

PAC-BAYESIAN INDUCTIVE AND TRANSDUCTIVE LEARNING

OLIVIER CATONI

May 28, 2006

ABSTRACT: We present here a PAC-Bayesian point of view on adaptive supervised classification. Using convex analysis on the set of posterior probability measures on the parameter space, we show how to get local measures of the complexity of the classification model involving the relative entropy of posterior distributions with respect to Gibbs posterior measures. We then discuss relative bounds, comparing the generalization error of two classification rules, showing how the margin assumption of Mammen and Tsybakov can be replaced with some empirical measure of the covariance structure of the classification model. We also show how to associate to any posterior distribution an *effective temperature* relating it to the Gibbs prior distribution with the same level of expected error rate, and how to estimate this effective temperature from data, resulting in an estimator whose expected error rate converges according to the best possible power of the sample size adaptively under any margin and parametric complexity assumptions. Then we introduce a PAC-Bayesian point of view on transductive learning and use it to improve on known Vapnik's generalization bounds, extending them to the case when the sample is made of independent non identically distributed pairs of patterns and labels. Eventually we review briefly the construction of Support Vector Machines and show how to derive generalization bounds for them, measuring the complexity either through the number of support vectors or through transductive or inductive margin estimates.

2000 MATHEMATICS SUBJECT CLASSIFICATION: 62H30, 68T05, 62B10.

KEYWORDS: Statistical learning theory, adaptive statistics, pattern recognition, PAC-Bayesian theorems, VC dimension, local complexity bounds, randomized estimators, Gibbs posterior distributions, effective temperature, Mammen and Tsybakov margin assumption, transductive inference, compression schemes, Support Vector Machines, margin bounds.

CNRS – Laboratoire de Probabilités et Modèles Aléatoires, Université Paris 6 (site Chevaleret), 4 place Jussieu – Case 188, 75 252 Paris Cedex 05.

CONTENTS

INTRODUCTION	4
1. INDUCTIVE PAC-BAYESIAN LEARNING	8
1.1. BASIC INEQUALITY	9
1.2. NON LOCAL BOUNDS	11
1.2.1. <i>Unbiased empirical bounds</i>	12
1.2.2. <i>Optimizing explicitly the exponential parameter λ</i>	16
1.2.3. <i>Non random bounds</i>	18
1.2.4. <i>Deviation bounds</i>	20
1.3. LOCAL BOUNDS	22
1.3.1. <i>Choice of the prior</i>	22
1.3.2. <i>Unbiased local empirical bounds</i>	24
1.3.3. <i>Non random local bounds</i>	27
1.3.4. <i>Local deviation bounds</i>	28
1.3.5. <i>Partially local bounds</i>	31
1.3.6. <i>Two step localization</i>	35
1.4. RELATIVE BOUNDS	40
1.4.1. <i>Basic inequalities</i>	41
1.4.2. <i>Non random bounds</i>	45
1.4.3. <i>Unbiased empirical bounds</i>	49
1.4.4. <i>Relative empirical deviation bounds</i>	53
1.5. BOUNDS RELATIVE TO A GIBBS DISTRIBUTION	60
1.5.1. <i>Comparing a posterior distribution with a Gibbs prior</i>	61
1.5.2. <i>The effective temperature of a posterior distribution</i>	65
1.5.3. <i>Analysis of an empirical bound for the effective temperature</i>	66
1.5.4. <i>Adaptation to parametric and margin assumptions</i>	72
1.5.5. <i>Estimating the divergence of a posterior with respect to a Gibbs prior</i>	79
1.5.6. <i>Comparing two posterior distributions</i>	80
1.5.7. <i>Two step localization of relative bounds</i>	83
1.5.8. <i>Analysis of the two step relative bound</i>	91
2. TRANSDUCTIVE PAC-BAYESIAN LEARNING	98
2.1. BASIC INEQUALITIES.	98
2.2. VAPNIK'S BOUNDS FOR TRANSDUCTIVE CLASSIFICATION	103
2.2.1. <i>With a shadow sample of arbitrary size</i>	103

2.2.2. When the shadow sample has the same size as the training sample	106
2.2.3. When moreover the distribution of the augmented sample is exchangeable	107
2.3. VAPNIK'S BOUNDS FOR INDUCTIVE CLASSIFICATION	110
2.3.1. Arbitrary shadow sample size	110
2.3.2. A better minimization with respect to the exponential parameter	112
2.3.3. Equal shadow and training sample sizes	115
2.3.4. Improvement on the equal sample size bound in the i.i.d. case	115
2.4. GAUSSIAN APPROXIMATION IN VAPNIK'S BOUNDS	117
2.4.1. Arbitrary shadow sample size	118
2.4.2. Equal sample sizes in the i.i.d. case	119
3. SUPPORT VECTOR MACHINES	121
3.1. HOW TO BUILD THEM	121
3.1.1. The canonical hyperplane	121
3.1.2. Computation of the canonical hyperplane	122
3.1.3. Support vectors	124
3.1.4. The non separable case	125
3.1.5. Support Vector Machines	130
3.1.6. Building kernels	132
3.2. BOUNDS FOR SUPPORT VECTOR MACHINES	134
3.2.1. Compression scheme bounds	134
3.2.2. The Vapnik Cervonenkis dimension of a family of subsets	136
3.2.3. VC dimension of linear rules with margin	139
3.2.4. Application to Support Vector Machines	142
3.2.5. Inductive margin bounds	143
4. APPENDIX: CLASSIFICATION BY THRESHOLDING	148
4.1. DESCRIPTION OF THE MODEL	148
4.2. COMPUTATION OF INDUCTIVE BOUNDS	149
4.3. TRANSDUCTIVE BOUNDS	152
REFERENCES	154

INTRODUCTION

Among the possible approaches to pattern recognition, statistical learning theory has received a lot of attention in the last few years. Although a realistic pattern recognition scheme involves data pre-processing and post-processing that need a theory of their own, a central role is often played by some kind of supervised learning algorithm. This central piece of work is the subject we are going to analyse in these notes.

Accordingly, we assume that we have prepared in some way or another a *sample* of N labelled patterns $(X_i, Y_i)_{i=1}^N$, where X_i ranges in some pattern space \mathcal{X} and Y_i ranges in some finite label set \mathcal{Y} . We also assume that we have devised our experiment in such a way that the couples of random variables (X_i, Y_i) are independent (but not necessarily equidistributed). Here, randomness should be understood to come from the way the statistician has planned his experiment. He may for instance have drawn the X_i s at random from some larger population of patterns the algorithm is meant to be applied to in a second stage. The labels Y_i may have been set with the help of some external expertise (which may itself be faulty or contain some amount of randomness, therefore we do not assume that Y_i is a function of X_i , and allow the couple of random variables (X_i, Y_i) to follow any kind of joint distribution). In practice, patterns will be extracted from some high dimensional and highly structured data, like digital images, speech signals, DNA sequences, etc. We will not discuss here this pre-processing stage (although it poses crucial problems dealing with segmentation and the choice of a representation).

To fix notations, let $(X_i, Y_i)_{i=1}^N$ be the canonical process on $\Omega = (\mathcal{X} \times \mathcal{Y})^N$ (which means the coordinate process). Let the pattern space be provided with a sigma-algebra \mathcal{B} turning it into a measurable space $(\mathcal{X}, \mathcal{B})$. On the finite label space \mathcal{Y} , we will consider the trivial algebra \mathcal{B}' made of all its subsets. Let $\mathcal{M}_+^1[(\mathcal{X} \times \mathcal{Y})^N, (\mathcal{B} \otimes \mathcal{B}')^{\otimes N}]$ be our notation for the set of probability measures (i.e. of positive measures of total mass equal to 1) on the measurable space $[(\mathcal{X} \times \mathcal{Y})^N, (\mathcal{B} \otimes \mathcal{B}')^{\otimes N}]$. Once some probability distribution $\mathbb{P} \in \mathcal{M}_+^1[(\mathcal{X} \times \mathcal{Y})^N, (\mathcal{B} \otimes \mathcal{B}')^{\otimes N}]$ is chosen, it turns $(X_i, Y_i)_{i=1}^N$ into the canonical realization of a stochastic process modeling the observed sample (also called the training set). We will assume that $\mathbb{P} = \bigotimes_{i=1}^N P_i$, where for each $i = 1, \dots, N$, $P_i \in \mathcal{M}_+^1(\mathcal{X} \times \mathcal{Y}, \mathcal{B} \otimes \mathcal{B}')$, to reflect the assumption that we observe independent pairs of patterns and labels. We will also assume that we are provided with some indexed set of possible classification rules

$$\mathcal{R}_\Theta = \{f_\theta : \mathcal{X} \rightarrow \mathcal{Y}; \theta \in \Theta\},$$

where (Θ, \mathcal{T}) is some measurable index set. Assuming some indexation of the classification rules is just a matter of presentation. Although it leads to longer notations, it allows to integrate over the space of classification rules as well as over Ω using the usual formalism of multiple integrals. For this matter, we will assume that $(\theta, x) \mapsto f_\theta(x) : (\Theta \times \mathcal{X}, \mathcal{B} \otimes \mathcal{T}) \rightarrow (\mathcal{Y}, \mathcal{B}')$ is a measurable function.

In many cases $\Theta = \bigcup_{i \in I} \Theta_i$ will be a finite (or more generally countable) union of subspaces, dividing the classification model $\mathcal{R}_\Theta = \bigcup_{i \in I} \mathcal{R}_{\Theta_i}$ into a union of submodels. The importance of introducing such a structure has been put forward by V. Vapnik, as a way to avoid making strong hypotheses on the distribution \mathbb{P} of the sample. If neither the distribution of the sample nor the set of classification rules were constrained, it is well known indeed that no kind of statistical inference would be possible. Considering a family of submodels is a way to provide for adaptive classification where the choice of the model depends on the observed sample. Restricting the set of classification rules is more realistic than restricting the distribution of patterns, since the classification rules are a processing tool left to the choice of the statistician, whereas the distribution of the patterns is not fully under his control, except for some planning of the learning experiment which may enforce some weak properties like independence, but not the precise shapes of the marginal distributions P_i which are as a rule unknown distributions on some high dimensional space.

In these notes, we will concentrate on general issues concerned with a natural measure of risk, namely the *expected error rate* of each classification rule f_θ , expressed as

$$R(\theta) = \frac{1}{N} \sum_{i=1}^N \mathbb{P}[f_\theta(X_i) \neq Y_i].$$

As this quantity is unobserved, we will be led to work with the corresponding *empirical error rate*

$$r(\theta, \omega) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[f_\theta(X_i) \neq Y_i].$$

This does not mean that practical learning algorithms will always try to minimize this criterion. They often on the contrary try to minimize some other criterion which is linked with the structure of the problem and has some nice additional properties (like smoothness and convexity, for example). Nevertheless, and independently from the precise form of the estimator

$\hat{\theta} : \Omega \rightarrow \Theta$ under study, the analysis of $R(\hat{\theta})$ is a natural question, and often corresponds to what is required in practice.

Answering this question is not straightforward because, although $R(\theta)$ is the expectation of $r(\theta)$, a sum of independent Bernoulli random variables, $R(\hat{\theta})$ is not the expectation of $r(\hat{\theta})$, because of the dependence of $\hat{\theta}$ on the sample, and neither is $r(\hat{\theta})$ a sum of independent random variables. To circumvent this unfortunate situation, some uniform control over the deviations of r with respect to R is needed.

The PAC-Bayesian approach to this problem, originated in the machine learning community and pioneered by D. McAllester [25, 26], can be seen as some variant of the more classical approach of M -estimators relying on empirical process theory (as exposed for instance in [36]).

It is built on three corner stones:

- One idea is to embed the set of estimators of the type $\hat{\theta} : \Omega \rightarrow \Theta$ into the larger set of regular conditional probability measures $\rho : (\Omega, (\mathcal{B} \otimes \mathcal{B}')^{\otimes N}) \rightarrow \mathcal{M}_+^1(\Theta, \mathcal{T})$. We will call these conditional probability measures *posterior distributions*, to follow a usual terminology.
- A second idea is to measure the fluctuations of ρ with respect to the sample, using some prior distribution $\pi \in \mathcal{M}_+^1(\Theta, \mathcal{T})$, and the Kullback divergence function $\mathcal{K}(\rho, \pi)$. The expectation $\mathbb{P}\{\mathcal{K}(\rho, \pi)\}$ measures the randomness of ρ . The optimal choice of π would be $\mathbb{P}(\rho)$, resulting in a measure of the randomness of ρ equal to the mutual information between the sample and the estimated parameter drawn from ρ . Anyhow, since $\mathbb{P}(\rho)$ is as a rule no more observed than \mathbb{P} , we will have to be content with some less concentrated prior distribution π , resulting in some looser measure of randomness, as shown by the identity $\mathbb{P}[\mathcal{K}(\rho, \pi)] = \mathbb{P}\{\mathcal{K}[\rho, \mathbb{P}(\rho)]\} + \mathcal{K}[\mathbb{P}(\rho), \pi]$.
- A third idea is to analyze the fluctuations of the random process $\theta \mapsto r(\theta)$ with respect to its mean process $\theta \mapsto R(\theta)$ through the log-Laplace transform

$$-\frac{1}{\lambda} \log \left\{ \iint \exp[-\lambda r(\theta, \omega)] \pi(d\theta) \mathbb{P}(d\omega) \right\},$$

as a physicist prone to statistical mechanics (where this is called the free energy) would do. This transform is well suited to relate $\min_{\theta \in \Theta} r(\theta)$ to $\inf_{\theta \in \Theta} R(\theta)$.

This monograph is divided into two sections. The first one deals with the inductive setting presented in these lines, the second one with the *transductive* setting, where, following Vapnik's seminal approach [37], a shadow

sample is considered.

In the first section, two types of bounds are shown. *Empirical bounds* can be used to choose between estimators or to build estimators. *Non random bounds* can be used to assess the speed of convergence of estimators, relating this speed to the speed of convergence of the Gibbs prior expected error rate $\beta \mapsto \pi_{\exp(-\beta R)}(R)$ towards $\text{ess inf}_{\pi} R$ as β goes to infinity, and to other quantities akin to the margin assumption of Mammen and Tsybakov in more sophisticated cases. We will progress from the most straightforward bounds to more elaborate ones, built to achieve a better asymptotic behaviour. We will thus introduce *local bounds* and *relative bounds*. From an asymptotic point of view, the culminating result of these notes is Theorem 1.39 (page 63). It is used in Proposition 1.40 to build a classification rule which is proved to be adaptive in all the parameters of the Mammen and Tsybakov margin assumption and of a parametric complexity assumption in Corollary 1.52 (page 78) of Theorem 1.50 (page 77). This opens the road to Theorem 1.59 (page 88) which performs two step localization on top of Theorem 1.39 in order to be able to achieve adaptive model selection with a decreased influence of the number of empirically inefficient models included in the comparison. The analysis of this bound is hinted at in subsequent pages, but not fully developed, since we are not sure the amount of technicalities it requires is worth it. Anyhow we would not like to induce the reader into thinking that each result in the first section is actually an *improvement* on the previous one, it is as a rule only an *asymptotic improvement*, and the price to pay for being asymptotically tighter is to get looser bounds for small sample sizes. What is a small sample size in practice is a question of ratio between the number of examples and the complexity (roughly speaking the number of parameters) of the model used to classify. Since our aim here is to describe classification methods suitable for complex data (images, speech, DNA, ...), we suspect that practitioners wanting to make use of our proposals will be confronted with small sample sizes more often than with large ones, and should try to make use of the simplest bounds first and see only afterwards whether the asymptotically better ones can bring them more for the size of samples their computers can handle and their data bases can provide. Let us advocate also that the results of this first section are not only of a theoretical nature for two reasons : the first one is that posterior parameter distributions can be computed effectively, using Monte Carlo techniques, there is a whole tradition about these computations in Bayesian statistics, proving that what we call here Gibbs estimators are not only a way to show that some optimal speeds of convergence can be reached in some theoretically well understood situations, but that they can

also be computed in practice. The second reason is that a traditional non randomized estimator $\hat{\theta} \in \Theta$ of the parameter can be approximated by a posterior distribution ρ which is supported by a fairly narrow neighborhood of $\hat{\theta} \in \Theta$, without spoiling excessively our bounds, resulting in a classification rule which is to provide a randomized answer only for a small amount of dubious examples and will most of the time issue the same deterministic answer as the classification rule indexed by $\hat{\theta}$ it is derived from. This is explained on page 14.

In the second section, we show first how we can transport all the results obtained in the inductive case to the transductive case, allowing to replace prior distributions by *partially exchangeable posterior distributions* depending on an extended sample where unlabelled shadow examples are added, with increased possibilities of adaptation to the data. We then focus on the small sample case, where local and relative bounds are not expected to be of great help. Using a fictitious (that is unobserved) shadow sample, we study Vapnik type generalization bounds, showing how to tighten and extend them using some original ideas, like making no Gaussian approximation to the log-Laplace of Bernoulli random variables, — using a shadow sample of arbitrary size, — shrinking from the use of any symmetrization trick — and using a subset of the group of permutations suitable to cover the case of independent non identically distributed data. The culminating result of the second section is Theorem 2.17 on page 114, subsequent bounds showing the separate influence of the above ideas and providing an easier comparison with Vapnik’s original results. Vapnik type generalization bounds have a broad applicability, not only through the concept of VC dimension, but also through the use of compression schemes [24], which are briefly described on page 105.

1. INDUCTIVE PAC-BAYESIAN LEARNING

The setting of inductive inference (as opposed to transductive inference to be discussed later) is the one described in the introduction.

When we will have to take the expectation of a random variable $Z : \Omega \rightarrow \mathbb{R}$ as well as of a function of the parameter $h : \Theta \rightarrow \mathbb{R}$ with respect to some probability measure, we will as a rule use functional short notations instead of resorting to the integral sign: thus we will write $\mathbb{P}(Z)$ for $\int_{\Omega} Z(\omega)\mathbb{P}(d\omega)$ and $\pi(h)$ for $\int_{\Theta} h(\theta)\pi(d\theta)$.

The PAC-Bayesian approach, in its simplest form, relies on some basic upper bound for the Laplace transform of $\sup_{\rho \in \mathcal{M}_+^1(\Theta)} [\rho(R) - \rho(r)]$, or more

technically on some penalized variant of it, as will be seen. This will be the subject of the next subsection, where we will start with the Laplace transform of $R(\theta) - r(\theta)$, for any $\theta \in \Theta$, before encompassing posterior distributions. As it is already easy to guess, the purpose of these preliminaries is to gain some uniform control on the lower deviations of the empirical error rate from the expected error rate under any posterior distribution.

1.1. BASIC INEQUALITY. In the setting described in the introduction, let us consider the Bernoulli random variables $\sigma_i(\theta) = \mathbb{1}[Y_i \neq f_\theta(X_i)]$. Using independence and the concavity of the logarithm function, it is readily seen that for any real constant λ

$$\begin{aligned} \log\left\{\mathbb{P}\left\{\exp[-\lambda r(\theta)]\right\}\right\} &= \sum_{i=1}^N \log\left\{\mathbb{P}\left[\exp\left(-\frac{\lambda}{N}\sigma_i\right)\right]\right\} \\ &\leq N \log\left\{\frac{1}{N} \sum_{i=1}^N \mathbb{P}\left[\exp\left(-\frac{\lambda}{N}\sigma_i\right)\right]\right\}. \end{aligned}$$

The right-hand side of this inequality is the log Laplace transform of a Bernoulli distribution with parameter $\frac{1}{N} \sum_{i=1}^N \mathbb{P}(\sigma_i) = R(\theta)$. As any Bernoulli distribution is fully defined by its parameter, this log Laplace transform is necessarily a function of $R(\theta)$. It can be expressed with the help of the family of functions

$$\Phi_a(p) = -a^{-1} \log\{1 - [1 - \exp(-a)]p\}, \quad a \in \mathbb{R}, p \in (0, 1).$$

It is immediately seen that Φ_a is an increasing one to one mapping of the unit interval unto itself, and that it is convex when $a > 0$, concave when $a < 0$ and can be defined by continuity to be the identity when $a = 0$. Moreover the inverse of Φ_a is given by the formula

$$\Phi_a^{-1}(q) = \frac{1 - \exp(-aq)}{1 - \exp(-a)}, \quad a \in \mathbb{R}, q \in (0, 1).$$

This formula may be used to extend Φ_a^{-1} to $q \in \mathbb{R}$, and we will use this extension without further notice when required.

Using these notations, the previous inequality becomes

$$\log\left\{\mathbb{P}\left\{\exp[-\lambda r(\theta)]\right\}\right\} \leq -\lambda \Phi_{\frac{\lambda}{N}}[R(\theta)], \quad \text{proving}$$

LEMMA 1.1. *For any real constant λ and any parameter $\theta \in \Theta$,*

$$\mathbb{P} \left\{ \exp \left\{ \lambda \left[\Phi_{\frac{\lambda}{N}} [R(\theta)] - r(\theta) \right] \right\} \right\} \leq 1.$$

In previous versions of this study, we had used some Bernstein bound, instead of this lemma. Anyhow, as it will turn out, keeping the log Laplace of a Bernoulli instead of approximating it provides simpler and tighter results.

Lemma 1.1 implies that for any constants $\lambda \in \mathbb{R}_+$ and $\epsilon \in]0, 1)$,

$$\mathbb{P} \left[\Phi_{\frac{\lambda}{N}} [R(\theta)] + \frac{\log(\epsilon)}{\lambda} \leq r(\theta) \right] \geq 1 - \epsilon.$$

Choosing $\bar{\lambda} \in \arg \max_{\mathbb{R}_+} \Phi_{\frac{\lambda}{N}} [R(\theta)] + \frac{\log(\epsilon)}{\lambda}$, we deduce

LEMMA 1.2. *For any $\epsilon \in]0, 1)$, any $\theta \in \Theta$,*

$$\mathbb{P} \left\{ R(\theta) \leq \inf_{\lambda \in \mathbb{R}_+} \Phi_{\frac{\lambda}{N}}^{-1} \left[r(\theta) - \frac{\log(\epsilon)}{\lambda} \right] \right\} \geq 1 - \epsilon.$$

We will illustrate throughout these notes the bounds we prove with a small numerical example: in the case where $N = 1000$, $\epsilon = 0.01$ and $r(\theta) = 0.2$, we get with a confidence level of 0.99 that $R(\theta) \leq .2402$, this being obtained for $\lambda = 234$.

Now, to proceed towards the analysis of posterior distributions, let us put for short $U_\lambda(\theta, \omega) = \lambda \left[\Phi_{\frac{\lambda}{N}} [R(\theta)] - r(\theta, \omega) \right]$, and let us consider $\log \left\{ \mathbb{P} \left[\pi \left[\exp(U_\lambda) \right] \right] \right\}$, where $\pi \in \mathcal{M}_+^1(\Theta, \mathcal{J})$ is some prior probability measure on the parameter space. Using Fubini's theorem for non negative functions, we see that

$$\log \left\{ \mathbb{P} \left[\pi \left[\exp(U_\lambda) \right] \right] \right\} = \log \left\{ \pi \left[\mathbb{P} \left[\exp(U_\lambda) \right] \right] \right\} \leq 0.$$

To relate this quantity to the expectation $\rho(U_\lambda)$ with respect to any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, we will use the properties of the Kullback divergence $\mathcal{K}(\rho, \pi)$ of ρ with respect to π , which is defined as

$$\mathcal{K}(\rho, \pi) = \begin{cases} \int \log \left(\frac{d\rho}{d\pi} \right) d\rho, & \text{when } \rho \ll \pi, \\ +\infty, & \text{otherwise.} \end{cases}$$

The following lemma shows in which sense the Kullback divergence function can be thought of as the dual of the log Laplace transform.

LEMMA 1.3. *For any bounded measurable function $h : \Theta \rightarrow \mathbb{R}$, and any probability distribution $\rho \in \mathcal{M}_+^1(\Theta)$ such that $\mathcal{K}(\rho, \pi) < \infty$,*

$$\log\{\pi[\exp(h)]\} = \rho(h) - \mathcal{K}(\rho, \pi) + \mathcal{K}(\rho, \pi_{\exp(h)}),$$

where by definition $\frac{d\pi_{\exp(h)}}{d\pi} = \frac{\exp[h(\theta)]}{\pi[\exp(h)]}$. Consequently

$$\log\{\pi[\exp(h)]\} = \sup_{\rho \in \mathcal{M}_+^1(\Theta)} \rho(h) - \mathcal{K}(\rho, \pi).$$

The proof is just a matter of writing down the definition of the quantities involved and using the fact that the Kullback divergence function is non negative. It can be found in [17, page 160]. In the duality between measurable functions and probability measures, we thus see that the log Laplace transform with respect to π is the Legendre transform of the Kullback divergence function with respect to π . Using this, we get

$$\mathbb{P}\left\{\exp\left\{\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \rho[U_\lambda(\theta)] - \mathcal{K}(\rho, \pi)\right\}\right\} \leq 1,$$

which, combined with the convexity of $\lambda\Phi_{\frac{\lambda}{N}}$, proves the basic inequality we were looking for.

THEOREM 1.4. *For any real constant λ ,*

$$\begin{aligned} & \mathbb{P}\left\{\exp\left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \lambda\left[\rho\left(\Phi_{\frac{\lambda}{N}} \circ R\right) - \rho(r)\right] - \mathcal{K}(\rho, \pi)\right]\right\} \\ & \leq \mathbb{P}\left\{\exp\left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \lambda\left[\Phi_{\frac{\lambda}{N}}[\rho(R)] - \rho(r)\right] - \mathcal{K}(\rho, \pi)\right]\right\} \leq 1. \end{aligned}$$

The following sections will show how to use this theorem.

1.2. NON LOCAL BOUNDS. At least three sorts of bounds can be deduced from Theorem 1.4.

The most interesting ones to build estimators and tune parameters, as well as the first that have been considered in the development of the PAC-Bayesian approach, are deviation bounds. They provide an empirical upper bound for $\rho(R)$ — that is a bound which can be computed from observed data — with some probability $1 - \epsilon$, where ϵ is a presumably small and tunable confidence level.

Anyhow, since most of the results about the convergence speed of estimators to be found in the statistical literature are concerned with the expectation $\mathbb{P}[\rho(R)]$, it is also enlightening to bound this quantity. In order to know at which rate it may be approaching $\inf_{\Theta} R$, a non random upper bound is required, which will relate the average of the expected risk $\mathbb{P}[\rho(R)]$ with the properties of the contrast function $\theta \mapsto R(\theta)$.

Since the values of constants do matter a lot when a bound is to be used to select between various estimators using classification models of various complexities, a third kind of bound, related to the first, may be considered for the sake of its hopefully better constants: we will call them *unbiased empirical bounds*, to stress the fact that they provide some empirical quantity whose expectation under \mathbb{P} can be proved to be an upper bound for $\mathbb{P}[\rho(R)]$, the average expected risk. The price to pay for these better constants is of course the lack of formal guarantee given by the bound : two random variables whose expectations are ordered in a certain way may very well be ordered in the reverse way with a large probability, so that basing the estimation of parameters or the selection of an estimator on some unbiased empirical bound is a hazardous business. Anyhow, since it is common practice to use the inequalities provided by mathematical statistical theory while replacing the proven constants with smaller values showing a better practical efficiency, considering unbiased empirical bounds akin to deviation bounds provides an indication about how much the constants may be decreased while not violating the theory too outrageously.

1.2.1. Unbiased empirical bounds. Let $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$ be some fixed (and arbitrary) posterior distribution, describing some randomized estimator of θ . As we already mentioned, in these notes a posterior distribution will always be a regular conditional probability measure. By this we mean that

- for any $A \in \mathcal{T}$, the map $\omega \mapsto \rho(\omega, A) : (\Omega, (\mathcal{B} \otimes \mathcal{B}')^{\otimes N}) \rightarrow \mathbb{R}_+$ is assumed to be measurable;
- for any $\omega \in \Omega$, the map $A \mapsto \rho(\omega, A) : \mathcal{T} \rightarrow \mathbb{R}_+$ is assumed to be a probability measure.

We will also assume without further notice that the σ -algebras we deal with are always countably generated. The technical implications of these assumptions are standard and discussed for instance in [17, pages 50-54] (where, among other things, a detailed proof of the decomposition of the Kullback Liebler divergence is given).

Let us restrict to the case when the constant λ is positive. We get from Theorem 1.4 that

$$\exp\left[\lambda\left\{\Phi_{\frac{\lambda}{N}}\left[\mathbb{P}[\rho(R)]\right] - \mathbb{P}[\rho(r)]\right\} - \mathbb{P}[\mathcal{K}(\rho, \pi)]\right] \leq 1, \quad (1.1)$$

where we have used the convexity of the exp function and of $\Phi_{\frac{\lambda}{N}}$. Since we have restricted our attention to positive values of the constant λ , Equation (1.1) can also be written

$$\mathbb{P}[\rho(R)] \leq \Phi_{\frac{\lambda}{N}}^{-1}\left\{\mathbb{P}\left[\rho(r) + \lambda^{-1}\mathcal{K}(\rho, \pi)\right]\right\},$$

leading to

THEOREM 1.5. *For any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, for any positive parameter λ ,*

$$\begin{aligned} \mathbb{P}[\rho(R)] &\leq \frac{1 - \exp\left[-N^{-1}\mathbb{P}\left[\lambda\rho(r) + \mathcal{K}(\rho, \pi)\right]\right]}{1 - \exp\left(-\frac{\lambda}{N}\right)} \\ &\leq \mathbb{P}\left\{\frac{\lambda}{N\left[1 - \exp\left(-\frac{\lambda}{N}\right)\right]}\left[\rho(r) + \frac{\mathcal{K}(\rho, \pi)}{\lambda}\right]\right\}. \end{aligned}$$

The last inequality provides the *unbiased empirical upper bound* for $\rho(R)$ we were looking for, meaning that the expectation of $\frac{\lambda}{N\left[1 - \exp\left(-\frac{\lambda}{N}\right)\right]}\left[\rho(r) + \frac{\mathcal{K}(\rho, \pi)}{\lambda}\right]$ is larger than the expectation of $\rho(R)$. Let us notice that $1 \leq \frac{\lambda}{N\left[1 - \exp\left(-\frac{\lambda}{N}\right)\right]} \leq \left[1 - \frac{\lambda}{2N}\right]^{-1}$ and therefore that this coefficient is close to 1 when λ is significantly smaller than N .

If we are ready to believe in this bound (although this belief is not mathematically well founded, as we already mentioned), we can use it to optimize λ and to choose ρ . While the optimal choice of ρ when λ is fixed is to take it equal to $\pi_{\exp(-\lambda r)}$, a Gibbs posterior distribution, as it is sometimes called, we may for computational reasons be more interested in choosing ρ in some other class of posterior distributions.

For instance, our real interest may be to select some deterministic estimator from a family $\hat{\theta}_m : \Omega \rightarrow \Theta_m$, $m \in M$, of possible ones, where Θ_m are measurable subsets of Θ and where M is an arbitrary (non necessarily countable) index set. We may for instance think of the case when

$\hat{\theta}_m \in \arg \min_{\Theta_m} r$. We may slightly randomize the estimators to start with, considering for any $\theta \in \Theta_m$ and any $m \in M$,

$$\Delta_m(\theta) = \left\{ \theta' \in \Theta_m : [f_{\theta'}(X_i)]_{i=1}^N = [f_{\theta}(X_i)]_{i=1}^N \right\},$$

and defining ρ_m by the formula

$$\frac{d\rho_m}{d\pi}(\theta) = \frac{\mathbf{1}[\theta \in \Delta_m(\hat{\theta}_m)]}{\pi[\Delta_m(\hat{\theta}_m)]}.$$

Our posterior is minimizing $\mathcal{K}(\rho, \pi)$ among those whose support is restricted to the values of θ in Θ_m for which the classification rule f_{θ} is identical to the estimated one $f_{\hat{\theta}_m}$ on the observed sample. Presumably, in many practical situations, $f_{\theta}(x)$ will be ρ_m almost surely identical to $f_{\hat{\theta}_m}(x)$ when θ is drawn from ρ_m , for the vast majority of the values of $x \in \mathcal{X}$ and all the submodels Θ_m not plagued with too much overfitting (since this is by construction the case when $x \in \{X_i : i = 1, \dots, N\}$). Therefore replacing $\hat{\theta}_m$ with ρ_m can be expected to be a minor change in many situations. This change by the way can be estimated in the (admittedly not so common) case when the distribution of the patterns $(X_i)_{i=1}^N$ is known. Indeed, introducing the pseudo distance

$$D(\theta, \theta') = \frac{1}{N} \sum_{i=1}^N \mathbb{P}[f_{\theta}(X_i) \neq f_{\theta'}(X_i)], \quad \theta, \theta' \in \Theta, \quad (1.2)$$

one immediately sees that $R(\theta') \leq R(\theta) + D(\theta, \theta')$, for any $\theta, \theta' \in \Theta$, and therefore that

$$R(\hat{\theta}_m) \leq \rho_m(R) + \rho_m[D(\cdot, \hat{\theta}_m)].$$

Let us notice also that in the case where $\Theta_m \subset \mathbb{R}^{d_m}$, and R happens to be convex on $\Delta_m(\hat{\theta}_m)$, then $\rho_m(R) \geq R[\int \theta \rho_m(d\theta)]$, and we can replace $\hat{\theta}_m$ with $\tilde{\theta}_m = \int \theta \rho_m(d\theta)$, and obtain bounds for $R(\tilde{\theta}_m)$. This is not a very heavy assumption about R , in the case where we consider $\hat{\theta}_m \in \arg \min_{\Theta_m} r$. Indeed, $\tilde{\theta}_m$, and therefore $\Delta_m(\tilde{\theta}_m)$, will be presumably close to $\arg \min_{\Theta_m} R$, and requiring a function to be convex in the neighborhood of its minima is not a very strong assumption.

Since $r(\hat{\theta}_m) = \rho_m(r)$, and $\mathcal{K}(\rho_m, \pi) = -\log\{\pi[\Delta_m(\hat{\theta}_m)]\}$, our unbiased empirical upper bound in this context reads as

$$\frac{\lambda}{N[1 - \exp(-\frac{\lambda}{N})]} \left\{ r(\hat{\theta}_m) - \frac{\log\{\pi[\Delta_m(\hat{\theta}_m)]\}}{\lambda} \right\}.$$

Let us notice that we obtain a complexity factor $-\log\{\pi[\Delta_m(\hat{\theta}_m)]\}$ which may be compared with the Vapnik Cervonenkis dimension. Indeed, in the case of binary classification, when using a classification model with VC dimension not greater than h_m , that is when any subset of \mathcal{X} which can be split in any arbitrary way by some classification rule f_θ of the model Θ_m has at most h_m points, then

$$\{\Delta_m(\theta) : \theta \in \Theta_m\}$$

is a partition of Θ_m with at most $\left(\frac{eN}{h_m}\right)^h$ components. Therefore

$$\inf_{\theta \in \Theta_m} -\log\{\pi[\Delta_m(\theta)]\} \leq h_m \log\left(\frac{eN}{h_m}\right) - \log[\pi(\Theta_m)].$$

Thus, if the model and prior distribution are well suited to the classification task, in the sense that there is more “room” (where room is measured with π) between the two clusters defined by $\hat{\theta}_m$ than between other partitions of the sample of patterns $(X_i)_{i=1}^N$, then we will have

$$-\log\{\pi[\Delta_m(\hat{\theta})]\} \leq h_m \log\left(\frac{eN}{h_m}\right) - \log[\pi(\Theta_m)].$$

An optimal value \hat{m} may be selected so that

$$\hat{m} \in \arg \min_{m \in M} \left\{ \inf_{\lambda \in \mathbb{R}_+} \frac{\lambda}{N[1 - \exp(-\frac{\lambda}{N})]} \left(r(\hat{\theta}_m) - \frac{\log\{\pi[\Delta_m(\hat{\theta}_m)]\}}{\lambda} \right) \right\}.$$

Since $\rho_{\hat{m}}$ is still another posterior distribution, we can be sure that

$$\begin{aligned} \mathbb{P}\left\{R(\hat{\theta}_{\hat{m}}) - \rho_{\hat{m}}[D(\cdot, \hat{\theta}_{\hat{m}})]\right\} &\leq \mathbb{P}[\rho_{\hat{m}}(R)] \\ &\leq \inf_{\lambda \in \mathbb{R}_+} \mathbb{P}\left\{\frac{\lambda}{N[1 - \exp(-\frac{\lambda}{N})]} \left(r(\hat{\theta}_{\hat{m}}) - \frac{\log\{\pi[\Delta_{\hat{m}}(\hat{\theta}_{\hat{m}})]\}}{\lambda} \right)\right\}. \end{aligned}$$

(Taking the infimum in λ inside the expectation with respect to \mathbb{P} would be possible at the price of some supplementary technicalities and a slight increase of the bound that we prefer to postpone to the discussion of deviation bounds, since they are the only ones to provide a rigorous mathematical foundation to the adaptive selection of estimators.)

1.2.2. *Optimizing explicitly the exponential parameter λ .* We would like to deal in this section with some technical issue we think helpful to the understanding of Theorem 1.5 (see page 13): namely to investigate how the upper bound it provides could be optimized, or at least approximately optimized, in λ . It turns out that this can be done quite explicitly.

So we will consider in this discussion the posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$ to be fixed, and our aim will be to eliminate the constant λ from the bound by choosing its value in some nearly optimal way as a function of $\mathbb{P}[\rho(r)]$, the average of the empirical risk, and of $\mathbb{P}[\mathcal{K}(\rho, \pi)]$, which controls overfitting.

Let the bound be written as

$$\varphi(\lambda) = [1 - \exp(-\frac{\lambda}{N})]^{-1} \left\{ 1 - \exp \left[-\frac{\lambda}{N} \mathbb{P}[\rho(r)] - N^{-1} \mathbb{P}[\mathcal{K}(\rho, \pi)] \right] \right\}.$$

We see that

$$N \frac{\partial}{\partial \lambda} \log[\varphi(\lambda)] = \frac{\mathbb{P}[\rho(r)]}{\exp \left[\frac{\lambda}{N} \mathbb{P}[\rho(r)] + N^{-1} \mathbb{P}[\mathcal{K}(\rho, \pi)] \right] - 1} - \frac{1}{\exp(\frac{\lambda}{N}) - 1}.$$

Thus, the optimal value for λ is such that

$$\left[\exp(\frac{\lambda}{N}) - 1 \right] \mathbb{P}[\rho(r)] = \exp \left[\frac{\lambda}{N} \mathbb{P}[\rho(r)] + N^{-1} \mathbb{P}[\mathcal{K}(\rho, \pi)] \right] - 1.$$

Assuming that $1 \gg \frac{\lambda}{N} \mathbb{P}[\rho(r)] \gg \frac{\mathbb{P}[\mathcal{K}(\rho, \pi)]}{N}$, and keeping only higher order terms, we are led to choose

$$\lambda = \sqrt{\frac{2N \mathbb{P}[\mathcal{K}(\rho, \pi)]}{\mathbb{P}[\rho(r)] \{1 - \mathbb{P}[\rho(r)]\}}},$$

obtaining

THEOREM 1.6. *For any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\mathbb{P}[\rho(R)] \leq \frac{1 - \exp \left\{ -\sqrt{\frac{2\mathbb{P}[\mathcal{K}(\rho, \pi)]\mathbb{P}[\rho(r)]}{N\{1 - \mathbb{P}[\rho(r)]\}}} - \frac{\mathbb{P}[\mathcal{K}(\rho, \pi)]}{N} \right\}}{1 - \exp \left\{ -\sqrt{\frac{2\mathbb{P}[\mathcal{K}(\rho, \pi)]}{N\mathbb{P}[\rho(r)]\{1 - \mathbb{P}[\rho(r)]\}}} \right\}}.$$

This result of course is not very useful in itself, since none of the two quantities $\mathbb{P}[\rho(r)]$ and $\mathbb{P}[\mathcal{K}(\rho, \pi)]$ are easy to evaluate. Anyhow it gives a hint that replacing them boldly with $\rho(r)$ and $\mathcal{K}(\rho, \pi)$ could produce something

close to a legitimate empirical upper bound for $\rho(R)$. We will see in the subsection about deviation bounds that this is indeed essentially true.

Let us remark that in the second section of these notes, we will see another way of bounding

$$\inf_{\lambda \in \mathbb{R}_+} \Phi_{\frac{\lambda}{N}}^{-1} \left(q + \frac{d}{\lambda} \right), \text{ leading to}$$

THEOREM 1.7. *For any prior distribution $\pi \in \mathcal{M}_+^1(\Theta)$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned} \mathbb{P}[\rho(R)] \leq & \left(1 + \frac{2\mathbb{P}[\mathcal{K}(\rho, \pi)]}{N} \right)^{-1} \left\{ \mathbb{P}[\rho(r)] + \frac{\mathbb{P}[\mathcal{K}(\rho, \pi)]}{N} \right. \\ & \left. + \sqrt{\frac{2\mathbb{P}[\mathcal{K}(\rho, \pi)]\mathbb{P}[\rho(r)]\{1 - \mathbb{P}[\rho(r)]\}}{N} + \frac{\mathbb{P}[\mathcal{K}(\rho, \pi)]^2}{N^2}} \right\}, \end{aligned}$$

$$\text{as soon as } \mathbb{P}[\rho(r)] + \sqrt{\frac{\mathbb{P}[\mathcal{K}(\rho, \pi)]}{2N}} \leq \frac{1}{2},$$

$$\text{and } \mathbb{P}[\rho(R)] \leq \mathbb{P}[\rho(r)] + \sqrt{\frac{\mathbb{P}[\mathcal{K}(\rho, \pi)]}{2N}} \text{ otherwise.}$$

This theorem enlightens the influence of three terms on the average expected risk :

- the average empirical risk, $\mathbb{P}[\rho(r)]$, which as a rule will decrease as the size of the classification model increases, acts as a *bias* term, grasping the ability of the model to account for the observed sample itself;
- a *variance* term $\mathbb{P}[\rho(r)]\{1 - \mathbb{P}[\rho(r)]\}$ is due to the random fluctuations of $\rho(r)$;
- a *complexity* term $\mathbb{P}[\mathcal{K}(\rho, \pi)]$, which as a rule will increase with the size of the classification model, eventually acts as a multiplier of the variance term.

We observed numerically that the bound provided by Theorem 1.6 is better than the more classical Vapnik's like bound of Theorem 1.7. For instance, when $N = 1000$, $\mathbb{P}[\rho(r)] = 0.2$ and $\mathbb{P}[\mathcal{K}(\rho, \pi)] = 10$, Theorem 1.6 gives a bound lower than 0.2604, whereas the more classical Vapnik's like approximation of Theorem 1.7 gives a bound larger than 0.2622. Numerical simulations tend to suggest the two bounds are always ordered in the same way, although this could be a little tedious to prove mathematically.

1.2.3. Non random bounds. It is time now to come to less tentative results and see how far is the average expected error rate $\mathbb{P}[\rho(R)]$ from its best possible value $\inf_{\Theta} R$.

Let us notice first that

$$\lambda\rho(r) + \mathcal{K}(\rho, \pi) = \mathcal{K}(\rho, \pi_{\exp(-\lambda r)}) - \log\left\{\pi[\exp(-\lambda r)]\right\}.$$

Let us remark moreover that $r \mapsto \log\left\{\pi[\exp(-\lambda r)]\right\}$ is a convex functional, a property which can be used in the following way:

$$\begin{aligned} \mathbb{P}\left\{\log\left[\pi[\exp(-\lambda r)]\right]\right\} &= \mathbb{P}\left\{\sup_{\rho \in \mathcal{M}_+^1(\Theta)} -\lambda\rho(r) - \mathcal{K}(\rho, \pi)\right\} \\ &\geq \sup_{\rho \in \mathcal{M}_+^1(\Theta)} \mathbb{P}\left\{-\lambda\rho(r) - \mathcal{K}(\rho, \pi)\right\} = \sup_{\rho \in \mathcal{M}_+^1(\Theta)} -\lambda\rho(R) - \mathcal{K}(\rho, \pi) \\ &= \log\left\{\pi[\exp(-\lambda R)]\right\} = -\int_0^\lambda \pi_{\exp(-\beta R)}(R) d\beta. \end{aligned} \quad (1.3)$$

These remarks applied to Theorem 1.5 lead to

THEOREM 1.8. *For any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, for any positive parameter λ ,*

$$\begin{aligned} \mathbb{P}[\rho(R)] &\leq \frac{1 - \exp\left\{-\frac{1}{N} \int_0^\lambda \pi_{\exp(-\beta R)}(R) d\beta - \frac{1}{N} \mathbb{P}[\mathcal{K}(\rho, \pi_{\exp(-\lambda r)})]\right\}}{1 - \exp(-\frac{\lambda}{N})} \\ &\leq \frac{1}{N[1 - \exp(-\frac{\lambda}{N})]} \left\{ \int_0^\lambda \pi_{\exp(-\beta R)}(R) d\beta + \mathbb{P}[\mathcal{K}(\rho, \pi_{\exp(-\lambda r)})] \right\}. \end{aligned}$$

This theorem is particularly well fitted for the case of the Gibbs posterior distribution $\rho = \pi_{\exp(-\lambda r)}$, where the entropy factor cancels and where $\mathbb{P}[\pi_{\exp(-\lambda r)}(R)]$ is shown to be bound to get close to $\inf_{\Theta} R$ when N goes to ∞ , as soon as λ/N goes to 0 while λ goes to $+\infty$.

We can elaborate on Theorem 1.8 and define a notion of dimension of (Θ, R) , with margin $\eta \geq 0$ putting

$$\begin{aligned} d_\eta(\Theta, R) &= \sup_{\beta \in \mathbb{R}_+} \beta \left[\pi_{\exp(-\beta R)}(R) - \operatorname{ess\,inf}_\pi R - \eta \right] \\ &\leq -\log\left\{\pi[R \leq \operatorname{ess\,inf}_\pi R + \eta]\right\}. \end{aligned} \quad (1.4)$$

This last inequality can be established by the chain of inequalities:

$$\begin{aligned} \beta \pi_{\exp(-\beta R)}(R) &\leq \int_0^\beta \pi_{\exp(-\gamma R)}(R) d\gamma = -\log \left\{ \pi \left[\exp(-\beta R) \right] \right\} \\ &\leq \beta \left(\operatorname{ess\,inf}_\pi R + \eta \right) - \log \left[\pi \left(R \leq \operatorname{ess\,inf}_\pi R + \eta \right) \right], \end{aligned}$$

where we have used successively the fact that $\lambda \mapsto \pi_{\exp(-\lambda R)}(R)$ is decreasing (because it is the derivative of the concave function $\lambda \mapsto -\log \left\{ \pi \left[\exp(-\lambda R) \right] \right\}$) and the fact that the exponential function takes positive values.

In typical “parametric” situations $d_0(\Theta, R)$ will be finite, and in all circumstances $d_\eta(\Theta, R)$ will be finite for any $\eta > 0$ (this is a direct consequence of the definition of the essential infimum). Using this notion of dimension, we see that

$$\begin{aligned} \int_0^\lambda \pi_{\exp(-\beta R)}(R) d\beta &\leq \lambda \left(\operatorname{ess\,inf}_\pi R + \eta \right) \\ &\quad + \int_0^\lambda \left[\frac{d_\eta}{\beta} \wedge (1 - \operatorname{ess\,inf}_\pi R - \eta) \right] d\beta \\ &= \lambda \left(\operatorname{ess\,inf}_\pi R + \eta \right) + d_\eta(\Theta, R) \log \left[\frac{e\lambda}{d_\eta(\Theta, R)} (1 - \operatorname{ess\,inf}_\pi R - \eta) \right]. \end{aligned}$$

This leads to

COROLLARY 1.9 *With the above notations, for any margin $\eta \in \mathbb{R}_+$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\mathbb{P}[\rho(R)] \leq \inf_{\lambda \in \mathbb{R}_+} \Phi_{\frac{\lambda}{N}}^{-1} \left[\operatorname{ess\,inf}_\pi R + \eta + \frac{d_\eta}{\lambda} \log \left(\frac{e\lambda}{d_\eta} \right) + \frac{\mathbb{P} \left\{ \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}] \right\}}{\lambda} \right].$$

If one is wanting a posterior distribution with a small support, the theorem can also be applied to the case when ρ is obtained by truncating $\pi_{\exp(-\lambda r)}$ to some level set to reduce its support: let $\Theta_p = \{\theta \in \Theta : r(\theta) \leq p\}$, and let us define for any $q \in (0, 1)$ the level $p_q = \inf\{p : \pi_{\exp(-\lambda r)}(\Theta_p) \geq q\}$, let us then define ρ_q by its density

$$\frac{d\rho_q}{d\pi_{\exp(-\lambda r)}}(\theta) = \frac{\mathbb{1}(\theta \in \Theta_{p_q})}{\pi_{\exp(-\lambda r)}(\Theta_{p_q})},$$

then $\rho_0 = \pi_{\exp(-\lambda r)}$ and for any $q \in (0, 1)$,

$$\begin{aligned} \mathbb{P}[\rho_q(R)] &\leq \frac{1 - \exp \left\{ -\frac{1}{N} \int_0^\lambda \pi_{\exp(-\beta R)}(R) d\beta - \frac{\log(q)}{N} \right\}}{1 - \exp(-\frac{\lambda}{N})} \\ &\leq \frac{1}{N [1 - \exp(-\frac{\lambda}{N})]} \left\{ \int_0^\lambda \pi_{\exp(-\beta R)}(R) d\beta - \log(q) \right\}. \end{aligned}$$

1.2.4. Deviation bounds. They provide results holding under the distribution \mathbb{P} of the sample with probability at least $1 - \epsilon$, for any given confidence level, set by the choice of $\epsilon \in]0, 1[$. Using them is the only way to be quite (i.e. with probability $1 - \epsilon$) sure to do the right thing, although this right thing may be overpessimistic, since deviation upper bounds are larger than corresponding non biased bounds.

Starting again from Theorem 1.4, and using Markov's inequality $\mathbb{P}[\exp(h) \geq 1] \leq \mathbb{P}[\exp(h)]$, we obtain

THEOREM 1.10. *For any positive parameter λ , with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned} \rho(R) &\leq \Phi_{\frac{\lambda}{N}}^{-1} \left\{ \rho(r) + \frac{\mathcal{K}(\rho, \pi) - \log(\epsilon)}{\lambda} \right\} \\ &= \frac{1 - \exp \left\{ -\frac{\lambda \rho(r)}{N} - \frac{\mathcal{K}(\rho, \pi) - \log(\epsilon)}{N} \right\}}{1 - \exp(-\frac{\lambda}{N})} \\ &\leq \frac{\lambda}{N [1 - \exp(-\frac{\lambda}{N})]} \left[\rho(r) + \frac{\mathcal{K}(\rho, \pi) - \log(\epsilon)}{\lambda} \right]. \end{aligned}$$

We see that for a fixed value of the parameter λ , the upper bound is optimized when the posterior is chosen to be the Gibbs distribution $\rho = \pi_{\exp(-\lambda r)}$.

Moreover we would like to be entitled to optimize the bound in λ . Gaining the required uniformity in λ can be done in the following way. Let us notice first that values of λ less than 1 are not interesting (because they provide a bound larger than one, at least as soon as $\epsilon \leq \exp(-1)$). Let us consider some real parameter $\alpha > 1$, and the set $\Lambda = \{\alpha^k; k \in \mathbb{N}\}$. Let us put on this set the probability measure $\nu(\alpha^k) = [(k+1)(k+2)]^{-1}$. Applying the previous theorem to $\lambda = \alpha^k$ at confidence level $1 - \frac{\epsilon}{(k+1)(k+2)}$, and using a union bound, we see that with probability at least $1 - \epsilon$, for any posterior distribution ρ ,

$$\rho(R) \leq \inf_{\lambda' \in \Lambda} \Phi_{\frac{\lambda'}{N}}^{-1} \left\{ \rho(r) + \frac{\mathcal{K}(\rho, \pi) - \log(\epsilon) + 2 \log \left[\frac{\log(\alpha^2 \lambda')}{\log(\alpha)} \right]}{\lambda'} \right\}.$$

Now we can remark that for any $\lambda \in (1, +\infty[$, there is $\lambda' \in \Lambda$ such that $\alpha^{-1} \lambda \leq \lambda' \leq \lambda$. Moreover, for any $q \in (0, 1)$, $\beta \mapsto \Phi_{\beta}^{-1}(q)$ is increasing on \mathbb{R}_+ . Thus with probability at least $1 - \epsilon$, for any posterior distribution ρ ,

$$\rho(R) \leq \inf_{\lambda \in (1, \infty[} \Phi_{\frac{\lambda}{N}}^{-1} \left\{ \rho(r) + \frac{\alpha}{\lambda} \left[\mathcal{K}(\rho, \pi) - \log(\epsilon) + 2 \log \left(\frac{\log(\alpha^2 \lambda)}{\log(\alpha)} \right) \right] \right\}$$

$$= \inf_{\lambda \in (1, \infty)} \frac{1 - \exp \left\{ -\frac{\lambda}{N} \rho(r) - \frac{\alpha}{N} \left[\mathcal{K}(\rho, \pi) - \log(\epsilon) + 2 \log \left(\frac{\log(\alpha^2 \lambda)}{\log(\alpha)} \right) \right] \right\}}{1 - \exp(-\frac{\lambda}{N})}.$$

Taking the approximately optimal value

$$\lambda = \sqrt{\frac{2N\alpha [\mathcal{K}(\rho, \pi) - \log(\epsilon)]}{\rho(r)[1 - \rho(r)]}},$$

we obtain

THEOREM 1.11. *With probability $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, putting $d(\rho, \epsilon) = \mathcal{K}(\rho, \pi) - \log(\epsilon)$,*

$$\begin{aligned} \rho(R) &\leq \inf_{k \in \mathbb{N}} \frac{1 - \exp \left\{ -\frac{\alpha^k}{N} \rho(r) - \frac{1}{N} \left[d(\rho, \epsilon) + \log[(k+1)(k+2)] \right] \right\}}{1 - \exp \left(-\frac{\alpha^k}{N} \right)} \\ &\leq \frac{1 - \exp \left\{ -\sqrt{\frac{2\alpha\rho(r)d(\rho, \epsilon)}{N[1 - \rho(r)]}} - \frac{\alpha}{N} \left[d(\rho, \epsilon) + 2 \log \left(\frac{\log(\alpha^2 \sqrt{\frac{2N\alpha d(\rho, \epsilon)}{\rho(r)[1 - \rho(r)]}})}{\log(\alpha)} \right) \right] \right\}}{1 - \exp \left[-\sqrt{\frac{2\alpha d(\rho, \epsilon)}{N\rho(r)[1 - \rho(r)]}} \right]}. \end{aligned}$$

Moreover with probability at least $1 - \epsilon$, for any posterior distribution ρ such that $\rho(r) = 0$,

$$\rho(R) \leq 1 - \exp \left[-\frac{\mathcal{K}(\rho, \pi) - \log(\epsilon)}{N} \right].$$

We can also elaborate on the results in an other direction by introducing the *empirical dimension*

$$d_e = \sup_{\beta \in \mathbb{R}_+} \beta \left[\pi_{\exp(-\beta r)}(r) - \operatorname{ess\,inf}_{\pi} r \right] \leq -\log \left[\pi(r = \operatorname{ess\,inf}_{\pi} r) \right]. \quad (1.5)$$

(There is no need to introduce a margin in this definition, since r takes at most N values, and therefore $\pi(r = \operatorname{ess\,inf}_{\pi} r)$ will be strictly positive.) This leads to

COROLLARY 1.12. *For any positive real constant λ , with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\rho(R) \leq \Phi_{\frac{\lambda}{N}}^{-1} \left[\operatorname{ess\,inf}_{\pi} r + \frac{d_e}{\lambda} \log \left(\frac{e\lambda}{d_e} \right) + \frac{\mathcal{K}[\rho, \pi_{\exp(-\lambda r)}] - \log(\epsilon)}{\lambda} \right].$$

We could then make the bound uniform in λ and optimize this parameter in a way similar to what was done to obtain Theorem 1.11.

1.3. LOCAL BOUNDS. In this subsection, better bounds will be achieved through a better choice of the prior distribution. This better prior distribution turns out to depend on the unknown sample distribution \mathbb{P} , and some work is required to circumvent this and obtain empirical bounds.

1.3.1. Choice of the prior. As mentioned in the introduction, if one is willing to minimize the bound in expectation provided by Theorem 1.5 (page 13), one is led to consider the optimal choice $\pi = \mathbb{P}(\rho)$. However, this is but an ideal choice, since \mathbb{P} is in all conceivable situations unknown. Nevertheless it shows that it is possible through Theorem 1.5 to measure the *complexity* of the classification model with $\mathbb{P}\{\mathcal{K}[\rho, \mathbb{P}(\rho)]\}$, which is nothing but the *mutual information* between the random sample $(X_i, Y_i)_{i=1}^N$ and the estimated parameter $\hat{\theta}$, when the sample is drawn according to \mathbb{P} and the estimated parameter knowing the sample is drawn according to ρ .

In practice, since we cannot choose $\pi = \mathbb{P}(\rho)$, we have to be content with a *flat* prior π , resulting in a bound measuring complexity according to $\mathbb{P}[\mathcal{K}(\rho, \pi)] = \mathbb{P}\{\mathcal{K}[\rho, \mathbb{P}(\rho)]\} + \mathcal{K}[\mathbb{P}(\rho), \pi]$ larger by the entropy factor $\mathcal{K}[\mathbb{P}(\rho), \pi]$ than the optimal one (we are still commenting on Theorem 1.5).

If we want to base the choice of π on Theorem 1.8 (page 18), and if we choose $\rho = \pi_{\exp(-\lambda r)}$ to optimize this bound, we will be inclined to choose some π such that

$$\frac{1}{\lambda} \int_0^\lambda \pi_{\exp(-\beta R)}(R) d\beta = -\frac{1}{\lambda} \log \left\{ \pi[\exp(-\lambda R)] \right\}$$

is as far as possible close to $\inf_{\theta \in \Theta} R(\theta)$ in all circumstances. To give some more specific example, in the case when the distribution of the design $(X_i)_{i=1}^N$ is known, one can introduce on the parameter space Θ the metric D already defined by equation (1.2, page 14) (or some available upper bound for this distance). In view of the fact that $R(\theta) - R(\theta') \leq D(\theta, \theta')$, for any $\theta, \theta' \in \Theta$, it can be meaningful, at least theoretically, to choose π as

$$\pi = \sum_{k=1}^{\infty} \frac{1}{k(k+1)} \pi_k,$$

where π_k is the uniform measure on some minimal (or close to minimal) 2^{-k} -net $\mathcal{N}(\Theta, D, 2^{-k})$ of the metric space (Θ, D) . With this choice

$$-\frac{1}{\lambda} \log \left\{ \pi \left[\exp(-\lambda R) \right] \right\} \leq \inf_{\theta \in \Theta} R(\theta) + \inf_k \left\{ 2^{-k} + \frac{\log(|\mathcal{N}(\Theta, D, 2^{-k})|) + \log[k(k+1)]}{\lambda} \right\}.$$

Another possibility, when we have to deal with real valued parameters, meaning that $\Theta \subset \mathbb{R}^d$, is to code each real component $\theta_i \in \mathbb{R}$ of $\theta = (\theta_i)_{i=1}^d$ to some precision and to use a prior μ which is atomic on dyadic numbers. More precisely let us parametrize the set of dyadic real numbers as

$$\mathcal{D} = \left\{ r[s, m, p, (b_j)_{j=1}^p] = s2^m \left(1 + \sum_{j=1}^p b_j 2^{-j} \right) : s \in \{-1, +1\}, m \in \mathbb{Z}, p \in \mathbb{N}, b_j \in \{0, 1\} \right\},$$

where, as can be seen, s codes the sign, m the order of magnitude, p the precision and $(b_j)_{j=1}^p$ the binary representation of the dyadic number $r[s, m, p, (b_j)_{j=1}^p]$. We can for instance consider on \mathcal{D} the probability distribution

$$\mu \{ r[s, m, p, (b_j)_{j=1}^p] \} = \left[3(|m| + 1)(|m| + 2)(p + 1)(p + 2)2^p \right]^{-1}, \quad (1.6)$$

and define $\pi \in \mathcal{M}_+^1(\mathbb{R}^d)$ as $\pi = \mu^{\otimes d}$. This kind of ‘‘coding’’ prior distribution can be used also to define a prior on the integers (by renormalizing the restriction of μ to integers to get a probability distribution). Using μ is somehow equivalent to picking up a representative of each dyadic interval, and makes it possible to restrict to the case when the posterior ρ is a Dirac mass without losing too much (when $\Theta = (0, 1)$, this approach is somewhat equivalent to considering as prior distribution the Lebesgue measure and using as posterior distributions the uniform probability measures on dyadic intervals, with the advantage of obtaining non randomized estimators). When one uses in this way an atomic prior and Dirac masses as posterior distributions, the bounds proven so far can be obtained through a simpler union bound argument. This is so true that some of the detractors of the PAC-Bayesian approach (which, as a newcomer, has sometimes received a suspicious greeting among statisticians) have argued that it cannot bring anything that elementary union bound arguments could not essentially provide. We do not share of course this derogatory opinion, and while we think that allowing for non atomic priors and posteriors is worthwhile, we

also would like to stress that next to come local and relative bounds could hardly be obtained with the only help of union bounds.

