

# Chapter 1

## Bernstein - von Mises theorem and misspecified models: a review

Natalia Bochkina

**Abstract** This is a review of asymptotic and non-asymptotic behaviour of Bayesian methods under model specification. In particular we focus on consistency, i.e. convergence of the posterior distribution to the point mass at the best parametric approximation to the true model, and conditions for it to be locally Gaussian around this point. For well specified regular models, variance of the Gaussian approximation coincides with the Fisher information, making Bayesian inference asymptotically efficient. In this review, we discuss how this is affected by model misspecification. In particular, we highlight contribution of Volodia Spokoiny to this area. We also discuss approaches to adjust Bayesian inference to make it asymptotically efficient under model misspecification.

**Abstract** This is a review of asymptotic and non-asymptotic behaviour of Bayesian methods under model specification. In particular we focus on consistency, i.e. convergence of the posterior distribution to the point mass at the best parametric approximation to the true model, and conditions for it to be locally Gaussian around this point. For well specified models, variance of the Gaussian approximation coincides with the Fisher information makes Bayesian inference asymptotically efficient which does not always hold under model misspecification. In particular, we highlight contribution of Volodia Spokoiny to this area. We also discuss approaches to adjust Bayesian inference so that it is asymptotically efficient under model misspecification.

### 1.1 Introduction

Consider a family of probability models  $P(\mathbf{Y} | \theta)$  indexed by parameter  $\theta \in \Theta$  for observations  $\mathbf{y}$ , and a prior distribution  $\pi$  on the parameter space  $\Theta$ . In a classical Bayesian approach, the posterior distribution  $p(\theta | \mathbf{Y})$  is used for statistical inference [Van der Vaart, 2000].

---

University of Edinburgh and Maxwell Institute, UK,  
e-mail: N.Bochkina@ed.ac.uk

Denote the true distribution of observations  $P_0$ , and we consider the case  $P_0 \notin \{P(\cdot | \theta), \theta \in \Theta\}$ , i.e. the model is misspecified. Such case arises in many applications, particularly in complex models where the numerical evaluation of posterior distribution under the ideal probability model takes a long time to run, leading to increased use of approximate models with faster computing time. A typical example is approximating complex dependence structure by pairwise dependence only [Besag, 1986], [Varin et al., 2011].

For well-specified regular models, the classical Bernstein - von Mises theorem states that for  $n$  independent identically distributed (iid) observations, for large  $n$ , the posterior distribution behaves approximately as a normal distribution centered on the true value of the parameter with a random shift, with both the posterior variance and the variance of the random shift being asymptotically equal to the inverse Fisher information, thus making the Bayesian inference asymptotically consistent and efficient. This was extended to locally asymptotically normal (LAN) models for  $\Theta \in \mathbb{R}^p$  with fixed  $p$  [Van der Vaart, 2000], for LAN models with growing  $p$ , etc.

A version of Bernstein - von Mises theorem is also available for nonregular models such as locally asymptotically exponential models ([Ibragimov and Hasminskij, 1981], [Chernozhukov and Hong, 2004]), models with parameter on the boundary of the parameter space [Bochkina and Green, 2014] which hold under misspecified models. Here, however, we focus on “regular” models where estimators are asymptotically Gaussian.

However, under model misspecification, Bayesian model is no longer optimal, as the posterior variance does not match the minimal lower bound on the variance of unbiased estimators [Kleijn and van der Vaart, 2012] and [Panov and Spokoiny, 2015]).

Therefore, the standard way of constructing a posterior distribution following the Bayes theorem may not be appropriate for particular purposes, e.g. inference or prediction [Müller, 2013]. Different ways to construct a distribution of  $\theta$  given  $\mathbf{Y}$  that produces inference appropriate for the purpose of the analysis have been proposed. A natural aim for such a method is to behave like a standard Bayesian method when the model is well specified, i.e. when  $P_0 \in \{P(\cdot | \theta), \theta \in \Theta\}$ , and to provide at least asymptotically optimal inference, from the frequentist perspective, under model misspecification.

In this review we focus on regular misspecified models. The review is organised as follows. We start with the summary of frequentist results for regular misspecified models (Section 1.2). In Section 1.3, we formulate classical Bernstein - von Mises theorem and in Section 1.4 we discuss the analogue of Bernstein - von Mises theorem under model misspecification, particularly the conditions when this local Gaussian approximation holds. In Section 1.5, we review the proposed methods to construct distribution  $p(\theta | Y)$  that results in improved inference under model misspecification compared to the standard Bayesian approach. We conclude with discussion and open questions.

*Definitions.* For a vector  $x \in \mathbb{R}^p$ ,  $\|x\|$  denotes the Euclidean norm of  $x$ , and for a matrix  $A \in \mathbb{R}^{p \times p}$ ,  $\|A\|$  denotes the spectral (operator) norm of  $A$ .

## 1.2 Frequentist results for misspecified models

### 1.2.1 Probability model

A probability-based set up is as follows. Consider a measurable space  $(\mathcal{Y}, \mathcal{A})$  and let  $\mathcal{P}$  be a set of probability distributions on  $(\mathcal{Y}, \mathcal{A})$ , and assume that  $\mathbf{Y} \sim P(\cdot | \theta)$ ,  $\theta \in \Theta \subseteq \mathbb{R}^p$  with finite  $p$  (which may or may not be allowed to grow with  $n$ ),  $\{P(\cdot | \theta), \theta \in \Theta\} \subset \mathcal{P}$ , where  $\mathbf{Y}$  is  $n$ -dimensional random variable,  $P(\mathbf{y} | \theta)$  is the probability density function (with respect to the Lebesgue or counting measure). The true distribution of  $\mathbf{Y}$  is denoted as  $P_0$ , with density  $p_0$ .

### 1.2.2 Best parameter

Given the parametric family and the true distribution of the data, the “best” parameter is defined by

$$\theta^* = \arg \min_{\theta \in \Theta} KL(P(\cdot | \theta), P_0) \quad (1.1)$$

where  $KL(P_0, P_1) = \int \log \left( \frac{dP_0}{dP_1} \right) dP_0$  is the Kullback-Leibler divergence between probability measures  $P_0$  and  $P_1$ .

Usually, it is either assumed that the model misspecification is such that  $\theta^*$  is the parameter of interest (e.g. in machine learning or quasi-likelihood approaches, this is often done by construction). However, there are alternative approaches when  $\theta^*$  differs from the parameter of interest  $\theta_0$  (e.g. [Miller and Dunson, 2019]).

### 1.2.3 Regular models

We consider a regular setting, under the following assumptions.

1.  $\Theta$  is an open set.
2. Maximum of  $E \log p(\theta | \mathbf{Y})$  over  $\theta \in \Theta$  is attained at a single point  $\theta^*$  ((1.1)).
3.  $D(\theta)$  and  $V(\theta)$  are finite positive definite matrices for all  $\theta$  in some neighbourhood of  $\theta^*$ , where

$$\begin{aligned} V(\theta) &= E[\nabla \log p(\theta | \mathbf{Y})(\nabla \log p(\theta | \mathbf{Y}))^T], \\ D(\theta) &= -E\nabla^2 \log p(\theta | \mathbf{Y}). \end{aligned} \quad (1.2)$$

Here (and throughout the paper) the expectation is taken with respect to the true distribution of the data,  $\mathbf{Y} \sim P_0$  (which is sometimes emphasised by writing  $E_{P_0}$ ), and  $\nabla$  is the differentiation operator with respect to  $\theta$ . If the model is correctly specified,  $D(\theta^*) = V(\theta^*)$ .

Typically, two main types of models are considered in theory: independent identically distributed (iid) models:  $\mathbf{Y} = (Y_1, \dots, Y_n)$  with independent  $Y_i$  with the same pdf or pmf  $p(\cdot | \theta)$ , and more generally locally asymptotically normal (LAN) models [Van der Vaart, 2000]. We give the definition of LAN models stated in [Kleijn and van der Vaart, 2006] that applies to a misspecified model.

**Definition 1** Stochastic local asymptotic normality (LAN) condition: given an interior point  $\theta^* \in \Theta$  and a rate  $\delta_n \rightarrow 0$ , there exist random vectors  $\Delta_{n,\theta^*}$  and a nonsingular matrix  $D_0$  such that the sequence  $\Delta_{n,\theta^*}$  is bounded in probability, and for every compact set  $K \in \mathbb{R}^p$ ,

$$\sup_{h \in K} \left| \frac{\log p_{\theta^* + \delta_n h}(Y_1, \dots, Y_n)}{\log p_{\theta^*}(Y_1, \dots, Y_n)} - h^T D_0 \Delta_{n,\theta^*} - 0.5 h^T D_0 h \right| \rightarrow 0$$

as  $n \rightarrow \infty$  in (outer)  $P_0^{(n)}$ -probability.

For iid model and iid true distribution, with possibly misspecified density,  $\delta_n = 1/\sqrt{n}$  and  $D_0$  is the limit of  $\delta_n^2 D(\theta^*)$  as  $n \rightarrow \infty$  where  $D(\theta^*)$  is defined by (1.2). Also,

$$\Delta_{n,\theta^*} = n^{-1/2} D_0^{-1} \sum_{i=1}^n \nabla \log p(Y_i | \theta^*)$$

which has mean 0 and variance  $D_0^{-1} V_0 D_0^{-1}$ , called the sandwich covariance, where  $V_0$  is the limit of  $\delta_n^2 V(\theta^*)$  as  $n \rightarrow \infty$ , and  $V(\theta^*)$  is defined by (1.2).

### 1.2.4 Nonasymptotic LAN condition

[Spokoiny, 2012] provides a non-asymptotic version of LAN expansion under model misspecification, and non-asymptotic bounds on consistency of the MLE and coverage of MLE-based and likelihood-based confidence regions. These conditions have been updated in [Panov and Spokoiny, 2015].

