



HAL
open science

JOINT RANK AND VARIABLE SELECTION FOR PARSIMONIOUS ESTIMATION IN A HIGH-DIMENSIONAL FINITE MIXTURE REGRESSION MODEL

Emilie Devijver

► **To cite this version:**

Emilie Devijver. JOINT RANK AND VARIABLE SELECTION FOR PARSIMONIOUS ESTIMATION IN A HIGH-DIMENSIONAL FINITE MIXTURE REGRESSION MODEL. 2017. hal-01099296v2

HAL Id: hal-01099296

<https://hal.science/hal-01099296v2>

Preprint submitted on 14 Feb 2017

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

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

JOINT RANK AND VARIABLE SELECTION FOR PARSIMONIOUS ESTIMATION IN A HIGH-DIMENSIONAL FINITE MIXTURE REGRESSION MODEL

EMILIE DEVIJVER

ABSTRACT. We study a dimensionality reduction technique for finite mixtures of high-dimensional multivariate response regression models. Both the dimension of the response and the number of predictors are allowed to exceed the sample size. We consider predictor selection and rank reduction to obtain lower-dimensional approximations. A class of estimators with a fast rate of convergence is introduced. We apply this result to a specific procedure, introduced in [11], where the relevant predictors are selected by the Group-Lasso.

Finite mixture regression ; Low rank estimation ; Non-asymptotic penalized criterion ; Oracle inequality ; Variable selection

62H30

CONTENTS

1. Introduction	1
2. The model and the collection of models	3
2.1. Mixture of linear regression models	3
2.2. Framework and collection of models	3
3. Oracle inequality	5
3.1. Kullback–Leibler and Jensen–Kullback–Leibler divergences	5
3.2. Oracle inequality for a general procedure to select relevant columns	6
4. Lasso-Rank procedure	7
4.1. Description of the main steps	8
4.2. Oracle inequality for relevant columns selected by the Group-Lasso estimator	9
4.3. Use of the slope heuristic	9
References	9
References	9
Appendix	10

1. INTRODUCTION

The popular multivariate response regression model $\mathbf{Y} = \beta\mathbf{X} + \varepsilon$ postulates a linear relationship between a $q \times n$ matrix \mathbf{Y} containing q responses for n subjects and a $p \times n$ matrix \mathbf{X} on p predictor variables. The responses for individual $i \in \{1, \dots, n\}$ are stored in the vector $\mathbf{Y}_i \in \mathbb{R}^q$ and the corresponding predictors are stored in the vector $\mathbf{X}_i \in \mathbb{R}^p$. In this model, the error term ε is an unobserved $q \times n$ matrix with independent columns $\varepsilon_1, \dots, \varepsilon_n \sim \mathcal{N}_q(0, \Sigma)$. The $q \times p$ regression matrix β and the $q \times q$ covariance matrix Σ are parameters that need to be estimated. In practice, however, this model is often too simple for heterogeneous data.

This paper considers a finite mixture of K linear multivariate regression models. If the subject i originates from the class $k \in \{1, \dots, K\}$, we assume that there exists an unknown matrix of coefficients $\beta_k \in \mathbb{R}^{q \times p}$ and an unknown covariance matrix $\Sigma_k \in \mathbb{S}_q^{++}$ such that the target variable \mathbf{Y}_i results from a regression model $\mathbf{Y}_i = \beta_k \mathbf{X}_i + \varepsilon_i$, with $\varepsilon_i \sim \mathcal{N}_q(0, \Sigma_k)$. If we use this model to study complex data with a high number of responses and predictors, the number of parameters could be quickly larger than the sample size. In order to deal with this issue, we propose to work with parsimonious models combining two well-known approaches.

The first approach consists in selecting relevant variables. The Lasso estimator, introduced by Tibshirani [22], is a classical tool. In multiple multivariate regression models, it seems reasonable to take into account the matrix structure. For this purpose, we focus on row sparsity using the Group-Lasso estimator, introduced by Yuan and Lin [23]. In our context, the groups correspond to the rows. Thus if, for $k \in \{1, \dots, K\}$, a matrix β_k has $|J|$ relevant rows, there are $|J|q$ coefficients to estimate instead of pq per cluster. If $|J| \ll p$, the number of parameters to estimate is then drastically reduced. Moreover, by selecting only the few relevant rows to describe the response, row sparsity may enhance interpretation.

The second approach is based on rank sparse models, introduced by Anderson [2]. A linear model is rank sparse if the regression matrix has low rank or can at least be well approximated by a low-rank matrix. In our model, for each $k \in \{1, \dots, K\}$, if the matrix β_k has rank \mathbf{R}_k , β_k is fully determined by $\mathbf{R}_k(p + q - \mathbf{R}_k)$ coefficients, which may be smaller than the sample size nq . These low-rank regression matrices appear in many applications, e.g., analysis of fMRI image data [13] and the analysis of EEG data decoding [2]. Moreover, note that low-rank estimation generalizes principal component analysis, often used to reduce the dimension of multivariate data.

Bunea et al. [8] mix these two estimators to provide a row and rank sparse estimator for the linear model. In practice, row and rank sparse approaches are combined to benefit from the two methods: fewer parameters to estimate, better interpretation thanks to the selection of relevant variables, and better estimation for data with low-rank structure. In [11], a method has been proposed, called the Lasso-Rank procedure, which combines both sparsity constraints for finite-mixture regression models. Heterogeneous data are then described by several clusters.

In this article, we focus on the theoretical aspects of combining row and rank sparse structures. Assuming such a structure for linear models, Bunea et al. [8] obtain substantial oracle rate improvements over the existing methods, especially when the estimated rank is low. In this paper, we give a similar result for finite-mixture regression models. Note that, in contrast to Bunea et al. [8], we do not need to assume that the covariance matrices $\Sigma_1, \dots, \Sigma_K$ are known. Therefore, our result generalizes theirs even in the case $K = 1$. Furthermore, our result does not assume strong assumptions on the design and on the true underlying density. We discuss its merits under classical assumptions and show that it can also serve to justify the model selection step of the Lasso-Rank procedure used to construct a collection of models.

For the general model, we construct a penalty satisfying a non asymptotic control of the risk of the model minimizing the penalized likelihood, among a random collection of row and rank sparse models. However, the penalty achieving the oracle type inequality is only known up to multiplicative constants. As these quantities depend on bounds on the parameters that are difficult to construct in practice, we propose to calibrate the required constants using slope heuristics.

Since the geometry of mixture models is complex, we use tools which are different from [8]. Our result exploits a generalization of a model selection theorem for maximum likelihood estimation in a regression framework developed by Cohen and Le Pennec in [9]. The main difficulty lies in the randomness of the collection of models under consideration. Indeed, as the number of responses

and predictors is potentially high, the whole collection of models is not tractable in practice. Thus, we do not look at every possible sparsity level but rather focus on a random subset constructed from the data. As an example, in the Lasso-Rank procedure, this random subset corresponding to the relevant rows is constructed using the Group-Lasso estimator along the regularization path. To solve this issue, we use an extension of the model selection theorem proved in [10].

The paper is organized as follows. In Section 2, we introduce finite mixtures of linear regression models, and we describe the collection of models we consider. In Section 3, we define the divergences we need, and we derive an oracle type inequality. Section 4 investigates the Lasso-Rank procedure, which is a specific procedure used to construct the random subcollection. All proofs are relegated to an Appendix.

2. THE MODEL AND THE COLLECTION OF MODELS

2.1. Mixture of linear regression models. Let $(\mathbf{X}_1, \mathbf{Y}_1), \dots, (\mathbf{X}_n, \mathbf{Y}_n)$ be a random sample, and let \mathbf{x}_i and \mathbf{y}_i denote the observed values of the random variables $\mathbf{X}_i \in \mathbb{R}^p$ and $\mathbf{Y}_i \in \mathbb{R}^q$, respectively. We assume that the response variable \mathbf{Y}_i depends on the set of explanatory variables \mathbf{X}_i through a regression-type model; its conditional probability density function, denoted s^* , is approximated by a mixture of K multivariate Gaussian regression models.