Although the choice of a *flat* prior seems at first glance to be the only alternative when nothing is known about the sample distribution \mathbb{P} , the previous discussion shows that this type of choice is lacking proper localisation, and namely that we loose a factor $\mathcal{K}\{\mathbb{P}[\pi_{\exp(-\lambda r)}], \pi\}$, the divergence between the bound-optimal prior $\mathbb{P}[\pi_{\exp(-\lambda r)}]$, which is concentrated near the minima of R in favourable situations, and the flat prior π . Fortunately, there are technical ways to get around this difficulty and to obtain more local empirical bounds.

1.3.2. Unbiased local empirical bounds. The idea is to start with some flat prior $\pi \in \mathcal{M}_+^1(\Theta)$, and the posterior distribution $\rho = \pi_{\exp(-\lambda r)}$ minimizing the bound of Theorem 1.5 (page 13), when π is used as a prior. To improve the bound, we would like to use $\mathbb{P}[\pi_{\exp(-\lambda r)}]$ instead of π , and we are going to make the guess that we could approximate it with $\pi_{\exp(-\beta R)}$ (we have replaced the parameter λ with some distinct parameter β to give some more freedom to our investigation, and also because, intuitively, $\mathbb{P}[\pi_{\exp(-\lambda r)}]$ may be expected to be less concentrated than each of the $\pi_{\exp(-\lambda r)}$ it is mixing, which suggests that the best approximation of $\mathbb{P}[\pi_{\exp(-\lambda r)}]$ by some $\pi_{\exp(-\beta R)}$ may be obtained for some parameter $\beta < \lambda$). We are then led to look for some empirical upper bound of $\mathcal{K}[\rho, \pi_{\exp(-\beta R)}]$. This is happily provided by the following computation

$$\begin{aligned} \mathbb{P}\{\mathcal{K}[\rho, \pi_{\exp(-\beta R)}]\} &= \mathbb{P}\{\mathcal{K}(\rho, \pi) + \beta\mathbb{P}[\rho(R)] + \log\{\pi[\exp(-\beta R)]\}\} \\ &= \mathbb{P}\{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]\} + \beta\mathbb{P}[\rho(R - r)] \\ &\quad + \log\{\pi[\exp(-\beta R)]\} - \mathbb{P}\{\log \pi[\exp(-\beta r)]\}. \end{aligned}$$

Using the convexity of $r \mapsto \log\{\pi[\exp(-\beta r)]\}$ as in equation (1.3) on page 18, we see that

$$0 \leq \mathbb{P}\{\mathcal{K}[\rho, \pi_{\exp(-\beta R)}]\} \leq \beta\mathbb{P}[\rho(R - r)] + \mathbb{P}\{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]\}.$$

This inequality has an interest of its own, since it provides a lower bound for $\mathbb{P}[\rho(R)]$. Moreover we can plug it into Theorem 1.5 (page 13) applied to the prior distribution $\pi_{\exp(-\beta R)}$ and obtain for any posterior distribution ρ and any positive parameter λ that

$$\Phi_{\frac{\lambda}{N}}\{\mathbb{P}[\rho(R)]\} \leq \mathbb{P}\left\{\rho(r) + \frac{\beta}{\lambda}\rho(R - r) + \frac{1}{\lambda}\mathbb{P}\{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]\}\right\}.$$

In view of this, it is convenient to introduce the function

$$\begin{aligned}\tilde{\Phi}_{a,b}(p) &= (1-b)^{-1}[\Phi_a(p) - bp] \\ &= -(1-b)^{-1}\left\{a^{-1}\log\{1-p[1-\exp(-a)]\} + bp\right\}, \\ &\quad p \in (0, 1), a \in (0, \infty), b \in (0, 1).\end{aligned}$$

This is a convex function of p , moreover

$$\tilde{\Phi}'_{a,b}(0) = \left\{a^{-1}[1 - \exp(-a)] - b\right\}(1-b)^{-1},$$

showing that it is an increasing one to one convex map of the unit interval onto itself as soon as $b \leq a^{-1}[1 - \exp(-a)]$. Its convexity, combined with the value of its derivative at the origin, shows that

$$\tilde{\Phi}_{a,b}(p) \geq \frac{a^{-1}[1 - \exp(-a)] - b}{1-b}p.$$

Using these notations and remarks, we can state

THEOREM 1.13. *For any positive real constants β and λ such that $0 \leq \beta < N[1 - \exp(-\frac{\lambda}{N})]$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned}\mathbb{P}\left\{\rho(r) - \frac{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]}{\beta}\right\} &\leq \mathbb{P}[\rho(R)] \\ &\leq \tilde{\Phi}_{\frac{\lambda}{N}, \frac{\beta}{\lambda}}^{-1}\left\{\mathbb{P}\left[\rho(r) + \frac{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]}{\lambda - \beta}\right]\right\} \\ &\leq \frac{\lambda - \beta}{N[1 - \exp(-\frac{\lambda}{N})] - \beta}\mathbb{P}\left[\rho(r) + \frac{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]}{\lambda - \beta}\right].\end{aligned}$$

Thus (taking $\lambda = 2\beta$), for any β such that $0 \leq \beta < \frac{N}{2}$,

$$\mathbb{P}[\rho(R)] \leq \frac{1}{1 - \frac{2\beta}{N}}\mathbb{P}\left\{\rho(r) + \frac{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]}{\beta}\right\}.$$

Note that the last inequality is obtained using the fact that $1 - \exp(-x) \geq x - \frac{x^2}{2}$, $x \in \mathbb{R}_+$.

COROLLARY 1.14. *For any $\beta \in (0, N)$,*

$$\begin{aligned}
\mathbb{P}[\pi_{\exp(-\beta r)}(r)] &\leq \mathbb{P}[\pi_{\exp(-\beta r)}(R)] \\
&\leq \inf_{\lambda \in (-N \log(1 - \frac{\beta}{N}), \infty)} \frac{\lambda - \beta}{N[1 - \exp(-\frac{\lambda}{N})] - \beta} \mathbb{P}[\pi_{\exp(-\beta r)}(r)] \\
&\leq \frac{1}{1 - \frac{2\beta}{N}} \mathbb{P}[\pi_{\exp(-\beta r)}(r)],
\end{aligned}$$

the last inequality holding only when $\beta < \frac{N}{2}$.

It is interesting to compare the upper bound provided by this corollary with Theorem 1.5 on page 13 when the posterior is a Gibbs measure $\rho = \pi_{\exp(-\beta r)}$. We see that we have succeeded to get rid of the entropy term $\mathcal{K}[\pi_{\exp(-\beta r)}, \pi]$, but at the price of an increase of the multiplicative factor, which for small values of $\frac{\beta}{N}$ grows from $(1 - \frac{\beta}{2N})^{-1}$ (when we take $\lambda = \beta$ in Theorem 1.5), to $(1 - \frac{2\beta}{N})^{-1}$. Therefore non localized bounds have an interest of their own, and are superseded by localized bounds only in favourable circumstances (presumably when the sample is large enough when compared with the complexity of the classification model).

Corollary 1.14 shows that when $\frac{2\beta}{N}$ is small, $\pi_{\exp(-\beta r)}(r)$ is a tight approximation of $\pi_{\exp(-\beta r)}(R)$ in the mean (since we have an upper bound and a lower bound which are close together).

Another corollary is obtained by optimizing the bound given by Theorem 1.13 in ρ , which is done by taking $\rho = \pi_{\exp(-\lambda r)}$.

COROLLARY 1.15. *For any positive real constants β and λ such that $0 \leq \beta < N[1 - \exp(-\frac{\lambda}{N})]$,*

$$\begin{aligned}
\mathbb{P}[\pi_{\exp(-\lambda r)}(R)] &\leq \tilde{\Phi}_{\frac{\lambda}{N}, \frac{\beta}{\lambda}}^{-1} \left\{ \mathbb{P} \left[\frac{1}{\lambda - \beta} \int_{\beta}^{\lambda} \pi_{\exp(-\gamma r)}(r) d\gamma \right] \right\} \\
&\leq \frac{1}{N[1 - \exp(-\frac{\lambda}{N})] - \beta} \mathbb{P} \left[\int_{\beta}^{\lambda} \pi_{\exp(-\gamma r)}(r) d\gamma \right].
\end{aligned}$$

Although this inequality gives by construction a better upper bound for $\inf_{\lambda \in \mathbb{R}_+} \mathbb{P}[\pi_{\exp(-\lambda r)}(R)]$ than Corollary 1.14, it is not easy to tell which one of the two inequalities is the best to bound $\mathbb{P}[\pi_{\exp(-\lambda r)}(R)]$ for a fixed (and possibly suboptimal) value of λ , because in this case, one factor is improved while the other is worsened.

Using the *empirical dimension* d_e defined by equation (1.5) on page 21, we see that

$$\frac{1}{\lambda - \beta} \int_{\beta}^{\lambda} \pi_{\exp(-\gamma r)}(r) d\gamma \leq \operatorname{ess\,inf}_{\pi} r + d_e \log \left(\frac{\lambda}{\beta} \right).$$

Therefore, in the case when we keep the ratio $\frac{\lambda}{\beta}$ bounded, we get a better dependence on the empirical dimension d_e than in Corollary 1.12 (page 21).

1.3.3. Non random local bounds. Let us come now to the localization of the non random upper bound given by Theorem 1.8 on page 18. According to Theorem 1.5 (page 13) applied to the localized prior $\pi_{\exp(-\beta R)}$,

$$\begin{aligned} \lambda \Phi_{\frac{\lambda}{N}} \{ \mathbb{P}[\rho(R)] \} &\leq \mathbb{P} \left\{ \lambda \rho(r) + \mathcal{K}(\rho, \pi) + \beta \rho(R) \right\} + \log \{ \pi[\exp(-\beta R)] \} \\ &= \mathbb{P} \left\{ \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}] - \log \{ \pi[\exp(-\lambda r)] \} + \beta \rho(R) \right\} + \log \{ \pi[\exp(-\beta R)] \} \\ &\leq \mathbb{P} \left\{ \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}] + \beta \rho(R) \right\} - \log \{ \pi[\exp(-\lambda R)] \} + \log \{ \pi[\exp(-\beta R)] \}, \end{aligned}$$

where we have used as previously inequality (1.3) (page 18). This proves

THEOREM 1.16. *For any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, for any real parameters β and λ such that $0 \leq \beta < N[1 - \exp(-\frac{\lambda}{N})]$,*

$$\begin{aligned} \mathbb{P}[\rho(R)] &\leq \tilde{\Phi}_{\frac{\lambda}{N}, \frac{\beta}{\lambda}}^{-1} \left\{ \frac{1}{\lambda - \beta} \int_{\beta}^{\lambda} \pi_{\exp(-\gamma R)}(R) d\gamma + \mathbb{P} \left[\frac{\mathcal{K}[\rho, \pi_{\exp(-\lambda r)}]}{\lambda - \beta} \right] \right\} \\ &\leq \frac{1}{N[1 - \exp(-\frac{\lambda}{N})] - \beta} \left\{ \int_{\beta}^{\lambda} \pi_{\exp(-\gamma R)}(R) d\gamma + \mathbb{P} \left\{ \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}] \right\} \right\}. \end{aligned}$$

Let us notice in particular that this theorem contains Theorem 1.8 (page 18) which corresponds to the case $\beta = 0$. As a corollary, we see also, taking $\rho = \pi_{\exp(-\lambda r)}$ and $\lambda = 2\beta$, and noticing that $\gamma \mapsto \pi_{\exp(-\gamma R)}(R)$ is decreasing, that

$$\begin{aligned} \mathbb{P}[\pi_{\exp(-\lambda r)}(R)] &\leq \inf_{\beta, \beta < N[1 - \exp(-\frac{\lambda}{N})]} \frac{\beta}{N[1 - \exp(-\frac{\lambda}{N})] - \beta} \pi_{\exp(-\beta R)}(R) \\ &\leq \frac{1}{1 - \frac{\lambda}{N}} \pi_{\exp(-\frac{\lambda}{2} R)}(R). \end{aligned}$$

We can use this inequality in conjunction with the notion of dimension with margin η introduced by equation (1.4) on page 18, to see that the Gibbs posterior achieves for a proper choice of λ and any margin parameter $\eta \geq 0$ (which can be chosen to be equal to zero in parametric situations)

$$\inf_{\lambda} \mathbb{P}[\pi_{\exp(-\lambda r)}(R)] \leq \operatorname{ess\,inf}_{\pi} R + \eta + \frac{4d_{\eta}}{N}$$

$$+ 2\sqrt{\frac{2d_\eta(\text{ess inf}_\pi R + \eta)}{N} + \frac{4d_\eta^2}{N^2}}. \quad (1.7)$$

Deviation bounds to come next will show that the optimal λ can be estimated from empirical data.

Let us propose a little numerical example as an illustration : assuming that $d_0 = 10$, $N = 1000$ and $\text{ess inf}_\pi R = 0.2$, we obtain from equation (1.7) that $\inf_\lambda \mathbb{P}[\pi_{\exp(-\lambda r)}(R)] \leq 0.373$.

1.3.4. Local deviation bounds. When it comes to deviation bounds, we will for technical reasons choose a slightly more involved change of prior distribution and apply Theorem 1.10 (page 20) to the prior $\pi_{\exp[-\beta\Phi_{-\frac{\beta}{N}} \circ R]}$. The advantage of tweaking R with the nonlinear function $\Phi_{-\frac{\beta}{N}}$ will appear in the search for an empirical upper bound of the local entropy term. Theorem 1.4 (page 11), used with the above mentioned local prior, shows that

$$\mathbb{P}\left\{\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \lambda \left\{ \rho(\Phi_{\frac{\lambda}{N}} \circ R) - \rho(r) \right\} - \mathcal{K}[\rho, \pi_{\exp(-\beta\Phi_{-\frac{\beta}{N}} \circ R)}] \right\} \leq 1. \quad (1.8)$$

Moreover

$$\begin{aligned} \mathcal{K}[\rho, \pi_{\exp[-\beta\Phi_{-\frac{\beta}{N}} \circ R]}] &= \mathcal{K}[\rho, \pi_{\exp(-\beta r)}] + \beta \rho \left[\Phi_{-\frac{\beta}{N}} \circ R - r \right] \\ &\quad + \log \left\{ \pi \left[\exp(-\beta\Phi_{-\frac{\beta}{N}} \circ R) \right] \right\} - \log \left\{ \pi \left[\exp(-\beta r) \right] \right\}, \end{aligned} \quad (1.9)$$

which is an invitation to find an upper bound for $\log \left\{ \pi \left[\exp[-\beta\Phi_{-\frac{\lambda}{N}} \circ R] \right] \right\} - \log \left\{ \pi \left[\exp(-\beta r) \right] \right\}$. Let us call for short $\bar{\pi}$ our localized prior distribution, thus defined as

$$\frac{d\bar{\pi}}{d\pi}(\theta) = \frac{\exp \left\{ -\beta\Phi_{-\frac{\beta}{N}}[R(\theta)] \right\}}{\pi \left\{ \exp[-\beta\Phi_{-\frac{\beta}{N}} \circ R] \right\}}.$$

Applying once again Theorem 1.4 (page 11), but this time to $-\beta$, we see that

$$\begin{aligned} &\mathbb{P}\left\{\exp \left[\log \left\{ \pi \left[\exp(-\beta\Phi_{-\frac{\beta}{N}} \circ R) \right] \right\} - \log \left\{ \pi \left[\exp(-\beta r) \right] \right\} \right] \right\} \\ &= \mathbb{P}\left\{\exp \left[\log \left\{ \pi \left[\exp(-\beta\Phi_{-\frac{\beta}{N}} \circ R) \right] \right\} + \inf_{\rho \in \mathcal{M}_+^1(\Theta)} \beta \rho(r) + \mathcal{K}(\rho, \pi) \right] \right\} \end{aligned}$$

$$\begin{aligned}
&\leq \mathbb{P} \left\{ \exp \left[\log \left\{ \pi \left[\exp(-\beta \Phi_{-\frac{\beta}{N}} \circ R) \right] \right\} + \beta \bar{\pi}(r) + \mathcal{K}(\bar{\pi}, \pi) \right] \right\} \\
&= \mathbb{P} \left\{ \exp \left[\beta \left[\bar{\pi}(r) - \bar{\pi}(\Phi_{-\frac{\beta}{N}} \circ R) \right] - \mathcal{K}(\bar{\pi}, \bar{\pi}) \right] \right\} \leq 1. \quad (1.10)
\end{aligned}$$

Combining equations (1.9) and (1.10) and using the concavity of $\Phi_{-\frac{\beta}{N}}$, we see that with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,

$$0 \leq \mathcal{K}(\rho, \bar{\pi}) \leq \mathcal{K}[\rho, \pi_{\exp(-\beta r)}] + \beta \left[\Phi_{-\frac{\beta}{N}}[\rho(R)] - \rho(r) \right] - \log(\epsilon).$$

We have proved a lower deviation bound:

THEOREM 1.17 *For any positive real constant β , with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\frac{\exp \left\{ \frac{\beta}{N} \left[\rho(r) - \frac{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}] - \log(\epsilon)}{\beta} \right] \right\} - 1}{\exp\left(\frac{\beta}{N}\right) - 1} \leq \rho(R).$$

Let us now seek for an upper bound. Using the Cauchy-Schwarz inequality to combine equations (1.8) and (1.10), we obtain

$$\begin{aligned}
&\mathbb{P} \left\{ \exp \left[\frac{1}{2} \sup_{\rho \in \mathcal{M}_+^1(\Theta)} \lambda \rho(\Phi_{\frac{\lambda}{N}} \circ R) - \beta \rho(\Phi_{-\frac{\beta}{N}} \circ R) - (\lambda - \beta) \rho(r) - \mathcal{K}[\rho, \pi_{\exp(-\beta r)}] \right] \right\} \\
&= \mathbb{P} \left\{ \exp \left[\frac{1}{2} \sup_{\rho \in \mathcal{M}_+^1(\Theta)} \left(\lambda \left\{ \rho(\Phi_{\frac{\lambda}{N}} \circ R) - \rho(r) \right\} - \mathcal{K}(\rho, \bar{\pi}) \right) \right] \right\} \\
&\quad \times \exp \left[\frac{1}{2} \left(\log \left\{ \pi \left[\exp(-\beta \Phi_{-\frac{\beta}{N}} \circ R) \right] \right\} - \log \left\{ \pi \left[\exp(-\beta r) \right] \right\} \right) \right] \\
&\leq \mathbb{P} \left\{ \exp \left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \left(\lambda \left\{ \rho(\Phi_{\frac{\lambda}{N}} \circ R) - \rho(r) \right\} - \mathcal{K}(\rho, \bar{\pi}) \right) \right] \right\}^{1/2} \\
&\quad \times \mathbb{P} \left\{ \exp \left[\left(\log \left\{ \pi \left[\exp(-\beta \Phi_{-\frac{\beta}{N}} \circ R) \right] \right\} - \log \left\{ \pi \left[\exp(-\beta r) \right] \right\} \right) \right] \right\}^{1/2} \leq 1. \quad (1.11)
\end{aligned}$$

Thus with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution ρ ,

$$\lambda \Phi_{\frac{\lambda}{N}}[\rho(R)] - \beta \Phi_{-\frac{\beta}{N}}[\rho(R)] \leq (\lambda - \beta) \rho(r) + \mathcal{K}(\rho, \pi_{\exp(-\beta r)}) - 2 \log(\epsilon).$$

(It would have been more straightforward to use a union bound on deviation inequalities instead of the Cauchy-Schwarz inequality on exponential moments, anyhow, this would have led to replace $-2\log(\epsilon)$ with the worse factor $2\log(\frac{2}{\epsilon})$.) Let us now remind that

$$\begin{aligned} \lambda\Phi_{\frac{\lambda}{N}}(p) - \beta\Phi_{-\frac{\beta}{N}}(p) &= -N\log\left\{1 - [1 - \exp(-\frac{\lambda}{N})]p\right\} \\ &\quad - N\log\left\{1 + [\exp(\frac{\beta}{N}) - 1]p\right\}, \end{aligned}$$

and let us put

$$\begin{aligned} B &= (\lambda - \beta)\rho(r) + \mathcal{K}[\rho, \pi_{\exp(-\beta r)}] - 2\log(\epsilon) \\ &= \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}] + \int_{\beta}^{\lambda} \pi_{\exp(-\xi r)}(r)d\xi - 2\log(\epsilon). \end{aligned}$$

Let us consider moreover the change of variables $\alpha = 1 - \exp(-\frac{\lambda}{N})$ and $\gamma = \exp(\frac{\beta}{N}) - 1$.

We obtain $[1 - \alpha\rho(R)][1 + \gamma\rho(R)] \geq \exp(-\frac{B}{N})$, leading to

THEOREM 1.18. *For any positive constants α, γ , such that $0 \leq \gamma < \alpha < 1$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, the bound*

$$\begin{aligned} M(\rho) &= -\frac{\log[(1 - \alpha)(1 + \gamma)]}{\alpha - \gamma}\rho(r) + \frac{\mathcal{K}(\rho, \pi_{\exp[-N\log(1+\gamma)r]}) - 2\log(\epsilon)}{N(\alpha - \gamma)} \\ &= \frac{\mathcal{K}[\rho, \pi_{\exp[N\log(1-\alpha)r]}] + \int_{N\log(1+\gamma)}^{-N\log(1-\alpha)} \pi_{\exp(-\xi r)}(r)d\xi - 2\log(\epsilon)}{N(\alpha - \gamma)}, \end{aligned}$$

is such that

$$\rho(R) \leq \frac{\alpha - \gamma}{2\alpha\gamma} \left(\sqrt{1 + \frac{4\alpha\gamma}{(\alpha - \gamma)^2} \{1 - \exp[-(\alpha - \gamma)M(\rho)]\}} - 1 \right) \leq M(\rho),$$

Using the *empirical dimension* d_e defined by equation (1.5) on page 21, we can use the inequality

$$\int_{\beta}^{\lambda} \pi_{\exp(-\xi r)}(r)d\xi \leq (\lambda - \beta) \operatorname{ess\,inf}_{\pi} r + d_e \log\left(\frac{\lambda}{\beta}\right),$$

to prove that

$$M(\rho) \leq \frac{\log[(1+\gamma)(1-\alpha)]}{\gamma-\alpha} \operatorname{ess\,inf}_{\pi} r + \frac{d_e \log\left[\frac{-\log(1-\alpha)}{\log(1+\gamma)}\right] + \mathcal{K}[\rho, \pi_{\exp[N \log(1-\alpha)r]}] - 2 \log(\epsilon)}{N(\alpha-\gamma)}.$$

Let us give a little numerical illustration : assuming that $d_e = 10$ and $N = 1000$, taking $\epsilon = 0.01$, $\alpha = 0.5$ and $\gamma = 0.1$, we obtain from Theorem 1.18 $\pi_{\exp[N \log(1-\alpha)r]}(R) \simeq \pi_{\exp(-693r)}(R) \leq 0.332 \leq 0.372$, where we have given respectively the non linear and the linear bound. This shows the practical interest of keeping the non-linearity. Let us also mention that optimizing the values of the parameters α and γ would not have yielded a significantly lower bound.

The following corollary is obtained by taking $\lambda = 2\beta$ and keeping only the linear bound, we give it for the sake of its simplicity:

COROLLARY 1.19. *For any positive real constant β such that $\exp(\frac{\beta}{N}) + \exp(-\frac{2\beta}{N}) < 2$, which is the case when $\beta < 0.48N$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned} \rho(R) &\leq \frac{\beta \rho(r) + \mathcal{K}[\rho, \pi_{\exp(-\beta r)}] - 2 \log(\epsilon)}{N[2 - \exp(\frac{\beta}{N}) - \exp(-\frac{2\beta}{N})]} \\ &= \frac{\int_{\beta}^{2\beta} \pi_{\exp(-\xi r)}(r) d\xi + \mathcal{K}[\rho, \pi_{\exp(-2\beta r)}] - 2 \log(\epsilon)}{N[2 - \exp(\frac{\beta}{N}) - \exp(-\frac{2\beta}{N})]}. \end{aligned}$$

Let us mention that this corollary applied to the above numerical example gives $\pi_{\exp(-200r)}(R) \leq 0.475$ (when we take $\beta = 100$, consistently with the choice $\gamma = 0.1$).

1.3.5. Partially local bounds. Local bounds are suitable when the lowest values of the empirical error rate r are reached only on a small part of the parameter set Θ . When Θ is the disjoint union of submodels of different complexities, the minimum of r will as a rule not be “localized” in a way that calls for the use of local bounds. Just think for instance of the case when $\Theta = \bigsqcup_{m=1}^M \Theta_m$, where the sets $\Theta_1 \subset \Theta_2 \subset \dots \subset \Theta_M$ are nested. In this case we will have $\inf_{\Theta_1} r \geq \inf_{\Theta_2} r \geq \dots \geq \inf_{\Theta_M} r$, although Θ_M may be too large to be the right model to use. In this situation, we do not want to localize the bound completely. Let us make a more specific fancyful

but typical pseudo computation. Just imagine we have a countable collection $(\Theta_m)_{m \in M}$ of submodels. Let us assume we are interested in choosing between the estimators $\hat{\theta}_m \in \arg \min_{\Theta_m} r$, maybe randomizing them (e.g. replacing them with $\pi_{\exp(-\lambda r)}^m$). Let us imagine moreover that we are in a typically parametric situation, where, for some priors $\pi^m \in \mathcal{M}_+^1(\Theta_m)$, $m \in M$, there is a “dimension” d_m such that $\lambda[\pi_{\exp(-\lambda r)}^m(r) - r(\hat{\theta}_m)] \simeq d_m$. Let $\mu \in \mathcal{M}_+^1(M)$ be some distribution on the index set M . It is easy to see that $(\mu\pi)_{\exp(-\lambda r)}$ will typically not be properly local, in the sense that typically

$$\begin{aligned} (\mu\pi)_{\exp(-\lambda r)}(r) &= \frac{\mu\left\{\pi_{\exp(-\lambda r)}(r)\pi[\exp(-\lambda r)]\right\}}{\mu\left\{\pi[\exp(-\lambda r)]\right\}} \\ &\simeq \frac{\sum_{m \in M} \left[(\inf_{\Theta_m} r) + \frac{d_m}{\lambda} \right] \exp\left[-\lambda(\inf_{\Theta_m} r) - d_m \log\left(\frac{e\lambda}{d_m}\right)\right] \mu(m)}{\sum_{m \in M} \exp\left[-\lambda(\inf_{\Theta_m} r) - d_m \log\left(\frac{e\lambda}{d_m}\right)\right] \mu(m)} \\ &\simeq \left\{ \inf_{m \in M} \left(\inf_{\Theta_m} r \right) + \frac{d_m}{\lambda} \log\left(\frac{e\lambda}{d_m \mu(m)}\right) \right\} \\ &\quad + \log\left\{ \sum_{m \in M} \exp\left[-d_m \log\left(\frac{\lambda}{d_m}\right)\right] \mu(m) \right\}. \end{aligned}$$

where we have used the estimate

$$\begin{aligned} -\log\left\{\pi[\exp(-\lambda r)]\right\} &= \int_0^\lambda \pi_{\exp(-\beta r)}(r) d\beta \\ &\simeq \int_0^\lambda \left(\inf_{\Theta_m} r \right) + \left[\frac{d_m}{\beta} \wedge 1 \right] d\beta \simeq \lambda(\inf_{\Theta_m} r) + d_m \left[\log\left(\frac{\lambda}{d_m}\right) + 1 \right]. \end{aligned}$$

Our approximations have no pretention to be rigorous or very accurate, but they nevertheless give the best order of magnitude we can expect in typical situations, and show that this order of magnitude is not what we are looking for: mixing different models with the help of μ spoils the localization, introducing a multiplier $\log\left(\frac{\lambda}{d_m}\right)$ to the dimension d_m which is precisely what we would have got if we had not localized at all the bound. What we would really like to do in such situations is to use a *partially localized* posterior distribution, such as $\mu_{\exp(-\lambda r)}^{\hat{m}}$, where \hat{m} is an estimator of the best submodel to be used. While the most straightforward way to do this is to use a union bound on results obtained for each submodel Θ_m , we are going

here to show how to allow arbitrary posterior distributions on the index set (corresponding to a randomization of the choice of \widehat{m}).

Let us consider the framework we just mentioned: let the measurable parameter set (Θ, \mathcal{T}) be a disjoint union of measurable submodels, $\Theta = \bigsqcup_{m \in M} \Theta_m$. Let the index set (M, \mathcal{M}) be some measurable space (most of the time it will be a countable set). Let $\mu \in \mathcal{M}_+^1(M)$ be a prior probability distribution on (M, \mathcal{M}) . Let $\pi : M \rightarrow \mathcal{M}_+^1(\Theta)$ be a regular conditional probability measure such that $\pi(m, \Theta_m) = 1$, for any $m \in M$. Let $\mu\pi \in \mathcal{M}_+^1(M \times \Theta)$ be the product probability measure defined by $\mu\pi(h) = \int_{m \in M} \left(\int_{\theta \in \Theta} h(m, \theta) \pi(m, d\theta) \right) \mu(dm)$, for any bounded measurable function $h : M \times \Theta \rightarrow \mathbb{R}$. Let $\pi_{\exp(h)} \in \mathcal{M}_+(M \times \Theta)$ be the regular conditional probability measure defined by

$$\frac{d\pi_{\exp(h)}}{d\pi}(m, \theta) = \frac{\exp[h(\theta)]}{\pi[m, \exp(h)]},$$

where consistently with previous notations $\pi(m, h) = \int_{\Theta} h(m, \theta) \pi(m, d\theta)$ (we will also often use the less explicit notation $\pi(h)$). Let for short

$$U(\theta, \omega) = \lambda \Phi_{\frac{\lambda}{N}}[R(\theta)] - \beta \Phi_{-\frac{\beta}{N}}[R(\theta)] - (\lambda - \beta)r(\theta, \omega).$$

Integrating with respect to μ equation (1.11) on page 29, written in each submodel Θ_m using the prior distribution $\pi(m, \cdot)$, we see that

$$\begin{aligned} & \mathbb{P} \left\{ \exp \left[\sup_{\nu \in \mathcal{M}_+^1(M)} \sup_{\rho: M \rightarrow \mathcal{M}_+^1(\Theta)} \frac{1}{2} \left[(\nu\rho)(U) - \nu \{ \mathcal{K}([\rho, \pi_{\exp(-\beta r)}] \} \right] - \mathcal{K}(\nu, \mu) \right] \right\} \\ & \leq \mathbb{P} \left\{ \exp \left[\sup_{\nu \in \mathcal{M}_+^1(M)} \frac{1}{2} \nu \left(\sup_{\rho: M \rightarrow \mathcal{M}_+^1(\Theta)} \rho(U) - \mathcal{K}(\rho, \pi_{\exp(-\beta r)}) \right) - \mathcal{K}(\nu, \mu) \right] \right\} \\ & = \mathbb{P} \left\{ \mu \left[\exp \left\{ \frac{1}{2} \sup_{\rho: M \rightarrow \mathcal{M}_+^1(\Theta)} \left[\rho(U) - \mathcal{K}[\rho, \pi_{\exp(-\beta r)}] \right] \right\} \right] \right\} \\ & = \mu \left\{ \mathbb{P} \left[\exp \left\{ \frac{1}{2} \sup_{\rho: M \rightarrow \mathcal{M}_+^1(\Theta)} \left[\rho(U) - \mathcal{K}[\rho, \pi_{\exp(-\beta r)}] \right] \right\} \right] \right\} \leq 1. \end{aligned}$$

This proves that

$$\begin{aligned} & \mathbb{P} \left\{ \exp \left[\frac{1}{2} \sup_{\nu \in \mathcal{M}_+^1(M)} \sup_{\rho: M \rightarrow \mathcal{M}_+^1(\Theta)} \lambda \Phi_{\frac{\lambda}{N}}[\nu\rho(R)] - \beta \Phi_{-\frac{\beta}{N}}[\nu\rho(R)] \right. \right. \\ & \quad \left. \left. - (\lambda - \beta)\nu\rho(r) - 2\mathcal{K}(\nu, \mu) - \nu \{ \mathcal{K}[\rho, \pi_{\exp(-\beta r)}] \} \right] \right\} \leq 1. \quad (1.12) \end{aligned}$$

Introducing the optimal value of r on each submodel $r^*(m) = \text{ess inf}_{\pi(m, \cdot)} r$ and the empirical dimensions

$$d_e(m) = \sup_{\xi \in \mathbb{R}_+} \xi [\pi_{\exp(-\xi r)}(m, r) - r^*(m)],$$

we can thus state

THEOREM 1.20. *For any positive real constants $\beta < \lambda$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\nu : \Omega \rightarrow \mathcal{M}_+^1(M)$, for any conditional posterior distribution $\rho : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\lambda \Phi_{\frac{\lambda}{N}}[\nu \rho(R)] - \beta \Phi_{-\frac{\beta}{N}}[\nu \rho(R)] \leq B_1(\nu, \rho),$$

$$\begin{aligned} \text{where } B_1(\nu, \rho) &= (\lambda - \beta)\nu \rho(r) + 2\mathcal{K}(\nu, \mu) + \nu\{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]\} - 2\log(\epsilon) \\ &= \nu \left[\int_{\beta}^{\lambda} \pi_{\exp(-\alpha r)}(r) d\alpha \right] + 2\mathcal{K}(\nu, \mu) + \nu\{\mathcal{K}[\rho, \pi_{\exp(-\lambda r)}]\} - 2\log(\epsilon) \\ &= 2\log \left\{ \mu \left[\exp \left(-\frac{1}{2} \int_{\beta}^{\lambda} \pi_{\exp(-\alpha r)}(r) d\alpha \right) \right] \right\} \\ &\quad + 2\mathcal{K} \left[\nu, \mu \left(\frac{\pi[\exp(-\lambda r)]}{\pi[\exp(-\beta r)]} \right)^{1/2} \right] + \nu\{\mathcal{K}[\rho, \pi_{\exp(-\lambda r)}]\} - 2\log(\epsilon), \end{aligned}$$

$$\text{and therefore } B_1(\nu, \rho) \leq \nu \left[(\lambda - \beta)r^* + \log \left(\frac{\lambda}{\beta} \right) d_e \right] + 2\mathcal{K}(\nu, \mu) + \nu\{\mathcal{K}[\rho, \pi_{\exp(-\lambda r)}]\} - 2\log(\epsilon),$$

$$\begin{aligned} \text{as well as } B_1(\nu, \rho) &\leq 2\log \left\{ \mu \left[\exp \left(-\frac{1}{2} r^* + \frac{1}{2} \log \left(\frac{\lambda}{\beta} \right) d_e \right) \right] \right\} \\ &\quad + 2\mathcal{K} \left[\nu, \mu \frac{\pi[\exp(-\lambda r)]}{\pi[\exp(-\beta r)]} \right] + \nu\{\mathcal{K}[\rho, \pi_{\exp(-\lambda r)}]\} - 2\log(\epsilon). \end{aligned}$$

Thus, for any real constants α and γ such that $0 \leq \gamma < \alpha < 1$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\nu : \Omega \rightarrow \mathcal{M}_+^1(M)$ and any conditional posterior distribution $\rho : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$, the bound

$$\begin{aligned} B_2(\nu, \rho) &= -\frac{\log[(1-\alpha)(1+\gamma)]}{\alpha-\gamma} \nu \rho(r) + \frac{2\mathcal{K}(\nu, \mu) + \nu\{\mathcal{K}[\rho, \pi_{(1+\gamma)^{-N} r}]\} - 2\log(\epsilon)}{N(\alpha-\gamma)} \\ &= \frac{2\mathcal{K} \left[\nu, \mu \left(\frac{\pi[(1-\alpha)^{N r}]}{\pi[(1+\gamma)^{-N r}] } \right)^{1/2} \right] + \nu\{\mathcal{K}[\rho, \pi_{(1-\alpha)^{N r}}]\}}{N(\alpha-\gamma)} \\ &\quad - \frac{2\log \left\{ \mu \left[\exp \left[-\frac{1}{2} \int_{N \log(1+\gamma)}^{-N \log(1-\alpha)} \pi_{\exp(-\xi r)}(\cdot, r) d\xi \right] \right] \right\} + 2\log(\epsilon)}{N(\alpha-\gamma)} \end{aligned}$$

satisfies

$$\nu\rho(R) \leq \frac{\alpha - \gamma}{2\alpha\gamma} \left(\sqrt{1 + \frac{4\alpha\gamma}{(\alpha - \gamma)^2} \left\{ 1 - \exp[-(\alpha - \gamma)B(\nu, \rho)] \right\}} - 1 \right) \leq B(\nu, \rho).$$

Let us remark that in the case when $\nu = \mu \left(\frac{\pi[(1-\alpha)Nr]}{\pi[(1+\gamma)-Nr]} \right)^{1/2}$ and $\rho = \pi_{(1-\alpha)Nr}$, we get as desired a bound that is adaptively local in all the Θ_m (at least when M is countable and μ is atomic):

$$\begin{aligned} B(\nu, \rho) &\leq -\frac{2}{N(\alpha-\gamma)} \log \left\{ \mu \left\{ \exp \left[\frac{N}{2} \log[(1+\gamma)(1-\alpha)] r^* \right. \right. \right. \\ &\quad \left. \left. \left. - \log \left(\frac{-\log(1-\alpha)}{\log(1+\gamma)} \right) \frac{d_e}{2} \right] \right\} \right\} - \frac{2 \log(\epsilon)}{N(\alpha-\gamma)} \\ &\leq \inf_{m \in M} \left\{ -\frac{\log[(1-\alpha)(1+\gamma)]}{\alpha-\gamma} r^*(m) \right. \\ &\quad \left. + \log \left(\frac{-\log(1-\alpha)}{\log(1+\gamma)} \right) \frac{d_e(m)}{N(\alpha-\gamma)} - 2 \frac{\log[\epsilon\mu(m)]}{N(\alpha-\gamma)} \right\}. \end{aligned}$$

The penalization by the *empirical dimension* $d_e(m)$ in each submodel is as desired linear in $d_e(m)$. Non random partially local bounds could be obtained in a way that is easy to imagine. We leave this investigation to the reader.

1.3.6. Two step localization. We have seen that the bound optimal choice of the posterior distribution ν on the index set in Theorem 1.20 (page 34) is such that

$$\frac{d\nu}{d\mu}(m) \sim \left(\frac{\pi[\exp(-\lambda r(m, \cdot))]}{\pi[\exp(-\beta r(m, \cdot))]} \right)^{\frac{1}{2}} = \exp \left[-\frac{1}{2} \int_{\beta}^{\lambda} \pi_{\exp(-\alpha r)}(m, r) d\alpha \right].$$

This suggests to replace the prior distribution μ with $\bar{\mu}$ defined by its density

$$\begin{aligned} \frac{d\bar{\mu}}{d\mu}(m) &= \frac{\exp[-h(m)]}{\mu[\exp(-h)]}, \\ \text{where } h(m) &= -\xi \int_{\beta}^{\gamma} \pi_{\exp(-\alpha\Phi_{-\frac{\eta}{N}} \circ R)}[\Phi_{-\frac{\eta}{N}} \circ R(m, \cdot)] d\alpha. \quad (1.13) \end{aligned}$$

The use of $\Phi_{-\frac{\eta}{N}} \circ R$ instead of R is motivated by technical reasons which will appear in subsequent computations. Indeed, we will need to bound

$$\nu \left[\int_{\beta}^{\lambda} \pi_{\exp(-\alpha\Phi_{-\frac{\eta}{N}} \circ R)}(\Phi_{-\frac{\eta}{N}} \circ R) d\alpha \right]$$

in order to handle $\mathcal{K}(\nu, \bar{\mu})$. In the spirit of equation (1.8, page 28), starting back from Theorem 1.4 (page 11), applied in each submodel Θ_m to the prior distribution $\pi_{\exp(-\gamma\Phi_{-\frac{\eta}{N}} \circ R)}$ and integrated with respect to $\bar{\mu}$, we see that for any positive real constants λ , γ and η , with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\nu : \Omega \rightarrow \mathcal{M}_+^1(M)$ on the index set and any conditional posterior distribution $\rho : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$,

$$\begin{aligned} \nu\rho(\lambda\Phi_{\frac{\lambda}{N}} \circ R - \gamma\Phi_{-\frac{\eta}{N}} \circ R) &\leq \lambda\nu\rho(r) \\ &+ \nu\mathcal{K}(\rho, \pi) + \mathcal{K}(\nu, \bar{\mu}) + \nu\left\{\log\left[\pi\left[\exp(-\gamma\Phi_{-\frac{\eta}{N}} \circ R)\right]\right]\right\} - \log(\epsilon). \end{aligned} \quad (1.14)$$

Since $x \mapsto f(x) \stackrel{\text{def}}{=} \lambda\Phi_{\frac{\lambda}{N}} - \gamma\Phi_{-\frac{\eta}{N}}(x)$ is a convex function, it is such that

$$f(x) \geq xf'(0) = xN\left\{\left[1 - \exp(-\frac{\lambda}{N})\right] + \frac{\gamma}{\eta}\left[\exp(\frac{\eta}{N}) - 1\right]\right\}.$$

Thus if we put

$$\gamma = \frac{\eta\left[1 - \exp(-\frac{\lambda}{N})\right]}{\exp(\frac{\eta}{N}) - 1}, \quad (1.15)$$

we obtain that $f(x) \geq 0$, $x \in \mathbb{R}$, and therefore that the left-hand side of equation (1.14) is non negative. We can moreover introduce the prior conditional distribution $\bar{\pi}$ defined by

$$\frac{d\bar{\pi}}{d\pi}(m, \theta) = \frac{\exp[-\beta\Phi_{-\frac{\eta}{N}} \circ R(\theta)]}{\pi\{m, \exp[-\beta\Phi_{-\frac{\eta}{N}} \circ R]\}}.$$

With \mathbb{P} probability at least $1 - \epsilon$, for any posterior distributions $\nu\Omega \rightarrow \mathcal{M}_+^1(M)$ and $\rho : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$,

$$\begin{aligned} \beta\nu\rho(r) + \nu[\mathcal{K}(\rho, \pi)] &= \nu\{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]\} - \nu\left[\log\left\{\pi\left[\exp(-\beta r)\right]\right\}\right] \\ &\leq \nu\{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]\} + \beta\nu\bar{\pi}(r) + \nu[\mathcal{K}(\bar{\pi}, \pi)] \\ &\leq \nu\{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]\} + \beta\nu\bar{\pi}(\Phi_{-\frac{\eta}{N}} \circ R) \\ &\quad + \frac{\beta}{\eta}[\mathcal{K}(\nu, \bar{\mu}) - \log(\epsilon)] + \nu[\mathcal{K}(\bar{\pi}, \pi)] \\ &= \nu\{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]\} - \nu\left\{\log\left[\pi\left[\exp(-\beta\Phi_{-\frac{\eta}{N}} \circ R)\right]\right]\right\} \\ &\quad + \frac{\beta}{\eta}[\mathcal{K}(\nu, \bar{\mu}) - \log(\epsilon)]. \end{aligned}$$

Thus, coming back to equation (1.14), we see that under condition (1.15), with \mathbb{P} probability at least $1 - \epsilon$,

$$0 \leq (\lambda - \beta)\nu\rho(r) + \nu\{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]\} \\ - \nu \left[\int_{\beta}^{\gamma} \pi_{\exp(-\alpha\Phi - \frac{\eta}{N} \circ R)}(\Phi - \frac{\eta}{N} \circ R) d\alpha \right] + (1 + \frac{\beta}{\eta})[\mathcal{K}(\nu, \bar{\mu}) + \log(\frac{2}{\epsilon})].$$

Noticing moreover that

$$(\lambda - \beta)\nu\rho(r) + \nu\{\mathcal{K}[\rho, \pi_{\exp(-\beta r)}]\} \\ = \nu\{\mathcal{K}[\rho, \pi_{\exp(-\lambda r)}]\} + \nu \left[\int_{\beta}^{\lambda} \pi_{\exp(-\alpha r)}(r) d\alpha \right],$$

and choosing $\rho = \pi_{\exp(-\lambda r)}$, we have proved

THEOREM 1.21 *For any positive real constants β , γ and η , such that $\gamma < \eta[\exp(\frac{\eta}{N}) - 1]^{-1}$, defining λ by condition (1.15), so that $\lambda = -N \log\{1 - \frac{\gamma}{\eta}[\exp(\frac{\eta}{N}) - 1]\}$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\nu : \Omega \rightarrow \mathcal{M}_+^1(M)$, any conditional posterior distribution $\rho : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\nu \left[\int_{\beta}^{\gamma} \pi_{\exp(-\alpha\Phi - \frac{\eta}{N} \circ R)}(\Phi - \frac{\eta}{N} \circ R) d\alpha \right] \\ \leq \nu \left[\int_{\beta}^{\lambda} \pi_{\exp(-\alpha r)}(r) d\alpha \right] + (1 + \frac{\beta}{\eta})[\mathcal{K}(\nu, \bar{\mu}) + \log(\frac{2}{\epsilon})].$$

Let us remark that this theorem does not require that $\beta < \gamma$, and thus provides both an upper and a lower bound for the quantity of interest:

COROLLARY 1.22 *For any positive real constants β , γ and η such that $\max\{\beta, \gamma\} < \eta[\exp(\frac{\eta}{N}) - 1]^{-1}$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distributions $\nu : \Omega \rightarrow \mathcal{M}_+^1(M)$ and $\rho : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\nu \left[\int_{-N \log\{1 - \frac{\beta}{\eta}[\exp(\frac{\eta}{N}) - 1]\}}^{\gamma} \pi_{\exp(-\alpha r)}(r) d\alpha \right] - (1 + \frac{\gamma}{\eta})[\mathcal{K}(\nu, \bar{\mu}) + \log(\frac{3}{\epsilon})] \\ \leq \nu \left[\int_{\beta}^{\gamma} \pi_{\exp(-\alpha\Phi - \frac{\eta}{N} \circ R)}(\Phi - \frac{\eta}{N} \circ R) d\alpha \right] \\ \leq \nu \left[\int_{\beta}^{-N \log\{1 - \frac{\gamma}{\eta}[\exp(\frac{\eta}{N}) - 1]\}} \pi_{\exp(-\alpha r)}(r) d\alpha \right] \\ + (1 + \frac{\beta}{\eta})[\mathcal{K}(\nu, \bar{\mu}) + \log(\frac{3}{\epsilon})].$$

We can then remember that

$$\mathcal{K}(\nu, \bar{\mu}) = \xi(\nu - \bar{\mu}) \left[\int_{\beta}^{\gamma} \pi_{\exp(-\alpha\Phi_{-\frac{\eta}{N}} \circ R)}(\Phi_{-\frac{\eta}{N}} \circ R) d\alpha \right] + \mathcal{K}(\nu, \mu) - \mathcal{K}(\bar{\mu}, \mu),$$

to conclude that, putting

$$G_{\eta}(\alpha) = -N \log \left\{ 1 - \frac{\alpha}{\eta} \left[\exp\left(\frac{\eta}{N}\right) - 1 \right] \right\} \geq \alpha, \quad \alpha \in \mathbb{R}_+, \quad (1.16)$$

and

$$\frac{d\hat{\nu}}{d\mu}(m) \stackrel{\text{def}}{=} \frac{\exp[-h(m)]}{\mu[\exp(-h)]} \text{ where } h(m) = \xi \int_{G_{\eta}(\beta)}^{\gamma} \pi_{\exp(-\alpha r)}(m, r) d\alpha, \quad (1.17)$$

the divergence of ν with respect to the local prior $\bar{\mu}$ is bounded by

$$\begin{aligned} & \left[1 - \xi \left(1 + \frac{\beta}{\eta} \right) \right] \mathcal{K}(\nu, \bar{\mu}) \\ & \leq \xi \nu \left[\int_{\beta}^{G_{\eta}(\gamma)} \pi_{\exp(-\alpha r)}(r) d\alpha \right] - \xi \bar{\mu} \left[\int_{G_{\eta}(\beta)}^{\gamma} \pi_{\exp(-\alpha r)}(r) d\alpha \right] \\ & \quad + \mathcal{K}(\nu, \mu) - \mathcal{K}(\bar{\mu}, \mu) + \xi \left(2 + \frac{\beta + \gamma}{\eta} \right) \log\left(\frac{3}{\epsilon}\right) \\ & \leq \xi \nu \left[\int_{\beta}^{G_{\eta}(\gamma)} \pi_{\exp(-\alpha r)}(r) d\alpha \right] + \mathcal{K}(\nu, \mu) \\ & \quad + \log \left\{ \mu \left[\exp \left(-\xi \int_{G_{\eta}(\beta)}^{\gamma} \pi_{\exp(-\alpha r)}(r) d\alpha \right) \right] \right\} \\ & \quad + \xi \left(2 + \frac{\beta + \gamma}{\eta} \right) \log\left(\frac{3}{\epsilon}\right) \\ & = \mathcal{K}(\nu, \hat{\nu}) + \xi \nu \left[\left(\int_{\beta}^{G_{\eta}(\beta)} + \int_{\gamma}^{G_{\eta}(\gamma)} \right) \pi_{\exp(-\alpha r)}(r) d\alpha \right] \\ & \quad + \xi \left(2 + \frac{\beta + \gamma}{\eta} \right) \log\left(\frac{3}{\epsilon}\right). \end{aligned}$$

We have proved

THEOREM 1.23. *For any positive constants β , γ and η such that $\max\{\beta, \gamma\} < \eta \left[\exp\left(\frac{\eta}{N}\right) - 1 \right]^{-1}$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\nu : \Omega \rightarrow \mathcal{M}_+^1(M)$ and any conditional posterior distribution $\rho : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\mathcal{K}(\nu, \bar{\mu}) \leq \left[1 - \xi \left(1 + \frac{\beta}{\eta} \right) \right]^{-1} \left\{ \mathcal{K}(\nu, \hat{\nu}) \right.$$

$$\begin{aligned}
& + \xi \nu \left[\left(\int_{\beta}^{G_{\eta}(\beta)} + \int_{\gamma}^{G_{\eta}(\gamma)} \right) \pi_{\exp(-\alpha r)}(r) d\alpha \right] \\
& \qquad \qquad \qquad + \xi \left(2 + \frac{\beta + \gamma}{\eta} \right) \log\left(\frac{3}{\epsilon}\right) \Big\} \\
\leq & \left[1 - \xi \left(1 + \frac{\beta}{\eta} \right) \right]^{-1} \left\{ \mathcal{K}(\nu, \widehat{\nu}) \right. \\
& + \xi \nu \left[[G_{\eta}(\gamma) - \gamma + G_{\eta}(\beta) - \beta] r^* + \log\left(\frac{G_{\eta}(\beta)G_{\eta}(\gamma)}{\beta\gamma}\right) d_e \right] \\
& \qquad \qquad \qquad \left. + \xi \left(2 + \frac{\beta + \gamma}{\eta} \right) \log\left(\frac{3}{\epsilon}\right) \right\},
\end{aligned}$$

where the local prior $\bar{\mu}$ is defined by equation (1.13) on page 35 and the local posterior $\widehat{\nu}$ and the function G_{η} are defined by equation (1.17) above.

We can then use this theorem to give a local version of Theorem 1.20 (page 34). To get something pleasing to read, we can apply Theorem 1.23 with constants β' , γ' and η chosen so that $\frac{2\xi}{1-\xi(1+\frac{\beta'}{\eta})} = 1$, $G_{\eta}(\beta') = \beta$ and $\gamma' = \lambda$, where β and λ are the constants appearing in Theorem 1.20. This gives

THEOREM 1.24. *For any positive real constants $\beta < \lambda$ and η such that $\lambda < \eta[\exp(\frac{\lambda}{\eta}) - 1]^{-1}$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\nu : \Omega \rightarrow \mathcal{M}_+^1(M)$, for any conditional posterior distribution $\rho : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned}
& \lambda \Phi_{\frac{\lambda}{\eta}}[\nu\rho(R)] - \beta \Phi_{-\frac{\beta}{\eta}}[\nu\rho(R)] \leq B_3(\nu, \rho), \text{ where} \\
B_3(\nu, \rho) = & \nu \left[\int_{G_{\eta}^{-1}(\beta)}^{G_{\eta}(\lambda)} \pi_{\exp(-\alpha r)}(r) d\alpha \right] \\
& + \left(3 + \frac{G_{\eta}^{-1}(\beta)}{\eta} \right) \mathcal{K} \left[\nu, \mu_{\exp\left[-\left(3 + \frac{G_{\eta}^{-1}(\beta)}{\eta}\right)^{-1} \int_{\beta}^{\lambda} \pi_{\exp(-\alpha r)}(r) d\alpha\right]} \right] \\
& \qquad \qquad \qquad + \nu \{ \mathcal{K}(\rho, \pi_{\exp(-\lambda r)}) \} + \left(4 + \frac{G_{\eta}^{-1}(\beta) + \lambda}{\eta} \right) \log\left(\frac{4}{\epsilon}\right) \\
\leq & \nu \left[[G_{\eta}(\lambda) - G_{\eta}^{-1}(\beta)] r^* + \log\left(\frac{G_{\eta}(\lambda)}{G_{\eta}^{-1}(\beta)}\right) d_e \right] \\
& + \left(3 + \frac{G_{\eta}^{-1}(\beta)}{\eta} \right) \mathcal{K} \left[\nu, \mu_{\exp\left[-\left(3 + \frac{G_{\eta}^{-1}(\beta)}{\eta}\right)^{-1} \int_{\beta}^{\lambda} \pi_{\exp(-\alpha r)}(r) d\alpha\right]} \right] \\
& \qquad \qquad \qquad + \nu \{ \mathcal{K}(\rho, \pi_{\exp(-\lambda r)}) \} + \left(4 + \frac{G_{\eta}^{-1}(\beta) + \lambda}{\eta} \right) \log\left(\frac{4}{\epsilon}\right),
\end{aligned}$$

and where the function G_{η} is defined by equation (1.16) on page 38.

A first remark: if we had the stamina to use Cauchy Schwarz inequalities (or more generally Hölder inequalities) on exponential moments instead of using weighted union bounds on deviation inequalities, we could have replaced $\log(\frac{4}{\epsilon})$ with $-\log(\epsilon)$ in the above inequalities.

We see that we have achieved the desired kind of localization of Theorem 1.20 (page 34), since the new empirical entropy term

$$\mathcal{K}[\nu, \mu_{\exp[-\xi \int_{\beta}^{\lambda} \pi_{\exp(-\alpha r)}(r) d\alpha]]$$

cancels for a value of the posterior distribution on the index set ν which is of the same form as the one minimizing the bound $B_1(\nu, \rho)$ of Theorem 1.20 (with a decreased constant, as could be expected). In a typical parametric setting, we will have

$$\int_{\beta}^{\lambda} \pi_{\exp(-\alpha r)}(r) d\alpha \simeq (\lambda - \beta) r^*(m) + \log\left(\frac{\lambda}{\beta}\right) d_e(m),$$

and therefore, if we choose for ν the Dirac mass at

$$\hat{m} \in \arg \min_{m \in M} r^*(m) + \frac{\log(\frac{\lambda}{\beta})}{\lambda - \beta} d_e(m),$$

and $\rho(m, \cdot) = \pi_{\exp(-\lambda r)}(m, \cdot)$, we will get, in the case when the index set M is countable,

$$\begin{aligned} B_3(\nu, \rho) &\lesssim \max \left\{ \left[G_{\eta}(\lambda) - G_{\eta}^{-1}(\beta) \right], (\lambda - \beta) \frac{\log\left[\frac{G_{\eta}(\lambda)}{G_{\eta}^{-1}(\beta)}\right]}{\log\left(\frac{\lambda}{\beta}\right)} \right\} \\ &\quad \times \left[r^*(\hat{m}) + \frac{\log(\frac{\lambda}{\beta})}{\lambda - \beta} d_e(\hat{m}) \right] \\ &+ \left(3 + \frac{G_{\eta}^{-1}(\beta)}{\eta} \right) \log \left\{ \sum_{m \in M} \frac{\mu(m)}{\mu(\hat{m})} \exp \left[- \left(3 + \frac{G_{\eta}^{-1}(\beta)}{\eta} \right)^{-1} \right. \right. \\ &\quad \left. \left. \times \left\{ (\lambda - \beta) [r^*(m) - r^*(\hat{m})] + \log\left(\frac{\lambda}{\beta}\right) [d_e(m) - d_e(\hat{m})] \right\} \right] \right\} \\ &\quad + \left(4 + \frac{G_{\eta}^{-1}(\beta) + \lambda}{\eta} \right) \log\left(\frac{4}{\epsilon}\right). \end{aligned}$$

Therefore, as long as there are not too many of them, we do not feel strongly in this bound the models for which the penalized minimum empirical risk $r^*(m) + \frac{\log(\frac{\lambda}{\beta})}{\lambda - \beta} d_e(m)$ is far from optimal.

1.4. RELATIVE BOUNDS. The behaviour of the minimum of the empirical process $\theta \mapsto r(\theta)$ is known to depend on the covariances between pairs

$[r(\theta), r(\theta')]$, $\theta, \theta' \in \Theta$. Accordingly, our previous study, based on the analysis of the variance of $r(\theta)$ (or technically on some exponential moment playing quite the same role), is missing some accuracy in some circumstances (namely when $\inf_{\Theta} R$ is not close enough to zero). In this subsection, instead of bounding the expected risk $\rho(R)$, we are going to upper bound the difference $\rho(R) - \inf_{\Theta} R$, and more generally $\rho(R) - R(\tilde{\theta})$, where $\tilde{\theta} \in \Theta$ is some fixed parameter value. Eventually in the next subsection we will analyze $\rho(R) - \pi_{\exp(-\beta R)}(R)$, allowing to compare the expected error rate of a posterior distribution ρ with the error rate of a Gibbs prior distribution. Thus relative bounds are not exactly of the same nature as previous ones: although it is not possible to estimate $\rho(R)$ with an order of precision higher than $(\rho(R)/N)^{1/2}$, it is still possible in some situations to reach a better precision for $\rho(R) - \inf_{\Theta} R$, as we will see. The study of PAC-Bayesian relative bounds stems from the second and third part of J. Y. Audibert's dissertation [3].

We will suggest two different kinds of applications of these bounds. The first more obvious one is to upper bound $\rho(R) - \inf_{\Theta} R$ to get an idea of the performance of the posterior distribution ρ .

The second application is to compare the classification model indexed by Θ with a submodel indexed by one of its measurable subsets $\Theta_1 \subset \Theta$. For this purpose we are going to compare $\rho(R)$, where $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$ is any posterior distribution, with $R(\tilde{\theta})$, where $\tilde{\theta} \in \Theta_1$ is some possibly unobservable value of the parameter in the submodel defined by Θ_1 . We will typically consider the case when $\tilde{\theta} \in \arg \min_{\Theta_1} R$. In this special case, a negative bound for $\rho(R) - R(\tilde{\theta}) = \rho(R) - \inf_{\Theta_1} R$ indicates that it is definitely worth using a randomized estimator ρ supported by the larger parameter set Θ instead of using only the classification model defined by the smaller set Θ_1 .

1.4.1. Basic inequalities. Relative bounds in this section are based on the control of $r(\theta) - r(\tilde{\theta})$, where $\theta, \tilde{\theta} \in \Theta$. These differences are related to the random variables

$$\psi_i(\theta, \tilde{\theta}) = \sigma_i(\theta) - \sigma_i(\tilde{\theta}) = \mathbb{1}[f_{\theta}(X_i) \neq Y_i] - \mathbb{1}[f_{\tilde{\theta}}(X_i) \neq Y_i].$$

Some supplementary technical difficulties, as compared to the previous sections, come from the fact that $\psi_i(\theta, \tilde{\theta})$ takes three values, whereas $\sigma_i(\theta)$ takes only two. Let $r'(\theta, \tilde{\theta}) = r(\theta) - r(\tilde{\theta})$ and $R'(\theta, \tilde{\theta}) = R(\theta) - R(\tilde{\theta})$. We have as usual from independence that

$$\begin{aligned} \log \left\{ \mathbb{P} \left[\exp[-\lambda r'(\theta, \tilde{\theta})] \right] \right\} &= \sum_{i=1}^N \log \left\{ \mathbb{P} \left[\exp \left[-\frac{\lambda}{N} \psi_i(\theta, \tilde{\theta}) \right] \right] \right\} \\ &\leq N \log \left\{ \frac{1}{N} \sum_{i=1}^N \mathbb{P} \left\{ \exp \left[-\frac{\lambda}{N} \psi_i(\theta, \tilde{\theta}) \right] \right\} \right\}. \end{aligned}$$

Let C_i be the distribution of $\psi_i(\theta, \tilde{\theta})$ under \mathbb{P} and let $\bar{C} = \frac{1}{N} \sum_{i=1}^N C_i \in \mathcal{M}_+^1(\{-1, 0, 1\})$. With these notations

$$\log \left\{ \mathbb{P} \left[\exp[-\lambda r'(\theta, \tilde{\theta})] \right] \right\} \leq N \log \left\{ \int \exp \left(-\frac{\lambda}{N} \psi \right) \bar{C}(d\psi) \right\}. \quad (1.18)$$

The right-hand side of this inequality is a function of \bar{C} . On the other hand, \bar{C} being a probability measure on a three point set, is defined by two parameters, that we may take equal to $\int \psi \bar{C}(d\psi)$ and $\int \psi^2 \bar{C}(d\psi)$. To this purpose, let us introduce

$$M'(\theta, \tilde{\theta}) = \int \psi^2 \bar{C}(d\psi) = \bar{C}(+1) + \bar{C}(-1) = \frac{1}{N} \sum_{i=1}^N \mathbb{P}[\psi_i^2(\theta, \tilde{\theta})], \quad \theta, \tilde{\theta} \in \Theta.$$

It is a pseudo distance (meaning that it is symmetric and satisfies the triangle inequality), since it can also be written as

$$M'(\theta, \tilde{\theta}) = \frac{1}{N} \sum_{i=1}^N \mathbb{P} \left\{ \left| \mathbb{1}[f_\theta(X_i) \neq Y_i] - \mathbb{1}[f_{\tilde{\theta}}(X_i) \neq Y_i] \right| \right\}, \quad \theta, \tilde{\theta} \in \Theta.$$

It is readily seen that

$$N \log \left\{ \int \exp \left(-\frac{\lambda}{N} \psi \right) \bar{C}(d\psi) \right\} = -\lambda \Psi_{\frac{\lambda}{N}} [R'(\theta, \tilde{\theta}), M'(\theta, \tilde{\theta})],$$

where

$$\begin{aligned} \Psi_a(p, m) &= -a^{-1} \log \left[(1-m) + \frac{m+p}{2} \exp(-a) + \frac{m-p}{2} \exp(a) \right] \\ &= -a^{-1} \log \left\{ 1 - \sinh(a) \left[p - m \tanh\left(\frac{a}{2}\right) \right] \right\}. \end{aligned}$$

Thus plugging this equality into inequality (1.18) we see that for any real parameter λ ,

$$\log \left\{ \mathbb{P} \left[\exp[-\lambda r'(\theta, \tilde{\theta})] \right] \right\} \leq -\lambda \Psi_{\frac{\lambda}{N}} [R'(\theta, \tilde{\theta}), M'(\theta, \tilde{\theta})],$$

To make a link with previous works initiated by Mammen and Tsybakov (see e.g. [28, 34]), we may consider the pseudo distance D on Θ defined on page 14 by equation (1.2). This distance only depends on the distribution of the patterns. It is often used to formulate margin assumptions (in the sense of Mammen and Tsybakov). Here we are going to work rather with M' : as it is dominated by D in the sense that $M'(\theta, \tilde{\theta}) \leq D(\theta, \tilde{\theta})$, $\theta, \tilde{\theta} \in \Theta$, with equality in the important case of binary classification, hypotheses formulated on D induce hypotheses on M' , and working with M' may only sharpen the results when compared to working with D .

