

CLUSTERING AND CLASSIFICATION RISKS IN NON-PARAMETRIC HIDDEN MARKOV AND I.I.D. MODELS

BY ÉLISABETH GASSIAT^{1,a}, IBRAHIM KADDOURI^{1,b} AND ZACHARIE NAULET^{1,c}

¹Université Paris-Saclay, CNRS, Laboratoire de mathématiques d'Orsay, 91405, Orsay, France,

^aelisabeth.gassiat@universite-paris-saclay.fr; ^bibrahim.kaddouri@universite-paris-saclay.fr;

^czacharie.naulet@universite-paris-saclay.fr

We provide an in-depth analysis of the Bayes risk of clustering and the Bayes risk of classification in the context of Hidden Markov and i.i.d. models. In both settings, we identify the situations where the two risks are comparable or not and those where the associated minimizers are related or not, as well as the key quantity measuring the difficulty of both tasks. Then, leveraging the nonparametric identifiability of HMMs, we control the excess risk of a plug-in clustering procedure. Simulations illustrate our findings.

1. Introduction. Clustering is the problem of organizing a collection of objects into groups, ensuring that elements within each group exhibit greater *similarity* to each other than to those in other groups. Clustering plays a crucial role in various domains such as machine learning, pattern recognition and image processing, helping to reveal the hidden structure of data without requiring prior labels of the observations. In contrast, classification, while closely related to clustering, seeks not only to group similar objects together but also to assign class labels to them.

Mixture models are a common framework in which the two problems can be formally defined. In these models, observations $\mathbf{Y} = (Y_1, Y_2, \dots)$ are independent conditional on unobserved random variables $\mathbf{X} = (X_1, X_2, \dots)$ taking values in $\mathbb{X} = \{1, \dots, J\}$ that represent the labels of the classes in which observations originated, with J being the total number of classes. In this work, we consider the settings where the latent variables \mathbf{X} are i.i.d. or form a Markov Chain. Since the i.i.d. setting is a strict subcase of the Markovian case, without loss of generality we consider the general model

$$(1) \quad \begin{aligned} Y_i | \mathbf{X} &\stackrel{\text{ind}}{\sim} F_{X_i} \quad i = 1, 2, \dots \\ \mathbf{X} &\sim \text{Markov}(\nu, Q) \end{aligned}$$

with parameter $\theta = (\nu, Q, (F_x)_{x \in \mathbb{X}})$ where ν is the initial distribution of the chain (in the next identified with probability vectors on \mathbb{X}), Q the transition matrix of the hidden chain and $(F_x)_{x \in \mathbb{X}}$ are the emission distributions (*aka.* populations). In particular, the case of i.i.d. latent variables is obtained by restricting the parameters of the previous model to transition matrices Q having identical lines equal to ν . The reason why we emphasize the i.i.d. subcase in the paper is not only because of its wide use, but also because some interesting phenomena emerge only in the absence of dependencies. More details on the model are given in Section 2.2.

On one hand classification is the task of identifying the hidden variables X_1, \dots, X_n based on the observed variables Y_1, \dots, Y_n , usually given the knowledge of some of the labels (supervised learning). The performance of a classifier h is naturally measured using the loss that counts the number of misclassified observations, with the associated risk of classification

MSC2020 subject classifications: Primary 62M05; secondary 62G99.

Keywords and phrases: Clustering, Nonparametric populations, Hidden Markov models, Mixture models.

denoted $\mathcal{R}_n^{\text{class}}(\theta, h)$ in the sequel. On the other hand, clustering is the task of recovering the partition $\{A_1, \dots, A_m\}$ of $\{1, \dots, n\}$ induced by the labels (*aka.* groups), namely $i, j \in A_k \iff X_i = X_j$. In contrast with classification, clustering can be done in an unsupervised setting. Measuring the loss incurred by a clusterer g requires defining a notion of similarity between two partitions, which can be achieved in many ways. Here we make the choice of using the misclassification error metric [31, 30] which is based on measuring the best overlap achievable by matching the elements of the two partitions. A formal definition of the loss is stated later in Section 2.3. We will denote the risk of clusterer g by $\mathcal{R}_n^{\text{clust}}(\theta, g)$.

The Bayes risk of classification or clustering is defined as the smallest risk that a procedure of classification or clustering can reach, that is $\inf_h \mathcal{R}_n^{\text{class}}(\theta, h)$ and $\inf_g \mathcal{R}_n^{\text{clust}}(\theta, g)$. The Bayes risk emerges as a crucial concept within the decision-theoretic framework. It establishes a benchmark for evaluating classification and clustering algorithms as it measures the best performance that can be attained by an oracle having access to the parameter θ . The minimizer of the risk of classification is called *the Bayes classifier* and the minimizer of the risk of clustering is called *the Bayes clusterer*. While the Bayes classifier has a closed formula and has been extensively studied (see [13] and references therein), the Bayes clusterer is less explored, more difficult to compute and in general has no explicit formula. To the best of our knowledge, the existence of a relationship between the two has never been studied. Some natural questions arise in this context:

1. Is there any clear link between the Bayes classifier and the Bayes clusterer? If so, under what condition?
2. When can the Bayes risk of clustering be comparable to the Bayes risk of classification?
3. What is the appropriate measure of separation quantifying the difficulty of the problems of classification and clustering?
4. Is it possible for a clusterer agnostic of the true parameter to have vanishing excess risk as the number of observation grows (*ie.* having clustering risk comparable to that of the Bayes clusterer)?

In this work, we address these questions in the context of the model (1) with a focus on both the i.i.d. and dependent subcases.

To answer Question 1, we note that a clusterer can be obtained from a classifier by forgetting the labels, *ie.* retaining only the partition induced by the groups. We show that surprisingly, the clustering using the Bayes classifier equals that of the Bayes clusterer if and only if observations are i.i.d. from a mixture of two components. In the remaining situations (more than two components or dependent HMMs), there is always a set of parameters for which the clusterer obtained from the Bayes classifier and the Bayes clusterer differ with non-zero probability. See Theorems 1, 3, 5 and 7.

The Question 2 is of interest because, thanks to the closed formula of the Bayes risk of classification, it can be easily analyzed and when the two risks are comparable, any bound on the risk of classification can be extrapolated to the risk of clustering. We show that the Bayes risk of clustering is comparable to the Bayes risk of classification in two situations: when there are two clusters or when the risk of classification does not decrease exponentially in n . These results are proved for i.i.d. and HMM observations. See Corollary 1 and Theorem 4 for the i.i.d. model, Theorem 6 and Theorem 8 for the HMM model.

Thanks to the relationship between the risks studied in Question 2, we are able to conduct a precise analysis of the Bayes risk of clustering relying on the simpler Bayes risk of classification and answer Question 3. Understanding the dependence of the Bayes risks with respect to the model parameters (mainly the population densities) is important because this will clearly identify the appropriate notion of separation measuring the difficulty of each problem. From a practical point of view, guaranteeing that the Bayes risk is of a certain magnitude will therefore translate into a simple condition on the separation between densities,

which is easily interpretable. In Theorem 9 and Corollary 4, we identify the key quantity driving the difficulty of clustering and classification which turns out to be

$$\Lambda := \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left(\sum_{x \neq x_0} f_x(y) \right) d\mathcal{L}(y)$$

in both HMMs and i.i.d. models. Here, $(f_x)_{x \in \mathbb{X}}$ are the population densities of the distributions $(F_x)_{x \in \mathbb{X}}$ with respect to a dominating measure \mathcal{L} over the observation space \mathbb{Y} , that is $dF_x = f_x d\mathcal{L}$. Notice that when $J = 2$, $\Lambda = 1 - \|F_1 - F_2\|_{\text{TV}}$, where $\|F_1 - F_2\|_{\text{TV}}$ is the total variation distance between the two distributions, unsurprisingly showing that the difficulty of the clustering tasks is governed by the difficulty of the hypothesis testing between F_0 and F_1 .

Finally, we turn our attention to Question 4. We focus on the situation where the emission densities $(f_x)_{x \in \mathbb{X}}$ are modeled non-parametrically. In many contexts, the nonparametric setting seems to be more realistic because, due to the complex nature of real-world data, populations may not always be modeled using parametric mixtures [10, 25] and sometimes, strong assumptions on the distributions should be relaxed. Unfortunately, mixture models suffer from non-identifiability in general non-parametric settings, leading sometimes to the failure of common clustering procedures that do not leverage the distribution of the data (k -means and its variants, hierarchical clustering, etc.). To address the issue of non-identifiability, various assumptions are often made about the mixture components. For instance, some approaches rely on assumptions about the dependence between observations [17, 15], symmetry of the mixture components [22], or separation between components [7]. A significant breakthrough in non-parametric mixtures has been achieved through the use of the Hidden Markov Model (HMM), which assumes a simple Markovian dependence between the hidden labels. Recent research has demonstrated the identifiability of non-parametric HMMs [16, 4]. Given that the transition matrix of the hidden Markov chain is non-singular and ergodic, all parameters – including the number J of hidden states and the emission distributions – can be identified from the distribution of K consecutive observations (Y_1, \dots, Y_K) , provided that the emission distributions are distinct and K is sufficiently large compared to J [4]. This allows complete flexibility on the emission distributions, provided a Markovian dependence of labels. Within the framework of this model, non-parametric estimation of the mixture components becomes possible without any restrictions on the population densities, apart from the fact that they are distinct, and various estimation procedures have been proposed [9, 12, 26, 27, 1, 2]. These estimates can then be used to develop clustering procedures. For example, the clustering procedure most commonly used in HMM involves two steps. First, the parameters of the HMM (initial distribution, transition matrix of the hidden Markov chain and population densities - usually called emission densities) are estimated, then, using these estimates, the celebrated forward-backward algorithm (cf. [8]) is run to compute the plug-in estimator of the a posteriori distribution of the hidden states given the observations. The clustering is then performed by soft assignment. See [24, 35, 18]. While widely used, this procedure does not have any theoretical justification. In fact, assuming the estimation of the parameters is perfect, this clustering would match the clustering using the Bayes classifier instead of the Bayes clusterer, which is not justified unless the Bayes clusterer can be derived from the Bayes classifier or at least, the corresponding clustering excess risk is small. This is why it is important to provide theoretical guarantees on the performance of this clustering procedure. In Theorem 10, we show that the excess risk of this plug-in procedure decays at the nonparametric estimation rate for a suitable estimator of the parameter, providing thus a theoretical justification for its use. Finally, numerical simulations demonstrate that HMM modeling can successfully cluster data where other standard methods fail to recover the hidden structure.

The assumptions of the model and the definitions of the various risks we shall use in this work are described in Section 2. We state our main results in Section 3, the proofs of which are detailed in Section 6. Section 4 is devoted to simulation experiments, and Section 5 to possible further work.

2. Setting and definitions.

2.1. Notations. We note for $i \leq j$ $X_{i:j} := (X_i, \dots, X_j)$ and $Y_{i:j} := (Y_i, \dots, Y_j)$. We denote $[n] := \{1, \dots, n\}$ and $\mathcal{P}[n]$ is the set of partitions of $[n]$. The set of permutations of $[J]$ will be denoted by \mathcal{S}_J , and for any $\tau \in \mathcal{S}_J$, we consider θ^τ , the model parameter when labels are permuted with τ , that is $\theta^\tau = (\nu^\tau, Q^\tau, (f_{\tau(x)})_{x \in \mathbb{X}})$ where $\nu^\tau = (\nu_{\tau(x)})_{x \in \mathbb{X}}$ and $Q^\tau = (Q_{\tau(x), \tau(x')})_{x, x' \in \mathbb{X}}$. For h a measurable function, $\|h\|_\infty$ denotes the essential supremum of h , possibly infinite. Frobenius norm is denoted by $\|\cdot\|_F$ and $\|\cdot\|$ stands for the operator norm. Given two sequences of positive numbers $(a_n)_n$ and $(b_n)_n$, $a_n \gtrsim b_n$ means that there exists an absolute constant $c > 0$ and $n_0 \in \mathbb{N}$ such that $(\forall n \geq n_0) \quad a_n \geq cb_n$. In the HMM modeling, for any parameter θ and any $\tau \in \mathcal{S}_J$, the distribution of the observations under \mathbb{P}_θ is the same as under \mathbb{P}_{θ^τ} . In other words, this distribution is invariant up to permutation of the labels given by the hidden states. This is known as the *label-switching* issue.

2.2. The model. We consider a Hidden Markov Model with J hidden states taking value in a set of labels $\mathbb{X} = \{1, \dots, J\}$, and observations in a Polish space endowed with its Borel σ -field $(\mathbb{Y}, \mathcal{Y})$. We denote $\mathbf{X} = (X_1, X_2, \dots)$ and $\mathbf{Y} = (Y_1, Y_2, \dots)$ respectively the sequence of hidden states forming the Markov chain and the observations. We assume that the emission distributions have densities f_x , $x = 1, \dots, J$, with respect to a dominating measure \mathcal{L} on $(\mathbb{Y}, \mathcal{Y})$. The HMM assumption boils down to:

$$\begin{aligned} \mathbf{X} &\sim \text{Markov}(\nu, Q) \\ \mathbb{P}_\theta(Y_i \in \cdot \mid \mathbf{X}) &= f_{X_i}(\cdot) d\mathcal{L} \end{aligned}$$

where ν is the initial distribution of the chain and Q the transition matrix of the hidden chain. Throughout the paper, it is assumed that only the beginning $Y_{1:n}$ of the sequence \mathbf{Y} is observed, and nothing else. We set $\theta = (\nu, Q, (f_x)_{x \in \mathbb{X}})$ the parameter of the model and Θ denotes the space of all valid parameters. The following assumption is made:

- *Independent case.* In this case we assume that all the lines of Q are identical and equal to the vector ν of weights forming its stationary distribution. This corresponds to the usual mixture model with independent latent variables. The set of these parameters will be denoted Θ^{ind} .
- *Dependent case.* In this case we assume the lines of the transition matrix Q not all equal, so that the $(X_i)_{i \geq 1}$ are not independent. The set of these parameters will be denoted Θ^{dep} .

Throughout this work, we shall consider the transition matrix being fixed, and we will be interested in how the separation of populations, understood as some quantity depending on the populations densities, drives the difficulty of the clustering task. We show that, under the HMM modelling, it is possible to cluster general populations without assuming they belong to some prescribed parametric family. While our primary interest is in the dependent case in which the emission densities can be identified without any further constraint, we obtain, however, results that have interest also in the widely used independent case, in particular the analysis of the Bayes risk and its minimizers. (see Section 3).

One main difference between the HMM and the i.i.d. situation in the analysis of the Bayes risk of classification is that in the HMM modeling, the probability of a label X_i given the observations Y_1, \dots, Y_n depends on all the observations, whereas in the i.i.d. case it depends

only on the associated observation Y_i . The vector of probabilities of the X_i 's given Y_1, \dots, Y_n are called the smoothing distributions, they depend on i, n and the observations. The vectors of probabilities of the X_i 's given the observations up to time i , Y_1, \dots, Y_i , are called the filtering distributions, they depend on i, n and the involved observations. Recursive formulas verified by the filtering and smoothing distributions make the computations tractable under the HMM modelling. See Section 3 of [8] for details.

2.3. The problem of clustering. For any $n \geq 1$, the finite sequence $X_{1:n} = (X_1, \dots, X_n)$ induces a random partition $\Pi_n = \{C_1, C_2, \dots\}$ of $[n]$ whose blocks – the so-called *clusters* – are the equivalence classes for the random equivalence relation $i \sim j \iff X_i = X_j$. The goal of clustering is to uncover this partition Π_n on the sole basis of the observation $Y_{1:n} = (Y_1, \dots, Y_n)$. We define a *clusterer*:

DEFINITION 1 (Clusterer). A n -clusterer is a measurable map $g : \mathbb{Y}^n \rightarrow \mathcal{P}[n]$. We denote by \mathcal{G}_n the set of all n -clusterers.

We measure the loss incurred by guessing $g(Y_{1:n})$ in place of Π_n via the misclassification error distance [31, 30]. For two partitions A and B of $[n]$, this loss is defined by:

$$(2) \quad \ell(A, B) = 1 - \frac{1}{n} \sup_{\substack{M \subseteq \mathcal{E}(A, B) \\ M \text{ is a matching}}} \sum_{\{C, C'\} \in M} \text{Card}(C \cap C')$$

where the supremum is taken over the set of matchings. To define a matching, we build the complete bipartite graph $(A, B, \mathcal{E}(A, B))$ on vertices A and B with edge set $\mathcal{E}(A, B) := \{\{C, C'\} : C \in A, C' \in B\}$. Then we recall that a matching M is a set $M \subseteq \mathcal{E}(A, B)$ of edges without common vertices (*i.e.* each block of A and B appears in at most one edge of the matching). An example of a matching is depicted in Figure 1. Then the risk of a clusterer g can be defined as the expected loss of the partition $g(Y_{1:n})$ with respect to the true partition Π_n which is summarized in the risk function $\mathcal{R}_n^{\text{clust}} : \Theta \times \mathcal{G}_n \rightarrow [0, 1]$

$$(3) \quad \mathcal{R}_n^{\text{clust}}(\theta, g) := \mathbb{E}_\theta \left[\ell(g(Y_{1:n}), \Pi_n) \right].$$

A closely related notion is that of a *classifier*:

DEFINITION 2 (Classifier). A n -classifier is a measurable map $h : \mathbb{Y}^n \rightarrow \mathbb{X}^n$. We denote by \mathcal{H}_n the set of all n -classifiers.

One may find our Definition 2 different from standard textbook definitions [13]: we explain the reason of this choice later in Remark 1. A classifier differs from a clusterer in that it not only seeks for the hidden partition, but also for the labels of the observations. Hence, the usage of a classifier only makes sense in a supervised framework where access to some labeled data is allowed in some way. In an unsupervised framework such as our model, \mathbf{Y} does not contain any information about the labels and recovering them better than a lucky guess is impossible. It is however true that, to any n -classifier $h \in \mathcal{H}_n$ corresponds a unique n -clusterer $g \in \mathcal{G}_n$ which can be built via the map $\pi_n : \mathbb{X}^n \rightarrow \mathcal{P}[n]$ such that

$$g(Y_{1:n}) = \pi_n \circ h(Y_{1:n}) = \{\{i : h_i(Y_{1:n}) = x\} : x \in \mathbb{X}\} \setminus \{\emptyset\}$$

and any clusterer can be represented that way by choosing a specific labelling of the clusters. For this reason, the notions of clusterer and classifier are very much often amalgamated in the literature. We argue that it would be better to define them separately in order to avoid




Figure 1: Example of a matching. Nodes on the left represent the clusters induced by the partition of Π_n ; those on the right are the clusters of $g(Y_{1:n})$. Edges form a matching between the two partitions.

confusions between the risk of clustering $\mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ h)$ and the risk of classification $\mathcal{R}_n^{\text{class}} : \Theta \times \mathcal{H}_n \rightarrow [0, 1]$ (relative to the loss counting number of misclassified observations)

$$(4) \quad \mathcal{R}_n^{\text{class}}(\theta, h) := \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq X_i} \right].$$

Here again we insist that the definition of $\mathcal{R}_n^{\text{class}}$ in Equation (4) – although mathematically correct – has no statistical interest in an unsupervised model where the classifier h is not allowed to see some of the labels. An easy exercise (see Lemma 14 in Section 6.13) shows that the risk of clustering of $\pi_n \circ h$ can be rewritten as

$$(5) \quad \mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ h) = \mathbb{E}_\theta \left[\min_{\tau \in \mathcal{S}_J} \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \tau(X_i)} \right]$$

and differs from (4). Note that it does not depend on the classifier h chosen to represent the clusterer thanks to the infimum over the permutations of the labels. Note also that the infimum inside the expectation reflects the ease of the clustering problem compared to the classification problem, because unlike classification, clustering does not seek to identify the true labels themselves.

It is customary to compare the performance of a given classifier h to the best performance attainable by an oracle classifier, namely the *Bayes risk of classification*:

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h).$$

In particular, it is well-known that the previous optimization problem is solved by the (albeit non necessarily unique) so-called Bayes classifier $h_\theta^* = (h_{\theta,1}^*, \dots, h_{\theta,n}^*)$ such that

$$\mathbb{P}_\theta(X_i = h_{\theta,i}^*(Y_{1:n}) \mid Y_{1:n}) = \max_{x \in \mathbb{X}} \mathbb{P}_\theta(X_i = x \mid Y_{1:n}), \quad i = 1, \dots, n.$$

In an unsupervised learning context, it makes sense to compare the risk of a clusterer to the *Bayes risk of clustering*:

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g).$$

As for the classification risk, the solution to the previous optimization problem exists and is obtained by the Bayes clusterer g_θ^* such that $g_\theta^*(Y_{1:n})$ is the partition that minimizes

$$(6) \quad g \mapsto \mathbb{E}_\theta \left[1 - \frac{1}{n} \sup_{\substack{M \subseteq \mathcal{E}(g(Y_{1:n}), \Pi_n) \\ M \text{ is a matching}}} \sum_{\{C, C'\} \in M} \text{Card}(C \cap C') \middle| Y_{1:n} \right].$$

In contrast with classification, however, the Bayes clusterer has usually no simple expression. It is to be noted that there is no reason that $g_\theta^* = \pi_n \circ h_\theta^*$. An analysis will be conducted to determine when the equality $g_\theta^* = \pi_n \circ h_\theta^*$ holds. Although the inequality $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$ is true and easily proved from (5), it is not guaranteed that $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g)$ and $\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$ are equivalent (*ie.* have comparable magnitude). In Sections 3.1 and 3.2 we show that the equivalence holds both in the independent and in the dependent scenarios as long as the Bayes risk of classification is not exponentially small in n or when there are only two clusters. Though we prove the two risks to be equivalent in some specific contexts, we provide a counter-example showing that this is not always true.

REMARK 1. It is common [13] to define a classifier as a function $h : \mathbb{Y} \rightarrow \mathbb{X}$ (as opposed to $\mathbb{Y}^n \rightarrow \mathbb{X}^n$ in our Definition 2). This usual definition is motivated by the fact that in the independent scenario the law of $X_i | Y_{1:n}$ is that of $X_1 | Y_1$. Hence if one is willing to classify the vector $Y_{1:n}$, the Bayes classifier rewrites as $h_\theta^* = (h_{\theta,1}^*(Y_1), \dots, h_{\theta,1}^*(Y_n))$ with $h_{\theta,1}(y)$ maximizing $x \mapsto \mathbb{P}_\theta(X_1 = x | Y_1 = y)$ and the Bayes risk of classification equals $\mathbb{P}_\theta(h_{\theta,1}^*(Y_1) \neq X_1)$. Thus classifying Y_1 or $Y_{1:n}$ is not very different. In the dependent case, however, the situation differs since $\mathbf{X} | \mathbf{Y}$ is a inhomogeneous Markov chain. This implies in particular that classifying Y_1 or $Y_{1:n}$ are different problems when $n \geq 2$, and the optimal solution for classifying $Y_{1:n}$ can not be obtained from the optimal solution of classifying Y_1 .

REMARK 2. In [29], the authors define the risk of clustering of a n -classifier h in a different way:

$$\mathcal{R}_n^{\text{MRSS}}(\theta, h) := \mathbb{E}_\theta \left[\inf_{\tau \in \mathcal{S}_J} \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\tau(X_i) \neq h_i(Y_{1:n})} \middle| Y_{1:n} \right] \right]$$

Similarly, in [28], the risk of clustering is defined by:

$$\tilde{\mathcal{R}}_n^{\text{clust}}(\theta, h) := \inf_{\tau \in \mathcal{S}_J} \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\tau(X_i) \neq h_i(Y_{1:n})} \right]$$

Their definition is mathematically convenient as one can easily show that (see Lemma 3):

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{MRSS}}(\theta, h) = \inf_{h \in \mathcal{H}_n} \tilde{\mathcal{R}}_n^{\text{clust}}(\theta, h) = \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$$

Hence, the Bayes clusterer relative to their risks is $g_\theta^* = \pi_n \circ h_\theta^*$ with h_θ^* the Bayes classifier relative to the risk $\mathcal{R}_n^{\text{class}}(\theta, \cdot)$. However, contrary to our definition, it seems that there is no suitable statistical interpretation of these risks of clustering.

3. Main results. The Bayes risk of classification offers an advantage due to its closed formula, stemming from the explicit identification of the Bayes classifier. In contrast, the Bayes risk of clustering lacks this straightforward formulation. However, thanks to Equation (5), the risk of a clustering procedure can be seen as the risk of classification of the associated classifier up to the best random permutation (the one minimizing the sum inside the

expectation in Equation (5)). This is why a common idea is that risk of clustering and risk of classification are closely related. Following this intuition, authors of [38, 14] propose an inequality linking the minimax (over a specific class) risk of clustering to the minimax risk of classification, in the context of community detection under the stochastic block model. Their inequality is applied in [33] in the context of the mixture of two Gaussian distributions and in [28] for general subGaussian mixtures. Although their argument is neat, it relies heavily on two key ingredients: (i) in their model the latent partition is deterministic and viewed as a parameter, and (ii) their point of view is minimax and thus, it is enough to consider some well-chosen parameters in order to lower bound the minimax risk. Thus their argument is not transposable to the situation where one is interested in bounding the Bayes risk of clustering at every possible value of the parameter with random labels.

3.1. I.I.D. case. The following theorem shows that in the i.i.d. setting and when there are only two classes ($J = 2$), the clustering resulting from the Bayes classifier coincides with that of the Bayes clusterer, almost-everywhere.

THEOREM 1. *In the case of independent labels, if $J = 2$, then for all $\theta \in \Theta^{\text{ind}}$ and all $n \geq 2$*

$$g_\theta^*(Y_{1:n}) = \pi_n \circ h_\theta^*(Y_{1:n}) \quad \mathbb{P}_\theta\text{-a.s.}$$

The proof of Theorem 1 can be found in Section 6.1. Thanks to this result, the difference between the two risks of clustering and classification can be bounded almost tightly, as shown in the next theorem. Let $\varepsilon_{n,\theta} = \frac{1}{2} - \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$.

THEOREM 2. *When $J = 2$ and $\theta \in \Theta^{\text{ind}}$, $\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) = 0$ if and only if $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) = 0$. If $\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) \neq 0$ then the difference between the Bayes risks satisfies*

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq \min \left(\frac{(1 - 4\varepsilon_{n,\theta}^2)^{\frac{n}{2}}}{\frac{n}{2} \log \left(\frac{1+2\varepsilon_{n,\theta}}{1-2\varepsilon_{n,\theta}} \right)}, \sqrt{\frac{\pi}{2n}} \right).$$

Furthermore, there exists a universal constant $B > 0$ such that for all $n \geq 100$ and all $\theta \in \Theta^{\text{ind}}$

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \geq B \min \left(\frac{(1 - 4\varepsilon_{n,\theta}^2)^{\frac{n}{2} \left(1 + \frac{6.8}{1 \vee \sqrt{n\varepsilon_{n,\theta}}} \right)}}{\frac{n}{2} \log \left(\frac{1+2\varepsilon_{n,\theta}}{1-2\varepsilon_{n,\theta}} \right)}, \frac{1}{\sqrt{n}} \right).$$

The proof of Theorem 2 is given in Section 6.2. As emphasized by the lower bound, the upper bound on the difference between the two Bayes risks in the previous theorem is essentially tight. The lower bound also shows that the Bayes risk of clustering is always strictly smaller than the Bayes risk of classification, unless they are both zero (which happens when the two emission densities have disjoint support). Also, the difference between the risks decays super-polynomially in n as soon as $\varepsilon_{n,\theta} \gg \sqrt{\log(n)/n}$, and polynomially if $\varepsilon_{n,\theta} = O(\sqrt{\log(n)/n})$ with worst-case rate $\asymp n^{-1/2}$ when $\varepsilon_{n,\theta} = O(n^{-1/2})$.

A direct consequence of this result is that the Bayes risk of classification is equivalent to the Bayes risk of clustering for n large enough as shown in the following corollary.

COROLLARY 1. *In the case of independent labels with $J = 2$, for all $\theta \in \Theta^{\text{ind}}$ and all $n \geq 2$*

$$(1 - \alpha_n) \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) \leq \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$$

where $\alpha_n = 2 \min \left(2(1 + \varepsilon_{n,\theta}) \frac{(1 - 4\varepsilon_{n,\theta}^2)^{\frac{n-2}{2}}}{n \log \left(\frac{1 + 2\varepsilon_{n,\theta}}{1 - 2\varepsilon_{n,\theta}} \right)}, \frac{1}{1 - 2\varepsilon_{n,\theta}} \sqrt{\frac{\pi}{2n}} \right)$.

The following proposition shows that, contrary to Corollary 1, the two Bayes risks are not equivalent in general.

PROPOSITION 1. *Whenever $J > 2$ and $n \geq 1$:*

$$\inf_{\theta \in \Theta^{\text{ind}}} \frac{\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g)}{\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)} = 0.$$

The proof of Proposition 1 is given in Section 6.9. It sheds light on why the previously established equivalence when there were only two classes no longer holds when the number of classes exceeds two. The proof is based on the following intuition. The observations which are misclustered or misclassified appear only in regions of overlap between the emission densities (cf. Theorem 9 and Corollary 4). Consider now the situation where n observations are derived from a mixture of $J = 3$ densities F_1 , F_2 and F_3 with weights $\nu_1 = 1 - 2\eta$, $\nu_2 = \eta$ and $\nu_3 = \eta$. Assume that the support of F_1 is disjoint from that of F_2 and F_3 and that F_2 and F_3 have overlapping supports over a small region of the space. When the weight η is small, two situations are possible: eventually only one observation belongs to the support of F_2 or F_3 , in which case the clustering error is null, or more observations appear in this region of the space which happens with very small probability because η is small. Combining these two insights, the magnitude of the risk of clustering is shown to be negligible with respect to the risk of classification. The use of compactly supported distributions is not essential to the previous argument. For instance, the argument still holds modulo supplementary technicalities if one takes $F_1 = \mathcal{N}(-\alpha x, 1)$, $F_2 = \mathcal{N}(-\alpha, 1)$ and $F_3 = \mathcal{N}(\alpha, 1)$ with $\alpha > 0$ and $x > 0$ large enough.