Define the stochastic term

$$\zeta(\theta) = \log p(\mathbf{Y} | \theta) - \mathbb{E} \log p(\mathbf{Y} | \theta).$$

1. There exist  $g > 0$  and positive-definite  $p \times p$  matrix  $V$  such that for any  $|\lambda| \leq g$ ,

$$\sup_{\gamma \in \mathbb{R}^p} \mathbb{E}_{P_0} \exp \left\{ \lambda \frac{\gamma^T \nabla \zeta(\theta^*)}{\|V\gamma\|} \right\} \leq e^{\lambda^2/2},$$

If such matrix  $V$  exists, it satisfies  $\text{Var}(\nabla \zeta(\theta^*)) \leq V$ . [Spokoiny and Panov, 2020] show that this holds as long as the following condition holds for some  $\tilde{g} > 0$  and  $C \in (0, \infty)$ :

$$\sup_{u \in \mathbb{R}^p: \|u\| \leq \tilde{g}} \mathbb{E} \exp \{u^T \nabla \zeta(\theta^*)\} \leq C.$$

2. There exists  $\omega > 0$  such that for any  $\|D^{1/2}(\theta - \theta^*)\| \leq r$  and  $|\lambda| \leq g$ ,

$$\sup_{\gamma \in \mathbb{R}^p} \mathbb{E} \exp \left\{ \lambda \frac{\gamma^T (\nabla \zeta(\theta) - \nabla \zeta(\theta^*))}{\omega \|V\gamma\|} \right\} \leq e^{\lambda^2/2}.$$

In later work, [Spokoiny, 2020] assumes that  $\nabla \zeta(\theta)$  is independent of  $\theta$  by introducing an augmented model in the context of an inverse problem (the approach the author refers to as calming) thus making this condition unnecessary.

3. Conditions on  $\mathbb{E} \log p(\mathbf{Y} | \theta)$ :  $\nabla \mathbb{E} \log p(\mathbf{Y} | \theta^*) = 0$  and that the second derivative is continuous in the neighbourhood  $\|D^{1/2}(\theta - \theta^*)\| \leq r$ :

$$\|D^{-1}D(\theta) - I_p\| \leq \delta(r)$$

where  $I_p$  is  $p \times p$  identity matrix,  $D(\theta) = -\nabla^2 \mathbb{E} \log p(\mathbf{Y} | \theta)$  and  $D = D(\theta^*)$ . This condition is taken from [Panov and Spokoiny, 2015], in [Spokoiny, 2012] this condition was written in terms of  $\mathbb{E} \log p(\mathbf{Y} | \theta) - \mathbb{E} \log p(\mathbf{Y} | \theta^*)$  rather than its second derivative  $D(\theta)$  which are similar due to  $\nabla \mathbb{E} \log p(\mathbf{Y} | \theta^*) = 0$ . In later papers this condition is rewritten in terms of the moments of the third and fourth derivatives of  $\mathbb{E} \log p(\mathbf{Y} | \theta)$  [Spokoiny and Panov, 2020].

Under the conditions for the stochastic terms, using results of [Spokoiny and Panov, 2020], for any  $p \times p$  matrix  $B$  such that  $\text{trace}(BVB^T) < \infty$ , the random term can be bounded nonasymptotically as follows:

$$P(\|B(\nabla \zeta(\theta) - \nabla \zeta(\theta^*))\| \geq c_0 \omega z(B, x)) \leq 2e^{-x}$$

where  $c_0$  is an absolute constant and

$$z(B, x) = \sqrt{\text{trace}(BVB^T) + 2x^{1/2}\text{trace}((BVB^T)^2) + 2x\|BVB^T\|}.$$

Note that if these conditions hold, then for  $\|D^{1/2}(\theta - \theta^*)\| \leq r$ , with probability at least  $1 - 2e^{-x}$ ,

$$\begin{aligned} & \left| \log \frac{p(\mathbf{Y} | \theta)}{p(\mathbf{Y} | \theta^*)} - (\theta - \theta^*)^T \nabla \zeta(\theta^*) - 0.5(\theta - \theta^*)^T D(\theta - \theta^*) \right| \\ & \leq 0.5\delta(r)r^2 + c_0 r \omega z(D^{-1/2}, x), \end{aligned}$$

which is a non-asymptotic analogue of LAN expansion under a possibly misspecified model with  $\delta_n = \|D\|^{-1/2}$ ,  $D_0 = \delta_n^2 D$  (or its limit as  $\delta_n \rightarrow 0$ ),  $h = (\theta - \theta^*)/\delta_n$  with  $\|D_0^{1/2}h\| \leq \delta_n r$  and  $\Delta_{n,\theta^*} = \delta_n D_0^{-1} \nabla \zeta(\theta^*)$ .

*Example 1* Consider a model with  $p$ -dimensional  $\theta$  and iid observations  $Y_1, \dots, Y_n$ . Assume that the true observations are also iid but they may have a different true distribution with finite positive definite  $V_0 = \mathbb{E} [ [\nabla \log p(Y_i | \theta^*)]^T \nabla \log p(Y_i | \theta^*) ]$  and  $D_0 = -\mathbb{E} \nabla^2 \log p(Y_i | \theta^*)$ . Therefore,  $D = D(\theta^*) = nD_0$  and  $V = V(\theta^*) = nV_0$ . Denote  $r_0 = r/\sqrt{n}$  the radius of the local neighbourhood  $\|D_0^{1/2}(\theta - \theta^*)\| \leq r_0$ .

If  $\nabla\zeta(\theta)$  does not depend on  $\theta$  (e.g. for  $p(\cdot | \theta)$  from an exponential family with natural parameter), then  $\omega = 0$  in Condition 2 and the upper bound in the LAN condition is  $0.5nr_0^3$  which tends to 0 if  $r_0 = o(n^{-1/3})$ , or equivalently if  $r = o(n^{1/6})$ .

### 1.2.5 Optimal variance for unbiased estimators

For regular models discussed in Section 1.2.3, the lower information bound for the variance of unbiased estimators of  $\theta^*$  when the true model is unknown, is the sandwich covariance [White, 1982], [Diong et al., 2017]:

$$\text{Var}(\hat{\theta}) \geq D^{-1}VD^{-1}$$

where  $V = V(\theta^*)$  and  $D = D(\theta^*)$  as defined by (1.2). This is an analogue of the Cramer-Rao inequality for regular misspecified models. In frequentist inference, the MLE is asymptotically unbiased and its variance is approximately the sandwich covariance, i.e. inference based on the MLE for misspecified models is asymptotically efficient [White, 1982].

When the true model is known but a misspecified model is used, e.g. for computational convenience, it is possible in principle to achieve the smallest variance, inverse Fisher information, using a misspecified model with additional adjustment [Diong et al., 2017]. We illustrate it on the model used in [Stoehr and Friel, 2015] in Section 1.5.5.

## 1.3 Bernstein- von Mises theorem for correctly specified models

For a correctly specified parametric model  $\{P(\cdot | \theta), \theta \in \Theta\}$  with a density  $p(\cdot | \theta)$  and a prior distribution with density  $\pi(\theta)$ , Bayesian inference is conducted using the posterior distribution

$$p(\theta | \mathbf{y}) = \frac{p(\mathbf{y} | \theta)\pi(\theta)}{\int_{\Theta} p(\mathbf{y} | \theta)\pi(\theta)d\theta}, \quad \theta \in \Theta.$$

We formulate the Bernstein - von Mises theorem in a regular setting defined in Section 1.2.3, under the additional assumption that prior density  $\pi(\theta)$  is continuous for in a neighbourhood of  $\theta^*$ , following [Van der Vaart, 2000].

**Theorem 1 (Bernstein - von Mises theorem)** *For a well-specified regular parametric model  $\{p(\mathbf{y} | \theta), \theta \in \Theta\}$  with  $P_0 = P_{\theta^*}$  under the regularity assumptions listed in Section 1.2.3, LAN condition with  $D_0$ , for a prior density  $\pi(\theta)$  continuous for in a neighbourhood of  $\theta^*$ , then*

$$\sup_A |P(\delta_n^{-1}(\theta - \theta^*) \in A | Y_1, \dots, Y_n) - N(A; \Delta_{n,\theta^*}, D_0^{-1})| \xrightarrow{P_0^\infty} 0$$

as  $n \rightarrow \infty$ , where the supremum is taken over measurable sets  $A$ , and  $\Delta_{n,\theta^*}$  weakly converges to  $N(0, D_0^{-1})$ .

Matching variances of the posterior distribution and of the random shift  $\Delta_{n,\theta^*}$ , that are equal to the inverse Fisher information, make Bayesian inference efficient asymptotically, from the frequentist perspective. As the random variable  $\Delta_{n,\theta^*}$  is bounded with high probability and  $\delta_n \rightarrow 0$ , this theorem also implies consistency of the posterior distribution of  $\theta$ .

## 1.4 Bernstein - von Mises theorem and model misspecification

### 1.4.1 Bayesian inference under model misspecification

Given a prior distribution with density  $\pi(\theta)$ , the posterior distribution is constructed as the conditional distribution of the parameter  $\theta$  given data  $\mathbf{y}$  using Bayes theorem. A more general approach, often referred to as a Gibbs posterior distribution, is where the posterior distribution is defined using a loss function  $\ell(\theta, \mathbf{y})$  and a prior distribution with density  $\pi(\theta)$ :

$$p(\theta \mid \mathbf{y}) = \frac{e^{-\ell(\theta, \mathbf{y})} \pi(\theta)}{\int_{\theta \in \Theta} e^{-\ell(\theta, \mathbf{y})} \pi(\theta) d\theta}, \quad \theta \in \Theta.$$

If the loss function  $\ell(\theta, \mathbf{y})$  is chosen to be  $-\log p(\mathbf{y} \mid \theta)$ , this approach leads to the usual posterior distribution. As well as differing in the interpretation, the key technical difference to the classical Bayesian approach is that function  $e^{-\ell(\theta, \mathbf{y})}$  does not integrate to 1 over  $\mathbf{y}$ . This approach is used in applications where only moment conditions are available ([Chib et al., 2018], Huber function can be used as a loss for robust inference, etc.