Using the same reasoning as in the previous section, we deduce

THEOREM 1.25. *For any real parameter λ , any $\tilde{\theta} \in \Theta$,*

$$\mathbb{P} \left\{ \exp \left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \lambda \left[\rho \left\{ \Psi_{\frac{\lambda}{N}} [R'(\cdot, \tilde{\theta}), M'(\cdot, \tilde{\theta})] \right\} - \rho[r'(\cdot, \tilde{\theta})] \right] - \mathcal{K}(\rho, \pi) \right] \right\} \leq 1.$$

We are now going to derive some variant of Theorem 1.25. In this theorem, we obtain an inequality comparing one observed quantity $\rho[r'(\cdot, \tilde{\theta})]$ with two unobserved ones, $\rho[R'(\cdot, \tilde{\theta})]$ and $\rho[M'(\cdot, \tilde{\theta})]$ (because of the convexity of the function $\lambda \Psi_{\frac{\lambda}{N}}$,

$$\lambda \rho \left\{ \Psi_{\frac{\lambda}{N}} [R'(\cdot, \tilde{\theta}), M'(\cdot, \tilde{\theta})] \right\} \geq \lambda \Psi_{\frac{\lambda}{N}} \left\{ \rho [R'(\cdot, \tilde{\theta})], \rho [M'(\cdot, \tilde{\theta})] \right\}.$$

This may be inconvenient when looking for an empirical bound for $\rho[R'(\cdot, \tilde{\theta})]$, and we are going now to seek an inequality comparing $\rho[R'(\cdot, \tilde{\theta})]$ with empirical quantities only. This is possible through a change of variables in the exponential inequality. Indeed, if we consider now random variables $\chi_i(\theta, \tilde{\theta})$, such that

$$1 - \frac{\lambda}{N} \psi_i = \exp \left(-\frac{\lambda}{N} \chi_i \right),$$

which is possible when $\frac{\lambda}{N} \in]-1, 1[$ (and leads to define

$$\chi_i = -\frac{N}{\lambda} \log \left(1 - \frac{\lambda}{N} \psi_i \right),$$

we obtain easily following the same reasoning as previously

$$\log \left\{ \mathbb{P} \left\{ \exp \left[\sum_{i=1}^N \log \left(1 - \frac{\lambda}{N} \psi_i \right) \right] \right\} \right\}$$

$$\leq \sum_{i=1}^N \log \left[1 - \frac{\lambda}{N} \mathbb{P}(\psi_i) \right] \leq N \log \left[1 - \frac{\lambda}{N} R'(\theta, \tilde{\theta}) \right].$$

Let us replace for simplicity λ/N with λ . Let us also introduce the random pseudo distance

$$\begin{aligned} m'(\theta, \tilde{\theta}) &= \frac{1}{N} \sum_{i=1}^N \psi_i(\theta, \tilde{\theta})^2 \\ &= \frac{1}{N} \sum_{i=1}^N \left| \mathbb{1}[f_\theta(X_i) \neq Y_i] - \mathbb{1}[f_{\tilde{\theta}}(X_i) \neq Y_i] \right|, \quad \theta, \tilde{\theta} \in \Theta. \end{aligned} \quad (1.19)$$

This is the empirical counter part of M' , since $\mathbb{P}(m') = M'$. Let us notice that

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \log [1 - \lambda \psi_i(\theta, \tilde{\theta})] &= \frac{\log(1 - \lambda) - \log(1 + \lambda)}{2} r'(\theta, \tilde{\theta}) \\ &\quad + \frac{\log(1 - \lambda) + \log(1 + \lambda)}{2} m'(\theta, \tilde{\theta}) \\ &= \frac{1}{2} \log \left(\frac{1 - \lambda}{1 + \lambda} \right) r'(\theta, \tilde{\theta}) + \frac{1}{2} \log(1 - \lambda^2) m'(\theta, \tilde{\theta}). \end{aligned}$$

With these notations, we can conveniently write the previous inequality as

$$\mathbb{P} \left\{ \exp \left[-N \log [1 - \lambda R'(\theta, \tilde{\theta})] \right. \right. \\ \left. \left. - \frac{N}{2} \log \left(\frac{1 + \lambda}{1 - \lambda} \right) r'(\theta, \tilde{\theta}) + \frac{N}{2} \log(1 - \lambda^2) m'(\theta, \tilde{\theta}) \right] \right\} \leq 1.$$

Integrating with respect to a prior probability measure $\pi \in \mathcal{M}_+^1(\Theta)$, we obtain

THEOREM 1.26. *For any real parameter $\lambda \in]-1, 1[$, for any $\tilde{\theta} \in \Theta$, for any prior probability distribution $\pi \in \mathcal{M}_+^1(\Theta)$,*

$$\mathbb{P} \left\{ \exp \left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \left\{ -N \rho \left\{ \log [1 - \lambda R'(\cdot, \tilde{\theta})] \right\} \right. \right. \right.$$

$$\left. \begin{aligned} & -\frac{N}{2} \log\left(\frac{1+\lambda}{1-\lambda}\right) \rho[r'(\cdot, \tilde{\theta})] \\ & + \frac{N}{2} \log(1-\lambda^2) \rho[m'(\cdot, \tilde{\theta})] - \mathcal{K}(\rho, \pi) \end{aligned} \right\} \leq 1.$$

1.4.2. Non random bounds. Let us first deduce a non random bound from Theorem 1.25. This theorem can be conveniently taken advantage of by throwing the non linearity into a localized prior, considering the prior probability measure μ defined by

$$\frac{d\mu}{d\pi}(\theta) = \frac{\exp\{-\lambda\Psi_{\frac{\lambda}{N}}[R'(\theta, \tilde{\theta}), M'(\theta, \tilde{\theta})] + \beta R'(\theta, \tilde{\theta})\}}{\pi\left\{\exp\{-\lambda\Psi_{\frac{\lambda}{N}}[R'(\cdot, \tilde{\theta}), M'(\cdot, \tilde{\theta})] + \beta R'(\cdot, \tilde{\theta})\}\right\}}.$$

Indeed, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,

$$\begin{aligned} \mathcal{K}(\rho, \mu) &= \mathcal{K}(\rho, \pi) + \lambda\rho\left\{\Psi_{\frac{\lambda}{N}}[R'(\cdot, \tilde{\theta}), M'(\cdot, \tilde{\theta})]\right\} - \beta\rho[R'(\cdot, \tilde{\theta})] \\ &+ \log\left\{\pi\left[\exp\{-\lambda\Psi_{\frac{\lambda}{N}}[R'(\cdot, \tilde{\theta}), M'(\cdot, \tilde{\theta})] + \beta R'(\cdot, \tilde{\theta})\}\right]\right\}. \end{aligned}$$

Plugging this into Theorem 1.25 and using the convexity of the exponential function, we see that for any posterior probability distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,

$$\begin{aligned} \beta\mathbb{P}\{\rho[R'(\cdot, \tilde{\theta})]\} &\leq \lambda\mathbb{P}\{\rho[r'(\cdot, \tilde{\theta})]\} + \mathbb{P}[\mathcal{K}(\rho, \pi)] \\ &+ \log\left\{\pi\left[\exp\{-\lambda\Psi_{\frac{\lambda}{N}}[R'(\cdot, \tilde{\theta}), M'(\cdot, \tilde{\theta})] + \beta R'(\cdot, \tilde{\theta})\}\right]\right\}. \end{aligned}$$

We can then recall that

$$\lambda\rho[r'(\cdot, \tilde{\theta})] + \mathcal{K}(\rho, \pi) = \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}] - \log\left\{\pi\left[\exp[-\lambda r'(\cdot, \tilde{\theta})]\right]\right\},$$

and notice moreover that

$$-\mathbb{P}\left\{\log\left\{\pi\left[\exp[-\lambda r'(\cdot, \tilde{\theta})]\right]\right\}\right\} \leq -\log\left\{\pi\left[\exp[-\lambda R'(\cdot, \tilde{\theta})]\right]\right\},$$

since $R' = \mathbb{P}(r')$ and $h \mapsto \log\left\{\pi\left[\exp(h)\right]\right\}$ is a convex functional. Putting these two remarks together, we obtain

THEOREM 1.27. *For any real positive parameter λ , for any prior distribution $\pi \in \mathcal{M}_+^1(\Theta)$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned} \mathbb{P}\{\rho[R'(\cdot, \tilde{\theta})]\} &\leq \frac{1}{\beta} \mathbb{P}[\mathcal{K}(\rho, \pi_{\exp(-\lambda r)})] \\ &\quad + \frac{1}{\beta} \log \left\{ \pi \left[\exp \left\{ -\lambda \Psi_{\frac{\lambda}{N}} [R'(\cdot, \tilde{\theta}), M'(\cdot, \tilde{\theta})] + \beta R'(\cdot, \tilde{\theta}) \right\} \right] \right\} \\ &\quad - \frac{1}{\beta} \log \left\{ \pi \left[\exp[-\lambda R'(\cdot, \tilde{\theta})] \right] \right\} \\ &\leq \frac{1}{\beta} \mathbb{P}[\mathcal{K}(\rho, \pi_{\exp(-\lambda r)})] \\ &\quad + \frac{1}{\beta} \log \left\{ \pi \left[\exp \left\{ -[N \sinh(\frac{\lambda}{N}) - \beta] R'(\cdot, \tilde{\theta}) \right. \right. \right. \\ &\quad \left. \left. \left. + 2N \sinh(\frac{\lambda}{2N})^2 M'(\cdot, \tilde{\theta}) \right\} \right] \right\} \\ &\quad - \frac{1}{\beta} \log \left\{ \pi \left[\exp[-\lambda R'(\cdot, \tilde{\theta})] \right] \right\}. \end{aligned}$$

It may be interesting to derive some more suggestive (but slightly weaker) bound in the important case when $\Theta_1 = \Theta$ and $R(\tilde{\theta}) = \inf_{\Theta} R$. In this case, it is convenient to introduce the *margin function*

$$\varphi(x) = \sup_{\theta \in \Theta} M'(\theta, \tilde{\theta}) - x R'(\theta, \tilde{\theta}), \quad x \in \mathbb{R}_+. \quad (1.20)$$

We see that φ is convex and nonnegative on \mathbb{R}_+ . Using the bound $M'(\theta, \tilde{\theta}) \leq x R'(\theta, \tilde{\theta}) + \varphi(x)$, we obtain

$$\begin{aligned} \mathbb{P}\{\rho[R'(\cdot, \tilde{\theta})]\} &\leq \frac{1}{\beta} \mathbb{P}[\mathcal{K}(\rho, \pi_{\exp(-\lambda r)})] \\ &\quad + \frac{1}{\beta} \log \left\{ \pi \left[\exp \left\{ -\{N \sinh(\frac{\lambda}{N}) [1 - x \tanh(\frac{\lambda}{2N})] - \beta\} R'(\cdot, \tilde{\theta}) \right\} \right] \right\} \\ &\quad + \frac{N \sinh(\frac{\lambda}{N}) \tanh(\frac{\lambda}{2N})}{\beta} \varphi(x) - \frac{1}{\beta} \log \left\{ \pi \left[\exp[-\lambda R'(\cdot, \tilde{\theta})] \right] \right\}. \end{aligned}$$

Let us make the change of variable $\gamma = N \sinh(\frac{\lambda}{N}) [1 - x \tanh(\frac{\lambda}{2N})] - \beta$ to obtain

COROLLARY 1.28. *For any real positive parameters x , γ and λ such that $x \leq \tanh(\frac{\lambda}{2N})^{-1}$ and $0 \leq \gamma < N \sinh(\frac{\lambda}{N}) [1 - x \tanh(\frac{\lambda}{2N})]$,*

$$\begin{aligned} \mathbb{P}[\rho(R)] - \inf_{\Theta} R &\leq \left\{ N \sinh\left(\frac{\lambda}{N}\right) \left[1 - x \tanh\left(\frac{\lambda}{2N}\right)\right] - \gamma \right\}^{-1} \\ &\times \left\{ \int_{\gamma}^{\lambda} [\pi_{\exp(-\alpha R)}(R) - \inf_{\Theta} R] d\alpha \right. \\ &\quad \left. + N \sinh\left(\frac{\lambda}{N}\right) \tanh\left(\frac{\lambda}{2N}\right) \varphi(x) + \mathbb{P}[\mathcal{K}(\rho, \pi_{\exp(-\lambda r)})] \right\}. \end{aligned}$$

Let us remark that these results, although well suited to study Mammen and Tsybakov's margin assumptions, hold in the general case: introducing the convex *expected margin function* φ is a substitute for making hypotheses about the relations between R and D .

Using the fact that $R'(\theta, \tilde{\theta}) \geq 0$, $\theta \in \Theta$ and that $\varphi(x) \geq 0$, $x \in \mathbb{R}_+$, we can weaken and simplify even more the preceding corollary to get

COROLLARY 1.29. *For any real parameters β , λ and x such that $x \geq 0$ and $0 \leq \beta < \lambda - x \frac{\lambda^2}{2N}$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned} \mathbb{P}[\rho(R)] &\leq \inf_{\Theta} R \\ &+ \left[\lambda - x \frac{\lambda^2}{2N} - \beta \right]^{-1} \left\{ \int_{\beta}^{\lambda} [\pi_{\exp(-\alpha R)}(R) - \inf_{\Theta} R] d\alpha \right. \\ &\quad \left. + \mathbb{P}\{\mathcal{K}[\rho, \pi_{\exp(-\lambda r)}]\} + \varphi(x) \frac{\lambda^2}{2N} \right\}. \end{aligned}$$

Let us apply this bound under the *margin assumption* first considered by Mammen and Tsybakov [28, 34], which tells that for some real positive constant c and some real exponent $\kappa \geq 1$,

$$R'(\theta, \tilde{\theta}) \geq cD(\theta, \tilde{\theta})^{\kappa}, \quad \theta \in \Theta. \quad (1.21)$$

In the case when $\kappa = 1$, then $\varphi(c^{-1}) = 0$, proving that

$$\begin{aligned} \mathbb{P}\{\pi_{\exp(-\lambda r)}[R'(\cdot, \tilde{\theta})]\} &\leq \frac{\int_{\beta}^{\lambda} \pi_{\exp(-\gamma R)}[R'(\cdot, \tilde{\theta})] d\gamma}{N \sinh\left(\frac{\lambda}{N}\right) \left[1 - c^{-1} \tanh\left(\frac{\lambda}{2N}\right)\right] - \beta} \\ &\leq \frac{\int_{\beta}^{\lambda} \pi_{\exp(-\gamma R)}[R'(\cdot, \tilde{\theta})] d\gamma}{\lambda - \frac{\lambda^2}{2cN} - \beta}. \end{aligned}$$

Taking for example $\lambda = \frac{cN}{2}$, $\beta = \frac{\lambda}{2} = \frac{cN}{4}$, we obtain

$$\begin{aligned} \mathbb{P}[\pi_{\exp(-2^{-1}cNr)}(R)] &\leq \inf R + \frac{8}{cN} \int_{\frac{cN}{4}}^{\frac{cN}{2}} \pi_{\exp(-\gamma R)}[R'(\cdot, \tilde{\theta})] d\gamma \\ &\leq \inf R + 2\pi_{\exp(-\frac{cN}{4}R)}[R'(\cdot, \tilde{\theta})]. \end{aligned}$$

If moreover the behaviour of the prior distribution π is parametric meaning that $\pi_{\exp(-\beta R)}[R'(\cdot, \tilde{\theta})] \leq \frac{d}{\beta}$, for some positive real constant d linked with the dimension of the classification model, then

$$\mathbb{P}\left[\pi_{\exp(-\frac{cN}{2}r)}(R)\right] \leq \inf R + \frac{8 \log(2)d}{cN} \leq \inf R + \frac{5.55 d}{cN}.$$

In the case when $\kappa > 1$,

$$\varphi(x) \leq (\kappa - 1)\kappa^{-\frac{\kappa}{\kappa-1}}(cx)^{-\frac{1}{\kappa-1}} = (1 - \kappa^{-1})(\kappa cx)^{-\frac{1}{\kappa-1}},$$

thus $\mathbb{P}\left\{\pi_{\exp(-\lambda r)}[R'(\cdot, \tilde{\theta})]\right\}$

$$\leq \frac{\int_{\beta}^{\lambda} \pi_{\exp(-\gamma R)}[R'(\cdot, \tilde{\theta})]d\gamma + (1 - \kappa^{-1})(\kappa cx)^{-\frac{1}{\kappa-1}} \frac{\lambda^2}{2N}}{\lambda - \frac{x\lambda^2}{2N} - \beta}.$$

Taking for instance $\beta = \frac{\lambda}{2}$, $x = \frac{N}{2\lambda}$, and putting $b = (1 - \kappa^{-1})(\kappa c)^{-\frac{1}{\kappa-1}}$, we obtain

$$\mathbb{P}\left[\pi_{\exp(-\lambda r)}(R)\right] - \inf R \leq \frac{4}{\lambda} \int_{\lambda/2}^{\lambda} \pi_{\exp(-\gamma R)}[R'(\cdot, \tilde{\theta})]d\gamma + b \left(\frac{2\lambda}{N}\right)^{\frac{\kappa}{\kappa-1}}.$$

In the *parametric* case when $\pi_{\exp(-\gamma R)}[R'(\cdot, \tilde{\theta})] \leq \frac{d}{\gamma}$, we get

$$\mathbb{P}\left[\pi_{\exp(-\lambda r)}(R)\right] - \inf R \leq \frac{4 \log(2)d}{\lambda} + b \left(\frac{2\lambda}{N}\right)^{\frac{\kappa}{\kappa-1}}.$$

Taking

$$\bar{\lambda} = 2^{-1} [8 \log(2)d]^{\frac{\kappa-1}{2\kappa-1}} (\kappa c)^{\frac{1}{2\kappa-1}} N^{\frac{\kappa}{2\kappa-1}},$$

we obtain

$$\mathbb{P}\left[\pi_{\exp(-\bar{\lambda} r)}(R)\right] - \inf R \leq (2 - \kappa^{-1})(\kappa c)^{-\frac{1}{2\kappa-1}} \left(\frac{8 \log(2)d}{N}\right)^{\frac{\kappa}{2\kappa-1}}.$$

We see that this formula coincides with the result for $\kappa = 1$. We can thus reduce the two cases to a single one and state

COROLLARY 1.30. *Let us assume that for some $\tilde{\theta} \in \Theta$, some positive real constant c , some real exponent $\kappa \geq 1$ and for any $\theta \in \Theta$, $R(\theta) \geq R(\tilde{\theta}) + cD(\theta, \tilde{\theta})^{\kappa}$. Let us also assume that for some positive real constant d and any positive real parameter γ , $\pi_{\exp(-\gamma R)}(R) - \inf R \leq \frac{d}{\gamma}$. Then*

$$\begin{aligned} & \mathbb{P} \left[\pi \left[\exp \left\{ -2^{-1} [8 \log(2)d]^{\frac{\kappa-1}{2\kappa-1}} (\kappa c)^{\frac{1}{2\kappa-1}} N^{\frac{\kappa}{2\kappa-1}} r \right\} (R) \right] \right] \\ & \leq \inf R + (2 - \kappa^{-1}) (\kappa c)^{-\frac{1}{2\kappa-1}} \left(\frac{8 \log(2)d}{N} \right)^{\frac{\kappa}{2\kappa-1}}. \end{aligned}$$

Let us remark that the exponent of N in this corollary is known to be the minimax exponent under these assumptions: it is unimprovable, whatever estimator is used in place of the Gibbs posterior shown here (at least in the worst case compatible with the hypotheses). The interest of the corollary is to show not only the minimax exponent in N , but also an explicit non asymptotic bound with reasonable and simple constants. It is also clear that we could have got slightly better constants if we had kept the full strength of Theorem 1.27 (page 46) instead of using the weaker Corollary 1.29 (page 47).

We will prove in the following empirical bounds showing how the constant λ can be estimated from the data instead of being chosen according to some margin and complexity assumptions.

1.4.3. Unbiased empirical bounds. We are going to provide an empirical counter part for the *expected margin function* φ . It will appear in empirical bounds having otherwise the same structure as the non random bound we just proved. Anyhow, we will not launch into trying to compare the behaviour of our proposed *empirical margin function* with the *expected margin function*, since the margin function involves taking a supremum which is not straightforward to handle.

Let us start as in the previous subsection with the inequality

$$\begin{aligned} \beta \mathbb{P} \left\{ \rho [R'(\cdot, \tilde{\theta})] \right\} & \leq \mathbb{P} \left\{ \lambda \rho [r'(\cdot, \tilde{\theta})] + \mathcal{K}(\rho, \pi) \right\} \\ & + \log \left\{ \pi \left[\exp \left\{ -\lambda \Psi_{\frac{\lambda}{N}} [R'(\cdot, \tilde{\theta}), M'(\cdot, \tilde{\theta})] + \beta R'(\cdot, \tilde{\theta}) \right\} \right] \right\}. \end{aligned}$$

We have already defined by equation (1.19) the empirical pseudo distance

$$m'(\theta, \tilde{\theta}) = \frac{1}{N} \sum_{i=1}^N \psi_i(\theta, \tilde{\theta})^2.$$

Recalling that $\mathbb{P} [m'(\theta, \tilde{\theta})] = M'(\theta, \tilde{\theta})$, and using the convexity of $h \mapsto \log \left\{ \pi [\exp(h)] \right\}$, leads to the following inequalities:

$$\begin{aligned}
& \log \left\{ \pi \left[\exp \left\{ -\lambda \Psi_{\frac{\lambda}{N}} \left[R'(\cdot, \tilde{\theta}), M'(\cdot, \tilde{\theta}) \right] + \beta R'(\cdot, \tilde{\theta}) \right\} \right] \right\} \\
& \leq \log \left\{ \pi \left[\exp \left\{ -N \sinh\left(\frac{\lambda}{N}\right) R'(\cdot, \tilde{\theta}) \right. \right. \right. \\
& \quad \left. \left. \left. + N \sinh\left(\frac{\lambda}{N}\right) \tanh\left(\frac{\lambda}{2N}\right) M'(\cdot, \tilde{\theta}) + \beta R'(\cdot, \tilde{\theta}) \right\} \right] \right\} \\
& \leq \mathbb{P} \left\{ \log \left\{ \pi \left[\exp \left\{ -\left[N \sinh\left(\frac{\lambda}{N}\right) - \beta \right] r'(\cdot, \tilde{\theta}) \right. \right. \right. \right. \\
& \quad \left. \left. \left. + N \sinh\left(\frac{\lambda}{N}\right) \tanh\left(\frac{\lambda}{2N}\right) m'(\cdot, \tilde{\theta}) \right\} \right] \right\} \right\}.
\end{aligned}$$

We may moreover remark that

$$\begin{aligned}
\lambda \rho[r'(\cdot, \tilde{\theta})] + \mathcal{K}(\rho, \pi) &= [\beta - N \sinh(\frac{\lambda}{N}) + \lambda] \rho[r'(\cdot, \tilde{\theta})] \\
& \quad + \mathcal{K}[\rho, \pi_{\exp\{-[N \sinh(\frac{\lambda}{N}) - \beta]r\}}] \\
& \quad - \log \left\{ \pi \left[\exp \left\{ -\left[N \sinh\left(\frac{\lambda}{N}\right) - \beta \right] r'(\cdot, \tilde{\theta}) \right\} \right] \right\}.
\end{aligned}$$

This ends to prove

THEOREM 1.31. *For any positive real parameters β and λ , for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned}
\mathbb{P}\{\rho[R'(\cdot, \tilde{\theta})]\} &\leq \mathbb{P} \left\{ \left[1 - \frac{N \sinh(\frac{\lambda}{N}) - \lambda}{\beta} \right] \rho[r'(\cdot, \tilde{\theta})] \right. \\
& \quad \left. + \frac{\mathcal{K}[\rho, \pi_{\exp\{-[N \sinh(\frac{\lambda}{N}) - \beta]r\}}]}{\beta} \right. \\
& \quad \left. + \beta^{-1} \log \left\{ \pi_{\exp\{-[N \sinh(\frac{\lambda}{N}) - \beta]r\}} \left[\exp \left[N \sinh\left(\frac{\lambda}{N}\right) \tanh\left(\frac{\lambda}{2N}\right) m'(\cdot, \tilde{\theta}) \right] \right] \right\} \right\}.
\end{aligned}$$

Taking $\beta = \frac{N}{2} \sinh(\frac{\lambda}{N})$, using the fact that $\sinh(a) \geq a$, $a \geq 0$ and expressing $\tanh(\frac{a}{2}) = a^{-1} [\sqrt{1 + \sinh(a)^2} - 1]$ and $a = \log[\sqrt{1 + \sinh(a)^2} + \sinh(a)]$, we deduce

COROLLARY 1.32. *For any positive real constant β and any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\mathbb{P}\{\rho[R'(\cdot, \tilde{\theta})]\} \leq \mathbb{P} \left\{ \underbrace{\left[\frac{N}{\beta} \log \left(\sqrt{1 + \frac{4\beta^2}{N^2}} + \frac{2\beta}{N} \right) - 1 \right]}_{\leq 1} \rho[r'(\cdot, \tilde{\theta})] \right\}$$

$$\begin{aligned}
& + \frac{1}{\beta} \left\{ \mathcal{K}[\rho, \pi_{\exp(-\beta r)}] \right. \\
& \quad \left. + \log \left[\pi_{\exp(-\beta r)} \left\{ \exp \left[N \left(\sqrt{1 + \frac{4\beta^2}{N^2}} - 1 \right) m'(\cdot, \tilde{\theta}) \right] \right\} \right] \right\}.
\end{aligned}$$

This theorem and its corollary are really analogous to Theorem 1.27 (page 46) and it could easily be proved that under Mammen and Tsybakov margin assumptions, we obtain an upper bound of the same order as Corollary 1.30 (page 48). Anyhow, in order to obtain an empirical bound, we are going now to take a supremum over all possible values of $\tilde{\theta}$, that is over Θ_1 . Although we believe that taking this supremum will not spoil the bound in cases when overfitting remains under control, we will not try to investigate precisely if and when this is actually true, and provide our empirical bound as such. Let us only say that on a qualitative ground, the values of the margin function quantify how steep is the contrast function R or its empirical counterpart r , and that the definition of the empirical margin function is obtained by substituting \mathbb{P} , the true sample distribution, with $\bar{\mathbb{P}} = \left(\frac{1}{N} \sum_{i=1}^N \delta_{(X_i, Y_i)}\right)^{\otimes N}$, the empirical sample distribution, in the definition of the expected margin function. Therefore, on qualitative grounds, it sounds like hopeless to presume that R is steep when r is not, or in other words that a classification model that would be unefficient at estimating a bootstrapped sample according to our non random bound would be by some miracle efficient at estimating the true sample distribution according to the same bound. To this extent, we feel that our empirical bounds bring a satisfactory counterpart of our non random bounds. Anyhow, we will also produce estimators which can be proved to be adaptive using PAC-Bayesian tools in the next subsection, at the price of a more sophisticated construction involving comparisons between a posterior distribution and a Gibbs prior distribution.

Let us restrict now to the important case when $\tilde{\theta} \in \arg \min_{\Theta_1} R$. To obtain an observable bound, let $\hat{\theta} \in \arg \min_{\theta \in \Theta} r(\theta)$ and let us introduce the *empirical margin functions*

$$\begin{aligned}
\bar{\varphi}(x) &= \sup_{\theta \in \Theta} m'(\theta, \hat{\theta}) - x[r(\theta) - r(\hat{\theta})], \quad x \in \mathbb{R}_+, \\
\tilde{\varphi}(x) &= \sup_{\theta \in \Theta_1} m'(\theta, \hat{\theta}) - x[r(\theta) - r(\hat{\theta})], \quad x \in \mathbb{R}_+.
\end{aligned}$$

Using the fact that $m'(\theta, \tilde{\theta}) \leq m'(\theta, \hat{\theta}) + m'(\hat{\theta}, \tilde{\theta})$, we get

COROLLARY 1.33. *For any positive real parameters β and λ , for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned}
\mathbb{P}[\rho(R)] - \inf_{\Theta_1} R &\leq \mathbb{P} \left\{ \left[1 - \frac{N \sinh(\frac{\lambda}{N}) - \lambda}{\beta} \right] [\rho(r) - r(\hat{\theta})] \right. \\
&\quad \left. + \frac{\mathcal{K}[\rho, \pi_{\exp\{-[N \sinh(\frac{\lambda}{N}) - \beta]r\}}]}{\beta} \right. \\
&\quad \left. + \beta^{-1} \log \left\{ \pi_{\exp\{-[N \sinh(\frac{\lambda}{N}) - \beta]r\}} \left[\exp \left[N \sinh(\frac{\lambda}{N}) \tanh(\frac{\lambda}{2N}) m'(\cdot, \hat{\theta}) \right] \right] \right\} \right. \\
&\quad \left. + \beta^{-1} N \sinh(\frac{\lambda}{N}) \tanh(\frac{\lambda}{2N}) \tilde{\varphi} \left[\frac{\beta}{N \sinh(\frac{\lambda}{N}) \tanh(\frac{\lambda}{2N})} \left(1 - \frac{N \sinh(\frac{\lambda}{N}) - \lambda}{\beta} \right) \right] \right\}.
\end{aligned}$$

Taking $\beta = \frac{N}{2} \sinh(\frac{\lambda}{N})$, we also obtain

$$\begin{aligned}
\mathbb{P}[\rho(R)] - \inf_{\Theta_1} R &\leq \mathbb{P} \left\{ \underbrace{\left[\frac{N}{\beta} \log \left(\sqrt{1 + \frac{4\beta^2}{N^2}} + \frac{2\beta}{N} \right) - 1 \right]}_{\leq 1} [\rho(r) - r(\hat{\theta})] \right. \\
&\quad \left. + \frac{1}{\beta} \left\{ \mathcal{K}[\rho, \pi_{\exp(-\beta r)}] \right. \right. \\
&\quad \left. \left. + \log \left[\pi_{\exp(-\beta r)} \left\{ \exp \left[N \left(\sqrt{1 + \frac{4\beta^2}{N^2}} - 1 \right) m'(\cdot, \hat{\theta}) \right] \right\} \right] \right\} \right. \\
&\quad \left. + \frac{N}{\beta} \left(\sqrt{1 + \frac{4\beta^2}{N^2}} - 1 \right) \tilde{\varphi} \left[\frac{\log \left(\sqrt{1 + \frac{4\beta^2}{N^2}} + \frac{2\beta}{N} \right) - \frac{\beta}{N}}{\left(\sqrt{1 + \frac{4\beta^2}{N^2}} - 1 \right)} \right] \right\}.
\end{aligned}$$

Note that we could also use the upper bound $m'(\theta, \hat{\theta}) \leq x[r(\theta) - r(\hat{\theta})] + \bar{\varphi}(x)$ and put $\alpha = N \sinh(\frac{\lambda}{N}) [1 - x \tanh(\frac{\lambda}{2N})] - \beta$, to obtain

COROLLARY 1.34. *For any non negative real parameters x , α and λ , such that $\alpha < N \sinh(\frac{\lambda}{N}) [1 - x \tanh(\frac{\lambda}{2N})]$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned}
&\mathbb{P}[\rho(R)] - \inf_{\Theta_1} R \\
&\leq \mathbb{P} \left\{ \left[1 - \frac{N \sinh(\frac{\lambda}{N}) [1 - x \tanh(\frac{\lambda}{2N})] - \lambda}{N \sinh(\frac{\lambda}{N}) [1 - x \tanh(\frac{\lambda}{2N})] - \alpha} \right] [\rho(r) - r(\hat{\theta})] \right. \\
&\quad \left. + \frac{\mathcal{K}[\rho, \pi_{\exp(-\alpha r)}]}{N \sinh(\frac{\lambda}{N}) [1 - x \tanh(\frac{\lambda}{2N})] - \alpha} \right. \\
&\quad \left. + \frac{N \sinh(\frac{\lambda}{N}) \tanh(\frac{\lambda}{2N})}{N \sinh(\frac{\lambda}{N}) [1 - x \tanh(\frac{\lambda}{2N})] - \alpha} \right\}
\end{aligned}$$

$$\times \left[\bar{\varphi}(x) + \tilde{\varphi} \left(\frac{\lambda - \alpha}{N \sinh(\frac{\lambda}{N}) \tanh(\frac{\lambda}{2N})} \right) \right].$$

Let us notice that in the case when $\Theta_1 = \Theta$, the upper bound provided by this corollary has the same general form as the upper bound provided by Corollary 1.28 (page 46), with the sample distribution \mathbb{P} replaced with the empirical distribution of the sample $\bar{\mathbb{P}} = \left(\frac{1}{N} \sum_{i=1}^N \delta_{(X_i, Y_i)} \right)^{\otimes N}$. Therefore, our empirical bound can be of a larger order of magnitude than our non random bound only in the case when our non random bound applied to the bootstrapped sample distribution $\bar{\mathbb{P}}$ would be of a larger order of magnitude than when applied to the true sample distribution \mathbb{P} . In other words, we can say that our empirical bound is close to our non random bound in every situation where the bootstrapped sample distribution $\bar{\mathbb{P}}$ is not harder to bound than the true sample distribution \mathbb{P} . Although this does not prove that our empirical bound is always of the same order as our non random bound, this is a good qualitative hint that this will be the case in most practical situations of interest, since in situations of “underfitting”, if they exist, it is likely that the choice of the classification model is inappropriate to the data and should be modified.

Another reassuring remark is that the empirical margin functions $\bar{\varphi}$ and $\tilde{\varphi}$ behave well in the case when $\inf_{\Theta} r = 0$. Indeed in this case $m'(\theta, \hat{\theta}) = r'(\theta, \hat{\theta}) = r(\theta)$, $\theta \in \Theta$, and thus $\bar{\varphi}(1) = \tilde{\varphi}(1) = 0$, and

$$\tilde{\varphi}(x) \leq -(x - 1) \inf_{\Theta_1} r, \quad x \geq 1.$$

This shows that we recover in this case the same accuracy as with non relative local empirical bounds. Thus the bound of Corollary 1.34 does not collapse in presence of massive overfitting in the larger model, causing $r(\hat{\theta}) = 0$, which is another hint that this may be an accurate bound in many situations.

1.4.4. Relative empirical deviation bounds. It is natural to make use of Theorem 1.26 on page 44 to obtain empirical deviation bounds, since this theorem provides an empirical variance term.

Theorem 1.26 is written in a way which exploits the fact that ψ_i takes only the three values -1, 0 and +1. However, it will be more convenient for the following computations to use it in its more general form, which only makes use of the fact that $\psi_i \in (-1, 1)$. With notations to be explained hereafter, it can indeed also be written as

$$\mathbb{P} \left\{ \exp \left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \left\{ -N\rho \left\{ \log \left[1 - \lambda P(\psi) \right] \right\} \right\} \right] \right\}$$

$$+ N\rho\left\{\overline{P}\left[\log(1 - \lambda\psi)\right] - \mathcal{K}(\rho, \pi)\right\} \leq 1. \quad (1.22)$$

We have used the following notations in this inequality. We have put

$$\overline{P} = \frac{1}{N} \sum_{i=1}^N \delta_{(X_i, Y_i)},$$

so that \overline{P} is our notation for the empirical distribution of the process $(X_i, Y_i)_{i=1}^N$. Moreover we have also used

$$P = \mathbb{P}(\overline{P}) = \frac{1}{N} \sum_{i=1}^N P_i,$$

where it should be remembered that the joint distribution of the process $(X_i, Y_i)_{i=1}^N$ is $\mathbb{P} = \bigotimes_{i=1}^N P_i$. We have considered $\psi(\theta, \tilde{\theta})$ as a function defined on $\mathcal{X} \times \mathcal{Y}$,

as $\psi(\theta, \tilde{\theta})(x, y) = \mathbb{1}[y \neq f_\theta(x)] - \mathbb{1}[y \neq f_{\tilde{\theta}}(x)]$, $(x, y) \in \mathcal{X} \times \mathcal{Y}$
so that it should be understood that

$$\begin{aligned} P(\psi) &= \frac{1}{N} \sum_{i=1}^N \mathbb{P}[\psi_i(\theta, \tilde{\theta})] \\ &= \frac{1}{N} \sum_{i=1}^N \mathbb{P}\left\{\mathbb{1}[Y_i \neq f_\theta(X_i)] - \mathbb{1}[Y_i \neq f_{\tilde{\theta}}(X_i)]\right\} = R'(\theta, \tilde{\theta}). \end{aligned}$$

In the same way

$$\overline{P}\left[\log(1 - \lambda\psi)\right] = \frac{1}{N} \sum_{i=1}^N \log[1 - \lambda\psi_i(\theta, \tilde{\theta})].$$

Moreover integration with respect to ρ bears on the index θ , so that

$$\begin{aligned} \rho\left\{\log\left[1 - \lambda P(\psi)\right]\right\} &= \int_{\theta \in \Theta} \log\left\{1 - \frac{\lambda}{N} \sum_{i=1}^N \mathbb{P}[\psi_i(\theta, \tilde{\theta})]\right\} \rho(d\theta), \\ \rho\left\{\overline{P}\left[\log(1 - \lambda\psi)\right]\right\} &= \int_{\theta \in \Theta} \left\{\frac{1}{N} \sum_{i=1}^N \log[1 - \lambda\psi_i(\theta, \tilde{\theta})]\right\} \rho(d\theta). \end{aligned}$$

We have chosen concise notations, as we did throughout these notes, in order to make the computations easier to follow.

To get an alternate version of empirical relative deviation bounds, we need to find some convenient way to localize the choice of the prior distribution π in equation (1.22, page 54). Here we propose to replace π with $\mu = \pi_{\exp\{-N \log[1+\beta P(\psi)]\}}$, which can also be written $\pi_{\exp\{-N \log[1+\beta R'(\cdot, \tilde{\theta})]\}}$. Indeed we see that

$$\begin{aligned} \mathcal{K}(\rho, \mu) &= N\rho \left\{ \log[1 + \beta P(\psi)] \right\} + \mathcal{K}(\rho, \pi) \\ &\quad + \log \left\{ \pi \left[\exp\{-N \log[1 + \beta P(\psi)]\} \right] \right\}. \end{aligned}$$

Moreover, we deduce from our deviation inequality applied to $-\psi$, that (as long as $\beta > -1$),

$$\mathbb{P} \left\{ \exp \left[N\mu \left\{ \bar{P}[\log(1 + \beta\psi)] \right\} - N\mu \left\{ \log[1 + \beta P(\psi)] \right\} \right] \right\} \leq 1.$$

Thus

$$\begin{aligned} &\mathbb{P} \left\{ \exp \left[\log \left\{ \pi \left[\exp\{-N \log[1 + \beta P(\psi)]\} \right] \right\} \right. \right. \\ &\quad \left. \left. - \log \left\{ \pi \left[\exp\{-N \bar{P}[\log(1 + \beta\psi)]\} \right] \right\} \right] \right\} \\ &\leq \mathbb{P} \left\{ \exp \left[-N\mu \left\{ \log[1 + \beta P(\psi)] \right\} - \mathcal{K}(\mu, \pi) \right. \right. \\ &\quad \left. \left. + N\mu \left\{ \bar{P}[\log(1 + \beta\psi)] \right\} + \mathcal{K}(\mu, \pi) \right] \right\} \leq 1. \end{aligned}$$

This can be used to handle $\mathcal{K}(\rho, \mu)$, making use of the Cauchy Schwarz inequality as follows

$$\begin{aligned} &\mathbb{P} \left\{ \exp \left[\frac{1}{2} \left[-N \log \left\{ (1 - \lambda\rho[P(\psi)]) (1 + \beta\rho[P(\psi)]) \right\} \right. \right. \right. \\ &\quad \left. \left. + N\rho \left\{ \bar{P}[\log(1 - \lambda\psi)] \right\} \right. \right. \\ &\quad \left. \left. - \mathcal{K}(\rho, \pi) - \log \left\{ \pi \left[\exp\{-N \bar{P}[\log(1 + \beta\psi)]\} \right] \right\} \right] \right\} \\ &\leq \mathbb{P} \left\{ \exp \left[-N \log \left\{ (1 - \lambda\rho[P(\psi)]) \right\} \right. \right. \\ &\quad \left. \left. + N\rho \left\{ \bar{P}[\log(1 - \lambda\psi)] \right\} - \mathcal{K}(\rho, \mu) \right] \right\}^{1/2} \end{aligned}$$

$$\begin{aligned} & \times \mathbb{P} \left\{ \exp \left[\log \left\{ \pi \left[\exp \left\{ -N \log [1 + \beta P(\psi)] \right\} \right] \right\} \right. \right. \\ & \quad \left. \left. - \log \left\{ \pi \left[\exp \left\{ -N \bar{P} [\log(1 + \beta \psi)] \right\} \right] \right\} \right] \right\}^{1/2} \leq 1. \end{aligned}$$

This implies that with \mathbb{P} probability at least $1 - \epsilon$,

$$\begin{aligned} & -N \log \left\{ \left(1 - \lambda \rho [P(\psi)] \right) \left(1 + \beta \rho [P(\psi)] \right) \right\} \\ & \leq -N \rho \left\{ \bar{P} [\log(1 - \lambda \psi)] \right\} \\ & \quad + \mathcal{K}(\rho, \pi) + \log \left\{ \pi \left[\exp \left\{ -N \bar{P} [\log(1 + \beta \psi)] \right\} \right] \right\} - 2 \log(\epsilon). \end{aligned}$$

It is now convenient to remember that

$$\bar{P} [\log(1 - \lambda \psi)] = \frac{1}{2} \log \left(\frac{1 - \lambda}{1 + \lambda} \right) r'(\theta, \tilde{\theta}) + \frac{1}{2} \log(1 - \lambda^2) m'(\theta, \tilde{\theta}).$$

We thus can write the previous inequality as

$$\begin{aligned} & -N \log \left\{ \left(1 - \lambda \rho [R'(\cdot, \tilde{\theta})] \right) \left(1 + \beta \rho [R'(\cdot, \tilde{\theta})] \right) \right\} \\ & \leq \frac{N}{2} \log \left(\frac{1 + \lambda}{1 - \lambda} \right) \rho [r'(\cdot, \tilde{\theta})] - \frac{N}{2} \log(1 - \lambda^2) \rho [m'(\cdot, \tilde{\theta})] + \mathcal{K}(\rho, \pi) \\ & \quad + \log \left\{ \pi \left[\exp \left\{ -\frac{N}{2} \log \left(\frac{1 + \beta}{1 - \beta} \right) r'(\cdot, \tilde{\theta}) \right. \right. \right. \\ & \quad \left. \left. \left. - \frac{N}{2} \log(1 - \beta^2) m'(\cdot, \tilde{\theta}) \right\} \right] \right\} - 2 \log(\epsilon). \end{aligned}$$

Let us assume now that $\tilde{\theta} \in \arg \min_{\Theta_1} R$. Let us introduce $\hat{\theta} \in \arg \min_{\Theta} r$. Decomposing $r'(\theta, \tilde{\theta}) = r'(\theta, \hat{\theta}) + r'(\hat{\theta}, \tilde{\theta})$ and considering that

$$m'(\theta, \tilde{\theta}) \leq m'(\theta, \hat{\theta}) + m'(\hat{\theta}, \tilde{\theta}),$$

we see that with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,

$$\begin{aligned}
& -N \log \left\{ \left(1 - \lambda \rho[R'(\cdot, \tilde{\theta})]\right) \left(1 + \beta \rho[R'(\cdot, \tilde{\theta})]\right) \right\} \\
& \leq \frac{N}{2} \log \left(\frac{1+\lambda}{1-\lambda} \right) \rho[r'(\cdot, \hat{\theta})] - \frac{N}{2} \log(1 - \lambda^2) \rho[m'(\cdot, \hat{\theta})] + \mathcal{K}(\rho, \pi) \\
& + \log \left\{ \pi \left[\exp \left\{ -\frac{N}{2} \log \left(\frac{1+\beta}{1-\beta} \right) [r'(\cdot, \hat{\theta})] - \frac{N}{2} \log(1 - \beta^2) m'(\cdot, \hat{\theta}) \right\} \right] \right\} \\
& \quad + \frac{N}{2} \log \left[\frac{(1+\lambda)(1-\beta)}{(1-\lambda)(1+\beta)} \right] [r(\hat{\theta}) - r(\tilde{\theta})] \\
& \quad - \frac{N}{2} \log [(1 - \lambda^2)(1 - \beta^2)] m'(\hat{\theta}, \tilde{\theta}) - 2 \log(\epsilon).
\end{aligned}$$

Let us now define for simplicity the posterior $\nu : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$ by the identity

$$\frac{d\nu}{d\pi}(\theta) = \frac{\exp \left\{ -\frac{N}{2} \log \left(\frac{1+\lambda}{1-\lambda} \right) r'(\theta, \hat{\theta}) + \frac{N}{2} \log(1 - \lambda^2) m'(\theta, \hat{\theta}) \right\}}{\pi \left[\exp \left\{ -\frac{N}{2} \log \left(\frac{1+\lambda}{1-\lambda} \right) r'(\cdot, \hat{\theta}) + \frac{N}{2} \log(1 - \lambda^2) m'(\cdot, \hat{\theta}) \right\} \right]}.$$

Let us also introduce the random bound

$$\begin{aligned}
B &= \frac{1}{N} \log \left\{ \nu \left[\exp \left[\frac{N}{2} \log \left[\frac{(1+\lambda)(1-\beta)}{(1-\lambda)(1+\beta)} \right] r'(\cdot, \hat{\theta}) \right. \right. \right. \\
& \quad \left. \left. \left. - \frac{N}{2} \log [(1 - \lambda^2)(1 - \beta^2)] m'(\cdot, \hat{\theta}) \right] \right] \right\} \\
& + \sup_{\theta \in \Theta_1} \frac{1}{2} \log \left[\frac{(1-\lambda)(1+\beta)}{(1+\lambda)(1-\beta)} \right] r'(\theta, \hat{\theta}) \\
& \quad - \frac{1}{2} \log [(1 - \lambda^2)(1 - \beta^2)] m'(\theta, \hat{\theta}) - \frac{2}{N} \log(\epsilon).
\end{aligned}$$

THEOREM 1.35. *Using the above notations, for any real constants $0 \leq \beta < \lambda < 1$, for any prior distribution $\pi \in \mathcal{M}_+^1(\Theta)$, for any subset $\Theta_1 \subset \Theta$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$-\log \left\{ \left(1 - \lambda [\rho(R) - \inf_{\Theta_1} R]\right) \left(1 + \beta [\rho(R) - \inf_{\Theta_1} R]\right) \right\} \leq \frac{\mathcal{K}(\rho, \nu)}{N} + B.$$

Therefore,

$$\begin{aligned}
\rho(R) - \inf_{\Theta_1} R &\leq \frac{\lambda - \beta}{2\lambda\beta} \left(\sqrt{1 + 4 \frac{\lambda\beta}{(\lambda - \beta)^2} \left[1 - \exp\left(-B - \frac{\mathcal{K}(\rho, \nu)}{N}\right)\right]} - 1 \right) \\
&\leq \frac{1}{\lambda - \beta} \left(B + \frac{\mathcal{K}(\rho, \nu)}{N} \right).
\end{aligned}$$

Let us define the posterior $\hat{\nu}$ by the identity

$$\frac{d\hat{\nu}}{d\pi}(\theta) = \frac{\exp\left[-\frac{N}{2} \log\left(\frac{1+\beta}{1-\beta}\right) r'(\theta, \hat{\theta}) - \frac{N}{2} \log(1 - \beta^2) m'(\theta, \hat{\theta})\right]}{\pi \left\{ \exp\left[-\frac{N}{2} \log\left(\frac{1+\beta}{1-\beta}\right) r'(\cdot, \hat{\theta}) - \frac{N}{2} \log(1 - \beta^2) m'(\cdot, \hat{\theta})\right] \right\}}.$$

It is useful to remark that

$$\begin{aligned}
&\frac{1}{N} \log \left\{ \nu \left[\exp\left[\frac{N}{2} \log\left(\frac{(1+\lambda)(1-\beta)}{(1-\lambda)(1+\beta)}\right) r'(\cdot, \hat{\theta}) - \frac{N}{2} \log[(1-\lambda^2)(1-\beta^2)] m'(\cdot, \hat{\theta})\right] \right] \right\} \\
&\leq \hat{\nu} \left\{ \frac{1}{2} \log\left(\frac{(1+\lambda)(1-\beta)}{(1-\lambda)(1+\beta)}\right) r'(\cdot, \hat{\theta}) - \frac{1}{2} \log[(1-\lambda^2)(1-\beta^2)] m'(\cdot, \hat{\theta}) \right\}.
\end{aligned}$$

Let us introduce as previously $\bar{\varphi}(x) = \sup_{\theta \in \Theta} m'(\theta, \hat{\theta}) - x r'(\theta, \hat{\theta})$, $x \in \mathbb{R}_+$. Let us moreover consider $\tilde{\varphi}(x) = \sup_{\theta \in \Theta_1} m'(\theta, \hat{\theta}) - x r'(\theta, \hat{\theta})$, $x \in \mathbb{R}_+$. These functions can be used to produce a result which is slightly weaker, but maybe easier to read and understand. Indeed, coming back a little while, we see that, for any $x \in \mathbb{R}_+$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution ρ ,

$$\begin{aligned}
& - N \log \left\{ \left(1 - \lambda \rho [R'(\cdot, \tilde{\theta})] \right) \left(1 + \beta \rho [R'(\cdot, \tilde{\theta})] \right) \right\} \\
& \leq \frac{N}{2} \log \left[\frac{(1+\lambda)}{(1-\lambda)(1-\lambda^2)^x} \right] \rho [r'(\cdot, \hat{\theta})] \\
& \quad - \frac{N}{2} \log [(1-\lambda^2)(1-\beta^2)] \bar{\varphi}(x) + \mathcal{K}(\rho, \pi) \\
& \quad + \log \left\{ \pi \left[\exp \left\{ -\frac{N}{2} \log \left[\frac{(1+\beta)}{(1-\beta)(1-\beta^2)^x} \right] r'(\cdot, \hat{\theta}) \right\} \right] \right\} \\
& \quad - \frac{N}{2} \log [(1-\lambda^2)(1-\beta^2)] \tilde{\varphi} \left(\frac{\log \left[\frac{(1+\lambda)(1-\beta)}{(1-\lambda)(1+\beta)} \right]}{-\log [(1-\lambda^2)(1-\beta^2)]} \right) \\
& \hspace{20em} - 2 \log(\epsilon) \\
& = \int_{\frac{N}{2} \log \left[\frac{(1+\beta)}{(1-\beta)(1-\beta^2)^x} \right]}^{\frac{N}{2} \log \left[\frac{(1+\lambda)}{(1-\lambda)(1-\lambda^2)^x} \right]} \pi_{\exp(-\alpha r)} [r'(\cdot, \hat{\theta})] d\alpha \\
& \quad + \mathcal{K}(\rho, \pi_{\exp\{-\frac{N}{2} \log \left[\frac{(1+\lambda)}{(1-\lambda)(1-\lambda^2)^x} \right] r\}}) - 2 \log(\epsilon) \\
& \quad - \frac{N}{2} \log [(1-\lambda^2)(1-\beta^2)] \left[\bar{\varphi}(x) + \tilde{\varphi} \left(\frac{\log \left[\frac{(1+\lambda)(1-\beta)}{(1-\lambda)(1+\beta)} \right]}{-\log [(1-\lambda^2)(1-\beta^2)]} \right) \right].
\end{aligned}$$

THEOREM 1.36. *With the previous notations, for any real constants $0 \leq \beta < \lambda < 1$, for any positive real constant x , for any prior probability distribution $\pi \in \mathcal{M}_+^1(\Theta)$, for any subset $\Theta_1 \subset \Theta$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, putting*

$$\begin{aligned}
B(\rho) &= \frac{1}{N(\lambda - \beta)} \int_{\frac{N}{2} \log \left[\frac{(1+\beta)}{(1-\beta)(1-\beta^2)^x} \right]}^{\frac{N}{2} \log \left[\frac{(1+\lambda)}{(1-\lambda)(1-\lambda^2)^x} \right]} \pi_{\exp(-\alpha r)} [r'(\cdot, \hat{\theta})] d\alpha \\
& \quad + \frac{\mathcal{K}(\rho, \pi_{\exp\{-\frac{N}{2} \log \left[\frac{(1+\lambda)}{(1-\lambda)(1-\lambda^2)^x} \right] r\}}) - 2 \log(\epsilon)}{N(\lambda - \beta)} \\
& \quad - \frac{1}{2(\lambda - \beta)} \log [(1-\lambda^2)(1-\beta^2)] \left[\bar{\varphi}(x) + \tilde{\varphi} \left(\frac{\log \left[\frac{(1+\lambda)(1-\beta)}{(1-\lambda)(1+\beta)} \right]}{-\log [(1-\lambda^2)(1-\beta^2)]} \right) \right] \\
& \leq \frac{1}{N(\lambda - \beta)} d_e \log \left(\frac{\log \left[\frac{(1+\lambda)}{(1-\lambda)(1-\lambda^2)^x} \right]}{\log \left[\frac{(1+\beta)}{(1-\beta)(1-\beta^2)^x} \right]} \right)
\end{aligned}$$

$$\begin{aligned}
& \mathcal{K}(\rho, \pi_{\exp\{-\frac{N}{2} \log[\frac{(1+\lambda)}{(1-\lambda)(1-\lambda^2)^x}]r\}}) - 2 \log(\epsilon) \\
& + \frac{1}{N(\lambda - \beta)} \\
& - \frac{1}{2(\lambda - \beta)} \log[(1 - \lambda^2)(1 - \beta^2)] \left[\bar{\varphi}(x) + \tilde{\varphi} \left(\frac{\log \left[\frac{(1+\lambda)(1-\beta)}{(1-\lambda)(1+\beta)} \right]}{-\log[(1 - \lambda^2)(1 - \beta^2)]} \right) \right],
\end{aligned}$$

the following bounds hold true:

$$\begin{aligned}
\rho(R) - \inf_{\Theta_1} R \\
\leq \frac{\lambda - \beta}{2\lambda\beta} \left(\sqrt{1 + \frac{4\lambda\beta}{(\lambda - \beta)^2} \left\{ 1 - \exp[-(\lambda - \beta)B(\rho)] \right\}} - 1 \right) \\
\leq B(\rho).
\end{aligned}$$

Let us remark that this alternative way of handling relative deviation bounds made it possible to carry on with non linear bounds up to the final result. (For instance, if $\lambda = 0.5$, $\beta = 0.2$ and $B(\rho) = 0.1$, the non linear bound gives $\rho(R) - \inf_{\Theta_1} R \leq 0.096$.)

1.5. BOUNDS RELATIVE TO A GIBBS DISTRIBUTION. The empirical bounds of the previous section involve taking suprema in $\theta \in \Theta$, and replacing the *margin function* φ by some empirical counter parts $\bar{\varphi}$ or $\tilde{\varphi}$, which may prove unsafe when using very complex classification models. Moreover, they are not easy to analyze with PAC-Bayesian tools. To remedy these weaknesses, we are going now to propose another type of relative bounds. We will first explain how to compare the expected error rate $\rho(R)$ of any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$ with $\pi_{\exp(-\beta R)}(R)$, the expected risk of a Gibbs prior distribution. We will then show how to analyze the behaviour of this bound. This will provide an estimator proven to reach adaptively the best possible asymptotic behaviour of the error rate under Mammen and Tsybakov margin assumptions and parametric complexity assumptions.

Then, we will provide an empirical bound for the Kullback divergence $\mathcal{K}(\rho, \pi_{\exp(-\beta R)})$ of a posterior distribution with respect to a Gibbs prior, making use of relative deviation inequalities.

To tackle the question of model selection, we will estimate the relative performance of one posterior distribution with respect to another, which is useful when the two posteriors are supported by different models.

Eventually, we will propose a more integrated approach to model selection, showing how to build a two step localization strategy, in which the

performance of the posterior distribution to be analyzed is compared with some *two step* Gibbs prior.

