



Universiteit
Leiden
The Netherlands

Topics in mathematical and applied statistics

Pas, S.L. van der

Citation

Pas, S. L. van der. (2017, February 28). *Topics in mathematical and applied statistics*. Retrieved from <https://hdl.handle.net/1887/46454>

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/46454>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/46454> holds various files of this Leiden University dissertation

Author: Pas, S.L. van der

Title: Topics in mathematical and applied statistics

Issue Date: 2017-02-28

5

The switch criterion in nested model selection

Abstract

We study the switch distribution, introduced by Van Erven et al. (2012), applied to model selection and subsequent estimation. While switching was known to be strongly consistent, here we show that it achieves minimax optimal parametric risk rates up to a $\log \log n$ factor when comparing two nested exponential families, partially confirming a conjecture by Lauritzen (2012) and Cavanaugh (2012) that switching behaves asymptotically like the Hannan-Quinn criterion. Moreover, like Bayes factor model selection but unlike standard significance testing, when one of the models represents a simple hypothesis, the switch criterion defines a robust null hypothesis test, meaning that its Type-I error probability can be bounded irrespective of the stopping rule. Hence, switching is consistent, insensitive to optional stopping and almost minimax risk optimal, showing that, Yang's (2005) impossibility result notwithstanding, it is possible to 'almost' combine the strengths of AIC and Bayes factor model selection.

5.1 Introduction

We consider the following standard model selection problem, where we have i.i.d. observations X_1, \dots, X_n and we wish to select between two nested parametric models,

$$\mathcal{M}_0 = \{p_\mu \mid \mu \in M_0\} \quad \text{and} \quad \mathcal{M}_1 = \{p_\mu \mid \mu \in M_1\}. \quad (5.1)$$

This chapter is set to appear in *Statistica Sinica*, as: S. van der Pas and P. Grünwald. Almost the best of three worlds: risk, consistency and optional stopping for the switch criterion in nested model selection. The central result of this paper, Theorem 5.5, already appeared in the Master's Thesis (Van der Pas, 2013) for the special case where $m_1 = 1$ and $m_0 = 0$. This research was supported by NWO VICI Project 639.073.04.

Here the X_i are random vectors taking values in some set \mathcal{X} , $M_1 \subseteq \mathbb{R}^{m_1}$ for some $m_1 > 0$ and $\mathcal{M}_0 = \{p_\mu : \mu \in M_0\} \subset \mathcal{M}_1$ represents an m_0 -dimensional submodel of \mathcal{M}_1 , where $0 \leq m_0 < m_1$. We may thus denote \mathcal{M}_0 as the ‘simple’ and \mathcal{M}_1 as the ‘complex’ model. We will assume that \mathcal{M}_1 is an exponential family, represented as a set of densities on \mathcal{X} with respect to some fixed underlying measure, so that p_μ represents the density of the observations, and we take it to be given in its mean-value parameterization. As the notation indicates, we require, without loss of generality, that the parameterizations of \mathcal{M}_0 and \mathcal{M}_1 coincide, that is $M_0 \subset M_1$ is itself a set of m_1 -dimensional vectors, the final $m_1 - m_0$ components of which are fixed to known values. We restrict ourselves to the case in which both \mathcal{M}_1 and the restriction of \mathcal{M}_0 to its first m_0 components are products of open intervals. Most model selection methods output not just a decision $\delta(X^n) \in \{0, 1\}$, but also an indication $r(X^n) \in \mathbb{R}$ of the strength of evidence, such as a p -value or a Bayes factor. As a result, such procedures can often be interpreted as methods for hypothesis testing, where \mathcal{M}_0 represents the *null* model and \mathcal{M}_1 the alternative; a very simple example of our setting is when the X_i consist of two components $X_i \equiv (X_{i1}, X_{i2})$, which according to \mathcal{M}_1 are independent Gaussians whereas under \mathcal{M}_2 they can have an arbitrary bivariate Gaussian distribution and hence can be dependent. Since we allow \mathcal{M}_0 to be a singleton, this setting also includes some very simple, classical yet important settings such as testing whether a coin is biased (\mathcal{M}_0 is the fair coin model, \mathcal{M}_1 contains all Bernoulli distributions).

We consider three desirable properties of model selection methods: (a) optimal worst-case risk rate of post-model selection estimation (with risk measured in terms of squared error loss, squared Hellinger distance, Rényi or Kullback-Leibler divergence); (b) consistency, and, (c) for procedures which also output a strength of evidence $r(X^n)$, whether the validity of the evidence is insensitive to optional stopping under the null model. We evaluate the recently introduced model selection criterion δ_{sw} based on the switch distribution (Van Erven et al., 2012) on properties (a), (b) and (c).

The switch distribution, introduced by Van Erven et al., (2007), was originally designed to address the *catch-up phenomenon*, which occurs when the best predicting model is not the same across sample sizes. The switch distribution can be interpreted as a modification of the Bayesian predictive distribution. It also has an MDL interpretation: if one corrects standard MDL approaches (Grünwald, 2007) to take into account that the best predicting method changes over time, one naturally arrives at the switch distribution. Lhéritier and Cazals (2015) describes a successful practical application for two-sample sequential testing, related to the developments in this paper but in a nonparametric context. We briefly give the definitions relevant to our setting in Section 5.2; for all further details we refer to Van Erven et al. (2012) and 5.7.5 in Section 5.7.

When evaluating any model selection method, there is a well-known tension between (a) and (b): the popular AIC method (Akaike, 1973) achieves the minimax optimal parametric rate of order $1/n$ in the problem above, but is inconsistent; the same holds for the many popular model selection methods that asymptotically tend to behave like AIC, such as k -fold and leave-one-out-cross-validation, the bootstrap and Mallows’s C_p in linear regression (Efron, 1986; Shao, 1997; Stone, 1977). On the other hand, BIC (Schwarz, 1978) is consistent in the sense that for large enough n , it will select the smallest model containing the ‘true’ μ ; but it misses the minimax parametric rate by a factor of $\log n$. The same holds for traditional Minimum Description Length (MDL) approaches (Grünwald, 2007)

and Bayes factor model selection (BFMS) (Kass and Raftery, 1995), of which BIC is an approximation. This might lead one to wonder if there exists a single method that is optimal in both respects. A key result by Yang (2005) shows that this is impossible: any consistent method misses the minimax optimal rate by a factor $g(n)$ with $\lim_{n \rightarrow \infty} g(n) = \infty$.

In Section 5.4.2 we show that, Yang's result notwithstanding, the switch distribution allows us to get very close to satisfying property (a) and (b) at the same time, at least in the problem defined above (Yang's result was shown in a nested linear regression rather than our exponential family context, but it does hold in our exponential family setting as well; see the discussion at the end of Section 5.3.3). We prove that in our setting, the switch model selection criterion δ_{sw} (a) misses the minimax optimal rate only by an exceedingly small $g_{\text{sw}}(n) \asymp \log \log n$ factor (Theorem 5.5). Property (b), strong consistency, was already shown by Van Erven et al. (2012). The factor $g_{\text{sw}}(n) \asymp \log \log n$ is an improvement over the extra factor resulting from Bayes factor model selection, which has $g_{\text{BFMS}}(n) \asymp \log n$. Indeed, as discussed in the introduction of Van Erven et al. (2012), the catch-up phenomenon that the switch distribution addresses is intimately related to the rate-suboptimality of Bayesian inference. Van Erven et al. (2012) show that, while model selection based on switching is consistent, sequential prediction based on model averaging with the switching method achieves minimax optimal *cumulative* risk rates in general parametric and nonparametric settings, where the cumulative risk at sample size n is obtained by summing the standard, instantaneous risk from 1 to n . In contrast, in nonparametric settings, standard Bayesian model averaging typically has a cumulative risk rate that is larger by a $\log n$ factor. Using the cumulative risk is natural in sequential prediction settings, but Van Erven et al. (2012) left open the question of how switching would behave for the more standard, instantaneous risk. In contrast to the cumulative setting, we cannot expect to achieve the optimal rate here by Yang's (2005) result, but it is interesting to see that switching gets so close.

We now turn to the third property, robustness to optional stopping. While consistency in the sense above is an asymptotic and even somewhat controversial notion (see Section 5.6), there exists a nonasymptotic property closely related to consistency that, while arguably much more important in practice, has received relatively little attention in the recent statistical literature. This is property (c) above, insensitivity to optional stopping. In statistics, the issue was thoroughly discussed, yet never completely resolved, in the 1960s; nowadays, it is viewed as a highly desirable feature of testing methods by, for example, psychologists; see (Sanborn and Hills, 2014; Wagenmakers, 2007). In particular, it is often argued (Wagenmakers, 2007) that the fixed stopping rule required by the classical Neyman-Pearson paradigm severely and unnecessarily restricts the application domain of hypothesis testing, invalidating much of the p -values reported in the psychological literature. Approximately 55% of psychologists admitted in a survey to deciding whether to collect more data after looking at their results to see if they were significant (John et al., 2012). We analyze property (c) in terms of *robust null hypothesis tests*, formally defined in Section 5.5. A method defines a robust null hypothesis test if (1) it outputs evidence $r(X^n)$ that does not depend on the stopping rule used to determine n , and (2) (some function of) $r(X^n)$ gives a bound on the Type-I error that is valid no matter what this stopping rule is. Standard (Neyman-Pearson) null hypothesis testing and tests derived from AIC-type methods are not robust in this sense. For example, such tests cannot be used if the stop-

ping rule is simply unknown, as is often the case when analyzing externally provided data — but this is just the tip of an iceberg of problems with nonrobust tests. For an exhaustive review of such problems we refer to Wagenmakers (2007) who builds on, amongst others, Berger and Wolpert (1988) and Pratt (1962).

Now, as first noted by Edwards et al. (1963), in simple versus composite testing (i.e. when \mathcal{M}_0 is a singleton), the output of BFMS, the Bayes factor, does provide a robust null hypothesis test. This is one of the main reasons why for example, in psychology, Bayesian testing is becoming more and more popular (Andrews and Baguley, 2012; Dienes, 2011), even among ‘frequentist’ researchers (Sanborn and Hills, 2014). Our third result, in Section 5.5, shows that the same holds for the switch criterion: if \mathcal{M}_0 is a singleton, so that the problem (5.1) reduces to a simple versus composite hypothesis test, then the evidence $r(X^n)$ associated with the switching criterion has the desired robustness property as well and thus in this sense behaves like the Bayes factor method. The advantage, from a frequentist point of view, of switching as compared to Bayes is then that switching is a lot more sensitive: our risk rate results directly imply that the Type II error ($1 - \text{power}$) of the switch criterion goes to 0 as soon as, at sample size n , the distance between the ‘true’ distribution μ_1 and the null model, i.e. $\inf_{\mu \in \mathcal{M}_0} \|\mu - \mu_1\|_2^2$ is of order $(\log \log n)/n$; for Bayes factor testing, in order for the Type-II error to reach 0, this distance must be of order $(\log n)/n$ (this was informally recognized by Lhéritier and Cazals (2015), who reported substantially larger power of switching as compared to the Bayes factor method in a sequential two-sample testing setting).

Thus, switching gives us ‘almost the best of three worlds’: minimax rate optimality up to a $\log \log n$ factor (in contrast to BFMS), consistency (in contrast to AIC-type methods) and nonasymptotic insensitivity to optional stopping (in contrast to standard Neyman-Pearson testing) in combination with a small Type-II error.

Organization This paper is organized as follows. The switch criterion is introduced in Section 5.2. In Section 5.3, we provide some preliminaries: we list the loss/risk functions for which our result holds, describe the sets in which the truth is assumed to lie, and discuss the tension between consistency and rate-optimality. Suitable post-model-selection estimators to be used in combination with the switch criterion are introduced in Section 5.4, after which our main result on the worst-case risk of the switch criterion is stated. We also go into the relationship between the switch criterion and the Hannan-Quinn criterion in that section. In Section 5.5 we define robust null hypothesis tests, give some examples, and show that testing by switching has the desired nonasymptotic robustness to optional stopping; in contrast, AIC does not satisfy such a property at all and the Hannan-Quinn criterion only satisfies an asymptotic analogue. We also provide some simulations that illustrate our results. Section 5.6 provides some additional discussion and ideas for future work. All proofs are given in Section 5.7.

Notations and conventions We use $x^n = x_1, \dots, x_n$ to denote n observations, each taking values in a sample space \mathcal{X} . For a set of parameters M , $\mu \in M$, and $x \in \mathcal{X}$, $p_\mu(x)$ invariably denotes the density or mass function of x under the distribution \mathbb{P}_μ of random variable X , taking values in \mathcal{X} . This is extended to n outcomes by independence, so that $p_\mu(x^n) := \prod_{i=1}^n p_\mu(x_i)$ and $\mathbb{P}_\mu(X^n \in A_n)$, abbreviated to $\mathbb{P}_\mu(A_n)$, denotes the probability

that $X^n \in A_n$ for $X^n = X_1, \dots, X_n$ i.i.d. $\sim \mathbb{P}_\mu$. Similarly, \mathbb{E}_μ denotes expectation under \mathbb{P}_μ . As is customary, we write $a_n \asymp b_n$ to denote $0 < \lim_{n \rightarrow \infty} \inf a_n/b_n \leq \lim_{n \rightarrow \infty} \sup a_n/b_n < \infty$. For notational simplicity we assume throughout this paper that whenever we refer to a sample size n , then $n \geq 3$ to ensure that $\log \log n$ is defined and positive.

Throughout the text, we refer to standard properties of exponential families without always giving an explicit reference; all desired properties can be found, in precise form, in (Barndorff-Nielsen, 1978) and, on a less formal level, in (Grünwald, 2007, Chapter 18,19).

5.2 Model selection by switching

The *switch distribution* (Van Erven et al., 2012; 2007) is a modification of the Bayesian predictive distribution, inspired by Dawid's (1984) 'prequential' approach to statistics and the related *Minimum Description Length* (MDL) Principle (Barron et al., 1998; Grünwald, 2007). The corresponding *switch criterion* can be thought of as Bayes factor model selection with a prior on meta-models, where each meta-model consists of a sequence of basic models and associated starting times: until time t_1 , follow model k_1 , from time t_1 to t_2 , follow model k_2 , and so on. The fact that we only need to select between two nested parametric models allows us to considerably simplify the set-up of Van Erven et al. (2012), who dealt with countably infinite sets of arbitrary models.

It is convenient to directly introduce the switch criterion as a modification of the Bayes factor model selection (BFMS). Assuming equal prior $1/2$ on each of the models \mathcal{M}_0 and \mathcal{M}_1 , BFMS associates each model \mathcal{M}_k , $k \in \{0, 1\}$, with a *marginal distribution* $p_{B,k}$ with

$$p_{B,k}(x^n) := \int_{\mu \in M_k} \omega_k(\mu) p_\mu(x^n) d\mu, \quad (5.2)$$

where ω_k is a prior density on M_k . It then selects model \mathcal{M}_1 if and only if $p_{B,1}(x^n) > p_{B,0}(x^n)$.

The basic idea behind MDL model selection is to generalize this in the sense that each model \mathcal{M}_k is associated with *some* 'universal' distribution $p_{U,k}$; one then picks the k for which $p_{U,k}(x^n)$ is largest. $p_{U,k}$ may be set to the Bayesian marginal distribution, but other choices may be preferable in some situations. Switching is an instance of this; in our simplified setting, it amounts to associating \mathcal{M}_0 with a Bayes marginal distribution $p_{B,0}$ as before. $p_{U,1}$ however is set to the *switch distribution* $p_{\text{sw},1}$. This distribution corresponds to a switch between models \mathcal{M}_0 and \mathcal{M}_1 at some sample point s , which is itself uncertain; before point s , the data are modelled as coming from \mathcal{M}_0 , using $p_{B,0}$; after point s , they are modelled as coming from \mathcal{M}_1 , using $p_{B,1}$. Formally, we denote the strategy that switches from the simple to the complex model after t observations by \bar{p}_t ; $p_{\text{sw},1}$ is then defined as the marginal distribution by averaging \bar{p}_t over t , with some probability mass function π (analogous to a Bayesian prior) over $t \in \{1, 2, \dots\}$:

$$\begin{aligned} \bar{p}_t(x^n) &= p_{B,0}(x^{t-1}) \cdot p_{B,1}(x_t, \dots, x_n \mid x^{t-1}) \\ p_{\text{sw},1}(x^n) &= \sum_{t=1}^{\infty} \pi(t) \bar{p}_t(x^n), \end{aligned}$$

where switching at $t = 1$ corresponds to predicting with $p_{B,1}$ at each data point, and switching at any $t > n$ to predicting with $p_{B,0}$. We remind the reader that even for i.i.d. models, $p_{B,1}(x_t, \dots, x_n \mid x^{t-1})$ usually depends on x^{t-1} — the Bayes predictive distribution learns from data. The model selection criterion δ_{sw} mapping sequences of arbitrary length to $k \in \{0, 1\}$ is then defined, for each n , as follows:

$$\delta_{\text{sw}}(x^n) = \begin{cases} 0 & \text{if } \frac{p_{\text{sw},1}(x^n)}{p_{B,0}(x^n)} \leq 1 \\ 1 & \text{if } \frac{p_{\text{sw},1}(x^n)}{p_{B,0}(x^n)} > 1 \end{cases}. \quad (5.3)$$

When defining $p_{\text{sw},1}$ it is sufficient to consider switching times that are equal to a power of two. Thus, we restrict attention to ‘priors’ π on switching time with support on $2^0, 2^1, 2^2, \dots$. For our subsequent results to hold, π should be such that $\pi(2^i)$ decays like $i^{-\kappa}$ for some $\kappa \geq 2$. An example of such a prior with $\kappa = 2$ is $\pi(2^i) = 1/((i+1)(i+2))$, $\pi(j) = 0$ for any j that is not a power of 2.

To prepare for Theorem 5.5, we instantiate the switch criterion to the problem (5.1). We define $p_{B,1}$ as any distribution of the form (5.2) where ω_1 is a continuous prior density on M_1 that is strictly positive on all $\mu \in M_1$. To define $p_{B,0}$ we need to take a slight detour, because we parameterized \mathcal{M}_0 in terms of an M_0 that has a fixed value on its final $m_1 - m_0$ components: it is an m_0 -dimensional family with an m_1 -dimensional parameterization, so one cannot easily express a prior on \mathcal{M}_0 as a density on M_0 . To overcome this, we distinguish between the case that $m_0 = 0$ and $m_0 > 0$. In the former case M_0 has a single element v , and we define $p_{B,0} = p_v$. In the latter case, we define $\Pi'_0 : M_0 \rightarrow \mathbb{R}^{m_0}$ as the projection of $\mu \in M_0$ on its first m_0 components, and $\Pi'_0(M_0) := \{\Pi'_0(\mu) : \mu \in M_0\}$. For $\mu \in M_0$, we define $p_{\Pi'_0(\mu)} = p_\mu$, and we then let ω_0 be a continuous strictly positive prior density on $\Pi'_0(M_0)$, and we define $p_{B,0}(x^n) := \int_{\mu' \in \Pi'_0(M_0)} \omega_0(\mu') p_{\mu'}(x^n) d\mu'$.