[Bissiri et al., 2016] provide a decision-theoretical justification of this approach, by showing that this distribution minimises the following loss function with respect to probability measure  $\nu$  on  $\Theta$ ,

$$\int_{\Theta} \ell(\mathbf{y}, \theta) \nu(d\theta) + KL(\nu, \pi). \quad (1.3)$$

The authors argue that for iid observations, this is a Bayesian equivalent of

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n \ell(y_i, \theta).$$

In the latter approach the interest is in a point estimator whereas in the former approach the interest is in a distribution over  $\theta$  given  $\mathbf{y}$ .

When applying Bayesian approach under model misspecification, the key question is whether Bayesian inference remains asymptotically efficient, i.e. whether

the Bernstein - von Mises theorem holds with the posterior variance being close asymptotically to the sandwich covariance.

### 1.4.2 Concentration

A necessary condition for a Bernstein - von Mises - type result is to prove that the posterior distribution concentrates in the limit at the point mass at  $\theta^*$ .

Consistency of the posterior distribution can be defined as follows. Given a distance  $d$  between the class of probability models  $\{P_\theta, \theta \in \Theta\}$  and the true distribution  $P_0$ , as for any  $\varepsilon > 0$ ,

$$P(d(P_\theta, P_0) > \varepsilon \mid Y_1, \dots, Y_n) \xrightarrow{P_0^\infty} 0 \quad \text{as } n \rightarrow \infty.$$

Here  $P_\theta$  is the probability distribution associated with probability density (mass) function  $p_\theta$ .

For a misspecified model, the distance between  $\{P_\theta, \theta \in \Theta\}$  and  $P_0$  may be positive, so it is not always possible to achieve consistency. However, it may be possible to prove concentration at the probability model with the best parameter  $\theta^*$ : as for any  $\varepsilon > 0$ ,

$$P(d(P_\theta, P_{\theta^*}) > \varepsilon \mid Y_1, \dots, Y_n) \xrightarrow{P_0^\infty} 0 \quad \text{as } n \rightarrow \infty.$$

This is referred to as posterior concentration. It is also often of interest to prove that the posterior distribution contracts at some rate (usually the corresponding minimax rate), namely that there exists a sequence  $\varepsilon_n$  such that for any sequence  $M_n$  growing to infinity,

$$P(d(P_\theta, P_{\theta^*}) > M_n \varepsilon_n \mid Y_1, \dots, Y_n) \xrightarrow{P_0^\infty} 0 \quad \text{as } n \rightarrow \infty. \quad (1.4)$$

The main paper on posterior contraction rate under model misspecification is [Kleijn and van der Vaart, 2006]. Their results apply to nonparametric models with  $\theta = P$ . One of their conditions is formulated in terms of the covering number for testing under misspecification.

**Definition 2** Given  $\varepsilon > 0$ , define  $N_t(\varepsilon, \mathcal{P}, d, P_0, P_{\theta^*})$ , the covering number for testing under misspecification, as the minimal number  $N$  of convex sets  $B_1, \dots, B_N$  of probability measures on  $(\mathcal{Y}, \mathcal{A})$  needed to cover the set  $\{P \in \mathcal{P} : \varepsilon < d(p_\theta, p_{\theta^*}) < 2\varepsilon\}$  such that, for every  $i$ ,

$$\inf_{P \in B_i} \sup_{0 < \eta < 1} -\log E_{P_0}[p/p_{\theta^*}]^\eta \geq \varepsilon^2/4$$

If there is no finite covering of this type, the covering number is defined to be infinite.

Then, their main result (Theorem 2.1) is a typical statement on Bayesian (nonparametric) rates of posterior contraction. Their Corollary 2.1 simplifies the statement for consistency, without the rate.

**Corollary 1 (Corollary 2.1 in [Kleijn and van der Vaart, 2006])** *For a given model  $\mathcal{P}$ , prior  $\Pi$  on  $\mathcal{P}$  some  $P_{\theta^*} \in \mathcal{P}$  and a semi-metric  $d$  on  $\mathcal{P}^2$ , assume that*

1.  $KL(p(\cdot|\theta^*), p_0) < \infty$
2.  $E(p(\mathbf{Y}|\theta)/p(\mathbf{Y}|\theta^*)) < \infty$  for all  $\theta \in \Theta$
3.  $p(y|\theta^*) > 0$  for all  $y \in \mathcal{Y}$

and for every  $\varepsilon > 0$ ,

$$\Pi(B(\varepsilon, P_{\theta^*}, P_0)) > 0$$

where

$$B(\varepsilon, P_{\theta^*}, P_0) = \left\{ \theta : -E_{P_0} \log \left[ \frac{p_\theta}{p_{\theta^*}} \right] \leq \varepsilon^2, -E_{P_0} \left[ \frac{p_\theta}{p_{\theta^*}} \right]^2 \leq \varepsilon^2 \right\},$$

and

$$\sup_{\eta > \varepsilon} N_t(\eta, \mathcal{P}, d, P_0, P_{\theta^*}) < \infty.$$

Then, for every  $\varepsilon > 0$ , as  $n \rightarrow \infty$ ,

$$P(d(P_\theta, P_{\theta^*}) > \varepsilon \mid Y_1, \dots, Y_n) \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

The authors also consider the case when the best approximation  $\theta^*$  is not unique. Their results are illustrated on consistency of density estimation using mixture models, and on nonparametric regression models using a convex set of prior models for the regression function.

[Grünwald, 2012] demonstrates that the convexity of  $\mathcal{P}$  is crucial. [Grünwald and van Ommen, 2017] show via simulations that if  $\{P_\theta, \theta \in \Theta\}$  is not convex, the posterior distribution does not concentrate on  $\theta^*$  but instead it concentrates on the best approximation of  $P_0$  in the convex hull of the class of the parametric models  $Conv(\{P_\theta, \theta \in \Theta\})$ :

$$\tilde{P} = \arg \min_{P \in Conv(\{P_\theta, \theta \in \Theta\})} KL(P_0, P). \quad (1.5)$$

We discuss their example in more detail in Section 1.4.4. In particular, the authors say that it is possible to achieve consistency, i.e. for the posterior distribution to converge to the point mass at  $\theta^*$  if for any  $\eta \in (0, 1]$ ,

$$E_{P_0} \left( \frac{dP(\mathbf{Y}|\theta)}{dP(\mathbf{Y}|\theta^*)} \right)^\eta \leq 1 \quad \text{for all } \theta \in \Theta. \quad (1.6)$$

Note that this condition is reminiscent of one of the conditions of [Kleijn and van der Vaart, 2006] who assume that the above condition holds for  $\eta = 1$ , and a similar condition is present in the definition of the covering numbers under model misspecification. Fur-

ther, [Grünwald and Mehta, 2020] relax this condition to be upper bounded by  $1 + u$  for some small  $u > 0$ .

[Bhattacharya et al., 2019] study the posterior contraction rate for a particular type of semi-metric  $d$  that is matched to the considered misspecified model. For such a matched semi-metric, they show that the posterior contraction rate is determined only by the prior mass condition of the posterior contraction theorem of Ghoshal and van der Vaart (2007), and does not involve the entropy condition. In their setting, the pseudo-likelihood is a power of a probability density, so condition (1.6) holds. The authors consider only examples of misspecified models with a convex parameter space. [Syring and Martin, 2020] study the concentration rate of Gibbs posteriors in a (semi-metric)  $d$  under more general losses and semi-metrics, focusing on iid models and iid true distribution, and discuss a setting where  $Y_i$ 's are independent but not necessarily identically distributed. Their assumptions also include the condition on the prior mass of KL neighbourhood but not the entropy; they use a different additional assumption instead. See Section 1.5.3.2 for details.

### 1.4.3 Bernstein - von Mises - type results under model misspecification

The first Bernstein - von Mises - type result under model misspecification was formulated by [Kleijn and van der Vaart, 2012]. The authors state that for misspecified LAN models (see Definition 1), under assumptions of Theorem 2.1 in [Kleijn and van der Vaart, 2006] with rate  $\delta_n$ ,

$$\sup_A |P(\delta_n^{-1}(\theta - \theta^*) \in A \mid Y_1, \dots, Y_n) - N(A; \Delta_{n, \theta^*}, D_0^{-1})| \xrightarrow{P_0^\infty} 0$$

as  $n \rightarrow \infty$ , i.e. the posterior distribution converges to the Gaussian distribution in the total variation distance.

[Panov and Spokoiny, 2015] state the BvM for semi-parametric possibly misspecified models, with flat or a Gaussian prior distribution, in a non-asymptotic setting. Here we only state conditions and statement for a parametric model.

In addition to the non-asymptotic LAN assumptions stated in Section 1.2.3, the following assumptions are made.