1.5.1. Comparing a posterior distribution with a Gibbs prior. Similarly to Theorem 1.26 we can prove that for any prior distribution $\tilde{\pi} \in \mathcal{M}_+^1(\Theta)$,

$$\mathbb{P} \left\{ \tilde{\pi} \otimes \tilde{\pi} \left\{ \exp \left[-N \log(1 - \lambda R') - \frac{N}{2} \log \left(\frac{1 + \lambda}{1 - \lambda} \right) r' + \frac{N}{2} \log(1 - \lambda^2) m' \right] \right\} \right\} \leq 1. \quad (1.23)$$

Replacing $\tilde{\pi}$ with $\pi_{\exp(-\beta R)}$ and considering the posterior distribution $\rho \otimes \pi_{\exp(-\beta R)}$, provides a starting point in the comparison of ρ with $\pi_{\exp(-\beta R)}$; we can indeed state with \mathbb{P} probability at least $1 - \epsilon$ that

$$\begin{aligned} & -N \log \left\{ 1 - \lambda \left[\rho(R) - \pi_{\exp(-\beta R)}(R) \right] \right\} \\ & \leq \frac{N}{2} \log \left(\frac{1 + \lambda}{1 - \lambda} \right) \left[\rho(r) - \pi_{\exp(-\beta R)}(r) \right] \\ & \quad - \frac{N}{2} \log(1 - \lambda^2) \rho \otimes \pi_{\exp(-\beta R)}(m') \\ & \quad + \mathcal{K}[\rho, \pi_{\exp(-\beta R)}] - \log(\epsilon). \end{aligned} \quad (1.24)$$

Using the parameter $\gamma = \frac{N}{2} \log \left(\frac{1 + \lambda}{1 - \lambda} \right)$, so that $\lambda = \tanh \left(\frac{\gamma}{N} \right)$ and $-\frac{N}{2} \log(1 - \lambda^2) = N \log [\cosh(\frac{\gamma}{N})]$, and noticing that

$$\begin{aligned} \mathcal{K}[\rho, \pi_{\exp(-\beta R)}] &= \beta \left[\rho(R) - \pi_{\exp(-\beta R)}(R) \right] \\ & \quad + \mathcal{K}(\rho, \pi) - \mathcal{K}[\pi_{\exp(-\beta R)}, \pi], \end{aligned} \quad (1.25)$$

makes a step further in the proper handling of the entropy term:

$$\begin{aligned} & -N \log \left\{ 1 - \tanh \left(\frac{\gamma}{N} \right) \left[\rho(R) - \pi_{\exp(-\beta R)}(R) \right] \right\} - \beta \left[\rho(R) - \pi_{\exp(-\beta R)}(R) \right] \\ & \leq \gamma \left[\rho(r) - \pi_{\exp(-\beta R)}(r) \right] + N \log [\cosh(\frac{\gamma}{N})] \rho \otimes \pi_{\exp(-\beta R)}(m') \\ & \quad + \mathcal{K}(\rho, \pi) - \mathcal{K}[\pi_{\exp(-\beta R)}, \pi] - \log(\epsilon). \end{aligned} \quad (1.26)$$

We can then decompose in the right-hand side $\gamma \left[\rho(r) - \pi_{\exp(-\beta R)}(r) \right]$ into $(\gamma - \lambda) \left[\rho(r) - \pi_{\exp(-\beta R)}(r) \right] + \lambda \left[\rho(r) - \pi_{\exp(-\beta R)}(r) \right]$ and use the fact that

$$\begin{aligned}
& \lambda[\rho(r) - \pi_{\exp(-\beta R)}(r)] + N \log[\cosh(\frac{\gamma}{N})] \rho \otimes \pi_{\exp(-\beta R)}(m') \\
& \quad + \mathcal{K}(\rho, \pi) - \mathcal{K}[\pi_{\exp(-\beta R)}, \pi] \\
& \leq \lambda\rho(r) + \mathcal{K}(\rho, \pi) + \log\left\{\pi\left[\exp\{-\lambda r + N \log[\cosh(\frac{\gamma}{N})] \rho(m')\}\right]\right\} \\
& = \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}] + \log\left\{\pi_{\exp(-\lambda r)}\left[\exp\{N \log[\cosh(\frac{\gamma}{N})] \rho(m')\}\right]\right\},
\end{aligned}$$

to get rid of the appearance of the unobserved Gibbs prior $\pi_{\exp(-\beta R)}$ in most places of the right-hand side of our inequality, leading to

THEOREM 1.37. *For any real constants β and γ , with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, for any real constant λ ,*

$$\begin{aligned}
& [N \tanh(\frac{\gamma}{N}) - \beta] [\rho(R) - \pi_{\exp(-\beta R)}(R)] \\
& \leq -N \log\left\{1 - \tanh(\frac{\gamma}{N}) [\rho(R) - \pi_{\exp(-\beta R)}(R)]\right\} \\
& \quad - \beta [\rho(R) - \pi_{\exp(-\beta R)}(R)] \\
& \leq (\gamma - \lambda) [\rho(r) - \pi_{\exp(-\beta R)}(r)] + \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}] \\
& \quad + \log\left\{\pi_{\exp(-\lambda r)}\left[\exp\{N \log[\cosh(\frac{\gamma}{N})] \rho(m')\}\right]\right\} - \log(\epsilon) \\
& = \mathcal{K}[\rho, \pi_{\exp(-\gamma r)}] \\
& \quad + \log\left\{\pi_{\exp(-\gamma r)}\left[\exp\{(\gamma - \lambda)r + N \log[\cosh(\frac{\gamma}{N})] \rho(m')\}\right]\right\} \\
& \quad - (\gamma - \lambda)\pi_{\exp(-\beta R)}(r) - \log(\epsilon).
\end{aligned}$$

We would like to have a fully empirical upper bound even in the case when $\lambda \neq \gamma$. This can be done by using the theorem twice. We will need a lemma

LEMMA 1.38 *For any probability distribution $\pi \in \mathcal{M}_+^1(\Theta)$, for any bounded measurable functions $g, h : \Theta \rightarrow \mathbb{R}$,*

$$\pi_{\exp(-g)}(g) - \pi_{\exp(-h)}(g) \leq \pi_{\exp(-g)}(h) - \pi_{\exp(-h)}(h).$$

PROOF. Let us notice that

$$\begin{aligned}
0 & \leq \mathcal{K}(\pi_{\exp(-g)}, \pi_{\exp(-h)}) = \pi_{\exp(-g)}(h) + \log\{\pi[\exp(-h)]\} + \mathcal{K}(\pi_{\exp(-g)}, \pi) \\
& = \pi_{\exp(-g)}(h) - \pi_{\exp(-h)}(h) - \mathcal{K}(\pi_{\exp(-h)}, \pi) + \mathcal{K}(\pi_{\exp(-g)}, \pi) \\
& = \pi_{\exp(-g)}(h) - \pi_{\exp(-h)}(h) - \mathcal{K}(\pi_{\exp(-h)}, \pi) - \pi_{\exp(-g)}(g) - \log\{\pi[\exp(-g)]\}.
\end{aligned}$$

Moreover

$$-\log\{\pi[\exp(-g)]\} \leq \pi_{\exp(-h)}(g) + \mathcal{K}(\pi_{\exp(-h)}, \pi),$$

which achieves the proof. \square

For any positive real constants β and λ , we can then apply Theorem 1.37 to $\rho = \pi_{\exp(-\lambda r)}$, and use the inequality

$$\frac{\lambda}{\beta} [\pi_{\exp(-\lambda r)}(r) - \pi_{\exp(-\beta R)}(r)] \leq \pi_{\exp(-\lambda r)}(R) - \pi_{\exp(-\beta R)}(R) \quad (1.27)$$

provided by the previous lemma. We thus obtain with \mathbb{P} probability at least $1 - \epsilon$

$$\begin{aligned} & -N \log \left\{ 1 - \tanh\left(\frac{\gamma}{N}\right) \frac{\lambda}{\beta} [\pi_{\exp(-\lambda r)}(r) - \pi_{\exp(-\beta R)}(r)] \right\} \\ & \quad - \gamma [\pi_{\exp(-\lambda r)}(r) - \pi_{\exp(-\beta R)}(r)] \\ & \leq \log \left\{ \pi_{\exp(-\lambda r)} \left[\exp \left\{ N \log \left[\cosh\left(\frac{\gamma}{N}\right) \right] \pi_{\exp(-\lambda r)}(m') \right\} \right] \right\} - \log(\epsilon). \end{aligned}$$

Let us introduce the convex function

$$F_{\gamma, \alpha}(x) = -N \log \left[1 - \tanh\left(\frac{\gamma}{N}\right) x \right] - \alpha x \geq [N \tanh\left(\frac{\gamma}{N}\right) - \alpha] x.$$

With \mathbb{P} probability at least $1 - \epsilon$,

$$\begin{aligned} -\pi_{\exp(-\beta R)}(r) & \leq \inf_{\lambda \in \mathbb{R}_+^*} \left\{ -\pi_{\exp(-\lambda r)}(r) \right. \\ & \quad \left. + \frac{\beta}{\lambda} F_{\gamma, \frac{\beta\gamma}{\lambda}}^{-1} \left[\log \left\{ \pi_{\exp(-\lambda r)} \left[\exp \left\{ N \log \left[\cosh\left(\frac{\gamma}{N}\right) \right] \pi_{\exp(-\lambda r)}(m') \right\} \right] \right\} \right. \right. \\ & \quad \left. \left. - \log(\epsilon) \right] \right\}. \end{aligned}$$

Since Theorem 1.37 holds uniformly for any posterior distribution ρ , we can apply it again to some arbitrary posterior distribution ρ . We can moreover make the result uniform in β and γ by considering some atomic measure $\nu \in \mathcal{M}_+^1(\mathbb{R})$ on the real line and using a union bound. This leads to

THEOREM 1.39. *For any atomic probability distribution on the positive real line $\nu \in \mathcal{M}_+^1(\mathbb{R}_+)$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, for any positive real constants β and γ ,*

$$\begin{aligned}
& [N \tanh(\frac{\gamma}{N}) - \beta] [\rho(R) - \pi_{\exp(-\beta R)}(R)] \\
& \leq F_{\gamma, \beta} [\rho(R) - \pi_{\exp(-\beta R)}(R)] \leq B(\rho, \beta, \gamma), \text{ where} \\
B(\rho, \beta, \gamma) = & \inf_{\substack{\lambda_1 \in \mathbb{R}_+, \lambda_1 \leq \gamma \\ \lambda_2 \in \mathbb{R}, \lambda_2 > \frac{\beta \gamma}{N} \tanh(\frac{\gamma}{N})^{-1}}} \left\{ \mathcal{K}[\rho, \pi_{\exp(-\lambda_1 r)}] \right. \\
& + (\gamma - \lambda_1) [\rho(r) - \pi_{\exp(-\lambda_2 r)}(r)] \\
& + \log \left\{ \pi_{\exp(-\lambda_1 r)} \left[\exp \{ N \log [\cosh(\frac{\gamma}{N})] \rho(m') \} \right] \right\} - \log [\epsilon \nu(\beta) \nu(\gamma)] \\
& + (\gamma - \lambda_1) \frac{\beta}{\lambda_2} F_{\gamma, \frac{\beta \gamma}{\lambda_2}}^{-1} \left[\log \left\{ \right. \right. \\
& \left. \left. \pi_{\exp(-\lambda_2 r)} \left[\exp \{ N \log [\cosh(\frac{\gamma}{N})] \pi_{\exp(-\lambda_2 r)}(m') \} \right] \right\} \right. \\
& \left. \left. - \log [\epsilon \nu(\beta) \nu(\gamma)] \right] \right\} \\
\leq & \inf_{\substack{\lambda_1 \in \mathbb{R}_+, \lambda_1 \leq \gamma \\ \lambda_2 \in \mathbb{R}, \lambda_2 > \frac{\beta \gamma}{N} \tanh(\frac{\gamma}{N})^{-1}}} \left\{ \mathcal{K}[\rho, \pi_{\exp(-\lambda_1 r)}] \right. \\
& + (\gamma - \lambda_1) [\rho(r) - \pi_{\exp(-\lambda_2 r)}(r)] \\
& + \log \left\{ \pi_{\exp(-\lambda_1 r)} \left[\exp \{ N \log [\cosh(\frac{\gamma}{N})] \rho(m') \} \right] \right\} \\
& + \frac{\beta}{\lambda_2} \frac{(1 - \frac{\lambda_1}{\gamma})}{[\frac{N}{\gamma} \tanh(\frac{\gamma}{N}) - \frac{\beta}{\lambda_2}]} \log \left\{ \pi_{\exp(-\lambda_2 r)} \left[\right. \right. \\
& \left. \left. \exp \{ N \log [\cosh(\frac{\gamma}{N})] \pi_{\exp(-\lambda_2 r)}(m') \} \right] \right\} \\
& \left. - \left\{ 1 + \frac{\beta}{\lambda_2} \frac{(1 - \frac{\lambda_1}{\gamma})}{[\frac{N}{\gamma} \tanh(\frac{\gamma}{N}) - \frac{\beta}{\lambda_2}]} \right\} \log [\epsilon \nu(\beta) \nu(\gamma)] \right\},
\end{aligned}$$

where we have written for short $\nu(\beta)$ and $\nu(\gamma)$ instead of $\nu(\{\beta\})$ and $\nu(\{\gamma\})$.

Let us notice that $B(\rho, \beta, \gamma) = +\infty$ when $\nu(\beta) = 0$ or $\nu(\gamma) = 0$, the uniformity in β and γ of the theorem therefore necessarily bears on a countable number of values of these parameters. We can typically choose for ν distributions such as the one used in Theorem 1.11 on page 21: namely we can put for some positive real ratio $\alpha > 1$

$$\nu(\alpha^k) = \frac{1}{(k+1)(k+2)}, \quad k \in \mathbb{N},$$

or alternatively, since we are interested in values of the parameters less than

N , we can prefer

$$\nu(\alpha^k) = \frac{\log(\alpha)}{\log(\alpha N)}, \quad 0 \leq k < \frac{\log(N)}{\log(\alpha)}.$$

We can also use such a coding distribution on dyadic numbers as the one defined by equation (1.6) on page 23.

1.5.2. The effective temperature of a posterior distribution. Using the parametric approximation $\pi_{\exp(-\alpha r)}(r) - \inf_{\Theta} r \simeq \frac{d_e}{\alpha}$, we get as an order of magnitude

$$\begin{aligned} B(\pi_{\exp(-\lambda_1 r)}, \beta, \gamma) &\lesssim -(\gamma - \lambda_1) d_e [\lambda_2^{-1} - \lambda_1^{-1}] \\ &+ 2d_e \log \frac{\lambda_1}{\lambda_1 - N \log[\cosh(\frac{\gamma}{N})] x} \\ &+ 2 \frac{\beta}{\lambda_2} \frac{(1 - \frac{\lambda_1}{\gamma})}{[\frac{N}{\gamma} \tanh(\frac{\gamma}{N}) - \frac{\beta}{\lambda_2}]} d_e \log \left(\frac{\lambda_2}{\lambda_2 - N \log[\cosh(\frac{\gamma}{N})] x} \right) \\ &+ 2N \log[\cosh(\frac{\gamma}{N})] \left[1 + \frac{\beta}{\lambda_2} \frac{(1 - \frac{\lambda_1}{\gamma})}{[\frac{N}{\gamma} \tanh(\frac{\gamma}{N}) - \frac{\beta}{\lambda_2}]} \right] \tilde{\varphi}(x) \\ &- \left\{ 1 + \frac{\beta}{\lambda_2} \frac{(1 - \frac{\lambda_1}{\gamma})}{[\frac{N}{\gamma} \tanh(\frac{\gamma}{N}) - \frac{\beta}{\lambda_2}]} \right\} \log[\nu(\beta)\nu(\gamma)\epsilon]. \end{aligned}$$

Therefore, if the empirical dimension d_e stays bounded when N increases, we are going to obtain a negative upper bound for any values of the constants $\lambda_1 > \lambda_2 > \beta$, as soon as γ and $\frac{N}{\gamma}$ are chosen to be large enough. This ability to obtain negative values for the bound $B(\pi_{\exp(-\lambda_1 r)}, \gamma, \beta)$, and more generally $B(\rho, \gamma, \beta)$, leads the way to introducing the new concept of the *effective temperature* of an estimator.

DEFINITION 1.1 For any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$ we define the *effective temperature* $T(\rho) \in \mathbb{R} \cup \{-\infty, +\infty\}$ of ρ by the equation

$$\rho(R) = \pi_{\exp(-\frac{R}{T(\rho)})}(R).$$

Note that $\beta \mapsto \pi_{\exp(-\beta R)}(R) : \mathbb{R} \cup \{-\infty, +\infty\} \rightarrow (0, 1)$ is continuous and strictly decreasing from $\text{ess sup}_{\pi} R$ to $\text{ess inf}_{\pi} R$ (as soon as these two bounds do not coincide). This shows that the effective temperature $T(\rho)$ is a well defined random variable.

Theorem 1.39 provides a bound for $T(\rho)$, indeed:

PROPOSITION 1.40. *Let*

$$\widehat{\beta}(\rho) = \sup\left\{\beta \in \mathbb{R}; \inf_{\gamma, N \tanh(\frac{\gamma}{N}) > \beta} B(\rho, \beta, \gamma) \leq 0\right\},$$

where $B(\rho, \beta, \gamma)$ is as in Theorem 1.39. Then with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, $T(\rho) \leq \widehat{\beta}(\rho)^{-1}$, or equivalently $\rho(R) \leq \pi_{\exp[-\widehat{\beta}(\rho)R]}(R)$.

This notion of *effective temperature* of a (randomized) estimator ρ is interesting for two reasons:

- the difference $\rho(R) - \pi_{\exp(-\beta R)}(R)$ can be estimated with a better accuracy than $\rho(R)$ itself, due to the use of relative deviation inequalities, leading to convergence rates up to $1/N$ in favourable situations, even when $\inf_{\Theta} R$ is not close to zero;
- and of course $\pi_{\exp(-\beta R)}(R)$ is a decreasing function of β , thus being able to estimate $\rho(R) - \pi_{\exp(-\beta R)}(R)$ with some given accuracy, means being able to discriminate between values of $\rho(R)$ with the same accuracy, although doing so through the parametrization $\beta \mapsto \pi_{\exp(-\beta R)}(R)$, which cannot be observed nor estimated with the same precision!

1.5.3. Analysis of an empirical bound for the effective temperature. We are now going to launch into a mathematically rigorous analysis of the bound $B(\pi_{\exp(-\lambda_1 r)}, \beta, \gamma)$ provided by Theorem 1.39, to show that $\inf_{\rho \in \mathcal{M}_+^1(\Theta)} \pi_{\exp[-\widehat{\beta}(\rho)R]}(R)$ converges indeed to $\inf_{\Theta} R$ at some unimprovable rates in favourable situations.

It is more convenient for this purpose to use deviation inequalities involving M' rather than m' . It is straightforward to extend Theorem 1.25 on page 43 to

THEOREM 1.41. *For any real constants β and γ , for any prior distribution $\mu \in \mathcal{M}_+^1(\Theta)$, with \mathbb{P} probability at least $1 - \eta$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\gamma \rho \otimes \pi_{\exp(-\beta R)} \left[\Psi_{\frac{\gamma}{N}}(R', M') \right] \leq \gamma \rho \otimes \pi_{\exp(-\beta R)}(r') + \mathcal{K}(\rho, \mu) - \log(\eta).$$

In order to transform the left-hand side into a linear expression and in the same time to localize this theorem, let us choose μ defined by its density

$$\begin{aligned} \frac{d\mu}{d\pi}(\theta_1) = C^{-1} \exp & \left[-\beta R(\theta_1) \right. \\ & - \gamma \int_{\Theta} \left\{ \Psi_{\frac{\gamma}{N}} [R'(\theta_1, \theta_2), M'(\theta_1, \theta_2)] \right. \\ & \left. \left. - \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) R'(\theta_1, \theta_2) \right\} \pi_{\exp(-\beta R)}(d\theta_2) \right], \end{aligned}$$

where C is such that $\mu(\Theta) = 1$. We get

$$\begin{aligned} \mathcal{K}(\rho, \mu) &= \beta \rho(R) + \gamma \rho \otimes \pi_{\exp(-\beta R)} \left[\Psi_{\frac{\gamma}{N}}(R', M') - \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) R' \right] + \mathcal{K}(\rho, \pi) \\ &+ \log \left\{ \int_{\Theta} \exp \left[-\beta R(\theta_1) \right. \right. \\ &\quad \left. \left. - \gamma \int_{\Theta} \left\{ \Psi_{\frac{\gamma}{N}} [R'(\theta_1, \theta_2), M'(\theta_1, \theta_2)] \right. \right. \right. \\ &\quad \left. \left. \left. - \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) R'(\theta_1, \theta_2) \right\} \pi_{\exp(-\beta R)}(d\theta_2) \right] \pi(d\theta_1) \right\} \\ &= \beta [\rho(R) - \pi_{\exp(-\beta R)}(R)] \\ &\quad + \gamma \rho \otimes \pi_{\exp(-\beta R)} \left[\Psi_{\frac{\gamma}{N}}(R', M') - \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) R' \right] \\ &\quad + \mathcal{K}(\rho, \pi) - \mathcal{K}(\pi_{\exp(-\beta R)}, \pi) \\ &+ \log \left\{ \int_{\Theta} \exp \left[-\gamma \int_{\Theta} \left\{ \Psi_{\frac{\gamma}{N}} [R'(\theta_1, \theta_2), M'(\theta_1, \theta_2)] \right. \right. \right. \\ &\quad \left. \left. \left. - \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) R'(\theta_1, \theta_2) \right\} \pi_{\exp(-\beta R)}(d\theta_2) \right] \pi_{\exp(-\beta R)}(d\theta_1) \right\}. \end{aligned}$$

Thus with \mathbb{P} probability at least $1 - \eta$,

$$\begin{aligned} & [N \sinh\left(\frac{\gamma}{N}\right) - \beta] [\rho(R) - \pi_{\exp(-\beta R)}(R)] \\ & \leq \gamma [\rho(r) - \pi_{\exp(-\beta R)}(r)] + \mathcal{K}(\rho, \pi) - \mathcal{K}(\pi_{\exp(-\beta R)}, \pi) - \log(\eta) + C(\beta, \gamma) \end{aligned}$$

$$\begin{aligned} \text{where } C(\beta, \gamma) &= \log \left\{ \int_{\Theta} \exp \left[-\gamma \int_{\Theta} \left\{ \Psi_{\frac{\gamma}{N}} [R'(\theta_1, \theta_2), M'(\theta_1, \theta_2)] \right. \right. \right. \\ & \quad \left. \left. \left. - \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) R'(\theta_1, \theta_2) \right\} \pi_{\exp(-\beta R)}(d\theta_2) \right] \pi_{\exp(-\beta R)}(d\theta_1) \right\}. \quad (1.28) \end{aligned}$$

Remarking that

$$\mathcal{K}[\rho, \pi_{\exp(-\beta R)}] = \beta [\rho(R) - \pi_{\exp(-\beta R)}(R)] + \mathcal{K}(\rho, \pi) - \mathcal{K}(\pi_{\exp(-\beta R)}, \pi),$$

we deduce from the previous inequality

THEOREM 1.42. *For any real constants β and γ , with \mathbb{P} probability at least $1 - \eta$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$N \sinh\left(\frac{\gamma}{N}\right) [\rho(R) - \pi_{\exp(-\beta R)}(R)] \leq \gamma [\rho(r) - \pi_{\exp(-\beta R)}(r)] + \mathcal{K}[\rho, \pi_{\exp(-\beta R)}] - \log(\eta) + C(\beta, \gamma).$$

We can also go into a slightly different direction, starting back again from equation (1.28) on page 67 and remarking that for any real constant λ ,

$$\begin{aligned} \lambda [\rho(r) - \pi_{\exp(-\beta R)}(r)] + \mathcal{K}(\rho, \pi) - \mathcal{K}(\pi_{\exp(-\beta R)}, \pi) \\ \leq \lambda \rho(r) + \mathcal{K}(\rho, \pi) + \log\{\pi[\exp(-\lambda r)]\} = \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}]. \end{aligned}$$

This leads to

THEOREM 1.43. *For any real constants β and γ , with \mathbb{P} probability at least $1 - \eta$, for any real constant λ ,*

$$\begin{aligned} [N \sinh\left(\frac{\gamma}{N}\right) - \beta] [\rho(R) - \pi_{\exp(-\beta R)}(R)] \\ \leq (\gamma - \lambda) [\rho(r) - \pi_{\exp(-\beta R)}(r)] + \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}] - \log(\eta) + C(\beta, \gamma), \end{aligned}$$

where the definition of $C(\beta, \gamma)$ is given by equation (1.28) on page 67.

We can now use this inequality in the case when $\rho = \pi_{\exp(-\lambda r)}$ and combine it with inequality (1.27) on page 63 to obtain

THEOREM 1.44 *For any real constants β and γ , with \mathbb{P} probability at least $1 - \eta$, for any real constant λ ,*

$$\left[\frac{N\lambda}{\beta} \sinh\left(\frac{\gamma}{N}\right) - \gamma\right] [\pi_{\exp(-\lambda r)}(r) - \pi_{\exp(-\beta R)}(r)] \leq C(\beta, \gamma) - \log(\eta).$$

We deduce from this theorem

PROPOSITION 1.45 *For any real positive constants β_1, β_2 and γ , with \mathbb{P} probability at least $1 - \eta$, for any real constants λ_1 and λ_2 , such that $\lambda_2 < \beta_2 \frac{\gamma}{N} \sinh\left(\frac{\gamma}{N}\right)^{-1}$ and $\lambda_1 > \beta_1 \frac{\gamma}{N} \sinh\left(\frac{\gamma}{N}\right)^{-1}$,*

$$\begin{aligned} \pi_{\exp(-\lambda_1 r)}(r) - \pi_{\exp(-\lambda_2 r)}(r) \leq \pi_{\exp(-\beta_1 R)}(r) - \pi_{\exp(-\beta_2 R)}(r) \\ + \frac{C(\beta_1, \gamma) + \log(2/\eta)}{\frac{N\lambda_1}{\beta_1} \sinh\left(\frac{\gamma}{N}\right) - \gamma} + \frac{C(\beta_2, \gamma) + \log(2/\eta)}{\gamma - \frac{N\lambda_2}{\beta_2} \sinh\left(\frac{\gamma}{N}\right)}. \end{aligned}$$

Moreover, $\pi_{\exp(-\beta_1 R)}$ and $\pi_{\exp(-\beta_2 R)}$ being prior distributions, with \mathbb{P} probability at least $1 - \eta$,

$$\begin{aligned} \gamma [\pi_{\exp(-\beta_1 R)}(r) - \pi_{\exp(-\beta_2 R)}(r)] \\ \leq \gamma \pi_{\exp(-\beta_1 R)} \otimes \pi_{\exp(-\beta_2 R)} [\Psi_{-\frac{\gamma}{N}}(R', M')] - \log(\eta). \end{aligned}$$

Hence

PROPOSITION 1.46 *For any positive real constants β_1, β_2 and γ , with \mathbb{P} probability at least $1 - \eta$, for any positive real constants λ_1 and λ_2 such that $\lambda_2 < \beta_2 \frac{\gamma}{N} \sinh(\frac{\gamma}{N})^{-1}$ and $\lambda_1 > \beta_1 \frac{\gamma}{N} \sinh(\frac{\gamma}{N})^{-1}$,*

$$\begin{aligned} \pi_{\exp(-\lambda_1 r)}(r) - \pi_{\exp(-\lambda_2 r)}(r) \\ \leq \pi_{\exp(-\beta_1 R)} \otimes \pi_{\exp(-\beta_2 R)} [\Psi_{-\frac{\gamma}{N}}(R', M')] \\ + \frac{\log(\frac{3}{\eta})}{\gamma} + \frac{C(\beta_1, \gamma) + \log(\frac{3}{\eta})}{\frac{N\lambda_1}{\beta_1} \sinh(\frac{\gamma}{N}) - \gamma} + \frac{C(\beta_2, \gamma) + \log(\frac{3}{\eta})}{\gamma - \frac{N\lambda_2}{\beta_2} \sinh(\frac{\gamma}{N})}. \end{aligned}$$

In order to achieve the analysis of the bound $B(\pi_{\exp(-\lambda_1 r)}, \beta, \gamma)$ given by Theorem 1.39, there remains now to bound quantities of the general form

$$\begin{aligned} \log \left\{ \pi_{\exp(-\lambda r)} \left[\exp \left\{ N \log \left[\cosh \left(\frac{\gamma}{N} \right) \right] \pi_{\exp(-\lambda r)}(m') \right\} \right] \right\} \\ = \sup_{\rho \in \mathcal{M}_+^1(\Theta)} N \log \left[\cosh \left(\frac{\gamma}{N} \right) \right] \rho \otimes \pi_{\exp(-\lambda)}(m') - \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}]. \end{aligned}$$

Let us consider the prior distribution $\mu \in \mathcal{M}_+^1(\Theta \times \Theta)$ on couples of parameters defined by its density

$$\frac{d\mu}{d(\pi \otimes \pi)}(\theta_1, \theta_2) = C^{-1} \exp \left\{ -\beta R(\theta_1) - \beta R(\theta_2) + \alpha \Phi_{-\frac{\alpha}{N}}[M'(\theta_1, \theta_2)] \right\},$$

where the normalizing constant C is such that $\mu(\Theta \times \Theta) = 1$. Since for fixed values of the parameters θ and $\theta' \in \Theta$, $m'(\theta, \theta')$, like $r(\theta)$, is a sum of independent Bernoulli random variables, we can easily adapt the proof of Theorem 1.4 on page 11, to establish that with \mathbb{P} probability at least $1 - \eta$, for any posterior distribution ρ and any real constant λ ,

$$\begin{aligned} \alpha \rho \otimes \pi_{\exp(-\lambda r)}(m') \leq \alpha \rho \otimes \pi_{\exp(-\lambda r)} \left[\Phi_{-\frac{\alpha}{N}}(M') \right] \\ + \mathcal{K}(\rho \otimes \pi_{\exp(-\lambda r)}, \mu) - \log(\eta) \end{aligned}$$

$$= \mathcal{K}[\rho, \pi_{\exp(-\beta R)}] + \mathcal{K}[\pi_{\exp(-\lambda r)}, \pi_{\exp(-\beta R)}] \\ + \log \left\{ \pi_{\exp(-\beta R)} \otimes \pi_{\exp(-\beta R)} \left[\exp\left(\alpha \Phi_{-\frac{\alpha}{N}} \circ M'\right) \right] \right\} - \log(\eta).$$

Thus for any real constant β and any positive real constants α and γ , with \mathbb{P} probability at least $1 - \eta$, for any real constant λ ,

$$\log \left\{ \pi_{\exp(-\lambda r)} \left[\exp\left\{ N \log \left[\cosh\left(\frac{\gamma}{N}\right) \right] \pi_{\exp(-\lambda r)}(m') \right\} \right] \right\} \\ \leq \sup_{\rho \in \mathcal{M}_+^1(\Theta)} \left(\frac{N}{\alpha} \log \left[\cosh\left(\frac{\gamma}{N}\right) \right] \left\{ \mathcal{K}[\rho, \pi_{\exp(-\beta R)}] + \mathcal{K}[\pi_{\exp(-\lambda r)}, \pi_{\exp(-\beta R)}] \right. \right. \\ \left. \left. + \log \left\{ \pi_{\exp(-\beta R)} \otimes \pi_{\exp(-\beta R)} \left[\exp\left(\alpha \Phi_{-\frac{\alpha}{N}} \circ M'\right) \right] \right\} \right. \right. \\ \left. \left. - \log(\eta) \right\} - \mathcal{K}[\rho, \pi_{\exp(-\lambda r)}] \right). \quad (1.29)$$

To conclude, we need some suitable upper bound for the entropy $\mathcal{K}[\rho, \pi_{\exp(-\beta R)}]$. This question can be handled in the following way: using Theorem 1.42 on page 68, we see that for any positive real constants γ and β , with \mathbb{P} probability at least $1 - \eta$, for any posterior distribution ρ ,

$$\mathcal{K}[\rho, \pi_{\exp(-\beta R)}] = \beta [\rho(R) - \pi_{\exp(-\beta R)}(R)] + \mathcal{K}(\rho, \pi) - \mathcal{K}(\pi_{\exp(-\beta R)}, \pi) \\ \leq \frac{\beta}{N \sinh\left(\frac{\gamma}{N}\right)} \left[\gamma [\rho(r) - \pi_{\exp(-\beta R)}(r)] \right. \\ \left. + \mathcal{K}[\rho, \pi_{\exp(-\beta R)}] - \log(\eta) + C(\beta, \gamma) \right] \\ + \mathcal{K}(\rho, \pi) - \mathcal{K}(\pi_{\exp(-\beta R)}, \pi) \\ \leq \mathcal{K}\left[\rho, \pi_{\exp\left(-\frac{\beta\gamma}{N \sinh\left(\frac{\gamma}{N}\right)} r\right)}\right] \\ + \frac{\beta}{N \sinh\left(\frac{\gamma}{N}\right)} \left\{ \mathcal{K}[\rho, \pi_{\exp(-\beta R)}] + C(\beta, \gamma) - \log(\eta) \right\}.$$

In other words,

THEOREM 1.47. *For any positive real constants β and γ such that $\beta < N \sinh\left(\frac{\gamma}{N}\right)$, with \mathbb{P} probability at least $1 - \eta$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\mathcal{K}[\rho, \pi_{\exp(-\beta R)}] \leq \frac{\mathcal{K}\left[\rho, \pi_{\exp\left[-\beta \frac{\gamma}{N} \sinh\left(\frac{\gamma}{N}\right)^{-1} r\right]}\right]}{1 - \frac{\beta}{N \sinh\left(\frac{\gamma}{N}\right)}} + \frac{C(\beta, \gamma) - \log(\eta)}{\frac{N \sinh\left(\frac{\gamma}{N}\right)}{\beta} - 1}.$$

Choosing in equation (1.29) on page 70 $\alpha = \frac{N \log[\cosh(\frac{\gamma}{N})]}{1 - \frac{\beta}{N \sinh(\frac{\gamma}{N})}}$ and $\beta = \lambda \frac{N}{\gamma} \sinh(\frac{\gamma}{N})$, so that $\alpha = \frac{N \log[\cosh(\frac{\gamma}{N})]}{1 - \frac{\lambda}{\gamma}}$, we obtain with \mathbb{P} probability at least $1 - \eta$,

$$\begin{aligned} & \log \left\{ \pi_{\exp(-\lambda r)} \left[\exp \left\{ N \log \left[\cosh \left(\frac{\gamma}{N} \right) \right] \pi_{\exp(-\lambda r)}(m') \right\} \right] \right\} \\ & \leq \frac{2\lambda}{\gamma} \left[C(\beta, \gamma) + \log \left(\frac{2}{\eta} \right) \right] \\ & \quad + \left(1 - \frac{\lambda}{\gamma} \right) \left[\log \left\{ \pi_{\exp(-\beta R)} \otimes \pi_{\exp(-\beta R)} \left[\exp \left(\alpha \Phi_{-\frac{\alpha}{N}} \circ M' \right) \right] \right\} \right. \\ & \quad \left. + \log \left(\frac{2}{\eta} \right) \right]. \end{aligned}$$

This proves

PROPOSITION 1.48. *For any positive real constants $\lambda < \gamma$, with \mathbb{P} probability at least $1 - \eta$,*

$$\begin{aligned} & \log \left\{ \pi_{\exp(-\lambda r)} \left[\exp \left\{ N \log \left[\cosh \left(\frac{\gamma}{N} \right) \right] \pi_{\exp(-\lambda r)}(m') \right\} \right] \right\} \\ & \leq \frac{2\lambda}{\gamma} \left[C \left(\frac{N\lambda}{\gamma} \sinh \left(\frac{\gamma}{N} \right), \gamma \right) + \log \left(\frac{2}{\eta} \right) \right] \\ & \quad + \left(1 - \frac{\lambda}{\gamma} \right) \log \left\{ \pi_{\exp \left[-\frac{N\lambda}{\gamma} \sinh \left(\frac{\gamma}{N} \right) R \right]}^{\otimes 2} \left[\right. \right. \\ & \quad \left. \left. \exp \left(\frac{N \log \left[\cosh \left(\frac{\gamma}{N} \right) \right]}{1 - \frac{\lambda}{\gamma}} \Phi_{-\frac{\log \left[\cosh \left(\frac{\gamma}{N} \right) \right]}{1 - \frac{\lambda}{\gamma}}} \circ M' \right) \right] \right\} \\ & \quad + \left(1 - \frac{\lambda}{\gamma} \right) \log \left(\frac{2}{\eta} \right). \end{aligned}$$

We are now ready to analyse the bound $B(\pi_{\exp(-\lambda_1 r)}, \beta, \gamma)$ of Theorem 1.39 on page 63.

THEOREM 1.49. *For any positive real constants $\lambda_1, \lambda_2, \beta_1, \beta_2, \beta$ and γ , such that*

$$\begin{aligned} \lambda_1 &< \gamma, & \beta_1 &< \frac{N\lambda_1}{\gamma} \sinh \left(\frac{\gamma}{N} \right), \\ \lambda_2 &< \gamma, & \beta_2 &> \frac{N\lambda_2}{\gamma} \sinh \left(\frac{\gamma}{N} \right), \\ & & \beta &< \frac{N\lambda_2}{\gamma} \tanh \left(\frac{\gamma}{N} \right), \end{aligned}$$

with \mathbb{P} probability $1 - \eta$, the bound $B(\pi_{\exp(-\lambda_1 r)}, \beta, \gamma)$ of Theorem 1.39 on page 63 satisfies

$$\begin{aligned}
& B(\pi_{\exp(-\lambda_1 r)}, \beta, \gamma) \\
& \leq (\gamma - \lambda_1) \left\{ \pi_{\exp(-\beta_1 R)} \otimes \pi_{\exp(-\beta_2 R)} [\Psi_{-\frac{\gamma}{N}}(R', M')] + \frac{\log(\frac{\gamma}{\eta})}{\gamma} \right. \\
& \quad \left. + \frac{C(\beta_1, \gamma) + \log(\frac{\gamma}{\eta})}{\frac{N\lambda_1}{\beta_1} \sinh(\frac{\gamma}{N}) - \gamma} + \frac{C(\beta_2, \gamma) + \log(\frac{\gamma}{\eta})}{\gamma - \frac{N\lambda_2}{\beta_2} \sinh(\frac{\gamma}{N})} \right\} \\
& \quad + \frac{2\lambda_1}{\gamma} \left[C\left(\frac{N\lambda_1}{\gamma} \sinh(\frac{\gamma}{N}), \gamma\right) + \log\left(\frac{\gamma}{\eta}\right) \right] \\
& + \left(1 - \frac{\lambda_1}{\gamma}\right) \log \left\{ \pi_{\exp[-\frac{N\lambda_1}{\gamma} \sinh(\frac{\gamma}{N}) R]}^{\otimes 2} \left[\right. \right. \\
& \quad \left. \left. \exp\left(\frac{N \log[\cosh(\frac{\gamma}{N})]}{1 - \frac{\lambda_1}{\gamma}} \Phi_{-\frac{\log[\cosh(\frac{\gamma}{N})]}{1 - \frac{\lambda_1}{\gamma}}} \circ M'\right) \right] \right\} \\
& \quad + \left(1 - \frac{\lambda_1}{\gamma}\right) \log\left(\frac{\gamma}{\eta}\right) - \log[\nu(\{\beta\})\nu(\{\gamma\})\epsilon] \\
& + (\gamma - \lambda_1) \frac{\beta}{\lambda_2} F_{\gamma, \frac{\beta\gamma}{\lambda_2}}^{-1} \left\{ \frac{2\lambda_2}{\gamma} \left[C\left(\frac{N\lambda_2}{\gamma} \sinh(\frac{\gamma}{N}), \gamma\right) + \log\left(\frac{\gamma}{\eta}\right) \right] \right. \\
& \quad \left. + \left(1 - \frac{\lambda_2}{\gamma}\right) \log \left\{ \pi_{\exp[-\frac{N\lambda_2}{\gamma} \sinh(\frac{\gamma}{N}) R]}^{\otimes 2} \left[\right. \right. \right. \\
& \quad \left. \left. \exp\left(\frac{N \log[\cosh(\frac{\gamma}{N})]}{1 - \frac{\lambda_2}{\gamma}} \Phi_{-\frac{\log[\cosh(\frac{\gamma}{N})]}{1 - \frac{\lambda_2}{\gamma}}} \circ M'\right) \right] \right\} \right\} \\
& \quad \left. + \left(1 - \frac{\lambda_2}{\gamma}\right) \log\left(\frac{\gamma}{\eta}\right) - \log[\nu(\{\beta\})\nu(\{\gamma\})\epsilon] \right\},
\end{aligned}$$

where the function $C(\beta, \gamma)$ is defined by equation (1.28) on page 67.

1.5.4. Adaptation to parametric and margin assumptions. To help understand the previous theorem, it may be useful to give linear upper-bounds to the factors appearing in the right-hand side of the previous inequality. Introducing $\tilde{\theta}$ such that $R(\tilde{\theta}) = \inf_{\Theta} R$ (assuming that such a parameter exists) and remembering that

$$\begin{aligned}
\Psi_{-a}(p, m) & \leq a^{-1} \sinh(a)p + 2a^{-1} \sinh\left(\frac{a}{2}\right)^2 m, & a \in \mathbb{R}_+, \\
\Phi_{-a}(p) & \leq a^{-1} [\exp(a) - 1] p, & a \in \mathbb{R}_+, \\
\Psi_a(p, m) & \geq a^{-1} \sinh(a)p - 2a^{-1} \sinh\left(\frac{a}{2}\right)^2 m, & a \in \mathbb{R}_+,
\end{aligned}$$

$$\begin{aligned}
M'(\theta_1, \theta_2) &\leq M'(\theta_1, \tilde{\theta}) + M'(\theta_2, \tilde{\theta}), & \theta_1, \theta_2 \in \Theta, \\
M'(\theta_1, \tilde{\theta}) &\leq xR'(\theta_1, \tilde{\theta}) + \varphi(x), & x \in \mathbb{R}_+, \theta_1 \in \Theta,
\end{aligned}$$

(the last inequality being rather a consequence of the definition of φ than a property of M'), we easily see that

$$\begin{aligned}
&\pi_{\exp(-\beta_1 R)} \otimes \pi_{\exp(-\beta_2 R)} [\Psi_{-\frac{\gamma}{N}}(R', M')] \\
&\leq \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) [\pi_{\exp(-\beta_1 R)}(R) - \pi_{\exp(-\beta_2 R)}(R)] \\
&\quad + \frac{2N}{\gamma} \sinh\left(\frac{\gamma}{2N}\right)^2 \pi_{\exp(-\beta_1 R)} \otimes \pi_{\exp(-\beta_2 R)}(M') \\
&\leq \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) [\pi_{\exp(-\beta_1 R)}(R) - \pi_{\exp(-\beta_2 R)}(R)] \\
&\quad + \frac{2xN}{\gamma} \sinh\left(\frac{\gamma}{2N}\right)^2 \left\{ \pi_{\exp(-\beta_1 R)} [R'(\cdot, \tilde{\theta})] + \pi_{\exp(-\beta_2 R)} [R'(\cdot, \tilde{\theta})] \right\} \\
&\quad + \frac{4N}{\gamma} \sinh\left(\frac{\gamma}{2N}\right)^2 \varphi(x).
\end{aligned}$$

$$\begin{aligned}
C(\beta, \gamma) &\leq \log \left\{ \pi_{\exp(-\beta R)} \left\{ \exp \left[2N \sinh\left(\frac{\gamma}{2N}\right)^2 \pi_{\exp(-\beta R)}(M') \right] \right\} \right\} \\
&\leq \log \left\{ \pi_{\exp(-\beta R)} \left\{ \exp \left[2N \sinh\left(\frac{\gamma}{2N}\right)^2 M'(\cdot, \tilde{\theta}) \right] \right\} \right\} \\
&\quad + 2N \sinh\left(\frac{\gamma}{2N}\right)^2 \pi_{\exp(-\beta R)} [M'(\cdot, \tilde{\theta})] \\
&\leq \log \left\{ \pi_{\exp(-\beta R)} \left\{ \exp \left[2xN \sinh\left(\frac{\gamma}{2N}\right)^2 R'(\cdot, \tilde{\theta}) \right] \right\} \right\} \\
&\quad + 2xN \sinh\left(\frac{\gamma}{2N}\right)^2 \pi_{\exp(-\beta R)} [R'(\cdot, \tilde{\theta})] + 4N \sinh\left(\frac{\gamma}{2N}\right)^2 \varphi(x) \\
&= \int_{\beta-2xN \sinh(\frac{\gamma}{2N})^2}^{\beta} \pi_{\exp(-\alpha R)} [R'(\cdot, \tilde{\theta})] d\alpha \\
&\quad + 2xN \sinh\left(\frac{\gamma}{2N}\right)^2 \pi_{\exp(-\beta R)} [R'(\cdot, \tilde{\theta})] + 4N \sinh\left(\frac{\gamma}{2N}\right)^2 \varphi(x) \\
&\leq 4xN \sinh\left(\frac{\gamma}{2N}\right)^2 \pi_{\exp[-(\beta-2xN \sinh(\frac{\gamma}{2N})^2)R]} [R'(\cdot, \tilde{\theta})] \\
&\quad + 4N \sinh\left(\frac{\gamma}{2N}\right)^2 \varphi(x).
\end{aligned}$$

$$\begin{aligned}
&\log \left\{ \pi_{\exp(-\beta R)}^{\otimes 2} \left[\exp \left(N\alpha \Phi_{-\alpha} \circ M' \right) \right] \right\} \\
&\leq 2 \log \left\{ \pi_{\exp(-\beta R)} \left[\exp \left(N [\exp(\alpha) - 1] M'(\cdot, \tilde{\theta}) \right) \right] \right\}
\end{aligned}$$

$$\begin{aligned} &\leq 2xN[\exp(\alpha) - 1]\pi_{\exp[-(\beta-xN[\exp(\alpha)-1])R]}[R'(\cdot, \tilde{\theta})] \\ &\quad + 2xN[\exp(\alpha) - 1]\varphi(x). \end{aligned}$$

Let us push further the investigation under the parametric assumption that for some positive real constant d

$$\lim_{\beta \rightarrow +\infty} \beta \pi_{\exp(-\beta R)}[R'(\cdot, \tilde{\theta})] = d, \quad (1.30)$$

This assumption will for instance hold true with $d = \frac{n}{2}$ when $R : \Theta \rightarrow (0, 1)$ is a smooth function defined on a compact subset Θ of \mathbb{R}^n that reaches its minimum value on a finite number of non degenerate (i.e. with a positive definite Hessian) interior points of Θ , and π is absolutely continuous with respect to the Lebesgue measure on Θ and has a smooth density.

In case of assumption (1.30), if we restrict to sufficiently large values of the constants β , β_1 , β_2 , λ_1 , λ_2 and γ (the smaller of which being as a rule β as we will see), we can use the fact that for some (small) positive constant δ , and some (large) positive constant A ,

$$\frac{d}{\alpha}(1 - \delta) \leq \pi_{\exp(-\alpha R)}[R'(\cdot, \tilde{\theta})] \leq \frac{d}{\alpha}(1 + \delta), \quad \alpha \geq A. \quad (1.31)$$

Under this assumption,

$$\begin{aligned} &\pi_{\exp(-\beta_1 R)} \otimes \pi_{\exp(-\beta_2 R)} [\Psi_{-\frac{\gamma}{N}}(R', M')] \\ &\quad \leq \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) \left[\frac{d}{\beta_1}(1 + \delta) - \frac{d}{\beta_2}(1 - \delta) \right] \\ &\quad \quad + \frac{2xN}{\gamma} \sinh\left(\frac{\gamma}{2N}\right)^2 (1 + \delta) \left[\frac{d}{\beta_1} + \frac{d}{\beta_2} \right] + \frac{4N}{\gamma} \sinh\left(\frac{\gamma}{2N}\right)^2 \varphi(x). \\ C(\beta, \gamma) &\leq d(1 + \delta) \log\left(\frac{\beta}{\beta - 2xN \sinh\left(\frac{\gamma}{2N}\right)^2}\right) \\ &\quad + 2xN \sinh\left(\frac{\gamma}{2N}\right)^2 \frac{(1 + \delta)d}{\beta} + 4N \sinh\left(\frac{\gamma}{2N}\right)^2 \varphi(x). \\ \log\left\{ \pi_{\exp(-\beta R)}^{\otimes 2} \left[\exp\left(N\alpha\Phi_{-\alpha} \circ M'\right) \right] \right\} \\ &\leq 2xN[\exp(\alpha) - 1] \frac{d(1 + \delta)}{\beta - xN[\exp(\alpha) - 1]} + 2N[\exp(\alpha) - 1]\varphi(x). \end{aligned}$$

Thus with \mathbb{P} probability at least $1 - \eta$,

$$\begin{aligned} B(\pi_{\exp(-\lambda_1 r)}, \beta, \gamma) &\leq -(\gamma - \lambda_1) \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) \frac{d}{\beta_2} (1 - \delta) \\ &\quad + (\gamma - \lambda_1) \left\{ \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) \frac{(1 + \delta)d}{\beta_1} \right. \\ &\quad \quad \left. + \frac{2xN}{\gamma} \sinh\left(\frac{\gamma}{2N}\right)^2 (1 + \delta) \left[\frac{d}{\beta_1} + \frac{d}{\beta_2} \right] + \frac{4N}{\gamma} \sinh\left(\frac{\gamma}{2N}\right)^2 \varphi(x) + \frac{\log\left(\frac{7}{\eta}\right)}{\gamma} \right\} \end{aligned}$$

$$\begin{aligned}
& + \frac{4xN \sinh(\frac{\gamma}{2N})^2 \frac{(1+\delta)d}{\beta_1 - 2xN \sinh(\frac{\gamma}{2N})^2} + 4N \sinh(\frac{\gamma}{2N})^2 \varphi(x) + \log(\frac{\gamma}{\eta})}{\frac{N\lambda_1}{\beta_1} \sinh(\frac{\gamma}{N}) - \gamma} \\
& + \frac{4xN \sinh(\frac{\gamma}{2N})^2 \frac{(1+\delta)d}{\beta_2 - 2xN \sinh(\frac{\gamma}{2N})^2} + 4N \sinh(\frac{\gamma}{2N})^2 \varphi(x) + \log(\frac{\gamma}{\eta})}{\gamma - \frac{N\lambda_2}{\beta_2} \sinh(\frac{\gamma}{N})} \Big\} \\
& + \frac{2\lambda_1}{\gamma} \left\{ 4xN \sinh(\frac{\gamma}{2N})^2 \frac{(1+\delta)d}{\frac{N\lambda_1}{\gamma} \sinh(\frac{\gamma}{N}) - 2xN \sinh(\frac{\gamma}{2N})^2} \right. \\
& \qquad \qquad \qquad \left. + 4N \sinh(\frac{\gamma}{2N})^2 \varphi(x) + \log(\frac{\gamma}{\eta}) \right\} \\
& + \left(1 - \frac{\lambda_1}{\gamma}\right) \left\{ 2d(1+\delta) \left(\frac{\lambda_1 \sinh(\frac{\gamma}{N})}{x\gamma \left[\exp\left(\frac{\log[\cosh(\frac{\gamma}{N})]}{1 - \frac{\lambda_1}{\gamma}}\right) - 1 \right]} - 1 \right)^{-1} \right. \\
& \qquad \qquad \qquad \left. + 2N \left[\exp\left(\frac{\log[\cosh(\frac{\gamma}{N})]}{1 - \frac{\lambda_1}{\gamma}}\right) - 1 \right] \varphi(x) \right\} \\
& \qquad \qquad \qquad + \left(1 - \frac{\lambda_1}{\gamma}\right) \log\left(\frac{\gamma}{\eta}\right) - \log[\nu(\{\beta\})\nu(\{\gamma\})\epsilon] \\
& + \frac{1 - \frac{\lambda_1}{\gamma}}{\frac{N\lambda_2}{\beta\gamma} \tanh(\frac{\gamma}{N}) - 1} \left\{ \frac{2\lambda_2}{\gamma} \left\{ 4xN \sinh(\frac{\gamma}{2N})^2 \frac{(1+\delta)d}{\frac{N\lambda_2}{\gamma} \sinh(\frac{\gamma}{N}) - 2xN \sinh(\frac{\gamma}{2N})^2} \right. \right. \\
& \qquad \qquad \qquad \left. \left. + 4N \sinh(\frac{\gamma}{2N})^2 \varphi(x) + \log(\frac{\gamma}{\eta}) \right\} \right. \\
& + \left(1 - \frac{\lambda_2}{\gamma}\right) \left[2d(1+\delta) \left(\frac{\lambda_2 \sinh(\frac{\gamma}{N})}{x\gamma \left[\exp\left(\frac{\log[\cosh(\frac{\gamma}{N})]}{1 - \frac{\lambda_2}{\gamma}}\right) - 1 \right]} - 1 \right)^{-1} \right. \\
& \qquad \qquad \qquad \left. + 2N \left[\exp\left(\frac{\log[\cosh(\frac{\gamma}{N})]}{1 - \frac{\lambda_2}{\gamma}}\right) - 1 \right] \varphi(x) \right] \\
& \qquad \qquad \qquad \left. + \left(1 - \frac{\lambda_2}{\gamma}\right) \log\left(\frac{\gamma}{\eta}\right) - \log[\nu(\beta)\nu(\gamma)\epsilon] \right\}.
\end{aligned}$$

Now let us choose for simplicity $\beta_2 = 2\lambda_2 = 4\beta$, $\beta_1 = \lambda_1/2 = \gamma/4$, and let us introduce the notations

$$\begin{aligned}
C_1 &= \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right), \\
C_2 &= \frac{N}{\gamma} \tanh\left(\frac{\gamma}{N}\right),
\end{aligned}$$

$$C_3 = \frac{N^2}{\gamma^2} [\exp(\frac{\gamma^2}{N^2}) - 1]$$

and
$$C_4 = \frac{2N^2(1 - \frac{2\beta}{\gamma})}{\gamma^2} \left[\exp\left(\frac{\gamma^2}{2N^2(1 - \frac{2\beta}{\gamma})}\right) - 1 \right],$$

to obtain

$$\begin{aligned} B(\pi_{\exp(-\lambda_1 r)}, \beta, \gamma) &\leq -\frac{C_1 \gamma}{8\beta} (1 - \delta) d \\ &+ \frac{C_1 \gamma}{2} \left\{ \frac{4(1+\delta)d}{\gamma} + x \frac{\gamma}{2N} (1 + \delta) \left[\frac{4d}{\gamma} + \frac{d}{4\beta} \right] + \frac{\gamma}{N} \varphi(x) \right\} + \frac{1}{2} \log\left(\frac{7}{\eta}\right) \\ &+ \frac{1}{2C_1 - 1} \left[(1 + \delta) d \left(\frac{N}{2x C_1 \gamma} - 1 \right)^{-1} + C_1 \frac{\gamma^2}{2N} \varphi(x) + \frac{1}{2} \log\left(\frac{7}{\eta}\right) \right] \\ &+ \frac{1}{2 - C_1} \left[2(1 + \delta) d \left(\frac{8N\beta}{x C_1 \gamma^2} - 1 \right)^{-1} + C_1 \frac{\gamma^2}{N} \varphi(x) + \log\left(\frac{7}{\eta}\right) \right] \\ &\quad + \frac{2x\gamma(1 + \delta)d}{N - x\gamma} + C_1 \frac{\gamma^2}{N} \varphi(x) + \log\left(\frac{7}{\eta}\right) \\ &+ d(1 + \delta) \frac{x\gamma}{N} \left(\frac{C_1}{2C_3} - \frac{x\gamma}{N} \right)^{-1} + \frac{\gamma^2}{N} C_3 \varphi(x) + \frac{\log\left(\frac{7}{\eta}\right)}{2} - \log[\nu(\beta)\nu(\gamma)\epsilon] \\ &+ (4C_2 - 2)^{-1} \left\{ \frac{4\beta}{\gamma} \left\{ x \frac{\gamma^2}{N} C_1 (1 + \delta) d \left(2\beta C_1 - x C_1 \frac{\gamma^2}{2N} \right)^{-1} \right. \right. \\ &\quad \left. \left. + \frac{\gamma^2}{N} \varphi(x) + \log\left(\frac{7}{\eta}\right) \right\} \right. \\ &+ \left(1 - \frac{2\beta}{\gamma} \right) \left\{ 2d(1 + \delta) \frac{x\gamma}{N} \left[\frac{4\beta C_1}{\gamma C_4} \left(1 - \frac{2\beta}{\gamma} \right) - \frac{x\gamma}{N} \right]^{-1} \right. \\ &\quad \left. + \frac{\gamma^2}{N(1 - \frac{2\beta}{\gamma})} C_4 \varphi(x) \right\} \\ &\quad \left. + \left(1 - \frac{2\beta}{\gamma} \right) \log\left(\frac{7}{\eta}\right) - \log[\nu(\beta)\nu(\gamma)\epsilon] \right\}. \end{aligned}$$

This simplifies to

$$\begin{aligned} B(\pi_{\exp(-\lambda_1 r)}, \beta, \gamma) &\leq -\frac{C_1}{8} (1 - \delta) d \frac{\gamma}{\beta} \\ &+ 2C_1 (1 + \delta) d + \log\left(\frac{7}{\eta}\right) \left[2 + \frac{3C_1}{(4C_1 - 2)(2 - C_1)} + \frac{1 + \frac{2\beta}{\gamma}}{4C_2 - 2} \right] \end{aligned}$$

$$\begin{aligned}
& - \left(1 + \frac{1}{4C_2-2}\right) \log[\nu(\beta)\nu(\gamma)\epsilon] \\
& + \frac{(1+\delta)d\gamma}{N} \left\{ C_1 + \frac{1}{2C_1-1} \left(\frac{1}{2C_1} - \frac{\gamma x}{N}\right)^{-1} \right. \\
& \quad \left. + 2\left(1 - \frac{\gamma x}{N}\right)^{-1} + \left(\frac{C_1}{2C_3} - \frac{\gamma x}{N}\right)^{-1} + \frac{4C_1\beta}{\gamma(4C_2-2)} \right\} \\
& + \frac{(1+\delta)d\gamma^2}{N\beta} \left\{ \frac{C_1}{16} + \frac{2}{2-C_1} \left(\frac{8}{C_1} - \frac{x\gamma^2}{N\beta}\right)^{-1} \right. \\
& \quad \left. + \left(1 - \frac{2\beta}{\gamma}\right) \frac{1}{2C_2-1} \left[\frac{4C_1}{C_4} \left(1 - \frac{2\beta}{\gamma}\right) - \frac{\gamma^2 x}{\beta N}\right]^{-1} \right\} \\
& + \frac{\gamma^2}{N} \varphi(x) \left\{ \frac{3C_1}{2} + \frac{C_1}{4C_1-2} + \frac{C_1}{2-C_1} + C_3 + \frac{4\beta}{\gamma(4C_2-2)} + \frac{C_4}{4C_2-2} \right\}.
\end{aligned}$$

This shows that there exist universal positive real constants $A_1, A_2, B_1, B_2, B_3,$ and B_4 such that as soon as $\frac{\gamma^{\max\{x,1\}}}{N} \leq A_1 \frac{\beta}{\gamma} \leq A_2,$

$$\begin{aligned}
B(\pi_{\exp(-\lambda_1 r)}, \beta, \gamma) & \leq -B_1(1-\delta)d\frac{\gamma}{\beta} + B_2(1+\delta)d \\
& \quad - B_3 \log[\nu(\beta)\nu(\gamma)\epsilon\eta] + B_4 \frac{\gamma^2}{N} \varphi(x).
\end{aligned}$$

Thus $\pi_{\exp(-\lambda_1 r)}(R) \leq \pi_{\exp(-\beta R)}(R) \leq \inf_{\Theta} R + \frac{(1+\delta)d}{\beta}$ as soon as moreover

$$\frac{\beta}{\gamma} \leq \frac{B_1}{B_2 \frac{(1+\delta)}{(1-\delta)} + \frac{B_4 \frac{\gamma^2}{N} \varphi(x) - B_3 \log[\nu(\beta)\nu(\gamma)\epsilon\eta]}{(1-\delta)d}}.$$

Choosing some real ratio $\alpha > 1,$ we can now make the above result uniform for any

$$\beta, \gamma \in \Lambda_\alpha \stackrel{\text{def}}{=} \left\{ \alpha^k; k \in \mathbb{N}, 0 \leq k < \frac{\log(N)}{\log(\alpha)} \right\}, \quad (1.32)$$

by substituting $\nu(\beta)$ and $\nu(\gamma)$ with $\frac{\log(\alpha)}{\log(\alpha N)}$ and $-\log(\eta)$ with $-\log(\eta) + 2 \log \left[\frac{\log(\alpha N)}{\log(\alpha)} \right].$

Taking moreover for simplicity $\eta = \epsilon,$ let us summarize the type of result we got by

THEOREM 1.50. *There exist positive real universal constants A, B_1, B_2, B_3 and B_4 such that for any positive real constants $\alpha > 1, d$ and $\delta,$ for any prior distribution $\pi \in \mathcal{M}_+^1(\Theta),$ with \mathbb{P} probability at least $1 - \epsilon,$ for any $\beta, \gamma \in \Lambda_\alpha$ (where Λ_α is defined by equation (1.32) above) such that*

$$\sup_{\beta' \in \mathbb{R}, \beta' \geq \beta} \left| \frac{\beta'}{d} \left[\pi_{\exp(-\beta' R)}(R) - \inf_{\Theta} R \right] - 1 \right| \leq \delta$$

and such that also for some positive real parameter x

$$\frac{\gamma \max\{x, 1\}}{N} \leq \frac{A\beta}{\gamma} \quad \text{and} \quad \frac{\beta}{\gamma} \leq \frac{B_1}{B_2 \frac{(1+\delta)}{(1-\delta)} + \frac{B_4 \frac{\gamma^2}{N} \varphi(x) - 2B_3 \log(\epsilon) + 4B_3 \log\left[\frac{\log(N)}{\log(\alpha)}\right]}{(1-\delta)d}},$$

the bound $B(\pi_{\exp(-\frac{\gamma}{2}r)}, \beta, \gamma)$ given by Theorem 1.39 on page 63 in the case where we have chosen ν to be the uniform probability measure on Λ_α , satisfies $B(\pi_{\exp(-\frac{\gamma}{2}r)}, \beta, \gamma) \leq 0$, proving that $\hat{\beta}(\pi_{\exp(-\frac{\gamma}{2}r)}) \geq \beta$ and therefore that

$$\pi_{\exp(-\frac{\gamma}{2}r)}(R) \leq \pi_{\exp(-\beta R)}(R) \leq \inf_{\Theta} R + \frac{(1+\delta)d}{\beta}.$$

What is important in this result is that we do not only bound $\pi_{\exp(-\frac{\gamma}{2}r)}(R)$, but also $B(\pi_{\exp(-\frac{\gamma}{2}r)}, \beta, \gamma)$, and that we do it uniformly on a grid of values of β and γ , showing that we can indeed set the constants β and γ adaptively using the empirical bound $B(\pi_{\exp(-\frac{\gamma}{2}r)}, \beta, \gamma)$.

Let us see what we get under the margin assumption (1.21) (see page 47). When $\kappa = 1$, $\varphi(c^{-1}) \leq 0$, leading to

COROLLARY 1.51. *Assuming that the margin assumption 1.21 (on page 47) is satisfied for $\kappa = 1$, that $R : \Theta \rightarrow (0, 1)$ is independent of N (which is the case for instance when $\mathbb{P} = P^{\otimes N}$), and is such that*

$$\lim_{\beta' \rightarrow +\infty} \beta' [\pi_{\exp(-\beta' R)}(R) - \inf_{\Theta} R] = d,$$

there are universal positive real constants B_5 and B_6 and $N_1 \in \mathbb{N}$ such that for any $N \geq N_1$, with \mathbb{P} probability at least $1 - \epsilon$

$$\pi_{\exp(-\hat{\gamma} \frac{r}{2})}(R) \leq \inf_{\Theta} R + \frac{B_5 d}{cN} \left[1 + \frac{B_6}{d} \log\left(\frac{\log(N)}{\epsilon}\right) \right]^2,$$

where $\hat{\gamma} \in \arg \max_{\gamma \in \Lambda_2} \max\{\beta \in \Lambda_2; B(\pi_{\exp(-\gamma \frac{r}{2})}, \beta, \gamma) \leq 0\}$ (where Λ_2 is defined by equation (1.32) on page 77).