If $(\mathbf{X}_i, \mathbf{Y}_i)$ originates from an individual i in class k , we assume that there exists a matrix of class-specific regression coefficients $\boldsymbol{\beta}_k \in \mathbb{R}^{q \times p}$, a $q \times q$ positive-definite matrix $\boldsymbol{\Sigma}_k$, and a random vector $\boldsymbol{\varepsilon}_i \sim \mathcal{N}_q(0, \boldsymbol{\Sigma}_k)$, such that

$$\mathbf{Y}_i = \boldsymbol{\beta}_k \mathbf{X}_i + \boldsymbol{\varepsilon}_i.$$

We further assume that whatever the values of $\mathbf{x}_1, \dots, \mathbf{x}_n$, the variables $\mathbf{Y}_1 | \mathbf{X}_1 = \mathbf{x}_1, \dots, \mathbf{Y}_n | \mathbf{X}_n = \mathbf{x}_n$ are mutually independent. For each $i \in \{1, \dots, n\}$, we take $\mathbf{Y}_i | \mathbf{X}_i = \mathbf{x}_i \sim s_\xi^K(\mathbf{y} | \mathbf{x}_i)$, with

$$(1) \quad s_\xi^K(\mathbf{y} | \mathbf{x}) = \sum_{k=1}^K \frac{\boldsymbol{\pi}_k}{(2\pi)^{q/2} \det(\boldsymbol{\Sigma}_k)^{1/2}} \exp \left\{ -\frac{(\mathbf{y} - \boldsymbol{\beta}_k \mathbf{x})^\top \boldsymbol{\Sigma}_k^{-1} (\mathbf{y} - \boldsymbol{\beta}_k \mathbf{x})}{2} \right\},$$

where $\xi = (\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_K, \boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_K, \boldsymbol{\Sigma}_1, \dots, \boldsymbol{\Sigma}_K) \in \Xi_K$ with

$$(2) \quad \Xi_K = \Pi_K \times (\mathbb{R}^{q \times p})^K \times (\mathbb{S}_q^{++})^K.$$

In the above, $\boldsymbol{\pi}_k$ denotes the proportion for class $k \in \{1, \dots, K\}$. For each such class and all $z \in \{1, \dots, q\}$ and $\mathbf{x} = (\zeta_1, \dots, \zeta_p) \in \mathbb{R}^p$, $[\boldsymbol{\beta}_k \mathbf{x}]_z = \sum_{j=1}^p [\boldsymbol{\beta}_k]_{z,j} \zeta_j$ is the z th component of the k th mixture component. We set $\Pi_K = \{(\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_K) : \boldsymbol{\pi}_1 + \dots + \boldsymbol{\pi}_K = 1\} \subset (0, 1)^K$ and let \mathbb{S}_q^{++} denote the collection of symmetric positive definite matrices on \mathbb{R}^q .

Moreover, we assume that, for all $k \in \{1, \dots, K\}$, the covariance matrices $\boldsymbol{\Sigma}_k$ are diagonal. As q is potentially large, it reduces drastically the number of parameters to estimate. From a theoretical viewpoint, if q is not too large, the current result works for full covariance matrices, focusing on the eigenvalues. However, if q is large, we are unable to estimate these covariance matrices unless we make some additional assumptions. One way to allow for correlations might be to construct block-diagonal covariance matrices as in [12]. This strategy, however, will not be further considered here.

2.2. Framework and collection of models. In what follows, a row will be said to be relevant if every coefficient of the row is equal to 0 in every cluster. Otherwise, the row is said to be irrelevant. We denote by J the set of indices of relevant rows.

Given a model with K clusters, to combine rank and row sparsity, we consider models with rank $\mathbf{R} = (\mathbf{R}_1, \dots, \mathbf{R}_K)$ and relevant rows indexed by $j \in J$. We denote by $S_{(K,J,\mathbf{R})}$ the set of conditional

probability density functions corresponding to a mixture of K regression models, when focusing on $|J|$ relevant predictors indexed by $j \in J$, and with matrix of regression coefficients β_k of rank \mathbf{R}_k in class k , viz.

$$(3) \quad S_{(K,J,\mathbf{R})} = \{s_\xi^{(K,J,\mathbf{R})}\},$$

where

$$s_\xi^{(K,J,\mathbf{R})}(\mathbf{y}|\mathbf{x}) = \sum_{k=1}^K \frac{\pi_k}{(2\pi)^{q/2} \det(\Sigma_k)^{1/2}} \times \exp \left[-\frac{1}{2} \{ \mathbf{y} - (\beta_k^{\mathbf{R}_k})^{[J]} \mathbf{x} \}^\top \Sigma_k^{-1} \{ \mathbf{y} - (\beta_k^{\mathbf{R}_k})^{[J]} \mathbf{x} \} \right]$$

with $\xi = (\pi_1, \dots, \pi_K, \beta_1^{\mathbf{R}_1}, \dots, \beta_K^{\mathbf{R}_K}, \Sigma_1, \dots, \Sigma_K) \in \Xi_{(K,J,\mathbf{R})}$, $\Xi_{(K,J,\mathbf{R})} = \Pi_K \times \Psi_{(K,J,\mathbf{R})} \times (\mathbb{S}_q^{++})^K$, and

$$\Psi_{(K,J,\mathbf{R})} = \left\{ ((\beta_1^{\mathbf{R}_1})^{[J]}, \dots, (\beta_K^{\mathbf{R}_K})^{[J]}) \in (\mathbb{R}^{q \times p})^K \text{ such that } \forall_{k \in \{1, \dots, K\}} \text{rank}(\beta_k^{\mathbf{R}_k}) = \mathbf{R}_k \right\}.$$

Here, $A^{[J]}$ stands for the matrix A whose coefficients indexed by $j \in J$ are the nonzero coefficients. The dimension of the model $S_{(K,J,\mathbf{R})}$ is

$$D_{(K,J,\mathbf{R})} = \sum_{k=1}^K \{ \mathbf{R}_k(p + q - \mathbf{R}_k) + q + 1 \} - 1.$$

The complete collection of models consists of models with K clusters, row sparsity defined by J describing all subsets of $\{1, \dots, p\}$, and vector of ranks for a fixed K and fixed J belonging to $\{1, \dots, |J| \wedge q\}^K$, viz.

$$(4) \quad \mathcal{S} = \bigcup_{K \in \{1, 2, \dots\}} \bigcup_{J \in \mathcal{P}(\{1, \dots, p\})} \bigcup_{\mathbf{R} \in \{1, \dots, |J| \wedge q\}^K} S_{(K,J,\mathbf{R})} = \bigcup_{(K,J,\mathbf{R}) \in \mathcal{M}} S_{(K,J,\mathbf{R})}.$$

This collection is generally large and hence not tractable in practice. Therefore, we bound the number of clusters by restricting our attention to a finite set $\mathcal{K} \subset \{1, 2, \dots\}$. We further focus on a (potentially random) subcollection $\tilde{\mathcal{J}}$ of $\mathcal{P}(\{1, \dots, p\})$ of moderate size, and we reduce the number of possible vectors of ranks considered by introducing $\mathcal{R}_{(K,J)} \subset \{1, \dots, |J| \wedge q\}^K$, where in general, $a \wedge b = \min(a, b)$. Accordingly, we consider

$$(5) \quad \tilde{\mathcal{S}} = \bigcup_{K \in \mathcal{K}} \bigcup_{J \in \tilde{\mathcal{J}}} \bigcup_{\mathbf{R} \in \mathcal{R}_{(K,J)}} S_{(K,J,\mathbf{R})} = \bigcup_{(K,J,\mathbf{R}) \in \tilde{\mathcal{M}}} S_{(K,J,\mathbf{R})}.$$

In order to derive theoretical results, we further assume that the parameters are bounded. Our bounds are expressed in terms of the singular values of β_k and the eigenvalues of the covariance matrices $\Sigma_1, \dots, \Sigma_K$. If we fix K , the number of components, the relevant columns J and the

vector of ranks \mathbf{R} , we define the bounded model by

$$(6) \quad S_{(K,J,\mathbf{R})}^{\mathcal{B}} = \left\{ s_{\xi}^{(K,J,\mathbf{R})} \in S_{(K,J,\mathbf{R})} : \xi = (\boldsymbol{\pi}, \boldsymbol{\beta}, \boldsymbol{\Sigma}), \text{ such that} \right. \\ \left. \begin{aligned} \forall_{k \in \{1, \dots, K\}} \boldsymbol{\Sigma}_k &= \text{diag}([\boldsymbol{\Sigma}_k]_{1,1}, \dots, [\boldsymbol{\Sigma}_k]_{q,q}), \forall_{z \in \{1, \dots, q\}} a_{\boldsymbol{\Sigma}} \leq [\boldsymbol{\Sigma}_k]_{z,z} \leq A_{\boldsymbol{\Sigma}}, \\ \forall_{k \in \{1, \dots, K\}} \boldsymbol{\beta}_k^{\mathbf{R}_k} &= \sum_{r=1}^{\mathbf{R}_k} [\sigma_k]_r [u_k]_{\bullet, r} [v_k]_{r, \bullet}, \forall_{r \in \{1, \dots, \mathbf{R}_k\}} [\sigma_k]_r < A_{\sigma} \end{aligned} \right\}.$$

In the above, for $k \in \{1, \dots, K\}$, $[\sigma_k]_1, \dots, [\sigma_k]_{\mathbf{R}_k}$ denote the singular values of $\boldsymbol{\beta}_k^{\mathbf{R}_k}$, with corresponding unit vectors u_k and v_k . We are then led to consider the following collection of models:

$$(7) \quad \tilde{\mathcal{S}}^{\mathcal{B}} = \bigcup_{K \in \mathcal{K}} \bigcup_{J \in \tilde{\mathcal{J}}} \bigcup_{\mathbf{R} \in \mathcal{R}_{(K,J)}} S_{(K,J,\mathbf{R})}^{\mathcal{B}}.$$

Finally, we assume that the covariates \mathbf{X} belong to an hypercube. Without loss of generality, this hypercube is taken to be $\in [0, 1]^p$.

3. ORACLE INEQUALITY

In this section, we first introduce the Kullback–Leibler divergence and some extensions thereof that are useful to compare two conditional probability density functions. We then derive an oracle type inequality which ensures that if we have penalized the log-likelihood in an appropriate way, we select a model which is as good as the oracle.