This dichotomy between $J = 2$ and $J > 2$ concerns also the minimizers of the risks. Contrary to Theorem 1 (when $J = 2$), in the i.i.d. setting with $J > 2$, one can always find some model parameters for which the result of clustering using the Bayes classifier differs from that of the Bayes clusterer with positive probability, as shown in the next theorem.

THEOREM 3. *If $J > 2$, then for all $\theta \in \Theta^{\text{ind}}$, if*

$$\mathbb{P}_\theta \left(\bigcup_{j=1}^J \left\{ 0 < \max_{l \neq j} \nu_l f_l(Y) < \nu_j f_j(Y) \leq \sum_{l \neq j} \nu_l f_l(Y) \right\} \right) > 0$$

then,

$$\forall n \geq 2, \quad \mathbb{P}_\theta (g_\theta^*(Y_{1:n}) \neq \pi_n \circ h_\theta^*(Y_{1:n})) > 0.$$

The condition above can be ensured easily for many distributions $(f_j)_{j \in \mathbb{X}}$ such as multinomials, mixtures of Gaussians, etc. For example, it is valid when $J = 3$, $(\nu_1, \nu_2, \nu_3) = (0.4, 0.4, 0.2)$, and the emission densities are Gaussians with variance $\sigma^2 = 1$ and means $(\mu_1, \mu_2, \mu_3) = (1, 2, 3)$. The proof of this theorem is given in Section 6.3.

Even though the Bayes risks of clustering and classification are not equivalent uniformly over Θ^{ind} by Proposition 1, we show in the next theorem that when $J > 2$, they become equivalent when $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \gtrsim J^2 e^{-n\beta/8}$ where $\beta = \min_{j \neq k} (\nu_j + \nu_k)$.

THEOREM 4. *For all $\theta \in \Theta^{\text{ind}}$ and all $n \geq 1$ the following bounds hold*

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \geq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - \sqrt{\frac{\log(J!)}{2n}},$$

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \geq (1 - \xi_n) \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - J^2 e^{-n\beta/8},$$

where $\beta = \min_{j \neq k} (\nu_j + \nu_k)$ and $\xi_n = \frac{4e}{\beta} [\sqrt{\log(J!)/(2n)}]^{1-4/(n\beta)}$.

The proof of Theorem 4 is given in Section 6.5. Even if the second lower bound obtained does not establish an equivalence between the Bayes risks of clustering and classification, it remains useful for analyzing phase transitions. Specifically, under a mild assumption on β , this bound results in similar phase transitions for both risks.

Even if clustering using the Bayes classifier differs sometimes from that of the Bayes clusterer (as shown in Theorem 3), Theorem 4 provides guarantees for the risk of clustering using the Bayes classifier as shown by the following corollary.

COROLLARY 2. *For all $\theta \in \Theta^{\text{ind}}$ and all $n \geq 1$ the following bounds hold*

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq \mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ h_\theta^*) \leq \frac{1}{1 - \xi_n} \left(\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) + J^2 e^{-n\beta/8} \right)$$

with ξ_n and β as defined in Theorem 4. When there are only two classes

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq \mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ h_\theta^*) \leq \frac{1}{1 - \alpha_n} \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g)$$

where α_n is defined in Corollary 1.

3.2. HMM case. The behavior of the Bayes risks under the HMM modeling exhibits similarities and differences with the i.i.d. case. As we will see, the Bayes risks keep having the same behavior under the HMM setting, while their minimizers, on the other hand, can be different. First, contrary to the i.i.d. case where the result of clustering using the Bayes classifier matches that of the Bayes clusterer, it is always possible to find a set of parameters for which they differ under the HMM setting, as shown in the next theorem.

THEOREM 5. *In the case of dependent labels and when $J = 2$, there exists a subset $\tilde{\Theta} \subset \Theta^{\text{dep}}$ such that*

$$\left(\forall \theta \in \tilde{\Theta} \right) (\forall n \geq 2), \quad \mathbb{P}_\theta (g_\theta^*(Y_{1:n}) \neq \pi_n \circ h_\theta^*(Y_{1:n})) > 0.$$

We illustrate this in the simple situation of $n = 2$. Assume one observes two consecutive observations Y_1 and Y_2 of a HMM with transition matrix

$$Q = \begin{pmatrix} 1-p & p \\ q & 1-q \end{pmatrix}$$

and emission densities f_1 and f_2 . Denoting for simplicity $\bar{p} = 1 - p$ and $\bar{q} = 1 - q$, one can easily check that the Bayes clusterer puts the two observations in the same cluster when

$$(7) \quad q\bar{p}f_1(Y_1)f_1(Y_2) + p\bar{q}f_2(Y_1)f_2(Y_2) \geq pq(f_2(Y_1)f_1(Y_2) + f_1(Y_1)f_2(Y_2))$$

On the other hand, the Bayes classifier puts the two observations in the same class when

$$(\mathbb{P}_\theta(X_1 = 2 \mid Y_{1:2}) - \mathbb{P}_\theta(X_1 = 1 \mid Y_{1:2})) (\mathbb{P}_\theta(X_2 = 2 \mid Y_{1:2}) - \mathbb{P}_\theta(X_2 = 1 \mid Y_{1:2})) \geq 0,$$

or equivalently,

$$(8) \quad (pqf_2(Y_1)f_1(Y_2) + p\bar{q}f_2(Y_1)f_2(Y_2) - q\bar{p}f_1(Y_1)f_1(Y_2) - qp f_1(Y_1)f_2(Y_2)) \\ \times (qp f_1(Y_1)f_2(Y_2) + p\bar{q}f_2(Y_1)f_2(Y_2) - q\bar{p}f_1(Y_1)f_1(Y_2) - pqf_2(Y_1)f_1(Y_2)) \geq 0.$$

The two conditions (7) and (8) are not equivalent. In the simple situation of Bernoulli emissions $(f_1, f_2) = (\mathcal{B}(\alpha_1), \mathcal{B}(\alpha_2))$, and $Y_{1:2} = (1, 1)$, the Bayes clustering puts the two observations in the same cluster when

$$q\bar{p}\alpha_1^2 + p\bar{q}\alpha_2^2 > 2pq\alpha_1\alpha_2$$

while the Bayes classifier puts them always in the same cluster since the corresponding condition becomes

$$(p\bar{q}\alpha_2^2 - q\bar{p}\alpha_1^2)^2 \geq 0.$$

Note that the first condition is not always ensured when p and q are chosen near 1. Note also that in the situation where $p + q = 1$, the dependence structure is lost and the two conditions become equivalent to

$$(qf_1(Y_1) - pf_2(Y_1))(qf_1(Y_2) - pf_2(Y_2)) \geq 0$$

which is coherent with Theorem 1 in the i.i.d. case. This highlights the strong difference between the dependent and independent setting. The proof of Theorem 5 can be found in Section 6.7.

We now establish the equivalence between the two risks under the Markovian dependence of the labels. We consider the following assumption.

ASSUMPTION 1. $\delta = \min_{x, x'} Q_{x, x'} > 0$ and $\min_x \nu_x \geq \delta$.

The positive lower bound δ introduced in Assumption 1 makes the hidden Markov chain irreducible. It will be used in proving deviation inequalities or when forgetting properties of the chain are needed. Even if the minimizers of the risk of classification and the risk of clustering might differ, we are still able to prove the equivalence between the two risks. The following theorem shows this is the case for n large enough.

THEOREM 6. *If $J = 2$, then for all $\theta \in \Theta^{\text{dep}}$ such that Assumption 1 holds and all $n \geq 1$*

$$(1 - \tilde{\alpha}_n) \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) \leq \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h),$$

where $\tilde{\alpha}_n = 2e\left(\frac{1-\delta}{\delta}\right)^4 [\sqrt{\log(2)/(2n)}]^{1-2/n}$.

Thanks to Theorem 6, it suffices to study the Bayes risk of classification, since any bound on this risk can be extrapolated to the risk of clustering, regardless of the magnitude of the Bayes risk of clustering.

Let us now consider the case $J > 2$. As in the i.i.d. setting where the minimizers of the risks do not always coincide, it is always possible to find a set of parameters for which the HMM observations are dependent, but the results of clustering using the Bayes clusterer and the partition induced by the Bayes classifier are not the same, as shown in the next theorem.

THEOREM 7. *In the case of dependent labels, for all $J > 2$ and all $n \geq 2$, there exists a subset $\tilde{\Theta}_{n,J} \subset \Theta^{\text{dep}}$ such that*

$$\forall \theta \in \tilde{\Theta}_{n,J}, \quad \mathbb{P}_\theta(g_\theta^*(Y_{1:n}) \neq \pi_n \circ h_\theta^*(Y_{1:n})) > 0.$$

The proof of Theorem 7 can be found in Section 6.8.

Finally, we establish a result similar to Theorem 4 under the HMM setting.

THEOREM 8. *For all $\theta \in \Theta^{\text{dep}}$ such that Assumption 1 holds and all $n \geq 1$, the following bounds are true*

$$\begin{aligned} \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) &\geq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - \frac{1}{1 - \rho_0} \sqrt{\frac{\log(J!)}{2n}}, \\ \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) &\geq (1 - \tilde{\xi}_n) \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - (J^2 + 1)e^{-2n(1-\rho_0)^2\beta^2/25}, \end{aligned}$$

where $\beta = \min_{i,j \neq k} \mathbb{P}_\theta(X_i \in \{j, k\})$, $\rho_0 = \frac{1-J\delta}{1-(J-1)\delta}$, and $\tilde{\xi}_n = \frac{5}{\beta(1-\rho_0)} \sqrt{\log(J!)/(2n)}$.

The proof of Theorems 6 and 8 is given in Section 6.6. Note that when the HMM is stationary, $\beta = \min_{j \neq k} (\nu_j + \nu_k)$ as in the i.i.d. case. Notice also that when $\delta = 1/J$, the observations are i.i.d. with uniform distribution over the set \mathbb{X} , and we recover the first inequality for i.i.d. observations. However, we do not recover the inequality for i.i.d. observations in general from that of HMM observations.

As for the i.i.d setting, even if clustering using the Bayes classifier differs sometimes from that of the Bayes clusterer (as shown in Theorem 5 and Theorem 7), Theorem 8 provides guarantees for the risk of clustering using the Bayes classifier as shown by the following corollary.

COROLLARY 3. *For all $\theta \in \Theta^{\text{dep}}$ and all $n \geq 1$ the following bounds hold*

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq \mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ h_\theta^*) \leq \frac{1}{1 - \tilde{\xi}_n} \left(\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) + (J^2 + 1)e^{-2n(1-\rho_0)^2\beta^2/25} \right)$$

where $\tilde{\xi}_n$, β and ρ_0 are as in Theorem 8. When there are only two classes

$$\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq \mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ h_\theta^*) \leq \frac{1}{1 - \tilde{\alpha}_n} \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g)$$

where $\tilde{\alpha}_n$ is as in Theorem 6.

3.3. A key quantity for the Bayes risk of clustering for both I.I.D. and HMM. We now state our main result which proves upper and lower bounds on the Bayes risks in function of a quantity measuring the separation between the emission densities up to constants depending on the transition matrix. These bounds translate into bounds on the risk of clustering thanks to the result above. Let $\Lambda := \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y)$.

THEOREM 9. *Under Assumption 1, for all $n \geq 1$, the Bayes risk of classification satisfies:*

$$\begin{aligned} \forall \theta \in \Theta^{\text{ind}}, \quad \delta \Lambda &\leq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) \leq (1 - (J-1)\delta) \Lambda \\ \forall \theta \in \Theta^{\text{dep}}, \quad \frac{\delta^2}{1 - (J-1)\delta} \Lambda &\leq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) \leq (1 - (J-1)\delta) \Lambda \end{aligned}$$

The proof of Theorem 9 is given in Section 6.10. Note that the bounds vanish when all the emission densities have disjoint supports. Note also that upper and lower bounds match when $\delta = \frac{1}{J}$ and the risk corresponds to Λ/J in this case. This situation corresponds to i.i.d. observations derived from a mixture with J components of equal weights. This proves in particular the tightness of the bounds which can not be improved by any absolute multiplicative constant without restricting the parameter space. Thanks to the results comparing the Bayes risks, Λ is the appropriate measure of the difficulty of clustering in many regimes as shown in the following corollary. Recall the definition of α_n from Corollary 1, of ξ_n from Theorem 4 of $\tilde{\alpha}_n$ from Theorem 6, of $\tilde{\xi}_n$ from Theorem 8, and of $\beta = \min_{i \in [n], j \neq k \in \mathbb{X}} \mathbb{P}_\theta(X_i \in \{j, k\})$.

COROLLARY 4. *Under Assumption 1, the following holds.*

- When $J = 2$:

$$\forall \theta \in \Theta^{\text{ind}}, \quad (1 - \alpha_n)\delta\Lambda \leq \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq (1 - \delta)\Lambda,$$

$$\forall \theta \in \Theta^{\text{dep}}, \quad \frac{\delta^2(1 - \tilde{\alpha}_n)}{1 - \delta}\Lambda \leq \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq (1 - \delta)\Lambda.$$

- When $J > 2$ and $\theta = (\nu, Q, (f_x)_{x \in \mathbb{X}}) \in \Theta^{\text{ind}}$ is such that $\delta\Lambda \geq 4J^2e^{-n\beta/8}$ and n is sufficiently large to have $\xi_n \leq \frac{1}{2}$:

$$\frac{\delta}{4}\Lambda \leq \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq (1 - (J - 1)\delta)\Lambda.$$

- When $J > 2$ and $\theta = (\nu, Q, (f_x)_{x \in \mathbb{X}}) \in \Theta^{\text{dep}}$ is such that $\delta^2\Lambda \geq 4(1 - (J - 1)\delta)(J^2 + 1)e^{-2n(1-\rho_0)^2\beta^2/15}$ and n is sufficiently large to have $\tilde{\xi}_n \leq \frac{1}{2}$:

$$\frac{\delta^2}{4(1 - \delta)}\Lambda \leq \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq (1 - (J - 1)\delta)\Lambda.$$

In other words, when there are only two classes, we obtain a tight characterization of the Bayes risk of clustering in terms of the separation between the two emission densities (i) covering all the regimes (ii) valid in the parametric and non-parametric setting (iii) without imposing any separation between the emission densities. In both i.i.d. and HMM settings, when there are only two classes, there exists a positive constant $\alpha(\delta)$ depending only on δ such that

$$\alpha(\delta) \int_{\mathbb{Y}} [f_1 \wedge f_2](y) d\mathcal{L}(y) \leq \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq (1 - \delta) \int_{\mathbb{Y}} [f_1 \wedge f_2](y) d\mathcal{L}(y).$$

For example, in the case of Gaussian emission distributions with two hidden states, this translates in a clear identification of the signal-to-noise ratio which drives the Bayes risk of clustering. For two Gaussian emission densities with means μ_0 and μ_1 and the same covariance matrix Σ , the Bayes risk of clustering ensures:

$$\frac{\alpha(\delta)}{2} \exp\left(-\frac{\text{SNR}}{4}\right) \leq \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \leq (1 - \delta) \exp\left(-\frac{\text{SNR}}{8}\right)$$

where $\text{SNR} = (\mu_0 - \mu_1)^\top \Sigma^{-1} (\mu_0 - \mu_1)$.

3.4. Reaching the Bayes risk.

3.4.1. *I.I.D. setting.* While there is no formal proof of this fact, it seems that there is no way of reaching the Bayes risk of clustering or classification without strong assumptions on the mixture components under the non-parametric i.i.d. mixture model. Without structural assumptions, the associated Bayes classifier can never be learnt nor approximated from the data in the i.i.d. case because the model is not identifiable and thus, the mixture components can not be estimated. The algorithms proposed in the literature perform well in clustering i.i.d. observations only when the clusters are assumed to be separated. This is the case for the k -means and its variants [20, 32] and the spectral algorithms [37, 23] to cite a few. We refer to [21] and Chapter 12 of [19] for a review of model-based clustering techniques.

3.4.2. *HMM setting.* Unlike the i.i.d. hypothesis, the HMM hypothesis allows us to identify the model without any assumptions about the emission distributions, apart from the fact that they are distinct [5]. This allows the construction of simple clustering procedures with risk comparable to the Bayes risk of clustering. The widely used method of clustering when the observations are drawn from a HMM consists in approaching the behavior of the Bayes classifier h_θ^* by plugging-in estimated parameters $\hat{\theta}$ and using the induced clusterer $\pi_n \circ h_{\hat{\theta}}^*$. Even though we have shown in Theorems 5 and 7 that the Bayes clusterer g_θ^* does not necessarily coincide with $\pi_n \circ h_{\hat{\theta}}^*$, the Corollary 3 justifies that it is not worth the effort of computing g_θ^* and use the plugging-in $g_{\hat{\theta}}^*$: the price to pay in the excess risk for trading the true θ by estimate $\hat{\theta}$ is most likely of several order of magnitude larger than the price to pay in using $\pi_n \circ h_{\hat{\theta}}^*$ in place of g_θ^* . Thus in this section, we focus on the clustering rule

$$g_{\hat{\theta}}^*(Y_{1:n}) := \pi_n \circ h_{\hat{\theta}}^*(Y_{1:n}) = \pi_n \left(\left(\arg \max_{x \in \mathbb{X}} \mathbb{P}_{\hat{\theta}}(X_i = x \mid Y_{1:n}) \right)_{1 \leq i \leq n} \right)$$

where $\hat{\theta} = \hat{\theta}(Y_{1:n})$ is an estimator constructed using the celebrated tensor method [6, 1]. The main advantage in using $\pi_n \circ h_{\hat{\theta}}^*$ is that in contrast with g_θ^* the classifier $h_{\hat{\theta}}^*$ is easily computed thanks to the recurrence formulas ensured by the Forward-Backward algorithm [8]. Theorem 10 below controls the excess risk of this clustering procedure. We will make the following assumption:

ASSUMPTION 2. The initial distribution ν is the stationary distribution of \mathbf{X} .

Notice however that under Assumption 2, the second part of Assumption 1 follows directly from the first part.

ASSUMPTION 3. Q is full-rank and aperiodic.

Under Assumptions 1, 2 and 3, the hidden Markov chain is stationary ergodic. The following assumption is sufficient to build estimators using the empirical distribution of the distribution of three consecutive observations.

ASSUMPTION 4. The emission densities $(f_x)_{x \in \mathbb{X}}$ are compactly supported and $C^* = \int \frac{dy}{\sum_{x \in \mathbb{X}} f_x(y)} < \infty$, the integral is over the union of the supports of the emission densities.

ASSUMPTION 5. $(f_x)_{x \in \mathbb{X}}$ are linearly independent and belong to $C^s(\mathbb{R})$ the space of locally Hölder continuous functions.

We now state the theorem. Its proof is given in Section 6.11.

THEOREM 10. *Let $D_n \rightarrow +\infty$ arbitrarily slowly. There exists a sequence of randomized estimators $(\hat{\theta}_n)_{n \geq 1}$ such that for all $\theta \in \Theta^{\text{dep}}$ satisfying Assumption 1 to 5*

$$\mathbb{E}[\mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ h_{\hat{\theta}}^*)] - \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) = \mathcal{O} \left(D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \right)$$

where the expectation $\mathbb{E}[\cdot]$ is understood with respect to the randomness of the algorithm.

The cornerstone of the proof is the analysis of how errors in estimation of the parameters propagate to errors in the filtering and smoothing distributions proved in [12]. To achieve this rate for the excess risk, we must specify an estimator. Here we make the choice of using a modified version of the spectral algorithm of [1] which relies on the tensor method developed in [6]. The full algorithm is given in details in Section 6.11. The main difference with the original algorithm of [1] is a modification guaranteeing that the algorithm outputs the same permutation for the estimation of the transition matrix \hat{Q} and the emission densities $\hat{f}_1, \dots, \hat{f}_J$. Indeed, in the previous works, the aim was more to get upper bounds on the risks up to label-switching, and the analysis was done for Q and the emission distributions separately which does not guarantee that the permutation used in the control of the error of estimation of Q and the emission densities is the same. Albeit the modification is rather natural, proof that it works require a substantial effort.

Notice that the rate on the excess risk is exactly the estimation rate of Hölder regular functions in sup-norm. However, since we do not observe realizations of each emission density separately, this rate is not a straightforward consequence of density estimation theory. See [1] where the usual non-parametric rate for the estimation of densities in sup-norm is obtained in the HMM context.

Though we analyse the question for the spectral estimator, we believe that results similar to Theorem 10 hold for most estimation procedures previously proposed in the literature, putting in the upper bounds the bounds on the estimation risk up to label-switching obtained in those works. This is why we shall use in Section 4 the least-squares estimation method for which a public and efficient code exists [11].

4. Numerical simulations. We present here the results of numerical simulations which leverage the added value of non-parametric clustering under hidden Markov modelling. We will consider two examples in which the hidden states will be generated through the same transition matrix

$$Q = \begin{pmatrix} 0.8 & 0.2 \\ 0.3 & 0.7 \end{pmatrix}.$$

EXAMPLE 1. A sample of $n = 5 \cdot 10^4$ observations of a HMM with transition matrix Q and emissions: $F_1 = \frac{1}{2} (\mathcal{N}(1.7, 0.2) + \mathcal{N}(7, 0.15))$ and $F_2 = \frac{1}{2} (\mathcal{N}(3.5, 0.2) + \mathcal{N}(5, 0.4))$.

EXAMPLE 2. A sample of $n = 10^5$ observations of a HMM with transition matrix Q and emissions: $F_1 = \frac{1}{2} (\mathcal{N}(3, 0.6) + \mathcal{N}(7, 0.4))$ and $F_2 = \frac{1}{2} (\mathcal{N}(5, 0.3) + \mathcal{N}(9, 0.4))$.

On these examples, we use the plug-in classifier whose clustering risk has been controlled in Theorem 10. Recall our procedure works as follows:

- First, the emission densities are estimated using the observations. We use the penalized least squares estimator proposed in [9] whose code has been made public in [11]. We use the histogram basis for the estimation.

- Second, the Forward-Backward algorithm is used to compute the a posteriori distributions of the hidden states under the estimated model parameters and given the observations.
- Third, the hidden states are estimated by maximizing the a posteriori distributions.

The results of clustering using the Forward-Backward algorithm will be compared to those using the k -means algorithm. Since we have access to the hidden states, the error of clustering can be estimated by choosing the best permutation. Figures 2, 3, and 4 display the results of estimation and clustering. Performance of Bayes classifier, plug-in classifier and k -means are reported in Table 1. If the observations were independent and the emissions

	Bayes classifier	Plug-in classifier	k -means algorithm	Λ
Example 1	1.56%	1.61%	46.7%	0.046
Example 2	6.42%	6.51%	47.3%	0.165


TABLE 1

Errors of clustering using three clustering rules: the Bayes classifier (using the true model parameters), the plug-in classifier (using the estimated parameters) and the k -means algorithm.


modelled non-parametrically, the unique quantity that could have been estimated consistently would be the stationary distribution. However, under the HMM assumption, the estimation of each emission density with the minimax rate is possible. This is due to the identifiability of the model which holds even if no assumption is made on the emission densities. This is not possible in the independent case. Figure 2 shows the estimation results of the emission densities and confirms the theoretical properties of the estimator. On the other hand, Corollary 4 proves that in the case of a HMM with two classes, clustering errors could appear only in zones where the two emission densities overlap. In Figures 3b, and 4b, misclustered observations appear only in the overlaps between the emission densities. Compared to the k -means algorithm which is purely geometric and does not exploit the distribution of the observations, the plug-in procedure allows combining together observations even if they are geometrically distant from each other, which is not possible with the k -means algorithm. In this context of Gaussian mixtures, the performance of the k -means algorithm is mediocre as depicted in Table 1 and does not improve significantly when the overlap between the emission densities is small. However, for the plug-in procedure, the more separated are the emission densities, the better are the results of clustering.

5. Discussions and Perspectives. This work focuses on an in-depth study of Bayes risks of clustering and classification. This analysis has led us to prove a form of equivalence between both risks. After identifying the key quantity which measures the difficulty of the classification task, it was extrapolated to the Bayes risk of clustering in several regimes. Finally, the excess risk of the plug-in procedure was studied. Although the analysis is sufficiently detailed to ensure a thorough understanding of both problems, there are still some interesting questions which were not covered by this work. We give a small overview below.

Lower-bound on entries of Q . Throughout our analysis of the Bayes risk of clustering, we have used Assumption 1 which was crucial in obtaining the lower-bounds. In the absence of such an assumption, the lower-bound of Theorem 9 no more matches the upper-bound and the magnitude of the Bayes risk of classification can not be precisely understood. The same thing applies to the Bayes risk of clustering. In addition, the control of the excess risk of the plug-in procedure is no more guaranteed.



(a) Estimation results for Example 1



(b) Estimation results for Example 2

Figure 2: Non-parametric penalized least squares density estimation using the histogram basis for Example 1 and Example 2

Approaching the frontier to independence. This situation happens when the emission distributions are nearly similar or when the transition matrix has almost equal lines. In this case, one can hope to improve the coefficient $\frac{\delta^2}{1-(J-1)\delta}$ which appears in the lower-bound on the risk of classification (Theorem 9) to δ as in the independent case. In fact, in this situation, the dependence between the observations is so weak and the effect of future and past observations on the current classification rule is so negligible that the magnitude of the Bayes risk of classification is the same as in the i.i.d. setting. On the other hand, as shown in [3, 2], estimation of the model parameters no more becomes possible when approaching the frontier. The plug-in procedure should not work as well.

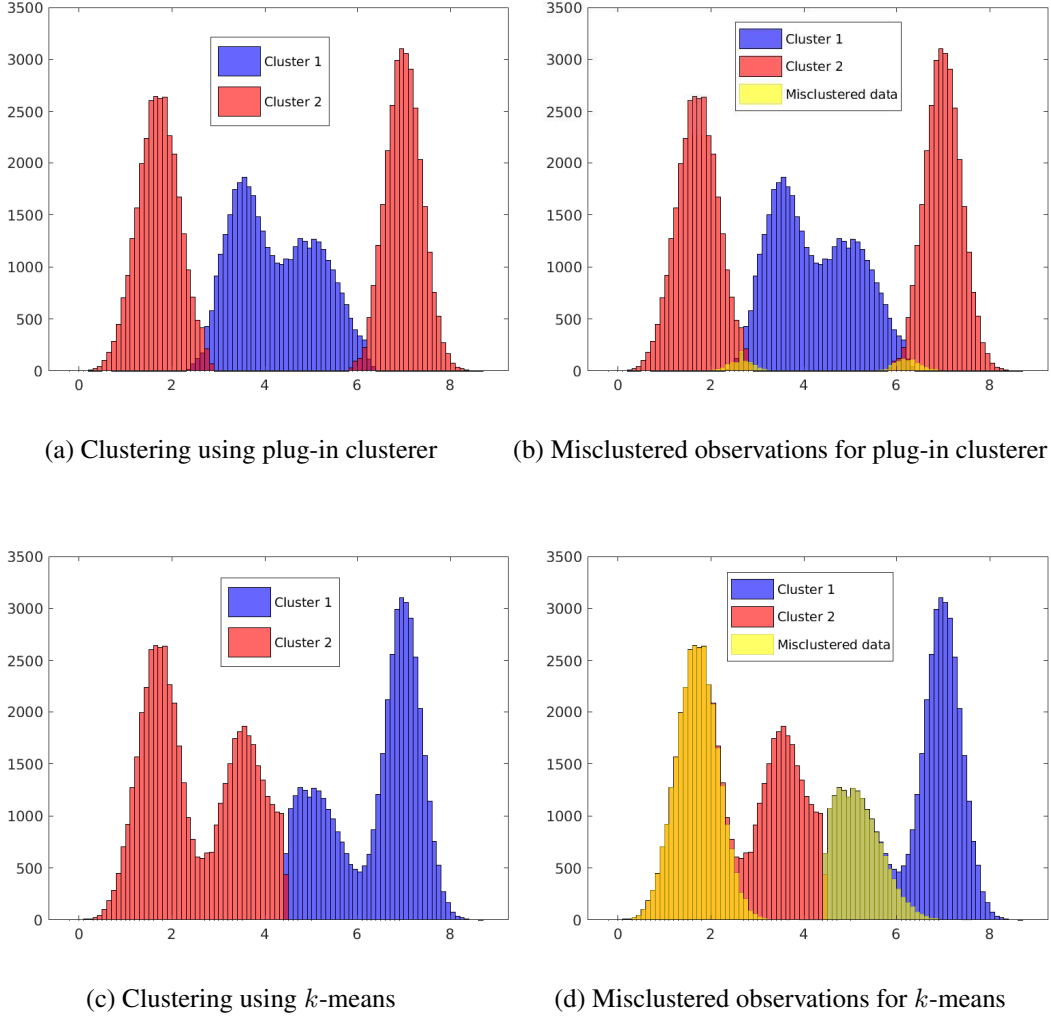


Figure 3: Histograms of clusters and clustering errors for Example 1

Lower bounds on the Bayes risk of clustering when it is very small. Theorems 2, 4, 6 and 8 establish lower bounds on the Bayes risk of clustering in terms of the Bayes risk of classification. When $J = 2$ these bounds are meaningful regardless of how small is the Bayes risk of clustering. When $J > 2$, however, these bounds can be vacuous if the Bayes risk of classification gets too small. This is not an artifact of our bounds since we have shown in Proposition 1 that the two Bayes risks are not uniformly comparable when $J > 2$. Whence from the current work we only know that the Bayes risk of clustering is driven by Λ in the region of parameters for which it is not exponentially small in n and that it can not be driven by Λ otherwise. If it was the case, then it would be equivalent to the Bayes risk of classification in contradiction with the Proposition 1. Understanding the Bayes risk of clustering in the region of extreme parameters is still an open question.