When $\kappa > 1$, $\varphi(x) \leq (1 - \kappa^{-1})(\kappa x)^{-\frac{1}{\kappa-1}}$, and we can choose γ and x such that $\frac{\gamma^2}{N} \varphi(x) \simeq d$ to prove

COROLLARY 1.52. *Assuming that the margin assumption (1.21) is satisfied for some exponent $\kappa > 1$, that $R : \Theta \rightarrow (0, 1)$ is independent of N (which is for instance the case when $\mathbb{P} = P^{\otimes N}$), and is such that*

$$\lim_{\beta' \rightarrow +\infty} \beta' [\pi_{\exp(-\beta' R)}(R) - \inf_{\Theta} R] = d,$$

there are universal positive constants B_7 and B_8 and $N_1 \in \mathbb{N}$ such that for any $N \geq N_1$, with \mathbb{P} probability at least $1 - \epsilon$,

$$\pi_{\exp(-\hat{\gamma}\frac{r}{2})}(R) \leq \inf_{\Theta} R + B_7 c^{-\frac{1}{2\kappa-1}} \left[1 + \frac{B_8}{d} \log\left(\frac{\log(N)}{\epsilon}\right) \right]^{\frac{2\kappa}{2\kappa-1}} \left(\frac{d}{N}\right)^{\frac{\kappa}{2\kappa-1}},$$

where $\hat{\gamma} \in \arg \max_{\gamma \in \Lambda_2} \max\{\beta \in \Lambda_2; B(\pi_{\exp(-\gamma\frac{r}{2}), \beta, \gamma) \leq 0\}$ (Λ_2 being defined by equation (1.32) on page 77).

We find the same rate of convergence as in Corollary 1.30 on page 48, but this time, we were able to provide an empirical posterior distribution $\pi_{\exp(-\hat{\gamma}\frac{r}{2})}$ which achieves this rate adaptively in all the parameters (meaning in particular that we do not need to know d , c or κ). Moreover, as already mentioned, the power of N in this rate of convergence is known to be unimprovable in the worst case (see [28, 34, 35], and more specifically in [3] — downloadable from its author's web page, — Theorem 3.3 on page 132).

1.5.5. Estimating the divergence of a posterior with respect to a Gibbs prior.

Another interesting question is to estimate $\mathcal{K}[\rho, \pi_{\exp(-\beta R)}]$ using relative deviation inequalities. We follow here an idea to be found first in Audibert [3, page 93]. Indeed, combining equation (1.24) with equation (1.25) on page 61, we see that for any positive real parameters β and λ , with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,

$$\begin{aligned} \mathcal{K}[\rho, \pi_{\exp(-\beta R)}] &\leq \frac{\beta}{N\lambda} \left\{ \frac{N}{2} \log\left(\frac{1+\lambda}{1-\lambda}\right) [\rho(r) - \pi_{\exp(-\beta R)}(r)] \right. \\ &\quad \left. - \frac{N}{2} \log(1-\lambda^2) \rho \otimes \pi_{\exp(-\beta R)}(m') \right. \\ &\quad \left. + \mathcal{K}[\rho, \pi_{\exp(-\beta R)}] - \log(\epsilon) \right\} + \mathcal{K}(\rho, \pi) - \mathcal{K}[\pi_{\exp(-\beta R)}, \pi] \\ &\leq \mathcal{K}[\rho, \pi_{\exp[-\frac{\beta}{2\lambda} \log(\frac{1+\lambda}{1-\lambda})r]}] + \frac{\beta}{N\lambda} \mathcal{K}[\rho, \pi_{\exp(-\beta R)}] - \frac{\beta}{N\lambda} \log(\epsilon) \\ &\quad + \log \left[\pi_{\exp[-\frac{\beta}{2\lambda} \log(\frac{1+\lambda}{1-\lambda})r]} \left\{ \exp \left[-\frac{\beta}{2\lambda} \log(1-\lambda^2) \rho(m') \right] \right\} \right]. \end{aligned}$$

Thus, putting $\gamma = \frac{N}{2} \log(\frac{1+\lambda}{1-\lambda})$, we obtain

THEOREM 1.53. *For any positive real constants β and γ such that $\beta < N \tanh(\frac{\gamma}{N})$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned} \mathcal{K}[\rho, \pi_{\exp(-\beta R)}] &\leq \left(1 - \frac{\beta}{N} \tanh\left(\frac{\gamma}{N}\right)^{-1}\right)^{-1} \\ &\quad \times \left\{ \mathcal{K}[\rho, \pi_{\exp[-\frac{\beta\gamma}{N} \tanh(\frac{\gamma}{N})^{-1}r]}] - \frac{\beta}{N \tanh(\frac{\gamma}{N})} \log(\epsilon) \right. \\ &\quad \left. + \log \left\{ \pi_{\exp[-\frac{\beta\gamma}{N} \tanh(\frac{\gamma}{N})^{-1}r]} \left[\exp\left\{ \beta \tanh\left(\frac{\gamma}{N}\right)^{-1} \log[\cosh(\frac{\gamma}{N})] \rho(m') \right\} \right] \right\} \right\}. \end{aligned}$$

This theorem provides another way of measuring overfitting, since it gives an upper bound for $\mathcal{K}[\pi_{\exp[-\frac{\beta\gamma}{N} \tanh(\frac{\gamma}{N})^{-1}r]}, \pi_{\exp(-\beta R)}]$. It may be used in combination with Theorem 1.10 on page 20 as an alternative to Theorem 1.18 on page 30. It will also be used in the next section.

An alternative parametrization of the same result providing a simpler right-hand side is also useful:

COROLLARY 1.54 *For any positive real constants β and γ such that $\beta < \gamma$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned} \mathcal{K}[\rho, \pi_{\exp[-N\frac{\beta}{\gamma} \tanh(\frac{\gamma}{N})R]}] &\leq \left(1 - \frac{\beta}{\gamma}\right)^{-1} \left\{ \mathcal{K}[\rho, \pi_{\exp(-\beta r)}] - \frac{\beta}{\gamma} \log(\epsilon) \right. \\ &\quad \left. + \log \left\{ \pi_{\exp(-\beta r)} \left[\exp\left\{ N\frac{\beta}{\gamma} \log[\cosh(\frac{\gamma}{N})] \rho(m') \right\} \right] \right\} \right\}. \end{aligned}$$

1.5.6. Comparing two posterior distributions. Estimating the effective temperature of an estimator provides an efficient way to tune parameters in a model with a parametric behaviour. On the other hand, it will not be fitted to choose between different models, especially in the case when they are nested (because as we already saw in the case when Θ is a union of nested models, the prior distribution $\pi_{\exp(-\beta R)}$ is not providing an efficient localization of the parameter in this case, in the sens that $\pi_{\exp(-\beta R)}(R)$ is not going down to $\inf_{\Theta} R$ at the desired rate when β goes to $+\infty$, requiring to resort to partial localization).

Once some estimator (in the form of a posterior distribution) has been chosen in each submodel, these estimators can be compared between themselves with the help of the relative bounds that we will establish in this section.

From equation (1.23) (slightly modified by replacing $\pi \otimes \pi$ with $\pi^1 \otimes \pi^2$), we obtain easily

THEOREM 1.55. *For any positive real constant λ , for any prior distributions $\pi^1, \pi^2 \in \mathcal{M}_+^1(\Theta)$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distributions ρ_1 and $\rho_2 : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned} -N \log \left\{ 1 - \tanh\left(\frac{\lambda}{N}\right) \left[\rho_2(R) - \rho_1(R) \right] \right\} &\leq \lambda [\rho_2(r) - \rho_1(r)] \\ &+ N \log \left[\cosh\left(\frac{\lambda}{N}\right) \right] \rho_1 \otimes \rho_2(m') \\ &+ \mathcal{K}(\rho_1, \pi^1) + \mathcal{K}(\rho_2, \pi^2) - \log(\epsilon). \end{aligned}$$

There enters into the game the entropy bound of the previous section, providing a localized version of Theorem 1.55. We will use the notation

$$\Xi_a(q) = \tanh(a)^{-1} [1 - \exp(-aq)] \leq \frac{a}{\tanh(a)} q, \quad a, q \in \mathbb{R}.$$

THEOREM 1.56. *For any sequence of prior distributions $(\pi^i)_{i \in \mathbb{N}} \in \mathcal{M}_+^1(\Theta)^{\mathbb{N}}$, any probability distribution μ on \mathbb{N} , any atomic probability distribution ν on \mathbb{R}_+ , with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distributions $\rho_1, \rho_2 : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned} \rho_2(R) - \rho_1(R) &\leq B(\rho_1, \rho_2), \text{ where} \\ B(\rho_1, \rho_2) &= \inf_{\lambda, \beta_1 < \gamma_1, \beta_2 < \gamma_2 \in \mathbb{R}_+, i, j \in \mathbb{N}} \Xi_{\frac{\lambda}{N}} \left\{ \left[\rho_2(r) - \rho_1(r) \right] \right. \\ &\quad \left. + \frac{N}{\lambda} \log \left[\cosh\left(\frac{\lambda}{N}\right) \right] \rho_1 \otimes \rho_2(m') \right. \\ &\quad \left. + \frac{1}{\lambda \left(1 - \frac{\beta_1}{\gamma_1}\right)} \left\{ \mathcal{K} \left[\rho_1, \pi_{\exp(-\beta_1 r)}^i \right] \right. \right. \\ &\quad \left. \left. + \log \left\{ \pi_{\exp(-\beta_1 r)}^i \left[\exp \left\{ \beta_1 \frac{N}{\gamma_1} \log \left[\cosh\left(\frac{\gamma_1}{N}\right) \right] \rho_1(m') \right\} \right] \right\} \right\} \right. \\ &\quad \left. + \frac{1}{\lambda \left(1 - \frac{\beta_2}{\gamma_2}\right)} \left\{ \mathcal{K} \left[\rho_2, \pi_{\exp(-\beta_2 r)}^j \right] \right. \right. \\ &\quad \left. \left. + \log \left\{ \pi_{\exp(-\beta_2 r)}^j \left[\exp \left\{ \beta_2 \frac{N}{\gamma_2} \log \left[\cosh\left(\frac{\gamma_2}{N}\right) \right] \rho_2(m') \right\} \right] \right\} \right\} \right. \\ &\quad \left. - \left[\left(\frac{\gamma_1}{\beta_1} - 1 \right)^{-1} + \left(\frac{\gamma_2}{\beta_2} - 1 \right)^{-1} + 1 \right] \right. \\ &\quad \left. \times \frac{\log \left[3^{-1} \nu(\beta_1) \nu(\beta_2) \nu(\gamma_1) \nu(\gamma_2) \nu(\lambda) \mu(i) \mu(j) \epsilon \right]}{\lambda} \right\}. \end{aligned}$$

The sequence of prior distributions $(\pi^i)_{i \in \mathbb{N}}$ should be understood to be typically supported by subsets of Θ corresponding to parametric submodels, that is submodels for which it is reasonable to expect that

$$\lim_{\beta \rightarrow +\infty} \beta [\pi_{\exp(-\beta R)}^i(R) - \operatorname{ess\,inf}_{\pi^i} R]$$

exists and is positive and finite. As there is no reason why the bound $B(\rho_1, \rho_2)$ provided by the previous theorem should be subadditive (in the sense that $B(\rho_1, \rho_3) \leq B(\rho_1, \rho_2) + B(\rho_2, \rho_3)$), it is adequate, at least from a theoretical point of view, to consider some workable subset $\mathcal{P} \subset \mathcal{M}_+^1(\Theta)$ of posterior distributions (for instance the distributions of the form $\pi_{\exp(-\beta r)}^i$, $i \in \mathbb{N}$, $\beta \in \mathbb{R}_+$, it is understood that \mathcal{P} is allowed to be a random subset of $\mathcal{M}_+^1(\Theta)$, as in this suggested example), and to define the subadditive chained bound

$$\tilde{B}(\rho, \rho') = \inf \left\{ \sum_{k=0}^{n-1} B(\rho_k, \rho_{k+1}); n \in \mathbb{N}^*, (\rho_k)_{k=0}^n \in \mathcal{P}^{n+1}, \right. \\ \left. \rho_0 = \rho, \rho_n = \rho' \right\}, \quad \rho, \rho' \in \mathcal{P}.$$

PROPOSITION 1.57. *With \mathbb{P} probability at least $1 - \epsilon$, for any posterior distributions $\rho_1, \rho_2 \in \mathcal{P}$, $\rho_2(R) - \rho_1(R) \leq \tilde{B}(\rho_1, \rho_2)$. Moreover for any posterior distribution $\rho_1 \in \mathcal{P}$, any posterior distribution $\rho_2 \in \mathcal{P}$ such that $\tilde{B}(\rho_1, \rho_2) = \inf_{\rho_3 \in \mathcal{P}} \tilde{B}(\rho_1, \rho_3)$ is unimprovable with the help of \tilde{B} in \mathcal{P} in the sense that $\inf_{\rho_3 \in \mathcal{P}} \tilde{B}(\rho_2, \rho_3) \geq 0$.*

PROOF. The first assertion is a direct consequence of the previous theorem, therefore only the second assertion requires a proof: for any $\rho_3 \in \mathcal{P}$, we deduce from the optimality of ρ_2 and the subadditivity of \tilde{B} that $\tilde{B}(\rho_1, \rho_2) \leq \tilde{B}(\rho_1, \rho_3) \leq \tilde{B}(\rho_1, \rho_2) + \tilde{B}(\rho_2, \rho_3)$. \square

This proposition provides a way to improve a posterior distribution $\rho_1 \in \mathcal{P}$ by choosing $\rho_2 \in \arg \min_{\rho \in \mathcal{P}} \tilde{B}(\rho_1, \rho)$ whenever $\tilde{B}(\rho_1, \rho_2) < 0$. This improvement process is proved according to Proposition 1.57 to be a one step process: the obtained improved posterior ρ_2 cannot be improved again using the same technique.

Let us give some example of possible starting distribution ρ_1 for this improvement scheme: ρ_1 may be chosen as the best posterior Gibbs distribution according to Proposition 1.40 on page 66. More precisely, we may build from the prior distributions π^i , $i \in \mathbb{N}$, a global prior $\pi = \sum_{i \in \mathbb{N}} \mu(i) \pi^i$. We can

then define the estimator of the inverse effective temperature as in Proposition 1.40 and choose $\rho_1 \in \arg \min_{\rho \in \mathcal{P}} \widehat{\beta}(\rho)$, where \mathcal{P} is as suggested above the set of posterior distributions

$$\mathcal{P} = \left\{ \pi_{\exp(-\beta r)}^i; i \in \mathbb{N}, \beta \in \mathbb{R}_+ \right\}.$$

(This starting point ρ_1 should already be pretty good, at least in an asymptotic perspective, the only gain in the rate of convergence to be expected bearing on spurious $\log(N)$ factors).

For more elaborate uses of relative bounds, we refer to the third section of the second chapter of Audibert [3], where an algorithm is proposed and analyzed, which allows to use relative bounds between two posterior distributions as a stand alone estimation tool.

1.5.7. Two step localization of relative bounds. Let us consider again in this section the case when we want to choose adaptively between a family of parametric models. Let us thus assume that the parameter set is a disjoint union of measurable submodels, so that we can write $\Theta = \sqcup_{m \in M} \Theta_m$, where M is some measurable index set. Let us choose some prior probability distribution on the index set $\mu \in \mathcal{M}_+^1(M)$, and some regular conditional prior distribution on (M, Θ) , $\pi : M \rightarrow \mathcal{M}_+^1(\Theta)$, such that $\pi(m, \Theta_m) = 1$, $m \in M$. Let us then study some arbitrary posterior distributions $\nu : \Omega \rightarrow \mathcal{M}_+^1(M)$ and $\rho : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$, such that $\rho(\omega, m, \Theta_m) = 1$, $\omega \in \Omega$, $m \in M$. We would like to compare $\nu \rho(R)$ with some doubly localized prior distribution $\mu_{\exp[-\frac{\beta}{1+\zeta_2} \pi_{\exp(-\beta R)}(R)]} [\pi_{\exp(-\beta R)}](R)$ (where ζ_2 is a positive parameter to be set as needed later on). We will define to ease notations two prior distributions (one being more precisely a conditional distribution) depending on the positive real parameters β and ζ_2 , putting

$$\overline{\pi} = \pi_{\exp(-\beta R)} \text{ and } \overline{\mu} = \mu_{\exp[-\frac{\beta}{1+\zeta_2} \overline{\pi}(R)]}. \quad (1.33)$$

Similarly to Theorem 1.26 on page 44 we can write for any positive real constants β and γ

$$\mathbb{P} \left\{ (\overline{\mu} \overline{\pi}) \otimes (\overline{\mu} \overline{\pi}) \left[\exp \left[-N \log \left[1 - \tanh\left(\frac{\gamma}{N}\right) R' \right] \right. \right. \right. \\ \left. \left. \left. - \gamma r' - N \log \left[\cosh\left(\frac{\gamma}{N}\right) m' \right] \right] \right] \right\} \leq 1,$$

and deduce, using Lemma 1.3 on page 11

$$\mathbb{P} \left\{ \exp \left[\sup_{\nu \in \mathcal{M}_+^1(M)} \sup_{\rho: M \rightarrow \mathcal{M}_+^1(\Theta)} \left\{ -N \log [1 - \tanh(\frac{\gamma}{N})(\nu\rho - \bar{\mu}\bar{\pi})(R)] \right. \right. \right. \\ \left. \left. \left. - \gamma(\nu\rho - \bar{\mu}\bar{\pi})(r) - N \log [\cosh(\frac{\gamma}{N})] (\nu\rho) \otimes (\bar{\mu}\bar{\pi})(m') \right. \right. \right. \\ \left. \left. \left. - \mathcal{K}(\nu, \bar{\mu}) - \nu[\mathcal{K}(\rho, \bar{\pi})] \right\} \right] \right\} \leq 1. \quad (1.34)$$

This will be our starting point in comparing $\nu\rho(R)$ with $\bar{\mu}\bar{\pi}(R)$. However, obtaining an empirical bound will require some supplementary efforts. For each $m \in M$, we can write in the same way

$$\mathbb{P} \left\{ \bar{\pi} \otimes \bar{\pi} \left[\exp \left[-N \log [1 - \tanh(\frac{\gamma}{N})R'] - \gamma r' - N \log [\cosh(\frac{\gamma}{N})] m' \right] \right] \right\} \leq 1.$$

Integrating this inequality with respect to $\bar{\mu}$ and using Fubini's lemma for positive functions, we get

$$\mathbb{P} \left\{ \bar{\mu}(\bar{\pi} \otimes \bar{\pi}) \left[\exp \left[-N \log [1 - \tanh(\frac{\gamma}{N})R'] - \gamma r' - N \log [\cosh(\frac{\gamma}{N})] m' \right] \right] \right\} \leq 1.$$

Let us make clear that $\bar{\mu}(\bar{\pi} \otimes \bar{\pi})$ is a probability measure on $M \times \Theta \times \Theta$, whereas $(\bar{\mu}\bar{\pi}) \otimes (\bar{\mu}\bar{\pi})$ considered previously is a probability measure on $(M \times \Theta) \times (M \times \Theta)$. We get as previously

$$\mathbb{P} \left\{ \exp \left[\sup_{\nu \in \mathcal{M}_+^1(M)} \sup_{\rho: M \rightarrow \mathcal{M}_+^1(\Theta)} \left\{ -N \log [1 - \tanh(\frac{\gamma}{N})\nu(\rho - \bar{\pi})(R)] \right. \right. \right. \\ \left. \left. \left. - \gamma\nu(\rho - \bar{\pi})(r) - N \log [\cosh(\frac{\gamma}{N})] \nu(\rho \otimes \bar{\pi})(m') \right. \right. \right. \\ \left. \left. \left. - \mathcal{K}(\nu, \bar{\mu}) - \nu[\mathcal{K}(\rho, \bar{\pi})] \right\} \right] \right\} \leq 1. \quad (1.35)$$

Let us eventually recall that

$$\mathcal{K}(\nu, \bar{\mu}) = \frac{\beta}{1+\zeta_2}(\nu - \bar{\mu})\bar{\pi}(R) + \mathcal{K}(\nu, \mu) - \mathcal{K}(\bar{\mu}, \mu), \quad (1.36)$$

$$\mathcal{K}(\rho, \bar{\pi}) = \beta(\rho - \bar{\pi})(R) + \mathcal{K}(\rho, \pi) - \mathcal{K}(\bar{\pi}, \pi). \quad (1.37)$$

From equations (1.34), (1.35) and (1.37) we deduce

PROPOSITION 1.58. *For any positive real constants β , γ and ζ_2 , with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\nu : \Omega \rightarrow \mathcal{M}_+^1(M)$ and any conditional posterior distribution $\rho : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\begin{aligned}
& -N \log \left[1 - \tanh\left(\frac{\gamma}{N}\right) (\nu\rho - \bar{\mu}\bar{\pi})(R) \right] - \beta\nu(\rho - \bar{\pi})(R) \\
& \leq \gamma(\nu\rho - \bar{\mu}\bar{\pi})(r) + N \log \left[\cosh\left(\frac{\gamma}{N}\right) \right] (\nu\rho) \otimes (\bar{\mu}\bar{\pi})(m') \\
& \quad + \mathcal{K}(\nu, \bar{\mu}) + \nu[\mathcal{K}(\rho, \pi)] - \nu[\mathcal{K}(\bar{\pi}, \pi)] + \log\left(\frac{2}{\epsilon}\right).
\end{aligned}$$

and

$$\begin{aligned}
& -N \log \left[1 - \tanh\left(\frac{\gamma}{N}\right) \nu(\rho - \bar{\pi})(R) \right] \\
& \leq \gamma\nu(\rho - \bar{\pi})(r) + N \log \left[\cosh\left(\frac{\gamma}{N}\right) \right] \nu(\rho \otimes \bar{\pi})(m') \\
& \quad + \mathcal{K}(\nu, \bar{\mu}) + \nu[\mathcal{K}(\rho, \bar{\pi})] + \log\left(\frac{2}{\epsilon}\right),
\end{aligned}$$

where the prior distribution $\bar{\mu}\bar{\pi}$ is defined by equation (1.33) on page 83 and depends on β and ζ_2 .

Let us put for short

$$T = \tanh\left(\frac{\gamma}{N}\right) \text{ and } C = N \log \left[\cosh\left(\frac{\gamma}{N}\right) \right].$$

We will use some entropy compensation strategy for which we need a couple of entropy bounds. Let us assume that $\beta < NT$. We have according to Proposition 1.58, with \mathbb{P} probability at least $1 - \epsilon$,

$$\begin{aligned}
\nu[\mathcal{K}(\rho, \bar{\pi})] &= \beta\nu(\rho - \bar{\pi})(R) + \nu[\mathcal{K}(\rho, \pi) - \mathcal{K}(\bar{\pi}, \pi)] \\
&\leq \frac{\beta}{NT} \left[\gamma\nu(\rho - \bar{\pi})(r) + C\nu(\rho \otimes \bar{\pi})(m') \right. \\
&\quad \left. + \mathcal{K}(\nu, \bar{\mu}) + \nu[\mathcal{K}(\rho, \bar{\pi})] + \log\left(\frac{2}{\epsilon}\right) \right] \\
&\quad + \nu[\mathcal{K}(\rho, \pi) - \mathcal{K}(\bar{\pi}, \pi)].
\end{aligned}$$

Similarly

$$\begin{aligned}
\mathcal{K}(\nu, \bar{\mu}) &= \frac{\beta}{1 + \zeta_2} (\nu - \bar{\mu})\bar{\pi}(R) + \mathcal{K}(\nu, \mu) - \mathcal{K}(\bar{\mu}, \mu) \\
&\leq \frac{\beta}{(1 + \zeta_2)NT} \left[\gamma(\nu - \bar{\mu})\bar{\pi}(r) + C(\nu\bar{\pi}) \otimes (\bar{\mu}\bar{\pi})(m') \right. \\
&\quad \left. + \mathcal{K}(\nu, \bar{\mu}) + \log\left(\frac{2}{\epsilon}\right) \right] + \mathcal{K}(\nu, \mu) - \mathcal{K}(\bar{\mu}, \mu).
\end{aligned}$$

Thus, for any positive real constants β , γ and ζ_i , $i = 1, \dots, 5$, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distributions $\nu, \nu_3 : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, any posterior conditional distributions $\rho, \rho_1, \rho_2, \rho_4, \rho_5 : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$,

$$\begin{aligned}
& -N \log[1 - T(\nu\rho - \bar{\mu}\bar{\pi})(R)] - \beta\nu(\rho - \bar{\pi})(R) \\
& \leq \gamma(\nu\rho - \bar{\mu}\bar{\pi})(r) + C(\nu\rho) \otimes (\bar{\mu}\bar{\pi})(m') \\
& \quad + \mathcal{K}(\nu, \bar{\mu}) + \nu[\mathcal{K}(\rho, \pi) - \mathcal{K}(\bar{\pi}, \pi)] + \log\left(\frac{2}{\epsilon}\right), \\
\zeta_1 \frac{NT}{\beta} \bar{\mu}[\mathcal{K}(\rho_1, \bar{\pi})] & \leq \zeta_1 \gamma \bar{\mu}(\rho_1 - \bar{\pi})(r) + \zeta_1 C \bar{\mu}(\rho_1 \otimes \bar{\pi})(m') \\
& \quad + \zeta_1 \bar{\mu}[\mathcal{K}(\rho_1, \bar{\pi})] + \zeta_1 \log\left(\frac{2}{\epsilon}\right) + \zeta_1 \frac{NT}{\beta} \bar{\mu}[\mathcal{K}(\rho_1, \pi) - \mathcal{K}(\bar{\pi}, \pi)], \\
\zeta_2 \frac{NT}{\beta} \nu[\mathcal{K}(\rho_2, \bar{\pi})] & \leq \zeta_2 \gamma \nu(\rho_2 - \bar{\pi})(r) + \zeta_2 C \nu(\rho_2 \otimes \bar{\pi})(m') \\
& \quad + \zeta_2 \mathcal{K}(\nu, \bar{\mu}) + \zeta_2 \nu[\mathcal{K}(\rho_2, \bar{\pi})] + \zeta_2 \log\left(\frac{2}{\epsilon}\right) \\
& \quad + \zeta_2 \frac{NT}{\beta} \nu[\mathcal{K}(\rho_2, \pi) - \mathcal{K}(\bar{\pi}, \pi)], \\
\zeta_3(1 + \zeta_2) \frac{NT}{\beta} \mathcal{K}(\nu_3, \bar{\mu}) & \leq \zeta_3 \gamma(\nu_3 - \bar{\mu})\bar{\pi}(r) \\
& \quad + \zeta_3 C[(\nu_3 \bar{\pi}) \otimes (\nu_3 \rho_1) + (\nu_3 \rho_1) \otimes (\bar{\mu}\bar{\pi})](m') + \zeta_3 \mathcal{K}(\nu_3, \bar{\mu}) + \zeta_3 \log\left(\frac{2}{\epsilon}\right) \\
& \quad + \zeta_3(1 + \zeta_2) \frac{NT}{\beta} [\mathcal{K}(\nu_3, \mu) - \mathcal{K}(\bar{\mu}, \mu)], \\
\zeta_4 \frac{NT}{\beta} \nu_3[\mathcal{K}(\rho_4, \bar{\pi})] & \leq \zeta_4 \gamma \nu_3(\rho_4 - \bar{\pi})(r) \\
& \quad + \zeta_4 C \nu_3(\rho_4 \otimes \bar{\pi})(m') + \zeta_4 \mathcal{K}(\nu_3, \bar{\mu}) + \zeta_4 \nu_3[\mathcal{K}(\rho_4, \bar{\pi})] + \zeta_4 \log\left(\frac{2}{\epsilon}\right) \\
& \quad + \zeta_4 \frac{NT}{\beta} \nu_3[\mathcal{K}(\rho_4, \pi) - \mathcal{K}(\bar{\pi}, \pi)], \\
\zeta_5 \frac{NT}{\beta} \bar{\mu}[\mathcal{K}(\rho_5, \bar{\pi})] & \leq \zeta_5 \gamma \bar{\mu}(\rho_5 - \bar{\pi})(r) + \zeta_5 C \bar{\mu}(\rho_5 \otimes \bar{\pi})(m') \\
& \quad + \zeta_5 \bar{\mu}[\mathcal{K}(\rho_5, \bar{\pi})] + \zeta_5 \log\left(\frac{2}{\epsilon}\right) + \zeta_5 \frac{NT}{\beta} \bar{\mu}[\mathcal{K}(\rho_5, \pi) - \mathcal{K}(\bar{\pi}, \pi)].
\end{aligned}$$

Adding these six inequalities and assuming that $\zeta_4 \leq \zeta_3[(1 + \zeta_2)\frac{NT}{\beta} - 1]$, we find

$$\begin{aligned}
& -N \log[1 - T(\nu\rho - \bar{\mu}\bar{\pi})(R)] - \beta(\nu\rho - \bar{\mu}\bar{\pi})(R) \\
& \leq -N \log[1 - T(\nu\rho - \bar{\mu}\bar{\pi})(R)] - \beta(\nu\rho - \bar{\mu}\bar{\pi})(R) \\
& \quad + \zeta_1 \left(\frac{NT}{\beta} - 1\right) \bar{\mu}[\mathcal{K}(\rho_1, \bar{\pi})] + \zeta_2 \left(\frac{NT}{\beta} - 1\right) \nu[\mathcal{K}(\rho_2, \bar{\pi})] \\
& \quad + [\zeta_3(1 + \zeta_2) \frac{NT}{\beta} - \zeta_3 - \zeta_4] \mathcal{K}(\nu_3, \bar{\mu}) \\
& \quad + \zeta_4 \left(\frac{NT}{\beta} - 1\right) \nu_3[\mathcal{K}(\rho_4, \bar{\pi})] + \zeta_5 \left(\frac{NT}{\beta} - 1\right) \bar{\mu}[\mathcal{K}(\rho_5, \bar{\pi})] \\
& \leq \gamma(\nu\rho - \bar{\mu}\bar{\pi})(r) + \zeta_1 \gamma \bar{\mu}(\rho_1 - \bar{\pi})(r) + \zeta_2 \gamma \nu(\rho_2 - \bar{\pi})(r)
\end{aligned}$$

$$\begin{aligned}
& + \zeta_3 \gamma (\nu_3 - \bar{\mu}) \bar{\pi}(r) + \zeta_4 \gamma \nu_3 (\rho_4 - \bar{\pi})(r) + \zeta_5 \gamma \bar{\mu} (\rho_5 - \bar{\pi})(r) \\
& \quad + C [(\nu \rho) \otimes (\bar{\mu} \bar{\pi}) + \zeta_1 \bar{\mu} (\rho_1 \otimes \bar{\pi}) + \zeta_2 \nu (\rho_2 \otimes \bar{\pi}) \\
& + \zeta_3 (\nu_3 \bar{\pi}) \otimes (\nu_3 \rho_1) + \zeta_3 (\nu_3 \rho_1) \otimes (\bar{\mu} \bar{\pi}) \\
& \quad + \zeta_4 \nu_3 (\rho_4 \otimes \bar{\pi}) + \zeta_5 \bar{\mu} (\rho_5 \otimes \bar{\pi})] (m') \\
& + (1 + \zeta_2) [\mathcal{K}(\nu, \mu) - \mathcal{K}(\bar{\mu}, \mu)] + \nu [\mathcal{K}(\rho, \pi) - \mathcal{K}(\bar{\pi}, \pi)] \\
& \quad + \zeta_1 \frac{NT}{\beta} \bar{\mu} [\mathcal{K}(\rho_1, \pi) - \mathcal{K}(\bar{\pi}, \pi)] + \zeta_2 \frac{NT}{\beta} \nu [\mathcal{K}(\rho_2, \pi) - \mathcal{K}(\bar{\pi}, \pi)] \\
& + \zeta_3 (1 + \zeta_2) \frac{NT}{\beta} [\mathcal{K}(\nu_3, \mu) - \mathcal{K}(\bar{\mu}, \mu)] + \zeta_4 \frac{NT}{\beta} \nu_3 [\mathcal{K}(\rho_4, \pi) - \mathcal{K}(\bar{\pi}, \pi)] \\
& \quad + \zeta_5 \frac{NT}{\beta} \bar{\mu} [\mathcal{K}(\rho_5, \pi) - \mathcal{K}(\bar{\pi}, \pi)] + (1 + \zeta_1 + \zeta_2 + \zeta_3 + \zeta_4 + \zeta_5) \log\left(\frac{2}{\epsilon}\right).
\end{aligned}$$

Let us now apply to $\bar{\pi}$ (we shall later do the same with $\bar{\mu}$) the following inequalities, holding for any random functions of the sample and the parameters $h : \Omega \times \Theta \rightarrow \mathbb{R}$ and $g : \Omega \times \Theta \rightarrow \mathbb{R}$,

$$\begin{aligned}
\bar{\pi}(g - h) - \mathcal{K}(\bar{\pi}, \pi) & \leq \sup_{\rho: \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)} \rho(g - h) - \mathcal{K}(\rho, \pi) \\
& = \log \{ \pi [\exp(g - h)] \} \\
& \quad = \log \{ \pi [\exp(-h)] \} + \log \{ \pi_{\exp(-h)} [\exp(g)] \} \\
& \quad = -\pi_{\exp(-h)}(h) - \mathcal{K}(\pi_{\exp(-h)}, \pi) + \log \{ \pi_{\exp(-h)} [\exp(g)] \}.
\end{aligned}$$

When h and g are observable, and h is not too far from $\beta r \simeq \beta R$, this gives a way to replace $\bar{\pi}$ with some satisfactory empirical approximation. We will apply this method, choosing ρ_1 and ρ_5 such that $\bar{\mu} \bar{\pi}$ is replaced either with $\bar{\mu} \rho_1$, when it comes from the first two inequalities or with $\bar{\mu} \rho_5$ otherwise, choosing ρ_2 such that $\nu \bar{\pi}$ is replaced with $\nu \rho_2$ and ρ_4 such that $\nu_3 \bar{\pi}$ is replaced with $\nu_3 \rho_4$. We will do so because it leads to a lot of helpful cancellations. For those to happen, we need to choose $\rho_i = \pi_{\exp(-\lambda_i r)}$, $i = 1, 2, 4$, where λ_1, λ_2 and λ_4 are such that

$$\begin{aligned}
(1 + \zeta_1) \gamma & = \zeta_1 \frac{NT}{\beta} \lambda_1, \\
\zeta_2 \gamma & = (1 + \zeta_2 \frac{NT}{\beta}) \lambda_2, \\
(\zeta_4 - \zeta_3) \gamma & = \zeta_4 \frac{NT}{\beta} \lambda_4, \\
\zeta_3 \gamma & = \zeta_5 \frac{NT}{\beta} \lambda_5,
\end{aligned}$$

and to assume that $\zeta_4 > \zeta_3$. We obtain that with \mathbb{P} probability at least $1 - \epsilon$,

$$-N \log [1 - T(\mu \rho - \bar{\mu} \bar{\pi})(R)] - \beta(\nu \rho - \bar{\mu} \bar{\pi})(R)$$

$$\begin{aligned}
&\leq \gamma(\nu\rho - \bar{\mu}\rho_1)(r) + \zeta_3\gamma(\nu_3\rho_4 - \bar{\mu}\rho_5)(r) \\
&+ \zeta_1\frac{NT}{\beta}\bar{\mu}\left\{\log\left[\rho_1\left\{\exp\left[C\frac{\beta}{NT\zeta_1}[\nu\rho + \zeta_1\rho_1](m')\right]\right\}\right]\right\} \\
&+ (1 + \zeta_2)\frac{NT}{\beta}\nu\left\{\log\left[\rho_2\left\{\exp\left[\frac{C}{1+\zeta_2}\frac{NT}{\beta}\zeta_2\rho_2(m')\right]\right\}\right]\right\} \\
&+ \zeta_4\frac{NT}{\beta}\nu_3\left\{\log\left[\rho_4\left\{\exp\left[C\frac{\beta}{NT\zeta_4}[\zeta_3\nu_3\rho_1 + \zeta_4\rho_4](m')\right]\right\}\right]\right\} \\
&+ \zeta_5\frac{NT}{\beta}\bar{\mu}\left\{\log\left[\rho_5\left\{\exp\left[C\frac{\beta}{NT\zeta_5}[\zeta_3\nu_3\rho_1 + \zeta_5\rho_5](m')\right]\right\}\right]\right\} \\
&+ (1 + \zeta_2)[\mathcal{K}(\nu, \mu) - \mathcal{K}(\bar{\mu}, \mu)] + \nu[\mathcal{K}(\rho, \pi) - \mathcal{K}(\rho_2, \pi)] \\
&+ \zeta_3(1 + \zeta_2)\frac{NT}{\beta}[\mathcal{K}(\nu_3, \mu) - \mathcal{K}(\bar{\mu}, \mu)] \\
&+ \left(1 + \sum_{i=1}^5 \zeta_i\right) \log\left(\frac{2}{\epsilon}\right).
\end{aligned}$$

In order to obtain more cancellations while replacing $\bar{\mu}$ by some posterior distribution, we will choose the constants such that $\lambda_5 = \lambda_4$, which can be done by choosing

$$\zeta_5 = \frac{\zeta_3\zeta_4}{\zeta_4 - \zeta_3}.$$

We can now replace $\bar{\mu}$ with $\mu_{\exp -\xi_1\rho_1(r) - \xi_4\rho_4(r)}$, where

$$\begin{aligned}
\xi_1 &= \frac{\gamma}{(1 + \zeta_2)\left(1 + \frac{NT}{\beta}\zeta_3\right)}, \\
\xi_4 &= \frac{\gamma\zeta_3}{(1 + \zeta_2)\left(1 + \frac{NT}{\beta}\zeta_3\right)}.
\end{aligned}$$

Choosing moreover $\nu_3 = \mu_{\exp -\xi_1\rho_1(r) - \xi_4\rho_4(r)}$, to induce some more cancellations, we get

THEOREM 1.59. *For any positive real constants satisfying the above mentioned constraints, with \mathbb{P} probability at least $1 - \epsilon$, for any posterior distribution $\nu : \Omega \rightarrow \mathcal{M}_+^1(M)$ and any conditional posterior distribution $\rho : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$-N \log[1 - T(\nu\rho - \bar{\mu}\bar{\pi})(R)] - \beta(\nu\rho - \bar{\mu}\bar{\pi})(R) \leq B(\nu, \rho, \beta),$$

$$\begin{aligned}
& \text{where } B(\nu, \rho, \beta) \stackrel{\text{def}}{=} \gamma(\nu\rho - \nu_3\rho_1)(r) \\
& + (1 + \zeta_2)\left(1 + \frac{NT}{\beta}\zeta_3\right) \\
& \times \log \left\{ \nu_3 \left[\rho_1 \left\{ \exp \left[C \frac{\beta}{NT\zeta_1} [\nu\rho + \zeta_1\rho_1](m') \right] \right\}^{\frac{\zeta_1 NT}{\beta(1+\zeta_2)(1+\frac{NT}{\beta}\zeta_3)}} \right. \right. \\
& \quad \left. \left. \times \rho_4 \left\{ \exp \left[C \frac{\beta}{NT\zeta_5} [\zeta_3\nu_3\rho_1 + \zeta_5\rho_4](m') \right] \right\}^{\frac{\zeta_5 NT}{\beta(1+\zeta_2)(1+\frac{NT}{\beta}\zeta_3)}} \right] \right\} \\
& + (1 + \zeta_2 \frac{NT}{\beta}) \nu \left\{ \log \left\{ \rho_2 \left\{ \exp \left[\frac{C}{1+\zeta_2} \frac{NT}{\beta} \zeta_2 \rho_2(m') \right] \right\} \right\} \right\} \\
& + \zeta_4 \frac{NT}{\beta} \nu_3 \left\{ \log \left[\rho_4 \left\{ \exp \left[C \frac{\beta}{NT\zeta_4} [\zeta_3\nu_3\rho_1 + \zeta_4\rho_4](m') \right] \right\} \right] \right\} \\
& + (1 + \zeta_2) [\mathcal{K}(\nu, \mu) - \mathcal{K}(\nu_3, \mu)] \\
& \quad + \nu [\mathcal{K}(\rho, \pi) - \mathcal{K}(\rho_2, \pi)] + \left(1 + \sum_{i=1}^5 \zeta_i \right) \log\left(\frac{2}{\epsilon}\right).
\end{aligned}$$

This theorem can be used to find the largest value $\hat{\beta}(\nu\rho)$ of β such that $B(\nu, \rho, \beta) \leq 0$, thus providing an estimator for $\beta(\nu\rho)$ defined as $\nu\rho(R) = \bar{\mu}_{\beta(\nu\rho)} \bar{\pi}_{\beta(\nu\rho)}(R)$, where we have mentioned explicitly the dependence of $\bar{\mu}$ and $\bar{\pi}$ in β , the constant ζ_2 staying fixed. The posterior distribution $\nu\rho$ may then be chosen to maximize $\hat{\beta}(\nu\rho)$ within some manageable subset of posterior distributions \mathcal{P} , thus gaining the assurance that $\nu\rho(R) \leq \bar{\mu}_{\hat{\beta}(\nu\rho)} \bar{\pi}_{\hat{\beta}(\nu\rho)}(R)$, with the largest parameter $\hat{\beta}(\nu\rho)$ that this approach can provide. Maximizing $\hat{\beta}(\nu\rho)$ is supported by the fact that $\lim_{\beta \rightarrow +\infty} \bar{\mu}_{\beta} \bar{\pi}_{\beta}(R) = \text{ess inf}_{\mu\pi} R$. Anyhow, there is no assurance (to our knowledge) that $\beta \mapsto \bar{\mu}_{\beta} \bar{\pi}_{\beta}(R)$ will be a decreasing function of β all the way, although this may be expected to be the case in many practical situations.

We can make the bound more explicit in several ways. One point of view is to put forward the optimal values of ρ and ν . We can thus remark that

$$\begin{aligned}
& \nu [\gamma\rho(r) + \mathcal{K}(\rho, \pi) - \mathcal{K}(\rho_2, \pi)] + (1 + \zeta_2)\mathcal{K}(\nu, \mu) \\
& = \nu \left[\mathcal{K}[\rho, \pi_{\exp(-\gamma r)}] + \lambda_2 \rho_2(r) + \int_{\lambda_2}^{\gamma} \pi_{\exp(-\alpha r)}(r) d\alpha \right] + (1 + \zeta_2)\mathcal{K}(\nu, \mu) \\
& = \nu \left\{ \mathcal{K}[\rho, \pi_{\exp(-\gamma r)}] \right\} + (1 + \zeta_2) \mathcal{K} \left[\nu, \mu_{\exp\left(-\frac{\lambda_2 \rho_2(r)}{1+\zeta_2} - \frac{1}{1+\zeta_2} \int_{\lambda_2}^{\gamma} \pi_{\exp(-\alpha r)}(r) d\alpha\right)} \right] \\
& \quad - (1 + \zeta_2) \log \left\{ \mu \left[\exp \left\{ -\frac{\lambda_2}{1 + \zeta_2} \rho_2(r) - \frac{1}{1 + \zeta_2} \int_{\lambda_2}^{\gamma} \pi_{\exp(-\alpha r)}(r) d\alpha \right\} \right] \right\}.
\end{aligned}$$

Thus

$$\begin{aligned}
B(\nu, \rho, \beta) &= (1 + \zeta_2) \left[\xi_1 \nu_3 \rho_1(r) + \xi_4 \nu_3 \rho_4(r) \right. \\
&\quad \left. + \log \left\{ \mu \left[\exp(-\xi_1 \rho_1(r) - \xi_4 \rho_4(r)) \right] \right\} \right] \\
&\quad - (1 + \zeta_2) \log \left\{ \mu \left[\exp \left\{ -\frac{\lambda_2}{1 + \zeta_2} \rho_2(r) - \frac{1}{1 + \zeta_2} \int_{\lambda_2}^{\gamma} \pi_{\exp(-\alpha r)}(r) d\alpha \right\} \right] \right\} \\
&\quad - \gamma \nu_3 \rho_1(r) + (1 + \zeta_2) \left(1 + \frac{NT}{\beta} \zeta_3 \right) \\
&\quad \times \log \left\{ \nu_3 \left[\rho_1 \left\{ \exp \left[C \frac{\beta}{NT \zeta_1} [\nu \rho + \zeta_1 \rho_1](m') \right] \right\}^{\frac{\zeta_1 NT}{\beta(1+\zeta_2)(1+\frac{NT}{\beta} \zeta_3)}} \right. \right. \\
&\quad \left. \left. \times \rho_4 \left\{ \exp \left[C \frac{\beta}{NT \zeta_5} [\zeta_3 \nu_3 \rho_1 + \zeta_5 \rho_4](m') \right] \right\}^{\frac{\zeta_5 NT}{\beta(1+\zeta_2)(1+\frac{NT}{\beta} \zeta_3)}} \right] \right\} \\
&\quad + (1 + \zeta_2 \frac{NT}{\beta}) \nu \left\{ \log \left\{ \rho_2 \left\{ \exp \left[\frac{C}{1+\zeta_2} \frac{NT}{\beta} \zeta_2 \rho_2(m') \right] \right\} \right\} \right\} \\
&\quad + \zeta_4 \frac{NT}{\beta} \nu_3 \left\{ \log \left[\rho_4 \left\{ \exp \left[C \frac{\beta}{NT \zeta_4} [\zeta_3 \nu_3 \rho_1 + \zeta_4 \rho_4](m') \right] \right\} \right] \right\} \\
&\quad + \nu \{ \mathcal{K}[\rho, \pi_{\exp(-\gamma r)}] \} \\
&\quad + (1 + \zeta_2) \mathcal{K} \left[\nu, \mu_{\exp(-\frac{\lambda_2 \rho_2(r)}{1+\zeta_2} - \frac{1}{1+\zeta_2} \int_{\lambda_2}^{\gamma} \pi_{\exp(-\alpha r)}(r) d\alpha)} \right] \\
&\quad + \left(1 + \sum_{i=1}^5 \zeta_i \right) \log \left(\frac{2}{\epsilon} \right).
\end{aligned}$$

This formula is better understood when thinking about the following upper bound for the two first lines in the expression of $B(\nu, \rho, \beta)$:

$$\begin{aligned}
&(1 + \zeta_2) \left[\xi_1 \nu_3 \rho_1(r) + \xi_4 \nu_3 \rho_4(r) + \log \left\{ \mu \left[\exp(-\xi_1 \rho_1(r) - \xi_4 \rho_4(r)) \right] \right\} \right] \\
&\quad - (1 + \zeta_2) \log \left\{ \mu \left[\exp \left\{ -\frac{\lambda_2}{1 + \zeta_2} \rho_2(r) \right. \right. \right. \\
&\quad \quad \left. \left. - \frac{1}{1 + \zeta_2} \int_{\lambda_2}^{\gamma} \pi_{\exp(-\alpha r)}(r) d\alpha \right\} \right] \right\} - \gamma \nu_3 \rho_1(r) \\
&\quad \leq \nu_3 \left[\lambda_2 \rho_2(r) + \int_{\lambda_2}^{\gamma} \pi_{\exp(-\alpha r)}(r) d\alpha - \gamma \rho_1(r) \right].
\end{aligned}$$

Another approach to understanding Theorem 1.59 is to put forward $\rho_0 = \pi_{\exp(-\lambda_0 r)}$, for some positive real constant $\lambda_0 < \gamma$, noticing that

$$\nu[\mathcal{K}(\rho_0, \pi) - \mathcal{K}(\rho_2, \pi)] = \lambda_0 \nu(\rho_2 - \rho_0)(r) - \nu[\mathcal{K}(\rho_2, \rho_0)].$$

Thus

$$\begin{aligned} B(\nu, \rho_0, \beta) &\leq \nu_3 [(\gamma - \lambda_0)(\rho_0 - \rho_1)(r) + \lambda_0(\rho_2 - \rho_1)(r)] \\ &+ (1 + \zeta_2) \left(1 + \frac{NT}{\beta} \zeta_3\right) \\ &\times \log \left\{ \nu_3 \left[\rho_1 \left\{ \exp \left[C \frac{\beta}{NT\zeta_1} [\nu\rho_0 + \zeta_1\rho_1](m') \right] \right\}^{\frac{\zeta_1 NT}{\beta(1+\zeta_2)(1+\frac{NT}{\beta}\zeta_3)}} \right. \right. \\ &\quad \left. \left. \times \rho_4 \left\{ \exp \left[C \frac{\beta}{NT\zeta_5} [\zeta_3\nu_3\rho_1 + \zeta_5\rho_4](m') \right] \right\}^{\frac{\zeta_5 NT}{\beta(1+\zeta_2)(1+\frac{NT}{\beta}\zeta_3)}} \right] \right\} \\ &+ (1 + \zeta_2 \frac{NT}{\beta}) \nu \left\{ \log \left\{ \rho_2 \left\{ \exp \left[\frac{C}{1+\zeta_2} \frac{NT}{\beta} \zeta_2 \rho_2(m') \right] \right\} \right\} \right\} \\ &+ \zeta_4 \frac{NT}{\beta} \nu_3 \left\{ \log \left[\rho_4 \left\{ \exp \left[C \frac{\beta}{NT\zeta_4} [\zeta_3\nu_3\rho_1 + \zeta_4\rho_4](m') \right] \right\} \right] \right\} \\ &+ (1 + \zeta_2) \mathcal{K} \left[\nu, \mu_{\exp\left(-\frac{(\gamma-\lambda_0)\rho_0(r)+\lambda_0\rho_2(r)}{1+\zeta_2}\right)} \right] \\ &\quad - \nu[\mathcal{K}(\rho_2, \rho_0)] + \left(1 + \sum_{i=1}^5 \zeta_i\right) \log\left(\frac{2}{\epsilon}\right). \end{aligned}$$

In the case when we want to select a single model $\hat{m}(\omega)$, and therefore to set $\nu = \delta_{\hat{m}}$, the previous inequality engages us to take

$$\hat{m} \in \arg \min_{m \in M} (\gamma - \lambda_0)\rho_0(m, r) + \lambda_0\rho_2(m, r).$$

In parametric situations where $\pi_{\exp(-\lambda r)}(r) \simeq r^*(m) + \frac{d_e(m)}{\lambda}$, we get

$$(\gamma - \lambda_0)\rho_0(m, r) - \lambda_0\rho_2(m, r) \simeq \gamma \left[r^*(m) + d_e(m) \left(\frac{1}{\lambda_0} + \frac{\lambda_0 - \lambda_2}{\gamma\lambda_2} \right) \right],$$

resulting in a linear penalization of the empirical dimension of the models.

1.5.8. Analysis of the two step relative bound. We will not state a formal result, but will nevertheless give some hints about how to establish one. We should start from Theorem 1.25, which gives a deterministic variance term. From Theorem 1.25, after a change of prior distribution, we obtain for any positive constants α_1 and α_2 , any prior distributions $\tilde{\mu}_1$ and $\tilde{\mu}_2 \in \mathcal{M}_+^1(M)$, for any prior conditional distributions $\tilde{\pi}_1$ and $\tilde{\pi}_2 : M \rightarrow \mathcal{M}_+^1(\Theta)$, with \mathbb{P} probability at least $1 - \eta$, for any posterior distributions $\nu_1\rho_1$ and $\nu_2\rho_2$,

$$\begin{aligned} \alpha_1(\nu_1\rho_1 - \nu_2\rho_2)(R) &\leq \alpha_2(\nu_1\rho_1 - \nu_2\rho_2)(r) \\ &\quad + \mathcal{K}[(\nu_1\rho_1) \otimes (\nu_2\rho_2), (\tilde{\mu}_1 \tilde{\pi}_1) \otimes (\tilde{\mu}_2 \tilde{\pi}_2)] \\ &\quad + \log \left\{ (\tilde{\mu}_1 \tilde{\pi}_1) \otimes (\tilde{\mu}_2 \tilde{\pi}_2) \left[\exp \left\{ -\alpha_2 \Psi_{\frac{\alpha_2}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} - \log(\eta). \end{aligned}$$

Applying this to $\alpha_1 = 0$, we get that

$$\begin{aligned} (\nu\rho - \nu_3\rho_1)(r) &\leq \frac{1}{\alpha_2} \left[\mathcal{K}[(\nu\rho) \otimes (\nu_3\rho_1), (\tilde{\mu} \tilde{\pi}) \otimes (\tilde{\mu}_3 \tilde{\pi}_1)] \right. \\ &\quad \left. + \log \left\{ (\tilde{\mu} \tilde{\nu}) \otimes (\tilde{\mu}_3 \tilde{\pi}_1) \left[\exp \left\{ \alpha_2 \Psi_{-\frac{\alpha_2}{N}}(R', M') \right\} \right] \right\} \right] - \log(\eta). \end{aligned}$$

In the same way, to bound quantities of the form

$$\begin{aligned} &\log \left\{ \nu_3 \left[\rho_1 \left\{ \exp \left[C_1(\nu\rho + \zeta_1\rho_1)(m') \right] \right\}^{p_1} \right. \right. \\ &\quad \left. \left. \times \rho_4 \left\{ \exp \left[C_2[\zeta_3\nu_3\rho_1 + \zeta_5\rho_4](m') \right] \right\}^{p_2} \right] \right\} \\ &= \sup_{\nu_5} \left\{ p_1 \sup_{\rho_5} \left\{ C_1[(\nu\rho) \otimes (\nu_5\rho_5) + \zeta_1\nu_5(\rho_1 \otimes \rho_5)](m') - \mathcal{K}(\rho_5, \rho_1) \right\} \right. \\ &\quad \left. + p_2 \sup_{\rho_6} \left\{ C_2[\zeta_3(\nu_3\rho_1) \otimes (\nu_5\rho_6) \right. \right. \\ &\quad \left. \left. + \zeta_5\nu_5(\rho_4 \otimes \rho_6)](m') - \mathcal{K}(\rho_6, \rho_4) \right\} - \mathcal{K}(\nu_5, \nu_3) \right\}, \end{aligned}$$

where C_1, C_2, p_1 and p_2 are positive constants, and similar terms, we need to use inequalities of the type: for any prior distributions $\tilde{\mu}_i \tilde{\pi}_i, i = 1, 2$, with \mathbb{P} probability at least $1 - \eta$, for any posterior distributions $\nu_i \rho_i, i = 1, 2$,

$$\begin{aligned} \alpha_3(\nu_1\rho_1) \otimes (\nu_2\rho_2)(m') &\leq \log \left\{ (\tilde{\mu}_1 \tilde{\pi}_1) \otimes (\tilde{\mu}_2 \tilde{\pi}_2) \exp \left[\alpha_3 \Phi_{-\frac{\alpha_3}{N}}(M') \right] \right\} \\ &\quad + \mathcal{K}[(\nu_1\rho_1) \otimes (\nu_2\rho_2), (\tilde{\mu}_1 \tilde{\pi}_1) \otimes (\tilde{\mu}_2 \tilde{\pi}_2)] - \log(\eta). \end{aligned}$$

We need also the variant: with \mathbb{P} probability at least $1 - \eta$, for any posterior distribution $\nu_1 : \Omega \rightarrow \mathcal{M}_+^1(M)$ and any conditional posterior distributions $\rho_1, \rho_2 : \Omega \times M \rightarrow \mathcal{M}_+^1(\Theta)$,

$$\begin{aligned} \alpha_3\nu_1(\rho_1 \otimes \rho_2)(m') &\leq \log \left\{ \tilde{\mu}_1(\tilde{\pi}_1 \otimes \tilde{\pi}_2) \exp \left[\alpha_3 \Phi_{-\frac{\alpha_3}{N}}(M') \right] \right\} \\ &\quad + \mathcal{K}(\nu_1, \tilde{\mu}_1) + \nu_1 \left\{ \mathcal{K}[\rho_1 \otimes \rho_2, \tilde{\pi}_1 \otimes \tilde{\pi}_2] \right\} - \log(\eta). \end{aligned}$$

We deduce that

$$\begin{aligned}
& \log \left\{ \nu_3 \left[\rho_1 \left\{ \exp \left[C_1(\nu\rho + \zeta_1\rho_1)(m') \right] \right\}^{p_1} \right. \right. \\
& \qquad \qquad \qquad \left. \left. \times \rho_4 \left\{ \exp \left[C_2[\zeta_3\nu_3\rho_1 + \zeta_5\rho_4](m') \right] \right\}^{p_2} \right] \right\} \\
& \leq \sup_{\nu_5} \left\{ p_1 \sup_{\rho_5} \left[\frac{C_1}{\alpha_3} \left\{ \log \left\{ (\tilde{\mu} \tilde{\pi}) \otimes (\tilde{\mu}_5 \tilde{\pi}_5) \exp \left[\alpha_3 \Phi_{-\frac{\alpha_3}{N}}(M') \right] \right\} \right. \right. \right. \\
& \qquad \qquad \qquad \left. \left. + \mathcal{K}[(\nu\rho) \otimes (\nu_5\rho_5), (\tilde{\mu} \tilde{\pi} \otimes (\tilde{\mu}_5 \tilde{\pi}_5))] + \log\left(\frac{2}{\eta}\right) \right. \right. \\
& \qquad \qquad \qquad \left. \left. + \zeta_1 \left[\log \left\{ \tilde{\mu}_5(\tilde{\pi}_1 \otimes \tilde{\pi}_5) \exp \left[\alpha_3 \Phi_{-\frac{\alpha_3}{N}}(M') \right] \right\} \right. \right. \right. \\
& \left. \left. + \mathcal{K}(\nu_5, \tilde{\mu}_5) + \nu_5 \{ \mathcal{K}[\rho_1 \otimes \rho_5, \tilde{\pi}_1 \otimes \tilde{\pi}_5] \} + \log\left(\frac{2}{\eta}\right) \right] \right\} - \mathcal{K}(\rho_5, \rho_1) \\
& \quad + p_2 \sup_{\rho_6} \left[\frac{C_1}{\alpha_3} \left\{ \log \left\{ (\tilde{\mu}_3 \tilde{\pi}_1) \otimes (\tilde{\mu}_5 \tilde{\pi}_6) \exp \left[\alpha_3 \Phi_{-\frac{\alpha_3}{N}}(M') \right] \right\} \right. \right. \\
& \qquad \qquad \qquad \left. \left. + \mathcal{K}[(\nu_3\rho_1) \otimes (\nu_5\rho_6), (\tilde{\mu}_3 \tilde{\pi}_1 \otimes (\tilde{\mu}_5 \tilde{\pi}_6))] + \log\left(\frac{2}{\eta}\right) \right. \right. \\
& \qquad \qquad \qquad \left. \left. + \zeta_1 \left[\log \left\{ \tilde{\mu}_5(\tilde{\pi}_4 \otimes \tilde{\pi}_6) \exp \left[\alpha_3 \Phi_{-\frac{\alpha_3}{N}}(M') \right] \right\} \right. \right. \right. \\
& \qquad \qquad \qquad \left. \left. + \mathcal{K}(\nu_5, \tilde{\mu}_5) + \nu_5 \{ \mathcal{K}[\rho_4 \otimes \rho_6, \tilde{\pi}_4 \otimes \tilde{\pi}_6] \} + \log\left(\frac{2}{\eta}\right) \right] \right\} \\
& \qquad \qquad \qquad \left. - \mathcal{K}(\rho_6, \rho_4) \right] - \mathcal{K}(\nu_5, \nu_3) \Big\}.
\end{aligned}$$

We are then left with the need to bound entropy terms like $\mathcal{K}(\nu_3\rho_1, \tilde{\mu}_3\tilde{\pi}_1)$, where we have the choice of $\tilde{\mu}_3$ and $\tilde{\pi}_1$, to obtain a useful bound. As could be expected, we decompose it into

$$\mathcal{K}(\nu_3\rho_1, \tilde{\mu}_3\tilde{\pi}_1) = \mathcal{K}(\nu_3, \tilde{\mu}_3) + \nu_3[\mathcal{K}(\rho_1, \tilde{\pi}_1)].$$

Let us look after the second term first, choosing $\tilde{\pi}_1 = \pi_{\exp(-\beta_1 R)}$:

$$\begin{aligned}
\nu_3[\mathcal{K}(\rho_1, \tilde{\pi}_1)] &= \nu_3[\beta_1(\rho_1 - \tilde{\pi}_1)(R) + \mathcal{K}(\rho_1, \pi) - \mathcal{K}(\tilde{\pi}_1, \pi)] \\
&\leq \frac{\beta_1}{\alpha_1} \left[\alpha_2\nu_3(\rho_1 - \tilde{\pi}_1)(r) + \mathcal{K}(\nu_3, \tilde{\mu}_3) + \nu_3[\mathcal{K}(\rho_1, \tilde{\pi}_1)] \right] \\
&\quad + \log \left\{ \tilde{\mu}_3(\tilde{\pi}_1^{\otimes 2}) \left[\exp \left\{ -\alpha_2 \Psi_{\frac{\alpha_2}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} - \log(\eta)
\end{aligned}$$

$$\begin{aligned}
& + \nu_3 [\mathcal{K}(\rho_1, \pi) - \mathcal{K}(\tilde{\pi}_1, \pi)] \\
\leq & \frac{\beta_1}{\alpha_1} \left[\mathcal{K}(\nu_3, \tilde{\mu}_3) + \nu_3 [\mathcal{K}(\rho_1, \tilde{\pi}_1)] \right. \\
& \left. + \log \left\{ \tilde{\mu}_3 (\tilde{\pi}_1^{\otimes 2}) \left[\exp \left\{ -\alpha_2 \Psi_{\frac{\alpha_2}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} - \log(\eta) \right] \\
& + \nu_3 \left\{ \mathcal{K} \left[\rho_1, \pi_{\exp(-\frac{\beta_1 \alpha_2}{\alpha_1} r)} \right] \right\}.
\end{aligned}$$

