d\beta_{\mathcal{T}^-}} \\ &= \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^-)} \sqrt{\frac{|\Sigma_{\mathcal{T}^-}|}{|\Sigma_{\mathcal{T}}|}} \sqrt{\frac{|nI_{K-1} + \Sigma_{\mathcal{T}^-}^{-1}|}{|nI_K + \Sigma_{\mathcal{T}}^{-1}|}} \frac{e^{n^2\mathbf{X}'_{\mathcal{T}}(nI_K + \Sigma_{\mathcal{T}}^{-1})^{-1}\mathbf{X}_{\mathcal{T}}/2}}{e^{n^2\mathbf{X}'_{\mathcal{T}^-}(nI_{K-1} + \Sigma_{\mathcal{T}^-}^{-1})^{-1}\mathbf{X}_{\mathcal{T}^-}/2}}. \end{aligned}$$

Since $X_{l_1 k_1}$ corresponds to the node (l, k) with the highest index $2^l + k$, one can write $\mathbf{X}_{\mathcal{T}} = (\mathbf{X}_{\mathcal{T}^-}, X_{l_1 k_1})'$.

We focus on the GW prior from Section 2.1.1 and on the independent prior $\Sigma_{\mathcal{T}} = I_K$ and present proofs for the remaining priors in Section 10.1. Using the expression (45) and since (l_1, k_1) is the deepest rightmost internal node in \mathcal{T} , and \mathcal{T} is of depth $d = d(\mathcal{T}) = l_1 + 1$, using the definition of the GW prior,

$$\frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^-)} = \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^-)} \prod_{(l,k) \in \mathcal{T}'_{int} \setminus \mathcal{T}'_{int}^-} \frac{e^{\frac{n^2}{2(n+1)} X_{lk}^2}}{\sqrt{n+1}} = \frac{p_{d-1}(1-p_d)^2}{1-p_{d-1}} \frac{e^{\frac{n^2}{2(n+1)} X_{l_1 k_1}^2}}{\sqrt{n+1}}.$$

Suppose \mathcal{T} has depth $d(\mathcal{T}) > \mathcal{L}_c$. Then $l_1 \geq \mathcal{L}_c$ and from the Hölder continuity (23), one gets $8|\beta_{l_1 k_1}| \leq \sqrt{\log n/n}$, where \mathcal{L}_c is as in (42). Then, conditionally on the event (41),

$$(46) \quad |X_{l_1 k_1}| \leq \frac{1}{\sqrt{n}} \left[\frac{1}{8} \sqrt{\log n} + \sqrt{2 \log n + \log 4} \right]$$

and thereby $2X_{l_1 k_1}^2 \leq 5 \log n/n$. Recall that, under the GW-prior, the split probability is $p_d = \Gamma^{-d}$. As $\Gamma > 2$, one has $p_d < 1/2$ and so, for any $d > \mathcal{L}_c$,

$$\frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^-)} \leq 2p_{d-1} \exp\left(\frac{5n \log n}{4(n+1)} - \frac{1}{2} \log(1+n)\right) < 2n^{3/4} p_{d-1}.$$

Going back to the ratio (44), we now bound, with $a(n, d) =: 2n^{3/4} p_{d-1}$,

$$\frac{\Pi[\mathbb{T}_d | X]}{a(n, d)} \leq \frac{\sum_{\mathcal{T} \in \mathbb{T}_d} W_X(\mathcal{T}^-)}{\sum_{\mathcal{T} \in \mathbb{T}} W_X(\mathcal{T})} \leq \frac{\sum_{\mathcal{T} \in \mathbb{T}_d^-} 2^{d(\mathcal{T}^-)} W_X(\mathcal{T})}{\sum_{\mathcal{T} \in \mathbb{T}} W_X(\mathcal{T})} \leq 2^d,$$

where \mathbb{T}_d^- is the image of \mathbb{T}_d under the map $\mathcal{T} \rightarrow \mathcal{T}^-$, and using that at most $2^{d(\mathcal{T}^-)}$ trees are mapped to the same \mathcal{T}^- . Using this bound one deduces that, on the event \mathcal{A} , with $L = L_{max} = \log_2 n$,

$$\begin{aligned} \Pi[d(\mathcal{T}) > \mathcal{L}_c | X] &= \sum_{d=\mathcal{L}_c+1}^L \Pi[\mathbb{T}_d | X] \leq 4n^{3/4} \sum_{d=\mathcal{L}_c+1}^L 2^{d-1} p_{d-1} \\ &< 4n^{3/4} L \exp[-\mathcal{L}_c \log(\Gamma/2)]. \end{aligned}$$

As $\mathcal{L}_c \asymp (\log n)/(1+2\alpha)$, the right hand side goes to zero as soon as, e.g. $\log(\Gamma/2) > 7(1+2\alpha)/8$ that is, for $\alpha \leq 1$, $\Gamma > 2e^3$. \square

6.2. Posterior Probability of Missing Signal. The next step is showing that the posterior probability of missing a node with large enough signal vanishes.

LEMMA 2. *Let us denote, for $A > 0$ to be chosen suitably large,*

$$(47) \quad S(f_0; A) = \left\{ (l, k) : |\beta_{lk}^0| \geq A \frac{\log n}{\sqrt{n}} \right\}.$$

Under the assumptions of Theorem 1, on the event \mathcal{A} from (41),

$$(48) \quad \Pi[\{\mathcal{T} : S(f_0; A) \not\subseteq \mathcal{T}\} | X] \rightarrow 0 \quad (n \rightarrow \infty).$$

PROOF. As before, we present the proof with the GW prior from Section 2.1.1 and for the independent prior with $\Sigma_{\mathcal{T}} = I_K$, referring to Section 10.1 for the g -prior. Let us first consider a given node $(l_S, k_S) \in S(f_0; A)$, for A to be specified below, and note that the Hölder condition on f_0 implies $l_S \leq \mathcal{L}_c$ (for n large enough). Let $\mathbb{T}_{\setminus(l_S, k_S)} = \{\mathcal{T} \in \mathbb{T} : (l_S, k_S) \notin \mathcal{T}_{int}\}$ denote the set of trees that miss the signal node in the sense that they *do not have a cut* at (l_S, k_S) . For any such tree $\mathcal{T} \in \mathbb{T}_{\setminus(l_S, k_S)}$ we then denote by \mathcal{T}^+ the smallest full binary tree (in terms of the number of nodes) that contains \mathcal{T} and that splits on (l_S, k_S) . Such a tree can be constructed from $\mathcal{T} \in \mathbb{T}_{\setminus(l_S, k_S)}$ as follows. Denote by $(l_0, k_0) \in \mathcal{T}_{ext} \cap [(0, 0) \leftrightarrow (l_S, k_S)]$ the external node of \mathcal{T} which is *closest* to (l_S, k_S) on the route from the root to (l_S, k_S) in a flat tree (denoted by $[(0, 0) \leftrightarrow (l_S, k_S)]$). Next, denote by \mathcal{T}^+ the extended tree obtained from \mathcal{T} by sequentially splitting all $(l, k) \in [(l_0, k_0) \leftrightarrow (l_S, k_S)]$. Similarly as for $\mathcal{T} \rightarrow \mathcal{T}^-$ above, the map $\mathcal{T} \rightarrow \mathcal{T}^+$ is not injective and we denote by $\mathbb{T}_{(l_S, k_S)}$ the set of all extended trees \mathcal{T}^+ obtained from some $\mathcal{T} \in \mathbb{T}_{\setminus(l_S, k_S)}$. Now, the posterior probability $\Pi[\mathbb{T}_{\setminus(l_S, k_S)} | X]$ of missing the signal node (l_S, k_S) equals

$$(49) \quad \frac{\sum_{\mathcal{T} \in \mathbb{T}_{\setminus(l_S, k_S)}} W_X(\mathcal{T})}{\sum_{\mathcal{T} \in \mathbb{T}} W_X(\mathcal{T})} \leq \frac{\sum_{\mathcal{T} \in \mathbb{T}_{\setminus(l_S, k_S)}} \frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^+)} W_X(\mathcal{T}^+)}{\sum_{\mathcal{T} \in \mathbb{T}_{(l_S, k_S)}} W_X(\mathcal{T})}.$$

Let us denote by $\mathcal{T}^{(j)}$ for $j = -1, \dots, s$ the sequence of nested trees obtained by extending one branch of \mathcal{T} towards (l_S, k_S) by splitting the nodes $[(l_0, k_0) \leftrightarrow (l_S, k_S)]$, where $\mathcal{T}^+ = \mathcal{T}^{(s)}$ and $\mathcal{T} = \mathcal{T}^{(-1)}$. Then

$$(50) \quad \frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^+)} = \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^+)} \prod_{j=0}^s \frac{N_X(\mathcal{T}^{(j-1)})}{N_X(\mathcal{T}^{(j)})}.$$

Under the GW process prior with $p_l = \Gamma^{-l}$ for some $\Gamma > 2$, the ratio of prior tree probabilities in the last expression satisfies

$$(51) \quad \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^+)} = \frac{1 - p_{l_0}}{p_{l_0}} \times \left(\prod_{l=l_0+1}^{l_S} \frac{1}{p_l(1 - p_l)} \right) \times \frac{1}{(1 - p_{l_S+1})^2}.$$

The first term is due to the fact that \mathcal{T}^+ splits the node (l_0, k_0) while \mathcal{T} does not. The second term in the denominator is the extra prior probability of \mathcal{T}^+ over \mathcal{T} that is due to the branch reaching out to (l_S, k_S) . Along this branch (note that this is the smallest possible branch), one splits *only* one daughter node for each layer l (thereby the term p_l) and not the other (thereby the term $1 - p_l$). The third term above is due to the fact that the two daughters of (l_S, k_S) are not split. The quantity (51) is bounded by $2^{l_S - l_0 + 2} \Gamma^{(l_0 + l_S)(l_S - l_0 + 1)/2} < 4\Gamma^{2l_S^2}$.

Assuming $\Sigma_{\mathcal{T}} = I_K$, we can write for any \mathcal{T} in $\mathbb{T}_{\setminus(l_S, k_S)}$

$$(52) \quad \frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^+)} = \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^+)} \prod_{(l,k) \in \mathcal{T}^+ \setminus \mathcal{T}} \frac{\sqrt{n+1}}{e^{\frac{n^2}{2(n+1)} X_{lk}^2}}.$$

Using the definition of the model and the inequality $2ab \geq -a^2/2 - 2b^2$ for $a, b \in \mathbb{R}$, we obtain $X_{l_S k_S}^2 \geq (\beta_{l_S k_S}^0)^2/2 - \varepsilon_{l_S k_S}^2/n$. On the event \mathcal{A} , one gets

$$\exp \left\{ -\frac{n^2}{2(n+1)} X_{l_S k_S}^2 \right\} \leq \exp \left\{ -\frac{n^2(\beta_{l_S k_S}^0)^2}{4(n+1)} + \frac{n(\log 2)(\log_2 n + 1)}{n+1} \right\}.$$

The term in (52) can be thus bounded, for any $\mathcal{T} \in \mathbb{T}_{\setminus(l_S, k_S)}$, by

$$\frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^+)} \leq C\Gamma^{2l_S^2} \exp \left\{ \frac{3(l_S - l_0 + 1)(\log_2 n + 1)}{2} - \frac{nA^2 \log^2 n}{4(n+1)} \right\} =: b(n, l_S).$$

We now continue to bound the ratio (49). For each given \mathcal{T}^+ , there are *at most* l_S trees $\tilde{\mathcal{T}} \in \mathbb{T}_{\setminus(l_S, k_S)}$ which have the same extended tree $\tilde{\mathcal{T}}^+ = \mathcal{T}^+$. This is because \mathcal{T}^+ is obtained by extending one given branch by adding no more than l_S nodes. Using this fact, (49), and the definition of $b(n, l_S)$ on the last display,

$$\frac{\Pi[\mathbb{T}_{\setminus(l_S, k_S)} | X]}{b(n, l_S)} \leq \frac{\sum_{\mathcal{T} \in \mathbb{T}_{\setminus(l_S, k_S)}} W_X(\mathcal{T}^+)}{\sum_{\mathcal{T} \in \mathbb{T}_{\setminus(l_S, k_S)}} W_X(\mathcal{T})} \leq l_S \frac{\sum_{\mathcal{T} \in \mathbb{T}_{\setminus(l_S, k_S)}} W_X(\mathcal{T})}{\sum_{\mathcal{T} \in \mathbb{T}_{\setminus(l_S, k_S)}} W_X(\mathcal{T})}.$$

By choosing $A = A(\Gamma) > 0$ large enough, this leads to

$$\Pi[\mathbb{T}_{\setminus(l_S, k_S)} | X] \lesssim e^{(3/2 + 3 \log \Gamma)(\log_2 n + 1)^2 - \frac{A^2}{8} \log^2 n} \lesssim e^{-\frac{A^2}{16} \log^2 n}.$$

Then the result follows as, on the event \mathcal{A} ,

$$\sum_{(l_S, k_S) \in S(f_0, A)} \Pi[\mathbb{T}_{\setminus(l_S, k_S)} | X] \lesssim 2^{\mathcal{L}_c + 1} e^{-\frac{A^2}{16} \log^2 n} \lesssim e^{-\frac{A^2}{32} \log^2 n} \rightarrow 0. \quad \square$$

6.3. *Posterior Concentration Around Signals.* Let us now show that the posterior does not distort large signals too much.

LEMMA 3. Let us denote, for \mathcal{L}_c as in (42) and $S(f_0; A)$ as in (47),

$$(53) \quad \mathbb{T} = \{\mathcal{T} : d(\mathcal{T}) \leq \mathcal{L}_c, S(f_0; A) \subset \mathcal{T}\}.$$

Then, on the event \mathcal{A} , for some $C' > 0$, uniformly over $\mathcal{T} \in \mathbb{T}$,

$$(54) \quad \int \max_{(l,k) \in \mathcal{T}'_{int}} |\beta_{lk} - \beta_{lk}^0| d\Pi[\beta_{\mathcal{T}} | \mathbf{X}_{\mathcal{T}}] < C' \sqrt{\frac{\log n}{n}},$$

with $\mathbf{X}_{\mathcal{T}} = (X_{lk} : (l, k) \in \mathcal{T}'_{int})'$ the ordered vector of active responses.

PROOF. For a given tree \mathcal{T} with $K = |\mathcal{T}_{ext}|$ leaves, we denote by $\beta_{\mathcal{T}} = (\beta_{lk} : (l, k) \in \mathcal{T}'_{int})'$ the vector of wavelet (internal node) coefficients, with $\mathbf{X}_{\mathcal{T}}$ the corresponding responses and with $\varepsilon_{\mathcal{T}}$ the white noise disturbances. It follows from (21) that, given $\mathbf{X}_{\mathcal{T}}$ (so for fixed ε_{lk}) and \mathcal{T} , the vector $\beta_{\mathcal{T}}$ has a Gaussian distribution $\beta_{\mathcal{T}} | \mathbf{X}_{\mathcal{T}} \sim \mathcal{N}(\mu_{\mathcal{T}}, \tilde{\Sigma}_{\mathcal{T}})$, where $\tilde{\Sigma}_{\mathcal{T}} = (nI_K + \Sigma_{\mathcal{T}}^{-1})^{-1}$ and $\mu_{\mathcal{T}} = n\tilde{\Sigma}_{\mathcal{T}} \left(\beta_{\mathcal{T}}^0 + \frac{1}{\sqrt{n}} \varepsilon_{\mathcal{T}} \right)$. Next, using Lemma 7, we have

$$(55) \quad \mathbb{E} [\|\beta_{\mathcal{T}} - \beta_{\mathcal{T}}^0\|_{\infty} | \mathbf{X}_{\mathcal{T}}] \leq \|\mu_{\mathcal{T}} - \beta_{\mathcal{T}}^0\|_{\infty} + \sqrt{2\bar{\sigma}^2 \log K} + 2\sqrt{2\pi\bar{\sigma}^2},$$

where $\bar{\sigma}^2 = \max \text{diag}(\tilde{\Sigma}_{\mathcal{T}})$. Focusing on the first term, we can write

$$(56) \quad \|\mu_{\mathcal{T}} - \beta_{\mathcal{T}}^0\|_{\infty} \leq \sqrt{n} \|\tilde{\Sigma}_{\mathcal{T}} \varepsilon_{\mathcal{T}}\|_{\infty} + \|(n\tilde{\Sigma}_{\mathcal{T}} - I_K) \beta_{\mathcal{T}}^0\|_{\infty}.$$

Using the fact $(I + B)^{-1} = I - (I + B^{-1})^{-1}$, we obtain $n\tilde{\Sigma}_{\mathcal{T}} - I_K = -(I_K + n\Sigma_{\mathcal{T}})^{-1}$. From now on, we focus on the simpler case $\Sigma_{\mathcal{T}} = I_K$ and refer to Section 10.1.3 (Supplement) for the proof for the g -prior. With $\Sigma_{\mathcal{T}} = I_K$ we can write $\|(n\tilde{\Sigma}_{\mathcal{T}} - I_K) \beta_{\mathcal{T}}^0\|_{\infty} = \frac{\|\beta_{\mathcal{T}}^0\|_{\infty}}{1+n} < C/n$.

Using the fact that $\|\varepsilon_{\mathcal{T}}\|_{\infty} \lesssim \sqrt{\log n}$ on the event \mathcal{A} , we obtain $\sqrt{n} \|\tilde{\Sigma}_{\mathcal{T}} \varepsilon_{\mathcal{T}}\|_{\infty} \lesssim \sqrt{\frac{\log n}{n}}$. The sum of the remaining two terms in (55) can be bounded by a multiple of $\sqrt{\log n/n}$ by noting that $\bar{\sigma}^2 = 1/(n+1)$. The statement (54) then follows from (55). \square

6.4. *Supremum-norm Convergence Rate.* Let us write $f_0 = f_0^{\mathcal{L}_c} + f_0^{\setminus \mathcal{L}_c}$, where $f_0^{\mathcal{L}_c}$ is the L^2 -projection of f_0 onto the first \mathcal{L}_c layers of wavelet coefficients. Under the Hölder condition the equality holds pointwise and $\|f_0^{\setminus \mathcal{L}_c}\|_{\infty} \leq \sum_{l > \mathcal{L}_c} 2^{l/2} 2^{-l(1/2+\alpha)} \lesssim (\log n/n)^{\alpha/(2\alpha+1)}$.

The following inequality bounds the supremum norm by the ℓ_{∞} -norm,

$$(57) \quad \begin{aligned} \|f - f_0\|_{\infty} &\leq \sum_{l \geq -1} \max_{0 \leq k < 2^l} |\beta_{lk} - \beta_{lk}^0| \cdot \left\| \sum_{0 \leq k < 2^{-l}} |\psi_{lk}| \right\|_{\infty} \\ &\leq |\langle f - f_0, \varphi \rangle| + \sum_{l \geq 0} 2^{l/2} \max_{0 \leq k < 2^l} |\beta_{lk} - \beta_{lk}^0| = \ell_{\infty}(f, f_0). \end{aligned}$$

We use the notation $S(f_0; A)$, \mathbb{T} as in (47) and (53) and

$$(58) \quad \mathcal{E} = \{f_{\mathcal{T}, \beta} : \mathcal{T} \in \mathbb{T}\}.$$

Using the definition of the event \mathcal{A} from (41), one can write

$$(59) \quad \begin{aligned} E_{f_0} \Pi[f_{\mathcal{T}, \beta} : \|f_{\mathcal{T}, \beta} - f_0\|_{\infty} > \varepsilon_n | X] &\leq P_{f_0}[\mathcal{A}^c] + E_{f_0} \Pi[\mathcal{E}^c | X] \\ &+ E_{f_0} \{\Pi[f_{\mathcal{T}, \beta} \in \mathcal{E} : \|f_{\mathcal{T}, \beta} - f_0\|_{\infty} > \varepsilon_n | X] \mathbb{I}_{\mathcal{A}}\}. \end{aligned}$$

By Markov's inequality and the previous bound (57),

$$\begin{aligned} \Pi[f_{\mathcal{T},\beta} \in \mathcal{E} : \|f_{\mathcal{T},\beta} - f_0\|_\infty > \varepsilon_n \mid X] \mathbb{I}_{\mathcal{A}} &\leq \varepsilon_n^{-1} \int_{\mathcal{E}} \|f_{\mathcal{T},\beta} - f_0\|_\infty d\Pi[f_{\mathcal{T},\beta} \mid X] \mathbb{I}_{\mathcal{A}} \\ &\leq \varepsilon_n^{-1} \sum_{l \leq \mathcal{L}_c} 2^{l/2} \left\{ \int_{\mathcal{E}} \max_{0 \leq k < 2^l} |\beta_{lk} - \beta_{lk}^0| d\Pi[f_{\mathcal{T},\beta} \mid X] \mathbb{I}_{\mathcal{A}} \right\} + \varepsilon_n^{-1} \|f_0^{\setminus \mathcal{L}_c}\|_\infty. \end{aligned}$$

With \mathbf{T} as in (53), the integral in the last display can be written, for $l \leq \mathcal{L}_c$,

$$\begin{aligned} \int_{\mathcal{E}} \max_{0 \leq k < 2^l} |\beta_{lk} - \beta_{lk}^0| d\Pi[f_{\mathcal{T},\beta} \mid X] &= \sum_{\mathcal{T} \in \mathbf{T}} \pi[\mathcal{T} \mid X] \int \max_{0 \leq k < 2^l} |\beta_{lk} - \beta_{lk}^0| d\Pi[\beta_{\mathcal{T}} \mid \mathbf{X}_{\mathcal{T}}] \\ &= \sum_{\mathcal{T} \in \mathbf{T}} \pi[\mathcal{T} \mid X] \int \max \left(\max_{0 \leq k < 2^l, (l,k) \notin \mathcal{T}'_{int}} |\beta_{lk}^0|, \max_{0 \leq k < 2^l, (l,k) \in \mathcal{T}'_{int}} |\beta_{lk} - \beta_{lk}^0| \right) d\Pi[\beta_{\mathcal{T}} \mid \mathbf{X}_{\mathcal{T}}] \\ &\leq \min \left(\max_{0 \leq k < 2^l} |\beta_{lk}^0|, A \frac{\log n}{\sqrt{n}} \right) + \sum_{\mathcal{T} \in \mathbf{T}} \pi[\mathcal{T} \mid X] \int \max_{0 \leq k < 2^l, (l,k) \in \mathcal{T}'_{int}} |\beta_{lk} - \beta_{lk}^0| d\Pi[\beta_{\mathcal{T}} \mid \mathbf{X}_{\mathcal{T}}], \end{aligned}$$

where we have used that on the set \mathcal{E} , selected trees cannot miss any true signal larger than $A \log n / \sqrt{n}$. This means that any node (l, k) that is *not* in a selected tree must satisfy $|\beta_{lk}^0| \leq A \log n / \sqrt{n}$.

Let $L^* = L^*(\alpha)$ be the integer closest to the solution of the equation in L given by $M2^{-L(\alpha+1/2)} = A \log n / \sqrt{n}$. Then, using that $f_0 \in \mathcal{H}(\alpha, M)$,

$$\begin{aligned} \sum_{l \leq \mathcal{L}_c} 2^{l/2} \min \left(\max_{0 \leq k < 2^l} |\beta_{lk}^0|, A \frac{\log n}{\sqrt{n}} \right) &\leq \sum_{l \leq L^*} 2^{l/2} A \frac{\log n}{\sqrt{n}} + \sum_{L^* < l \leq \mathcal{L}_c} 2^{l/2} M2^{-l(\frac{1}{2}+\alpha)} \\ (60) \quad &\leq C2^{L^*/2} A \frac{\log n}{\sqrt{n}} + C2^{-L^*\alpha} \leq \tilde{C}2^{-L^*\alpha} \leq c(n^{-1} \log^2 n)^{\frac{\alpha}{2\alpha+1}}. \end{aligned}$$

Using $P_{f_0}[\mathcal{A}^c] + E_{f_0} \Pi[\mathcal{E}^c \mid X] = o(1)$ and Lemma 3, one obtains

$$\begin{aligned} E_{f_0} \Pi[f_{\mathcal{T},\beta} : \|f_{\mathcal{T},\beta} - f_0\|_\infty > \varepsilon_n \mid X] &\leq o(1) + \\ &\varepsilon_n^{-1} \sum_{l \leq \mathcal{L}_c} 2^{l/2} \left[\min \left(\max_{0 \leq k < 2^l} |\beta_{lk}^0|, A \frac{\log n}{\sqrt{n}} \right) + C' \sqrt{\frac{\log n}{n}} \right] + \varepsilon_n^{-1} \|f_0^{\setminus \mathcal{L}_c}\|_\infty \\ &\leq o(1) + \varepsilon_n^{-1} \left[c \left(\frac{\log^2 n}{n} \right)^{\frac{\alpha}{2\alpha+1}} + 2C' \sqrt{\frac{2^{\mathcal{L}_c} \log n}{n}} \right] + \varepsilon_n^{-1} \|f_0^{\setminus \mathcal{L}_c}\|_\infty \\ &\leq o(1) + \varepsilon_n^{-1} \left[c(\log n)^{\alpha/(2\alpha+1)} + 2C' \right] \left(\frac{\log n}{n} \right)^{\frac{\alpha}{2\alpha+1}} + \varepsilon_n^{-1} \|f_0^{\setminus \mathcal{L}_c}\|_\infty \end{aligned}$$

for some $C' > 0$. Choosing $\varepsilon_n = M_n ((\log^2 n)/n)^{\frac{\alpha}{2\alpha+1}}$, the right hand side goes to zero for any arbitrarily slowly increasing sequence $M_n \rightarrow \infty$. \square

SUPPLEMENT TO “UNCERTAINTY QUANTIFICATION FOR BAYESIAN CART”

BY ISMAËL CASTILLO[‡] AND VERONIKA ROČKOVÁ[§]

Sorbonne Université[‡]
University of Chicago[§]

*Sorbonne Université & Institut Universitaire de France
Laboratoire de Probabilités, Statistique et Modélisation, LPSM,
4, Place Jussieu, 75005 Paris cedex 05, France
ismael.castillo@upmc.fr*

*University of Chicago, Booth School of Business
5807 S. Woodlawn Avenue
Chicago, IL, 60637, USA
veronika.rockova@chicagobooth.edu*

This supplementary file contains additional material, including results for nonparametric regression, a simulation study, an adaptive nonparametric Bernstein–von Mises theorem, lower bound results, and details on tensor–multivariate versions of the considered prior distributions. It also contains all remaining proofs for the results stated in the main paper.

7. Additional results.

7.1. *Nonparametric regression: $\|\cdot\|_\infty$ –rate and bands.* Assume outcomes $Y = (Y_1, \dots, Y_n)'$ arising from

$$(61) \quad Y_i = f_0(t_i) + \varepsilon_i, \quad \varepsilon_i \stackrel{iid}{\sim} \mathcal{N}(0, 1), \quad i = 1, \dots, n = 2^{L_{max}+1}$$

where f_0 is an unknown regression function and $\{t_i \in [0, 1] : 1 \leq i \leq n\}$ are fixed design points. For simplicity we consider a regular grid, i.e. $t_i = i/n$ for $1 \leq i \leq n$ and assume n is a power of 2. Irregularly spaced design points could also be considered with more technical proofs. Below we show that many results for Bayesian CART posteriors obtained in the white noise model in the main paper carry over to the regression model (61). Although the white noise and regression models can be shown to be ‘asymptotically equivalent’ in the Le Cam sense under smoothness assumptions, this does not enable to transfer results for posterior distributions from one model to the other.

To derive a supremum norm contraction rate in this setting, we follow an approach close in spirit to the practically used Haar wavelet transform. Below, we put a prior on the *empirical* wavelet coefficients of the regression function f . Indeed, those coefficients can be directly related to the values $f(t_i)$.

Let $X = (x_{ij})$ denote the $(n \times p)$ regression matrix of $p = 2^{L_{max}} = n/2$ regressors constructed from Haar wavelets ψ_{lk} up to the maximal resolution L_{max} , i.e., for $1 \leq i \leq n$,

$$x_{ij} = \begin{cases} \psi_{-10}(t_i) = 1 & \text{for } j = 1 \\ \psi_{lk}(t_i) & \text{for } j = 2^l + k + 1. \end{cases}$$

AMS 2000 subject classifications: Primary 62G20, 62G15.

Keywords and phrases: Bayesian CART, Posterior Concentration, Nonparametric Bernstein–von Mises theorem, Recursive Partitioning, Regression Trees.