3.1. Kullback–Leibler and Jensen–Kullback–Leibler divergences. We first need to define divergences used to compare conditional probability density functions. The Kullback–Leibler divergence is defined, for any two probability density functions s and t , by

$$(8) \quad \text{KL}(s, t) = \begin{cases} \int_{\mathbb{R}^q} \ln \left\{ \frac{s(y)}{t(y)} \right\} s(y) dy & \text{if } s dy \ll t dy, \\ \infty & \text{otherwise.} \end{cases}$$

To deal with regression data, for observed covariates $\mathbf{x}_1, \dots, \mathbf{x}_n$, we define

$$(9) \quad \text{KL}^{\otimes n}(s, t) = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \text{KL}\{s(\cdot | \mathbf{x}_i), t(\cdot | \mathbf{x}_i)\} \right]$$

for any two conditional probability density functions s and t .

For any two probability density functions s and t , we also define the Jensen–Kullback–Leibler divergence, first introduced in [9], by

$$\text{JKL}_{\rho}(s, t) = \frac{1}{\rho} \text{KL}\{s, (1 - \rho)s + \rho t\}$$

for any $\rho \in (0, 1)$. The tensorized version is defined by

$$\text{JKL}_{\rho}^{\otimes n}(s, t) = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \text{JKL}_{\rho}\{s(\cdot | \mathbf{x}_i), t(\cdot | \mathbf{x}_i)\} \right]$$

for any two conditional probability density functions s and t . We use the Jensen–Kullback–Leibler divergence rather than the Kullback–Leibler divergence because it is bounded. This boundedness turns out to be crucial to control the loss of the penalized maximum likelihood estimator under mild assumptions on the complexity of the model and on the parameters.

Note that the above divergences are not metrics because they do not satisfy the triangular inequality and they are not symmetric. Nevertheless, they are very useful for comparing two probability density functions.

3.2. Oracle inequality for a general procedure to select relevant columns. Among the collection of models $\tilde{\mathcal{S}}^{\mathcal{B}}$ defined in (7), with various ranks and various level of sparsities, we want to select a model which is close to the best one. The oracle is by definition the model belonging to the collection which minimizes the contrast with the true model (here, the contrast considered is the log-likelihood).

In practice, we cannot identify the true model and hence the oracle is unknown. Nevertheless, we could select a model close to the oracle. More precisely, we aim at selecting, among $\tilde{\mathcal{S}}^{\mathcal{B}}$, the optimal set of predictors J , the optimal number of clusters K , and the optimal vector of ranks \mathbf{R} . First, for each model defined by $(K, J, \mathbf{R}) \in \tilde{\mathcal{M}}$, we consider the conditional probability density function $\hat{s}^{(K, J, \mathbf{R})}$ where the parameters are estimated by the maximum likelihood estimator. Among all $(K, J, \mathbf{R}) \in \tilde{\mathcal{M}}$, we want to select the density $\hat{s}^{(\hat{K}, \hat{J}, \hat{\mathbf{R}})}$ which is the closest one to the true distribution s^* according to the divergences defined in Section 3.1.

Even if the oracle is not attainable in practice, it turns out that we can do almost as well as the oracle. This result, which is our main finding, is formally stated below and proved in the Appendix.

Theorem 3.1. *Let $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n) \in [0, 1]^p \times \mathbb{R}^q$ be the observations arising from an unknown conditional probability density function s^* . Let \mathcal{M} denotes the set of every potential model and $\tilde{\mathcal{M}}$ the restricted set, as defined in (4) and (5), respectively.*

Assume that there exist $\tau > 0$ and $\delta_{\text{KL}} > 0$ such that, for all $(K, J, \mathbf{R}) \in \mathcal{M}$, one can find $\bar{s}^{(K, J, \mathbf{R})} \in S_{(K, J, \mathbf{R})}^{\mathcal{B}}$ such that

$$(10) \quad \text{KL}^{\otimes n}(s^*, \bar{s}^{(K, J, \mathbf{R})}) \leq \inf_{t \in S_{(K, J, \mathbf{R})}^{\mathcal{B}}} \text{KL}^{\otimes n}(s^*, t) + \frac{\delta_{\text{KL}}}{n} \quad \text{and} \quad \bar{s}^{(K, J, \mathbf{R})} \geq e^{-\tau} s^*.$$

For $(K, J, \mathbf{R}) \in \mathcal{M}$, let $\hat{s}^{(K, J, \mathbf{R})}$ denote the maximum likelihood estimator for the model $S_{(K, J, \mathbf{R})}^{\mathcal{B}}$, satisfying

$$\hat{s}^{(K, J, \mathbf{R})} = \underset{s_{\xi}^{(K, J, \mathbf{R})} \in S_{(K, J, \mathbf{R})}^{\mathcal{B}}}{\text{argmin}} \left[-\frac{1}{n} \sum_{i=1}^n \ln \left\{ s_{\xi}^{(K, J, \mathbf{R})}(\mathbf{y}_i | \mathbf{x}_i) \right\} \right].$$

Then the following statements hold true.

a) *There exist absolute constants $\kappa > 0$ and C such that whenever*

$$(11) \quad \begin{aligned} & \text{pen}(K, J, \mathbf{R}) \\ & \geq \kappa \frac{D(K, J, \mathbf{R})}{n} \left[2B^2(A_{\sigma}, A_{\Sigma}, a_{\Sigma}) - \ln \left\{ \frac{D(K, J, \mathbf{R})}{n} B^2(A_{\sigma}, A_{\Sigma}, a_{\Sigma}) \wedge 1 \right\} \right] \\ & \quad + (1 \vee \tau) \ln \left\{ \frac{4epq}{(K(1 + |J|) - q^2) \wedge pq} + \sum_{k=1}^K \mathbf{R}_k \right\} \end{aligned}$$

for every $(K, J, \mathbf{R}) \in \mathcal{M}$, with $B(A_\sigma, A_\Sigma, a_\Sigma)$ a constant depending only on parameter bounds, there exists a random vector $(\hat{K}, \hat{J}, \hat{\mathbf{R}}) \in \tilde{\mathcal{M}}$ such that

$$(\hat{K}, \hat{J}, \hat{\mathbf{R}}) = \operatorname{argmin}_{(K, J, \mathbf{R}) \in \tilde{\mathcal{M}}} \left[-\frac{1}{n} \sum_{i=1}^n \ln \{ \hat{s}^{(K, J, \mathbf{R})}(\mathbf{y}_i | \mathbf{x}_i) \} + \operatorname{pen}(K, J, \mathbf{R}) \right].$$

b) Whatever the true density s^* ,

$$(12) \quad \mathbb{E} \{ \operatorname{JKL}_{\rho^n}^{\otimes n}(s^*, \hat{s}^{(\hat{K}, \hat{J}, \hat{\mathbf{R}})}) \} \\ \leq \mathbb{C} \mathbb{E} \left[\inf_{(K, J, \mathbf{R}) \in \tilde{\mathcal{M}}} \left\{ \inf_{t \in S_{(K, J, \mathbf{R})}^{\mathcal{B}}} \operatorname{KL}^{\otimes n}(s^*, t) \right. \right. \\ \left. \left. + \operatorname{pen}(K, J, \mathbf{R}) \right\} + \frac{(1 \vee \tau)}{n} \right].$$

This theorem is deduced from an adaptation of a model selection theorem for maximum likelihood estimators developed by Massart [18]. This adaptation, derived in [11], allows us to focus on regression models (see also [9]) and to focus on a random subcollection of the whole collection of models.

In Theorem 3.1, the critical assumptions are the control of the bracketing entropy of each model in the whole collection of models and the construction of weights for each model to control the complexity of the collection of models. Assumption (10) on the true density s^* is needed because we consider a random subcollection of models $\tilde{\mathcal{S}}^{\mathcal{B}}$ from the whole collection of models $\mathcal{S}^{\mathcal{B}}$. Thanks to this assumption, we can invoke Bernstein's inequality to control the additional randomness. The parameter τ depends on the true unknown density s^* and cannot be explicitly determined for this reason. Specific cases can be discussed, e.g., the true conditional probability density function s^* should not be bounded along \mathbf{Y} to allow (10). Moreover, if the true density is Gaussian, we need to bound the variance of the models from below, which is related to the collection of models constructed in (6).

While we could make appropriate assumptions on the true density s^* in order to determine τ explicitly, we choose not to do so. Under the assumption that the Kullback–Leibler divergence and the Hellinger distance are equivalent, we can explicitly determine τ . Then, the result is still meaningful as the equivalence between the Kullback–Leibler divergence and Hellinger's distance is well known. Note that the parameter τ appears in the oracle type inequality (12), but also in the penalty term (11). We may construct a larger penalty independent of τ , achieving an oracle-type inequality, but the rate will be larger. Details are discussed in the Appendix.

For all $(K, J, \mathbf{R}) \in \mathcal{M}$, the mixture parameters are bounded in order to construct brackets over $S_{(K, J, \mathbf{R})}^{\mathcal{B}}$, and thus to bound the entropy from above.