Fast rates. Theorem 10 has interest mainly in the situation where $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) \gtrsim (\log(n)/n)^{\frac{s}{2s+1}}$. In this regime, Theorem 10 tells us that the plug-in procedure has a risk of the same magnitude. However, in the situation where the magnitude of the Bayes risk of clustering is much smaller, that is when the emission distributions are very separated, one can hope to obtain faster rates. For example, when $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) = \mathcal{O}(e^{-cn})$ for c a

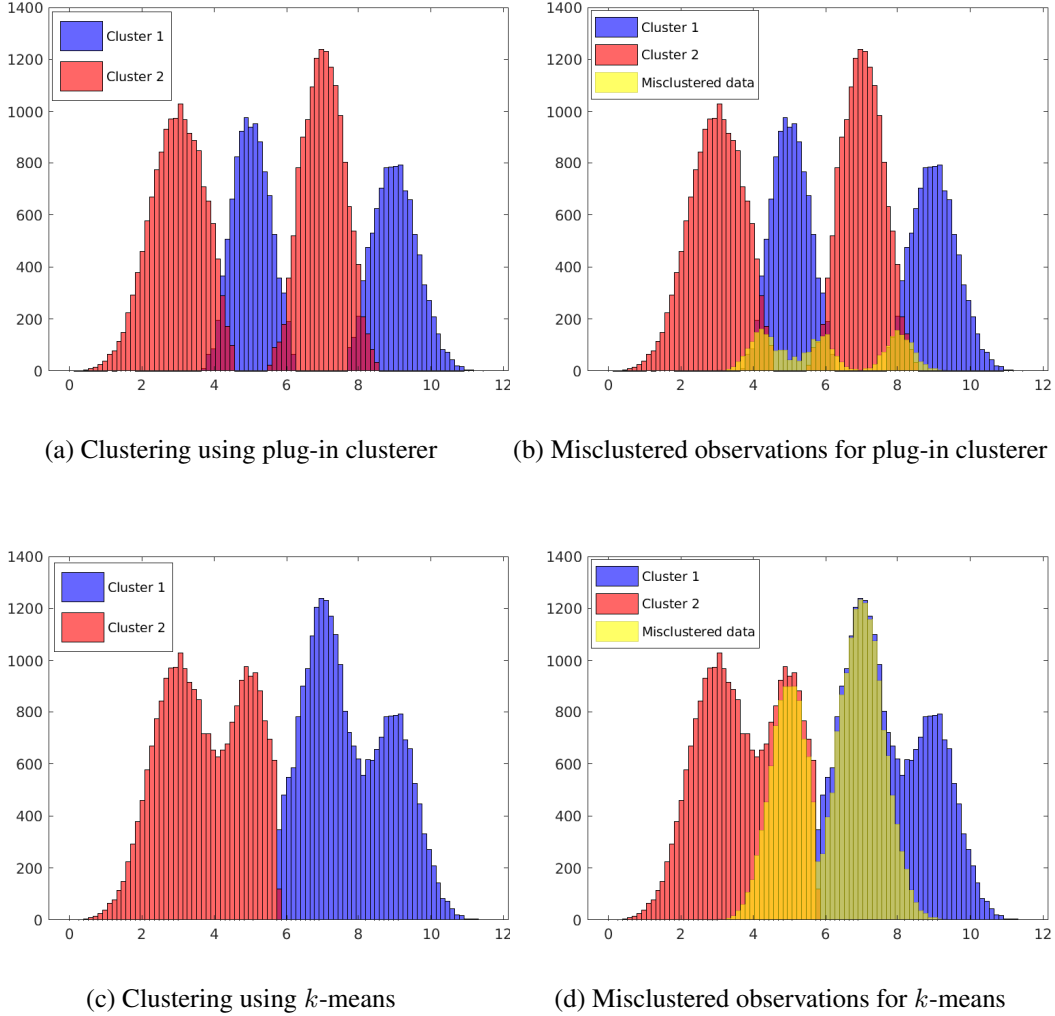


Figure 4: Histograms of clusters and clustering errors for Example 2

positive absolute constant, one can hope to show that the risk of the plug-in procedure is exponentially small in n . The following lemma represents a first step for the proof of such a result.

LEMMA 1. *For all $0 < \gamma < 1/2$ and all $\theta \in \Theta$*

$$(9) \quad \mathcal{R}_n^{\text{class}}(\theta, h_{\hat{\theta}}) \leq \frac{\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)}{1/2 - \gamma} + \frac{1}{n} \sum_{i=1}^n \mathbb{P}_{\theta} \left(\left\| \phi_{\theta, i|n} - \phi_{\hat{\theta}, i|n} \right\|_{\text{TV}} > \gamma \right).$$

where $\phi_{\theta, i|n} = \mathbb{P}_{\theta}(X_i \in \cdot \mid Y_{1:n})$ and $h_{\hat{\theta}}$ is the plug-in classifier defined in Section 3.4.2.

The proof of Lemma 1 is given in Section 6.12.

Observe that the second term of the rhs of (9) is a large deviation term which may eventually decrease exponentially fast in n . The only price to pay to obtain large deviation type of decay is a constant factor of at least 2 in front of $\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$. In situations where $\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$ is small, this might be advantageous compared to the bound in Theorem 10. In Proposition 2.2 of [12], the authors prove the inequality

$$(10) \quad \|\phi_{\theta,i|n}(\cdot, Y_{1:n}) - \phi_{\hat{\theta},i|n}(\cdot, Y_{1:n})\|_{\text{TV}} \leq \frac{4(1-\delta)}{\delta^2} \left(\rho^{i-1} \|\nu - \hat{\nu}\|_2 \right. \\ \left. + (1/(1-\rho) + 1/(1-\hat{\rho})) \|Q - \hat{Q}\|_{\text{F}} + \sum_{l=1}^n \delta \frac{(\hat{\rho} \vee \rho)^{|l-i|}}{c^*(Y_l)} \max_{x \in \mathbb{X}} |f_x(Y_l) - \hat{f}_x(Y_l)| \right)$$

where $c^*(y) = \min_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} Q(x, x') f_{x'}(y)$, $\rho = 1 - \delta/(1-\delta)$ and $\hat{\rho} = 1 - \hat{\delta}/(1-\hat{\delta})$ where $\hat{\delta}$ is an estimator of δ . (10) can be used to control the second term in the rhs of (9). This would require to derive large deviation inequalities for all the terms involved in (10), which turns out to be a rather challenging problem.

Optimal excess risk. Although we obtain upper bounds on the excess risk of clustering of the plug-in procedure (Theorem 10), we do not know the optimal rate of decay of the excess risk. In particular, it is unknown if the plug-in procedure achieves optimal excess risk. The Lemma 1 suggests that when the Bayes risk is smaller than $\mathcal{O}(n^{-s/(2s+1)})$ then our upper bounds on the excess risk of the plug-in could be improved. Yet without optimality guarantees. Determining the optimal excess risk is an open and interesting question to investigate.

Alternatives to plug-in. Under the hidden Markov modeling, the most straightforward way to take advantage of the identifiability of the model is to estimate the model parameters and use them for clustering through the plug-in procedure. Unlike the i.i.d. case where algorithms such as k -means can be used as an alternative, we do not know of any alternative to the plug-in in the HMM case. It would be very interesting to find clustering procedures that leverage the nonparametric identifiability of HMM without relying on estimating the parameters first.

6. Proofs.

6.1. *Proof of Theorem 1.* Let $\theta \in \Theta^{\text{ind}}$. Thanks to Equation (5), the Bayes clusterer g_θ^* can be defined as the partition $g_\theta^* = \pi_n \circ \tilde{h}_\theta$ where $\tilde{h}_\theta = (\tilde{h}_{\theta,i})_{i \in [n]}$ is the classifier minimizing:

$$(11) \quad h = (h_i)_{1 \leq i \leq n} \mapsto \mathbb{E}_\theta \left[\min_{\tau \in \mathcal{S}_2} \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \tau(X_i)} \right].$$

Let $Y_{1:n}$ be n i.i.d. observations of the mixture with parameters θ . We have $g_\theta^*(Y_{1:n}) = \pi_n \circ \tilde{h}_\theta(Y_{1:n})$ a.e, where

$$\tilde{h}_\theta(Y_{1:n}) \in \arg \min_{h=(h_i)_{i \in [n]}} \mathbb{E}_\theta \left[\min_{\tau \in \mathcal{S}_2} \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \tau(X_i)} \middle| Y_{1:n} \right].$$

Given that:

$$\mathbb{E}_\theta \left[\min_{\tau \in \mathcal{S}_2} \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \tau(X_i)} \middle| Y_{1:n} \right] = \mathbb{E}_\theta \left[\min \left(\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq X_i}, 1 - \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq X_i} \right) \middle| Y_{1:n} \right] \\ = \frac{1}{2} - \frac{1}{n} \mathbb{E}_\theta \left[\left| \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq X_i} - \frac{n}{2} \right| \middle| Y_{1:n} \right]$$

one gets $\tilde{h}_\theta(Y_{1:n}) \in \arg \max_{h=(h_i)_{i \in [n]}} \mathbb{E}_\theta \left[\left| \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq X_i} - \frac{n}{2} \right| \middle| Y_{1:n} \right]$. We now consider the following lemma. Its proof is due to Ziv Scully and can be found in [36]. We will detail its proof in Section 6.14 for the sake of completeness.

LEMMA 2. *Let $(Z_i)_{i \in [n]}$ be a sequence of independent Bernoulli random variables such that $Z_i \sim \mathcal{B}(p_i)$ where $p_i \in \{\alpha_i, 1 - \alpha_i\}$. Then the maximum $\max_{(p_i)_{i \in [n]}} \mathbb{E} \left[\left| \sum_{i \in [n]} Z_i - \frac{n}{2} \right| \right]$ is reached at $(p_i)_{i \in [n]} = (\alpha_i \wedge (1 - \alpha_i))_{i \in [n]}$ and $(p_i)_{i \in [n]} = (\alpha_i \vee (1 - \alpha_i))_{i \in [n]}$.*

We apply this lemma to the random variables $(\mathbf{1}_{h_i(Y_{1:n}) \neq X_i})_{i \in [n]}$ which are independent conditionally to $Y_{1:n}$ and ensure :

$$\mathbf{1}_{h_i(Y_{1:n}) \neq X_i} \mid Y_{1:n} \sim \mathcal{B}(p_i(Y_i))$$

such that $p_i(Y_i) \in \{\mathbb{P}_\theta(X_i = 1 \mid Y_i), \mathbb{P}_\theta(X_i = 2 \mid Y_i)\}$. Two cases occur:

- $(\forall i \in [n]) \mathbb{P}_\theta(X_i = 1 \mid Y_i) \neq 1/2$, then the Bayes classifier is unique and Lemma 2 allows us to conclude that $\tilde{h}_\theta(Y_{1:n}) = (\arg \max_{x=1,2} \mathbb{P}_\theta(X_i = x \mid Y_i))_{i \in [n]} = h_\theta^*(Y_{1:n})$. Consequently:

$$g_\theta^*(Y_{1:n}) = \pi_n \circ h_\theta^*(Y_{1:n}).$$

- $(\exists i \in [n]) \mathbb{P}_\theta(X_i = 1 \mid Y_i) = 1/2$, then the same argument yields

$$g_\theta^*(Y_{1:n}) = \pi_n \circ h_\theta^*(Y_{1:n}).$$

where we abuse the notation of $h_\theta^*(Y_{1:n})$ to refer not only to a specific Bayes classifier but to the set of all the Bayes classifiers since it is not unique. The same is done for the Bayes clusterer $g_\theta^*(Y_{1:n})$. The permutation π_n is then applied to the set of all Bayes classifiers.

6.2. *Proof of Theorem 2.* We first prove the upper bound. Let $h_{\theta,i}^*(Y_i) = \arg \max_{a \in \{1,2\}} \mathbb{P}_\theta(X_i = a \mid Y_i)$. Let $Z_n = \sum_{i=1}^n \mathbf{1}_{X_i \neq h_{\theta,i}^*(Y_i)}$. One gets:

$$\begin{aligned} \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) &= \mathbb{E}_\theta \left[\left(\frac{2}{n} Z_n - 1 \right) \mathbf{1}_{\frac{2}{n} Z_n > 1} \right] \\ &= \int_0^1 \mathbb{P}_\theta \left(\left(\frac{2}{n} Z_n - 1 \right) \mathbf{1}_{\frac{2}{n} Z_n > 1} > x \right) dx. \end{aligned}$$

Namely,

$$(12) \quad \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) - \inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g) = \int_0^1 \mathbb{P}_\theta \left(Z_n > \frac{n}{2} (x + 1) \right) dx =: J_n.$$

Chernoff bound yields:

$$\begin{aligned} \mathbb{P}_\theta \left(Z_n > \frac{n}{2} (x + 1) \right) &\leq \inf_{\lambda} \left\{ \exp \left(-\frac{n\lambda}{2} (x + 1) \right) \mathbb{E}_\theta [\exp(\lambda Z_n)] \right\} dx \\ &\leq \left(\left(\frac{1+x}{1-2\varepsilon_{n,\theta}} \right)^{\frac{1+x}{2}} \left(\frac{1-x}{1+2\varepsilon_{n,\theta}} \right)^{\frac{1-x}{2}} \right)^{-n}. \end{aligned}$$

Let $g(t) = \frac{1+t}{2} \log \left(\frac{1+t}{1-2\varepsilon_{n,\theta}} \right) + \frac{1-t}{2} \log \left(\frac{1-t}{1+2\varepsilon_{n,\theta}} \right)$. Then, $g'(t) = \frac{1}{2} \log \left(\frac{1+t}{1-t} \right) + \frac{1}{2} \log \left(\frac{1+2\varepsilon_{n,\theta}}{1-2\varepsilon_{n,\theta}} \right)$ and $g''(t) = \frac{1}{1-t^2}$. Deduce that $g(t) \geq g(0) + \max(g'(0)t, \frac{t^2}{2})$ for all $t \in (0, 1)$. Then,

$$\begin{aligned} J_n &\leq e^{-ng(0)} \int_0^1 e^{-n \max(g'(0)t, \frac{t^2}{2})} dt \\ &\leq e^{-ng(0)} \min \left(\int_0^1 e^{-ng'(0)t} dt, \int_0^1 e^{-nt^2/2} dt \right) \\ &\leq \min \left(\frac{e^{-ng(0)}}{ng'(0)}, \sqrt{\frac{\pi}{2n}} \right) \end{aligned}$$

We now derive the more challenging lower bound. We assume throughout that $n \geq 100$. We also assume that $\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) > 0$ otherwise the lower bound is zero and holds trivially. Suppose first that $0 \leq \varepsilon_{n,\theta} \leq \frac{1}{10\sqrt{n}}$. Then,

$$\begin{aligned} J_n &\geq \int_0^{\frac{1}{10\sqrt{n}}} P\left(S_n > \frac{n(x+1)}{2}\right) dx \\ &\geq \frac{1}{10\sqrt{n}} P\left(S_n > \frac{n}{2} + \frac{\sqrt{n}}{20}\right) \\ &= \frac{1}{10\sqrt{n}} P\left(\frac{S_n - n(\frac{1}{2} - \varepsilon_{n,\theta})}{\sqrt{n(1 - \varepsilon_{n,\theta}^2)/4}} > \frac{\frac{\sqrt{n}}{20} + n\varepsilon_{n,\theta}}{\sqrt{n(1 - \varepsilon_{n,\theta}^2)/4}}\right) \\ &\geq \frac{1}{10\sqrt{n}} P\left(\frac{S_n - n(\frac{1}{2} - \varepsilon_{n,\theta})}{\sqrt{n(1 - \varepsilon_{n,\theta}^2)/4}} > \sqrt{\frac{10}{111}}\right) \end{aligned}$$

because $n \geq 100$ and $\varepsilon_{n,\theta} \leq \frac{1}{10\sqrt{n}} \leq \frac{1}{100}$. By the theorem of Berry and Esseen,

$$J_n \geq \frac{1}{\sqrt{n}} \left(1 - \Phi\left(\sqrt{\frac{10}{111}}\right) - \frac{0.4748}{\sqrt{n(1 - \varepsilon_{n,\theta}^2)/4}}\right) \geq \frac{0.2870}{\sqrt{n}}$$

since $n \geq 100$. Finally, because $\varepsilon_{n,\theta} \leq \frac{1}{10\sqrt{n}}$ it must be that

$$\begin{aligned} \frac{\exp\left(-ng(0)\left[1 + \frac{6.8}{1\sqrt{10}\sqrt{n\varepsilon_{n,\theta}^2}}\right]\right)}{ng'(0)} &\geq \frac{1}{\sqrt{n}} \frac{\exp\left(-\frac{1}{2}\log(1 - 0.04/n)[1 + 6.8]\right)}{\frac{1}{2\sqrt{n}}\log\left(\frac{1+0.2/\sqrt{n}}{1-0.2/\sqrt{n}}\right)} \\ &\geq \frac{500}{\sqrt{n}}. \end{aligned}$$

Deduce that for a universal constant $B > 0$

$$(13) \quad J_n \geq B \min\left(\frac{\exp\left(-ng(0)\left[1 + \frac{6.8}{1\sqrt{10}\sqrt{n\varepsilon_{n,\theta}^2}}\right]\right)}{ng'(0)}, \frac{1}{\sqrt{n}}\right).$$

Now suppose $\frac{1}{10\sqrt{n}} < \varepsilon_{n,\theta} < \frac{1}{2}$. We first lower bound,

$$J_n \geq \int_0^{\frac{1}{10\sqrt{n}}} P\left(S_n > \frac{n(x+1)}{2}\right) dx.$$

We lower bound the probability $P(S_n > \frac{n}{2}(x+1))$ using Cramér's technique. In the next $0 \leq x \leq \frac{1}{10\sqrt{n}}$. Then for all $\lambda > 0$ and all $0 < \delta < \frac{n}{2} - \frac{\sqrt{n}}{20}$ (which guarantees that $\frac{n(1+x)}{2} + \delta < n$) we have

$$\begin{aligned} P\left(S_n > \frac{n}{2}(x+1)\right) &= \sum_{\frac{n(x+1)}{2} < y \leq n} \binom{n}{y} r^y (1-r)^{n-y} \\ &\geq \sum_{\frac{n(x+1)}{2} < y < \frac{n(x+1)}{2} + \delta} \binom{n}{y} e^{-\lambda y + n\psi_r(\lambda)} \frac{(re^\lambda)^y (1-r)^{n-y}}{\exp(n\psi_r(\lambda))} w \end{aligned}$$

where $\psi_r(\lambda) = \log(1 - r + re^\lambda)$ is the cumulant generating function of the Bernoulli distribution with parameter r ; where $r \equiv \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$ for simplicity. Here one notice that $y \mapsto \binom{n}{y} \frac{(pe^\lambda)^y (1-p)^{n-y}}{\exp(n\psi_r(\lambda))}$ is the density of the Binomial distribution with parameters (n, q_λ) where $q_\lambda = \frac{re^\lambda/(1-r)}{1+re^\lambda/(1-r)}$. Letting $\tilde{S}_n \sim \text{Binomial}(n, q_\lambda)$, it is seen that

$$P\left(S_n > \frac{n}{2}(x+1)\right) \geq e^{-\lambda(\frac{n(x+1)}{2} + \delta) + n\psi_r(\lambda)} P\left(\frac{n}{2}(x+1) < \tilde{S}_n < \frac{n}{2}(x+1) + \delta\right).$$

Now we make the choice that $q_\lambda = \frac{1+x}{2} + \delta/n$, which corresponds to $\lambda = -\log(1 - \frac{1+x+2\delta/n}{2}) + \log(\frac{1-r}{r})$. Observe that λ is well-defined and positive since by assumption $0 < r \leq \frac{1}{2}$ and $1+x+2\delta/n < 2$; this also guarantees that $0 < q_\lambda < 1$. Then,

$$\begin{aligned} P\left(S_n > \frac{n}{2}(x+1)\right) &\geq e^{-nI_r(\frac{1+x}{2} + \frac{\delta}{n})} P\left(\frac{n(1+x)}{2} < \tilde{S}_n < \frac{n(1+x)}{2} + \delta\right) w \\ &= e^{-ng(x+2\delta/n)} P\left(-\frac{\delta}{\sqrt{nq_\lambda(1-q_\lambda)}} < \frac{\tilde{S}_n - nq_\lambda}{\sqrt{nq_\lambda(1-q_\lambda)}} < 0\right) \end{aligned}$$

By the theorem of Berry and Esseen, we can conclude that

$$\begin{aligned} P\left(S_n > \frac{n}{2}(x+1)\right) &\geq e^{-ng(x+2\delta/n)} \left(\Phi(0) - \Phi\left(-\frac{\delta}{\sqrt{nq_\lambda(1-q_\lambda)}}\right) - 2\frac{0.4748}{\sqrt{nq_\lambda(1-q_\lambda)}} \right) \\ &\geq e^{-ng(x+2\delta/n)} \left(\frac{1}{2} - \Phi\left(-\frac{2\delta}{\sqrt{n}}\right) - \frac{1.8992}{\sqrt{n(1-\kappa_x^2)}} \right) \end{aligned}$$

where $\kappa_x := x + \frac{2\delta}{n}$. We choose $\delta = -\frac{\sqrt{n}}{2}\Phi^{-1}(1/4)$. This implies that $\kappa_x \leq \frac{0.1-\Phi^{-1}(1/4)}{\sqrt{n}} \leq \frac{0.7745}{\sqrt{n}}$ and $\Phi(-2\delta/\sqrt{n}) = \frac{1}{4}$. Consequently for $n \geq 100$,

$$P\left(S_n > \frac{n}{2}(x+1)\right) \geq 0.0595 \cdot e^{-ng(x+2\delta/n)}, \text{ and, } J_n \geq 0.0595 \int_0^{\frac{1}{10\sqrt{n}}} e^{-ng(x+2\delta/n)} dx.$$

A Taylor expansion of g near zero yields the existence of $u \in (0, x+2\delta/n)$ such that $g(x+2\delta/n) = g(0) + g'(0)(x+2\delta/n) + g''(u)(x+2\delta/n)^2/2$. But when $x \in (0, \frac{1}{10\sqrt{n}})$ it must be that $0 \leq x+2\delta/n \leq \frac{0.1-\Phi^{-1}(1/4)}{\sqrt{n}} \leq \frac{0.7745}{\sqrt{n}}$ and thus $g''(u) = \frac{1}{1-u^2} \leq 1.0061$ when $n \geq 100$. Hence,

$$\begin{aligned} J_n &\geq 0.0595 \cdot e^{-ng(0)-2\delta g'(0)-\frac{1.0061 \cdot 0.7745^2}{2}} \int_0^{\frac{1}{10\sqrt{n}}} e^{-ng'(0)x} dx \\ &= 0.0440 \cdot \left(1 - e^{-\frac{g'(0)\sqrt{n}}{10}}\right) \frac{e^{-ng(0)[1+\frac{2\delta g'(0)}{ng(0)}]}}{ng'(0)}. \end{aligned}$$

Here recall that $g(0) = -\frac{1}{2}\log(1-4\varepsilon_{n,\theta}^2)$ and $g'(0) = \frac{1}{2}\log(\frac{1+2\varepsilon_{n,\theta}}{1-2\varepsilon_{n,\theta}})$. It follows that the function $\varepsilon_{n,\theta} \mapsto g'(0) - \frac{g(0)}{\varepsilon_{n,\theta}}$ admits the derivative $\varepsilon_{n,\theta} \mapsto \frac{-(1-4\varepsilon_{n,\theta}^2)\log(1-4\varepsilon_{n,\theta}^2)-4\varepsilon_{n,\theta}^2}{\varepsilon_{n,\theta}^2(1-4\varepsilon_{n,\theta}^2)}$ which is negative for $\varepsilon_{n,\theta} > 0$. Deduce that $g'(0) - \frac{g(0)}{\varepsilon_{n,\theta}} \leq 0$, or equivalently $\frac{g'(0)}{g(0)} \leq \frac{1}{\varepsilon_{n,\theta}}$. Therefore,

$$\frac{2\delta g'(0)}{ng(0)} \leq \frac{-\Phi^{-1}(1/4)}{\sqrt{n}\varepsilon_{n,\theta}} \leq \frac{6.8}{\max(1, 10\sqrt{n\varepsilon_{n,\theta}^2})}.$$

Similarly since $\varepsilon_{n,\theta} \mapsto g'(0)$ is monotonically increasing, $\frac{g'(0)\sqrt{n}}{10} \geq \frac{\frac{1}{2} \log(\frac{1+0.4/\sqrt{n}}{1-0.4/\sqrt{n}})\sqrt{n}}{10} \geq 0.04$. Hence,

$$J_n \geq 0.0017 \cdot \frac{\exp\left(-ng(0)\left[1 + \frac{6.8}{1\vee 10\sqrt{n\varepsilon_{n,\theta}^2}}\right]\right)}{ng'(0)}$$

Finally, since $\varepsilon_{n,\theta} > \frac{1}{10\sqrt{n}}$ the above computations show that $ng'(0) \geq 0.4\sqrt{n}$. Therefore there is a universal constant $B > 0$ such that (13) is also satisfied when $\frac{1}{10\sqrt{n}} < \varepsilon_{n,\theta} < \frac{1}{2}$.

6.3. Proof of Theorem 3. Given two partitions A and B , we recall the clustering loss defined in (2) :

$$\ell(A, B) = 1 - \frac{1}{n} \sup_{\substack{M \subseteq \mathcal{E}(A, B) \\ M \text{ is a matching}}} \sum_{\{C, C'\} \in M} \text{Card}(C \cap C')$$

We define

$$(\forall i \in [n]) (\forall k \in \mathbb{X}) \quad \alpha_k^{(Y_i)} = \frac{\nu_k f_k(Y_i)}{\sum_{j=1}^J \nu_j f_j(Y_i)}$$

and $h_{\theta}^*(Y_{1:n}) = \left(h_{\theta,i}^*(Y_i)\right)_{i \in [n]}$ where $(\forall i \in [n]) \quad h_{\theta,i}^*(Y_i) = \arg \max_{k \in \mathbb{X}} \alpha_k^{(Y_i)}$.