Two important remarks are in order: first, the fact that we associate \mathcal{M}_1 with a distribution incorporating a ‘switch’ from \mathcal{M}_0 to \mathcal{M}_1 *does not mean* that we really believe that data were sampled, until some point t , according to \mathcal{M}_0 and afterwards according to \mathcal{M}_1 . Rather, it is suggested by prequential and MDL considerations, which suggest that one should pick the model that performs best in sequentially predicting data; and if the data are sampled from a distribution in \mathcal{M}_1 that is not in \mathcal{M}_0 , but quite close to it in KL divergence, then $p_{B,1}$ is suboptimal for sequential prediction, and can be substantially outperformed by $p_{\text{sw},1}$. This is explained at length by Van Erven et al. (2012), and Figure 1 in that paper especially illustrates the point. The same paper also explains how one can use dynamic programming to arrive at an implementation that has the same computational efficiency as computation of the standard Bayes model selection decision.

Second, the criterion (5.3) as defined here is not 100% equivalent to the special case of the construction of Van Erven et al. (2012) specialized to two models, but rather an easily-explained simplification thereof. Yet, all our results continue to hold if we were to follow the original construction, as explained in Section 5.7.5; and conversely, the strong consistency result for the construction of Van Erven et al. (2012) trivially continues to hold for the criterion (5.3) used in the present paper.

5.3 Rate-optimality of post-model selection estimators

This section contains some background to our main result, Theorem 5.5. In Section 5.3.1, we first list the loss functions for which our main result holds, and define the CINECSI sets in which the truth assumed to lie. We then discuss the minimax parametric risk for our model selection problem in Section 5.3.2. This section ends with a discussion on the generality of the impossibility result of Yang (2005) in Section 5.3.3.

5.3.1 Loss functions and CINECSI sets

Let $\mathcal{M} = \{p_\mu \mid \mu \in M\}$ be an exponential family given in its mean-value parameterization with $M \subset \mathbb{R}^m$ a product of m open, possibly but not necessarily unbounded intervals for some $m > 0$; see Section 5.7 for a formal definition of exponential families and mean-value parameterizations. Note that we do not require the family to be ‘full’; for example, the Bernoulli model with success probability $\mu \in M_1 = (0.2, 0.4)$ counts as an exponential family in our (standard) definition.

Suppose that we measure the quality of a density $p_{\mu'}$ as an approximation to p_μ by a loss function $L : M \times M \rightarrow \mathbb{R}$. The standard definition of the (instantaneous) *risk* of estimator $\check{\mu} : \bigcup_{i \geq 0} \mathcal{X}^i \rightarrow M$ at sample size n , as defined relative to loss L , is given by its expected loss,

$$R(\mu, \check{\mu}, n) = \mathbb{E}_\mu [L(\mu, \check{\mu}(X^n))],$$

where \mathbb{E}_μ denotes expectation over X_1, \dots, X_n i.i.d. $\sim \mathbb{P}_\mu$. Popular loss functions are:

1. The *squared error loss*: $d_{SQ}(\mu', \mu) = \|\mu' - \mu\|_2^2$;
2. The *standardized squared error loss* which is a version of the *squared Mahalanobis distance*, defined as

$$d_{ST}(\mu' \parallel \mu) := (\mu - \mu')^T I(\mu') (\mu - \mu'), \quad (5.4)$$

where T denotes transpose, $I(\cdot)$ is the Fisher information matrix, and we view μ and μ' as column vectors;

3. The Rényi divergence of order $1/2$, defined as

$$d_R(\mu', \mu) = -2 \log \mathbb{E}_{\mu'} \left[\left(p_\mu(X) / p_{\mu'}(X) \right)^{\frac{1}{2}} \right];$$

4. The squared Hellinger distance $d_{H^2}(\mu', \mu) = 2 \left(1 - \mathbb{E}_{\mu'} \left[\left(p_\mu(X) / p_{\mu'}(X) \right)^{\frac{1}{2}} \right] \right)$;
5. The KL (Kullback-Leibler) divergence $D(p_{\mu'} \parallel p_\mu)$, henceforth abbreviated to $D(\mu' \parallel \mu)$.

We note that there is a direct relationship between the Rényi divergence and squared Hellinger distance:

$$d_{H^2}(\mu', \mu) = 2 \left(1 - e^{-d_R(\mu', \mu)/2} \right). \quad (5.5)$$

In fact, as we show below, these loss functions are all equivalent (equal up to universal constants) on CINECSI sets. Such sets will play an important role in the sequel. They are defined as follows:

Definition 5.1 (CINECSI). A *CINECSI* (Connected, Interior-Non-Empty-Compact-Subset-of-Interior) subset of a set M is a connected subset of the interior of M that is itself compact and has nonempty interior.

The following proposition is proved in Section 5.7.

Proposition 5.2. Let M be the mean-value parameter space of an exponential family as above, and let M' be an CINECSI subset of M . Then there exist positive constants c_1, c_2, \dots, c_6 such that for all $\mu, \mu' \in M'$,

$$c_1 \|\mu' - \mu\|_2^2 \leq c_2 \cdot d_{ST}(\mu' \|\mu) \leq d_{H^2}(\mu', \mu) \leq d_R(\mu', \mu) \leq D(\mu' \|\mu) \leq c_3 \|\mu' - \mu\|_2^2. \quad (5.6)$$

and for all $\mu' \in M', \mu \in M$ (i.e. μ is now not restricted to lie in M'),

$$d_{H^2}(\mu', \mu) \leq c_4 \|\mu' - \mu\|_2^2 \leq c_5 \cdot d_{ST}(\mu' \|\mu) \leq c_6 \|\mu' - \mu\|_2^2. \quad (5.7)$$

CINECSI subsets are a variation on the INECCSI sets of (Grünwald, 2007). Our main result, Theorem 5.5, holds for all of the above loss functions, and for general ‘sufficiently efficient’ estimators. While the equivalence of the losses above on CINECSI sets is a great help in the proofs, we emphasize that we never require these estimators to be restricted to CINECSI subsets of M — although, since we require M to be open, every ‘true’ $\mu \in M$ will lie in *some* CINECSI subset M' of M , a statistician who employs the model \mathcal{M} cannot know what this M' is, so such a requirement would be unreasonably strong.

5.3.2 Minimax parametric risk

We say that a quantity f_n converges at rate g_n if $f_n \asymp g_n$. We say that an estimator $\check{\mu}$ is *minimax rate optimal* relative to a model $\mathcal{M} = \{p_\mu \mid \mu \in M\}$ restricted to a subset $M' \subset M$ if

$$\sup_{\mu \in M'} R(\mu, \check{\mu}, n)$$

converges at the same rate as

$$\inf_{\dot{\mu}} \sup_{\mu \in M'} R(\mu, \dot{\mu}, n), \quad (5.8)$$

where $\dot{\mu}$ ranges over all estimators of μ at sample size n , that is, all measurable functions from \mathcal{X}^n to M .

For most parametric models encountered in practice, the minimax risk (5.8) is of order $1/n$ when R is defined relative to any of the loss measures defined in Section 5.3.1 and M' is an arbitrary CINECSI subset of M (Van der Vaart, 1998). In particular this holds if \mathcal{M} is an exponential family. For this reason, from now on we refer to $1/n$ as the *minimax parametric rate*. Note that, crucially, the restriction $\mu \in M'$ is imposed only on the data-generating distribution, not on the estimators, and, since we will require models with open parameter sets M such that for every $\delta > 0$, there is a CINECSI subset M'_δ of M with $\sup_{\mu \in M} \inf_{\mu' \in M'_\delta} \|\mu - \mu'\|_2^2 < \delta$, every possible $\mu \in M$ will also lie in some CINECSI subset M'_δ that ‘nearly’ covers M_δ . This makes the restriction to CINECSI M' in the definition above a mild one. Still, it is necessary: at least for the squared error loss, for most exponential families (the exception being the Gaussian location family), we have

$\inf_{\mu} \sup_{\mu \in M'_\delta} R(\mu, \mu, n) = C_\delta/n$ for some constant $C_\delta > 0$, but the smallest constant for which this holds may grow arbitrarily large as $\delta \rightarrow 0$, the reason being that the determinant of the Fisher information may tend to 0 or ∞ as $\delta \rightarrow 0$.

Now consider a model selection criterion $\delta : \bigcup_{i>0} X^i \rightarrow \{0, 1, \dots, K-1\}$ that selects, for given data x^n of arbitrary length n , one of a finite number K of parametric models $\mathcal{M}_0, \dots, \mathcal{M}_{K-1}$ with respective parameter sets M_0, \dots, M_{K-1} . One way to evaluate the quality of δ is to consider the risk attained after first selecting a model and then estimating the parameter vector μ using an estimator $\check{\mu}_k$ associated with each model \mathcal{M}_k . This *post-model selection estimator* (Leeb and Pötscher, 2005) will be denoted by $\check{\mu}_{\check{k}}(x^n)$, where \check{k} is the index of the model selected by δ . The risk of a model selection criterion δ is thus $R(\mu, \delta, n) = \mathbb{E}_\mu [L(\mu, \check{\mu}_{\check{k}}(X^n))]$, where L is a given loss function, and its worst-case risk relative to μ restricted to $M'_k \subset M_k$ is given by

$$\sup_{\mu \in M'_k} R(\mu, \delta, n) = \sup_{\mu \in M'_k} \mathbb{E}_\mu [L(\mu, \check{\mu}_{\check{k}}(X^n))]. \quad (5.9)$$

We are now ready to define what it means for a model selection criterion to achieve the minimax parametric rate.

Definition 5.3. A model selection criterion δ achieves the minimax parametric rate if there exist estimators $\check{\mu}_k$, one for each \mathcal{M}_k under consideration, such that, for every CINECSI subset M'_k of M :

$$\sup_{\mu \in M'_k} R(\mu, \delta, n) \asymp 1/n.$$

Just as in the fixed-model case, the restriction $\mu \in M'_k$ is imposed only on the data-generating distribution, not on the estimators.

5.3.3 The result of Yang (2005) transplanted to our setting

In this paper, as stated in the introduction, we further specialize the setting above to problem (5.1) where we select between two nested exponential families, which we shall always assume to be given in their mean-value parameterization. To be precise, the ‘complex’ model \mathcal{M}_1 contains distributions from an exponential family parametrized by an m_1 -dimensional mean vector μ , and the ‘simple’ model \mathcal{M}_0 contains distributions with the same parametrization, where the final $m_1 - m_0$ components are fixed to values $v_{m_0+1}, \dots, v_{m_1}$. We introduce some notation to deal with the assumption that \mathcal{M}_1 and its restriction of \mathcal{M}_0 to its first m_0 components are products of open intervals. Formally, we require that \mathcal{M}_1 and \mathcal{M}_0 are of the form

$$\begin{aligned} \mathcal{M}_1 &= (\zeta_{1,1}, \eta_{1,1}) \times \dots \times (\zeta_{1,m_1}, \eta_{1,m_1}) \\ \mathcal{M}_0 &= (\zeta_{0,1}, \eta_{0,1}) \times \dots \times (\zeta_{0,m_0}, \eta_{0,m_0}) \times \{v_{m_0+1}\} \times \dots \times \{v_{m_1}\} \end{aligned} \quad (5.10)$$

where, for $j = 1, \dots, m_0$, we have $-\infty \leq \zeta_{1,j} \leq \zeta_{0,j} < \eta_{0,j} \leq \eta_{1,j} \leq \infty$; and for $j = m_0 + 1, \dots, m_1$, we have $-\infty \leq \zeta_{1,j} < v_j < \eta_{1,j} \leq \infty$.

For example, \mathcal{M}_1 could contain all normal distributions with mean μ and variance σ^2 , with mean value parameters $\mu_1 = \mu^2 + \sigma^2$ and $\mu_2 = \mu$, and $\mathcal{M}_1 = (0, \infty) \times (-\infty, \infty)$, while

\mathcal{M}_0 could contain all normal distributions with mean zero and unknown variance σ^2 , so $M_0 = (0, \infty) \times \{0\}$.

Yang (2005) showed in a linear regression context that a model selection criterion cannot both achieve the minimax optimal parametric rate and be consistent; a practitioner is thus forced to choose between a rate-optimal method such as AIC and a consistent method such as BIC. Inequality (5.12) below provides some insight into why this *AIC-BIC dilemma* can occur. A similar inequality appears in Yang's paper for his linear regression context, but it is still valid in our exponential family setting, and the derivation — which we now give — is essentially equivalent.

To state the inequality, we need to relate $\mu_1 \in M_1$ to a component in M_0 . For any given $\mu_1 = (\mu_{1,1}, \dots, \mu_{1,m_1})^T \in M_1$, we will define

$$\Pi_0(\mu_1) := (\mu_{1,1}, \dots, \mu_{1,m_0}, \nu_{m_0+1}, \dots, \nu_{m_1})^T \quad (5.11)$$

to be the *projection* of μ_1 on M_0 . The difference between Π_0 of (5.11) and Π'_0 in Section 5.2 is that Π_0 is a function from \mathbb{R}^{m_1} to \mathbb{R}^{m_1} , whereas Π'_0 is a function from \mathbb{R}^{m_1} to \mathbb{R}^{m_0} ; $\Pi_0(\mu_1)$ and $\Pi'_0(\mu_1)$ agree in the first m_0 components. Note that $\Pi_0(\mu_1)$ obviously minimizes, among all $\mu \in M_0$, the squared Euclidean distance $\|\mu - \mu_1\|_2^2$ to μ_1 ; somewhat less obviously it also minimizes, among $\mu \in M_0$, the KL divergence $D(p_{\mu_1} \| p_{\mu})$ (Grünwald, 2007, Chapter 19); we may thus think of it as the ‘best’ approximation of the ‘true’ μ_1 within M_0 ; we will usually abbreviate $\Pi_0(\mu_1)$ to μ_0 .

Let A_n be the event that the complex model is selected at sample size n . Since \mathcal{M}_1 is an exponential family, the MLE $\widehat{\mu}_1$ is unbiased and $\widehat{\mu}_0$ coincides with $\widehat{\mu}_1$ in the first m_0 components, so that $\mathbb{E}_{\mu_1} [\mu_0 - \widehat{\mu}_0(X^n)] = 0$, and hence we can rewrite, for any $\mu_1 \in M_1$, the squared error risk as

$$\begin{aligned} R(\mu_1, \delta, n) &= \mathbb{E}_{\mu_1} \left[\mathbf{1}_{A_n} \|\mu_1 - \widehat{\mu}_1(X^n)\|_2^2 + \mathbf{1}_{A_n^c} \|\mu_1 - \widehat{\mu}_0(X^n)\|_2^2 \right] \\ &= \mathbb{E}_{\mu_1} \left[\mathbf{1}_{A_n} \|\mu_1 - \widehat{\mu}_1(X^n)\|_2^2 + \mathbf{1}_{A_n^c} \|\mu_0 - \widehat{\mu}_0(X^n)\|_2^2 + \mathbf{1}_{A_n^c} \|\mu_1 - \mu_0\|_2^2 \right] \\ &\leq \mathbb{E}_{\mu_1} \left[\|\mu_1 - \widehat{\mu}_1(X^n)\|_2^2 + \|\mu_0 - \widehat{\mu}_0(X^n)\|_2^2 \right] + \mathbb{P}(A_n^c) \|\mu_1 - \mu_0\|_2^2 \\ &\leq 2R(\mu_1, \widehat{\mu}_1, n) + \mathbb{P}(A_n^c) \|\mu_1 - \mu_0\|_2^2. \end{aligned} \quad (5.12)$$

The first part of the proof of our main result, Theorem 5.5, extends this decomposition to general estimators and loss functions.

The first term on the right of (5.12) is of order $1/n$. The second term depends on the ‘Type-II error’, i.e. the probability of selecting the simple model when it is not actually true. A low worst-case risk is attained if this probability is small, even if the true parameter is close to μ_0 . This does leave the possibility for a risk-optimal model selection criterion to incorrectly select the complex model with high probability. In other words, a risk-optimal model selection method may not be consistent if the simple model is correct. The theorem by Yang (2005), arguing from decomposition (5.12), essentially demonstrates that it cannot be. Due to the general nature of (5.12), it seems likely that his result holds in much more general settings: a procedure attains a low worst-case risk by selecting the complex model with high probability, which is excellent if the complex model is indeed true, but leads to inconsistency if the simple model is correct. Indeed, we have shown in earlier work that the dilemma is not restricted to linear regression, but occurs in our exponential family

problem (5.1) as well as long as $\mathcal{M}_0 = \{v\}$ is a singleton (see Van der Pas (2013) for the proof, which is a simple adaptation of Yang's proof that, we suspect, can be extended to nonsingleton \mathcal{M}_0 as well). Hence, as the switch criterion is strongly consistent Van Erven et al. (2012), we know that the worst-case risk rate of the switch criterion cannot be of the order $1/n$ in general.

5.4 Main result

We perform model selection by using the switch criterion, as specified in Section 5.2. After the model selection, we estimate the underlying parameter μ . We discuss post-model selection estimators suitable to our problem in Section 5.4.1. We are then ready to present our main result, Theorem 5.5 in Section 5.4.2, stating that the worst-case risk for the switch criterion under the loss functions listed in Section 5.3.1 attains the minimax parametric rate up to a $\log \log n$ factor.