Thus, when the constraint $\lambda_1 = \frac{\beta_1 \alpha_2}{\alpha_1}$ is satisfied,

$$\begin{aligned}
\nu_3 [\mathcal{K}(\rho_1, \tilde{\pi}_1)] & \leq \left(1 - \frac{\beta_1}{\alpha_1} \right)^{-1} \frac{\beta_1}{\alpha_1} \left[\mathcal{K}(\nu_3, \tilde{\mu}_3) \right. \\
& \left. + \log \left\{ \tilde{\mu}_3 (\tilde{\pi}_1^{\otimes 2}) \left[\exp \left\{ -\alpha_2 \Psi_{\frac{\alpha_2}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} - \log(\eta) \right].
\end{aligned}$$

We can further specialize the constants, choosing $\alpha_1 = N \sinh(\frac{\alpha_2}{N})$, so that

$$-\alpha_2 \Psi_{\frac{\alpha_2}{N}}(R', M') + \alpha_1 R' \leq 2N \sinh\left(\frac{\alpha_2}{2N}\right)^2 M'.$$

We can for instance choose $\alpha_2 = \gamma$, $\alpha_1 = N \sinh(\frac{\gamma}{N})$, and $\beta_1 = \lambda_1 \frac{N}{\gamma} \sinh(\frac{\gamma}{N})$, leading to

PROPOSITION 1.60. *With the notations of Theorem 1.59, the constants being set as explained above, putting $\tilde{\pi}_1 = \pi_{\exp(-\lambda_1 \frac{N}{\gamma} \sinh(\frac{\gamma}{N}) R)}$, with \mathbb{P} probability at least $1 - \eta$,*

$$\begin{aligned}
\nu_3 [\mathcal{K}(\rho_1, \tilde{\pi}_1)] & \leq \left(1 - \frac{\lambda_1}{\gamma} \right)^{-1} \frac{\lambda_1}{\gamma} \left[\mathcal{K}(\nu_3, \tilde{\mu}_3) \right. \\
& \left. + \log \left\{ \tilde{\mu}_3 (\tilde{\pi}_1^{\otimes 2}) \left[\exp \left\{ 2N \sinh\left(\frac{\gamma}{2N}\right)^2 M' \right\} \right] \right\} - \log(\eta) \right].
\end{aligned}$$

More generally

$$\begin{aligned}
\nu_3 [\mathcal{K}(\rho, \tilde{\pi}_1)] & \leq \left(1 - \frac{\lambda_1}{\gamma} \right)^{-1} \frac{\lambda_1}{\gamma} \left[\mathcal{K}(\nu_3, \tilde{\mu}_3) \right. \\
& \left. + \log \left\{ \tilde{\mu}_3 (\tilde{\pi}_1^{\otimes 2}) \left[\exp \left\{ 2N \sinh\left(\frac{\gamma}{2N}\right)^2 M' \right\} \right] \right\} - \log(\eta) \right] \\
& + \left(1 - \frac{\lambda_1}{\gamma} \right)^{-1} \nu_3 [\mathcal{K}(\rho, \rho_1)].
\end{aligned}$$

In a similar way, let us choose now $\tilde{\mu}_3 = \mu_{\exp[-\alpha_3\bar{\pi}(R)]}$. We can write

$$\begin{aligned} \mathcal{K}(\nu, \tilde{\mu}_3) &= \alpha_3(\nu - \tilde{\mu}_3)\bar{\pi}(R) + \mathcal{K}(\nu, \mu) - \mathcal{K}(\tilde{\mu}_3, \mu) \\ &\leq \frac{\alpha_3}{\alpha_1} \left[\alpha_2(\nu - \tilde{\mu}_3)\bar{\pi}(r) + \mathcal{K}(\nu, \tilde{\mu}_3) \right. \\ &\quad \left. + \log \left\{ (\tilde{\mu}_3\bar{\pi}) \otimes (\tilde{\mu}_3\bar{\pi}) \left[\exp \left\{ -\alpha_2 \Psi_{\frac{\alpha_2}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} - \log(\eta) \right] \\ &\quad + \mathcal{K}(\nu, \mu) - \mathcal{K}(\tilde{\mu}_3, \mu). \end{aligned}$$

Let us choose $\alpha_2 = \gamma$, $\alpha_1 = N \sinh(\frac{\gamma}{N})$, and let us add some other entropy inequalities to get rid of $\bar{\pi}$ in a suitable way, the approach of entropy compensation being quite the same as the one used to obtain the empirical bound of Theorem 1.59. This results with \mathbb{P} probability at least $1 - \eta$ in

$$\begin{aligned} \left(1 - \frac{\alpha_3}{\alpha_1}\right) \mathcal{K}(\nu, \tilde{\mu}_3) &\leq \frac{\alpha_3}{\alpha_1} \left[\gamma(\nu - \tilde{\mu}_3)\bar{\pi}(r) \right. \\ &\quad \left. + \log \left\{ (\tilde{\mu}_3\bar{\pi}) \otimes (\tilde{\mu}_3\bar{\pi}) \left[\exp \left\{ -\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} + \log\left(\frac{2}{\eta}\right) \right] \\ &\quad + \mathcal{K}(\nu, \mu) - \mathcal{K}(\tilde{\mu}_3, \mu), \\ \zeta_6 \left(1 - \frac{\beta}{\alpha_1}\right) \tilde{\mu}_3 [\mathcal{K}(\rho_6, \bar{\pi})] &\leq \zeta_6 \frac{\beta}{\alpha_1} \left[\gamma \tilde{\mu}_3(\rho_6 - \bar{\pi})(r) \right. \\ &\quad \left. + \log \left\{ \tilde{\mu}_3(\bar{\pi}^{\otimes 2}) \left[\exp \left\{ -\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} + \log\left(\frac{2}{\eta}\right) \right] \\ &\quad + \zeta_6 \tilde{\mu}_3 [\mathcal{K}(\rho_6, \pi) - \mathcal{K}(\bar{\pi}, \pi)], \\ \zeta_7 \left(1 - \frac{\beta}{\alpha_1}\right) \tilde{\mu}_3 [\mathcal{K}(\rho_7, \bar{\pi})] &\leq \zeta_7 \frac{\beta}{\alpha_1} \left[\gamma \tilde{\mu}_3(\rho_7 - \bar{\pi})(r) \right. \\ &\quad \left. + \log \left\{ \tilde{\mu}_3(\bar{\pi}^{\otimes 2}) \left[\exp \left\{ -\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} + \log\left(\frac{2}{\eta}\right) \right] \\ &\quad + \zeta_7 \tilde{\mu}_3 [\mathcal{K}(\rho_7, \pi) - \mathcal{K}(\bar{\pi}, \pi)], \\ \zeta_8 \left(1 - \frac{\beta}{\alpha_1}\right) \nu [\mathcal{K}(\rho_8, \bar{\pi})] &\leq \zeta_8 \frac{\beta}{\alpha_1} \left[\gamma \nu(\rho_8 - \bar{\pi})(r) + \mathcal{K}(\nu, \tilde{\mu}_3) \right. \\ &\quad \left. + \log \left\{ \tilde{\mu}_3(\bar{\pi}^{\otimes 2}) \left[\exp \left\{ -\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} + \log\left(\frac{2}{\eta}\right) \right] \\ &\quad + \zeta_8 \nu [\mathcal{K}(\rho_8, \pi) - \mathcal{K}(\bar{\pi}, \pi)], \\ \zeta_9 \left(1 - \frac{\beta}{\alpha_1}\right) \nu [\mathcal{K}(\rho_9, \bar{\pi})] &\leq \zeta_9 \frac{\beta}{\alpha_1} \left[\gamma \nu(\rho_9 - \bar{\pi})(r) + \mathcal{K}(\nu, \tilde{\mu}_3) \right. \end{aligned}$$

$$\begin{aligned}
& + \log \left\{ \tilde{\mu}_3(\bar{\pi}^{\otimes 2}) \left[\exp \left\{ -\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} + \log \left(\frac{2}{\eta} \right) \\
& \qquad \qquad \qquad + \zeta_9 \nu [\mathcal{K}(\rho_9, \pi) - \mathcal{K}(\bar{\pi}, \pi)],
\end{aligned}$$

where we have introduced a bunch of constants, assumed to be positive, that we will more precisely set to

$$\begin{aligned}
x_8 + x_9 &= 1, \\
(\zeta_6 \beta + x_8 \alpha_3) \frac{\gamma}{\alpha_1} &= \lambda_6, \\
(\zeta_7 \beta + x_9 \alpha_3) \frac{\gamma}{\alpha_1} &= \lambda_7, \\
(\zeta_8 \beta - x_8 \alpha_3) \frac{\gamma}{\alpha_1} &= \lambda_8, \\
(\zeta_9 \beta - x_9 \alpha_3) \frac{\gamma}{\alpha_1} &= \lambda_9.
\end{aligned}$$

We get with \mathbb{P} probability at least $1 - \eta$,

$$\begin{aligned}
& \left(1 - \frac{\alpha_3}{\alpha_1} - (\zeta_8 + \zeta_9) \frac{\beta}{\alpha_1} \right) \mathcal{K}(\nu, \tilde{\mu}_3) \leq \\
& \quad \frac{\alpha_3}{\alpha_1} \left[\gamma [\nu(x_8 \rho_8 + x_9 \rho_9)(r) - \tilde{\mu}_3(x_8 \rho_6 + x_9 \rho_7)(r)] \right. \\
& \quad \left. + \frac{\alpha_3}{\alpha_1} \log \left\{ (\tilde{\mu}_3 \bar{\pi}) \otimes (\tilde{\mu}_3 \bar{\pi}) \left[\exp \left\{ -\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} \right. \\
& \quad \left. + (\zeta_6 + \zeta_7 + \zeta_8 + \zeta_9) \frac{\beta}{\alpha_1} \log \left\{ \tilde{\mu}_3(\bar{\pi}^{\otimes 2}) \left[\exp \left\{ -\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R' \right\} \right] \right\} \right. \\
& \quad \left. + \mathcal{K}(\nu, \mu) - \mathcal{K}(\tilde{\mu}_3, \mu) + \left(\frac{\alpha_3}{\alpha_1} + (\zeta_6 + \zeta_7 + \zeta_8 + \zeta_9) \frac{\beta}{\alpha_1} \right) \log \left(\frac{2}{\eta} \right) \right].
\end{aligned}$$

Let us choose the constants so that $\lambda_1 = \lambda_7 = \lambda_9$, $\lambda_4 = \lambda_6 = \lambda_8$, $\alpha_3 x_9 \frac{\gamma}{\alpha_1} = \xi_1$ and $\alpha_3 x_8 \frac{\gamma}{\alpha_1} = \xi_4$. This is done by setting

$$\begin{aligned}
x_8 &= \frac{\xi_4}{\xi_1 + \xi_4}, \\
x_9 &= \frac{\xi_1}{\xi_1 + \xi_4}, \\
\alpha_3 &= \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) (\xi_1 + \xi_4), \\
\zeta_6 &= \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) \frac{(\lambda_4 - \xi_4)}{\beta}, \\
\zeta_7 &= \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) \frac{(\lambda_1 - \xi_1)}{\beta},
\end{aligned}$$

$$\zeta_8 = \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) \frac{(\lambda_4 + \xi_4)}{\beta},$$

$$\zeta_9 = \frac{N}{\gamma} \sinh\left(\frac{\gamma}{N}\right) \frac{(\lambda_1 + \xi_1)}{\beta}.$$

The inequality $\lambda_1 > \xi_1$ is always satisfied. The inequality $\lambda_4 > \xi_4$ is required for the above choice of constants, and will be satisfied for a suitable choice of ζ_3 and ζ_4 .

Under these assumptions, we obtain with \mathbb{P} probability at least $1 - \eta$

$$\begin{aligned} \left(1 - \frac{\alpha_3}{\alpha_1} - (\zeta_8 + \zeta_9) \frac{\beta}{\alpha_1}\right) \mathcal{K}(\nu, \tilde{\mu}_3) &\leq (\nu - \tilde{\mu}_3)(\xi_1 \rho_1 + \xi_4 \rho_4)(r) \\ &+ \frac{\alpha_3}{\alpha_1} \log \left\{ (\tilde{\mu}_3 \bar{\pi}) \otimes (\tilde{\mu}_3 \bar{\pi}) \left[\exp\{-\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R'\} \right] \right\} \\ &+ (\zeta_6 + \zeta_7 + \zeta_8 + \zeta_9) \frac{\beta}{\alpha_1} \log \left\{ \tilde{\mu}_3 (\bar{\pi}^{\otimes 2}) \left[\exp\{-\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R'\} \right] \right\} \\ &+ \mathcal{K}(\nu, \mu) - \mathcal{K}(\tilde{\mu}_3, \mu) + \left(\frac{\alpha_3}{\alpha_1} + (\zeta_6 + \zeta_7 + \zeta_8 + \zeta_9) \frac{\beta}{\alpha_1} \right) \log\left(\frac{2}{\eta}\right). \end{aligned}$$

This proves

PROPOSITION 1.61. *The constants being set as explained above, with \mathbb{P} probability at least $1 - \eta$, for any posterior distribution $\nu : \Omega \rightarrow \mathcal{M}_+^1(M)$,*

$$\begin{aligned} \mathcal{K}(\nu, \tilde{\mu}_3) &\leq \left(1 - \frac{\alpha_3}{\alpha_1} - (\zeta_8 + \zeta_9) \frac{\beta}{\alpha_1}\right)^{-1} \left[\mathcal{K}(\nu, \nu_3) \right. \\ &+ \frac{\alpha_3}{\alpha_1} \log \left\{ (\tilde{\mu}_3 \bar{\pi}) \otimes (\tilde{\mu}_3 \bar{\pi}) \left[\exp\{-\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R'\} \right] \right\} \\ &+ (\zeta_6 + \zeta_7 + \zeta_8 + \zeta_9) \frac{\beta}{\alpha_1} \log \left\{ \tilde{\mu}_3 (\bar{\pi}^{\otimes 2}) \left[\exp\{-\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R'\} \right] \right\} \\ &\left. + \left(\frac{\alpha_3}{\alpha_1} + (\zeta_6 + \zeta_7 + \zeta_8 + \zeta_9) \frac{\beta}{\alpha_1} \right) \log\left(\frac{2}{\eta}\right) \right]. \end{aligned}$$

Thus

$$\begin{aligned} \mathcal{K}(\nu_3 \rho_1, \tilde{\mu}_3 \tilde{\pi}_1) &\leq \frac{1 + \left(1 - \frac{\lambda_1}{\gamma}\right)^{-1} \frac{\lambda_1}{\gamma}}{1 - \frac{\alpha_3}{\alpha_1} - (\zeta_8 + \zeta_9) \frac{\beta}{\alpha_1}} \\ &\times \left[\frac{\alpha_3}{\alpha_1} \log \left\{ (\tilde{\mu}_3 \bar{\pi}) \otimes (\tilde{\mu}_3 \bar{\pi}) \left[\exp\{-\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R'\} \right] \right\} \right. \\ &\left. + (\zeta_6 + \zeta_7 + \zeta_8 + \zeta_9) \frac{\beta}{\alpha_1} \log \left\{ \tilde{\mu}_3 (\bar{\pi}^{\otimes 2}) \left[\exp\{-\gamma \Psi_{\frac{\gamma}{N}}(R', M') + \alpha_1 R'\} \right] \right\} \right] \end{aligned}$$

$$\begin{aligned}
& + \left(\frac{\alpha_3}{\alpha_1} + (\zeta_6 + \zeta_7 + \zeta_8 + \zeta_9) \frac{\beta}{\alpha_1} \right) \log\left(\frac{2}{\eta}\right) \Big] \\
& + \left(1 - \frac{\lambda_1}{\gamma}\right)^{-1} \frac{\lambda_1}{\gamma} \left[\log \left\{ \tilde{\mu}_3(\tilde{\pi}_1^{\otimes 2}) \left[\exp\{2N \sinh\left(\frac{\gamma}{2N}\right)^2 M'\} \right] \right\} - \log\left(\frac{2}{\eta}\right) \right].
\end{aligned}$$

We will not go further, lest it may become tedious, but we hope we have given sufficient hints to state informally that the bound $B(\nu, \rho, \beta)$ of Theorem 1.59 is upper bounded with \mathbb{P} probability close to one by a bound of the same flavour where the empirical quantities r and m' have been replaced with their expectations R and M' .

2. TRANSDUCTIVE PAC-BAYESIAN LEARNING

2.1. BASIC INEQUALITIES. In this section the observed sample $(X_i, Y_i)_{i=1}^N$ will be supplemented with a *shadow sample* $(X_i, Y_i)_{i=N+1}^{(k+1)N}$. This point of view, called *transductive classification*, has been introduced by V. Vapnik. It may be justified in different ways.

On the practical side, one interest of the transductive setting is that it is often a lot easier to collect examples than it is to label them, so that it is not unrealistic to assume that we indeed have two training samples, one labelled and one unlabelled. It also covers the case when a batch of patterns is to be classified and we are allowed to observe the whole batch before issuing the classification.

On the mathematical side, considering a shadow sample proves technically fruitful. Indeed, when introducing the VC entropy and VC dimension concepts, as well as when dealing with compression schemes, albeit the *inductive* setting is our final concern, the transductive setting is a useful detour. In this second scenario, intermediate technical results involving the shadow sample are integrated with respect to unobserved random variables in a second stage of the proofs.

Let us describe now the changes to be made to previous notations to adapt them to the transductive setting. The distribution \mathbb{P} will be a probability measure on the canonical space $\Omega = (\mathcal{X} \times \mathcal{Y})^{(k+1)N}$, and $(X_i, Y_i)_{i=1}^{(k+1)N}$ will be the canonical process on this space (that is the coordinate process). Unless explicitly mentioned, the parameter k indicating the size of the shadow sample will remain fixed. Assuming the shadow sample size is a multiple of the training sample size is convenient without significantly restricting the generality. For a while, we will use a weaker assumption than independence,

assuming that \mathbb{P} is *partially exchangeable*, since this is all what we need in the proofs.

DEFINITION 2.1. For $i = 1, \dots, N$, let $\tau_i : \Omega \rightarrow \Omega$ be defined for any $\omega = (\omega_j)_{j=1}^{(k+1)N} \in \Omega$ by

$$\begin{cases} \tau_i(\omega)_{i+jN} = \omega_{i+(j-1)N}, & j = 1, \dots, k, \\ \tau_i(\omega)_i = \omega_{i+kN}, \\ \text{and } \tau_i(\omega)_{m+jN} = \omega_{m+jN}, & m \neq i, m = 1, \dots, N, j = 0, \dots, k. \end{cases}$$

Clearly, if we arrange the $(k+1)N$ samples in a $N \times (k+1)$ array, τ_i performs a circular permutation of $k+1$ entries on the i th row, letting the other rows unchanged. Moreover, all the circular permutations of the i th row have the form τ_i^j , j ranging from 0 to k .

The probability distribution \mathbb{P} is said to be partially exchangeable if for any $i = 1, \dots, N$, $\mathbb{P} \circ \tau_i^{-1} = \mathbb{P}$.

This means equivalently that for any bounded measurable function $h : \Omega \rightarrow \mathbb{R}$, $\mathbb{P}(h \circ \tau_i) = \mathbb{P}(h)$.

In the same way a function h defined on Ω will be said to be partially exchangeable if $h \circ \tau_i = h$ for any $i = 1, \dots, N$. Accordingly a posterior distribution $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta, \mathcal{T})$ will be said to be partially exchangeable when $\rho(\omega, A) = \rho[\tau_i(\omega), A]$, for any $\omega \in \Omega$, any $i = 1, \dots, N$ and any $A \in \mathcal{T}$.

For any bounded measurable function h , let us define $T_i(h) = \frac{1}{k+1} \sum_{j=0}^k h \circ \tau_i^j$. Let $T(h) = T_N \circ \dots \circ T_1(h)$. For any partially exchangeable probability distribution \mathbb{P} , and for any bounded measurable function h , $\mathbb{P}[T(h)] = \mathbb{P}(h)$. Let us put

$$\begin{aligned} \sigma_i(\theta) &= \mathbb{1}[f_\theta(X_i) \neq Y_i], && \text{indicating the success or failure of } f_\theta \\ &&& \text{to predict } Y_i \text{ from } X_i, \\ r_1(\theta) &= \frac{1}{N} \sum_{i=1}^N \sigma_i(\theta), && \text{the empirical error rate of } f_\theta \\ &&& \text{on the observed sample,} \\ r_2(\theta) &= \frac{1}{kN} \sum_{i=N+1}^{(k+1)N} \sigma_i(\theta), && \text{the error rate of } f_\theta \text{ on the shadow sample,} \\ \bar{r}(\theta) &= \frac{r_1(\theta) + kr_2(\theta)}{k+1} = \frac{1}{(k+1)N} \sum_{i=1}^{(k+1)N} \sigma_i(\theta), && \text{the global error} \\ &&& \text{rate of } f_\theta, \end{aligned}$$

$R_i(\theta) = \mathbb{P}[f_\theta(X_i) \neq Y_i]$, the expected error
rate of f_θ on the i th input,

$$R(\theta) = \frac{1}{N} \sum_{i=1}^N R_i(\theta) = \mathbb{P}[r_1(\theta)] = \mathbb{P}[r_2(\theta)], \quad \text{the average expected}$$

error rate of f_θ on all inputs.

We will allow for posterior distributions $\rho : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$ depending on the shadow sample. The most interesting ones will anyhow be independent of the shadow labels $Y_{N+1}, \dots, Y_{(k+1)N}$. We will be interested in the conditional expected error rate of the randomized classification rule described by ρ on the shadow sample, given the observed sample, which reads as $\mathbb{P}[\rho(r_2)|(X_i, Y_i)_{i=1}^N]$.

Let us comment on the case when \mathbb{P} is invariant by any permutations of the rows, meaning that

$$\mathbb{P}[h(\omega \circ s)] = \mathbb{P}[h(\omega)] \quad \text{for all } s \in \mathfrak{S}(\{i + jN; j = 0, \dots, k\})$$

and all $i = 1, \dots, N$ (where $\mathfrak{S}(A)$ is the set of permutations of A , extended to $\{1, \dots, (k+1)N\}$ so as to be the identity outside of A). In this case, if ρ is invariant by permutations of the rows of the shadow sample, meaning that $\rho(\omega \circ s) = \rho(\omega) \in \mathcal{M}_+^1(\Theta)$, $s \in \mathfrak{S}(\{i + jN; j = 1, \dots, k\})$, $i = 1, \dots, N$, then $\mathbb{P}[\rho(r_2)|(X_i, Y_i)_{i=1}^N] = \frac{1}{N} \sum_{i=1}^N \mathbb{P}[\rho(\sigma_{i+N})|(X_i, Y_i)_{i=1}^N]$, meaning that the expectation can be taken on a restricted shadow sample of the same size as the observed sample. If moreover the rows are equidistributed (meaning that their marginal distributions are equal), then

$$\mathbb{P}[\rho(r_2)|(X_i, Y_i)_{i=1}^N] = \mathbb{P}[\rho(\sigma_{N+1})|(X_i, Y_i)_{i=1}^N].$$

This means that under these quite commonly fulfilled assumptions, the expectation can be taken on a single new object to be classified, our study thus covers the case when only one of the patterns from the shadow sample is to be labelled and one is interested in the expected error rate of this single labelling. Of course, in the case when \mathbb{P} is i.i.d. and ρ depends only on the training sample $(X_i, Y_i)_{i=1}^N$, we fall back on the usual criterion of performance $\mathbb{P}[\rho(r_2)|(Z_i)_{i=1}^N] = \rho(R) = \rho(R_1)$.

Let us recall the notation $\Phi_a(p) = -a^{-1} \log\{1 - p[1 - \exp(-a)]\}$.

Using an obvious factorization, and considering for the moment a fixed value of θ and any partially exchangeable positive real measurable function $\lambda : \Omega \rightarrow \mathbb{R}_+$, we can compute the log Laplace transform of r_1 under T , which acts like a conditional probability distribution:

$$\log\left\{T[\exp(-\lambda r_1)]\right\} = \sum_{i=1}^N \log\left\{T_i[\exp(-\frac{\lambda}{N}\sigma_i)]\right\}$$

$$\leq N \log \left\{ \frac{1}{N} \sum_{i=1}^N T_i \left[\exp \left(-\frac{\lambda}{N} \sigma_i \right) \right] \right\} = -\lambda \Phi_{\frac{\lambda}{N}}(\bar{r}).$$

Remarking that $T \left\{ \exp \left[\lambda \left[\Phi_{\frac{\lambda}{N}}(\bar{r}) - r_1 \right] \right] \right\} = \exp \left[\lambda \Phi_{\frac{\lambda}{N}}(\bar{r}) \right] T \left[\exp(-\lambda r_1) \right]$ we obtain

LEMMA 2.1. *For any $\theta \in \Theta$ and any partially exchangeable positive real measurable function $\lambda : \Omega \rightarrow \mathbb{R}_+$,*

$$T \left\{ \exp \left[\lambda \left\{ \Phi_{\frac{\lambda}{N}}[\bar{r}(\theta)] - r_1(\theta) \right\} \right] \right\} \leq 1.$$

We deduce from this lemma a result analogous to the inductive case:

THEOREM 2.2. *For any partially exchangeable positive real measurable function $\lambda : \Omega \times \Theta \rightarrow \mathbb{R}_+$, for any partially exchangeable posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\mathbb{P} \left\{ \exp \left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \rho \left[\lambda \left[\Phi_{\frac{\lambda}{N}}(\bar{r}) - r_1 \right] \right] - \mathcal{K}(\rho, \pi) \right] \right\} \leq 1.$$

The proof is deduced from the previous lemma, using the fact that π is partially exchangeable :

$$\begin{aligned} & \mathbb{P} \left\{ \exp \left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \rho \left[\lambda \left[\Phi_{\frac{\lambda}{N}}(\bar{r}) - r_1 \right] \right] - \mathcal{K}(\rho, \pi) \right] \right\} \\ &= \mathbb{P} \left\{ \pi \left\{ \exp \left[\lambda \left[\Phi_{\frac{\lambda}{N}}(\bar{r}) - r_1 \right] \right] \right\} \right\} = \mathbb{P} \left\{ T \pi \left\{ \exp \left[\lambda \left[\Phi_{\frac{\lambda}{N}}(\bar{r}) - r_1 \right] \right] \right\} \right\} \\ &= \mathbb{P} \left\{ \pi \left\{ T \exp \left[\lambda \left[\Phi_{\frac{\lambda}{N}}(\bar{r}) - r_1 \right] \right] \right\} \right\} \leq 1. \end{aligned}$$

Introducing in the same way

$$\begin{aligned} m'(\theta, \theta') &= \frac{1}{N} \sum_{i=1}^N \left| \mathbb{1}[f_\theta(X_i) \neq Y_i] - \mathbb{1}[f_{\theta'}(X_i) \neq Y_i] \right| \\ \text{and } \bar{m}(\theta, \theta') &= \frac{1}{(k+1)N} \sum_{i=1}^{(k+1)N} \left| \mathbb{1}[f_\theta(X_i) \neq Y_i] - \mathbb{1}[f_{\theta'}(X_i) \neq Y_i] \right|, \end{aligned}$$

we could prove along the same line of reasoning

THEOREM 2.3. *For any real parameter λ , any $\tilde{\theta} \in \Theta$, any partially exchangeable posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\mathbb{P} \left\{ \exp \left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \lambda \left[\rho \left\{ \Psi_{\frac{\lambda}{N}} [\bar{r}(\cdot) - \bar{r}(\tilde{\theta}), \bar{m}(\cdot, \tilde{\theta})] \right\} - [\rho(r_1) - r_1(\tilde{\theta})] \right] - \mathcal{K}(\rho, \pi) \right] \right\} \leq 1.$$

THEOREM 2.4. *For any real constant γ , for any $\tilde{\theta} \in \Theta$, for any partially exchangeable posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\mathbb{P} \left\{ \exp \left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \left\{ -N \rho \left\{ \log \left[1 - \tanh\left(\frac{\gamma}{N}\right) [\bar{r}(\cdot) - \bar{r}(\tilde{\theta})] \right] \right\} - \gamma [\rho(r_1) - r_1(\tilde{\theta})] - N \log \left[\cosh\left(\frac{\gamma}{N}\right) \right] \rho [m'(\cdot, \tilde{\theta})] - \mathcal{K}(\rho, \pi) \right\} \right] \right\} \leq 1.$$

This last theorem can be generalized to give

THEOREM 2.5. *For any real constant γ , for any partially exchangeable posterior distributions $\pi^1, \pi^2 : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\mathbb{P} \left\{ \exp \left[\sup_{\rho_1, \rho_2 \in \mathcal{M}_+^1(\Theta)} \left\{ -N \log \left\{ 1 - \tanh\left(\frac{\gamma}{N}\right) [\rho_1(\bar{r}) - \rho_2(\bar{r})] \right\} - \gamma [\rho_1(r_1) - \rho_2(r_1)] - N \log \left[\cosh\left(\frac{\gamma}{N}\right) \right] \rho_1 \otimes \rho_2(m') - \mathcal{K}(\rho_1, \pi^1) - \mathcal{K}(\rho_2, \pi^2) \right\} \right] \right\} \leq 1.$$

To conclude this section, we see that the basic theorems of transductive PAC-Bayesian classification have exactly the same form as the basic inequalities of inductive classification, Theorems 1.4, 1.25 and 1.26 *with $R(\theta)$ replaced with $\bar{r}(\theta)$, $r(\theta)$ replaced with $r_1(\theta)$ and $M'(\theta, \tilde{\theta})$ replaced with $\bar{m}(\theta, \tilde{\theta})$.*

Thus all the results of the first section remain true under the hypotheses of transductive classification, with $R(\theta)$ replaced with $\bar{r}(\theta)$, $r(\theta)$ replaced with $r_1(\theta)$ and $M'(\theta, \tilde{\theta})$ replaced with $\bar{m}(\theta, \tilde{\theta})$.

Consequently, in the case when the unlabelled shadow sample is observed, it is possible to improve on Vapnik's bounds to be discussed hereafter by using an explicit partially exchangeable posterior distribution π and resorting

to localized or to relative bounds (in the case at least of unlimited computing resources, which of course may still be unrealistic in many real world situations, and with the caveat, to be recalled in the conclusion of this article, that for small sample sizes and comparatively complex classification models, the improvement may not be so decisive).

Let us notice also that the transductive setting when experimentally available, has the advantage that

$$\begin{aligned} \bar{d}(\theta, \theta') &= \frac{1}{(k+1)N} \sum_{i=1}^{(k+1)N} \mathbf{1}[f_{\theta'}(X_i) \neq f_{\theta}(X_i)] \\ &\geq \bar{m}(\theta, \theta') \geq \bar{r}(\theta) - \bar{r}(\theta'), \quad \theta, \theta' \in \Theta, \end{aligned}$$

is observable in this context, providing an empirical upper bound for the difference $\bar{r}(\hat{\theta}) - \rho(\bar{r})$ for any non randomized estimator $\hat{\theta}$ and any posterior distribution ρ , namely

$$\bar{r}(\hat{\theta}) \leq \rho(\bar{r}) + \rho[\bar{d}(\cdot, \hat{\theta})].$$

Thus in the setting of transductive statistical experiments, the PAC-Bayesian framework provides fully empirical bounds for the error rate of non randomized estimators $\hat{\theta} : \Omega \rightarrow \Theta$, even when using a non atomic prior π (or more generally a non atomic partially exchangeable posterior distribution π), when Θ is not a vector space and $\theta \mapsto R(\theta)$ cannot be proved to be convex on the support of some useful posterior distribution ρ .

2.2. VAPNIK'S BOUNDS FOR TRANSDUCTIVE CLASSIFICATION. In this section, we are going to stick to plain unlocalized non relative bounds. As we have already mentioned, (and as it was put forward by Vapnik himself in his seminal works), these bounds are not always superseded by the asymptotically better ones, and deserve all our efforts since they deal in many situations better with small samples.

2.2.1. *With a shadow sample of arbitrary size.* The great thing with the transductive setting is that we are manipulating only r_1 and \bar{r} which can take but a finite number of values and therefore are piecewise constant on Θ . To make use of this, let us consider for any value $\theta \in \Theta$ of the parameter the subset $\Delta(\theta) \subset \Theta$ of parameters θ' such that the classification rule $f_{\theta'}$ answers the same on the extended sample $(X_i)_{i=1}^{(k+1)N}$ as f_{θ} . Namely, let us put for any $\theta \in \Theta$

$$\Delta(\theta) = \{\theta' \in \Theta; f_{\theta'}(X_i) = f_{\theta}(X_i), i = 1, \dots, (k+1)N\}.$$

We see immediately that $\Delta(\theta)$ is an exchangeable parameter subset on which r_1 and r_2 (and therefore also \bar{r}) take a constant value. Thus for any $\theta \in \Theta$ we may consider the posterior ρ_θ defined by

$$\frac{d\rho_\theta}{d\pi}(\theta') = \mathbf{1}[\theta' \in \Delta(\theta)]\pi[\Delta(\theta)]^{-1},$$

and use the fact that $\rho_\theta(r_1) = r_1(\theta)$ and $\rho_\theta(\bar{r}) = \bar{r}(\theta)$, to prove that

LEMMA 2.6. *For any partially exchangeable positive real measurable function $\lambda : \Omega \times \Theta \rightarrow \mathbb{R}$ such that*

$$\lambda(\omega, \theta') = \lambda(\omega, \theta), \quad \theta \in \Theta, \theta' \in \Delta(\theta), \omega \in \Omega, \quad (2.1)$$

and any partially exchangeable posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,

$$\Phi_{\frac{\lambda}{N}}[\bar{r}(\theta)] + \frac{\log\{\epsilon\pi[\Delta(\theta)]\}}{\lambda(\theta)} \leq r_1(\theta).$$

We can then remark that for any value of λ independent of ω , the left-hand side of the previous inequality is a partially exchangeable function of $\omega \in \Omega$. Thus this left-hand side is maximized by some partially exchangeable function λ , namely

$$\arg \max_{\lambda} \Phi_{\frac{\lambda}{N}}[\bar{r}(\theta)] + \frac{\log\{\epsilon\pi[\Delta(\theta)]\}}{\lambda}$$

is partially exchangeable as depending only on partially exchangeable quantities. Moreover this choice of $\lambda(\omega, \theta)$ satisfies also condition (2.1) stated in the previous lemma of being constant on $\Delta(\theta)$, proving

LEMMA 2.7. *For any partially exchangeable posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$ and any $\lambda \in \mathbb{R}_+$,*

$$\Phi_{\frac{\lambda}{N}}[\bar{r}(\theta)] + \frac{\log\{\epsilon\pi[\Delta(\theta)]\}}{\lambda} \leq r_1(\theta).$$

Writing $\bar{r} = \frac{r_1 + kr_2}{k+1}$ and rearranging terms we obtain

THEOREM 2.8. *For any partially exchangeable posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,*

$$r_2(\theta) \leq \frac{k+1}{k} \inf_{\lambda \in \mathbb{R}_+} \frac{1 - \exp\left(-\frac{\lambda}{N}r_1(\theta) + \frac{\log\{\epsilon\pi[\Delta(\theta)]\}}{N}\right)}{1 - \exp\left(-\frac{\lambda}{N}\right)} - \frac{r_1(\theta)}{k}.$$

Let us remind the reader that in the case when we have a set of binary classification rules $\{f_\theta; \theta \in \Theta\}$ whose VC dimension is not greater than h , then we can choose π such that $\pi[\Delta(\theta)]$ is independent of θ and not less than $\left(\frac{h}{e(k+1)N}\right)^h$.

Another important case when the complexity term $-\log\{\pi[\Delta(\theta)]\}$ can easily be controlled is the setting of *compression schemes*, introduced by Littlestone et Warmuth [24]. In this case, we are given for each labelled subsample $(X_i, Y_i)_{i \in J}$, $J \subset \{1, \dots, N\}$, an estimator of the parameter

$$\hat{\theta}[(X_i, Y_i)_{i \in J}] = \hat{\theta}_J, \quad J \subset \{1, \dots, N\}, |J| \leq h,$$

where

$$\hat{\theta} : \bigsqcup_{k=1}^N (\mathcal{X} \times \mathcal{Y})^k \rightarrow \Theta$$

is an exchangeable function providing estimators for subsamples of arbitrary size. Let us assume that $\hat{\theta}$ is exchangeable, meaning that for any $k = 1, \dots, N$ and any permutation σ of $\{1, \dots, k\}$

$$\hat{\theta}[(x_i, y_i)_{i=1}^k] = \hat{\theta}[(x_{\sigma(i)}, y_{\sigma(i)})_{i=1}^k], \quad (x_i, y_i)_{i=1}^k \in (\mathcal{X} \times \mathcal{Y})^k.$$

In this situation, we can introduce the exchangeable subset

$$\left\{ \hat{\theta}_J; J \subset \{1, \dots, (k+1)N\}, |J| \leq h \right\} \subset \Theta,$$

which is seen to contain at most $\sum_{j=0}^h \binom{(k+1)N}{j} \leq \left(\frac{e(k+1)N}{h}\right)^h$ classification rules (as will be proved later on in Theorem 3.14 on page 138). Note that we had to extend the range of J to all the subsets of the extended sample, although we will use for estimation only those of the training sample, on which the labels are observed. Thus in this case also we can find a partially exchangeable posterior distribution π such that $\pi[\Delta(\hat{\theta}_J)] \geq \left(\frac{h}{e(k+1)N}\right)^h$.

We see that the size of the compression scheme plays the same role in this complexity bound as the VC dimension for VC classes.

In these two cases of binary classification with VC dimension not greater than h and compression schemes depending on a compression set with at most h points, we get a bound of

$$r_2(\theta) \leq \frac{k+1}{k} \inf_{\lambda \in \mathbb{R}_+} \frac{1 - \exp\left(-\frac{\lambda}{N} r_1(\theta) - \frac{h \log\left(\frac{\epsilon(k+1)N}{h}\right) - \log(\epsilon)}{N}\right)}{1 - \exp\left(-\frac{\lambda}{N}\right)} - \frac{r_1(\theta)}{k}.$$

Let us make some numerical application: when $N = 1000$, $h = 10$, $\epsilon = 0.01$, and $\inf_{\Theta} r_1 = r_1(\hat{\theta}) = 0.2$, we find that $r_2(\hat{\theta}) \leq 0.4093$, for k between 15 and 17, and values of λ equal respectively to 965, 968 and 971. For $k = 1$, we find only $r_2(\hat{\theta}) \leq 0.539$, showing the interest of allowing k to be larger than 1.

2.2.2. When the shadow sample has the same size as the training sample. In the case when $k = 1$, we can improve Theorem 2.2 by taking advantage of the fact that $T_i(\sigma_i)$ can take only 3 values, namely 0, 0.5 and 1. We see thus that $T_i(\sigma_i) - \Phi_{\frac{\lambda}{N}}[T_i(\sigma_i)]$ can take only two values, 0 and $\frac{1}{2} - \Phi_{\frac{\lambda}{N}}(\frac{1}{2})$, because $\Phi_{\frac{\lambda}{N}}(0) = 0$ and $\Phi_{\frac{\lambda}{N}}(1) = 1$. Thus

$$T_i(\sigma_i) - \Phi_{\frac{\lambda}{N}}[T_i(\sigma_i)] = [1 - |1 - 2T_i(\sigma_i)|] \left[\frac{1}{2} - \Phi_{\frac{\lambda}{N}}\left(\frac{1}{2}\right)\right].$$

This shows that in the case when $k = 1$,

$$\begin{aligned} \log\left\{T[\exp(-\lambda r_1)]\right\} &= -\lambda \bar{r} + \frac{\lambda}{N} \sum_{i=1}^N T_i(\sigma_i) - \Phi_{\frac{\lambda}{N}}[T_i(\sigma_i)] \\ &= -\lambda \bar{r} + \frac{\lambda}{N} \sum_{i=1}^N [1 - |1 - 2T_i(\sigma_i)|] \left[\frac{1}{2} - \Phi_{\frac{\lambda}{N}}\left(\frac{1}{2}\right)\right] \\ &\leq -\lambda \bar{r} + \lambda \left[\frac{1}{2} - \Phi_{\frac{\lambda}{N}}\left(\frac{1}{2}\right)\right] [1 - |1 - 2\bar{r}|]. \end{aligned}$$

Noticing that $\frac{1}{2} - \Phi_{\frac{\lambda}{N}}(\frac{1}{2}) = \frac{N}{\lambda} \log[\cosh(\frac{\lambda}{2N})]$, we obtain

THEOREM 2.9. *For any partially exchangeable function $\lambda : \Omega \times \Theta \rightarrow \mathbb{R}_+$, for any partially exchangeable posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,*

$$\mathbb{P}\left\{\exp\left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \rho\left[\lambda(\bar{r} - r_1) - N \log\left[\cosh\left(\frac{\lambda}{2N}\right)\right] (1 - |1 - 2\bar{r}|)\right] - \mathcal{K}(\rho, \pi)\right]\right\} \leq 1.$$

As a consequence, reasoning as previously, we deduce

THEOREM 2.10. *In the case when $k = 1$, for any partially exchangeable posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$ and any $\lambda \in \mathbb{R}_+$,*

$$\bar{r}(\theta) - \frac{N}{\lambda} \log \left[\cosh \left(\frac{\lambda}{2N} \right) \right] (1 - |1 - 2\bar{r}(\theta)|) + \frac{\log \{ \epsilon \pi [\Delta(\theta)] \}}{\lambda} \leq r_1(\theta);$$

and consequently for any $\theta \in \Theta$,

$$r_2(\theta) \leq 2 \inf_{\lambda \in \mathbb{R}_+} \frac{r_1(\theta) - \frac{\log \{ \epsilon \pi [\Delta(\theta)] \}}{\lambda}}{1 - \frac{2N}{\lambda} \log \left[\cosh \left(\frac{\lambda}{2N} \right) \right]} - r_1(\theta).$$

In the case of binary classification using a VC class of VC dimension not greater than h , we can choose π such that $-\log \{ \pi [\Delta(\theta)] \} \leq h \log \left(\frac{2eN}{h} \right)$ and obtain the following numerical illustration of this theorem : for $N = 1000$, $h = 10$, $\epsilon = 0.01$ and $\inf_{\Theta} r_1 = r_1(\hat{\theta}) = 0.2$, we find an upper bound $r_2(\hat{\theta}) \leq 0.5033$, which improves on Theorem 2.8 but still is not under the significance level $\frac{1}{2}$ (achieved by blind random classification). This indicates that considering shadow samples of arbitrary sizes brings in some noisy situations a significant improvement on bounds obtained with a shadow sample of the same size as the training sample.

2.2.3. When moreover the distribution of the augmented sample is exchangeable. In the case when $k = 1$ and \mathbb{P} is exchangeable meaning that for any bounded measurable function $h : \Omega \rightarrow \mathbb{R}$ and any permutation $s \in \mathfrak{S}(\{1, \dots, 2N\})$ $\mathbb{P}[h(\omega \circ s)] = \mathbb{P}[h(\omega)]$, then we can still improve the bound as follows. Let

$$T'(h) = \frac{1}{N!} \sum_{s \in \mathfrak{S}(\{N+1, \dots, 2N\})} h(\omega \circ s).$$

Then we can write

$$1 - |1 - 2T_i(\sigma_i)| = (\sigma_i - \sigma_{i+N})^2 = \sigma_i + \sigma_{i+N} - 2\sigma_i\sigma_{i+N}.$$

Using this identity, we get for any exchangeable function $\lambda : \Omega \times \Theta \rightarrow \mathbb{R}_+$,

$$T \left\{ \exp \left[\lambda(\bar{r} - r_1) - \log \left[\cosh \left(\frac{\lambda}{2N} \right) \right] \sum_{i=1}^N (\sigma_i + \sigma_{i+N} - 2\sigma_i\sigma_{i+N}) \right] \right\} \leq 1.$$

Let us put

$$A(\lambda) = \frac{2N}{\lambda} \log \left[\cosh \left(\frac{\lambda}{2N} \right) \right], \quad (2.2)$$

$$v(\theta) = \frac{1}{2N} \sum_{i=1}^N (\sigma_i + \sigma_{i+N} - 2\sigma_i \sigma_{i+N}). \quad (2.3)$$

With these notations

$$T \left\{ \exp \left\{ \lambda [\bar{r} - r_1 - A(\lambda)v] \right\} \right\} \leq 1.$$

Let notice now that

$$T' [v(\theta)] = \bar{r}(\theta) - r_1(\theta)r_2(\theta).$$

Let $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$ be any given exchangeable posterior distribution. Using the exchangeability of \mathbb{P} and π and the exchangeability of the exponential function, we get

$$\begin{aligned} \mathbb{P} \left\{ \pi \left[\exp \left\{ \lambda [\bar{r} - r_1 - A(\bar{r} - r_1 r_2)] \right\} \right] \right\} &= \mathbb{P} \left\{ \pi \left[\exp \left\{ \lambda [\bar{r} - r_1 - AT'(v)] \right\} \right] \right\} \\ &\leq \mathbb{P} \left\{ \pi \left[T' \exp \left\{ \lambda [\bar{r} - r_1 - Av] \right\} \right] \right\} = \mathbb{P} \left\{ T' \pi \left[\exp \left\{ \lambda [\bar{r} - r_1 - Av] \right\} \right] \right\} \\ &= \mathbb{P} \left\{ \pi \left[\exp \left\{ \lambda [\bar{r} - r_1 - Av] \right\} \right] \right\} = \mathbb{P} \left\{ T \pi \left[\exp \left\{ \lambda [\bar{r} - r_1 - Av] \right\} \right] \right\} \\ &= \mathbb{P} \left\{ \pi \left[T \exp \left\{ \lambda [\bar{r} - r_1 - Av] \right\} \right] \right\} \leq 1. \end{aligned}$$

We are thus ready to state

THEOREM 2.11. *In the case when $k = 1$, for any exchangeable probability distribution \mathbb{P} , for any exchangeable posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, for any exchangeable function $\lambda : \Omega \times \Theta \rightarrow \mathbb{R}_+$,*

$$\mathbb{P} \left\{ \exp \left[\sup_{\rho \in \mathcal{M}_+^1(\Theta)} \rho \left\{ \lambda [\bar{r} - r_1 - A(\lambda)(\bar{r} - r_1 r_2)] \right\} - \mathcal{K}(\rho, \pi) \right] \right\} \leq 1,$$

where $A(\lambda)$ is defined by equation (2.2) above.

We then deduce as previously

COROLLARY 2.12. *For any exchangeable posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, for any exchangeable probability measure $\mathbb{P} \in \mathcal{M}_+^1(\Omega)$, for any measurable exchangeable function $\lambda : \Omega \times \Theta \rightarrow \mathbb{R}_+$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,*

$$\bar{r}(\theta) \leq r_1(\theta) + A(\lambda) [\bar{r}(\theta) - r_1(\theta)r_2(\theta)] - \frac{\log \{ \epsilon \pi [\Delta(\theta)] \}}{\lambda},$$

where $A(\lambda)$ is defined by equation (2.2) on page 108.

In order to deduce an empirical bound from this theorem, we have to make some choice for $\lambda(\omega, \theta)$. Fortunately, it is easy to show that the bound indeed holds uniformly in λ . This is the case because the inequality can be rewritten as a function of only one non exchangeable quantity, namely $r_1(\theta)$. Indeed, since $r_2 = 2\bar{r} - r_1$, we see that the inequality can be written as

$$\bar{r}(\theta) \leq r_1(\theta) + A(\lambda) [\bar{r}(\theta) - 2\bar{r}(\theta)r_1(\theta) + r_1(\theta)^2] - \frac{\log\{\epsilon\pi[\Delta(\theta)]\}}{\lambda}.$$

It can be solved in $r_1(\theta)$, to get

$$r_1(\theta) \geq f\left(\lambda, \bar{r}(\theta), -\log\{\epsilon\pi[\Delta(\theta)]\}\right),$$

where namely

$$f(\lambda, \bar{r}, d) = [2A(\lambda)]^{-1} \left\{ 2\bar{r}A(\lambda) - 1 + \sqrt{[1 - 2\bar{r}A(\lambda)]^2 + 4A(\lambda)\left\{\bar{r}[1 - A(\lambda)] - \frac{d}{\lambda}\right\}} \right\}.$$

Thus we can find some exchangeable function $\lambda(\omega, \theta)$, such that

$$f\left(\lambda(\omega, \theta), \bar{r}(\theta), -\log\{\epsilon\pi[\Delta(\theta)]\}\right) = \sup_{\beta \in \mathbb{R}_+} f\left(\beta, \bar{r}(\theta), -\log\{\epsilon\pi[\Delta(\theta)]\}\right).$$

Applying Corollary 2.12 to that choice of λ , we see that

THEOREM 2.13. *For any exchangeable probability measure $\mathbb{P} \in \mathcal{M}_+^1(\Omega)$, for any exchangeable posterior probability distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$, for any $\lambda \in \mathbb{R}_+$,*

$$\bar{r}(\theta) \leq r_1(\theta) + A(\lambda) [\bar{r}(\theta) - r_1(\theta)r_2(\theta)] - \frac{\log\{\epsilon\pi[\Delta(\theta)]\}}{\lambda},$$

where $A(\lambda)$ is defined by equation (2.2) on page 108.

Solving the previous inequality in $r_2(\theta)$, we get

COROLLARY 2.14. *Under the same assumptions as in the previous theorem, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,*

$$r_2(\theta) \leq \inf_{\lambda \in \mathbb{R}_+} \frac{r_1(\theta) \left\{ 1 + \frac{2N}{\lambda} \log \left[\cosh\left(\frac{\lambda}{2N}\right) \right] \right\} - \frac{2 \log\{\epsilon\pi[\Delta(\theta)]\}}{\lambda}}{1 - \frac{2N}{\lambda} \log \left[\cosh\left(\frac{\lambda}{2N}\right) \right] [1 - 2r_1(\theta)]}.$$

Applying this to our usual numerical example of a binary classification model with VC dimension not greater than $h = 10$, when $N = 1000$, $\inf_{\Theta} r_1 = r_1(\hat{\theta}) = 10$ and $\epsilon = 0.01$, we obtain that $r_2(\hat{\theta}) \leq 0.4450$.

2.3. VAPNIK'S BOUNDS FOR INDUCTIVE CLASSIFICATION.

2.3.1. Arbitrary shadow sample size. We assume in this section that

$$\mathbb{P} = \left(\bigotimes_{i=1}^N P_i \right)^{\otimes \infty} \in \mathcal{M}_+^1 \left\{ [(\mathcal{X} \times \mathcal{Y})^N]^{\mathbb{N}} \right\},$$

where $P_i \in \mathcal{M}_+^1(\mathcal{X} \times \mathcal{Y})$: we consider an infinite i.i.d. sequence of independent *not* identically distributed samples of size N , the first one only being observed. The shadow samples will only appear in the proofs. The aim of this section is to prove better Vapnik's bounds, generalizing them in the same time to the independent non i.i.d. setting, which to our knowledge had not been done before.

Let us introduce the notation $\mathbb{P}'[h(\omega)] = \mathbb{P}[h(\omega) | (X_i, Y_i)_{i=1}^N]$, where h may be any suitable (e.g. bounded) random variable, let us also put $\Omega = [(\mathcal{X} \times \mathcal{Y})^N]^{\mathbb{N}}$.

DEFINITION 2.2. For any subset $A \subset \mathbb{N}$ of integers, let $\mathfrak{C}(A)$ be the set of circular permutations of the totally ordered set A , extended to a permutation of \mathbb{N} by taking it to be the identity on the complement $\mathbb{N} \setminus A$ of A . We will say that a random function $h : \Omega \rightarrow \mathbb{R}$ is k -partially exchangeable if

$$h(\omega \circ s) = h(\omega), \quad s \in \mathfrak{C}(\{i + jN; j = 0, \dots, k\}), i = 1, \dots, N.$$

In the same way, we will say that a posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$ is k -partially exchangeable if

$$\pi(\omega \circ s) = \pi(\omega) \in \mathcal{M}_+^1(\Theta), \quad s \in \mathfrak{C}(\{i + jN; j = 0, \dots, k\}), i = 1, \dots, N.$$

Note that \mathbb{P} itself is k -partially exchangeable for any k in the sense that for any bounded measurable function $h : \Omega \rightarrow \mathbb{R}$

$$\mathbb{P}[h(\omega \circ s)] = \mathbb{P}[h(\omega)], \quad s \in \mathfrak{C}(\{i + jN; j = 0, \dots, k\}), i = 1, \dots, N.$$

Let $\Delta_k(\theta) = \left\{ \theta' \in \Theta; [f_{\theta'}(X_i)]_{i=1}^{(k+1)N} = [f_{\theta}(X_i)]_{i=1}^{(k+1)N} \right\}$, $\theta \in \Theta, k \in \mathbb{N}^*$,

and let also $\bar{r}_k(\theta) = \frac{1}{(k+1)N} \sum_{i=1}^{(k+1)N} \mathbf{1}[f_{\theta}(X_i) \neq Y_i]$. Theorem 2.2 shows

that for any positive real parameter λ and any k -partially exchangeable posterior distribution $\pi_k : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,

$$\mathbb{P} \left\{ \exp \left[\sup_{\theta \in \Theta} \lambda [\Phi_{\frac{\lambda}{N}}(\bar{r}_k) - r_1] + \log \{ \epsilon \pi_k [\Delta_k(\theta)] \} \right] \right\} \leq \epsilon.$$

Using the general fact that

$$\mathbb{P} [\exp(h)] = \mathbb{P} \left\{ \mathbb{P}' [\exp(h)] \right\} \geq \mathbb{P} \left\{ \exp [\mathbb{P}'(h)] \right\},$$

and the fact that the expectation of a supremum is larger than the supremum of an expectation, we see that with \mathbb{P} probability at most $1 - \epsilon$, for any $\theta \in \Theta$,

$$\mathbb{P}' \left\{ \Phi_{\frac{\lambda}{N}} [\bar{r}_k(\theta)] \right\} \leq r_1(\theta) - \frac{\mathbb{P}' \left\{ \log \{ \epsilon \pi_k [\Delta_k(\theta)] \} \right\}}{\lambda}.$$

Let us put for short

$$\begin{aligned} \bar{d}_k(\theta) &= -\log \{ \epsilon \pi_k [\Delta_k(\theta)] \}, \\ d'_k(\theta) &= -\mathbb{P}' \left\{ \log \{ \epsilon \pi_k [\Delta_k(\theta)] \} \right\}, \\ d_k(\theta) &= -\mathbb{P} \left\{ \log \{ \epsilon \pi_k [\Delta_k(\theta)] \} \right\}. \end{aligned}$$

We can use the convexity of $\Phi_{\frac{\lambda}{N}}$ and the fact that $\mathbb{P}'(\bar{r}_k) = \frac{r_1 + kR}{k+1}$, to see that

$$\mathbb{P}' \left\{ \Phi_{\frac{\lambda}{N}} [\bar{r}_k(\theta)] \right\} \geq \Phi_{\frac{\lambda}{N}} \left[\frac{r_1(\theta) + kR(\theta)}{k+1} \right].$$

We have proved

THEOREM 2.15. *Using the above hypotheses and notations, for any sequence $\pi_k : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, where π_k is a k -partially exchangeable posterior distribution, for any positive real constant λ , any positive integer k , with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,*

$$\Phi_{\frac{\lambda}{N}} \left[\frac{r_1(\theta) + kR(\theta)}{k+1} \right] \leq r_1(\theta) + \frac{d'_k(\theta)}{\lambda}.$$

We can make as we did with Theorem 1.10 on page 20 the result of this theorem uniform in $\lambda \in \{\alpha^j ; j \in \mathbb{N}^*\}$ and $k \in \mathbb{N}^*$ (considering on k the prior $\frac{1}{k(k+1)}$ and on j the prior $\frac{1}{j(j+1)}$), and obtain

THEOREM 2.16. *For any real parameter $\alpha > 1$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,*

$$R(\theta) \leq \inf_{k \in \mathbb{N}^*, j \in \mathbb{N}^*} \frac{1 - \exp \left\{ -\frac{\alpha^j}{N} r_1(\theta) - \frac{1}{N} \left\{ d'_k(\theta) + \log [k(k+1)j(j+1)] \right\} \right\}}{\frac{k}{k+1} \left[1 - \exp \left(-\frac{\alpha^j}{N} \right) \right]} - \frac{r_1(\theta)}{k}.$$

Note that as a special case we can choose π_k such that $\log \{ \pi_k [\Delta_k(\theta)] \}$ is independent of θ and equal to $\log(\mathfrak{N}_k)$, where $\mathfrak{N}_k = |\{ [f_\theta(X_i)]_{i=1}^{(k+1)N}; \theta \in \Theta \}|$ is the size of the trace of the classification model on the extended sample of size $(k+1)N$. With this choice, we obtain a bound involving a new flavour of conditional Vapnik's entropy, namely

$$d'_k(\theta) = \mathbb{P} [\log(\mathfrak{N}_k) | (Z_i)_{i=1}^N] - \log(\epsilon).$$

In the case of binary classification using a VC class of VC dimension not greater than $h = 10$, when $N = 1000$, $\inf_{\Theta} r_1 = r_1(\hat{\theta}) = 0.2$ and $\epsilon = 0.01$, choosing $\alpha = 1.1$, we obtain $R(\hat{\theta}) \leq 0.4271$ (for an optimal value of $\lambda = 1071.8$, and an optimal value of $k = 16$).

2.3.2. *A better minimization with respect to the exponential parameter.* If we are not pleased with the fact of optimizing λ on a discrete subset of the real line, we can use a slightly different approach. From Theorem 2.2, we see that for any positive integer k , for any k -partially exchangeable positive real measurable function $\lambda : \Omega \times \Theta \rightarrow \mathbb{R}_+$ satisfying equation (2.1) on page 104 (with $\Delta(\theta)$ replaced with $\Delta_k(\theta)$), for any $\epsilon \in]0, 1)$ and $\eta \in]0, 1)$,

$$\mathbb{P} \left\{ \mathbb{P}' \left[\exp \left[\sup_{\theta} \lambda \left[\Phi_{\frac{\lambda}{N}}(\bar{r}_k) - r_1 \right] + \log \{ \epsilon \eta \pi_k [\Delta_k(\theta)] \} \right] \right] \right\} \leq \epsilon \eta,$$

therefore with \mathbb{P} probability at least $1 - \epsilon$,

$$\mathbb{P}' \left\{ \exp \left[\sup_{\theta} \lambda \left[\Phi_{\frac{\lambda}{N}}(\bar{r}_k) - r_1 \right] + \log \{ \epsilon \eta \pi_k [\Delta_k(\theta)] \} \right] \right\} \leq \eta,$$

and consequently, with \mathbb{P} probability at least $1 - \epsilon$, with \mathbb{P}' probability at least $1 - \eta$, for any $\theta \in \Theta$,

$$\Phi_{\frac{\lambda}{N}}(\bar{r}_k) + \frac{\log \{ \epsilon \eta \pi_k [\Delta_k(\theta)] \}}{\lambda} \leq r_1.$$

Now we are entitled to choose

$$\lambda(\omega, \theta) \in \arg \max_{\lambda' \in \mathbb{R}_+} \Phi_{\frac{\lambda'}{N}}(\bar{r}_k) + \frac{\log\{\epsilon \eta \pi_k[\Delta_k(\theta)]\}}{\lambda'}.$$

This shows that with \mathbb{P} probability at least $1 - \epsilon$, with \mathbb{P}' probability at least $1 - \eta$, for any $\theta \in \Theta$,

$$\sup_{\lambda \in \mathbb{R}_+} \Phi_{\frac{\lambda}{N}}(\bar{r}_k) - \frac{\bar{d}_k(\theta) - \log(\eta)}{\lambda} \leq r_1,$$

which can also be written

$$\Phi_{\frac{\lambda}{N}}(\bar{r}_k) - r_1 - \frac{\bar{d}_k(\theta)}{\lambda} \leq -\frac{\log(\eta)}{\lambda}, \quad \lambda \in \mathbb{R}_+.$$

Thus with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$, any $\lambda \in \mathbb{R}_+$,

$$\mathbb{P}' \left[\Phi_{\frac{\lambda}{N}}(\bar{r}_k) - r_1 - \frac{\bar{d}_k(\theta)}{\lambda} \right] \leq -\frac{\log(\eta)}{\lambda} + \left[1 - r_1 + \frac{\log(\eta)}{\lambda} \right] \eta.$$

On the other hand, $\Phi_{\frac{\lambda}{N}}$ being a convex function,

$$\begin{aligned} \mathbb{P}' \left[\Phi_{\frac{\lambda}{N}}(\bar{r}_k) - r_1 - \frac{\bar{d}_k(\theta)}{\lambda} \right] &\geq \Phi_{\frac{\lambda}{N}}[\mathbb{P}'(\bar{r}_k)] - r_1 - \frac{d'_k}{\lambda} \\ &= \Phi_{\frac{\lambda}{N}}\left(\frac{kR + r_1}{k + 1}\right) - r_1 - \frac{d'_k}{\lambda}. \end{aligned}$$

Thus with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,

$$\frac{kR + r_1}{k + 1} \leq \inf_{\lambda \in \mathbb{R}_+} \Phi_{\frac{\lambda}{N}}^{-1} \left[r_1(1 - \eta) + \eta + \frac{d'_k - \log(\eta)(1 - \eta)}{\lambda} \right].$$

We can generalize this approach by considering a finite decreasing sequence $\eta_0 = 1 > \eta_1 > \eta_2 > \dots > \eta_J > \eta_{J+1} = 0$, and the corresponding sequence of levels

$$\begin{aligned} L_j &= -\frac{\log(\eta_j)}{\lambda}, \quad 0 \leq j \leq J, \\ L_{J+1} &= 1 - r_1 - \frac{\log(J) - \log(\epsilon)}{\lambda}. \end{aligned}$$

Taking a union bound in j , we see that with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$, for any $\lambda \in \mathbb{R}_+$,

$$\mathbb{P}' \left[\Phi_{\frac{\lambda}{N}}(\bar{r}_k) - r_1 - \frac{\bar{d}_k + \log(J)}{\lambda} \geq L_j \right] \leq \eta_j, \quad j = 0, \dots, J + 1,$$

and consequently

$$\begin{aligned} & \mathbb{P}' \left[\Phi_{\frac{\lambda}{N}}(\bar{r}_k) - r_1 - \frac{\bar{d}_k + \log(J)}{\lambda} \right] \\ & \leq \int_0^{L_{J+1}} \mathbb{P}' \left[\Phi_{\frac{\lambda}{N}}(\bar{r}_k) - r_1 - \frac{\bar{d}_k + \log(J)}{\lambda} \geq \alpha \right] d\alpha \leq \sum_{j=1}^{J+1} \eta_{j-1} (L_j - L_{j-1}) \\ & = \eta_J \left[1 - r_1 - \frac{\log(J) - \log(\epsilon) - \log(\eta_J)}{\lambda} \right] - \frac{\log(\eta_1)}{\lambda} + \sum_{j=1}^{J-1} \frac{\eta_j}{\lambda} \log \left(\frac{\eta_j}{\eta_{j+1}} \right). \end{aligned}$$

Let us put

$$\begin{aligned} d_k''[\theta, (\eta_j)_{j=1}^J] &= d_k'(\theta) + \log(J) - \log(\eta_1) \\ & \quad + \sum_{j=1}^{J-1} \eta_j \log \left(\frac{\eta_j}{\eta_{j+1}} \right) + \log \left(\frac{\epsilon \eta_J}{J} \right) \eta_J. \end{aligned}$$

We have proved that for any decreasing sequence $(\eta_j)_{j=1}^J$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,

$$\frac{kR + r_1}{k + 1} \leq \inf_{\lambda \in \mathbb{R}_+} \Phi_{\frac{\lambda}{N}}^{-1} \left[r_1(1 - \eta_J) + \eta_J + \frac{d_k''[\theta, (\eta_j)_{j=1}^J]}{\lambda} \right].$$

REMARK 2.1. We can for instance choose $J = 2$, $\eta_2 = \frac{1}{10N}$, $\eta_1 = \frac{1}{\log(10N)}$, resulting in

$$d_k'' = d_k' + \log(2) + \log \log(10N) + 1 - \frac{\log \log(10N)}{\log(10N)} - \frac{\log \left(\frac{20N}{\epsilon} \right)}{10N}.$$

In the case when $N = 1000$ and for any $\epsilon \in (0, 1)$, we get $d_k'' \leq d_k' + 3.7$, in the case when $N = 10^6$, we get $d_k'' \leq d_k' + 4.4$, and in the case $N = 10^9$, we get $d_k'' \leq d_k' + 4.7$.