Consider the event :

$$A_n = \bigcup_{j=1}^J \bigcap_{i=1}^n \left\{ \max_{k \neq j} \nu_k f_k(Y_i) < \nu_j f_j(Y_i) \right\}.$$

Then,

$$A_n \subset \left\{ \pi_n \left((h_{\theta,i}^*(Y_i))_{i \in [n]} \right) = \pi_n((1, \dots, 1)) \right\}.$$

Let $\theta \in \Theta^{\text{ind}}$, such that

$$\mathbb{P}_{\theta} \left(\bigcup_{j=1}^J \left\{ 0 < \max_{l \neq j} \nu_l f_l(Y) < \nu_j f_j(Y) \leq \sum_{l \neq j} \nu_l f_l(Y) \right\} \right) > 0.$$

Then $\mathbb{P}_{\theta}(A_n) > 0$. Since:

$$\begin{aligned} A_n &\cap \left\{ \mathbb{E}_{\theta} \left[\ell(\pi_n(X_{1:n}), \pi_n((1, \dots, 1))) \middle| Y_{1:n} \right] > \mathbb{E}_{\theta} \left[\ell(\pi_n(X_{1:n}), \pi_n((1, \dots, 1, 2))) \middle| Y_{1:n} \right] \right\} \\ &\subset \{g_{\theta}^*(Y_{1:n}) \neq \pi_n \circ h_{\theta}^*(Y_{1:n})\}, \end{aligned}$$

it suffices then to show that

$$\mathbb{P}_{\theta} \left(\mathbb{E}_{\theta} \left[\ell(\pi_n(X_{1:n}), \pi_n((1, \dots, 1))) \middle| Y_{1:n} \right] > \mathbb{E}_{\theta} \left[\ell(\pi_n(X_{1:n}), \pi_n((1, \dots, 1, 2))) \middle| Y_{1:n} \right] \middle| A_n \right) > 0.$$

Let:

$$k_n^{(1)}(x_{1:n}) = \arg \max_{i \in \mathbb{X}} \sum_{k=1}^n \mathbf{1}_{x_k=i}$$

$$N_n^{(1)}(x_{1:n}) = \max_{i \in \mathbb{X}} \sum_{k=1}^n \mathbf{1}_{x_k=i}$$

$$N_n^{(2)}(x_{1:n}) = \max_{i \neq k_n^{(1)}(x_{1:n})} \sum_{k=1}^n \mathbf{1}_{x_k=i}$$

$$N_{n,i}(x_{1:n}) = \sum_{k=1}^n \mathbf{1}_{x_k=i}$$

First, note that for $x_{1:n} \in \mathbb{X}^n$:

$$\ell(\pi_n(x_{1:n}), \pi_n(1, \dots, 1, 2)) = \begin{cases} n - N_n^{(1)}(x_{1:n}) - 1 & \text{if } x_n \neq k_n^{(1)}(x_{1:n}) \\ n - N_n^{(1)}(x_{1:n}) + 1 & \text{if } N_n^{(2)}(x_{1:n}) < N_n^{(1)}(x_{1:n}) - 1, x_n = k_n^{(1)}(x_{1:n}) \\ n - N_n^{(1)}(x_{1:n}) & \text{if } N_n^{(2)}(x_{1:n}) = N_n^{(1)}(x_{1:n}) - 1, x_n = k_n^{(1)}(x_{1:n}) \\ n - N_n^{(1)}(x_{1:n}) - 1 & \text{if } N_n^{(2)}(x_{1:n}) = N_n^{(1)}(x_{1:n}), x_n = k_n^{(1)}(x_{1:n}) \end{cases}$$

The inequality

$$(14) \quad \mathbb{E}_\theta [\ell(\pi_n(X_{1:n}), \pi_n((1, \dots, 1))) | Y_{1:n}] > \mathbb{E}_\theta [\ell(\pi_n(X_{1:n}), \pi_n((1, \dots, 1, 2))) | Y_{1:n}]$$

is equivalent to:

$$\sum_{x_{1:n} \in \mathbb{X}^n} \ell(\pi_n(x_{1:n}), \pi_n((1, \dots, 1))) \prod_{i=1}^n \alpha_{x_i}^{(Y_i)} > \sum_{x_{1:n} \in \mathbb{X}^n} \ell(\pi_n(x_{1:n}), \pi_n((1, \dots, 1, 2))) \prod_{i=1}^n \alpha_{x_i}^{(Y_i)}$$

Conditionally to $Y_{1:n}$, X_1, \dots, X_n are independent multinomial random variables such that

$X_i | Y_i \sim \left(\alpha_k^{(Y_i)} \right)_{k \in \mathbb{X}}$. Using the expression of $\ell(\pi_n(x_{1:n}), \pi_n((1, \dots, 1, 2)))$ and the fact that

$\ell(\pi_n(x_{1:n}), \pi_n((1, \dots, 1))) = n - N_n^{(1)}(x_{1:n})$, one obtains:

$$(14) \Leftrightarrow \sum_{\substack{x_{1:n} \in \mathbb{X}^n \\ x_n \neq k_n^{(1)}(x_{1:n}) \text{ or } \\ x_n = k_n^{(1)}(x_{1:n}) \text{ and } N_n^{(2)}(x_{1:n}) = N_n^{(1)}(x_{1:n})}} \prod_{i=1}^n \alpha_{x_i}^{(Y_i)} > \sum_{\substack{x_{1:n} \in \mathbb{X}^n \\ x_n = k_n^{(1)}(x_{1:n}) \text{ and } N_n^{(2)}(x_{1:n}) < N_n^{(1)}(x_{1:n}) - 1}} \prod_{i=1}^n \alpha_{x_i}^{(Y_i)}$$

$$\Leftrightarrow \mathbb{P}_\theta \left(X_n \neq k_n^{(1)}(X_{1:n}) \middle| Y_{1:n} \right) + \mathbb{P}_\theta \left(X_n = k_n^{(1)}(X_{1:n}), N_n^{(2)}(X_{1:n}) = N_n^{(1)}(X_{1:n}) \middle| Y_{1:n} \right)$$

$$> \mathbb{P}_\theta \left(X_n = k_n^{(1)}(X_{1:n}), N_n^{(2)}(X_{1:n}) < N_n^{(1)}(X_{1:n}) - 1 \middle| Y_{1:n} \right)$$

$$\Leftrightarrow \mathbb{P}_\theta \left(X_n \neq k_n^{(1)}(X_{1:n}) \middle| Y_{1:n} \right) - \mathbb{P}_\theta \left(X_n = k_n^{(1)}(X_{1:n}), N_n^{(2)}(X_{1:n}) < N_n^{(1)}(X_{1:n}) - 1 \middle| Y_{1:n} \right)$$

$$+ \mathbb{P}_\theta \left(X_n = k_n^{(1)}(X_{1:n}), N_n^{(2)}(X_{1:n}) = N_n^{(1)}(X_{1:n}) \middle| Y_{1:n} \right) > 0$$

By marginalization over the possible values of X_n , one gets:

$$\begin{aligned}
(14) &\Leftrightarrow \sum_{j=1}^J \left[\mathbb{P}_\theta \left(X_n \neq j, N_{n,j}(X_{1:n}) \geq \max_{k \neq j} N_{n,k}(X_{1:n}) \middle| Y_{1:n} \right) \right. \\
&\quad - \mathbb{P}_\theta \left(X_n = j, N_{n,j}(X_{1:n}) > \max_{k \neq j} N_{n,k}(X_{1:n}) + 1 \middle| Y_{1:n} \right) \\
&\quad \left. + \mathbb{P}_\theta \left(X_n = j, N_{n,j}(X_{1:n}) = \max_{k \neq j} N_{n,k}(X_{1:n}) \middle| Y_{1:n} \right) \right] > 0 \\
&\Leftrightarrow \sum_{j=1}^J \sum_{l \neq j} \left[\alpha_l^{(Y_n)} \mathbb{P}_\theta \left(N_{n-1,j}(X_{1:n-1}) \geq \max_{k \neq j, l} N_{n-1,k}(X_{1:n-1}) \vee (N_{n-1,l}(X_{1:n-1}) + 1) \middle| Y_{1:n-1} \right) \right. \\
&\quad - \frac{\alpha_l^{(Y_n)} \alpha_j^{(Y_n)}}{1 - \alpha_j^{(Y_n)}} \mathbb{P}_\theta \left(N_{n-1,j}(X_{1:n-1}) > \max_{k \neq j} N_{n-1,k}(X_{1:n-1}) \middle| Y_{1:n-1} \right) \\
&\quad \left. - \frac{\alpha_l^{(Y_n)} \alpha_j^{(Y_n)}}{1 - \alpha_j^{(Y_n)}} \mathbb{P}_\theta \left(N_{n-1,j}(X_{1:n-1}) = \max_{k \neq j} N_{n-1,k}(X_{1:n-1}) - 1 \middle| Y_{1:n-1} \right) \right] > 0 \\
&\Leftrightarrow \sum_{j=1}^J \sum_{l \neq j} \left[\alpha_l^{(Y_n)} \left\{ \mathbb{P}_\theta \left(N_{n-1,j}(X_{1:n-1}) \geq \max_{k \neq j, l} N_{n-1,k}(X_{1:n-1}) \vee (N_{n-1,l}(X_{1:n-1}) + 1) \middle| Y_{1:n-1} \right) \right. \right. \\
&\quad \left. - \frac{\alpha_j^{(Y_n)}}{1 - \alpha_j^{(Y_n)}} \mathbb{P}_\theta \left(N_{n-1,j}(X_{1:n-1}) > \max_{k \neq j} N_{n-1,k}(X_{1:n-1}) \middle| Y_{1:n-1} \right) \right\} \\
&\quad \left. + \sum_{j=1}^J \alpha_j^{(Y_n)} \mathbb{P}_\theta \left(N_{n-1,j}(X_{1:n-1}) = \max_{k \neq j} N_{n-1,k}(X_{1:n-1}) - 1 \middle| Y_{1:n-1} \right) \right] > 0.
\end{aligned}$$

On the one hand, since:

$$\begin{aligned}
&\left\{ N_{n-1,j}(X_{1:n-1}) > \max_{k \neq j} N_{n-1,k}(X_{1:n-1}) \right\} \\
&\subset \left\{ N_{n-1,j}(X_{1:n-1}) \geq \max_{k \neq j, l} N_{n-1,k}(X_{1:n-1}) \vee (N_{n-1,l}(X_{1:n-1}) + 1) \right\}
\end{aligned}$$

and $\alpha_k^{(Y_n)} \leq \frac{1}{2} \iff \frac{\alpha_k^{(Y_n)}}{1 - \alpha_k^{(Y_n)}} \leq 1$, one gets:

$$\begin{aligned}
&\bigcap_{k \in \mathbb{X}} \left\{ \alpha_k^{(Y_n)} \leq \frac{1}{2} \right\} \\
&\subset \left\{ \sum_{j=1}^J \sum_{l \neq j} \alpha_l^{(Y_n)} \left[\mathbb{P}_\theta \left(N_{n-1,j}(X_{1:n-1}) \geq \max_{k \neq j, l} N_{n-1,k}(X_{1:n-1}) \vee (N_{n-1,l}(X_{1:n-1}) + 1) \middle| Y_{1:n-1} \right) \right. \right. \\
&\quad \left. \left. - \frac{\alpha_j^{(Y_n)}}{1 - \alpha_j^{(Y_n)}} \mathbb{P}_\theta \left(N_{n-1,j}(X_{1:n-1}) > \max_{k \neq j} N_{n-1,k}(X_{1:n-1}) \middle| Y_{1:n-1} \right) \right] \geq 0 \right\}.
\end{aligned}$$

On the other hand,

$$\bigcup_{l \neq j \in \mathbb{X}} \bigcap_{i=1}^n \left\{ \alpha_l^{(Y_i)} \wedge \alpha_j^{(Y_i)} > 0 \right\} \subset \left\{ \sum_{j=1}^J \alpha_j^{(Y_n)} \mathbb{P}_\theta \left(N_{n-1,j}(X_{1:n-1}) = \max_{k \neq j} N_{n-1,k}(X_{1:n-1}) - 1 \middle| Y_{1:n-1} \right) > 0 \right\}$$

because for $k \in \mathbb{X}$, $N_{n-1,k}(X_{1:n-1})$ is a sum of Bernoulli random variables. Consequently,

$$\bigcup_{l \neq j \in \mathbb{X}} \bigcap_{i=1}^n \left\{ \alpha_l^{(Y_i)} \wedge \alpha_j^{(Y_i)} > 0 \right\} \cap \bigcap_{k \in \mathbb{X}} \left\{ \alpha_k^{(Y_n)} \leq \frac{1}{2} \right\} \subset \{(14)\}$$

Consequently,

$$\begin{aligned} & \mathbb{P}_\theta \left(\bigcup_{l \neq j \in \mathbb{X}} \bigcap_{i=1}^n \left\{ \alpha_l^{(Y_i)} \wedge \alpha_j^{(Y_i)} > 0 \right\} \cap \bigcap_{k \in \mathbb{X}} \left\{ \alpha_k^{(Y_n)} \leq \frac{1}{2} \right\} \middle| A_n \right) \\ & \leq \mathbb{P}_\theta \left(\mathbb{E}_\theta \left[\ell(\pi_n(X_{1:n}), \pi_n((1, \dots, 1))) \middle| Y_{1:n} \right] > \mathbb{E}_\theta \left[\ell(\pi_n(X_{1:n}), \pi_n((1, \dots, 1, 2))) \middle| Y_{1:n} \right] \middle| A_n \right) \end{aligned}$$

To conclude, we only need to prove that the conditional probability in the lower-bound is positive. Finally,

$$\begin{aligned} & \bigcup_{1 \leq j \neq k \leq J} \bigcap_{i=1}^n \left\{ \max_{l \neq j} \nu_l f_l(Y_i) < \nu_j f_j(Y_i) \leq \sum_{l \neq j} \nu_l f_l(Y_i), \nu_k f_k(Y_i) > 0 \right\} \\ & \subset \bigcup_{j=1}^J \bigcap_{i=1}^n \left\{ \max_{l \neq j} \nu_l f_l(Y_i) < \nu_j f_j(Y_i) \leq \frac{1}{2} \sum_{l=1}^J \nu_l f_l(Y_i) \right\} \cap \bigcup_{1 \leq l \neq j \leq J} \bigcap_{i=1}^n \{ \nu_l f_l(Y_i) > 0, \nu_j f_j(Y_i) > 0 \} \\ & \subset \bigcup_{j=1}^J \bigcap_{i=1}^n \left\{ \max_{l \neq j} \nu_l f_l(Y_i) < \nu_j f_j(Y_i) \right\} \cap \bigcup_{1 \leq l \neq j \leq J} \bigcap_{i=1}^n \{ \nu_l f_l(Y_i) > 0, \nu_j f_j(Y_i) > 0 \} \cap \bigcap_{k=1}^J \left\{ \alpha_k^{(Y_n)} < \frac{1}{2} \right\} \\ & \subset A_n \cap \bigcup_{1 \leq l \neq j \leq J} \bigcap_{i=1}^n \left\{ \alpha_l^{(Y_i)} \wedge \alpha_j^{(Y_i)} > 0 \right\} \cap \bigcap_{k=1}^J \left\{ \alpha_k^{(Y_n)} \leq \frac{1}{2} \right\} \end{aligned}$$

Given that the observations $Y_{1:n}$ are i.i.d. following the stationary distribution $\sum_{k=1}^J \nu_k f_k$ and that

$$\begin{aligned} & \bigcup_{j=1}^J \left\{ 0 < \max_{l \neq j} \nu_l f_l(Y) < \nu_j f_j(Y) \leq \sum_{l \neq j} \nu_l f_l(Y) \right\} \\ & \subset \bigcup_{1 \leq j \neq k \leq J} \left\{ \max_{l \neq j} \nu_l f_l(Y) < \nu_j f_j(Y) \leq \sum_{l \neq j} \nu_l f_l(Y), \nu_k f_k(Y) > 0 \right\} \end{aligned}$$

and that by assumption,

$$\mathbb{P}_\theta \left(\bigcup_{j=1}^J \left\{ 0 < \max_{l \neq j} \nu_l f_l(Y) < \nu_j f_j(Y) \leq \sum_{l \neq j} \nu_l f_l(Y) \right\} \right) > 0,$$

the result follows.

6.4. *Common elements to the proof of Theorems 4, 6 and 8.* Recall the definition of $\mathcal{R}_n^{\text{MRSS}}$ in Remark 2: for all $(\theta, h) \in \Theta \times \mathcal{H}_n$,

$$(15) \quad \mathcal{R}_n^{\text{MRSS}}(\theta, h) = \mathbb{E}_\theta \left[\min_{\tau \in \mathcal{S}_J} \mathbb{E}_\theta [U_{n,\tau}(h) \mid Y_{1:n}] \right]$$

where $U_{n,\tau}(h) := \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\tau(X_i) \neq h_i(Y_{1:n})}$. We also make use of the notation $\hat{p}_\tau(h) := \mathbb{E}_\theta [U_{n,\tau}(h) \mid Y_{1:n}]$. Let $\hat{\tau}_h$ denote a $Y_{1:n}$ -measurable permutation satisfying:

$$\hat{p}_{\hat{\tau}_h}(h) = \mathbb{E}_\theta [U_{n,\hat{\tau}_h}(h) \mid Y_{1:n}] = \min_{\tau} \mathbb{E}_\theta [U_{n,\tau}(h) \mid Y_{1:n}] = \min_{\tau} \hat{p}_\tau(h).$$

Instead of comparing $\mathcal{R}_n^{\text{clust}}$ and $\mathcal{R}_n^{\text{class}}$, we compare $\mathcal{R}_n^{\text{clust}}$ and $\mathcal{R}_n^{\text{MRSS}}$, which is enough to obtain the result thanks to the following easy lemma.

LEMMA 3. *For all $\theta \in \Theta$ and all $n \geq 1$*

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) = \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{MRSS}}(\theta, h).$$

PROOF. The optimal permutation $\hat{\tau}_h$ such that $\hat{p}_{\hat{\tau}_h}(h) = \min_{\tau \in \mathcal{S}_J} \mathbb{E}_\theta [U_{n,\tau}(h) \mid Y_{1:n}]$ is a $Y_{1:n}$ -measurable permutation valued random variable. Since any $h \in \mathcal{H}_n$ is also $Y_{1:n}$ -measurable, the result is immediate. \square

In the next we then focus on comparing $\mathcal{R}_n^{\text{clust}}$ and $\mathcal{R}_n^{\text{MRSS}}$. We let $N_j := \sum_{i=1}^n \mathbf{1}_{\{X_i=j\}}$ and $N_{(1)} \leq N_{(2)} \leq \dots \leq N_{(J)}$ denote the order statistics of (N_1, \dots, N_J) .

PROPOSITION 2. *A generic lower bound that works for any latent model (i.i.d. or HMM or whatever). For all classifiers h , all ε , all η and all $\theta \in \Theta$*

$$\begin{aligned} \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \right] &\geq \mathbb{E}_\theta \left[\min_{\tau} \mathbb{E}_\theta [U_{n,\tau}(h) \mid Y_{1:n}] \right] \\ &\quad - \mathbb{E}_\theta \left[\mathbb{E}_\theta \left[\max_{\tau} (-U_{n,\tau}(h) + \hat{p}_\tau(h)) \mid Y_{1:n} \right] \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}} \right] \\ &\quad - \mathbb{E}_\theta \left[\mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}} \right] \\ &\quad - \mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta). \end{aligned}$$

PROOF. For $\varepsilon \in [0, 1]$, we decompose

$$(16) \quad \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \right] = \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}} \right] + \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}} \right].$$

The first term in the rhs of (16) is handled via a small deviation principle, remarking that on the event $\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}$

$$\begin{aligned} \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \mid Y_{1:n} \right] &= \mathbb{E}_\theta \left[\min_{\tau} (\hat{p}_\tau(h) + U_{n,\tau}(h) - \hat{p}_\tau(h)) \mid Y_{1:n} \right] \\ &\geq \min_{\tau} \hat{p}_\tau(h) - \mathbb{E}_\theta \left[\max_{\tau} (-U_{n,\tau}(h) + \hat{p}_\tau(h)) \mid Y_{1:n} \right]. \end{aligned}$$

The second term in (16) is handled via a large deviation principle. By Lemma 4, on the event $\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}$

$$\begin{aligned} \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \mid Y_{1:n} \right] &\geq \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \mathbf{1}_{\{U_{n,\hat{\tau}_h}(h) \leq (N_{(1)} + N_{(2)})/(2n)\}} \mid Y_{1:n} \right] \\ &= \mathbb{E}_\theta \left[U_{n,\hat{\tau}_h}(h) \mathbf{1}_{\{U_{n,\hat{\tau}_h}(h) \leq (N_{(1)} + N_{(2)})/(2n)\}} \mid Y_{1:n} \right] \\ &\geq \hat{p}_{\hat{\tau}_h}(h) - \mathbb{P}_\theta \left(U_{n,\hat{\tau}_h}(h) > \frac{N_{(1)} + N_{(2)}}{2n} \mid Y_{1:n} \right). \end{aligned}$$

But,

$$\mathbb{P}_\theta \left(U_{n,\hat{\tau}_h}(h) > \frac{N_{(1)} + N_{(2)}}{2n} \mid Y_{1:n} \right) \leq \mathbb{P}_\theta (U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) + \mathbb{P}_\theta (N_{(1)} + N_{(2)} < 2n\eta \mid Y_{1:n}).$$

The generic lower bound follows. \square

LEMMA 4. *If $U_{n,\tau'}(h) \leq \frac{N_{(1)} + N_{(2)}}{2n}$ then $\min_\tau U_{n,\tau}(h) = U_{n,\tau'}(h)$.*

PROOF. Let $U_{n,\tau'}(h) \leq \frac{N_{(1)} + N_{(2)}}{2n}$ and suppose $\min_\tau U_{n,\tau}(h) < U_{n,\tau'}(h)$. Then, there exists a permutation $\tau'' \neq \tau'$ such that $U_{n,\tau''}(h) < U_{n,\tau'}(h)$. But then, letting $I = \{i \in \{1, \dots, n\} : \tau'(X_i) = h_i(Y_{1:n})\}$:

$$\begin{aligned} n(U_{n,\tau'}(h) - U_{n,\tau''}(h)) &= \sum_{i=1}^n \left(\mathbf{1}_{\{\tau'(X_i) \neq h_i(Y_{1:n})\}} - \mathbf{1}_{\{\tau''(X_i) \neq h_i(Y_{1:n})\}} \right) \\ &= \sum_{i=1}^n \left(\mathbf{1}_{\{\tau''(X_i) = h_i(Y_{1:n})\}} - \mathbf{1}_{\{\tau'(X_i) = h_i(Y_{1:n})\}} \right) \\ &= - \sum_{i \in I} \mathbf{1}_{\{\tau'(X_i) \neq \tau''(X_i)\}} + \sum_{i \in I^c} \mathbf{1}_{\{\tau'(X_i) = h_i(Y_{1:n})\}} \\ &= - \sum_{i=1}^n \mathbf{1}_{\{\tau'(X_i) \neq \tau''(X_i)\}} + \sum_{i \in I^c} \left(\mathbf{1}_{\{\tau'(X_i) = h_i(Y_{1:n})\}} + \mathbf{1}_{\{\tau'(X_i) \neq \tau''(X_i)\}} \right) \\ &\leq -(N_{(1)} + N_{(2)}) + 2|I^c| \\ &= -(N_{(1)} + N_{(2)}) + 2nU_{n,\tau'}(h) \end{aligned}$$

where we have used that since $\tau' \neq \tau''$, it must be that $\sum_{i=1}^n \mathbf{1}_{\{\tau'(X_i) \neq \tau''(X_i)\}} \geq N_{(1)} + N_{(2)}$. Rearranging the previous:

$$U_{n,\tau''}(h) \geq \frac{N_{(1)} + N_{(2)}}{n} - U_{n,\tau'}(h) \geq U_{n,\tau'}(h)$$

which contradicts that $U_{n,\tau''}(h) < U_{n,\tau'}(h)$. Hence $\min_\tau U_{n,\tau}(h) \geq U_{n,\tau'}(h)$. \square

6.5. *Proof of Theorem 4 (independent scenario).* Here we apply the result of Proposition 2 to the i.i.d. case.

When $J > 2$, the first trivial bound is obtained by choosing $\varepsilon = \eta = 0$. With this choice, Proposition 2 gives for $J \geq 2$:

$$\begin{aligned} \mathbb{E}_\theta \left[\min_\tau U_{n,\tau}(h) \right] &\geq \mathbb{E}_\theta \left[\min_\tau \mathbb{E}_\theta [U_{n,\tau}(h) \mid Y_{1:n}] \right] - \mathbb{E}_\theta \left[\mathbb{E}_\theta \left[\max_\tau (-U_{n,\tau}(h) + \hat{p}_\tau(h)) \mid Y_{1:n} \right] \right] \\ &\geq \mathbb{E}_\theta \left[\min_\tau \mathbb{E}_\theta [U_{n,\tau}(h) \mid Y_{1:n}] \right] - \sqrt{\frac{\log(J!)}{2n}} \end{aligned}$$

by Lemma 5 below.

When $J > 2$, Lemma 7 can be used to find that

$$\mathbb{P}_\theta (N_{(1)} + N_{(2)} < 2n\eta) \leq J^2 e^{-\frac{n(\beta - 2\eta)^2}{2\beta}}$$

and the bound is obtained by choosing $\eta = \beta/4$ and $\varepsilon = \frac{\beta}{4e} [\log(J!)/(2n)]^{2/(n\beta)}$.

LEMMA 5. For all $\theta \in \Theta^{\text{ind}}$, \mathbb{P}_θ -almost-surely

$$\mathbb{E}_\theta \left[\max_{\tau} (-U_{n,\tau}(h) + \hat{p}_\tau(h)) \mid Y_{1:n} \right] \leq \sqrt{\frac{\log(J!)}{2n}}.$$

PROOF. For any $\lambda > 0$,

$$\begin{aligned} \mathbb{E}_\theta \left[\max_{\tau} \{-U_{n,\tau}(h) + \hat{p}_\tau(h)\} \right] &\leq \frac{1}{\lambda} \log \left(\mathbb{E}_\theta \left[\exp \left(\sup_{\tau} \{-\lambda(U_{n,\tau}(h) - \hat{p}_\tau(h))\} \right) \mid Y_{1:n} \right] \right) \\ &= \frac{1}{\lambda} \log \left(\mathbb{E}_\theta \left[\sup_{\tau} \exp(-\lambda(U_{n,\tau}(h) - \hat{p}_\tau(h))) \mid Y_{1:n} \right] \right) \\ &\leq \frac{1}{\lambda} \log \left(J! \sup_{\tau} \mathbb{E}_\theta \left[\exp(-\lambda(U_{n,\tau}(h) - \hat{p}_\tau(h))) \mid Y_{1:n} \right] \right) \\ &\leq \frac{1}{\lambda} \log \left(J! \exp \left(\frac{\lambda^2}{8} \times n \times \left(\frac{1}{n} \right)^2 \right) \right) \text{ (Hoeffding's lemma)} \\ &\leq \inf_{\lambda > 0} \left\{ \frac{\log(J!)}{\lambda} + \frac{\lambda}{8n} \right\} \\ &\leq \sqrt{\frac{\log(J!)}{2n}}. \end{aligned}$$

Hoeffding's lemma applies because conditionally to the sequence of observations $Y_{1:n}$, the labels $X_{1:n}$ are still independent. \square

LEMMA 6. For all $\theta \in \Theta^{\text{ind}}$, \mathbb{P}_θ -almost-surely

$$\mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) \leq \hat{p}_{\hat{\tau}_h}(h) \cdot \frac{e}{\eta} \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta-1} e^{-n\hat{p}_{\hat{\tau}_h}(h)}.$$

PROOF. By Chernoff's bound (with $q_i(h) = \mathbb{P}_\theta(h_i(Y_{1:n}) \neq \hat{\tau}_h(X_i) \mid Y_{1:n})$):

$$\begin{aligned} \mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) &= \mathbb{P}_\theta \left(\sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \hat{\tau}_h(X_i)} > n\eta \mid Y_{1:n} \right) \\ &\leq \inf_{\lambda > 0} \exp \left(-\lambda n\eta + \sum_{i=1}^n \log(q_i(h)e^\lambda + 1 - q_i(h)) \right) \\ &\leq \inf_{\lambda > 0} \exp \left(-\lambda n\eta + n\hat{p}_{\hat{\tau}_h}(h)(e^\lambda - 1) \right) \\ &= \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta} e^{-n\hat{p}_{\hat{\tau}_h}(h)} \\ &\leq \hat{p}_{\hat{\tau}_h}(h) \cdot \frac{e}{\eta} \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta-1} e^{-n\hat{p}_{\hat{\tau}_h}(h)}. \end{aligned}$$

\square

LEMMA 7. For $\theta \in \Theta^{\text{ind}}$, let $\beta = \min_{j \neq k}(\nu_j + \nu_k)$. If $J > 2$, then

$$\mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta) \leq J^2 e^{-\frac{n(\beta-2\eta)^2}{2\beta}}.$$

PROOF. If $J = 2$, remark that $N_{(1)} + N_{(2)} = n$. Now we assume that $J > 2$. It holds that

$$\begin{aligned} \mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta) &= P_\theta(\exists j \neq k, N_j + N_k \leq 2n\eta) \\ &\leq J^2 \max_{j \neq k} \mathbb{P}_\theta(N_j + N_k \leq 2n\eta). \end{aligned}$$

Then observe that $N_j + N_k = \sum_{i=1}^n (\mathbf{1}_{X_i=j} + \mathbf{1}_{X_i=k}) = \sum_{i=1}^n \mathbf{1}_{X_i \in \{j,k\}}$ whenever $j \neq k$. In other words, when $j \neq k$ the random variables $N_j + N_k$ has a Binomial distribution with parameters $(n, \nu_j + \nu_k)$ under \mathbb{P}_θ . The conclusion follows using Chernoff's bound on the Binomial distribution (recall the Binomial distribution is subGaussian on the left-tail). \square

6.6. Proof of Theorems 6 and 8 (dependent scenario).

6.6.1. *Preliminary.* We first recall basic results for HMMs that can be found in [8] about the distribution of the hidden states given a set of observations. For any parameter θ , any integers $k, i \leq j$, the distribution of X_k given $Y_{i:j}$ under \mathbb{P}_θ will be denoted $\phi_{\theta,k|i:j}(\cdot, Y_{i:j})$. For any integers $i \leq n$, we shall simplify the so-called filtering distribution $\phi_{\theta,n|i:n}(\cdot, Y_{i:n})$ to $\phi_{\theta,n}(\cdot, Y_{i:n})$.

Conditional on observations $Y_{i:n}$, the sequence of the hidden states is an inhomogeneous Markov chain, with transition matrices called *forward kernels*. For each $k \leq n-1$, the forward kernel is denoted $(F_{\theta,k|n}[Y_{k+1:n}])$ to emphasize that it only depends on $Y_{k+1:n}$. When $k \geq n$, the kernel does not depend on the observations and is equal to the transition matrix Q , so that $F_{\theta,k|n}[Y_{k+1:n}] := Q$ for $k \geq n$. In other words, for any $n \in \mathbb{N}$, for any index $i \leq n$ and $k \geq i$ and any real-valued function f on \mathbb{X} (understood as a vector in \mathbb{R}^J),

$$\mathbb{E}_\theta[f(X_{k+1}) | X_{i:k}, Y_{i:n}] = F_{\theta,k|n}[Y_{k+1:n}]f = \sum_{x \in \mathbb{X}} F_{\theta,k|n}[Y_{k+1:n}](X_k, x)f(x).$$

Conditional on observations $Y_{i:n}$, the reverse time sequence of hidden states is also an inhomogeneous Markov chain with transition matrices $(B_{\theta,k}[Y_{i:k}])_{k \leq n-1}$ called *backward kernels*. In other words, for any $n \in \mathbb{N}$, $i \leq k \leq n-1$ and any function f on \mathbb{X} :

$$\mathbb{E}_\theta[f(X_k) | X_{k+1:n}, Y_{i:n}] = B_{\theta,k}[Y_{i:k}]f = \sum_{x \in \mathbb{X}} B_{\theta,k}[Y_{i:k}](X_{k+1}, x)f(x).$$

Here, the backward kernel $B_{\theta,k}[Y_{i:k}]$ depends only on the observations up to time k . It is given by:

$$(17) \quad B_{\theta,k}[Y_{i:k}](\tilde{x}, x) = \frac{\phi_{\theta,k}(x, Y_{i:k})Q(x, \tilde{x})}{\sum_{x' \in \mathbb{X}} \phi_{\theta,k}(x', Y_{i:k})Q(x', \tilde{x})}.$$

Note that the denominator is always positive thanks to Assumption 1.

