



HAL
open science

Efficient Bayesian estimation and use of cut posterior in semiparametric hidden Markov models

Daniel Moss, Judith Rousseau

► **To cite this version:**

Daniel Moss, Judith Rousseau. Efficient Bayesian estimation and use of cut posterior in semiparametric hidden Markov models. *Electronic Journal of Statistics*, 2023, 18 (1), pp.1815-1886. 10.1214/23-EJS2201 . hal-04386010

HAL Id: hal-04386010

<https://hal.science/hal-04386010v1>

Submitted on 11 Sep 2024

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

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



Distributed under a Creative Commons Attribution 4.0 International License

Efficient Bayesian estimation and use of cut posterior in semiparametric hidden Markov models

Daniel Moss¹ and Judith Rousseau¹

¹*Department of Statistics, University of Oxford, United Kingdom,
e-mail: daniel.moss@stats.ox.ac.uk; judith.rousseau@stats.ox.ac.uk*

Abstract: We consider the problem of estimation in Hidden Markov models with finite state space and nonparametric emission distributions. Efficient estimators for the transition matrix are exhibited, and a semiparametric Bernstein-von Mises result is deduced. Following from this, we propose a modular approach using the cut posterior to jointly estimate the transition matrix and the emission densities. We first derive a general theorem on contraction rates for this approach. We then show how this result may be applied to obtain a contraction rate result for the emission densities in our setting; a key intermediate step is an inversion inequality relating L^1 distance between the marginal densities to L^1 distance between the emissions. Finally, a contraction result for the smoothing probabilities is shown, which avoids the common approach of sample splitting. Simulations are provided which demonstrate both the theory and the ease of its implementation.

Keywords and phrases: Bernstein-von Mises, contraction rates, cut posterior, efficiency, Hidden Markov models, inversion inequality, semiparametric estimation.

Received June 2022.

1. Introduction

Hidden Markov models (HMMs) are a broad and widely used class of statistical models, enjoying applications in a variety of settings where observed data is linked to some ordered process, for which an assumption of independently distributed data would be both inappropriate and uninformative. Specific applications include modelling of weather [3, 35], circadian rhythms [34], animal behaviour [18, 38], finance [43], information retrieval [47, 54], biomolecular dynamics [32], genomics [60] and speech recognition [50].

In this paper, we consider inference in finite state space HMMs. Such models are characterised by an unobserved (latent) Markov chain $(X_t)_{t \geq 1}$ taking values in $[R] = \{1, 2, \dots, R\}$ with $R < \infty$, evolving according to a transition matrix Q . Conditionally on $X_t = j$, $Y_t \sim F_j$ where F_j is the *emission distribution* with associated density f_j . These models generalise independent mixture models, which are obtained as a special case when the X_t are independent and identically distributed. Here we assume that R is known so that the parameters in such models are then Q and $\mathbf{F} = (F_1, \dots, F_R)$.

Most of the work on HMMs considers parametric models, where the emissions are assumed to admit densities in a parametric class $\{f_\theta : \theta \in \Theta\}$ where $\Theta \subset \mathbb{R}^d$ for some $d < \infty$, see for instance [11, 20, 23]. However such an assumption leads to inference which is strongly influenced by the choice of the parametric family $\{f_\theta : \theta \in \Theta\}$. This problem has been often discussed in the literature, see for instance [50], [21] or [60], especially but not solely in relation to clustering or segmentation. In the seminal paper [24], the authors show identifiability under weak assumptions on the emissions, provided the transition matrix is of full rank, paving the way for estimation of semiparametric models. In [4], it is shown further that the number R of hidden states may be identified. In [6, 33] and [16], frequentist estimators of Q and $\mathbf{f} = (f_1, \dots, f_R)$ respectively have been proposed using spectral methods, showing in particular that Q can be estimated at the rate $1/\sqrt{n}$. Frequentist estimation of the emission densities has also been addressed using penalised least squares approaches [15] and spectral methods [16]. However, no results exist on the asymptotic distribution of frequentist estimators for Q , nor on efficient estimation for Q .

Although Bayesian nonparametric estimation methods have been considered in practice in Hidden Markov models, see for instance [60] or [21], little is known about their theoretical properties. While [57, 58] established posterior consistency under general conditions on the prior and refined the analysis to derive posterior concentration rates on the marginal density of ℓ successive observations, $g_{Q, \mathbf{F}}^\ell$, no results exist regarding the properties of Bayesian procedures when seeking to recover the parameters Q and \mathbf{F} , or other functionals of Q, \mathbf{F} which are often of interests in HMMs. For instance, when interests lie in clustering or segmentation, the quantities of interest are the smoothing probabilities, being the conditional distribution of the latent states given the observations. In [16], the authors obtain rates of convergence for frequentist estimators of the smoothing probabilities, but their result requires either a sup-norm convergence for the estimator $\hat{\mathbf{f}}$ or splitting the data into two parts, with estimation of \mathbf{f} based on one part of the data and estimation of the smoothing probabilities based on the other part of the observations. In this paper we intend to bridge this gap, concentrating on Bayesian semiparametric methods while also exhibiting non-Bayesian, semiparametric efficient estimators of Q .

We first construct a family of priors Π_1 on (Q, \mathbf{F}) which we show in Theorem 2 leads to an asymptotically normal posterior distribution for $\sqrt{n}(Q - \hat{Q})$, of variance V detailed therein. Here \hat{Q} is a frequentist estimator, exhibited in Theorem 1, for which $\sqrt{n}(\hat{Q} - Q^*)$ converges in distribution to $N(0, V)$, with Q^* being the frequentist true parameter under which the data is assumed to be generated. Consequently, Bayesian point estimates associated to such posteriors, such as the posterior mean, enjoy parametric convergence rates to Q^* and, importantly, credible regions for Q are asymptotically confidence regions. In Theorem 3, we then refine this construction to obtain a scheme for which V is the optimal variance for semiparametric estimation of Q , i.e. it is the inverse efficient Fisher information. Semiparametric Bernstein-von Mises properties are highly non trivial results and correspondingly sparse in the literature, moreover in [13, 14, 22, 51] a number of counterexamples are exhibited, where semiparametric

Bernstein-von Mises results do not hold. Since this property is crucial to ensure that credible regions are also confidence regions and thus robust to the choice of the prior distribution, it is important to study and verify it.

Our approach to obtaining these results follows the ideas of [27], extending their work on mixtures to the more complex HMM setting. In particular, the construction of \hat{Q} (together with the prior on Q, \mathbf{F}) is based on a parametric approximation to the nonparametric emission model, with the property that estimation in this model with appropriately ‘coarsened’ data leads to a well-specified parametric model for estimation of Q . Once we reduce to such a parametric model, we have asymptotic normality of the corresponding MLE as in [9] as well as Bernstein-von Mises for the posterior as in [17], although the intuition that our ‘coarsened’ data is less informative translates to an inefficient asymptotic variance V . We then show that we can construct efficient estimators of Q by adjusting the degree of coarsening with the sample size, similar to the approach of [27]. The proof techniques in our context are however significantly more complex, since the notions of Fisher information and score functions are much less tractable in hidden Markov models (see for instance [19]).

The prior distributions $\Pi_1(dQ, d\mathbf{F})$ considered above, which lead to the Bernstein-von Mises property of the posterior on Q , rely on a crude modelling of the emission densities and are therefore not well behaved as far as the estimation of \mathbf{F} is concerned. To overcome this problem, we adapt in Section 4 the cut posterior approach (see [8, 12, 36, 42, 49]) to our semiparametric setting. Cut posteriors were originally proposed in the context of modular approaches to modelling, where a model is assembled from a number of constituent models, each with its own parameters θ_i and data Y_i . In the usual construction, the cut posterior has the effect of ‘cutting feedback’ of one of the (less reliable) data sources Y_i on the other parameters (associated to more reliable data).

Our approach, though also modular, departs from this setting in that we use a single data source but wish to choose different priors for different parts of the parameter. We consider a conditional prior $\Pi_2(d\mathbf{F}|Q)$, then combine $\Pi_1(dQ|Y_{1:n})$ with the conditional posterior $\Pi_2(d\mathbf{F}|Q, Y_{1:n})$ to produce a joint distribution over Q and \mathbf{F} . In this way, we construct a distribution over the parameters which is not a proper Bayesian posterior, but which simultaneously satisfies a Bernstein-von Mises theorem for the posterior marginal on Q , and is well-concentrated on the emission distributions and other functionals such as the smoothing probabilities. Through this construction, we manage to combine the “best of both worlds” in the estimation of the parametric and nonparametric parts. We believe that this idea could be used more generally in other semiparametric models.

As previously mentioned, the existing posterior concentration result in the semiparametric HMMs covered only the marginal density $g_{Q, \mathbf{F}}^\ell$ of a fixed number ℓ of consecutive observations, see [57]. A key step in obtaining posterior contraction rates on the emission distributions \mathbf{F} is an inversion inequality allowing us to deduce L^1 concentration of the posterior distribution of the emission distributions from concentration (at the same rate) of the marginal distribution of the observations. This is established in Theorem 4, from which we derive

contraction rates for the cut-posterior on \mathbf{f} in Theorem 5. This inversion inequality is of independent interest and can be used outside the framework of Bayesian inference. We finally show in Theorem 6 that these results lead to posterior concentration of the smoothing distributions, which are the conditional distributions of the hidden states given the observations, building on [16] but refining the analysis so that we require neither sup-norm contraction rate on the emissions \mathbf{f} , nor a splitting of the data set in two parts.

Organisation of the paper The paper is organised as follows. In Section 2, we introduce the model and the notions involved, together with a general strategy for inference on Q and \mathbf{f} based on the cut posterior approach. In Section 3, we study the estimation of Q , proving asymptotic normality of the posterior and asymptotic efficiency. In Section 4, the cut posterior approach is studied, posterior contraction rates for \mathbf{f} are derived together with posterior contraction rates for the smoothing probabilities. Theoretical results from both sections are then illustrated in Section 5. The proofs for the efficiency of the estimation of Q (Proposition 1), for the inversion inequality (Theorem 4), and for the convergence of the smoothing probabilities (Theorem 6), are presented in Section 6. The other proofs are deferred to the appendix. The appendix also contains a general contraction rate theorem for cut posteriors, together with an exposition of relevant results in the literature and further details of simulations.

2. Inference in semiparametric finite state space HMMs

In this section we present our strategy to estimate Q and $\mathbf{F} = (F_1, \dots, F_R)$, based on a modular approach which consists of first constructing the marginal posterior distribution of Q based on a first prior Π_1 on (Q, \mathbf{F}) , and then combining it with the conditional posterior distribution of \mathbf{F} given Q and $Y_{1:n}$ based on a different prior $\Pi_2(d\mathbf{F}|Q)$.

2.1. Model and notation

Hidden Markov models (HMMs) are latent variable models where one observes $Y_{1:n} = (Y_t)_{1 \leq t \leq n}$ whose distribution is modelled via latent (non-observed) variables $X_{1:n} = (X_t)_{1 \leq t \leq n} \in [R]^n$ which form a Markov chain. In this work we consider finite state space HMMs:

$$Y_t | X_t \stackrel{\text{ind}}{\sim} F_{X_t}, \quad \mathbb{P}_Q(X_t = r | X_{t-1} = s) = Q_{rs}, \quad r, s \leq R, \quad t \geq 1 \quad (1)$$

and the number of states R is assumed to be known throughout the paper.

The parameters are then the transition matrix $Q = (Q_{rs})_{r,s \leq R}$ of the Markov chain and the emission distributions $\mathbf{F} = \{F_r\}_{r \leq R}$, which represent the conditional distribution of the observations given the latent states. For a transition matrix Q we denote by p_Q its invariant distribution (when it exists).

Throughout the paper we denote by $\mathbf{F}^* = \{F_r^*\}_{r \leq R}$ and Q^* respectively the true emission distributions and the true transition matrix. The aim is to




FIG 1. Visual representation of data generating process of HMM.

make inference on \mathbf{F}^* , Q^* , and some functionals of these parameters, using likelihood based methods and in particular Bayesian methods. We assume that the distributions F_j^* , $j = 1, \dots, R$ are absolutely continuous with respect to some measure λ on $\mathcal{Y} \subset \mathbb{R}^d$, $d \geq 1$ and we denote by f_1^*, \dots, f_R^* their corresponding densities.

When the latent states $(X_t)_{t \geq 1}$ are independent and identically distributed on $[R]$, the parameters are not identifiable unless some strong assumptions are made on the F_j 's, see [5]. However, in [24], it is proved that under weak assumptions on the data generating process, both Q and \mathbf{F} are identifiable up to label swapping (or label switching). More precisely, let $\mathbb{P}_{Q, \mathbf{F}}^{(\ell)}$ be the marginal distribution of ℓ consecutive observations from model 1 with parameters Q and \mathbf{F} , so that

$$\mathbb{P}_{Q, \mathbf{F}}^{(\ell)}(dy_{1:\ell}) = \sum_{i_1, \dots, i_\ell=1}^R p_Q(i_1) \prod_{t=1}^{\ell-1} (Q_{i_t, i_{t+1}} F_{i_t}(dy_t)) F_{i_\ell}(dy_\ell). \quad (2)$$

Under assumptions of ergodicity, p_Q exists and is unique. For example, this holds under the assumption $Q_{ij} > 0$ for all $i, j \in [R]$, see for instance equation (1) in [25]. Denote by $\mathcal{Q} = \Delta_R^R \subset [0, 1]^{R \times R}$ the R -fold product of the $R - 1$ dimensional simplex and let \mathcal{P} be the set of probability measures on \mathbb{R}^d . Consider the following assumption:

- Assumption 1.** (i) The latent chain $(X_t)_{t \geq 1}$ has true transition matrix $Q^* = (Q_{ij}^*)$ satisfying $\det Q^* \neq 0$.
 (ii) The true emission distributions $\{F_r^*\}_{r=1}^R$ are linearly independent.

From [24], if **Assumption 1** holds, then for any $Q \in \mathcal{Q}$ and any $F_1, \dots, F_R \in \mathcal{P}$, if $\mathbb{P}_{Q, \mathbf{F}}^{(3)} = \mathbb{P}_{Q^*, \mathbf{F}^*}^{(3)}$ then $Q = Q^*$ and $\mathbf{F} = \mathbf{F}^*$ up to label swapping. By “up to label swapping,” we precisely mean that there exists a permutation τ of $[R]$ such that ${}^\tau Q = Q^*$ and ${}^\tau \mathbf{F} = \mathbf{F}^*$, where ${}^\tau Q = (Q_{\tau(r), \tau(s)})$, $r, s \in [R]$ and ${}^\tau \mathbf{F} = (F_{\tau(1)}, \dots, F_{\tau(R)})$. The requirement for such a permutation τ is unavoidable, since the labelling of the hidden states is fundamentally arbitrary. Correspondingly, the results which follow will always be given up to label swapping. In a slight abuse of notation, we will sometimes interchange \mathbf{F} and $\mathbf{f} = (f_1, \dots, f_R)$,

the latter being the densities of \mathbf{F} with respect to some measure λ in a dominated model.

The likelihood associated to model (1), when $(F_r)_{r \in [R]}$ are dominated by a measure λ , is then given by

$$L_n(Q, \mathbf{F}) = g^{(n)}(Q, \mathbf{F})(y_1, \dots, y_n) = \sum_{i_1, \dots, i_n=1}^R p_Q(i_1) \prod_{t=1}^{n-1} Q_{i_t, i_{t+1}} f_{i_t}(y_t) f_{i_n}(y_n). \tag{3}$$

Extension to initial distributions different from the stationary one is straightforward, under the exponential forgetfulness of the Markov chain which holds under our assumptions below, see Section C.5.

If Π is a prior on $\mathcal{Q} \times \mathcal{F}^R$, with $\mathcal{F} = \{f : \mathbb{R}^d \rightarrow \mathbb{R}_+ : \int f(y) d\lambda(y) = 1\}$ being the set of densities on \mathbb{R}^d with respect to λ , then the Bayesian posterior $\Pi(\cdot | Y_{1:n})$ is defined as follows: for any Borel subset A of $\mathcal{Q} \times \mathcal{F}^R$, we have

$$\Pi(A | Y_{1:n}) = \frac{\int_A L_n(Q, \mathbf{F}) \Pi(dQ, d\mathbf{F})}{\int_{\mathcal{Q} \times \mathcal{F}^R} L_n(Q, \mathbf{F}) \Pi(dQ, d\mathbf{F})}. \tag{4}$$

This posterior is well defined as soon as p_Q exists Π -almost surely, which holds for instance when $\Pi(\min_{i,j} Q_{ij} > 0) = 1$, which implies that the transition matrix is ergodic Π -almost surely by the earlier remarks. Throughout the paper we consider the parameterisation of Q given by $\tilde{Q} = (Q_{rs}, r \in [R], s \in [R-1])$ so that for each r we have $Q_{rR} = 1 - \sum_{s < R} Q_{rs}$. Hence, specification of the matrix Q amounts to specification of the $(R) \times (R-1)$ matrix $\tilde{Q} \in \tilde{\mathcal{Q}}$ for which $\sum_{s=1}^{R-1} Q_{rs} \leq 1$ for all r , and we will identify Q with \tilde{Q} (and \mathcal{Q} with $\tilde{\mathcal{Q}}$) when making statements about asymptotic distributions.

It will be helpful to consider that the stationary HMM $(X_t, Y_t)_{t \geq 1}$ arises as the restriction to $\{t \geq 1\}$ of a process $(X_t, Y_t)_{t \in \mathbb{Z}}$ indexed by $t \in \mathbb{Z}$, where the $(X_t)_{t \in \mathbb{Z}}$ and the $(Y_t)_{t \leq 0}$ are not observed. This is possible to define by considering the reversal of the latent chain.

We use \mathbb{P}_* to denote the joint law of the variables $(X_t, Y_t)_{t \in \mathbb{Z}}$ under the stationary distribution associated with the parameters Q^*, \mathbf{F}^* . Estimators, given a sample of size n , are then understood to be random variables which are measurable with respect to the σ -algebra generated by the observations $(Y_t)_{1 \leq t \leq n}$.

2.2. Cut posterior inference: A general strategy for joint inference on Q and \mathbf{f}

In Section 3.1 below we construct a family of prior distributions Π_1 on $\mathcal{Q} \times \mathcal{F}^R$ such that the associated marginal posterior distribution of Q satisfies a Bernstein-von Mises theorem centred at an asymptotically normal regular estimator \hat{Q} . Specifically, this means that with probability going to 1 under \mathbb{P}_* ,

$$\|\Pi_1(\min_{\tau \in \mathcal{S}_R} \sqrt{n}(\tau Q - \hat{Q}) \in (\cdot) | Y_{1:n}) - N(0, V)\|_{TV} = o_{\mathbb{P}_*}(1) \tag{5}$$

and

$$\min_{\tau \in \mathcal{S}_R} \sqrt{n}(\tau \hat{Q} - Q^*) \Rightarrow^* N(0, V), \tag{6}$$

where \Rightarrow^* denotes convergence in distribution under \mathbb{P}_* and where \mathcal{S}_R is the set of permutations of $[R]$.

However, the choice of Π_1 which we require for this control over Q leads to a posterior distribution on \mathbf{F} which is badly behaved, see Section 3. In order to jointly estimate well both Q and \mathbf{F} , we propose the following cut posterior approach. Consider a second conditional prior distribution $\Pi_2(d\mathbf{F}|Q)$ which leads to a conditional posterior distribution of \mathbf{F} given $Q, Y_{1:n}$ of the form

$$\Pi_2(d\mathbf{F}|Q, Y_{1:n}) = \frac{L_n(Q, \mathbf{F})\Pi_2(d\mathbf{F}|Q)}{\int_{\mathcal{F}^R} L_n(Q, \mathbf{F})\Pi_2(d\mathbf{F}|Q)}.$$

The cut posterior is then the probability distribution on $\mathcal{Q} \times \mathcal{F}^R$ given by

$$\Pi_{cut}(d(Q, \mathbf{F})|Y_{1:n}) = \Pi_1(dQ|Y_{1:n})\Pi_2(d\mathbf{F}|Q, Y_{1:n}). \tag{7}$$

Note that, if $\Pi_2(d\mathbf{F}|Q) = \Pi_1(d\mathbf{F}|Q)$, then $\Pi_{cut}(d(Q, \mathbf{F})|Y_{1:n})$ is a proper posterior distribution and is equal to $\Pi_1(d(Q, \mathbf{F})|Y_{1:n})$. The motivation behind the use of the cut posterior $\Pi_{cut}(d(Q, \mathbf{F})|Y_{1:n})$ is to retain the good behaviour of $\Pi_1(dQ|Y_{1:n})$ in terms of estimating Q , while being flexible in the modelling of \mathbf{F} to ensure that the posterior distribution over both \mathbf{F} and Q (and functionals of these parameters) are well behaved.

Adapting the proof techniques from [29] to posterior concentration rates for cut posteriors, we derive in Section 4 contraction rates for cut posteriors in terms of the L^1 norm of $g_{Q, \mathbf{F}}^{(3)} - g_{Q^*, \mathbf{F}^*}^{(3)}$, where $g_{Q, \mathbf{F}}^{(3)}$ is the density of $\mathbb{P}_{Q, \mathbf{F}}^{(3)}$ with respect to the dominating measure λ . We show that, under suitable conditions and choice of $\epsilon_n = o(1)$, $\Pi_{cut}(\|g_{Q, \mathbf{F}}^{(3)} - g_{Q^*, \mathbf{F}^*}^{(3)}\|_1 \leq \epsilon_n | Y_{1:n}) = 1 + o_{\mathbb{P}_*}(1)$.

To derive cut posterior contraction rates in terms of the L^1 norm of $\mathbf{f} - \mathbf{f}^*$, we prove in Section 4.2 an inversion inequality in the form

$$\sum_{r=1}^R \|f_r - f_r^*\|_1 \lesssim \|Q - Q^*\| + \|g_{Q, \mathbf{F}}^{(3)} - g_{Q^*, \mathbf{F}^*}^{(3)}\|_{L^1},$$

which is also of independent interest.

We also derive cut-posterior contraction rates for the smoothing probabilities $(p_{Q, \mathbf{F}}(X_i = \cdot | Y_{1:n}))_{i=1, \dots, n}$ in Section 4.3. In contrast with [16], concentration rates for $(p_{Q, \mathbf{F}}(X_i = \cdot | Y_{1:n}))_{i=1, \dots, n}$ do not require a splitting of the data into 2 groups, nor do they require to have a control of $f_r - f_r^*$ in sup-norm. We can avoid these difficulties thanks to the Bayesian approach as is explained in Section 4.3.

In our implementation of the cut posterior, we adopt a nested MCMC approach of the kind detailed in [12] and [36], see Section 5 for details.

In the following section, we present Π_1 and show that the associated marginal posterior distribution $\Pi_1(dQ|Y_{1:n})$ is asymptotically normal in the sense of Equation (5).

3. Semi - parametric estimation of Q : Bernstein-von Mises property and efficient estimation

The prior Π_1 is based on a simple histogram prior on the f_1, \dots, f_R . For the sake of simplicity, we present the case of univariate data; the multivariate case can be treated similarly. Without loss of generality, we assume that the observations belong to $[0, 1]$ and note that, if $\mathcal{Y} = \mathbb{R}$, we can transform the data to $[0, 1]$ via some C^1 diffeomorphism ϕ prior to the analysis. The prior relies on a partition of the space $[0, 1]$ into finitely many intervals $\{I_1, I_2, \dots\}$, and transforming the data is equivalent to constructing a prior based on the corresponding partition $\{\phi^{-1}(I_1), \phi^{-1}(I_2), \dots\}$ of \mathbb{R} . The construction of Π_1 is analogous to the construction considered in [27].

Let $M \in \mathbb{N}$ with $M \geq R$ and consider $\mathcal{I}_M = (I_m^{(M)}, m \leq \kappa_M)$ a partition of $[0, 1]$ into κ_M bins, with $\kappa_{(\cdot)} : \mathbb{N} \rightarrow \mathbb{N}$ a strictly increasing sequence. Given \mathcal{I}_M , we consider the model of piece-wise constant densities as the set, \mathcal{F}_M , of densities with respect to Lebesgue measure, in the form:

$$f_\omega = \sum_{m=1}^{\kappa_M} \frac{\omega_m}{|I_m^{(M)}|} \mathbb{1}_{I_m^{(M)}}, \quad \min_m \omega_m \geq 0, \quad \sum_{m=1}^{\kappa_M} \omega_m = 1, \quad (8)$$

where $|I_m^{(M)}|$ denotes the length of interval $I_m^{(M)}$. For instance, one could consider a sequence of dyadic partitions with $\kappa_M = 2^M$. Such partitions are admissible for sufficiently large M in the sense we detail below, and are used for the empirical investigation of Section 5.

The parameters for this model are then Q and $\underline{\omega}_{(M)} = (\omega_{mr})_{m=1, \dots, \kappa_M}^{r=1, \dots, R}$, the latter of which varies in the set

$$\Omega_M = \{\underline{\omega}_{(M)} \in [0, 1]^{\kappa_M \times R} ; \sum_{m=1}^{\kappa_M} \omega_{mr} = 1, \forall r \in [R]\}.$$

Through (8), we identify each $\underline{\omega}_{(M)} \in \Omega_M$ with a vector of emission densities $\mathbf{f}_{\underline{\omega}_{(M)}} \in \mathcal{F}_M^R$, and thus a prior distribution Π_M over the parameter space $\mathcal{Q} \times \Omega_M$ is identified with a prior distribution Π over $\mathcal{Q} \times \mathcal{F}_M^R$. The corresponding posterior distribution is denoted $\Pi_M(\cdot | Y_{1:n})$ and is defined through (4).

Throughout this section, we write $\underline{\omega}_r := (\omega_{mr}, m \leq \kappa_M)$ and we denote by $\mathcal{M}(\mathcal{I}_M)$ the hidden Markov model associated with densities of the form (8). Note that, for all $M > 0$, $\mathcal{M}(\mathcal{I}_M)$ is of dimension $(\kappa_M - 1)^R R(R - 1) < \infty$.

A key argument used in [24], to identify Q^* from $g_{Q^*, \mathbf{f}^*}^{(3)}$, is to find a partition \mathcal{I}_M for some $M > 0$ such that the matrix

$$\mathbf{F}^*(\mathcal{I}_M) := \begin{pmatrix} F_1^*(I_1^{(M)}), & \dots, & F_1^*(I_{\kappa_M}^{(M)}) \\ F_2^*(I_1^{(M)}), & \dots, & F_2^*(I_{\kappa_M}^{(M)}) \\ \vdots & \dots & \vdots \\ F_R^*(I_1^{(M)}), & \dots, & F_R^*(I_{\kappa_M}^{(M)}) \end{pmatrix}$$

has full rank. We call such a partition an *admissible* partition.

Definition 1. A partition \mathcal{I}_M is said to be *admissible* for f_1^*, \dots, f_R^* if the rank of $\mathbf{F}^*(\mathcal{I}_M)$ is equal to R .

Remark 1. Although we are using piece-wise constant functions to model the emission densities, we do not assume that the f_r^* , $r \in [R]$ are piece-wise constant and the simplified models M are not meant to lead to good approximation of the emissions densities. However, as far as the parameter Q is concerned, the likelihood induced by such a model is not misspecified but instead corresponds to a coarsening of the data. Specifically, it corresponds to the likelihood associated to observations of $Y_i^{(M)} = (\mathbb{1}_{I_1^{(M)}}(Y_i), \dots, \mathbb{1}_{I_{\kappa_M}^{(M)}}(Y_i))$; for such observations this simplified model leads to a well-specified likelihood. Note also that, although we are modelling densities with respect to Lebesgue measure, we do not require F_r^* to have density with respect to Lebesgue measure since the quantities of importance are the probabilities $F_r^*(I_m^{(M)})$, $r \in [R]$, $m \in [\kappa_M]$.

This particular coarsening was introduced in the context of mixtures by [27]. It is not at all obvious how other types of coarsening can be found, which are associated with well-specified parametric models for the coarsened data and the parameter of interest Q .

3.1. Asymptotic normality and Bernstein-von Mises

In this section, we study the asymptotic behaviour of the marginal posterior distribution of Q under model $\mathcal{M}(\mathcal{I}_M)$, together with the asymptotic normality of the maximum likelihood estimator $\hat{Q}_{n,M}$ in this model. As mentioned in Remark 1, the likelihood associated to model $\mathcal{M}(\mathcal{I}_M)$ is given by

$$L_n(Q, \omega_{(M)}) = \sum_{i_1, \dots, i_n=1}^R p_Q(i_1) \left(\prod_{t=1}^{n-1} Q_{i_t, i_{t+1}} f_{\omega_{i_t}}(Y_t^{(M)}) \right) f_{\omega_{i_n}}(Y_n^{(M)}), \quad (9)$$

with the abuse of notation $f_{\omega_{i_t}}(Y_t^{(M)}) = \sum_{m=1}^{\kappa_M} \frac{\omega_{m i_t} Y_t^{(M)}(m)}{|I_m^{(M)}|}$ for $t = 1, \dots, n$.

In other words, our likelihood becomes one of a hidden Markov model with finite state space and multinomial emission distributions. Under \mathbb{P}_* , the coarsened observations $Y_{1:n}^{(M)}$ arise from a hidden Markov model with multinomial emission distributions, governed by parameters $\omega_{(M)}^*$ and transition matrix Q^* , with $\omega_{mr}^* = F_r^*(I_m^{(M)})$. We write $\theta = \theta^{(M)} := (Q, \omega_{(M)})$ for the parameter of $\mathcal{M}(\mathcal{I}_M)$, and suppress the superscript on θ when the dependence on M is clear.

Asymptotic normality of the maximum likelihood estimator (MLE) of parametric finite state space hidden Markov models was considered for instance in [9], who showed that the MLE was asymptotically normal with covariance matrix given by the inverse of the Fisher information matrix, which is given by the limiting covariance matrix of the score statistics; see Lemma 1 of [9].

Let $J_M(\theta)$ be the Fisher information matrix associated to the likelihood (9):

$$J_M(\theta^*) = \lim_{n \rightarrow \infty} \frac{\mathbb{E}_* \left[-D^2 \log L_n(Q^*, \omega_{(M)}^*) \right]}{n},$$

as defined in [9]. We also write $J_M[Q, Q]$, $J_M[\omega, \omega]$, and $J_M[Q, \omega]$ to denote the sub-matrices corresponding to the second derivatives with respect to (Q, Q) , (ω, ω) and (Q, ω) respectively.

The following theorem demonstrates that asymptotically normal, parametric-rate estimators of the transition matrix exist, although such estimators may not have optimal asymptotic variance.

Theorem 1. *Let $M > 0$ and \mathcal{I}_M be an admissible partition for \mathbf{F}^* , and let $\hat{Q}_{n,M} = \hat{Q}_n^{(M)}(Y_{1:n})$ be the MLE in model $\mathcal{M}(\mathcal{I}_M)$, given observations $Y_{1:n}$. Grant **Assumption 1**, and assume also that the transition matrix Q^* is irreducible and aperiodic, and that $\omega_{mr}^* > 0$ for all $m \in [\kappa_M]$ and all $r \in [R]$. We then have (up to label-swapping)*

$$\sqrt{n}(\hat{Q}_{n,M} - Q^*) \Rightarrow^* N(0, \tilde{J}_M^{-1}),$$

where \tilde{J}_M^{-1} is positive definite and

$$\tilde{J}_M^{-1} = (J_M^{-1})[Q, Q] = J_M[Q, Q] - J_M[Q, \omega]J_M[\omega, \omega]^{-1}J_M[\omega, Q].$$

The main difficulty in the proof of Theorem 1 is showing that the Fisher information matrix J_M is invertible.