5.4.1 Post-model selection: sufficiently efficient estimators

Our goal is to determine the worst-case rate for the switch criterion applied to two nested exponential families, which we combine with an estimator as follows: if the simple model is selected, μ will be estimated by an estimator $\check{\mu}_0$ with range M_0 . If the complex model is selected, the estimate of μ will be provided by another estimator $\check{\mu}_1$ with range M_1 . Our result will hold for all estimators $\check{\mu}_0$ and $\check{\mu}_1$ that are *sufficiently efficient*:

Definition 5.4 (sufficiently efficient). The estimators $\{\check{\mu}_k \rightarrow M_k \mid k \in \{0, 1\}\}$ are *sufficiently efficient* with respect to a divergence measure $d_{\text{gen}}(\cdot \parallel \cdot)$ if (with $\mu_0 = \Pi_0(\mu_1)$) as in (5.11)), for every CINECSI subset M'_1 of M_1 , there exists a constant $C > 0$ such that for all n ,

$$\sup_{\mu_1 \in M'_1} \mathbb{E}_{\mu_1}[d_{\text{gen}}(\mu_0 \parallel \check{\mu}_0)] \leq C \cdot \sup_{\mu_1 \in M'_1} \mathbb{E}_{\mu_1}[d_{\text{gen}}(\mu_1 \parallel \check{\mu}_1)] \leq \frac{C}{n}. \quad (5.13)$$

Example 1. [Sufficient efficiency for MLE's for squared (standardized) error and Hellinger] If \mathcal{M}_0 and \mathcal{M}_1 are exponential families given in their mean-value parameterization with $M_0 \subset M_1$ as in (5.10), then \mathcal{M}_1 has sufficient statistic $\varphi \equiv (\varphi_1, \dots, \varphi_{m_1})^T : \mathcal{X} \rightarrow \mathbb{R}^{m_1}$ (see Section 5.7 for the formal definition). Now by standard properties of exponential families, if

$$n^{-1} \sum \varphi(X_i) \in M_1, \quad (5.14)$$

then the ML estimator for model \mathcal{M}_k is equal to $n^{-1} \sum_{i=1}^n \varphi(X_i)$. For many full families such as the full (multivariate) Gaussians, Gamma and many others, (5.14) holds μ -almost surely for each n , for all $\mu \in M_1$. Then the MLE is almost surely well-defined for M_1 . We can then take $\check{\mu}_1 := \widehat{\mu}_1$ to be the MLE for \mathcal{M}_1 , and $\check{\mu}_0$ to be its projection on the first m_0 coordinates (usually (5.14) will still hold for M_0 and then this $\check{\mu}_0$ will be the MLE for \mathcal{M}_0). This pair of estimators will be sufficiently efficient for (standardized) squared error and squared Hellinger distance, i.e. (5.13) holds for these three losses. To show this, note that from (5.27) in Proposition 5.2, we see that it is sufficient to show that (5.13) holds for the squared error loss. Since the j -th component of $\widehat{\mu}_1$ is equal to $n^{-1} \sum_{i=1}^n \varphi_j(X_i)$ and

$\mathbb{E}_{\mu_1} [n^{-1} \sum_{i=1}^n \varphi_j(X_i)] = \mu_{1,j}$ and $\text{var}_{\mu_1} [n^{-1} \sum_{i=1}^n \varphi_j(X_i)] = n^{-1} \text{var}_{\mu_1} [\varphi_j(X_1)]$, it suffices to show that

$$\sup_{\mu_1 \in M'_1} \sup_{j=1, \dots, m_1} \text{var}_{\mu_1} [\varphi_j(X_1)] = O(1),$$

which is indeed the case since M'_1 is an CINECSI set, so that the variance of all φ_j 's is uniformly bounded on M'_1 (Barndorff-Nielsen, 1978).

Example 2. [Other sufficiently efficient estimators for squared (standardized) error and Hellinger] For models such as the Bernoulli or multinomial, (5.14) may fail to hold with positive probability: the full Bernoulli exponential family does not contain the distributions with $P(X_1 = 1) = 1$ and $P(X_1 = 0) = 1$, so if after n examples, only zeros or only ones have been observed, the MLE is undefined. We can then go either of three ways. The first way, which we shall not pursue in detail here, is to work with so-called 'aggregate' exponential families, which are extensions of full families to their limit points. For models with finite support (such as the multinomial) these are well-defined (Barndorff-Nielsen, 1978, page 154–158) and then the MLE's for these extended families are almost surely well-defined again, and the MLE's are sufficiently efficient by the same reasoning as above. Another approach that works in some cases (e.g. multinomial) is to take $\check{\mu}_1$ to be a truncated MLE, that, at sample size n , maps X^n to the MLE within some CINECSI subset $M_1^{(n)}$ of M_1 , where $M_1^{(n)}$ converges to M_1 as n increases in the sense that $\sup_{\mu \in M_1^{(n)}, \mu' \in M_1 \setminus M_1^{(n)}} \|\mu - \mu'\|_2^2 = O(1/n)$. The resulting truncated MLE, and its projection on M_0 (usually itself a truncated MLE) will then again be sufficiently efficient. This approach also works if the models M_0 and M_1 are not full but restricted families to begin with. For full families though, a more elegant approach than truncating MLE's is to work with Bayesian posterior MAP estimates with conjugate priors. For steep exponential families (nearly all families one encounters in practice are steep), one can always find conjugate priors such that the Bayes MAP estimates based on these priors exist and take a value in M_1 almost surely (Grünwald and de Rooij, 2005). They then take the form $\check{\mu}_1 = \sum_{i=1}^n (\varphi(X_i) + \lambda_0 \mu_1^\circ) / (n + \lambda_0)$, where $\lambda_0 > 0$ and $\mu_1^\circ \in M_1$ are determined by the prior. $\check{\mu}_0$ can then again be taken to be the projection of $\check{\mu}_1$ onto M_0 . Under the assumption that μ_1 is contained in a CINECSI set M'_1 , one can now again show, using the same arguments as in Example 1, that such estimators are sufficiently efficient for squared (standardized) error and Hellinger loss.

Example 3. [Sufficient efficiency for Rényi and KL divergence] As is well-known, for the multivariate Gaussian model with fixed covariance matrix, the squared error risk and KL divergence are identical up to constant factors, so the unrestricted MLE's will still be sufficiently efficient for KL divergence. For other models, though, the MLE will not always be sufficiently efficient. For example, with the Bernoulli model and other models with finite support, to make the unrestricted MLE's well-defined, we would have to extend the family to its boundary points as indicated in Example 1. Since, however, for any $0 < \mu < 1$ and $\mu' = 0$, the KL divergence $D(\mu \| \mu') = \infty$ and $\mathbb{P}_\mu(\hat{\mu}(X^n) = \mu') > 0$, the unrestricted MLE in the full Bernoulli model including the boundaries will have infinite risk and thus will not be sufficiently efficient. The MAP estimators tend to behave better though: Grünwald and de Rooij (2005) implicitly show that for 1-dimensional families, under weak conditions on the family (Condition 1 underneath Theorem 1 in their paper) —

which were shown to hold for a number of families such as Bernoulli, Poisson, geometric — sufficient efficiency for the KL divergence still holds for MAP estimators of the form above. We conjecture that a similar result can be shown for multidimensional families, but will not attempt to do so here.

5.4.2 Main result: risk of the switch criterion

We now present our main result, which states that for the exponential family problem under consideration, the worst-case instantaneous risk rate of δ_{sw} is of order $(\log \log n)/n$. Hence, the worst-case instantaneous risk of δ_{sw} is very close to the lower bound of $1/n$, while the criterion still maintains consistency.

The theorem holds for any of the loss functions listed in Section 5.3.1. We denote this by using the generic loss function d_{gen} , which can be one of the following loss functions: squared error loss, standardized squared error loss, KL divergence, Rényi divergence of order $1/2$, or squared Hellinger distance.

Theorem 5.5. Let $\mathcal{M}_0 = \{p_\mu \mid \mu \in M_0\}$ and $\mathcal{M}_1 = \{p_\mu \mid \mu \in M_1\}$ be nested exponential families in their mean-value parameterization, where $M_0 \subseteq M_1$ are of the form (5.10). Assume:

1. $\check{\mu}_0$ and $\check{\mu}_1$ are sufficiently efficient estimators relative to the chosen loss d_{gen} ;
2. δ_{sw} is constructed with $p_{B,0}$ and $p_{B,1}$ defined as in Section 5.2 with priors ω_k that admit a strictly positive, continuous density;
3. and $p_{\text{sw},1}$ is defined relative to a prior π with support on $\{0, 1, 2, 4, 8, \dots\}$ and $\pi(2^i) \propto i^{-\kappa}$ for some $\kappa \geq 2$.

Then for every CINECSI subset M'_1 of M_1 , we have:

$$\sup_{\mu_1 \in M'_1} R(\mu_1, \delta_{\text{sw}}, n) = O\left(\frac{\log \log n}{n}\right),$$

for $R(\mu, \delta_{\text{sw}}, n)$ the risk at sample size n defined relative to the chosen loss d_{gen} .

Example 4. [Our setting vs. Yang’s] Yang (2005) considers model selection between two nested linear regression models with fixed design, where the errors are Gaussian with fixed variance. The risk is measured as the in-model squared error risk (‘in-model’ means that the loss is measured conditional on a randomly chosen design point that already appeared in the training sample). Within this context he shows that every model selection criterion that is (weakly) consistent cannot achieve the $1/n$ minimax rate. The exponential family result above leads one to conjecture that the switch distribution achieves $O((\log \log n)/n)$ risk in Yang’s setting as well. We suspect that this is so, but actually showing this would require substantial additional work. Compared to our setting, Yang’s setting is easier in some and harder in other respects: under the fixed-variance, fixed design regression model, the Fisher information is constant, making asymptotic results hold nonasymptotically, which would greatly facilitate our proofs (and obliterate any need to consider CINECSI sets or

undefined MLE's). On the other hand, evaluating the risk conditional on a design point is not something that can be directly embedded in our proofs.

Example 5. [Switching vs. Hannan-Quinn] In their comments on Van Erven et al. (2012), Lauritzen (2012) and Cavanaugh (2012) suggested a relationship between the switch model selection criterion and the criterion due to Hannan and Quinn (1979). For the exponential family models under consideration, the Hannan-Quinn criterion with parameter c , denoted as HQ, selects the simple model, i.e. $\delta_{\text{HQ}}(x^n) = 0$, if

$$-\log p_{\hat{\mu}_0}(x^n) < -\log p_{\hat{\mu}_1}(x^n) + c \log \log n,$$

and the complex model otherwise. In their paper, Hannan and Quinn show that this criterion is strongly consistent for $c > 1$.

As shown by Barron et al. (1999), under some regularity conditions, penalized maximum likelihood criteria achieve worst-case quadratic risk of the order of their penalty divided by n . One can show (details omitted) that this is also the case in our specific setting and hence, that the worst-case risk rate of HQ for our problem is of order $(\log \log n)/n$. Our main result, Theorem 5.5, shows that the same risk rate is achieved by the switch distribution, thus partially confirming the conjecture of Lauritzen (2012) and Cavanaugh (2012): HQ achieves the same risk rate as the switch distribution and, for the right choice of c , is also strongly consistent. This suggests that the switch distribution and HQ, at least for some specific value c_0 , may behave asymptotically indistinguishably. The earlier results of Van der Pas (2013) suggest that this is indeed the case if \mathcal{M}_0 is a singleton; if \mathcal{M}_0 has dimensionality larger than 0, this appears to be a difficult question which we will not attempt to resolve here — in this sense the conjecture of Lauritzen (2012) and Cavanaugh (2012) has only been partially resolved.

Because HQ and δ_{sw} have been shown to be both strongly consistent and achieve the same rates for this problem, one may wonder whether one criterion is to be preferred over the other. For this parametric problem, HQ has the advantage of being simpler to analyze and implement. The criterion δ_{sw} can however, be used to define a robust hypothesis test as in Section 5.5 below. As we shall see there, HQ is insensitive to optional stopping in an asymptotic sense only, whereas robust tests such as the switch criterion are insensitive to optional stopping in a much stronger, nonasymptotic sense. Except for the normal location model, for which the asymptotics are precise, the HQ criterion cannot be easily adapted to define such a robust, nonasymptotic test. Another advantage of switching is that it can be combined with arbitrary priors and applied much more generally, for example when the constituting models are themselves nonparametric (Lhéritier and Cazals, 2015), are so irregular that standard asymptotics such as the law of the iterated logarithm are no longer valid, or are represented by black-box predictors such that ML estimators and the like cannot be calculated. In all of these cases the switch criterion can still be defined and — given the explanation in the introduction of Van Erven et al. (2012) — one may still expect it to perform well.

5.5 Robust null hypothesis tests

Bayes factor model selection, the switch criterion, AIC, BIC, HQ and most model selection methods used in practice are really based on thresholding the output of a more informative *model comparison method*. This is defined as a function from data of arbitrary size to the nonnegative reals. Given data x^n , it outputs a number $r(x^n)$ between 0 and ∞ that is a deterministic function of the data x^n . Every model comparison method r and threshold t has an associated model selection method $\delta_{r,t}$ that outputs 1 (corresponding to selecting model \mathcal{M}_1) if $r(x^n) \leq t$, and 0 otherwise. As explained below, such model comparison methods can often be viewed as performing a null hypothesis with \mathcal{M}_0 the null hypothesis, \mathcal{M}_1 the alternative hypothesis and t akin to a significance level.

Example 1 (BFMS): The output of the Bayes factor model comparison method is the posterior odds ratio $r_{\text{Bayes}}(x^n) = \mathbb{P}(\mathcal{M}_0|x^n)/\mathbb{P}(\mathcal{M}_1|x^n)$. The associated model selection method (BFMS) with threshold t selects model \mathcal{M}_1 if and only if $r_{\text{Bayes}}(x^n) \leq t$.

Example 2 (AIC): Standard AIC selects model \mathcal{M}_1 if $\log(p_{\hat{\mu}_1}(x^n)/p_{\hat{\mu}_0}(x^n)) > m_1 - m_0$. We may however consider more conservative versions of AIC that only select \mathcal{M}_1 if

$$\log(p_{\hat{\mu}_1}(x^n)/p_{\hat{\mu}_0}(x^n)) - (m_1 - m_0) \geq -\log t. \quad (5.15)$$

We may thus think of AIC as a model comparison method that outputs the left-hand side of (5.15), and that becomes a model selection method when supplied with a particular t .

Now classical Neyman-Pearson null hypothesis testing requires the *sampling plan*, or equivalently, the *stopping rule*, to be determined in advance to ensure the validity of the subsequent inference. In the important special case of (generalized) likelihood ratio tests, this even means that the sample size n has to be fixed in advance. In practice, greater flexibility in choosing the sample size n is desirable (Wagenmakers (2007) provides sophisticated examples and discussion). Below, we discuss hypothesis tests that allow such flexibility by virtue of the property that their Type I-error probability remains bounded irrespective of the stopping rule used. These *robust* null hypothesis tests are defined below. As will be shown, whenever the null hypothesis $\mathcal{M}_0 = \{p_{\mu_0}\}$ is ‘simple’, i.e. a singleton (simple vs. composite testing), both Bayes factor model selection (BFMS) and the switch distribution define such robust null hypothesis tests, whereas AIC does not and HQ does so only in an asymptotic sense. As we argue in Section 5.5.3, the advantage of switching over BFMS is then that, while both share the robustness Type-I error property, switching has significantly smaller Type-II error (larger power) than BFMS when the ‘truth’ is close to \mathcal{M}_0 , which is a direct consequence of it having a smaller risk under the alternative \mathcal{M}_1 . To make this point concrete, and to indicate what may happen if \mathcal{M}_0 is not a singleton, we provide a simulation study in Section 5.5.4.

5.5.1 Bayes factors with singleton \mathcal{M}_0 are robust under optional stopping

In many cases, for each $0 < \alpha < 1$ there is an associated threshold $t(\alpha)$, which is a strictly increasing function of α , such that for every $t \leq t(\alpha)$ we have that $\delta_{r,t}$ becomes a null hypothesis significance test (NHST) with type-I error probability bounded by α . In particular, then $\delta_{r,t(\alpha)}$ is a standard NHST with type-I error bounded by α . For example,

for AIC with $M_0 = \{0\}$ and $M_1 = \mathbb{R}$ representing the normal family of distributions with unit variance, we may select $t(\alpha) = \exp(-2/z_{\alpha/2}^2)$, where $z_{\alpha/2}$ is the upper $(\alpha/2)$ -quantile of the standard normal distribution. This results in the generalized likelihood ratio test at significance level α .

We say that model comparison method r defines a *robust null hypothesis test* for null hypothesis \mathcal{M}_0 and significance level α if for all $\mu_0 \in M_0$,

$$P_{\mu_0}(\exists n : \delta_{r,t(\alpha)}(X^n) = 1) \leq \alpha. \quad (5.16)$$

Hence, a test that satisfies (5.16) is a valid NHST test at significance level α , independently of the stopping rule used. If a researcher can obtain a maximum of n observations, the probability of incorrectly selecting the complex model will remain bounded away from one, regardless of the actual number of observations made.

It is well-known that Bayes factor model selection provides a robust null hypothesis test with $t(\alpha) = \alpha$ for *all* fixed α between 0 and 1, as long as \mathcal{M}_0 is a singleton. In other words, we may view the output of BFMS as a ‘robust’ variation of the p -value. This was already noted by Edwards et al. (1963) and interpreted as a frequentist justification for BFMS; it also follows immediately from the following result.

Theorem 5.6 (Special Case of Eq. (2) of Shafer et al. (2011)). Let $\mathcal{M}_0, \mathcal{M}_1, M_0$ and M_1 be as in Theorem 5.5 with common support $\mathcal{X} \subset \mathbb{R}^d$ for some $d > 0$. Let (X_1, X_2, \dots) be an infinite sequence of random vectors all with support \mathcal{X} , and fix two distributions, $\bar{\mathbb{P}}_0$ and $\bar{\mathbb{P}}_1$ on \mathcal{X}^∞ (so that under both $\bar{\mathbb{P}}_0$ and $\bar{\mathbb{P}}_1$, (X_1, X_2, \dots) constitutes a random process). Let, for each n , $\bar{p}_j^{(n)}$ represent the marginal density of (X_1, \dots, X_n) for the first n outcomes under distribution $\bar{\mathbb{P}}_j$, relative to some product measure ρ^n on $(\mathbb{R}^d)^n$ (we assume $\bar{\mathbb{P}}_0$ and $\bar{\mathbb{P}}_1$ to be such that these densities exist). Then for all $\alpha \geq 0$,

$$\bar{\mathbb{P}}_0 \left(\exists n : \frac{\bar{p}_0^{(n)}(X^n)}{\bar{p}_1^{(n)}(X^n)} \leq \alpha \right) \leq \alpha.$$

We first apply this result for Bayes factor model selection, with model priors $\pi_0 = \pi_1 = 1/2$, so that $r_{\text{Bayes}}(x^n) = \mathbb{P}(\mathcal{M}_0|x^n)/\mathbb{P}(\mathcal{M}_1|x^n) = p_{B,0}(x^n)/p_{B,1}(x^n)$. We immediately see that If $M_0 = \{\mu_0\}$ represents a singleton null model, then Bayes factor model selection constitutes a robust hypothesis test for null hypothesis \mathcal{M}_0 .