Because we consider a random sub-collection of models which is small enough, our estimator $\hat{s}^{(K, J, \mathbf{R})}$ is tractable in practice. Moreover, we remark that this is a finite sample result, which allows us to study cases for which p increases with n .

4. LASSO-RANK PROCEDURE

In this section, we describe briefly the Lasso-Rank procedure introduced in [11], and we derive a corollary of Theorem 3.1 in this specific context.

4.1. Description of the main steps. The Lasso-Rank procedure is a method used to construct a collection of models as defined in (5). We recall here the main steps of this procedure. To this end, we assume that we have access to a set of potential number of components \mathcal{K} .

A) Row selection. Fix $K \in \mathcal{K}$. To detect the relevant rows, the empirical contrast is penalized by a group-penalty on the conditional mean parameters proportional to

$$\|\mathbf{P}_k^{-1}\boldsymbol{\beta}_k\|_{\text{gr}} = \sum_{j=1}^p \sqrt{q} \sqrt{\sum_{z=1}^q [\mathbf{P}_k^{-1}\boldsymbol{\beta}_k]_{z,j}},$$

where the Cholesky decomposition $\mathbf{P}_k^\top \mathbf{P}_k = \boldsymbol{\Sigma}_k^{-1}$ defines \mathbf{P}_k for all $k \in \{1, \dots, K\}$. We then consider, for some choice of regularization parameter $\lambda > 0$,

$$(13) \quad \hat{\xi}^{\text{GL}}(\lambda) = \underset{\xi=(\boldsymbol{\pi}, \boldsymbol{\beta}, \boldsymbol{\Sigma}) \in \Xi_K}{\text{argmin}} \left[-\frac{1}{n} \sum_{i=1}^n \ln\{s_\xi^K(\mathbf{y}_i|\mathbf{x}_i)\} + \lambda \sum_{k=1}^K \pi_k \|\mathbf{P}_k^{-1}\boldsymbol{\beta}_k\|_{\text{gr}} \right]$$

with Ξ_K defined in (2).

We remark that the penalty takes into account the mixture weight as, e.g., introduced in [15]. This is in line with the common practice in relating the penalty to the sample size. We refer the interested reader to Section 7.1.4 of [21] for an interesting empirical discussion about a penalty similar to this one, focusing on this mixture weight.

To get models that are more or less sparse, we compute the Group-Lasso estimator defined in (13) with several regularization parameters λ , and deduce sets of relevant rows, denoted by $\hat{J}_{(K,\lambda)}$ for all λ . We denote by $\mathcal{J}^{\text{GL}} = \cup_\lambda \hat{J}_{(K,\lambda)}$.

Remark 1. Notice that here, the Group-Lasso penalty is used to select a subset of variables for one choice of regularization parameter in the Lasso-Rank procedure. One could also use the Group-Lasso penalty to get a ranking of the variables, as done, e.g., in Bach [4].

B) Refitting. We estimate the parameters by maximum likelihood estimation under rank constraint on the restricted set of relevant rows.

C) Model selection. Varying $K \in \mathcal{K}$, $J \in \mathcal{J}^{\text{GL}}$, $\mathbf{R} \in \mathcal{R}_{(K,J)}$, we obtain a collection of models

$$\mathcal{S} = \bigcup_{K \in \mathcal{K}} \bigcup_{J \in \mathcal{J}^{\text{GL}}} \bigcup_{\mathbf{R} \in \mathcal{R}_{(K,J)}} S_{(K,J,\mathbf{R})}.$$

In practice, in [11], the slope heuristic is used, which is a data-driven non-asymptotic criterion. It considers a penalized criterion, where the penalty is proportional to the dimension. There are two ways to calibrate this proportionality constant, but both consider the minimal coefficient κ_{\min} , the optimal one κ_{opt} being approximated by twice the minimal coefficient κ_{\min} .

One calibration method is the *Slope Heuristic Dimension Jump* (SHDJ): κ_{\min} corresponds to the largest dimension jump on the graph representing the model dimension as a function of the coefficient κ . Another method is the *Slope Heuristic Robust Regression* (SHRR): κ_{\min} corresponds to the slope of a robust regression performed between the log-likelihood and the model dimension for complex models.

The two methods are derived from the same heuristic and they offer two different visual checks of the adequacy of the model selection procedure to the data. They should select the same model. For an implementation of this approach, see Baudry et al. [6] and the `Capushe` package.

4.2. Oracle inequality for relevant columns selected by the Group-Lasso estimator.

In this section, we comment the use of Theorem 3.1 for this specific Lasso-Rank procedure. In general, we are not able to combine this theorem with results on the Group-Lasso to bound the oracle term, and then to bound the divergence between the model selected and the true conditional probability density function. However, in some specific contexts, we can control the oracle.

For example, if s^* corresponds to a linear regression with sparse rank and a few selected row, then under some conditions on the design matrix, one can prove that this procedure is adaptive to row and rank sparsity. This kind of result is discussed in Section 2 of Bunea et al. [8], and we refer to Lounici et al. [17] for results on the Group-Lasso estimator. However, our theorem stands non asymptotically and without any assumptions on the design. We have proved that we do the best possible within the collection, although specific conditions are required to prove that the best model within the random collection is indeed objectively good. There is here a dirth of theoretical tools and this challenging topic deserves further consideration.

4.3. Use of the slope heuristic. From a theoretical point of view, we want to ensure that the slope heuristic which penalizes the log-likelihood with the model dimension selects a good model. We choose a penalty

$$\text{pen}(K, J, R) = C \frac{D_{(K,J,R)}}{n} + \tilde{C} \frac{D_{(K,J,R)}}{n} \ln \left(\frac{D_{(K,J,R)}}{n} \right)$$

with explicit constants C and \tilde{C} depending on absolute constants and on some bounds on the parameters. In practice, however, these bounds are not tractable, and so we prefer to calibrate C and \tilde{C} from the data using the slope heuristic. Moreover, we consider a simple version of the penalty term, viz.

$$\text{pen}(K, J, R) = C \frac{D_{(K,J,R)}}{n}$$

as already proposed in [12] and [16]. The calibration of C from data has originally been proposed and validated in the context of heteroscedastic regression with fixed design [5, 7], and for homoscedastic regression with fixed design [3]. In this paper, we proved that the shape of an optimal penalty is close to the one used in the slope heuristic, up to a logarithmic term.

Acknowledgments. I am grateful to Pascal Massart and Christophe Giraud for stimulating discussions. I also thank the referees, the Associate Editor and the Editor-in-Chief, Christian Genest, for comments and suggestions which led to substantial improvements. This research was supported in part by the Interuniversity Attraction Poles Program (IAP-network P7/06), Belgian Science Policy Office, whose funding is gratefully acknowledged.

REFERENCES

REFERENCES

- [1] H. Akaike, A new look at the statistical model identification, *IEEE Trans. Automatic Control* 19 (1974) 716–723.
- [2] C.W. Anderson, E.A. Stolz, S. Shamsunder, Multivariate autoregressive models for classification of spontaneous electroencephalographic signals during mental tasks, *IEEE Trans. Biomedical Eng.* 45 (1998) 277–286.
- [3] S. Arlot, P. Massart, Data-driven calibration of penalties for least-squares regression, *J. Machine Learning Res.* 10 (2009) 245–279.
- [4] F.R. Bach, Bolasso: Model consistent lasso estimation through the bootstrap. In: *Proceedings of the 25th International Conference on Machine Learning, ICML’08*, pp. 33–40, New York, 2008.

- [5] Y. Baraud, C. Giraud, S. Huet, Gaussian model selection with an unknown variance, *Ann. Statist.* 37 (2009) 630–672.
- [6] J-P Baudry, C. Maugis, B. Michel, Slope heuristics: Overview and implementation, *Statist. Comput.* 22 (2012) 455–470.
- [7] L. Birgé and P. Massart, Minimal penalties for Gaussian model selection, *Prob. Theory Relat. Fields* 138 (2007) 33–73.
- [8] F. Bunea, Y. She, M.H. Wegkamp, Joint variable and rank selection for parsimonious estimation of high-dimensional matrices, *Ann. Statist.* 40 (2012) 2359–2388.
- [9] S. Cohen, E. Le Pennec, Partition-based conditional density estimation, *ESAIM: Probab. Statist.* 17 (2013) 672–697.
- [10] E. Devijver, Finite mixture regression: A sparse variable selection by model selection for clustering, *Electr. J. Statist.* 9 (2015) 2642–2674.
- [11] E. Devijver, Model-based clustering for high-dimensional data: Application to functional data, *Adv. Data Anal. Class.*, in press.
- [12] E. Devijver, M. Gallopin, Block-diagonal covariance selection for high-dimensional gaussian graphical models, *J. Amer. Statist. Assoc.*, in press.
- [13] K.J. Friston, L. Harrison, W.D. Penny. Dynamic causal modelling, *NeuroImage* 19 (2003) 1273–1302.
- [14] C.R. Genovese, L.A. Wasserman, Rates of convergence for the gaussian mixture sieve, *Ann. Statist.* 28 (2000) 1105–1127.
- [15] A. Khalili, J. Chen, Variable selection in finite mixture of regression models, *J. Amer. Statist. Assoc.* 102 (2007) 1025–1038.
- [16] E. Lebarbier, T. Mary-Huard, Le critère BIC: Fondements théoriques et interprétation, *J. Soc. Fr. Statist.* 147 (2006) 39–57.
- [17] K. Lounici, M. Pontil, S. van de Geer, A.B. Tsybakov, Oracle inequalities and optimal inference under group sparsity, *Ann. Statist.* 39 (2011) 2164–2204.
- [18] P. Massart, Concentration inequalities and model selection, In: *Ecole d’été de probabilités de Saint-Flour XXXIII–2003*, 2007.
- [19] C. Maugis, B. Michel, A non asymptotic penalized criterion for gaussian mixture model selection, *ESAIM: Probab. Statist.* 15 (2012) 41–68.
- [20] G. Schwarz, Estimating the dimension of a model, *Ann. Statist.* 6 (1978) 461–464.
- [21] N. Städler, P. Bühlmann, S. van de Geer, ℓ_1 -penalization for mixture regression models, *Test* 19 (2010) 209–256.
- [22] R.J. Tibshirani, Regression shrinkage and selection via the lasso, *J. Roy. Statist. Soc. Ser. B* 58 (1996) 267–288.
- [23] M. Yuan, Y. Lin, Model selection and estimation in regression with grouped variables, *J. Roy. Statist. Soc. Ser. B* 68 (2006) 49–67.