The columns have been ordered according to the index $2^l + k$ (from smallest to largest). Note that the matrix X is orthogonal and $X'X = nI_n$. In the sequel we denote $F_0 = (f_0(t_1), \dots, f_0(t_n))'$ the vector of realized values of the true regression function at the design points. Also, we set, for two functions f, g defined on $[0, 1]$,

$$\|f - g\|_{\infty, n} = \max_{1 \leq i \leq n} |f(t_i) - g(t_i)|.$$

For given indexes l, k , the empirical wavelet coefficient b_{lk}^0 of f_0 is

$$(62) \quad b_{lk}^0 = n^{-1} \sum_{i=1}^n f_0(t_i) \psi_{lk}(t_i).$$

Let $\mathbf{b}^0 = (b_{lk}^0)$ denote the vector of ordered empirical coefficients. Since $X^{-1} = X'/n$, we see that $\mathbf{b}^0 = X^{-1}F_0$ or $F_0 = X\mathbf{b}^0$, so model 61 can be rewritten, with $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)'$,

$$(63) \quad Y = X\mathbf{b}^0 + \varepsilon.$$

Setting $Z = X^{-1}Y$ and $\eta = \sqrt{n}X^{-1}\varepsilon$, we have

$$(64) \quad Z = \mathbf{b}^0 + \eta/\sqrt{n}.$$

Since $\eta/\sqrt{n} \sim \mathcal{N}(0, X'X/n)$ is a vector of iid $\mathcal{N}(0, 1)$ variables (as $X'X = nI$), the vector Z follows a Gaussian sequence model truncated at the n th observation. On the other hand, note that the likelihood in model (63) is proportional to $\exp(-\|Y - X\mathbf{b}^0\|^2/2) = \exp(-n\|Z - \mathbf{b}^0\|^2/2)$. Therefore, the posterior distribution induced by putting a prior distribution on \mathbf{b}^0 in model (63) is the same as the one obtained by choosing the same prior on \mathbf{b}^0 in model (64). Let us denote by $\Pi_b[\cdot | Y]$ this posterior distribution on vectors \mathbf{b} . Theorem 5 below states that its convergence rate in terms of the maximum norm $\|f - f_0\|_{\infty, n}$ is the same as the rate obtained in the main paper. If one rather wishes to control $\|f - f_0\|_{\infty}$, it is possible to produce a procedure that yields a rate for the whole function f by interpolation as follows.

Let \mathcal{I} be the map that takes a vector of values $\varphi := (f(t_i))$ of size n and maps it to the piecewise-linear function $\mathcal{I}(\varphi)$ on $[0, 1]$ that linearly interpolates between the values $f(t_i)$ (for definiteness, assume $\mathcal{I}(\varphi)$ takes the constant value $f(t_1)$ on $[0, t_1]$). Further define, for χ the map $\chi(\mathbf{b}) = \mathcal{I}(X\mathbf{b})$,

$$(65) \quad \bar{\Pi}_Y = \Pi_b[\cdot | Y] \circ \chi^{-1}.$$

In words, $\bar{\Pi}$ is simply the distribution on functions on $[0, 1]$ induced as follows: sampling from the posterior $\Pi_b[\cdot | Y]$ induces a posterior on vectors $(f(t_i))$, from which one obtains a distribution on functions f 's by the linear interpolation \mathcal{I} .

THEOREM 5. *Let Π_b denote the prior distribution on empirical wavelet coefficients in model (61)/(63) defined in the same way as the prior in Theorem 1, with ε_n the rate as in that statement. Then, for $\Pi_b[\cdot | Y]$ the corresponding posterior distribution, for any $\alpha \in (0, 1]$, $M > 0$ and any sequence $M_n \rightarrow \infty$,*

$$\sup_{f_0 \in \mathcal{H}_M^\alpha} E_{f_0} \Pi_b[f : \|f - f_0\|_{\infty, n} > M_n \varepsilon_n | Y] \rightarrow 0.$$

and, for $\bar{\Pi}_Y$ the distribution defined in (65),

$$\sup_{f_0 \in \mathcal{H}_M^\alpha} E_{f_0} \bar{\Pi}_Y[f : \|f - f_0\|_{\infty} > M_n \varepsilon_n] \rightarrow 0.$$

The proposed approach uses the relationship between function values and empirical wavelet coefficients given by the Haar transform. It is worth noticing that it naturally yields a result on the canonical empirical max-loss $\|\cdot\|_{\infty,n}$ and avoids regularity conditions on the true f_0 (such as a minimal smoothness level $\alpha > 1/2$). This approach could be extended to handle non-equally spaced designs and/or a smoother wavelet basis, under appropriate conditions on the matrix $X'X$. This is beyond the scope of this paper, but we refer to the work [61] for results in this vein for spike-and-slab priors.

REMARK 1. *Instead of a continuous linear interpolation \mathcal{I} as above, one may use instead a piecewise constant interpolation: defining $\mathcal{J}((f(t_i)))$ to be the histogram that takes the value $f(t_i)$ on $[t_i, t_{i+1})$, Theorem 5 holds for $\tilde{\Pi}_Y$ defined in a similar way as $\tilde{\Pi}_Y$ but with \mathcal{J} in place of \mathcal{I} , as can easily be seen from the proof of Theorem 5.*

REMARK 2. *A seemingly different approach to estimating f consists in putting a prior distribution directly on the function f via putting a prior distribution on its (Haar-) wavelet coefficients. Note however that the induced prior on $f(t_i)$ is the same as the one above, since $f(t_i) = \sum_{l,k} \beta_{lk} \psi_{lk}(t_i)$ can be rewritten $F = X\beta$ since by definition the prior does not put mass on β_{lk} for $l > L_{max}$. One may also note that for a piecewise constant function over intervals $[t_i, t_{i+1})$, empirical Haar-wavelet coefficients (the b_{lk} 's) and Haar-wavelet coefficients (the β_{lk} 's) coincide. This implies that the posterior distributions induced on the vector $(f(t_i))$ through both approaches coincide. Since posterior samples are piecewise constant on dyadic intervals, the posterior distribution on f 's obtained from using the prior on β_{lk} 's as above coincides with the posterior $\tilde{\Pi}_Y$ from Remark 1.*

We now turn to the problem of construction of confidence bands for f in the regression model (61). We follow the approach above and model the empirical wavelet coefficients \mathbf{b}^0 in (63)–(64) via the tree prior Π_b as in Theorem 5. Recall Definition 4 of the median tree associated with a posterior over trees. Let $\mathcal{T}_{\tilde{Y}}$ denote the median tree associated with the posterior distribution $\Pi_b[\cdot | Y]$. Given the observed noisy empirical wavelet coefficients Z_{lk} as in (64), let us define an *empirical median tree estimator* \tilde{f}_T over gridpoints (t_i) as

$$(66) \quad \tilde{f}_T(t_i) = \sum_{(l,k) \in \mathcal{T}_{\tilde{Y}}} Z_{lk} \psi_{lk}(t_i).$$

Define, for some $v_n \rightarrow \infty$ to be chosen,

$$(67) \quad \tilde{\sigma}_n = \tilde{\sigma}_n(Y) = \max_{1 \leq i \leq n} \sum_{l=0}^{L_{max}} v_n \sqrt{\frac{\log n}{n}} \sum_{k=0}^{2^l-1} \mathbb{I}_{(l,k) \in \mathcal{T}_{\tilde{Y}}} |\psi_{lk}(t_i)|.$$

A credible band with radius $\tilde{\sigma}_n(Y)$ as in (67) and center \tilde{f}_T as in (66) is

$$(68) \quad \mathcal{C}_n^e = \left\{ f : \|f - \tilde{f}_T\|_{\infty,n} \leq \tilde{\sigma}_n(Y) \right\}.$$

We obtain in the next statement the analogue of Theorem 2 for regression. The confidence band is in terms of the natural empirical norm $\|\cdot\|_{\infty,n}$. We slightly update the definition of self-similar functions by restricting to f 's in \mathcal{H}_M^α in Definition 3, instead of $\mathcal{H}(\alpha, M)$ defined from wavelet coefficients. This is for technical convenience because $\mathcal{H}_{\alpha,M}$ is more natural for controlling empirical wavelet coefficients. We denote by $\mathcal{H}'_{SS}(\alpha, M, \varepsilon)$ the corresponding class (note that the classes are nearly the same; we could also have opted for taking \mathcal{H}_M^α in Definition 3, which would have avoided the distinction).

THEOREM 6. *Let $0 < \alpha_1 \leq \alpha_2 \leq 1$, $M \geq 1$ and $\varepsilon > 0$. Let Π_b be a prior as in the statement of Theorem 5. Let $\tilde{\sigma}_n$ be as in (67) with v_n such that $(\log n)^{1/2} = o(v_n)$ and let \tilde{f}_T denote the median tree estimator (66). Then for \mathcal{C}_n^e defined in (68), uniformly over $\alpha \in [\alpha_1, \alpha_2]$, as $n \rightarrow \infty$,*

$$\inf_{f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)} P_{f_0}(f_0 \in \mathcal{C}_n^e) \rightarrow 1.$$

For every $\alpha \in [\alpha_1, \alpha_2]$ and uniformly over $f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)$, the diameter $|\mathcal{C}_n^e|_{\infty, n} = \sup_{f, g \in \mathcal{C}_n^e} \|f - g\|_{\infty, n}$ and the credibility of the band verify, as $n \rightarrow \infty$,

$$(69) \quad |\mathcal{C}_n^e|_{\infty, n} = O_{P_{f_0}}((n/\log n)^{-\alpha/(2\alpha+1)} v_n),$$

$$(70) \quad \Pi[\mathcal{C}_n^e | Y] = 1 + o_{P_{f_0}}(1).$$

The confidence band \mathcal{C}_n^e is slightly conservative in the sense that its coverage and credibility go to 1 as $n \rightarrow \infty$. By intersecting this band with an appropriate ball, using a nonparametric BvM theorem, one can build a band with desired prescribed coverage $1 - \gamma$, $\gamma > 0$, as demonstrated in Section 7.3. The latter ‘intersection’-band is actually the one we implement in simulations in the next section and illustrated in Figure 7.2.

7.2. Numerical Examples. We highlight the practicality of the confidence bands presented in Section 3.2 (and the following Section 7.3.5) on numerical examples. We implement Dyadic Bayesian CART as in [20] using the standard Metropolis-Hastings algorithm with a proposal distribution consisting of two steps: grow (splitting a randomly chosen bottom node) and prune (collapsing two children bottom nodes into one). The implementation is fairly straightforward due to immediate access to the posterior tree probabilities through a variant of (22) (these closed-form calculations can be easily updated for non-parametric regression).

We illustrate uncertainty quantification for f_0 using the intersection band (75). This construction yields exact asymptotic coverage; it uses up more posterior information and in this sense can be viewed as more Bayesian in spirit. Kolyan Ray [47], Section 6, implemented the intersection-credible set in the case of spike-and-slab priors and the white noise model. Here we provide an implementation in the regression model for tree priors. For the computation of the empirical median tree estimator in (66), one can easily identify nodes (l, k) that have occurred in at least 50% of posterior samples. Regarding the computation of R_n in (74), the radius can be approximated using the 95% quantile of the posterior samples of the multiscale norm $\|\cdot\|_{\mathcal{M}(w)}$ -norm (acting on coefficients up to level L_{max}). Therefore, we see that while the confidence band in (75) may at first look computationally cumbersome, it can be readily obtained from posterior samples.

We generate $n = 2^{11}$ observations from (61) with $f_0(x) = 2\sqrt{10x}$ (top row in Figure 6) and $f_0(x) = (4x - 1)\mathbb{I}(x \leq 1/2) + (-2x + 2)\mathbb{I}(x > 1/2)$ (bottom row in Figure 6). We choose $j_0(n) = 4$ and $w_l = l^{1/2+\epsilon}$ (the multiscale weighting sequence) with $\epsilon = 0.01$ and $\Gamma = 1.01$ (the splitting probability parameter of the GW process prior). We run 2000 iterations of the MH sampler with 500 burnin samples. One tempting approach to uncertainty quantification is computing the *pointwise* 95% credible intervals for each given t_i . These intervals are readily available from the posterior samples of the bottom node coefficients (transformed from the samples of the wavelet coefficients via the pinball formula (13)) and are portrayed in Figure 6 on the left. For both functions f_0 these intervals are too narrow to uniformly capture f_0 (depicted in a green dashed line). In comparison, the multiscale 0.95% credible band in (75) (Figure 6 on the right) is more successful at containing the true function uniformly. Figure 6 plots the intersection band which has exact asymptotic coverage. Note, however, that this

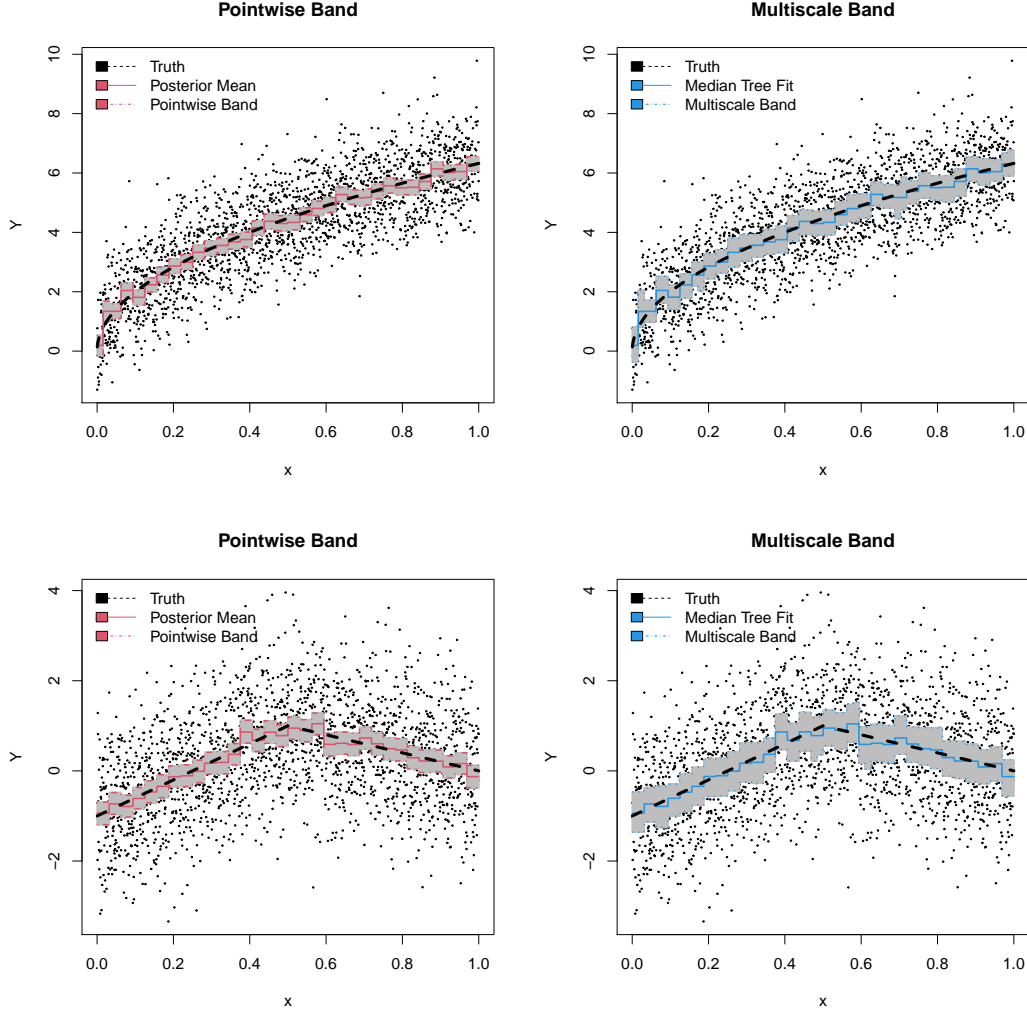


FIG 6. (Left) Pointwise 0.95% credible intervals. (Right) Multiscale 0.95% credible band (31) obtained with $v_n = 1$ and $w_l = l^{1/2+0.01}$. (Upper) The true function is $f_0(x) = 2\sqrt{10x}$. (Lower) The true function is $f_0(x) = (4x - 1)\mathbb{I}(x \leq 1/2) + (-2x + 2)\mathbb{I}(x > 1/2)$.

intersecting band is smaller than the band analyzed in Theorem 6 which yields even better coverage. The centering in Figure 6 (right) is the median tree estimator.

To illustrate the usefulness of Theorem 3, we plot the histogram of posterior samples (together with 2.5% and 97.5% quantiles) of the primitive functional $F(x) = \sum_{t_i \leq x} f(t_i)$ for $x = 0.8$ for the piecewise linear model. The true value is marked with a red solid line. We can see that the 95%-credible interval covers the true value. We also construct the confidence band (34) (Figure 6 on the right).

7.3. Adaptive nonparametric BvM and applications.

7.3.1. From $\|\cdot\|_\infty$ to BvM. Let us now formalize the notion of a nonparametric BvM theorem in multiscale spaces following [16] (we refer also to [15] for more background and discussion of the, different, L^2 -type setting). Such spaces are defined through the speed of decay of multiscale coefficients $\beta_{lk} = \langle f, \psi_{lk} \rangle$. For a monotone increasing weighting se-




FIG 7. 0.95% credible interval for the primitive functional $F(x)$ with $x = 0.8$ (left) and the confidence band (34) (right) obtained for $f_0(x) = (4x - 1)\mathbb{I}(x \leq 1/2) + (-2x + 2)\mathbb{I}(x > 1/2)$.

quence $w = (w_l)_{l=0}^\infty$, with $w_l \geq 1$ and $w_l/\sqrt{l} \rightarrow \infty$ as $l \rightarrow \infty$ (such a $w = (w_l)_{l=0}^\infty$ is called *admissible*) we define the following *multiscale sequence space*

$$\mathcal{M}(w) = \left\{ x = (x_{lk}) : \|x\|_{\mathcal{M}(w)} := \sup_l \frac{\max_k |x_{lk}|}{w_l} < \infty \right\}.$$

We consider a separable closed subspace of $\mathcal{M}(w)$ defined as

$$\mathcal{M}_0(w) = \left\{ x \in \mathcal{M}(w) : \lim_{l \rightarrow \infty} \max_k \frac{|x_{lk}|}{w_l} = 0 \right\}.$$

Defining random variables $g_{lk} = \int_0^1 \psi_{lk}(t) dW(t) \sim \mathcal{N}(0, 1)$, according to Proposition 2 in [16], the Gaussian white noise $\mathbb{W} = (g_{lk})$ defines a tight Gaussian Borel measure in the space $\mathcal{M}_0(w)$ for admissible sequences w . The convergence in distribution of random variables in the multiscale space $\mathcal{M}_0(w)$ is metrised via the bounded Lipschitz metric $\beta_{\mathcal{M}_0(w)}$ defined below. For μ, η probability measures on a metric space (S, d) define

$$\beta_S(\mu, \eta) = \sup_{F: \|F\|_{BL} \leq 1} \left| \int_S F(x) (d\mu(x) - d\eta(x)) \right|,$$

$$\|F\|_{BL} = \sup_{x \in S} |F(x)| + \sup_{x \neq y, x, y \in S} \frac{|F(x) - F(y)|}{d(x, y)}.$$

Denote with $\mathbb{X} = \mathbb{X}^{(n)} = (X_{lk} : l \in \mathbb{N}_0, 0 \leq k < 2^l)$, where X_{lk} satisfy (2). Let $\tilde{\Pi}_n = \Pi_n \circ \tau_{\mathbb{X}}^{-1}$ be the image measure of $\Pi(\cdot | X)$ under $\tau_{\mathbb{X}} : f \rightarrow \sqrt{n}(f - \mathbb{X})$. Namely, for any Borel set B we have

$$(71) \quad \tilde{\Pi}_n(B) = \Pi(\sqrt{n}(f - \mathbb{X}) \in B | X).$$

The following Theorem characterizes the *adaptive* nonparametric Bernstein-von Mises behavior of posteriors under the Bayesian Dyadic CART. In the result below, one assumes that trees sampled from $\Pi_{\mathbb{T}}$ contain all nodes (j, k) for all $j \leq j_0(n) \rightarrow \infty$ slowly. Note that this constraint is easy to accommodate in the construction: for the GW process, one starts stopping splits only after depth $j_0(n)$, while for priors (5), it suffices to constrain the indicator $\mathbb{I}_{\mathcal{T} \in \mathbb{T}}$ to trees that fill all first $j_0(n)$ layers.

THEOREM 7. (Adaptive nonparametric BvM) Let $\mathcal{M}_0 = \mathcal{M}_0(w)$ for some admissible sequence $w = (w_l)$. Assume the Bayesian CART priors $\Pi_{\mathbb{T}}$ from Theorem 1 constrained to trees that fit $j_0(n)$ layers, i.e. $\gamma_{lk} = 1$ for $l \leq j_0(n)$, for some strictly increasing sequence $j_0(n) \rightarrow \infty$ that satisfies $w_{j_0(n)} \geq c \log n$ for some $c > 0$. Consider tree-shaped priors as in Theorem 1 (or using an S -regular wavelet basis, $S \geq 1$). Then the posterior distribution satisfies the weak Bernstein-von Mises phenomenon in \mathcal{M}_0 in the sense that

$$E_{f_0} \beta_{\mathcal{M}_0}(\tilde{\Pi}_n, \mathcal{N}) \rightarrow 0 \quad \text{as } n \rightarrow \infty,$$

where \mathcal{N} is the law of \mathbb{W} in \mathcal{M}_0 .

This result states an *adaptive* nonparametric BvM result, in the sense that the prior it considers also leads to an adaptive nonparametric convergence rate in L^∞ (optimal up to log terms). It is only the second result of this kind after the one derived by Ray in [47]. This statement, proved in Section 12.3, can be shown, for example, by verifying the conditions in Proposition 6 of [16] (appropriate ‘tightness’ and convergence of finite dimensional distributions). The first tightness condition pertains to contraction in the \mathcal{M}_0 -space, which can be obtained from our $\|\cdot\|_\infty$ -results. In order to attain BvM, we need to modify the prior to always include a few coarsest dense layers in the tree (similarly as [47]). Such trees are semi-dense, where sparsity kicks in only deeper in the tree after $j_0(n)$ dense layers have already been fitted. This enables one to derive the convergence of finite dimensional distributions to suitable Gaussian distributions. For the independent wavelet prior, the last point follows easily from results in [15]. For the g -prior on trees corresponding to actual Bayesian CART, it requires a completely new argument based on the conditional posteriors given possible trees.

7.3.2. Application 1: multiscale confidence sets. First, let us consider multiscale credible balls for f_0 , which we will use in the next subsection for refining the band construction used in the main paper. Such multiscale balls consist of functions f that simultaneously satisfy multi-scale linear constraints (see e.g. (5) in [16]):

$$(72) \quad \mathcal{B}_n = \{f : \|f - \mathbb{X}\|_{\mathcal{M}(w)} \leq R_n / \sqrt{n}\},$$

where R_n is chosen such that $\Pi[\mathcal{B}_n | X] = 1 - \gamma$ (or the smallest radius such that $\Pi[\mathcal{B}_n | X] \geq 1 - \gamma$), i.e. \mathcal{B}_n is a credible set of level $1 - \gamma$.

PROPOSITION 1. Let $f_0 \in \mathcal{H}_M^\alpha$ for some $\alpha \in (0, 1]$, $M > 0$. Then for \mathcal{B}_n as in (72),

$$P_{f_0}(f_0 \in \mathcal{B}_n) \rightarrow 1 - \gamma, \quad (\text{as } n \rightarrow \infty).$$

The proof of the Proposition 1 is a consequence of the fact that the nonparametric BvM in the multiscale space holds: one combines Theorem 7 and Theorem 5 in [16].

7.3.3. Application 2: BvM for functionals. As a second application of Theorem 7, we derive confidence bands for $F(t) := \int_0^t f(x)dx$, $0 \leq t \leq 1$: those result from taking quantile credible bands in the following limiting distribution result. This application is described in Theorem 3 and the (short) proof is in Section 10.4.

7.3.4. Application 3: multiscale confidence bands in the white noise model. Let us first consider the case of the white noise model. For R_n as in (72), $\sigma_n(X)$ as in (30) and \hat{f}_T as in (29), let us set

$$(73) \quad \mathcal{C}_n^{\mathcal{M}} = \left\{ f : \|f - \mathbb{X}\|_{\mathcal{M}(w)} \leq R_n / \sqrt{n}, \|f - \hat{f}_T\|_\infty \leq \sigma_n(X) \right\}.$$

Let us recall the definition of the self-similarity class $\mathcal{H}_{SG}(\alpha, M, \varepsilon)$ from Definition 3.

COROLLARY 1. *Let $0 < \alpha_1 \leq \alpha_2 \leq 1$, $M \geq 1$, $\gamma \in (0, 1)$ and $\varepsilon > 0$. Let the prior Π and the sequence (w_l) be as in the statement of Theorem 7. Take R_n as in (72), σ_n as in (30) with v_n such that $(\log n)^{1/2} = o(v_n)$ and let \tilde{f}_T denote the median tree estimator (29). Then for $\mathcal{C}_n^{\mathcal{M}}$ defined in (73), uniformly over $\alpha \in [\alpha_1, \alpha_2]$,*

$$\sup_{f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)} |P_{f_0}(f_0 \in \mathcal{C}_n^{\mathcal{M}}) - (1 - \gamma)| \rightarrow 0,$$

as $n \rightarrow \infty$. In addition, for every $\alpha \in [\alpha_1, \alpha_2]$ and uniformly over $f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)$, the diameter $|\mathcal{C}_n^{\mathcal{M}}|_{\infty} = \sup_{f, g \in \mathcal{C}_n^{\mathcal{M}}} \|f - g\|_{\infty}$ and the credibility of the band verify, as $n \rightarrow \infty$,

$$|\mathcal{C}_n^{\mathcal{M}}|_{\infty} = O_{P_{f_0}}((n/\log n)^{-\alpha/(2\alpha+1)} v_n),$$

$$\Pi[\mathcal{C}_n^{\mathcal{M}} | X] = 1 - \gamma + o_{P_{f_0}}(1).$$

This result states that by intersecting the confidence set \mathcal{C}_n in (31) with the ball \mathcal{B}_n from (72), one obtains a set with confidence (and credibility) at the prescribed level $1 - \gamma$. It directly follows by applying Proposition 1 (which guarantees that \mathcal{B}_n has confidence level $1 - \gamma$ and hence also $\mathcal{C}_n^{\mathcal{M}}$, since \mathcal{C}_n has confidence that goes to 1) and the fact that by definition \mathcal{B}_n has credibility $1 - \gamma$ (and hence also $\mathcal{C}_n^{\mathcal{M}}$ asymptotically).

7.3.5. Application 4: multiscale confidence bands in regression. Now consider the regression setting (61). Let us define a discrete analogue for the multiscale confidence ball (72). Recall that in this setting the observations are the Y 's and that X here denotes the matrix of $\psi_{lk}(t_i)$'s as introduced in Section 7.1. Denote $\tilde{F} = (\tilde{f}_T(t_i))_{1 \leq i \leq n}$ for \tilde{f}_T the empirical median tree estimator. Set

$$(74) \quad \mathcal{B}_n^b = \{b = (b_{lk})_{l \leq L_{max}, k} : \|b - X^{-1} \tilde{F}\|_{\mathcal{M}(w)} \leq R_n / \sqrt{n}\},$$

where R_n is defined in such a way that $\Pi[\mathcal{B}_n^b | Y] = 1 - \gamma$ (or possibly $\geq 1 - \gamma$) and where in slight abuse of notation $\|\cdot\|_{\mathcal{M}(w)}$ stands for the multiscale norm acting on coefficients up to level L_{max} only (i.e. the supremum over l in the definition of $\|\cdot\|_{\mathcal{M}(w)}$ is replaced by the maximum over $l \leq L_{max}$). For \mathcal{C}_n^e as in (68), recalling the notation $F = (f(t_i))'_{1 \leq i \leq n}$, define

$$(75) \quad \tilde{\mathcal{C}}_n^{\mathcal{M}} = \mathcal{C}_n^e \cap \left\{ f : X^{-1} F \in \mathcal{B}_n^b \right\}.$$

COROLLARY 2. *Consider a prior as in Theorem 7, and let \mathcal{C}_n^e be constructed as in Theorem 6. For $\gamma > 0$, let $\tilde{\mathcal{C}}_n^{\mathcal{M}}$ be defined as in (75). The set $\tilde{\mathcal{C}}_n^{\mathcal{M}}$ verifies the properties stated in Theorem 6 except that both confidence and credibility go to the nominal level $1 - \gamma$ as $n \rightarrow \infty$.*

This result is obtained in a similar way as for Corollary 1: first, one notes that a BvM result similar to Theorem 7 in white noise holds (details are omitted, the proof being similar). This implies that the ball \mathcal{B}_n^b in (74) has confidence going to nominal level $1 - \gamma$. One concludes by using Theorem 6, that ensures that \mathcal{C}_n^e , and then in turn $\tilde{\mathcal{C}}_n^{\mathcal{M}}$, has the desired properties in terms of coverage, confidence and credibility.