Therefore, for any practical purpose we could take $d_k'' = d_k' + 4.7$ and $\eta_J = \frac{1}{10N}$ in the above inequality.

Taking moreover a weighted union bound in k , we get

THEOREM 2.17. *For any $\epsilon \in (0, 1)$, any sequence $1 > \eta_1 > \dots > \eta_J > 0$, any sequence $\pi_k : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, where π_k is a k -partially exchangeable posterior distribution, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,*

$$R(\theta) \leq \inf_{k \in \mathbb{N}^*} \frac{k+1}{k} \inf_{\lambda \in \mathbb{R}_+} \Phi_{\frac{\lambda}{N}}^{-1} \left[r_1(\theta) + \eta_J [1 - r_1(\theta)] \right. \\ \left. + \frac{d_k''[\theta, (\eta_j)_{j=1}^J] + \log[k(k+1)]}{\lambda} \right] - \frac{r_1(\theta)}{k}.$$

COROLLARY 2.18. *For any $\epsilon \in]0, 1)$, for any $N \leq 10^9$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,*

$$R(\theta) \leq \inf_{k \in \mathbb{N}^*} \inf_{\lambda \in \mathbb{R}_+} \frac{k+1}{k} [1 - \exp(-\frac{\lambda}{N})]^{-1} \left\{ 1 - \exp \left[-\frac{\lambda}{N} [r_1(\theta) + \frac{1}{10N}] \right. \right. \\ \left. \left. - \frac{\mathbb{P}'[\log(\mathfrak{N}_k) | (Z_i)_{i=1}^N] - \log(\epsilon) + \log[k(k+1)] + 4.7]}{N} \right] \right\} - \frac{r_1(\theta)}{k}.$$

Let us end this section with a numerical example: in the case of binary classification with a VC class of dimension not greater than 10, when $N = 1000$, $\inf_{\Theta} r_1 = r_1(\hat{\theta}) = 0.2$ and $\epsilon = 0.01$, we get a bound $R(\hat{\theta}) \leq 0.4211$ (for optimal values of $k = 15$ and of $\lambda = 1010$).

2.3.3. *Equal shadow and training sample sizes.* In the case when $k = 1$, we can use Theorem 2.10, and replace $\Phi_{\frac{\lambda}{N}}^{-1}(q)$ with $\{1 - \frac{2N}{\lambda} \log[\cosh(\frac{\lambda}{2N})]\}^{-1} q$, resulting in

THEOREM 2.19. *For any $\epsilon \in]0, 1)$, any $N \leq 10^9$, any 1-partially exchangeable posterior distribution $\pi_1 : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,*

$$R(\theta) \leq \inf_{\lambda \in \mathbb{R}_+} \frac{\left\{ 1 + \frac{2N}{\lambda} \log[\cosh(\frac{\lambda}{2N})] \right\} r_1(\theta) + \frac{1}{5N} + 2 \frac{d_1'(\theta) + 4.7}{\lambda}}{1 - \frac{2N}{\lambda} \log[\cosh(\frac{\lambda}{2N})]}.$$

2.3.4. *Improvement on the equal sample size bound in the i.i.d. case.* Eventually, in the case when \mathbb{P} is i.i.d., meaning that all the P_i are equal, we can improve the previous bound. For any partially exchangeable function $\lambda : \Omega \times \Theta \rightarrow \mathbb{R}_+$, we saw in the discussion preceding Theorem 2.11 on page 108 that

$$T \left[\exp[\lambda(\bar{r}_k - r_1) - A(\lambda)v] \right] \leq 1,$$

with the notations introduced therein. Thus for any partially exchangeable positive real measurable function $\lambda : \Omega \times \Theta \rightarrow \mathbb{R}_+$ satisfying equation (2.1)

on page 104, any 1-partially exchangeable posterior distribution $\pi_1 : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$,

$$\mathbb{P} \left\{ \exp \left[\sup_{\theta \in \Theta} \lambda [\bar{r}_k(\theta) - r_1(\theta) - A(\lambda)v(\theta)] + \log[\epsilon \pi_1[\Delta(\theta)]] \right] \right\} \leq 1.$$

Therefore with \mathbb{P} probability at least $1 - \epsilon$, with \mathbb{P}' probability $1 - \eta$,

$$\bar{r}_k(\theta) \leq r_1(\theta) + A(\lambda)v(\theta) + \frac{1}{\lambda} [\bar{d}_1(\theta) - \log(\eta)]$$

We can then choose $\lambda(\omega, \theta) \in \arg \min_{\lambda' \in \mathbb{R}_+} A(\lambda')v(\theta) + \frac{\bar{d}_1(\theta) - \log(\eta)}{\lambda'}$, which satisfies the required conditions, to show that with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$, with \mathbb{P}' probability at least $1 - \eta$, for any $\lambda \in \mathbb{R}_+$,

$$\bar{r}_k(\theta) \leq r_1(\theta) + A(\lambda)v(\theta) + \frac{\bar{d}_1(\theta) - \log(\eta)}{\lambda}.$$

We can then take a union bound on a decreasing sequence of J values $\eta_1 \geq \dots \geq \eta_J$ of η . Weakening a little the order of quantifiers, we then obtain the following statement: with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$, for any $\lambda \in \mathbb{R}_+$, for any $j = 1, \dots, J$

$$\mathbb{P}' \left[\bar{r}_k(\theta) - r_1(\theta) - A(\lambda)v(\theta) - \frac{\bar{d}_1(\theta) + \log(J)}{\lambda} \geq -\frac{\log(\eta_j)}{\lambda} \right] \leq \eta_j.$$

Consequently for any $\lambda \in \mathbb{R}_+$,

$$\begin{aligned} & \mathbb{P}' \left[\bar{r}_k(\theta) - r_1(\theta) - A(\lambda)v(\theta) - \frac{\bar{d}_1(\theta) + \log(J)}{\lambda} \right] \\ & \leq -\frac{\log(\eta_1)}{\lambda} + \eta_J \left[1 - r_1(\theta) - \frac{\log(J) - \log(\epsilon) - \log(\eta_J)}{\lambda} \right] \\ & \quad + \sum_{j=1}^{J-1} \frac{\eta_j}{\lambda} \log \left(\frac{\eta_j}{\eta_{j+1}} \right). \end{aligned}$$

Moreover $\mathbb{P}'[v(\theta)] = \frac{r_1 + R}{2} - r_1 R$, (this is where we need equidistribution) thus proving that

$$\frac{R - r_1}{2} \leq \frac{A(\lambda)}{2} [R + r_1 - 2r_1 R] + \frac{d_1''[\theta, (\eta_j)_{j=1}^J]}{\lambda} + \eta_J [1 - r_1(\theta)].$$

Keeping track of quantifiers, we obtain

THEOREM 2.20. *For any decreasing sequence $(\eta_j)_{j=1}^J$, any $\epsilon \in (0, 1)$, any 1-partially exchangeable posterior distribution $\pi : \Omega \rightarrow \mathcal{M}_+^1(\Theta)$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,*

$$R(\theta) \leq \inf_{\lambda \in \mathbb{R}_+} \frac{\left\{ 1 + \frac{2N}{\lambda} \log \left[\cosh \left(\frac{\lambda}{2N} \right) \right] \right\} r_1(\theta) + \frac{2d_1''[\theta, (\eta_j)_{j=1}^J]}{\lambda} + 2\eta_J [1 - r_1(\theta)]}{1 - \frac{2N}{\lambda} \log \left[\cosh \left(\frac{\lambda}{2N} \right) \right] [1 - 2r_1(\theta)]}.$$

2.4. GAUSSIAN APPROXIMATION IN VAPNIK'S BOUNDS. To obtain formulas which could be easily compared with original Vapnik's bounds, we may replace $p - \Phi_a(p)$ with a Gaussian upper bound:

LEMMA 2.21. *For any $p \in (0, \frac{1}{2})$, any $a \in \mathbb{R}_+$,*

$$p - \Phi_a(p) \leq \frac{a}{2} p(1 - p).$$

For any $p \in (\frac{1}{2}, 1)$,

$$p - \Phi_a(p) \leq \frac{a}{8}.$$

PROOF. Let us notice that for any $p \in (0, 1)$,

$$\begin{aligned} \frac{\partial}{\partial a} [-a\Phi_a(p)] &= -\frac{p \exp(-a)}{1 - p + p \exp(-a)}, \\ \frac{\partial^2}{\partial a^2} [-a\Phi_a(p)] &= \frac{p \exp(-a)}{1 - p + p \exp(-a)} \left(1 - \frac{p \exp(-a)}{1 - p + p \exp(-a)} \right) \\ &\leq \begin{cases} p(1 - p) & p \in (0, \frac{1}{2}), \\ \frac{1}{4} & p \in (\frac{1}{2}, 1). \end{cases} \end{aligned}$$

Thus taking a Taylor expansion of order one with integral remainder :

$$-a\Phi(a) \leq \begin{cases} -ap + \int_0^a p(1 - p)(a - b)db \\ \qquad \qquad \qquad = -ap + \frac{a^2}{2} p(1 - p), & p \in (0, \frac{1}{2}), \\ -ap + \int_0^a \frac{1}{4}(a - b)db = -ap + \frac{a^2}{8}, & p \in (\frac{1}{2}, 1). \end{cases}$$

This ends the proof of our lemma. \square

LEMMA 2.22. *Let us consider the bound*

$$B(q, d) = \left(1 + \frac{2d}{N}\right)^{-1} \left[q + \frac{d}{N} + \sqrt{\frac{2dq(1-q)}{N} + \frac{d^2}{N^2}} \right], \quad q \in \mathbb{R}_+, d \in \mathbb{R}_+.$$

Let us also put

$$\bar{B}(q, d) = \begin{cases} B(q, d) & B(q, d) \leq \frac{1}{2}, \\ q + \sqrt{\frac{d}{2N}} & \text{otherwise.} \end{cases}$$

For any positive real parameters q and d

$$\inf_{\lambda \in \mathbb{R}_+} \Phi_{\frac{\lambda}{N}}^{-1} \left(q + \frac{d}{\lambda} \right) \leq \bar{B}(q, d).$$

PROOF. Let $p = \inf_{\lambda} \Phi_{\frac{\lambda}{N}}^{-1} \left(q + \frac{d}{\lambda} \right)$. For any $\lambda \in \mathbb{R}_+$,

$$p - \frac{\lambda}{2N} (p \wedge \frac{1}{2}) [1 - (p \wedge \frac{1}{2})] \leq \Phi_{\frac{\lambda}{N}}(p) \leq q + \frac{d}{\lambda}.$$

Thus

$$\begin{aligned} p &\leq q + \inf_{\lambda \in \mathbb{R}_+} \frac{\lambda}{2N} (p \wedge \frac{1}{2}) [1 - (p \wedge \frac{1}{2})] + \frac{d}{\lambda} \\ &= q + \sqrt{\frac{2d(p \wedge \frac{1}{2}) [1 - (p \wedge \frac{1}{2})]}{N}} \leq q + \sqrt{\frac{d}{2N}}. \end{aligned}$$

Then let us remark that $B(q, d) = \sup \left\{ p' \in \mathbb{R}_+ ; p' \leq q + \sqrt{\frac{2dp'(1-p')}{N}} \right\}$.

If moreover $\frac{1}{2} \geq B(q, d)$, then according to this remark $\frac{1}{2} \geq q + \sqrt{\frac{d}{2N}} \geq p$.

Therefore $p \leq \frac{1}{2}$, and consequently $p \leq q + \sqrt{\frac{2dp(1-p)}{N}}$, implying that $p \leq B(q, d)$. \square

2.4.1. Arbitrary shadow sample size. This lemma combined with Corollary 2.18 on page 115 implies

COROLLARY 2.23. *For any $\epsilon \in (0, 1)$, any integer $N \leq 10^9$, with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,*

$$R(\theta) \leq \inf_{k \in \mathbb{N}^*} \frac{k+1}{k} \left\{ \bar{B} \left[r_1(\theta) + \frac{1}{10N}, d'_k(\theta) + \log[k(k+1)] + 4.7 \right] \right\} - \frac{r_1(\theta)}{k}.$$

2.4.2. *Equal sample sizes in the i.i.d. case.* To make a link with Vapnik's result, it is useful to work out the Gaussian approximation to Theorem 2.20 on page 117. Indeed, using the upper bound $A(\lambda) \leq \frac{\lambda}{4N}$, where $A(\lambda)$ is defined by equation (2.2) on page 108, we get with \mathbb{P} probability at least $1 - \epsilon$

$$R - r_1 - 2\eta_J \leq \inf_{\lambda \in \mathbb{R}_+} \frac{\lambda}{4N} [R + r_1 - 2r_1 R] + \frac{2d_1''}{\lambda} = \sqrt{\frac{2d_1''(R + r_1 - 2r_1 R)}{N}},$$

which can be solved in R to obtain

COROLLARY 2.24. *With \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$,*

$$\begin{aligned} R(\theta) &\leq r_1(\theta) + \frac{d_1''(\theta)}{N} [1 - 2r_1(\theta)] + 2\eta_J \\ &+ \sqrt{\frac{4d_1''(\theta)[1 - r_1(\theta)]r_1(\theta)}{N} + \frac{d_1''(\theta)^2}{N^2} [1 - 2r_1(\theta)]^2 + \frac{4d_1''(\theta)}{N} [1 - 2r_1(\theta)]\eta_J}. \end{aligned}$$

This is to be compared with Vapnik's result, as proved in [37, page 138]:

THEOREM 2.25 (VAPNIK). *For any i.i.d. probability distribution \mathbb{P} , with \mathbb{P} probability at least $1 - \epsilon$, for any $\theta \in \Theta$, putting*

$$d_V = \log[\mathbb{P}(\mathfrak{N}_1)] + \log(4/\epsilon),$$

$$R(\theta) \leq r_1(\theta) + \frac{2d_V}{N} + \sqrt{\frac{4d_V r_1(\theta)}{N} + \frac{4d_V^2}{N^2}}.$$

Recalling that we can choose $(\eta_j)_{j=1}^2$ such that $\eta_J = \frac{1}{10N}$ (which is negligible by all means) and such that for any $N \leq 10^9$,

$$d_1''(\theta) \leq \mathbb{P}[\log(\mathfrak{N}_1) | (Z_i)_{i=1}^N] - \log(\epsilon) + 4.7,$$

we see that our complexity term is somehow more satisfactory than Vapnik's, since it is integrated outside the logarithm, with a little larger additional constant (remember that $\log(4) \simeq 1.4$, which is better than our 4.7, which could presumably be improved by working out a better sequence η_j , but not down to $\log(4)$). Our variance term is better, since we get $r_1(1 - r_1)$ as we should, instead of only r_1 . We also have $\frac{d_1''}{N}$ instead of $2\frac{d_V}{N}$, coming from the fact that we do not use any symmetrization trick.

Let us illustrate these bound on a numerical example (corresponding to a situation where the sample is noisy or the classification model is weak).

Let us assume that $N = 1000$, $\inf_{\Theta} r_1 = r_1(\hat{\theta}) = 0.2$, that we are performing binary classification with a model with VC dimension not greater than $h = 10$, and that we work at level of confidence $\epsilon = 0.01$. Vapnik's theorem provides an upper bound for $R(\hat{\theta})$ not smaller than 0.610, whereas Corollary 2.24 gives $R(\hat{\theta}) \leq 0.461$ (using the bound $d_1'' \leq d_1' + 3.7$ when $N = 1000$). Now if we go for Theorem 2.20 and do not make a Gaussian approximation, we get $R(\hat{\theta}) \leq 0.453$. It is interesting to remark that this bound is achieved for $\lambda = 1195 > N = 1000$. This explains why the Gaussian approximation in Vapnik's bound can be improved: for such a large value of λ , $\lambda r_1(\theta)$ does not behave like a Gaussian random variable.

Let us remind in conclusion that the best bound is provided by Theorem 2.17, giving $R(\hat{\theta}) \leq 0.4211$, (that is approximately 2/3 of Vapnik's bound), for optimal values of $k = 15$, and of $\lambda = 1010$. This bound can be seen to take advantage of the fact that Bernoulli random variables are not Gaussian (its Gaussian approximation, Corollary 2.23, gives a bound $R(\theta) \simeq 0.4325$, still with an optimal $k = 15$), and of the fact that the optimal size of the shadow sample is significantly larger than the size of the observed sample. Moreover, Theorem 2.17 does not assume that the sample is i.i.d., but only that it is independent, thus generalizing Vapnik's bounds to inhomogeneous data (this will presumably be the case when data are collected from different places where the experimental conditions may not be expected to be the same, although they may reasonably be assumed to be independent). We would like also to emphasize that our little numerical example shows that Vapnik's bounds can be expected to be appropriate when dealing with moderate sample sizes. More sophisticated bounds obviously have a better asymptotic behaviour as shown in the first section. Nevertheless the numerical illustration of Theorem 1.18 given on page 30 suggests that Vapnik's bounds are not doing so bad for small to medium ratios between the sample size and the dimension of the classification model (with local bounds, we could only get down to 0.332, although using a quite stronger dimension assumption).

We chose on purpose an example where it is non trivial to decide whether the chosen classifier does better than the 0.5 error rate of blind random classification. We think that this situation of weak learning is of practical interest, since "significant" weak learning rules may afterwards be aggregated or combined in various ways to achieve better classification rates.

3. SUPPORT VECTOR MACHINES

3.1. HOW TO BUILD THEM.

3.1.1. The canonical hyperplane. Support Vector Machines, of widely spread use and renown, were introduced by V. Vapnik [37]. Before introducing them, we will study as a prerequisite the separation of points by hyperplanes in a finite dimensional Euclidean space. Support Vector Machines perform the same kind of linear separation after an implicit change of pattern space. The preceding PAC-Bayesian results provide a fit framework to analyze their generalization properties.

We will deal in this section with the classification of points in \mathbb{R}^d in two classes. Let $Z = (x_i, y_i)_{i=1}^N \in (\mathbb{R}^d \times \{-1, +1\})^N$ be some set of labelled examples (called the training set hereafter). Let us split the set of indices $I = \{1, \dots, N\}$ according to the labels into two subsets

$$\begin{aligned} I_+ &= \{i \in I : y_i = +1\}, \\ I_- &= \{i \in I : y_i = -1\}. \end{aligned}$$

Let us then consider the set of admissible separating directions

$$A_Z = \{w \in \mathbb{R}^d : \sup_{b \in \mathbb{R}} \inf_{i \in I} (\langle w, x_i \rangle - b)y_i \geq 1\},$$

which can also be written as

$$A_Z = \{w \in \mathbb{R}^d : \max_{i \in I_-} \langle w, x_i \rangle + 2 \leq \min_{i \in I_+} \langle w, x_i \rangle\}.$$

As it is easily seen, the optimal value of b for a fixed value of w , in other words the value of b which maximizes $\inf_{i \in I} (\langle w, x_i \rangle - b)y_i$, is equal to

$$b_w = \frac{1}{2} \left[\max_{i \in I_-} \langle w, x_i \rangle + \min_{i \in I_+} \langle w, x_i \rangle \right].$$

LEMMA 3.1. *When $A_Z \neq \emptyset$, $\inf\{\|w\|^2 : w \in A_Z\}$ is reached for only one value w_Z of w .*

PROOF. Let $w_0 \in A_Z$. The set $A_Z \cap \{w \in \mathbb{R}^d : \|w\| \leq \|w_0\|\}$ is a compact convex set and $w \mapsto \|w\|^2$ is strictly convex and therefore has a unique minimum on this set, which is also obviously its minimum on A_Z . \square

DEFINITION 3.1. When $A_Z \neq \emptyset$, the training set Z is said to be linearly separable. The hyperplane

$$H = \{x \in \mathbb{R}^d : \langle w_Z, x \rangle - b_Z = 0\},$$

where

$$\begin{aligned} w_Z &= \arg \min \{\|w\| : w \in A_Z\}, \\ b_Z &= b_{w_Z}, \end{aligned}$$

is called the canonical separating hyperplane of the training set Z . The quantity $\|w_Z\|^{-1}$ is called the margin of the canonical hyperplane.

Note that as $\min_{i \in I_+} \langle w_Z, x_i \rangle - \max_{i \in I_-} \langle w_Z, x_i \rangle = 2$, the margin is also equal to half the distance between the projections on the direction w_Z of the positive and negative patterns.

3.1.2. *Computation of the canonical hyperplane.* Let us consider the convex hulls X_+ and X_- of the positive and negative patterns:

$$\begin{aligned} \mathcal{X}_+ &= \left\{ \sum_{i \in I_+} \lambda_i x_i : (\lambda_i)_{i \in I_+} \in \mathbb{R}_+^{I_+}, \sum_{i \in I_+} \lambda_i = 1 \right\}, \\ \mathcal{X}_- &= \left\{ \sum_{i \in I_-} \lambda_i x_i : (\lambda_i)_{i \in I_-} \in \mathbb{R}_+^{I_-}, \sum_{i \in I_-} \lambda_i = 1 \right\}. \end{aligned}$$

Let us introduce the closed convex set

$$\mathcal{V} = \mathcal{X}_+ - \mathcal{X}_- = \{x_+ - x_- : x_+ \in \mathcal{X}_+, x_- \in \mathcal{X}_-\}.$$

As $v \mapsto \|v\|^2$ is strictly convex, with compact lower level sets, there is a unique vector v^* such that

$$\|v^*\|^2 = \inf_{v \in \mathcal{V}} \{\|v\|^2 : v \in \mathcal{V}\}.$$

LEMMA 3.2. *The set A_Z is non empty (i.e. the training set Z is linearly separable) if and only if $v^* \neq 0$. In this case*

$$w_Z = \frac{2}{\|v^*\|^2} v^*,$$

and the margin of the canonical hyperplane is equal to $\frac{1}{2}\|v^\|$.*

PROOF. Let us assume first that $v^* = 0$, or equivalently that $\mathcal{X}_+ \cap \mathcal{X}_- \neq \emptyset$. As for any vector $w \in \mathbb{R}^d$,

$$\begin{aligned}\min_{i \in I_+} \langle w, x_i \rangle &= \min_{x \in \mathcal{X}_+} \langle w, x \rangle, \\ \max_{i \in I_-} \langle w, x_i \rangle &= \max_{x \in \mathcal{X}_-} \langle w, x \rangle,\end{aligned}$$

we see that necessarily $\min_{i \in I_+} \langle w, x_i \rangle - \max_{i \in I_-} \langle w, x_i \rangle \leq 0$, which shows that w cannot be in A_Z and therefore that A_Z is empty.

Let us assume now that $v^* \neq 0$, or equivalently that $\mathcal{X}_+ \cap \mathcal{X}_- = \emptyset$. Let us put $w^* = \frac{2}{\|v^*\|^2} v^*$. Let us remark first that

$$\begin{aligned}\min_{i \in I_+} \langle w^*, x_i \rangle - \max_{i \in I_-} \langle w^*, x_i \rangle &= \inf_{x \in \mathcal{X}_+} \langle w^*, x \rangle - \sup_{x \in \mathcal{X}_-} \langle w^*, x \rangle \\ &= \inf_{x_+ \in \mathcal{X}_+, x_- \in \mathcal{X}_-} \langle w^*, x_+ - x_- \rangle \\ &= \frac{2}{\|v^*\|^2} \inf_{v \in \mathcal{V}} \langle v^*, v \rangle.\end{aligned}$$

Let us now prove that $\inf_{v \in \mathcal{V}} \langle v^*, v \rangle = \|v^*\|^2$. Some arbitrary $v \in \mathcal{V}$ being fixed, consider the function

$$\beta \mapsto \|\beta v + (1 - \beta)v^*\|^2 : [0, 1] \rightarrow \mathbb{R}.$$

By definition of v^* , it reaches its minimum value for $\beta = 0$, and therefore has a non negative derivative at this point. Computing this derivative, we find that $\langle v - v^*, v^* \rangle \geq 0$, as claimed. We have proved that

$$\min_{i \in I_+} \langle w^*, x_i \rangle - \max_{i \in I_-} \langle w^*, x_i \rangle = 2,$$

and therefore that $w^* \in A_Z$. On the other hand, any $w \in A_Z$ is such that

$$2 \leq \min_{i \in I_+} \langle w, x_i \rangle - \max_{i \in I_-} \langle w, x_i \rangle = \inf_{v \in \mathcal{V}} \langle w, v \rangle \leq \|w\| \inf_{v \in \mathcal{V}} \|v\| = \|w\| \|v^*\|.$$

This proves that $\|w^*\| = \inf\{\|w\| : w \in A_Z\}$, and therefore that $w^* = w_Z$ as claimed. \square One way to compute w_Z would be therefore to compute v^* by minimizing

$$\left\{ \left\| \sum_{i \in I} \lambda_i y_i x_i \right\|^2 : (\lambda_i)_{i \in I} \in \mathbb{R}_+^I, \sum_{i \in I} \lambda_i = 2, \sum_{i \in I} y_i \lambda_i = 0 \right\}.$$

Although this is a tractable quadratic programming problem, a direct computation of w_Z through the following proposition is usually preferred.

PROPOSITION 3.3. *The canonical direction w_Z can be expressed as*

$$w_Z = \sum_{i=1}^N \alpha_i^* y_i x_i,$$

where $(\alpha_i^*)_{i=1}^N$ is obtained by minimizing

$$\inf\{F(\alpha) : \alpha \in \mathcal{A}\},$$

where

$$\mathcal{A} = \left\{ (\alpha_i)_{i \in I} \in \mathbb{R}_+^I, \sum_{i \in I} \alpha_i y_i = 0 \right\},$$

and

$$F(\alpha) = \left\| \sum_{i \in I} \alpha_i y_i x_i \right\|^2 - 2 \sum_{i \in I} \alpha_i.$$

PROOF. Let $w(\alpha) = \sum_{i \in I} \alpha_i y_i x_i$ and let $S(\alpha) = \frac{1}{2} \sum_{i \in I} \alpha_i$. We can express the function $F(\alpha)$ as $F(\alpha) = \|w(\alpha)\|^2 - 4S(\alpha)$. Moreover it is important to notice that for any $s \in \mathbb{R}_+$ $\{w(\alpha) : \alpha \in \mathcal{A}, S(\alpha) = s\} = s\mathcal{V}$. This shows that for any $s \in \mathbb{R}_+$, $\inf\{F(\alpha) : \alpha \in \mathcal{A}, S(\alpha) = s\}$ is reached and that for any $\alpha_s \in \{\alpha \in \mathcal{A} : S(\alpha) = s\}$ reaching this infimum, $w(\alpha_s) = sv^*$. As $s \mapsto s^2 \|v^*\|^2 - 4s : \mathbb{R}_+ \rightarrow \mathbb{R}$ reaches its infimum for only one value s^* of s , namely at $s^* = \frac{2}{\|v^*\|^2}$, this shows that $F(\alpha)$ reaches its infimum on \mathcal{A} , and that for any $\alpha^* \in \mathcal{A}$ such that $F(\alpha^*) = \inf\{F(\alpha) : \alpha \in \mathcal{A}\}$, $w(\alpha^*) = \frac{2}{\|v^*\|^2} v^* = w_Z$. \square

3.1.3. Support vectors.

DEFINITION 3.2. The set of support vectors \mathfrak{S} is defined by

$$\mathfrak{S} = \{x_i : \langle w_Z, x_i \rangle - b_Z = y_i\}.$$

PROPOSITION 3.4. *Any α^* minimizing $F(\alpha)$ on \mathcal{A} is such that*

$$\{x_i : \alpha_i^* > 0\} \subset \mathfrak{S}.$$

This implies that the representation $w_Z = w(\alpha^)$ involves in general only a limited number of non zero coefficients and that $w_Z = w_{Z'}$, where $Z' = \{(x_i, y_i) : x_i \in \mathfrak{S}\}$.*

PROOF. Let us consider any given $i \in I_+$ and $j \in I_-$, such that $\alpha_i^* > 0$ and $\alpha_j^* > 0$ (there exists at least one such index in each set I_- and I_+ , since the sum of the components of α^* on each of these sets are equal and since $\sum_{k \in I} \alpha_k^* > 0$). For any $t \in \mathbb{R}$, consider

$$\alpha_k(t) = \alpha_k^* + t\mathbb{1}(k \in \{i, j\}), \quad k \in I.$$

The vector $\alpha(t)$ is in \mathcal{A} for any value of t in some neighborhood of 0, therefore $\frac{\partial}{\partial t}|_{t=0} F[\alpha(t)] = 0$. Computing this derivative, we find that

$$y_i \langle w(\alpha^*), x_i \rangle + y_j \langle w(\alpha^*), x_j \rangle = 2.$$

As $y_i = -y_j$, this can also be written as

$$y_i [\langle w(\alpha^*), x_i \rangle - b_Z] + y_j [\langle w(\alpha^*), x_j \rangle - b_Z] = 2.$$

As $w(\alpha^*) \in A_Z$,

$$y_k [\langle w(\alpha^*), x_k \rangle - b_Z] \geq 1, \quad k \in I,$$

which implies necessarily as claimed that

$$y_i [\langle w(\alpha^*), x_i \rangle - b_Z] = y_j [\langle w(\alpha^*), x_j \rangle - b_Z] = 1.$$

□

3.1.4. The non separable case. In the case when the training set $Z = (x_i, y_i)_{i=1}^N$ is not linearly separable, we can define a noisy canonical hyperplane as follows. We can choose $w \in \mathbb{R}^d$ and $b \in \mathbb{R}$ to minimize

$$C(w, b) = \sum_{i=1}^N [1 - (\langle w, x_i \rangle - b)y_i]_+ + \frac{1}{2} \|w\|^2, \quad (3.1)$$

where for any real number r , $r_+ = \max\{r, 0\}$ is the positive part of r .

THEOREM 3.5. *Let us introduce the dual criterion*

$$F(\alpha) = \sum_{i=1}^N \alpha_i - \frac{1}{2} \left\| \sum_{i=1}^N y_i \alpha_i x_i \right\|^2$$

and the domain $\mathcal{A}' = \left\{ \alpha \in \mathbb{R}_+^N : \alpha_i \leq 1, i = 1, \dots, N, \sum_{i=1}^N y_i \alpha_i = 0 \right\}$. Let

$\alpha^* \in \mathcal{A}'$ be such that $F(\alpha^*) = \sup_{\alpha \in \mathcal{A}'} F(\alpha)$. Let $w^* = \sum_{i=1}^N y_i \alpha_i^* x_i$. There is a threshold b^* (whose construction will be detailed in the proof), such that

$$C(w^*, b^*) = \inf_{w \in \mathbb{R}^d, b \in \mathbb{R}} C(w, b).$$

COROLLARY 3.6. (SCALED CRITERION) *For any positive real parameter λ let us consider the criterion*

$$C_\lambda(w, b) = \lambda^2 \sum_{i=1}^N [1 - (\langle w, x_i \rangle - b)y_i]_+ + \frac{1}{2} \|w\|^2$$

and the domain $\mathcal{A}'_\lambda = \left\{ \alpha \in \mathbb{R}_+^N : \alpha_i \leq \lambda^2, i = 1, \dots, N, \sum_{i=1}^N y_i \alpha_i = 0 \right\}$. For any solution α^* of the minimization problem $F(\alpha^*) = \sup_{\alpha \in \mathcal{A}'_\lambda} F(\alpha)$, the vector $w^* = \sum_{i=1}^N y_i \alpha_i^* x_i$ is such that

$$\inf_{b \in \mathbb{R}} C_\lambda(w^*, b) = \inf_{w \in \mathbb{R}^d, b \in \mathbb{R}} C_\lambda(w, b).$$

Let us remark that in the separable case, the scaled criterion is minimized by the canonical hyperplane for λ large enough. This extension of the canonical hyperplane computation in dual space is often called *the box constraint*, for obvious reasons.

PROOF. The corollary is a straightforward consequence of the scale property $C_\lambda(w, b, x) = \lambda^2 C(\lambda^{-1}w, b, \lambda x)$, where we have made the dependence of the criterion in $x \in \mathbb{R}^{dN}$ explicit. Let us come now to the proof of the theorem.

The minimization of $C(w, b)$ can be performed in dual space extending the couple of parameters (w, b) to $\bar{w} = (w, b, \gamma) \in \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}_+^N$ and introducing the dual multipliers $\alpha \in \mathbb{R}_+^N$ and the criterion

$$G(\alpha, \bar{w}) = \sum_{i=1}^N \gamma_i + \sum_{i=1}^N \alpha_i \{ [1 - (\langle w, x_i \rangle - b)y_i] - \gamma_i \} + \frac{1}{2} \|w\|^2.$$

We see that

$$C(w, b) = \inf_{\gamma \in \mathbb{R}_+^N} \sup_{\alpha \in \mathbb{R}_+^N} G[\alpha, (w, b, \gamma)],$$

and therefore, putting $\bar{\mathcal{W}} = \{(w, b, \gamma) : w \in \mathbb{R}^d, b \in \mathbb{R}, \gamma \in \mathbb{R}_+^N\}$, we are led to solve the minimization problem

$$G(\alpha_*, \bar{w}_*) = \inf_{\bar{w} \in \bar{\mathcal{W}}} \sup_{\alpha \in \mathbb{R}_+^N} G(\alpha, \bar{w}),$$

whose solution $\bar{w}_* = (w_*, b_*, \gamma_*)$ is such that $C(\bar{w}_*, b_*) = \inf_{(w, b) \in \mathbb{R}^{d+1}} C(w, b)$, according to the preceding identity. As for any value of $\alpha' \in \mathbb{R}_+^N$,

$$\inf_{\bar{w} \in \bar{\mathcal{W}}} \sup_{\alpha \in \mathbb{R}_+^N} G(\alpha, \bar{w}) \geq \inf_{\bar{w} \in \bar{\mathcal{W}}} G(\alpha', \bar{w}),$$

it is immediately seen that

$$\inf_{\bar{w} \in \bar{\mathcal{W}}} \sup_{\alpha \in \mathbb{R}_+^N} G(\alpha, \bar{w}) \geq \sup_{\alpha \in \mathbb{R}_+^N} \inf_{\bar{w} \in \bar{\mathcal{W}}} G(\alpha, \bar{w}).$$

We are going to show that there is no duality gap, meaning that this inequality is indeed an equality. More importantly, we will do so by exhibiting a saddle point, which, solving the dual minimization problem will also solve the original one.

Let us first make explicit the solution of the dual problem (the interest of this dual problem precisely lies in the fact that it can more easily be solved explicitly). Introducing the admissible set of values of α ,

$$\mathcal{A}' = \left\{ \alpha \in \mathbb{R}^N : 0 \leq \alpha_i \leq 1, i = 1, \dots, N, \sum_{i=1}^N y_i \alpha_i = 0 \right\},$$

it is elementary to check that

$$\inf_{\bar{w} \in \bar{\mathcal{W}}} G(\alpha, \bar{w}) = \begin{cases} \inf_{w \in \mathbb{R}^d} G[\alpha, (w, 0, 0)], & \alpha \in \mathcal{A}', \\ -\infty, & \text{otherwise.} \end{cases}$$

As

$$G[\alpha, (w, 0, 0)] = \frac{1}{2} \|w\|^2 + \sum_{i=1}^N \alpha_i (1 - \langle w, x_i \rangle y_i),$$

we see that $\inf_{w \in \mathbb{R}^d} G[\alpha, (w, 0, 0)]$ is reached at

$$w_\alpha = \sum_{i=1}^N y_i \alpha_i x_i.$$

This proves that

$$\inf_{\bar{w} \in \bar{\mathcal{W}}} G(\alpha, \bar{w}) = F(\alpha).$$

The continuous map $\alpha \mapsto \inf_{\bar{w} \in \bar{\mathcal{W}}} G(\alpha, \bar{w})$ reaches a (non necessarily unique) maximum α^* on the compact convex set \mathcal{A}' . We are now going to exhibit a choice of $\bar{w}^* \in \bar{\mathcal{W}}$ such that (α^*, \bar{w}^*) is a *saddle point*. This means that we are going to show that

$$G(\alpha^*, \bar{w}^*) = \inf_{\bar{w} \in \bar{\mathcal{W}}} G(\alpha^*, \bar{w}) = \sup_{\alpha \in \mathbb{R}_+^N} G(\alpha, \bar{w}^*).$$

It will imply that

$$\inf_{\bar{w} \in \overline{\mathcal{W}}} \sup_{\alpha \in \mathbb{R}_+^d} G(\alpha, \bar{w}) \leq \sup_{\alpha \in \mathbb{R}_+^N} G(\alpha, \bar{w}^*) = G(\alpha^*, \bar{w}^*)$$

on the one hand and that

$$\inf_{\bar{w} \in \overline{\mathcal{W}}} \sup_{\alpha \in \mathbb{R}_+^d} G(\alpha, \bar{w}) \geq \inf_{\bar{w} \in \overline{\mathcal{W}}} G(\alpha^*, \bar{w}) = G(\alpha^*, \bar{w}^*)$$

on the other hand, proving that

$$G(\alpha^*, \bar{w}^*) = \inf_{\bar{w} \in \overline{\mathcal{W}}} \sup_{\alpha \in \mathbb{R}_+^N} G(\alpha, \bar{w})$$

as required.

CONSTRUCTION OF \bar{w}^* .

- Let us put $w^* = w_{\alpha^*}$.
- If there is $j \in \{1, \dots, N\}$ such that $0 < \alpha_j^* < 1$, let us put

$$b^* = \langle x_j, w^* \rangle - y_j.$$

Otherwise, let us put

$$b^* = \sup\{\langle x_i, w^* \rangle - 1 : \alpha_i^* > 0, y_i = +1, i = 1, \dots, N\}.$$

- Let us then put

$$\gamma_i^* = \begin{cases} 0, & \alpha_i^* < 1, \\ 1 - (\langle w^*, x_i \rangle - b^*)y_i, & \alpha_i^* = 1. \end{cases}$$

If we can prove that

$$1 - (\langle w^*, x_i \rangle - b^*)y_i \begin{cases} \leq 0, & \alpha_i^* = 0, \\ = 0, & 0 < \alpha_i^* < 1, \\ \geq 0, & \alpha_i^* = 1, \end{cases} \quad (3.2)$$

it will show that $\gamma^* \in \mathbb{R}_+^N$ and therefore that $\bar{w}^* = (w^*, b^*, \gamma^*) \in \overline{\mathcal{W}}$. It will also show that

$$G(\alpha, \bar{w}^*) = \sum_{i=1}^N \gamma_i^* + \sum_{i, \alpha_i^*=0} \alpha_i [1 - (\langle \bar{w}^*, x_i \rangle - b^*)y_i] + \frac{1}{2} \|\bar{w}^*\|^2,$$

proving that $G(\alpha^*, \bar{w}^*) = \sup_{\alpha \in \mathbb{R}_+^N} G(\alpha, \bar{w}^*)$. As obviously $G(\alpha^*, \bar{w}^*) = G[\alpha^*, (w^*, 0, 0)]$, we already know that $G(\alpha^*, \bar{w}^*) = \inf_{\bar{w} \in \bar{\mathcal{W}}} G(\alpha^*, \bar{w})$. This will show that (α^*, \bar{w}^*) is the saddle point we were looking for, thus ending the proof of the theorem.

PROOF OF EQUATION (3.2): Let us deal first with the case when there is $j \in \{1, \dots, N\}$ such that $0 < \alpha_j^* < 1$.

For any $i \in \{1, \dots, N\}$ such that $0 < \alpha_i^* < 1$, there is $\epsilon > 0$ such that for any $t \in (-\epsilon, \epsilon)$, $\alpha^* + ty_i e_i - ty_j e_j \in \mathcal{A}'$, where $(e_k)_{k=1}^N$ is the canonical base of \mathbb{R}^N . Thus $\frac{\partial}{\partial t}|_{t=0} F(\alpha^* + ty_i e_i - ty_j e_j) = 0$. Computing this derivative, we obtain

$$\begin{aligned} \frac{\partial}{\partial t}|_{t=0} F(\alpha^* + ty_i e_i - ty_j e_j) &= y_i - \langle w^*, x_i \rangle + \langle w^*, x_j \rangle - y_j \\ &= y_i [1 - (\langle w, x_i \rangle - b^*) y_i]. \end{aligned}$$

Thus $1 - (\langle w, x_i \rangle - b^*) y_i = 0$, as required. This shows also that the definition of b^* does not depend on the choice of j such that $0 < \alpha_j^* < 1$.

For any $i \in \{1, \dots, N\}$ such that $\alpha_i^* = 0$, there is $\epsilon > 0$ such that for any $t \in (0, \epsilon)$, $\alpha^* + te_i - ty_i y_j e_j \in \mathcal{A}'$. Thus $\frac{\partial}{\partial t}|_{t=0} F(\alpha^* + te_i - ty_i y_j e_j) \leq 0$, showing that $1 - (\langle w^*, x_i \rangle - b^*) y_i \leq 0$ as required.

For any $i \in \{1, \dots, N\}$ such that $\alpha_i^* = 1$, there is $\epsilon > 0$ such that $\alpha^* - te_i + ty_i y_j e_j \in \mathcal{A}'$. Thus $\frac{\partial}{\partial t}|_{t=0} F(\alpha^* - te_i + ty_i y_j e_j) \leq 0$, showing that $1 - (\langle w^*, x_i \rangle - b^*) y_i \geq 0$ as required. This ends to prove that (α^*, \bar{w}^*) is a saddle point in this case.

Let us deal now with the case where $\alpha^* \in \{0, 1\}^N$. If we are not in the trivial case where the vector $(y_i)_{i=1}^N$ is constant, the case $\alpha^* = 0$ is ruled out. Indeed, in this case, considering $\alpha^* + te_i + te_j$, where $y_i y_j = -1$, we would get the contradiction $2 = \frac{\partial}{\partial t}|_{t=0} F(\alpha^* + te_i + te_j) \leq 0$.

Thus there are values of j such that $\alpha_j^* = 1$, and since $\sum_{i=1}^N \alpha_i y_i = 0$, both classes are present in the set $\{j : \alpha_j^* = 1\}$.

Now for any $i, j \in \{1, \dots, N\}$ such that $\alpha_i^* = \alpha_j^* = 1$ and such that $y_i = +1$ and $y_j = -1$, $\frac{\partial}{\partial t}|_{t=0} F(\alpha^* - te_i - te_j) = -2 + \langle w^*, x_i \rangle - \langle w^*, x_j \rangle \leq 0$. Thus

$$\sup\{\langle w^*, x_i \rangle - 1 : \alpha_i^* = 1, y_i = +1\} \leq \inf\{\langle w^*, x_j \rangle + 1 : \alpha_j^* = 1, y_j = -1\},$$

showing that

$$1 - (\langle w^*, x_k \rangle - b^*) y_k \geq 0, \alpha_k^* = 1.$$

Eventually, for any i such that $\alpha_i^* = 0$, for any j such that $\alpha_j^* = 1$ and $y_j = y_i$

$$\frac{\partial}{\partial t}|_{t=0} F(\alpha^* + te_i - te_j) = y_i \langle w^*, x_i - x_j \rangle \leq 0,$$

showing that $1 - (\langle w^*, x_i \rangle - b^*)y_i \leq 0$. This ends to prove that (α^*, \bar{w}^*) is in all circumstances a saddle point.

3.1.5. Support Vector Machines.

DEFINITION 3.3. The symmetric measurable kernel $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is said to be positive (or more precisely positive semi-definite) if for any $n \in \mathbb{N}$, any $(x_i)_{i=1}^n \in \mathcal{X}^n$,

$$\inf_{\alpha \in \mathbb{R}^n} \sum_{i=1}^n \sum_{j=1}^n \alpha_i K(x_i, x_j) \alpha_j \geq 0.$$

Let $Z = (x_i, y_i)_{i=1}^N$ be some training set. Let us consider as previously

$$\mathcal{A} = \left\{ \alpha \in \mathbb{R}_+^N : \sum_{i=1}^N \alpha_i y_i = 0 \right\}.$$

Let

$$F(\alpha) = \sum_{i=1}^N \sum_{j=1}^N \alpha_i y_i K(x_i, x_j) y_j \alpha_j - 2 \sum_{i=1}^N \alpha_i.$$

DEFINITION 3.4. Let K be a positive symmetric kernel. The training set Z is said to be K -separable if

$$\inf \{ F(\alpha) : \alpha \in \mathcal{A} \} > -\infty.$$

LEMMA 3.7. When Z is K -separable, $\inf \{ F(\alpha) : \alpha \in \mathcal{A} \}$ is reached.

PROOF. Consider the training set $Z' = (x'_i, y_i)_{i=1}^N$, where

$$x'_i = \left\{ \left[\left\{ K(x_k, x_\ell) \right\}_{k=1, \ell=1}^N \right]^{1/2} (i, j) \right\}_{j=1}^N \in \mathbb{R}^N.$$

We see that $F(\alpha) = \left\| \sum_{i=1}^N \alpha_i y_i x'_i \right\|^2 - 2 \sum_{i=1}^N \alpha_i$. We have proved in the previous section that Z' is linearly separable if and only if $\inf \{ F(\alpha) : \alpha \in \mathcal{A} \} > -\infty$, and that the infimum is reached in this case. \square

PROPOSITION 3.8. *Let K be a symmetric positive kernel and let $Z = (x_i, y_i)_{i=1}^N$ be some K -separable training set. Let $\alpha^* \in \mathcal{A}$ be such that $F(\alpha^*) = \inf\{F(\alpha) : \alpha \in \mathcal{A}\}$. Let*

$$I_-^* = \{i \in \mathbb{N} : 1 \leq i \leq N, y_i = -1, \alpha_i^* > 0\}$$

$$I_+^* = \{i \in \mathbb{N} : 1 \leq i \leq N, y_i = +1, \alpha_i^* > 0\}$$

$$b^* = \frac{1}{2} \left\{ \sum_{j=1}^N \alpha_j^* y_j K(x_j, x_{i_-}) + \sum_{j=1}^N \alpha_j^* y_j K(x_j, x_{i_+}) \right\}, \quad i_- \in I_-^*, i_+ \in I_+^*,$$

where the value of b^* does not depend on the choice of i_- and i_+ . The classification rule $f : \mathcal{X} \rightarrow \mathcal{Y}$ defined by the formula

$$f(x) = \text{sign} \left(\sum_{i=1}^N \alpha_i^* y_i K(x_i, x) - b^* \right)$$

is independent of the choice of α^* and is called the support vector machine defined by K and Z . The set $\mathcal{S} = \{x_j : \sum_{i=1}^N \alpha_i^* y_i K(x_i, x_j) - b^* = y_j\}$ is called the set of support vectors. For any choice of α^* , $\{x_i : \alpha_i^* > 0\} \subset \mathcal{S}$.

An important consequence of this proposition is that the support vector machine defined by K and Z is also the support vector machine defined by K and $Z' = \{(x_i, y_i) : \alpha_i^* > 0, 1 \leq i \leq N\}$, since this restriction of the index set contains the value α^* where the minimum of F is reached.

PROOF. The independence from the choice of α^* , which is not necessarily unique, is seen as follows. Let $(x_i)_{i=1}^N$ and $x \in \mathcal{X}$ be fixed. Let us put for ease of notations $x_{N+1} = x$. Let M be the $(N+1) \times (N+1)$ symmetric semi-definite matrix defined by $M(i, j) = K(x_i, x_j)$, $i = 1, \dots, N+1$, $j = 1, \dots, N+1$. Let us consider the mapping $\Psi : \{x_i : i = 1, \dots, N+1\} \rightarrow \mathbb{R}^{N+1}$ defined by

$$\Psi(x_i) = [M^{1/2}(i, j)]_{j=1}^{N+1} \in \mathbb{R}^{N+1}. \quad (3.3)$$

Let us consider the training set $Z' = [\Psi(x_i), y_i]_{i=1}^N$. Then Z' is linearly separable,

$$F(\alpha) = \left\| \sum_{i=1}^N \alpha_i y_i \Psi(x_i) \right\|^2 - 2 \sum_{i=1}^N \alpha_i,$$

and we have proved that for any choice of $\alpha^* \in \mathcal{A}$ minimizing $F(\alpha)$, $w_{Z'} = \sum_{i=1}^N \alpha_i^* y_i \Psi(x_i)$. Thus the support vector machine defined by K and Z can also be expressed by the formula

$$f(x) = \text{sign} \left[\langle w_{Z'}, \Psi(x) \rangle - b_{Z'} \right]$$

which does not depend on α^* . The definition of \mathcal{S} is such that $\Psi(\mathcal{S})$ is the set of support vectors defined in the linear case, where its stated property has already been proved. \square

We can in the same way use the box constraint and show that any solution $\alpha^* \in \arg \min\{F(\alpha) : \alpha \in \mathcal{A}, \alpha_i \leq \lambda^2, i = 1, \dots, N\}$ minimizes

$$\inf_{b \in \mathbb{R}} \lambda^2 \sum_{i=1}^N \left[1 - \left(\sum_{j=1}^N y_j \alpha_j K(x_j, x_i) - b \right) y_i \right]_+ + \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j y_i y_j K(x_i, x_j). \quad (3.4)$$

3.1.6. Building kernels. The results of this section (except the last one) are drawn from [18]. We have no reference for the last proposition of this section, although we believe it is well known. We include them for the convenience of the reader.

PROPOSITION 3.9. *Let K_1 and K_2 be positive symmetric kernels on \mathcal{X} . Then for any $a \in \mathbb{R}_+$*

$$(aK_1 + K_2)(x, x') \stackrel{\text{def}}{=} aK_1(x, x') + K_2(x, x')$$

$$\text{and } (K_1 \cdot K_2)(x, x') \stackrel{\text{def}}{=} K_1(x, x')K_2(x, x')$$

are also positive symmetric kernels. Moreover, for any measurable function $g : \mathcal{X} \rightarrow \mathbb{R}$, $K_g(x, x') \stackrel{\text{def}}{=} g(x)g(x')$ is also a positive symmetric kernel.

PROOF. It is enough to prove the proposition in the case when \mathcal{X} is finite and kernels are just ordinary symmetric matrices. Thus we can assume without loss of generality that $\mathcal{X} = \{1, \dots, n\}$. Then for any $\alpha \in \mathbb{R}^N$, using usual matrix notations,

$$\begin{aligned} \langle \alpha, (aK_1 + K_2)\alpha \rangle &= a\langle \alpha, K_1\alpha \rangle + \langle \alpha, K_2\alpha \rangle \geq 0, \\ \langle \alpha, (K_1 \cdot K_2)\alpha \rangle &= \sum_{i,j} \alpha_i K_1(i, j) K_2(i, j) \alpha_j \\ &= \sum_{i,j,k} \alpha_i K_1^{1/2}(i, k) K_1^{1/2}(k, j) K_2(i, j) \alpha_j \\ &= \sum_k \underbrace{\sum_{i,j} [K_1^{1/2}(k, i) \alpha_i] K_2(i, j) [K_1^{1/2}(k, j) \alpha_j]}_{\geq 0} \geq 0, \end{aligned}$$

$$\langle \alpha, K_g \alpha \rangle = \sum_{i,j} \alpha_i g(i) g(j) \alpha_j = \left(\sum_i \alpha_i g(i) \right)^2 \geq 0.$$

□

PROPOSITION 3.10. *Let K be some positive symmetric kernel on \mathcal{X} . Let $p : \mathbb{R} \rightarrow \mathbb{R}$ be a polynomial with positive coefficients. Let $g : \mathcal{X} \rightarrow \mathbb{R}^d$ be a measurable function. Then*

$$\begin{aligned} p(K)(x, x') &\stackrel{\text{def}}{=} p[K(x, x')], \\ \exp(K)(x, x') &\stackrel{\text{def}}{=} \exp[K(x, x')] \\ \text{and } G_g(x, x') &\stackrel{\text{def}}{=} \exp(-\|g(x) - g(x')\|^2) \end{aligned}$$

are all positive symmetric kernels.

PROOF. The first assertion is a direct consequence of the previous proposition. The second one comes from the fact that the exponential function is the pointwise limit of a sequence of polynomial functions with positive coefficients. The third one is seen from the second one and the decomposition

$$G_g(x, x') = \left[\exp(-\|g(x)\|^2) \exp(-\|g(x')\|^2) \right] \exp[2\langle g(x), g(x') \rangle]$$

□

PROPOSITION 3.11. *With the notations of the previous proposition, any training set $Z = (x_i, y_i)_{i=1}^N \in (\mathcal{X} \times \{-1, +1\})^N$ is G_g -separable as soon as $g(x_i)$, $i = 1, \dots, N$ are distinct points of \mathbb{R}^d .*

PROOF. It is clearly enough to prove the case when $\mathcal{X} = \mathbb{R}^d$ and g is the identity. Let us consider some other generic point $x_{N+1} \in \mathbb{R}^d$ and define Ψ as in (3.3). It is enough to prove that $\Psi(x_1), \dots, \Psi(x_N)$ are affine independent, since the simplex, and therefore any affine independent set of points can be shattered by affine half-spaces. Let us assume that (x_1, \dots, x_N) are affine dependent, this means that for some $(\lambda_1, \dots, \lambda_N) \neq 0$ such that $\sum_{i=1}^N \lambda_i = 0$,

$$\sum_{i=1}^N \sum_{j=1}^N \lambda_i G(x_i, x_j) \lambda_j = 0.$$

Thus, $(\lambda_i)_{i=1}^{N+1}$, where we have put $\lambda_{N+1} = 0$ is in the kernel of the symmetric positive semi-definite matrix $G(x_i, x_j)_{i,j \in \{1, \dots, N+1\}}$. Therefore

$$\sum_{i=1}^N \lambda_i G(x_i, x_{N+1}) = 0,$$

for any $x_{N+1} \in \mathbb{R}^d$. This would mean that the functions $x \mapsto \exp(-\|x - x_i\|^2)$ are linearly dependent, which can be easily proved to be false. Indeed, let $n \in \mathbb{R}^d$ be such that $\|n\| = 1$ and $\langle n, x_i \rangle$, $i = 1, \dots, N$ are distinct (such a vector exists, because it has to be outside the union of a finite number of hyperplanes, which is of zero Lebesgue measure on the sphere). Let us assume for a while that for some $(\lambda_i)_{i=1}^N \in \mathbb{R}^N$, for any $x \in \mathbb{R}^d$,

$$\sum_{i=1}^N \lambda_i \exp(-\|x - x_i\|^2) = 0.$$

Considering $x = tn$, for $t \in \mathbb{R}$, we would get

$$\sum_{i=1}^N \lambda_i \exp(2t\langle n, x_i \rangle - \|x_i\|^2) = 0, \quad t \in \mathbb{R}.$$

Letting t go to infinity, we see that this is only possible if $\lambda_i = 0$ for all values of i . \square

3.2. BOUNDS FOR SUPPORT VECTOR MACHINES.

3.2.1. Compression scheme bounds. We can use Support Vector Machines in the framework of compression schemes and apply Theorem 2.17 on page 114. More precisely, given some positive symmetric kernel K on \mathcal{X} , we may consider for any training set $Z' = (x'_i, y'_i)_{i=1}^h$ the classifier $\hat{f}_{Z'} : \mathcal{X} \rightarrow \mathcal{Y}$ which is equal to the Support Vector Machine defined by K and Z' whenever Z' is K -separable, and which is equal to some constant classification rule otherwise (we take this convention to stick to the framework described on page 105, we will only use $\hat{f}_{Z'}$ in the K -separable case, so this extension of the definition is just a matter of presentation). In the application of Theorem 2.17 in the case when the observed sample $(X_i, Y_i)_{i=1}^N$ is K -separable, a natural (if not always optimal) choice of Z' is to choose for (x'_i) the set of support vectors defined by $Z = (X_i, Y_i)_{i=1}^N$ and to choose for (y'_i) the corresponding values of Y . This is justified by the fact that $\hat{f}_Z = \hat{f}_{Z'}$, as shown in Proposition 3.8 (page 131). In the case when Z is not K -separable,

we can train a Support Vector Machine with the box constraint, then remove all the errors to obtain a K -separable subsample $Z' = \{(X_i, Y_i) : \alpha_i^* < \lambda^2, 1 \leq i \leq N\}$, (using the same notations as in equation (3.4) on page 132) and then consider its support vectors as the compression set. Still using the notations of page 132, this means we have to compute successively $\alpha^* \in \arg \min\{F(\alpha) : \alpha \in \mathcal{A}, \alpha_i \leq \lambda^2\}$, and $\alpha^{**} \in \arg \min\{F(\alpha) : \alpha \in \mathcal{A}, \alpha_i = 0 \text{ when } \alpha_i^* = \lambda^2\}$, to keep eventually the compression set indexed by $J = \{i : 1 \leq i \leq N, \alpha_i^{**} > 0\}$, and the corresponding Support Vector Machine \hat{f}_J . Different values of λ can be used at this stage, producing different candidate compression sets : when λ increases, the number of errors should decrease, on the other hand when λ decreases, the margin $\|w\|^{-1}$ of the separable subset Z' increases, supporting the hope for a smaller set of support vectors, thus we can use λ to monitor the number of errors on the training set we accept from the compression scheme. As we can use whatever heuristic we want while selecting the compression set, we can also try to threshold in the previous construction α_i^{**} at different levels $\eta \geq 0$, to produce candidate compression sets $J_\eta = \{i : 1 \leq i \leq N, \alpha_i^{**} > \eta\}$ of various sizes.

As the size $|J|$ of the compression set is random in this construction, we have to use a version of Theorem 2.17 (page 114) which handles compression sets of arbitrary sizes. This is done by choosing for each k a k -partially exchangeable posterior distribution π_k which weights the compression sets of all dimensions. We immediately see that we can choose π_k such that $-\log[\pi_k(\Delta_k(J))] \leq \log[|J|(|J| + 1)] + |J| \log\left[\frac{(k+1)\epsilon N}{|J|}\right]$.

If we observe the shadow sample patterns, and if computer resources permit, we can of course use more elaborate bounds than Theorem 2.17, such as the transductive correspondent to Theorem 1.24 (page 39) (where we may consider the submodels made of all the compression sets of the same size). Theorems based on relative bounds, such as Theorem 1.59 (page 88) can also be used. Gibbs distributions can be approximated by Monte Carlo techniques, where a Markov chain with the proper invariant measure consists in suitable local perturbations of the compression set.