For any transition kernel T , we denote $\delta(T)$ is the Dobrushin coefficient of T defined by:

$$\delta(T) = \sup_{(x, x') \in \mathbb{X} \times \mathbb{X}} \|T(x, \cdot) - T(x', \cdot)\|_{\text{TV}}$$

where $\|\cdot\|_{\text{TV}}$ is the total variation norm. We recall the following two lemmas which can be found in [8]. To end with, notice that under Assumption 1, for any subset A of \mathbb{X} ,

$$J\delta\gamma(A) \leq \sum_{x' \in A} Q(x, x') \leq J(1 - (J-1)\delta)\gamma(A)$$

with γ the uniform distribution over \mathbb{X} . Using Lemma 4.3.13 in [8], this leads to the following lemma.

LEMMA 8. *Under Assumption 1, for any integers k and n , the Dobrushin coefficient of the forward kernel $F_{\theta,k|n}$ satisfies:*

$$\delta(F_{\theta,k|n}) \leq \begin{cases} \rho_0 & k < n \\ \rho_1 & k \geq n \end{cases}$$

with $\rho_0 = 1 - \frac{J\delta}{J(1-(J-1)\delta)} = \frac{1-J\delta}{1-(J-1)\delta}$ and $\rho_1 = 1 - J\delta$.

Using Equation (17) we get that, under Assumption 1, for any (possibly non positive) integers $i \leq k$,

$$\forall x \in \mathbb{X}, \quad B_{\theta,k}[Y_{i:k}](x, \cdot) \geq \frac{\delta}{1 - (J-1)\delta} \phi_{\theta,k}(\cdot, Y_{i:k}),$$

so that applying Lemma 4.3.13 in [8], one gets

LEMMA 9. *Under Assumption 1, for any (possibly non positive) integers $i \leq k \leq n-1$, the Dobrushin coefficient of the backward kernel $B_{\theta,k}[Y_{i:k}]$ satisfies:*

$$\delta(B_{\theta,k}[Y_{i:k}]) \leq 1 - \frac{\delta}{1 - (J-1)\delta} = \rho_0.$$

6.6.2. *Proofs.* We apply the result of Proposition 2 to the HMM case. As in the i.i.d. case, when $J = 2$ it must be that $\mathbb{P}_\theta(N_{(1)} + N_{(2)} < n) = 0$. On the one hand, using Lemma 10 and Markov's inequality as in the independent case, one gets for $\theta \in \Theta^{\text{dep}}$

$$\begin{aligned} \mathbb{E}_\theta \left[\mathbb{E}_\theta \left[\max_{\tau} (-U_{n,\tau}(h) + \hat{p}_\tau(h)) \mid Y_{1:n} \right] \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}} \right] \\ \leq \frac{1}{1 - \rho_0} \sqrt{\frac{\log(J!)}{2n}} \mathbb{P}_\theta(\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon) \\ \leq \frac{1}{\varepsilon(1 - \rho_0)} \sqrt{\frac{\log(J!)}{2n}} \mathbb{E}_\theta \left[\min_{\tau} \mathbb{E}_\theta[U_{n,\tau}(h) \mid Y_{1:n}] \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon\}} \right] \end{aligned}$$

and then using Lemma 11, we establish that for all $\varepsilon, \eta > 0$

$$\begin{aligned} \mathbb{E}_\theta \left[\mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}} \right] \\ \leq \mathbb{E}_\theta \left[\hat{p}_{\hat{\tau}_h}(h) \cdot \frac{e}{\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta-1} e^{-n\hat{p}_{\hat{\tau}_h}(h)} \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}} \right] \\ \leq \frac{e}{\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} \left(\frac{e\varepsilon}{\eta} \right)^{n\eta-1} \mathbb{E}_\theta \left[\min_{\tau} \mathbb{E}_\theta[U_{n,\tau}(h) \mid Y_{1:n}] \mathbf{1}_{\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}} \right]. \end{aligned}$$

Then, the bound in this case follows by taking $\eta \rightarrow \frac{1}{2}$ (by below) and $\varepsilon = \frac{1}{2e} \left(\frac{\delta}{1-\delta} \right)^4 [\log(J!)/(2n)]^{1/n}$.

When $J > 2$, the first trivial bound is obtained by choosing $\varepsilon = \eta = 0$. Proposition 2 yields:

$$\begin{aligned} \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \right] &\geq \mathbb{E}_\theta \left[\min_{\tau} \mathbb{E}_\theta(U_{n,\tau}(h) \mid Y_{1:n}) \right] - \mathbb{E}_\theta \left[\mathbb{E}_\theta \left[\max_{\tau} (\hat{p}_\tau(h) - U_{n,\tau}(h)) \mid Y_{1:n} \right] \right] \\ &\geq \mathbb{E}_\theta \left[\min_{\tau} \mathbb{E}_\theta(U_{n,\tau}(h) \mid Y_{1:n}) \right] - \frac{1}{1 - \rho_0} \sqrt{\frac{\log(J!)}{2n}} \end{aligned}$$

by Lemma 10 below. The inequality follows by taking the infimum over h on both sides.

For the remaining inequality, when $J > 2$, Lemma 12 ensures:

$$\mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta) \leq J^2 e^{-2n(1-\rho_1)^2(\beta-2\eta)^2}.$$

On the other hand, we have by Lemma 10:

$$\mathbb{E}_\theta \left[\max_{\tau} (\hat{p}_\tau(h) - U_{n,\tau}(h)) \mid Y_{1:n} \right] \leq \frac{1}{1-\rho_0} \sqrt{\frac{\log(J!)}{2n}}$$

On the event $\{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon\}$, for $\varepsilon < \eta$:

$$\begin{aligned} \mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) &= \mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) - \hat{p}_{\hat{\tau}_h}(h) > \eta - \hat{p}_{\hat{\tau}_h}(h) \mid Y_{1:n}) \\ &\leq \mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) - \hat{p}_{\hat{\tau}_h}(h) > \eta - \varepsilon \mid Y_{1:n}) \\ &\leq e^{-\lambda(\eta-\varepsilon)} \mathbb{E}_\theta \left[e^{\lambda\{U_{n,\hat{\tau}_h}(h) - \hat{p}_{\hat{\tau}_h}(h)\}} \mid Y_{1:n} \right] \\ &\leq e^{-\lambda(\eta-\varepsilon)} e^{\frac{1}{8n} \left(\frac{\lambda}{1-\rho_0} \right)^2} \end{aligned}$$

where the last inequality is due to the argument using Marton coupling as shown in the proof of Lemma 10. Taking the minimum over λ , one gets

$$\mathbb{E}_\theta \left[\mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) \mathbf{1}_{\hat{p}_{\hat{\tau}_h}(h) < \varepsilon} \right] \leq e^{-2n(1-\rho_0)^2(\eta-\varepsilon)^2}.$$

Using Lemma 12, the final bound reads:

$$\begin{aligned} \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \right] &\geq \mathbb{E}_\theta \left[\min_{\tau} \mathbb{E}_\theta[U_{n,\tau}(h) \mid Y_{1:n}] \right] \\ &\quad - \frac{1}{\varepsilon(1-\rho_0)} \sqrt{\frac{\log(J!)}{2n}} \mathbb{E}_\theta \left[\hat{p}_{\hat{\tau}_h}(h) \mathbf{1}_{\hat{p}_{\hat{\tau}_h}(h) \geq \varepsilon} \right] \\ &\quad - e^{-2n(1-\rho_0)^2(\eta-\varepsilon)^2} - J^2 e^{-2n(1-\rho_1)^2(\beta-2\eta)^2} \end{aligned}$$

Choosing $\varepsilon = \frac{\eta}{2}$ and $\eta = \frac{2}{5}\beta$ and noting that $\rho_1 < \rho_0$ one obtains:

$$\begin{aligned} \mathbb{E}_\theta \left[\min_{\tau} U_{n,\tau}(h) \right] &\geq \left[1 - \frac{5}{\beta(1-\rho_0)} \sqrt{\frac{\log(J!)}{2n}} \right] \mathbb{E}_\theta \left[\min_{\tau} \mathbb{E}_\theta[U_{n,\tau}(h) \mid Y_{1:n}] \right] \\ &\quad - (J^2 + 1)e^{-cn(1-\rho_0)^2\beta^2}. \end{aligned}$$

The result follows by taking the infimum over h .

LEMMA 10. Under Assumptions 1 and 2, for all $\theta \in \Theta^{\text{dep}}$, \mathbb{P}_θ -almost-surely

$$\mathbb{E}_\theta \left[\max_{\tau} (\hat{p}_\tau(h) - U_{n,\tau}(h)) \mid Y_{1:n} \right] \leq \frac{1}{1-\rho_0} \sqrt{\frac{\log(J!)}{2n}}$$

where $\rho_0 = \frac{1-J\delta}{1-(J-1)\delta}$.

PROOF. Given that for any $\lambda > 0$,

$$\mathbb{E}_\theta \left[\max_{\tau} (\hat{p}_\tau(h) - U_{n,\tau}(h)) \right] \leq \frac{1}{\lambda} \log \left(J! \max_{\tau} \mathbb{E}_\theta \left[\exp(-\lambda(U_{n,\tau}(h) - \hat{p}_\tau(h))) \mid Y_{1:n} \right] \right)$$

we shall exhibit an upper bound of the rhs term by applying Theorem 2.9 of [34], conditional on $Y_{1:n}$. So that for now we consider $Y = Y_{1:n}$ as fixed. Define $f_Y^{h,\tau}$ for any $x = x_{1:n}$ by:

$$f_Y^{h,\tau}(x) = -\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\tau(x_i) \neq h_i(Y_{1:n})}.$$

Then, for any $x = x_{1:n}$ and $x' = x'_{1:n}$,

$$f_Y^{h,\tau}(x) - f_Y^{h,\tau}(x') \leq \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{x_i \neq x'_i}.$$

We thus may apply (2.5) in Theorem 2.9 of [34] to get

$$(18) \quad \mathbb{E}_\theta \left[e^{\lambda(f_Y^{h,\tau}(X) - \mathbb{E}_\theta[f_Y^{h,\tau}(X)|Y_{1:n}])} \middle| Y_{1:n} \right] \leq e^{\frac{\lambda^2}{8n^2} \sum_{i=1}^n (\sum_{j=i}^n \Gamma_{i,j})^2},$$

where Γ comes from a Marton coupling (see Definition 2.1 in [34]) and is given by:

$$\Gamma_{j,i} := 0, \quad \Gamma_{i,j} := \sup_{x_1, \dots, x_i, x'_i \in \mathbb{X}} \mathbb{P}_\theta \left(X_{1,j}^{(x_1, \dots, x_i, x'_i, Y_{1:n})} \neq X_{2,j}^{(x_1, \dots, x_i, x'_i, Y_{1:n})} \middle| Y_{1:n} \right)$$

for $1 \leq i < j \leq n$. Now,

$$\begin{aligned} \Gamma_{i,j} \leq \sup_{x_{1:i-1} \in \mathbb{X}^{i-1}} & \left\| \mathbb{P}_\theta \left(X_{1,j}^{(x_{1:i-1}, \tilde{a}, \tilde{b}, Y_{1:n})} \in \cdot \middle| X_{1:i-1} = x_{1:i-1}, Y_{1:n} \right) \right. \\ & \left. - \mathbb{P}_\theta \left(X_{2,j}^{(x_{1:i-1}, \tilde{a}, \tilde{b}, Y_{1:n})} \in \cdot \middle| X_{1:i-1} = x_{1:i-1}, Y_{1:n} \right) \right\|_{\text{TV}}, \end{aligned}$$

ie.,

$$\begin{aligned} \Gamma_{i,j} \leq \sup_{x_{1:i-1} \in \mathbb{X}^{i-1}} & \left\| \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_{1:i-1} = x_{1:i-1}, X_i = \tilde{a}, Y_{1:n} \right) \right. \\ & \left. - \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_{1:i-1} = x_{1:i-1}, X_i = \tilde{b}, Y_{1:n} \right) \right\|_{\text{TV}}, \end{aligned}$$

ie.,

$$\Gamma_{i,j} \leq \left\| \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_i = \tilde{a}, Y_{1:n} \right) - \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_i = \tilde{b}, Y_{1:n} \right) \right\|_{\text{TV}}$$

since conditional on $Y_{1:n}$, the hidden states form an inhomogeneous Markov chain with transition kernels $(F_{k|n}[Y_{k+1:n}])$. Exponential forgetting of the smoothing distributions in HMMs (Proposition 4.3.26 in [8]) allows to conclude that

$$\left\| \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_i = \tilde{a}, Y_{1:n} \right) - \mathbb{P}_\theta \left(X_j \in \cdot \middle| X_i = \tilde{b}, Y_{1:n} \right) \right\|_{\text{TV}} \leq \rho_0^{j-i}$$

where $\rho_0 = \frac{1-J\delta}{1-(J-1)\delta}$ (see also Lemma 8). By inequality (18):

$$\max_{\tau} \mathbb{E}_\theta \left[e^{\lambda(f_Y^{h,\tau}(X) - \mathbb{E}_\theta[f_Y^{h,\tau}(X)|Y_{1:n}])} \middle| Y_{1:n} \right] \leq e^{\frac{1}{8n} \left(\frac{\lambda}{1-\rho_0} \right)^2}.$$

Thus,

$$\begin{aligned} \mathbb{E}_\theta \left[\max_{\tau} (\hat{p}_\tau(h) - U_{n,\tau}(h)) \middle| Y_{1:n} \right] & \leq \inf_{\lambda > 0} \left\{ \frac{\log(J!)}{\lambda} + \frac{\lambda}{8n(1-\rho_0)^2} \right\} \\ & = \frac{1}{1-\rho_0} \sqrt{\frac{\log(J!)}{2n}} \end{aligned}$$

□

LEMMA 11. For all $\theta \in \Theta^{\text{dep}}$, \mathbb{P}_θ -almost-surely

$$\mathbb{P}_\theta(U_{n,\hat{\tau}_h}(h) > \eta \mid Y_{1:n}) \leq \hat{p}_{\hat{\tau}_h}(h) \cdot \frac{e}{\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta-1} e^{-n\hat{p}_{\hat{\tau}_h}(h)}.$$

PROOF. Let $S = \sum_{i=1}^n \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}$. We consider the following operators defined on $L^\infty(\{0,1\})$, for $i \in [n]$:

$$(M_{\theta,i} \cdot f)(x) := f(0)e^{\lambda \mathbf{1}_{\hat{\tau}_h(0) \neq h_i(Y_{1:n})}} B_{\theta,i}(x, 0) + f(1)e^{\lambda \mathbf{1}_{\hat{\tau}_h(1) \neq h_i(Y_{1:n})}} B_{\theta,i}(x, 1)$$

where $B_{\theta,i}$ is the Backward kernel defined by:

$$B_{\theta,i}(x, y) = \frac{\phi_{\theta,i}(y)Q(y, x)}{\sum_{y'} \phi_{\theta,i}(y')Q(y', x)}.$$

Then observe that,

$$\begin{aligned} \mathbb{E}_\theta[e^{\lambda S} \mid Y_{1:n}] &= \mathbb{E}_\theta \left[e^{\lambda \sum_{i=2}^n \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}} \mathbb{E}_\theta \left[e^{\lambda \mathbf{1}_{\hat{\tau}_h(X_1) \neq h_1(Y_{1:n})}} \mid X_{2:n}, Y_{1:n} \right] \mid Y_{1:n} \right] \\ &= \mathbb{E}_\theta \left[e^{\lambda \sum_{i=2}^n \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}} (M_{\theta,1} \cdot \mathbf{1})(X_2) \mid Y_{1:n} \right] \end{aligned}$$

Repeating inductively the same trick leads to

$$\mathbb{E}_\theta[e^{\lambda S} \mid Y_{1:n}] = \mathbb{E}_\theta[(M_{\theta,n} \dots M_{\theta,1} \cdot \mathbf{1})(X_{n+1}) \mid Y_{1:n}]$$

Hence

$$\mathbb{E}_\theta[e^{\lambda S} \mid Y_{1:n}] \leq \|(M_{\theta,n} \dots M_{\theta,1} \cdot \mathbf{1})\|_\infty \leq \prod_{i=1}^n \|M_{\theta,i}\|_\infty$$

where

$$\begin{aligned} \|M_{\theta,i}\|_\infty &:= \sup_{f, \|f\|_\infty=1} \|M_{\theta,i} \cdot f\|_\infty \\ &= \max((M_{\theta,i} \cdot f)(0), (M_{\theta,i} \cdot f)(1)) \\ &\leq \max \left(\sum_{z \in \{0,1\}} e^{\lambda \mathbf{1}_{\hat{\tau}_h(z) \neq h_i(Y_{1:n})}} B_{\theta,i}(0, z), \sum_{z \in \{0,1\}} e^{\lambda \mathbf{1}_{\hat{\tau}_h(z) \neq h_i(Y_{1:n})}} B_{\theta,i}(1, z) \right). \end{aligned}$$

But, given that

$$B_{\theta,i}(x, y) = \frac{\phi_{\theta,i}(y)Q(y, x)}{\sum_{y'} \phi_{\theta,i}(y')Q(y', x)} \leq \frac{1 - (J-1)\delta}{\delta} \phi_{\theta,i}(y)$$

and that

$$\phi_{\theta,i|n}(y) = \phi_{\theta,i+1|n} B_{\theta,i}(y) = \frac{\sum_x \phi_{\theta,i+1|n}(x) \phi_{\theta,i}(y) Q(y, x)}{\sum_{y'} \phi_{\theta,i}(y') Q(y', x)} \geq \frac{\delta \phi_{\theta,i}(y)}{1 - (J-1)\delta}$$

one obtains

$$B_{\theta,i}(x, y) \leq \left(\frac{1 - (J-1)\delta}{\delta} \right)^2 \phi_{\theta,i|n}(y).$$

Thus,

$$\begin{aligned} \|M_{\theta,i}\|_\infty &\leq \left(\frac{1 - (J-1)\delta}{\delta} \right)^2 \left(e^{\lambda \mathbf{1}_{\hat{\tau}_h(0) \neq h_i(Y_{1:n})}} \phi_{i|n}(0) + e^{\lambda \mathbf{1}_{\hat{\tau}_h(1) \neq h_i(Y_{1:n})}} \phi_{i|n}(1) \right) \\ &= \left(\frac{1 - (J-1)\delta}{\delta} \right)^2 \mathbb{E}_\theta \left[e^{\lambda \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}} \mid Y_{1:n} \right]. \end{aligned}$$

Finally,

$$\mathbb{E}_\theta[\exp(\lambda S) \mid Y_{1:n}] \leq \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} \prod_{i=1}^n \mathbb{E}_\theta \left[e^{\lambda \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}} \mid Y_{1:n} \right].$$

One can then use Chernoff's bound (with $q_i(h) = \mathbb{P}_\theta(h_i(Y_{1:n}) \neq \hat{\tau}_h(X_i) \mid Y_{1:n})$):

$$\begin{aligned} \mathbb{P}_\theta(U_{n, \hat{\tau}_h}(h) > \eta \mid Y_{1:n}) &= \mathbb{P}_\theta \left(\sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \hat{\tau}_h(X_i)} > n\eta \mid Y_{1:n} \right) \\ &\leq \inf_{\lambda > 0} e^{-\lambda n\eta} \mathbb{E}_\theta \left[e^{\lambda S} \mid Y_{1:n} \right] \\ &\leq \inf_{\lambda > 0} e^{-\lambda n\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} \prod_{i=1}^n \mathbb{E}_\theta \left[e^{\lambda \mathbf{1}_{\hat{\tau}_h(X_i) \neq h_i(Y_{1:n})}} \mid Y_{1:n} \right] \\ &\leq \inf_{\lambda > 0} e^{-\lambda n\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} e^{\sum_{i=1}^n \log(q_i(h) e^{\lambda} + 1 - q_i(h))} \\ &\leq \inf_{\lambda > 0} e^{-\lambda n\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} e^{n\hat{p}_{\hat{\tau}_h}(h)(e^\lambda - 1)} \\ &\leq \hat{p}_{\hat{\tau}_h}(h) \cdot \frac{e}{\eta} \left(\frac{1 - (J-1)\delta}{\delta} \right)^{2n} \left(\frac{e\hat{p}_{\hat{\tau}_h}(h)}{\eta} \right)^{n\eta-1} e^{-n\hat{p}_{\hat{\tau}_h}(h)} \end{aligned}$$

□

LEMMA 12. Let $\beta = \min_{i,j \neq k} \mathbb{P}_\theta(X_i \in \{j, k\})$ and $\rho_1 = 1 - J\delta$. If $J \geq 3$, then for $\eta < \frac{\beta}{2}$ and $\theta \in \Theta^{\text{dep}}$

$$\mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta) \leq J^2 e^{-2n(1-\rho_1)^2(\beta-2\eta)^2}.$$

PROOF. Let $\lambda > 0$, $\eta < \frac{\beta}{2}$, $j \neq k$, $\beta_i(j, k) = \mathbb{P}_\theta(X_i \in \{j, k\})$.

$$\begin{aligned} \mathbb{P}_\theta(N_j + N_k < 2n\eta) &= \mathbb{P}_\theta \left(\sum_{i=1}^n (\mathbf{1}_{X_i \in \{j, k\}} - \beta_i(j, k)) < 2n\eta - \sum_{i=1}^n \beta_i(j, k) \right) \\ &= \mathbb{P}_\theta \left(e^{\lambda \sum_{i=1}^n (\beta_i(j, k) - \mathbf{1}_{X_i \in \{j, k\}})} > e^{\lambda (\sum_{i=1}^n \beta_i(j, k) - 2n\eta)} \right) \\ &\leq e^{-\lambda (\sum_{i=1}^n \beta_i(j, k) - 2n\eta)} \mathbb{E}_\theta \left[e^{\lambda \sum_{i=1}^n (\beta_i(j, k) - \mathbf{1}_{X_i \in \{j, k\}})} \right] \end{aligned}$$

$\mathbb{E}_\theta \left[e^{\lambda \sum_{i=1}^n (\beta_i(j, k) - \mathbf{1}_{X_i \in \{j, k\}})} \right]$ can be controlled by the same technique using the Marton coupling that was used in the proof of Lemma 10:

$$\mathbb{E}_\theta \left[e^{\lambda \sum_{i=1}^n (\beta_i(j, k) - \mathbf{1}_{X_i \in \{j, k\}})} \right] \leq e^{\frac{(n\lambda)^2}{8n(1-\rho_1)^2}} = e^{\frac{n\lambda^2}{8(1-\rho_1)^2}}$$

It follows that

$$\begin{aligned} \mathbb{P}_\theta(N_j + N_k < 2n\eta) &\leq \inf_{\lambda > 0} e^{-\lambda (\sum_{i=1}^n \beta_i(j, k) - 2n\eta) + \frac{n\lambda^2}{8(1-\rho_1)^2}} \\ &\leq e^{-2n(1-\rho_1)^2 \left(\frac{\sum_{i=1}^n \beta_i(j, k)}{n} - 2\eta \right)^2} \\ &\leq e^{-2n(1-\rho_1)^2(\beta-2\eta)^2} \end{aligned}$$

Thus,

$$\begin{aligned}\mathbb{P}_\theta(N_{(1)} + N_{(2)} < 2n\eta) &\leq J^2 \max_{j \neq k} \mathbb{P}_\theta(N_j + N_k < 2n\eta) \\ &\leq J^2 e^{-2n(1-\rho_1)^2(\beta-2\eta)^2}.\end{aligned}$$

□

6.7. Proof of Theorem 5. Let $a \in \{1, 2\}$ and $n \in \mathbb{N}$. Let $Q = \begin{pmatrix} 1-p & p \\ q & 1-q \end{pmatrix}$ be the transition matrix and assume $0 < q < \frac{1}{2} < p$ and $p+q < 1$. Assume also that the initial distribution is the stationary distribution, that is $\nu = \left(\frac{q}{p+q}, \frac{p}{p+q}\right)$.

$$\begin{aligned}\mathbb{P}_\theta(X_1 = a \mid Y_{1:n}) &= \sum_{x_{2:n} \in \{1,2\}} \mathbb{P}_\theta(X_1 = a, X_{2:n} = x_{2:n} \mid Y_{1:n}) \\ &\propto \sum_{x_{2:n} \in \{1,2\}} \nu(a) Q_{a,x_2} \dots Q_{x_{n-1},x_n} f_a(Y_1) \dots f_{x_n}(Y_n) \\ \mathbb{P}_\theta(X_2 = a \mid Y_{1:n}) &= \sum_{x_1, x_{3:n} \in \{1,2\}} \mathbb{P}_\theta(X_1 = x_1, X_2 = a, X_{3:n} = x_{3:n} \mid Y_{1:n}) \\ &\propto \sum_{x_1, x_{3:n} \in \{1,2\}} \nu(x_1) Q_{x_1,a} \dots Q_{x_{n-1},x_n} f_{x_1}(Y_1) f_a(Y_2) \dots f_{x_n}(Y_n)\end{aligned}$$

The Bayes classifier puts the two first observations in the same cluster when:

$$\begin{aligned}&\left[\mathbb{P}_\theta(X_1 = 2 \mid Y_{1:n}) - \mathbb{P}_\theta(X_1 = 1 \mid Y_{1:n}) \right] \times \left[\mathbb{P}_\theta(X_2 = 2 \mid Y_{1:n}) - \mathbb{P}_\theta(X_2 = 1 \mid Y_{1:n}) \right] \geq 0 \\ &\iff \left[\sum_{x_{2:n} \in \{1,2\}} (\nu(2) Q_{2,x_2} f_2(Y_1) - \nu(1) Q_{1,x_2} f_1(Y_1)) Q_{x_2,x_3} \dots Q_{x_{n-1},x_n} f_{x_2}(Y_2) \dots f_{x_n}(Y_n) \right] \\ &\times \left[\sum_{x_1, x_{3:n} \in \{1,2\}} (Q_{x_1,2} Q_{2,x_3} f_2(Y_2) - Q_{x_1,1} Q_{1,x_3} f_1(Y_2)) \nu(x_1) Q_{x_3,x_4} \dots Q_{x_{n-1},x_n} f_{x_1}(Y_1) f_{x_3}(Y_3) \dots f_{x_n}(Y_n) \right] \geq 0\end{aligned}$$

A sufficient condition for this to be ensured is:

$$\left(\frac{f_1}{f_2}(Y_1) < \frac{p \min(q, 1-q)}{q \max(p, 1-p)} \text{ and } \frac{f_1}{f_2}(Y_2) < \frac{\min(p, 1-q) \min(q, 1-p)}{\max(q, 1-p) \max(p, 1-p)} \right)$$

or

$$\left(\frac{f_2}{f_1}(Y_1) < \frac{q \min(p, 1-p)}{p \max(q, 1-q)} \text{ and } \frac{f_2}{f_1}(Y_2) < \frac{\min(q, 1-p) \min(p, 1-p)}{\max(p, 1-q) \max(q, 1-q)} \right)$$

Since $q < \frac{1}{2} < p$ and $p+q < 1$, the condition simplifies to:

$$\left(\frac{f_1}{f_2}(Y_1) < 1 \text{ and } \frac{f_1}{f_2}(Y_2) < \frac{q}{1-p} \right) \text{ or } \left(\frac{f_2}{f_1}(Y_1) < \frac{q(1-p)}{p(1-q)} \text{ and } \frac{f_2}{f_1}(Y_2) < \frac{q(1-p)}{(1-q)^2} \right).$$

We consider the event:

$$A = \left\{ \frac{f_1}{f_2}(Y_1) < 1, \frac{f_1}{f_2}(Y_2) < \frac{q}{1-p} \right\} \cup \left\{ \frac{f_2}{f_1}(Y_1) < \frac{q(1-p)}{p(1-q)}, \frac{f_2}{f_1}(Y_2) < \frac{q(1-p)}{(1-q)^2} \right\}$$

In what follows, we seek a sufficient condition under which the Bayes clusterer puts the two first observations in two different clusters. The Bayes clusterer is a partition F_n^* that

minimizes $\mathbb{E}_\theta [\ell(\pi_n(X_{1:n}), F_n) \mid Y_{1:n}]$. Let $L(Y_{1:n})$ be the likelihood of the observations $Y_{1:n}$. Consider the event

$$B_n = \{(\forall i \in \llbracket 3, n \rrbracket) \quad Y_i \notin \text{Supp}(f_2)\}$$

Assume B_n has positive probability. Since the hidden Markov chain is mixing, this happens for example when f_1 and f_2 do not have the same support. On this event,

$$\begin{aligned} \mathbb{E}_\theta [\ell(\pi_n(X_{1:n}), F_n) \mid Y_{1:n}] &= \sum_{x_{1:n} \in \{1,2\}^n} \ell(\pi_n(x_{1:n}), F_n) \mathbb{P}_\theta(X_{1:n} = x_{1:n} \mid Y_{1:n}) \\ &= \sum_{x_{1:2} \in \{1,2\}^2} \ell(\pi_n((x_1, x_2, 1, \dots, 1), F_n) \mathbb{P}_\theta(X_{1:2} = x_{1:2}, X_{3:n} = 1 \mid Y_{1:n}) \\ &= \frac{1}{L(Y_{1:n})} \sum_{x_{1:2} \in \{1,2\}^2} \ell(\pi_n((x_1, x_2, 1, \dots, 1), F_n) \nu(x_1) Q_{x_1, x_2} Q_{x_2, 1} Q_{1,1}^{n-3} f_{x_1}(Y_1) f_{x_2}(Y_2) \prod_{i=3}^n f_1(Y_i) \\ &\propto \left(\ell(\pi_n((1, \dots, 1), F_n) \nu(1) Q_{1,1}^2 f_1(Y_1) f_1(Y_2) + \ell(\pi_n((2, 2, 1, \dots, 1), F_n) \nu(2) Q_{2,2} Q_{2,1} f_2(Y_1) f_2(Y_2) \right. \\ &\quad \left. + \ell(\pi_n((1, 2, 1, \dots, 1), F_n) \nu(1) Q_{1,2} Q_{2,1} f_1(Y_1) f_2(Y_2) + \ell(\pi_n((2, 1, \dots, 1), F_n) \nu(2) Q_{2,1} Q_{1,1} f_2(Y_1) f_1(Y_2) \right) \\ &\propto \left(\ell(\pi_n((1, \dots, 1), F_n) q(1-p)^2 f_1(Y_1) f_1(Y_2) + \ell(\pi_n((2, 2, 1, \dots, 1), F_n) pq(1-q) f_2(Y_1) f_2(Y_2) \right. \\ &\quad \left. + \ell(\pi_n((1, 2, 1, \dots, 1), F_n) pq^2 f_1(Y_1) f_2(Y_2) + \ell(\pi_n((2, 1, \dots, 1), F_n) pq(1-p) f_2(Y_1) f_1(Y_2) \right) \end{aligned}$$

For $n \geq 5$, F_n^* is necessarily of the form $F_n^* = \pi_n((y_1^*, y_2^*, 1, \dots, 1))$ with y_1^* and y_2^* in $\{1, 2\}$.