What happens if \mathcal{M}_0 is not singleton? Full robustness would require that (5.16) holds for all $\mu_0 \in M_0$. The simulations below show that this will in general not be the case for Bayes factor model selection. Yet, the same reasoning as used above implies that we still have some type of robustness in a much weaker sense, which one might call “robustness in prior expectation” relative to prior ω_0 on M_0 . Namely, we have for all $0 \leq \alpha \leq 1$:

$$\mathbb{P}_{B,0}(\exists n : \delta_{r,t(\alpha)}(X^n) = 1) \leq \alpha, \quad (5.17)$$

where $\mathbb{P}_{B,0}$ is the Bayes marginal distribution under prior ω_0 . In other words, if the beliefs of a Bayesian who adopts prior ω_0 on model \mathcal{M}_0 were accurate, then the BFMS method would still give robust p -values, independently of the stopping rule. While for a subjective Bayesian, such a weak form of robustness might perhaps still be acceptable, we will stick to the stronger definition instead, equating ‘robust hypothesis tests’ with tests satisfying (5.17) uniformly for all $\mu_0 \in M_0$.

5.5.2 AIC is not, and HQ is only asymptotically robust

The situation for AIC is quite different from that for BFMS and switching: for every function $t : (0, 1) \rightarrow \mathbb{R}_{>0}$, we have, even for every *single* $0 < \alpha < 1$, that $\delta_{AIC, t(\alpha)}$ is *not* a robust null hypothesis test for significance level α . Hence AIC cannot be transformed into a robust test in this sense. This can immediately be seen when comparing a 0-dimensional (fixed mean μ_0) with a 1-dimensional Gaussian location family \mathcal{M}_1 (extension to general multivariate exponential families is straightforward but involves tedious manipulations with the Fisher information). Evaluating the left hand side of (5.15) yields that $\delta_{AIC, t(\alpha)}$ will select the complex model if

$$\left| \sum_{i=1}^n \tilde{X}_i \right| \geq \frac{\sqrt{2n}}{t(\alpha)}, \quad (5.18)$$

where the \tilde{X}_i are variables with mean 0 and variance 1 if \mathcal{M}_0 is correct. Hence, as a consequence of the law of the iterated logarithm (see for example Van der Vaart (1998)), with probability one, infinitely many n exist such that the complex model will be favored, even though it is incorrect.

It is instructive to compare this to the HQ criterion, which, in this example, using the same notation as in (5.18), selects the complex model if

$$\left| \sum_{i=1}^n \tilde{X}_i \right| \geq \sqrt{2cn \log \log n}.$$

If $c > 1$ (the case in which HQ is strongly consistent), then this inequality will almost surely not hold for infinitely many n , as again follows from the law of the iterated logarithm. The reasoning can again be extended to other exponential families, and we find that the HQ criterion with $c > 1$ is robust to optional stopping in the crude, asymptotic sense that the probability that there exist infinitely many sample sizes such that the simple model is incorrectly rejected is zero. Yet HQ does not define a robust hypothesis test in the sense above: to get the numerically precise Type I-error bound (5.16) we would need to define $t(\alpha)$ in a model-dependent manner, which is quite complicated in all cases except the Gaussian location families where the asymptotics hold precisely. We note that the same type of asymptotic robustness holds for the BIC criterion as well.

5.5.3 Switching with singleton \mathcal{M}_0 is robust under optional stopping

The main insight of this section is simply that, just like BFMS, switching can be used as a robust null hypothesis test as well, as long as \mathcal{M}_0 is a singleton: we can view the switch distribution as a model comparison method that outputs odds ratio $r_{sw}(x^n) = p_{B,0}(x^n)/p_{sw,1}(x^n)$. Until now, we used it to select model 1 if $r_{sw}(x^n) \leq 1$. If instead we fix a significance level α and select model 1 if $r_{sw}(x^n) \leq \alpha$, then we immediately see, by applying Theorem 5.6 in the same way as for the Bayes factor case, that r_{sw} constitutes a robust null hypothesis test as long as \mathcal{M}_0 is a singleton model. Similarly – at least if the priors involved in the switch criterion are chosen independently of the stopping rule –

just like BFMS, the result $r_{\text{sw}}(x^n)$ of model comparison by switching does not depend on the ‘sampling intentions’ of the analyst, thus addressing the two most problematic issues with Neyman-Pearson testing. Yet, from a frequentist perspective, switching is preferable to BFMS, since it has substantially better power (type-II error) properties. As could already be seen from Yang’s decomposition (5.12), there is an intimate connection between Type-II error and the risk rate achieved by any model comparison method. Formally, we have the following result, a direct corollary of Theorem 5.14 of Section 5.7, which is itself a major building block of our main result Theorem 5.5 (plug in $\gamma = \alpha^{-1}$ into (5.46) to get the corollary):

Corollary 5.7. Using the same notations and under the same conditions as Theorem 5.5, for any $\alpha > 0$, there exist constants $C_1, C_2 > 0$ such that, for every CINECSI subset M'_1 of M_1 , for every sequence $\mu_1^{(1)}, \mu_1^{(2)}, \dots$ of elements of M'_1 with for all n , $\inf_{\mu_0 \in M_0} \|\mu_1^{(n)} - \mu_0\|_2^2 \geq C_1(\log \log n)/n$, we have

$$\mathbb{P}_{\mu_1^{(n)}}(r_{\text{sw}}(x^n) \geq \alpha) \leq \frac{C_2}{\log n}. \quad (5.19)$$

Hence, for any fixed significance level, the power of testing by switching goes to 1 as long as the data are sampled from a distribution $\mu_1^{(n)}$ in M_1 that is farther away from M_0 than order $(\log \log n)/n$; for BFMS, the power only goes to 1 if $\mu^{(n)}$ is farther away than order $O((\log n)/n)$.

Robustness to optional stopping (and hence ‘almost the best of three worlds’) only holds if M_0 is a singleton; if M_0 is composite, then — using again the same argument as for the Bayes factor case — we immediately see from Theorem 5.6 that the much weaker ‘prior expected robustness’ property (5.17) still holds. But, the simulations below show that full robustness does fail if μ_0 is ‘atypical’, i.e. if it resides far out in the tails of the prior ω_0 . A major question for future work is now obviously whether there exist versions of the switch criterion that give a truly robust null hypothesis test even under a composite null hypothesis M_0 . We return to this question in Section 5.6.

5.5.4 Simulation study

We now provide a simulation to illustrate the differences between AIC, BIC, HQ and the switch criterion in terms of consistency, strong consistency and robustness to optional stopping, illustrating the insights of the previous subsections. In each setting, two of the following three models are compared:

- $M_0 = \{\mathcal{N}(0, 1)\}$.
- $M_1 = \{\mathcal{N}(\mu, 1), \mu \in \mathbb{R}\}$, with a normal prior with mean zero and variance equal to 100 on μ .
- $M_2 = \{\mathcal{N}(\mu, \sigma^2), \mu \in \mathbb{R}, \sigma \in \mathbb{R}_{>0}\}$, with a normal-inverse-gamma prior: $\mu|\sigma^2 \sim \mathcal{N}(0, C\sigma^2), \sigma^2 \sim IG(\alpha, \beta)$, with $C = 100, \alpha = 1, \beta = 1$.

To illustrate standard consistency, M_1 and M_2 are considered. In the first setting, M_1 is true. $N = 1000$ data sets of length $n = 2500$ are generated from a standard normal

distribution, and AIC, BIC, HQ with $c = 1.05$ and δ_{sw} are evaluated at each sample size. The average selected model index (0 for \mathcal{M}_1 , 1 for \mathcal{M}_2) is given in Figure 5.1.

In the second setting, \mathcal{M}_2 is true. The data is generated from a normal distribution with mean 0 and a variance that is varied. For each value of σ , $N = 1000$ datasets of length $n = 2500$ are generated, and the four model selection criteria are evaluated at that sample size. The average selected model index is given in Figure 5.2.

The results are as expected. When the complex model is true, AIC is most likely to select it, at the cost of inconsistency when the simple model is true. BIC is the slowest to correctly select the complex model and the first to correctly select the simple model. HQ and δ_{sw} show intermediate behaviour, HQ being slightly more likely to select the complex model.

To illustrate strong consistency and optional stopping, three scenarios are considered:

1. \mathcal{M}_0 vs \mathcal{M}_1 , data from a standard normal distribution (“scenario 1” — Theorem 5.6 implies that switching defines a test that is robust with respect to optional stopping).
2. \mathcal{M}_1 vs \mathcal{M}_2 , data from a standard normal distribution (“scenario 2”, Theorem 5.6 does not only imply robustness, because null model is composite).
3. \mathcal{M}_1 vs \mathcal{M}_2 , data from a normal distribution with mean 35 and variance 1 (“scenario 3”, Theorem 5.6 again does not imply robustness).

We create $N = 1000$ data sets of length $n_{\text{max}} = 10000$ in each scenario. We select the complex model when δ_{sw} is larger than 20 (in terms of the robust p -value interpretation of Theorem 5.6, this corresponds to a significance level of 0.05). We estimate two probabilities at each sample size n :

- The probability that there will ever be a model index after n at which the complex model will be selected (Figure 5.3), approximated by checking whether the complex model is selected at any sample size between n and $3n_{\text{max}}$.
- The probability that there exists a model index before n at which the complex model would have been selected (Figure 5.4).

Figure 5.3 can be interpreted as a check whether strong consistency holds — if it does, then the probabilities should converge to 0 as $n \rightarrow \infty$. Van Erven et al.’s (2007) theorem implies that strong consistency holds in all three scenarios, and the graphs confirm this — even though for scenario 3, in which data comes from a $\mu \in M_0$ that is ‘atypical’ under the prior, it takes a bit longer — illustrating that strong consistency is not a uniform notion. The graph also illustrates that strong consistency can be viewed as an asymptotic, nonuniform version of robustness to optional stopping — it implies that from some sample size (which may be very large though) onwards, one will never again falsely reject no matter how long one keeps sampling.

Figure 5.4 refers to nonasymptotic optional stopping: in scenario 1, the conditions from Theorem 5.6 hold, and indeed the figure shows that the probability that the complex model is *ever* incorrectly selected even when optional stopping is used, is bounded by 0.05 (the observed bound is 0.015). In scenarios 2 and 3, the conditions from Theorem 5.6 do not hold. In scenario 2, the behaviour of the switch criterion is similar to scenario 1. However,

in scenario 3, the probability of a false rejection opportunity before sample size n is not bounded by 0.05, but quickly goes to 0.15. We clearly see that δ_{sw} is not robust to optional stopping in scenario 3.

When the simplest model is not a singleton, the choice of prior on the model parameters (in scenarios 2 and 3 on μ in \mathcal{M}_1 and on (μ, σ^2) in \mathcal{M}_2) affects the results. In both scenario 2 and 3, δ_{sw} must still satisfy the weak, prior-expected version of robustness (5.17), as we have seen in Section 5.5.3. In scenario 2, the prior is centered at the data-generating value of zero and we do observe actual robustness. In scenario 3 however, the prior is centered at zero while the data is generated with a mean of 35, 3.5 standard deviations away from the prior mean — thus μ is ‘atypical’ under the prior, and, as the figure shows, nonasymptotic robustness is violated.

5.6 Discussion and future work

In this paper we showed that switching combines near-rate optimality, consistency and, for singleton \mathcal{M}_0 , robustness to optional stopping. We end the paper by highlighting three issues which, we feel, need additional discussion: first, the desirability of consistency; second, whether there is anything ‘special’ to the switch criterion as opposed to other possible trade-offs between risk optimality and consistency; and third, the limitations of switching in its current form.

Consistency Since the desirability of consistency, in the sense of finding the smallest model containing the true distribution, is somewhat controversial, let us discuss it a bit further. The main argument against consistency is made by those adhering to Box’s maxim ‘Essentially, all models are wrong, but some are useful’ (Box and Draper, 1987). According to some, the goal of model selection should therefore not be to select a non-existing ‘true’ model, but to obtain the best predictive inference or best inference about a parameter (Burnham and Anderson, 2004; Forster, 2000). Another issue with consistency is that it is a ‘nonuniform’ notion, which in our context means that — as is indeed easy to see — it is impossible to give a bound on the probability under \mathbb{P}_μ of selecting the wrong model at sample size n that converges to 0 uniformly for all $\mu \in M$. This nonuniformity implies that consistency is of little practical consequence for post-model selection inference (Leeb and Pötscher, 2005).

As to the first argument, one can reply that there do exist situations in which a model can be correct, for example in the field of extrasensory perception (Bem, 2011), in which it seems exceedingly likely that the null model (expressing that no such thing exists) is correct; another example is genetic linkage (Gusella et al., 1983; Tsui et al., 1985). The second argument is more convincing, but only to argue that even if consistency holds, a method may not be very useful in practice. It does not contradict that consistency can sometimes be a highly desirable (but never the only highly desirable) property — we feel that this is the case whenever we are not purely interested in prediction but instead are also seeking to find out whether a certain structural relationship (e.g. dependence between variables) holds or not.

Going one step further, it seems a good idea to study model selection methods not in

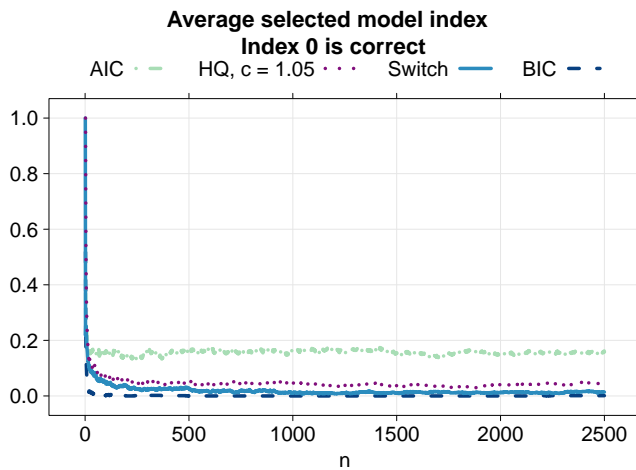


Figure 5.1: $N = 1000$ data sets of length $n = 2500$ are generated from a standard normal distribution and the criteria are evaluated at each sample size. The figure shows the average selected model index (0 for \mathcal{M}_1 , 1 for \mathcal{M}_2). The true index is 0.

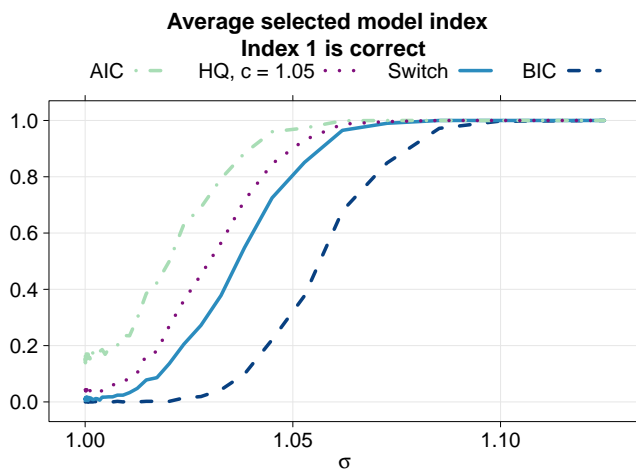


Figure 5.2: $N = 1000$ data sets of length $n = 2500$ are generated from a normal distribution with mean 0 and variance σ^2 for a range of values of σ . The criteria are evaluated at $n = 2500$. The figure shows the average selected model index (0 for \mathcal{M}_1 , 1 for \mathcal{M}_2). The true index is 1.

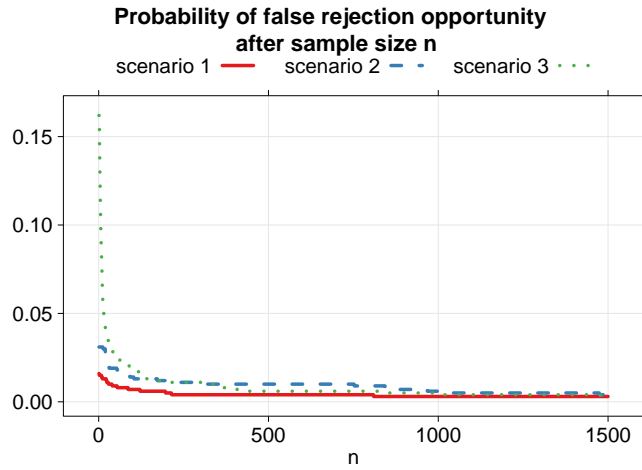


Figure 5.3: $N = 1000$ data sets of length $n_{\max} = 10000$ in each scenario, from the simple model. The complex model is selected when $\delta_{\text{sw}}(x^n) > 20$. Estimated probability that there exists a model index after n at which the complex model will be selected. Results shown up to $n = 1500$ for clarity. After $n = 1500$, the three curves are indistinguishable and all very close to zero.

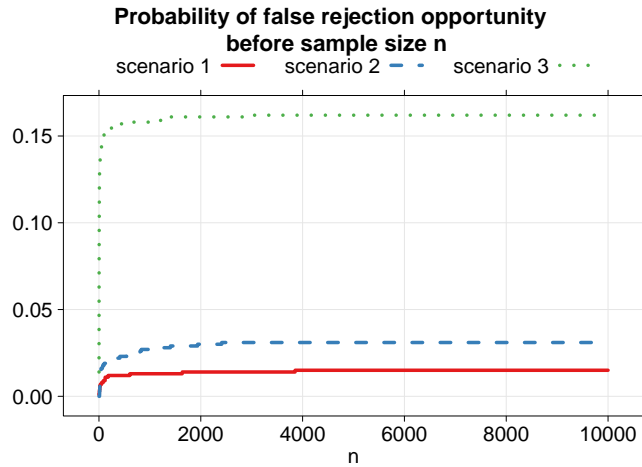


Figure 5.4: Setting as Figure 5.3. Estimated probability that there exists a model index before n at which the complex model would have been selected.

terms of the asymptotic, nonuniform notion of consistency but instead by a more tangible finite-sample analogue. For the case of just two models, Type-I and Type-II errors provide exactly this analogue — note that if both errors go to 0 as $n \rightarrow \infty$, this implies consistency. Thus, the *practical* importance of the present work, for us, is mostly that model comparison by switching defines, like Bayes, a robust null hypothesis test — providing Type-I errors irrespective of the stopping rule and thus more in line with actual practice — yet has better Type-II error behaviour, allowing the Type-II error to become small (i.e. the power to go to 1) whenever the true distribution sits at a distance of order $\sqrt{(\log \log n)/n}$ rather than $\sqrt{(\log n)/n}$, as with Bayes. We only showed robustness for singleton \mathcal{M}_0 , however, and our simulations show that it may fail for composite \mathcal{M}_0 , so *the* major goal for future work is therefore, to come up with methods that are robust to optional stopping also under composite \mathcal{M}_0 .