Let us mention also that the use of compression schemes based on Support Vector Machines can be tailored to perform some kind of *feature aggregation*. Imagine that the kernel K is defined as the scalar product in $L_2(\pi)$, where $\pi \in \mathcal{M}_+^1(\Theta)$. More precisely let us consider for some set of soft classification rules $\{f_\theta : \mathcal{X} \rightarrow \mathbb{R}; \theta \in \Theta\}$ the kernel

$$K(x, x') = \int_{\theta \in \Theta} f_\theta(x) f_\theta(x') \pi(d\theta).$$

In this setting, the Support Vector Machine applied to the training set $Z =$

$(x_i, y_i)_{i=1}^N$ has the form

$$f_Z(x) = \text{sign} \left(\int_{\theta \in \Theta} f_\theta(x) \sum_{i=1}^N y_i \alpha_i f_\theta(x_i) \pi(d\theta) - b \right)$$

and, may it be too burdening to compute, we can replace it with some finite approximation

$$\tilde{f}_Z(x) = \text{sign} \left(\sum_{k=1}^m f_{\theta_k}(x) w_k - b \right),$$

where the set $\{\theta_k, k = 1, \dots, m\}$ and the weights $\{w_k, k = 1, \dots, m\}$ are computed in some suitable way from $Z' = (x_i, y_i)_{i, \alpha_i > 0}$, the set of support vectors of f_Z . For instance, we can draw $\{\theta_k, k = 1, \dots, m\}$ at random according to the probability distribution proportional to

$$\left| \sum_{i=1}^N y_i \alpha_i f_\theta(x_i) \right| \pi(d\theta),$$

define the weights w_k by

$$w_k = \text{sign} \left(\sum_{i=1}^N y_i \alpha_i f_{\theta_k}(x_i) \right) \int_{\theta \in \Theta} \left| \sum_{i=1}^N y_i \alpha_i f_\theta(x_i) \right| \pi(d\theta),$$

and choose the smallest value of m for which this approximation still classifies Z' without errors. Let us remark that we have built \tilde{f}_Z in such a way that

$$\lim_{m \rightarrow +\infty} \tilde{f}_Z(x_i) = f_Z(x_i) = y_i, \quad \text{a.s.}$$

for any support index i such that $\alpha_i > 0$.

Alternatively, given Z' , we can select a finite set of features $\Theta' \subset \Theta$ such that Z' is $K_{\Theta'}$ separable, where $K_{\Theta'}(x, x') = \sum_{\theta \in \Theta'} f_\theta(x) f_\theta(x')$ and consider the Support Vector Machines $f_{Z'}$ built with the kernel $K_{\Theta'}$. As soon as Θ' is chosen as a function of Z' only, Theorem 2.17 (page 114) applies and provides some level of confidence for the risk of $f_{Z'}$.

3.2.2. The Vapnik Cervonenkis dimension of a family of subsets. Let us consider some set X and some set $S \subset \{0, 1\}^X$ of subsets of X . Let $h(S)$ be the VC dimension of S , defined as

$$h(S) = \max\{|A| : A \text{ finite and } A \cap S = \{0, 1\}^A\},$$

where by definition $A \cap S = \{A \cap B : B \in S\}$. Let us notice that this definition does not depend on the choice of the reference set X . Indeed X can be chosen to be $\bigcup S$, the union of all the sets in S or any bigger set. Let us notice also that for any set B , $h(B \cap S) \leq h(S)$, the reason being that $A \cap (B \cap S) = B \cap (A \cap S)$.

This notion of VC dimension is useful because it can, as we will see about Support Vector Machines, be computed in some important special cases. Let us prove here as an illustration that $h(S) = d + 1$ when $X = \mathbb{R}^d$ and S is made of all the half spaces :

$$S = \{A_{w,b} : w \in \mathbb{R}^d, b \in \mathbb{R}\}, \text{ where } A_{w,b} = \{x \in X : \langle w, x \rangle \geq b\}.$$

PROPOSITION 3.12. *With the previous notations, $h(S) = d + 1$.*

PROOF. Let $(e_i)_{i=1}^{d+1}$ be the canonical base of \mathbb{R}^{d+1} , and let X be the affine subspace it generates, which can be identified with \mathbb{R}^d . For any $(\epsilon_i)_{i=1}^{d+1} \in \{-1, +1\}^{d+1}$, let $w = \sum_{i=1}^{d+1} \epsilon_i e_i$ and $b = 0$. The half space $A_{w,b} \cap X$ is such that $\{e_i ; i = 1, \dots, d+1\} \cap (A_{w,b} \cap X) = \{e_i ; \epsilon_i = +1\}$. This proves that $h(S) \geq d + 1$.

To prove that $h(S) \leq d+1$, we have to show that for any set $A \subset \mathbb{R}^d$ of size $|A| = d+2$, there is $B \subset A$ such that $B \not\subset (A \cap S)$. This will obviously be the case if the convex hulls of B and $A \setminus B$ have a non empty intersection : indeed if a hyperplane separates two sets of points, it also separates their convex hulls. As $|A| > d + 1$, A is affine dependent : there is $(\lambda_x)_{x \in A} \in \mathbb{R}^{d+2} \setminus \{0\}$ such that $\sum_{x \in A} \lambda_x x = 0$ and $\sum_{x \in A} \lambda_x = 0$. The set $B = \{x \in A : \lambda_x > 0\}$ is non-empty, as well as its complement $A \setminus B$, because $\sum_{x \in A} \lambda_x = 0$ and $\lambda \neq 0$. Moreover $\sum_{x \in B} \lambda_x = \sum_{x \in A \setminus B} -\lambda_x > 0$. The relation

$$\frac{1}{\sum_{x \in B} \lambda_x} \sum_{x \in B} \lambda_x x = \frac{1}{\sum_{x \in B} \lambda_x} \sum_{x \in A \setminus B} -\lambda_x x$$

shows that the convex hulls of B and $A \setminus B$ have a non void intersection. \square

Let us introduce the function of two integers

$$\Phi_n^h = \sum_{k=0}^h \binom{n}{k}$$

Let us notice that Φ can alternatively be defined by the relations :

$$\Phi_n^h = \begin{cases} 2^n & \text{when } n \leq h, \\ \Phi_{n-1}^{h-1} + \Phi_{n-1}^h & \text{when } n > h. \end{cases}$$

THEOREM 3.13. *Whenever $\bigcup S$ is finite,*

$$|S| \leq \Phi \left(\left| \bigcup S \right|, h(S) \right).$$

THEOREM 3.14. *For any $h \leq n$,*

$$\Phi_n^h \leq \exp\left(nH\left(\frac{h}{n}\right)\right) \leq \exp\left[h\left(\log\left(\frac{n}{h}\right) + 1\right)\right],$$

where $H(p) = -p \log(p) - (1-p) \log(1-p)$ is the Shannon entropy of the Bernoulli distribution with parameter p .

PROOF OF THEOREM 3.13. Let us prove this theorem by induction on $|\bigcup S|$. It is easy to check that it holds true when $|\bigcup S| = 1$. Let $X = \bigcup S$, let $x \in X$ and $X' = X \setminus \{x\}$. Define (Δ denoting the symmetric difference of two sets)

$$\begin{aligned} S' &= \{A \in S : A \Delta \{x\} \in S\}, \\ S'' &= \{A \in S : A \Delta \{x\} \notin S\}. \end{aligned}$$

Clearly, \sqcup denoting the disjoint union, $S = S' \sqcup S''$ and $S \cap X' = (S' \cap X') \sqcup (S'' \cap X')$. Moreover $|S'| = 2|S' \cap X'|$ and $|S''| = |S'' \cap X'|$. Thus $|S| = |S'| + |S''| = 2|S' \cap X'| + |S''| = |S \cap X'| + |S' \cap X'|$. Obviously $h(S \cap X') \leq h(S)$. Moreover $h(S' \cap X') = h(S') - 1$, because if $A \subset X'$ is shattered by S' (or equivalently by $S' \cap X'$), then $A \cup \{x\}$ is shattered by S' (we say that A is shattered by S when $S \cap A = \{0, 1\}^A$). Using the induction hypothesis, we then see that $|S \cap X'| \leq \Phi_{|X'|}^{h(S)} + \Phi_{|X'|}^{h(S)-1}$. But as $|X'| = |X| - 1$, the righthand side of this inequality is equal to $\Phi_{|X|}^{h(S)}$, according to the recurrence equation satisfied by Φ .

PROOF OF THEOREM 3.14: This is the well known Chernoff bound for the deviation of sums of Bernoulli r.v.: let $(\sigma_1, \dots, \sigma_n)$ be i.i.d. Bernoulli r.v. with parameter $1/2$. Let us notice that

$$\Phi_n^h = 2^n \mathbb{P} \left(\sum_{i=1}^n \sigma_i \leq h \right).$$

For any positive real number λ ,

$$\begin{aligned} \mathbb{P} \left(\sum_{i=1}^n \sigma_i \leq h \right) &\leq \exp(\lambda h) \mathbb{E} \left[\exp \left(-\lambda \sum_{i=1}^n \sigma_i \right) \right] \\ &= \exp \left\{ \lambda h + n \log \left\{ \mathbb{E} \left[\exp(-\lambda \sigma_1) \right] \right\} \right\}. \end{aligned}$$

Differentiating the right-hand side in λ shows that its minimal value is $\exp[-n\mathcal{K}(\frac{h}{n}, \frac{1}{2})]$, where $\mathcal{K}(p, q) = p \log(\frac{p}{q}) + (1-p) \log(\frac{1-p}{1-q})$ is the Kullback divergence function between two Bernoulli distributions B_p and B_q of parameters p and q . Indeed the optimal value λ^* of λ is such that $h = n \frac{\mathbb{E}[\sigma_1 \exp(-\lambda^* \sigma_1)]}{\mathbb{E}[\exp(-\lambda^* \sigma_1)]} = nB_{h/n}(\sigma_1)$. Therefore (using the fact that two Bernoulli distributions with the same expectations are equal)

$$\log\{\mathbb{E}[\exp(-\lambda^* \sigma_1)]\} = -\lambda^* B_{h/n}(\sigma_1) - \mathcal{K}(B_{h/n}, B_{1/2}) = -\lambda^* \frac{h}{n} - \mathcal{K}(\frac{h}{n}, \frac{1}{2}).$$

The announced result then follows from the identity

$$\begin{aligned} H(p) &= \log(2) - \mathcal{K}(p, \frac{1}{2}) \\ &= p \log(p^{-1}) + (1-p) \log(1 + \frac{p}{1-p}) \leq p[\log(p^{-1}) + 1]. \end{aligned}$$

3.2.3. VC dimension of linear rules with margin. The proof of the following theorem has been suggested to us by a similar proof presented in [18].

THEOREM 3.15. *Consider a family of points (x_1, \dots, x_n) in some Euclidean vector space E and a family of affine functions*

$$\mathcal{H} = \{g_{w,b} : E \rightarrow \mathbb{R}; w \in E, \|w\| = 1, b \in \mathbb{R}\},$$

where

$$g_{w,b}(x) = \langle w, x \rangle - b, \quad x \in E.$$

Assume that there is a set of thresholds $(b_i)_{i=1}^n \in \mathbb{R}^n$ such that for any $(y_i)_{i=1}^n \in \{-1, +1\}^n$, there is $g_{w,b} \in \mathcal{H}$ such that

$$\inf_{i=1}^n (g_{w,b}(x_i) - b_i) y_i \geq \gamma.$$

Let us also introduce the empirical variance of $(x_i)_{i=1}^n$,

$$\text{Var}(x_1, \dots, x_n) = \frac{1}{n} \sum_{i=1}^n \left\| x_i - \frac{1}{n} \sum_{j=1}^n x_j \right\|^2.$$

In this case and with these notations,

$$\frac{\text{Var}(x_1, \dots, x_n)}{\gamma^2} \geq \begin{cases} n-1 & \text{when } n \text{ is even,} \\ (n-1) \frac{n^2-1}{n^2} & \text{when } n \text{ is odd.} \end{cases} \quad (3.5)$$

Moreover, equality is reached when γ is optimal, $b_i = 0$, $i = 1, \dots, n$ and (x_1, \dots, x_n) is a regular simplex (i.e. when 2γ is the minimum distance between the convex hulls of any two subsets of $\{x_1, \dots, x_n\}$ and $\|x_i - x_j\|$ does not depend on $i \neq j$).

PROOF. Let $(s_i)_{i=1}^n \in \mathbb{R}^n$ be such that $\sum_{i=1}^n s_i = 0$. Let σ be a uniformly distributed random variable with values in \mathfrak{S}_n , the set of permutations of the first n integers $\{1, \dots, n\}$. By assumption, for any value of σ , there is an affine function $g_{w,b} \in \mathcal{H}$ such that

$$\min_{i=1, \dots, n} [g_{w,b}(x_i) - b_i] [2\mathbf{1}(s_{\sigma(i)} > 0) - 1] \geq \gamma.$$

As a consequence

$$\begin{aligned} \left\langle \sum_{i=1}^n s_{\sigma(i)} x_i, w \right\rangle &= \sum_{i=1}^n s_{\sigma(i)} (\langle x_i, w \rangle - b - b_i) + \sum_{i=1}^n s_{\sigma(i)} b_i \\ &\geq \sum_{i=1}^n \gamma |s_{\sigma(i)}| + s_{\sigma(i)} b_i. \end{aligned}$$

Therefore, using the fact that the map $x \mapsto (\max\{0, x\})^2$ is convex,

$$\begin{aligned} \mathbb{E} \left(\left\| \sum_{i=1}^n s_{\sigma(i)} x_i \right\|^2 \right) &\geq \mathbb{E} \left[\left(\max \left\{ 0, \sum_{i=1}^n \gamma |s_{\sigma(i)}| + s_{\sigma(i)} b_i \right\} \right)^2 \right] \\ &\geq \left(\max \left\{ 0, \sum_{i=1}^n \gamma \mathbb{E}(|s_{\sigma(i)})| + \mathbb{E}(s_{\sigma(i)}) b_i \right\} \right)^2 = \gamma^2 \left(\sum_{i=1}^n |s_i| \right)^2, \end{aligned}$$

where \mathbb{E} is the expectation with respect to the random permutation σ . On the other hand

$$\mathbb{E} \left(\left\| \sum_{i=1}^n s_{\sigma(i)} x_i \right\|^2 \right) = \sum_{i=1}^n \mathbb{E}(s_{\sigma(i)}^2) \|x_i\|^2 + \sum_{i \neq j} \mathbb{E}(s_{\sigma(i)} s_{\sigma(j)}) \langle x_i, x_j \rangle.$$

Moreover

$$\mathbb{E}(s_{\sigma(i)}^2) = \frac{1}{n} \mathbb{E} \left(\sum_{i=1}^n s_{\sigma(i)}^2 \right) = \frac{1}{n} \sum_{i=1}^n s_i^2.$$

In the same way, for any $i \neq j$,

$$\begin{aligned} \mathbb{E}(s_{\sigma(i)} s_{\sigma(j)}) &= \frac{1}{n(n-1)} \mathbb{E} \left(\sum_{i \neq j} s_{\sigma(i)} s_{\sigma(j)} \right) \\ &= \frac{1}{n(n-1)} \sum_{i \neq j} s_i s_j \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{n(n-1)} \left[\underbrace{\left(\sum_{i=1}^n s_i \right)^2}_{=0} - \sum_{i=1}^n s_i^2 \right] \\
&= -\frac{1}{n(n-1)} \sum_{i=1}^n s_i^2.
\end{aligned}$$

Thus

$$\begin{aligned}
\mathbb{E} \left(\left\| \sum_{i=1}^n s_{\sigma(i)} x_i \right\|^2 \right) &= \left(\sum_{i=1}^n s_i^2 \right) \left[\frac{1}{n} \sum_{i=1}^n \|x_i\|^2 - \frac{1}{n(n-1)} \sum_{i \neq j} \langle x_i, x_j \rangle \right] \\
&= \left(\sum_{i=1}^n s_i^2 \right) \left[\left(\frac{1}{n} + \frac{1}{n(n-1)} \right) \sum_{i=1}^n \|x_i\|^2 \right. \\
&\quad \left. - \frac{1}{n(n-1)} \left\| \sum_{i=1}^n x_i \right\|^2 \right] \\
&= \frac{n}{n-1} \left(\sum_{i=1}^n s_i^2 \right) \text{Var}(x_1, \dots, x_n).
\end{aligned}$$

We have proved that

$$\frac{\text{Var}(x_1, \dots, x_n)}{\gamma^2} \geq \frac{(n-1) \left(\sum_{i=1}^n |s_i| \right)^2}{n \sum_{i=1}^n s_i^2}.$$

This can be used with $s_i = \mathbf{1}(i \leq \frac{n}{2}) - \mathbf{1}(i > \frac{n}{2})$ in the case when n is even and $s_i = \frac{2}{(n-1)} \mathbf{1}(i \leq \frac{n-1}{2}) - \frac{2}{n+1} \mathbf{1}(i > \frac{n-1}{2})$ in the case when n is odd to establish the first inequality (3.5) of the theorem.

Checking that equality is reached for the simplex is an easy computation when the simplex $(x_i)_{i=1}^n \in (\mathbb{R}^n)^n$ is parametrized in such a way that

$$x_i(j) = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

Indeed the distance between the convex hulls of any two subsets of the simplex is the distance between their mean values (i.e. centers of mass). \square

3.2.4. Application to Support Vector Machines. We are going to apply Theorem 3.15 (page 139) to Support Vector Machines in the transductive case. So let us consider $(X_i, Y_i)_{i=1}^{(k+1)N}$ distributed according to some partially exchangeable distribution \mathbb{P} and assume that $(X_i)_{i=1}^{(k+1)N}$ and $(Y_i)_{i=1}^N$ are observed. Let us consider some positive kernel K on \mathcal{X} . For any K -separable training set of the form $Z' = (X_i, y'_i)_{i=1}^{(k+1)N}$, where $(y'_i)_{i=1}^{(k+1)N} \in \mathcal{Y}^{(k+1)N}$, let $\hat{f}_{Z'}$ be the Support Vector Machine defined by K and Z' and let $\gamma(Z')$ be its margin. Let

$$R^2 = \max_{i=1, \dots, (k+1)N} K(X_i, X_i) + \frac{1}{(k+1)^2 N^2} \sum_{j=1}^{(k+1)N} \sum_{k=1}^{(k+1)N} K(X_j, X_k) - \frac{2}{(k+1)N} \sum_{j=1}^{(k+1)N} K(X_i, X_j).$$

(This is an easily computable upper-bound for the radius of some ball containing the image of $(X_1, \dots, X_{(k+1)N})$ in feature space.)

Let us define for any integer h the margins

$$\gamma_{2h} = (2h-1)^{-1/2} \text{ and } \gamma_{2h+1} = \left[2h \left(1 - \frac{1}{(2h+1)^2} \right) \right]^{-1/2}. \quad (3.6)$$

Let us consider for any $h = 1, \dots, N$ the exchangeable model

$$\mathcal{R}_h = \{ \hat{f}_{Z'} : Z' = (X_i, y'_i)_{i=1}^{(k+1)N} \text{ is } K\text{-separable and } \gamma(Z') \geq R\gamma_h \}.$$

The family of models \mathcal{R}_h , $h = 1, \dots, N$ is nested, and we know from Theorem 3.15 (page 139) and Theorems 3.13 (page 138) and 3.14 (page 138) that

$$\log(|\mathcal{R}_h|) \leq h \log\left(\frac{(k+1)eN}{h}\right).$$

We can then consider on the large model $\mathcal{R} = \bigsqcup_{h=1}^N \mathcal{R}_h$ (the disjoint union of the submodels) an exchangeable prior π which is uniform on each \mathcal{R}_h and is such that $\pi(\mathcal{R}_h) \geq \frac{1}{h(h+1)}$. Applying Theorem 2.8 (page 104) we get

PROPOSITION 3.16. *With \mathbb{P} probability at least $1 - \epsilon$, for any $h = 1, \dots, N$, any Support Vector Machine $f \in \mathcal{R}_h$,*

$$r_2(f) \leq \frac{k+1}{k} \inf_{\lambda \in \mathbb{R}_+} \frac{1 - \exp\left[-\frac{\lambda}{N} r_1(f) - \frac{h}{N} \log\left(\frac{e(k+1)N}{h}\right) - \frac{\log[h(h+1)] - \log(\epsilon)}{N}\right]}{1 - \exp\left(-\frac{\lambda}{N}\right)} - \frac{r_1(f)}{k}.$$

Searching the whole model \mathcal{R}_h may be unfeasible, nonetheless any heuristic can be applied to choose f . For instance, a Support Vector Machine f' can be trained from the training set $(X_i, Y_i)_{i=1}^N$ and then $(y'_i)_{i=1}^{(k+1)N}$ can be set to $y'_i = \text{sign}(f'(X_i))$, $i = 1, \dots, (k+1)N$.

3.2.5. Inductive margin bounds for Support Vector Machines. In order to establish inductive margin bounds, we will need a different combinatorial lemma. It is due to [1]. We will reproduce their proof with some tiny improvements on the values of constants.

Let us consider the finite case when $\mathcal{X} = \{1, \dots, n\}$, $\mathcal{Y} = \{1, \dots, b\}$ and $b \geq 3$ (the question we will study would be meaningless in the case when $b \leq 2$). Assume as usual that we are dealing with a prescribed set of classification rules

$\mathcal{R} = \{f : \mathcal{X} \rightarrow \mathcal{Y}\}$. Let us say that a pair (A, s) , where $A \subset \mathcal{X}$ is a non empty set of shapes and $s : A \rightarrow \{2, \dots, b-1\}$ a threshold function, is *shattered* by the set of functions $F \subset \mathcal{R}$ if for any $(\sigma_x)_{x \in A} \in \{-1, +1\}^A$, there exists some $f \in F$ such that $\min_{x \in A} \sigma_x [f(x) - s(x)] \geq 1$.

DEFINITION 3.5. Let the *fat shattering dimension* of $(\mathcal{X}, \mathcal{R})$ be the maximal size $|A|$ of the first component of the pairs which are shattered by \mathcal{R} .

Let us say that a subset of classification rules $F \subset \mathcal{Y}^{\mathcal{X}}$ is *separated* whenever for any pair $(f, g) \in F^2$ such that $f \neq g$, $\|f - g\|_{\infty} = \max_{x \in \mathcal{X}} |f(x) - g(x)| \geq 2$. Let $\mathfrak{M}(\mathcal{R})$ be the maximum size $|F|$ of separated subsets F of \mathcal{R} . Note that if F is a separated subset of \mathcal{R} such that $|F| = \mathfrak{M}(\mathcal{R})$, then it is a 1-net for the \mathcal{L}_{∞} distance: for any function $f \in \mathcal{R}$ there exists $g \in F$ such that $\|f - g\|_{\infty} \leq 1$ (otherwise f could be added to F to create a larger separated set).

LEMMA 3.17. *With the above notations, whenever the fat shattering dimension of $(\mathcal{X}, \mathcal{R})$ is not greater than h ,*

$$\begin{aligned} \log[\mathfrak{M}(\mathcal{R})] &< \log[(b-1)(b-2)n] \left\{ \frac{\log[\sum_{i=1}^h \binom{n}{i} (b-2)^i]}{\log(2)} + 1 \right\} + \log(2) \\ &\leq \log[(b-1)(b-2)n] \left\{ \left[\log\left[\frac{(b-2)n}{h}\right] + 1 \right] \frac{h}{\log(2)} + 1 \right\} + \log(2). \end{aligned}$$

PROOF. For any set of functions $F \subset \mathcal{Y}^{\mathcal{X}}$, let $t(F)$ be the number of pairs (A, s) shattered by F . Let $t(m, n)$ be the minimum of $t(F)$ over all *separated* sets of functions $F \subset \mathcal{Y}^{\mathcal{X}}$ of size $|F| = m$ (n is here to recall that the shape

space \mathcal{X} is made of n shapes). For any m such that $t(m, n) > \sum_{i=1}^h \binom{n}{i} (b-2)^i$, it is clear that any separated set of functions of size $|F| \geq m$ shatters at least one pair (A, s) such that $|A| > h$. Indeed, $t(m, n)$ is clearly from its definition a non decreasing function of m , so that $t(|F|, n) > \sum_{i=1}^h \binom{n}{i} (b-2)^i$. Moreover there are only $\sum_{i=1}^h \binom{n}{i} (b-2)^i$ pairs (A, s) such that $|A| \leq h$. As a consequence, whenever the fat shattering dimension of $(\mathcal{X}, \mathcal{R})$ is not greater than h we have $\mathfrak{M}(\mathcal{R}) < m$.

It is clear that for any $n \geq 1$, $t(2, n) = 1$.

LEMMA 3.18. *For any $m \geq 1$, $t[mn(b-1)(b-2), n] \geq 2t[m, n-1]$, and therefore $t[2n(n-1) \dots (n-r+1)(b-1)^r(b-2)^r, n] \geq 2^r$.*

PROOF. Let $F = \{f_1, \dots, f_{mn(b-1)(b-2)}\}$ be some separated set of functions of size $mn(b-1)(b-2)$. For any pair (f_{2i-1}, f_{2i}) , $i = 1, \dots, mn(b-1)(b-2)/2$, there is $x_i \in \mathcal{X}$ such that $|f_{2i-1}(x_i) - f_{2i}(x_i)| \geq 2$. Since $|\mathcal{X}| = n$, there is $x \in \mathcal{X}$ such that $\sum_{i=1}^{mn(b-1)(b-2)/2} \mathbf{1}(x_i = x) \geq m(b-1)(b-2)/2$. Let $I = \{i : x_i = x\}$. Since there are $(b-1)(b-2)/2$ pairs $(y_1, y_2) \in \mathcal{Y}^2$ such that $1 \leq y_1 < y_2 - 1 \leq b-1$, there is some pair (y_1, y_2) , such that $1 \leq y_1 < y_2 \leq b$ and such that $\sum_{i \in I} \mathbf{1}(\{y_1, y_2\} = \{f_{2i-1}(x), f_{2i}(x)\}) \geq m$. Let $J = \{i \in I : \{f_{2i-1}(x), f_{2i}(x)\} = \{y_1, y_2\}\}$. Let

$$\begin{aligned} F_1 &= \{f_{2i-1} : i \in J, f_{2i-1}(x) = y_1\} \cup \{f_{2i} : i \in J, f_{2i}(x) = y_1\}, \\ F_2 &= \{f_{2i-1} : i \in J, f_{2i-1}(x) = y_2\} \cup \{f_{2i} : i \in J, f_{2i}(x) = y_2\}. \end{aligned}$$

Obviously $|F_1| = |F_2| = |J| = m$. Moreover the restrictions of the functions of F_1 to $\mathcal{X} \setminus \{x\}$ are separated, and it is the same with F_2 . Thus F_1 strongly shatters at least $t(m, n-1)$ pairs (A, s) such that $A \subset \mathcal{X} \setminus \{x\}$ and it is the same with F_2 . Eventually, if the pair (A, s) where $A \subset \mathcal{X} \setminus \{x\}$ is both shattered by F_1 and F_2 , then $F_1 \cup F_2$ shatters also $(A \cup \{x\}, s')$ where $s'(x') = s(x')$ for any $x' \in A$ and $s'(x) = \lfloor \frac{y_1+y_2}{2} \rfloor$. Thus $F_1 \cup F_2$, and therefore F , shatters at least $2t(m, n-1)$ pairs (A, s) . \square

Resuming the proof of lemma 3.17, let us choose for r the smallest integer such that $2^r > \sum_{i=1}^h \binom{n}{i} (b-2)^i$, which is not greater than

$$\left\{ \frac{\log \left[\sum_{i=1}^h \binom{n}{i} (b-2)^i \right]}{\log(2)} + 1 \right\}.$$

In the case when $1 \leq n \leq r$,

$$\log(\mathfrak{M}(\mathcal{R})) < |\mathcal{X}| \log(|\mathcal{Y}|) = n \log(b) \leq r \log(b) \leq r \log[(b-1)(b-2)n] + \log(2),$$

which proves the lemma. In the remaining case $n > r$,

$$\begin{aligned}
& t[2n^r(b-1)^r(b-2)^r, n] \\
& \geq t[2n(n-1)\dots(n-r+1)(b-1)^r(b-2)^r, n] \\
& > \sum_{i=1}^h \binom{n}{i} (b-2)^i.
\end{aligned}$$

Thus $|\mathfrak{M}(\mathcal{R})| < 2[(b-2)(b-1)n]^r$ as claimed. \square

In order to apply this combinatorial lemma to Support Vector Machines, let us consider now the case of separating hyperplanes in \mathbb{R}^d (the generalization to Support Vector Machines being straightforward). Assume that $\mathcal{X} = \mathbb{R}^d$ and $\mathcal{Y} = \{-1, +1\}$. For any sample $(X)_{i=1}^{(k+1)N}$, let

$$R(X_1^{(k+1)N}) = \max\{\|X_i\| : 1 \leq i \leq (k+1)N\}.$$

Let us consider the set of parameters

$$\Theta = \{(w, b) \in \mathbb{R}^d \times \mathbb{R} : \|w\| = 1\}.$$

For any $(w, b) \in \Theta$, let $g_{w,b}(x) = \langle w, x \rangle - b$. Let h be some fixed integer and let $\gamma = R(X_1^{(k+1)N})\gamma_h$, where γ_h is defined by equation (3.6) on page 142.

Let us define $\zeta : \mathbb{R} \rightarrow \mathbb{Z}$ by

$$\zeta(r) = \begin{cases} -5 & \text{when } r \leq -4\gamma, \\ -3 & \text{when } -4\gamma < r \leq -2\gamma, \\ -1 & \text{when } -2\gamma < r \leq 0, \\ +1 & \text{when } 0 < r \leq 2\gamma, \\ +3 & \text{when } 2\gamma < r \leq 4\gamma, \\ +5 & \text{when } 4\gamma < r. \end{cases}$$

Let $G_{w,b}(x) = \zeta[g_{w,b}(x)]$. The fat shattering dimension (as defined in 3.5) of

$$\left(X_1^{(k+1)N}, \{(G_{w,b} + 7)/2 : (w, b) \in \Theta\}\right)$$

is not greater than h (according to Theorem 3.15, page 139), therefore there is some set \mathcal{F} of functions from $X_1^{(k+1)N}$ to $\{-5, -3, -1, +1, +3, +5\}$ such that

$$\log(|\mathcal{F}|) \leq \log[20(k+1)N] \left\{ \frac{h}{\log(2)} \left[\log \left(\frac{4(k+1)N}{h} \right) + 1 \right] + 1 \right\} + \log(2).$$

and for any $(w, b) \in \Theta$, there is $f_{w,b} \in \mathcal{F}$ such that $\sup\{|f_{w,b}(X_i) - G_{w,b}(X_i)| : i = 1, \dots, (k+1)N\} \leq 2$. Moreover, the choice of $f_{w,b}$ may be required to depend on $(X_i)_{i=1}^{(k+1)N}$ in an exchangeable way. Similarly to Theorem 2.8 (page 104), it can be proved that for any partially exchangeable probability distribution $\mathbb{P} \in \mathcal{M}_+^1(\Omega)$, with \mathbb{P} probability at least $1 - \epsilon$, for any $f_{w,b} \in \mathcal{F}$,

$$\begin{aligned} & \frac{1}{kN} \sum_{i=N+1}^{(k+1)N} \mathbf{1}[f_{w,b}(X_i)Y_i \leq 1] \\ & \leq \frac{k+1}{k} \inf_{\lambda \in \mathbb{R}_+} [1 - \exp(-\frac{\lambda}{N})]^{-1} \left\{ 1 - \right. \\ & \quad \left. \exp\left[-\frac{\lambda}{N^2} \sum_{i=1}^N \mathbf{1}[f_{w,b}(X_i)Y_i \leq 1] - \frac{\log(|\mathcal{F}|) - \log(\epsilon)}{N}\right]\right\} \\ & \quad - \frac{1}{kN} \sum_{i=1}^N \mathbf{1}[f_{w,b}(X_i)Y_i \leq 1]. \end{aligned}$$

Let us remark that

$$\mathbf{1}\left\{2\mathbf{1}[g_{w,b}(X_i) \geq 0] - 1 \neq Y_i\right\} = \mathbf{1}[G_{w,b}(X_i)Y_i < 0] \leq \mathbf{1}[f_{w,b}(X_i)Y_i \leq 1]$$

and

$$\mathbf{1}[f_{w,b}(X_i)Y_i \leq 1] \leq \mathbf{1}[G_{w,b}(X_i)Y_i \leq 3] \leq \mathbf{1}[g_{w,b}(X_i)Y_i \leq 4\gamma].$$

This proves the following theorem.

THEOREM 3.19. *With \mathbb{P} probability at least $1 - \epsilon$, for any $(w, b) \in \Theta$,*

$$\begin{aligned} & \frac{1}{kN} \sum_{i=N+1}^{(k+1)N} \mathbf{1}\left\{2\mathbf{1}[g_{w,b}(X_i) \geq 0] - 1 \neq Y_i\right\} \\ & \leq \frac{k+1}{k} \inf_{\lambda \in \mathbb{R}_+, h \in \mathbb{N}^*} [1 - \exp(-\frac{\lambda}{N})]^{-1} \left\{ 1 - \right. \\ & \quad \exp\left[-\frac{\lambda}{N^2} \sum_{i=1}^N \mathbf{1}[g_{w,b}(X_i)Y_i \leq 4R\gamma h] \right. \\ & \quad \left. \left. - \frac{\log[20(k+1)N] \left\{\frac{h}{\log(2)} \log\left(\frac{4e(k+1)N}{h}\right) + 1\right\} + \log\left[\frac{2h(h+1)}{\epsilon}\right]}{N}\right]\right\} \end{aligned}$$

$$- \frac{1}{kN} \sum_{i=1}^N \mathbb{1}[g_{w,b}(X_i)Y_i \leq 4R\gamma_h].$$

As a consequence, we obtain with \mathbb{P} probability at least $1 - \epsilon$, for any $(w, b) \in \Theta$ such that

$$\gamma = \min_{i=1, \dots, N} g_{w,b}(X_i)Y_i > 0,$$

$$\begin{aligned} & \frac{1}{kN} \sum_{i=N+1}^{(k+1)N} \mathbb{1}[g_{w,b}(X_i)Y_i < 0] \\ & \leq \frac{k+1}{k} \left\{ 1 - \exp \left[- \frac{\log[20(k+1)N]}{N} \left\{ \frac{16R^2 + 2\gamma^2}{\log(2)\gamma^2} \log \left(\frac{\epsilon(k+1)N\gamma^2}{4R^2} \right) + 1 \right\} \right. \right. \\ & \qquad \qquad \qquad \left. \left. + \frac{1}{N} \log \left(\frac{\epsilon}{2} \right) \right] \right\}. \end{aligned}$$

This inequality compares favourably with similar inequalities in [18], which moreover do not extend to the margin quantile case as this one.

Let us also remark that it is easy to circumvent the fact that R is not observed when the test set $X_{N+1}^{(k+1)N}$ is not observed.

Indeed, we can consider the sample obtained by projecting $X_1^{(k+1)N}$ on some ball of fixed radius R_{\max} , putting

$$t_{R_{\max}}(X_i) = \min \left\{ 1, \frac{R_{\max}}{\|X_i\|} \right\} X_i.$$

We can further consider an atomic prior distribution $\nu \in \mathcal{M}_+^1(\mathbb{R}_+)$ bearing on R_{\max} , to obtain a uniform result through a union bound. As a consequence of the previous theorem indeed,

COROLLARY 3.20. *For any atomic prior $\nu \in \mathcal{M}_+^1(\mathbb{R}_+)$, for any partially exchangeable probability measure $\mathbb{P} \in \mathcal{M}_+^1(\Omega)$, with \mathbb{P} probability at least $1 - \epsilon$, for any $(w, b) \in \Theta$, any $R_{\max} \in \mathbb{R}_+$,*

$$\begin{aligned}
& \frac{1}{kN} \sum_{i=N+1}^{(k+1)N} \mathbb{1} \left\{ 2\mathbb{1} [g_{w,b} \circ t_{R_{\max}}(X_i) \geq 0] - 1 \neq Y_i \right\} \\
& \leq \frac{k+1}{k} \inf_{\lambda \in \mathbb{R}_+, h \in \mathbb{N}^*} [1 - \exp(-\frac{\lambda}{N})]^{-1} \left\{ 1 - \right. \\
& \quad \exp \left[-\frac{\lambda}{N^2} \sum_{i=1}^N \mathbb{1} [g_{w,b} \circ t_{R_{\max}}(X_i) Y_i \leq 4R_{\max} \gamma h] \right. \\
& \quad \left. \left. - \frac{\log [20(k+1)N] \left\{ \frac{h}{\log(2)} \log \left(\frac{4e(k+1)N}{h} \right) + 1 \right\} + \log \left[\frac{2h(h+1)}{\epsilon \nu(R_{\max})} \right]}{N} \right] \right\} \\
& \quad \left. - \frac{1}{kN} \sum_{i=1}^N \mathbb{1} [g_{w,b} \circ t_{R_{\max}}(X_i) Y_i \leq 4R_{\max} \gamma h] \right\}.
\end{aligned}$$

4. APPENDIX: CLASSIFICATION BY THRESHOLDING

In this appendix, we show how the bounds given in the first section of this monograph can be computed in practice on a simple example: the case when the classification is performed by comparing a series of measurements to threshold values. Let us mention that our description covers the case when the same measurement is compared to several thresholds, since it is enough to repeat a measurement in the list of measurements describing a pattern to cover this case.

4.1. DESCRIPTION OF THE MODEL. Let us assume that the patterns we want to classify are described through h real valued measurements normalized in the range $(0, 1)$. In this setting the pattern space can thus be defined as $\mathcal{X} = (0, 1)^h$.

Consider the threshold set $\mathcal{T} = (0, 1)^h$ and the response set $\mathcal{R} = \mathcal{Y}^{\{0,1\}^h}$. For any $t \in (0, 1)^h$ and any $a : \{0, 1\}^h \rightarrow \mathcal{Y}$, let

$$f_{(t,a)}(x) = a \left\{ \left[\mathbb{1}(x^j \geq t_j) \right]_{j=1}^h \right\}, \quad x \in \mathcal{X},$$

where x^j is the j th coordinate of $x \in \mathcal{X}$. Thus our parameter set here is $\Theta = \mathcal{T} \times \mathcal{R}$. Let us consider on \mathcal{T} the Lebesgue measure L and on \mathcal{R} the uniform probability distribution U . Let our prior distribution be $\pi = L \otimes U$. Let us define for any threshold sequence $t \in \mathcal{T}$

$$\Delta_t = \left\{ t' \in \mathcal{T} : \overline{(t'_j, t_j)} \cap \{X_i^j; i = 1, \dots, N\} = \emptyset, j = 1, \dots, h \right\},$$

where X_i^j is the j th coordinate of the sample pattern X_i , and where the interval $\overline{(t'_j, t_j)}$ of the real line is defined as the convex hull of the two point set $\{t'_j, t_j\}$, whether $t'_j \leq t_j$ or not. We see that Δ_t is the set of thresholds giving the same response as t on the training patterns. Let us consider for any $t \in \mathcal{T}$ the middle

$$m(\Delta_t) = \frac{\int_{\Delta_t} t' L(dt')}{L(\Delta_t)}$$

of Δ_t . The set Δ_t being a product of intervals, its middle is the point whose coordinates are the middle of these intervals. Let us introduce the finite set T composed of the middles of the cells Δ_t , which can be defined as

$$T = \{t \in \mathcal{T} : t = m(\Delta_t)\}.$$

It is easy to see that $|T| \leq (N+1)^h$ and that $|\mathcal{R}| = |\mathcal{Y}|^{2^h}$.

4.2. COMPUTATION OF INDUCTIVE BOUNDS. For any parameter $(t, a) \in \mathcal{T} \times \mathcal{R} = \Theta$, let us consider the posterior distribution defined by its density

$$\frac{d\rho_{(t,a)}}{d\pi}(t', a') = \frac{\mathbf{1}(t' \in \Delta_t) \mathbf{1}(a' = a)}{\pi(\Delta_t \times \{a\})}.$$

Let us notice that we are in fact considering a finite number of posterior distributions, since $\rho_{(t,a)} = \rho_{(m(\Delta_t), a)}$, where $m(\Delta_t) \in T$. Let us also mention that for any exchangeable sample distribution $\mathbb{P} \in \mathcal{M}_+^1[(\mathcal{X} \times \mathcal{Y})^{N+1}]$ and any thresholds $t \in \mathcal{T}$,

$$\mathbb{P} \left[\overline{(X_{N+1}^j, t_j)} \cap \{X_i^j, i = 1, \dots, N\} = \emptyset \right] \leq \frac{2}{N+1}.$$

Thus, for any $(t, a) \in \Theta$,

$$\mathbb{P} \left\{ \rho_{(t,a)}[f.(X_{N+1})] \neq f_{(t,a)}(X_{N+1}) \right\} \leq \frac{2h}{N+1},$$

showing that the classification produced by $\rho_{(t,a)}$ on new examples is most of the time non random (this result is only indicative, since it is concerned with a non random choice of (t, a)).

Let us then compute the various quantities needed to apply the results of the first section, focussing our attention of Theorem 1.39 (page 63):

It is to be noted first of all that $\rho_{(t,a)}(r) = r[(t, a)]$. The entropy term is such that

$$\mathcal{K}(\rho_{t,a}, \pi) = -\log[\pi(\Delta_t \times \{r\})] = -\log[L(\Delta_t)] + 2^h \log(|\mathcal{Y}|).$$

Let us notice accordingly that

$$\min_{(t,a) \in \Theta} \mathcal{K}(\rho_{(t,a)}, \pi) \leq h \log(N+1) + 2^h \log(|\mathcal{Y}|).$$

Let us introduce the counters

$$b_y^t(c) = \frac{1}{N} \sum_{i=1}^N \mathbf{1} \left\{ Y_i = y \text{ and } [\mathbf{1}(X_i^j \geq t_j)]_{j=1}^h = c \right\},$$

$$t \in T, c \in \{0, 1\}^h, y \in \mathcal{Y},$$

$$b^t(c) = \sum_{y \in \mathcal{Y}} b_y^t(c) = \frac{1}{N} \sum_{i=1}^N \mathbf{1} \left\{ [\mathbf{1}(X_i^j \geq t_j)]_{j=1}^h = c \right\}, \quad t \in T, c \in \{0, 1\}^h.$$

Since

$$r[(t, a)] = \sum_{c \in \{0, 1\}^h} [b^t(c) - b_{a(c)}^t(c)],$$

the partition function of the Gibbs estimator can be computed as

$$\begin{aligned} \pi[\exp(-\lambda r)] &= \sum_{t \in T} L(\Delta_t) \sum_{a \in \mathcal{R}} \frac{1}{|\mathcal{Y}|^{2h}} \exp \left[-\lambda \sum_{i=1}^N \mathbf{1}[Y_i \neq f_{(t,a)}(X_i)] \right] \\ &= \sum_{t \in T} L(\Delta_t) \sum_{a \in \mathcal{R}} \frac{1}{|\mathcal{Y}|^{2h}} \exp \left[-\lambda \sum_{c \in \{0, 1\}^h} [b^t(c) - b_{a(c)}^t(c)] \right] \\ &= \sum_{t \in T} L(\Delta_t) \prod_{c \in \{0, 1\}^h} \left[\frac{1}{|\mathcal{Y}|} \sum_{y \in \mathcal{Y}} \exp \left(-\lambda [b^t(c) - b_y^t(c)] \right) \right]. \end{aligned}$$

We see that the number of operations needed to compute $\pi[\exp(-\lambda r)]$ is proportional to $|T| \times 2^h \times |\mathcal{Y}| \leq (N+1)^h 2^h |\mathcal{Y}|$. An exact computation will therefore be feasible only for small values of N and h . For higher values, a Monte Carlo approximation of this sum will have to be performed instead.

If we want to compute the bound provided by Theorem 1.39 (page 63), we need also to compute, for any fixed parameter $\theta \in \Theta$, quantities of the type

$$\pi_{\exp(-\lambda r)} \left\{ \exp[\xi m'(\cdot, \theta)] \right\} = \pi_{\exp(-\lambda r)} \left\{ \exp[\xi \rho_\theta(m')] \right\}, \quad \lambda, \xi \in \mathbb{R}_+.$$

To this purpose we need to introduce

$$\bar{b}_y^t(\theta, c) = \frac{1}{N} \sum_{i=1}^N \left| \mathbf{1}[f_\theta(X_i) \neq Y_i] - \mathbf{1}(y \neq Y_i) \right| \mathbf{1} \left\{ [\mathbf{1}(X_i^j \geq t_j)]_{j=1}^h = c \right\}.$$

Similarly to what has been done previously, we obtain

$$\begin{aligned} & \pi \left\{ \exp[-\lambda r + \xi m'(\cdot, \theta)] \right\} \\ &= \sum_{t \in T} L(\Delta_t) \prod_{c \in \{0,1\}^h} \left[\frac{1}{|Y|} \sum_{y \in Y} \exp \left(-\lambda [b^t(c) - b_y^t(c)] + \xi \bar{b}_y^t(\theta, c) \right) \right]. \end{aligned}$$

We can then compute

$$\begin{aligned} \pi_{\exp(-\lambda r)}(r) &= -\frac{\partial}{\partial \lambda} \log \left\{ \pi [\exp(-\lambda r)] \right\}, \\ \pi_{\exp(-\lambda r)} \left\{ \exp[\xi \rho_\theta(m')] \right\} &= \frac{\pi \left\{ \exp[-\lambda r + \xi m'(\cdot, \theta)] \right\}}{\pi [\exp(-\lambda r)]}, \\ \pi_{\exp(-\lambda r)} [m'(\cdot, \theta)] &= \frac{\partial}{\partial \xi} \Big|_{\xi=0} \log \left[\pi \left\{ \exp[-\lambda r + \xi m'(\cdot, \theta)] \right\} \right]. \end{aligned}$$

This is all we need to compute $B(\rho_\theta, \beta, \gamma)$ (and also $B(\pi_{\exp(-\lambda r)}, \beta, \gamma)$) in Theorem 1.39 (page 63), using the approximation

$$\begin{aligned} & \log \left\{ \pi_{\exp(-\lambda_1 r)} \left[\exp \left\{ \xi \pi_{\exp(-\lambda_2 r)}(m') \right\} \right] \right\} \\ & \leq \log \left\{ \pi_{\exp(-\lambda_1 r)} \left[\exp \left\{ \xi m'(\cdot, \theta) \right\} \right] \right\} + \xi \pi_{\exp(-\lambda_2 r)} [m'(\cdot, \theta)], \quad \xi \geq 0. \end{aligned}$$

Let us also explain how to apply the posterior distribution $\rho_{(t,a)}$, in other words our randomized estimated classification rule, to a new pattern X_{N+1} :

$$\begin{aligned} \rho_{(t,a)} [f.(X_{N+1}) = y] &= L(\Delta_t)^{-1} \int_{\Delta_t} \mathbf{1} \left[a \left\{ [\mathbf{1}(X_{N+1}^j \geq t'_j)]_{j=1}^h \right\} = y \right] L(dt') \\ &= L(\Delta_t)^{-1} \sum_{c \in \{0,1\}^h} L \left(\left\{ t' \in \Delta_t : [\mathbf{1}(X_{N+1}^j \geq t'_j)]_{j=1}^h = c \right\} \right) \mathbf{1} [a(c) = y]. \end{aligned}$$

Let us define for short

$$\Delta_t(c) = \left\{ t' \in \Delta_t : [\mathbf{1}(X_{N+1}^j \geq t'_j)]_{j=1}^h = c \right\}, \quad c \in \{0,1\}^h.$$

With this notation

$$\rho_{(t,a)} [f.(X_{N+1}) = y] = L(\Delta_t)^{-1} \sum_{c \in \{0,1\}^h} L[\Delta_t(c)] \mathbf{1} [a(c) = y].$$

We can compute in the same way the probabilities for the label of the new pattern under the Gibbs posterior distribution:

$$\begin{aligned}
& \pi_{\exp(-\lambda r)}[f.(X_{N+1}) = y'] \\
&= \left\{ \sum_{t \in T} \prod_{c \in \{0,1\}^h} \left[\frac{1}{|\mathcal{Y}|} \sum_{y \in \mathcal{Y}} \exp\left(-\lambda[b^t(c) - b_y^t(c)]\right) \right] \right. \\
&\quad \times \left. \sum_{c \in \{0,1\}^h} L[\Delta_t(c)] \frac{\sum_{y \in \mathcal{Y}} \mathbf{1}(y = y') \exp\{-\lambda[b^t(c) - b_y^t(c)]\}}{\sum_{y \in \mathcal{Y}} \exp\{-\lambda[b^t(x) - b_y^t(c)]\}} \right\} \\
&\quad \times \left\{ \sum_{t \in T} L(\Delta_t) \prod_{c \in \{0,1\}^h} \left[\frac{1}{|\mathcal{Y}|} \sum_{y \in \mathcal{Y}} \exp\left(-\lambda[b^t(c) - b_y^t(c)]\right) \right] \right\}^{-1}.
\end{aligned}$$

4.3. TRANSDUCTIVE BOUNDS. In the case when we observe the patterns of a shadow sample $(X_i)_{i=N+1}^{(k+1)N}$ on top of the training sample $(X_i, Y_i)_{i=1}^N$, we can introduce the set of thresholds responding as t on the extended sample $(X_i)_{i=1}^{(k+1)N}$

$$\bar{\Delta}_t = \left\{ t' \in \mathcal{T} : \overline{(t'_j, t_j)} \cap \{X_i^j; i = 1, \dots, (k+1)N\} = \emptyset, j = 1, \dots, h \right\},$$

consider the set

$$\bar{T} = \{t \in \mathcal{T} : t = m(\bar{\Delta}_t)\},$$

of the middle points of the cells $\bar{\Delta}_t$, $t \in \mathcal{T}$, and replace the Lebesgue measure $L \in \mathcal{M}_{\perp}^1[(0,1)^h]$ of the previous section with the uniform probability measure \bar{L} on \bar{T} . We can then consider $\pi = \bar{L} \otimes U$, where U is as previously the uniform probability measure on \mathcal{R} . This gives obviously an exchangeable posterior distribution and therefore qualifies π for transductive bounds. Let us notice that $|\bar{T}| \leq [(k+1)N+1]^h$, and therefore that $\pi(t, a) \geq [(k+1)N+1]^{-h} |\mathcal{Y}|^{-2^h}$, for any $(t, a) \in \bar{T} \times \mathcal{R}$.

For any $(t, a) \in \mathcal{T} \times \mathcal{R}$ we may similarly to the inductive case consider the posterior distribution $\rho_{(t,a)}$ defined by

$$\frac{d\rho_{(t,a)}}{d\pi}(t', a') = \frac{\mathbf{1}(t' \in \Delta_t) \mathbf{1}(a' = a)}{\pi(\Delta_t \times \{a\})},$$

but we may also consider $\delta_{(m(\bar{\Delta}_t), a)}$, which is such that $r_i\{[m(\bar{\Delta}_t), a]\} = r_i[(t, a)]$, $i = 1, 2$, whereas only $\rho_{(t,a)}(r_1) = r_1[(t, a)]$, while

$$\rho_{(t,a)}(r_2) = \frac{1}{|\bar{T} \cap \Delta_t|} \sum_{t' \in \bar{T} \cap \Delta_t} r_2[(t', a)].$$

We get

$$\begin{aligned} \mathcal{K}(\rho_{(t,a)}, \pi) &= -\log[\bar{L}(\Delta_t)] + 2^h \log(|\mathcal{Y}|) \\ &\leq \log(|\bar{T}|) + 2^h \log(|\mathcal{Y}|) = \mathcal{K}(\delta_{[m(\bar{\Delta}_t), a]}, \pi) \\ &\leq h \log[(k+1)N+1] + 2^h \log(|\mathcal{Y}|), \end{aligned}$$

whereas we had no such uniform bound in the inductive case. Similarly to the inductive case

$$\pi[\exp(-\lambda r_1)] = \sum_{t \in T} \bar{L}(\Delta_t) \prod_{c \in \{0,1\}^h} \left[\frac{1}{|\mathcal{Y}|} \sum_{y \in \mathcal{Y}} \exp\left(-\lambda[b^t(c) - b_y^t(c)]\right) \right].$$

Moreover, for any $\theta \in \Theta$,

$$\begin{aligned} \pi\{\exp[-\lambda r_1 + \xi \rho_\theta(m')]\} &= \pi\{\exp[-\lambda r_1 + \xi m'(\cdot, \theta)]\} \\ &= \sum_{t \in T} \bar{L}(\Delta_t) \prod_{c \in \{0,1\}^h} \left[\frac{1}{|\mathcal{Y}|} \sum_{y \in \mathcal{Y}} \exp\left(-\lambda[b^t(c) - b_y^t(c)] + \xi \bar{b}(\theta, c)\right) \right]. \end{aligned}$$

The bound for the transductive counter part to Theorem 1.39 (page 63), obtained as explained page 102, can be computed as in the inductive case, from these two partitions functions and the above entropy estimates.

Let us mention eventually that, using the same notations as in the inductive case,

$$\begin{aligned} &\pi_{\exp(-\lambda r_1)}[f.(X_{N+1}) = y'] \\ &= \left\{ \sum_{t \in T} \prod_{c \in \{0,1\}^h} \left[\frac{1}{|\mathcal{Y}|} \sum_{y \in \mathcal{Y}} \exp\left(-\lambda[b^t(c) - b_y^t(c)]\right) \right] \right\} \\ &\quad \times \sum_{c \in \{0,1\}^h} \bar{L}[\Delta_t(c)] \frac{\sum_{y \in \mathcal{Y}} \mathbb{1}(y = y') \exp\{-\lambda[b^t(c) - b_y^t(c)]\}}{\sum_{y \in \mathcal{Y}} \exp\{-\lambda[b^t(x) - b_y^t(c)]\}} \left\{ \sum_{t \in T} \bar{L}(\Delta_t) \prod_{c \in \{0,1\}^h} \left[\frac{1}{|\mathcal{Y}|} \sum_{y \in \mathcal{Y}} \exp\left(-\lambda[b^t(c) - b_y^t(c)]\right) \right] \right\}^{-1}. \end{aligned}$$

REFERENCES

- [1] N. Alon, S. Ben-David, N. Cesa-Bianchi, D. Haussler, Scale sensitive dimensions, uniform convergence and learnability, *J. of ACM* **44**(4):615-631, 1997.
- [2] J. Y. Audibert, Aggregated estimators and empirical complexity for least square regression, *Ann. Inst. H. Poincaré Probab. Statist.*, **40**, 6, pp. 685-736, 2004.
- [3] J. Y. Audibert, *PAC-Bayesian Statistical Learning Theory*, PhD, University Paris 6, June 2004.
- [4] A. Barron (1987) Are Bayes Rules Consistent in Information ? *Open Problems in Communication and Computation*, T. M. Cover and B. Gopinath Ed., Springer Verlag 1987.
- [5] A. Barron and Y. Yang (1995) Information Theoretic Determination of Minimax Rates of Convergence, *Ann. Statist.* **27** (1999), no. 5, 1564–1599.
- [6] A. Barron, L. Birgé and P. Massart, (1995) Risk bounds for model selection via penalization, *Probab. Theory Related Fields*, **113** (1999), no. 3, 301–413.
- [7] G. Blanchard, The “progressive mixture” estimator for regression trees, *Annales de l’I.H.P.*, **35**(6):793-820, 1999.
- [8] G. Blanchard, A new algorithm for Bayesian MCMC CART sampling, *preprint*, 2000.
- [9] G. Blanchard, *Mixture and aggregation of estimators for pattern recognition. Application to decision trees [Méthodes de mélange et d’agrégation d’estimateurs en reconnaissance de formes. Application aux arbres de décision.]* in English with an introduction in French, PhD dissertation, Université Paris XIII, January 2001.
- [10] L. Birgé and P. Massart (1997) From model selection to adaptive estimation, *Festschrift for Lucien Le Cam*, 55–87, Springer, New York.
- [11] L. Birgé and P. Massart (1995) Minimum contrast estimators on sieves, *Bernoulli* **4** (1998), no. 3, 329–375.
- [12] L. Birgé and P. Massart, A generalized C_p criterion for Gaussian model selection, *preprint*, 2001.
- [13] L. Birgé and P. Massart, Gaussian model selection, *J. Eur. Math. Soc.*, 2001.
<http://www.springer.de/link>
- [14] A. Blum and J. Langford, PAC-MDL Bounds, *Computational Learning Theory and Kernel Machines, 16th Annual Conference on Com-*

- putational Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003, Washington, DC, USA, August 24-27, 2003, Proceedings*, Lecture Notes in Computer Science 2777, Springer 2003, pp. 344-357.
- [15] O. Catoni, Laplace transform estimates and deviation inequalities *Ann. Inst. H. Poincaré Probab. Statist.* **39** (2003), no. 1, 1–26.
- [16] O. Catoni, Data compression and adaptive histograms. In F. Cucker and J.M. Rojas, editors, *Foundations of Computational Mathematics, Proceedings of the Smalefest 2000*, pages 35–60. World Scientific, 2002.
- [17] O. Catoni, Statistical learning theory and stochastic optimization, *Ecole d’Eté de Probabilités de Saint-Flour XXXI - 2001*, Springer Lecture Notes in Mathematics, **1851**, pp 1 - 270, 2004.
- [18] N. Cristianini and J. Shawe Taylor, *An introduction to Support Vector Machines and other kernel based learning methods*, Cambridge University Press, 2000.
- [19] M. Feder and N. Merhav, Hierarchical Universal Coding, *IEEE Trans. Inform. Theory*, vol 42, no 5, Sept, 1996.
- [20] T. Hastie, R. Tibshirani and J. Friedman, *The elements of statistical learning*, Springer, New York, 2001.
- [21] J. Langford and D. McAllester, Computable Shell Decomposition Bounds, *Journal of Machine Learning Research*, **5**, 2004, pp. 529-547 (communicated at COLT 2000).
- [22] J. Langford and M. Seeger, Bounds for Averaging Classifiers, *technical report CMU-CS-01-102*, Carnegie Mellon University, jan. 2001, www.cs.cmu.edu/~jcl.
- [23] J. Langford, M. Seeger and N. Megiddo, An Improved Predictive Accuracy Bound for Averaging Classifiers, *International Conference on Machine Learning* **18** (2001), 290-297.
- [24] N. Littlestone and M. Warmuth, Relating data compression and learnability. *Technical report*, University of California, Santa Cruz, 1986.
- [25] D. A. McAllester, Some PAC-Bayesian Theorems, *Proceedings of the Eleventh Annual Conference on Computational Learning Theory (Madison, WI, 1998)*, 230–234 (electronic), ACM, New York, 1998;
- [26] D. A. McAllester, PAC-Bayesian Model Averaging, *Proceedings of the Twelfth Annual Conference on Computational Learning Theory (Santa Cruz, CA, 1999)*, 164–170 (electronic), ACM, New York, 1999;
- [27] C. McDiarmid, Concentration *Probabilistic Methods for Algorithmic Discrete Mathematics*, Habib M., McDiarmid C. and Reed B. Eds.,

- Springer, 1998.
- [28] E. Mammen and A. Tsybakov, Smooth discrimination analysis. *Ann. Statist.* **27** (1999), 1808-1829.
 - [29] B. Y. Ryabko, Twice-universal coding, *Probl. Inform. Transm.*, vol 20, no 3, pp. 24-28, July-Sept 1984.
 - [30] M. Seeger, PAC-Bayesian Generalization Error Bounds for Gaussian Process Classification, *Journal of Machine Learning Research* **3** (2002), 233–269.
 - [31] M. Seeger, PAC-Bayesian Generalization Error Bounds for Gaussian Process Classification, *Informatics Research Report, Division of Informatics, University of Edinburgh EDI-INF-RR-0094* (2002), 1–42. <http://www.informatics.ed.ac.uk>
 - [32] J. Shawe-Taylor, P. L. Bartlett, R. C. Williamson, M. Anthony, Structural risk minimization over data-dependent hierarchies, *IEEE Trans. Inform. Theory* **44** (1998), no. 5, 1926–1940.
 - [33] J. Shawe-Taylor and N. Cristianini, On the generalization of soft margin algorithms, *IEEE Trans. Inform. Theory* **48** (2002), no. 10, 2721–2735.
 - [34] A. Tsybakov, Optimal aggregation of classifiers in statistical learning, *Ann. Statist.* **32**(2004), no. 1, 135-166.
 - [35] A. Tsybakov, S. Van de Geer, Square root penalty: adaptation to the margin in classification and in edge estimation, *Ann. Statist.*, **33**, pp 1203-1224.
 - [36] S. Van de Geer, *Applications of empirical process theory*. Cambridge Series in Statistical and Probabilistic Mathematics, 6. Cambridge University Press, Cambridge, 2000. xii+286 pp.
 - [37] V. N. Vapnik, *Statistical learning theory*, Wiley, New York, 1998.
 - [38] J.-P. Vert, Double mixture and universal inference, *preprint*, 2000, <http://cg.ensmp.fr/~vert/publi/> .
 - [39] J.-P. Vert, Adaptive context trees and text clustering, *IEEE Trans. Inform. Theory*, **47** (2001), no. 5, 1884-1901.
 - [40] J.-P. Vert, Text categorization using adaptive context trees, *Proceedings of the CICLing-2001 conference*, A. Gelbukh (Ed.), LNCS 2004, Springer-Verlag, pp. 423-436, 2001.
 - [41] F. M. J. Willems, Y. M. Shtarkov and T. J. Tjalkens, The Context-Tree Weighting Method: Basic Properties, *IEEE Trans. Inform. Theory*, vol 41, no 3, May, 1995.
 - [42] F. M. J. Willems, Y. M. Shtarkov and T. J. Tjalkens, Context Weighting for General Finite-Context Sources, *IEEE Trans. Inform. Theory*, vol 42, no 5, Sept, 1996.

-
- [43] T. Zhang, From epsilon-entropy to KL-entropy: analysis of minimum information complexity density estimation, *Ann. Statist.* (2006), to appear.
 - [44] T. Zhang, Information Theoretical Upper and Lower Bounds for Statistical Estimation, *IEEE. Trans. Inform. Theory*, (2006), to appear.