1. *Small bias condition:* the norm of the bias of the penalised estimator  $\|\theta_\pi^* - \theta^*\|$  is small where  $\theta_\pi^*$  is defined by

$$\begin{aligned} \theta_\pi^* &= \arg \min_{\theta \in \Theta} [KL(p(\cdot | \theta), p_0) - \log \pi(\theta)] \\ &= \arg \min_{\theta \in \Theta} [-E \log p(\mathbf{Y} | \theta) - \log \pi(\theta)]. \end{aligned} \quad (1.7)$$

2. *Identifiability:*  $\|D^{-1}V\| \leq a^2 \in (0, \infty)$ .
3. *Global deterministic condition:* for  $\|D^{1/2}(\theta - \theta^*)\| > r$ ,

$$E \log p(\mathbf{Y} | \theta) - E \log p(\mathbf{Y} | \theta^*) \geq -\|D^{1/2}(\theta - \theta^*)\|b(r)$$

with  $b(r)$  growing to infinity as  $r$  grows to infinity. [Panov and Spokoiny, 2015] have a stronger condition however, it is possible to show that this condition is sufficient to bound the tail of the posterior distribution on  $\|D^{1/2}(\theta - \theta^*)\| > r$  for large  $r$ .

4. *Global stochastic condition:* for any  $r > 0$  there exists  $g(r) > 0$  such that for any  $|\lambda| \leq g(r)$ ,

$$\sup_{\|D^{1/2}(\theta - \theta^*)\| \leq r} \sup_{\gamma \in \mathbb{R}^p} \mathbb{E} \exp \left\{ \lambda \frac{\gamma^T \nabla \zeta(\theta)}{\|V\gamma\|} \right\} \leq e^{\lambda^2/2},$$

where the expectation is taken with respect to  $Y \sim P_0$ .

Under these assumptions, [Panov and Spokoiny, 2015] formulated a non-asymptotic version of Bernstein - von Mises theorem for misspecified models.

**Theorem 2 (Theorem 1 in [Panov and Spokoiny, 2015])** *Suppose that the assumptions stated in this section hold, and consider a flat prior  $\pi(\theta) = 1$  for all  $\theta \in \Theta$ .*

*Then, for any measurable  $A$ , with probability at least  $1 - 4e^{-x}$ ,*

$$|P(\tilde{D}^{1/2}(\theta - \theta^* - \Delta_{\star,n}) \in A \mid \mathbf{Y}) - N(A; 0, I_p)| \leq e^{\tilde{\Delta}(r)} - 1, \quad (1.8)$$

where  $\hat{\theta}$  is the MLE of  $\theta$ ,  $\tilde{D} = D(\hat{\theta})$ ,  $\Delta_{\star,n} = \tilde{D}^{-1} \nabla \zeta(\theta^*)$  and  $\tilde{\Delta}(r) = r^2(\delta(r) + 6\omega z(\tilde{D}^{-1/2}, x)) + 8e^{-x}$ .

The authors also prove a similar result with posterior mean and posterior precision matrix instead of  $\theta^* + \Delta_{\star,n}$  and  $\tilde{D}$ .

In their Theorem 2, the authors extend this result to a Gaussian prior  $\theta \sim N(0, G^{-2})$  such that

$$\|G^2 D^{-1}\| \leq \epsilon < 1/2, \quad \text{trace}[(G^2 D^{-1})^2] \leq \delta^2, \\ \|(D + G^2)^{-1} G^2 \theta^*\| \leq \beta.$$

Then, the posterior is approximated by the Gaussian distribution, namely, equation (1.8) holds with the upper bound replaced by  $e^{2\tilde{\Delta}(r)+8e^{-x}}(1 + \tau) + e^{-x} - 1$  where

$$\tau = 0.5 \left[ (1 + \epsilon)(3\beta + \epsilon z(D^{-1/2}, x))^2 + \delta^2 \right]^{1/2}.$$

In particular, the authors show that for a high dimensional parameter, the upper bound is small if  $x$  is large and  $p^3/n$  is small. [Spokoiny and Panov, 2020] have shown a similar result under assumption that the stochastic term is a constant, with posterior distribution centered either at the MLE  $\hat{\theta}$  or at the posterior mean. The authors also apply these results to nonparametric problems.

[Chib et al., 2018] study Bayesian exponentially-tilted empirical likelihood posterior distributions, which are defined by moment conditions rather than by a likelihood or loss function. The authors show the BvM result for well-specified and misspecified models under fairly general conditions.

#### 1.4.4 Example: misspecified linear model

Now we consider the example of a misspecified model given in [Grünwald and van Ommen, 2017] where the Bayesian approach considered by the authors fails, and we apply theory of [Panov and Spokoiny, 2015] to analyse it. In particular, we check whether the small bias condition holds, i.e. whether  $\theta_\pi^*$  defined by (1.7) is close to  $\theta^*$  defined by (1.1).

The authors considered the following linear model

$$Y_i = \beta_0 + \sum_{j=1}^p \beta_j X_{ij} + \epsilon_i, \quad \epsilon_i \sim N(0, \sigma^2), \quad i = 1, \dots, n \quad (1.9)$$

independently, with a conjugate prior distribution on its parameters:

$$\beta = (\beta_0, \beta_1, \dots, \beta_p)^T \sim N(0, c^{-1} \sigma^2 G), \quad \tau := \sigma^{-2} \sim \Gamma(a, b) \quad (1.10)$$

independently, with  $G = I_{p+1}$ . The values of the hyperparameters were chosen to be  $c = 1$ ,  $a = 1$ ,  $b = 40$ .

The true distribution of the data, i.e. the data generating mechanism, is as follows:

$$\begin{aligned} X_{ij} &\sim N(0, 1), \quad Z_i \sim \text{Bern}(0.5), \\ Y_{ij} | Z_i &\sim N\left(\sum_{j=1}^p \beta_{true,j} X_{ij}, \sigma_{true}^2(1 + Z_i)\right), \end{aligned} \quad (1.11)$$

independently for  $i = 1, \dots, n$  and  $j = 1, \dots, p$ . The true values were taken as  $n = 100$ ,  $p = 40$ ,  $\sigma_{true}^2 = 1/40$ ,  $\beta_{true,j} = 0.1$  for  $j = 1, 2, 3, 4$  and  $\beta_{true,j} = 0$  otherwise. Note that  $\beta_{true,0} = 0$ , i.e. there is no intercept, and that the variance of  $Y_i$  given  $X$  is  $\sigma_{true}^2(0.5 \cdot 1 + 0.5 \cdot 2) = 1.5\sigma_{true}^2$ .

Now we work out the best parameter for this model defined by (1.1) and the point at which posterior distribution concentrates asymptotically (1.7), and whether they are close or not.

The log likelihood for the considered model is

$$L(\beta, \tau) = -0.5\tau(Y - X\beta)^T(Y - X\beta) + 0.5n \log \tau,$$

and logarithm of the posterior distribution of  $\theta = (\beta, \tau)$  is

$$L_\pi(\beta, \tau) = L(\beta, \tau) - 0.5c\tau\beta^T\beta + 0.5p \log \tau + (a - 1) \log \tau - b\tau.$$

Negative Kullback-Leibler distance  $KL(p_\theta, p_0)$  (up to an additive constant independent of unknown parameters), is

$$EL(\beta, \tau) = 0.5n \log \tau - 0.5\tau \|\beta - \beta_{true}\|_2^2 - \frac{3}{4}n\sigma_{true}^2\tau \quad (1.12)$$

where the expectation is taken under the true model, using  $E(Y^T Y | X) = \beta_{true}^T X^T X \beta_{true} + 1.5n\sigma_{true}^2$  and  $EX^T X = I_{p+1}$ . Then, the best parameter, i.e. the parameter maximising expression (1.12) is

$$\beta^* = \beta_{true}, \quad \tau^{*-1} = \sigma^{*2} = \frac{3}{2}\sigma_{true}^2, \quad (1.13)$$

as stated in [Grünwald and van Ommen, 2017].

Now we study the value of the parameters where the posterior concentrates which minimises

$$EL_{\pi}(\beta, \tau) = (0.5(n+p) + a - 1) \log \tau - \tau b - \frac{3}{4}n\sigma_{true}^2 \tau \quad (1.14)$$

$$-0.5\tau(1/c + 1)^{-1} \|\beta_{true}\|_2^2 \quad (1.15)$$

$$-0.5(1+c)\tau \|\beta - \beta_{true}/(1+c)\|_2^2$$

and which are equal to

$$\beta_{\pi}^* = (1+c)^{-1} \beta_{true}, \quad (1.16)$$

$$\tau_{\pi}^{*-1} = \sigma_{\pi}^{*2} = \frac{1.5\sigma_{true}^2 + 2b/n + c(1+c)^{-1}n^{-1} \|\beta_{true}\|_2^2}{1 + p/n + 2(a-1)/n}.$$

Hence,  $(\beta_{\pi}^*, \sigma_{\pi}^{*2})$  is close to  $(\beta^*, \sigma^{*2})$  if the following conditions hold:

1.  $c = o(1)$
2.  $b/n = o(\sigma_{true}^2)$
3.  $c(1+c)^{-1}n^{-1} \|\beta_{true}\|_2^2 = o(\sigma_{true}^2)$
4.  $p/n = o(1)$
5.  $(a-1)/n = o(1)$ .

The choice of the parameters given in [Grünwald and van Ommen, 2017] is the following:

1.  $c = 1$
2.  $b/n = 0.4$ ,  $\sigma_{true}^2 = 0.025$
3.  $c(1+c)^{-1}n^{-1} \|\beta_{true}\|_2^2 = 0.0002$ ,
4.  $p/n = 0.5$
5.  $(a-1)/n = 0$

i.e. conditions 3 and 5 hold whereas the remaining conditions do not hold. So, it is possible to tune hyperparameters so that all conditions, except condition 4, hold, e.g. by taking small  $b$  and  $c$  leading to weakly informative priors with large variances. Condition  $p/n = o(1)$  is due to the choice of the conjugate prior for  $\beta$  with its prior variance proportional to the variance of the noise; if the prior variance of  $\beta$  does not depend on the noise variance, then this condition is not necessary.