How special is the switch distribution? Since Yang proved that in general, the conflict between consistency and risk-optimality is not resolvable, one might argue that any model selection rule just picks some position in the spectrum of behaviours of consistency vs. risk-optimality. For example, one might have a modified HQ criterion which picks \mathcal{M}_1 if, using the same setup and notation as in (5.18),

$$\left| \sum_{i=1}^n \tilde{X}_i \right| \geq \sqrt{n \log \log \log n}. \quad (5.20)$$

By the central limit theorem, such a method will be consistent, yet when combined with an efficient estimator will achieve the minimax estimation rate up to a $\log \log \log n$ factor, improving on the switch criterion by an additional logarithm. Note however that both the switch distribution and HQ (with $c > 1$) achieve *strong* consistency. The meaning of strong consistency is illustrated in Figure 5.3 above: it means that, from some n onward, the wrong model will never be selected any more, no matter how long one keeps sampling. It is easy to see from the law of the iterated logarithm that any strongly consistent method can have rate no faster than order $(\log \log n)/n$ — in particular, (5.20) is not strongly consistent. Thus, in this sense both switching and HQ do take a special place in the consistency vs. risk-optimality spectrum as obtaining the fastest rates compatible with strong consistency, which may be viewed as asymptotic robustness to optional stopping. While this may mostly be of theoretical interest, the switch distribution also takes a special place in terms of its nonasymptotic robustness to optional stopping: again, the law of the iterated logarithm implies that any model comparison method that defines a robust hypothesis test cannot achieve estimation rate better than order $(\log \log n)/n$. Again, the main open question here is whether one can modify it so that robustness for composite \mathcal{M}_0 is achieved as well.

Future work — limitations of the switch distribution and our results Whereas the results in this paper all apply to the original switch distribution as defined by Van Erven et al. (2007) and a simplification thereof, for full robustness to optional stopping with composite \mathcal{M}_0 , some substantial changes have to be made, as suggested by the results in Figure 5.4. Initial research suggests that such a modification of the switch distribution

might indeed be constructed, based on techniques in Ramdas and Balasubramani (2015); whereas, compared to Bayes factor testing, in the current switch criterion, $p_{B,1}$ is modified to another distribution and $p_{B,0}$ can remain the same, in this new version we would also have to change $p_{B,0}$ — the resulting distribution would not have a Bayesian interpretation any more. While this work is still under development, to avoid the nonrobustness seen in Figure 5.4 as much as possible, for the time being we recommend using flat priors (but in this case, not completely flat - Jeffreys' prior on μ is improper, in which case Theorem 5.6 holds in none of the scenarios and simulations — not reported here — show that optional stopping robustness is violated).

Another limitation lies not in the switch distribution, but in our results: these are restricted to two nested exponential family models. It would be interesting to extend them to more than two models — highlighting the distinction between model selection and testing — and going beyond exponential families. We are hopeful that switching still behaves well in such contexts — we note that the risk rate convergence results of Van Erven et al. (2012) were for countable, possibly infinite collections of completely general models — but they invariably dealt with the cumulative risk. While all our experiments suggest that small cumulative risk usually goes together with small instantaneous risk, formal analysis of the switch criterion's instantaneous risk is far more difficult, and the present paper heavily relies on sufficiency to do so — so extension of our results beyond exponential families would be difficult.

Before doing so, we would prefer to modify the switch distribution further, since the present version has a drawback when used in nonsequential settings: the precise results it gives are dependent on the order of the data, even if all the models under consideration are i.i.d. Thus, it would be interesting and challenging to design an alternative, order-independent method that, like the switch distribution, is strongly consistent, near rate- and power-optimal, and is robust to optional stopping under composite \mathcal{M}_0 . Such a method would essentially truly achieve the best of the three worlds we considered in this paper — and this is the method we aim for in our future research.

Acknowledgements

The central result of this paper, Theorem 5.5, already appeared in the Master's Thesis (Van der Pas, 2013) for the special case where $m_1 = 1$ and $m_0 = 0$, but the proof supplied there contained an error. We are grateful to Tim van Erven for pointing this out to us.

5.7 Proofs

In this appendix, we start by listing some well-known properties of exponential families which we will repeatedly use in the proofs. Then, in Section 5.7.4, we provide a sequence of technical lemmata that lead up to the proof of our main result, Theorem 5.5. Finally, in Section 5.7.5, we compare the switch distribution and criterion as defined here to the original switch distribution and criterion of Van Erven et al. (2012).

Additional notation Our results will often involve displays involving several constants. The following abbreviation proves useful: when we write ‘for positive constants \vec{c} , we have ...’, we mean that there exist some $(c_1, \dots, c_N) \in \mathbb{R}^N$, with $c_1, \dots, c_N > 0$, such that ... holds; here N is left unspecified but it will always be clear from the application what N is. Further, for positive constants $\vec{b} = (b_1, b_2, b_3)$, we define $\text{small}_{\vec{b}}(n)$ as

$$\text{small}_{\vec{b}}(n) = \begin{cases} 1 & \text{if } n < b_1 \\ b_2 e^{-b_3 n} & \text{if } n \geq b_1, \end{cases}$$

and we frequently use the following fact. Suppose that $\mathcal{E}_1, \mathcal{E}_2, \dots$ is a sequence of events such that $\mathbb{P}(\mathcal{E}_n) \leq \text{small}_{\vec{b}}(n)$. Then we also have, for any event \mathcal{A} , and for all n ,

$$\mathbb{P}(\mathcal{A}, \mathcal{E}_n^c) \geq \mathbb{P}(\mathcal{A}) - \text{small}_{\vec{b}}(n), \quad (5.21)$$

as is immediate from $\mathbb{P}(\mathcal{A}, \mathcal{E}_n^c) = \mathbb{P}(\mathcal{A}) - \mathbb{P}(\mathcal{A}, \mathcal{E}_n) \geq \mathbb{P}(\mathcal{A}) - \mathbb{P}(\mathcal{E}_n)$.

The components of a vector $\mu \in \mathbb{R}^n$ are given by $(\mu_1, \mu_2, \dots, \mu_n)$. If the vector already has an index, we add a comma, for example $\mu_1 = (\mu_{1,1}, \mu_{1,2}, \dots, \mu_{1,n})$. A sequence of vectors is denoted by $\mu^{(1)}, \mu^{(2)}, \dots$.

5.7.1 Definitions concerning and properties of exponential families

The following definitions and properties can all be found in the standard reference (Barndorff-Nielsen, 1978) and, less formally, in (Grünwald, 2007, Chapters 18 and 19).

A k -dimensional exponential family is a set of distributions on \mathcal{X} , which we invariably represent by the corresponding set of densities $\{p_\theta \mid \theta \in \Theta\}$, where $\Theta \subset \mathbb{R}^k$, such that any member p_θ can be written as

$$p_\theta(x) = \frac{1}{z(\theta)} e^{\theta^T \varphi(x)} r(x) = e^{\theta^T \varphi(x) - \psi(\theta)} r(x), \quad (5.22)$$

where $\varphi(x) = (\varphi_1(x), \dots, \varphi_k(x))$ is a *sufficient statistic*, r is a non-negative function called the *carrier*, z the *partition function* and $\psi(\theta) = \log z(\theta)$. We assume the representation (5.7.1) to be *minimal*, meaning that the components of $\varphi(x)$ are linearly independent.

The parameterization in (5.22) is referred to as the *canonical* or *natural parameterization*; we only consider families for which the set Θ is open and connected. Every exponential family can alternatively be parameterized in terms of its *mean-value parameterization*, where the family is parameterized by the mean $\mu = \mathbb{E}_\theta[\varphi(X)]$, with μ taking values in $M \subset \mathbb{R}^k$, where μ as a function of θ is smooth and strictly increasing; as a consequence, the set M of mean-value parameters corresponding to an open and connected set Θ is itself also open and connected. Whenever for data x_1, \dots, x_n , we have $\frac{1}{n} \sum_{i=1}^n \varphi(x_i) \in M$, then the maximum likelihood is uniquely achieved by the μ that is itself equal to this value,

$$\widehat{\mu}(x^n) = \frac{1}{n} \sum_{i=1}^n \varphi(x_i). \quad (5.23)$$

We thus define the maximum likelihood estimator (MLE) to be equal to (5.23) whenever $\frac{1}{n} \sum_{i=1}^n \varphi(x_i) \in M$. Since the result below which directly involves the MLE (Lemma 5.11)

does not depend on its value for x^n with $\frac{1}{n} \sum_{i=1}^n \varphi(x_i) \notin M$, we can leave $\widehat{\mu}(x^n)$ undefined for such values. However, if we want to use the MLE as a ‘sufficiently efficient’ estimator as used in the statement of Theorem 5.5, we need to define $\widehat{\mu}(x^n)$ for such values in such a way that (5.13) is satisfied, as illustrated in Example 1.

A standard property of exponential families says that, for any $\mu \in M$, any distribution \mathbb{Q} on \mathcal{X} with $\mathbb{E}_{\mathcal{X} \sim \mathbb{Q}}[\varphi(X)] = \mu$, any $\mu' \in M$, we have

$$\mathbb{E}_{\mathcal{X} \sim \mathbb{Q}} \left[\log \frac{p_\mu(X)}{p_{\mu'}(X)} \right] = \mathbb{E}_{\mathcal{X} \sim \mathbb{P}_\mu} \left[\log \frac{p_\mu(X)}{p_{\mu'}(X)} \right] = D(\mu \| \mu'), \quad (5.24)$$

the final equality being just the definition of $D(\cdot \| \cdot)$. Now fix an arbitrary sample x^n . By taking \mathbb{Q} to be the empirical distribution on \mathcal{X} corresponding to sample x^n , it follows from (5.24) that if $\widehat{\mu}(x^n) \in M$ then also the following relationship holds for any $\mu' \in M$:

$$\frac{1}{n} \log \frac{p_{\widehat{\mu}(x^n)}(x^n)}{p_{\mu'}(x^n)} = D(\widehat{\mu}(x^n) \| \mu'). \quad (5.25)$$

(5.24) and (5.25) are a direct consequence of the sufficiency of $\widehat{\mu}_1(X^n)$, and folklore among information theorists. For a proof of (5.24) and more details on (5.25), see e.g. (Grünwald, 2007, Chapter 19), who calls this the *robustness property* of the KL divergence for exponential families.

We are now in a position to prove Proposition 5.2, which we repeat for convenience.

Proposition 5.2 Let M , a product of open intervals, be the mean-value parameter space of an exponential family, and let M' be a CINECSI subset of M . Then there exist positive constants \vec{c} such that for all $\mu, \mu' \in M'$,

$$c_1 \|\mu' - \mu\|_2^2 \leq c_2 \cdot d_{ST}(\mu' \| \mu) \leq d_{H^2}(\mu', \mu) \leq d_R(\mu', \mu) \leq D(\mu' \| \mu) \leq c_3 \|\mu' - \mu\|_2^2. \quad (5.26)$$

and for all $\mu' \in M', \mu \in M$ (i.e. μ is now not restricted to lie in M'),

$$d_{H^2}(\mu', \mu) \leq c_4 \|\mu' - \mu\|_2^2 \leq c_5 \cdot d_{ST}(\mu' \| \mu) \leq c_6 \|\mu' - \mu\|_2^2. \quad (5.27)$$

Proof. We start with (5.26). The third and fourth inequality are immediate by using $-\log x \geq 1 - x$ and Jensen’s inequality, respectively. From standard properties of Fisher information for exponential families (Barndorff-Nielsen, 1978) we have that, for any CINECSI (hence compact and bounded away from the boundaries of M) subset M' of M , there exists positive \vec{C} with

$$0 < C_1 = \inf_{\mu \in M'} \det I(\mu) < \sup_{\mu \in M'} \det I(\mu) = C_2 < \infty, \quad (5.28)$$

from which we infer that for all $\mu' \in M', \mu, \mu'' \in \mathbb{R}^m$,

$$C_3 \|\mu - \mu''\|_2^2 \leq (\mu - \mu'')^T I(\mu') (\mu - \mu'') \leq C_4 \|\mu - \mu''\|_2^2, \quad (5.29)$$

for some $0 < C_3 \leq C_4 < \infty$. Using (5.29), the first inequality is immediate, and the final inequality follows straightforwardly from a second-order Taylor approximation of KL divergence as in (Grünwald, 2007, Chapter 4). It only remains to establish the second

inequality. Now, since M' is CINECSI and hence compact the fifth (rightmost) inequality implies that there is a $C_5 < \infty$ such that $\sup_{\mu, \mu' \in M'} D(\mu' \| \mu) < C_5$ and hence, via the fourth inequality, that $\sup_{\mu, \mu' \in M'} d_R(\mu', \mu) < C_5$. Equality (5.5) now implies that there is a C_6 such that

$$\sup_{\mu, \mu' \in M'} d_R(\mu', \mu) / d_{H^2}(\mu', \mu) < C_6. \quad (5.30)$$

Using again (5.28), a second order Taylor approximation as in Van Erven and Harremoës (2014) now gives that for some constant $C_7 > 0$, $\|\mu - \mu'\|_2^2 \leq C_7 d_R(\mu', \mu)$ for all $\mu, \mu' \in M'$. The first result, (5.26), now follows upon combining this with (5.30).

As to (5.27), the second and third inequality are immediate from (5.29). For the first inequality, note that, since M' is CINECSI and we assume M to be a product of open intervals, there must exist another CINECSI subset M'' of M strictly containing M' such that $\inf_{\mu' \in M', \mu \in M \setminus M''} \|\mu' - \mu\|_2^2 = \delta$ for some $\delta > 0$. We now distinguish between μ in (5.27) being an element of (a) M'' or (b) $M \setminus M''$. For case (a) (5.26), with M'' in the role of M' , gives that there is a constant C_8 such that for all $\mu \in M''$, $d_{H^2}(\mu', \mu) \leq C_8 \|\mu' - \mu\|_2^2$. For case (b), $\mu \in M \setminus M''$, we have $\|\mu' - \mu\|_2^2 \geq \delta$ and, using that squared Hellinger distance for any pair of distributions is bounded by 2, we have $d_{H^2}(\mu', \mu) \leq (2/\delta) \|\mu' - \mu\|_2^2$. Thus, by taking $c_4 = \max\{C_8, 2/\delta\}$, case (a) and (b) together establish the first inequality in (5.27). \square

5.7.2 Preparation for proof of main result: results on large deviations

Let \mathcal{M}_1 and M_1 be as in Theorem 5.5. For the following result, Lemma 5.8, we set $\widehat{\mu}'_1(X^n) := n^{-1} \sum \varphi(X_i)$, so that $\widehat{\mu}'_1(X^n) = \widehat{\mu}_1(X^n)$ whenever $n^{-1} \sum \varphi(X_i) \in M_1$. It is essentially a multidimensional extension of a standard information-theoretic result, with KL divergence replaced by squared error loss. The result states the following: whenever \mathcal{M}_1 is a single-parameter exponential family (that is, $m_1 = 1$), then for any $\mu \in M_1$, all $a, a' > 0$ with $\mu + a \in M_1$, $\mu - a' \in M_1$,

$$\mathbb{P}_\mu(\widehat{\mu}'_1(X^n) \geq \mu + a) \leq e^{-nD(\mu+a\|\mu)}; \quad \mathbb{P}_\mu(\widehat{\mu}'_1(X^n) \leq \mu - a') \leq e^{-nD(\mu-a'\|\mu)}. \quad (5.31)$$

For a simple proof, see (Grünwald, 2007, Section 19.4.2); for discussion see (Csiszár, 1984) — the latter reference gives a multidimensional extension of (5.31) but of a very different kind than Lemma 5.8 below. To prepare for the lemma, let \mathcal{M}_1 and M_1 be as in Theorem 5.5 and, for any $\mu \in M_1$ and any $\vec{a}, \vec{b} \in \mathbb{R}_{>0}^{m_1}$, define the ℓ_∞ -rectangle $R_\infty(\mu, \vec{a}, \vec{b}) = \{\mu' \in \mathbb{R}^{m_1} : \forall j = 1, \dots, m_1, -b_j \leq \mu'_j - \mu_j \leq a_j\}$.