7.4. Lower bound: the rate of Theorem 1 is sharp. The following result shows that the rate ε_n in (25) in Theorem 1 is in fact sharp. We state the result for smooth-wavelet priors, which enable to cover arbitrarily large regularities, but a similar result could also be formulated for the Haar basis.

THEOREM 8. *Let $\Pi_{\mathbb{T}}$ be one of the Bayesian CART priors from Theorem 1. Consider the tree-shaped wavelet prior (19) with $\pi(\beta_{\mathcal{T}}) \sim \mathcal{N}(0, \Sigma_{\mathcal{T}})$, where $\Sigma_{\mathcal{T}}$ is $I_{|\mathcal{T}_{ext}|}$ and $\{\psi_{lk}\}$ an S -regular wavelet basis, $S \geq 1$. Let ε_n be the rate defined in (25) for a given $0 < \alpha \leq S$. Let the parameters of $\Pi_{\mathbb{T}}$ verify either $\Gamma \geq \Gamma_0(S)$ a large enough constant, or $c \geq c_0 > 0$ large enough. For any $M > 0$, there exists $m > 0$ such that, as $n \rightarrow \infty$,*

$$(76) \quad \inf_{f_0 \in \mathcal{H}(\alpha, M)} E_{f_0} \Pi [\ell_{\infty}(f_{\mathcal{T}, \beta}, f_0) \leq m \varepsilon_n \mid X] \rightarrow 0.$$

In other words, there exists a sequence of elements of $\mathcal{H}(\alpha, M)$ along which the posterior convergence rate is *slower* than $m \varepsilon_n$ in terms of the ℓ_{∞} -metric. In particular, the upper-bound rate of Theorem 1 *cannot* hold uniformly over $\mathcal{H}(\alpha, M)$ with a rate faster than ε_n , which shows that the obtained rate is sharp (note the reversed inequality in (76) with respect to (26); we refer to [11] for more details on the notion of posterior rate lower bound). For simplicity we focused in the paper on the ℓ_{∞} -norm rather than the $\|\cdot\|_{\infty}$ -norm. Upper-bound results in ℓ_{∞} are stronger, since $\ell_{\infty}(f, f_0)$ dominates $\|f - f_0\|_{\infty}$. An interesting question is whether (76) still holds with the supremum norm. The answer is positive: by a slight adaptation of the proof below it is possible to show that the key bound $\ell_{\infty}(f_{\mathcal{T}, \beta}, f_0) \geq 2^{L^*/2} |\beta_{L^*0}|$ also holds up to a constant for $\|f_{\mathcal{T}, \beta} - f_0\|_{\infty}$ for the considered ‘least favorable’ sequence of functions.

7.5. Lower bound: flat trees are (grossly) suboptimal for the $\|\cdot\|_{\infty}$ -loss. Recall that the spike-and-slab prior achieves the *actual* ℓ_{∞} -minimax rate *without* any additional factor. Interestingly, the very same prior misses the ℓ_2 -minimax rate by a log factor [38]. This illustrates that ℓ_2 and ℓ_{∞} adaptations require different desiderata when constructing priors. Product priors that correspond to separable rules *do not* yield adaptation with exact rates in the ℓ_2 sense [10]. Mixture priors that are adaptive in ℓ_2 , on the other hand, may not yield ℓ_{∞} adaptation. We now provide one example of this phenomenon in the context of flat (complete binary) trees.

The *flat* tree of depth $d = d(\mathcal{T})$ is the binary tree which contains all possible nodes until level d , i.e. $\gamma_{lk} = \mathbb{I}_{l < d}$. An example of a flat tree with $d = 3$ layers is in Figure 2. The simplest possible prior on tree topologies (confined to symmetric trees) is just the Dirac mass at a given flat tree of fixed depth $d = D$; an adaptive version thereof puts a prior D and samples from the set of all flat trees. Such priors coincide with so-called *sieve* priors, where the sieve spans the expansion basis (e.g. Haar) up to level D .

Flat dyadic trees only keep Haar wavelet coefficients at resolutions smaller than some $d > 0$ (i.e. $\gamma_{lk} = 0$ for $l \geq d$). The implied prior on $(\beta_{lk})_{lk}$ can be written as, with $\pi(\beta_{lk}) \propto \sigma_l^{-1} \phi(\beta_{lk}/\sigma_l)$,

$$(77) \quad (\beta_{lk}) \mid d \sim \bigotimes_{l < d, k} \pi(\beta_{lk}) \otimes \bigotimes_{l \geq d, k} \delta_0(\beta_{lk}),$$

where $\phi(\cdot)$ is some bounded density that is strictly positive on \mathbb{R} and σ_l are fixed positive scalars. The sequence (σ_l) is customarily chosen so as it decays with the resolution index l , e.g. $\sigma_l = 2^{-l(\beta+1/2)}$ for some $0 < \beta \leq \alpha$. This “undersmoothing” prior requires the knowledge of (a lower bound on) α and yields a *non-adaptive* non-parametric BvM behavior [16].

A tempting strategy to manufacture adaptation is to treat the threshold d as random through a prior $\pi(d)$ on integers (and take constant σ_l), which corresponds to the hierarchical prior on regular regression histograms [49, 54]. It is not hard to check that the flat-tree prior (77) with random d has a marginal mixture distribution similar to the one of the spike-and-slab prior on each coordinate (l, k) . Despite marginally similar, the probabilistic structure of these two priors is very different. Zeroing out signals internally, the spike-and-slab prior (10) *is*




FIG 8. A plot of tensor Haar wavelets. The top figure plots $\Psi_{0(0,0)'}'$ and the bottom figures are $\Psi_{1(0,0)'}$, $\Psi_{0(1,0)'}$, $\Psi_{0(0,1)'}$, $\Psi_{0(1,1)'}$ (from left to right).

ℓ_∞ adaptive [38]. The flat tree prior (77), on the other hand, fits a few dense layers *without* internal sparsity and is ℓ_2 adaptive (up to a log term) [54]. However, as shown in the following Theorem, flat trees fall short of ℓ_∞ -adaptation.

THEOREM 9. Assume the flat tree prior (77) with random d , where $\pi(d)$ is non-increasing and where the active wavelet coefficients β_{lk} are Gaussian iid $\mathcal{N}(0, 1)$. Moreover, assume $\{\psi_{lk}\}$ is an S -regular wavelet basis for some $S \geq 1$. For any $0 < \alpha \leq S$ and $M > 0$, there exists $f_0 \in \mathcal{H}(\alpha, M)$ such that

$$E_{f_0} \Pi[\ell_\infty(f_{\mathcal{T}, \beta}, f_0) < \zeta_n \mid X] \rightarrow 0,$$

where the lower-bound rate ζ_n is given by $\zeta_n = \left(\frac{\log n}{n}\right)^{\frac{\alpha}{2\alpha+2}}$.

Theorem 9 can be applied to standard priors $\pi(d)$ with exponential decrease, proportional to e^{-d} or $e^{-d \log d}$, or to a uniform prior over $\{1, \dots, L_{max}\}$. In [1], a negative result is also derived for sieve-type priors, but only for the posterior mean and for Sobolev classes instead of the, here arguably more natural, Hölder classes for supremum losses (which leads to different rates for estimating the functional-at-a-point). Here, we show that when the target is the ℓ_∞ -loss for Hölder classes the sieve-prior is severely sub-optimal.

7.6. Multi-dimensional extensions. Our tree-shaped wavelet reconstruction generalizes to the multivariable case, where a fixed number $d \geq 1$ of covariate directions are available for split. We outline one such generalization using the tensor product of Haar basis functions ψ_{lk} from (3) defined as

$$\Psi_{l\mathbf{k}}(\mathbf{x}) := \psi_{lk_1}(x_1) \cdots \psi_{lk_d}(x_d)$$

for $l \geq 0$ and $\mathbf{k} = (k_1, \dots, k_d)'$ with $0 \leq k_i \leq 2^l - 1$ for $i = 1, \dots, d$, where $\Psi_{-1\mathbf{0}}(\mathbf{x}) = \mathbb{I}_{(0,1]^d}(\mathbf{x})$. These wavelet tensor products can be associated with d -ary trees (as opposed to binary trees), where each internal node has 2^d children. The nodes in a d -ary tree satisfy a hierarchical constraint: $(l, \mathbf{k}) \in \mathcal{T}, l \geq 1 \Rightarrow (l-1, \lfloor \mathbf{k}/2 \rfloor) \in \mathcal{T}$, where the floor operation is applied element-wise. This intuition can be gleaned from Figure 7 which organizes tensor wavelets with $l = 0, 1$ and $d = 2$ in a flat 4-ary tree. We assume that f_0 belongs to α -Hölder functions on $[0, 1]^d$ for $0 < \alpha \leq 1$ defined as

$$(78) \quad \mathcal{H}_M^{\alpha, d} := \left\{ f \in \mathcal{C}([0, 1]^d) : \|f\|_\infty + \sup_{\mathbf{x} \neq \mathbf{y}} \frac{|f(\mathbf{x}) - f(\mathbf{y})|}{\|\mathbf{x} - \mathbf{y}\|^\alpha} \leq M \right\}.$$

The multiscale coefficients $\beta_{-1\mathbf{0}} = \langle f_0, \Psi_{-1\mathbf{0}} \rangle$ and

$$\beta_{l\mathbf{k}} = \langle f_0, \Psi_{l\mathbf{k}} \rangle = \int_{[0,1]^d} f_0(\mathbf{x}) \Psi_{l\mathbf{k}}(\mathbf{x}) d\mathbf{x}.$$

can be verified to satisfy, for some universal constant $C > 0$,

$$(79) \quad |\beta_{l\mathbf{k}}| \leq C 2^{-l(\frac{1}{2} + \alpha)d}.$$

Similarly as in Section 2.2, denoting with \mathcal{T}'_{int} the collection of internal nodes (l, \mathbf{k}) in a d -ary tree (including the node $(-1, \mathbf{0})$), one then obtains a wavelet reconstruction $f_{\mathcal{T}, \beta}(\mathbf{x}) = \sum_{(l, \mathbf{k}) \in \mathcal{T}'_{int}} \beta_{l\mathbf{k}} \Psi_{l\mathbf{k}}(\mathbf{x})$, where coefficients $\beta_{l\mathbf{k}}$ can be assigned, for instance, a Gaussian independent product prior. There are several options for defining the d -dimensional version of the prior $\Pi_{\mathbb{T}}$. Restricting to Galton-Watson type priors, the most direct extension, for each node (l, \mathbf{k}) to be potentially split, either does not split it with probability $1 - \Gamma^{-l}$, or splits it into 2^d children, leading to a full 2^d -ary tree. Another, more flexible option, is to split (l, \mathbf{k}) into a random number of children inbetween 0 and 2^d , where a split in each specific direction occurs with probability Γ^{-l} , for Γ a large enough constant.

Assuming that d is fixed as $n \rightarrow \infty$, the general proving strategy of Theorem 1 can still be applied to conclude ℓ_∞ -posterior convergence at the rate $\varepsilon_n = (\log n/n)^{\alpha/(2\alpha+d)} \log^\delta n$ for some $\delta > 0$. The proof requires the threshold \mathcal{L}_c in (42) to be modified as satisfying $2^{\mathcal{L}_c} \asymp (n/\log n)^{1/(2\alpha+d)}$.

8. Basic Lemmata.

8.1. *Properties of the pinball matrix (14).* While $A'_{\mathcal{T}} A_{\mathcal{T}}$ is not proportional to an identity matrix (for trees other than flat trees), it *does* have a nested sparse structure which will be exploited in our analysis.

PROPOSITION 2. *Denote with (l_1, k_1) the deepest rightmost internal node in the tree \mathcal{T} , i.e. the node $(l, k) \in \mathcal{T}_{int}$ with the highest index $2^l + k$. Let \mathcal{T}^- be a tree obtained from \mathcal{T} by turning (l_1, k_1) into a terminal node. Then*

$$(80) \quad A'_{\mathcal{T}} A_{\mathcal{T}} = \begin{pmatrix} A'_{\mathcal{T}^-} A_{\mathcal{T}^-} + \mathbf{v}\mathbf{v}' & \mathbf{0} \\ \mathbf{0}' & 2^{l_1+1} \end{pmatrix}$$

for a vector of zeros $\mathbf{0} \in \mathbb{R}^{|\mathcal{T}_{ext}|-1}$ and a vector $\mathbf{v} \in \mathbb{R}^{|\mathcal{T}_{ext}|-1}$ obtained from $A_{\mathcal{T}}$ by first deleting its last column and then transposing the last row of this reduced matrix.

PROOF. The index (l_1, k_1) , by definition, corresponds to the last entry in the vector $\beta_{\mathcal{T}}$. We note that $\mathcal{T}_{int}^- = \mathcal{T}_{int} \setminus \{(l_1, k_1)\}$ and $\mathcal{T}_{ext}^- = \mathcal{T}_{ext} \setminus \{(l_1 + 1, 2k_1), (l_1 + 1, 2k_1 + 1)\} \cup \{(l_1, k_1)\}$. The matrix $A_{\mathcal{T}^-}$ can be obtained from $A_{\mathcal{T}}$ by deleting the last column of $A_{\mathcal{T}}$ and then deleting the last row, further denoted with \mathbf{v}' . The desired statement (80) is obtained by noting that the last column of $A_{\mathcal{T}}$ (associated with β_{l_1, k_1}) is orthogonal to all the other columns. This is true because (a) this column has only two nonzero entries that correspond to the last two siblings $\{(l_1 + 1, 2k_1), (l_1 + 1, 2k_1 + 1)\}$, (b) the last two rows of $A_{\mathcal{T}}$ differ only in the sign of the last entry because $\{(l_1 + 1, 2k_1), (l_1 + 1, 2k_1 + 1)\}$ are siblings and share the same ancestry with the same weights up to the sign of their immediate parent. Finally, the entry 2^{l_1+1} follows from (13). \square

COROLLARY 3. *Under the prior (17), the coefficient $\beta_{l\mathbf{k}}$ of any internal node (l, \mathbf{k}) which has terminal descendants is independent of all the remaining internal coefficients.*

PROOF. Follows directly from Prop. 2 after reordering the nodes. \square

The following proposition characterizes the eigenspectrum of $A'_{\mathcal{T}}A_{\mathcal{T}}$ which will be exploited in our proofs.

PROPOSITION 3. *The eigenspectrum of $A'_{\mathcal{T}}A_{\mathcal{T}}$ consists of the diagonal entries of $\mathbf{D} = \text{diag}(\tilde{d}_{lk, lk}) = A_{\mathcal{T}}A'_{\mathcal{T}}$ in (15). Moreover, the diagonal entries $\text{diag}(A'_{\mathcal{T}}A_{\mathcal{T}}) = \{d_{lk, lk}\}_{lk \in \mathcal{T}_{int}}$ satisfy $d_{-10, -10} = |\mathcal{T}_{ext}|$ and $d_{lk, lk} = \sum_{j=l+1}^{d(\mathcal{T})} 2^j \sum_{m=0}^{2^{j-1}} \mathbb{I}[\beta_{lk} \in [(0, 0) \leftrightarrow (j, m)]_{\mathcal{T}}]$ with $[(0, 0) \leftrightarrow (l, k)]_{\mathcal{T}} := \{(0, 0), (1, \lfloor k/2^{l-1} \rfloor), \dots, (l-1, \lfloor k/2 \rfloor)\}$.*

PROOF. The first statement follows from (15) and the fact that $A'_{\mathcal{T}}A_{\mathcal{T}}$ and $A_{\mathcal{T}}A'_{\mathcal{T}}$ have the same spectrum, and the second statement from (13). \square

8.2. Other lemmata.

LEMMA 4. *Assume that a square matrix A is diagonally dominant by rows (i.e., $a_{kk} > \sum_{j \neq k} |a_{kj}|$). Then*

$$\|A\|_{\infty} < \frac{1}{\min_k (|a_{kk}| - \sum_{j \neq k} |a_{kj}|)}.$$

PROOF. Theorem 1 in Varah [55]. \square

LEMMA 5. *For an invertible matrix $M \in \mathbb{R}^{p \times p}$ and $\mathbf{v} \in \mathbb{R}^p$ we have*

$$(M^{-1} + \mathbf{v}\mathbf{v}'/g_n)^{-1} = M - \frac{M\mathbf{v}\mathbf{v}'M}{g_n + \mathbf{v}'M\mathbf{v}} \quad \text{for } g_n > 0.$$

PROOF. Follows immediately by direct computation. \square

LEMMA 6. *Let \mathbb{C}_K denote the number of full binary trees with $K+1$ leaves. Then*

$$\mathbb{C}_K = \frac{(2K)!}{(K+1)!K!} \asymp 4^K / K^{3/2}.$$

PROOF. The number \mathbb{C}_K is the Catalan number (see e.g. [53]), which verifies the identity. The second assertion follows from Stirling's formula. \square

LEMMA 7. *Let $\mathbf{Y} \sim \mathcal{N}_K(\boldsymbol{\mu}, \Sigma)$ be a Gaussian random vector. Denote with $\{\sigma_i\}_{i=1}^K = \text{diag}(\Sigma)$, with $\bar{\mu} = \max_{1 \leq i \leq K} \mu_i$ and with $\bar{\sigma}^2 = \max_{1 \leq i \leq K} \sigma_i^2$ the maximal mean and variance. Then*

$$(81) \quad \mathbb{E} \left[\max_{1 \leq i \leq K} |Y_i| \right] \leq \bar{\mu} + \sqrt{2\bar{\sigma}^2 \log K} + 2\sqrt{2\pi\bar{\sigma}^2}.$$

PROOF. We start by noting that $|Y_i| \leq \bar{\mu} + |Y_i - \mu_i|$. Next, one can use the formula, valid for any real μ_i , $c > 0$ and real random variables Y_i ,

$$(82) \quad \mathbb{E} \left[\max_{1 \leq i \leq K} |Y_i - \mu_i| \right] \leq c + \sum_{i=1}^K \int_c^{\infty} \mathbb{P}(|Y_i - \mu_i| > x) dx.$$

Assuming the Gaussian distribution, the integral is of order $\int_c^{\infty} 2e^{-x^2/2\sigma_i^2} dx \leq \sqrt{2\pi\sigma_i^2} e^{-c^2/2\sigma_i^2}$. Then (81) follows from (82) by choosing $c = \sqrt{2\bar{\sigma}^2 \log K}$. \square

LEMMA 8 (see, e.g., [24]). *For a positive integer d , let $\mu, \mu_1, \mu_2 \in \mathbb{R}^d$ and let $\Sigma, \Sigma_1, \Sigma_2$ be positive definite $d \times d$ matrices. Then there exist universal constants $C_1, C_2 > 0$ such that, for TV the total variation distance,*

$$\begin{aligned} TV(\mathcal{N}(\mu, \Sigma_1), \mathcal{N}(\mu, \Sigma_2)) &\leq C_1 \|\Sigma_1^{-1} \Sigma_2 - I_d\|_F \\ TV(\mathcal{N}(\mu_1, \Sigma), \mathcal{N}(\mu_2, \Sigma)) &\leq C_2 \frac{\|\mu_1 - \mu_2\|^2}{\sqrt{(\mu_1 - \mu_2)' \Sigma (\mu_1 - \mu_2)}}, \end{aligned}$$

where $\|\cdot\|_F$ denotes the Frobenius norm.

PROOF. The first inequality follows from Theorem 1.1 in [24] and the second by Theorem 1.2 in [24] (by setting $\Sigma = \Sigma_1 = \Sigma_2$ in their statement). \square

9. Non-dyadic Bayesian CART: properties and examples.

9.1. Basic properties and examples.

LEMMA 9. *For a set A of admissible nodes (l, k) as above, let us define $\beta_{lk}^B = \langle f, \psi_{lk}^B \rangle$, where ψ_{lk}^B is the unbalanced Haar wavelet in (37) and where $f \in \mathcal{H}_M^\alpha$ was defined in (24). Then*

$$(83) \quad |\beta_{lk}^B| \leq M 2^{-1/2} \max\{|L_{lk}|, |R_{lk}|\}^{\alpha+1/2}.$$

PROOF. Let us denote by $\bar{C} \equiv 1/\sqrt{|L_{lk}|^{-1} + |R_{lk}|^{-1}}$. We have

$$\begin{aligned} |\beta_{lk}^B| &= \left| \bar{C} \left\{ \int_{L_{lk}} \frac{f(x)}{|L_{lk}|} dx - \int_{R_{lk}} \frac{f(x)}{|R_{lk}|} dx \right\} \right| \\ &\leq \frac{\bar{C}}{|L_{lk}|} \int_0^{|L_{lk}|} \left| f(x + l_{lk}) - f\left(b_{lk} + x \frac{|R_{lk}|}{|L_{lk}|}\right) \right| dx. \end{aligned}$$

Next, from α -Hölder continuity (24), we have

$$\left| f(x + l_{lk}) - f\left(b_{lk} + x \frac{|R_{lk}|}{|L_{lk}|}\right) \right| \leq M \left| |L_{lk}| + x \left(\frac{|R_{lk}|}{|L_{lk}|} - 1 \right) \right|^\alpha.$$

Since $\bar{C} \leq 2^{-1/2} \max\{|L_{lk}|, |R_{lk}|\}^{1/2}$ we have

$$|\beta_{lk}^B| \leq M \bar{C} \max\{|L_{lk}|, |R_{lk}|\}^\alpha \leq M 2^{-1/2} \max\{|L_{lk}|, |R_{lk}|\}^{\alpha+1/2}.$$

This follows from the fact that for $x \in (0, |L_{lk}|)$ we have

$$M \left| |L_{lk}| + x \left(\frac{|R_{lk}|}{|L_{lk}|} - 1 \right) \right|^\alpha \leq \max\{|L_{lk}|, |R_{lk}|\}^\alpha. \quad \square$$

For the classical Haar basis (3), one obtains (23) from (83) by noting $\max\{|L_{lk}|, |R_{lk}|\} = 2^{-(l+1)}$. [27] points out that the computational complexity of the discrete UH transform could be unnecessarily large and imposes the balancing requirement $\max\{|L_{lk}|, |R_{lk}|\} \leq E(|L_{lk}| + |R_{lk}|) \forall (l, k) \in A$, for some $1/2 \leq E < 1$. In order to control the combinatorial complexity of the basis system, we also require that the UH wavelets are not too imbalanced. To this end, we introduce the notion of *weakly balanced* wavelets.

EXAMPLE 2. To glean insights into the balancing condition (38), we first consider an example of UH system which is not weakly balanced for some given n, D , say $n = 2^4$ and $D = 2$. If we choose $b_{00} = 1/2$, $b_{10} = 1/2 - 1/n$ and $b_{11} = 3/4$, we have

$$\begin{aligned} L_{10} &= (0, 1/2 - 1/n], & R_{10} &= (1/2 - 1/n, 1/2], \\ L_{11} &= (1/2, 3/4], & R_{11} &= (3/4, 1]. \end{aligned}$$

While the node $(1, 1)$ is seen to satisfy (38) with $E = 5$, we note that $\max\{|L_{10}|, |R_{10}|\} = (n - 2)/(2n) = 7/16$. However, $7/16$ cannot be written as $M_{10}/2^{D+1} = M_{10}/8$ for any integer M_{10} . This is why the split points b_{00}, b_{10} and b_{11} do not belong to any weakly balanced UH wavelet system with balancing constant $D = 2$. Weakly balanced systems can be built by choosing splits in such a way that the “granularity” does not increase too rapidly throughout the branching process. With granularity $R(l, \Psi_A^B)$ of the l^{th} layer we mean the smallest integer $R \geq 1$ such that $\min_{0 \leq k < 2^l} \min\{|L_{lk}|, |R_{lk}|\} = j/2^R$ for some $j \in \{1, 2, \dots, 2^{R-1}\}$. For instance, setting $D = 2$ and $E = 3$ one can build weakly balanced wavelets by first picking b_{00} from values $\{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$. If, e.g. $b_{00} = 3/4$ (i.e. $R(0, \Psi_A^B) = 2$), the next split b_{10} can be selected from $\{\frac{1}{4}, \frac{3}{8}, \frac{1}{2}\}$, while b_{11} has to be set as $7/8$.

Our theoretical development relies in part on combinatorial properties of weakly balanced UH systems and on the speed of decay of the multiscale functionals $\beta_{lk}^B = \langle f, \psi_{lk}^B \rangle$ as the layer index $l \in \mathbb{N}$ increases. These two fundamental properties are encapsulated in Lemma 11 which is vital to the proof of the upcoming Theorem 4.

Note that in the actual BART implementation, the splits are chosen from sample quantiles to ensure balancedness (similar to our condition (38)). Quantile splits are a natural way to generate many weakly balanced systems, providing a much increased flexibility compared to dyadic splits, which correspond to uniform quantiles.

EXAMPLE 3 (Quantile Splits). Denote with G a c.d.f with a density g on $[0, 1]$ that satisfies $\|g\|_\infty \leq 2^{D-1}/(2E)$ for $E, D > 0$ chosen below and $\|1/g\|_\infty \leq C_q$ for some $C_q > 0$. Let us define a dyadic projection of G as

$$G_l^{-1}(x) := 2^{-l} \lfloor 2^l G^{-1}(x) \rfloor,$$

and next define the breakpoints, for $l \leq L_{max}$ and $0 \leq k < 2^l$, as

$$(84) \quad b_{lk} = G_{L_{max}+D}^{-1}[(2k+1)/2^{l+1}].$$

The system Ψ_A^B obtained from steps (a) and (b) with splits (84) is weakly balanced for $E = 2 + 3C_q 2^{D-1}$. This is verified in Lemma 12 in the Appendix (Section 9.4). Moreover, Figure 8 in Section 9.4 illustrates the implementation of the quantile system, where splits are placed more densely in areas where $G(x)$ changes more rapidly.

9.2. Granularity lemma.

LEMMA 10. Denote with Ψ_A^B a weakly balanced UH system according to Definition 5. Then for any $(l, k) \in A$,

$$\min\{|L_{lk}|, |R_{lk}|\} = \frac{m_{lk}}{2^{l+D}} \quad \text{for some } m_{lk} \in \{1, \dots, C + l\}.$$

PROOF. We prove the statement by induction. First, from the definition of weak balancedness, we have $\min\{|L_{00}|, |R_{00}|\} = 1 - M_{00}/2^D = j/2^D$ (for $j = 2^D - M_{00}$) and by definition this is less than $M_{00}/2^D \leq C/2^D$, so $j \leq C$. Assume now that the statement holds for

$l - 1 \geq 0$ and consider a node $(l, k) \in A$ for some $0 \leq k < 2^l$. The union $L_{lk} \cup R_{lk}$ is either $L_{l-1 \lfloor k/2 \rfloor}$ or $R_{l-1 \lfloor k/2 \rfloor}$; without loss of generality, suppose it is $R_{l-1 \lfloor k/2 \rfloor}$. Then, from weak balancedness, we find

$$(85) \quad \min\{|L_{lk}|, |R_{lk}|\} = |R_{l-1 \lfloor k/2 \rfloor}| - M_{lk}/2^{l+D}.$$

If $|R_{l-1 \lfloor k/2 \rfloor}| \leq |L_{l-1 \lfloor k/2 \rfloor}|$, we use induction to find $|R_{l-1 \lfloor k/2 \rfloor}| = j_1/2^{l-1+D}$ for some $j_1 \in \{1, \dots, C + l - 1\}$ and thereby (85) equals $j/2^{l+D}$ for $j = 2j_1 - M_{lk}$. As this is at most $M_{lk}/2^{l+D} = \max\{|L_{lk}|, |R_{lk}|\}$, one deduces $M_{lk} \geq j_1$ and then $j/2^{l+D} \leq j_1/2^{l+D}$ with $j_1 \leq C + l - 1 \leq C + l$. If $|R_{l-1 \lfloor k/2 \rfloor}| > |L_{l-1 \lfloor k/2 \rfloor}|$, we again use weak balancedness to write (85) as $j/2^{l+D}$ with $j = 2M_{l-1 \lfloor k/2 \rfloor} - M_{lk} \leq M_{lk}$, using again $M_{lk}/2^{l+D} = \max\{|L_{lk}|, |R_{lk}|\}$, so that j is again less than $C + l$. The result follows by induction. \square

9.3. Complexity lemma.

LEMMA 11. Consider a weakly balanced UH wavelet system $\Psi_A^B = \{\Psi_{lk}^B : (l, k) \in A\}$ according to (38) and let $f \in \mathcal{H}_M^\alpha$. Then the following conditions hold for $\delta = 3$, with constants independent of B : for any $(l, k) \in A$

- (B1) the basis function ψ_{lk}^B can be expressed as a linear combination of at most $C_0 l^\delta$ Haar functions ψ_{jk} for $j \leq l + D$ and some $C_0 > 0$, and
- (B2) there exists $C_1 > 0$ (depending only on E, D from (38)) such that $|\beta_{lk}^B| \leq C_1 M l^{\delta/2} 2^{-(\alpha+1/2)}$.