APPENDIX

This Appendix contains the details of the proof of Theorem 3.1. It is derived from a general model selection theorem stated in Part A. The proof of Theorem 3.1 can be derived by checking Assumptions H_m , Sep_m and K described in Parts B and C.

A. A general oracle type inequality for model selection

Model selection can be performed using either the AIC or the BIC criterion [1, 20]. An important limitation of these criteria, however, is that they are asymptotic. Their use in high-dimensional settings is thus ad hoc. The slope heuristic is a method which overcomes this difficulty. It was introduced by Birgé and Massart [7].

In [18], Massart proposed a model selection theorem which guarantees that a penalized criterion leads to a good model selection, the penalty being defined by the model complexity. It provides support for the slope heuristic approach in a finite sample setting. Cohen and Le Pennec [9] generalized this theorem to a regression framework. Here, we need a generalization of their result detailed in [10] because we consider a random collection of models $(S_m)_{m \in \mathcal{M}}$, indexed by \mathcal{M} .

Let $(\mathbf{X}, \mathbf{Y}) \in \mathcal{X} \times \mathcal{Y}$. We observe $\mathbf{x}_1, \dots, \mathbf{x}_n$. Let us start with a description of the required assumptions. First, we impose a structural assumption. It is a bracketing entropy condition on the model with respect to the tensorized Hellinger divergence

$$(d_H^{\otimes n})^2(s, t) = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n d_H^2\{s(\cdot|\mathbf{x}_i), t(\cdot|\mathbf{x}_i)\} \right]$$

for two conditional probability density functions s and t . A bracket $[\ell, u]$ is a pair of functions such that for all $(\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$, $\ell(\mathbf{y}, \mathbf{x}) \leq s(\mathbf{y}|\mathbf{x}) \leq u(\mathbf{y}, \mathbf{x})$. For $\varepsilon > 0$, the bracketing entropy $\mathcal{H}_{[\cdot]}(\varepsilon, S, d_H^{\otimes n})$ of a set S is defined as the logarithm of the minimum number of brackets $[\ell, u]$ of width $d_H^{\otimes n}(\ell, u)$ smaller than ε such that every probability density function of S belongs to one of these brackets. In what follows, m denotes any fixed element of \mathcal{M} .

Assumption H_m . *There exists a non-decreasing function ψ_m such that $\varpi \mapsto 1/\varpi\psi_m(\varpi)$ is non-increasing on $(0, \infty)$. Furthermore, for all $\varpi \in \mathbb{R}^+$ and $s_m \in S_m$,*

$$\int_0^\varpi \sqrt{\mathcal{H}_{[\cdot]}(\varepsilon, S_m(s_m, \varpi), d_H^{\otimes n})} d\varepsilon \leq \psi_m(\varpi),$$

where $S_m(s_m, \varpi) = \{t \in S_m : d_H^{\otimes n}(t, s_m) \leq \varpi\}$. The model complexity \mathcal{D}_m is then defined by $n\varpi_m^2$ with ϖ_m standing for the unique solution of

$$(14) \quad \frac{1}{\varpi} \psi_m(\varpi) = \sqrt{n}\varpi.$$

Observe that the model complexity does not depend on the bracketing entropies of the global models S_m , but rather on those of the smaller localized sets $S_m(s_m, \varpi)$. This is a weaker assumption.

For technical reasons, a separability assumption is also required. In what follows, λ denotes Lebesgue's measure.

Assumption Sep_m . *There exists a countable subset S'_m of S_m and a set \mathcal{Y}'_m with $\lambda(\mathcal{Y} \setminus \mathcal{Y}'_m) = 0$ such that for every $t \in S_m$, there exists a sequence $(t_\ell)_{\ell \geq 1}$ of elements of S'_m such that for every \mathbf{x} and every $\mathbf{y} \in \mathcal{Y}'_m$, $\ln\{t_\ell(\mathbf{y}|\mathbf{x})\} \rightarrow \ln\{t(\mathbf{y}|\mathbf{x})\}$ as $\ell \rightarrow \infty$.*

This assumption allows us to work with a countable subset. We also need an information theory type assumption on our collection of models. We assume the existence of a Kraft-type inequality for the collection.

Assumption K . *There exists a family $(w_m)_{m \in \mathcal{M}}$ of non-negative numbers such that*

$$\sum_{m \in \mathcal{M}} e^{-w_m} \leq \Omega < \infty.$$

We are now in a position to state our global theorem for a random subcollection of models in a regression framework.

Theorem A . *Assume we observe $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n) \in \mathcal{X} \times \mathcal{Y}$ with unknown conditional probability density function s^* . Let $\mathcal{S} = (S_m)_{m \in \mathcal{M}}$ be an at most countable collection of conditional probability density function sets. Suppose that Assumption K holds and that Assumptions H_m and*

Sep_m are valid for every model $S_m \in \mathcal{S}$. Further assume that there exist $\tau > 0$ and $\delta_{\text{KL}} > 0$ such that for all $m \in \mathcal{M}$, there is an element \bar{s}_m of S_m such that

$$(15) \quad \text{KL}^{\otimes n}(s^*, \bar{s}_m) \leq \inf_{t \in S_m} \text{KL}^{\otimes n}(s^*, t) + \frac{\delta_{\text{KL}}}{n} \quad \text{and} \quad \bar{s}_m \geq e^{-\tau} s^*.$$

Introduce $(S_m)_{m \in \tilde{\mathcal{M}}}$ some random sub-collection of $(S_m)_{m \in \mathcal{M}}$. For $m \in \tilde{\mathcal{M}}$, consider the collection \hat{s}_m of η -log-likelihood minimizer in S_m , satisfying

$$\sum_{i=1}^n -\ln\{\hat{s}_m(\mathbf{y}_i|\mathbf{x}_i)\} \leq \inf_{s_m \in S_m} \left[\sum_{i=1}^n -\ln\{s_m(\mathbf{y}_i|\mathbf{x}_i)\} \right] + \eta.$$

Then, for any $\rho \in (0, 1)$ and any $C_1 > 1$, there are two constants κ_0 and C_2 depending only on ρ and C_1 such that, as soon as for every index $m \in \mathcal{M}$,

$$(16) \quad \text{pen}(m) \geq \kappa\{\mathcal{D}_m + (1 \vee \tau)w_m\}$$

with $\kappa > \kappa_0$, and where the model complexity \mathcal{D}_m is defined through the equation (14), the penalized likelihood estimate $\hat{s}_{\hat{m}}$ with $\hat{m} \in \tilde{\mathcal{M}}$ such that

$$\sum_{i=1}^n -\ln\{\hat{s}_{\hat{m}}(\mathbf{y}_i|\mathbf{x}_i)\} + \text{pen}(\hat{m}) \leq \inf_{m \in \tilde{\mathcal{M}}} \left[\sum_{i=1}^n -\ln\{\hat{s}_m(\mathbf{y}_i|\mathbf{x}_i)\} + \text{pen}(m) \right] + \eta'$$

satisfies

$$(17) \quad \begin{aligned} \mathbb{E}\{\text{JKL}_{\rho}^{\otimes n}(s^*, \hat{s}_{\hat{m}})\} &\leq C_1 \mathbb{E} \left\{ \inf_{m \in \tilde{\mathcal{M}}} \inf_{t \in S_m} \text{KL}^{\otimes n}(s^*, t) + 2 \frac{\text{pen}(m)}{n} \right\} \\ &+ C_2 (1 \vee \tau) \frac{\Omega^2}{n} + \frac{\eta' + \eta}{n}. \end{aligned}$$

Under the conditions of this theorem, we are thus able to choose, within a random subcollection of models, a representative which is as good as the oracle, up to a constant C_1 , and some additive terms of order $1/n$. This result is non-asymptotic, and gives a theoretical penalty to select this model.