Proof. The model $\mathcal{M}(\mathcal{I}_M)$ is a regular parametric HMM, hence using Theorem 1 of [9], we establish asymptotic normality of the MLE for the parameter $\theta = (Q, \underline{\omega}_{(M)})$ as soon as J_M is invertible. We prove invertibility of J_M in Section C.1, then projecting onto the Q co-ordinates gives the result. \square

To derive the Bernstein-von Mises Theorem associated to model $\mathcal{M}(\mathcal{I}_M)$, we need the following assumption on the prior distributions on Q and $\underline{\omega}_{(M)}$:

- Assumption 2.** (i) The prior Π_Q on Q has positive and continuous density on \mathcal{Q} .
 (ii) The prior on $\underline{\omega}_{(M)}$ has positive and continuous density with respect to Lebesgue measure.

Theorem 2. *Let $M > 0$, \mathcal{I}_M be an admissible partition for \mathbf{f}^* , and let Π_M be a prior satisfying **Assumption 2**. Grant **Assumption 1**, and assume also that the transition matrix Q^* is irreducible and aperiodic, and that $\omega_{mr}^* > 0$ for all $m \in [\kappa_M]$ and all $r \in [R]$. We then have (up to label-swapping)*

$$\|\Pi_M(\sqrt{n}(Q - \hat{Q}_{n,M}) \in (\cdot) | Y_{1:n}) - N(0, \tilde{J}_M^{-1})\|_{TV} = o_{\mathbb{P}_*}(1).$$

As with the proof of Theorem 1, parametric results apply as soon as the Fisher information matrix is seen to be invertible.

Proof. Inspecting the proof of Theorem 1 of [9], we see (for $\theta = (Q, \underline{\omega}_{(M)})$, $J_M = J_M(\theta^*)$ the Fisher information and $l_n = \log g_n$ the log-likelihood) that

$$\sqrt{n}(\hat{\theta} - \theta^*) = n^{-\frac{1}{2}}J_M^{-1}\nabla_{\theta}l_n(\theta^*) + o_{\mathbb{P}_*}(1),$$

which up to $o_{\mathbb{P}_*}(1)$ is equal to the T_n of [17] at which their Bernstein-von Mises result (Theorem 2.1) is centred. Then this result implies the total variation convergence, in \mathbb{P}_* -probability, of the posterior to the given normal distribution, by considering the marginal posterior on Q . We remark further that the MLE $\hat{\theta}$ is regular, which follows from the characterisations of Fisher information in Lemmas 1 and 2 of [9], the expansion of the MLE above, and an application of Le Cam's third lemma (Example 6.7 of [55]) along the lines of the proof of Lemma 8.14 in [55]. \square

An interesting feature of Theorems 1 and 2 is that they essentially only require that \mathcal{I}_M is an admissible partition of \mathbf{F}^* . For a given partition, this is an assumption on \mathbf{F}^* , and indeed the choice of the partition is important. Note, however, that under Assumption 1 and, for instance, if the (Lebesgue) densities f_r^* are positive and continuous, then for all sequences of embedded partitions $(\mathcal{I}_M)_M$ with radius going to 0 there exists an M such that \mathcal{I}_M is admissible for \mathbf{F}^* . More discussion is provided in Section 5.

Combining Theorems 1 and 2, we see that credible regions for Q , based on the posterior associated to Π_M , are also asymptotic confidence regions. Their size may not be optimal however, even asymptotically. To ensure that such credible regions have optimal size while being asymptotic confidence regions, we would require that \tilde{J}_M^{-1} is the best possible (asymptotic) variance, but this is not true in general; while \tilde{J}_M^{-1} is the efficient covariance matrix for the estimation of Q in model $\mathcal{M}(\mathcal{I}_M)$, it is not necessarily the semiparametric efficient covariance matrix for the estimation of Q in model (3). The existence of an efficient estimator of Q in the semiparametric hidden Markov model, with likelihood (3), has not been established, although the fact that \sqrt{n} -convergent estimates of Q exist in the literature (see Section C.1) indicate that semiparametric efficient estimation of Q should be possible.

In the following section we construct an efficient estimator for Q and a prior leading to an efficient Bernstein-von Mises theorem.

3.2. Efficient estimation

In [27], in the context of semiparametric mixture models with finite state space, the authors also consider the prior model (8) for the emission distributions. They derive for fixed M a Bernstein-von Mises theorem similar to Theorem 2 and they show that, if we let $M = M_n$ increase to infinity sufficiently slowly as $n \rightarrow \infty$, and if the corresponding partitions \mathcal{I}_{M_n} are embedded with radius going to 0, then \tilde{J}_{M_n} converges to the efficient semiparametric Fisher matrix for the mixture weights. This allows them to define an efficient estimator, whose value given n observations is the maximum likelihood estimator in the histogram model with M_n bins.

In this section we prove a similar result, however the proof is significantly more involved than in the case of mixture models. To study the theory on efficient semiparametric estimation of Q in the semiparametric HMMs, we follow the approach of [46]. We first prove the LAN expansion in local sub-models,

which allows us to describe the tangent space, and then we prove that \tilde{J}_M converges to \tilde{J} as M goes to infinity where \tilde{J} is the semiparametric efficient Fisher information matrix for Q . Throughout Section 3.2 we assume that the F_r^* have density f_r^* with respect to Lebesgue measure on $[0, 1]$ and that the following holds:

Assumption 3. For all $i, j \in [R]$, $Q_{ij}^* > 0$.

3.2.1. Scores and tangent space in the semiparametric model

We begin by exhibiting the local asymptotic normality (LAN) expansion for our model, following the framework of [46]. As is usual in semiparametric efficiency arguments, this involves identifying the score functions and LAN norm along one-dimensional sub-models passing through the true parameter; see Chapter 25 of [55] for discussion in the iid setting. Since these sub-models are themselves parametric, this identification can be made by following the framework of [19], who considered asymptotic normality in the context of parametric HMMs. A more thorough treatment, alongside a recollection of the relevant definitions and results in [46], is provided in [48], we aim to give an overview.

The first step is to exhibit a LAN expansion for the semiparametric model we consider. The parameter space is $\Theta = \mathcal{Q} \times \mathcal{F}^R$ and we consider the tangent space, in the sense of Definition 1 of [46], at the parameter (Q^*, \mathbf{f}^*) as

$$\mathcal{H} = \{(a, \mathbf{h}), a \in \mathbb{R}^{R \times (R-1)}, \mathbf{h} = (h_r)_{r=1}^R, h_r \in \mathcal{H}_r\},$$

where $\mathcal{H}_r = \mathcal{H}_r(f_r^*) = \{h_r \in L^2(f_r^*) : \int h_r f_r^* = 1, \|h_r\|_\infty < \infty\}$, and where we use the parameterisation \tilde{Q} defined in Section 2.1. Write

$$\theta_n(a, \mathbf{h}) = \left(Q^* + \frac{a}{\sqrt{n}}, f_1^* \left(1 + \frac{h_1}{\sqrt{n}}\right), \dots, f_R^* \left(1 + \frac{h_R}{\sqrt{n}}\right) \right).$$

Then $\theta_n(a, \mathbf{h})$ is a perturbation of the parameter $\theta^* = (Q^*, \mathbf{f}^*)$ along the path characterised by a given element $(a, \mathbf{h}) \in \mathcal{H}$. Consider the sub-model

$$\Theta_{a, \mathbf{h}} = \{\theta_t = (Q^* + ta, (f_r^*(1 + th_r) : r \in [R])) , |t| \leq \epsilon\},$$

for some ϵ (depending on a, \mathbf{h}) sufficiently small that $\Theta_{a, \mathbf{h}} \subset \Theta$. Then, for given a, \mathbf{h} and for sufficiently large n , the perturbed parameters $\theta_n(a, \mathbf{h})$ are elements of $\Theta_{a, \mathbf{h}}$. This means that, for each a, \mathbf{h} , we can make an asymptotic expansion of the log-likelihood ratio between $\theta_n(a, \mathbf{h})$ and θ^* by considering the likelihood in the sub-model $\Theta_{a, \mathbf{h}}$. To this end, we expand the gradient of the log-likelihood $\ell_n^{(a, \mathbf{h})}(t)$, in $\Theta_{a, \mathbf{h}}$ at $t_0 = 0$ as

$$\frac{1}{\sqrt{n}} \nabla_t \ell_n^{(a, \mathbf{h})}(t_0) = \frac{1}{\sqrt{n}} \sum_{k=1}^n \Delta_{k, \infty}^{(a, \mathbf{h})} + o_{L^2(\mathbb{P}_*)}(1),$$

where

$$\begin{aligned} \Delta_{k,\infty}^{(a,\mathbf{h})} &= \Delta_{k,\infty}^{(a,\mathbf{h})}(Y_{-\infty:k}) = \mathbb{E}_*[\phi(X_{k-1}, X_k, Y_k)|Y_{-\infty:k}] \\ &\quad + \sum_{i=-\infty}^{k-1} (\mathbb{E}_*[\phi(X_{i-1}, X_i, Y_i)|Y_{-\infty:k}] - \mathbb{E}_*[\phi(X_{i-1}, X_i, Y_i)|Y_{-\infty:k-1}]), \end{aligned} \tag{10}$$

and where, writing $a_{rR} = -\sum_s a_{rs}$ and $Q_{r,R} = 1 - \sum_s Q_{rs}$,

$$\phi(X_{k-1}, X_k, Y_k) = \sum_{r=1}^R \sum_{s=1}^R \mathbb{1}\{X_{k-1} = r, X_k = s\} \frac{a_{r,s}}{Q_{r,s}^*} + \sum_{s=1}^R \mathbb{1}\{X_k = s\} h_s(Y_k). \tag{11}$$

These formulae arise as an application of the results of Section 6.1 of [19] to the (parametric) model $\Theta_{a,\mathbf{h}}$.

We note that the contribution of the first term of the right hand side of (11), which may be rewritten as

$$\sum_{r=1}^R \sum_{s=1}^{R-1} \left(\frac{\mathbb{1}\{X_{k-1} = r, X_k = s\}}{Q_{rs}^*} - \frac{\mathbb{1}\{X_{k-1} = r, X_k = R\}}{Q_{rR}^*} \right) a_{rs},$$

to the expression defined in (10) is precisely the score function for estimation in the model $\Theta_{a,\mathbf{0}}$, in which the emission densities are fixed and known. We see that this is equal to $a^T S_{Q^*}(Y_{-\infty:k})$, for S_{Q^*} the score at Q^* in the $R \times (R - 1)$ -dimensional parametric model with known emissions and unknown Q . Specifically, S_Q takes the form

$$\begin{aligned} S_{Q^*}(r, s)(Y_{-\infty:k}) &= \frac{\mathbb{P}_*(X_{k-1} = r, X_k = s|Y_{-\infty:k})}{Q_{rs}^*} - \frac{\mathbb{P}_*(X_{k-1} = r, X_k = R|Y_{-\infty:k})}{Q_{rR}^*} \\ &\quad + \sum_{i=-\infty}^{k-1} \left\{ \frac{\mathbb{P}_*(X_{i-1} = r, X_i = s|Y_{-\infty:k}) - \mathbb{P}_*(X_{i-1} = r, X_i = s|Y_{-\infty:k-1})}{Q_{rs}^*} \right. \\ &\quad \left. - \frac{\mathbb{P}_*(X_{i-1} = r, X_i = R|Y_{-\infty:k}) - \mathbb{P}_*(X_{i-1} = r, X_i = R|Y_{-\infty:k-1})}{Q_{rR}^*} \right\}, \end{aligned} \tag{12}$$

where we index with $(r, s) \in [R] \times [R - 1]$ for convenience sake, but consider S_{Q^*} as a vector of length $R(R - 1)$.

We then rewrite (10), substituting also the expression in (11) as

$$\begin{aligned} \Delta_{k,\infty}^{(a,\mathbf{h})} &= a^T S_{Q^*} + \sum_{r=1}^R \left\{ \mathbb{P}_*(X_k = r|Y_{-\infty:k}) h_r(Y_k) \right. \\ &\quad \left. + \sum_{i=-\infty}^{k-1} (\mathbb{P}_*(X_i = r|Y_{-\infty:k}) - \mathbb{P}_*(X_i = r|Y_{-\infty:k-1})) h_r(Y_i) \right\}. \end{aligned} \tag{13}$$

To set notation for what follows, we will denote the r^{th} summand of the above display, for a given direction h_r , as

$$\begin{aligned}
 H_r(h_r)(Y_{-\infty:k}) &:= \mathbb{P}_*(X_k = r|Y_{-\infty:k})h_r(Y_k) \\
 &+ \sum_{i=-\infty}^{k-1} (\mathbb{P}_*(X_i = r|Y_{-\infty:k}) - \mathbb{P}_*(X_i = r|Y_{-\infty:k-1}))h_r(Y_i).
 \end{aligned}
 \tag{14}$$

The above discussion means that, through a Taylor expansion, we can write

$$\begin{aligned}
 \ell_n(\theta_n((a, \mathbf{h}))) - \ell_n(\theta_*) &= \frac{1}{\sqrt{n}} \sum_{k=1}^{\infty} \Delta_{k,\infty}^{(a,\mathbf{h})} - \frac{1}{2n} \left(-\frac{d^2}{dt^2} \Big|_{t=t_1} \ell_n^{(a,\mathbf{h})}(t) \right) \\
 &= \Delta_{n,(a,\mathbf{h})} - \frac{1}{2} \|(a, \mathbf{h})\|_{\mathcal{H}}^2 + o_{\mathbb{P}_*}(1),
 \end{aligned}$$

with $t_1 \in (0, n^{-1/2})$ and by choosing $\|(a, \mathbf{h})\|_{\mathcal{H}}^2$ the Fisher information at $t = 0$ in the model $\Theta_{a,\mathbf{h}}$, which is defined in [19] as the L^2 norm¹ of $\Delta_{0,\infty}^{(a,\mathbf{h})}$. We have $\Delta_{n,(a,\mathbf{h})} = \frac{1}{\sqrt{n}} \sum_{k=1}^{\infty} \Delta_{k,\infty}^{(a,\mathbf{h})}$, is linear in (a, \mathbf{h}) and satisfies $\Delta_{n,(a,\mathbf{h})} \rightarrow N(0, \|(a, \mathbf{h})\|_{\mathcal{H}}^2)$, by the discussion directly preceding Theorem 2 in [19]. We also used above the local uniform convergence of the second derivative of the score to the Fisher information matrix at $t = 0$ in $\Theta_{a,\mathbf{h}}$, which is guaranteed by Theorem 3 of [19].

The preceding discussion shows that our model is LAN, which is the first step in understanding efficient estimators of a parameter of interest. In our case, the parameter of interest will be $v_n(P_{n,\theta_n(a,\mathbf{h})}) = Q$, which has ‘derivative’ $\dot{v}(a, \mathbf{h}) = a \in \mathbb{R}^p$ in the sense of [19], with $p = R(R - 1)$.

Following what precedes, we apply the convolution theorem also described in [46], and originally proven in [56]. This theorem essentially states that the limiting law of a regular estimator is lower bounded by that of a Gaussian random variable, whose covariance matrix is the covariance of the *efficient influence function*, which is itself characterised by the tangent space and the parameter ‘derivative’ \dot{v} .

The arguments are similar to those used in the iid setting, see for instance Section 25 of [55]. The finer details of the application of [55] to our context are given in [48]. This essentially involves identifying influence functions \dot{v}_b^T for estimation of one-dimensional functions $b^T Q$ at $b^T Q^*$, as we vary $b \in \mathbb{R}^p$. By considering the elements for which $a = 0$, we first find that \dot{v}_b^T is in the orthogonal complement (with respect to the LAN norm on \mathcal{H}) of the linear span of the scores in the models $\Theta_{0,\mathbf{h}}$, which is the span of the $H_r(h_r)$ as we vary r and $h_r \in \mathcal{H}_r$. Write \mathcal{A} for the projection onto this space, and write

$$\tilde{S}_Q = S_Q - \mathcal{A}S_Q
 \tag{15}$$

¹The L^2 norms of all $\Delta_{k,\infty}^{(a,\mathbf{h})}$ coincide by stationarity.

for the projection of the score function S_Q onto the orthogonal complement of this space, which we call the efficient score. By making further standard arguments, we find that the influence function \dot{v}_b^T has variance $b^T \tilde{J}^{-1} b$, where \tilde{J} is the covariance matrix of the efficient score \tilde{S}_{Q^*} , which is the efficient information matrix. By varying b , this characterises the optimal limiting covariance matrix as \tilde{J}^{-1} , the inverse efficient information matrix.

In the following section, we show that \tilde{J}_M converges to \tilde{J} when M goes to infinity and that, if M_n is a sequence going to infinity slowly enough, the sequence of estimators $(\hat{Q}_{n, M_n})_{n \geq 1}$ is asymptotically efficient and the posterior distributions $\Pi_{M_n}(dQ|Y_{1:n})$ satisfy the efficient Bernstein-von Mises theorem.

3.2.2. Approximation by the models $\mathcal{M}(\mathcal{I}_M)$

To prove that \tilde{J}_M converges to \tilde{J} , we require some additional assumptions on the true generating process and the partitions:

Assumption 4. For all $r = 1, \dots, R$, the densities f_r^* are continuous, (Lebesgue) almost-everywhere positive, linearly independent and

$$\forall r, s \leq R \quad 0 < \inf_x \frac{f_r^*(x)}{f_s^*(x)} < \sup_x \frac{f_r^*(x)}{f_s^*(x)} < \infty.$$

We also consider a sequence of partitions with vanishing radius.

Assumption 5. The sequence $(\mathcal{I}_M)_{M \geq R}$ of partitions is embedded, and

$$\max_{m \leq \kappa_M} \text{Vol}(I_m^{(M)}) \rightarrow 0, \quad \text{as } M \rightarrow \infty,$$

where $\text{Vol}(A)$ denotes the Lebesgue measure of A .

Assumption 5 is directly under the control of the practitioner, and is verified for instance by considering nested dyadic partitions, as is done in Section 5. **Assumption 4** is an assumption on the data generating process, although is very common in the context of semiparametric efficiency. Note also that, under **Assumptions 4** and **5**, there exists $M_0 > 0$ for which \mathcal{I}_{M_0} is admissible for \mathbf{f}^* .

In Proposition 1, we show that the efficient Fisher information matrix is precisely the limit of \tilde{J}_M . To do this, we first define the efficient scores $\tilde{S}_{Q^*}^{(M)}$ for the histogram models, whose covariance matrix will be \tilde{J}_M , analogously to what was done in Section 3.2.1 in the context of the full semiparametric model.

Recall, from Remark 1, that the likelihood in the model with piece-wise constant emissions corresponds to a likelihood of a HMM with multinomial emission distributions and observations $Y_i^{(M)} = (\mathbb{1}_{I_1^{(M)}}(Y_i), \dots, \mathbb{1}_{I_{\kappa_M}^{(M)}}(Y_i))$, and so we emphasize in our notation that the scores in $\mathcal{M}(\mathcal{I}_M)$ depend only on these summaries. A straightforward adaptation of the earlier presentation then leads

us to define the score functions for Q in $\mathcal{M}(\mathcal{I}_M)$ by

$$\begin{aligned}
 S_{Q^*}^{(M)}(r, s)(Y_{-\infty:k}^{(M)}) &= \frac{\mathbb{P}_*(X_{k-1} = r, X_k = s|Y_{-\infty:k}^{(M)})}{Q_{rs}^*} - \frac{\mathbb{P}_*(X_{k-1} = r, X_k = R|Y_{-\infty:k}^{(M)})}{Q_{rR}^*} \\
 &+ \sum_{i=-\infty}^{k-1} \left\{ \frac{\mathbb{P}_*(X_{i-1} = r, X_i = s|Y_{-\infty:k}^{(M)}) - \mathbb{P}_*(X_{i-1} = r, X_i = s|Y_{-\infty:k-1}^{(M)})}{Q_{rs}^*} \right. \\
 &\quad \left. - \frac{\mathbb{P}_*(X_{i-1} = r, X_i = R|Y_{-\infty:k}^{(M)}) - \mathbb{P}_*(X_{i-1} = r, X_i = R|Y_{-\infty:k-1}^{(M)})}{Q_{rR}^*} \right\}. \tag{16}
 \end{aligned}$$

for $(r, s) \in [R] \times [R - 1]$. For the model \mathcal{I}_M , the perturbations on the emissions vary in

$$\mathcal{H}_{r,M} := \{h = \sum_{m=1}^{\kappa_M} \alpha_m \mathbf{1}_{I_m}(y), \alpha_m \in \mathbb{R}, \int h f_{r,M}^*(y) dy = 0\}.$$

and we then define, for $h_r \in \mathcal{H}_{r,M}$,

$$\begin{aligned}
 H_{r,M}(h_r)(Y_{-\infty:k}^{(M)}) &= \mathbb{P}_*(X_j = r|Y_{-\infty:k}^{(M)})h_r(Y_k) \\
 &+ \sum_{i=-\infty}^{k-1} \left(\mathbb{P}_*(X_i = r|Y_{-\infty:k}^{(M)}) - \mathbb{P}_*(X_i = r|Y_{-\infty:k-1}^{(M)}) \right) h_r(Y_i). \tag{17}
 \end{aligned}$$

Following again the arguments of Section 3.2.1, we write \mathcal{A}_M for the projection onto the space spanned by the $H_{r,M}(h_r)$, and finally define

$$\tilde{S}_{Q^*}^{(M)} = S_{Q^*}^{(M)} - \mathcal{A}_M S_{Q^*}^{(M)},$$

whose covariance matrix is \tilde{J}_M , the efficient Fisher information matrix for the model $\mathcal{M}(\mathcal{I}_M)$.

The convergence of $\tilde{S}_{Q^*}^{(M)}$ to \tilde{S}_{Q^*} is then established by a martingale argument. To show convergence of the projection operators \mathcal{A}_M to \mathcal{A} is more involved; we use a deconvolution argument which shows that boundedness in the space of the nuisance scores $H_r(h_r)$ implies boundedness of the h_r in the index space \mathcal{H}_r (and likewise for $H_{r,M}(h_r), \mathcal{H}_{r,M}$). These intermediate arguments, and the proof of the following results, are in Section 6.

Proposition 1. *Grant Assumptions 1, 3–5. Then*

$$\tilde{S}_{Q^*}^{(M)} \rightarrow \tilde{S}_{Q^*}, \quad M \rightarrow \infty$$

where the convergence is in $L^2(\mathbb{P}_*)$. Moreover, we have $\tilde{J}_M \rightarrow \tilde{J}$ as $M \rightarrow \infty$ where $\tilde{J} \geq \tilde{J}_M$ is the efficient information for estimating Q in the full data model, and is invertible.

From Proposition 1, we deduce the following result.

Theorem 3. *Let $\hat{Q}_n = \hat{Q}_n^{(M_n)}$ where $M_n \rightarrow \infty$ sufficiently slowly. Then under **Assumptions 1, 3-5**, \hat{Q}_n is a \mathbb{P}_* -regular estimator and satisfies (up to label-swapping)*

$$\sqrt{n}(\hat{Q}_n - Q^*) \Rightarrow^* N(0, \tilde{J}^{-1})$$

where \tilde{J} is the variance of the efficient score function, as defined in Proposition 1. In particular, \hat{Q}_n is an efficient estimator of Q in the full semiparametric model. Moreover, for Π_n a sequence of priors placing its mass on models $\mathcal{M}(\mathcal{I}_{M_n})$ respectively, and satisfying **Assumption 2**, we have (up to label-swapping) that

$$\left\| \Pi_n(\sqrt{n}(Q - \hat{Q}_n) | Y_{1:n}) - \mathcal{N}(0, \tilde{J}^{-1}) \right\|_{TV} = o_{\mathbb{P}_*}(1).$$

4. Cut posterior contraction

In what follows, we study contraction rates for the cut posterior $\Pi_{cut}(\cdot | Y_{1:n})$ defined in Section 2.

4.1. Concentration of marginal densities $g_{Q,\mathbf{f}}^{(\ell)}$

In this section, we present Proposition 2 which controls $\Pi_{cut}(\|g_{Q,\mathbf{f}}^{(\ell)} - g_{Q^*,\mathbf{f}^*}^{(\ell)}\| \leq \epsilon_n | Y_{1:n})$. This result follows from Theorem 7, which is an adaptation of the general approach of [29] and is of independent interest; see Section B for full details.

For the sake of simplicity, we consider a prior in the form $\Pi_2(d\mathbf{f} | Q) = \Pi_2(d\mathbf{f})$; extension to the case where the prior Π_2 depends on Q is straightforward from Section B. Note that, in any case, the conditional posterior on \mathbf{f} given Q depends on Q through the likelihood.

Hence, similarly to [57], we consider the following assumptions used to verify the Kullback-Leibler condition.

Assumption 6. Let $\epsilon_n, \tilde{\epsilon}_n > 0$ denote two sequences such that $\tilde{\epsilon}_n \leq \epsilon_n = o(1)$ and $n\tilde{\epsilon}_n^2 \rightarrow \infty$. Assume that the prior Π_2 on \mathcal{F}^R satisfies the following conditions:

A. There exists $C_{\Pi_2} > 0$, depending on the choice of prior Π_2 and a sequence $S_n \subset \mathcal{F}^R$, such that $\Pi_2(S_n) \gtrsim \exp(-C_{\Pi_2} n \tilde{\epsilon}_n^2)$ and such that, for all $\mathbf{f} \in S_n$, there exists a set $S \subset \mathcal{Y}$ and functions $\tilde{f}_1, \dots, \tilde{f}_R$ satisfying, for all $1 \leq i \leq R$ and for $C_{R,Q} = 4 + \log(\frac{2R}{\min_{i,j} Q_{ij}^*}) + 10^4 \frac{R^2}{(\min_{i,j} Q_{ij}^*)^5}$,

$$\int_S \left[\frac{|f_i^*(y) - f_i(y)|^2}{f_i^*(y)} + \frac{|f_i^*(y) - \tilde{f}_i(y)|^2}{\tilde{f}_i(y)} \right] \lambda(dy) \leq \tilde{\epsilon}_n^2;$$

$$\int_S f_i^*(y) \max_{1 \leq j \leq R} \log \left(\frac{\tilde{f}_j(y)}{f_j(y)} \right) \lambda(dy) \leq \tilde{\epsilon}_n^2;$$

$$\int_{S^c} f_i^*(y) \max_{1 \leq j \leq R} \log \left(\frac{f_j^*(y)}{f_j(y)} \right) \lambda(dy) \leq C_{R,Q} \tilde{\epsilon}_n^2;$$

$$\int_{S^c} [\tilde{f}_i(y) + f_i^*(y)] \lambda(dy) \leq \tilde{\epsilon}_n^2.$$

B. For all constants $C > 0$, there exists a sequence $(\mathcal{F}_n)_{n \geq 1}$ of subsets of \mathcal{F}^R and a constant $C' > 0$ such that

$$\Pi_2(\mathcal{F}_n^c) = o(\exp(-Cn\tilde{\epsilon}_n^2)), \quad \text{and} \quad N\left(\frac{\epsilon_n}{12}, \mathcal{F}_n, d\right) \lesssim \exp(C'n\epsilon_n^2),$$

where $N\left(\frac{\epsilon_n}{12}, \mathcal{F}_n, d\right)$ is the $\frac{\epsilon_n}{12}$ -covering number of \mathcal{F}_n with respect to $d(f, \tilde{f}) = \max_{i=1, \dots, R} \|f_i - \tilde{f}_i\|_{L^1}$.

Proposition 2. Let $(Y_t)_{t \geq 1}$ be observations from a finite state space HMM with transition matrix Q and emission densities $\mathbf{f} = (f_r)_{r=1, \dots, R}$. Grant **Assumptions 1–5**, and consider the cut posterior obtained by choosing Π_1 as the prior Π_M of Theorem 2 associated to the admissible partition \mathcal{I}_M , and Π_2 such that **Assumption 6** is verified for suitable $\epsilon_n, \tilde{\epsilon}_n$ satisfying $n\epsilon_n^2 \gtrsim \log n$. Then, for any $K_n \rightarrow \infty$,

$$\Pi_{cut}(\{(Q, f) : \|g_{Q, \mathbf{f}}^{(3)} - g_{Q^*, \mathbf{f}^*}^{(3)}\|_{L^1} > K_n \epsilon_n\} | Y_n) = o_{\mathbb{P}_*}(1),$$

up to label-swapping.

Remark 2. Given Proposition 2, we may use the results of [57] to derive posterior contraction rates for a number of priors Π_2 . We explore the case that Π_2 is a (product of) Dirichlet process mixture of normals in Section 4.4, as this is a popular choice for density estimation.

Remark 3. We could also choose $\Pi_1 = \Pi_{1,n}$ equal to Π_{M_n} for some $M_n \rightarrow \infty$ sufficiently slowly to give the refined control over the marginal cut posterior on Q as described in Theorem 3.

4.2. Concentration of emission distributions

Proposition 2 applies the general contraction result of Theorem 7 to obtain an estimation result for the marginal distribution of the observations. Given that we already have control over the transition matrix in this setting when using Π_1 as a histogram prior, it remains to establish concentration rates for the emission distributions. Theorem 4 allows us to translate a rate on a marginal distribution into a corresponding rate on the emission distribution.

Theorem 4. Let the HMM satisfy the assumptions of Proposition 2, and let $g_{Q^*, \mathbf{f}}^{(3)}$ be the marginal density for three consecutive observations under the parameters (Q^*, \mathbf{f}) . Then there exists a constant $C = C(f^*, Q^*) > 0$ such that, for sufficiently small $\|g_{Q^*, \mathbf{f}}^{(3)} - g_*^{(3)}\|_{L^1}$,

$$\sum_{r=1}^R \|f_r - f_r^*\|_{L^1} \leq C \|g_{Q^*, \mathbf{f}}^{(3)} - g_*^{(3)}\|_{L^1}.$$

Remark 4. Theorem 4 provides an inversion inequality from L^1 to L^1 , which has interest in both Bayesian and frequentist estimation of emission densities. It is the first such result for the L^1 distance, with Theorem 6 of [15] establishing a similar inequality in the L^2 case. Given that the testing assumptions of Theorem 7 (and more generally, of results based on [29]) are much more straightforward to verify for the L^1 distance, our result has particular interest in the context of Bayesian (or pseudo-Bayesian) settings.

Using Proposition 2 and Theorem 4 together with Theorem 2, we easily deduce the following.

Theorem 5. *Let $Y_{1:n}$ be distributed according to the HMM with parameters Q^* and \mathbf{f}^* , and grant **Assumptions 1–5**. Let $\Pi_1 = \Pi_M$ and let Π_2 satisfy **Assumption 6**. Then, for ϵ_n as in Proposition 2 and any $K_n \rightarrow \infty$, we have (up to label-swapping)*

$$\Pi_{cut}(\{\|Q - Q^*\| > \frac{K_n}{\sqrt{n}}, \max_{r=1, \dots, R} \|f_r - f_r^*\|_{L^1} > K_n \epsilon_n\} | Y_{1:n}) = o_{\mathbb{P}_*}(1),$$

and

$$\|\Pi_{cut}(\sqrt{n}(Q - \hat{Q}_{n,M}) \in (\cdot) | Y_{1:n}) - N(0, \tilde{J}_M^{-1})\|_{TV} = o_{\mathbb{P}_*}(1).$$

4.3. Concentration of smoothing distributions

When clustering data, the *smoothing distribution* $\mathbb{P}_*(X_k = x | Y_{1:n})$ is often of interest. Our final main result concerns recovery of these probabilities using a (cut) Bayesian approach, establishing contraction of the posterior distribution over these smoothing distributions in total variation. We do so by combining novel arguments with the inequality given in Proposition 2.2 of [16]. We recall the notation $\theta = (Q, \mathbf{f})$.