For instance, it is easy to show using the same technique, that considering a non-conjugate prior  $\beta \sim N(0, \tau_0^{-1}\Sigma_0)$  and  $\tau \sim \Gamma(a, b)$ , with  $\|\Sigma_0\| = 1$ , under the following conditions

1.  $\tau_0 = o(1)$
2.  $b/n = o(\sigma_{true}^2)$
3.  $(a - 1)/n = o(1)$

leads to  $\theta_\pi^*$  being close to  $\theta^*$ . These conditions are satisfied e.g. with small  $\tau_0$ , small  $b$  and  $a = 1$ , provided  $\sigma_{true}^2$  is not much smaller than  $1/n$ .

## 1.5 “Optimising” Bayesian inference under model misspecification

### 1.5.1 Asymptotic risk of parameter estimation under a misspecified model

[Müller, 2013] showed that the asymptotic frequentist risk associated with misspecified Bayesian estimators is inferior to that of an artificial posterior which is normally distributed, centred at the maximum likelihood estimator and with the sandwich covariance matrix.

This provided theoretical justification for constructing a (quasi-) posterior distribution based on a misspecified model such that its posterior variance is approximately the sandwich covariance. Several such approaches have been used that we discuss below.

### 1.5.2 Composite likelihoods

Composite likelihoods (also known as pseudo-likelihoods) have been studied by [Lindsay, 1988], and they are defined as follows. Denote by  $\{A_1, \dots, A_K\}$  a set of marginal or conditional events with associated likelihoods  $L_k(\theta; y) \propto P(y \in A_k; \theta)$ . Then, a composite likelihood is the weighted product

$$L_c(\theta; y) = \prod_{k=1}^K [L_k(\theta; y)]^{w_k},$$

where  $w_k$  are nonnegative weights to be chosen. It is often used to simplify the model for dependence structure in time series and in spatial models, with one of the most famous examples given by [Besag, 1986] of approximating spatial dependence by the product of conditional densities of a single observation given its neighbours. Selection of unequal weights to improve efficiency in the context of particular applications and a review of frequentist inference for this approach is discussed by [Varin et al., 2011]. For the discussion of connection of the choice of weights with the Bayesian inference under empirical likelihood see [Schennach, 2005]. A typical example is when  $L_k(\theta; y)$  is the marginal likelihood for  $y_k$  (with  $K = n$ ). Unless

more information is available, it is generally difficult to estimate individual weights from the sample however this formulation gave rise to a number of approaches with randomised weights  $(w_1, \dots, w_n)$ .

The idea of composite likelihood is used to sample the powers (weights of the contributions of individual samples  $w_i$ ) from some probability distribution. The typical choice of a joint Dirichlet distribution for the weights corresponds to Bayesian bootstrap and is discussed in Section 1.5.4. Other choices of weights and their effect on the corresponding posterior inference are discussed in [Wang et al., 2017]. As far as I am aware, currently there are no BvM results for other randomisation schemes, apart from a joint Dirichlet distribution.

Choosing the same weight  $w_k = w$  for all  $k$  leads to fractional or tempered posterior distributions discussed in Section 1.5.3.

## 1.5.3 Generalised (Gibbs) posterior distribution

### 1.5.3.1 Definition and interpretation

Let  $\ell(\mathbf{y}, \theta)$  be some loss function. Then, generalised posterior distribution is given by

$$p_\eta(\theta \mid y_1, \dots, y_n) = \frac{\exp\{-\eta\ell(\mathbf{y}, \theta)\}\pi(\theta)}{\int \exp\{-\eta\ell(\mathbf{y}, \theta)\}\pi(\theta)d\theta} \quad (1.17)$$

where  $\eta$  is the parameter that adjusts for misspecification. This parameter is called the learning rate (in machine learning), inverse temperature. Taking  $\ell(\mathbf{y}, \theta) = \sum_{i=1}^n \ell_i(y_i, \theta)$  for some loss  $\ell_i$  associated with observation  $y_i$  given parameter  $\theta$  corresponds to the assumption that observations  $y_i$  are independent. Taking  $\ell(\mathbf{y}, \theta) = -\log p(\mathbf{y} \mid \theta)$  and  $\eta = 1$  leads to the classical Bayesian inference. Different functions  $\ell$  may be used for different types of model misspecification and different inference purposes, e.g. Huber function for robust parameter estimation.

This is also known as a Gibbs posterior in Bayesian literature, exponential weighting in frequentist literature [Dalalyan and Tsybakov, 2008], typically with  $\ell_i = \|y_i - \theta\|_2^2$ , and it is used as a model for PAC-Bayesian approach in machine learning. Lately it has also been referred to as a fractional posterior and as a tempered posterior.

[Grünwald and van Ommen, 2017] argue that if there exists  $\bar{\eta} \leq 1$  such that for all  $0 < \eta \leq \bar{\eta}$

$$\int p_0(\mathbf{y}) \left( \frac{p(\mathbf{y} \mid \theta)}{p(\mathbf{y} \mid \theta^*)} \right)^\eta d\mathbf{y} \leq 1 \quad \text{for all } \theta \in \Theta,$$

then the generalised posterior with  $\eta < \bar{\eta}$  is asymptotically consistent, i.e. converges to the point mass at  $p_{\theta^*}$ . For  $\eta$  such that condition (1.6) holds, the authors interpret the generalised posterior as a posterior distribution based on the reweighted true likelihood:

$$\tilde{p}_{\eta, \theta^*}(\mathbf{y} | \theta) = p_0(\mathbf{y}) \left( \frac{p(\mathbf{y} | \theta)}{p(\mathbf{y} | \theta^*)} \right)^\eta$$

which is interpreted as a density on the probability space augmented by an unobserved event if this function integrates to a positive value less than 1. This is due to the following: if this density was used as a density of  $y$  given  $\theta$  to construct the likelihood, then this would correspond to a correctly specified model, since the KL distance between  $\tilde{p}_{\eta, \theta^*}(\mathbf{y} | \theta)$  and  $p_0$  is minimised at  $\theta^*$  and  $\tilde{p}_{\eta, \theta^*}(\cdot | \theta^*) = p_0(\cdot)$ , and the corresponding posterior would be a proper posterior and it coincides with the generalised posterior. This is done for interpretation only, as it is not possible to use  $\tilde{p}_{\eta, \theta^*}$  for inference in practice due to unknown  $p_0$  and  $\theta^*$ .

In the following section we discuss known results about concentration of the Gibbs posterior distribution.

### 1.5.3.2 Concentration and posterior contraction rate

[Bhattacharya et al., 2019] study the posterior contraction rate for a particular type of semi-metric  $d$  that is matched to the considered misspecified model. They consider generalised Bayesian approach

$$p_\eta(\theta | \mathbf{y}) = \frac{[p(\mathbf{y} | \theta)]^\eta \pi(\theta)}{\int [p(\mathbf{y} | \theta)]^\eta \pi(\theta) d\theta}$$

with  $p(y | \theta)$  being a density of a probability measure with respect to some measure  $\mu$ , and the semi-metric based on Renyi divergence with *matching* index  $\eta$ :

$$D_\eta^{(n)}(\theta, \theta^*) = -\frac{1}{1-\eta} \log A_\eta^{(n)}(\theta, \theta^*)$$

where  $A_\eta^{(n)}(\theta, \theta^*)$  is the integral defined in (1.6) for all  $n$  observation  $y = (y_1, \dots, y_n)$ :

$$A_\eta^{(n)}(\theta, \theta^*) = E_{p_0} \left[ \left( \frac{p(\mathbf{Y} | \theta)}{p(\mathbf{Y} | \theta^*)} \right)^\eta \right].$$

The authors show that since  $p(y | \theta)$  is a density of a probability measure, condition  $A_\eta^{(n)}(\theta, \theta^*) \leq 1$  (equation (1.6)) holds. Also, the authors show that for  $\eta \rightarrow 1-$ , the generalised posterior converges to the corresponding posterior distribution. Therefore, their approach is not shown to apply to so called Gibbs posteriors where other loss functions (rather than a negative log density) can be used to specify the (pseudo-)likelihood.

Following [Grünwald and van Ommen, 2017], one may argue that in the setting considered by [Bhattacharya et al., 2019], it is not necessary to use  $\eta < 1$  to adjust the inference to achieve posterior consistency (it may be necessary e.g. to achieve asymptotic efficiency). Under the assumptions of [Bhattacharya et al., 2019],  $A_\eta^{(n)}(\theta, \theta^*) \leq 1$  for all  $\theta$  and  $\eta \in (0, 1)$ , and, due to convergence of the generalised

posterior to the posterior as  $\eta \rightarrow 1-$ , the posterior distribution (with  $\eta = 1$ ) is asymptotically consistent.

[Syring and Martin, 2020] study the concentration rate of a Gibbs posterior (1.17) in (semi-metric)  $d$  under the following fairly general assumptions. The authors focus on iid models and iid true distribution, and discuss a setting where  $Y_i$ 's are independent but not necessarily identically distributed.