Thus,

$$F_n^* \in \arg \min \mathbb{E}_\theta [\ell(\pi_n(X_{1:n}), F_n) \mid Y_{1:n}] \iff (y_1^*, y_2^*) \in \arg \min H(y_1, y_2)$$

where

$$\begin{aligned} H(y_1, y_2) &= (y_1 + y_2 - 2)(1-p)^2 f_1(Y_1) f_1(Y_2) + (4 - y_1 - y_2)p(1-q) f_2(Y_1) f_2(Y_2) \\ &\quad + (1 + y_1 - y_2)pq f_1(Y_1) f_2(Y_2) + (1 - y_1 + y_2)p(1-p) f_2(Y_1) f_1(Y_2) \\ &= y_1 [(1-p)^2 f_1(Y_1) f_1(Y_2) + pq f_1(Y_1) f_2(Y_2) - p(1-q) f_2(Y_1) f_2(Y_2) - p(1-p) f_2(Y_1) f_1(Y_2)] \\ &\quad + y_2 [(1-p)^2 f_1(Y_1) f_1(Y_2) - pq f_1(Y_1) f_2(Y_2) - p(1-q) f_2(Y_1) f_2(Y_2) + p(1-p) f_2(Y_1) f_1(Y_2)] \\ &\quad - 2(1-p)^2 f_1(Y_1) f_1(Y_2) + 4p(1-q) f_2(Y_1) f_2(Y_2) + pq f_1(Y_1) f_2(Y_2) + p(1-p) f_2(Y_1) f_1(Y_2) \\ y_1^* \neq y_2^* &\iff [(1-p)^2 f_1(Y_1) f_1(Y_2) + pq f_1(Y_1) f_2(Y_2) - p(1-q) f_2(Y_1) f_2(Y_2) - p(1-p) f_2(Y_1) f_1(Y_2)] \\ &\quad \times [(1-p)^2 f_1(Y_1) f_1(Y_2) - pq f_1(Y_1) f_2(Y_2) - p(1-q) f_2(Y_1) f_2(Y_2) + p(1-p) f_2(Y_1) f_1(Y_2)] < 0 \\ &\iff |(1-p)^2 f_1(Y_1) f_1(Y_2) - p(1-q) f_2(Y_1) f_2(Y_2)| < p|q f_1(Y_1) f_2(Y_2) - (1-p) f_2(Y_1) f_1(Y_2)| \end{aligned}$$

Finally, consider the event:

$$C_n = \{ |(1-p)^2 f_1(Y_1) f_1(Y_2) - p(1-q) f_2(Y_1) f_2(Y_2)| < p|q f_1(Y_1) f_2(Y_2) - (1-p) f_2(Y_1) f_1(Y_2)| \}$$

We finally have:

$$A \cap B_n \cap C_n \subset \{g_\theta^*(Y_{1:n}) \neq \pi_n \circ h_\theta^*(Y_{1:n})\}.$$

By choosing appropriately the parameters p and q , one can ensure that $\mathbb{P}_\theta(A \cap B_n \cap C_n)$ is positive for many emission densities not having the same support. For example, one can

ensure that for $p = 0.58$, $q = 0.35$, $f_1(Y_1) = 5$, $f_2(Y_1) = 2.5$, $f_1(Y_2) = 1.8$, $f_2(Y_2) = 1.2$, both inequalities involved in the definition of event C_n and A are ensured. Choosing smooth densities having not exactly the same support, the event $A \cap B_n \cap C_n$ can be ensured to have positive probability. The counterexamples for $n \in \{3, 4\}$ can be proved as for the case $n = 2$ presented in Section 3.2.

6.8. *Proof of Theorem 7.* Let $n \in \mathbb{N}$ and $\theta^* = (\nu^*, Q^*, (f_x^*)_{x \in \mathbb{X}}) \in \Theta^{\text{ind}}$ such that the assumption of Theorem 3 is ensured. It follows that $\mathbb{P}_{\theta^*}(g_{\theta^*}^*(Y_{1:n}) \neq \pi_n \circ h_{\theta^*}^*(Y_{1:n})) > 0$. We also assume the emission densities $(f_x^*)_{x \in \mathbb{X}}$ to be uniformly continuous. We denote $\tilde{\Theta} \subset \Theta$ the subset of parameters of the form $\theta = (\nu, Q, f_1^*, \dots, f_J^*)$. Consider the two functions

$$\begin{aligned} H: \mathcal{H}_n \times \Theta \times \mathcal{Y}^n &\longrightarrow [0, 1] \\ (h, \theta, y_{1:n}) &\longmapsto \mathbb{E}_{\theta} \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq X_i} \middle| Y_{1:n} = y_{1:n} \right] \\ G: \mathcal{G}_n \times \Theta \times \mathcal{Y}^n &\longrightarrow [0, 1] \\ (g, \theta, y_{1:n}) &\longmapsto \mathbb{E}_{\theta} \left[\ell(\pi_n(X_{1:n}), g(Y_{1:n})) \middle| Y_{1:n} = y_{1:n} \right] \end{aligned}$$

By uniform continuity of $(f_x^*)_{x \in \mathbb{X}}$, it follows that for all $g \in \mathcal{G}_n$ and $h \in \mathcal{H}_n$, $H(h, \cdot, \cdot)$ and $G(g, \cdot, \cdot)$ are uniformly continuous. Since \mathcal{H}_n and \mathcal{G}_n are finite, there exists $\mathcal{V}(\theta^*) \subset \tilde{\Theta}$ a neighborhood of θ^* and A_n an open subset of \mathcal{Y}^n such that $\mathbb{P}_{\theta^*}(A_n) > 0$ and $(\forall \theta \in \mathcal{V}(\theta^*)) (\forall y_{1:n} \in A_n)$

$$\arg \min_h H(h, \theta, y_{1:n}) = \arg \min_h H(h, \theta^*, y_{1:n}) \text{ and } \arg \min_g G(g, \theta, y_{1:n}) = \arg \min_g G(g, \theta^*, y_{1:n})$$

Or equivalently $g_{\theta^*}^*(y_{1:n}) = g_{\theta}^*(y_{1:n})$ and $h_{\theta^*}^*(y_{1:n}) = h_{\theta}^*(y_{1:n})$. On the other hand, using exactly the same arguments as those of Theorem 3, one could also have chosen $\theta^* \in \Theta^{\text{ind}}$ such that not only $\mathbb{P}_{\theta^*}(g_{\theta^*}^*(Y_{1:n}) \neq \pi_n \circ h_{\theta^*}^*(Y_{1:n})) > 0$ but also $\mathbb{P}_{\theta^*}(\{g_{\theta^*}^*(Y_{1:n}) \neq \pi_n \circ h_{\theta^*}^*(Y_{1:n})\} \cap A_n) > 0$. It follows that

$$(\forall \theta \in \mathcal{V}(\theta^*)) \mathbb{P}_{\theta}(\{g_{\theta}^*(Y_{1:n}) \neq \pi_n \circ h_{\theta}^*(Y_{1:n})\} \cap A_n) = \mathbb{P}_{\theta}(\{g_{\theta^*}^*(Y_{1:n}) \neq \pi_n \circ h_{\theta^*}^*(Y_{1:n})\} \cap A_n)$$

which is positive when θ approaches θ^* by continuity of the map

$$\theta \mapsto \mathbb{P}_{\theta}(\{g_{\theta}^*(Y_{1:n}) \neq \pi_n \circ h_{\theta}^*(Y_{1:n})\} \cap A_n).$$

The result follows.

6.9. *Proof of Proposition 1.* The result is straightforward when $n = 1$. We assume in what follows $n \geq 2$. We prove the proposition by showing that when the probability of having small clusters is high, the two risks are not necessarily equivalent; and $\inf_{g \in \mathcal{G}_n} \mathcal{R}_n^{\text{clust}}(\theta, g)$ may be much smaller than $\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h)$.

Consider $J = 3$ (similar examples can be constructed for any $J \geq 3$) with $\nu_1 = 1 - 2\eta$, $\nu_2 = \nu_3 = \eta$. We take $F_1 = U(0, 1/2)$, $F_2 = U(3/4, 1)$ and $F_3 = U(3/4 - \varepsilon, 1 - \varepsilon)$ for some $0 < \varepsilon < 1/4$ where $U(a, b)$ is the uniform distribution on the interval (a, b) . In this case,

$$\begin{aligned} \mathbb{P}_{\theta}(X_i \in \cdot \mid Y_i \in (0, 1/2)) &= \delta_1(\cdot) \\ \mathbb{P}_{\theta}(X_i \in \cdot \mid Y_i \in (3/4 - \varepsilon, 3/4)) &= \delta_3(\cdot) \\ \mathbb{P}_{\theta}(X_i \in \cdot \mid Y_i \in (1 - \varepsilon, 1)) &= \delta_2(\cdot) \\ \mathbb{P}_{\theta}(X_i \in \cdot \mid Y_i \in (3/4, 1 - \varepsilon)) &= \frac{1}{2}\delta_2(\cdot) + \frac{1}{2}\delta_3(\cdot). \end{aligned}$$

So in this case

$$\begin{aligned}\mathbb{E}_\theta [U_{n,\tau}(h) \mid Y_{1:n}] &= \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \tau(X_i)} \mid Y_{1:n} \right] \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{P}_\theta(h_i(Y_{1:n}) \neq \tau(X_i) \mid Y_{1:n}).\end{aligned}$$

A Bayes “classifier” minimizing $h \mapsto \mathcal{R}_n^{\text{class}}(\theta, h)$ in this case is given by $h_\theta^* = (h_{\theta,i}^*)_{i \in [n]}$ where

$$h_{\theta,i}^* = \begin{cases} 1 & \text{if } Y_i \in (0, 1/2) \\ 2 & \text{if } Y_i \in (1 - \varepsilon, 1) \\ 3 & \text{if } Y_i \in (3/4 - \varepsilon, 1 - \varepsilon). \end{cases}$$

With this choice, the optimal permutation is identity and

$$\begin{aligned}\min_\tau \mathbb{E}_\theta [U_{n,\tau}(h_\theta^*) \mid Y_{1:n}] &= \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (0, 1/2)} \mathbb{P}_\theta(1 \neq X_i \mid Y_i) \\ &\quad + \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4 - \varepsilon, 3/4)} \mathbb{P}_\theta(3 \neq X_i \mid Y_i) \\ &\quad + \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (1 - \varepsilon, 1)} \mathbb{P}_\theta(2 \neq X_i \mid Y_i) \\ &\quad + \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1 - \varepsilon)} \mathbb{P}_\theta(3 \neq X_i \mid Y_i),\end{aligned}$$

ie.,

$$\min_\tau \mathbb{E}_\theta [U_{n,\tau}(h_\theta^*) \mid Y_{1:n}] = \frac{1}{2n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1 - \varepsilon)}.$$

Thus,

$$\begin{aligned}\mathbb{E}_\theta \left[\min_\tau \mathbb{E}_\theta (U_{n,\tau}(h_\theta^*) \mid Y_{1:n}) \right] &= \frac{1}{2} \mathbb{P}_\theta(Y_1 \in (3/4, 1 - \varepsilon)) \\ &= \frac{1}{2} (\eta \cdot (1 - 4\varepsilon) + \eta \cdot (1 - 4\varepsilon)) \\ &= \eta(1 - 4\varepsilon).\end{aligned}$$

We now investigate $\mathbb{E}_\theta [\min_\tau U_{n,\tau}(h_\theta^*)]$, for the previous Bayes classifier h_θ^* (which does not necessarily minimize $h \mapsto \mathbb{E}_\theta [\min_\tau U_{n,\tau}(h)]$). We rewrite,

$$\begin{aligned}\mathbb{E}_\theta \left[\min_\tau U_{n,\tau}(h_\theta^*) \right] &= \mathbb{E}_\theta \left[\min_\tau U_{n,\tau}(h_\theta^*) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4 - \varepsilon, 3/4) \cup (1 - \varepsilon, 1)} = 0} \right] \\ &\quad + \sum_{m=1}^n \mathbb{E}_\theta \left[\min_\tau U_{n,\tau}(h_\theta^*) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4 - \varepsilon, 3/4) \cup (1 - \varepsilon, 1)} = m} \right]\end{aligned}$$

Let first consider $\mathbb{E}_\theta [\min_\tau U_{n,\tau}(h_\theta^*) \mid Y_{1:n}]$ on the event that $\{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4 - \varepsilon, 3/4) \cup (1 - \varepsilon, 1)} = 0\}$. Let define $N = \sum_{i=1}^n \mathbf{1}_{X_i \neq 1}$:

- if $N = 0$ this means that all the Y_i are in $(0, 1/2)$ and our classifier h_θ^* combined with the identity permutation will make zero error, *i.e.* $\min_\tau U_{n,\tau}(h_\theta^*) = 0$ on this event;
- if $N = 1$ this means that there is only one $j \in \{1, \dots, n\}$ such that $Y_j \in (3/4, 1 - \varepsilon)$ [by assumption it can not be in $(3/4 - \varepsilon, 3/4)$ or $(1 - \varepsilon, 1)$]. Our classifier will predict $h_{\theta,i}^* = 1$ for all $i \neq j$ and $h_{\theta,j}^* = 3$. Now, necessarily $X_i = 1$ for $i \neq j$. If $X_j = 3$ then $h_\theta^* \circ \text{Id}$ will have loss zero, and if $X_j = 2$ then $h_\theta^* \circ \begin{pmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{pmatrix}$ will have loss zero. So in the event that $\{N = 1\}$ we also have that $\min_\tau U_{n,\tau}(h_\theta^*) = 0$.
- If $N \geq 2$ our classifier will still classify perfectly all the $Y_i \in (0, 1/2)$ so the loss can not exceed $\min_\tau U_{n,\tau}(h_\theta^*) \leq N/n$ in this case.

So on the event $\{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4 - \varepsilon, 3/4) \cup (1 - \varepsilon, 1)} = 0\}$:

$$\begin{aligned} \mathbb{E}_\theta \left[\min_\tau U_{n,\tau}(h_\theta^*) \mid Y_{1:n} \right] &\leq \mathbb{E}_\theta \left[\frac{N}{n} \mathbf{1}_{N \geq 2} \mid Y_{1:n} \right] \\ &\leq \sum_{k=2}^n \frac{k}{n} \mathbb{P}_\theta(N = k \mid Y_{1:n}) \end{aligned}$$

But under the law of $X_{1:n} \mid Y_{1:n}$ in the considered event, we have that N is almost-surely equal to the number of $Y_i \in (3/4, 1 - \varepsilon)$, so

$$\begin{aligned} \mathbb{E}_\theta \left[\min_\tau U_{n,\tau}(h_\theta^*) \mid Y_{1:n} \right] &\leq \sum_{k=2}^n \frac{k}{n} \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1 - \varepsilon)} = k} \\ &= \left(\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1 - \varepsilon)} \right) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1 - \varepsilon)} \geq 2}. \end{aligned}$$

Deduce that,

$$\begin{aligned} &\mathbb{E}_\theta \left[\min_\tau U_{n,\tau}(h_\theta^*) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4 - \varepsilon, 3/4) \cup (1 - \varepsilon, 1)} = 0} \right] \\ &\leq \mathbb{E}_\theta \left[\left(\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1 - \varepsilon)} \right) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1 - \varepsilon)} \geq 2} \mid \sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4 - \varepsilon, 3/4) \cup (1 - \varepsilon, 1)} = 0 \right] \end{aligned}$$

Conditional on $\{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4 - \varepsilon, 3/4) \cup (1 - \varepsilon, 1)} = 0\}$, the random variable $\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4, 1 - \varepsilon)}$ has a Binomial($n, 2\eta$) distribution. Then,

$$\begin{aligned} \mathbb{E}_\theta \left[\min_\tau U_{n,\tau}(h_\theta^*) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4 - \varepsilon, 3/4) \cup (1 - \varepsilon, 1)} = 0} \right] &\leq \frac{1}{n} \sum_{k=2}^n \binom{n}{k} k \cdot (2\eta)^k (1 - 2\eta)^{n-k} \\ &= \frac{2\eta(1 - (1 - 2\eta)^n - 2\eta)}{1 - 2\eta} \\ &\asymp 4(n - 1)\eta^2 \end{aligned}$$

when $\eta \ll 1/n$.

Next (remark that this can be largely improved, but this is indeed for our purpose),

$$\sum_{m=1}^n \mathbb{E}_\theta \left[\min_\tau U_{n,\tau}(h_\theta^*) \mathbf{1}_{\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4 - \varepsilon, 3/4) \cup (1 - \varepsilon, 1)} = m} \right]$$

$$\begin{aligned}
&\leq \mathbb{P}_\theta \left(\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} \geq 1 \right) \\
&= 1 - \mathbb{P}_\theta \left(\sum_{i=1}^n \mathbf{1}_{Y_i \in (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1)} = 0 \right) \\
&= 1 - \mathbb{P}_\theta \left(\forall i, Y_i \notin (3/4-\varepsilon, 3/4) \cup (1-\varepsilon, 1) \right) \\
&= 1 - \left((1-2\eta) + 2\eta(1-4\varepsilon) \right)^n \\
&= 1 - (1-8\eta\varepsilon)^n \\
&\asymp 8n\eta\varepsilon
\end{aligned}$$

when $\eta \ll n$. So by choosing $\varepsilon \asymp \eta$, we have shown that whenever $\eta \ll 1/n$

$$\inf_h \mathbb{E}_\theta \left[\min_\tau U_{n,\tau}(h) \right] \leq \mathbb{E}_\theta \left[\min_\tau U_{n,\tau}(h_\theta^*) \right] \lesssim n\eta^2$$

but

$$\inf_h \mathbb{E}_\theta \left[\min_\tau \mathbb{E}_\theta [U_{n,\tau}(h) \mid Y_{1:n}] \right] = \mathbb{E}_\theta \left[\min_\tau \mathbb{E}_\theta [U_{n,\tau}(h_\theta^*) \mid Y_{1:n}] \right] \sim \eta$$

so that

$$\frac{\inf_h \mathbb{E}_\theta [\min_\tau U_{n,\tau}(h)]}{\inf_h \mathbb{E}_\theta [\min_\tau \mathbb{E}_\theta [U_{n,\tau}(h) \mid Y_{1:n}]]} \lesssim n\eta$$

which goes to zero as $\eta \rightarrow 0$.

6.10. Proof of Theorem 9. Simple computations lead to the expression of the Bayes risk of classification:

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \mathbb{P}_\theta (X_i \neq x_0 \mid Y_{1:n}) \right] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\sum_{x \neq x_0} \phi_{\theta, i|n}(x) \right) \right].$$

6.10.1. Bounds for the independent scenario. First, for $\theta \in \Theta^{\text{ind}}$,

$$\begin{aligned}
\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) &= \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \mathbb{P}_\theta (X_1 \neq x_0 \mid Y_1) \right] \\
&= \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \frac{\sum_{x \neq x_0} \nu_x f_x(Y_1)}{\sum_x \nu_x f_x(Y_1)} \right] \\
&= \int \min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \nu_x f_x(y) d\mathcal{L}(y)
\end{aligned}$$

so that using Assumption 1,

$$\begin{aligned}
(19) \quad \delta \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y) &\leq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) \\
&\leq (1 - (J-1)\delta) \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y).
\end{aligned}$$

6.10.2. *Bounds for the dependent scenario.* Let $\theta \in \Theta^{\text{dep}}$. We first have

$$\inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta, i|1:n}(x) \right] \leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta, i}(x) \right].$$

Now, for any $i \leq n-1$, using the Backward recursions (see Proposition 3.3.9 in [8]),

$$\mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta, i|1:n}(x) \right] = \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \sum_{x' \in \mathbb{X}} \phi_{\theta, i+1|1:n}(x') B_{\theta, i}[Y_{1:i}](x', x) \right].$$

Now, using Assumption 1 and Equation (17), we get

$$B_{\theta, i}[Y_{1:i}](x', x) = \frac{\phi_{\theta, i}(x) Q(x, x')}{\sum_{\tilde{x} \in \mathbb{X}} \phi_{\theta, i}(\tilde{x}) Q(\tilde{x}, x')} \geq \frac{\phi_{\theta, i}(x) Q(x, x')}{1 - (J-1)\delta}$$

so that

$$\begin{aligned} \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta, i|1:n}(x) \right] &\geq \frac{\delta}{1 - (J-1)\delta} \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x'} \phi_{\theta, i+1}(x') \sum_{x \neq x_0} \phi_{\theta, i}(x) \right] \\ &\geq \frac{\delta}{1 - (J-1)\delta} \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta, i}(x) \right] \\ &= \frac{\delta}{1 - (J-1)\delta} \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta, i}(x) \right]. \end{aligned}$$

Now, for $i = n$, the same inequality obviously holds, so that we get

$$\frac{\delta}{1 - (J-1)\delta} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta, i}(x) \right] \leq \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) \leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta, i}(x) \right].$$

It suffices then to exhibit upper and lower bounds on $\sum_{i=1}^n \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \sum_{x \neq x_0} \phi_{\theta, i}(x) \right]$. Using the Forward recursions (see Equation (3.22) in Proposition 3.2.5 of [8]), for any $i \geq 2$,

$$\mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\sum_{x \neq x_0} \phi_{\theta, i}(x) \right) \right] = \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\frac{\sum_{x \neq x_0} \sum_{x' \in \mathbb{X}} Q_{x', x} \phi_{\theta, i-1}(x') f_x(Y_i)}{\sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} Q_{x', x} \phi_{\theta, i-1}(x') f_x(Y_i)} \right) \right].$$

Let $A \in \mathcal{Y}$. One has:

$$\begin{aligned}
\mathbb{P}_\theta(Y_i \in A \mid Y_{1:i-1}) &= \sum_{x \in \mathbb{X}} \mathbb{P}_\theta(Y_i \in A, X_i = x \mid Y_{1:i-1}) \\
&= \sum_{x \in \mathbb{X}} \mathbb{P}_\theta(Y_i \in A \mid X_i = x, Y_{1:i-1}) \mathbb{P}_\theta(X_i = x \mid Y_{1:i-1}) \\
&= \sum_{x \in \mathbb{X}} \mathbb{P}_\theta(X_i = x \mid Y_{1:i-1}) \int_A f_x(y) d\mathcal{L}(y) \\
&= \sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} \mathbb{P}_\theta(X_{i-1} = x', X_i = x \mid Y_{1:i-1}) \int_A f_x(y) d\mathcal{L}(y) \\
&= \sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} \mathbb{P}_\theta(X_i = x \mid X_{i-1} = x', Y_{1:i-1}) \phi_{\theta, i-1}(x') \int_A f_x(y) d\mathcal{L}(y) \\
&= \int_A \left(\sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} Q_{x', x} \phi_{\theta, i-1}(x') f_x(y) \right) d\mathcal{L}(y),
\end{aligned}$$

so that, conditionally on $Y_{1:i-1}$, Y_i has density $\sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} Q_{x', x} \phi_{\theta, i-1}(x') f_x$ with respect to the dominating measure \mathcal{L} . We thus get :

$$\begin{aligned}
\mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\sum_{x \neq x_0} \phi_{\theta, i}(x) \right) \right] &= \mathbb{E}_\theta \left[\mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\frac{\sum_{x \neq x_0} \sum_{x' \in \mathbb{X}} Q_{x', x} \phi_{\theta, i-1}(x') f_x(Y_i)}{\sum_{x \in \mathbb{X}} \sum_{x' \in \mathbb{X}} Q_{x', x} \phi_{\theta, i-1}(x') f_x(Y_i)} \right) \mid Y_{1:i-1} \right] \right] \\
&= \mathbb{E}_\theta \left[\int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left(\sum_{x \neq x_0} \sum_{x' \in \mathbb{X}} Q_{x', x} \phi_{\theta, i-1}(x') f_x(y) \right) d\mathcal{L}(y) \right].
\end{aligned}$$

Then, under Assumption 1, for any $i \geq 2$,

$$\begin{aligned}
(20) \quad \delta \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y) &\leq \mathbb{E}_\theta \left[\min_{x_0 \in \mathbb{X}} \left(\sum_{x \neq x_0} \phi_{\theta, i}(x) \right) \right] \\
&\leq (1 - (J-1)\delta) \int_{\mathbb{Y}} \min_{x_0 \in \mathbb{X}} \left[\sum_{x \neq x_0} f_x(y) \right] d\mathcal{L}(y),
\end{aligned}$$

The result follows.

6.11. *Proof of Theorem 10.* We first control the excess risk of classification by the errors made in the estimation of the parameters.

PROPOSITION 3. For all $\theta \in \Theta^{\text{dep}}$ satisfying assumptions 1, 2 and 4 and for all $n \geq 1$,

$$\begin{aligned}
\mathcal{R}_n^{\text{class}}(\theta, \hat{h}) - \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) &\leq C \mathbb{E}_\theta \left[\frac{1}{n} \|\nu - \hat{\nu}\|_2 + \left(1 + \frac{\delta}{\hat{\delta}} \right) \|Q - \hat{Q}\|_F \right] \\
&\quad + \frac{\delta C \sqrt{C^*}}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[(\hat{\rho} \vee \rho)^{2|l-i|} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2},
\end{aligned}$$

where $C = \frac{8(1-\delta)}{\delta^3}$, $\rho = \frac{1-2\delta}{1-\delta}$, $\hat{\rho} = \frac{1-2\hat{\delta}}{1-\hat{\delta}}$ and $\hat{\delta} = \min_{x, x'} \hat{Q}_{x, x'}$.