Theorem 6. *Grant the assumptions of Theorem 5, together with*

$$\max_{r \in [R]} \|\sqrt{f_r^*}\|_1 < \infty, \tag{18}$$

and assume that

$$\Pi_2(\max_i \|f_i\|_{L_2} > e^{\gamma n \epsilon_n^2} | Q) \leq C' e^{-2\gamma n \epsilon_n^2}; \quad \max_i \|f_i^*\|_{L_2} \leq e^{\gamma n \epsilon_n^2}; \tag{19}$$

for some constants $\gamma, C' > 0$. Then, for ϵ_n as in Proposition 2 for which $n\epsilon_n^3 \rightarrow 0$ and for any $K_n \rightarrow \infty$, we have (up to label swapping)

$$\Pi_{cut}(\|\mathbb{P}_\theta(X_k = x | Y_{1:n}) - \mathbb{P}_*(X_k = x | Y_{1:n})\|_{TV} > K_n \epsilon_n | Y_{1:n}) = o_{\mathbb{P}_*}(1).$$

Remark 5. The requirement that $n\epsilon_n^3 \rightarrow 0$ is used in the proof, although we expect that this condition is not fundamental and it may be possible to weaken with appropriate proof techniques. It is not clear if assumption (19) is crucial, although it is anyway very weak and satisfied, for instance, by choosing a prior Π_2 for which $E_{\Pi_2}(\|f_r\|_{L_2}^2) < \infty$ for all r .

4.4. Example: Dirichlet process mixtures of Gaussians

In this section, we show that **Assumption 6** is verified when Π_2 is a Dirichlet process mixture of Gaussians. We use the results of Section 4 of [57], in which **Assumption 6** is verified for $\mathbf{f}^* = (f_r^*)_{r=1}^R$, with each f_r^* in the class

$$\mathcal{P}(\beta, L, \gamma) = \{f \in \mathcal{F} : \log f \text{ is locally } \beta - \text{H\"older with derivatives } l_j = (\log f)^{(j)} \text{ and } |l_{k_\beta}(y) - l_{k_\beta}(x)| \leq r!L(y)|y - x|^{\beta - k_\beta} \text{ when } |x - y| \leq \gamma\},$$

and satisfying weak tail assumptions which we give in **Assumption 8**. Here $\beta > 0$, L is a polynomial function, $\gamma > 0$ and $k_\beta = \lceil \beta \rceil - 1$ for $\lceil \cdot \rceil$ the usual ceiling operator.

The following result, which is a corollary of Theorem 4.3 of [57], shows that, for prior choices satisfying **Assumption 9**, the conditions of Theorem 5 are verified with

$$\epsilon_n = n^{-\frac{\beta}{2\beta+1}} \log(n)^t, \tag{20}$$

where $t > t_0 \geq \left(2 + \frac{2}{\gamma} + \frac{1}{\beta}\right) \left(\frac{1}{\beta} + 2\right)^{-1}$.

Corollary 1. *Let $Y_{1:n}$ be distributed according to the HMM with parameters Q^* and \mathbf{f}^* , and grant **Assumptions 1–5**. Let $\Pi_1 = \Pi_M$ and let the prior Π_2 on \mathcal{F}^R take the form of an R -fold product of Dirichlet process mixtures of Gaussians, in which the base measure α and variance prior Π_σ satisfy **Assumption 9**. Suppose further that the true emission distributions $(f_i^*)_{i=1}^R$ satisfy **Assumption 8** and that, for each i , $f_i^* \in \mathcal{P}(\beta, L, \gamma)$. Then Theorems 5 and 6 hold with ϵ_n as in (20), where $t > t_0 \geq \left(2 + \frac{2}{\gamma} + \frac{1}{\beta}\right) \left(\frac{1}{\beta} + 2\right)^{-1}$.*

Proof. The proof follows from [57] who showed that **Assumption 6** is verified under the stated assumptions with $\lambda(dy)$ the Lebesgue measure, appropriately chosen B_n , $\tilde{\epsilon}_n = n^{-\frac{\beta}{2\beta+1}} \log(n)^{t_0}$ with t_0 as in the statement, and ϵ_n as in the statement. □

The rate $n^{-\frac{\beta}{2\beta+1}}$ is minimax optimal for the classes $\mathcal{P}(\beta, L, \gamma)$ under iid sampling assumptions, see [45]. Although not proved here, we strongly believe the rate ϵ_n to be minimax up to $\log n$ factors, since an iid sampling assumption corresponds to estimating f_1, \dots, f_R after observing $X_{1:n}, Y_{1:n}$ which is easier than estimating f_1, \dots, f_R from $Y_{1:n}$ only. We thank the anonymous reviewer for pointing this out to us.

Remark 6. It is also possible to verify the conditions of Proposition 2 in the case of a countable observation space, again by following the example set out in [57]. In this case, **Assumption 6** is verified with $\lambda(dy)$ the counting measure on \mathbb{N} by using a Dirichlet process prior on the emissions whose base measure satisfies a tail condition, and under a tail assumption on the true emissions. In this case, a parametric rate up to log factors is obtained.

5. Practical considerations and simulation study

In this section, we discuss the practical implication of the method and results described in the previous sections². As described in Section 2.2, we consider a cut posterior approach where Π_1 is based on a histogram prior on the emission distributions and where Π_2 is a Dirichlet process mixture of normals for the emission densities. We first describe the implementation of $\Pi_1(Q|Y_{1:n})$.

5.1. MCMC algorithm for Q

We first define the construction of the partition \mathcal{I}_M . In Section 3, it is defined as a partition on $[0, 1]$ (or more generally on $[0, 1]^d$), but this is easily generalized to a partition of \mathbb{R} (or \mathbb{R}^d) via an appropriate C_1 diffeomorphism ϕ . In this section we restrict ourselves to dyadic partitions, i.e. $I_m^{(M)} = (\phi^{-1}((m-1)2^{-M}), \phi^{-1}(m2^{-M}))$, $m = 0, \dots, 2^M - 1$, where $\phi : \mathbb{R} \rightarrow [0, 1]$ is a C_1 diffeomorphism.

Recall that, for Theorems 1 and 2 to be valid, we merely require that $\mathcal{I}^{(M)} = (I_m^{(M)}, m \leq 2^M - 1)$ is admissible and that, for all $r \leq R, m \leq 2^M - 1$, $F_r^*(I_m^{(M)}) > 0$. Interestingly Theorem 3 does not restrict the choice of ϕ , but restricts the choice of M . In practice, however, the choice of M and ϕ matters and more details are provided below.

Once \mathcal{I}_M is chosen, we consider a prior on $Q, \underline{\omega}_M$ (suppressing the M henceforth) and, for the sake of computational simplicity, we consider the following family of Dirichlet priors:

$$\begin{aligned} \forall i \leq R, \quad Q_{i\cdot} &= (Q_{ij}, j \in [R]) \sim \mathcal{D}(\gamma_{i1}, \dots, \gamma_{iR}) \\ \omega_i &= (\omega_{im}, m \in [2^M]) \sim \mathcal{D}(\beta_{i1}, \dots, \beta_{i2^M}). \end{aligned}$$

We then use a Gibbs sampler on $(Q, \underline{\omega}, \mathbf{X})$ where given $\mathbf{X}, Y_{1:n}$,

$$\begin{aligned} Q_{i\cdot} &\overset{\text{ind}}{\sim} \mathcal{D}(\gamma_{i1} + n_{i1}, \dots, \gamma_{iR} + n_{iR}), \quad n_{ij} = \sum_{t=2}^n \mathbb{1}_{X_{t-1}=i, X_t=j} \\ \omega_i &\overset{\text{ind}}{\sim} \mathcal{D}(\beta_{i1} + N_0^{(i)}, \dots, \beta_{i2^M} + N_{2^M-1}^{(i)}), \quad N_m^{(i)} = \sum_{t=1}^n \mathbb{1}_{Y_t \in I_m^{(M)}, X_t=i}, \end{aligned}$$

and the conditional distribution of X given $Q, \underline{\omega}, Y$ is derived using the forward-backward algorithm (see [44] or [23]). To overcome the usual label-switching issue in mixtures and HMMs, we take the approach of Chapter 6 of [44] which deals with MCMC in the mixtures setting, in which the authors propose relabelling relative to the posterior mode as a post-processing step (with likelihood computed also with forward-backward).

²Code is provided at <https://github.com/dm1729/HMMs> .

In our simulation study, we have considered $R = 2$ hidden states with transition matrix and emission distributions

$$Q^* = \begin{pmatrix} 0.7 & 0.3 \\ 0.2 & 0.8 \end{pmatrix}, \quad F_1^* \sim \mathcal{N}(-1, 1), F_2^* \sim \mathcal{N}(1, 1).$$

We have considered data of size $n = 1000, 2500, 5000, 10000$, obtained by restricting a single simulated data set of size 10000. To study the effect of M we run our MCMC algorithm targeting $\Pi_1(\cdot|Y_{1:n})$ with $\kappa_M = 2, 4, 8, 16, 64, 128$ bins.

We took $\gamma_{ij} = \beta_{ij} = 1$ so that independent uniform priors over the simplex were used for rows of the transition matrix and the histogram weights, and chose

$$\phi(y) = \begin{cases} (1 + e^{-y})^{-1} & |y| > 3 \\ \zeta + \eta y & |y| \leq 3 \end{cases}$$

with $\zeta, \eta \in \mathbb{R}$ chosen so that ϕ is continuous. The sigmoid $y \mapsto (1 + e^{-y})^{-1}$ is a popular transformation from \mathbb{R} to $[0, 1]$. However, its steep gradient near $y = 0$ means that a uniform partition of $[0, 1]$ is mapped to a partition of \mathbb{R} concentrated tightly around $y = 0$. This in turn leads to an over-coarsening of the data, with many data points lying in the two unbounded bins of the partition which extend to $\pm\infty$. The linear interpolation in the range $|y| \leq 3$ is thus intended to provide more discrimination between data points lying in this interval, which includes the vast majority (approx 97.5%) of the data.

We ran the MCMC for 150000 iterations for each binning, discarding 10000 iterations as burn-in and retaining one in every twenty of the remaining draws, for a total of 7000 posterior draws.

The fitted distributions for Q under the priors Π_1 , with varying M , are shown in Figures 2, 3, 4 and 5, and demonstrate Theorems 2 and 3 as we detail below. To make an additional comparison with a typical Bayesian nonparametric approach, we also fitted a model with a prior Π' , which uses Dirichlet priors on the transition matrix as in Π_1 and Dirichlet process mixtures of Gaussians to model the emission densities. We defer further details of Π' to Section 5.2 as it is more relevant as a comparison with Π_2 , since both have a much higher computational cost in comparison to Π_1 . In Figure 6, we compare distributions for the transition matrix under $\Pi'(\cdot|Y_{1:n})$ with the distribution under $\Pi_1(\cdot|Y_{1:n})$ for a selection of values of M , but we emphasize the large difference in compute time.

Even when $\kappa_M = 2$, the posterior distribution under the prior Π_1 , as n increases, looks increasingly like a Gaussian, though its variance is quite large. For slightly larger values of κ_M , this Gaussian shape is preserved but with lower variance, demonstrating Proposition 8 which states that the Fisher information grows as we refine the partition. However, we can also see that taking κ_M too large leads to erroneous posterior inference, which may simply be biased (for instance when $n = 5000$ and we take $\kappa_M = 64$) or lose its Gaussian shape entirely (for instance when $n = 1000$ and $\kappa_M \in \{64, 128\}$). This demonstrates that

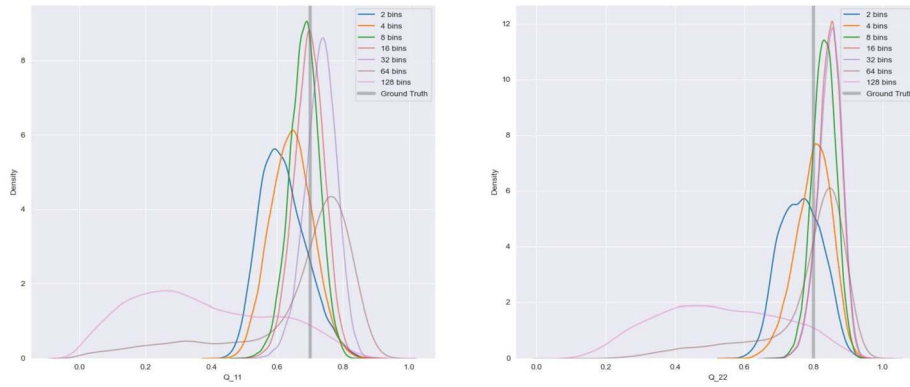


FIG 2. Posterior draws for Q under priors Π_1 with $\kappa_M \in \{2, 4, \dots, 128\}$ bins when $n = 1000$.

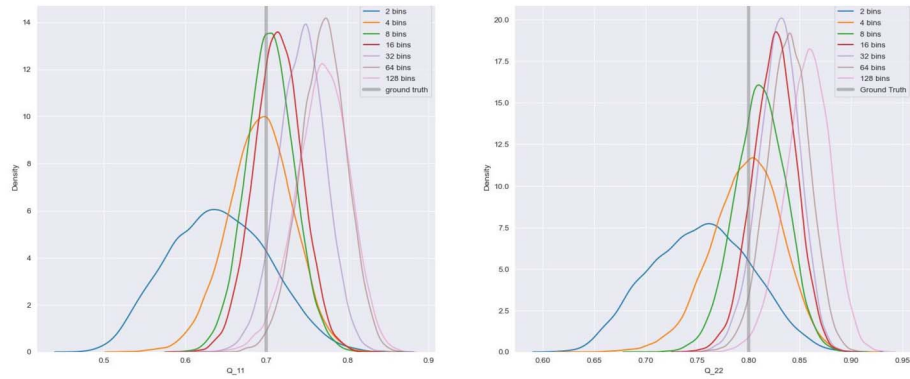


FIG 3. Posterior draws for Q under priors Π_1 with $\kappa_M \in \{2, 4, \dots, 128\}$ bins when $n = 2500$.




FIG 4. Posterior draws for Q under priors Π_1 with $\kappa_M \in \{2, 4, \dots, 128\}$ bins when $n = 5000$. The dotted orange line is the reference mean, taken when $\kappa_M = 4$. The dashed green lines are the bounds of the 90% credible set when $\kappa_M = 8$.




FIG 5. Posterior draws for Q under priors Π_1 when $n = 10000$, when using $\kappa_M \in \{2, 4, \dots, 128\}$ bins.

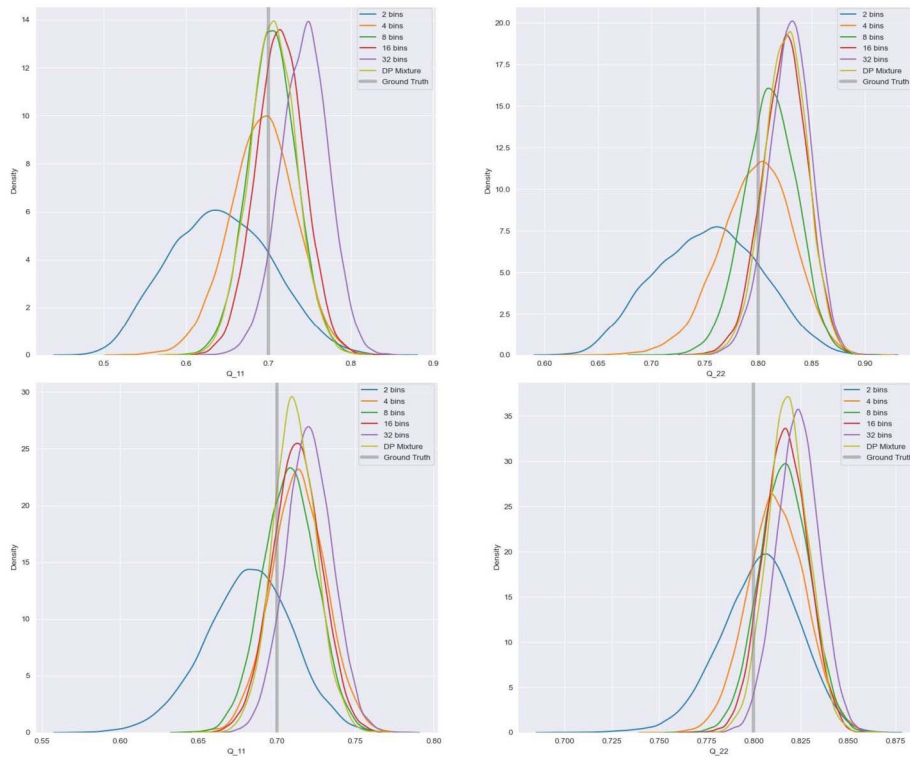


FIG 6. Posterior draws for Q under priors Π_1 (for $\kappa_M \in \{2, 4, 8, 16, 32\}$) and Π' (DP Mixture line) when $n = 2500$ (top) and $n = 10000$ (bottom).

the requirement of Theorem 3, that $M_n \rightarrow \infty$ sufficiently slowly, has practical consequences and is not merely a theoretical artefact.

While the latter issue (of a non-Gaussian shaped posterior) is easily diagnosed by eye, the former issue (of a biased, Gaussian-like posterior) is somewhat more worrisome, as the bias cannot be so easily identified when one does not have access to the true data generating process. In Figure 6, we see preliminary empirical evidence that this problem of bias may also be present when taking a fully Bayesian approach based on the prior Π' , indicating that our approach may work better than a fully Bayes approach if M_n can be tuned well. We emphasize however that this simulation study is very limited in scope, and is only intended to demonstrate our results rather than to make conclusive comparisons.

For the histogram prior Π_1 , we propose the following heuristic to tune κ_M after computing the posterior for a range of different numbers of bins. For a small number of bins (say $\kappa_M = 4$), take the posterior mean as a reference estimate (say \hat{Q}_0), which should have low bias and moderate variance. For higher κ_M values, compare \hat{Q}_0 κ_M -specific posterior mean and the $1 - \alpha$ credible sets C_α , for $\alpha = 0.05, 0.1$ say. We consider that κ_M is not too large if \hat{Q}_0 is well within the bounds of C_α . We have added some further markings to Figure 4 which illustrate this approach.

We emphasize that it is most important not to refine the partition too quickly, so even when adopting this heuristic one may wish to favour lower values of κ_M . When fitting Π_2 as discussed in Section 5.2, we used posterior draws based on one possible refinement, in which $\kappa_M = 4$ for $n = 1000$, $\kappa_M = 8$ for $n \in \{2500, 5000\}$ and $\kappa_M = 16$ for $n = 10000$ – see Figure 7.

5.2. MCMC algorithm for $\{f_r\}_{r=1}^R$

After simulating draws from the marginal posterior Π_1 of Q , we then use a Gibbs sampler to target the cut-posterior distribution $\Pi_2(f|Q, Y)$, for which the prior $\Pi_2(f|Q)$ is a Dirichlet mixture as follows. For each hidden state r (whose labels are fixed after relabelling in Section 5.1), we use independent dirichlet Process mixtures of normals to model f_r :

$$f_r(y) = \int \phi_{\sigma^{(r)}}(y - \mu) P^{(r)}(d\mu), \quad P^{(r)} \stackrel{iid}{\sim} DP(M_0, \mathcal{N}(\mu_c, \sigma_c^2)), \quad (21)$$

with $M_0 = 1$, $\mu_c = 0$ and $\sigma_c^2 = 1$. The kernels ϕ_σ are centred normal distributions with variance $v = \sigma^2$, and $v^{(r)}$ is equipped with an InvGamma($\alpha_\sigma, \beta_\sigma$) prior³. Note that the scale parameter σ is fixed across μ 's for each r . The MCMC procedure will rely on the following stick breaking representation of (21) involving mixture allocation variables $s^{(r)} = (s_i^{(r)})_{i \leq n}$ and stick breaks $(V_j^{(r)})_{j \in \mathbb{N}}$

³We remark that Proposition 1 verifies the conditions of Proposition 2 under an inverse gamma prior on the standard deviation. We have instead adopted in our simulations the common practice of placing an inverse gamma prior on the variance for computational convenience.

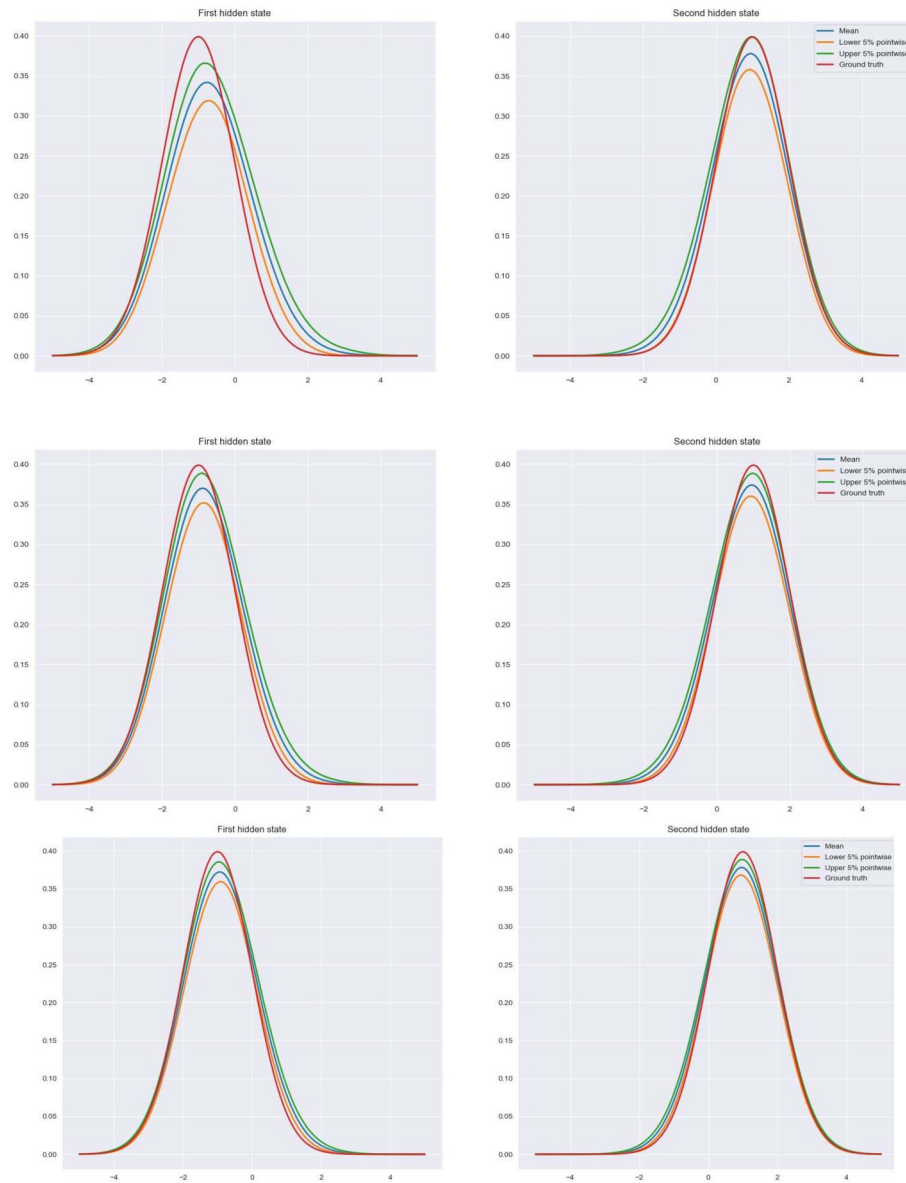


FIG 7. Cut posterior mean and pointwise 90% credible bands for emissions when refining the partition with $n = 2500$ (top), $n = 5000$ (middle) and $n = 10000$ (bottom), each with $C = 10$ interior iterations. See top plot of Figure 8 for $n = 1000$.

associated to a vector of weights $(W_j^{(r)})_{j \in \mathbb{N}}$:

$$\begin{aligned}
 Y_i | \mathbf{X}, \mu, s, \sigma, W &\stackrel{iid}{\sim} \phi_{\sigma(r)}(\cdot - \mu_{s_i}^{(X_i)}); & s_i | V, \mathbf{X}, \mu, \sigma^2 &\stackrel{indep}{\sim} W^{(X_i)}, \\
 W_j^{(r)} &= V_j^{(r)} \prod_{i < j} (1 - V_i^{(r)}); & \mathbf{X} &\sim \text{MarkovChain}(Q) \\
 V_j^{(r)} &\stackrel{iid}{\sim} \beta(1, M_0); & \mu_j^{(r)} &\sim \mathcal{N}(\mu_c, \sigma_c^2); & v^{(r)} &\sim \text{InvGamma}(\alpha_\sigma, \beta_\sigma).
 \end{aligned}
 \tag{22}$$

In order to approximately sample from the corresponding posterior, we replace the above prior with the Dirichlet-Multinomial process (see Theorem 4.19 of [30]) with truncation level $S_{\max} = \lfloor \sqrt{n} \rfloor$, in which W is instead sampled from a Dirichlet distribution with parameter

$$\alpha = \left(\frac{M_0}{S_{\max}}, \frac{M_0}{S_{\max}}, \dots, \frac{M_0}{S_{\max}} \right).
 \tag{23}$$

This truncation level is suggested as a rule of thumb in the remark at the end of Section 5.2 of [30]. Further details on implementation can be found in Section E.

Nested MCMC For the i^{th} draw from the cut posterior, we would ideally sample first Q_i from $\Pi_1(\cdot | Y)$, and then $f_i = (f_{ir})_r$ from $\Pi_2(\cdot | Q_i, Y)$. For the first step, we used Algorithm SA1 with burn in and thinning. However, we encounter difficulty when simulating from $\Pi_2(\cdot | Q_i, Y)$, as the Q_i changes when i changes, and so an MCMC approach effectively needs to mix i by i . In order to achieve this, we run a nested MCMC approach (see [49]) in which, for each i , an interior chain of length C is run, taking the final draw from the interior chain as our i^{th} global draw.

Concerns about the computational cost of nested MCMC have been expressed, especially when there is strong dependence between the two modules (in our case, representing the transition and emissions) as in Section 4.2 of [40]. However, we found little improvement beyond $C = 10$ interior iterations, see Figure 8. We expect that this is down to the well localised posterior $\Pi_1(\cdot | Y)$ on Q which means that the Q_i , and hence the targets $\Pi_2(\cdot | Q_i, Y)$, don't vary too much i by i . Plots of the posterior mean when $C = 10$ are detailed in Figure 7. Since we used the thinned draws of Q from Π_1 , we ran 7000 such interior chains.

When using such a small number of iterations for the interior chain, nested MCMC is not costly compared to other cut sampling schemes (see e.g. Table 1 of [40]). We further suggest that the computational cost compared to fully Bayesian approaches with fixed targets is not as high as it may seem, given that the use of such interior chains should, at least partially, subsume the need for thinning. Indeed running such interior chains with a fixed target for each i would be precisely the same as thinning. A more detailed development of this idea can be found in Section 4.4 of the supplement to [12].

Comparison to fully Bayesian approach As mentioned in Section 5.1, we also consider the fully Bayesian model with prior Π' as a means of comparison. The prior Π' independently places a Dirichlet prior on the rows of the transition

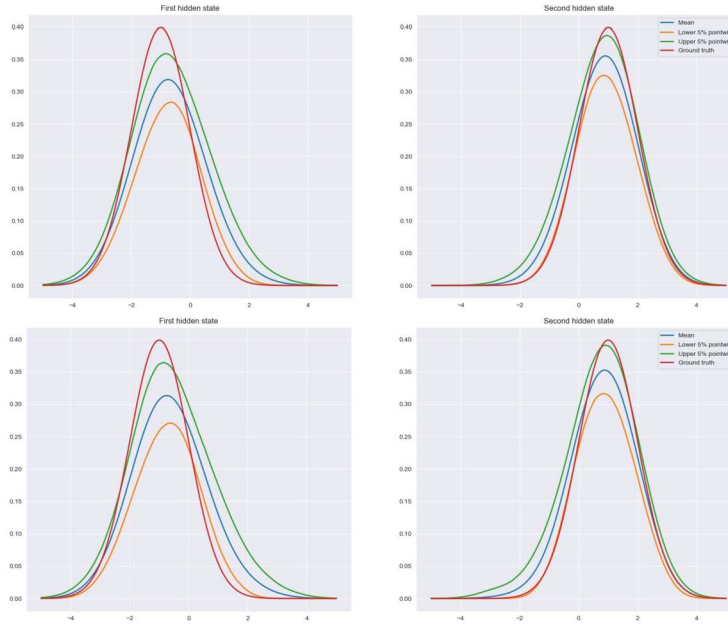


FIG 8. Cut posterior mean and pointwise 90% credible bands for emissions with $n = 1000$ and four bins. Top: $C = 10$ interior iterations; Bottom: $C = 100$ interior iterations.

matrix as discussed beforehand Π_1 , as well as a Dirichlet process mixture prior over the emission densities as discussed earlier in this section for Π_2 .

In the implementation of Π' , we ran the MCMC for 70000 iterations, discarding the first 10000 as burn-in and thinning at a rate of one in ten observations. This was chosen so that we had an approximate matching with the computational cost of 7000 iterations of the cut posterior, each with 10 interior iterations. We plot the resulting emission densities in Figure 9. In comparison to Figure 7, the pointwise bands seem to capture the ground truth more accurately. We remark however that Theorem 5 only provides guarantees on L^1 concentration and so the plots will not entirely reflect the theory; L^1 credible sets are rather less easily visualised. We also emphasize our earlier comment from Section 5.1, that the simulation study is limited in scope, and is not intended to provide conclusive comparisons with other approaches.

6. Proofs of main results

We first prove Proposition 1.

6.1. Proof of Proposition 1: Convergence of the scores

Throughout this section, Assumptions 1–5 are assumed to hold.




FIG 9. Posterior mean and pointwise 90% credible bands for emissions, with $n = 2500$ (top) and $n = 10000$ (bottom) under fully Bayesian posterior with prior Π' .

Recall that

$$\mathcal{H}_r = \{h \in L^2(f_r^* dx), \int h(x) f_r^* dx = 0\},$$

and

$$\mathcal{H}_{r,M} = \{h = \sum_{m=1}^{\kappa_M} \alpha_m \mathbb{1}_{I_m}(y), \alpha_m \in \mathbb{R}, \int h f_{r,M}^*(y) dy = 0\}.$$

Define \mathcal{P} to be the space spanned by the nuisance scores $\{H(h) : h = (h_r)_{r \leq R}, h_r \in \mathcal{H}_r\}$ in the full semiparametric model, and also \mathcal{P}_M to be the space spanned by the nuisance scores $\{H_M(h_M) : h_M = (h_{r,M})_{r \leq R}, h_{r,M} \in \mathcal{H}_{r,M}\}$ in the histogram models $\mathcal{M}(\mathcal{I}_M)$. Here $H(h) = (H_r(h_r) : r \in [R])$ with H_r as in Equation (14), and also $H_M(h_M) = (H_{r,M}(h_r) : r \in [R])$ with $H_{r,M}$ as in Equation (17). Denote also $\mathcal{G}_M^\ell = \sigma(Y_t^{(M)} : t \leq \ell)$ and $\mathcal{G}^\ell = \sigma(Y_t : t \leq \ell)$, for $\ell \in \mathbb{Z}$.