**Condition 1.** *There exist  $\bar{\eta}, K, r > 0$  such that for all  $\eta \in (0, \bar{\eta})$  and for all sufficiently small  $\delta > 0$ , for  $\theta \in \Theta$ ,*

$$d(\theta, \theta^*) > \delta \Rightarrow \log E \exp(-\eta(\ell(Y_1; \theta) - \ell(Y_1; \theta^*; Y_1))) < -K\eta\delta^r.$$

**KL neighbourhood condition.** *For a sequence  $(\varepsilon_n)$  such that  $\varepsilon_n \rightarrow 0$  and  $n\varepsilon_n^r \rightarrow \infty$  as  $n \rightarrow \infty$ , there exists  $C_1 \in (0, \infty)$  such that for all  $n$  large enough,*

$$\log \Pi(B_{KL}(\varepsilon^r)) > -C_1 n \varepsilon_n^r,$$

where

$$B_{KL}(R) = \left\{ \theta \in \Theta : \begin{aligned} -E[\ell(\mathbf{Y}; \theta^*) - \ell(\mathbf{Y}; \theta)] &\leq R \\ \text{Var}[\ell(\mathbf{Y}; \theta^*) - \ell(\mathbf{Y}; \theta)] &\leq R \end{aligned} \right\}.$$

**Theorem 3 (Theorem 3.2, [Syring and Martin, 2020])** *Under Condition 1 and KL neighbourhood condition, for a fixed  $\eta$ , the Gibbs posterior distribution defined by (1.17) satisfies (1.4) with asymptotic concentration rate  $\varepsilon_n$ .*

The authors show that this also holds for  $\eta_n \rightarrow 0$  as long as  $\eta_n n \varepsilon_n^r \rightarrow \infty$  and in the KL neighbourhood condition  $C_1 n \varepsilon_n^r$  is replaced by  $C_1 \eta_n n \varepsilon_n^r$ . The also how that this holds for a random  $\hat{\eta}$  as long as with high probability  $c^{-1} \eta_n \leq \hat{\eta} \leq c \eta_n$  for  $\eta_n \rightarrow 0$  and some  $c \geq 1$ .

The authors also discuss that condition (1.6) can be relaxed to hold on  $\Theta_n = \{\theta \in \Theta : \|\theta\| \leq \Delta_n\}$  for a sequence  $(\Delta_n)$  increasing to infinity, under stronger conditions (see Theorem 4.1 in [Syring and Martin, 2020]). The authors also discuss that conditions of this theorem are related to the entropy condition of [Kleijn and van der Vaart, 2006] and verify this condition for convex  $\ell$  as a function of  $\theta$ .

In the iid setting, Condition 1 combines several conditions of [Spokoiny, 2012] for a single observation  $Y_1$  since

$$\begin{aligned} \log E \exp(-\eta(\ell(Y_1; \theta) - \ell(Y_1; \theta^*))) &= \\ -\eta(E\ell(Y_1; \theta) - E\ell(Y_1; \theta^*)) + \log E \exp(\eta(\zeta_1(\theta) - \zeta_1(\theta^*))), \end{aligned} \quad (1.18)$$

where  $\zeta_1(\theta) = E\ell(Y_1, \theta) - \ell(Y_1, \theta)$ , except that the authors assume that this condition holds for all  $\theta \in \Theta$  whereas in [Spokoiny, 2012] the conditions are split into local (in

a neighbourhood of  $\theta^*$ ) and global (for all  $\theta \in \Theta$ ), with the global conditions being weaker.

As the authors discuss, their Condition 1 can hold if  $\zeta_1(\theta) - \zeta_1(\theta^*)$  has sub-Gaussian tails, and if for  $\eta$  small enough the first term (which is negative) in (1.5.3.2) is sufficiently greater in absolute value than the second term. More specifically, assume that there exist  $r, K_1 > 0$  such that

$$-[\mathbb{E}\ell(Y_1; \theta) - \mathbb{E}\ell(Y_1; \theta^*)] < -K_1[d(\theta, \theta^*)]^r, \quad \theta, \theta^* \in \Theta,$$

and that the sub-Gaussian tail condition holds with some  $b > 0$

$$\log \mathbb{E} \exp(\eta(\zeta_1(\theta) - \zeta_1(\theta^*))) \leq b\eta^2 \|\theta - \theta^*\|_2^2/2$$

which can be verified through Conditions 1 and 2 [Spokoiny, 2012], and there exist  $K_2 > 0$  such that  $\|\theta - \theta^*\|_2^2 \leq K_2[d(\theta, \theta^*)]^r$  for all  $\theta, \theta^* \in \Theta$ . Then,

$$\begin{aligned} \log \mathbb{E} \exp(-\eta(\ell(Y_1; \theta) - \ell(Y_1; \theta^*))) &\leq -\eta[d(\theta, \theta^*)]^r [K_1 - b\eta K_2/2] \\ &< -K\eta[d(\theta, \theta^*)]^r \end{aligned}$$

if  $\eta < \bar{\eta} = \min(1 + o(1), 2K_1/(bK_2))$  and  $K = K_1 - b\bar{\eta}K_2/2$ . Since the inequality  $\eta < \bar{\eta}$  is strict, as long as  $2K_1/(bK_2) > 1$ , we can take  $\eta = 1$ .

The authors suggest that case  $r = 2$  corresponds to regular problems, i.e. where  $\nabla \mathbb{E}\ell(\theta^*) = 0$  and  $\nabla^2 \mathbb{E}\ell(\theta^*)$  is positive and continuous in the neighbourhood of  $\theta^*$ , and the sub-Gaussian tails condition, whereas nonregular problems may require other values of  $r$ , e.g.  $r = 1$  if  $\theta^*$  is on the boundary of the parameter space [Bochkina and Green, 2014], or if there is a finite jump at  $\theta$  [Chernozhukov and Hong, 2004]. We illustrate the latter on an example.

*Example 2* Now we check if Condition 1 holds for a density with jump. Consider a density  $p(y|\theta)$  that is 0 for  $y < \theta$  and the right hand side limit  $\lim_{y \rightarrow \theta^+} p(y|\theta) = \lambda > 0$ , for instance with  $p(y|\theta) = e^{-(y-\theta)}$  for  $y > \theta$ , and the true density  $p_0(y)$  such that  $p_0(y) = 0$  for  $y < \theta_0$  and  $\lim_{y \rightarrow \theta_0^+} p_0(y) = c_0 > 0$ . Then,

$$\begin{aligned} \mathbb{E} \log p(Y | \theta) &= - \int_{\theta_0}^{\infty} (y - \theta) p_0(y) dy + \log(0) I(\theta > \theta_0) \\ &= \theta - \mathbb{E}Y + \log(0) I(\theta > \theta_0) \end{aligned}$$

which is minimised at  $\theta^* = \theta_0$ . This implies that for  $\theta \in \Theta_0 = (-\infty, \theta^*]$ ,

$$\mathbb{E} \log p(Y | \theta) - \mathbb{E} \log p(Y | \theta^*) = \theta - \theta^*$$

So if  $d(x, y) = |x - y|$  then this holds with  $K_1 = 1$  and  $r = 1$ .

The stochastic term is  $\zeta(\theta) = Y - \mathbb{E}Y$  for  $\theta \in \Theta_0$ , and for  $\theta \in \Theta_0$

$$\log \mathbb{E} \exp(\eta(\zeta(\theta) - \zeta(\theta^*))) = 0,$$

and Condition 1 holds for  $\theta \in (-\infty, \theta^*]$  with  $K = 1$  and  $r = 1$ .

### 1.5.3.3 Estimation of $\eta$

There are various approaches to estimation of  $\eta$  that lead to the posterior distribution concentrating at  $\theta^*$ , e.g. [Grünwald and Mehta, 2020] and [Ribatet et al., 2012]; see a review [Wu and Martin, 2020]. Here I will give a very brief discussion. There are two key issues: firstly, this parameter models misspecification so it cannot be estimated in a usual Bayesian way (e.g. by putting a hyperprior), and secondly, a relevant estimator depends on the aim of the inference.

1. When predictive inference is of interest, the Safe-Bayes estimator of [Grünwald, 2012] further explored in [Grünwald and Mehta, 2020], may be appropriate:

$$\hat{\eta} = \arg \min \left[ - \sum_{i=1}^n \mathbb{E} \log p(y_i | \theta) p_{\eta}(\theta | y_{1:(i-1)}) d\theta \right]$$

where  $p_{\eta}(\theta | y_{1:(i-1)})$  is the generalised posterior distribution with parameter  $\eta$  based on  $i - 1$  samples (if  $i = 1$  then it is the prior distribution).

2. Now we discuss estimators of  $\eta$  when estimation of  $\theta$  is of interest, in particular frequentist coverage of credible posterior regions  $C_{\alpha}$  such that  $P_{\eta}(\theta \in C_{\alpha} | \mathbf{Y}) = 1 - \alpha$ .

Under the conditions of Gaussian approximation of the posterior, if asymptotic coverage of asymptotic credible balls  $C_{\alpha}$  is of interest, then it is sufficient to check that asymptotic credible balls

$$(\theta - \hat{\theta})^T D_{\eta}(\theta - \hat{\theta})^T \leq \chi_p^2(\alpha)$$

are inside the frequentist confidence balls with the sandwich covariance

$$(\theta - \hat{\theta})^T DV^{-1}D(\theta - \hat{\theta})^T \leq \chi_p^2(\alpha),$$

i.e. it is sufficient to check that the largest eigenvalue of  $D_{\eta} = \eta D$  is not smaller than the largest eigenvalue of  $DV^{-1}D$ .

Therefore,  $\eta$  is chosen so that the largest eigenvalue of the posterior precision matrix  $\eta D$  matches that largest eigenvalue of the sandwich precision matrix  $DV^{-1}D$ , i.e. the ‘‘oracle’’ value is

$$\eta^* = \frac{\|DV^{-1}D\|}{\|D\|},$$

and it can be estimated if estimates of  $V = V(\theta^*)$  and  $D = D(\theta^*)$  are available.

[Holmes and Walker, 2017] used the Fisher information number to calibrate this parameter:

$$\eta^* = \frac{\text{trace}(DV^{-1}D)}{\text{trace}(D)}.$$

As the authors say, it is the sum of the marginal Fisher information for each dimension, which can be used as a summary for the amount of information in a sample about parameters.