PROOF. Recall that $\phi_{\theta,i|n}(\cdot) = \mathbb{P}_\theta(X_i \in \cdot \mid Y_{1:n})$. Then,

$$\begin{aligned}
\mathcal{R}_n^{\text{class}}(\theta, \hat{h}) - \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) &= \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{X_i \neq \hat{h}_i} \right] - \mathbb{E}_\theta \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{X_i \neq h_{\theta,i}^*} \right] \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\mathbb{P}_\theta(X_i = h_{\theta,i}^* \mid Y_{1:n}) - \mathbb{P}_\theta(X_i = \hat{h}_i \mid Y_{1:n}) \right] \\
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\mathbb{P}_\theta(X_i = h_{\theta,i}^* \mid Y_{1:n}) - \mathbb{P}_{\hat{\theta}}(X_i = \hat{h}_i \mid Y_{1:n}) + \|\phi_{\hat{\theta},i|n} - \phi_{\theta,i|n}\|_{\text{TV}} \right] \\
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\mathbb{P}_\theta(X_i = h_{\theta,i}^* \mid Y_{1:n}) - \mathbb{P}_{\hat{\theta}}(X_i = h_{\theta,i}^* \mid Y_{1:n}) + \|\phi_{\hat{\theta},i|n} - \phi_{\theta,i|n}\|_{\text{TV}} \right] \\
&\leq \frac{2}{n} \sum_{i=1}^n \mathbb{E}_\theta \left[\|\phi_{\hat{\theta},i|n} - \phi_{\theta,i|n}\|_{\text{TV}} \right]
\end{aligned}$$

where the penultimate line follows because by definition \hat{h}_i maximizes $x \mapsto \mathbb{P}_{\hat{\theta}}(X_i = x \mid Y_{1:n})$. Then, under Assumption 4 and by application of Proposition 2.2 of [12] (Equation (10)), one has

$$\begin{aligned}
\mathcal{R}_n^{\text{class}}(\theta, \hat{h}) - \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) &\leq \frac{8(1-\delta)}{\delta^2} \mathbb{E}_\theta \left[\frac{1}{n\delta} \|\nu - \hat{\nu}\|_2 + \left(\frac{1}{\delta} + \frac{1}{\hat{\delta}} \right) \|Q - \hat{Q}\|_{\text{F}} + \frac{1}{n} \sum_{i=1}^n \sum_{l=1}^n \delta \frac{(\hat{\rho} \vee \rho)^{|l-i|}}{c^*(Y_l)} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty \right] \\
&\leq C \mathbb{E}_\theta \left[\frac{1}{n} \|\nu - \hat{\nu}\|_2 + \left(1 + \frac{\delta}{\hat{\delta}} \right) \|Q - \hat{Q}\|_{\text{F}} + \frac{\delta^2}{n} \sum_{i=1}^n \sum_{l=1}^n \frac{(\hat{\rho} \vee \rho)^{|l-i|}}{c^*(Y_l)} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty \right] \\
&\leq C \mathbb{E}_\theta \left[\frac{1}{n} \|\nu - \hat{\nu}\|_2 + \left(1 + \frac{\delta}{\hat{\delta}} \right) \|Q - \hat{Q}\|_{\text{F}} \right] + \frac{\delta^2 C}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[\frac{(\hat{\rho} \vee \rho)^{|l-i|}}{c^*(Y_l)} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty \right] \\
&\leq C \mathbb{E}_\theta \left[\frac{1}{n} \|\nu - \hat{\nu}\|_2 + \left(1 + \frac{\delta}{\hat{\delta}} \right) \|Q - \hat{Q}\|_{\text{F}} \right] + \frac{\delta^2 C}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[\frac{1}{c^*(Y_l)^2} \right]^{1/2} \mathbb{E}_\theta \left[(\hat{\rho} \vee \rho)^{2|l-i|} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2} \\
&\leq C \mathbb{E}_\theta \left[\frac{1}{n} \|\nu - \hat{\nu}\|_2 + \left(1 + \frac{\delta}{\hat{\delta}} \right) \|Q - \hat{Q}\|_{\text{F}} \right] + \frac{\delta C \sqrt{C^*}}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[(\hat{\rho} \vee \rho)^{2|l-i|} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2}
\end{aligned}$$

□

In order to obtain a rate on the excess risk, we make use of Algorithm 1 which will yield the estimates used in the statement of the corollary. This algorithm merges the spectral algorithms of [12, 1] with some slight modifications. Note that all the expectations and probabilities of this proof are with respect to the observations and the random unit matrices. Also note that the algorithm outputs estimates of the densities that are not necessarily *bona-fide* densities.

Algorithm 1 Non-parametric spectral estimation of the transition matrix and the emission laws

Input

- Number of states J , integers D and r .
- Data $(Y_i)_{i \leq n+2}$ drawn from a HMM with J states.
- Functions $(\varphi_d)_{d \in \mathbb{N}}$ **uniformly bounded** such that $O = (\mathbb{E}_\theta[\varphi_d(Y_1) \mid X_1 = j])_{1 \leq d \leq D, 1 \leq j \leq J}$ is of rank J with $\sigma_J(O)$ **bounded away from 0 uniformly** in D , at least for D large enough.
- K a Lipschitz-continuous kernel

Output

- Spectral estimators \hat{Q} and $(\hat{f}_j)_{1 \leq j \leq J}$

Estimation

[Step 1] For all $a, b, c \in \llbracket 1, D \rrbracket$, consider the following empirical estimators:

$$\begin{aligned}\hat{L}(a) &= \frac{1}{n} \sum_{s=1}^n \varphi_a(Y_s) \\ \hat{N}(a, b) &= \frac{1}{n} \sum_{s=1}^n \varphi_a(Y_s) \varphi_b(Y_{s+1}) \\ \hat{P}(a, c) &= \frac{1}{n} \sum_{s=1}^n \varphi_a(Y_s) \varphi_c(Y_{s+2}) \\ \hat{M}(a, b, c) &= \frac{1}{n} \sum_{s=1}^n \varphi_a(Y_s) \varphi_b(Y_{s+1}) \varphi_c(Y_{s+2}) \\ \hat{M}^{x,L}(a, b) &= \frac{1}{n} \sum_{s=1}^n \varphi_a(Y_s) K_L(x, Y_{s+1}) \varphi_b(Y_{s+2})\end{aligned}$$

where L is chosen such that: $2^L \asymp (n/\log(n))^{1/(2s+1)}$ and s is the smoothness of the emissions.

[Step 2] Let \hat{V} be the $D \times J$ matrix of orthonormal right singular vectors of \hat{P} corresponding to its top J singular values.

[Step 3] For all $d \in \llbracket 1, D \rrbracket$, set $\hat{B}(d) = (\hat{V}^\top \hat{P} \hat{V})^{-1} \hat{V}^\top \hat{M}(\cdot, d, \cdot) \hat{V}$

[Step 4] Generate Ω a $J \times J$ unit matrix uniformly drawn, for all $x \in \llbracket 1, J \rrbracket$, $\hat{C}(x) = \sum_{d=1}^D (\hat{V} \Omega)(d, x) \hat{B}(d)$

[Step 5] Compute \hat{R}_1 a $J \times J$ unit Euclidean norm columns matrix that diagonalizes the matrix $\hat{C}(1)$:

$$\hat{R}_1^{-1} \hat{C}(1) \hat{R}_1 = \text{Diag}[(\hat{\Lambda}(1, 1), \dots, \hat{\Lambda}(1, J))]$$

[Step 6] For all $x, x' \in \llbracket 1, J \rrbracket$, $\hat{\Lambda}(x, x') = (\hat{R}_1^{-1} \hat{C}(x) \hat{R}_1)_{x', x'}$.

[Step 7] Repeat steps 4 to 6 r times and take Ω_r maximizing $i \mapsto \min_{k \leq J} \min_{k_1 \neq k_2} |\hat{\Lambda}_i(k, k_1) - \hat{\Lambda}_i(k, k_2)|$

[Step 8] Set $\hat{O} = \hat{V} \Omega_r \hat{\Lambda}$, $\tilde{\nu} = (\hat{V}^\top \hat{O})^{-1} \hat{V}^\top \hat{L}$ and $\hat{Q} = \Pi_{TM} \left((\hat{V}^\top \hat{O} \text{Diag}[\tilde{\nu}])^{-1} \hat{V}^\top \hat{N} \hat{V} (\hat{O}^\top \hat{V})^{-1} \right)$

where Π_{TM} denotes the projection (with respect to the scalar product given by the Frobenius norm) onto the convex set of transition matrices.

[Step 9] For $x \in \mathbb{R}$, set $\hat{B}^x = \hat{B}^{x,D,L} = (\hat{V}^\top \hat{P} \hat{V})^{-1} \hat{V}^\top \hat{M}^{x,L} \hat{V}$

[Step 10] Set $\hat{R}_2 = \hat{Q} \hat{O}^\top \hat{V}$ and take $\tilde{f}_j(x) = (\hat{R}_2 \hat{B}^x \hat{R}_2^{-1})_{j,j}$

[Step 11] $\hat{f}_j(x) = \begin{cases} \tilde{f}_j(x) & \text{si } |\tilde{f}_j(x)| \leq n^\beta \\ n^\beta \text{sign}(\tilde{f}_j(x)) & \text{otherwise} \end{cases}$ for $\beta > 0$ fixed (but arbitrary).

This is not problematic for the plug-in procedure as one typically uses the Forward-Backward algorithm [8] which works even if the emissions are not correctly normalized.

First, we start by controlling $\mathbb{E}_\theta \left[\|Q - \hat{Q}\|_F^2 \right]$, $\mathbb{E}_\theta \left[\frac{1}{\delta^2} \right]$ and $\mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \right]$ using the estimates yielded by the algorithm. Thanks to step 7 of the algorithm, one can obtain a slightly different version of Theorem 3.1 of [12] (Note that this version is used in the proof of Corollary 3.2 in [12]). It ensures the existence of positive constants C, x_0, y_0, D_0

and n_1 such that for all $D \geq D_0$ there exist a permutation $\tau_D \in \mathcal{S}_J$ such that for all $n \geq n_1 \eta_3^2(\varphi_D) x(y + \log(r)) e^{y/r}$, $x \geq x_0$, $y \geq y_0$ and $r \geq 1$, with probability at least $1 - 4e^{-x} - 2e^{-y}$:

$$(21) \quad \begin{aligned} \|Q^{\tau_D} - \hat{Q}\|_F^2 &\leq C \eta_3^2(\varphi_D) x(y + \log(r)) e^{y/r} / n \\ \max_{x \in \mathbb{X}} \|f_{D, \tau_D(x)} - \hat{f}_{D, x}^{(r)}\|_2^2 &\leq C \eta_3^2(\varphi_D) x(y + \log(r)) e^{y/r} / n \end{aligned}$$

where:

- $(\varphi_k)_{k \in \mathbb{N}}$ is the basis used in Algorithm 1
- $\eta_3^2(\varphi_D) = \sup_{y, y' \in \mathcal{Y}^3} \sum_{a, b, c=1}^D (\varphi_a(y_1) \varphi_b(y_2) \varphi_c(y_3) - \varphi_a(y'_1) \varphi_b(y'_2) \varphi_c(y'_3))^2$
- $f_{D, x} = \sum_{d=1}^D \langle f_x, \varphi_d \rangle \varphi_d$ the projection of the density f_x on the subspace spanned by the first D components of the basis.
- $\hat{f}_{D, x}^{(r)} = \sum_{d=1}^D \hat{O}_{d, x} \varphi_d$ where \hat{O} is the matrix constructed at step 8 of Algorithm 1.

As detailed in [12], when using a wavelet basis or trigonometric polynomials basis, $\eta_3(\varphi_D)$ ensures for a constant $C_\eta > 0$:

$$(22) \quad \eta_3(\varphi_D) \leq C_\eta D^{3/2} \text{ and } \max_{l \in \mathbb{N}} \|\varphi_l\|_\infty < \infty$$

We assume a similar basis is used.

It is important to note that the estimator $\hat{f}_{D, x}^{(r)}$ is not the one yielded by the algorithm 1 but it is rather the one used in [12]. We do not use it for the estimation because it does not allow obtaining the appropriate rate in infinite norm (see [12] for more details). However, we will use it in our proof because $\|\hat{O}(\cdot, k) - O(\cdot, \tau_{D_n}(k))\|_2 = \|\hat{f}_{D_n, k}^{(r)} - f_{D_n, \tau_{D_n}(k)}\|_2$.

Assume that the parameters $x = x_n$, $y = y_n$, $D = D_n$ and $r = r_n$ are increasing with respect to n and that $n \geq n_1 C_\eta^2 D_n^3 x_n (y_n + \log(r_n)) e^{y_n/r_n}$.

Control of $\mathbb{E}_\theta [\|Q^{\tau_{D_n}} - \hat{Q}\|_F^2]$. The control of $\|Q^{\tau_{D_n}} - \hat{Q}\|_F^2$ in expectation is already proved in Corollary 3.2 of [12] using Inequality (21). The proof chooses $r_n \propto \log(n)$ and $\eta_3(\varphi_{D_n}) = o(\sqrt{n}/\log(n))$ and yields:

$$\mathbb{E}_\theta [\|Q^{\tau_{D_n}} - \hat{Q}\|_F^2] = \mathcal{O}(\eta_3^2(\varphi_{D_n}) \log(n)/n) = \mathcal{O}(D_n^3 \log(n)/n) \quad (\text{by (22)})$$

for a sequence $(\tau_{D_n})_n$ of permutations. We will keep the same values of r_n and D_n in what follows.

One of the advantages of this algorithm with respect to the previous versions is that it allows obtaining the appropriate rate on the errors in the estimation of all the model parameters. This is done thanks to the use of the kernel estimator of the emission densities for which the error of approximation is tuned (through the parameter L) independently of the error of estimation of the transition matrix. The shortcoming of the algorithm proposed in [12] is that it does not allow controlling the rate on the emission densities without altering that of the transition matrix as it is clear in Corollary 3.3 of that paper.

Control of $\mathbb{E}_\theta [\frac{1}{\hat{\delta}^2}]$. Let $\tilde{\delta} = \frac{\delta}{2}$. Then,

$$\begin{aligned} \mathbb{E}_\theta \left[\frac{1}{\hat{\delta}^2} \right] &= \mathbb{E}_\theta \left[\frac{1}{\hat{\delta}^2} \mathbf{1}_{\hat{\delta} < \tilde{\delta}} \right] + \mathbb{E}_\theta \left[\frac{1}{\hat{\delta}^2} \mathbf{1}_{\hat{\delta} \geq \tilde{\delta}} \right] \\ &\leq \mathbb{E}_\theta \left[\frac{1}{\hat{\delta}^2} \mathbf{1}_{\hat{\delta} < \tilde{\delta}} \right] + \frac{1}{\tilde{\delta}^2} \end{aligned}$$

On the other hand:

$$\tilde{\delta} - \hat{\delta} = \delta - \hat{\delta} - \frac{\delta}{2} \leq \left| \delta - \hat{\delta} \right| - \frac{\delta}{2} \leq \max_{i,j} |Q_{i,j}^{\tau_{D_n}} - \hat{Q}_{i,j}| - \frac{\delta}{2} \leq \|Q^{\tau_{D_n}} - \hat{Q}\|_F - \frac{\delta}{2}$$

where we have used the inequality $|\min_{i,j} Q_{i,j} - \min_{i,j} \hat{Q}_{i,j}| \leq \max_{i,j} |Q_{i,j} - \hat{Q}_{i,j}|$.

We assume that all the entries of \hat{Q} are between $n^{-\alpha/2}$ and $1 - n^{-\alpha/2}$ for $\alpha \geq 2$. If it is not the case, modifying the entries of \hat{Q} to obtain a similar property induces an error of order $n^{-\alpha/2}$ which is negligible with respect to the rates we seek and all the subsequent results remain unchanged. It follows that $\hat{\delta} \geq n^{-\alpha/2}$ and:

$$\mathbb{E}_\theta \left[\frac{1}{\hat{\delta}^2} \right] \leq n^\alpha \mathbb{P}_\theta \left(\|Q^{\tau_{D_n}} - \hat{Q}\|_F > \frac{\delta}{2} \right) + \frac{1}{\hat{\delta}^2}$$

Choosing for example $x_n = y_n = r_n = \alpha \log(n)$, then one obtains for n large enough:

$$\begin{aligned} \mathbb{P}_\theta \left(\|Q^{\tau_{D_n}} - \hat{Q}\|_F > \frac{\delta}{2} \right) &\leq \mathbb{P}_\theta \left(\|Q^{\tau_{D_n}} - \hat{Q}\|_F > C\eta_3^2(\varphi_{D_n})x_n(y_n + \log(r_n))e/n \right) \\ &\leq 4e^{-x_n} + 2e^{-y_n} \end{aligned}$$

It follows that $\mathbb{E}_\theta \left[\frac{1}{\hat{\delta}^2} \right]$ is upper-bounded by an absolute constant.

The values x_n, y_n and r_n will be kept the same in what follows.

Control of $\mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_{\tau_{D_n}(x)} - \hat{f}_x\|_\infty^2 \right]$. The difficulty of the control of this quantity lies in ensuring that the same permutation τ_{D_n} used for the control of Q still works for the control of the emission densities. In the spectral algorithm of [1], the matrix \hat{R} is chosen independently of Q or \hat{Q} (In fact, this algorithm does not even estimate \hat{Q}). Had we used this matrix, there would not be any reason for which the same permutation τ_{D_n} works for the control of the emission densities. To solve this problem, we choose a matrix \hat{R} that depends explicitly on \hat{Q} so that the same permutation that works for the control of Q works also for that of the emission densities. We follow here the steps of the proof of Theorem 5 in [1].

Let $M^{x,L}, P, O$ be the quantities estimated by $\hat{M}^{x,L}, \hat{P}, \hat{O}$ and construct \tilde{f}_j, \tilde{f}_j using Algorithm 1. Let E_n be the event with probability greater than $1 - 4e^{-x_n} - 2e^{-y_n}$ on which the control of (21) holds when $x_n = y_n = r_n = \alpha \log(n)$ and $D = D_n$. For $\gamma > 0$ there exists $c = c(\gamma)$ such that the event:

$$\mathcal{A}_n = \left\{ \|\hat{P} - P\| \leq cD_n \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}}, \quad \sup_{x \in \mathbb{R}} \|\hat{M}^{x,L} - M^{x,L}\| \leq cD_n^2 \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \right\}$$

is measurable and has probability $n^{-\gamma}$ (Cf. Lemma 25.a of [1]). Given that E_n has probability greater than $1 - 4e^{-x_n} - 2e^{-y_n} = 1 - 6n^{-\alpha}$, it follows that $\mathcal{A}_n \cap E_n$ has probability greater than $1 - n^{-\gamma} - 6n^{-\alpha}$. Note that the difference with the original proof is that we use the event $\mathcal{A}_n \cap E_n$ instead of the event \mathcal{A}_n . This is compulsory to control the errors of \hat{Q} and \hat{f} simultaneously.

On the event $\mathcal{A}_n \cap E_n$, and at step 10 of Algorithm 1, instead of using the matrix \hat{R} appearing in the spectral algorithm in [1], we rather use the matrix $\hat{R}_2 = \hat{Q}\hat{O}^\top \hat{V}$ where the components \hat{V}, \hat{Q} and \hat{O} are constructed in the Algorithm 1. On the other hand, since the columns of \hat{R} are not normalized, we choose $\tilde{R}_2 = QO^\top \hat{V}$ on the contrary to what is done in the proof of Theorem 5 in [1]. By denoting $Q^{\tau_{D_n}} = P_{\tau_{D_n}} Q P_{\tau_{D_n}}^{-1}$, one obtains:

$$\hat{R}_2 - P_{\tau_{D_n}} \tilde{R}_2 = \hat{Q}\hat{O}^\top \hat{V} - P_{\tau_{D_n}} QO^\top \hat{V} = \left(\hat{Q}(\hat{O} - O P_{\tau_{D_n}}^\top)^\top + (\hat{Q} - Q^{\tau_{D_n}})P_{\tau_{D_n}} O^\top \right) \hat{V}$$

It follows by using operator norm:

$$(23) \quad \|\hat{R}_2 - P_{\tau_{D_n}} \tilde{R}_2\| \leq \left(\|\hat{Q}\| \|\hat{O} - OP_{\tau_{D_n}}^\top\| + \|P_{\tau_{D_n}} O^\top\| \|\hat{Q} - Q^{\tau_{D_n}}\| \right) \|\hat{V}\|$$

First, note that on the event $\mathcal{A}_n \cap E_n$:

$$\begin{aligned} \|\hat{Q}\| &\leq 1 \\ \|\hat{V}\| &\leq \|\hat{V}\|_F = J^{1/2} \quad (\text{columns are normalized}) \\ \|\hat{Q} - Q^{\tau_{D_n}}\| &\leq \|\hat{Q} - Q^{\tau_{D_n}}\|_F \\ &\leq \left(CC_\eta^2 D_n^3 x_n (y_n + \log(r_n)) e^{y_n/r_n} / n \right)^{1/2} \quad (\text{by (21) and (22)}) \\ \|P_{\tau_{D_n}} O^\top\| = \|O\| &= \sup_{\|v\|=1} \left(\sum_{j=1}^J \left(\sum_{l=1}^{D_n} v_l \langle f_j, \varphi_l \rangle \right)^2 \right)^{1/2} \\ &\leq (JD_n)^{1/2} \max_{1 \leq l \leq D_n} \|\varphi_l\|_\infty \\ \|\hat{O} - OP_{\tau_{D_n}}^\top\| &\leq \|\hat{O} - OP_{\tau_{D_n}}^\top\|_F = \left(\sum_{k=1}^K \|\hat{O}(\cdot, k) - O(\cdot, \tau_{D_n}(k))\|_2^2 \right)^{1/2} \\ &= \left(\sum_{k=1}^K \|\hat{f}_{D_n, k}^{(r)} - f_{D_n, \tau_{D_n}(k)}\|_2^2 \right)^{1/2} \\ &\leq \left(KCC_\eta^2 D_n^3 x_n (y_n + \log(r_n)) e^{y_n/r_n} / n \right)^{1/2} \quad (\text{by (21) and (22)}). \end{aligned}$$

By keeping the previous choices of x_n, y_n and r_n then by Inequality (23), there exists a constant C' such that:

$$\mathbf{1}_{\mathcal{A}_n \cap E_n} \|\hat{R}_2 - P_{\tau_{D_n}} \tilde{R}_2\| \leq C' D_n^{1/2} D_n^{3/2} \frac{\log(n)}{\sqrt{n}}.$$

By Lemma 25.b of [1], for n large enough, \hat{P} has rank J , $(\hat{V}^T \hat{P} \hat{V})$ and $(\hat{V}^T P \hat{V})$ are invertible and the matrices $(\hat{B}(d))_{1 \leq d \leq D_n}$ appearing in algorithm 1 are then well-defined. By Lemma 11 and 25.b of [1], $\tilde{B}^x = (\hat{V}^T P \hat{V})^{-1} \hat{V}^T M^x \hat{V}$ satisfies:

$$(24) \quad \tilde{B}^x = (QO^T \hat{V})^{-1} D^x (QO^T \hat{V}) = \tilde{R}_2^{-1} D^x \tilde{R}_2$$

where $D^x = (K_L[f_j](x))_{j \leq J}$. Using the fact that $\sigma_J(\hat{V}) = 1$ (the columns of \hat{V} are orthonormal) and $\sigma_J(QO^T) \geq \sigma_J(Q)\sigma_J(O) > 0$ (because Q in full rank and $\sigma_J(O)$ is bounded from below by an absolute constant by assumption of the algorithm), it follows that $\|\tilde{R}_2^{-1}\|^{-1}$ is bounded from below by an absolute constant because:

$$\begin{aligned} (25) \quad \|\tilde{R}_2^{-1}\|^{-1} &= \frac{1}{\sigma_1(\tilde{R}_2^{-1})} = \sigma_J(\tilde{R}_2) \\ &= \sigma_J(QO^T \hat{V}) \geq \sigma_J(QO^T) \sigma_J(\hat{V}) = \sigma_J(QO^T) \geq \sigma_J(Q) \sigma_J(O) > 0. \end{aligned}$$

Thus, for n large enough, the assumption of Lemma 13 (stated below) is verified with $A_t = \tilde{B}^t$, $\hat{A}_t = \hat{B}^t$, $R = \tilde{R}_2 = QO^T \hat{V}$. This ensures that:

$$\begin{aligned} & \mathbf{1}_{\mathcal{A}_n \cap E_n} \max_{x \in \mathbb{X}} \|\tilde{f}_x - K_L[f_{\tau_{D_n}}(x)]\|_\infty \\ & \leq 4\kappa(\tilde{R}_2) \left[\mathbf{1}_{\mathcal{A}_n \cap E_n} \sup_t \|\tilde{B}^t - \hat{B}^t\| + \lambda_{\max} \kappa(\tilde{R}_2) \|\tilde{R}_2^{-1}\| \mathbf{1}_{\mathcal{A}_n \cap E_n} \|\hat{R}_2 - P_{\tau_{D_n}} \tilde{R}_2\| \right] \end{aligned}$$

where $\lambda_{\max} = \sup_t \max_j |\lambda_{t,j}| = \sup_t \max_j |K_L[f_j]| < \infty$ and $\lambda_{t,j}$ is the j -th eigenvalue of \tilde{B}^t . By Lemma 35.c of [1], one has for some constant C :

$$\mathbf{1}_{\mathcal{A}_n \cap E_n} \sup_t \|\tilde{B}^t - \hat{B}^t\| \leq C D_n^2 \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}}.$$

By Lemma 35.b in [1], one obtains: $\kappa(\tilde{R}_2) \leq \tilde{C} D_n^{1/2}$ and $\|\tilde{R}_2^{-1}\|$ upper-bounded by an absolute constant by (25). We finally obtain that:

$$\begin{aligned} \mathbf{1}_{\mathcal{A}_n \cap E_n} \max_{x \in \mathbb{X}} \|\tilde{f}_x - K_L[f_{\tau_{D_n}}(x)]\|_\infty & \leq c' D_n^{1/2} \left[D_n^2 \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} + D_n^{1/2} D_n^{1/2} D_n^{3/2} \frac{\log(n)}{\sqrt{n}} \right] \\ & \leq c'' D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \end{aligned}$$

for some constants c', c'' . The choice of L (cf. Algorithm 1) allows obtaining then:

$$\mathbf{1}_{\mathcal{A}_n \cap E_n} \max_{x \in \mathbb{X}} \|\tilde{f}_x - f_{\tau_{D_n}}(x)\|_\infty \leq c'' D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}}.$$

Given that for n large enough $\|f_j\|_\infty \leq n^\beta$, it follows that: $\|\hat{f}_j - f_{\tau_{D_n}(j)}\|_\infty \leq \|\hat{f}_j^L - f_{\tau_{D_n}(j)}\|_\infty$.

Finally, thanks to the truncation of the emission densities, it is possible to obtain the same rate in expectation:

$$\begin{aligned} \mathbb{E}_\theta \left[\|\hat{f}_j - f_{\tau_{D_n}(j)}\|_\infty^2 \right] & \leq c'' D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{2s}{2s+1}} + 2n^\beta \mathbb{P}((\mathcal{A}_n \cap E_n)^c) \\ & \leq c'' D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{2s}{2s+1}} + 2n^\beta (n^{-\gamma} + 6n^{-\alpha}). \end{aligned}$$

By choosing α and γ sufficiently large, one obtains:

$$\mathbb{E}_\theta \left[\|\hat{f}_j - f_{\tau_{D_n}(j)}\|_\infty^2 \right] = \mathcal{O} \left(D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{2s}{2s+1}} \right).$$

Control of $\frac{1}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[(\rho \vee \hat{\rho})^{2|l-i|} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2}$. Let $\tilde{\rho} = \frac{1-2\tilde{\delta}}{1-\tilde{\delta}}$ where $\tilde{\delta} = \frac{\delta}{2}$.

$$\begin{aligned} \mathbb{E}_\theta \left[(\rho \vee \hat{\rho})^{2|l-i|} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2} & \leq \tilde{\rho}^{|l-i|} \mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2} \\ & \quad + \mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \mathbf{1}_{\hat{\rho} \geq \tilde{\rho}} \right]^{1/2}. \end{aligned}$$

The term $\mathbb{E}_\theta \left[\max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \mathbf{1}_{\hat{\rho} \geq \tilde{\rho}} \right]^{1/2}$ can be made of order $n^{-\alpha}$ for arbitrary $\alpha > 0$ by the same large deviation argument used in the control of $\mathbb{E}_\theta \left[\frac{1}{\delta^2} \right]$. Then, summing up over

i and n yields:

$$\frac{1}{n} \sum_{i=1}^n \sum_{l=1}^n \mathbb{E}_\theta \left[(\rho \vee \hat{\rho})^{2|l-i|} \max_{x \in \mathbb{X}} \|f_x - \hat{f}_x\|_\infty^2 \right]^{1/2} = \mathcal{O} \left(D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \right)$$

Finally, using Proposition 3 and the previous controls of the errors of estimation of the model parameters, one gets:

$$\mathbb{E}[\mathcal{R}_n^{\text{class}}(\theta, \hat{h}^{\tau_n})] - \inf_{h \in \mathcal{H}_n} \mathcal{R}_n^{\text{class}}(\theta, h) = \mathcal{O} \left(D_n^{5/2} \left(\frac{\log(n)}{n} \right)^{\frac{s}{2s+1}} \right).$$

A similar rate holds for the excess risk of clustering thanks to the relationship between the Bayes risk of classification and the Bayes risk of clustering established in Theorems 6 and 8.