PROOF. First, the function ψ_{lk}^B belongs to $\text{Vect}\{\mathbb{I}_{I_{(l+D)m}} : 0 \leq m < 2^{l+D}\}$ and the support of ψ_{lk}^B spans at most $2(E + l)$ of the indicators $\mathbb{I}_{I_{(l+D)m}}$. These indicators can be expressed in terms of at most $l + D$ of ψ_{lk} 's (one per level above $l + D$), which yields an upper bound $2(E + l)(l + D) \asymp l^2$ and thereby (B1) with $\delta = 2$. Second, the balancing condition (38) gives $\max\{|L_{lk}|, |R_{lk}|\} \leq (E + l)2^{-l-D}$ which, combined with Lemma 9 implies

$$|\beta_{lk}^B| \leq M 2^{\alpha-1} (E + l)^{\alpha+1/2} 2^{-(l+D)(\alpha+1/2)} \leq C_1 M l^{3/2} 2^{-l(\alpha+1/2)},$$

by taking the worst case $\alpha = 1$, which proves (B2) with $\delta = 3$. \square

9.4. The Quantile Example.

LEMMA 12. The quantile system Ψ_A^B from Example 3 is weakly balanced in the sense of Definition 5 for balancing constants satisfying $E = 2 + 3C_q 2^{D-1}$, where $\|1/g\|_\infty \leq C_q$ and $\|g\|_\infty < 2^{D-1}/(2E)$.

PROOF. Let us start by writing explicitly the intervals L_{lk}, R_{lk} . Assuming without loss of generality that k is odd, i.e. (l, k) is the right child node,

$$\begin{aligned} |L_{lk}| &= b_{lk} - b_{(l-1) \lfloor k/2 \rfloor} = G_{L_{max}+D}^{-1}[(2k+1)/2^{l+1}] - G_{L_{max}+D}^{-1}[(2 \lfloor k/2 \rfloor + 1)/2^l], \\ |R_{lk}| &= b_{(l-2) \lfloor k/4 \rfloor} - b_{lk} = G_{L_{max}+D}^{-1}[(2 \lfloor k/4 \rfloor + 1)/2^{l-1}] - G_{L_{max}+D}^{-1}[(2k+1)/2^{l+1}]. \end{aligned}$$

We first show by contradiction that $\max\{|L_{lk}|, |R_{lk}|\} \geq E/2^{l+D}$ for $E \geq 1$. Let us denote $y_1 = G^{-1}[(2k+1)/2^{l+1}]$, $y_2 = G^{-1}[(2 \lfloor k/2 \rfloor + 1)/2^l]$ and $y_3 = G^{-1}[(2 \lfloor k/4 \rfloor + 1)/2^{l-1}]$. Assuming $|L_{lk}| < E/2^{l+D}$, one obtains

$$\lfloor 2^{L_{max}+D} y_1 \rfloor - \lfloor 2^{L_{max}+D} y_2 \rfloor < E 2^{L_{max}-l},$$

and thereby $y_1 - y_2 < E2^{-l-D+1}$. Next, using the fact that k is odd,

$$\begin{aligned} \frac{1}{2^{l+1}} &= |(2k+1)/2^{l+1} - (2\lfloor k/2 \rfloor + 1)/2^l| \\ &= |G(y_1) - G(y_2)| \leq \|g\|_\infty |y_1 - y_2| \leq \|g\|_\infty 2E2^{-l-D}, \end{aligned}$$

which yields a contradiction when $\|g\|_\infty < 2^{D-1}/(2E)$. Similarly, when $|R_{lk}| < E/2^{l+D}$, one obtains

$$\begin{aligned} 1/2^{l+1} &< |(2k+1)/2^{l+1} - (2\lfloor k/4 \rfloor + 1)/2^{l-1}| \\ &= |G(y_3) - G(y_1)| \leq \|g\|_\infty |y_3 - y_1| < \|g\|_\infty 2E3 \cdot 2^{-l-D}. \end{aligned}$$

Next, we note that for $\|1/g\|_\infty \leq C_q$ and $E := (2 + 3C_q2^{D-1})$,

$$\begin{aligned} |R_{lk}| &= \frac{1}{2^{L_{max}+D}} [\lfloor 2^{L_{max}+D} y_1 \rfloor - \lfloor 2^{L_{max}+D} y_3 \rfloor] \\ &\leq \frac{2}{2^{L_{max}+D}} + \left\| \frac{1}{g} \right\|_\infty \frac{3}{2^{l+1}} \leq \frac{E}{2^{l+D}}. \end{aligned}$$

Similarly, one obtains $|L_{lk}| < E/2^{l+D}$, which concludes that the quantile system is weakly balanced. \square




FIG 9. Example of quantile splits for a uniform density $g(x)$ and a non-uniform beta density $g(x)$ using $L_{max} = 6$.

10. Remaining proofs for Section 3.

10.1. *Proof of Theorem 1: remaining settings.* In this section, we provide the proof of Theorem 1 under more general prior settings. We consider different tree priors $\Pi_{\mathbb{T}}(\mathcal{T})$ from Section 2.1 (i.e. the conditionally uniform prior in (5) and the exponential-type prior mentioned in Extension 2). In addition, we consider the prior

$$(86) \quad \pi(\beta_{\mathcal{T}}) \sim \mathcal{N}(0, \Sigma_{\mathcal{T}})$$

with $\Sigma_{\mathcal{T}} = g_n(A'_{\mathcal{T}}A_{\mathcal{T}})^{-1}$ (the g -prior) as well as more general covariance structures $\Sigma_{\mathcal{T}}$ that satisfy

$$(87) \quad \lambda_{\min}(\Sigma_{\mathcal{T}}) \gtrsim 1/\sqrt{\log n} \quad \text{and} \quad \lambda_{\max}(\Sigma_{\mathcal{T}}) \lesssim n^a \quad \text{for some } a > 0.$$

First, we show that Lemma 1 holds under this setting.

10.1.1. *Proof of Lemma 1 (general setting).* Recall the ratio (45) from Section 6.1

$$(88) \quad \frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^-)} = \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^-)} \frac{\sqrt{|I + n\Sigma_{\mathcal{T}^-}|}}{\sqrt{|I + n\Sigma_{\mathcal{T}}|}} \frac{e^{n^2 \mathbf{X}'_{\mathcal{T}}(nI + \Sigma_{\mathcal{T}^-})^{-1} \mathbf{X}_{\mathcal{T}}/2}}{e^{n^2 \mathbf{X}'_{\mathcal{T}^-}(nI + \Sigma_{\mathcal{T}^-})^{-1} \mathbf{X}_{\mathcal{T}^-}/2}}.$$

The g -prior. We first focus on bounding this ratio assuming $\Sigma_{\mathcal{T}} = g_n(A'_{\mathcal{T}}A_{\mathcal{T}})^{-1}$. Proposition 2 implies

$$(89) \quad nI_K + \Sigma_{\mathcal{T}}^{-1} = \begin{pmatrix} nI_{K-1} + \Sigma_{\mathcal{T}^-}^{-1} + \frac{1}{g_n} \mathbf{v}\mathbf{v}' & \mathbf{0} \\ \mathbf{0}' & n + \frac{2^{l_1+1}}{g_n} \end{pmatrix},$$

where the vector $\mathbf{v} \in \mathbb{R}^{|\mathcal{T}_{ext}|-1}$ (defined in Proposition 2) is obtained from $A_{\mathcal{T}}$ by first deleting its last column and then transposing the last row of this reduced matrix. Using the determinant formula $|A + \mathbf{u}\mathbf{u}'| = |A|(1 + \mathbf{u}'A^{-1}\mathbf{u})$ for A invertible (see Lemma 1.1 in [?]), and setting $M = (nI_{K-1} + \Sigma_{\mathcal{T}^-}^{-1})^{-1}$, one gets

$$(90) \quad \frac{|nI_{K-1} + \Sigma_{\mathcal{T}^-}^{-1}|}{|nI_K + \Sigma_{\mathcal{T}}^{-1}|} = \frac{1}{(n + 2^{l_1+1}/g_n)[1 + \mathbf{v}'M\mathbf{v}/g_n]} < \frac{1}{(n + 2^{l_1+1}/g_n)}.$$

Using again the determinant formula and Proposition 2, we can write

$$\frac{|\Sigma_{\mathcal{T}^-}|}{|\Sigma_{\mathcal{T}}|} = \frac{|A'_{\mathcal{T}}A_{\mathcal{T}}|}{|A'_{\mathcal{T}^-}A_{\mathcal{T}^-}|} = \frac{2^{l_1+1}}{g_n} (1 + \mathbf{v}'(A'_{\mathcal{T}^-}A_{\mathcal{T}^-})^{-1}\mathbf{v}) \leq \frac{2^{l_1+1}}{g_n} (1 + 2^{l_1}).$$

The inequality above uses two facts. First, we have $\|\mathbf{v}\|_2^2 = 1 + \sum_{l=0}^{l_1-1} 2^l = 2^{l_1}$ which follows from the definition of \mathbf{v} which describes its entries as amplitudes of the Haar wavelets in the ancestry of the deepest rightmost internal node (l_1, k_1) (using the notation from Lemma 1). Second, from Proposition 3 we have $\lambda_{\max}[(A'_{\mathcal{T}^-}A_{\mathcal{T}^-})^{-1}] = 1/\lambda_{\min}(A'_{\mathcal{T}^-}A_{\mathcal{T}^-}) \leq 1$. Indeed, Proposition 3 (combined with (15)) shows that the smallest eigenvalue of $A'_{\mathcal{T}}A_{\mathcal{T}}$ equals 2^d where d is the depth of the most shallow leaf node in \mathcal{T}_{ext} . Let us now set

$$(91) \quad D := \mathbf{X}'_{\mathcal{T}}(nI + \Sigma_{\mathcal{T}}^{-1})^{-1} \mathbf{X}_{\mathcal{T}} - \mathbf{X}'_{\mathcal{T}^-}(nI + \Sigma_{\mathcal{T}^-}^{-1})^{-1} \mathbf{X}_{\mathcal{T}^-}.$$

Combining with (89), it follows from a variant of the Sherman–Morrison’s matrix inversion formula (Lemma 5) that

$$(nI_K + \Sigma_{\mathcal{T}}^{-1})^{-1} = \begin{pmatrix} M - \frac{M\mathbf{v}\mathbf{v}'M}{g_n + \mathbf{v}'M\mathbf{v}} & \mathbf{0} \\ \mathbf{0}' & 1/(n + 2^{l_1+1}/g_n) \end{pmatrix},$$

from which one deduces that

$$(92) \quad D = \frac{X_{l_1 k_1}^2}{n + 2^{l_1+1}/g_n} - \frac{\mathbf{X}'_{\mathcal{T}^-}M\mathbf{v}\mathbf{v}'M\mathbf{X}_{\mathcal{T}^-}}{g_n + \mathbf{v}'M\mathbf{v}} < \frac{X_{l_1 k_1}^2}{n + 2^{l_1+1}/g_n}.$$

Since for $l_1 > \mathcal{L}_c$ we have $2X_{l_1 k_1}^2 \leq 5 \log n/n$, we can write

$$\frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^-)} \frac{\Pi_{\mathbb{T}}(\mathcal{T}^-)}{\Pi_{\mathbb{T}}(\mathcal{T})} < \sqrt{\frac{2^{l_1+1}(1 + 2^{l_1})}{g_n(n + 2^{l_1+1}/g_n)}} e^{\frac{X_{l_1 k_1}^2 n^2}{2(n + 2^{l_1+1}/g_n)}} < \sqrt{\frac{2^{2(l_1+1)}}{ng_n}} n^{5/4}.$$

For $g_n = n$, and for $\mathcal{T} \in \mathbb{T}_d$, so that $2^{l_1} \lesssim 2^d$, the last display is bounded by a constant times $n^{-1/4} 2^d p_d$, and the argument can be completed in similar vein as before, with now $\Pi[d(\mathcal{T}) > \mathcal{L}_c | X] = o_P(1)$ if $\Gamma > 2e^5$.

General Covariance $\Sigma_{\mathcal{T}}$. We now show how the proof can be modified by assuming a general covariance matrix $\Sigma_{\mathcal{T}}$ on the internal wavelet coefficients. Recall again the ratio (45) from Section 6.1. We use the Cauchy’s interlace theorem for eigenvalues of Hermitian matrices which states that the eigenvalues of a principal submatrix are interlaced within eigenvalues of the original matrix (Theorem 8.1.7 of [34]). Since $\Sigma_{\mathcal{T}^-}$ is a $(K-1) \times (K-1)$ dimensional submatrix of a $(K \times K)$ dimensional matrix $\Sigma_{\mathcal{T}}$, we have (denoting with $\lambda_l(\Sigma)$ the l^{th} largest eigenvalue of Σ)

$$\lambda_1(\Sigma_{\mathcal{T}}) \geq \lambda_1(\Sigma_{\mathcal{T}^-}) \geq \lambda_2(\Sigma_{\mathcal{T}}) \geq \cdots \geq \lambda_{K-1}(\Sigma_{\mathcal{T}}) \geq \lambda_{K-1}(\Sigma_{\mathcal{T}^-}) \geq \lambda_K(\Sigma_{\mathcal{T}})$$

and thereby

$$\frac{|I + n\Sigma_{\mathcal{T}^-}|}{|I + n\Sigma_{\mathcal{T}}|} = \frac{\prod_{l=1}^{K-1} [1 + n\lambda_l(\Sigma_{\mathcal{T}^-})]}{\prod_{l=1}^K [1 + n\lambda_l(\Sigma_{\mathcal{T}})]} \leq \frac{1}{1 + n\lambda_{\min}(\Sigma_{\mathcal{T}})}.$$

Using the matrix inversion formula $(I + B)^{-1} = I - (I + B^{-1})^{-1}$ (a variant of Sherman-Morrison-Woodbury formula), we get

$$(nI + \Sigma_{\mathcal{T}}^{-1})^{-1} = \frac{1}{n} [I - (I + n\Sigma_{\mathcal{T}})^{-1}]$$

and thereby

$$\mathbf{X}'_{\mathcal{T}}(nI + \Sigma_{\mathcal{T}}^{-1})^{-1} \mathbf{X}_{\mathcal{T}} = \frac{1}{n} \|\mathbf{X}_{\mathcal{T}}\|_2^2 - \frac{1}{n} \mathbf{X}'_{\mathcal{T}}(I + n\Sigma_{\mathcal{T}})^{-1} \mathbf{X}_{\mathcal{T}}.$$

Writing $\mathbf{X}_{\mathcal{T}} = (\mathbf{X}_{\mathcal{T}^-}, X_{l_1 k_1})'$, where (l_1, k_1) is the deepest rightmost internal node in \mathcal{T} (as in Section 6.1), and using the definition of D in (91), we have

$$\begin{aligned} D &= \frac{X_{l_1 k_1}^2}{n} - \frac{1}{n} [\mathbf{X}'_{\mathcal{T}}(I + n\Sigma_{\mathcal{T}})^{-1} \mathbf{X}_{\mathcal{T}} - \mathbf{X}'_{\mathcal{T}^-}(I + n\Sigma_{\mathcal{T}^-})^{-1} \mathbf{X}_{\mathcal{T}^-}] \\ &< \frac{X_{l_1 k_1}^2}{n} \left(1 - \frac{1}{1 + n\lambda_{\max}(\Sigma_{\mathcal{T}})} \right) \\ &\quad + \frac{\|\mathbf{X}_{\mathcal{T}^-}\|_2^2}{n} \left(\frac{1}{1 + n\lambda_{\min}(\Sigma_{\mathcal{T}^-})} - \frac{1}{1 + n\lambda_{\max}(\Sigma_{\mathcal{T}})} \right). \end{aligned}$$

This inequality follows from the fact that

$$\mathbf{X}'_{\mathcal{T}}(I + n\Sigma_{\mathcal{T}})^{-1} \mathbf{X}_{\mathcal{T}} \geq \|\mathbf{X}_{\mathcal{T}}\|_2^2 \lambda_{\min}[(I + n\Sigma_{\mathcal{T}})^{-1}] = \frac{\|\mathbf{X}_{\mathcal{T}^-}\|_2^2 + X_{l_1 k_1}^2}{1 + n\lambda_{\max}(\Sigma_{\mathcal{T}})}.$$

It follows from the proof of Lemma 1 that $X_{l_1 k_1}^2 \lesssim \log n/n$ and $\|\mathbf{X}_{\mathcal{T}^-}\|_2^2 \leq C_1$ (as was shown in (94)). Moreover, from our assumption (86) we have $\lambda_{\min}(\Sigma_{\mathcal{T}^-}) \geq 1/\sqrt{\log n}$ and thereby

$$\begin{aligned} \frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^-)} &< \sqrt{\frac{\log n}{n}} \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^-)} \exp \left\{ \frac{nX_{l_1 k_1}^2}{2} + \frac{n\|\mathbf{X}_{\mathcal{T}^-}\|_2^2}{2(1 + n\lambda_{\min}(\Sigma_{\mathcal{T}^-}))} \right\} \\ &< \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^-)} e^{C_2 \log n}. \end{aligned}$$

Proceeding as in the proof of Lemma 1, one can show (43) for a suitably large $\Gamma > 0$.

Other Tree Priors $\Pi_{\mathbb{T}}(\mathcal{T})$. The only modification needed to carry over the proof to the other two priors is the bound for the ratio $\Pi_{\mathbb{T}}(\mathcal{T})/\Pi_{\mathbb{T}}(\mathcal{T}^-)$. Consider the conditionally uniform prior from Section 2.1.1 defined in (5). Denoting $K = |\mathcal{T}_{ext}|$ and \mathbb{C}_K the number of full binary trees with $K + 1$ leaves, we have

$$\frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^-)} = \frac{\pi(K)}{\pi(K-1)} \frac{\mathbb{C}_{K-1}}{\mathbb{C}_{K-2}},$$

and Lemma 6 now implies, for a universal constant $C > 0$,

$$\frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^-)} \lesssim \frac{\lambda}{K} \frac{4^{K-1}(K-2)^{3/2}}{4^{K-2}(\{K-1\} \vee 1)^{3/2}} \leq C\lambda/K.$$

Choosing $\lambda = 1/n^c$ for some $c > 7/4$, it follows that

$$\Pi[d(\mathcal{T}) > \mathcal{L}_c | X] \leq 4\lambda n^{3/4} \sum_{d=\mathcal{L}_c+1}^L 2^d/K \leq 4\lambda n^{3/4} 2^L \rightarrow 0.$$

Finally, for the exponential-type prior mentioned in Extension 2), one has $\Pi_{\mathbb{T}}(\mathcal{T})/\Pi_{\mathbb{T}}(\mathcal{T}^-) = 1/n^c$, so one can argue similarly.

10.1.2. *Proof of Lemma 2 (general setting).* We recall the ratio

$$(93) \quad \frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^+)} = \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^+)} \prod_{j=0}^s \frac{N_X(\mathcal{T}^{(j-1)})}{N_X(\mathcal{T}^{(j)})}.$$

in (50) and find an upper bound assuming different priors.

The g-prior. We now modify the proof of Lemma 2 for the g -prior obtained with $\Sigma_{\mathcal{T}} = g_n(A'_{\mathcal{T}}A_{\mathcal{T}})^{-1}$. Denoting with $K_j = |\mathcal{T}_{ext}^{(j)}|$ and because $\mathcal{T}^{(j-1)}$ is obtained from $\mathcal{T}^{(j)}$ by removing two children nodes (or, equivalently, an internal node $X_{l_j k_j}$), we can apply Proposition 2 to obtain the following upper bound for $N_X(\mathcal{T}^{(j-1)})/N_X(\mathcal{T}^{(j)})$. Namely, going back to (91) and (92), we invoke again the matrix determinant lemma $|A + \mathbf{u}\mathbf{u}'| = |A|(1 + \mathbf{u}'A^{-1}\mathbf{u})$ and the matrix inversion lemma (Lemma 5) as in (90) to obtain

$$\begin{aligned} \frac{N_X(\mathcal{T}^{(j-1)})}{N_X(\mathcal{T}^{(j)})} &\leq \sqrt{\frac{(g_n n + 2^{l_j+1})(g_n + \mathbf{v}'M\mathbf{v})}{g_n^2}} \\ &\quad \times \exp \left\{ -\frac{n^2 X_{l_j k_j}^2}{2(n + 2^{l_j+1}/g_n)} + n^2 \frac{\mathbf{X}'_{\mathcal{T}^{(j-1)}} M \mathbf{v} \mathbf{v}' M \mathbf{X}_{\mathcal{T}^{(j-1)}}}{2(g_n + \mathbf{v}'M\mathbf{v})} \right\}, \end{aligned}$$

where $M = (nI_{K_{j-1}} + \Sigma_{\mathcal{T}^{(j-1)}}^{-1})^{-1}$ for $\mathbf{v} \in \mathbb{R}^{|\mathcal{T}^{(j-1)}|}$ which depends on $\mathcal{T}^{(j-1)}$ according to Proposition 2. Next, for $C > 0$ a large enough constant and l_{j-1} the depth of the deepest internal node in $\mathcal{T}^{(j-1)}$,

$$(94) \quad \|\mathbf{X}_{\mathcal{T}^{(j-1)}}\|_2^2 \leq X_{-10}^2 + \sum_{l=0}^{l_{j-1}} \sum_{k=0}^{2^l-1} X_{lk}^2 \leq C \left[1 + \sum_{l=0}^{l_{j-1}} \left(2^{-2l\alpha} + \frac{2^l}{n} \log n \right) \right] \leq C_4.$$

Moreover, using the fact that $\mathbf{v}'M\mathbf{v} = \mathbf{v}'M^{1/2}M^{1/2}\mathbf{v} \leq \lambda_{\max}(M)\mathbf{v}'M\mathbf{v}$ we obtain

$$\begin{aligned} \mathbf{X}'_{\mathcal{T}^{(j-1)}} M \mathbf{v} \mathbf{v}' M \mathbf{X}_{\mathcal{T}^{(j-1)}} &\leq \|\mathbf{X}_{\mathcal{T}^{(j-1)}}\|_2^2 \lambda_{\max}(M \mathbf{v} \mathbf{v}' M) \leq \|\mathbf{X}_{\mathcal{T}^{(j-1)}}\|_2^2 \text{tr}(M \mathbf{v} \mathbf{v}' M) \\ &\leq \|\mathbf{X}_{\mathcal{T}^{(j-1)}}\|_2^2 \mathbf{v}' M M \mathbf{v} \leq \|\mathbf{X}_{\mathcal{T}^{(j-1)}}\|_2^2 \lambda_{\max}(M) \mathbf{v}' M \mathbf{v}, \end{aligned}$$

and

$$\lambda_{\max}(M) = \frac{1}{n + \lambda_{\min}(A_{\mathcal{T}^{(j-1)'}} A_{\mathcal{T}^{(j-1)}})/g_n} < \frac{1}{n},$$

one can write

$$\begin{aligned} n^2 \frac{\mathbf{X}'_{\mathcal{T}^{(j-1)}} M \mathbf{v} \mathbf{v}' M \mathbf{X}_{\mathcal{T}^{(j-1)}}}{2(g_n + \mathbf{v}' M \mathbf{v})} &< \frac{n^2 \|\mathbf{X}_{\mathcal{T}^{(j-1)}}\|_2^2 \lambda_{\max}(M)}{2g_n/(\mathbf{v}' M \mathbf{v}) + 2} \leq \frac{n^2 C_4 \|\mathbf{v}\|_2^2 \lambda_{\max}^2(M)}{2g_n} \\ &\leq C_4 \frac{2^{l_j}}{2g_n}, \end{aligned}$$

where we used the fact that $\|\mathbf{v}\|_2^2 = 2^{l_j}$ (as explained previously in Section 10.1.1). Finally, because $X_{l_s k_s}^2 \geq C_5 A^2 \log^2 n/n$ for some $C_5 > 0$ we use the above bounds in the expression (93) to obtain

$$\begin{aligned} \prod_{j=0}^s \frac{N_X(\mathcal{T}^{(j-1)})}{N_X(\mathcal{T}^{(j)})} &< \left(\frac{n(g_n + 2)}{g_n} \right)^{s+1} \exp \left\{ C_4 \sum_{j=0}^s \frac{2^{l_j-1}}{g_n} - \frac{n C_5 A^2 \log^2 n}{2(n + 2^{l_j+1}/g_n)} \right\} \\ &< \exp \left\{ (s+1) \log(3n) + C_4 \frac{(s+1) 2^{l_s-1}}{g_n} - C_5 A^2 \log^2 n/2 \right\}. \end{aligned}$$

With $g_n = n$, the exponent is dominated by the last term. One then proceeds with (50) as before in the proof of Lemma 2.

General Covariance $\Sigma_{\mathcal{T}}$. We again deploy the interlacing eigenvalue theorem in the expression (50) to obtain the following upper bound for $\frac{N_X(\mathcal{T}^{(j-1)})}{N_X(\mathcal{T}^{(j)})}$ (using matrix determinant and inversion lemmata as before)

$$\sqrt{1 + n \lambda_{\max}(\Sigma_{\mathcal{T}^{(j)}})} \exp \left\{ -\frac{n X_{l_j k_j}^2}{2} \frac{n \lambda_{\min}(\Sigma_{\mathcal{T}^{(j)}})}{1 + n \lambda_{\min}(\Sigma_{\mathcal{T}^{(j)}})} + \frac{n \|\mathbf{X}_{\mathcal{T}^{(j-1)}}\|_2^2}{2(1 + n \lambda_{\min}(\Sigma_{\mathcal{T}^{(j)}}))} \right\}.$$

Using the expression (50) and assumptions $\lambda_{\max}(\Sigma_{\mathcal{T}}) \lesssim n^a$ for some $a \geq 1$ and $\lambda_{\min}(\Sigma_{\mathcal{T}^{(j)}}) \geq 1/\sqrt{\log n}$, we obtain for $C_2, C_3 > 0$

$$\frac{W_X(\mathcal{T})}{W_X(\mathcal{T}^+)} < \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^+)} \exp \left\{ C_2(s+1) \sqrt{\log n} - C_3 A^2 \log^2 n \right\}.$$

Using this bound, one can proceed as in the proof of Lemma 2 and show (113).

Other Tree Priors $\Pi_{\mathbb{T}}(\mathcal{T})$. As before, the only modification needed is the bound for $\Pi_{\mathbb{T}}(\mathcal{T})/\Pi_{\mathbb{T}}(\mathcal{T}^+)$. Denote by $K^+ = |\mathcal{T}_{ext}^+|$ and $K = |\mathcal{T}_{ext}|$ and note that $K^+ - K = l_s - l_0$. For the conditionally uniform prior from Section 2.1.1 we then have

$$\begin{aligned} \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}^+)} &= \lambda^{-(l_s-l_0)} \frac{K^+! \mathbb{C}_{K^+-1}}{K! \mathbb{C}_{K-1}} \lesssim \left(\frac{\lambda}{4} \right)^{-(l_s-l_0)} \frac{(K^+)!}{K!} \\ &\lesssim \left(\frac{\lambda}{4} \right)^{-(l_s-l_0)} e^{(l_s-l_0)\{1+\log[K+(l_s-l_0)]\}}. \end{aligned}$$

With $\lambda = n^{-c}$ and $K \leq 2\mathcal{L}^{c+1}$, this is bounded from above by $C e^{C \log^2 n}$ for some $C > 0$ and the proof is completed as before. For the exponential-type prior, one similarly uses $\Pi_{\mathbb{T}}(\mathcal{T})/\Pi_{\mathbb{T}}(\mathcal{T}^+) = e^{c(l_s-l_0) \log n} \leq e^{c \log^2 n}$.