To prove Proposition 1, we prove convergence of sequences of nuisance scores in \mathcal{P}_M to nuisance scores in \mathcal{P} , and also convergence of the score functions $S_{Q^*}^{(M)}$ to S_{Q^*} . Recall that the scores S_{Q^*} in the model with known emissions are given by (12), with the scores $S_{Q^*}^{(M)}$ associated to data $Y^{(M)}$ being given by (16).

Proof of Proposition 1. Write \mathcal{A}_M for the orthogonal projection onto \mathcal{P}_M and

\mathcal{A} for the orthogonal projection onto \mathcal{P} . Define

$$\tilde{S}_{Q^*}^{(M)} := S_{Q^*}^{(M)} - \mathcal{A}_M S_{Q^*}^{(M)} \quad , \text{ and } \quad \tilde{S}_{Q^*} := S_{Q^*} - \mathcal{A} S_{Q^*} .$$

To prove Proposition 1, it is enough to prove that $\tilde{S}_{Q^*}^{(M)}$ converges to \tilde{S}_{Q^*} . A major difficulty here is to prove the convergence of $\mathcal{A}_M S$ to $\mathcal{A} S$ in $L^2(\mathbb{P}_*)$, with $S \in L^2(\mathbb{P}_*)$, because, contrary to the mixture model case considered in [27], the sets \mathcal{P}_M are not embedded.

In Lemma 3, we prove using a martingale argument that $S_{Q^*}^{(M)}$ converges to S_{Q^*} in $L^2(\mathbb{P}_*)$ as M goes to infinity. Then

$$\begin{aligned} \|\tilde{S}_{Q^*}^{(M)} - \tilde{S}_{Q^*}\|_{L^2} &\leq \|S_{Q^*}^{(M)} - S_{Q^*}\|_{L^2} + \|\mathcal{A}_M(S_{Q^*} - S_{Q^*}^{(M)})\|_{L^2} + \|(\mathcal{A} - \mathcal{A}_M)S_{Q^*}\|_{L^2} \\ &= \|(\mathcal{A} - \mathcal{A}_M)S_{Q^*}\|_{L^2} + o(1) . \end{aligned}$$

We now prove that $\|(\mathcal{A} - \mathcal{A}_M)S\|_{L^2} = o(1)$, for all $S \in L^2(\mathbb{P}_*)$. As mentioned earlier, the non trivial part of this proof comes from the fact that the sets \mathcal{P}_M are not embedded. Hence we first prove that the set \mathcal{P}_M is close to the set $\tilde{\mathcal{P}}_M$, where

$$\tilde{\mathcal{P}}_M = \{H(h_M), h_M = (h_{r,M})_{r \leq R}, h_{r,M} \in \mathcal{H}_{r,M}\} .$$

More precisely, we prove in Lemma 1 below that $\mathcal{A}_M S - \tilde{\mathcal{A}}_M S = o_{L^2}(1)$, where $\tilde{\mathcal{A}}_M$ is the projection onto $\tilde{\mathcal{P}}_M$. Since the partitions are nested, the $\tilde{\mathcal{P}}_M$ are nested; reasoning as in [27], we then obtain that $(\tilde{\mathcal{A}}_M S)_{M \in \mathbb{N}}$ is Cauchy, and so converges to some $\tilde{\mathcal{A}} S$. Then, arguing as in [27], we identify $\tilde{\mathcal{A}} S$ with $\mathcal{A} S$ in Lemma 5. The argument involves first showing that $\tilde{\mathcal{A}}$ is a projection onto a subspace of \mathcal{P} , and then showing that elements of \mathcal{P} are well approximated by elements of $\tilde{\mathcal{P}}$, by approximating the corresponding $h_r \in L^2_0(f_r^*)$ by histograms $(h_{r,M})_{M \in \mathbb{N}}$. This terminates the proof of Proposition 1. \square

Lemma 1. For any element $S \in L^2(\mathbb{P}_*)$,

$$\|\mathcal{A}_M S - \tilde{\mathcal{A}}_M S\|_2 = o(1) .$$

Proof. As will appear in the proof, a key step is Lemma 2 which says that, for any sequence h_M bounded in L^2 , if $H_M(h_M) = o_{L^2}(1)$, then $h_M = o_{L^2}(1)$. Lemma 4 also applies for L^2 bounded sequences, and implies that, for any bounded sequence h_M , $(H_M - H)(h_M) = o_{L^2}(1)$.

To prove that $\mathcal{A}_M S - \tilde{\mathcal{A}}_M S = o_{L^2}(1)$, we decompose

$$\mathcal{A}_M S - \tilde{\mathcal{A}}_M S = \mathcal{A}_M S - \tilde{\mathcal{A}}_M \mathcal{A}_M S + \tilde{\mathcal{A}}_M \mathcal{A}_M S - \tilde{\mathcal{A}}_M S .$$

We first prove that $\tilde{\mathcal{A}}_M \mathcal{A}_M S - \mathcal{A}_M S = o_{L^2}(1)$, and then that $\tilde{\mathcal{A}}_M (\mathcal{A}_M S - S) = o_{L^2}(1)$. Define h_M such that $\mathcal{A}_M S = H_M(h_M)$; we have that $\|H_M(h_M)\|_2 \leq \|S\|_2$. Also, by Lemma 4 and using the fact that orthogonal projections onto a space minimize the L^2 distance to that space (which includes $H(h_M)$),

$$\|\tilde{\mathcal{A}}_M \mathcal{A}_M S - \mathcal{A}_M S\|_2 \leq \|(H - H_M)(h_M)\| .$$

Suppose now that, along a subsequence, $\|h_M\|_2 \rightarrow \infty$. Then, along this subsequence, $\frac{H_M(h_M)}{\|h_M\|_2} = H_M(\frac{h_M}{\|h_M\|_2}) \rightarrow 0$ in L^2 and we can apply Lemma 2 to conclude that $\frac{h_M}{\|h_M\|_2} \rightarrow 0$ in L^2 along this subsequence, when in fact this subsequence has constant unit L^2 norm. We thus get a contradiction and conclude that $\sup_M \|h_M\|_2$ is bounded. Then, by Lemma 4, we have that $\mathcal{A}_M S = H_M(h_M) = H(h_M) + o_{L^2}(1)$. Thus $\|(H_M - H)(h_M)\|_2 \rightarrow 0$ and $\|\tilde{\mathcal{A}}_M \mathcal{A}_M S - \tilde{\mathcal{A}}_M S\|_2 = o(1)$.

We now control $\tilde{\mathcal{A}}_M(\mathcal{A}_M S - S)$. Considering $\tilde{\mathcal{A}}_M(S - \mathcal{A}_M S) = H(\tilde{h}_M)$ for some $\tilde{h}_M = (\tilde{h}_{r,M}, r \in [R])$ and $\tilde{h}_{r,M} \in L^2_0(f_r^*)$ in model M , we have

$$\begin{aligned} \|\tilde{\mathcal{A}}_M S - \tilde{\mathcal{A}}_M \mathcal{A}_M S\|_{L^2}^2 &= \langle S - \mathcal{A}_M S, \tilde{\mathcal{A}}_M(S - \mathcal{A}_M S) \rangle \\ &= \langle S - \mathcal{A}_M S, H(\tilde{h}_M) \rangle = \underbrace{\langle S - \mathcal{A}_M S, H_M(\tilde{h}_M) \rangle}_{=0 \text{ as } (S - \mathcal{A}_M S) \perp H_M(\tilde{h}_M)} + R_M, \end{aligned}$$

where

$$|R_M| \leq \|S\|_2 \|(H - H_M)(\tilde{h}_M)\|_2 \rightarrow 0,$$

with the convergence following as before using Lemma 2 and the boundedness of $\|H(\tilde{h}_M)\|_{L^2}$ by $\|S\|_{L^2}$. Thus $\tilde{\mathcal{A}}_M(\mathcal{A}_M S - S) = o_{L^2}(1)$ and Lemma 1 is proved. \square

Lemma 2. Fix $r \in \{1, \dots, R\}$. Let $h_M = (h_{r,M}, r \in [R])$ be a sequence of step functions on the partitions of model $\mathcal{M}(\mathcal{I}_M)$ with $\|h_M\| = \sum_r \|h_{r,M}\|_{L^2} \leq 1$ and $h_{r,M} \in \mathcal{H}_{r,M}$. Then

$$H_M(h_M) \xrightarrow{L^2(\mathbb{P}^*)} 0 \implies h_M \xrightarrow{L^2(\mathbb{P}^*)} 0.$$

Proof. Assuming $H_M(h_M) \rightarrow 0$, Lemma 6 implies that, for all $k \geq 0$,

$$D_k(h_M) := \sum_{r=1}^R \sum_{k'=-k}^0 h_{r,M}(y_{k'}) \mathbb{P}_*(X_{k'} = r | Y_{-k:0}) = o_{L^2}(1).$$

For shortness sake we write $g^* = g_{Q^*, \mathbf{F}^*}^{(2)}$ and we will abuse notation and write $y_j := y_{-j}$ for $j \in \mathbb{N}$, so y_1 is y_{-1} etc. Expanding $D_1(h_M)$, we thus have that

$$\begin{aligned} &\sum_r h_{r,M}(y_0) f_r^*(y_0) \sum_s \frac{p_s^* Q_{sr}^* f_s^*(y_1)}{g^*(y_0, y_1)} \\ &+ \sum_r \frac{h_{r,M}(y_1) f_r^*(y_1)}{g^*(y_0, y_1)} p_r^* \sum_s f_s^*(y_0) Q_{rs}^* \xrightarrow{L^2(\mathbb{P}^*)} 0. \end{aligned}$$

Since, for any function H ,

$$\mathbb{E}_* \left[\frac{H^2(Y_0, Y_1)}{g^*(Y_0, Y_1)^2} \right] = \int_{[0,1]^2} \frac{H^2(y_0, y_1)}{g^*(y_0, y_1)} dy_1 dy_0 \geq \frac{\|H\|_2^2}{\|g^*\|_\infty},$$

and, since $\|g^*\|_\infty \leq \max_r \|f_r^*\|_\infty < \infty$ by continuity, we obtain

$$\begin{aligned} & \sum_r h_{r,M}(y_0) f_r^*(y_0) \sum_s p_s^* Q_{sr}^* f_s^*(y_1) \\ & + \sum_r h_{r,M}(y_1) f_r^*(y_1) p_r^* \sum_s f_s^*(y_0) Q_{rs}^* \xrightarrow{L^2(dy^2)} 0. \end{aligned} \tag{24}$$

Note further that the family of functions of y_1 given by $y_1 \mapsto (\sum_s p_s^* Q_{sr}^* f_s^*(y_1))_{r=1}^k$ is linearly independent. Consider now a partition $\bar{I}_1, \bar{I}_2, \dots, \bar{I}_R$ of $[0, 1]$ such that

$$F_{rj} = \int_{\bar{I}_j} f_r^*(y_1) dy_1, \quad r, j \in \{1, \dots, R\},$$

defines a matrix $F = (F_{rj}^*)$ of rank R . Such a partition always exists by Lemma 17. Define also

$$A_{rj} = \int_{I_j} \sum_s p_s^* Q_{sr}^* f_s^*(y_1) dy_1, \quad r, j \in \{1, \dots, R\},$$

with $A = (A_{rj})$, define $D_p = \text{diag}(p_1^*, \dots, p_R^*)$ (which is invertible under **Assumption 3**) and recall $Q^* = (Q_{sr}^*)$. Then we can write $A^T = F^T D_p Q^*$ and note that A also has rank R . Now, integrating the y_1 coordinate out of Equation (24), we get, for all $j = 1, \dots, k$,

$$\begin{aligned} \sum_r \underbrace{h_{r,M}(y_0) f_r^*(y_0)}_{\tilde{h}_{r,M}} A_{rj} &= - \sum_r \underbrace{\int_{\bar{I}_j} h_{r,M}(y_1) f_r^*(y_1) p_r^* dy_1}_{B_M(r,j)} \sum_s f_s^*(y_0) Q_{rs}^* \\ &+ o_{L^2(dy_0)}(1). \end{aligned}$$

Let $\tilde{h}_M = (\tilde{h}_{r,M})_{r=1}^R$, $h_M = (h_{r,M})_{r=1}^R$ and $B_M = (B_M(r, j))_{rj}$. Then the preceding display may be rewritten as

$$\tilde{h}_M(y_0) = \underbrace{-(A^T)^{-1} B_M^T Q^*}_{B_{0,M}} \mathbf{f}^*(y_0) + o_{L^2(dy_0)}(1). \tag{25}$$

Now, from its definition, $(B_M)_M$ is a bounded sequence in a finite dimensional space and so converges to some B along a subsequence. Working on this subsequence, we then get the limit

$$\tilde{h}_M(y_0) \xrightarrow{L^2(dy_0)} \underbrace{-(A^T)^{-1} B^T Q^*}_{B_0} \mathbf{f}^*(y_0) := \tilde{h}(y_0).$$

Consider now $D_2(h_M)$, which again vanishes under our assumptions and by Lemma 6, and take the limit in M along the subsequence where we just established convergence. As before, abuse notation by replacing Y_{-k} by Y_k for $k \in \mathbb{N}$. Expanding and multiplying through by joint marginals, we get at the limit

$$\sum_r (B_0 \mathbf{f}^*(y_0))_r \sum_{s_1, s_2} p_{s_2}^* Q_{s_2 s_1}^* Q_{s_1 r}^* f_{s_2}^*(y_2) f_{s_1}^*(y_1)$$

$$\begin{aligned}
& + \sum_r (B_0 \mathbf{f}^*(y_1))_r \sum_{s_1, s_2} p_{s_2}^* Q_{s_2 r}^* Q_{r s_1}^* f_{s_1}^*(y_0) f_{s_2}^*(y_2) \\
& + \sum_r (B_0 \mathbf{f}^*(y_2))_r \sum_{s_1, s_2} p_r^* Q_{r s_1}^* Q_{s_1 s_2}^* f_{s_1}^*(y_1) f_{s_2}^*(y_0) = 0. \tag{26}
\end{aligned}$$

Rewriting (26) into matrix products, we obtain

$$\begin{aligned}
& \sum_{s_1} [\mathbf{f}^*(y_2)^T D_p Q^*]_{s_1} f_{s_1}^*(y_1) (Q^* B_0 \mathbf{f}^*(y_0))_{s_1} \\
& + \sum_{s_1} [B_0 \mathbf{f}^*(y_1)]_{s_1} (Q^* \mathbf{f}^*(y_0))_{s_1} (\mathbf{f}^*(y_2)^T D_p Q^*)_{s_1} \\
& + \sum_{s_1} [\mathbf{f}^*(y_2)^T B_0^T D_p Q^*]_{s_1} f_{s_1}^*(y_1) (Q^* \mathbf{f}^*(y_0))_{s_1} = 0.
\end{aligned}$$

By linear independence of the the set of functions $\{f_s^*(y_1)\}_{s=1, \dots, R}$, we have that, for all s_1 , the coefficients vanish so that, for all s_1 , we get

$$\begin{aligned}
& [\mathbf{f}^*(y_2)^T D_p Q^*]_{s_1} (Q^* B_0 \mathbf{f}^*(y_0))_{s_1} + \sum_{r'} B_{0r's_1} (Q^* \mathbf{f}^*(y_0))_{r'} (\mathbf{f}^*(y_2)^T D_p Q^*)_{r'} \\
& + [\mathbf{f}^*(y_2)^T B_0^T D_p Q^*]_{s_1} (Q^* \mathbf{f}^*(y_0))_{s_1} = 0.
\end{aligned}$$

We can now repeat the argument, instead viewing the expression as a linear combination of the linearly independent set of functions $\{f_s^*(y_2)\}_{s=1, \dots, R}$. Expanding what precedes for each s_1 , we get

$$\begin{aligned}
& \sum_r f_r^*(y_2) p_r^* Q_{r s_1}^* (Q^* B_0 \mathbf{f}^*(y_0))_{s_1} + \sum_{r, r'} B_{0r's_1} (Q^* \mathbf{f}^*(y_0))_{r'} (f_r^*(y_2)) p_r^* Q_{r r'}^* \\
& + \sum_r f_r^*(y_2) (B_0^T D_p Q^*)_{r s_1} (Q^* \mathbf{f}^*(y_0))_{s_1} = 0,
\end{aligned}$$

so that, for each r, s_1 we get that

$$\begin{aligned}
& p_r^* Q_{r s_1}^* (Q^* B_0 \mathbf{f}^*(y_0))_{s_1} + \sum_{r'} B_{0r's_1} (Q^* \mathbf{f}^*(y_0))_{r'} p_r^* Q_{r r'}^* \\
& + (B_0^T D_p Q^*)_{r s_1} (Q^* \mathbf{f}^*(y_0))_{s_1} = 0.
\end{aligned}$$

Expanding in terms of the linearly independent functions $\{(Q^* \mathbf{f}^*)_r(y_0)\}_{r=1, \dots, R}$, we get

$$\begin{aligned}
& \sum_{r'} p_r^* Q_{r s_1}^* (Q^* B_0 Q^{*-1})_{s_1 r'} (Q^* \mathbf{f}^*)_{r'}(y_0) + \sum_{r'} B_{0r's_1} (Q^* \mathbf{f}^*(y_0))_{r'} p_r^* Q_{r r'}^* \\
& + \sum_{r'} \mathbb{1}_{s_1=r'} (B_0^T D_p Q^*)_{r r'} (Q^* \mathbf{f}^*(y_0))_{r'} = 0,
\end{aligned}$$

which then gives, for each r, r', s_1 , that

$$p_r^* Q_{r s_1}^* (Q^* B_0 Q^{*-1})_{s_1 r'} + B_{0r's_1} p_r^* Q_{r r'}^* + \mathbb{1}_{s_1=r'} (B_0^T D_p Q^*)_{r r'} = 0.$$

In particular, setting $s_1 = r'$ we get that, for all r, r' ,

$$p_r^* Q_{rr'}^* (Q^* B_0 Q^{*-1})_{r'r'} + B_{0r'r'} p_r^* Q_{rr'}^* + (B_0^T D_p Q^*)_{rr'} = 0.$$

By **Assumption 3**, $Q_{ij}^* > 0 \forall i, j$ and so we can divide through to get

$$\frac{(B_0^T D_p Q^*)_{rr'}}{Q_{rr'}^*} = p_r^* a(r'),$$

for some a depending on r' alone. Write D_a for the diagonal matrix whose diagonals are the $a(r')$. Then

$$(B_0^T D_p Q^*)_{rr'} = p_r^* Q_{rr'}^* a(r') = (D_p Q^* D_a)_{rr'}.$$

Equivalently, we can write $B_0^T = (D_p Q^*) D_a (D_p Q^*)^{-1}$. Recalling the earlier definition of B_0 , we had

$$B_0 = -(A^T)^{-1} B^T Q^*,$$

with each column of B expressible as the integral over \bar{I}_j , with respect to y_1 , of $D_p \tilde{h}(y_1) = D_p B_0 \mathbf{f}^*(y_1)$, and thus expressible as $B = D_p B_0 F$. We can then substitute into the previous display to see that

$$B_0 = -(A^T)^{-1} F^T B_0^T D_p Q^*.$$

The expression of A as $A^T = F^T D_p Q^*$ allows us to write

$$B_0 = -Q^{*-1} D_p^{-1} (F^T)^{-1} F^T B_0^T D_p Q^* = -Q^{*-1} D_p^{-1} B_0^T D_p Q^* = -D_a.$$

This means that $B_0^T = -D_a = (D_p Q^*) D_a (D_p Q^*)^{-1}$. So $Q^{-1} D_a Q^* = -D_a$ which gives

$$a(r) Q_{rs}^* = -Q_{rs}^* a(s).$$

Taking $r = s$ and dividing through by Q_{rr}^* gives $a_r = 0$ for all r , hence $B_0 = 0$ and $\tilde{h} = 0$. This means $\tilde{h}_m(y_0) \rightarrow 0$ in $L^2(dy_0)$ and so $h_{r,M} \rightarrow 0$ in $L^2(f_r(y_0) dy_0)$, and so in $L^2(\mathbb{P}_*)$ under **Assumption 4**, which is what we wanted to show. \square

Lemma 3. For $S_{Q^*}^{(M)}, S_{Q^*}$, as in (16) and (12), and for any choice of j , we have that

$$S_{Q^*}^{(M)}(Y_{-\infty:j}^{(M)}) \rightarrow S_{Q^*}(Y_{-\infty:j})$$

in $L^2(\mathbb{P}_*)$.

Proof. Without loss of generality we take $j = 0$. We first show that we can, uniformly in M , truncate the series in (16) to an arbitrary degree of accuracy. We prove that, uniformly in M ,

$$\sum_{j \leq -J} \mathbb{P}_*(X_{j-1} = a, X_j = b | \mathcal{G}_M^0) - \mathbb{P}_*(X_{j-1} = a, X_j = b | \mathcal{G}_M^{-1}) = o_{L^2(\mathbb{P}_*)}(1), \tag{27}$$

as J goes to infinity. Using Lemma 19, we have that there exists $\rho < 1$ and $C > 0$ such that, for all $j \leq -1$ and for all M ,

$$\left\| \mathbb{P}_*(X_{j-1} = a, X_j = b | \mathcal{G}_M^0) - \mathbb{P}_*(X_{j-1} = a, X_j = b | \mathcal{G}_M^{-1}) \right\|_{L^2(\mathbb{P}_*)} \leq C\rho^{-j}.$$

Analogous results hold with $\mathcal{G}^0, \mathcal{G}^1$ in place of $\mathcal{G}_M^0, \mathcal{G}_M^1$. Therefore, for all $\epsilon > 0$, there exists $J(\epsilon) > 0$ for which

$$\begin{aligned} \sum_{j \leq -J(\epsilon)} & \left\| \mathbb{P}_*(X_{j-1} = a, X_j = b | \mathcal{G}_M^0) - \mathbb{P}_*(X_{j-1} = a, X_j = b | \mathcal{G}_M^{-1}) \right\|_{L^2(\mathbb{P}_*)} \\ & + \left\| \mathbb{P}_*(X_{j-1} = a, X_j = b | \mathcal{G}^0) - \mathbb{P}_*(X_{j-1} = a, X_j = b | \mathcal{G}^{-1}) \right\|_{L^2(\mathbb{P}_*)} < \epsilon. \end{aligned}$$

It now suffices to prove for all j that

$$\mathbb{P}_*(X_{j-1} = a, X_j = b | \mathcal{G}_M^l) \xrightarrow{L^2(\mathbb{P}_*)} \mathbb{P}_*(X_{j-1} = a, X_j = b | \mathcal{G}^l). \quad (28)$$

Firstly, we show that, for each $\ell = -1, 0$, the \mathcal{G}_M^ℓ form an increasing sequence of σ -algebras $\mathcal{G}_M^\ell \subset \mathcal{G}_{M'}^\ell$ for $M < M'$, and that the limiting σ -algebra $\mathcal{G}_\infty^\ell := \sigma(\bigcup_{M \in \mathbb{N}} \mathcal{G}_M^\ell)$ is equal to \mathcal{G}^ℓ . We suppress the dependence on ℓ , since the proofs for each case are the same. For the sake of presentation, we code the coarsened observations $Y_i^{(M)}$ as their projection onto the left endpoint of whichever interval of the partition \mathcal{I}_M contains Y_i : $Y_i^{(M)} = e_j$ if $Y_i \in I_j^{(M)} = (e_j, e_{j+1})$.

To see that the \mathcal{G}_M form a filtration, we wish to show that every set in \mathcal{G}_M is also in \mathcal{G}_{M+1} . Every set in these σ -algebras is generated by the preimages of the left endpoints of intervals under Y_i^M for some i . Suppose the partition \mathcal{I}_{M+1} is obtained from \mathcal{I}_M by splitting one bin. The events $\{Y_i^M = e_j\}$, for $j = 1, \dots, M$ and e_j the endpoint of bin j , generate \mathcal{G}_M . For the split bin, if we split bin M in \mathcal{I}_M into bins M and $M + 1$ in \mathcal{I}_{M+1} , then we can express $\{Y_i^M = e_j\} = \{Y_i^{M+1} = e_j\}$ for $j < M$ and $\{Y_i^M = e_M\} = \{Y_i^{M+1} = e_M\} \cup \{Y_i^{M+1} = e_{M+1}\}$. Hence we can express these events as events in \mathcal{G}_{M+1} , and by induction the argument remains true for any nested sequence of intervals.

To see that each $\mathcal{G}_M \subset \mathcal{G}$ it suffices to note that $\{Y_i^M = e_j\} = \{Y_i \in [e_j, e_{j+1})\} \in \mathcal{G}$. This shows that $\mathcal{G}_\infty \subset \mathcal{G}$. To show that $\mathcal{G} \subset \mathcal{G}_\infty$, note that the Borel σ -algebra on $[0, 1]$ is generated by the sets $(t, 1]$. So it suffices to check that the sets $\{Y_i > t\}$ are in \mathcal{G}_∞ , as these sets generate \mathcal{G} . But, for all $t \in (0, 1)$, there exists a sequence $e_{j(t, M)} \leq t$ converging to t and such that $e_{j(t, M)+1} > t$ also converges to t . Thus $\{Y_i^M > t\} = \{Y_i \geq e_{j(t, M)+1}\}$ increases to $\{Y_i > t\}$ and $\{Y_i > t\} = \lim_M \{Y_i^M > t\} \subset \mathcal{G}_\infty$. Therefore, the conditional probabilities appearing in the histogram score expression (16) form a bounded martingale in L^2 with $\mathcal{G}_\infty^\ell = \mathcal{G}^\ell$, $\ell = -1, 0$. The L^2 martingale convergence theorem then implies that (28) holds for any j, ℓ , and Lemma 3 is proved. \square

The proofs of the following results, which are used in the proofs of Lemmas 1 and 2, are deferred to Section C of the appendix.

Lemma 4. *If $h_M = (h_{r,M})_{r \leq R}$ with $\sup_M \sum_r \|h_{r,M}\|_{L^2(f_r^*)} \leq L < \infty$, then*

$$H_M(h_M) = H(h_M) + o_{L^2(\mathbb{P}_*)}(1).$$

Lemma 5. *Let $\tilde{\mathcal{A}}_M, \tilde{\mathcal{P}}_M, \mathcal{A}, \mathcal{P}$ be as in the statement of Lemma 1. Then, if $\tilde{\mathcal{A}} := \lim_{M \rightarrow \infty} \tilde{\mathcal{A}}_M$, we have $\tilde{\mathcal{A}} = \mathcal{A}$.*

Lemma 6. *Let h_M be a sequence of step functions on the partitions of model $\mathcal{M}(\mathcal{I}_M)$ with $\|h_M\| = \sum_r \|h_{r,M}\|_{L^2} \leq 1$, such that $h_{r,M} \in L^2_0(f_r^*)$. For any $k \geq 0$,*

$$H_M(h_M) \xrightarrow{L^2(\mathbb{P}_*)} 0 \implies \sum_{r=1}^R \sum_{k'=-k}^0 h_M(y_{k'}) \mathbb{P}_*(X_{k'} = r | Y_{-k:0}) \xrightarrow{L^2(\mathbb{P}_*)} 0.$$

6.2. Proof of Theorem 4

The proof of Theorem 4 works by contradiction. If there exists no $C > 0$ such that, for all \mathbf{f} with $\|g_{Q,\mathbf{f}}^{(3)} - g_{Q,\mathbf{f}^*}^{(3)}\|_{L^1}$ sufficiently small,

$$\frac{\|g_{Q,\mathbf{f}}^{(3)} - g_{Q,\mathbf{f}^*}^{(3)}\|_{L^1}}{\sum_{r=1}^R \|f_r - f_r^*\|_{L^1}} \geq C,$$

then there exists a sequence of emission densities $\mathbf{f}^{(n)}$ such that $\|g_{Q,\mathbf{f}^{(n)}}^{(3)} - g_{Q,\mathbf{f}^*}^{(3)}\|_{L^1} = o(1)$ and

$$\liminf_n \frac{\|g_{Q,\mathbf{f}^{(n)}}^{(3)} - g_{Q,\mathbf{f}^*}^{(3)}\|_{L^1}}{\sum_{r=1}^R \|f_r^{(n)} - f_r^*\|_{L^1}} = 0.$$

For $j = 2, 3$, $r \in [R]$, and $\mathbf{f}^{(n)}$ as above, set

$$\Delta_{n,j} = \frac{\|g_{Q,\mathbf{f}^{(n)}}^{(j)} - g_{Q,\mathbf{f}^*}^{(j)}\|_{L^1}}{\sum_{r=1}^R \|f_r^{(n)} - f_r^*\|_{L^1}}; \quad h_r^n = \frac{f_r^{(n)} - f_r^*}{\sum_{r=1}^R \|f_r^{(n)} - f_r^*\|_{L^1}},$$

and write $\mathbf{h}^n = (h_r^n : r \in [R])$.

Since $\liminf_n \Delta_{n,3} = 0$, there exists a subsequence $\phi_1(n)$ such that $\Delta_{\phi_1(n),3} \rightarrow 0$, which implies also that $\Delta_{\phi_1(n),2} \rightarrow 0$ since $\Delta_{n,2} \leq \Delta_{n,3}$. We have by assumption that $\|g_{Q,\mathbf{f}^{(\phi_1(n))}}^{(j)} - g_{Q,\mathbf{f}^*}^{(j)}\|_{L^1} \rightarrow 0$, and by applying Lemma 7 we can find a further subsequence $\phi_2(\phi_1(n))$ such that $\mathbf{f}^{(\phi_2(\phi_1(n)))} \rightarrow \mathbf{f}^*$ in L^1 . For notational convenience we write $\Delta_{n,j}$ for the terms $\Delta_{\phi_2(\phi_1(n)),j}$ along this subsequence, and in general any index with n is now interpreted as being along this subsequence. Since $\sum_{r=1}^R \|f_r^n - f_r^*\|_{L^1} \rightarrow 0$, $\Delta_{n,2} \rightarrow 0$ implies that

$$\sum_{r,s} (PQ^*)_{rs} (f_r^*(y_1) h_s^n(y_2) + h_r^n(y_1) f_s^*(y_2)) = o_{L^1}(1),$$

where $P = \text{diag}(p_1^*, \dots, p_R^*)$. Partition I_1, \dots, I_R such that \mathbf{F}^* , as defined in Lemma 7, has rank R . Define also $\mathbf{H}^n = (H_{ir}^n)_{ir} = \int_{I_i} h_r^n(y) dy$. We have, for any integrable function G and any $i \in [R]$, that

$$\int_{\mathbb{R}} \int_{\mathbb{R}} |G(y_1, y_2)| dy_1 dy_2 \geq \int_{\mathbb{R}} \int_{I_i} |G(y_1, y_2)| dy_1 dy_2 \geq \int_{\mathbb{R}} \left| \int_{I_i} G(y_1, y_2) dy_1 \right| dy_2,$$

which implies in particular that

$$\sum_{r,s} (PQ^*)_{rs} (F_{ir}^* h_s^n(y_2) + H_{ir}^n f_s^*(y_2)) = o_{L^1}(1).$$

This gives us

$$(\mathbf{F}^*)^T (PQ^*) \mathbf{h}^n = -\mathbf{H}_n^T (PQ) \mathbf{f}^* + o_{L^1}(1).$$

The sequence \mathbf{H}_n is bounded, and so converges to some \mathbf{H}_0 along a subsequence $\phi_3(n)$ which we pass to. Then, since \mathbf{F}^* is of rank R , we may write

$$\mathbf{h}^n = -\underbrace{((\mathbf{F}^*)^T (PQ))^{-1} (\mathbf{H}_0^T (PQ))}_{=B_1} \mathbf{f}^* + o_{L^1}(1),$$

and so $h_r^n \rightarrow (B_1 \mathbf{f}^*)_r$ in L^1 for each $r = 1, \dots, R$. Next, since we have that $\Delta_{n,3} \rightarrow 0$ in L^1 , we get

$$\begin{aligned} & \sum_{r,s,t} (PQ)_{rs} Q_{st} (f_r^*(y_1) f_s^*(y_2) h_t^n(y_3) + f_r^*(y_1) h_s^n(y_2) f_t^*(y_3)) \\ & + h_r^n(y_1) f_s^*(y_2) f_t^*(y_3) = o_{L^1}(1). \end{aligned}$$

Replacing h_r^n by its limit $(B_1 \mathbf{f}^*)_r$, we get that, for Lebesgue-almost-all y_1, y_2, y_3

$$\begin{aligned} & \sum_{r,s,t} (PQ)_{rs} Q_{st} (f_r^*(y_1) f_s^*(y_2) (B_1 \mathbf{f}^*)_t(y_3) + f_r^*(y_1) (B_1 \mathbf{f}^*)_s(y_2) f_t^*(y_3)) \\ & + (B_0 \mathbf{f}^*)_r(y_1) f_s^*(y_2) f_t^*(y_3) = 0 \end{aligned}$$

Under **Assumption 4**, we have continuity of f^* and so in fact the above display holds for all y_1, y_2, y_3 . This expression is of the form of Equation (26) in the proof of Lemma 2, and the same proof techniques show that $B_1 = 0$ and hence $\|\mathbf{h}^n\| = o_{L^1}(1)$. This yields the desired contradiction as $\sum_{r=1}^R \|h_r^n\|_{L^1} = 1$ for all $n \in \mathbb{N}$. We conclude that no subsequence can exist for which $\Delta_{n,3} \rightarrow 0$, and hence that $\liminf_{n \rightarrow \infty} \Delta_{n,3} > 0$, which terminates the proof of Theorem 4.