LEMMA 13. *Suppose $(A_t, t \in \mathbb{R})$ are $J \times J$ matrices simultaneously diagonalized by a matrix R :*

$$RA_tR^{-1} = \text{diag}(\lambda_{t,1}, \dots, \lambda_{t,J}), t \in \mathbb{R}.$$

Let \hat{R} be a matrix such that for some permutation τ of $\{1, \dots, J\}$, we have:

$$\|\hat{R} - P_\tau R\| = \varepsilon_R \leq \frac{\|R^{-1}\|^{-1}}{2}$$

Assume $\lambda_{\max} = \sup_t \max_j |\lambda_{t,j}| < \infty$. For matrices $(\hat{A}_t)_{t \in \mathbb{R}}$, write $\varepsilon_A = \sup_t \|A_t - \hat{A}_t\|$ and define

$$\hat{\lambda}_{t,j} = e_j^T \hat{R} \hat{A}_t \hat{R}^{-1} e_j \quad \text{and} \quad \lambda_{t,\tau(j)} = e_j^T (P_\tau R) A_t (P_\tau R)^{-1} e_j.$$

Then

$$\sup_t \max_j |\hat{\lambda}_{t,j} - \lambda_{t,\tau(j)}| \leq 4\kappa(R) [\varepsilon_A + \lambda_{\max} \kappa(R) \|R^{-1}\| \varepsilon_R].$$

PROOF. Let $\hat{\zeta}_j^T = e_j^T \hat{R}$, let $\hat{\xi}_j = \hat{R}^{-1} e_j$ and define $\zeta_j^T = e_j^T P_\tau R$ and $\xi_j = (P_\tau R)^{-1} e_j$. Then, $\lambda_{t,\tau(j)} = \zeta_j^T A_t \xi_j$, $\hat{\lambda}_{t,j} = \hat{\zeta}_j^T \hat{A}_t \hat{\xi}_j$ and we have:

$$\begin{aligned} |\hat{\lambda}_{t,j} - \lambda_{t,\tau(j)}| &= |\hat{\zeta}_j^T \hat{A}_t \hat{\xi}_j - \zeta_j^T A_t \xi_j| \\ &= |\hat{\zeta}_j^T \hat{A}_t (\hat{\xi}_j - \xi_j) + (\hat{\zeta}_j^T - \zeta_j^T) A_t \xi_j + \hat{\zeta}_j^T (\hat{A}_t - A_t) \xi_j| \\ &\leq \|\hat{\zeta}_j^T\| \|\hat{A}_t\| \|\hat{\xi}_j - \xi_j\| + \|\hat{\zeta}_j^T - \zeta_j^T\| \|A_t \xi_j\| + \|\hat{\zeta}_j^T\| \|\xi_j\| \varepsilon_A \\ \|\zeta_j^T\| &= \|e_j^T P_\tau R\| \leq \|P_\tau R\| = \|R\| \\ \|\hat{\zeta}_j^T - \zeta_j^T\| &= \|e_j^T (\hat{R} - P_\tau R)\| \leq \|\hat{R} - P_\tau R\| = \varepsilon_R \\ \|\xi_j\| &= \|(P_\tau R)^{-1} e_j\| \leq \|(P_\tau R)^{-1}\| = \|R^{-1}\| \\ \|\hat{\xi}_j - \xi_j\| &= \|(\hat{R}^{-1} - (P_\tau R)^{-1}) e_j\| \leq \|\hat{R}^{-1} - (P_\tau R)^{-1}\| \\ &\leq \frac{\|R^{-1}\|^2 \varepsilon_R}{1 - \varepsilon_R \|R^{-1}\|} \text{ (by Lemma 37 of [1])} \\ \|A_t\| &= \|R^{-1} \text{diag}(\lambda_{t,\cdot}) R\| \leq \lambda_{\max} \kappa(R) \\ \|A_t \xi_j\| &= \|\lambda_{t,\tau(j)} \xi_j\| \leq \lambda_{\max} \|R^{-1}\| \\ \|\hat{\zeta}_j^T\| &= \|\hat{\zeta}_j^T - \zeta_j^T + \zeta_j^T\| \leq \|\hat{\zeta}_j^T - \zeta_j^T\| + \|\zeta_j^T\| \leq \varepsilon_R + \|R\| \\ \|\hat{A}_t\| &\leq \varepsilon_A + \|A_t\| \leq \varepsilon_A + \lambda_{\max} \kappa(R). \end{aligned}$$

Thus, by the inequalities above:

$$\begin{aligned}
|\hat{\lambda}_{t,j} - \lambda_{t,\tau(j)}| &\leq (\varepsilon_R + \|R\|)(\varepsilon_A + \lambda_{\max}\kappa(R)) \frac{\|R^{-1}\|^2 \varepsilon_R}{1 - \varepsilon_R \|R^{-1}\|} \\
&\quad + \lambda_{\max} \varepsilon_R \|R^{-1}\| + \|R^{-1}\|(\varepsilon_R + \|R\|)\varepsilon_A \\
&\leq 2(\varepsilon_R + \|R\|)(\varepsilon_A + \lambda_{\max}\kappa(R))\|R^{-1}\|^2 \varepsilon_R \\
&\quad + \lambda_{\max}\|R^{-1}\|\varepsilon_R + \|R^{-1}\|(\varepsilon_R + \|R\|)\varepsilon_A \\
&\leq 2\|R^{-1}\|^2 \varepsilon_R^2 \varepsilon_A + 2\kappa(R)\|R^{-1}\|\varepsilon_A \varepsilon_R + 2\lambda_{\max}\kappa(R)\|R^{-1}\|^2 \varepsilon_R^2 \\
&\quad + 2\lambda_{\max}\kappa(R)^2\|R^{-1}\|\varepsilon_R + \lambda_{\max}\|R^{-1}\|\varepsilon_R + \kappa(R)\varepsilon_A + \|R^{-1}\|\varepsilon_R \varepsilon_A \\
&\leq 2\|R^{-1}\|^2 \varepsilon_R^2 \varepsilon_A + (2\kappa(R) + 1)\|R^{-1}\|\varepsilon_A \varepsilon_R \\
&\quad + 2\lambda_{\max}\kappa(R)((\|R^{-1}\|\varepsilon_R)^2 + \kappa(R)(\|R^{-1}\|\varepsilon_R)) + \lambda_{\max}\|R^{-1}\|\varepsilon_R + \kappa(R)\varepsilon_A \\
&\leq 2\left(\frac{1}{2}\right)^2 \varepsilon_A + \frac{3}{2}\kappa(R)\varepsilon_A + 2\lambda_{\max}\kappa(R)\|R^{-1}\|\varepsilon_R \left(\kappa(R) + \frac{1}{2}\right) \\
&\quad + \lambda_{\max}\|R^{-1}\|\varepsilon_R + \kappa(R)\varepsilon_A \\
&\leq 3\kappa(R)\varepsilon_A + \lambda_{\max}(1 + 3\kappa(R)^2)\|R^{-1}\|\varepsilon_R \\
&\leq 4\kappa(R) [\varepsilon_A + \lambda_{\max}\kappa(R)\|R^{-1}\|\varepsilon_R]
\end{aligned}$$

where we have used $\|R^{-1}\|\varepsilon_R \leq 1/2$ and $\kappa(R) \geq 1$. \square

6.12. *Proof of Lemma 1.* Let define $p_{n,i}(\theta) := \max_k \mathbb{P}_\theta(X_i = k \mid Y_{1:n}) = \mathbb{P}_\theta(X_i = h_{\theta,i}^* \mid Y_{1:n})$; defining $h_{\theta,i}^*$ realizing the maximum. Suppose $p_{n,i}(\theta) \geq \frac{1}{2} + \gamma$ for some $0 < \gamma \leq 1/2$. Then,

$$\begin{aligned}
p_{n,i}(\theta) - \max_{k \neq h_{\theta,i}^*} \mathbb{P}_\theta(X_i = k \mid Y_{1:n}) &\geq p_{n,i}(\theta) - \sum_{k \neq h_{\theta,i}^*} \mathbb{P}_\theta(X_i = k \mid Y_{1:n}) \\
&= p_{n,i}(\theta) - [1 - p_{n,i}(\theta)] \\
&= 2p_{n,i}(\theta) - 1 \\
&\geq 2\gamma.
\end{aligned}$$

Consequently if $p_{n,i}(\theta) \geq \frac{1}{2} + \gamma$,

$$\begin{aligned}
\mathbb{P}_{\hat{\theta}}(X_i = h_{\hat{\theta},i}^* \mid Y_{1:n}) &\geq p_{n,i}(\theta) - \left\| \phi_{\theta,i|n} - \phi_{\hat{\theta},i|n} \right\|_{\text{TV}} \\
&\geq \max_{k \neq h_{\theta,i}^*} \mathbb{P}_\theta(X_i = k \mid Y_{1:n}) + 2\gamma - \left\| \phi_{\theta,i|n} - \phi_{\hat{\theta},i|n} \right\|_{\text{TV}} \\
&\geq \max_{k \neq h_{\theta,i}^*} \mathbb{P}_{\hat{\theta}}(X_i = k \mid Y_{1:n}) + 2\gamma - 2 \left\| \phi_{\theta,i|n} - \phi_{\hat{\theta},i|n} \right\|_{\text{TV}}
\end{aligned}$$

We have shown that on the intersection of the two events

$$E_{n,i} := \left\{ p_{n,i}(\theta) \geq \frac{1}{2} + \gamma \right\}, \quad F_{n,i} := \left\{ \left\| \phi_{\theta,i|n} - \phi_{\hat{\theta},i|n} \right\|_{\text{TV}} < \gamma \right\}$$

the plug-in rule $h_{\hat{\theta}}^*$ maximizing $k \mapsto \mathbb{P}_{\hat{\theta}}(X_i = k \mid Y_{1:n})$ is unique and it must be that $h_{\hat{\theta},i}^* = h_{\theta,i}^*$. Then, we bound the risk as follows,

$$\begin{aligned} \mathcal{R}_n^{\text{class}}(\theta, h_{\hat{\theta}}^*) &\leq \mathbb{E}_{\theta} \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{h_{\hat{\theta},i}^* \neq X_i\}} \mathbf{1}_{E_{n,i}} \cap F_{n,i} \right] + \mathbb{E}_{\theta} \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{h_{\hat{\theta},i}^* \neq X_i\}} \mathbf{1}_{E_{n,i}^c} \right] \\ &\quad + \mathbb{E}_{\theta} \left[\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{h_{\hat{\theta},i}^* \neq X_i\}} \mathbf{1}_{F_{n,i}^c} \right] \\ &\leq \mathbb{E}_{\theta} \left[\frac{1}{n} \sum_{i=1}^n \mathbb{P}_{\theta}(h_{\hat{\theta},i}^* \neq X_i) \mathbf{1}_{E_{n,i}} \right] + \frac{1}{n} \sum_{i=1}^n \mathbb{P}_{\theta}(E_{n,i}^c) + \frac{1}{n} \sum_{i=1}^n \mathbb{P}_{\theta}(F_{n,i}^c) \end{aligned}$$

Finally, notice that,

$$\begin{aligned} \mathbb{P}_{\theta}(E_{n,i}^c) &= \mathbb{P}_{\theta} \left(p_{n,i}(\theta) < \frac{1}{2} + \gamma \right) \\ &= \mathbb{P}_{\theta} \left(\mathbb{P}_{\theta}(X_i = h_{\hat{\theta},i}^* \mid Y_{1:n}) < \frac{1}{2} + \gamma \right) \\ &= \mathbb{P}_{\theta} \left(\mathbb{P}_{\theta}(X_i \neq h_{\hat{\theta},i}^* \mid Y_{1:n}) > \frac{1}{2} - \gamma \right) \\ &\leq \frac{1}{1/2 - \gamma} \mathbb{E}_{\theta} \left(\mathbb{P}_{\theta}(h_{\hat{\theta},i}^* \neq X_i \mid Y_{1:n}) \mathbf{1}_{E_{n,i}^c} \right). \end{aligned}$$

Hence the result.

6.13. Equivalence of the definitions of the risk of clustering.

LEMMA 14. *The risk of clustering of $\pi_n \circ h$ can be rewritten as*

$$(26) \quad \mathcal{R}_n^{\text{clust}}(\theta, \pi_n \circ h) := \mathbb{E}_{\theta} \left[\min_{\tau \in \mathcal{S}_J} \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) \neq \tau(X_i)} \right]$$

PROOF. It suffices to show that

$$\sup_{\substack{M \subseteq \mathcal{E}(\pi_n \circ h(Y_{1:n}), \Pi_n) \\ M \text{ is a matching}}} \sum_{\{C, C'\} \in M} \text{Card}(C \cap C') = \sup_{\tau \in \mathcal{S}_J} \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{h_i(Y_{1:n}) = \tau(X_i)}$$

Let $C_k = \{i \in [n] \mid h_i(Y_{1:n}) = k\}$ and $C'_k = \{i \in [n] \mid X_i = k\}$. Since the two partitions Π_n and $\pi_n \circ h(Y_{1:n})$ have the same number of clusters J (with possibly empty clusters), the supremum is reached on matchings with J edges. Using this fact, it follows that the matching reaching the supremum is of the form:

$$M = \left\{ (C_k, C'_{\tau(k)}) \mid 1 \leq k \leq J \right\}$$

where τ is a permutation of $\{1, \dots, J\}$. One obtains:

$$\begin{aligned}
 (27) \quad & \sup_{\substack{M \subseteq \mathcal{E}(\pi_n \circ h(Y_{1:n}), \Pi_n) \\ M \text{ is a matching}}} \sum_{\{C, C'\} \in M} \text{Card}(C \cap C') = \sup_{\tau \in \mathcal{S}_J} \sum_{k=1}^J \text{Card}(C_k \cap C'_{\tau(k)}) \\
 &= \sup_{\tau \in \mathcal{S}_J} \sum_{k=1}^J \sum_{i=1}^n \mathbf{1}_{\tau^{-1}(X_i)=h_i(Y_{1:n})=k} \\
 &= \sup_{\tau \in \mathcal{S}_J} \sum_{i=1}^n \mathbf{1}_{\tau(X_i)=h_i(Y_{1:n})}
 \end{aligned}$$

□

6.14. *Proof of Lemma 2.* Without loss of generality, let $\alpha_i \geq \frac{1}{2}$ for all i . When $p_i = \alpha_i$, we will say that i is given positive bias, and similarly when $p_i = 1 - \alpha_i$ we say it has negative bias. We show that unless we give all positive bias or all negative bias, we can flip the bias of some i to increase the expectation. This suffices to conclude. Let $(Z_i)_{i \in [n]}$ be a sequence of independent Bernoulli random variables such that $Z_n \sim \mathcal{B}(\alpha_i)$, and let $\beta_i \in \{-1, +1\}$ be the bias we give i . We then let

$$Y_i = 2Z_i - 1 \in \{-1, +1\} \text{ and } X_i = \frac{1 + \beta_i Y_i}{2} \in \{0, 1\}.$$

Consequently, $X_i \sim \mathcal{B}(\alpha_i \mathbf{1}_{\beta_i=1} + (1 - \alpha_i) \mathbf{1}_{\beta_i=-1})$ and $\sum_{i=1}^n X_i - \frac{n}{2} = \frac{1}{2} \sum_{i=1}^n \beta_i Y_i$. Letting $S_n = \sum_{i=1}^n \beta_i Y_i$, we intend to choose the β_i to maximize $\mathbb{E}[|S_n|]$. Let $S_{\neq k} = \sum_{i \neq k} \beta_i Y_i$ and define:

$$\text{sign}(x) = \begin{cases} 1 & \text{if } x < 0 \\ 0 & \text{if } x = 0 \\ -1 & \text{if } x > 0 \end{cases}$$

Then,

$$\begin{aligned}
 |S_n| &= S_n \text{sign}(S_n) \\
 &= S_n \text{sign}(S_{\neq k}) + S_n (\text{sign}(S_n) - \text{sign}(S_{\neq k})) \\
 &= S_{\neq k} \text{sign}(S_{\neq k}) + \beta_k Y_k \text{sign}(S_{\neq k}) + S_n (\text{sign}(S_n) - \text{sign}(S_{\neq k})) \\
 &= |S_{\neq k}| + \beta_k Y_k \text{sign}(S_{\neq k}) + \mathbf{1}_{S_{\neq k}=0}
 \end{aligned}$$

By the computation above and the fact that $S_{\neq k}$ and Y_k are independent:

$$\mathbb{E}[|S_n|] = \mathbb{E}[|S_{\neq k}|] + \mathbb{P}(S_{\neq k} = 0) + \mathbb{E}[\text{sign}(S_{\neq k})] \beta_k \mathbb{E}[Y_k].$$

Since $\mathbb{E}[Y_k] \geq 0$, we conclude that if we fix the values of $(\beta_i)_{i \neq k}$, then the value of β_k that maximizes $\mathbb{E}[|S_n|]$ is:

$$(28) \quad \beta_k = \text{sign}(\mathbb{E}[\text{sign}(S_{\neq k})]).$$

Assume for the moment that all the α_i are distinct and obey $\alpha_i > \frac{1}{2}$. This assumption will guarantee that for non-same-sign biases, there is always at least one bias we can flip to strictly increase $\mathbb{E}[|S_n|]$. We will remove this assumption at the end of the proof. Consider any assignment of the biases β_1, \dots, β_n .

LEMMA 15. *If all the α_i are distinct and ensure $\alpha_i > \frac{1}{2}$, the values $(\mathbb{E}[\text{sign}(S_{\neq j})])_{j \in [n]}$ are distinct, so there exists k such that $\mathbb{E}[\text{sign}(S_{\neq k})] \neq 0$.*

LEMMA 16. *Suppose $\beta_j = 1$ and $\beta_k = -1$. Then $\mathbb{E}[\text{sign}(S_{\neq k})] \geq \mathbb{E}[\text{sign}(S_{\neq j})]$.*

Both facts will be proved at the end of this section. With these facts in hand, consider any non-same-sign biases β_1, \dots, β_n . By Lemma 15, there exists k such that:

$$\mathbb{E}[\text{sign}(S_{\neq k})] \neq 0.$$

Without loss of generality, suppose $\beta_k = -1$. There are two cases to consider.

- If $\mathbb{E}[\text{sign}(S_{\neq k})] > 0$, then β_k disobeys the condition (28), so swapping to $\beta_k = 1$ increases $\mathbb{E}_\theta[|S_n|]$.
- If $\mathbb{E}[\text{sign}(S_{\neq k})] < 0$, then β_k obeys the condition (28), so we need to find another bias to swap. Because the assignment is non-same-sign, there exists j such that $\beta_j = 1$. And by Lemma 16, $\mathbb{E}[\text{sign}(S_{\neq j})] < 0$. This means β_j disobeys the condition 28, so swapping to $\beta_j = -1$ increases $\mathbb{E}[|S_n|]$.

We have shown that any non-same-sign bias assignment is suboptimal for maximizing $\mathbb{E}[|S_n|]$, so only the same-sign cases can be optimal. And it is clear by symmetry that they both yield the same $\mathbb{E}[|S_n|]$ value, so both are optimal. Note that $\mathbb{E}[|S_n|]$ is a polynomial in the parameters $(\alpha_i)_{i \in [n]}$. When the assumption that all α_i are distinct and obey $\alpha_i > \frac{1}{2}$ does not hold, the result is still true thanks to continuity with respect to $(\alpha_i)_{i \in [n]}$.

6.14.1. *Proof of Lemma 15.* Let $S_{\neq \{j,k\}} = \sum_{i \notin \{j,k\}} \beta_i Y_i$ and $j \neq k \in [n]$.

$$\begin{aligned} \mathbb{E}[\text{sign}(S_{\neq j})] - \mathbb{E}[\text{sign}(S_{\neq k})] &= \mathbb{P}(S_{\neq j} > 0) - \mathbb{P}(S_{\neq j} < 0) - \mathbb{P}(S_{\neq k} > 0) + \mathbb{P}(S_{\neq k} < 0) \\ &= \mathbb{P}(S_{\neq \{j,k\}} + \beta_k Y_k > 0) - \mathbb{P}(S_{\neq \{j,k\}} + \beta_k Y_k < 0) - \mathbb{P}(S_{\neq \{j,k\}} + \beta_j Y_j > 0) + \mathbb{P}(S_{\neq \{j,k\}} + \beta_j Y_j < 0) \\ &= \alpha_k \mathbb{P}(S_{\neq \{j,k\}} > -\beta_k) + (1 - \alpha_k) \mathbb{P}(S_{\neq \{j,k\}} > \beta_k) - \alpha_k \mathbb{P}(S_{\neq \{j,k\}} < -\beta_k) - (1 - \alpha_k) \mathbb{P}(S_{\neq \{j,k\}} < \beta_k) \\ &\quad - \alpha_j \mathbb{P}(S_{\neq \{j,k\}} > -\beta_j) - (1 - \alpha_j) \mathbb{P}(S_{\neq \{j,k\}} > \beta_j) + \alpha_j \mathbb{P}(S_{\neq \{j,k\}} < -\beta_j) + (1 - \alpha_j) \mathbb{P}(S_{\neq \{j,k\}} < \beta_j) \end{aligned}$$

When $\beta_j = \beta_k = 1$,

$$\mathbb{E}[\text{sign}(S_{\neq j})] - \mathbb{E}[\text{sign}(S_{\neq k})] = (\alpha_k - \alpha_j) (\mathbb{P}(S_{\neq \{j,k\}} \in \{0, 1\}) + \mathbb{P}(S_{\neq \{j,k\}} \in \{0, -1\})) \neq 0.$$

The remaining cases can be analyzed similarly.

6.14.2. *Proof of Lemma 16.* We want to show that:

$$\mathbb{E}[\text{sign}(S_{\neq \{j,k\}} + Y_j)] \geq \mathbb{E}[\text{sign}(S_{\neq \{j,k\}} - Y_k)]$$

The key observation is that we have a stochastic dominance $+Y_j \geq_{st} -Y_k$. This means that for any nondecreasing function f :

$$\mathbb{E}[f(Y_j)] \geq \mathbb{E}[f(-Y_k)].$$

The desired result follows by applying the above to

$$f(x) = \mathbb{E}[\text{sign}(S_{\neq \{j,k\}} + x)]$$

which is clearly nondecreasing.

Acknowledgements. Élisabeth Gassiat is supported by *Institut Universitaire de France*. All the authors are supported by the *Agence Nationale de la Recherche* under projects ANR-21-CE23-0035-02 and ANR-23-CE40-0018-02. The authors would like to thank Ziv Scully for providing a proof for Lemma 2 and the associate editor for pointing out the question of the relationship between the Bayes classifier and the Bayes clusterer.

REFERENCES

- [1] ABRAHAM, K., CASTILLO, I. and GASSIAT, E. (2022). Multiple testing in nonparametric hidden Markov models: An empirical Bayes approach. *J. Mach. Learn. Res.* **23** 4061–4117.
- [2] ABRAHAM, K., GASSIAT, E. and NAULET, Z. (2023). Frontiers to the learning of nonparametric hidden Markov models. *arXiv preprint*. <https://doi.org/10.48550/arXiv.2306.16293>
- [3] ABRAHAM, K., GASSIAT, E. and NAULET, Z. (2023). Fundamental limits for learning hidden Markov model parameters. *IEEE Trans. Inform. Theory* **69** 1777–1794. <https://doi.org/10.1109/tit.2022.3213429> MR4564681
- [4] ALEXANDROVICH, G., HOLZMANN, H. and LEISTER, A. (2016). Nonparametric identification and maximum likelihood estimation for hidden Markov models. *Biometrika* **103** 423–434. <https://doi.org/10.1093/biomet/asw001> MR3509896
- [5] ALEXANDROVICH, G., HOLZMANN, H. and LEISTER, A. (2016). Nonparametric identification and maximum likelihood estimation for hidden Markov models. *Biometrika* **103** 423–434. <https://doi.org/10.1093/biomet/asw001>
- [6] ANANDKUMAR, A., GE, R., HSU, D. J., KAKADE, S. M., TELGARSKY, M. et al. (2014). Tensor decompositions for learning latent variable models. *J. Mach. Learn. Res.* **15** 2773–2832.
- [7] ARAGAM, B., DAN, C., P. XING, E. and RAVIKUMAR, P. (2020). Identifiability of Non-parametric mixture models and Bayes Optimal clustering. *Ann. Statist.* **48** 2277–2302. <https://doi.org/10.1214/19-AOS1887>
- [8] CAPPÉ, O., MOULINES, E. and RYDEN, T. (2005). *Inference in Hidden Markov Models (Springer Series in Statistics)*. Springer-Verlag.
- [9] CASTRO, Y. D., GASSIAT, É. and LACOUR, C. (2016). Minimax Adaptive Estimation of Nonparametric Hidden Markov Models. *J. Mach. Learn. Res.* **17** 1–43.
- [10] COUVREUR, L. and COUVREUR, C. (2000). Wavelet-based non-parametric HMM's: theory and applications. In *2000 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.00CH37100)* **1** 604–607 vol.1. <https://doi.org/10.1109/ICASSP.2000.862054>
- [11] DE CASTRO, Y. (2016). Nonparametric HMM. <https://github.com/ydecastro/nonparametric-hmm>.
- [12] DE CASTRO, Y., GASSIAT, É. and LE CORFF, S. (2017). Consistent Estimation of the Filtering and Marginal Smoothing Distributions in Nonparametric Hidden Markov Models. *IEEE Trans. Inform. Theory* **63** 4758–4777. <https://doi.org/10.1109/TIT.2017.2696959>
- [13] DEVROYE, L., GYÖRFI, L. and LUGOSI, G. (1996). *A Probabilistic Theory of Pattern Recognition* **31**. <https://doi.org/10.1007/978-1-4612-0711-5>
- [14] GAO, C., MA, Z., ZHANG, A. Y. and ZHOU, H. H. (2018). Community detection in degree-corrected block models. *Ann. Statist.* **46** 2153–2185. <https://doi.org/10.1214/17-AOS1615> MR3845014
- [15] GASSIAT, E. (2019). Mixtures of nonparametric components and hidden Markov models. In *Handbook of mixture analysis. Chapman & Hall/CRC Handb. Mod. Stat. Methods* 343–360. CRC Press, Boca Raton, FL. MR3889699
- [16] GASSIAT, E., CLEYNEN, A. and ROBIN, S. (2016). Inference in finite state space non parametric hidden Markov models and applications. *Stat. Comput.* **26** 61–71. <https://doi.org/10.1007/s11222-014-9523-8> MR3439359
- [17] GASSIAT, E. and ROUSSEAU, J. (2016). Non parametric finite translation mixtures with dependent regime. *Bernoulli* **22**. <https://doi.org/10.3150/14-BEJ631>
- [18] GHASSEMPOUR, S., GIROSI, F. and MAEDER, A. (2014). Clustering Multivariate Time Series Using Hidden Markov Models. *International Journal of Environmental Research and Public Health* **3**. <https://doi.org/10.3390/ijerph110302741>
- [19] GIRAUD, C. (2021). *Introduction to high-dimensional statistics*, Second ed. Chapman and Hall/CRC.
- [20] GIRAUD, C. and VERZELEN, N. (2018). Partial recovery bounds for clustering with the relaxed K-means. *Math. Stat. Learn.* **3** 317–374.
- [21] GRÜN, B. (2019). Model-based clustering. In *Handbook of mixture analysis. Chapman & Hall/CRC Handb. Mod. Stat. Methods* 157–192. CRC Press, Boca Raton, FL. MR3889693
- [22] HUNTER, D. R., WANG, S. and HETTMANSPERGER, T. P. (2007). Inference for mixtures of symmetric distributions. *Ann. Statist.* **35** 224 – 251. <https://doi.org/10.1214/009053606000001118>
- [23] KANNAN, R., SALMASIAN, H. and VEMPALA, S. (2005). The spectral method for general mixture models. In *International conference on computational learning theory* 444–457. Springer.
- [24] KHIATANI, D. and GHOSE, U. (2017). Weather forecasting using Hidden Markov Model. In *2017 International Conference on Computing and Communication Technologies for Smart Nation (IC3TSN)* 220–225. <https://doi.org/10.1109/IC3TSN.2017.8284480>

- [25] LAMBERT, M. F., WHITING, J. P. and METCALFE, A. V. (2003). A non-parametric hidden Markov model for climate state identification. *Hydrology and Earth System Sciences* **7** 652–667. <https://doi.org/10.5194/hess-7-652-2003>
- [26] LEHÉRICY, L. (2018). State-by-state minimax adaptive estimation for nonparametric hidden Markov models. *J. Mach. Learn. Res.* **19** Paper No. 39, 46. [MR3862446](https://doi.org/10.1214/21-ejs1890)
- [27] LEHÉRICY, L. (2021). Nonasymptotic control of the MLE for misspecified nonparametric hidden Markov models. *Electron. J. Stat.* **15** 4916–4965. <https://doi.org/10.1214/21-ejs1890> [MR4347381](https://doi.org/10.1214/21-ejs1890)
- [28] LU, Y. and H. ZHOU, H. (2016). Statistical and Computational Guarantees of Lloyd’s Algorithm and Its Variants. *arXiv preprint*. <https://doi.org/10.48550/arXiv.1612.02099>
- [29] MARANDON, A., REBAFKA, T. and ROQUAIN, E. (2023). False clustering rate control in mixture models. *arXiv preprint*. <https://doi.org/10.48550/arXiv.2203.02597>
- [30] MEILĂ, M. (2005). Comparing clusterings: an axiomatic view. In *Proceedings of the 22nd international conference on Machine learning* 577–584.
- [31] MEILĂ, M. and HECKERMAN, D. (2001). An experimental comparison of model-based clustering methods. *Machine learning* **42** 9–29.
- [32] MIXON, D. G., VILLAR, S. and WARD, R. (2017). Clustering subgaussian mixtures by semidefinite programming. *Information and Inference: A Journal of the IMA* **6** 389–415.
- [33] NDAOUD, M. (2022). Sharp optimal recovery in the two component Gaussian mixture model. *Ann. Statist.* **50** 2096 – 2126. <https://doi.org/10.1214/22-AOS2178>
- [34] PAULIN, D. (2015). Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electron. J. Prob.* **20** 1 – 32. <https://doi.org/10.1214/EJP.v20-4039>
- [35] SCHLIEP, A., SCHÖNHUTH, A. and STEINHOFF, C. (2003). Using hidden Markov models to analyze gene expression time course data. *Bioinformatics*. <https://doi.org/10.1093/bioinformatics/btg1036>
- [36] SCULLY, Z. (2024). Maximizing the expected deviation of sum of independent Bernoulli. *Mathematics Stack exchange*.
- [37] SHI, T., BELKIN, M. and YU, B. (2009). Data spectroscopy: Eigenspaces of convolution operators and clustering. *The Annals of Statistics* 3960–3984.
- [38] ZHANG, A. Y. and ZHOU, H. H. (2016). Minimax rates of community detection in stochastic block models. *Ann. Statist.* **44** 2252–2280. <https://doi.org/10.1214/15-AOS1428> [MR3546450](https://doi.org/10.1214/15-AOS1428)