10.1.3. *Proof of Lemma 3 (general setting).* We now show the proof of Lemma 3 for the g -prior with $\Sigma_{\mathcal{T}} = g_n(A'_{\mathcal{T}}A_{\mathcal{T}})^{-1}$ and a general covariance matrix under the assumptions (87). Recall expressions (55) and (56) in Section 6.3 and the fact that $(I + B)^{-1} = I - (I + B^{-1})^{-1}$ which yields $n\tilde{\Sigma}_{\mathcal{T}} - I_K = -(I_K + n\Sigma_{\mathcal{T}})^{-1}$ for $\tilde{\Sigma}_{\mathcal{T}} = (nI_K + \Sigma_{\mathcal{T}}^{-1})^{-1}$ (recalling that $K = |\mathcal{T}_{ext}| \leq n$). Focusing on the g -prior, we have $\lambda_{\max}(\tilde{\Sigma}_{\mathcal{T}}) < 1/n$ and (using Proposition 3 which yields $\lambda_{\max}(A'_{\mathcal{T}}A_{\mathcal{T}}) < 2^d \leq n$ where d is the depth of the deepest node)

$$\lambda_{\max}(I_K + n\Sigma_{\mathcal{T}})^{-1} < \lambda_{\max}(A'_{\mathcal{T}}A_{\mathcal{T}})/(ng_n) < 1/g_n.$$

Assuming $g_n = n$ we can thus write

$$(95) \quad \|(n\tilde{\Sigma}_{\mathcal{T}} - I_K)\beta_{\mathcal{T}}^0\|_{\infty} \leq \frac{\|\beta_{\mathcal{T}}^0\|_{\infty}\sqrt{K}}{1 + n\lambda_{\min}(\Sigma_{\mathcal{T}})} \leq \frac{C\sqrt{K}\lambda_{\max}(A'_{\mathcal{T}}A_{\mathcal{T}})}{ng_n} \leq C/\sqrt{n}.$$

Next, we note that $\tilde{\Sigma}_{\mathcal{T}}^{-1}$ is strictly diagonally dominant. Indeed, with $g_n = n$ we have $\tilde{\Sigma}_{\mathcal{T}}^{-1} = nI_K + \frac{1}{g_n}A'_{\mathcal{T}}A_{\mathcal{T}}$ and

$$\frac{1}{g_n}\|A'_{\mathcal{T}}A_{\mathcal{T}}\|_{\infty} \leq \frac{\sqrt{K}}{g_n}\lambda_{\max}(A'_{\mathcal{T}}A_{\mathcal{T}}) < \sqrt{n},$$

where $\|A\|_{\infty} = \max_{1 \leq i \leq m} \sum_{j=1}^n |a_{ij}|$ is defined as the maximum absolute row sum of an $(m \times n)$ matrix A . Writing $A'_{\mathcal{T}}A_{\mathcal{T}} = (a_{ij})_{i,j}^{K,K}$, it then follows from Lemma 4 (Theorem 1 in [55]) that

$$(96) \quad \|\tilde{\Sigma}_{\mathcal{T}}\|_{\infty} \leq \frac{1}{n + \frac{1}{g_n} \min_{1 \leq k \leq K} \Delta_k}, \quad \text{where} \quad \Delta_k = |a_{kk}| - \sum_{j \neq k} |a_{kj}|.$$

Since $\Delta_k/g_n > -\frac{1}{g_n}\|A'_{\mathcal{T}}A_{\mathcal{T}}\|_{\infty} > -\sqrt{n}$ and using the fact that $\|\varepsilon_{\mathcal{T}}\|_{\infty} \lesssim \sqrt{\log n}$ on the event \mathcal{A} , we obtain

$$(97) \quad \sqrt{n}\|\tilde{\Sigma}_{\mathcal{T}}\varepsilon_{\mathcal{T}}\|_{\infty} \leq \sqrt{n}\|\tilde{\Sigma}_{\mathcal{T}}\|_{\infty}\|\varepsilon_{\mathcal{T}}\|_{\infty} \lesssim \sqrt{\frac{\log n}{n}}.$$

The sum of the remaining two terms in (55) can be bounded by a multiple of $\sqrt{\log n/n}$ by noting that $\bar{\sigma}^2 \leq \|\tilde{\Sigma}_{\mathcal{T}}\|_{\infty} \lesssim 1/n$. The statement (54) then follows from (55).

For the general covariance matrix $\Sigma_{\mathcal{T}}$ we find (similarly as in (95)) that when $\lambda_{\min}(\Sigma_{\mathcal{T}}) \geq 1/\sqrt{\log n}$ we have $\|(n\tilde{\Sigma}_{\mathcal{T}} - I_K)\beta_{\mathcal{T}}^0\|_{\infty} \leq C\sqrt{\log n/n}$. Next, because

$$\|\Sigma_{\mathcal{T}}^{-1}\|_{\infty} \leq \sqrt{K}\lambda_{\max}(\Sigma_{\mathcal{T}}^{-1}) \leq \sqrt{K\log n} < \sqrt{n\log n}$$

the matrix $\tilde{\Sigma}_{\mathcal{T}} = (nI_K + \Sigma_{\mathcal{T}}^{-1})^{-1}$ is diagonally dominant and thereby (using Lemma 4)

$$\|\tilde{\Sigma}_{\mathcal{T}}\|_{\infty} \leq \frac{1}{n + \min_{1 \leq k \leq K} \Delta_k} \quad \text{where} \quad \Delta_k = |\sigma_{kk}| - \sum_{j \neq k} |\sigma_{kj}|$$

and where $\Sigma_{\mathcal{T}}^{-1} = (\sigma_{jk})_{j,k=1}^{K,K}$. Since $\Delta_k > -\sqrt{n\log n}$ for all $k = 1, \dots, K$, the inequalities (97) and (54) hold.

10.1.4. *End of proof of Theorem 1 (general setting).* The rest of the proof can now be completed using similar arguments as in Section 6.4.

10.2. *Smooth Wavelets.* The strategy of the proof of Theorem 1 can be directly applied for an S -regular wavelet basis, also noting that the corresponding $\mathcal{H}(\alpha, M)$ space in (23) contains the usual Hölder-space for $\alpha \leq S$, using the properties of S -regular wavelets (see e.g. Section 2.2 in [12]).

10.3. *Proof of Theorem 2: confidence bands.* In the proof, we repeatedly use the properties of the median tree \mathcal{T}_X^* established in Lemma 14. We denote by \mathcal{E} the event from Lemma 14. We first show the diameter statement (69). The depth of the median tree estimator \hat{f}_T verifies condition (i) of Lemma 14 on the event \mathcal{E} . For any $f, g \in \mathcal{C}_n$, by definition of \mathcal{C}_n , we have, for $C(\psi)$ a constant depending on the wavelet basis only, on the event \mathcal{E} ,

$$\begin{aligned} \|f - g\|_\infty &\leq \|f - \hat{f}_T\|_\infty + \|\hat{f}_T - g\|_\infty \\ &\leq 2 \sup_{x \in [0,1]} \sum_{l=0}^{L_{max}} v_n \sqrt{\frac{\log n}{n}} \sum_{k=0}^{2^l-1} \mathbb{I}_{(l,k) \in \mathcal{T}_X^*} |\psi_{lk}(x)| \\ &\leq 2v_n C(\psi) \sqrt{\frac{\log n}{n}} \sum_{l: 2^l \leq C_1 2^{\mathcal{L}_c}} 2^{l/2} \leq C' v_n \sqrt{\frac{\log n}{n}} 2^{\mathcal{L}_c}. \end{aligned}$$

We now turn to the confidence statement. First, one shows that the median estimator (29) is (nearly) rate optimal. Denote with $\hat{f}_{T,lk} = \langle \hat{f}_T, \psi_{lk} \rangle$ and let $\mathcal{S} = \{(l, k) : |\beta_{lk}^0| \geq A \log n / \sqrt{n}\}$. Let us consider the event

$$(98) \quad B_n = \{\hat{f}_{T,lk} \neq 0, \forall (l, k) \in \mathcal{S}\} \cap \{\hat{f}_{T,lk} = 0, \forall (l, k) : 2^l \geq C_1 2^{\mathcal{L}_c}\} \cap \mathcal{A},$$

where the noise-event \mathcal{A} is defined in (41). Lemma 14 together with $P_{f_0}(\mathcal{A}) = 1 + o(1)$ imply that $P_{f_0}(B_n) = 1 + o(1)$. On the event B_n , we have

$$\begin{aligned} \|\hat{f}_T - f_0\|_\infty &\leq \sum_{l: 2^l \leq C_1 2^{\mathcal{L}_c}} 2^{l/2} \max \left(\max_{0 \leq k < 2^l: (l,k) \in \mathcal{S}} |X_{lk} - \beta_{lk}^0|, \max_{0 \leq k < 2^l: (l,k) \notin \mathcal{S}} \{|\beta_{lk}^0|\} \right) \\ &\quad + \sum_{l: 2^l > C_1 2^{\mathcal{L}_c}} 2^{l/2} \max_{0 \leq k < 2^l} |\beta_{lk}^0| \\ &\lesssim 2^{\mathcal{L}_c/2} \sqrt{\frac{\log n}{n}} + \sum_{l: 2^l \leq C_1 2^{\mathcal{L}_c}} 2^{l/2} \min \left(\max_{0 \leq k < 2^l} |\beta_{lk}^0|, A \frac{\log n}{\sqrt{n}} \right) + 2^{-\alpha \mathcal{L}_c}, \end{aligned}$$

where we have used the definition of \mathcal{S} , that \hat{f}_T equals 0 or X_{lk} , that f_0 belongs to $\mathcal{H}(\alpha, M)$ and $\max(a, b) \leq a + b$ (note also that the term with the minimum in the last display is an upper bound of the maximum over $(l, k) \notin \mathcal{S}$ on the first line of the display). This shows that the median tree estimator is rate-optimal up to a logarithmic factor, in probability under P_{f_0} . In particular, on B_n , we have for some $C > 0$

$$(99) \quad \|\hat{f}_T - f_0\|_\infty \leq C(\log^2 n / n)^{\alpha/(2\alpha+1)},$$

where we used the inequality in (60) in the case of smooth wavelets. Second, we now show that σ_n is appropriately large. By the proof of Proposition 3 of [37], we have for $f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)$, for $l_n \geq j_0$ suitable sequence chosen later

$$\sup_{(l,k): l \geq l_n} |\beta_{lk}^0| \geq C(M, \psi, \alpha, \varepsilon) 2^{-l_n(\alpha+1/2)},$$

for some constant $C(M, \psi, \alpha, \varepsilon)$ depending on α, M , the wavelet basis and ε (as in (2.12) of [37]). Let $\Lambda_n(\alpha)$ be defined by, for $\eta > 0$ to be chosen below,

$$(100) \quad \eta(n / \log^2 n)^{1/(2\alpha+1)} \leq 2^{\Lambda_n(\alpha)} \leq 2\eta(n / \log^2 n)^{1/(2\alpha+1)}$$

Combining the previous two displays leads to, for $f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)$,

$$\sup_{(l,k): l \geq \Lambda_n(\alpha)} |\beta_{lk}^0| \geq C(M, \psi, \alpha, \varepsilon) \eta^{-\alpha-1/2} \frac{\log n}{\sqrt{n}}.$$

By taking η small enough, one obtains that there exists (λ, κ) with $\lambda \geq \Lambda_n(\alpha)$ verifying $|\beta_{\lambda\kappa}^0| \geq A \log n / \sqrt{n}$ and thus, in turn, $\hat{f}_{T, \lambda\kappa} \neq 0$, by the second part of Lemma 14. One deduces that the term $(l, k) = (\lambda, \kappa)$ in the sum defining σ_n is nonzero on the event B_n , so that

$$\sigma_n \geq v_n c(\psi) \sqrt{\frac{\log n}{n}} \|\psi_{\lambda\kappa}\|_\infty \geq v_n c(\psi) \sqrt{\frac{\log n}{n}} 2^{\Lambda_n(\alpha)/2}.$$

This leads, on the event B_n , to

$$(101) \quad \sigma_n / \left(\frac{\log^2 n}{n} \right)^{\alpha/(2\alpha+1)} \geq c' v_n (\log n)^{-1/2}.$$

The ratio in the last display goes to infinity for v_n of larger order than $\log^{1/2} n$. Now, on the event B_n , one can thus write $\|\hat{f}_T - f_0\|_\infty \leq \sigma_n/2$ for large enough n , uniformly over $f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)$, implying that $B_n \subset \{\|\hat{f}_T - f_0\|_\infty \leq \sigma_n\}$. This implies the desired coverage property, since $P_{f_0}(f_0 \in \mathcal{C}_n) \geq P_{f_0}(B_n) = 1 + o(1)$.

For the credibility statement, we note that the posterior distribution (from Theorem 1) and the median estimator \hat{f}_T (from (99)) converge at a rate strictly faster than σ_n on the event B_n , using again the lower bound on σ_n in (101). In particular, because $B_n \subset \{\|\hat{f}_T - f_0\|_\infty \leq \sigma_n/2\}$, one can write

$$\Pi[\|f - \hat{f}_T\|_\infty \leq \sigma_n | X] \geq \Pi[\|f - f_0\|_\infty \leq \sigma_n/2 | X] \mathbb{I}_{B_n} + o_{P_{f_0}}(1).$$

The right side converges to 1 in P_{f_0} -probability, which concludes the proof of the theorem.

LEMMA 13. *The set of nodes \mathcal{T}_X^* in (28) P_{f_0} -almost surely defines a binary tree.*

PROOF. Let us recall that \mathbb{T} is the set of all admissible trees that can be obtained by sampling from the prior $\Pi_{\mathbb{T}}$ and with depth at most L_{max} . For any given node (l_1, k_1) with $0 \leq l_1 \leq L_{max}$, one can write

$$\begin{aligned} \Pi[(l_1, k_1) \in \mathcal{T} | X] &= \sum_{\mathcal{T}_1 \in \mathbb{T}} \Pi_{\mathbb{T}}[\mathcal{T}_1 | X] \times \Pi[(l_1, k_1) \in \mathcal{T}_1 | X, \mathcal{T} = \mathcal{T}_1] \\ &= \sum_{\mathcal{T}_1 \in \mathbb{T}: (l_1, k_1) \in \mathcal{T}_1} \Pi_{\mathbb{T}}[\mathcal{T}_1 | X]. \end{aligned}$$

Let $(l_1 - 1, k_1^-)$ denote the parent node of (l_1, k_1) in \mathcal{T}_1 , where $k_1^- = \lfloor k_1/2 \rfloor$. Any (full-)binary tree that contains (l_1, k_1) must also contain $(l_1 - 1, k_1^-)$, so that, using the formula in the last display, $\Pi[(l_1, k_1) \in \mathcal{T} | X] \leq \Pi[(l_1 - 1, k_1^-) \in \mathcal{T} | X]$. This implies, by definition of \mathcal{T}_X^* , that if a given node (l_1, k_1) belongs to \mathcal{T}_X^* , so does the node $(l_1 - 1, k_1^-)$. Therefore \mathcal{T}_X^* is a tree. \square

LEMMA 14. *Consider a prior distribution Π as in Theorem 1. There exists an event \mathcal{E} such that $P_{f_0}[\mathcal{E}] = 1 + o(1)$ on which the tree \mathcal{T}_X^* defined in (28) has the following properties: there exists a constant $C_1 > 0$ such that*

- (i) *the depth of the tree satisfies $2^{d(\mathcal{T}_X^*)} \leq C_1 2^{\mathcal{L}_c} \asymp (n/\log n)^{\alpha/(2\alpha+1)}$, where \mathcal{L}_c is as in (42),*
- (ii) *the tree contains as interior nodes all nodes (l, k) that satisfy $|\beta_{lk}^0| \geq A \log n / \sqrt{n}$, for some $A > 0$.*

PROOF. We focus on the GW-prior, the proof for the other two priors $\Pi_{\mathbb{T}}$ being similar. Let $\mathbb{T}^{(1)}$, respectively $\mathbb{T}^{(2)}$, denote the set of binary trees that satisfy condition (i), respectively (ii), in the statement of the lemma. By the proof of Theorem 1, $\Pi[\mathbb{T}^{(1)} | X]$ and $\Pi[\mathbb{T}^{(2)} | X]$ both tend to 1 in probability under P_{f_0} , hence also $\Pi[\mathbb{T}^{(1)} \cap \mathbb{T}^{(2)} | X]$. In fact, it also follows from the proof of Theorem 1 that, for $\mathbb{T}^{(1)}$, we also have the stronger estimate $\Pi[d(\mathcal{T}) > d | X] \leq 2^{-c_1 d \log \Gamma}$ for some $c_1 > 0$ under the GW process prior, uniformly over $\mathcal{L}_c < d \leq L_{max}$, on an event \mathcal{A} of P_{f_0} -probability going to 1. The latter probability is $o(2^{-d})$ provided Γ is chosen large enough, which will be used below. Defining $\mathcal{E} = \{\Pi[\mathbb{T}^{(1)} \cap \mathbb{T}^{(2)} | X] \geq 3/4\}$, we have $P_{f_0}[\mathcal{E}] \rightarrow 1$ as $n \rightarrow \infty$. For any node (l_2, k_2) such that $|\beta_{l_2 k_2}^0| \geq C \log n / \sqrt{n}$, we have

$$\Pi[(l_2, k_2) \in \mathcal{T}_{int} | X] = \sum_{\mathcal{T}_2 \in \mathbb{T}: (l_2, k_2) \in \mathcal{T}_{2int}} \Pi[\mathcal{T}_2 | X] \geq \Pi[\mathbb{T}^{(2)} | X],$$

where we used that any tree in $\mathbb{T}^{(2)}$ must, by definition, contain (l_2, k_2) . As $\Pi[\mathbb{T}^{(2)} | X] \geq 3/4 > 1/2$, we deduce that (l_2, k_2) belongs to \mathcal{T}_X^* on the event \mathcal{E} . In other words, \mathcal{T}_X^* verifies the second property (ii) of the lemma on \mathcal{E} . To conclude the proof of the lemma, one observes that on \mathcal{E} , for a given node (l_3, k_3) with $2^{l_3} > C_1 2^{\mathcal{L}_c}$,

$$\Pi[(l_3, k_3) \in \mathcal{T}_{int} | X] \leq \Pi[d(\mathcal{T}) > l_3 | X],$$

Recall that $\Pi[d(\mathcal{T}) > l_3 | X] \leq C 2^{-c_1 l_3 \log \Gamma}$ on \mathcal{A} (which holds uniformly over $l_3 \in [\mathcal{L}_c, L_{max}]$). Then, on the event \mathcal{A} , we can write

$$\begin{aligned} P_{f_0}[\{\mathcal{T}_X^* \notin \mathbb{T}^{(1)}\} \cap \mathcal{A}] &\leq P_{f_0}[\{\exists (l_3, k_3) : 2^{l_3} > C_1 2^{\mathcal{L}_c}, (l_3, k_3) \in \mathcal{T}_X^* \} \cap \mathcal{A}] \\ &\leq \sum_{l_3: 2^{l_3} > C_1 2^{\mathcal{L}_c}}^{L_{max}} \sum_{k_3=0}^{2^{l_3}-1} P_{f_0}[\{\Pi[(l_3, k_3) \in \mathcal{T}_{int} | X] \geq 1/2\} \cap \mathcal{A}] \\ &\leq \sum_{l_3: 2^{l_3} > C_1 2^{\mathcal{L}_c}}^{L_{max}} 2^{l_3+1} E_{f_0}[\Pi[d(\mathcal{T}) > l_3 | X] \mathbb{I}_{\mathcal{A}}] = o(1). \end{aligned}$$

Using that $P_{f_0}[\mathcal{A}]$ goes to 1, one obtains $P_{f_0}[\mathcal{T}_X^* \notin \mathbb{T}^{(1)}] = o(1)$, which concludes the proof. \square

10.4. *Proof of Theorem 3: smooth functionals.* For the second point, it suffices to combine Theorem 7 with the choice $w_l = l^2$ (which satisfies the condition $w_{j_0(n)} \gtrsim \log n$ given the assumption $j_0(n) \asymp \sqrt{\log n}$) and Theorem 4 in [16] (the condition $\sum_l w_l 2^{-l/2}$ holds), from which one indeed obtains

$$\beta_{C([0,1])} \left(\mathcal{L}(\sqrt{n}(F(\cdot) - \int_0^\cdot dX^{(n)} | X), \mathcal{L}(G)) \right) \rightarrow 0$$

in P_{f_0} -probability, where $(G(t) : t \in [0, 1])$ is a Brownian motion.

For the first point, we argue as in the paragraph preceding the statement of Theorem 4 in [16]: proceeding as in the proof of that Theorem, one notes that for a BvM to hold for the functional $\psi_b(f) = \int_0^1 f(u)b(u)du$, it suffices that, for some sequence (c_l) of positive numbers, both $\sum c_l w_l < \infty$ and $\sum_k |\langle b, \psi_{lk} \rangle| \leq c_l$ hold. With the choice $w_l = l^2$, this gives the condition of the present theorem and concludes the proof.

We note that other choices of $j_0(n)$ are possible, providing w_l is chosen appropriately: indeed $w_{j_0(n)} \gtrsim \log n$ is enough for Theorem 7 to hold. Then the BvM result for the linear functional ψ_b holds whenever $\sum c_l w_l < \infty$ and $\sum_k |\langle b, \psi_{lk} \rangle| \leq c_l$ hold, and for the functional BvM for $F(\cdot)$, it suffices that $\sum 2^{-l/2} w_l < \infty$. \square

11. Proof of Theorem 4: non-dyadic Bayesian CART. As the breakpoints verify the balancing condition (38), they verify the properties (B1)–(B2) in the complexity Lemma 11 for $\delta = 3$. The Gaussian white noise model projects onto the Haar system $\Psi_A^B = \{\psi_{-10}^B, \psi_{lk}^B : (l, k) \in A\}$ as follows:

$$(102) \quad X_{lk}^B = \beta_{lk}^{0B} + \frac{1}{\sqrt{n}} \varepsilon_{lk}^B,$$

where $X_{lk}^B = \langle X, \psi_{lk}^B \rangle$, $\beta_{lk}^{0B} = \langle f_0, \psi_{lk}^B \rangle$ and $\varepsilon_{lk}^B = \langle W, \psi_{lk}^B \rangle$. As the functions ψ_{lk}^B form an orthonormal system, the variables ε_{lk}^B are iid standard Gaussian given B . The observations here are viewed as the collection of X_{lk}^B variables which depend on B . We regard the breakpoints B as one extra “variable” in the model. Given the breakpoints B , we use the same notation $\mathbf{X}_{\mathcal{T}}^B$ and $\boldsymbol{\varepsilon}_{\mathcal{T}}^B$ for the ordered responses and noise variables (as in the proof of Theorem 1). Similarly, $\boldsymbol{\beta}_{\mathcal{T}} = \boldsymbol{\beta}_{\mathcal{T}}^B$ are the ordered internal wavelet coefficients.

As the priors on breakpoints B and trees \mathcal{T} are *independent*, the tree posterior remains relatively tractable where the amount of signal at each location (l, k) now depends on B , which requires a separate “uniform in B ” treatment.

The Multiscale Posterior Distribution. To determine the posterior distribution on f , it is enough to consider the posterior on wavelet coefficients (β_{lk}) , which then induces a posterior on f via

$$(103) \quad f_{\mathcal{T}, \tilde{\boldsymbol{\beta}}}^B(x) = \sum_{(l,k) \in \mathcal{T}_{ext}} \tilde{\beta}_{lk}^B I_{lk}^B(x) = \sum_{(l,k) \in \mathcal{T}'_{int}} \beta_{lk}^B \psi_{lk}^B(x).$$

Again, the *internal unbalanced* Haar wavelet coefficients $\boldsymbol{\beta}_{\mathcal{T}} = (\beta_{lk}^B : (l, k) \in \mathcal{T}'_{int})$ are linked to the *external* histogram coefficients $\tilde{\boldsymbol{\beta}}_{\mathcal{T}} = (\tilde{\beta}_{lk}^B : (l, k) \in \mathcal{T}_{ext})$ through $\tilde{\boldsymbol{\beta}}_{\mathcal{T}} = A_{\mathcal{T}} \boldsymbol{\beta}_{\mathcal{T}}$ for some sparse matrix $A_{\mathcal{T}} \in \mathbb{R}^{|\mathcal{T}_{ext}| \times |\mathcal{T}'_{int}|}$ (a generalization of (14)). This section describes the posterior distribution over coefficients (β_{lk}) driven by the prior distribution

$$(104) \quad (B, \mathcal{T}) \sim \Pi_{\mathbb{B}} \otimes \Pi_{\mathbb{T}} \\ (\beta_{lk})_{l \leq L, k} | B, \mathcal{T} \sim \pi(\boldsymbol{\beta}_{\mathcal{T}}) \otimes \bigotimes_{(l,k) \notin \mathcal{T}'_{int}} \delta_0(\beta_{lk}),$$

where $L = L_{max} = \lfloor \log_2 n \rfloor$. From the white noise model, we have, given B ,

$$\mathbf{X}_{\mathcal{T}}^B = \boldsymbol{\beta}_{\mathcal{T}} + \frac{1}{\sqrt{n}} \boldsymbol{\varepsilon}_{\mathcal{T}}^B, \quad \text{with} \quad \boldsymbol{\varepsilon}_{\mathcal{T}}^B \sim \mathcal{N}(0, I_{|\mathcal{T}_{ext}|}).$$

The joint density of $(B, \mathcal{T}, (\beta_{lk})_{l \leq L, k}, X)$ arising from the above distributions equals

$$\begin{aligned} & \Pi_{\mathbb{T}}(\mathcal{T}) \Pi_{\mathbb{B}}(B) \pi(\boldsymbol{\beta}_{\mathcal{T}}) \left[\prod_{(l,k) \in \mathcal{T}'_{int}} \phi_{\frac{1}{\sqrt{n}}}(X_{lk}^B - \beta_{lk}) \right] \left[\prod_{(l,k) \notin \mathcal{T}'_{int}} \phi_{\frac{1}{\sqrt{n}}}(X_{lk}^B - \beta_{lk}) \mathbb{I}_0(\beta_{lk}) \right] \\ &= \Pi_{\mathbb{T}}(\mathcal{T}) \Pi_{\mathbb{B}}(B) \left[\prod_{l \leq L, k} \phi_{\frac{1}{\sqrt{n}}}(X_{lk}^B) \right] \left[\prod_{(l,k) \notin \mathcal{T}'_{int}} \mathbb{I}_0(\beta_{lk}) \right] e^{-\frac{n}{2} \|\boldsymbol{\beta}_{\mathcal{T}}\|_2^2 + n \mathbf{X}_{\mathcal{T}}^{B'} \boldsymbol{\beta}_{\mathcal{T}}} \pi(\boldsymbol{\beta}_{\mathcal{T}}). \end{aligned}$$

Integrating out (β_{lk}) , one obtains the marginal density of (B, \mathcal{T}, X) as

$$(105) \quad \left[\Pi_{\mathbb{B}}(B) \prod_{l \leq L, k} \phi_{\frac{1}{\sqrt{n}}}(X_{lk}^B) \right] \Pi_{\mathbb{T}}(\mathcal{T}) N_X^B(\mathcal{T}),$$

where

$$N_X^B(\mathcal{T}) = \int \prod_{(l,k) \in \mathcal{T}'_{int}} e^{nX_{lk}^B \beta_{lk} - n\beta_{lk}^2/2} d\pi(\beta_{\mathcal{T}}).$$

The first bracket in (105) only depends on B and X , from which one deduces that the posterior distribution of \mathcal{T} , given B and X , satisfies

$$\Pi[\mathcal{T} | B, X] = \frac{W_X^B(\mathcal{T})}{\sum_{\mathcal{T} \in \mathbb{T}_L} W_X^B(\mathcal{T})}, \quad \text{with} \quad W_X^B(\mathcal{T}) = \Pi_{\mathbb{T}}(\mathcal{T}) N_X^B(\mathcal{T}).$$