The proof of Theorem 4 uses the following result, which we prove in Section C.6 of the appendix.

Lemma 7. Consider a sequence of densities $\mathbf{f}^{(n)}$ such that $g_{Q^*, \mathbf{f}^{(n)}}^{(3)} \xrightarrow{L^1(dy^3)}$ $g_{Q^*, \mathbf{f}^*}^{(3)}$. Then $\mathbf{f}^{(n)} \xrightarrow{L^1(dy)}$ \mathbf{f}^* along a subsequence.

6.3. Proof of Theorem 6

To prove Theorem 6, we make use of Lemma 21 which allows us to approximate the MLE by an estimator which does not depend on the observation for which we control the respective smoothing probability. Intuitively, this result says that a single observation does not influence our MLE too much.

Proof of Theorem 6. The starting point of the proof is Proposition 2.2 of [16]. Let $\theta = (Q, \mathbf{f})$ and $\|Q - Q^*\| \leq M_n/\sqrt{n}$ so that $Q_{i,j} \geq \bar{q}/2 > 0$ with $\bar{q} = \min_{i,j} Q_{i,j}^* < \frac{1}{2}$, under **Assumption 3**. Set $\rho = 1 - \bar{q}/(1 - \bar{q})$, then

$$\Delta_k := \|\mathbb{P}_\theta(X_k = \cdot | Y_{1:n}) - \mathbb{P}_*(X_k = \cdot | Y_{1:n})\|_{TV} \leq \frac{C^*}{\bar{q}} [\rho^k \|\mu_Q - \mu^*\| + \|Q - Q^*\| + \sum_{\ell=1}^n \frac{\bar{q}\rho^{|\ell-k|}}{c_*(Y_\ell)} \max_s |f_s(y_\ell) - f_s^*(Y_\ell)|],$$

where $c^*(y) = \min_{r \in [R]} \sum_{s \in [R]} Q_{rs}^* f_s^*(y) \geq \bar{q} \sum_s f_s^*(y)$, and C^* is a constant depending only on Q^*, \mathbf{f}^* . Throughout the proof, we reuse C^* to denote a generic constant depending only on Q^*, \mathbf{f}^* . First, for any $K_n \rightarrow \infty$, if

$$B_n = \{\theta, \|Q^* - \hat{Q}\| \leq \frac{K_n}{\sqrt{n}}; \max_s \|f_s^* - \hat{f}_s\|_{L^1} \leq K_n \epsilon_n\} \cap \{\max_r \|f_r\|_2 \leq e^{\gamma n \epsilon_n^2}\},$$

for γ defined in (19), then Theorem 5 together with condition (19) gives $\Pi_{cut}(B_n^c | Y_{1:n}) = o_{\mathbb{P}_*}(1)$. Recall also that $n\epsilon_n^2 \rightarrow \infty$, so that, for $z_n > 0$,

$$\Pi_{cut}(\Delta_k > z_n \epsilon_n | Y_{1:n}) \leq \Pi_{cut}(\{\Delta_k > z_n \epsilon_n\} \cap B_n | Y_{1:n}) + o_{\mathbb{P}_*}(1).$$

Since $\|\Pi_1(dQ | Y_{1:n}) - \phi_n(Q | Y_{1:n})dQ\|_{TV} \rightarrow 0$, where $\phi_n(Q | Y_{1:n})$ is a Gaussian density centred at the MLE of variance \tilde{J}^{-1}/n , we have $\|\Pi_{cut} - \tilde{\Pi}_{cut}\|_{TV} = o_{\mathbb{P}_*}(1)$, where $\tilde{\Pi}_{cut}$ is the measure $\phi_n(Q | Y_{1:n})dQ\Pi_2(d\mathbf{f} | Q, Y_{1:n})$. Moreover, on B_n ,

$$\Delta_k \leq \frac{C^* K_n}{\sqrt{n}} + \sum_{\ell=1}^n \frac{\bar{q}\rho^{|\ell-k|}}{c_*(Y_\ell)} \max_s |f_s(Y_\ell) - f_s^*(Y_\ell)|. \quad (29)$$

To prove Theorem 6, we thus need to control the sum in (29) under the posterior distribution. In [1], the authors use either a control in sup-norm of $\hat{f}_s - f_s^*$, or split their data into a *training set* used to estimate \mathbf{f}^* , and a *test set* used to estimate the smoothing probabilities. We show here that, with the Bayesian approach, we do not need to split the data in two parts nor do we need a sup-norm control. We believe that our technique of proof might still be useful for other approaches.

We split the sum in (29) into $|\ell - k| \leq T_n$ and $|\ell - k| > T_n$, where $T_n \epsilon_n$ converges arbitrarily slowly to 0. We first control, for all $r \in [R]$ and any $z_n > 0$ going to infinity,

$$\tilde{\Pi}_{cut} \left[B_n \cap \left\{ \sum_{|\ell-k| \leq T_n} \frac{\rho^{|\ell-k|} |f_r(Y_\ell) - f_r^*(Y_\ell)|}{c_*(Y_\ell)} \geq z_n \epsilon_n \sum_{|\ell-k| \leq T_n} \rho^{|\ell-k|/2} \right\} \middle| Y_{1:n} \right]$$

$$\leq \sum_{|\ell-k| \leq T_n} \tilde{\Pi}_{cut} \left[B_n \cap \left\{ \frac{|f_r(Y_\ell) - f_r^*(Y_\ell)|}{c_*(Y_\ell)} \geq z_n \epsilon_n \rho^{-|\ell-k|/2} \right\} \middle| Y_{1:n} \right]. \tag{30}$$

To bound (30), we need to study the quantities $\pi_{n,\ell}(z) = \Pi_{cut} [B_n \cap A_n(z; \ell) | Y_{1:n}]$, where $z > 0$ and

$$A_n(z; \ell) = \left\{ \frac{|f_r(Y_\ell) - f_r^*(Y_\ell)|}{\sum_s f_s^*(Y_\ell)} > z \epsilon_n \right\}.$$

Write $L_n(Y_{1:n}, \theta) = \sum_s f_s(Y_\ell) L_{-\ell}(\theta, s)$, where $L_{-\ell}(\theta, s)$ depends on $Y_{1:n}^{-\ell}$ and is given by

$$L_{-\ell}(\theta, s) = L(y_{1:\ell-1}, \theta) \mathbb{P}_\theta(X_\ell = s | Y_{1:\ell-1}) L(y_{\ell+1:n} | X_\ell = s, \theta).$$

Note that

$$\begin{aligned} L(y_{\ell+1:n} | X_\ell = s, \theta) &= \sum_{s'} Q_{ss'} L(y_{\ell+1:n} | X_{\ell+1} = s', \theta) \\ &\in \left[\frac{\bar{q}}{2} \sum_{s'} L(y_{\ell+1:n} | X_{\ell+1} = s', \theta), \sum_{s'} L(y_{\ell+1:n} | X_{\ell+1} = s', \theta) \right], \end{aligned}$$

and also that $\mathbb{P}_\theta(X_\ell = s | Y_{1:\ell-1}) \in [\frac{\bar{q}}{2}, 1]$, which together implies, denoting $\bar{L}_{-\ell}(\theta) = \max_s L_{-\ell}(\theta, s)$, that

$$L_{-\ell}(\theta, s) \in \left[\frac{\bar{q}^2}{4} \bar{L}_{-\ell}(\theta), \bar{L}_{-\ell}(\theta) \right].$$

This implies that the cut posterior mass of $A_n(z, \ell) \cap B_n$ is bounded as

$$\begin{aligned} &\Pi_{cut}(A_n(z; \ell) \cap B_n | Y_{1:n}) \\ &\leq \frac{4}{\bar{q}^2} \frac{\int_{B_n} \mathbb{1}_{A_n(z; \ell)}(f_r) \bar{L}_{-\ell}(\theta) (\sum_s f_s(Y_\ell)) \Pi_2(d\mathbf{f} | Q) \phi_n(Q | Y_{1:n}) dQ}{\int_{B_n} \bar{L}_{-\ell}(\theta) (\sum_s f_s(Y_\ell)) \Pi_2(d\mathbf{f} | Q) \phi_n(Q | Y_{1:n}) dQ}. \end{aligned} \tag{31}$$

Applying Lemma 21, we also have, uniformly in $\ell \in [k - T_n, k + T_n] \cap \mathbb{N}$, $\|\phi_n(Q | Y_{1:n}) dQ - \phi_n(Q | Y_{1:n}^{-\ell}) dQ\|_{TV} \rightarrow 0$, where $\phi_n(Q | Y_{1:n}^{-\ell})$ is a Gaussian centred at the estimator defined in Lemma 21 and also of variance \tilde{J}^{-1}/n , which depends only on $Y_{1:n}^{-\ell}$. Hence, with

$$\tilde{\Pi}_{cut}^{(\ell)}(A | Y_{1:n}) = \frac{\int_A L_n(Y_{1:n}, \theta) \Pi_2(d\mathbf{f}) \phi(Q | Y_{1:n}^{-\ell}) dQ}{D_n(\ell)},$$

where $D_n(\ell) = \int_{B_n} \bar{L}_{-\ell}(\theta) (\sum_s f_s(Y_\ell)) \Pi_2(d\mathbf{f}) \phi_n(Q | Y_{1:n}^{-\ell}) dQ$, we have

$$\max_{|\ell-k| \leq T_n} \|\tilde{\Pi}_{cut}^{(\ell)}(\cdot | Y_{1:n}) - \tilde{\Pi}_{cut}(\cdot | Y_{1:n})\|_{TV} = o_{\mathbb{P}^*}(1).$$

Hence, it suffices to control each summand of (30) with $\tilde{\Pi}_{cut}$ replaced by $\tilde{\Pi}_{cut}^{(\ell)}$. We first focus on lower bounding the denominator $D_n(\ell)$. Writing $\Omega_n = \{\sum_s f_s(Y_\ell) \geq \sum f_s^*(Y_\ell)/2\}$, it is immediate that

$$D_n(\ell) \geq \int_{B_n} \bar{L}_{-\ell}(\theta) \mathbb{1}_{\Omega_n} \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ \left(\sum_s f_s^*(Y_\ell)/2 \right).$$

We will show that we can replace the above bound with a bound in probability, in which the integrand does not feature the indicator $\mathbb{1}_{\Omega_n}$. This is similar in spirit to concentration results based around the framework of [29] (including our own Theorem 7), where a key ingredient of the proof is to lower bound the denominator by a suitable quantity. We have that

$$\begin{aligned} & \mathbb{P}_* \left(\int_{B_n} \bar{L}_{-\ell}(\theta) \mathbb{1}_{\Omega_n} \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ < \frac{1}{2} \int_{B_n} \bar{L}_{-\ell}(\theta) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ \right) \\ &= \mathbb{P}_* \left(\int_{B_n} \bar{L}_{-\ell}(\theta) \mathbb{1}_{\Omega_n^c} \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ \geq \frac{1}{2} \int_{B_n} \bar{L}_{-\ell}(\theta) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ \right) \\ &\leq 2\mathbb{E}_* \left[\mathbb{E}_* \left[\frac{\int_{B_n} \bar{L}_{-\ell}(\theta) \mathbb{1}_{\Omega_n^c} \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ}{\int_{B_n} \bar{L}_{-\ell}(\theta) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ} \middle| Y_{1:n}^{-\ell} \right] \right] \\ &\leq 2\mathbb{E}_* \left[\frac{\int_{B_n} \bar{L}_{-\ell}(\theta) \mathbb{P}(\Omega_n^c | Y_{1:n}^{-\ell}) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ}{\int_{B_n} \bar{L}_{-\ell}(\theta) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ} \right]. \end{aligned}$$

Now, we can bound the probability in the integrand by

$$\begin{aligned} \mathbb{P}_*(\Omega_n^c | Y_{1:n}^{-\ell}) &= \mathbb{P}_* \left(\sum_s f_s(Y_\ell) < \frac{1}{2} \sum_s f_s^*(Y_\ell) \middle| Y_{1:n}^{-\ell} \right) \\ &= \mathbb{P}_* \left(\sum_s (f_s^*(Y_\ell) - f_s(Y_\ell)) \geq \sum_s f_s^*(Y_\ell)/2 \middle| Y_{1:n}^{-\ell} \right) \\ &\leq 2 \int \frac{\sum_s |f_s^*(y_\ell) - f_s(y_\ell)|}{\sum_s f_s^*(y_\ell)} f^*(y_\ell | Y_{1:n}^{-\ell}) dy_\ell. \end{aligned}$$

We can bound the conditional density by $f^*(y_\ell | Y_{1:n}^{-\ell}) \leq \sum_s f_s^*(y_\ell)$, so that

$$\mathbb{P}_*(\Omega_n^c | Y_{1:n}^{-\ell}) \leq 2 \sum_s \|f_s^* - f_s\|_{L_1} \leq 2RK_n \epsilon_n,$$

which in turns implies that

$$\begin{aligned} & \mathbb{P}_* \left(\int_{B_n} \bar{L}_{-\ell}(\theta) \mathbb{1}_{\Omega_n} \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ < \right. \\ & \qquad \left. \frac{1}{2} \int_{B_n} \bar{L}_{-\ell}(\theta) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ \right) \leq 4RK_n \epsilon_n. \end{aligned}$$

Thus, on this event of probability at least $1 - 4RK_n\epsilon_n$, we can bound the denominator as

$$D_n(\ell) \geq \frac{1}{4} \sum_s f_s^*(y_\ell) \underbrace{\int_{B_n} \bar{L}_{-\ell}(\theta) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ}_{=: M(Y_{1:n}^{-\ell})}, \quad (32)$$

and so, on this event,

$$\begin{aligned} & \tilde{\Pi}_{cut}^{(\ell)}(A_n(z; \ell) \cap B_n | Y_{1:n}) \\ & \leq \frac{16}{\bar{q}^2} \frac{\int_{B_n} \mathbb{1}_{A_n(z; \ell)} (\sum_s f_s(Y_\ell)) \bar{L}_{-\ell}(\theta) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ}{\sum_s f_s^*(Y_\ell) M(Y_{1:n}^{-\ell})} \\ & \leq \frac{16}{\bar{q}^2} \left[\underbrace{\frac{\int_{B_n} \mathbb{1}_{A_n(z; \ell)} (\sum_s |f_s(Y_\ell) - f_s^*(Y_\ell)|) \bar{L}_{-\ell}(\theta) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ}{\sum_s f_s^*(Y_\ell) M(Y_{1:n}^{-\ell})}}_{J_1} \right. \\ & \quad \left. + \underbrace{\frac{\int_{B_n} \mathbb{1}_{A_n(z; \ell)} \bar{L}_{-\ell}(\theta) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ}{M(Y_{1:n}^{-\ell})}}_{J_2} \right]. \end{aligned}$$

Now, using $\mathbb{E}_*[J_1] = \mathbb{E}_*[\mathbb{E}_*[J_1|Y_{1:n}^{-\ell}]]$, together with the fact that $f(y_\ell|Y_{1:n}^{-\ell}) dy_\ell \leq \sum_s f_s^*(y_\ell) dy_\ell$, as well as the bound on $\|f_s - f_s^*\|_{L^1}$ on B_n , we obtain

$$\mathbb{E}_*[J_1] \leq \mathbb{E}_* \left[\frac{\int_{B_n} \sum_s \|f_s - f_s^*\|_1 \bar{L}_{-\ell}(\theta) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ}{M(Y_{1:n}^{-\ell})} \right] \leq RK_n\epsilon_n.$$

To bound J_2 , note that $\mathbb{1}_{A_n(z; \ell)} \leq |f_r(Y_\ell) - f_r^*(Y_\ell)| / (z\epsilon_n \sum_r f_r^*(Y_\ell))$, so that

$$J_2 \leq \frac{\int_{B_n} |f_r(Y_\ell) - f_r^*(Y_\ell)| \bar{L}_{-\ell}(\theta) \Pi_2(d\mathbf{f}) \phi_n(Q|Y_{1:n}^{-\ell}) dQ}{z\epsilon_n M(Y_{1:n}^{-\ell}) \sum_s f_s^*(Y_\ell)},$$

and bounding as before with J_1 , we obtain $\mathbb{E}_*[J_2] \leq \frac{RK_n}{z}$. This implies that

$$\begin{aligned} & \mathbb{P}_* \left[\sum_{|\ell-k| \leq T_n} \tilde{\Pi}_{cut}^{(\ell)}(A_n(\rho^{-|\ell-k|/2} z_n; \ell) \cap B_n | Y_{1:n}) > \epsilon \right] \\ & \leq \sum_{|\ell-k| \leq T_n} \mathbb{P}_* \left[D_n(\ell) \leq \frac{c_*(Y_\ell) M(Y_{1:n}^{-\ell})}{4} \right] + T_n \frac{16RK_n\epsilon_n}{\epsilon \bar{q}^2} + \sum_{|\ell-k| \leq T_n} \frac{\rho^{|\ell-k|/2} K_n}{\epsilon z_n} \\ & \lesssim \frac{T_n K_n \epsilon_n}{\epsilon} + \frac{K_n}{z_n \epsilon}. \end{aligned}$$

Recall that $T_n \leq \epsilon^2 [K_n \epsilon_n]^{-1}$ and $z_n \geq K_n / \epsilon^2$, so that the above bound becomes

$$\mathbb{P}_* \left[\sum_{|\ell-k| \leq T_n} \tilde{\Pi}_{cut}^{(\ell)}(A_n(\rho^{-|\ell-k|/2} z_n; \ell) \cap B_n | Y_{1:n}) > \epsilon \right] \leq c_1 \epsilon, \quad (33)$$

for some constant c_1 independent on n and ϵ .

Returning to the original goal of bounding (29), and having now bounded the sum over $|\ell - k| \leq T_n$, we study the sum over $|\ell - k| > T_n$. For this, we will bound the sum under $\tilde{\Pi}_{cut}$ defined previously. First note that

$$\begin{aligned} \mathbb{P}_* \left[\exists \ell, c_*(Y_\ell) \leq e^{-L_n n \epsilon_n^2} \right] &\leq n \int_{c_*(y) \leq e^{-L_n n \epsilon_n^2}} \sum_{s=1}^R p_s^* f_s^*(y) dy \\ &\leq n \int_{c_*(y) \leq e^{-L_n n \epsilon_n^2}} \sqrt{c^*(y)} e^{-L_n n \epsilon_n^2 / 2} dy \\ &\leq C^* n e^{-L_n n \epsilon_n^2 / 2}, \end{aligned}$$

where we use that $c_*(y) \geq \sum_s p_s^* f_s^*(y)$, and that $\sqrt{c^*}$ is L^1 bounded by some C^* from condition (18). As soon as $L_n n \epsilon_n^2 \geq L_0 \log n$ with L_0 large enough, which is possible since $n \epsilon_n^2 \gtrsim \log n$ by assumption, we have $\mathbb{P}_* \left[\exists \ell, c_*(Y_\ell) \leq e^{-L_n n \epsilon_n^2} \right] = o(1)$ and we can bound $c_*(Y_\ell) \geq e^{-L_n n \epsilon_n^2}$ for all ℓ . Next, we define

$$\bar{A}_n(\ell) = \left\{ \theta; \sum_{|\ell-k| > T_n} \rho^{|\ell-k|} |f_r(Y_\ell) - f_r^*(Y_\ell)| \leq z_n \epsilon_n e^{-L_n n \epsilon_n^2} \right\}.$$

We write

$$\begin{aligned} \tilde{\Pi}_{cut}(B_n \cap \bar{A}_n^c(\ell) | Y_{1:n}) &= \frac{\int_{B_n \cap \bar{A}_n^c(\ell)} e^{\ell_n(\theta) - \ell_n(\theta^*)} \Pi_2(d\mathbf{f}) \phi_n(Q | Y_{1:n}) dQ}{\int_{B_n} e^{\ell_n(\theta) - \ell_n(\theta^*)} \Pi_2(d\mathbf{f}) \phi_n(Q | Y_{1:n}) dQ} \\ &=: \frac{N_n(\bar{A}_n(\ell)^c)}{D_n}, \end{aligned}$$

where we write the denominator as D_n and write $N_n(\bar{A}_n(\ell)^c)$ for the numerator. Using Lemma 9 which shows that $\mathbb{P}_*[D_n < e^{-L_n n \epsilon_n^2}] \lesssim L_n^{-1}$, together with $\phi_n(Q | Y_{1:n}) \lesssim n^{R(R-1)/2}$, we obtain, with probability greater than $1 - O(1/L_n)$, and with L_n going arbitrarily slowly to infinity,

$$\begin{aligned} \tilde{\Pi}_{cut}(B_n \cap \bar{A}_n(\ell)^c | Y_{1:n}) &\leq e^{L_n n \epsilon_n^2} \int_{B_n} \mathbf{1}_{\bar{A}_n(\ell)^c} e^{\ell_n(\theta) - \ell_n(\theta^*)} \Pi_2(d\mathbf{f}) \phi_n(Q | Y_{1:n}) dQ \\ &\leq \frac{C_* e^{2L_n n \epsilon_n^2} n^{R(R-1)/2}}{z_n \epsilon_n} \sum_{|\ell-k| > T_n} \rho^{|\ell-k|} \int_{B_n} |f_r(Y_\ell) - f_r^*(Y_\ell)| e^{\ell_n(\theta) - \ell_n(\theta^*)} \Pi_2(d\mathbf{f}) dQ \\ &\leq C_* e^{3L_n n \epsilon_n^2} \sum_{|\ell-k| > T_n} \rho^{|\ell-k|} \int_{B_n} |f_r(Y_\ell) - f_r^*(Y_\ell)| e^{\ell_n(\theta) - \ell_n(\theta^*)} \Pi_2(d\mathbf{f}) dQ, \end{aligned}$$

where, in the final line, we again use that $\log n \lesssim n \epsilon_n^2$ to consolidate the exponential. This implies that $\tilde{\Pi}_{cut}(B_n \cap \bar{A}_n(\ell)^c | Y_{1:n})$ is controlled under \mathbb{P}_* as

$$\begin{aligned} \mathbb{P}_* \left[\tilde{\Pi}_{cut}(B_n \cap \bar{A}_n(\ell)^c | Y_{1:n}) \geq 2/L_n \right] &\leq o(1) \\ &+ \mathbb{P}_* \left[\sum_{|\ell-k| > T_n} \rho^{|\ell-k|} \int_{B_n} |f_r(Y_\ell) - f_r^*(Y_\ell)| e^{\ell_n(\theta) - \ell_n(\theta^*)} \Pi_2(d\mathbf{f}) dQ > \frac{e^{-3L_n n \epsilon_n^2}}{L_n} \right] \end{aligned}$$

$$\begin{aligned}
 &\leq \mathbb{P}_* \left[\sum_{|\ell-k|>T_n} \int_{B_n} |f_r(Y_\ell) - f_r^*(Y_\ell)| e^{\ell_n(\theta) - \ell_n(\theta^*)} \Pi_2(d\mathbf{f}) dQ > \frac{\rho^{-T_n} e^{-3L_n n \epsilon_n^2}}{L_n} \right] \\
 &\quad + o(1) \\
 &\leq \frac{C_* L_n \rho^{T_n} e^{3L_n n \epsilon_n^2}}{L_n} \sum_{|\ell-k|>T_n} \int_{B_n} \mathbb{E}_{Q, \mathbf{f}} [|f_r(Y_\ell) - f_r^*(Y_\ell)|] \Pi_2(d\mathbf{f}) dQ + o(1) \\
 &\leq C_* L_n \rho^{T_n} e^{3L_n n \epsilon_n^2} \sum_{s=1}^R \int_{B_n} \int_{B_n} f_s(y) (f_r(y) + f_r^*(y)) dy \Pi_2(d\mathbf{f}) + o(1) \\
 &\leq C_* L_n \rho^{T_n} e^{3L_n n \epsilon_n^2} \sum_s \int_{B_n} \|f_s\|_2 (\|f_r\|_2 + \|f_r^*\|_2) \Pi_2(d\mathbf{f}) + o(1) \\
 &\leq C_* \rho^{T_n} e^{4L_n n \epsilon_n^2} + o(1) = o(1), \tag{34}
 \end{aligned}$$

with the final equality holding as soon as $T_n \log(1/\rho) \geq 5nL_n \epsilon_n^2$. We recall that, in order to establish Equation (33) earlier in the proof, we required that $T_n \epsilon_n \rightarrow 0$ arbitrarily slowly. We thus use the assumption $n \epsilon_n^3 \rightarrow 0$ to choose T_n so that both conditions may hold simultaneously.

Finally, combining (34) with (33) and (29), and since K_n and L_n can be chosen to go arbitrarily slowly to infinity and $n \epsilon_n^2 \gtrsim \log n$, we obtain that, for any z_n going to infinity,

$$\Pi_{cut}(\Delta_k > z_n \epsilon_n | Y_{1:n}) = o_{\mathbb{P}_*}(1),$$

which terminates the proof. □

7. Conclusion and discussion

In this paper we use the cut posterior approach of [36] for inference in semiparametric models, which we apply to the nonparametric Hidden Markov models. A difficulty with the Bayesian approach in high or infinite dimension is that it is difficult (if not impossible) to construct priors on these complex models which allow for *good* simultaneous inference on a collection of parameters of interest, where *good* means having good frequentist properties (and thus leading to some robustness with respect to the choice of the prior). As mentioned in Section 1, a number of examples have been exhibited in the literature where *reasonable* priors lead to poorly behaved posterior distributions for some functionals of the parameters. We believe that the cut posterior approach is an interesting direction to pursue in order to address this general problem, and we demonstrate in the special case of semiparametric Hidden Markov models that it leads to interesting properties of the posterior distribution and is computationally tractable.

Our approach is based on a very simple prior Π_1 on Q, \mathbf{f} used for the estimation of Q based on finite histograms with a small number of bins. This enjoys a Bernstein-von Mises property, so that credible regions for Q are also asymptotic confidence regions. Moreover, by choosing M large (but not too large), we

obtain efficient estimators. Proving efficiency for semiparametric HMMs is non trivial and our proof has independent interest. Another original and important contribution of our paper is an inversion inequality (stability estimate) comparing the L^1 distance between $g_{Q, \mathbf{f}_1}^{(3)}$ and $g_{Q, \mathbf{f}_2}^{(3)}$ and the L^1 distance between \mathbf{f}_1 and \mathbf{f}_2 . Finally, another interesting contribution is our control of the error of the estimates of the smoothing probabilities using a Bayesian approach, which is based on a control under the posterior distribution of

$$\frac{1}{n} \sum_{i=1}^n \frac{|f_r(y_i) - f_r^*(y_i)|}{\sum_r f_r^*(y_i)},$$

despite the double use of the data (in the posterior on f_r and in the empirical distance above). It is a rather surprising result, which does not hold if \hat{f}_r is replaced with \hat{f}_r constructed using y_1, \dots, y_n , unless a sup-norm bound on $\hat{f}_r - f_r^*$ is obtained.

Section 5 demonstrates clearly the estimation procedures and highlights the importance of choosing a small number of bins in practical situations. Here, there remains open the question of how to choose this number in a principled way, and an interesting extension of the work would follow along the lines of the third section of [27], in which the authors produce an oracle inequality to justify a cross-validation scheme for choosing the number of bins.

Finally, the paper deals with the case where R is known. This assumption is rather common both in theory (see for instance the works of [26], [16] or [7]) and in practice (for instance in genomic applications as in [60]). The identifiability results of [4] and [24] do show that R can be identified as well, leaving open the possibility to jointly estimate R , together with Q and \mathbf{f} . We leave this direction of research for future work.

Appendix A: Proofs of main results

In the following section, we detail the proofs of Theorem 3 and Proposition 2.

A.1. Proof of Theorem 3

The proof of Theorem 3 follows from Proposition 1 and Lemma 8.

Lemma 8. *Let $(\mathcal{I}_M)_M$ be a sequence of embedded partitions, let M be such that \mathcal{I}_M is admissible for \mathbf{f}^* , and grant Assumptions 1, 3 and 4. Then the matrix*

$$\tilde{J}_{M+1} - \tilde{J}_M$$

is positive semi-definite.

Proof. Consider the MLE for $\theta = (Q(\theta), \underline{\omega}_M(\theta))$ for parameters in $\mathcal{Q} \times \Omega_M$ in the model with partition \mathcal{I}_M .

By the remarks at the end of the proof of Theorem 1, the MLE is regular. This means that, for all parameter sequences θ_n of the form $\theta_0 + n^{-\frac{1}{2}}h$ with

$h \in [-H, H]$ for some fixed H , $\sqrt{n}(\hat{\theta}^{(n)} - \theta_n)$ converges in distribution under \mathbb{P}_{θ_n} to a variable, say Z , of fixed distribution, say μ_Z . Explicitly, this means that for Borel sets A we have

$$\mathbb{P}_{\theta_n} \left(\sqrt{n}(\hat{\theta}^{(n)} - \theta_n) \in A \right) \rightarrow \mu_Z(A) \text{ as } n \rightarrow \infty.$$

Here θ_n is a sequence of parameters in $\mathcal{Q} \times \Omega_M$. However, any likelihood-based estimator depends on the observations only through the probabilities of each bin assignment, by considering the multinomial model with the count data. Thus, we can replace \mathbb{P}_{θ_n} with any law from the full semiparametric model for which the transition matrix Q_n is the same and the functions f_n have the corresponding $\omega_M(\theta_n)$ as the bin assignment probabilities, and the above convergence remains true.