Lemma 5.8. Let \mathcal{M}_1 and M_1 be as in Theorem 5.5 and fix an arbitrary CINECSI subset M'_1 of M_1 . Then there is a $c > 0$ (depending on M'_1) such that, for all $\mu \in M_1$, all n , all $\vec{a}, \vec{b} \in \mathbb{R}_{>0}^{m_1}$ such that $R_\infty(\mu, \vec{a}, \vec{b}) \subset M'_1$,

$$\mathbb{P}_\mu(\widehat{\mu}'_1(X^n) \notin R_\infty(\mu, \vec{a}, \vec{b})) \leq 2m_1 e^{-nc \cdot (\min_j \min\{a_j, b_j\})^2}. \quad (5.32)$$

Proof. For $j = 1, \dots, m_1$, $d \in \mathbb{R}$, let \vec{e}_j represent the j th standard basis vector, such that $\mu + d\vec{e}_j = (\mu_1, \dots, \mu_{j-1}, \mu_j + d, \mu_{j+1}, \dots, \mu_{m_1})$, and let $D_{\mu+d\vec{e}_j} := D(\mu + d\vec{e}_j \| \mu)$. We now have that

there exist constants $c_{a,1}, \dots, c_{a,m_1}, c_{b,1}, \dots, c_{b,m_1} > 0$ such that for $c := \min\{c_{a,1}, \dots, c_{a,m_1}, c_{b,1}, \dots, c_{b,m_1}\}$, all n ,

$$\begin{aligned} \mathbb{P}_\mu(\widehat{\mu}_1(X^n) \notin R_\infty(\mu, \vec{a}, \vec{b})) &\leq \sum_{j=1}^{m_1} \mathbb{P}_\mu(\widehat{\mu}_{1,j}(X_n) \geq \mu_j + a_j) + \sum_{j=1}^{m_1} \mathbb{P}_\mu(\widehat{\mu}_{1,j}(X^n) \leq \mu_j - b_j) \\ &\leq \sum_{j=1}^{m_1} \left(e^{-nD_{\mu+a_j\vec{e}_j}} + e^{-nD_{\mu-b_j\vec{e}_j}} \right) \leq \sum_{j=1}^{m_1} \left(e^{-nc_{a,j}a_j^2} + e^{-nc_{b,j}b_j^2} \right) \\ &\leq 2m_1 e^{-nc \cdot (\min_j \min\{a_j, b_j\})^2}, \end{aligned}$$

Here the first inequality follows from the union bound, and the second follows by applying, for each of the $2m_1$ terms, (5.31) above to the one-dimensional exponential sub-family $\{p_\mu \mid \mu \in M_1 \cap \{\mu : \mu = \mu + d\vec{e}_j \text{ for some } d \in \mathbb{R}\}\}$. The third follows by Proposition 5.2 together with the equivalence of the ℓ_2 and sup norms on \mathbb{R}^{m_1} , and the final inequality is immediate. \square

Lemma 5.9. Under conditions and notations as in Theorem 5.5, let μ, μ' be elements of M_1 and suppose $X^N = (X_{n_1}, \dots, X_{n_2})$ is a sequence of i.i.d. observations of length N from p_μ . Then, for any $A \in \mathbb{R}$:

$$\mathbb{P}_\mu \left(\log \frac{p_\mu(X^N)}{p_{\mu'}(X^N)} < A \right) \leq e^{\frac{1}{2}A} e^{-\frac{N}{2} d_R(\mu', \mu)}. \quad (5.33)$$

Proof. For any A , by Markov's inequality:

$$\begin{aligned} \mathbb{P}_\mu \left(\log \frac{p_\mu(X^N)}{p_{\mu'}(X^N)} < A \right) &= \mathbb{P}_\mu \left(\left(\frac{p_{\mu'}(X^N)}{p_\mu(X^N)} \right)^{\frac{1}{2}} > e^{-\frac{1}{2}A} \right) \leq e^{\frac{1}{2}A} \mathbb{E}_\mu \left[\left(\frac{p_{\mu'}(X^N)}{p_\mu(X^N)} \right)^{\frac{1}{2}} \right] \\ &= e^{\frac{1}{2}A} \left(\mathbb{E}_\mu \left[\left(\frac{p_{\mu'}(X_{n_1})}{p_\mu(X_{n_1})} \right)^{\frac{1}{2}} \right] \right)^N = e^{\frac{1}{2}A} e^{\log \left(\mathbb{E}_\mu \left[\left(\frac{p_{\mu'}(X_{n_1})}{p_\mu(X_{n_1})} \right)^{\frac{1}{2}} \right] \right)^N} \\ &= e^{\frac{1}{2}A} e^{-\frac{N}{2} \left(-\frac{1}{1-1/2} \log \mathbb{E}_\mu \left[\left(\frac{p_{\mu'}(X_{n_1})}{p_\mu(X_{n_1})} \right)^{\frac{1}{2}} \right] \right)} = e^{\frac{1}{2}A} e^{-\frac{N}{2} d_R(\mu, \mu')}. \end{aligned} \quad (5.34)$$

\square

Proposition 5.10. Let $\mathcal{M}_0, \mathcal{M}_1, M_0, M_1$ be as in Theorem 5.5 and let M'_1 be a CINECSI subset of M_1 . Then there exists another, larger, CINECSI subset M''_1 of M_1 and positive constants \vec{b} such that M'_1 is itself a CINECSI subset of M''_1 and for both $j \in \{0, 1\}$, the ML estimator $\widehat{\mu}_j(x^n)$ satisfies

$$\sup_{\mu \in M'_1} P_\mu(\widehat{\mu}_j(X^n) \notin M''_1) \leq \text{small}_{\vec{b}}(n).$$

Proof. M_1 can be written as in (5.10), and hence we can define a set

$$M_1'' = [\zeta_{1,1}^*, \eta_{1,1}^*] \times \dots \times [\zeta_{1,m_1}^*, \eta_{1,m_1}^*]$$

for values $\zeta_{1,j}^*, \eta_{1,j}^* \in \mathbb{R}$ such that M_1'' is a CINECSI subset of M_1 . Since M_1' is connected with compact closure in interior of M_1 and M_1'' is a subset of M_1 , we can choose the $\zeta_{1,j}^*, \eta_{1,j}^* \in \mathbb{R}$ such that M_1' is itself a CINECSI subset of M_1'' . Since M_1' is connected and its closure is in the interior of M_1'' which is itself compact, it follows that there is some $\delta > 0$ such that, for all $\mu_1' \in M_1', \mu_1'' \notin M_1'',$ all $j \in \{1, \dots, m_1\}$, it holds $|\mu_{1,j}' - \mu_{1,j}''| > \delta$. It now follows from Lemma 5.8, applied with \vec{a} chosen such that $R_\infty(\mu', \vec{a}) = M_1'',$ that for every $\mu' \in M_1',$ all $n,$

$$\mathbb{P}_{\mu'}(\widehat{\mu}_1(X^n) \notin M_1'') \leq C_1 e^{-nC_2\delta^2}$$

for some constants C_1, C_2 . Here we used that by construction, each entry of \vec{a} must be at least as large as δ . Since $\widehat{\mu}_{1,j}(x^n)$ and $\widehat{\mu}_{0,j}(x^n)$ coincide for $0 < j \leq m_0$ and $\widehat{\mu}_{0,j}(x^n)$ is constant for $m_0 < j \leq m_1,$ the result follows for $\widehat{\mu}_0(x^n)$ as well. \square

5.7.3 Preparation for proof of main result: results on Bayes factor model selection

Lemma 5.11. Let $\mathcal{M}_0, \mathcal{M}_1, M_0, M_1$ be as in Theorem 5.5 and let, for $j \in \{0, 1\}, M_j'$ be a CINECSI subset of M_j . For both $j \in \{0, 1\},$ there exist positive constants \vec{c}, \vec{b} such that for all $\mu_1 \in M_1',$

$$c_1 \leq n^{-m_j/2} \cdot \frac{p_{\widehat{\mu}_j(X^n)}(X^n)}{p_{B,j}(X^n)} \leq c_2, \quad (5.35)$$

with \mathbb{P}_{μ_1} -probability at least $1 - \text{small}_B(n).$

Proof. For a Bayesian marginal distribution p_B defined relative to m -dimensional exponential family \mathcal{M} given in its mean-value parameterization $M,$ with a prior $\omega(\cdot)$ that is continuous and strictly positive on $M,$ we have as a consequence of the familiar Laplace approximation of the Bayesian marginal distribution of exponential families as in e.g. (Kass and Raftery, 1995),

$$p_B(x^n) \sim \left(\frac{n}{2\pi}\right)^{-m/2} \cdot \frac{\omega(\widehat{\mu}(x^n))}{\sqrt{\det I(\widehat{\mu}(x^n))}} p_{\widehat{\mu}(x^n)}(x^n).$$

As shown in Theorem 8.1 in (Grünwald, 2007), this statement holds uniformly for all sequences x^n with ML estimators in any fixed CINECSI subset M' of M . By compactness of $M',$ and by positive definiteness and continuity of Fisher information for exponential families, the quantity $\omega(\widehat{\mu})/\sqrt{\det I(\widehat{\mu})}$ will be bounded away from zero and infinity on such sequences, and, applying the result to both the families \mathcal{M}_0 and \mathcal{M}_1 it follows that there exist $c_1, c_2 > 0$ such that for all n larger than some $n_0,$ uniformly for all sequences x^n with $\widehat{\mu}_j(x^n) \in M_j',$ we have:

$$c_1 \leq n^{-m_j/2} \cdot \frac{p_{\widehat{\mu}_j(X^n)}(x^n)}{p_{B,j}(x^n)} \leq c_2. \quad (5.36)$$

The result now follows by combining this statement with Proposition 5.10. \square

Lemma 5.12. Let $\mathcal{M}_0, \mathcal{M}_1, M_0, M_1$ and the Bayesian marginal distribution $p_{B,0}$ be as in Theorem 5.5. Let M'_1 be a CINECSI subset of M_1 . Then there exist positive constants \vec{c} and \vec{b} such that for all n , all $\mu_1 \in M'_1$, all $A \in \mathbb{R}$,

$$\mathbb{P}_{\mu_1} \left(\log \frac{p_{B,1}(X^n)}{p_{B,0}(X^n)} < A \right) \leq n^{m_1/2} \cdot c_1 \cdot e^{\frac{1}{2}c_2 A} e^{-\frac{n}{2}c_3 \|\mu_1 - \mu_0\|_2^2} + \text{small}_{\vec{b}}(n),$$

where for each $\mu_1, \mu_0 = \Pi_0(\mu_1)$ as in (5.11).

Proof. Fix constants C_1, C_2 such that they are smaller and larger respectively than the constants c_1, c_2 from Lemma 5.11 and define

$$\mathcal{E}_n = \left\{ X^n : C_1 \leq n^{-m_1/2} \frac{p_{\widehat{\mu}_1}(X^n)}{p_{B,1}(X^n)} \leq C_2 \right\}.$$

Using Lemma 5.11, we have that there exists positive \vec{b} such that for all $A \in \mathbb{R}$,

$$\begin{aligned} & \mathbb{P}_{\mu_1} \left(\log \frac{p_{B,1}(X^n)}{p_{B,0}(X^n)} < A \right) \\ &= \mathbb{P}_{\mu_1} \left(\log \frac{p_{B,1}(X^n)}{p_{B,0}(X^n)} < A, \mathcal{E}_n \right) + \mathbb{P}_{\mu_1} \left(\log \frac{p_{B,1}(X^n)}{p_{B,0}(X^n)} < A, \mathcal{E}_n^c \right) \\ &\leq \mathbb{P}_{\mu_1} \left(\log \frac{C_2^{-1} n^{-m_1/2} p_{\widehat{\mu}_1}(X^n)}{p_{B,0}(X^n)} < A, \mathcal{E}_n \right) + \text{small}_{\vec{b}}(n) \\ &\leq \mathbb{P}_{\mu_1} \left(\log \frac{C_2^{-1} n^{-m_1/2} p_{\mu_1}(X^n)}{p_{B,0}(X^n)} < A \right) + \text{small}_{\vec{b}}(n) \\ &= \mathbb{P}_{\mu_1} \left(\log \frac{p_{\mu_1}(X^n)}{p_{B,0}(X^n)} < A + \log C_2 n^{m_1/2} \right) + \text{small}_{\vec{b}}(n). \end{aligned} \quad (5.37)$$

To bound this probability further, we need to relate $p_{B,0}$ to $p_{B',0}$, the Bayesian marginal likelihood under model M_0 under a prior with support restricted to a compact set M'_0 . To define M'_0 , note first that there must exist a CINECSI subset, say M''_1 , of M_1 such that M'_1 is itself a CINECSI subset of M''_1 . Take any such M''_1 and let M'_0 be the closure of $M''_1 \cap M_0$. Given ω , the prior density on $\Pi'(M_0)$ used in the definition of $p_{B,0}$, define $\omega'(v) = \omega(v) / \int_{v \in \Pi'(M'_0)} \omega(v) dv$ as the prior density restricted to and normalized on $\Pi'(M'_0)$ and let $p_{B',0}$ be the corresponding Bayesian marginal density on X^n .

To continue bounding (5.37), define

$$\mathcal{E}'_n = \left\{ X^n : C_3 \leq n^{-m_0/2} \frac{p_{\widehat{\mu}_0}(X^n)}{p_{B,0}(X^n)} \leq C_4 \text{ and } C_3 \leq n^{-m_0/2} \frac{p_{\widehat{\mu}_0}(X^n)}{p_{B',0}(X^n)} \leq C_4 \right\},$$

with C_3 and C_4 smaller and larger respectively than the constants c_1 and c_2 resulting from Lemma 5.11 (note that Lemma 5.11 can be applied to $p_{B',0}$ as well, by taking M_0 in that lemma to be the interior of M'_0 as defined here). Set $C_5 > C_4/C_3$, and note that for any $A_1 \in \mathbb{R}$, abbreviating $\mathbb{P}_{\mu_1} \left(\log \frac{p_{\mu_1}(X^n)}{C_5 p_{B',0}(X^n)} < A_1 \right)$ to p^* , we have

$$\mathbb{P}_{\mu_1} \left(\log \frac{p_{\mu_1}(X^n)}{p_{B,0}(X^n)} < A_1 \right)$$

$$\begin{aligned}
&= \mathbb{P}_{\mu_1} \left(\log \frac{p_{\mu_1}(X^n)}{p_{B,0}(X^n)} < A_1, \frac{p_{B_0}(X^n)}{p_{B',0}(X^n)} < C_5 \right) + \mathbb{P}_{\mu_1} \left(\log \frac{p_{\mu_1}(X^n)}{p_{B,0}(X^n)} < A_1, \frac{p_{B_0}(X^n)}{p_{B',0}(X^n)} \geq C_5 \right) \\
&\leq \mathbb{P}_{\mu_1} \left(\log \frac{p_{\mu_1}(X^n)}{C_5 p_{B',0}(X^n)} < A_1 \right) + \mathbb{P}_{\mu_1} (p_{B,0}(X^n) \geq C_5 p_{B',0}(X^n)) \\
&= p^* + \mathbb{P}_{\mu_1} (p_{B,0}(X^n) \geq C_5 p_{B',0}(X^n)) \\
&\leq p^* + \mathbb{P}_{\mu_1} \left(\frac{p_{B,0}(X^n)}{p_{B',0}(X^n)} \geq C_5, \mathcal{E}'_n \right) + \mathbb{P}_{\mu_1} \left(\frac{p_{B,0}(X^n)}{p_{B',0}(X^n)} \geq C_5, (\mathcal{E}'_n)^c \right) \\
&\leq p^* + 0 + \text{small}_{\tilde{c}}(n).
\end{aligned} \tag{5.38}$$

Now it only remains to bound p^* . To this end, let

$$C_6 := \int_{v \in \Pi'(M'_0)} \sqrt{\omega(v)} dv. \tag{5.39}$$

Since M'_0 has compact closure in the interior of M_0 and we are assuming that ω has full support on M_0 , we have that $C_6 < \infty$.

Now using Markov's inequality as in the proof of Lemma 5.9, that is, the first line of (5.34) with $p_{B',0}$ in the role of $p_{\mu'}$, gives, for any $A_2 \in \mathbb{R}$,

$$\mathbb{P}_{\mu_1} \left(\log \frac{p_{\mu_1}(X^n)}{p_{B',0}(X^n)} < A_2 \right) \leq e^{\frac{1}{2}A_2} \mathbb{E}_{\mu_1} \left[\left(\frac{p_{B',0}(X^n)}{p_{\mu_1}(X^n)} \right)^{\frac{1}{2}} \right]. \tag{5.40}$$

The expectation on the right can be further bounded, defining $\omega'' = \sqrt{\omega}/C_6$ and noting that ω'' is a probability density, as

$$\begin{aligned}
\mathbb{E}_{\mu_1} \left[\left(\frac{p_{B',0}(X^n)}{p_{\mu_1}(X^n)} \right)^{\frac{1}{2}} \right] &\leq \mathbb{E}_{\mu_1} \left[\left(\frac{\int_{v \in \Pi'(M'_0)} \omega(v)^{1/2} p_v(X^n)^{1/2} dv}{p_{\mu_1}(X^n)^{1/2}} \right) \right] \\
&= C_6 \cdot \mathbb{E}_{\mu \sim \omega''} \mathbb{E}_{\mu_1} \left[\left(\frac{p_{\mu}(X^n)}{p_{\mu_1}(X^n)} \right)^{\frac{1}{2}} \right] \leq C_6 \cdot \mathbb{E}_{\mu_1} \left[\left(\frac{p_{\mu^\circ}(X^n)}{p_{\mu_1}(X^n)} \right)^{\frac{1}{2}} \right],
\end{aligned}$$

where $\mu^\circ \in M'_0$ achieves the supremum of $E_{\mu_1} \left[\left(\frac{p_{\mu^\circ}(X^n)}{p_{\mu_1}(X^n)} \right)^{\frac{1}{2}} \right]$ within M'_0 . By compactness of M'_0 and continuity, this supremum is achieved. The final term can be rewritten, following the same steps as in the second and third line of (5.34), as

$$\mathbb{E}_{\mu_1} \left[\left(\frac{p_{\mu^\circ}(X^n)}{p_{\mu_1}(X^n)} \right)^{\frac{1}{2}} \right] = e^{-\frac{n}{2} d_R(\mu_1, \mu^\circ)}. \tag{5.41}$$

Since M'_0 and M'_1 are both CINECSI, it now follows from Proposition 5.2 that for some fixed $C_7 > 0$,

$$d_R(\mu_1, \mu^\circ) \geq C_7 \|\mu_1 - \mu^\circ\|_2^2 \geq C_7 \|\mu_1 - \mu_0\|_2^2, \tag{5.42}$$

where the latter inequality follows by the definition of $\mu_0 = \Pi_0(\mu_1)$, see the explanation below (5.11). Combining (5.40), (5.41) and (5.42), we have thus shown that for all n , all

$\mu_1 \in M_1$, all $A_2 \in \mathbb{R}$,

$$\mathbb{P}_{\mu_1} \left(\log \frac{p_{\mu_1}(X^n)}{p_{B',0}(X^n)} < A_2 \right) \leq C_6 e^{\frac{1}{2} A''} e^{-\frac{n}{2} C_7 \|\mu_1 - \mu_0\|_2^2}. \quad (5.43)$$

The result now follows by combining (5.37), (5.38) and (5.43). \square

5.7.4 Proof of main result, Theorem 5.5

Proof Idea The proof is based on analyzing what happens if X_1, X_2, \dots, X_n are sampled from $p_{\mu_1^{(n)}}$, where $\mu_1^{(1)}, \mu_1^{(2)}, \dots$ are a sequence of parameters in M'_1 . We consider three regimes, depending on how fast (if at all) $\mu_1^{(n)}$ converges to $\mu_0^{(n)}$ as $n \rightarrow \infty$. Here $\mu_0^{(n)} = \Pi_0(\mu_1^{(n)})$ is the projection of $\mu_1^{(n)}$ onto M_0 , i.e. the distribution in M_0 defined, for each n , as in (5.11), with μ_1 and μ_0 in the role of $\mu_1^{(n)}$ and $\mu_0^{(n)}$, respectively. Our regimes are defined in terms of the function f given by

$$f(n) := \frac{\|\mu_1^{(n)} - \mu_0^{(n)}\|_2^2}{\frac{\log \log n}{n}} = \frac{n \cdot \|\mu_1^{(n)} - \mu_0^{(n)}\|_2^2}{\log \log n}, \quad (5.44)$$

which indicates how fast $d_{SQ}(\mu_1^{(n)}, \mu_0^{(n)})$ grows relative to the best possible rate $(\log \log n)/n$. We fix appropriate constants Γ_1 and Γ_2 , and we distinguish, for all n with $\Gamma_2 \log n \geq \Gamma_1$, the cases:

$$f(n) \in \begin{cases} [0, \Gamma_1] & \text{Case 1} \\ [\Gamma_1, \Gamma_2 \log n] & \text{Case 2 (Theorem 5.14)} \\ [\Gamma_2 \log n, \infty] & \text{Case 3 (Theorem 5.13)}. \end{cases}$$

For Case 1, the rate is easily seen to be upper bounded by $O((\log \log n)/n)$, as shown inside the proof of Theorem 5.5. In Case 3, Theorem 5.14 establishes that the probability that model M_0 is chosen is at most of order $1/(\log n)$, which, as shown inside the proof of Theorem 5.5, again implies an upper-bound on the rate-of convergence of $O((\log \log n)/n)$. Theorem 5.13 shows that in Case 3, which includes the case that $\|\mu_1^{(n)} - \mu_0^{(n)}\|_2^2$ does not converge at all, the probability that model M_0 is chosen is at most of order $1/n$, which, as again shown inside the proof of Theorem 5.5, again implies an upper-bound on the rate-of convergence of $O((\log \log n)/n)$.