Next, the posterior distribution on B , given X , is given by

$$\Pi[B | X] \propto \Pi_{\mathbb{B}}(B) \prod_{l \leq L, k} \phi_{\frac{1}{\sqrt{n}}}(X_{lk}^B) \left\{ \sum_{\mathcal{T} \in \mathbb{T}} W_X^B(\mathcal{T}) \right\}.$$

Also, we have

$$(106) \quad (\beta_{lk})_{l \leq L, k} | (X_{lk})_{l \leq L, k}, \mathcal{T}, B \sim \pi(\beta_{\mathcal{T}} | \mathbf{X}_{\mathcal{T}}^B) \otimes \bigotimes_{(l,k) \notin \mathcal{T}'_{int}} \delta_0(\beta_{lk}),$$

where the posterior density on the selected coefficients on \mathcal{T} is (in slight abuse of notation writing in the same way the distribution and its density)

$$(107) \quad \pi(\beta_{\mathcal{T}} | \mathbf{X}_{\mathcal{T}}^B) = \frac{e^{-\frac{n}{2} \|\beta_{\mathcal{T}}\|_2^2 + n \mathbf{X}_{\mathcal{T}}^{B'} \beta_{\mathcal{T}}} \pi(\beta_{\mathcal{T}})}{N_X^B(\mathcal{T})}.$$

Controlling the Noise. Similarly as in the proof of Theorem 1, we will condition on a set of large probability, where the noise level is relatively small. Denote with \mathbb{B} the set of *all* breakpoints B that can be obtained by performing steps (a) and (b) in Section 4.1 and that yield a system Ψ_B^A satisfying conditions (B1)–(B2) from Lemma 11. Recall $L = L_{max} = \lfloor \log_2 n \rfloor$ and $\varepsilon_{lk}^B = \int_0^1 \psi_{lk}^B(u) dW(u)$, and let δ be as in (B2). We define

$$(108) \quad \mathcal{A}_{\mathbb{B}} = \left\{ \max_{B \in \mathbb{B}} \max_{l \in [0, L], k \in [0, 2^l - 1]} (\varepsilon_{lk}^B)^2 \leq D_1 \log^{1+\delta} n \right\}$$

for some $D_1 > 0$. Using assumption (B1), one can express every single ψ_{lk}^B for $l \leq L$ in terms of a number $C_0 l^\delta$ of ψ_{jm} 's for $j \leq l + D$, where ψ_{jm} are the regular Haar wavelet functions from (3). That is, with \mathcal{X}_{lk}^B the set of such pairs (j, m) and $\text{Card}(\mathcal{X}_{lk}^B) \leq C_0 l^\delta$, we have

$$\psi_{lk}^B = \sum_{(j,m) \in \mathcal{X}_{lk}^B} p_{jm}^B \psi_{jm},$$

for some real numbers p_{jm}^B that satisfy $\sum_{(j,m) \in \mathcal{X}_{lk}^B} (p_{jm}^B)^2 = 1$ (since ψ_{lk}^B has a unit L^2 -norm and ψ_{lk} 's are orthonormal in $L^2[0, 1]$). Next, we have

$$\varepsilon_{lk}^B = \sum_{(j,m) \in \mathcal{X}_{lk}^B} p_{jm}^B \varepsilon_{jm}.$$

This itself implies the following, by the Cauchy-Schwarz inequality,

$$|\varepsilon_{lk}^B| \leq \max_{l \leq L, k} \left\{ \text{Card}(\mathcal{X}_{lk}^B) \max_{l \leq L+D, k} \varepsilon_{lk}^2 \right\}^{1/2} \leq C_0^{1/2} L^{\delta/2} \max_{l \leq L+D, k} |\varepsilon_{lk}|.$$

Using $L \leq \log_2 n$ and denoting

$$\mathcal{A} := \left\{ \max_{l \in [0, L+D], k \in [0, 2^l-1]} \varepsilon_{lk}^2 \leq 2 \log(2^{L+D+1}) \right\},$$

one obtains the inclusion $\mathcal{A} \subset \mathcal{A}_{\mathbb{B}}$, provided that D_1 is chosen larger than a universal constant (in particular it is independent of B). This implies $P_{f_0}(\mathcal{A}_{\mathbb{B}}^c) \leq P_{f_0}(\mathcal{A}^c) \leq c_0/\sqrt{\log(2^{L+D+1})}$. Next, we follow the structure of the proof of Theorem 1.

Posterior Probability of Too Deep Trees. For a given tree \mathcal{T} , we again denote with \mathcal{T}^- the pruned subtree obtained by turning the deepest rightmost internal node $(l_1, k_1) \in \mathcal{T}_{int}$ into a leaf. Given $B \in \mathbb{B}$, we proceed as in the proof of Lemma 1 and evaluate the ratio $W_X^B(\mathcal{T})/W_X^B(\mathcal{T}^-)$. When $l_1 > \mathcal{L}_c$, with \mathcal{L}_c as in (42), Lemma 11 leads to (B2), that is $|\beta_{l_1 k_1}^B| \lesssim (\log n)^{\delta/2} \sqrt{\log n/n}$ for large enough n . Similarly as in (46), we can write for $l_1 > \mathcal{L}_c$ and some $C_2 > 0$ (depending on E, D only), on the set $\mathcal{A}_{\mathbb{B}}$ from (108),

$$(X_{l_1 k_1}^B)^2 \leq \frac{C_2}{n} (\log n)^{1+\delta}.$$

Under the Galton-Watson process prior from Section 2.1.1 with $p_l \leq 1/2$ and the independent prior with $\Sigma_{\mathcal{T}} = I_{|\mathcal{T}_{ext}|}$ this gives for all $B \in \mathbb{B}$ and $d \geq \mathcal{L}_c$,

$$\frac{W_X^B(\mathcal{T})}{W_X^B(\mathcal{T}^-)} \leq 2 p_{d-1} e^{(C_2/2)(\log n)^{1+\delta}},$$

from which one deduces that

$$(109) \quad \Pi[d(\mathcal{T}) > \mathcal{L}_c \mid B, X] \leq 4 e^{(C_2/2)(\log n)^{1+\delta}} \sum_{d=\mathcal{L}_c+1}^{L_{max}} 2^{d-1} p_{d-1}.$$

The right side goes to 0 at rate $e^{-C(\log n)^{1+\delta}}$ if, e.g., p_d is of the order $(1/\Gamma)^{d^{1+\delta}}$ for some large enough $\Gamma > 0$. This also holds for a variant of the tree prior $\pi(\mathcal{T}) \propto e^{-c|\mathcal{T}_{ext}|\log^{1+\delta} n}$ and the conditionally uniform prior from Remark 2 using $\pi(K) \propto e^{-cK \log^{1+\delta} n}$ where $K = |\mathcal{T}_{ext}|$. A statement similar to (109) can also be obtained for the general prior $\pi(\beta_{\mathcal{T}}) \sim \mathcal{N}(0, \Sigma_{\mathcal{T}})$ where $\lambda_{min}(\Sigma_{\mathcal{T}}) > \sqrt{1/(\log n)^{1+\delta}}$.

Posterior Probability of Missing a Significant Node. We show a variant of Lemma 3 assuming instead that a signal node (l_S, k_S) satisfies

$$(110) \quad l_S \leq \mathcal{L}_c, \quad |\beta_{l_S k_S}^{0B}| \geq \frac{A(\log n)^{1+\frac{\delta}{2}}}{\sqrt{n}},$$

for some $A > 0$ to be chosen below. As before, for a tree $\mathcal{T} \in \mathbb{T}_{\setminus(l_S, k_S)}$ that does not have a cut at (l_S, k_S) , we denote with \mathcal{T}^+ the smallest full binary tree (in terms of number of nodes) that contains \mathcal{T} and cuts at (l_S, k_S) . Using similar arguments as in the proof of Lemma 3, we use the fact $(X_{l_S k_S}^B)^2 \geq (\beta_{l_S k_S}^{0B})^2/2 - (\varepsilon_{l_S k_S}^B)^2/n$ to find that on the event $\mathcal{A}_{\mathbb{B}}$ and for $A > 0$ large enough,

$$(111) \quad \frac{W_X^B(\mathcal{T})}{W_X^B(\mathcal{T}^+)} \leq \Gamma^{l_S^{1+\delta}(l_S+1)} e^{\frac{3}{2}D_1(l_S+1)\log^{1+\delta} n - \frac{A^2}{8}\log^{2+\delta} n} \leq e^{-\frac{A^2}{16}\log^{2+\delta} n}$$

under the independent Gaussian prior on $\beta_{\mathcal{T}}$ and the Galton-Watson process prior from Section 2.1.1 with $p_l \asymp (1/\Gamma)^{l^{1+\delta}}$. Following the steps in the proof of Lemma 2, one can show

similarly that $\Pi[(l_S, k_S) \notin \mathcal{T}_{int} \mid X, B] \rightarrow 0$ for each $B \in \mathbb{B}$ sufficiently quickly. More precisely, if

$$(112) \quad S^B(f_0; A) = \left\{ (l, k) : |\beta_{lk}^{0B}| \geq A \frac{(\log n)^{1+\frac{\delta}{2}}}{\sqrt{n}} \right\},$$

where \mathcal{L}_c is defined in (42), we have, on the event $\mathcal{A}_{\mathbb{B}}$ and for A large enough,

$$(113) \quad \Pi[\{\mathcal{T} : S^B(f_0; A) \not\subset \mathcal{T}\} \mid X] \leq e^{-C(\log n)^{1+\delta}}.$$

uniformly in $B \in \mathbb{B}$. This statement can be obtained also for the general prior $\pi(\beta_{\mathcal{T}}) \sim \mathcal{N}(0, \Sigma_{\mathcal{T}})$ with $\lambda_{\max}(\Sigma_{\mathcal{T}}) \lesssim n^a$ for some $a \geq 1$ and for other tree priors from Section 10.1.

Putting Pieces Together. Let us also set

$$(114) \quad \mathbb{T}^B = \{\mathcal{T} : d(\mathcal{T}) \leq \mathcal{L}_c, S^B(f_0; A) \subset \mathcal{T}\}, \quad \mathcal{E}^B = \{f_{\mathcal{T}, \beta} : \mathcal{T} \in \mathbb{T}^B\}.$$

From the two previous subsections one obtains that for some constant $C > 0$

$$\Pi[\mathcal{T} \notin \mathbb{T}^B \mid X, B] \leq e^{-C(\log n)^{1+\delta}},$$

for any possible set of breakpoints $B \in \mathbb{B}$ (that satisfy the balancing conditions). The uniformity in B is essential in the next bounds.

Using the definition of the event $\mathcal{A}_{\mathbb{B}}$ from (108), one can bound

$$E_{f_0} \Pi[\|f_{\mathcal{T}, \beta} - f_0\|_{\infty} > \varepsilon_n \mid X] \leq P_{f_0}[\mathcal{A}_{\mathbb{B}}^c] + E_{f_0} \{\Pi[\|f_{\mathcal{T}, \beta} - f_0\|_{\infty} > \varepsilon_n \mid X] \mathbb{I}_{\mathcal{A}_{\mathbb{B}}}\}.$$

By decomposing the posterior along B and \mathcal{T} and using Markov's inequality one obtains, on the event $\mathcal{A}_{\mathbb{B}}$,

$$\begin{aligned} \Pi[\|f_{\mathcal{T}, \beta} - f_0\|_{\infty} > \varepsilon_n \mid X] &= \sum_B \Pi[B \mid X] \sum_{\mathcal{T}} \Pi[\mathcal{T} \mid X, B] \Pi[\|f_{\mathcal{T}, \beta} - f_0\|_{\infty} > \varepsilon_n \mid X, \mathcal{T}, B] \\ &\leq \sum_B \Pi[B \mid X] \Pi[\mathcal{T} \notin \mathbb{T}^B \mid X, B] + \sum_B \Pi[B \mid X] \sum_{\mathcal{T} \in \mathbb{T}^B} \Pi[\mathcal{T} \mid B, X] \Pi[\|f_{\mathcal{T}, \beta} - f_0\|_{\infty} > \varepsilon_n \mid X, \mathcal{T}, B] \\ &\leq e^{-C(\log n)^{1+\delta}} + \sum_B \Pi[B \mid X] \sum_{\mathcal{T} \in \mathbb{T}^B} \Pi[\mathcal{T} \mid B, X] \varepsilon_n^{-1} \int \|f_{\mathcal{T}, \beta} - f_0\|_{\infty} d\Pi[f_{\mathcal{T}, \beta} \mid X, \mathcal{T}, B]. \end{aligned}$$

Let us now turn to bounding $\|f_{\mathcal{T}, \beta} - f_0\|_{\infty}$. First, note that unlike the traditional Haar basis, the UH basis system is never built up until $L = \infty$ because, by construction, we stop splitting when there are no x_i are available (i.e. we do not split nodes that are not *admissible*). In result, the very high frequencies are not covered by the system, which might induce some unwanted bias. This is, however, *not an issue* with our *weakly balanced* UH wavelets. The following Lemma shows that in weakly balanced UH systems, all nodes at levels $l \leq \Lambda := \lfloor \log_2(n/\log^c n) \rfloor$ for any $c > 0$ are admissible.

LEMMA 15. *Consider a weakly balanced UH wavelet system Ψ_A^B , where A is the set of admissible nodes (l, k) in the sense that $\mathcal{X} \cap (l_{lk}, r_{lk}] \neq \emptyset$ with $\mathcal{X} = \{x_i : x_i = 1/n, 1 \leq i \leq n\}$. Let $c > 0$, then for $\Lambda = \Lambda(c) = \lfloor \log_2(n/\log^c n) \rfloor$, we have*

$$A \supset \{(l, k) : l \leq \Lambda\}.$$

PROOF. The proof follows from the fact that the granularity of weakly balanced UH systems is very close to l . In Example 2 we defined the granularity $R(l, \Psi_A^B)$ of the l^{th} layer as the smallest integer $R \geq 1$ such that $\min_{0 \leq k < 2^l} \min\{|L_{lk}|, |R_{lk}|\} = j/2^R$ for some

$j \in \{1, 2, \dots, 2^{R-1}\}$. From Lemma 10, the granularity of weakly balanced systems Ψ_A^B is no larger than $l + D$. This means that for $l < \Lambda$, $0 \leq k < 2^l$, any $c > 0$ and n large enough

$$\min\{|L_{lk}|, |R_{lk}|\} \geq 1/2^{l+D} > \frac{\log^c n}{2^D n} > 1/n.$$

This implies that $\mathcal{X} \cap (l_{lk}, r_{lk}] \neq \emptyset$ for any (l, k) with $l \leq \Lambda$, where we used the fact that $(l_{lk}, r_{lk}]$ is either $R_{l-1 \lfloor k/2 \rfloor}$ (for when (l, k) is the right child) or $L_{l-1 \lfloor k/2 \rfloor}$ (for when (l, k) is the left child). \square

Next, we show that the weakly balanced UH systems are indeed rich enough to approximate f_0 well.

LEMMA 16. *Consider the weakly balanced UH system Ψ_A^B . Let f_0^Λ denote the L^2 -projection of $f_0 \in \mathcal{H}_M^\alpha$ onto $\text{Vect}\{\psi_{lk}^B : l \leq \Lambda\}$ for $\Lambda = \lfloor \log_2(n/\log^c n) \rfloor$ with some $c > 0$. Then*

$$\|f_0 - f_0^\Lambda\|_\infty \lesssim |\Lambda 2^{-\Lambda}|^\alpha \lesssim (\log^{c+1} n/n)^\alpha.$$

PROOF. The L^2 -projection is a step function $f_0^\Lambda = \sum_m \mathbb{I}_{\Omega_m} \tilde{\beta}_m$ supported on the pieces $\Omega_m \in \{L_{\Lambda k}, R_{\Lambda k} : 0 \leq k < 2^\Lambda\}$ where the jump sizes equal $\tilde{\beta}_m = |\Omega_m|^{-1} \int_{\Omega_m} f_0(x) dx$. From the Hölder continuity in (24) we have $|f_0(x) - f_0^\Lambda(x)| \leq M |\Omega_m|^\alpha$ for $x \in \Omega_m$. From the definition of weakly balanced UH systems, we have $\max_m |\Omega_m| \leq \frac{C+\Lambda}{2^{\Lambda+D}}$. The rest follows from the definition of Λ . \square

We can now write the following decomposition. For $f_{\mathcal{T}, \beta}$ in \mathcal{E}^B , we have

$$\begin{aligned} \|f_{\mathcal{T}, \beta} - f_0\|_\infty &\lesssim \sum_{l \leq \mathcal{L}_c} 2^{l/2} \max_{0 \leq k < 2^l} |\beta_{lk}^B - \beta_{lk}^{0B}| \\ (115) \quad &+ \sum_{\mathcal{L}_c \leq l \leq \Lambda} 2^{l/2} \max_{0 \leq k < 2^l} |\beta_{lk}^{0B}| + \|f_0 - f_0^\Lambda\|_\infty. \end{aligned}$$

In the last display, we have used the fact that for weakly balanced UH systems one has

$$\max_{0 \leq k < 2^l} \|\psi_{lk}^B\|_\infty < \max_{0 \leq k < 2^l} \left[\left(\frac{1}{|L_{lk}|} \vee \frac{1}{|R_{lk}|} \right) \frac{1}{\sqrt{|L_{lk}|^{-1} + |R_{lk}|^{-1}}} \right] < 2^{(l+D)/2}.$$

The second term in (115) can be upper-bounded by $(\log n)^{1+\delta/2} (\log n/n)^{\alpha/(2\alpha+1)}$ by using (B2) and the definition of \mathcal{L}_c . Using Lemma 16, the term $\|f_0 - f_0^\Lambda\|_\infty$ is always of smaller order than the previous one (as the bound decreases as $n^{-\alpha}$ up to a logarithmic factor).

Regarding the first term, one obtains for $\mathcal{T} \in \mathbb{T}^B$

$$\begin{aligned} &\int \max_{0 \leq k < 2^l} |\beta_{lk}^B - \beta_{lk}^{0B}| d\Pi[f_{\mathcal{T}, \beta} | X, \mathcal{T}, B] \\ &= \int \max \left(\max_{0 \leq k < 2^l, (l,k) \notin \mathcal{T}_{int}} |\beta_{lk}^{0B}|, \max_{0 \leq k < 2^l, (l,k) \in \mathcal{T}_{int}} |\beta_{lk}^B - \beta_{lk}^{0B}| \right) d\Pi[f_{\mathcal{T}, \beta} | X, \mathcal{T}, B] \\ &\leq A \frac{(\log n)^{1+\frac{\delta}{2}}}{\sqrt{n}} + \int \max_{0 \leq k < 2^l, (l,k) \in \mathcal{T}_{int}} |\beta_{lk}^B - \beta_{lk}^{0B}| d\Pi[f_{\mathcal{T}, \beta} | X, \mathcal{T}, B], \end{aligned}$$

where we have used that on the set \mathcal{E}^B , selected trees cannot miss any true signal larger than $A(\log n)^{1+\delta/2}/\sqrt{n}$. This means that any node (l, k) that is not in a selected tree must satisfy $|\beta_{lk}^{0B}| \leq A(\log n)^{1+\delta/2}/\sqrt{n}$.

We now focus on the independent prior $\beta_{\mathcal{T}} \sim \mathcal{N}(0, I_{|\mathcal{T}_{ext}|})$. We have seen above that, given X, B (so for fixed ε_{lk}^B) and \mathcal{T} , if (l, k) belongs to \mathcal{T}_{int} , the difference $\beta_{lk}^B - \beta_{lk}^{0B}$ has a Gaussian distribution Q_{lk} given by

$$Q_{lk} \stackrel{\mathcal{L}}{=} X_{lk} - \beta_{lk}^{0B} + \mathcal{N}\left(0, \frac{1}{n+1}\right) = -\frac{\beta_{lk}^{0B}}{n+1} + \frac{\sqrt{n}\varepsilon_{lk}^B}{n+1} + \mathcal{N}\left(0, \frac{1}{n+1}\right).$$

If Z_{lk} are arbitrary random variables distributed according to Q_{lk} , and \mathcal{Z}_{lk} arbitrary $\mathcal{N}(0, 1)$ random variables,

$$\mathbb{E} \left[\max_{0 \leq k < 2^l} |Z_{lk}| \right] \leq \max_{0 \leq k < 2^l} \frac{|\beta_{lk}^{0B}|}{n} + \max_{0 \leq k < 2^l} \frac{|\varepsilon_{lk}^B|}{\sqrt{n}} + \frac{1}{\sqrt{n}} \mathbb{E} \left[\max_{0 \leq k < 2^l} |\mathcal{Z}_{lk}| \right].$$

On the event $\mathcal{A}_{\mathbb{B}}$ from (108), the sum of the first two terms on the last display is bounded by $M/n + C\sqrt{(\log n)^{1+\delta}/n}$ while the last expectation is at most $C\sqrt{l/(n+1)}$ by Lemma 7. This implies

$$\int \max_{0 \leq k < 2^l, (l,k) \in \mathcal{T}_{int}} |\beta_{lk}^B - \beta_{lk}^{0B}| d\Pi[f_{\mathcal{T}, \beta} | X, \mathcal{T}, B] \leq C' \sqrt{\frac{(\log n)^{1+\delta}}{n}}$$

uniformly over B and \mathbb{T}^B , where we have used $l \leq C \log n$. Putting the various pieces together and using the fact that $P_{f_0}[\mathcal{A}_{\mathbb{B}}^c] = o(1)$, we obtain

$$\begin{aligned} & E_{f_0} \Pi[\|f_{\mathcal{T}, \beta} - f_0\|_{\infty} > \varepsilon_n | X] \\ & \leq o(1) + \varepsilon_n^{-1} \left\{ \sum_{l \leq \mathcal{L}_c} 2^{l/2} \left[A \frac{(\log n)^{1+\delta/2}}{\sqrt{n}} + C' \sqrt{\frac{\log^{1+\delta} n}{n}} \right] + 2(\log n)^{1+\frac{\delta}{2}} \left(\frac{\log n}{n} \right)^{\frac{\alpha}{2\alpha+1}} \right\} \\ & \leq o(1) + \varepsilon_n^{-1} \left\{ \left[A(\log n)^{\frac{1+\delta}{2}} + C'(\log n)^{\frac{\delta}{2}} \right] 2\sqrt{\frac{2^{\mathcal{L}_c} \log n}{n}} + 2(\log n)^{1+\frac{\delta}{2}} \left(\frac{\log n}{n} \right)^{\frac{\alpha}{2\alpha+1}} \right\} \\ & \leq o(1) + \varepsilon_n^{-1} 2C' \left[A(\log n)^{\frac{1+\delta}{2}} + 4(\log n)^{1+\frac{\delta}{2}} \right] \left(\frac{\log n}{n} \right)^{\frac{\alpha}{2\alpha+1}}. \end{aligned}$$

This means that one can set

$$\varepsilon_n = (\log n)^{1+\delta/2} \left(\frac{\log n}{n} \right)^{\frac{\alpha}{2\alpha+1}}$$

and this is the obtained posterior rate in terms of the supremum norm. A similar conclusion can be obtained for the general prior $\pi(\beta_{\mathcal{T}}) \sim \mathcal{N}(0, \Sigma_{\mathcal{T}})$ using Lemma 3 under the assumption $\lambda_{\min}(\Sigma_{\mathcal{T}}) \gtrsim \sqrt{1/\log^{1+\delta} n}$.

12. Proofs for additional results.

12.1. Proof of Theorem 5.

PROOF. For the first statement of Theorem 5, it suffices to note that the same exact proof of Theorem 1 can be used. Indeed, first, by Lemma 17, the maximum norm $\|f - f_0\|_{\infty, n}$ is bounded in a similar way as $\|f - f_0\|_{\infty}$ in terms of empirical coefficients $\mathbf{b} - \mathbf{b}^0$ (instead of original wavelet coefficients). Second, the Hölder regularity of f_0 induces a decrease of the order $2^{-l(1/2+\alpha)}$ in terms of empirical wavelet coefficients \mathbf{b}^0 (Lemma 18). Third, under P_{f_0} , the distribution of the observed Z_{lk} in (64) is $\mathcal{N}(b_{lk}^0, 1/n)$, identical to

that of $X_{lk} \sim \mathcal{N}(\beta_{lk}^0, 1/n)$, up to replacing β_{lk}^0 by b_{lk}^0 . It is then enough to prove that $\sum_{l \leq L_{max}} 2^{l/2} \max_k |b_{lk} - b_{lk}^0| \lesssim \varepsilon_n$ under the posterior distribution, which follows by the same proof as Theorem 1.

For the second statement of Theorem 5, for a given $t \in [0, 1]$, let i_t/n denote the closest (leftmost) design point $t_i = i/n$. For a given function f on $[0, 1]$, let us denote $\bar{f} = \mathcal{I}((f(t_i))_{1 \leq i \leq n})$ as a shorthand. Then for any $t \in [0, 1]$ (assuming $i_t \neq n$; if $i_t = n$ one adapts the argument) and $f_0 \in \mathcal{H}_\alpha^M$,

$$\begin{aligned} |\bar{f}(t) - f_0(t)| &\leq |f_0(t) - f_0(i_t/n)| + |f_0(i_t/n) - \bar{f}(i_t/n)| + |\bar{f}(i_t/n) - \bar{f}(t)| \\ &\leq Mn^{-\alpha} + \|\bar{f} - f_0\|_{\infty, n} + |\bar{f}(i_t/n) - \bar{f}((i_t + 1)/n)|, \end{aligned}$$

where the first line uses the triangle inequality and the second that $f_0 \in \mathcal{H}_\alpha^M$ as well as the fact that when linearly interpolating, the difference of two function values inbetween breakpoints is always the largest when taken at the two different breakpoints. Using the triangle inequality again,

$$|\bar{f}(i_t/n) - \bar{f}((i_t + 1)/n)| \leq 2\|\bar{f} - f_0\|_{\infty, n} + |f_0((i_t + 1)/n) - f_0(i_t/n)|.$$

Using again the Hölder property of f_0 , and combining with the previous bounds one gets, noting that $\|\bar{f} - f_0\|_{\infty, n} = \|f - f_0\|_{\infty, n}$,

$$\|\bar{f} - f_0\|_{\infty} \leq 2Mn^{-\alpha} + 3\|f - f_0\|_{\infty, n}.$$

This shows that if $\|f - f_0\|_{\infty, n} \leq \zeta_n$ and if $n^{-\alpha} = o(\zeta_n)$, then $\|\bar{f} - f_0\|_{\infty} \lesssim \zeta_n$. As $n^{-\alpha} = o(\varepsilon_n)$, this shows that a posterior rate on f in the $\|\cdot\|_{\infty, n}$ norm translates into the same rate for the distribution $\bar{\Pi}_Y$ in terms of the supremum norm. \square

LEMMA 17. *Given two functions f, f_0 defined on $[0, 1]$ with empirical wavelet coefficients $\mathbf{b} = (b_{lk})$, $\mathbf{b}^0 = (b_{lk}^0)$ given as in (62),*

$$\|f - f_0\|_{\infty, n} \leq \sum_{l \leq L_{max}} 2^{l/2} \max_k |b_{lk} - b_{lk}^0|.$$

PROOF. Using that $F = X\mathbf{b}$ and next that $\sum_k |\psi_{lk}(t_i)| = 2^{l/2}$,

$$\begin{aligned} \max_i |f(t_i) - f_0(t_i)| &= \max_i \left| \sum_{l \leq L_{max}, k} (b_{lk} - b_{lk}^0) \psi_{lk}(t_i) \right| \\ &\leq \sum_{l \leq L_{max}} \max_k |b_{lk} - b_{lk}^0| \max_i \sum_k |\psi_{lk}(t_i)| \\ &\leq \sum_{l \leq L_{max}} 2^{l/2} \max_k |b_{lk} - b_{lk}^0|. \quad \square \end{aligned}$$

LEMMA 18. *Suppose f_0 is in $\mathcal{H}_M^\alpha[0, 1]$ as in (24) with $\alpha \in (0, 1]$, $M > 0$, with empirical Haar wavelet coefficients $\mathbf{b}^0 = (b_{lk}^0)$ as in (62). Then for any $l \leq L_{max}$ and k ,*

$$|b_{lk}^0| \leq C2^{-l(1/2+\alpha)}.$$

PROOF. By splitting the support I_{lk} of ψ_{lk} in two halves I_{lk}^+, I_{lk}^- ,

$$\begin{aligned} n^{-1} \left| \sum_{i=1}^n f_0(t_i) \psi_{lk}(t_i) \right| &= \frac{2^{l/2}}{n} \left| \sum_{i: t_i \in I_{lk}^+} f_0(t_i) - \sum_{i: t_i \in I_{lk}^-} f_0(t_i) \right| \\ &\leq \frac{2^{l/2}}{n} \text{Card}(I_{lk}^+) \max_i |t_i - t'_i|^\alpha \leq \frac{2^{l/2}}{n} \frac{1}{2} \frac{n}{2^l} C2^{-l\alpha}. \quad \square \end{aligned}$$

LEMMA 19. *In the setting of Lemma 18, let $(\beta_{lk}^0 = \langle f_0, \psi_{lk} \rangle)$ denote the Haar wavelet coefficients of f_0 . Then for any $l \leq L_{max}$ and k ,*

$$|b_{lk}^0 - \beta_{lk}^0| \leq C2^{-l/2}n^{-\alpha}.$$

Further assume that $f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)$, i.e. belongs to \mathcal{H}_M^α and is self-similar as in Definition 3. Then for any diverging sequence $l_n \rightarrow \infty$ with $2^{l_n} \leq n$,

$$\sup_{(l,k): l \geq l_n} |b_{lk}^0| \geq C2^{-l_n(1/2+\alpha)}.$$

PROOF. For the first part, let f_0^D be the piecewise constant function that equals $f_0(t_i)$ on $[t_i, t_{i+1})$. Then observing that

$$b_{lk}^0 = 2^{l/2} \left(\int_{I_{lk}^+} f_0^D - \int_{I_{lk}^-} f_0^D \right),$$

it is enough to show that $|\int_{I_{lk}^+} (f_0^D - f_0)| \lesssim 2^{-l}n^{-\alpha}$, since then a symmetric bound similarly holds on I_{lk}^- . This follows from $\|f_0^D - f_0\|_\infty \lesssim n^{-\alpha}$, using the Hölder property of f_0 .