Since \mathcal{I}_{M+1} is made by splitting bins in \mathcal{I}_M , a collection of weights $\omega^{(M+1)} = (\omega_m^{(M+1)} : m \leq \kappa_{M+1})$ for the partition \mathcal{I}_{M+1} is identified with weights $\omega^{(M)}$ for the partition \mathcal{I}_M . We set $\omega_m^{(M+1)} = \sum_{m' \in \text{Ch}_M(m)} \omega_{m'}^{(M)}$, where $\text{Ch}_M(m)$ satisfies $I_m^{(M)} = \bigcup_{m' \in \text{Ch}_M(m)} I_{m'}^{(M+1)}$. This implies that the MLE for the model $\mathcal{M}(\mathcal{I}_M)$ may be obtained by maximising the likelihood in model $\mathcal{M}(\mathcal{I}_{M+1})$ subject to the constraint that the quantities $\frac{\omega_{m'}^{(M+1)}}{|I_{m'}^{(M+1)}|}$ coincide for all $m' \in \text{Ch}(m)$. Regularity of the constrained MLE within \mathcal{I}_{M+1} is then inherited from its regularity as a global MLE in \mathcal{I}_M , by the preceding discussion.

Moreover, its asymptotic variance, which is given by the inverse Fisher information \tilde{J}_M^{-1} for the model with partition \mathcal{I}_M , is lower bounded by the inverse Fisher information for the model with partition \mathcal{I}_{M+1} by Lemma 11, hence $\tilde{J}_M^{-1} - \tilde{J}_{M+1}^{-1}$ is positive semi definite. \square

Proof of Theorem 3. We have, by inspecting the part of the proof of Lemma 8 concerning regularity, that under any sequence of the form $\mathbb{P}_*^n = \mathbb{P}_{\theta^*(a, \mathbf{h})}$, where $\theta^*(a, \mathbf{h}) = (Q^* + \frac{a}{\sqrt{n}}, f_r^*(1 + \frac{1}{\sqrt{n}}h_r) : r \in [R])$ with $a \in \mathbb{R}^{R(R-1)}$ and $h_r \in L_0^2(f_r^*)$,

$$R_{M,n} := d_{BL} \left(\mathbb{P}_*^n \left(\sqrt{n}(\hat{Q}_n^{(M)} - (Q^* + \frac{a}{\sqrt{n}})) \in \cdot \right), \mu_M \right) \rightarrow 0,$$

where μ_M is Gaussian of covariance equal to \tilde{J}_M^{-1} by Theorem 1. Here d_{BL} is the bounded Lipschitz metric which metrizes weak convergence. By taking $M_n \rightarrow \infty$ sufficiently slowly, we have $R_{M_n,n} \rightarrow 0$ also. Since, for any $t \in \mathbb{R}^{R(R-1)}$, $(t^T \tilde{J}_{M_n} t)^{-1}$ is decreasing by Lemma 8 and bounded below by $(t^T \tilde{J} t)^{-1}$ the efficient variance, it converges and the limit is equal to $(t^T \tilde{J} t)^{-1}$ by Proposition 1. This implies weak convergence of the measures μ_{M_n} to μ , the centred Gaussian measure whose variance is \tilde{J}^{-1} , and applying the triangle inequality in the metric d_{BL} proves the first claim.

For the Bernstein-von Mises result, we have from Theorem 2 that

$$\left\| \Pi_M(\sqrt{n}(Q - \hat{Q}_n)|Y_{1:n}) - \mathcal{N}(0, \tilde{J}_M^{-1}) \right\|_{TV} = \tilde{R}_{M,n},$$

where $\tilde{R}_{M,n} \rightarrow 0$ in \mathbb{P}_* -probability as $n \rightarrow \infty$. Refining the sequence from before, so that $M_n \rightarrow \infty$ sufficiently slowly that we additionally have $\tilde{R}_{M,n} \rightarrow 0$, and noting that $\tilde{J}_{M_n} \rightarrow \tilde{J}$, we get

$$\begin{aligned} & \left\| \Pi_{M_n}(\sqrt{n}(Q - \hat{Q}_n) | Y_{1:n}) - \mathcal{N}(0, \tilde{J}_\infty^{-1}) \right\|_{TV} \\ & \leq \left\| \Pi_{M_n}(\sqrt{n}(Q - \hat{Q}_n) | Y_{1:n}) - \mathcal{N}(0, \tilde{J}_{M_n}^{-1}) \right\|_{TV} + \left\| \mathcal{N}(0, \tilde{J}^{-1}) - \mathcal{N}(0, \tilde{J}_{M_n}^{-1}) \right\|_{TV} \\ & = o_{\mathbb{P}_*}(1). \end{aligned}$$

□

A.2. Proof of Proposition 2

The proof of Proposition 2 is a direct consequence of the following result, which shows that Theorem 7 applies to semiparametric HMMs of the type described in Section 2.2. For the construction of Π_1 , a C_1 diffeomorphism ϕ is implicitly used (if necessary) in order to consider histograms on $[0, 1]$

Proposition 3. *Let $(Y_t)_{t \geq 1}$ be observations from a finite state space HMM with transition matrix Q and emission densities $\mathbf{f} = (f_r)_{r=1, \dots, R}$. Consider the cut posterior based on Π_1 associated to the partition \mathcal{I}_M and Π_2 .*

(i) *Under **Assumptions 1-5**, **Assumption 7** is satisfied with $\mathcal{T}_n = \{Q; \|Q - Q^*\| \leq z_n/\sqrt{n}\}$, with $z_n \rightarrow \infty$ sufficiently slowly and ϕ_n the restriction to \mathcal{T}_n of the Gaussian density centered at $\hat{Q}_{n,M}$ with variance \tilde{J}_M^{-1}/n .*

(ii) *Choosing Π_2 such that **Assumption 6** is verified for suitable $\epsilon_n, \tilde{\epsilon}_n$ with $n\epsilon_n^2 \gtrsim \log n$, the assumptions of Theorem 7 are verified for the same ϵ_n and any $K_n \rightarrow \infty$.*

Proof. (i) The idea of the proof is as follows: We know from Theorem 2 that $\Pi_1(\cdot | Y_{1:n})$ is close in total variation distance to the normal distribution centred at the estimator \hat{Q} . Then the restriction of this distribution to a ball of radius z_n , centred at Q^* , where $M_n \rightarrow \infty$ slowly, will be close in total variation to the original normal distribution with high probability as $n \rightarrow \infty$.

Let $p_{1,n}$ be the normal density of variance \tilde{J}_M^{-1}/n centred at \hat{Q} , where \hat{Q} is the estimator from Theorem 2. Let $p_{2,n}$ be the restriction of this normal density to $\{Q : \|Q - Q^*\| \leq z_n/\sqrt{n}\}$ where $z_n \rightarrow \infty$. Fix $\epsilon > 0$ and choose n large enough that $\mathbb{P}_*(\|\hat{Q} - Q^*\| > cz_n'/\sqrt{n}) < \epsilon$ where $1 > c > 0$ is a constant. Let $p_{3,n}$ be the density obtained by restriction of $p_{1,n}$ to $\{Q : \|\hat{Q} - Q^*\| \leq (1-c)z_n/\sqrt{n}\}$. Then, for sufficiently large n , we have with probability exceeding $1 - \epsilon$ that the support of $p_{3,n}$ is a subset of the support of p_2 , in which case

$$\|p_{1,n} - p_{2,n}\|_{L^1} \leq \|p_{1,n} - p_{3,n}\|_{L^1}.$$

But the variance of $p_{1,n}$ is of the order $O(n^{-1})$ and $z_n \rightarrow \infty$ so, as $n \rightarrow \infty$, the support of $p_{3,n}$ approaches the support of $p_{1,n}$ and $\|p_{1,n} - p_{3,n}\|_{L^1} \rightarrow 0$. But then $\|p_{1,n} - p_{2,n}\|_{L^1} \rightarrow 0$. Thus, with probability exceeding $1 - \epsilon$, we have

that $\|p_{1,n} - p_{2,n}\|_{L^1} \rightarrow 0$ as $n \rightarrow \infty$. Since $\epsilon > 0$ is arbitrary, we conclude $\|p_{1,n} - p_{2,n}\|_{L^1} = o_{\mathbb{P}_*}(1)$.

Taking $\phi_n = p_{2,n}$ as above, for any $z_n \rightarrow \infty$, we verify Assumption 7 for any ϵ_n for which $n\epsilon_n^2 \gtrsim \log n$.

(ii) This is a consequence of Theorem 3.1 in [57]. The choice of \mathcal{T}_n also provides verification of the hypothesis of Lemma 3.2 of [57] and establishes the required control on KL divergence described in (35). To satisfy the testing assumption (36), it suffices to note that, for large enough n , \mathcal{T}_n is a subset of $\{Q : Q_{i,j} \geq \frac{Q_{ij}^*}{2} \forall i, j \leq R\}$ under Assumption 3. \square

Appendix B: General theorem for cut posterior contraction

In this section, we present a general theory for contraction of cut posteriors, which is developed in the style of the usual theory for Bayesian posteriors of [29]. The main result of this section is Theorem 7, from which Proposition 2 follows.

Consider a general semiparametric model in which there is a finite dimensional parameter ϑ and an infinite-dimensional parameter η . Suppose we wish to estimate the pair (ϑ, η) governing the law $\mathbb{P}_{\vartheta, \eta}^n$ of a random sample $Y_{1:n} \in \mathcal{Y}^n$. Assume the data is generated by some true distribution $\mathbb{P}_*^n = \mathbb{P}_{\vartheta^*, \eta^*}^n$ and consider the following two models:

Model 1: Consider a model $\mathcal{T} \times \mathcal{F}_1$ on pairs (ϑ, η) such that $\vartheta^* \in \mathcal{T}$ but for which we do not require that $\eta^* \in \mathcal{F}_1$. Consider a joint prior Π_1 over this space yielding a marginal posterior on ϑ given by

$$\Pi_1(d\vartheta|Y_{1:n}) = \int_{\mathcal{F}_1} \Pi_1(d\vartheta, d\eta|Y_{1:n}).$$

Model 2: Consider a model on η , conditional on ϑ , with a prior $\Pi_2(\cdot|\vartheta)$. We denote the parameter set for η by \mathcal{F}_2 and we assume that $\eta^* \in \mathcal{F}_2$. We obtain in this model, a conditional posterior distribution $\Pi_2(d\eta|Y_{1:n}, \vartheta)$.

Definition 2. The *cut posterior* on (ϑ, η) is given by

$$\Pi_{\text{cut}}(d\vartheta, d\eta|Y_{1:n}) = \Pi_1(d\vartheta|Y_{1:n})\Pi_2(d\eta|\vartheta, Y_{1:n}).$$

Write $\mathbb{P}_{\vartheta, \eta}^n$ for the law of the observations $Y_{1:n}$. Define for some $\ell > 0$ and some loss function $d(\cdot, \cdot)$ on $\mathbb{P}_{\vartheta, \eta}^\ell$,

$$B_\ell(\epsilon) = \{(\vartheta, \eta) : d(\mathbb{P}_{\vartheta, \eta}^\ell, \mathbb{P}_{\vartheta^*, \eta^*}^\ell) < \epsilon\}, \quad M, \epsilon > 0,$$

with $B_\ell^c(\epsilon)$ its complement. Define also the Kullback-Leibler neighbourhoods of $\mathbb{P}_*^n = \mathbb{P}_{\vartheta^*, \eta^*}^n$ as

$$V_n(\mathbb{P}_*^n, \epsilon) = \{(\vartheta, \eta) : \mathcal{K}(\mathbb{P}_{\vartheta^*, \eta^*}^n | \mathbb{P}_{\vartheta, \eta}^n) < n\epsilon^2\},$$

where $\mathcal{K}(\mathbb{P}_{\vartheta^*, \eta^*}^n | \mathbb{P}_{\vartheta, \eta}^n) = \mathbb{E}_*^n \log \frac{d\mathbb{P}_{\vartheta^*, \eta^*}^n}{d\mathbb{P}_{\vartheta, \eta}^n}$ is the Kullback-Leibler divergence.

We now present a general theorem to characterize cut-posterior contraction result, in the spirit of the now classical result of [29]. Our main additional assumption for the cut setup is that we have sufficiently good control over Π_1 , similar to the kind established in Section 2 in the HMM setting, which we detail now.

Denote by $\pi_1(\cdot|Y_{1:n})$ the marginal posterior density of ϑ with respect to some measure μ_1 on \mathcal{T} , associated to the prior Π_1 on ϑ, η .

Assumption 7. For all sequences $z_n \rightarrow \infty$ there exist $\mathcal{T}_n \subset \mathcal{T}$ with $\vartheta^* \in \mathcal{T}_n$, $\epsilon_n = o(1)$, and a non-negative, random function ϕ_n on \mathcal{T}_n with $\text{supp}(\phi_n) \subset \mathcal{T}_n$, such that

$$\|\pi_1(\cdot|Y_{1:n}) - \phi_n(\cdot|Y_{1:n})\|_{L^1(\mu)} = o_{\mathbb{P}_*^n}(1),$$

and

$$\mu(\mathcal{T}_n) \mathbb{E}_*^n \left(\int_{\mathcal{T}_n} \phi_n^2(\vartheta) \mu(d\vartheta) \right) = O(z_n),$$

and $\exists K > 0$ such that

$$\mathbb{P}_*^n \left[\sup_{\vartheta \in \mathcal{T}_n} \phi_n(\vartheta) > e^{K n \epsilon_n^2} \right] = o(1)$$

Assumption 7 is mild, and holds for instance if ϕ_n is a Gaussian distribution with mean $\hat{\vartheta}$ and variance \mathbf{i}_*/n for some semi definite matrix \mathbf{i}_* , where $\hat{\vartheta} - \vartheta^* = O_{\mathbb{P}_*^n}(1/\sqrt{n})$. In this case, $\mathcal{T}_n = \{\|\vartheta - \vartheta^*\| \leq K_n/n\}$ with $K_n \rightarrow \infty$ arbitrarily slowly, since we have

$$\begin{aligned} \mathbb{E}_*^n \left(\int_{\mathcal{T}_n} \phi_n^2(\vartheta) \mu(d\vartheta) \right) &\lesssim n^{d/2}, \quad \text{vol}(\mathcal{T}_n) \lesssim n^{-d/2} K_n^{d/2}, \\ \sup_{\vartheta \in \mathcal{T}_n} \phi_n(\vartheta) &\lesssim n^{d/2} = o(e^{K_n n \epsilon_n^2}), \end{aligned}$$

as soon as $n \epsilon_n^2 \gtrsim \log n$.

Theorem 7. Let **Assumption 7** hold with $n \epsilon_n^2 \rightarrow \infty$, and assume that there exists $C > 0$ such that, for any $\vartheta \in \mathcal{T}_n$, sets $S_n(\vartheta) \subset \mathcal{F}_2$ satisfying

$$\bigcup_{\vartheta \in \mathcal{T}_n} \{\vartheta\} \times S_n(\vartheta) \subset V_n(\mathbb{P}_*^n, \epsilon_n), \quad \inf_{\vartheta \in \mathcal{T}_n} \Pi_2(S_n(\vartheta)|\vartheta) \geq e^{-C n \epsilon_n^2}. \quad (35)$$

Assume also that there exist $L_n \rightarrow \infty$, $\mathcal{F}_{2,n}(\vartheta) \subset \mathcal{F}_2$ and $\psi_n : \mathcal{Y}^n \rightarrow \{0, 1\}$ such that

$$\begin{aligned} \sup_{\vartheta \in \mathcal{T}_n} \Pi_2(\mathcal{F}_{2,n}(\vartheta)^c|\vartheta) &\leq e^{-L_n n \epsilon_n^2}, \quad \text{set } \mathcal{U}_n = \bigcup_{\vartheta \in \mathcal{T}_n} \{\vartheta\} \times \mathcal{F}_{2,n}(\vartheta), \\ \mathbb{E}_*^n \psi_n &= o(1), \quad \sup_{\substack{(\vartheta, \eta) \in \mathcal{U}_n \\ d(\mathbb{P}_{\vartheta, \eta}^n, \mathbb{P}_{\vartheta^*, \eta^*}^n) > K_n \epsilon_n}} \mathbb{E}_{\vartheta, \eta}^n (1 - \psi_n) &\leq C' e^{-2L_n n \epsilon_n^2}, \end{aligned} \quad (36)$$

for some $K_n, K > 0$. Then, as $n \rightarrow \infty$,

$$\Pi_{\text{cut}}(B_\ell^c(K_n \epsilon_n)|Y_{1:n}) = o_{\mathbb{P}_*^n}(1).$$

Remark 7. Typically $K_n = O(\sqrt{L_n})$, when the tests (36) are constructed as a supremum of local tests. Given some $L_n \rightarrow \infty$ for which the conditions are satisfied, we can choose $L_n \rightarrow \infty$ arbitrarily slowly, and consequently we can choose $K_n \rightarrow \infty$ arbitrarily slowly.

The proof of Theorem 7 which we present below is a rather simple adaptation of [29]. As in [30], we have simplified the common Kullback-Leibler neighbourhood assumption involving variances of the log-likelihood ratio using the technique of Lemma 6.26 therein.

Remark 8. By placing an additional assumption on neighbourhoods for higher-order Kullback-Leibler variations, the assumption on the existence of the sequence $L_n \rightarrow \infty$ can be replaced with an assumption that $L_n = L$ for a constant $L > 1 + C$ – the proof is similar but uses a different (though standard) technique for proving the evidence lower bound. In this case, we can choose K_n a constant.

Proof of Theorem 7. The proof is an adaptation of the proof on posterior contraction rates, as in [29]. Under **Assumption 7**, it suffices to prove the claim when we replace the cut posterior Π_{cut} with the distribution

$$\tilde{\Pi}(d\vartheta, d\eta) = \phi_n(\vartheta)\mu(d\vartheta)\Pi_2(d\eta|\vartheta, Y_{1:n}).$$

Write also $\bar{B}_\ell = B_\ell^c(K_n\epsilon_n) \cap (\mathcal{T}_n \times \mathcal{F}^R)$, the subset of $B_\ell^c(K_n\epsilon_n)$ on which $\tilde{\Pi}$ is supported, and $\delta_n = \mu(\mathcal{T}_n)$.

Writing $J_n(\vartheta)$ for the random variable $\Pi_2(\bar{B}_\ell|\vartheta, Y_{1:n})$, we have $J_n(\vartheta) = \frac{N_n(\vartheta)}{D_n(\vartheta)}$ where

$$N_n(\vartheta) = \int_{\bar{B}_\ell} e^{l_n(\vartheta, \eta) - l_n(\vartheta^*, \eta^*)} \Pi_2(d\eta|\vartheta);$$

$$D_n(\vartheta) = \int_{\mathcal{F}_2} e^{l_n(\vartheta, \eta) - l_n(\vartheta^*, \eta^*)} \Pi_2(d\eta|\vartheta).$$

Write also $\tilde{\Pi}(\bar{B}_\ell|Y_{1:n}) = \tilde{\Pi}(\bar{B}_\ell|Y_{1:n})\psi_n + \tilde{\Pi}(\bar{B}_\ell|Y_{1:n})(1 - \psi_n)\mathbb{1}_{\Omega_n^c} + \tilde{\Pi}(\bar{B}_\ell|Y_{1:n})(1 - \psi_n)\mathbb{1}_{\Omega_n}$, where $\Omega_n = \{\sup_{\vartheta \in \mathcal{T}_n} \phi_n(\vartheta) \leq e^{K_n\epsilon_n^2}\}$. By the testing assumption, the first term vanishes in probability, while the second term vanishes under Assumption 7. For the remaining term, we have, for $\tilde{L}_n \rightarrow \infty$ to be chosen later, that

$$\begin{aligned} & \mathbb{1}_{\Omega_n}(1 - \psi_n) \int_{\mathcal{T}_n} J_n(\vartheta)\phi_n(\vartheta)\mu(d\vartheta) \\ & \leq \int_{\mathcal{T}_n} \mathbb{1}_{\{D_n(\vartheta) \leq e^{-(C+\tilde{L}_n)n\epsilon_n^2}\}} \phi_n(\vartheta)\mu(d\vartheta) \tag{I_1} \\ & + (1 - \psi_n) \int_{\mathcal{T}_n} \mathbb{1}_{\Omega_n} N_n(\vartheta)e^{(C+\tilde{L}_n)n\epsilon_n^2} \phi_n(\vartheta)\mu(d\vartheta), \tag{I_2} \end{aligned}$$

using the fact that $J_n(\vartheta) \leq 1$. We can bound $\mathbb{E}_*[I_1]$ as

$$\mathbb{E}_*[I_1] \leq \left(\int_{\mathcal{T}_n} \mathbb{P}_* \{D_n(\vartheta) \leq e^{-(C+\tilde{L}_n)n\epsilon_n^2}\} \mu(d\vartheta) \right)^{\frac{1}{2}} \left[E_* \left(\int_{\mathcal{T}_n} \phi_n^2(\vartheta) \mu(d\vartheta) \right) \right]^{\frac{1}{2}},$$

by an application of Cauchy-Schwartz. Now, noting that by Lemma 9 and under the assumption of sufficient prior mass on the $S_n(\vartheta)$, we have for any $\tilde{L}_n \rightarrow \infty$ that

$$\mathbb{P}_* \{D_n(\vartheta) \leq e^{-(\tilde{L}_n+C)n\epsilon_n^2}\} = O(\tilde{L}_n^{-1}),$$

and so

$$\mathbb{E}_*[I_1] \lesssim \sqrt{\delta_n} \tilde{L}_n^{-1/2} \left(\int_{\mathcal{T}_n} \phi_n^2(\vartheta) \mu(d\vartheta) \right)^{\frac{1}{2}} = o(1)$$

under **Assumption 7**, by taking $z_n = o(\tilde{L}_n)$. We bound $\mathbb{E}_*[I_2]$ by

$$\begin{aligned} \mathbb{E}_*[I_2] &\leq e^{(C+\tilde{L}_n)n\epsilon_n^2} \mathbb{E}_* \left(\int_{\mathcal{T}_n} \int_{\tilde{B}_\ell} [e^{l_n(\vartheta, \eta) - l_n(\vartheta^*, \eta^*)} (1 - \psi_n) \phi_n] \Pi_2(d(\eta|\vartheta) \mu(d\vartheta)) \right) \\ &\leq e^{(C+\tilde{L}_n)n\epsilon_n^2} \left(\int_{\mathcal{T}_n} \int_{\tilde{B}_\ell} \mathbb{E}_{\vartheta, \eta} [(1 - \psi_n) \phi_n \mathbf{1}_{\Omega_n}] \Pi_2(d(\eta|\vartheta) \mu(d\vartheta)) \right). \end{aligned}$$

Using Fubini's theorem again, alongside the assumption on type II errors and on the sieves, and the deterministic bound on $\phi_n(\vartheta)$ over Ω_n from **Assumption 7**, we bound what precedes by

$$\begin{aligned} &\leq e^{(K+C+\tilde{L}_n)n\epsilon_n^2} \left(\int_{\mathcal{T}_n} \int_{\tilde{B}_\ell} \mathbb{E}_{\vartheta, \eta} [(1 - \psi_n) (\mathbf{1}_{\eta \in \mathcal{F}_{2,n}(\vartheta)} + \mathbf{1}_{\eta \in \mathcal{F}_{2,n}(\vartheta)^c})] \Pi_2(d(\eta|\vartheta) \mu(d\vartheta)) \right) \\ &\leq e^{(K+C+\tilde{L}_n)n\epsilon_n^2} \left\{ \left(\int_{\mathcal{T}_n} \int_{\tilde{B}_\ell} C' e^{-L_n n \epsilon_n^2} \Pi_2(d(\eta|\vartheta) \mu(d\vartheta)) \right) + \int_{\mathcal{T}_n} e^{-L_n n \epsilon_n^2} \mu(d\vartheta) \right\} \\ &\leq e^{(K+C+\tilde{L}_n)n\epsilon_n^2} (C' e^{-L_n n \epsilon_n^2} + e^{-L_n n \epsilon_n^2}). \end{aligned}$$

Choosing $\tilde{L}_n = o(L_n)$, we get the required convergence. □

Lemma 9 provides the lower bound on the denominator used in the proof of Theorem 7. The proof is standard but we include it for completeness, it follows almost exactly the proof of Lemma 6.26 in [30].

Lemma 9. *Let $A_n(\vartheta) = \{D_n(\vartheta) \geq \Pi_2(S_n(\vartheta)|\vartheta) e^{-n\tilde{L}_n\epsilon_n^2}\}$ where $\tilde{L}_n \rightarrow \infty$ and $S_n(\vartheta)$ is such that there exists \mathcal{T}_n with*

$$\bigcup_{\vartheta \in \mathcal{T}_n} \{\vartheta\} \times S_n(\vartheta) \subset V_0(P_0, \epsilon_n).$$

Then

$$\sup_{\vartheta \in \mathcal{T}_n} \mathbb{P}_*(A_n(\vartheta)^c) = O(\tilde{L}_n^{-1}) = o(1)$$

as $n \rightarrow \infty$.

Proof. We show that, for each ϑ , with \mathbb{P}_* -probability tending to one,

$$\int_{\mathcal{F}_2} e^{l_n(\vartheta, \eta)} - e^{l_n(\vartheta^*, \eta^*)} \Pi_2(d(\eta|\vartheta)) \geq \Pi_2(S_n(\vartheta)|\vartheta) e^{-nM_n \epsilon_n^2},$$

for any $M_n \rightarrow \infty$. It suffices to show the above equation holds when we restrict the integral to $S_n(\vartheta) \subset \mathcal{F}_2$, for which we have $\{\vartheta\} \times S_n(\vartheta) \subset V_n(\mathbb{P}_*, \epsilon_n)$. By dividing both sides by $\Pi_2(S_n(\vartheta)|\vartheta)$, we see that it suffices to show that, for any probability measure ν supported on $S_n = S_n(\vartheta)$, that

$$\int_{S_n} e^{l_n(\vartheta, \eta)} - e^{l_n(\vartheta^*, \eta^*)} d\nu(\eta) \geq e^{-nM_n \epsilon_n^2}.$$

By applying the logarithm to both sides and using Jensen's inequality, it suffices to show that, with high \mathbb{P}_* probability,

$$Z := \int_{S_n} l_n(\vartheta, \eta) - l_n(\vartheta^*, \eta^*) d\nu(\eta) \geq -nM_n \epsilon_n^2.$$

Arguing as in [30], we obtain

$$\mathbb{P}_*(Z < -nM_n \epsilon_n^2) = O(M_n^{-1}) = o(1)$$

as $n \rightarrow \infty$ for any $M_n \rightarrow \infty$. Since the bound does not depend on ϑ and uses only that $\{\vartheta\} \times S_n(\vartheta) \subset V_0(P_0, \epsilon_n)$, we obtain the desired bound on the supremum. \square

Appendix C: Proofs of technical results

Section C is devoted to a number of technical proofs which are required the main results, but are reasonably standard in their approach.

In Section C.1, we prove that the Fisher information matrix is invertible for general, discrete state-space, discrete observation HMMs. This is necessary to apply the results of [9] and [17] in the proofs of Theorems 1 and 2.

In Section C.2, we gather some properties of the Fisher information matrix, showing a Cramér-Rao bound for estimation in HMMs, and showing a local uniform convergence result for the expected information from n observations.

In Section C.3, we collect technical lemmas used for the deconvolution argument which forms the key part of the proof of Theorem 3.

In Section C.4, we state a result which implies the existence of an admissible partition.

In Section C.5, we gather some results from [19] concerning the ‘forgetting’ of the hidden chain.

In Section C.6, we collect technical lemmas which feature in the proofs of Theorems 4 and 6

C.1. Non-singularity of Fisher Information

In what follows, we establish invertibility of the Fisher Information matrix for general multinomial Hidden Markov models.

Proposition 4. *Consider a multinomial HMM with latent states $X_t \in \{1, \dots, R\}$ and discrete observations Y_t taking values in the set of basis vectors of \mathbb{R}^k , which we denote $\{e_1, \dots, e_M\}$. Denote by $Q \in \mathbb{R}^{R \times R}$ the transition matrix of (X_t) , and let $\Omega \in \mathbb{R}^{k \times R}$ be the matrix whose R columns $\omega_r = (\omega_{mr})_{m=1}^M$ are the emission probabilities for the r^{th} state, with $r \in [R]$.*

Denote $\theta = (Q, \Omega) \in \mathbb{R}^p$ the HMM parameter, and write $J(\theta)$ for the Fisher Information matrix with entries given by

$$J(\theta)_{ij} = - \lim_{n \rightarrow \infty} \mathbb{E}_\theta \left[\underbrace{\frac{\partial^2}{\partial \theta_i \partial \theta_j} \log p_\theta(Y_1, \dots, Y_n)}_{-J_n(\theta)_{ij}} \right].$$

Then, if Q, Ω have rank R , J is non-singular.

The idea of the proof is to exhibit estimators with L^2 risk of order $n^{-1/2}$. We then show that the local asymptotic minimax result of [28] implies that the existence of such estimators guarantees a non-singular Fisher information. We use the spectral estimators proposed in [6], the control of which is established in [48].

Proof of Proposition 4. By an application of the van Trees inequality, analogously to the derivation preceding Theorem 16 of [28], we obtain for the HMM that (for p the parameter dimension)

$$\begin{aligned} & \int_{B_p(0,1)} \mathbb{E}_{\theta^* + ch/\sqrt{n}}^{(n)} \left[\left(U^T (\hat{\theta}^{(n)} - (\theta^* + \frac{ch}{\sqrt{n}})) \right)^2 \right] q(h) dh \\ & \geq U^T \left(\frac{n}{c^2} J_q + n \int_{B_p(0,1)} \frac{1}{n} J_n(\theta^* + \frac{ch}{\sqrt{n}}) q(h) dh \right)^{-1} U, \end{aligned}$$

which holds for any vector U . Here J_n is the joint Fisher information for n observations as in Proposition 12, and q is a density on \mathbb{R}^p such that $J_q := \int_{\mathbb{R}^p} \nabla q \nabla q^T \frac{\mathbb{1}_{\{q>0\}}}{q} dx$ is non-singular. Rescaling, we get that, for any vector U ,

$$\begin{aligned} & \int_{B_p(0,1)} \mathbb{E}_{\theta^* + ch/\sqrt{n}}^{(n)} \left[\left(\sqrt{n} U^T (\hat{\theta}^{(n)} - (\theta^* + \frac{ch}{\sqrt{n}})) \right)^2 \right] q(h) dh \\ & \geq U^T \left(\frac{1}{c^2} J_q + \int_{B_p(0,1)} \frac{1}{n} J_n(\theta^* + \frac{ch}{\sqrt{n}}) q(h) dh \right)^{-1} U. \end{aligned}$$

Taking a limit in n and applying Lemma 12 stated below gives that the limit inferior of the left hand side is at least

$$\lim_n U^T \left(\frac{1}{c^2} J_q + \int_{B_p(0,1)} \frac{1}{n} J_n(\theta^* + \frac{ch}{\sqrt{n}}) q(h) dh \right)^{-1} U = U^T \left(\frac{1}{c^2} J_q + J(\theta^*) \right)^{-1} U.$$

Call the matrix on the right hand side which we invert J_c . It is indeed invertible for sufficiently large c , since J_q is invertible and the set of invertible matrices is open. Denote its matrix square root by $J_c^{\frac{1}{2}}$. Now suppose $\exists V^*$ such that $(V^*)^T J(\theta^*) V^* = 0$. Then, by writing $V = J_c^{-\frac{1}{2}} U$, we get

$$\begin{aligned} \sup_{\|U\|=1} U^T J_c^{-1} U &= \sup_{\|U\| \neq 0} \frac{U^T J_c^{-1} U}{U^T U} \\ &= \sup_{\|V\| \neq 0} \frac{V^T V}{V^T J_c V} \geq \frac{V^{*T} V^*}{V^{*T} J_c V^*} = A c^2, \end{aligned}$$

with A a fixed constant not depending on c . Taking the limit as $c \rightarrow \infty$ gives that (upper bounding also the averaging over the law qdh by the supremum over h)

$$\liminf_{c \rightarrow \infty} \sup_{\|U\|=1} \liminf_{n \rightarrow \infty} \sup_{\|h\| < 1} \mathbb{E}_{\theta^* + ch/\sqrt{n}}^{(n)} \left[\left(\sqrt{n} U^T (\hat{\theta}^{(n)} - (\theta^* + \frac{ch}{\sqrt{n}})) \right)^2 \right] = +\infty,$$

or equivalently,

$$\liminf_{c \rightarrow \infty} \sup_{\|U\|=1} \liminf_{n \rightarrow \infty} \sup_{\|\theta - \theta^*\| < \frac{c}{\sqrt{n}}} \mathbb{E}_{\theta}^{(n)} \left[\left(\sqrt{n} U^T (\hat{\theta}^{(n)} - \theta) \right)^2 \right] = +\infty,$$

which contradicts the local uniform bound of Proposition 5. □

The following proposition establishes the existence of an estimator with suitable risk, as required for the arguments of Proposition C.1.