[Pauli et al., 2012] propose  $\hat{\eta} = \text{trace}(V^{-1}(\hat{\theta}_c)D(\hat{\theta}_c))/\dim(\theta)$  which asymptotically is the average of the mutual eigenvalues of  $D(\hat{\theta}_c)$  with respect to  $V(\hat{\theta}_c)$ .

*Remark 1* Suppose that conditions of [Panov and Spokoiny, 2015] hold for some pseudo likelihood  $p(\mathbf{y} | \theta) = \prod_{i=1}^n p(y_i | \theta)$  and  $\theta \in \Theta \subseteq \mathbb{R}$ . Then, these conditions hold for  $p(\mathbf{y} | \theta)^\eta = \prod_{i=1}^n p(y_i | \theta)^\eta$  with  $V_\eta = \eta^2 V$  and  $D_\eta = \eta D$ ,  $r_\eta = r\eta$  and  $b_\eta$  such that  $b_\eta(r_\eta) = b(r)$ .

Therefore, the posterior variance  $D_\eta^{-1}$  coincides with the sandwich variance if  $\eta D = D^2 V^{-1}$ , i.e. if  $\eta = DV^{-1}$ . See also [de Heide et al., 2020] in the context of linear regression.

## 1.5.4 Nonparametric model for uncertainty in $p_0$ and bootstrap posterior

### 1.5.4.1 Nonparametric model and connection to bootstrap

[Lyddon et al., 2019] proposed to take into the account uncertainty about the parametric model by modelling the distribution of the data nonparametrically, e.g. using a Dirichlet process prior with the base model being the considered parametric model:

$$y_i \sim F, \quad i = 1, \dots, F \sim DP(\alpha, p(\cdot | \theta)).$$

For iid observations  $y_i$ , when  $\alpha \rightarrow 0$ , this approach corresponds to Bayesian bootstrap [Rubin, 1981], with the following sampling of  $(\theta^{(j)})_{j=1}^B$  from the bootstrap posterior:

$$\theta^{(j)} = \theta(F_j) \text{ with } F_j(x) = \sum_{i=1}^n \alpha_{ji} \delta_{y_i}(x) \quad (1.19)$$

$$\text{and } \alpha_j = (\alpha_{j1}, \dots, \alpha_{jn}) \sim \text{Dirichlet}(1, \dots, 1), \quad (1.20)$$

where

$$\theta(F) = \arg \min_{\theta \in \Theta} \int \ell(\theta, y) dF(y) \quad (1.21)$$

with  $\ell(\theta, y) = -\log p(y | \theta)$ . More generally, for a possibly different loss function, the authors refer to this as the loss-likelihood (LL) bootstrap approach. The authors argue that using this procedure induces a prior distribution on  $\theta$  defined as  $P(\theta \in A) = P(F : \theta(F) \in A)$ .

### 1.5.4.2 Asymptotic normality of bootstrap posterior

[Lyddon et al., 2019] show that the sample from the loss-likelihood bootstrap has asymptotically normal distribution with sandwich covariance matrix, weakly, under the following assumptions.

1.  $\Theta$  is a compact and convex subset of a  $p$ -dimensional Euclidean space.
2. The loss function  $\ell : \Theta \times \mathbb{R} \rightarrow \mathbb{R}$  is a measurable bounded from below function, with

$$\int \ell(\theta, y)p_0(y)dy < \infty \quad \text{for all } \theta \in \Theta$$

3. (Identifiability). There exists a unique minimizing parameter value

$$\theta^* = \arg \min_{\theta \in \Theta} \int \ell(\theta, y)p_0(y)dy,$$

and for all  $\delta > 0$  where exists  $\epsilon > 0$  such that

$$\liminf_n P\left(\sup_{|\theta - \theta^*| > \delta} \frac{1}{n} \sum_{i=1}^n [\ell(\theta, y_i) - \ell(\theta^*, y_i)] > \epsilon\right) = 1$$

4. Smoothness of loss: there exists an open ball  $B$  containing  $\theta^*$  such that

$$E[|\nabla_{I_k}^k \ell(\theta, Y)|] < \infty \text{ and } E[|\nabla_j \ell(\theta, Y) \nabla_{km}^2 \ell(\theta, Y)|] < \infty$$

for  $k = 1, 2, 3$  and for all corresponding indices  $I_k$ , i.e.  $I_1 \in \{1 : p\}$ ,  $I_2 \in \{1 : p\}^2$ ,  $(j, k, m) = I_3 \in \{1 : p\}^3$ , where  $Y \sim p_0$ .

5. For  $\theta \in B$ , the corresponding information matrices  $V(\theta)$  and  $D(\theta)$  are positive definite with all elements being finite, where

$$V(\theta) = \int \nabla \ell(\theta, y) \nabla^T \ell(\theta, y) p_0(y) dy,$$

$$D(\theta) = \int \nabla^2 \ell(\theta, y) p_0(y) dy.$$

**Theorem 4 (Theorem 1 in [Lyddon et al., 2019])** *Let  $\tilde{\theta}_n$  be a loss-likelihood bootstrap sample of a parameter defined by (1.19) and (1.21) with loss function  $\ell$ , given  $n$  iid observations  $(y_1, \dots, y_n)$ , and let  $P_{LL}$  be its probability measure. Under the above assumptions, for any Borel set  $A \in \mathbb{R}^p$ , as  $n \rightarrow \infty$ ,*

$$P_{LL}(n^{1/2}(\tilde{\theta}_n - \hat{\theta}_n \in A)) \rightarrow P(Z \in A)$$

where  $Z \sim N_p(0, DV^{-1}D)$ ,  $\hat{\theta}_n = \arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n \ell(\theta, y_i)$  and

$$V = V(\theta^*), \quad D = D(\theta^*, y).$$

Therefore, the inference approach based on the loss-likelihood bootstrap is asymptotically efficient in the case the true distribution of the observations is unknown. Strictly speaking, this is not a Bernstein - von Mises theorem, since the convergence is not in the total variation distance, and hence it does not guarantee approximation of  $\Pi(\theta \in A \mid \mathbf{Y})$  by the corresponding Gaussian probabilities for all Borel sets  $A$ . Also, assumption of compactness of  $\Theta$  is not present in other results on posterior concentration, so it should be possible to relax this assumption.

Another interesting problem is how to modify this approach to take into the account a given a prior  $\pi$  that results in coherent and efficient inference about parameter  $\theta$ . [Newton et al., 2021] proposed such a solution, by replacing the loss function in the optimisation problem (1.21) by the loss function penalised by negative log prior, however for their choice of weights, the authors give a heuristic argument that their method approximates the target posterior with posterior covariance  $D^{-1}$  rather than with the sandwich covariance.

### 1.5.4.3 Other bootstrap-based approaches

Another approach is called bagged posterior or ‘‘BayesBag’’ which applies bagging proposed by[Breiman, 1996] to the Bayesian posterior [Waddell et al., 2002]. The idea is to select subsets of data as in bootstrap, compute posterior distribution for each of these subsets of data and average these posteriors. Formally, the bagged posterior is defined by

$$p_{\text{BayesBag}}(\theta \mid \mathbf{Y}) = \frac{1}{|I|} \sum_{\mathbf{Y}_{(i,N)} \in I} \pi(\theta \mid \mathbf{Y}_{(i,N)})$$

of the original data  $\mathbf{Y} = (Y_1, \dots, Y_n)$  and bootstrap data sets  $\mathbf{Y}_{(i,N)}$  of size  $N$  as the observed data. [Huggins and Miller, 2020] show that under a range of conditions, for iid true distribution of the data and iid model, bagged posterior distribution of  $\sqrt{n}(\theta - E(\theta \mid \mathbf{Y}))$  converges weakly to a Gaussian distribution centered at 0 with covariance matrix  $D^{-1}/c + D^{-1}VD^{-1}/c$  where  $c = \lim(N/n)$ . Hence, this approach does better than the usual posterior distribution with e.g.  $N = n - n_0$  for some small finite constant  $n_0$  leading to  $c = 1$ , however it is still not efficient.

### 1.5.5 Curvature adjustment

The generalised posterior approach uses a single parameter  $\eta$  to adjust for model misspecification. In general, it is possible to use this approach to obtain variance adjustment - and hence asymptotically optimal and valid posterior inference - only for one-dimensional parameter  $\theta$ . In the case of higher dimensions, [Ribatet et al., 2012] proposed to use curvature adjustment in the following way. For a possibly misspecified parametric family  $\{p(\cdot \mid \theta), \theta \in \Theta\}$  and prior  $\pi(\theta)$ , consider the following family

of posterior distributions:

$$\pi_A(\theta|y) = p(A\theta | y) \propto p(y | A\theta)\pi(A\theta).$$

Then, the idea is to find an estimator of the “oracle” matrix  $A$  in this class of admissible transforms determined by the condition that the posterior variance of  $A\theta$  is optimal, i.e. under the condition

$$A^T D^{-1} A = D^{-1} V D^{-1}$$

in the case the true parametric model is unknown and  $V = V(\theta^*)$  and  $V = V(\theta^*)$  defined by (1.2), and under condition  $A^T D^{-1} A = V_{true}^{-1}$  if the true parametric model is known with  $V_{true}$  being the Fisher information under the true parametric model. For the case of the unknown parametric model,  $D$  can be estimated by the posterior precision matrix of  $\theta$ , and estimation of  $V$  is usually more challenging.

[Stoehr and Friel, 2015] applied an affine version of the transform, i.e. they considered  $p(A\theta + b | y)$  misspecified models with known true parametric model, estimating  $b$  and  $A$  so that the posterior mean and the posterior variance of this distribution coincide with the posterior mean and the posterior variance under the posterior distribution with the true parametric model.