The proof of this theorem is detailed in [10]. Nevertheless, we sketch the main ideas here in order to give the reader a better grasp of the assumptions. From Assumptions H_m and Sep_m , one can call on maximal inequalities to control, except on a set whose probability is less than $e^{-w_{m'}-w}$ for all w , the ratio of the centered empirical process of $\ln(\hat{s}_{m'})$ over the Hellinger distance between s^* and $\hat{s}_{m'}$; the bound is of the order of $1/n$. Thanks to Bernstein's inequality, which is valid in view of inequality (15), we thus control the process over the random subcollection. Thanks to Assumption K, we sum up each event over the collection of models. It leads to the oracle type inequality (17).

Now, to prove Theorem 3.1, we have to satisfy Assumptions H_m and K. As for Assumption Sep_m , it holds for our conditional probability density functions.

B. Assumption H_m

B1. Decomposition of the mixture conditional probability density function. Proceeding as in [9], we can decompose the entropy by

$$(18) \quad \mathcal{H}_{[\cdot]}(\varepsilon, \mathcal{S}_{(K,J,\mathbf{R})}^{\mathcal{B}}, d_H^{\otimes n}) \leq \mathcal{H}_{[\cdot]}(\varepsilon, \Pi_K, d_H^{\otimes n}) + \sum_{k=1}^K \mathcal{H}_{[\cdot]}(\varepsilon, \mathcal{F}_{(J,\mathbf{R}_k)}, d_H^{\otimes n}),$$

where

$$\mathcal{S}_{(K,J,\mathbf{R})}^{\mathcal{B}} = \left\{ \begin{array}{l} \mathbf{y} \in \mathbb{R}^q | \mathbf{x} \in \mathbb{R}^p \mapsto s_{\xi}^{(K,J,\mathbf{R})}(\mathbf{y} | \mathbf{x}), \\ s_{\xi}^{(K,J,\mathbf{R})}(\mathbf{y} | \mathbf{x}) = \sum_{k=1}^K \pi_k \varphi(\mathbf{y} | (\boldsymbol{\beta}_k^{\mathbf{R}_k})^{[J]} \mathbf{x}, \boldsymbol{\Sigma}_k), \\ \xi = (\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_K, (\boldsymbol{\beta}_1^{\mathbf{R}_1})^{[J]}, \dots, (\boldsymbol{\beta}_K^{\mathbf{R}_K})^{[J]}, \boldsymbol{\Sigma}_1, \dots, \boldsymbol{\Sigma}_K), \\ \xi \in \Xi_{(K,J,\mathbf{R})} = \Pi_K \times \tilde{\Psi}_{(K,J,\mathbf{R})} \times ([a_{\boldsymbol{\Sigma}}, A_{\boldsymbol{\Sigma}}]^q)^K \end{array} \right\},$$

$$\Psi_{(K,J,\mathbf{R})} = \{((\boldsymbol{\beta}_1^{\mathbf{R}_1})^{[J]}, \dots, (\boldsymbol{\beta}_K^{\mathbf{R}_K})^{[J]}) \in (\mathbb{R}^{q \times p})^K : \text{rank}(\boldsymbol{\beta}_k) = \mathbf{R}_k\},$$

$$\tilde{\Psi}_{(K,J,\mathbf{R})} = \left\{ \begin{array}{l} ((\boldsymbol{\beta}_1^{\mathbf{R}_1})^{[J]}, \dots, (\boldsymbol{\beta}_K^{\mathbf{R}_K})^{[J]}) \in \Psi_{(K,J,\mathbf{R})} : \text{for all } k \in \{1, \dots, K\}, \\ \boldsymbol{\beta}_k^{\mathbf{R}_k} = \sum_{r=1}^{\mathbf{R}_k} \sigma_r [\mathbf{u}]_{r,\bullet}^{\top} [\mathbf{v}]_{r,\bullet}, \text{ with } \sigma_r < A_{\sigma} \text{ for all } r \in \{1, \dots, \mathbf{R}_k\} \end{array} \right\},$$

$$\Pi_K = \{(\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_K) : \boldsymbol{\pi}_1 + \dots + \boldsymbol{\pi}_K = 1\} \subset (0, 1)^K \text{ and}$$

$$\mathcal{F}_{(J,R)} = \left\{ \begin{array}{l} \varphi(\cdot | (\boldsymbol{\beta}^R)^{[J]} \mathbf{X}, \boldsymbol{\Sigma}) : \boldsymbol{\beta}^R = \sum_{r=1}^R \sigma_r [\mathbf{u}]_{r,\bullet}^{\top} [\mathbf{v}]_{r,\bullet}, \text{ with } \sigma_r < A_{\sigma}, \\ \boldsymbol{\Sigma} = \text{diag}(\boldsymbol{\Sigma}_{1,1}, \dots, \boldsymbol{\Sigma}_{q,q}) \in [a_{\boldsymbol{\Sigma}}, A_{\boldsymbol{\Sigma}}]^q \end{array} \right\},$$

where φ denotes the Gaussian density. For the proportions, it is known (see, e.g., [14]) that

$$\mathcal{H}_{[\cdot]}(\varepsilon, \Pi_K, d_H^{\otimes n}) \leq \ln \left\{ K(2\pi e)^{K/2} \left(\frac{3}{\varepsilon} \right)^{K-1} \right\}.$$

B.2 Entropy of the Gaussian probability density functions with low rank. We need to bound the integrated entropy. To this end, first we have to construct some brackets to recover $\mathcal{F}_{(J,R)}$ for fixed rank R and arbitrary set J of relevant rows. Fix $f \in \mathcal{F}_{(J,R)}$. We are looking for functions ℓ and u such that $\ell \leq f \leq u$. Because f is Gaussian, we will choose ℓ and u as Gaussian dilatations. Therefore, we need only to determine the mean, the variance and the dilatation coefficients of ℓ and u . To construct these brackets, we rely on the following two lemmas; see [19] for details.

Lemma A *Let $\varphi(\cdot | \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ and $\varphi(\cdot | \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$ be two Gaussian probability density functions. Suppose that their covariance matrices are diagonal, with $\boldsymbol{\Sigma}_a = \text{diag}([\boldsymbol{\Sigma}_a]_{1,1}, \dots, [\boldsymbol{\Sigma}_a]_{q,q})$ for $a \in \{1, 2\}$, and*

that $[\Sigma_2]_{z,z} > [\Sigma_1]_{z,z} > 0$ for all $z \in \{1, \dots, q\}$. Then, for all $x \in \mathbb{R}^q$,

$$\begin{aligned} \frac{\varphi(x|\boldsymbol{\mu}_1, \Sigma_1)}{\varphi(x|\boldsymbol{\mu}_2, \Sigma_2)} &\leq \prod_{z=1}^q \frac{\sqrt{[\Sigma_2]_{z,z}}}{\sqrt{[\Sigma_1]_{z,z}}} \\ &\quad \times e^{\frac{1}{2}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^\top \text{diag}\left(\frac{1}{[\Sigma_2]_{1,1} - [\Sigma_1]_{1,1}}, \dots, \frac{1}{[\Sigma_2]_{q,q} - [\Sigma_1]_{q,q}}\right)(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)}. \end{aligned}$$

Lemma B. *The Hellinger distance between two Gaussian probability density functions with diagonal covariance matrices is given by*

$$\begin{aligned} d_H^2\{\varphi(\cdot|\boldsymbol{\mu}_1, \Sigma_1), \varphi(\cdot|\boldsymbol{\mu}_2, \Sigma_2)\} &= 2 - 2 \left(\prod_{z=1}^q \frac{2\sqrt{[\Sigma_1]_{z,z}[\Sigma_2]_{z,z}}}{[\Sigma_1]_{z,z} + [\Sigma_2]_{z,z}} \right)^{1/2} \\ &\quad \times \exp \left[-\frac{1}{4}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^\top \text{diag} \left\{ \left(\frac{1}{[\Sigma_1]_{z,z} + [\Sigma_2]_{z,z}} \right)_{z \in \{1, \dots, q\}} \right\} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \right]. \end{aligned}$$

To get an ε -bracket for the probability density functions, we have to construct a δ -net for the variance and the mean for some value of δ to be specified later.

Step 1: Construction of a net for the variance. Fix $\varepsilon \in (0, 1]$ and $\delta = \varepsilon/\sqrt{2}q$. Let $b_j = (1 + \delta)^{1-j/2}A_\Sigma$. For each $j \in \{2, \dots, N\}$, we have $[a_\Sigma, A_\Sigma] = [b_N, b_{N-1}] \cup \dots \cup [b_3, b_2]$, where N is chosen to recover everything. We want

$$\begin{aligned} a_\Sigma = (1 + \delta)^{1-N/2}A_\Sigma &\Leftrightarrow \ln \frac{a_\Sigma}{A_\Sigma} = \left(1 - \frac{N}{2}\right) \ln(1 + \delta) \\ &\Leftrightarrow N = \frac{2 \ln\left(\frac{A_\Sigma}{a_\Sigma} \sqrt{1 + \delta}\right)}{\ln(1 + \delta)}. \end{aligned}$$

As we want N to be an integer, we choose

$$N = \left\lceil \frac{2 \ln\left(\frac{A_\Sigma}{a_\Sigma} \sqrt{1 + \delta}\right)}{\ln(1 + \delta)} \right\rceil,$$

where $\lceil \cdot \rceil$ denotes the ceiling function. We then get a regular net for the variance. We could let $\mathbf{B} = \text{diag}(b_{i(1)}, \dots, b_{i(q)})$, close to Σ (and deterministic, independent of the values of Σ), where i is a permutation such that $b_{i(z)+1} \leq [\Sigma]_{z,z} \leq b_{i(z)}$ for all $z \in \{1, \dots, q\}$.