Proposition 5. *Let $\theta = (Q, \Omega)$ be the parameter for the HMM described in Proposition 4 and let θ^* be such that the identifiability conditions of **Assumption 1** hold. Then there exists an estimator $\hat{\theta}$ such that, for any $\epsilon > 0$ sufficiently small, we have, as $n \rightarrow \infty$,*

$$\sup_{\|\theta - \theta^*\| < \epsilon} \mathbb{E}_{\theta} \|\hat{\theta} - \theta\|_2^2 \leq \frac{C}{n} (1 + o(1)),$$

up to label-swapping.

Proof. Let $\hat{\theta}$ be the spectral estimator constructed in [48]. We have that

$$\mathbb{E}_{\theta} \|\hat{\theta} - \theta\|_2^2 = \int_0^{\infty} t \mathbb{P}_{\theta}(\|\hat{\theta} - \theta\| > t) dt.$$

Setting $t = \frac{Cx}{\sqrt{n}}$ we get

$$\mathbb{E}_\theta \|\hat{\theta} - \theta\|_2^2 = \int_0^\infty \frac{C^2 x}{n} \mathbb{P}_\theta(\|\hat{\theta} - \theta\| > \frac{Cx}{\sqrt{n}}) dx.$$

The bound of Lemma 10 is valid for $x \geq 1$, and so

$$\mathbb{E}_\theta \|\hat{\theta} - \theta\|_2^2 \leq \frac{C^2}{n} + \frac{C^2}{n} \int_0^\infty x e^{-x^2} dx \leq \frac{3C^2}{2n}.$$

Since the bound is valid on a small neighbourhood around θ , it holds in supremum over that neighbourhood. \square

C.1.1. Construction of spectral estimators

Spectral estimation of HMMs has been addressed by a number of previous works [1, 2, 6, 7, 16, 33]. In [1], the authors exhibit estimators for emissions in parametric HMMs with a \sqrt{n} rate in probability, but we require convergence in expectation. In [16], a convergence in expectation is shown but the rate contains a logarithmic factor which is not sufficient for our use in the proof of Proposition C.1. In [2], the authors exhibit estimators for which the concentration is sub-Gaussian, which would permit integration to an in-expectation bound, but their work only concerns two-state HMMs. We defer the details of the construction to [48]. The final result which we arrive at in the reference is as follows:

Lemma 10 (Spectral Estimators). *Recall $\theta = (Q, \Omega)$ and let $\hat{\theta} = (\hat{Q}, \hat{\Omega})$ to be the spectral estimator detailed in [48]. Then there exists a $\tau \in \mathcal{S}_R$ such that, for all $\epsilon > 0$ small enough,*

$$\sup_{\|\theta - \theta^*\| < \epsilon} \mathbb{P}_\theta(\max_{i=1, \dots, R} \|\hat{\theta}_i - \tau \theta_i\| > \frac{Cx}{\sqrt{n}}) \leq e^{-x^2},$$

for some $C = C(R, Q)$.

C.2. Fisher Information and asymptotic lower bound

The following result gives a lower bound on the asymptotic variance of regular estimators and is used in the proof of Lemma 8.

Lemma 11. *Let $\hat{\theta}_n$ be a regular estimator in the histogram model $\mathcal{M}(\mathcal{I}_M)$ with M fixed. Let Z denote the random variable whose law is equal to the limit distribution of the scaled and centred estimates $\sqrt{n}(\hat{\theta}_n - \theta^*)$ under \mathbb{P}_* as $n \rightarrow \infty$. Then*

$$\text{Cov}(Z) \geq J(\theta^*),$$

where $J(\theta)$ is the Fisher information matrix at the parameter θ .

Proof. Consider estimation of the one-dimensional parameter $\lambda^T \theta$. Recall from Lemma 12 that for any sequence $\theta_n \rightarrow \theta$, we have $J(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} J_n(\theta_n)$, with J_n the joint Fisher information for n observations, as defined in Lemma 12.

We follow the arguments of Gill and Levit [31]. Let π be a fixed prior density on $[-1, 1]$ and $J(\pi) = \mathbb{E}[(\log \pi(\theta))'^2]$. For a given $H > 0$, let $\pi(H, n)$ be the rescaling of this prior to the interval $A = [\theta_0 - n^{-\frac{1}{2}}H, \theta_0 + n^{-\frac{1}{2}}H]$ for given $H > 0$. Then, applying the van Trees inequality, we get

$$\mathbb{E}[\lambda^T \hat{\theta}^n - \lambda^T \theta]^2 \geq \frac{1}{\mathbb{E} \lambda^T J_n(\theta) \lambda + n J(\pi) / H^2},$$

where the left hand expectation is taken over the joint law of the parameter and the data given the parameter (with the parameter distributed according to $\pi(H, n)$) and the expectation in the denominator is over θ having that law also. Dividing through gives

$$\mathbb{E}[\sqrt{n}(\hat{\theta}^n - \theta)]^2 \geq \frac{1}{\frac{1}{n} \mathbb{E} \lambda^T J_n(\theta) \lambda + J(\pi) / H^2}.$$

Taking first $n \rightarrow \infty$, then $H \rightarrow \infty$, we obtain the result, applying also Lemma 12 to get convergence of the expectation term in the denominator to the Fisher information. \square

The previous result requires the following convergence property of the Fisher information matrices. We also use it in the proof of Proposition C.1.

Lemma 12. *Let $\theta_n \rightarrow \theta$. Denote $J_n = J_n^{(M)}$ the Fisher information for n observations in the model with M bins given by*

$$J_n(\theta) = \mathbb{E}_\theta[-\nabla_\theta^2 \log L_n(\theta)],$$

and J the Fisher information for the model. Then

$$\frac{1}{n} J_n(\theta_n) \rightarrow J(\theta).$$

When $\theta_n = \theta$ the result is simply the definition of J . The interest is hence in establishing local uniform convergence.

Proof. Theorem 3 of [19] establishes the result in the case of observed information⁴,

$$-\frac{1}{n} \nabla_\theta^2 l_n(\theta_n) \rightarrow J(\theta),$$

a.s. under \mathbb{P}_θ . Moreover, the Fisher and Louis identities [41] show that boundedness of the complete observed information and complete scores (which take simple forms) implies boundedness of the observed information. Write $\mathcal{Y} = [M]^{\mathbb{N}}$ and μ for the counting measure on this space with associated densities $p_\theta(y)$. We

⁴The result as stated in the reference conditions on the value $X_0 = x_0$ - the general statement follows and can be found explicitly in Theorem 13.24 in [20].

have since the observed information is bounded and $p_{\theta_n}(y) \rightarrow p_{\theta}(y)$ pointwise and so in $L^1(\mu)$ by Scheffé’s lemma [52],

$$\begin{aligned} \frac{1}{n} J_n(\theta_n) &= \int_{\mathcal{Y}} -\frac{1}{n} \nabla_{\theta}^2 l_n(\theta_n) p_{\theta_n}(y) d\mu(y) \\ &= \int_{\mathcal{Y}} -\frac{1}{n} \nabla_{\theta}^2 l_n(\theta_n) p_{\theta}(y) d\mu(y) + o(1) \\ &= J(\theta^*) + o(1). \end{aligned}$$

□

C.3. Technical results for the deconvolution argument

Here, we collect the technical results required to prove Lemma 2.

Let us recall some notation for what follows. For $r \in \{1, \dots, R\}$, and $g \in L^2(f_r^*)$, define

$$H_{r,M}(g) = g(y_0) \mathbb{P}_*(X_0 = r | \mathcal{G}_m^0) + \sum_{j=-\infty}^{-1} g(y_j) (\mathbb{P}_*(X_j = r | \mathcal{G}_M^0) - \mathbb{P}_*(X_j = r | \mathcal{G}_M^{-1})),$$

and define $H_r(g)$ similarly but with the conditioning being on the sigma algebras \mathcal{G}^0 and \mathcal{G}^{-1} . Then define, for $g = (g_1, \dots, g_r)$,

$$H_M(g) = \sum_r H_{r,M}(g_r); \quad H(g) = \sum_r H_r(g_r).$$

These are the score functions in the sub-models with fixed transition matrix, and emission densities varying along the path characterised by g .

Our first technical result allows us to eliminate certain terms when we make the main deconvolution argument. We recall $\mathcal{G}^\ell = \sigma(Y_{-\infty:\ell})$.

Lemma 13. *For $g \in \mathcal{H}_r$ and $j < k \leq 0$, we have*

$$\mathbb{E}_*[g(Y_j) \mathbb{P}_*(X_j = r | \mathcal{G}^0) | Y_{k:0}] = 0.$$

Proof. We first note that

$$\begin{aligned} \mathbb{E}_*[g(Y_j) \mathbb{P}(X_j = r | \mathcal{G}^0) | Y_{k:0}] &= \mathbb{E}_*[g(Y_j) \mathbb{1}_{X_j=r} | Y_{k:0}] \\ &= \mathbb{E}_*[\mathbb{E}_*[g(Y_j) \mathbb{1}_{X_j=r} | Y_{k:0}, X_j] | Y_{k:0}], \end{aligned}$$

by using usual properties of conditional expectations. Given X_j, Y_j and X_j are conditionally independent of $Y_{k:0}$ and thus $\mathbb{E}_*[g(Y_j) \mathbb{1}_{X_j=r} | Y_0, X_j] = \mathbb{E}_*[g(Y_j) \mathbb{1}_{X_j=r} | X_j]$. Since X_j has values in the discrete set $[R]$, we can explicitly write the conditional expectation as

$$\begin{aligned} \mathbb{E}_*[g(Y_j) \mathbb{1}_{X_j=r} | X_j] &= \sum_{r'=1}^R \mathbb{E}_*[g(Y_j) \mathbb{1}_{X_j=r} | X_j = r'] \mathbb{1}_{X_j=r'} \\ &= \mathbb{E}_*[g(Y_j) | X_j = r] \mathbb{1}_{X_j=r} \end{aligned}$$

But, since $g \in L^2(f_r^* dx)$ has zero expectation against the emission distribution for state r , the right hand side vanishes and hence so does its expectation given $Y_{k:0}$. \square

In the following section, we relate the score functions H_M in the binned model to score functions H in the full model, where the perturbation g is in the ‘direction’ of a histogram.

C.3.1. On the difference $H_M - H$

Lemma 14. *If $h_M = (h_{r,M})_{r \leq R}$ with $\sup_M \sum_r \|h_{r,M}\|_{L^2(f_r^*)} \leq L < \infty$, then*

$$H_M(h_M) = H(h_M) + o_{L^2}(1).$$

Proof. We have, for each r ,

$$\begin{aligned} H_{r,M}(h_{r,M}) &= h_{r,M}(y_0) \mathbb{P}_*(X_0 = r | \mathcal{G}_M^0) \\ &\quad + \sum_{j=-\infty}^{-1} h_{r,M}(y_j) (\mathbb{P}_*(X_j = r | \mathcal{G}_M^0) - \mathbb{P}_*(X_j = r | \mathcal{G}_M^{-1})); \\ H_r(h_{r,M}) &= h_{r,M}(y_0) \mathbb{P}_*(X_0 = r | \mathcal{G}^0) \\ &\quad + \sum_{j=-\infty}^{-1} h_{r,M}(y_j) (\mathbb{P}_*(X_j = r | \mathcal{G}^0) - \mathbb{P}_*(X_j = r | \mathcal{G}^{-1})). \end{aligned} \tag{37}$$

We will show that $H_{r,M}(h_{r,M}) - H_r(h_{r,M}) = o_{L^2}(1)$. Let $\epsilon > 0$ be arbitrary, and $J = J(L, \epsilon) > 0$ be such that $2L\rho^{J-1} < \frac{\epsilon}{3}(1 - \rho)$. Then Lemma 19 implies that

$$\begin{aligned} &\left\| \sum_{j=-\infty}^{-J} h_{r,M}(y_j) (\mathbb{P}_*(X_j = r | \mathcal{G}^0) - \mathbb{P}_*(X_j = r | \mathcal{G}^{-1})) \right. \\ &\quad \left. + \sum_{j=-\infty}^{-J} h_{r,M}(y_j) (\mathbb{P}_*(X_j = r | \mathcal{G}_M^0) - \mathbb{P}_*(X_j = r | \mathcal{G}_M^{-1})) \right\|_{L^2(\mathbb{P}_*)} \leq \frac{2\epsilon}{3}. \end{aligned}$$

We have now reduced the problem to the case where the sums in (37) are over $j > -J$. Since $\|h_M\|_{L^2} = O(1)$, to show convergence of the remaining difference between $H_{r,M}$ and H_r , it suffices to show that $\mathbb{P}_*(X_j = r | \mathcal{G}_M^\ell) \rightarrow \mathbb{P}_*(X_j = r | \mathcal{G}^\ell)$ in $L^\infty(\mathbb{P}_*)$ for $\ell = 0, 1$ and $-J \leq j \leq 0$.

We use Lemma 18. Let $K > J$. We then obtain, for all $j \geq -J$ and $\ell = 0, -1$,

$$\begin{aligned} \mathbb{P}_*(X_j = r | \mathcal{G}^\ell) &= \sum_{s=1}^R \mathbb{P}_*(X_j = r | Y_{-K:\ell}, X_K = s) \mathbb{P}_*(X_{-K} = s | \mathcal{G}_\infty^\ell) \\ &= \sum_{s=1}^R \mathbb{P}_*(X_j = r | Y_{-K:0}) (1 + O(\rho^{j+K})) \mathbb{P}_*(X_{-K} = s | \mathcal{G}_\infty^\ell) \end{aligned}$$

$$\begin{aligned} &= (1 + O(\rho^{j+K})) \sum_{s=1}^R \mathbb{P}_*(X_j = r|Y_{-K:0})\mathbb{P}_*(X_{-K} = s|\mathcal{G}_\infty^\ell) \\ &= (1 + O(\rho^{j+K})) \mathbb{P}_*(X_j = r|Y_{-K:0}). \end{aligned}$$

Similarly, $\mathbb{P}_*(X_j = r|\mathcal{G}_M^\ell) = (1 + O(\rho^{j+K}))\mathbb{P}_*(X_j = r|Y_{-K:0}^{(M)})$. Then

$$\begin{aligned} &\sup_{Y_{-\infty:0}} |\mathbb{P}_*(X_j = r|\mathcal{G}_M^\ell) - \mathbb{P}_*(X_j = r|\mathcal{G}^\ell)| \\ &\leq \sup_{Y_{-\infty:0}} |\mathbb{P}_*(X_j = r|\mathcal{G}_M^\ell) - \mathbb{P}_*(X_j = r|Y_{-K:\ell}^{(M)})| \\ &\quad + \sup_{Y_{-\infty:0}} |\mathbb{P}_*(X_j = r|Y_{-K:\ell}^{(M)}) - \mathbb{P}_*(X_j = r|Y_{-K:\ell})| \\ &\quad + \sup_{Y_{-\infty:0}} |\mathbb{P}_*(X_j = r|Y_{-K:\ell}) - \mathbb{P}_*(X_j = r|\mathcal{G}^\ell)|. \end{aligned}$$

By choosing K sufficiently larger than J , we have seen that the first and third terms are small, and bounded above by $\frac{\epsilon}{9JL}$ uniformly in $j \geq -J$. It remains to bound the second term. We have

$$\begin{aligned} &|\mathbb{P}_*(X_j = r|Y_{-K:\ell}^{(M)}) - \mathbb{P}_*(X_j = r|Y_{-K:\ell})| \\ &= \mathbb{P}_*(X_j = r) \left| \frac{\mathbb{P}_*(Y_{-K:\ell}^{(M)}|X_j = r)}{\mathbb{P}_*(Y_{-K:\ell}^{(M)})} - \frac{\mathbb{P}_*(Y_{-K:\ell}|X_j = r)}{\mathbb{P}_*(Y_{-K:\ell})} \right|. \end{aligned}$$

We write

$$\mathbb{P}_*(Y_{-K:\ell}^{(M)}|X_j = r) = \sum_{x_{-K}, \dots, x_\ell} \mathbf{1}_{x_j=r} \prod_{t=-K}^\ell f_{\omega_{x_t}^*}(Y_t)\mathbb{P}_*(X_{-K:\ell} = x_{-K:\ell}|X_j = r),$$

and similarly for $Y_{-K:\ell}$. The true emissions f_r^* are continuous functions on the compact set $[0, 1]$, hence they are uniformly continuous and hence are approximated in $L^\infty([0, 1])$ by the histograms $f_{\omega_r^*}$ with $\omega_r^* = (\int_{I_m} f_r^*(y)dy : m \in [\kappa_M])$. Since there are a fixed number of such terms, we can choose M large enough that, for each j , the preceding display is bounded above by $\frac{\epsilon}{9JL}$. Putting everything together and using Hölder’s inequality, we get that

$$\|H_{r,M}(h_{r,M}) - H_r(h_{r,M})\|_{L^2(\mathbb{P}_*)} \leq \epsilon.$$

The general case follows by summing over r , which concludes the proof. □

C.3.2. *Deconvolution argument*

The next lemma essentially provides a reduction of the deconvolution argument.

Lemma 15. *Let h_M be a sequence of step functions on the partition \mathcal{I}_M with $\|h_M\| = \sum_r \|h_{r,M}\|_{L^2} \leq 1$, such that $h_{r,M} \in \mathcal{H}_{r,M}$. For any $k \geq 0$,*

$$H_M(h_M) \xrightarrow{L^2(\mathbb{P}_*)} 0 \implies \sum_{r=1}^R \sum_{k'=-k}^0 h_M(y_{k'})\mathbb{P}_*(X_{k'} = r|Y_{-k:0}) \xrightarrow{L^2(\mathbb{P}_*)} 0.$$

Proof. For notational simplicity, we will write, for any index j and σ -algebra \mathcal{G} ,

$$h_M(Y_j)\mathbb{P}_*(X_0|\mathcal{G}) = \sum_r h_{r,M}(Y_j)\mathbb{P}_*(X_0 = r|\mathcal{G}).$$

By Lemma 14, we have $H_M(h_M) \xrightarrow{L^2(\mathbb{P}_*)} 0 \iff H(h_M) \xrightarrow{L^2(\mathbb{P}_*)} 0$. Denote by Y^k the vector $(Y_{-\infty:k})$. Note that $H(h_M)(Y_{-\infty:0}) = G_M(Y^0)$ has the same distribution as $G_M(Y^j) \forall j \leq 0$, by stationarity of the latent chain. Define

$$\begin{aligned} D_0^{(M)}(h^{(M)}, J) &= \sum_{j=J}^0 G_M(Y^j) = \sum_{j=J}^0 h_M(Y_j)\mathbb{P}_*(X_j|\mathcal{G}^0) \\ &\quad + \sum_{j=-\infty}^{J-1} h_M(Y_j)(\mathbb{P}_*(X_j|\mathcal{G}^0) - \mathbb{P}_*(X_j|\mathcal{G}^{J-1})) \\ &= o_{L^2}(1). \end{aligned}$$

Using Lemma 13 with $k = 0$ to eliminate the terms which condition on \mathcal{G}^0 , we obtain

$$\begin{aligned} \mathbb{E}_*(D_0^{(M)}(h^{(M)}, J)|Y_0) &= h_M(Y_0)\mathbb{P}_*(X_0|Y_0) - \sum_{j=-\infty}^{J-1} E(h_M(Y_j)\mathbb{P}_*(X_j|\mathcal{G}^{J-1})|Y_0) \\ &= o_{L^2}(1). \end{aligned}$$

Using Lemma 19, we obtain, for a particular $\rho < 1$ which does not depend on M , that

$$\begin{aligned} &\sum_{j=-\infty}^{J-1} \mathbb{E}_*(h_M(y_j)\mathbb{P}_*(X_j|\mathcal{G}^{J-1})|Y_0) \\ &= \underbrace{\sum_{j=-\infty}^{J-1} \mathbb{E}_*(h_M(y_j)\mathbb{P}_*(X_j|\mathcal{G}^{J-1}))}_{=0} + O(\rho^J) = O(\rho^J), \end{aligned}$$

with the right-hand sum vanishing because the $\mathbb{E}_*(h_M(y_j)\mathbb{P}_*(X_j|\mathcal{G}^{J-1})) = 0$, by Lemma 13. The term $O(\rho^J)$ is uniform in M since $\|h_M\|_{L^2(\mathbb{P}_*)} \leq 1$, and by choosing J large enough this term can be made arbitrarily small. We conclude that

$$h_M(y_0)\mathbb{P}_*(X_0|Y_0) = o_{L^2}(1),$$

as $M \rightarrow \infty$. By instead applying Lemma 13 with $k < 0$, we can argue similarly that, for all $k \geq 0$,

$$D_k^{(M)}(h_M) := \sum_{r=1}^R \sum_{k'=-k}^0 h_M(y_{k'})\mathbb{P}_*(X_{k'} = r|Y_{-k:0}) = o_{L^2}(1).$$

□

C.3.3. On the convergence on $\tilde{\mathcal{A}}_M$

The proof of the following result is very similar to the proof of Lemma 1 in [27], which used that the spaces of score functions for mixture models were nested to show convergence.

Lemma 16. *Let $\tilde{\mathcal{A}}_M \subset \tilde{\mathcal{P}}_M$, \mathcal{A} , \mathcal{P} be as in the statement of Lemma 1. Then $\mathcal{A} = \tilde{\mathcal{A}} := \lim_{M \rightarrow \infty} \tilde{\mathcal{A}}_M$.*

Proof. We have by the arguments of [27] that for any $S \in L^2(\mathbb{P}_*)$ $(\tilde{\mathcal{A}}_M S)_M$ is Cauchy, as the spaces $\tilde{\mathcal{P}}_M$ are nested. By completeness of $L^2(\mathbb{P}_*)$, we then establish convergence to some element $\tilde{\mathcal{A}}S \in L^2(\mathbb{P}_*)$. Since $\tilde{\mathcal{A}}_M$ are projections, $\tilde{\mathcal{A}}$ is also a projection onto its image, which is some subspace

$$\tilde{\mathcal{P}} \subset \text{Cl} \left(\bigcup_{M \in \mathbb{N}} \tilde{\mathcal{P}}_M \right).$$

Recall that \mathcal{P} is defined as the closure of the linear span of the $H(h)$ functions, as in the statement of Lemma 1. By definition, \mathcal{P} is closed and, since $\tilde{\mathcal{P}}_M \subset \mathcal{P}$ for all M , we must have $\text{Cl}(\bigcup_{M \in \mathbb{N}} \tilde{\mathcal{P}}_M) \subset \mathcal{P}$. We show that any element in $H(h) \in \mathcal{P}$ can be written as the L^2 limit of some sequence in $\bigcup_{M \in \mathbb{N}} \tilde{\mathcal{P}}_M$. First, choose a sequence $h_M \rightarrow h$ in $L^2(\mathbb{P}_*)$ such that $\|h_M\|_{L^2} \leq 2\|h\|_{L^2}$ for all M , which is possible under **Assumption 5**. We now wish to show that, for all $r \in [R]$,

$$\begin{aligned} & (h_{r,M}(y_0) - h_r(y_0))\mathbb{P}_*(X_0 = r|\mathcal{G}^0) \\ & + \sum_{j=-\infty}^{-1} (h_r(y_j) - h_{r,M}(y_j))(\mathbb{P}_*(X_j = r|\mathcal{G}^0) - \mathbb{P}_*(X_j = r|\mathcal{G}^{-1})) \end{aligned} \quad (38)$$

vanishes in L^2 . By Lemma 19, we can choose J sufficiently large negative that

$$\sum_{j=-\infty}^J (h_r(y_j) - h_{r,M}(y_j))(\mathbb{P}_*(X_j = r|\mathcal{G}^0) - \mathbb{P}_*(X_j = r|\mathcal{G}^{-1})) < \frac{\epsilon}{2}.$$

Since $\|h_{r,M}\|_{L^2} \leq 2\|h_r\|_{L^2}$, J can be chosen so that the above holds uniformly in M .

To control the finitely many remaining terms, it suffices to note the L^∞ boundedness of the probabilities and the L^2 convergence of h_M to h , and so, for M sufficiently large, the remaining finite sum is bounded by $\frac{\epsilon}{2}$, hence (38) has L^2 norm at most ϵ . Since ϵ was arbitrary, we conclude that $\mathcal{P} \subset \tilde{\mathcal{P}}$, and so the two spaces coincide and $\mathcal{A} = \tilde{\mathcal{A}}$ as claimed. \square

C.4. Admissible partitions

The following lemma is well-known in the identifiability literature. We recall it here for completeness.

Lemma 17. *Let F_1^*, \dots, F_R^* be linearly independent distribution functions on $[0, 1]$. Then there exists a partition I_1, \dots, I_R of $[0, 1]$ which is admissible in the sense of Definition 1.*

C.5. Forgetting of the hidden chain

The following results are from [19]. Recall that $\bar{q} = \min_{ij} Q_{ij}^* > 0$ under Assumption 3. The first result quantifies the exponential forgetting of the hidden chain.

Lemma 18 (Douc et al. 2004, [19]). *For all $s \in [R]$, $k, k' \in \mathbb{Z}$, and $j \geq k$, we have*

$$\sup_{Y_{k:k'}} \|\mathbb{P}_*(X_j = \cdot | X_k = s, Y_{k:k'}) - \mathbb{P}_*(X_j = \cdot | Y_{k:k'})\|_{TV} \leq \rho^{j-k},$$

where $\rho = \frac{1-2\bar{q}}{1-\bar{q}}$.

The second of these results follows from applying Equation (20) of [19] to both the known-emission model and the submodels on the emissions. It in particular ensures that the score functions we consider are well-defined.

Lemma 19 (Douc et al. 2004, [19]). *For $h \in L^2(\mathbb{P}_*)$, $j \leq 0$, $r, s \in [R]$, and for \mathbb{P}_* -almost all $Y_{-\infty:0}$,*

$$|\mathbb{P}_*(X_{j-1} = s, X_j = r | Y_{-\infty:0}) - \mathbb{P}_*(X_{j-1} = s, X_j = r | Y_{-\infty:-1})| \leq \rho^{-j-1}$$

and

$$|h(Y_j)\{\mathbb{P}_*(X_j = r | Y_{-\infty:0}) - \mathbb{P}_*(X_j = r | Y_{-\infty:-1})\}| \leq |h(Y_j)|\rho^{-j-1},$$

where $\rho = \frac{1-2\bar{q}}{1-\bar{q}}$. In particular, by integrating the square of both sides, we have

$$\|h(Y_j)\{\mathbb{P}_*(X_j = r | Y_{-\infty:0}) - \mathbb{P}_*(X_j = r | Y_{-\infty:-1})\}\|_{L^2(\mathbb{P}_*)} \leq \|h\|_{L^2(\mathbb{P}_*)}\rho^{-j-1}.$$

Lemmas 18 and 19 remain true when replacing Y by $Y^{(M)}$, the coarsened data. The choice of ρ is made independently of M , as it only depends on the transition matrix Q^* , which is the same in all histogram models as it is in the semiparametric model.

C.6. Technical results for the proofs of Theorems 4 and 6

We employ the following lemma (Lemma 7 of the main text) in the proof of Theorem 4 to eliminate certain terms.

Lemma 20. *Consider a sequence of emissions $\mathbf{f}^{(n)}$ such that $g_{Q^*, \mathbf{f}^{(n)}}^{(3)} \rightarrow g_{Q^*, \mathbf{f}^*}^{(3)}$ in $L^1(dx^3)$. Then $\mathbf{f}^{(n)} \xrightarrow{L^1} \mathbf{f}^*$ along a subsequence.*

Proof. Write $P = \text{diag}(p^*)$. Since $\|g_{Q^*, \mathbf{f}^{(n)}}^{(3)} - g_{Q^*, \mathbf{f}^*}^{(3)}\|_{L^1} = o(1)$, $\|g_{Q^*, \mathbf{f}^{(n)}}^{(2)} - g_{Q^*, \mathbf{f}^*}^{(2)}\|_{L^1} = o(1)$ and

$$\sum_{r,s} (PQ^*)_{rs} (f_r^{(n)}(y_1) f_s^n(y_2) - f_r^*(y_1) f_s^*(y_2)) = o_{L^1}(1).$$

Construct a partition I_1, \dots, I_R of \mathbb{R} such that $F_{ir}^* = \int_{I_i} f_r^* dy$ is of rank R and define $F_{ir}^{(n)} = \int_{I_i} f_r^{(n)}(y) dy$. Then we get, for all i ,

$$\sum_{r,s} (PQ^*)_{rs} (F_{ri}^{(n)} f_s^{(n)}(y) - F_{ri}^* f_s^*(y)) = o_{L^1}(1).$$

Since $g_{Q^*, \mathbf{f}^{(n)}}^{(3)} \rightarrow g_{Q^*, \mathbf{f}^*}^{(3)}$, Theorem 2.3 of [58] establishes weak convergence of the emissions, which then implies $\mathbf{F}^n \rightarrow \mathbf{F}^*$ and so

$$(\mathbf{F}^*)^T (PQ^*) \mathbf{f}^{(n)} = (\mathbf{F}^*)^T (PQ^*) \mathbf{f}^* + o_{L^1}(1),$$

which implies $\mathbf{f}^{(n)} = \mathbf{f}^* + o_{L^1}(1)$. \square

In the proof of Theorem 6, we use the fact that the approximating distribution for the posterior $\Pi_1(\cdot | Y_{1:n})$ can be chosen to not depend on the observation Y_ℓ . The difficulty in the proof compared to the iid setting arises from the dependence structure, but, by the exponential forgetting (Lemma 18), we can prove that the $1/\sqrt{n}$ scaled joint score functions coincide up to $o(1)$ terms, and hence so do the \sqrt{n} -scaled MLEs.