To prove the second part of the lemma, the proof of Proposition 3 in [37] gives, for $f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)$, for any diverging sequence l_n ,

$$\sup_{(l,k): l \geq l_n} |\beta_{lk}^0| \geq C2^{-l_n(1/2+\alpha)}.$$

The result now follow by combining the triangle inequality, the first part of the lemma, and $n^{-\alpha} \leq 2^{-l_n\alpha}$ (using $2^{l_n} \leq n$ by assumption). \square

12.2. *Proof of Theorem 6.* The proof follows the lines of that of Theorem 2, here we only highlight the few differences. First, one shows that the empirical median tree \mathcal{T}_X^\sim verifies the properties stated in Lemma 14 for the median tree \mathcal{T}_X in white noise: the depth of the tree satisfies $2^{d(\mathcal{T}_X^\sim)} \lesssim 2^{\mathcal{L}_c}$, and \mathcal{T}_X^\sim contains all nodes (l, k) such that $|b_{lk}^0| \geq A \log n / \sqrt{n}$, for some $A > 0$. The proof is the same as that of Lemma 14. From this one deduces that the diameter of the credible set \mathcal{C}_n^e is as announced, by the same argument as in the proof of Theorem 2.

Second, one shows that the empirical median tree estimator \tilde{f}_T converges at rate $(\log^2/n)^{\alpha/(2\alpha+1)}$ in terms of the $\|\cdot\|_{\infty, n}$ -norm: one adapts the proof in white noise. The event B_n and set \mathcal{S} are defined similarly, with \hat{f}_T replaced by \tilde{f}_T , β^0 by b_0 and the noise sequence ε_{lk} in the event \mathcal{A} by the noise sequence ζ_{lk} from (64). On the corresponding B_n , we have

$$\begin{aligned} \|\tilde{f}_T - f_0\|_{\infty, n} &\leq \max_i \left| \sum_{l \leq L_{max}, k} (Z_{lk} \mathbb{I}_{(l,k) \in \mathcal{T}_X^\sim} - b_{lk}^0) \psi_{lk}(t_i) \right| \\ &\leq \sum_{l: 2^l \leq C_1 2^{\mathcal{L}_c}} 2^{l/2} \max \left(\max_{0 \leq k < 2^l: (l,k) \in \mathcal{S}} |Z_{lk} - b_{lk}^0|, \max_{0 \leq k < 2^l: (l,k) \notin \mathcal{S}} \{|b_{lk}^0|\} \right) \\ &\quad + \sum_{l: n \geq 2^l > C_1 2^{\mathcal{L}_c}} 2^{l/2} \max_{0 \leq k < 2^l} |b_{lk}^0| \\ &\lesssim 2^{\mathcal{L}_c/2} \sqrt{\frac{\log n}{n}} + \sum_{l: 2^l \leq C_1 2^{\mathcal{L}_c}} 2^{l/2} \min \left(\max_{0 \leq k < 2^l} |b_{lk}^0|, A \frac{\log n}{\sqrt{n}} \right) + 2^{-\alpha \mathcal{L}_c}, \\ &\lesssim (\log^2 n/n)^{\alpha/(2\alpha+1)}, \end{aligned}$$

where we proceed as in the proof of Theorem 2 but this time using the bound on empirical coefficients from Lemma 18.

Third, one shows that $\tilde{\sigma}_n$ as in (30) is appropriately large. To do so, using that $f_0 \in \mathcal{H}_{SS}(\alpha, M, \varepsilon)'$ and arguing as in the proof of Theorem 2, for $\Lambda_n(\alpha)$ as in (100),

$$\sup_{(l,k): l \geq \Lambda_n(\alpha)} |\beta_{lk}^0| \geq C(M, \psi, \alpha, \varepsilon) \eta^{-\alpha-1/2} \frac{\log n}{\sqrt{n}}.$$

Now the same lower bound up to a different constant is obtained for b_{lk}^0 , using Lemma 19. From there on the proof is the same as for Theorem 2, which concludes the proof.

12.3. Proof of Theorem 7. This BvM statement can be shown, for example, by verifying the conditions in Proposition 6 of [16] or by proceeding as in the proof of Theorem 3.5 of [47]. The first requirement is the “tightness condition” (Proposition 6 of [16]) summarized by the following lemma.

LEMMA 20. *Under the assumptions of Theorem 7, we have*

$$E_{f_0} \Pi(\|f_{\mathcal{T},\beta} - f_0\|_{\mathcal{M}(w)} \geq M_n n^{-1/2} \mid X) \rightarrow 0.$$

PROOF. Similarly as in Section 6.4, for $j \in \mathbb{N}$ and $f \in L^2[0, 1]$ we denote with f^j the L^2 projection onto the first j layers of wavelet coefficients and write $f = f^j + f^{\setminus j}$. Similarly as in the proof of Theorem 1, we denote with \mathcal{A} the event (41) and with $S(f_0; A)$ the set (47). Recall also the notation

$$\mathcal{T} = \{\mathcal{T} : d(\mathcal{T}) \leq \mathcal{L}_c, S(f_0; A) \subset \mathcal{T}\} \quad \text{and} \quad \mathcal{E} = \{f_{\mathcal{T},\beta} : \mathcal{T} \in \mathcal{T}\}$$

from (58), where \mathcal{E} is the subset of tree-based functions $f_{\mathcal{T},\beta}$ with up to \mathcal{L}_c leaves that do not miss any signal. Similarly as in the proof of Theorem 1, we will condition on the event \mathcal{A} and focus on the set \mathcal{E} (as in (59)). For simplicity, we will write $j_0 = j_0(n)$. Following [47], one can write for some suitably chosen $D = D(\eta) > 0$, where $\eta > 0$ is a fixed small constant.

$$\begin{aligned} (116) \quad & E_{f_0} \left\{ \Pi(f_{\mathcal{T},\beta} : \|f_{\mathcal{T},\beta} - f_0\|_{\mathcal{M}(w)} \geq M_n n^{-1/2} \mid X) \right\} \leq o(1) \\ & + E_{f_0} \left\{ \Pi(f_{\mathcal{T},\beta} \in \mathcal{E} : \|f_{\mathcal{T},\beta}^{j_0} - f_0^{j_0}\|_{\mathcal{M}(w)} \geq D n^{-1/2} \mid X) \mathbb{I}_{\mathcal{A}} \right\} \\ (117) \quad & + E_{f_0} \left\{ \Pi(f_{\mathcal{T},\beta} \in \mathcal{E} : \|f_{\mathcal{T},\beta}^{\setminus j_0} - f_0^{\setminus j_0}\|_{\mathcal{M}(w)} \geq \widetilde{M}_n n^{-1/2} \mid X) \mathbb{I}_{\mathcal{A}} \right\}, \end{aligned}$$

where $\widetilde{M}_n = M_n - D \rightarrow \infty$ as $n \rightarrow \infty$. We have for $f_{\mathcal{T},\beta} \in \mathcal{E}$

$$(118) \quad \|f_{\mathcal{T},\beta}^{\setminus j_0} - f_0^{\setminus j_0}\|_{\mathcal{M}(w)} \leq \sup_{j_0 < l \leq \mathcal{L}_c} \frac{\max_k |\beta_{lk} - \beta_{lk}^0|}{w_l} + \|f_0^{\setminus \mathcal{L}_c}\|_{\mathcal{M}(w)}.$$

From the Hölder property (23) and the definition of \mathcal{L}_c in (42) we have

$$(119) \quad \|f_0^{\setminus \mathcal{L}_c}\|_{\mathcal{M}(w)} = \max_{l > \mathcal{L}_c} \frac{\max_k |\beta_{lk}^0|}{w_l} \lesssim \frac{2^{-\mathcal{L}_c(\alpha+0.5)}}{\sqrt{\mathcal{L}_c}} \leq C/\sqrt{n},$$

where we used the fact that $\{w_l\}$ is *admissible* in the sense that $w_l/\sqrt{l} \rightarrow \infty$ as $l \rightarrow \infty$. Using Markov's inequality and bounds (118) and (119), the term (117) can be bounded with

$$(120) \quad E_{f_0} \left\{ \frac{\sqrt{n}}{\widetilde{M}_n w_{j_0}} \int_{\mathcal{E}} \max_{j_0 < l \leq \mathcal{L}_c} \max_{0 \leq k \leq 2^l} |\beta_{lk} - \beta_{lk}^0| d\Pi[f_{\mathcal{T},\beta} \mid X] \mathbb{I}_{\mathcal{A}} \right\} + C/\widetilde{M}_n.$$

Using similar arguments as in Section 6.4 and using Lemma 3, we can upper bound the integral above on the event \mathcal{A} by

$$\begin{aligned} \sum_{\mathcal{T} \in \mathbb{T}} \pi[\mathcal{T} | X] \int \max_{j_0 < l \leq \mathcal{L}_c} \max_{0 \leq k \leq 2^l} |\beta_{lk} - \beta_{lk}^0| d\Pi[\beta_{\mathcal{T}} | X] \\ \leq \left(A \frac{\log n}{\sqrt{n}} + C' \sqrt{\frac{\log n}{n}} \right) \lesssim \frac{\log n}{\sqrt{n}}. \end{aligned}$$

For $w_{j_0} \geq c \log n$ for some $c > 0$, the term (120) goes to zero. Now we focus on the first term (116). By Markov's inequality and using the notation $\mathbb{W} = (g_{lk})$ for the white noise from Section 1 and $\mathbb{X} = (X_{lk})$ for the observation sequence, we find the following upper bound

$$(121) \quad \frac{\sqrt{n}}{D} \left\{ E_{f_0} \int_{\mathcal{E}} \|f_{\mathcal{T},\beta}^{j_0} - f_0^{j_0}\|_{\mathcal{M}(w)} d\Pi[f_{\mathcal{T},\beta} | X] \mathbb{I}_{\mathcal{A}} \right\} \leq \frac{E_{f_0} \{ \|\mathbb{W}^{j_0}\|_{\mathcal{M}(w)} \mathbb{I}_{\mathcal{A}} \}}{D}$$

$$(122) \quad + \frac{\sqrt{n}}{D} E_{f_0} \left\{ \int_{\mathcal{E}} \|\mathbb{X}^{j_0} - f_{\mathcal{T},\beta}^{j_0}\|_{\mathcal{M}(w)} d\Pi[f_{\mathcal{T},\beta} | X] \mathbb{I}_{\mathcal{A}} \right\}.$$

We can write the second term as

$$(123) \quad \sum_{\mathcal{T} \in \mathbb{T}} \pi[\mathcal{T} | X] E_{f_0} \int_{\mathcal{E}} \left(\sup_{l \leq j_0} l^{-1/2} \max_{0 \leq k < 2^l} \frac{\sqrt{n}}{D} |X_{lk} - \beta_{lk}| \right) d\Pi[\beta_{\mathcal{T}} | \mathbf{X}_{\mathcal{T}}].$$

Note that all trees $\mathcal{T} \in \mathbb{T}$ fit j_0 layers and under both the g -prior and the independent prior, the coefficients β_{lk} for $0 \leq l \leq j_0$ are a-priori (and a-posteriori) independent given \mathcal{T} . Similarly as in the proof of Theorem 2 in [16], we can show that the term (123) is bounded by a constant by first showing that for each $l \leq j_0$ and $0 \leq k < 2^l$

$$(124) \quad E_{f_0} \left\{ \int e^{t\sqrt{n}(\beta_{lk} - X_{lk})} d\Pi[\beta_{\mathcal{T}} | \mathbf{X}_{\mathcal{T}}] \mathbb{I}_{\mathcal{A}} \right\} \leq C e^{t^2/2}.$$

This follows from [16]. The second term $E_{f_0} \{ \|\mathbb{W}^{j_0}\|_{\mathcal{M}(w)} \mathbb{I}_{\mathcal{A}} \}$ is also bounded by C^*/D for some $C^* > 0$. Choosing $D = D(\eta) > 0$ large enough, the term on the left side of (121) can be made smaller than $\eta/2$. \square

The second step in the proof of Theorem 7 is showing convergence of finite-dimensional distributions (as in Proposition 6 of [16]). Similarly as in the proof of Theorem 2 of [16], convergence of the finite-dimensional distributions can be established by showing BvM for the projected posterior distribution onto $V_j = \text{Vect}\{\psi_{lk}, l \leq j\}$ for any fixed $j \in \mathbb{N}$. Denote with $\beta_j = (\beta_{-10}, \beta_{00}, \dots, \beta_{j2^j-1})'$ the Haar wavelet coefficients up to the level j . The prior on β_j consists of $\beta_j \sim \mathcal{N}(0, \Sigma_j)$, where Σ_j is the submatrix of Σ that corresponds to coefficients up to level j .

Let us first consider the case of the independent prior $\Sigma_{\mathcal{T}} = I_K$. Because $j_0(n) \rightarrow \infty$, for large enough n we have an *independent product prior* on β_j when $\Sigma_j = I$. Then one is exactly in the setting of Theorem 7 of [15] which derives finite-dimensional BvM for product priors (see the paragraph below the statement of Theorem 7 in [15] for two different arguments).

The case of the g -prior $\Sigma_{\mathcal{T}} = g_n(A'_{\mathcal{T}} A_{\mathcal{T}})^{-1}$ is more involved, as the induced prior distribution on the first coordinates is not of product form. Nevertheless, one can express the posterior distribution on coefficients as a mixture over trees (all containing the first $j_0(n)$ layers) of certain \mathcal{T} -dependent Gaussian distributions (complemented by zeroes for the coefficients outside the tree \mathcal{T}), and study each individual mixture component separately. Let P_{V_j} be the $n \times n$ projection matrix onto V_j and $P_{V_j}^{\mathcal{T}}$ the $|\mathcal{T}_{ext}| \times |\mathcal{T}_{ext}|$ projection matrix onto

V_j , projecting the coordinates corresponding to nodes in \mathcal{T} only (recalling that by definition of the prior, all nodes of V_j are in trees \mathcal{T} sampled from the prior). We also denote by I_{V_j} the identity matrix on V_j . It is enough for our needs to show, if $\text{TV}(P, Q)$ denotes the total variation distance between the probability distributions P and Q , that

$$(125) \quad \text{TV} \left(\Pi[\cdot | X] \circ P_{V_j}^{-1}, R_j^X \right) = o_P(1),$$

where $R_j^X := \mathcal{N}(P_{V_j} X, I_{V_j}/n)$. From the expression of the posterior (21),

$$(126) \quad \beta_{\mathcal{T}} | \mathcal{T}, X \sim \mathcal{N}(\mu_{\mathcal{T}}(X), \tilde{\Sigma}_{\mathcal{T}}) =: Q_{\mathcal{T}}^X,$$

where $\mu_{\mathcal{T}}(X) := n\tilde{\Sigma}_{\mathcal{T}} X_{\mathcal{T}}$ and $\tilde{\Sigma}_{\mathcal{T}} = (nI_K + \Sigma_{\mathcal{T}})^{-1}$. Further, the coefficients β_{lk} for $(l, k) \notin \mathcal{T}'_{int}$ are zero, which together gives a prior on $\beta_{L-1} \in \mathbb{R}^{2^L} = \mathbb{R}^n$. By definition of the prior distribution, all trees \mathcal{T} sampled from the prior contain the nodes $(l, k), l \leq j_0(n)$, in particular all nodes corresponding to V_j , so (identifying in slight abuse of notation a matrix with its corresponding linear map) the projected posterior $\Pi[\cdot | X, \mathcal{T}] \circ P_{V_j}^{-1}$ coincides with $\mathcal{N}(\mu_{\mathcal{T}}(X), \tilde{\Sigma}_{\mathcal{T}}) \circ P_{V_j}^{-1} =: Q_{\mathcal{T},j}^X$. Then

$$\begin{aligned} \text{TV} \left(\Pi[\cdot | X] \circ P_{V_j}^{-1}, R_j^X \right) &= \text{TV} \left(\sum_{\mathcal{T}} \Pi[\mathcal{T} | X] Q_{\mathcal{T},j}^X, R_j^X \right) \\ &= \text{TV} \left(\sum_{\mathcal{T}} \Pi[\mathcal{T} | X] Q_{\mathcal{T},j}^X, \sum_{\mathcal{T}} \Pi[\mathcal{T} | X] R_j^X \right) \leq \sum_{\mathcal{T}} \Pi[\mathcal{T} | X] \text{TV} (Q_{\mathcal{T},j}^X, R_j^X) \\ &\leq \max_{\mathcal{T}} \text{TV} \left(\mathcal{N}(\mu_{\mathcal{T}}(X), \tilde{\Sigma}_{\mathcal{T}}) \circ P_{V_j}^{-1}, \mathcal{N}(X_{\mathcal{T}}, I_K/n) \circ P_{V_j}^{-1} \right) \\ &\leq \max_{\mathcal{T}} \text{TV} \left(\mathcal{N}(\mu_{\mathcal{T}}(X), \tilde{\Sigma}_{\mathcal{T}}), \mathcal{N}(X_{\mathcal{T}}, I_K/n) \right), \end{aligned}$$

where sums and maxima in the last display span over trees that fill in the first $j_0(n)$ layers of nodes, and where the last line uses that the total variation distance can only decrease after projecting onto V_j (one restricts to marginal probabilities in the definition of the t.v. distance). In order to obtain (125), one now needs to bound individual distances given the tree \mathcal{T} , in a uniform way with respect to \mathcal{T} . By the triangle inequality, for any \mathcal{T} as above,

$$\begin{aligned} &\text{TV} \left(\mathcal{N}(\mu_{\mathcal{T}}(X), \tilde{\Sigma}_{\mathcal{T}}), \mathcal{N}(X_{\mathcal{T}}, I_K/n) \right) \\ &\leq \text{TV} \left(\mathcal{N}(\mu_{\mathcal{T}}(X), \tilde{\Sigma}_{\mathcal{T}}), \mathcal{N}(\mu_{\mathcal{T}}(X), \frac{I_K}{n}) \right) + \text{TV} \left(\mathcal{N}(\mu_{\mathcal{T}}(X), \frac{I_K}{n}), \mathcal{N}(X_{\mathcal{T}}, \frac{I_K}{n}) \right). \end{aligned}$$

Both terms on the right hand side of the last display can be bounded using Lemma 8, where one sets $d = K = |\mathcal{T}_{ext}|$, $\mu = \mu_{\mathcal{T}}(X) = \mu_1$, $\mu_2 = X_{\mathcal{T}}$, and $\Sigma = I_K/n = \Sigma_1$, $\Sigma_2 = \tilde{\Sigma}_{\mathcal{T}}$. Then $\Sigma_1^{-1}\Sigma_2 - I_K = n(nI_K + \Sigma_{\mathcal{T}}^{-1})^{-1} - I_K = -(I_K + n\Sigma_{\mathcal{T}})^{-1}$, using the formula $(I + B)^{-1} = I - (I + B^{-1})^{-1}$ for B invertible. Setting $M_{\mathcal{T}} := (I_K + n\Sigma_{\mathcal{T}})^{-1}$, the first and second inequalities of Lemma 8 lead to

$$\begin{aligned} \text{TV} \left(\mathcal{N}(\mu_{\mathcal{T}}(X), \tilde{\Sigma}_{\mathcal{T}}), \mathcal{N}(\mu_{\mathcal{T}}(X), \frac{I_K}{n}) \right) &\lesssim \|M_{\mathcal{T}}\|_F \\ \text{TV} \left(\mathcal{N}(\mu_{\mathcal{T}}(X), \frac{I_K}{n}), \mathcal{N}(X_{\mathcal{T}}, \frac{I_K}{n}) \right) &\lesssim \frac{\|M_{\mathcal{T}} X_{\mathcal{T}}\|^2}{\sqrt{\frac{1}{n} \|M_{\mathcal{T}} X_{\mathcal{T}}\|^2}} = \sqrt{n} \|M_{\mathcal{T}} X_{\mathcal{T}}\|. \end{aligned}$$

One now notes $\|M_{\mathcal{T}}\|_F \leq \sqrt{K} \lambda_{\max}(M_{\mathcal{T}}) = \sqrt{K} / \lambda_{\min}(I_K + n\Sigma_{\mathcal{T}})$. By Proposition 3, we have $\lambda_{\min}((A'_{\mathcal{T}}A_{\mathcal{T}})^{-1})$ is at least $2^{-L} \geq 1/n$, and one deduces that $\lambda_{\min}(I_K + n\Sigma_{\mathcal{T}}) \gtrsim 1 + ng_n/n \geq 1 + g_n$, so that $\|M_{\mathcal{T}}\|_F \lesssim \sqrt{K}/g_n \lesssim 2^{L/2}/g_n \lesssim 1/\sqrt{n} = o(1)$, uniformly over \mathcal{T} . On the other hand, we have, as $\lambda_{\max}(M_{\mathcal{T}}) \leq 1/(1 + g_n)$ as seen above and $X_{\mathcal{T}} = \beta_{\mathcal{T}}^0 + \varepsilon_{\mathcal{T}}/\sqrt{n}$, so that, working on the event \mathcal{A} from (41),

$$\begin{aligned} \|M_{\mathcal{T}}X_{\mathcal{T}}\|^2 &\leq \lambda_{\max}(M_{\mathcal{T}})^2 \|X_{\mathcal{T}}\|^2 \lesssim g_n^{-2} (\|\beta_{\mathcal{T}}^0\|^2 + \|\varepsilon_{\mathcal{T}}\|^2/n) \\ &\lesssim g_n^{-2} (1 + n(\log n)/n) \lesssim (\log n)/g_n^2, \end{aligned}$$

where we have used that $\|\beta^0\|^2 = \|f^0\|^2$ is bounded and $\|\varepsilon_{\mathcal{T}}\|^2 \lesssim n \log n$ on the event \mathcal{A} . The previous bounds together imply that the total variation distance between $\mathcal{N}(\mu_{\mathcal{T}}(X), \tilde{\Sigma}_{\mathcal{T}})$ and $\mathcal{N}(X_{\mathcal{T}}, I_K/n)$ goes to 0 uniformly in \mathcal{T} on the event \mathcal{A} . As $P[\mathcal{A}^c]$ vanishes, this proves (125).

12.4. Proof of Theorem 8.

PROOF. Define a sequence

$$L^* = \left\lceil \log_2 \left[M^{1/(\alpha+1/2)} (n/\log^2 n)^{1/(2\alpha+1)} \right] \right\rceil,$$

so that $2^{L^*} \asymp (n/\log^2 n)^{1/(2\alpha+1)}$. Define the sequence of functions f_0^n (below we write simply f_0 for simplicity) by its sequence of wavelet coefficients as follows: set all coefficients β_{lk}^0 to 0 except for $\beta_{L^*0}^0 = M2^{-L^*(1/2+\alpha)}$. By definition, f_0 belongs to $\mathcal{H}(\alpha, M)$. Let us also note that if $(L^*, 0)$ does not belong to the tree \mathcal{T} , one can bound from below

$$\ell_{\infty}(f_{\mathcal{T},\beta}, f_0) \geq 2^{L^*/2} |\beta_{L^*0}| = M2^{-L^*\alpha} \geq C' \varepsilon_n.$$

So, to prove the result, it is enough to show that $\Pi[(L^*, 0) \notin \mathcal{T}_{int} | X] \rightarrow 1$, i.e. the node $(L^*, 0)$ does not belong to a tree sampled from the posterior with probability going to 1, or equivalently, if \mathbb{T}_{L^*0} denotes the set of all full binary trees (of depth at most L_{\max}) that contain $(L^*, 0)$ as an internal node, that $\Pi[\mathbb{T}_{L^*0} | X] = o_P(1)$. To prove this, let us consider a given tree $\mathcal{T} \in \mathbb{T}_{L^*0}$. As it contains the node $(L^*, 0)$, it must also contain all nodes $(\lambda, 0)$ with $0 \leq \lambda \leq L^*$, in particular $(L_1, 0)$, where $L_1 = \lceil L^*/2 \rceil$, say. We note that $L^* \asymp L^* - L_1 \asymp \log n$. Let τ^* be the maximal subtree of \mathcal{T} that has $(L_1, 0)$ as its root. Next, let \mathcal{T}_-^* denote the remainder tree built from \mathcal{T} by erasing all of τ^* except for the node $(L_1, 0)$ (so that \mathcal{T}_-^* still has a full-binary tree structure). So, \mathcal{T}_-^* and τ^* have only the node $(L_1, 0)$ in common, and the union of their nodes gives the original tree \mathcal{T} . Let us now write

$$\Pi[\mathbb{T}_{L^*0} | X] = \frac{\sum_{\mathcal{T} \in \mathbb{T}_{L^*0}} W_X(\mathcal{T})}{\sum_{\mathcal{T} \in \mathbb{T}} W_X(\mathcal{T})} = \frac{\sum_{\mathcal{T} \in \mathbb{T}_{L^*0}} \frac{W_X(\mathcal{T})}{W_X(\mathcal{T}_-^*)} W_X(\mathcal{T}_-^*)}{\sum_{\mathcal{T} \in \mathbb{T}} W_X(\mathcal{T})}.$$

Let $q = q(\tau^*)$ denote the number of internal nodes τ_{int}^* of the subtree τ^* . From the Galton-Watson prior, we obtain

$$(127) \quad \frac{\Pi_{\mathbb{T}}[\mathcal{T}]}{\Pi_{\mathbb{T}}[\mathcal{T}_-^*]} = \prod_{(l,k) \in \tau_{int}^*} p_l \prod_{(l,k) \in \tau_{ext}^*} (1 - p_l) \frac{1}{1 - p_{L_1}} \leq 2 \prod_{(l,k) \in \tau_{int}^*} \Gamma^{-l}.$$

Then, by definition of \mathcal{T}_-^* and τ^* ,

$$\begin{aligned} \frac{W_X(\mathcal{T})}{W_X(\mathcal{T}_-^*)} &= \frac{\Pi_{\mathbb{T}}(\mathcal{T})}{\Pi_{\mathbb{T}}(\mathcal{T}_-^*)} \prod_{(l,k) \in \tau_{int}^*} \frac{\exp\{(n+1)X_{lk}^2/2\}}{\sqrt{n+1}} \\ (128) \quad &\leq 2(n+1)^{-q/2} \cdot \prod_{(l,k) \in \tau_{int}^*} \Gamma^{-l} \exp\{(n+1)X_{lk}^2/2\}. \end{aligned}$$