The two theorems take into account that $\mu_1^{(n)}$ is not just a fixed function of n , but may in reality be chosen by nature in a worst-case manner, and that $f(n)$ may actually fluctuate between regions for different n . Combining these two results, we finally prove the main theorem, Theorem 5.5.

Theorem 5.13. Let M_0, M_1, M'_1 and $p_{\text{sw},1}(x^n)$ be as in Theorem 5.5. Then there exist positive constants \vec{b}, \vec{c} such that for all $\mu_1 \in M'_1$, all n ,

$$\mathbb{P}_{\mu_1}(\delta_{\text{sw}}(X^n) = 0) \leq c_1 \cdot n^{m_1/2} \cdot e^{-c_2 n \|\mu_1^{(n)} - \mu_0^{(n)}\|_2^2} + \text{small}_{\vec{c}}(n), \quad (5.45)$$

where $\mu_0^{(n)} = \Pi_0(\mu_1^{(n)})$ is as in (5.11). As a consequence, with $\Gamma_2 := c_2^{-1}(1 + m_1/2)$, we have the following: for every sequence $\mu_1^{(1)}, \mu_1^{(2)}, \dots$ with $f(n)$ as in (5.44) larger than $\Gamma_2 \log n$, we have

$$\mathbb{P}_{\mu_1^{(n)}}(\delta_{\text{sw}}(X^n) = 0) \leq \frac{c_1}{n} + \text{small}_b(n).$$

Proof. We can bound the probability of selecting the simple model by:

$$\begin{aligned} \mathbb{P}_{\mu_1^{(n)}}(\delta_{\text{sw}}(X^n) = 0) &= \mathbb{P}_{\mu_1^{(n)}}\left(\frac{p_{\text{sw},1}(X^n)}{p_{B,0}(X^n)} \leq 1\right) = \mathbb{P}_{\mu_1^{(n)}}\left(\frac{\sum_{i=0}^{\infty} \pi(2^i) \bar{p}_{2^i}(X^n)}{p_{B,0}(X^n)} \leq 1\right) \\ &\leq \mathbb{P}_{\mu_1^{(n)}}\left(\frac{\pi(1)p_{B,1}(X^n)}{p_{B,0}(X^n)} \leq 1\right). \end{aligned}$$

Now (5.45) follows directly by applying Lemma 5.12 to the rightmost probability. For the second part, set $\Gamma_2 = c_2^{-1}(1 + m_1/2)$. By assumption $f(n) > \Gamma_2 \log n$, we have $\|\mu_1^{(n)} - \mu_0^{(n)}\|_2^2 > \Gamma_2(\log n)(\log \log n)/n$. Applying (5.45) now gives the desired result. \square

Theorem 5.14. Let f be as in (5.44) and M'_1 be as in Theorem 5.5. For any $\gamma > 0$, there exist constants $\Gamma_1, \Gamma_3 > 0$ such that, for every sequence $\mu_1^{(1)}, \mu_1^{(2)}, \dots$ of elements of M'_1 with for all n , $f(n) > \Gamma_1$, we have

$$\mathbb{P}_{\mu_1^{(n)}}\left(\frac{p_{\text{sw},1}(X^n)}{p_{B,0}(X^n)} \leq \gamma\right) \leq \frac{\Gamma_3}{\log n}. \quad (5.46)$$

In particular, by taking $\gamma = 1$, we have

$$\mathbb{P}_{\mu_1^{(n)}}(\delta_{\text{sw}}(X^n) = 0) \leq \frac{\Gamma_3}{\log n}.$$

The probabilities thus converge uniformly at rate $O(1/(\log n))$ for all such sequences $\mu_1^{(1)}, \mu_1^{(2)}, \dots$.

Proof. We specify Γ_1 later. By assumption, we have $\pi(2^i) \gtrsim (\log n)^{-\kappa}$ for $i \in \{0, \dots, \lfloor \log_2 n \rfloor\}$. We can restrict our attention to the strategy that switches to the complex model at the penultimate switching index, due to the following inequality: for any fixed γ , there exist positive constants \vec{C} such that for all large n :

$$\begin{aligned} \mathbb{P}_{\mu_1^{(n)}}\left(\frac{p_{\text{sw},1}(X^n)}{p_{B,0}(X^n)} \leq \gamma\right) &\leq \mathbb{P}_{\mu_1^{(n)}}\left(\frac{\sum_{i=0}^{\lfloor \log_2 n \rfloor} \pi(2^i) \bar{p}_{2^i}(X^n)}{p_{B,0}(X^n)} \leq \gamma\right) \\ &\leq \mathbb{P}_{\mu_1^{(n)}}\left(\frac{\sum_{i=0}^{\lfloor \log_2 n \rfloor} \bar{p}_{2^i}(X^n)}{p_{B,0}(X^n)} \leq C_1(\log n)^\kappa\right) \\ &\leq \mathbb{P}_{\mu_1^{(n)}}\left(\frac{\bar{p}_{2^{\lfloor \log_2 n \rfloor - 1}}(X^n)}{p_{B,0}(X^n)} \leq C_1(\log n)^\kappa\right) \\ &= \mathbb{P}_{\mu_1^{(n)}}\left(\log \frac{\bar{p}_{2^{\lfloor \log_2 n \rfloor - 1}}(X^n)}{p_{B,0}(X^n)} \leq \kappa \log \log n + C_2\right). \end{aligned} \quad (5.47)$$

For the remainder of this proof, we will denote the penultimate switching index by n^* , that is: $n^* = 2^{\lfloor \log_2 n \rfloor - 1}$. Now apply Lemma 5.11 twice, which gives that there exist C_3, C_4 such that, with probability at least $1 - \text{small}_{\vec{b}}(n)$,

$$\begin{aligned} \log \bar{p}_{n^*}(X^n) &= \log p_{B,0}(X^{n^*}) + \log p_{B,1}(X^n | X^{n^*}) = \\ &= \log p_{B,0}(X^{n^*}) + \log p_{B,1}(X^n) - \log p_{B,1}(X^{n^*}) \\ &\geq \log p_{B,0}(X^{n^*}) + \log p_{\widehat{\mu}_1(X^n)}(X^n) - \log p_{\widehat{\mu}_1(X^{n^*})}(X^{n^*}) + \frac{m_1}{2} \log \frac{n^*}{n} - C_3 \\ &\geq \log p_{B,0}(X^{n^*}) + \log \frac{p_{\widehat{\mu}_1(X^n)}(X^n)}{p_{\widehat{\mu}_1(X^{n^*})}(X^{n^*})} - C_4, \end{aligned} \quad (5.48)$$

where we used that $\log \frac{n^*}{n}$ is of the order of a constant, because n^* is between $\frac{n}{4}$ and $\frac{n}{2}$. From this, applying again Lemma 5.11 twice, it follows that there exists \vec{b} and C_5, C_6 such that for all n , with probability at least $1 - \text{small}_{\vec{b}}(n)$,

$$\begin{aligned} \log \frac{\bar{p}_{n^*}(X^n)}{p_{B,0}(X^n)} &\geq \log \frac{p_{B,0}(X^{n^*})}{p_{B,0}(X^n)} + \log \frac{p_{\widehat{\mu}_1(X^n)}(X^n)}{p_{\widehat{\mu}_1(X^{n^*})}(X^{n^*})} - C_4 \\ &= -\log \frac{p_{\widehat{\mu}_0(X^n)}(X^n)}{p_{\widehat{\mu}_0(X^{n^*})}(X^{n^*})} - \frac{m_0}{2} \log \frac{n^*}{n} + \log \frac{p_{\widehat{\mu}_1(X^n)}(X^n)}{p_{\widehat{\mu}_1(X^{n^*})}(X^{n^*})} - C_5 \\ &\geq -\log \frac{p_{\widehat{\mu}_0(X^n)}(X^n)}{p_{\widehat{\mu}_0(X^{n^*})}(X^{n^*})} + \log \frac{p_{\widehat{\mu}_1(X^n)}(X^n)}{p_{\widehat{\mu}_1(X^{n^*})}(X^{n^*})} - C_6 \end{aligned} \quad (5.49)$$

where we again used that $\log \frac{n^*}{n}$ can be bounded by constants. Let \mathcal{B}_n be the event that (5.49) holds. By (5.47) and (5.49), for all large n , all $\beta \geq 1$,

$$\begin{aligned} \mathbb{P}_{\mu_1^{(n)}} \left(\frac{p_{\text{sw},1}(X^n)}{p_{B,0}(X^n)} \leq \gamma \right) &\leq \mathbb{P}_{\mu_1^{(n)}} \left(\log \frac{\bar{p}_{n^*}(X^n)}{p_{B,0}(X^n)} \leq \kappa \log \log n + C_2 \right) \\ &\leq \mathbb{P}_{\mu_1^{(n)}} \left(\log \frac{\bar{p}_{n^*}(X^n)}{p_{B,0}(X^n)} \leq \kappa \log \log n + C_2, \mathcal{B}_n \right) + \mathbb{P}_{\mu_1^{(n)}}(\mathcal{B}_n^c) \\ &\leq \mathbb{P}_{\mu_1^{(n)}} \left(-\log \frac{p_{\widehat{\mu}_0(X^n)}(X^n)}{p_{\widehat{\mu}_0(X^{n^*})}(X^{n^*})} + \log \frac{p_{\widehat{\mu}_1(X^n)}(X^n)}{p_{\widehat{\mu}_1(X^{n^*})}(X^{n^*})} - C_6 \leq \kappa \log \log n + C_2 \right) + \text{small}_{\vec{b}}(n) \\ &= \mathbb{P}_{\mu_1^{(n)}}(\mathcal{E}_n^{(1)}) + \text{small}_{\vec{b}}(n) \leq \mathbb{P}_{\mu_1^{(n)}}(\mathcal{E}_n^{(\beta)}) + \text{small}_{\vec{b}}(n), \end{aligned} \quad (5.50)$$

where we defined

$$\mathcal{E}_n^{(\beta)} = \left\{ \log \frac{p_{\widehat{\mu}_1(X^n)}(X^n)}{p_{\widehat{\mu}_1(X^{n^*})}(X^{n^*})} \cdot \frac{p_{\widehat{\mu}_0(X^{n^*})}(X^{n^*})}{p_{\widehat{\mu}_0(X^n)}(X^n)} \leq A_n^{(\beta)} \right\} \quad (5.51)$$

and, for $\beta \geq 1$, we set $A_n^{(\beta)} = \beta \kappa \log \log n + C_2 + C_6$.

Below, if a sample is split up into two parts x_1, \dots, x_{n^*} and x_{n^*+1}, \dots, x_n , these partial samples will be referred to as x^{n^*} and $x^{>n^*}$ respectively. We also suppress in our notation

the dependency of A_n , \mathcal{E}_n and $\mathcal{D}_{j,n}$ as defined below on β ; all results below hold, with the same constants, for any $\beta \geq 1$.

We will now bound the right-hand side of (5.50) further. Define the events

$$\begin{aligned}\mathcal{D}_{1,n} &= \left\{ \log \frac{p_{\mu_1^{(n)}}(x^n)}{p_{\mu_1^{(n)}}(x^{n*})} \leq \log \frac{p_{\widehat{\mu}_1(X^n)}(x^n)}{p_{\widehat{\mu}_1(X^{n*})}(x^{n*})} + A_n \right\} \\ \mathcal{D}_{0,n} &= \left\{ \log \frac{p_{\mu_0^{(n)}}(x^n)}{p_{\mu_0^{(n)}}(x^{n*})} \geq \log \frac{p_{\widehat{\mu}_0(X^n)}(x^n)}{p_{\widehat{\mu}_0(X^{n*})}(x^{n*})} - A_n \right\}.\end{aligned}$$

The probability in (5.50) can be bounded, for all $\beta \geq 1$, as

$$\begin{aligned}\mathbb{P}_{\mu_1^{(n)}}(\mathcal{E}_n) &= \mathbb{P}_{\mu_1^{(n)}}(\mathcal{E}_n, \mathcal{D}_{0,n} \cap \mathcal{D}_{1,n}) + \mathbb{P}_{\mu_1^{(n)}}(\mathcal{E}_n, (\mathcal{D}_{0,n} \cap \mathcal{D}_{1,n})^c) + \text{small}_{\vec{b}}(n) \\ &\leq \mathbb{P}_{\mu_1^{(n)}}(\mathcal{E}_n, \mathcal{D}_{0,n}, \mathcal{D}_{1,n}) + \mathbb{P}_{\mu_1^{(n)}}(\mathcal{D}_{1,n}^c) + \mathbb{P}_{\mu_1^{(n)}}(\mathcal{D}_{0,n}^c) + \text{small}_{\vec{b}}(n).\end{aligned}\quad (5.52)$$

We first consider the first probability in (5.52): there are constants \vec{C} such that, for all large n ,

$$\begin{aligned}\mathbb{P}_{\mu_1^{(n)}}(\mathcal{E}_n, \mathcal{D}_{0,n}, \mathcal{D}_{1,n}) &\leq \mathbb{P}_{\mu_1^{(n)}} \left(\log \frac{p_{\mu_1^{(n)}}(X^n)}{p_{\mu_1^{(n)}}(X^{n*})} - A_n + \log \frac{p_{\mu_0^{(n)}}(X^n)}{p_{\mu_0^{(n)}}(X^{n*})} - A_n \leq A_n \right) \\ &= \mathbb{P}_{\mu_1^{(n)}} \left(\log \frac{p_{\mu_1^{(n)}}(X^{>n*})}{p_{\mu_0^{(n)}}(X^{>n*})} \leq 3A_n \right) \\ &\leq e^{\frac{3}{2}A_n} e^{-\frac{n}{4}d_R(\mu_1^{(n)}, \mu_0^{(n)})} \leq e^{(3/2)\beta\kappa \log \log n + C_7} e^{-C_8 n \|\mu_1^{(n)} - \mu_0^{(n)}\|_2^2} = e^{C_7(\log n)^{(3/2)\beta\kappa - \Gamma_1 \cdot C_8}},\end{aligned}\quad (5.53)$$

where Γ_1 is as in the statement of the theorem, the second inequality follows by Lemma 5.9 and noting $n^* < \frac{n}{2}$, we used Proposition 5.2.

We now consider the second probability in (5.52). Using $p_{\widehat{\mu}_1(X^n)}(x^n) \geq p_{\mu_1^{(n)}}(x^n)$ we have the following, where we define the event $\mathcal{F}_n = \{\widehat{\mu}_1(X^{n*}) \in M'_1\}$ with M'_1 the CINECSI subset of M_1 mentioned in the theorem statement: there is $C_9, C_{10} > 0$ such that for all large n ,

$$\begin{aligned}\mathbb{P}_{\mu_1^{(n)}}(\mathcal{D}_{1,n}^c) &= \mathbb{P}_{\mu_1^{(n)}} \left(\log \frac{p_{\mu_1^{(n)}}(X^n)}{p_{\mu_1^{(n)}}(X^{n*})} > \log \frac{p_{\widehat{\mu}_1(X^n)}(X^n)}{p_{\widehat{\mu}_1(X^{n*})}(X^{n*})} + A_n \right) \\ &\leq \mathbb{P}_{\mu_1^{(n)}} \left(\log \frac{p_{\widehat{\mu}_1(X^n)}(X^n)}{p_{\mu_1^{(n)}}(X^{n*})} > \log \frac{p_{\widehat{\mu}_1(X^n)}(X^n)}{p_{\widehat{\mu}_1(X^{n*})}(X^{n*})} + A_n \right) \\ &\leq \mathbb{P}_{\mu_1^{(n)}} \left(\log \frac{p_{\widehat{\mu}_1(X^{n*})}(X^{n*})}{p_{\mu_1^{(n)}}(X^{n*})} > A_n, \mathcal{F}_n \right) + \mathbb{P}_{\mu_1^{(n)}}(\mathcal{F}_n^c)\end{aligned}$$

$$\begin{aligned}
&\leq \mathbb{P}_{\mu_1^{(n)}} \left(D(\widehat{\mu}_1(X^{n^*}) \| \mu_1^{(n)}) > A_n, \mathcal{F}_n \right) + \text{small}_{\tilde{b}}(n) \\
&\leq \mathbb{P}_{\mu_1^{(n)}} \left(\|\widehat{\mu}_1(X^{n^*}) - \mu_1^{(n)}\|_2^2 > C_9 A_n, \mathcal{F}_n \right) + \text{small}_{\tilde{b}}(n) \\
&\leq \mathbb{P}_{\mu_1^{(n)}} \left(\|\widehat{\mu}_1(X^{n^*}) - \mu_1^{(n)}\|_\infty > \sqrt{C_9 A_n / m_1} \right) + \text{small}_{\tilde{b}}(n) \tag{5.54}
\end{aligned}$$

$$\leq e^{-C_{10} A_n} = e^{-C_{10}(C_2 - C_6)} \frac{1}{(\log n)^{C_{10} \beta \kappa}}, \tag{5.55}$$

where we used the KL robustness property (5.25), Proposition 5.2 and Lemma 5.8.