Lemma 21. *Grant Assumptions 1–5, let $\hat{\theta} = (\hat{Q}, \hat{\omega})$ denote the MLE in the model with admissible partition \mathcal{I}_M and define $\hat{\theta}^{-\ell} = (\hat{Q}^{-\ell}, \hat{\omega}^{-\ell})$ by*

$$\hat{\theta}^{-\ell} = \theta^* + \frac{J_M^{-1}}{\sqrt{n}} l_\theta^{(-\ell)}(Y_{1:n}^{-\ell}), \quad (39)$$

where $Y_{1:n}^{-\ell} = (Y_1, \dots, Y_{\ell-1}, Y_{\ell+1}, \dots, Y_n)$ and

$$l_\theta^{-\ell}(Y_{1:n}^{(-\ell)}) = \sum_{i=1}^{\ell-1} l_\theta(Y_i | Y_{1:i-1}) + \sum_{i=\ell+1}^n l_\theta(Y_i | Y_{1:i-1}^{(-\ell)}). \quad (40)$$

Let $\mathcal{L}_n \subset [n]$ be an index set whose cardinality satisfies $|\mathcal{L}_n| \log |\mathcal{L}_n| = o(n)$. Then, for all $\epsilon > 0$,

$$\mathbb{P}_*(\exists l \in \mathcal{L}_n : \sqrt{n} |\hat{Q}^{(-l)} - Q^*| > \epsilon) \rightarrow 0,$$

as $n \rightarrow \infty$

Proof. We may write the log-likelihood for the parameter $\theta = (Q, \omega)$ as

$$l_\theta(Y_1, \dots, Y_n) = \sum_{i=1}^n l_\theta(Y_i | Y_{1:i-1}), \quad (41)$$

where the $i = 1$ term is the unconditional likelihood. Note that (40) corresponds to the log-likelihood based on observing $Y_{1:n}^{(-\ell)}$. Standard arguments for showing asymptotic normality (see e.g. [9]), using Taylor expansions of the derivative of the expression (41), give

$$\sqrt{n}(\hat{\theta} - \theta^*) = \frac{J_M^{-1}}{\sqrt{n}} \nabla_{\theta^*} l_{\theta^*}(Y_{1:n}) + R_n,$$

for $R_n = O_{\mathbb{P}_*}(\sup_{\hat{\theta} \in [\theta^*, \hat{\theta}]} \|\frac{1}{n} D_{\hat{\theta}}^2 l_{\hat{\theta}}(Y_{1:n}) + J_M\|)$, the error in approximating the Fisher Information with the negative log-likelihood. The uniform law of large numbers for the observed information as given in Theorem 3 of [19], together with the consistency of the MLE (see e.g. [39]), then implies that $R_n = o_{\mathbb{P}_*}(1)$.

For the estimator $\hat{\theta}^{(-\ell)}$, we have, for each $l \in \mathcal{L}$, by the definition in (39), that

$$\sqrt{n}(\hat{\theta}^{(-\ell)} - \theta^*) = \frac{J_M^{-1}}{\sqrt{n}} \nabla_{\theta^*} l_{\theta^*}^{-\ell}(Y_{1:n}^{-\ell}).$$

Hence

$$\sqrt{n}(\hat{\theta}^{(-\ell)} - \hat{\theta}) = \frac{J_M^{-1}}{\sqrt{n}} (\nabla_{\theta^*} l_{\theta^*}^{-\ell}(Y_{1:n}^{-\ell}) - \nabla_{\theta^*} l_{\theta^*}(Y_{1:n})) - R_n$$

Note that the R_n term does not depend on ℓ , and vanishes in \mathbb{P}_* -probability. We will show that, for each $\ell \in \mathcal{L}_n$, $\mathbb{P}_* \left[\frac{1}{\sqrt{n}} (\nabla_{\theta^*} l_{\theta^*}^{-\ell}(Y_{1:n}^{-\ell}) - \nabla_{\theta^*} l_{\theta^*}(Y_{1:n})) > \epsilon \right] = o(|\mathcal{L}_n|^{-1})$, so that a union bound implies

$$\mathbb{P}_*(\exists \ell \in \mathcal{L}_n : \frac{1}{\sqrt{n}} (\nabla_{\theta^*} l_{\theta^*}^{-\ell}(Y_{1:n}^{-\ell}) - \nabla_{\theta^*} l_{\theta^*}(Y_{1:n})) > \epsilon) \leq |\mathcal{L}_n| o_{\mathbb{P}_*}(|\mathcal{L}_n|^{-1}) = o_{\mathbb{P}_*}(1),$$

from which the result follows. The argument is based on the exponential forgetting properties, and the expansion of the score considered in [19]. For the full log-likelihood (41), we have through the Fisher identity as in [19] that the log-likelihood of θ given $Y_{1:n}$ and $X_0 = x$, denoted $l_n(\theta, X_0 = x)$, has gradient at $\theta = \theta^*$ given by

$$\frac{1}{\sqrt{n}} \nabla_{\theta^*} l_n(\theta, X_0 = x) = \frac{1}{\sqrt{n}} \sum_{k=1}^n \Delta_{k,0,x},$$

with

$$\begin{aligned} \Delta_{k,0,x} = & \mathbb{E}_{\theta^*} \left[\sum_{i=1}^k \nabla_{\theta^*} \log(Q_{X_{i-1}, X_i} \omega_{X_i, Y_i}) | Y_{0:k}, X_0 = x \right] \\ & - \mathbb{E}_{\theta^*} \left[\sum_{i=1}^{k-1} \nabla_{\theta^*} \log(Q_{X_{i-1}, X_i} \omega_{X_i, Y_i}) | Y_{0:k-1}, X_0 = x \right], \end{aligned}$$

which can be seen by writing the log-likelihood as a telescoping sum and using the Fisher identity to write each term as a conditional expectation of the full

likelihood. We can do the same for the likelihood without Y^ℓ : define

$$\begin{aligned} \Delta_{k,0,x}^{(-\ell)} &= \mathbb{E}_{\theta^*} \left[\sum_{i=1}^k \nabla_{\theta^*} \log(Q_{X_{i-1}, X_i} \omega_{X_i, Y_i}) | Y_{0:k}^{-\ell}, X_0 = x \right] \\ &\quad - \mathbb{E}_{\theta^*} \left[\sum_{i=1}^{k-1} \nabla_{\theta^*} \log(Q_{X_{i-1}, X_i} \omega_{X_i, Y_i}) | Y_{0:k-1}^{(-\ell)}, X_0 = x \right]. \end{aligned}$$

We remark that the term $\Delta_{l,0,x}^{(-\ell)} = 0$ which corresponds to the missing contribution from Y^l in the expansion of the likelihood. Then we can write

$$\frac{1}{\sqrt{n}} \nabla_{\theta^*} l_n^{(-\ell)}(\theta, X_0 = x) = \frac{1}{\sqrt{n}} \sum_{k < \ell-1} \Delta_{k,0,x}^{(-\ell)} + \frac{1}{\sqrt{n}} \sum_{k=\ell+1}^{\ell+1+L_n} \Delta_{k,0,x}^{(-\ell)} + \frac{1}{\sqrt{n}} \sum_{k > L_n} \Delta_{k,0,x}^{(-\ell)}.$$

The first sum remains unchanged compared to the expression without the missing data point. Lemma 8 of [19] shows that we may replace the $\Delta_{k,0,x}$ by $\Delta_{k,0}$ where the latter is defined conditioning only on the $Y_{0:k}$ – the same argument made there shows the analogous result for the $\Delta_{k,0,x}^{(-\ell)}$ with respect to the analogously defined $\Delta_{k,0}^{(-\ell)}$. We note that, under **Assumptions 3** and **4**, the $\Delta_{k,0}$ and $\Delta_{k,0}^{-\ell}$ are bounded uniformly in $X_{(\cdot)}, Y_{(\cdot)}$, and so, for any $L_n = o(\sqrt{n})$ to be chosen later, we have that for some constant $C > 0$ depending on Q^* and w^* ,

$$\mathbb{P}^* \left[\frac{1}{\sqrt{n}} \sum_{k=\ell+1}^{\ell+1+L_n} \Delta_{k,0}^{(-\ell)}(\theta) - \frac{1}{\sqrt{n}} \sum_{k=\ell}^{\ell+1+L_n} \Delta_{k,0}(\theta) > CL_n/\sqrt{n} \right] = 0$$

It remains to control the difference between the remaining contributions for suitably chosen L_n . Similarly to what is done in [19], we define $\Delta_{k,m}$ by conditioning instead on $Y_{m:k}$. Write $\Delta_{k,m}$ as a telescoping sum:

$$\begin{aligned} \Delta_{k,m} &= \mathbb{E}_{\theta^*} [\nabla_{\theta^*} \log(Q_{X_{k-1}, X_k} \omega_{X_k, Y_k}) | Y_{m:k}] \\ &\quad + \sum_{i=m+1}^{k-1} \mathbb{E}_{\theta^*} [\nabla_{\theta^*} \log(Q_{X_{i-1}, X_i} \omega_{X_i, Y_i}) | Y_{m:k}] \\ &\quad - \mathbb{E}_{\theta^*} [\nabla_{\theta^*} \log(Q_{X_{i-1}, X_i} \omega_{X_i, Y_i}) | Y_{m:k-1}]. \end{aligned}$$

Then their equation (19) states that, for $k > m > m'$,⁵

$$\begin{aligned} &\| \mathbb{E}_{\theta^*} [\nabla_{\theta^*} \log(Q_{X_{i-1}, X_i} \omega_{X_i, Y_i}) | Y_{m:k}] - \mathbb{E}_{\theta^*} [\nabla_{\theta^*} \log(Q_{X_{i-1}, X_i} \omega_{X_i, Y_i}) | Y_{m':k}] \| \\ &\leq 2C\rho_*^{i-m-1}, \end{aligned}$$

where $\rho^* \in (0, 1)$ is a constant which depends on θ^* alone. Furthermore, their equation (20) states that

$$\begin{aligned} &\| \mathbb{E}_{\theta^*} [\nabla_{\theta^*} \log(Q_{X_{i-1}, X_i} \omega_{X_i, Y_i}) | Y_{m:k}] - \mathbb{E}_{\theta^*} [\nabla_{\theta^*} \log(Q_{X_{i-1}, X_i} \omega_{X_i, Y_i}) | Y_{m:k-1}] \| \\ &\leq 2C\rho_*^{k-i-1}. \end{aligned}$$

⁵Their result is only stated for $m < 0$, but by stationarity extends to all $m \in \mathbb{Z}$.

Combining the above two displays in the same way as Lemma 10 of [19], we have, for a suitable constant $C(Q^*, \omega^*)$ that

$$\|\Delta_{k,0} - \Delta_{k,\ell+1}\|_{L^2(\mathbb{P}_*)} \leq C(Q^*, \omega^*) \frac{\rho_*^{(k-\ell-1)/2-1}}{1 - \rho_*}.$$

Arguing similarly for the case of $\Delta_{k,0}^{-\ell}$, we see that

$$\|\Delta_{k,0}^{-\ell} - \Delta_{k,\ell+1}^{-\ell}\|_{L^2(\mathbb{P}_*)} \leq C(Q^*, \omega^*) \frac{\rho_*^{(k-\ell-1)/2-1}}{1 - \rho_*}.$$

The $\Delta_{k,\ell+1}^{-\ell} = \Delta_{k,\ell+1}$ as the conditioning excludes Y_ℓ in either case, and so the above displays combine to give

$$\|\Delta_{k,0}^{-\ell} - \Delta_{k,0}\|_{L^2(\mathbb{P}_*)} \leq 2C(Q^*, \omega^*) \frac{\rho_*^{(k-\ell-1)/2-1}}{1 - \rho_*}.$$

Applying Minkowski's inequality then gives that

$$\left\| \sum_{k>\ell+1+L_n} (\Delta_{k,0}^{-\ell} - \Delta_{k,0}) \right\|_{L^2(\mathbb{P}_*)} = O(\rho_*^{L_n/2}).$$

Choosing $L_n = 2M \log |\mathcal{L}_n|$ for sufficiently large $M > 0$ we have $\rho_*^{L_n/2} = e^{(\log |\mathcal{L}_n|)(M \log \rho_*)} = |\mathcal{L}_n|^{M \log \rho_*} = o(\frac{1}{n})$. Therefore we finally obtain that for all $l = o(n/\log n)$

$$\mathbb{P}_* \left(\frac{1}{\sqrt{n}} (\nabla_{\theta^*} l_{\theta^*}^{-\ell}(Y_{1:n}^{-\ell}) - \nabla_{\theta^*} l_{\theta^*}(Y_{1:n})) > \epsilon \right) \leq 1/n$$

and, since $|\mathcal{L}_n| = o(n)$, the claim is proved. \square

Appendix D: Assumptions required for the application of Theorem 5 to Dirichlet process mixtures

In this section, we detail the assumptions required for Proposition 1.

We will need the following assumptions on the behaviour of the emissions, which are assumptions (T1)–(T3) of [57].

Assumption 8. (i) There exist positive constants M_0, τ_0, γ_0 such that, for all $1 \leq i \leq R$ and all $y \in \mathbb{R}$,

$$f_i^*(y) \leq M_0 \exp(-\tau_0 |y|^{\gamma_0}).$$

(ii) For all $1 \leq i, j \leq R$ there exist constants $T_{i,j}, M_{i,j}, \tau_{i,j}, \gamma_{i,j} < \gamma_0$ such that

$$f_i^*(y) \leq f_j^*(y) M_{i,j} \exp(\tau_{i,j} |y|^{\gamma_{i,j}}), \quad |y| \geq T_{i,j}.$$

(iii) For all $1 \leq i \leq R$, f_i^* is positive and there exists $c_i > 0$ and $y_i^{low} < y_i^{high}$ such that f_i^* is non-decreasing on $(-\infty, y_i^{low})$, bounded below by $c > 0$ on (y_i^{low}, y_i^{high}) , and non-increasing on (y_i^{high}, ∞) .

We will further require the following assumptions on the choice of prior, which are Assumptions (G1) and (S1)–(S3) of [57]. These assumptions are verified by standard choices of Gaussian base measure and inverse gamma prior on the standard deviation.

Assumption 9. (i) $\alpha(\mathbb{R} \setminus [-y, y]) \lesssim \exp(-C_1 y^{a_1})$ for y large enough and some $C_1, a_1 > 0$

(ii) **B.** $\sigma \sim \Pi_\sigma$ where Π_σ is an Inverse Gamma distribution.

Condition (ii) of Assumption 9 on Π_σ can be relaxed; we essentially Π_σ to admit a Lebesgue density π_σ which behaves like an inverse – Gamma near 0 and has tails bounded by some power of $1/\sigma$ near infinity (see for instance [37]).

Appendix E: Details on MCMC algorithms used in Section 5

E.1. MCMC Algorithm for Π_1

In this section we describe the algorithm (Algorithm SA1) used for the simulation study in Section 5. S_R denotes the symmetric group of order R and where we use subscripts to refer to algorithm iterations. In our R implementation, used the RHmm package [53] for the forward-backward algorithm in the simulation of the hidden states and computation of the log likelihood, and the gtools package [59] for generation of permutations and Dirichlet draws. In our python implementation, we used the JAX package [10] to write both the forward-backward algorithm and the MCMC⁶.

Algorithm SA1: Algorithm for MCMC draws targeting posterior of Q .

Input : Binned data $Y^{(M)} \in [\kappa_M]^n$, number of hidden states R , prior parameters γ, β , iterations I , initial hidden states value \mathbf{X}_{init}

Output: List of I draws of Q

Initialise $\mathbf{X}_1 = \mathbf{X}_{init}$

for $i = 1, \dots, I + b$ **do**

 | Draw transition matrix $Q_i \sim P(Q|\mathbf{X}_i)$, histogram weights $\underline{\omega}_i \sim P(\underline{\omega}|\mathbf{X}_i, Y)$,
 | hidden states $\mathbf{X}_{i+1} \sim P(\mathbf{X}|Y, Q_i, \underline{\omega}_i)$

end

 Compute $i_{MAP} = \arg \max_i \log \Pi_1(Q_i, \underline{\omega}_i|Y)$

for $i = 1, \dots, I$ **do**

 | $\tau_i = \arg \min_{\tau \in S_R} d\left(\binom{(\tau)}{Q_i}, \binom{(\tau)}{\underline{\omega}_i}, (Q_{i_{MAP}}, \underline{\omega}_{i_{MAP}})\right)$

end

return $\{\binom{(\tau_i)}{Q_i}\}_{i=1}^I$

E.2. MCMC Algorithm for Π_2

In Algorithm SA2, we detail the MCMC procedure we implement in R and python/JAX. Once again, we make use of the forward-backward algorithm to

⁶All code is documented at <https://github.com/dm1729/>.

simulate latent states, allowing us to exploit the simple structure of the full likelihood. One key difference is that our algorithm requires allocation to a bivariate latent space (with total number of states RS_{\max}) involving both the HMM hidden state and the Dirichlet mixture component. Since the forward backward algorithm is $O(NK^2)$ with K the number of states and N the number of samples, and since we take $S_{\max} = \lfloor \sqrt{N} \rfloor$, the overall implementation is $O(N^2)$. The GPU capability of JAX code is of particular help here, as the computations for the forward-backward algorithm with a large number of states involve large matrices.

E.3. MCMC Algorithm for Π'

The MCMC algorithm of Π' is obtained by modifying Algorithm SA2 to remove the interior loop over $c = 1, \dots, C$ and sampling a transition matrix Q from the conditional distribution given the latent states, as in Algorithm SA1, at the beginning of each outer loop over i . As with Algorithm SA2, the complexity is $O(N^2)$ and so using it to target the marginal posterior on Q is much slower than using Π_1 , as discussed in Section 5.

Algorithm SA2: Algorithm for MCMC draws targeting cut posterior of emissions. Within the loop over c , we adopt the convention $W_i^0 = W_{i-1}$ and $W_{i-1}^0 = W_0$ etc.

Input : Data $Y \in \mathbb{R}^N$, Thinned transition matrix list $(Q_i)_{i=1, \dots, \bar{I}}$, number of hidden states R , prior parameters $M_0, \alpha_\sigma, \beta_\sigma, \mu_c, \sigma_c^2$, interior iterations C , initial hidden state value \mathbf{X}_{init} , truncation level S_{max} .

Output: List of \bar{I} draws of $\{f_r\}_{r=1}^R$

Initialise $\mathbf{X}_0 = \mathbf{X}_{init}$; variances $v^{(r)} \stackrel{iid}{\sim} \text{InvGamma}(\alpha_\sigma, \beta_\sigma)$; weights $W^{(r)} = (W_1^{(r)}, \dots, W_{S_{max}}^{(r)}) \sim \text{Dir}(\alpha)$ with $\alpha = \alpha(M_0, S_{max})$ as in (23); allocation variables $S_0 = (S_{01}, \dots, S_{0n})$ as in (22)

for $i = 1, \dots, \bar{I}$ **do**

for $c = 1, \dots, C$ **do**

Sample locations $(\mu_{i1}^{c,(r)}, \dots, \mu_{iS_{max}}^{c,(r)})_{r=1, \dots, R}$ from $P(\mu_i^c | S_i^{c-1}, \mathbf{X}_i^{c-1}, Y)$

Sample variances $(v_i^{c,(r)})_{r=1, \dots, R}$ from $P(v_i^c | S_i^{c-1}, \mathbf{X}_i^{c-1}, \mu_i^c, Y)$

Sample weights $(W_{i,1}^{c,(r)}, \dots, W_{i,S_{max}}^{c,(r)})_{r=1, \dots, R}$ from $P(W_i^c | S_i^{c-1}, \mathbf{X}_i^{c-1})$

Sample latents (S_i^c, \mathbf{X}_i^c) from $P(S_i^c, \mathbf{X}_i^c | \mu_i^c, v_i^c, W_i^c, Y, Q_i)$

end

Store $(\mu_i, v_i, W_i, S_i) = (\mu_i^C, v_i^C, W_i^C, S_i^C)$

end

for $i = 1, \dots, \bar{I}$ **do**

Compute $f_i^{(r)}$ from (μ_i, v_i, W_i, S_i) by evaluating (an approximation of) (21)

end

return $\{f_i^{(r)}\}_{i=1}^{\bar{I}}$

Acknowledgments

The project leading to this work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 834175). The project is also partially funded by the EPSRC via the CDT StatML.

References

- [1] Kweku Abraham, Ismaël Castillo, and Elisabeth Gassiat. Multiple testing in nonparametric hidden markov models: An empirical bayes approach. *Journal of Machine Learning Research*, 23(94):1–57, 2022. [MR4576679](#)
- [2] Kweku Abraham, Elisabeth Gassiat, and Zacharie Naulet. Fundamental limits for learning hidden markov model parameters. *IEEE Transactions on Information Theory*, 2022. [MR4564681](#)
- [3] Pierre Ailliot, Craig Thompson, and Peter Thomson. Space–time modelling of precipitation by using a hidden markov model and censored gaussian distributions. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 58(3):405–426, 2009. [MR2750013](#)
- [4] Grigory Alexandrovich, Hajo Holzmann, and Anna Leister. Nonparametric identification and maximum likelihood estimation for hidden markov models. *Biometrika*, 103(2):423–434, 2016. [MR3509896](#)
- [5] Elizabeth S. Allman, Catherine Matias, and John A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *Annals of Statistics*, 37(6 A):3099–3132, 2009. [MR2549554](#)
- [6] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of machine learning research*, 15:2773–2832, 2014. [MR3270750](#)
- [7] Animashree Anandkumar, Daniel Hsu, and Sham M Kakade. A method of moments for mixture models and hidden markov models. In *Conference on Learning Theory*, pages 33–1. JMLR Workshop and Conference Proceedings, 2012.
- [8] James Bennett and Jon Wakefield. Errors-in-variables in joint population pharmacokinetic/pharmacodynamic modeling. *Biometrics*, 57(3):803–812, 2001. [MR1863449](#)
- [9] Peter J Bickel, Ya’acov Ritov, and Tobias Ryden. Asymptotic normality of the maximum-likelihood estimator for general hidden markov models. *The Annals of Statistics*, 26(4):1614–1635, 1998. [MR1647705](#)
- [10] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018.
- [11] Olivier Cappé, Eric Moulines, and Tobias Rydén. Inference in hidden markov models. In *Proceedings of EUSFLAT conference*, pages 14–16, 2009. [MR2159833](#)

- [12] Christian Carmona and Geoff Nicholls. Semi-modular inference: enhanced learning in multi-modular models by tempering the influence of components. In *International Conference on Artificial Intelligence and Statistics*, pages 4226–4235. PMLR, 2020.
- [13] Ismaël Castillo. Semiparametric bernstein–von mises theorem and bias, illustrated with gaussian process priors. *Sankhya A*, 74(2):194–221, 2012. [MR3021557](#)
- [14] Ismaël Castillo and Judith Rousseau. A bernstein–von mises theorem for smooth functionals in semiparametric models. *The Annals of Statistics*, 43(6):2353–2383, 2015. [MR3405597](#)
- [15] Johann De Castro, Élisabeth Gassiat, and Claire Lacour. Minimax adaptive estimation of nonparametric hidden markov models. *The Journal of Machine Learning Research*, 17(1):3842–3884, 2016. [MR3543517](#)
- [16] Johann De Castro, Elisabeth Gassiat, and Sylvain Le Corff. Consistent estimation of the filtering and marginal smoothing distributions in nonparametric hidden markov models. *IEEE Transactions on Information Theory*, 63(8):4758–4777, 2017. [MR3683535](#)
- [17] M. C. M. de Gunst and O Shcherbakova. Asymptotic behavior of Bayes estimators for hidden Markov models with application to ion channels. *Mathematical Methods of Statistics*, 17(4):342–356, 2008. [MR2483462](#)
- [18] Stacy L. DeRuiter, Roland Langrock, Tomas Skirbutas, Jeremy A. Goldbogen, John Calambokidis, Ari S. Friedlaender, and Brandon L. Southall. A multivariate mixed hidden Markov model for blue whale behaviour and responses to sound exposure. *The Annals of Applied Statistics*, 11(1):362 – 392, 2017. [MR3634328](#)
- [19] Randal Douc, Éric Moulines, and Tobias Rydén. Asymptotic properties of the maximum likelihood estimator in autoregressive models with Markov regime. *Annals of Statistics*, 32(5):2254–2304, oct 2004. [MR2102510](#)
- [20] Randal Douc, Eric Moulines, and David Stoffer. *Nonlinear time series: Theory, methods and applications with R examples*. CRC press, 2014. [MR3289095](#)
- [21] Emily B Fox, Erik B Sudderth, Michael I Jordan, and Alan S Willsky. The sticky hdp-hmm: Bayesian nonparametric hidden markov models with persistent states. *Arxiv preprint*, 2007.
- [22] David Freedman. Wald lecture: On the bernstein-von mises theorem with infinite-dimensional parameters. *The Annals of Statistics*, 27(4):1119–1141, 1999. [MR1740119](#)
- [23] Sylvia Frühwirth-Schnatter. *Finite mixture and Markov switching models*. Springer Science & Business Media, 2006. [MR2265601](#)
- [24] Élisabeth Gassiat, Alice Cleynen, and Stephane Robin. Inference in finite state space non parametric hidden markov models and applications. *Statistics and Computing*, 26(1-2):61–71, 2016. [MR3439359](#)
- [25] Elisabeth Gassiat and Judith Rousseau. About the posterior distribution in hidden markov models with unknown number of states. *Bernoulli*, 20(4):2039–2075, 2014. [MR3263098](#)

- [26] Elisabeth Gassiat and Judith Rousseau. Nonparametric finite translation hidden Markov models and extensions. *Bernoulli*, 22(1):193–212, 2016. [MR3449780](#)
- [27] Elisabeth Gassiat, Judith Rousseau, and Elodie Vernet. Efficient semiparametric estimation and model selection for multidimensional mixtures. *Electronic Journal of Statistics*, 12(1):703–740, 2018. [MR3769193](#)
- [28] Élisabeth Gassiat and Gilles Stoltz. The van trees inequality in the spirit of hájek and le cam. *arXiv preprint arXiv:2402.06431*, 2024.
- [29] Subhashis Ghosal and Aad Van Der Vaart. Convergence rates of posterior distributions for noniid observations. *Annals of Statistics*, 35(1):192–223, 2007. [MR2332274](#)
- [30] Subhashis Ghosal and Aad Van der Vaart. *Fundamentals of nonparametric Bayesian inference*, volume 44. Cambridge University Press, 2017. [MR3587782](#)
- [31] Richard D Gill and Boris Y Levit. Applications of the van trees inequality: a bayesian cramér-rao bound. *Bernoulli*, 1(1-2):59–79, 1995. [MR1354456](#)
- [32] Illia Horenko and Christof Schütte. Likelihood-based estimation of multidimensional langevin models and its application to biomolecular dynamics. *Multiscale Modeling & Simulation*, 7(2):731–773, 2008. [MR2443010](#)
- [33] Daniel Hsu, Sham M Kakade, and Tong Zhang. A spectral algorithm for learning hidden markov models. *Journal of Computer and System Sciences*, 78(5):1460–1480, 2012. [MR2926144](#)
- [34] Qi Huang, Dwayne Cohen, Sandra Komarzynski, Xiao-Mei Li, Pasquale Innominato, Francis Lévi, and Bärbel Finkenstädt. Hidden markov models for monitoring circadian rhythmicity in telemetric activity data. *Journal of The Royal Society Interface*, 15(139):20170885, 2018.
- [35] James P Hughes, Peter Guttorp, and Stephen P Charles. A non-homogeneous hidden markov model for precipitation occurrence. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 48(1):15–30, 1999.
- [36] Pierre E Jacob, Lawrence M Murray, Chris C Holmes, and Christian P Robert. Better together? statistical learning in models made of modules. *arXiv preprint arXiv:1708.08719*, 2017.
- [37] Willem Kruijer, Judith Rousseau, and Aad Van Der Vaart. Adaptive bayesian density estimation with location-scale mixtures. *Electronic journal of statistics*, 4:1225–1257, 2010. [MR2735885](#)
- [38] Roland Langrock, Thomas Kneib, Alexander Sohn, and Stacy L DeRuiter. Nonparametric inference in hidden markov models using p-splines. *Biometrics*, 71(2):520–528, 2015. [MR3366256](#)
- [39] Brian G Leroux. Maximum-likelihood estimation for hidden markov models. *Stochastic processes and their applications*, 40(1):127–143, 1992. [MR1145463](#)
- [40] Yang Liu and Robert JB Goudie. Stochastic approximation cut algorithm for inference in modularized bayesian models. *Statistics and Computing*, 32(1):1–15, 2022. [MR4350200](#)
- [41] Thomas A Louis. Finding the Observed Information Matrix When Us-

- ing the EM Algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 44(2):226–233, 1982. [MR0676213](#)
- [42] David J Lunn, Andrew Thomas, Nicky Best, and David Spiegelhalter. Winbugs—a bayesian modelling framework: concepts, structure, and extensibility. *Statistics and computing*, 10(4):325–337, 2000.
- [43] Rogemar S Mamon and Robert James Elliott. *Hidden Markov models in finance*, volume 4. Springer, 2007. [MR2407726](#)
- [44] Jean-Michel Marin and Christian Robert. *Bayesian core: a practical approach to computational Bayesian statistics*. Springer Science & Business Media, 2007. [MR2289769](#)
- [45] Cathy Maugis-Rabusseau and Bertrand Michel. Adaptive density estimation for clustering with gaussian mixtures. *ESAIM: Probability and Statistics*, 17:698–724, 2013. [MR3126158](#)
- [46] Brad McNeney and Jon A Wellner. Application of convolution theorems in semiparametric models with non-iid data. *Journal of Statistical Planning and Inference*, 91(2):441–480, 2000. [MR1814795](#)
- [47] David RH Miller, Tim Leek, and Richard M Schwartz. A hidden markov model information retrieval system. In *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 214–221, 1999.
- [48] Daniel Moss and Judith Rousseau. *Bayesian Modelling of Dependent Data*. PhD thesis, University of Oxford, 2024. Chapter 2: Efficient Bayesian Estimation and use of cut posterior in semiparametric hidden Markov models.
- [49] Martyn Plummer. Cuts in bayesian graphical models. *Statistics and Computing*, 25(1):37–43, 2015. [MR3304902](#)
- [50] Lawrence R Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [51] Vincent Rivoirard and Judith Rousseau. Bernstein–von mises theorem for linear functionals of the density. *The Annals of Statistics*, 40(3):1489–1523, 2012. [MR3015033](#)
- [52] Henry Scheffé. A useful convergence theorem for probability distributions. *The Annals of Mathematical Statistics*, 18(3):434–438, 1947. [MR0021585](#)
- [53] Ollivier Taramasco and Sebastian Bauer. *RHmm: Hidden Markov Models Simulations and Estimations*, 2018. R package version 2.1.9/r443.
- [54] Yee Whye Teh, Michael I Jordan, Matthew J Beal, and David M Blei. Hierarchical dirichlet processes. *Journal of the american statistical association*, 101(476):1566–1581, 2006. [MR2279480](#)
- [55] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998. [MR1652247](#)
- [56] Aad Van Der Vaart. On differentiable functionals. *The Annals of Statistics*, pages 178–204, 1991. [MR1091845](#)
- [57] Elodie Vernet. Non parametric hidden markov models with finite state space: Posterior concentration rates. *arXiv preprint [arXiv:1511.08624](#)*, 2015. [MR3331855](#)
- [58] Elodie Vernet. Posterior consistency for nonparametric hidden markov

- models with finite state space. *Electronic Journal of Statistics*, 9(1):717–752, 2015. [MR3331855](#)
- [59] Gregory R. Warnes, Ben Bolker, and Thomas Lumley. *gtools: Various R Programming Tools*, 2020. R package version 3.8.2.
- [60] Christopher Yau, Omiros Papaspiliopoulos, Gareth O Roberts, and Christopher Holmes. Bayesian non-parametric hidden markov models with applications in genomics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(1):37–57, 2011. [MR2797735](#)