We bound the data-dependent part in the previous line by using $(a + b)^2 \leq 2a^2 + 2b^2$. Furthermore, noting that the noise variables $|\varepsilon_{lk}|$ are uniformly bounded for $l + 1 \leq L_{max}$, $0 \leq k < 2^l$, by $2 \log n$ on an event of overwhelming probability, we can upper-bound $(n + 1) \sum_{(l,k) \in \tau_{int}^*} X_{lk}^2/2$ by

$$\begin{aligned} & (n + 1) \sum_{(l,k) \in \tau_{int}^*} \left[(\beta_{L^*0}^0)^2 \mathbb{I}_{(l,k)=(L^*,0)} + \frac{1}{n} \max_{l+1 \leq L_{max}, k} \varepsilon_{lk}^2 \right] \\ & \leq (n + 1) (\beta_{L^*0}^0)^2 + 2 \frac{n+1}{n} (\log n) q \leq \frac{n+1}{n} \log^2 n + 2 \frac{n+1}{n} (\log n) q. \end{aligned}$$

Now, using that for $(l, k) \in \tau_{int}^*$, we have $l \geq L_1 \geq \frac{1}{2\alpha+1} \log_2 \left(\frac{M^2 n}{\log n} \right) =: c(\alpha, M, n)$, one notes that

$$\sum_{(l,k) \in \tau_{int}^*} l \geq \max \left(c(\alpha, M, n) q, \sum_{l=L_1}^{L^*} l \right) \geq c(\alpha, M, n) \max \left(q, \frac{3}{2} c(\alpha, M, n) \right),$$

which is bounded from below by $c(\alpha, M, n) q$, where we have used that $q \geq L^* - L_1 + 1 \geq L_1 \geq c(\alpha, M, n)$ and $L_1 + L^* > 3c(\alpha, M, n)$. One then deduces that the product of terms Γ^{-l} dominates (128), as long as $\log(\Gamma)$ is large enough (noting also that $1/(2\alpha + 1) \geq 1/(2S + 1)$), in the control of $W_X(\mathcal{T})/W_X(\mathcal{T}_-^*)$. That is, for some constant $C > 0$,

$$\frac{W_X(\mathcal{T})}{W_X(\mathcal{T}_-^*)} \leq \exp\{-C(\log n)q\} =: b_q,$$

where the last bound only depends on the number of internal nodes q of τ^* . By coming back to the above bound on the posterior $\Pi[\mathbb{T}_{L^*0} | X]$, let us split the sum on the numerator as follows. Let $\mathbb{T}_{L^*0}^q$ denote the set of trees $\mathcal{T} = \mathcal{T}_-^* \cup \tau^*$ in \mathbb{T}_{L^*0} such that $|\tau_{int}^*| = q$. Then $\Pi[\mathbb{T}_{L^*0} | X]$ is bounded by

$$\sum_{q \geq 1} \frac{\sum_{\mathcal{T} \in \mathbb{T}_{L^*0}^q} b_q W_X(\mathcal{T}_-^*)}{\sum_{\mathcal{T} \in \mathbb{T}} W_X(\mathcal{T})} \leq \sum_{q \geq 1} \frac{\sum_{\mathcal{T}_1 \in \mathbb{T}_-^{*q}} a_q b_q W_X(\mathcal{T}_1)}{\sum_{\mathcal{T}_1 \in \mathbb{T}_-^{*q}} W_X(\mathcal{T}_1)},$$

where \mathbb{T}_-^{*q} denotes the set of all possible \mathcal{T}_-^* that can be obtained from $\mathcal{T} \in \mathbb{T}_{L^*0}^q$ and where a_q denotes the number of different possible trees τ^* such that $|\tau_{int}^*| = q$. To obtain the last bound, we also used that each $\mathcal{T} \in \mathbb{T}_{L^*0}$ is uniquely characterised by a pair $(\mathcal{T}_-^*, \tau^*)$, so that the sum over \mathcal{T} can be rewritten as a double sum over \mathcal{T}_-^* and τ^* . One deduces that, as q cannot be larger than 2^L ,

$$\Pi[\mathbb{T}_{L^*0} | X] \leq \sum_{q=1}^{2^L} a_q b_q.$$

As a_q is less (because of the restriction $|\mathcal{T}| \leq L$) or equal to the number of full binary trees with q internal nodes, i.e. with $2q + 1$ nodes in total, we have $a_q \leq \mathbb{C}_{2q}$, which is bounded from above by 4^{2q} by Lemma 6. We conclude that $\Pi[\mathbb{T}_{L^*0} | X]$ is bounded above by $\exp(-C \log^2 n)$ for some $C > 0$, on an event of overwhelming probability, which concludes the proof for the Galton-Watson prior. Similarly, for the exponential prior, we replace (127) directly with

$$\frac{\Pi_{\mathbb{T}}[\mathcal{T}]}{\Pi_{\mathbb{T}}[\mathcal{T}_-^*]} \propto e^{-c(|\mathcal{T}_{ext}| - |\mathcal{T}_-^{*ext}|) \log n} = e^{-cq \log n},$$

where we used the fact that $|\mathcal{T}_{ext}| - |\mathcal{T}_{ext}^*| = |\tau_{ext}^*| - 1 = |\tau_{int}^*| = q$. For the conditionally uniform prior, we have for $\lambda = 1/n^c$ for some $c > 0$

$$\frac{\Pi_{\mathbb{T}}[\mathcal{T}]}{\Pi_{\mathbb{T}}[\mathcal{T}^*]} = \frac{\pi(|\mathcal{T}_{ext}|)}{\pi(|\mathcal{T}_{ext}^*|)} \frac{\mathbb{C}_{|\mathcal{T}_{int}^*|}}{\mathbb{C}_{|\mathcal{T}_{int}|}} \lesssim \lambda^q e^{-q \log |\mathcal{T}_{ext}|} \lesssim e^{-cq \log n},$$

and the end of the proof is then the same as for the GW prior. \square

12.5. Proof of Theorem 9.

PROOF. Denote with \mathcal{T}_D^F the flat tree of depth $D + 1$ (i.e. all β_{lk} 's for $l \leq D$ are active). The formula (22) gives

$$\begin{aligned} \Pi[\mathcal{T}_D^F | X] &\propto W_X(\mathcal{T}_D^F) = \Pi_{\mathbb{T}}(\mathcal{T}_D^F) \prod_{(l,k) \in \mathcal{T}_{D+1}^{F'}} \frac{e^{\frac{n^2}{2(n+1)} X_{lk}^2}}{\sqrt{n+1}} \\ &\propto \exp \left\{ -\log \Pi_{\mathbb{T}}(\mathcal{T}_D^F) - 2^D \log(n+1) + \frac{n^2}{2(n+1)} \|\mathbf{X}^{(D)}\|_2^2 \right\}, \end{aligned}$$

where $\|\mathbf{X}^{(D)}\|_2^2 = \sum_{l \leq D, k} X_{lk}^2$ is the squared L^2 -norm of the signal, truncated at the level D . Next, we have

$$\begin{aligned} \|\mathbf{X}^{(D)}\|_2^2 &= \sum_{l \leq D, k} (\beta_{lk}^0)^2 + \sum_{l \leq D, k} \frac{2}{\sqrt{n}} \varepsilon_{lk} \beta_{lk}^0 + \sum_{l \leq D, k} \frac{1}{n} \varepsilon_{lk}^2, \\ &= C_n - \sum_{D < l \leq L_{max}, k} (\beta_{lk}^0)^2 - \sum_{D < l \leq L_{max}, k} \frac{2}{\sqrt{n}} \varepsilon_{lk} \beta_{lk}^0 + \sum_{l \leq D, k} \frac{1}{n} \varepsilon_{lk}^2, \end{aligned}$$

where $C_n = C(n, \{\varepsilon_{lk}\}, f_0)$ does not depend on D . We can also write

$$(129) \quad \|\mathbf{X}^{(D)}\|^2 = C_n - \sum_{D < l \leq L_{max}, k} (\beta_{lk}^0)^2 + \frac{2^{D+1}}{n} - \frac{2}{\sqrt{n}} Z(D) + \frac{1}{n} Q(D),$$

where we have used $\sum_{l \leq D, k} 1 = 2^{D+1}$ and have set

$$Z(D) = \sum_{D < l \leq L_{max}, k} \beta_{lk}^0 \varepsilon_{lk}, \quad Q(D) = \sum_{l \leq D, k} (\varepsilon_{lk}^2 - 1).$$

Let D^* be an integer defined as, for f_0 to be chosen below,

$$D^* = \operatorname{argmin}_{0 \leq D \leq n} \left[2^D \log(n+1) + \frac{n}{2} \sum_{l=D+1}^{L_{max}} \sum_k (\beta_{lk}^0)^2 \right].$$

Consider the following true signal $f_0 = f_0^*$ which belongs to $\mathcal{H}(\alpha, M)$ with $M = 1$ (which we can assume without loss of generality) and which is characterized by the following wavelet coefficients

$$(130) \quad \beta_{lk}^{0*} = \begin{cases} 2^{-l(\frac{1}{2} + \alpha)} & \text{if } k = 0, \\ 0 & \text{otherwise.} \end{cases}$$

For such a signal, D^* above has the following behavior

$$(131) \quad 2^{D^*} \asymp \left(\frac{n}{\log n} \right)^{\frac{1}{2\alpha+2}}.$$

With the maximum-type norm ℓ_∞ defined in (57), we use the decomposition $\ell_\infty(f, f_0) = \ell_\infty(f, f_0^D) + \ell_\infty(f_0^D, f_0)$, where f_0^D is the L^2 -projection of f_0 onto the first D levels of wavelet coefficients. Moreover, using $\ell_\infty(f_0^D, f_0) = c2^{-D\alpha}$ for some $c > 0$, we can write, for $\rho_n = (\log n/n)^{\alpha/(2\alpha+2)}$

$$\begin{aligned} \Pi[\ell_\infty(f, f_0) < \mu\rho_n \mid X] &\leq \Pi[\ell_\infty(f_0, f_0^D) < \mu\rho_n \mid X] \\ &= \Pi[c2^{-D\alpha} < \mu\rho_n \mid X] = \Pi[2^D > (c\mu^{-1}\rho_n^{-1})^{1/\alpha} \mid X] \\ &\leq \Pi[2^D > (c\mu^{-1})^{1/\alpha} 2^{D^*} \mid X]. \end{aligned}$$

To conclude, it is enough to show that for $B = \{2^D > (c\mu^{-1})^{1/\alpha} 2^{D^*}\}$, where $\mu > 0$ is a small enough constant, we have $\Pi[B \mid X] = o(1)$ or, equivalently, $\Pi[B \mid X] = o(\Pi[B^c \mid X])$ (possibly on an event of vanishing probability). Rewriting $B = \{D : D > cD^*\}$ for $c = c(\mu) \geq 1$ (up to taking μ small enough), and using the above expression of $\Pi[\mathcal{T}_D^F \mid X]$, one obtains

$$\begin{aligned} \frac{\Pi[B \mid X]}{\Pi[B^c \mid X]} &= \frac{\sum_{D > cD^*} \exp \left\{ -\log \Pi(\mathcal{T}_D^F) - 2^D \log(n+1) + \frac{n^2}{2(n+1)} \|\mathbf{X}^{(D)}\|_2^2 \right\}}{\sum_{D \leq cD^*} \exp \left\{ -\log \Pi(\mathcal{T}_D^F) - 2^D \log(n+1) + \frac{n^2}{2(n+1)} \|\mathbf{X}^{(D)}\|_2^2 \right\}} \\ &\leq \frac{\sum_{D > cD^*} \exp \left\{ -\log \Pi(\mathcal{T}_D^F) - 2^D \log(n+1) + \frac{n^2}{2(n+1)} \|\mathbf{X}^{(D)}\|_2^2 \right\}}{\exp \left\{ -\log \Pi(\mathcal{T}_{D^*}^F) - 2^{D^*} \log(n+1) + \frac{n^2}{2(n+1)} \|\mathbf{X}^{(D^*)}\|_2^2 \right\}}. \end{aligned}$$

Since $c \geq 1$ we have $D \geq D^* + 1$ for any $D > cD^*$ and from the monotonicity assumption on the prior we obtain $\log \Pi(\mathcal{T}_{D^*}^F) - \log \Pi(\mathcal{T}_D^F) \leq 0$ on B . In addition, note that $2^{D^*} - 2^D \leq -2^D/2$ on B , which implies

$$(2^{D^*} - 2^D) \log(n+1) \leq -\frac{1}{2} 2^D \log(n+1).$$

Going further, using the decomposition of $\|\mathbf{X}^{(D)}\|_2^2$ in (129) we have for $Z = \|\mathbf{X}^{(D)}\|_2^2 - \|\mathbf{X}^{(D^*)}\|_2^2$ the following

$$\begin{aligned} Z &= \sum_{D^* < l \leq D, k} (\beta_{lk}^0)^2 + \frac{1}{n} (2^{D+1} - 2^{D^*+1}) - \frac{2}{\sqrt{n}} (Z(D) - Z(D^*)) + \frac{1}{n} (Q(D) - Q(D^*)) \\ &\leq \sum_{D^* < l \leq L_{max}, k} (\beta_{lk}^0)^2 + \frac{2^{D+1}}{n} + \frac{2}{\sqrt{n}} (|Z(D)| + |Z(D^*)|) + \frac{1}{n} (|Q(D)| + |Q(D^*)|). \end{aligned}$$

We now provide bounds for the stochastic terms Z and Q . First, for any $D > D^*$, denoting $\sigma_D^2 := \sum_{D^* < l \leq L_{max}, k} (\beta_{lk}^0)^2$, we have

$$|Z(D)| \leq \sigma_D \max_{D^* \leq D \leq L_{max}} \sigma_D^{-1} |Z(D)|.$$

The variables $Z(D)/\sigma_D$ are standard normal, which implies that, on some event A_1 such that $P_{f_0}(A_1^c) = o(1)$, we have, uniformly in $D \in B$,

$$|Z(D)| \leq \sigma_D \sqrt{2 \log L_{max}}.$$

To bound the term $Q(D)$, one can use the following standard concentration bound for chi-square distributions. Namely, for ξ_q standard normal variables and any $t > 0$, we can write

$$\mathbb{P} \left[\sum_{q=1}^Q (\xi_q^2 - 1) \geq t \right] \leq \exp \left\{ -\frac{t^2}{4(Q+t)} \right\}.$$

Applying this bound for the noise variables ε_{lk} and choosing $t = t_D := (D2^D)^{1/2}$ leads to

$$\mathbb{P} \left[\sum_{l \leq D, k} (\varepsilon_{lk}^2 - 1) > t_D \right] \leq \exp \left\{ -\frac{D2^D}{4(2^{D+1} + t_D)} \right\}.$$

For $D \geq D^*$, one has $t_D \leq 2^{D+1}$ so the last display is bounded from above by $\exp\{-C_1 D\}$. Let us consider the event, with t_D as above,

$$A_2 = \bigcap_{D=D^*}^{L_{max}} \left\{ \sum_{l \leq D} \sum_{k=0}^{2^l-1} (\varepsilon_{lk}^2 - 1) \leq t_D \right\}.$$

A union bound gives $P_{f_0}[A_2^c] \leq C \exp(-c_1 D^*)$, which is a $o(1)$ using the previous bound. Now let us choose μ small enough in such a way that $C_2 2^{D^*} \leq 2^D/2$ for any D in the set B defined above (this is possible by definition of B) and thereby

$$\frac{n}{2} \sum_{D^* < l \leq L_{max}, k} (\beta_{lk}^0)^2 \leq \frac{2^D}{4} \log(n+1)$$

for any D in B . This in particular implies that $\sigma_D \leq (2^D \log(n+1)/n)^{1/2}$. Now, on the event $A_1 \cap A_2$, we have

$$\begin{aligned} \frac{\Pi[B|X]}{\Pi[B^c|X]} &\leq \sum_{D > cD^*} \exp \left\{ -\frac{1}{2} 2^D \log(n+1) + \frac{2^D}{4} \log(n+1) + 2^D \right. \\ &\quad \left. + 2\sqrt{n}\sigma_D \sqrt{2 \log L_{max}} + (D2^D)^{1/2} \right\} \\ &\leq \sum_{D > cD^*} \exp \left\{ -\frac{1}{8} 2^D \log(n+1) \right. \\ &\quad \left. + \left[2^D + 2\sqrt{2^D \log(n+1) 2 \log L_{max}} + (D2^D)^{1/2} - \frac{1}{8} 2^D \log(n+1) \right] \right\} \\ &\leq \sum_{D > cD^*} \exp \left\{ -\frac{1}{8} 2^D \log(n+1) \right\} \leq \exp \left\{ -C 2^{D^*} \log(n+1) \right\}, \end{aligned}$$

where we have used that the term under brackets in the second inequality is negative for large enough n , as $2^D \gtrsim 2^{D^*}$ goes to infinity. This shows that the last display goes to 0, which concludes the proof. \square

REFERENCES

- [1] ARBEL, J., GAYRAUD, G. and ROUSSEAU, J. (2013). Bayesian Optimal Adaptive Estimation Using a Sieve Prior. *Scandinavian Journal of Statistics* **40** 549–570.
- [2] BARANIUK, R. G., CEVHER, V., DUARTE, M. F. and HEGDE, C. (2010). Model-based compressive sensing. *IEEE Trans. Inform. Theory* **56** 1982–2001. [MR2654489](#)
- [3] BARBIERI, M. M. and BERGER, J. O. (2004). Optimal Predictive Model Selection. *The Annals of Statistics* **32** 870–897.
- [4] BARBIERI, M. M., BERGER, J. O., GEORGE, E. I. and ROČKOVÁ, V. (2018). The Median Probability Model and Correlated Variables. arXiv preprint 1807.08336.
- [5] BAYARRI, M. J., BERGER, J. O., FORTE, A. and GARCÍA-DONATO, G. (2012). Criteria for Bayesian model choice with application to variable selection. *The Annals of Statistics* **40** 1550–1577.
- [6] BIAU, G. (2012). Analysis of a random forests model. *The Journal of Machine Learning Research* **13** 1063–1095.
- [7] BLANCHARD, G., SCHÄFER, C. and ROZENHOLC, Y. (2004). Oracle bounds and exact algorithm for dyadic classification trees. In *Learning theory. Lecture Notes in Comput. Sci.* **3120** 378–392. Springer, Berlin. [MR2177922](#)
- [8] BLEICH, J., KAPELNER, A., GEORGE, E. I. and JENSEN, S. (2014). Variable selection for BART: an application to gene regulation. *The Annals of Applied Statistics* **8** 1750–1781.
- [9] BREIMAN, L., FRIEDMAN, J. H., OLSHEN, R. A. and STONE, C. J. (1984). *Classification and Regression Trees*. Wadsworth and Brooks.
- [10] CAI, T. T. (2008). On information pooling, adaptability and superefficiency in nonparametric function estimation. *Journal of Multivariate Analysis* **99** 421–436.
- [11] CASTILLO, I. (2008). Lower bounds for posterior rates with Gaussian process priors. *Electronic Journal of Statistics* **2** 1281–1299.
- [12] CASTILLO, I. (2014). On Bayesian supremum norm contraction rates. *The Annals of Statistics* **42** 2058–2091.
- [13] CASTILLO, I. (2017). Pólya tree posterior distributions on densities. *Ann. Inst. Henri Poincaré Probab. Stat.* **53** 2074–2102.
- [14] CASTILLO, I. and MISMER, R. (2019). Spike and slab Pólya tree posterior densities: adaptive inference. arXiv preprint 1911.12106.
- [15] CASTILLO, I. and NICKL, R. (2013). Nonparametric Bernstein–von Mises theorems in Gaussian white noise. *The Annals of Statistics* **41** 1999–2028.
- [16] CASTILLO, I. and NICKL, R. (2014). On the Bernstein–von Mises theorem for nonparametric Bayes procedures. *The Annals of Statistics* **42** 1941–1969.
- [17] CASTILLO, I. and ROČKOVÁ, V. (2019). Supplement to “Multiscale Analysis of Bayesian CART”.
- [18] CHIPMAN, H., GEORGE, E. I., MCCULLOCH, R. E. and SHIVELY, T. (2016). High-dimensional nonparametric monotone function estimation using BART. arXiv preprint 1612.01619.
- [19] CHIPMAN, H., GEORGE, E. I. and MCCULLOCH, R. (2010). BART: Bayesian additive regression trees. *Annals of Applied Statistics* **4** 266–298.
- [20] CHIPMAN, H., GEORGE, E. I. and MCCULLOCH, R. E. (1997). Bayesian CART Model Search. *Journal of the American Statistical Association* **93** 935–960.
- [21] CHIPMAN, H. A., KOLACZYK, E. D. and MCCULLOCH, R. E. (1997). Adaptive Bayesian Wavelet Shrinkage. *Journal of the American Statistical Association* **92** 1413–1421.
- [22] COHEN, A., DAUBECHIES, I. and VIAL, P. (1993). Wavelets on the interval and fast wavelet transforms. *Appl. Comput. Harmon. Anal.* **1** 54–81.
- [23] DENISON, D., MALLICK, B. and SMITH, A. (1998). A Bayesian CART Algorithm. *Biometrika* **85** 363–377.
- [24] DEVROYE, L., MEHRABIAN, A. and REDDAD, T. (2018). The total variation distance between high-dimensional Gaussians. arXiv preprint 1810.08693.
- [25] DONOHO, D. L. (1997). CART and best-ortho-basis: a connection. *Ann. Statist.* **25** 1870–1911. [MR1474073](#)
- [26] ENGEL, J. (1994). A Simple Wavelet Approach to Nonparametric Regression from Recursive Partitioning Schemes. *Journal of Multivariate Analysis* **49** 242–254.
- [27] FRYZLEWICZ, P. (2007). Unbalanced Haar technique for nonparametric function estimation. *Journal of the American Statistical Association* **480** 1318–1327.
- [28] GEY, S. and NÉDÉLEC, E. (2005). Model selection for CART regression trees. *IEEE Trans. Inf. Th.* **51** 658–670.
- [29] GHOSAL, S., GHOSH, J. and VAN DER VAART, A. (2000). Convergence rates of posterior distributions. *Annals of Statistics* **28** 500–5311.

- [30] GHOSAL, S. and VAN DER VAART, A. (2017). *Fundamentals of nonparametric Bayesian inference*. Cambridge Series in Statistical and Probabilistic Mathematics **44**. Cambridge University Press, Cambridge. [MR3587782](#)
- [31] GINÉ, E. and NICKL, R. (2011). Rates of contraction for posterior distributions in L^r -metrics, $1 \leq r \leq \infty$. *Ann. Statist.* **39** 2883–2911. [MR3012395](#)
- [32] GINÉ, E. and NICKL, R. (2016). *Mathematical foundations of infinite-dimensional statistical models*. Cambridge Series in Statistical and Probabilistic Mathematics, [40]. Cambridge University Press, New York. [MR3588285](#)
- [33] GIRARD, M. and SWELDENS, W. (1997). A new class of unbalanced Haar wavelets that form an unconditional basis for L_p on general measure spaces. *Journal of Fourier Analysis and Applications* **3** 457–474.
- [34] GOLUB, G. H. and VAN LOAN, C. F. (1996). *Matrix Computations*. The John Hopkins University Press.
- [35] HAHN, P. R., MURRAY, J. S. and CARVALHO, C. (2017). Bayesian regression tree models for causal inference: regularization, confounding, and heterogeneous effects. arXiv preprint 1706.09523.
- [36] HILL, J. (2011). Bayesian Nonparametric Modeling for Causal Inference. *Journal of Computational and Graphical Statistics* **20** 217–240.
- [37] HOFFMANN, M. and NICKL, R. (2011). On adaptive inference and confidence bands. *Annals of Statistics* **39** 2383–2409.
- [38] HOFFMANN, M., ROUSSEAU, J. and SCHMIDT-HIEBER, J. (2015). On Adaptive Posterior Concentration Rates. *The Annals of Statistics* **43** 2259–2295.
- [39] HOOKER, G. and MENTCH, L. (2016). Quantifying uncertainty in random forests via confidence intervals and hypothesis tests. *The Journal of Machine Learning Research* **17** 841–881.
- [40] LINERO, A. (2016). Bayesian regression trees for high dimensional prediction and variable selection. *Journal of the American Statistical Association* **113** 626–636.
- [41] LINERO, A. and YANG, Y. (2018). Bayesian Regression Tree Ensembles that Adapt to Smoothness and Sparsity. *Journal of the Royal Statistical Association* **80** 1087–1110.
- [42] LIU, Y., ROČKOVÁ, V. and WANG, Y. (2018). ABC variable selection with Bayesian forests. arXiv preprint 1806.02304.
- [43] MURRAY, J. (2017). Log-Linear Bayesian Additive Regression Trees for Categorical and Count Responses. arXiv preprint 1706.09523.
- [44] NAULET, Z. (2018). Adaptive Bayesian density estimation in sup-norm. arXiv preprint 1805.05816.
- [45] NICKL, R. and RAY, K. (2019). Nonparametric statistical inference for drift vector fields of multi-dimensional diffusions. *The Annals of Statistics*. to appear.
- [46] NICKL, R. and SÖHL, J. (2019). Bernstein–von Mises theorems for statistical inverse problems II: compound Poisson processes. *Electronic Journal of Statistics* **13** 3513–3571.
- [47] RAY, K. (2017). Adaptive Bernstein–von Mises Theorems In Gaussian White Noise. *The Annals of Statistics* **45** 2511–2536.
- [48] ROČKOVÁ, V. and SAHA, E. (2019). On Theory for BART. *Proceedings of Machine Learning Research: 22nd International Conference on Artificial Intelligence and Statistics* **89** 2839–2848.
- [49] ROČKOVÁ, V. and VAN DER PAS, S. (2017). Posterior Concentration for Bayesian Regression Trees and Forests. *The Annals of Statistics* 1–40. to appear.
- [50] SCORNET, E., BIAU, G. and VERT, J. P. (2015). Consistency of Random Forests. *Annals of Statistics* **43** 1716–1741.
- [51] SCOTT, C. and NOWAK, R. D. (2006). Minimax-optimal classification with dyadic decision trees. *IEEE Trans. Inform. Theory* **52** 1335–1353. [MR2241192](#)
- [52] SCRICCIOLO, C. (2014). Adaptive Bayesian density estimation in L^p -metrics with Pitman-Yor or normalized inverse-Gaussian process kernel mixtures. *Bayesian Anal.* **9** 475–520. [MR3217004](#)
- [53] STANLEY, R. P. (1999). *Enumerative combinatorics. Vol. 2*. Cambridge Studies in Advanced Mathematics **62**. Cambridge University Press, Cambridge. [MR1676282](#)
- [54] VAN DER PAS, S. and ROČKOVÁ, V. (2017). Bayesian Dyadic Trees and Histograms for Regression. *Proceedings of the 31st International Conference on Neural Information Processing Systems* 2086–2096.
- [55] VARAH, J. M. (1975). A Lower Bound for the Smallest Singular Value of a Matrix. *Linear Algebra and Its Applications* **11** 3–5.
- [56] WAGER, S. and ATHEY, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* **113** 1228–1242.
- [57] WAGER, S. and WALTHER, G. (2015). Adaptive concentration of regression trees, with application to random forests. arXiv preprint 1503.06388.
- [58] WILLETT, R. and NOWAK, R. (2007). Multiscale Poisson Intensity and Density Estimation. *IEEE Transactions on Information Theory* **53** 3171–3187.

- [59] WONG, W. H. and MA, L. (2010). Optional Pólya tree and Bayesian inference. *The Annals of Statist.* **38** 1433–1459. [MR2662348](#)
- [60] YOO, W., RIVOIRARD, R. and ROUSSEAU, J. (2017). Adaptive Supremum Norm Posterior Contraction: wavelet Spike-and-Slab and Anisotropic Besov Spaces. arXiv preprint 1708.01909.
- [61] YOO, W., RIVOIRARD, V. and ROUSSEAU, J. (2018). Adaptive Supremum Norm Posterior Contraction: Wavelet Spike-and-Slab and Anisotropic Besov Spaces. 1–50. arXiv preprint 1708.01909.
- [62] YOO, W. and VAN DER VAART, A. (2017). The Bayes Lepski’s Method and Credible Bands through Volume of Tubular Neighborhoods. arXiv preprint 1711.06926.
- [63] YOO, W. W. and GHOSAL, S. (2016). Supremum norm posterior contraction and credible sets for nonparametric multivariate regression. *The Annals of Statistics* **44** 1069–1102. [MR3485954](#)
- [64] ZELLNER, A. (1986). On assessing prior distributions and Bayesian regression analysis with g -prior distributions. *Bayesian inference and decision techniques: Essays in Honor of Bruno De Finetti* **6** 233–243.