Step 2: Construction of a net for the mean vectors. We use the singular decomposition of $\boldsymbol{\beta}$, $\boldsymbol{\beta} = \sum_{r=1}^R \sigma_r [\mathbf{u}]_{r,\bullet}^\top [\mathbf{v}]_{r,\bullet}$, with $(\sigma_r)_{1 \leq r \leq R}$ the singular values, and $([\mathbf{u}]_{r,\bullet})_{1 \leq r \leq R}$ and $([\mathbf{v}]_{r,\bullet})_{1 \leq r \leq R}$ unit vectors. These vectors are also bounded.

We are looking for ℓ and u such that $d_H(\ell, u) \leq \varepsilon$, and $\ell \leq f \leq u$. We use a dilatation of a Gaussian to construct such an ε -bracket of φ . We let

$$\begin{aligned} \ell(\mathbf{x}, \mathbf{y}) &= (1 + \delta)^{-(p^2qR+3q/4)} \varphi\{\mathbf{y} | \boldsymbol{\nu}_{J,R\mathbf{x}}, (1 + \delta)^{-1/4} \mathbf{B}^1\}, \\ u(\mathbf{x}, \mathbf{y}) &= (1 + \delta)^{p^2qR+3q/4} \varphi\{\mathbf{y} | \boldsymbol{\nu}_{J,R\mathbf{x}}, (1 + \delta) \mathbf{B}^2\}, \end{aligned}$$

where \mathbf{B}^1 and \mathbf{B}^2 are such that, for all $z \in \{1, \dots, q\}$, $[\mathbf{B}^1]_{z,z} \leq \Sigma_{z,z} \leq [\mathbf{B}^2]_{z,z}$ (see step 1).

The means $\boldsymbol{\nu}_{J,R} \in \mathbb{R}^{q \times p}$ will be specified later. Just remark that J is the set of the relevant columns, and R the rank of $\boldsymbol{\nu}_{J,R}$: we decompose $\boldsymbol{\nu}_{J,R} = \sum_{r=1}^R \tilde{\sigma}_r [\tilde{\mathbf{u}}]_{r,\bullet}^\top [\tilde{\mathbf{v}}]_{r,\bullet}$, with $\tilde{\mathbf{u}} \in \mathbb{R}^{R \times |J|}$ and $\tilde{\mathbf{v}} \in \mathbb{R}^{R \times q}$. We have

$$\ell(\mathbf{x}, \mathbf{y}) \leq f(\mathbf{y}|\mathbf{x}) \leq u(\mathbf{x}, \mathbf{y}) \quad \Leftrightarrow \quad \|\boldsymbol{\beta}\mathbf{x} - \boldsymbol{\nu}_{J,R}\mathbf{x}\|_2^2 \leq p^2 q R \frac{\delta^2}{2} a_\Sigma^2 (1 - 2^{-1/4}).$$

Note that $\|\boldsymbol{\beta}\mathbf{x} - \boldsymbol{\nu}_{J,R}\mathbf{x}\|_2^2 \leq p \|\boldsymbol{\beta} - \boldsymbol{\nu}_{J,R}\|_2^2 \|\mathbf{x}\|_\infty$. Thus we need

$$(19) \quad \|\boldsymbol{\beta} - \boldsymbol{\nu}_{J,R}\|_2^2 \leq pqR \frac{\delta^2}{2} a_\Sigma^2 (1 - 2^{-1/4}).$$

According to [10], $d_H(\ell, u) \leq 2(p^2 q R + 3q/4)^2 \delta^2$ and hence by choosing $\delta = \varepsilon / \{\sqrt{2}(pqR + 3/4q)\}$, we get the desired bound. Let us now see how to construct $\boldsymbol{\nu}_{J,R}$ to get (19). First write

$$\begin{aligned} \|\boldsymbol{\beta} - \boldsymbol{\nu}_{J,R}\|_2^2 &= \sum_{j=1}^p \sum_{z=1}^q \left| \sum_{r=1}^R \sigma_r [\mathbf{u}]_{r,j} [\mathbf{v}]_{r,z} - \tilde{\sigma}_r [\tilde{\mathbf{u}}]_{r,j} [\tilde{\mathbf{v}}]_{r,z} \right|^2 \\ &= \sum_{j,z} \left| \sum_{r=1}^R (\sigma_r - \tilde{\sigma}_r) [\mathbf{u}]_{r,j} [\mathbf{v}]_{r,z} \right. \\ &\quad \left. - \tilde{\sigma}_r ([\tilde{\mathbf{u}}]_{r,j} - [\mathbf{u}]_{r,j}) [\tilde{\mathbf{v}}]_{r,z} - \tilde{\sigma}_r [\mathbf{u}]_{r,j} ([\tilde{\mathbf{v}}]_{r,z} - [\mathbf{v}]_{r,z}) \right|^2. \end{aligned}$$

Then

$$\begin{aligned} \|\boldsymbol{\beta} - \boldsymbol{\nu}_{J,R}\|_2^2 &\leq \sum_{j=1}^p \sum_{z=1}^q \left| \sum_{r=1}^R (\sigma_r - \tilde{\sigma}_r) + A_\sigma ([\tilde{\mathbf{u}}]_{r,j} - [\mathbf{u}]_{r,j}) + A_\sigma ([\mathbf{v}]_{r,z} - [\tilde{\mathbf{v}}]_{r,z}) \right|^2 \\ &\leq 2pqR (\max_r |\sigma_r - \tilde{\sigma}_r|^2 + A_\sigma \max_{r,j} |[\tilde{\mathbf{u}}]_{r,j} - [\mathbf{u}]_{r,j}|^2 \\ &\quad + A_\sigma \max_{r,z} |[\tilde{\mathbf{v}}]_{r,z} - [\mathbf{v}]_{r,z}|^2). \end{aligned}$$

If we now choose $\tilde{\sigma}_r$, $[\tilde{\mathbf{u}}]_{r,j}$ and $[\tilde{\mathbf{v}}]_{r,z}$ such that

$$\begin{aligned} |\sigma_r - \tilde{\sigma}_r| &\leq \frac{\delta}{\sqrt{12}} a_\Sigma \sqrt{1 - 2^{-1/4}}, \\ |[\mathbf{u}]_{r,j} - [\tilde{\mathbf{u}}]_{r,j}| &\leq \frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}}, \\ |[\mathbf{v}]_{r,z} - [\tilde{\mathbf{v}}]_{r,z}| &\leq \frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}}, \end{aligned}$$

we can conclude, as desired, that

$$\|\boldsymbol{\beta} - \boldsymbol{\nu}_{J,R}\|_2^2 \leq pqR \frac{\delta^2}{2} a_\Sigma^2 (1 - 2^{-1/4}).$$

To accomplish this, let $\lfloor \cdot \rfloor$ the floor function and, for all $r \in \{1, \dots, R\}, j \in \{1, \dots, p\}, z \in \{1, \dots, q\}$, set

$$\begin{aligned} S &= \mathbb{Z} \cap \left[0, \left\lfloor \frac{A_\sigma}{\frac{\delta}{\sqrt{12}} a_\Sigma \sqrt{1 - 2^{-1/4}}} \right\rfloor \right], \\ \tilde{\sigma}_r &= \operatorname{argmin}_{\varsigma \in S} \left| \sigma_r - \frac{\delta}{\sqrt{12}} a_\Sigma \sqrt{1 - 2^{-1/4}} \varsigma \right|, \\ U &= \mathbb{Z} \cap \left[0, \left\lfloor \frac{A_\sigma}{\frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}}} \right\rfloor \right], \\ [\tilde{\mathbf{u}}]_{r,j} &= \operatorname{argmin}_{\mu \in U} \left| [\mathbf{u}]_{r,j} - \frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}} \mu \right|, \\ V &= \mathbb{Z} \cap \left[0, \left\lfloor \frac{A_\sigma}{\frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}}} \right\rfloor \right], \\ [\tilde{\mathbf{v}}]_{r,z} &= \operatorname{argmin}_{\nu \in V} \left| [\mathbf{v}]_{r,z} - \frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}} \nu \right|. \end{aligned}$$

Note that we just need to determine vectors $(([\tilde{\mathbf{u}}]_{r,j})_{1 \leq j \leq J-r})_{1 \leq r \leq R}$ and $(([\tilde{\mathbf{v}}]_{r,z})_{1 \leq z \leq q-r})_{1 \leq r \leq R}$ because these are unit vectors. Then we let, for all $z \in \{1, \dots, p\}$,

$$[\boldsymbol{\nu}_{J,R}]_{z,j} = \begin{cases} \sum_{r=1}^R \tilde{\sigma}_r [\tilde{\mathbf{u}}]_{r,j} [\tilde{\mathbf{v}}]_{r,z} & \text{if } j \in J, \\ 0 & \text{if } j \in J^c. \end{cases}$$