The third probability in (5.52) is considered in a similar way. Using $p_{\widehat{\mu}_0(X^{n^*})}(X^{n^*}) \geq p_{\mu_0^{(n)}}(X^{n^*})$ we have $C_{11}, C_{12} > 0$ such that:

$$\begin{aligned}
\mathbb{P}_{\mu_1^{(n)}}(\mathcal{D}_{0,n}^c) &= \mathbb{P}_{\mu_1^{(n)}} \left(\log \frac{p_{\mu_0^{(n)}}(X^n)}{p_{\mu_0^{(n)}}(X^{n^*})} < \log \frac{p_{\widehat{\mu}_0(X^n)}(X^n)}{p_{\widehat{\mu}_0(X^{n^*})}(X^{n^*})} - \frac{1}{3} A_n \right) \\
&\leq \mathbb{P}_{\mu_1^{(n)}} \left(\log \frac{p_{\mu_0^{(n)}}(X^n)}{p_{\widehat{\mu}_0(X^{n^*})}(X^{n^*})} < \log \frac{p_{\widehat{\mu}_0(X^n)}(X^n)}{p_{\widehat{\mu}_0(X^{n^*})}(X^{n^*})} - \frac{1}{3} A_n \right) \\
&= \mathbb{P}_{\mu_1^{(n)}} \left(\log \frac{p_{\widehat{\mu}_0(X^n)}(X^n)}{p_{\mu_0^{(n)}}(X^n)} > \frac{1}{3} A_n \right) \\
&\leq C_{11} \frac{1}{(\log n)^{C_{12} \beta \kappa}} \tag{5.56}
\end{aligned}$$

where we omitted the last few steps which are exactly as in (5.54).

We now finish the proof by combining (5.52), (5.53), (5.54) and (5.56), which gives that, if we choose $\beta \geq \max\{1/(\kappa C_{10}), 1/(\kappa C_{12})\}$ and, for this choice of β , we choose Γ_1 as in (5.53) as $\Gamma_1 \geq (1 + (3/2)\beta\kappa)/C_8$, then we have $\mathbb{P}_{\mu_1^{(n)}}(\mathcal{E}_n) \leq \Gamma_4/(\log n)$ for some constant Γ_4 independent of n ; the result now follows from (5.50). \square

Proof of Theorem 5.5

Proof. We show the result in two stages. In Stage 1 we provide a tight upper bound on the risk, based on an extension of the decomposition of the risk (5.12) to general families and estimators $\check{\mu}_0$ and $\check{\mu}_1$ that are sufficiently efficient, i.e. that satisfy (5.13), and to losses $d_{\text{gen}}(\cdot \| \cdot)$ equal to squared error loss, standardized squared error loss and KL divergence (it is not sufficient to refer to Proposition 5.2 and prove the result only for squared error loss, because the equivalence result of Proposition 5.2 only holds on CINECSI sets and our estimators may take values outside of these; we do not need to consider Rényi and squared Hellinger divergences though, because these are uniformly upper bounded by KL divergence even for μ outside any CINECSI set). In Stage 2 we show how the bound implies the result.

Stage 1: Decomposition of upper bound on the risk Let A_n be the event that \mathcal{M}_1 is selected, as in Section 5.3.3. We will now show that, under the assumptions of Theorem 5.5,

we have for the constant C appearing in (5.13), for all $\mu_1 \in M'_1$,

$$R(\mu_1, \delta, n) \leq \frac{3C}{n} + 2\mathbb{P}(A_n^c) d_{\text{gen}}(\mu_1 \| \mu_0), \quad (5.57)$$

where the left inequality holds for all divergence measures mentioned in the theorem, and the right inequality holds for $d_{\text{gen}}(\cdot \| \cdot)$ set to any of the squared error, the standardized squared error or the KL divergence.

To prove (5.57), we use that for the three divergences of interest, for any $\mu_1 \in M_1, \mu \in M_0$, with $\mu_0 \in M_0$ as in (5.11), we have

$$d_{\text{gen}}(\mu_1 \| \mu) \leq 2(d_{\text{gen}}(\mu_1 \| \mu_0) + d_{\text{gen}}(\mu_0 \| \mu)), \quad (5.58)$$

For $d_{\text{gen}}(\cdot \| \cdot)$ the KL divergence, this follows because

$$\begin{aligned} D(\mu_1 \| \mu) &= \mathbb{E}_{\mu_1} \left[-\log \frac{p_\mu(X)}{p_{\mu_1}(X)} \right] = \mathbb{E}_{\mu_1} \left[-\log \frac{p_\mu(X)}{p_{\mu_0}(X)} \right] + \mathbb{E}_{\mu_1} \left[-\log \frac{p_{\mu_0}(X)}{p_{\mu_1}(X)} \right] \\ &= \mathbb{E}_{\mu_0} \left[-\log \frac{p_\mu(X)}{p_{\mu_0}(X)} \right] + \mathbb{E}_{\mu_1} \left[-\log \frac{p_{\mu_0}(X)}{p_{\mu_1}(X)} \right], \end{aligned} \quad (5.59)$$

where the last line follows by the robustness property of exponential families (5.24), since μ and μ_0 are both in M_0 .

For $d_{\text{gen}}(\cdot \| \cdot)$ the squared and standardized squared error case we show (5.58) as follows: Fix a matrix-valued function $J : M_1 \rightarrow \mathbb{R}^{m^2}$ that maps each $\mu \in M_1$ to a positive definite matrix J_μ . We can write

$$d_{\text{gen}}(\mu \| \mu') = (\mu - \mu')^T J_\mu (\mu - \mu'). \quad (5.60)$$

where J_μ is the identity matrix for the squared error case, and J_μ is the Fisher information matrix for the standardized squared error case. (5.58) follows since we can write, for any function J_μ of the above type including these two:

$$\begin{aligned} (\mu_1 - \mu)^T J_{\mu_1} (\mu_1 - \mu) &= (\mu_1 - \mu_0 + \mu_0 - \mu)^T J_{\mu_1} (\mu_1 - \mu_0 + \mu_0 - \mu) \\ &= (\mu_1 - \mu_0)^T J_{\mu_1} (\mu_1 - \mu_0) + (\mu_0 - \mu)^T J_{\mu_1} (\mu_0 - \mu) + 2(\mu_1 - \mu_0)^T J_{\mu_1} (\mu_0 - \mu) \\ &\leq 2 \left((\mu_1 - \mu_0)^T J_{\mu_1} (\mu_1 - \mu_0) + (\mu_0 - \mu)^T J_{\mu_1} (\mu_0 - \mu) \right), \end{aligned}$$

where the last line follows because for general positive definite $m \times m$ matrices J and m -component column vectors a and b , $(b - a)^T J (b - a) \geq 0$ so that $b^T J (b - a) \geq a^T J (b - a)$ and, after rearranging, $b^T J b + a^T J a \geq 2a^T J b$.

We have thus shown (5.58). It now follows that

$$\begin{aligned} R(\mu_1, \delta, n) &= \mathbb{E}_{\mu_1} \left[\mathbf{1}_{A_n} d_{\text{gen}}(\mu_1 \| \check{\mu}_1(X^n)) + \mathbf{1}_{A_n^c} d_{\text{gen}}(\mu_1 \| \check{\mu}_0(X^n)) \right] \\ &\leq \mathbb{E}_{\mu_1} \left[d_{\text{gen}}(\mu_1 \| \check{\mu}_1(X^n)) + 2 \cdot \mathbf{1}_{A_n^c} \left(d_{\text{gen}}(\mu_0 \| \check{\mu}_0(X^n)) + d_{\text{gen}}(\mu_1 \| \mu_0) \right) \right] \\ &\leq \frac{3C}{n} + 2\mathbb{P}(A_n^c) d_{\text{gen}}(\mu_1 \| \mu_0), \end{aligned} \quad (5.61)$$

where we used (5.58) and our condition (5.13) on $\check{\mu}_0$ and $\check{\mu}_1$. We have thus shown (5.57).

Stage 2 We proceed to prove our risk upper bound for the squared error loss, standardized squared error loss and KL divergence, for which the right inequality in (5.57) holds; the result then follows for squared Hellinger and Rényi divergence because these are upper bounded by KL divergence. From (5.57) we see that it is sufficient to show that for all n larger than some n_0 ,

$$\sup_{\mu_1 \in M'_1} \{\mathbb{P}_{\mu_1}(A_n^c) d_{\text{gen}}(\mu_1 \| \mu_0)\} = O\left(\frac{\log \log n}{n}\right), \quad (5.62)$$

for our three choices of $d_{\text{gen}}(\cdot \| \cdot)$. We first note that, since M'_1 is CINECSI, $\sup_{\mu_1 \in M'_1} d_{\text{gen}}(\mu_1 \| \mu_0)$ is bounded by some constant C_1 . It thus follows by Proposition 5.10 that there exists some CINECSI subset M''_1 of M_1 such that, with $B_n^c \subset A_n^c$ defined as $B_n^c = \{x^n : \delta(x^n) = 0; \widehat{\mu}_1(X^n) \in M''_1\}$, we have

$$\begin{aligned} \sup_{\mu_1 \in M'_1} \{\mathbb{P}_{\mu_1}(A_n^c) d_{\text{gen}}(\mu_1 \| \mu_0)\} &= \sup_{\mu_1 \in M'_1} \{(\mathbb{P}_{\mu_1}(B_n^c) + \mathbb{P}_{\mu_1}(A_n^c \setminus B_n^c)) d_{\text{gen}}(\mu_1 \| \mu_0)\} \\ &= \sup_{\mu_1 \in M'_1} \{\mathbb{P}_{\mu_1}(B_n^c) d_{\text{gen}}(\mu_1 \| \mu_0)\} + C_1 \cdot \mathbb{P}_{\mu_1}(\widehat{\mu}^{(1)} \notin M''_1) \\ &= \sup_{\mu_1 \in M'_1} \{\mathbb{P}_{\mu_1}(B_n^c) d_{\text{gen}}(\mu_1 \| \mu_0)\} + \text{small}_b(n), \end{aligned}$$

so that it is sufficient if we can show (5.62) with B_n^c instead of A_n^c . But on the set B_n^c , all three divergence measures considered are within constant factors of each other, so that it is sufficient if we can show that there is a constant C_2 such that for all n larger than some n_0 ,

$$\sup_{\mu_1 \in M'_1} \{\mathbb{P}_{\mu_1}(B_n^c) \cdot \|\mu_1 - \mu_0\|_2^2\} \leq C_2 \cdot \frac{\log \log n}{n}. \quad (5.63)$$

Now, fix some $\mu_1 \equiv \mu_1^{(n)}$ and consider $f(n)$ as in (5.44). By Theorem 5.13, $\mathbb{P}_{\mu_1}(B_n^c) \leq C_3/n$ for some constant C_3 that can be chosen uniformly for all $\mu_1 \in M'_1$ whenever $f(n) > \Gamma_2 \log n$ with Γ_2 as in that theorem. Using also that $\|\mu_1 - \mu_0\|_2^2$ is bounded by C_1 as above, it follows that (5.63) holds whenever $f(n) > \Gamma_2 \log n$ and $(C_1 C_3)/n \leq C_2 (\log \log n)/n$, i.e. whenever $f(n) > \Gamma_2 \log n$ and $C_2 \geq C_1 C_3 / (\log \log n)$.

Second, suppose that $\Gamma_1 < f(n) \leq \Gamma_2 \log n$ with Γ_1 as in Theorem 5.14. Then by that theorem, uniformly for all $\mu_1^{(n)}$ with such $f(n)$, we have, with Γ_3 as in that theorem,

$$\begin{aligned} \|\mu_1^{(n)} - \mu_0^{(n)}\|_2^2 \cdot \mathbb{P}_{\mu_1^{(n)}}(\delta_{\text{sw}}(X^n) = 0) &= f(n) \cdot \frac{\log \log n}{n} \cdot \mathbb{P}_{\mu_1^{(n)}}(\delta_{\text{sw}}(X^n) = 0) \leq \\ \Gamma_2 \cdot (\log n) \cdot \frac{\log \log n}{n} \cdot \mathbb{P}_{\mu_1^{(n)}}(\delta_{\text{sw}}(X^n) = 0) &\leq \Gamma_2 \Gamma_3 \cdot \frac{\log \log n}{n}, \end{aligned}$$

where $\mu_0^{(n)} = \Pi_0(\mu_1^{(n)})$ is defined as in (5.11), so that (5.63) holds again whenever $C_2 \geq \Gamma_2 \Gamma_3$.

Finally, suppose that $f(n) \leq \Gamma_1$ with Γ_1 as in Theorem 5.14. Then (5.63) holds whenever $C_2 \geq \Gamma_1$. Combining the three cases we find that (5.63) holds whenever $C_2 \geq \max\{\Gamma_1, \Gamma_2 \Gamma_3, C_1 C_3 / (\log \log n)\}$; the result is proved. \square

5.7.5 Switching as in Van Erven et al. (2012)

The basic building block of the switch distribution and criterion as formulated by Van Erven et al. (2012) is a countable set of *sequential prediction strategies* (also known as ‘prequential forecasting systems’ (Dawid, 1984)) $\{p_k \mid k \in \mathcal{K}\}$, where \mathcal{K} is a finite or countable set indexing the basic models under consideration. Thus, each model is associated with a corresponding prediction strategy, where a prediction strategy p is a function from $\bigcup_{i \geq 0} \mathcal{X}^i$ to the set of densities on \mathcal{X} , where $p(\cdot \mid x^{n-1})$ denotes the density on \mathcal{X} that x^{n-1} maps to, and $p(x_n \mid x^{n-1})$ is to be interpreted as the probabilistic prediction that strategy p makes for outcome X_n upon observation of the first $n - 1$ outcomes, $X^{n-1} = x^{n-1}$. For example, for a parametric model $\{p_\theta \mid \theta \in \Theta\}$ one can base p_k on a Bayesian marginal likelihood, $p_B(x^n) := \int_\Theta \omega(\theta) p_\theta(x^n) d\theta$, where ω is a prior density on Θ . The corresponding prediction strategy could then be defined by setting $p_k(x_n \mid x^{n-1}) := p_B(x^n)/p_B(x^{n-1})$, the standard Bayesian predictive distribution. In this paper, the basic strategies p_k were always Bayesian predictive distributions, but, in the spirit of Dawid (1984), one may consider other choices as well.

After constructing the set of basic prediction strategies, a new family of prediction strategies that switch between the strategies in the set $\{p_k \mid k \in \mathcal{K}\}$ is defined. Formally, let \mathbb{S} be the set

$$\mathbb{S} = \{((t_1, k_1), \dots, (t_m, k_m)) \in (\mathbb{N} \times \mathcal{K})^m \mid m \in \mathbb{N}, 1 = t_1 < t_2 < \dots < t_m\}. \quad (5.64)$$

Each $s \in \mathbb{S}$ specifies the times t_1, \dots, t_m at which a switch is made between the prediction strategies from the original set, identified by the indices k_1, \dots, k_m . The new family $Q = \{q_s \mid s \in \mathbb{S}\}$ is then defined by setting, for all $n, x^n \in \mathcal{X}^n$:

$$q_s(x_n \mid x^{n-1}) = p_{k_j}(x_n \mid x^{n-1}), \quad t_j \leq n < t_{j+1}, \quad (5.65)$$

with $t_{m+1} = \infty$ by convention. We now define $q_s(x^n) = \prod_{i=1}^n q_s(x_i \mid x^{i-1})$; one easily verifies that this defines a joint probability density on \mathcal{X}^n .

We now place a prior mass function π' on \mathbb{S} and define, for each n , the *switch distribution* in terms of its joint density for \mathcal{X}^n and \mathbb{S} :

$$p_{\text{sw}}(x^n, s) = q_s(x^n) \pi'(s), \quad p_{\text{sw}}(x^n) = \sum_{s \in \mathbb{S}} p_{\text{sw}}(x^n, s) = \sum_{s \in \mathbb{S}} q_s(x^n) \pi'(s).$$

If the p_k are defined as Bayesian predictive distributions as above, then, as explained by Van Erven et al. (2012), the density $p_{\text{sw}}(x^n)$ can be interpreted as a Bayesian marginal density of x^n under the prior π' on meta-models (model sequences) in \mathbb{S} .

The switch distribution can be used to define a model selection criterion δ'_{sw} by selecting the model with highest posterior probability under the switch distribution. This is done by defining the random variable $K_{n+1}(s)$ on \mathbb{S} to be the index of the prediction strategy that is used by q_s to predict the $(n + 1)$ th outcome. The model selection criterion is then:

$$\delta'_{\text{sw}}(x^n) = \arg \max_k p_{\text{sw}}(K_{n+1} = k \mid x^n) = \arg \max_k \frac{\sum_{s: K_{n+1}(s)=k} p_{\text{sw}}(x^n, s)}{p_{\text{sw}}(x^n)}$$

$$= \arg \max_k \frac{\sum_{s: K_{n+1}(s)=k} q_s(x^n) \pi'(s)}{\sum_{s \in \mathbb{S}} q_s(x^n) \pi'(s)}, \quad (5.66)$$

with ties resolved in any way desired.

In our nested two-model case, one might use, for example, a prior π' with support on

$$\mathbb{S}' = \{(1,0), (1,1), ((1,0), (2,1)), ((1,0), (4,1)), ((1,0), (8,1)), ((1,0), (16,1)), \dots\}.$$

Such a prior expresses that at time 1, for the first prediction, one can either switch to (i.e., start with), model 0, and keep predicting according to its Bayes predictive distribution — this strategy gets weight $\pi((1,0))$. Or one can start with model 1, and keep predicting according to its Bayes predictive distribution — this strategy gets weight $\pi((1,1))$. Or one can start with model 0 and switch to model 1 after 2^i observations and then stick with 1 forever — this strategy gets weight $\pi((1,0), (2^i, 1))$. If we now start with a prior π on $\{1, 2, \dots\}$ as in the main text and define $\pi'((1,0)) = 1/2$, $\pi'((1,1)) = (1/2) \cdot \pi(1)$, and for $i \geq 1$, $\pi'((1,0), (2^i, 1)) = (1/2) \cdot \pi(2^i)$, then $\sum_{s \in \mathbb{S}'} \pi'(s) = 1$, so π' is a probability mass function. A simple calculation gives that (5.66) based on switch prior π' now chooses model 1 if

$$\sum_{1 \leq t < n} \bar{p}_t(x^n) \pi(t) > (1 + g(n)) \cdot p_{B,0}(x^n), \quad (5.67)$$

where $g(n) = \sum_{t \geq n} \pi(t)$; note that $g(n)$ is decreasing and converges to 0 with increasing n . (5.67) is thus an instance of the switch criterion of Van Erven et al. (2012). Comparing this to (5.3), the criterion used in this paper, after rearranging we see that it chooses model 1 if

$$\sum_{1 \leq t < n} \bar{p}_t(x^n) \pi(t) > (1 - g(n)) \cdot p_{B,0}(x^n),$$

which is more likely by constant factor to select model \mathcal{M}_0 , the factor however tending to 1 with increasing n . It is completely straightforward to check that Theorem 5.5 and all other results in this paper still hold if δ_{sw} with prior π as in the main text is replaced by δ'_{sw} with corresponding prior π' as defined here; thus our results carry over to the original definitions of Van Erven et al. (2012). Similarly, the proof for the strong consistency of δ'_{sw} given by Van Erven et al. (2012) carries through for δ_{sw} , needing only trivial modifications.