Step 3: Upper bound of the number of ε -brackets for $\mathcal{F}_{(J,R)}$. We have defined the brackets. Let $c = (1 - 2^{-1/4})/12$. We want to control the entropy. We denote by $\mathcal{B}_\varepsilon(\mathcal{F}_{(J,R)})$ the minimum number of brackets of width $d_H^{\otimes n}$ smaller than ε such that every density of $\mathcal{F}_{(J,R)}$ belongs to one of those brackets. We have

$$\begin{aligned} |\mathcal{B}_\varepsilon(\mathcal{F}_{(J,R)})| &\leq \sum_{l=2}^N \left(\frac{A_\sigma}{\delta a_\Sigma \sqrt{c}} \right)^R \left(\frac{A_\sigma^2}{\delta a_\Sigma \sqrt{c}} \right)^{R \left(\frac{2J-R-1}{2} + \frac{2q-R-1}{2} \right)} \\ &\leq (N-1) \left(\frac{A_\sigma^2}{\delta a_\Sigma \sqrt{c}} \right)^{R(J+q-R)} A_\sigma^{-R} \\ &\leq C(a_\Sigma, A_\Sigma, A_\sigma, J, R) \delta^{-D_{(J,R)}-1} \end{aligned}$$

with

$$C(a_\Sigma, A_\Sigma, A_\sigma, J, R) = 2 \left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2} \right) \left(\frac{A_\sigma^2}{a_\Sigma \sqrt{c}} \right)^{R(J+q-R)} A_\sigma^{-R}.$$

B.3 For the mixture. We begin by computing the bracketing entropy: according to (18),

$$\mathcal{H}_{[\cdot]}(\varepsilon, \mathcal{S}_{(K,J,R)}^B, d_H) \leq \ln \left\{ C \left(\frac{1}{\varepsilon} \right)^{D_{(K,J,R)}} \right\},$$

where

$$(20) \quad C = 2^K K (2\pi e)^{K/2} 3^{K-1} \left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2} \right)^K \times \left(\frac{A_\sigma \sqrt{12}}{a_\Sigma^2 \sqrt{1 - 2^{-1/4}}} \right)^{D(K, J, \mathbf{R})} A_\sigma^{-\sum_{k=1}^K \mathbf{R}_k},$$

and $D_{(K, J, \mathbf{R})} = \sum_{k=1}^K \mathbf{R}_k (|J| + q - \mathbf{R}_k)$. We must determine $\psi_{(K, J, \mathbf{R})}$ such that

$$(21) \quad \int_0^\varpi \sqrt{\mathcal{H}_{[\cdot]}(\varepsilon, \mathcal{S}_{(K, J, \mathbf{R})}^{\mathcal{B}}(s^{(K, J, \mathbf{R})}, \varpi), d_H^{\otimes n})} d\varepsilon \leq \psi_{(K, J, \mathbf{R})}(\varpi).$$

To this end, we compute the integral, viz.

$$\begin{aligned} & \int_0^\varpi \sqrt{\mathcal{H}_{[\cdot]}(\varepsilon, \mathcal{S}_{(K, J, \mathbf{R})}^{\mathcal{B}}(s^{(K, J, \mathbf{R})}, \varpi), d_H^{\otimes n})} d\varepsilon \\ & \leq \varpi \sqrt{\ln(C)} + \sqrt{D_{(K, J, \mathbf{R})}} \int_0^\varpi \sqrt{\ln\left(\frac{1}{\varepsilon}\right)} d\varepsilon \\ & \leq \sqrt{D_{(K, J, \mathbf{R})}} \varpi \left\{ \sqrt{\pi} + \sqrt{\frac{\ln(C)}{D_{(K, J, \mathbf{R})}}} + \sqrt{\ln\left(\frac{1}{\varpi \wedge 1}\right)} \right\} \end{aligned}$$

with, according to (20),

$$\begin{aligned} \ln(C) &= K \ln(2) + \frac{K}{2} \ln(2\pi e) + (K-1) \ln(3) + K \ln\left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2}\right) \\ &+ D_{(K, J, \mathbf{R})} \ln\left(\frac{A_\sigma^2 \sqrt{12}}{a_\Sigma^2 \sqrt{1 - 2^{-1/4}}}\right) + \ln(K) + \sum_{k=1}^K \mathbf{R}_k \ln\left(\frac{1}{A_\sigma}\right) \\ &\leq D_{(K, J, \mathbf{R})} \left\{ \ln(2) + \ln(2\pi e) + \ln 3 + 1 + \ln\left(\frac{A_\sigma^2 \sqrt{12}}{a_\Sigma^2 \sqrt{1 - 2^{-1/4}}}\right) \right\} \\ &+ D_{(K, J, \mathbf{R})} \left\{ \ln\left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2}\right) \right\} \\ &\leq D_{(K, J, \mathbf{R})} \left[\ln\left(\frac{12\sqrt{12}\pi e}{\sqrt{1 - 2^{-1/4}}}\right) + \ln\left\{\frac{A_\sigma^2}{a_\Sigma} \left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2}\right)\right\} \right]. \end{aligned}$$

Then

$$\begin{aligned} & \int_0^\varpi \sqrt{\mathcal{H}_{[\cdot]}(\varepsilon, \mathcal{S}_{(K, J, \mathbf{R})}^{\mathcal{B}}(s^{(K, J, \mathbf{R})}, \varpi), d_H^{\otimes n})} d\varepsilon \\ & \leq \sqrt{D_{(K, J, \mathbf{R})}} \left\{ 2 + \sqrt{\ln\left\{\frac{A_\sigma^2}{a_\Sigma} \left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2}\right)\right\}} + \sqrt{\ln\left(\frac{1}{\varpi \wedge 1}\right)} \right\}. \end{aligned}$$

Consequently, by putting

$$B = 2 + \sqrt{\ln\left\{\frac{A_\sigma^2}{a_\Sigma} \left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2}\right)\right\}},$$

we find that the function $\psi_{(K,J,\mathbf{R})}$ defined on $(0, \infty)$ by

$$\psi_{(K,J,\mathbf{R})}(\varpi) = \sqrt{D_{(K,J,\mathbf{R})}} \varpi \left\{ B + \sqrt{\ln \left(\frac{1}{\varpi \wedge 1} \right)} \right\}$$

satisfies (21). Besides, $\psi_{(K,J,\mathbf{R})}$ is non-decreasing and $\varpi \mapsto \psi_{(K,J,\mathbf{R})}(\varpi)/\varpi$ is non-increasing, so $\psi_{(K,J,\mathbf{R})}$ is convenient.

Finally, we need to find an upper bound on ϖ_* satisfying $\psi_{(K,J,\mathbf{R})}(\varpi_*) = \sqrt{n\varpi_*^2}$. Consider ϖ_* such that $\psi_{(K,J,\mathbf{R})}(\varpi_*) = \sqrt{\varpi_*^2}$. This is equivalent to solving

$$\varpi_* = \sqrt{\frac{D_{(K,J,\mathbf{R})}}{n}} \left\{ B + \sqrt{\ln \left(\frac{1}{\varpi_* \wedge 1} \right)} \right\}.$$

We then choose

$$\varpi_*^2 \leq \frac{D_{(K,J,\mathbf{R})}}{n} \left\{ 2B^2 + \ln \left(\frac{1}{1 \wedge \frac{D_{(K,J,\mathbf{R})}}{n} B^2} \right) \right\}.$$

C Assumption K. Let

$$w_{(K,J,\mathbf{R})} = \tilde{D}_{(K,J)} \ln \left\{ \frac{4epq}{(\tilde{D}_{(K,J)} - q^2) \wedge pq} \right\} + \sum_{k \in \{1, \dots, K\}} \mathbf{R}_k,$$

where $\tilde{D}_{(K,J)} = K(1 + |J|)$. Then

$$\begin{aligned} \sum_{(K,J,\mathbf{R})} e^{-w_{(K,J,\mathbf{R})}} &\leq \sum_{K \geq 1} \left(\sum_{R \geq 1} e^{-R} \right)^K \left\{ \sum_{1 \leq |J| \leq p} e^{-\tilde{D}_{(K,J)} \ln \left(\frac{4epq}{(\tilde{D}_{(K,J)} - q^2) \wedge pq} \right)} \right\} \\ &\leq \sum_{K \geq 1} 1^K \left\{ \sum_{1 \leq |J| \leq p} e^{-\tilde{D}_{(K,J)} \ln \left(\frac{4epq}{(\tilde{D}_{(K,J)} - q^2) \wedge pq} \right)} \right\} \leq 2, \end{aligned}$$

where the last inequality stems from Proposition 4.5 in [10].

We observe that the weights depend on (K, J, \mathbf{R}) , which leads to a sharper control but forces the penalty to depend on τ , the parameter related to Bernstein's inequality, according to (16). Larger weights can be used to avoid this dependence, however. We prefer to use the present weights to get the thinnest control possible.

DEPARTMENT OF MATHEMATICS AND LEUVEN STATISTICS RESEARCH CENTER (LSTAT), KU LEUVEN, LEUVEN, BELGIUM