## 1.6 Discussion and open questions

The approach of [Panov and Spokoiny, 2015] and [Spokoiny and Panov, 2020] allows to address numerically the approximation properties of misspecified Bayesian inference and to verify whether it is close to being efficient, or whether a further adjustment is needed. While the authors have the assumption of a flat or Gaussian prior, for many model it is fairly straightforward to extend this to a larger class of continuous priors, in some cases with a continuous second derivative of the log likelihood.

There are many other interesting aspects of inference under model misspecification that are not considered here, such as optimality of predictive inference, model selection, etc.

It would be interesting to explore the connection between PAC-Bayesian inequalities and the conditional distribution  $p(\theta | y)$  defined as the solution of the optimisation problem (1.3).

Other interesting approaches include BvM for Variational Bayes under model misspecification [Wang and Blei, 2019], BvM for median and quantiles under classical and Gibbs posterior [Bhattacharya and Ghosal, 2019], [Bhattacharya and Martin, 2020]. Interestingly, [Zhang and Nalisnick, 2021] show that Bayesian neural networks show inconsistency similar to that discussed in [Grünwald and van Ommen, 2017], applying variational Bayes leads to BNN becoming consistent; it would be interesting to study whether it is possible to achieve asymptotic efficiency. Another version of robust Bayes-like estimation is proposed by [Baraud and Birgé, 2020] that does not

involve Kullback - Leibler distance but is based only on Hellinger distance between the true distribution and the parametric family.

Construction (asymptotically) efficient more general Bayesian inference under model misspecification (which is also computationally tractable) is a very active research area, with several promising solutions such as bootstrap posterior and curvature adjustment, however there is still no general unifying framework to encompass these approaches or to provide a coherent general framework. Fractional posterior allows a potentially simpler procedure for model correction which involves a single tuning parameter even if the parameter is multivariate which may result in conservative inference which can be sufficient for some problems but it is unlikely to be efficient in general. Linear curvature adjustment appear to work in practice and it is applicable to the models with no independence structure but there is no decision - theoretic justification for this is available yet; such justification is likely to involve geometry of the model space and its local linear adjustment. The open question in bootstrap-based posterior inference is the use of a given prior and its extension to data without independence structure which is likely to come from its Bayesian nonparametric interpretation. [Knoblauch et al., 2021] proposed an approach combining generalised variational inference, PAC-Bayes and other approaches into a single principled framework; they give conditions for consistency of their approach but not for efficiency. [Fong et al., 2021] propose a novel view to constructing a generalised posterior distribution, so it would be interesting to study its efficiency.

**Acknowledgements** This review was in part motivated by the discussion of the author with Peter Grünwald, Pierre Jacob and Jeffrey Miller during a Research in Groups meeting sponsored by the International Centre for Mathematical Sciences in Edinburgh, UK.

## References

- [Baraud and Birgé, 2020] Baraud, Y. and Birgé, L. (2020). Robust Bayes-like estimation: Rho-Bayes estimation. *Ann. Statist.*, 48(6):3699–3720.
- [Besag, 1986] Besag, J. (1986). On the statistical analysis of dirty pictures (with discussion). *J. Roy. Statist. Soc. B*, 48:259–302.
- [Bhattacharya et al., 2019] Bhattacharya, A., Pati, D., and Yang, Y. (2019). Bayesian fractional posteriors. *Ann. Statist.*, 47(1):39–66.
- [Bhattacharya and Ghosal, 2019] Bhattacharya, I. and Ghosal, S. (2019). Bayesian inference on multivariate medians and quantiles. *Statistica Sinica*.
- [Bhattacharya and Martin, 2020] Bhattacharya, I. and Martin, R. (2020). Gibbs posterior inference on multivariate quantiles. *arxiv: 2002.01052*.
- [Bissiri et al., 2016] Bissiri, P. G., Holmes, C., and Walker, S. G. (2016). A general framework for updating belief distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.
- [Bochkina and Green, 2014] Bochkina, N. A. and Green, P. J. (2014). The Bernstein-von Mises theorem and nonregular models. *Ann. Statist.*, 42(5):1850–1878.
- [Breiman, 1996] Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24:123 – 140.
- [Chernozhukov and Hong, 2004] Chernozhukov, V. and Hong, H. (2004). Likelihood estimation and inference in a class of nonregular econometric models. *Econometrica*, 72:1445–1480.

- [Chib et al., 2018] Chib, S., Shin, M., and Simoni, A. (2018). Bayesian estimation and comparison of moment condition models. *Journal of the American Statistical Association*, 113(524):1656–1668.
- [Dalalyan and Tsybakov, 2008] Dalalyan, A. and Tsybakov, A. B. (2008). Aggregation by exponential weighting, sharp PAC-Bayesian bounds and sparsity. *Machine Learning*, 72:39 – 61.
- [de Heide et al., 2020] de Heide, R., Kirichenko, A., Mehta, N., and Grünwald, P. (2020). Safe-Bayesian generalized linear regression. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2623–2633. PMLR.
- [Diong et al., 2017] Diong, M. L., Chaumette, E., and Vincent, F. (2017). On the efficiency of maximum-likelihood estimators of misspecified models. In *25th European Signal Processing Conference (EUSIPCO)*.
- [Fong et al., 2021] Fong, E., Holmes, C., and Walker, S. G. (2021). Martingale posterior distributions. *arXiv:2103.15671*.
- [Grünwald, 2012] Grünwald, P. (2012). The safe Bayesian. In *International Conference on Algorithmic Learning Theory*, pages 169–183. Springer.
- [Grünwald and van Ommen, 2017] Grünwald, P. and van Ommen, T. (2017). Inconsistency of Bayesian inference for misspecified linear models, and a proposal for repairing it. *Bayesian Anal.*, 12(4):1069–1103.
- [Grünwald and Mehta, 2020] Grünwald, P. D. and Mehta, N. A. (2020). Fast rates for general unbounded loss functions: From ERM to generalized Bayes. *Journal of Machine Learning Research*, 21(56):1–80.
- [Holmes and Walker, 2017] Holmes, C. and Walker, S. (2017). Assigning a value to a power likelihood in a general Bayesian model. *Biometrika*, 104(2):497–503.
- [Huggins and Miller, 2020] Huggins, J. and Miller, J. (2020). Robust and reproducible model selection using bagged posteriors. *arXiv:2007.14845*.
- [Ibragimov and Hasminskij, 1981] Ibragimov, I. and Hasminskij, R. (1981). *Statistical Estimation: Asymptotic Theory*. Springer.
- [Kleijn and van der Vaart, 2006] Kleijn, B. J. K. and van der Vaart, A. W. (2006). Misspecification in infinite-dimensional Bayesian statistics. *Ann. Statist.*, 34(2):837–877.
- [Kleijn and van der Vaart, 2012] Kleijn, B. J. K. and van der Vaart, A. W. (2012). The Bernstein-von-Mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354–381.
- [Knoblauch et al., 2021] Knoblauch, J., Jewson, J., and Damoulas, T. (2021). Generalized variational inference: Three arguments for deriving new posteriors. *arXiv:1904.02063*.
- [Lindsay, 1988] Lindsay, B. (1988). Composite likelihood methods. *Contemporary Mathematics*, 80:221 – 239.
- [Lyddon et al., 2019] Lyddon, S., Holmes, C., and Walker, S. (2019). General Bayesian updating and the loss-likelihood bootstrap. *Biometrika*, 106:465 – 478.
- [Miller and Dunson, 2019] Miller, J. W. and Dunson, D. B. (2019). Robust Bayesian inference via coarsening. *Journal of the American Statistical Association*, 114(527):1113–1125.
- [Müller, 2013] Müller, U. K. (2013). Risk of Bayesian inference in misspecified models, and the sandwich covariance matrix. *Econometrica*, 81(5):1805–1849.
- [Newton et al., 2021] Newton, M. A., Polson, N. G., and Xu, J. (2021). Weighted Bayesian bootstrap for scalable posterior distributions. *Canadian Journal of Statistics*.
- [Panov and Spokoiny, 2015] Panov, M. and Spokoiny, V. (2015). Finite sample Bernstein - von Mises theorem for semiparametric problems. *Bayesian Anal.*, 10(3):665–710.
- [Pauli et al., 2012] Pauli, F., Racugno, W., and Ventura, L. (2012). Bayesian composite marginal likelihoods. *Statistica Sinica*, 21:149–164.
- [Ribatet et al., 2012] Ribatet, M., Cooley, D., and Davison, A. C. (2012). Bayesian inference from composite likelihoods, with an application to spatial extremes. *Statistica Sinica*, 22:813–845.
- [Rubin, 1981] Rubin, D. B. (1981). The Bayesian bootstrap. *Ann. Statist.*, 9:130 – 134.
- [Schennach, 2005] Schennach, S. M. (2005). Bayesian exponentially tilted empirical likelihood. *Biometrika*, 92(1):31 – 46.
- [Spokoiny, 2012] Spokoiny, V. (2012). Parametric estimation. Finite sample theory. *Ann. Statist.*, 40(6):2877–2909.

- [Spokoiny, 2020] Spokoiny, V. (2020). Bayesian inference for nonlinear inverse problems. *arXiv:1912.12694*.
- [Spokoiny and Panov, 2020] Spokoiny, V. and Panov, M. (2020). Accuracy of Gaussian approximation in nonparametric Bernstein - von Mises. *arXiv:1910.06028*.
- [Stoehr and Friel, 2015] Stoehr, J. and Friel, N. (2015). Calibration of conditional composite likelihood for Bayesian inference on Gibbs random fields. In Lebanon, G. and Vishwanathan, S. V. N., editors, *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, volume 38 of *Proceedings of Machine Learning Research*, pages 921–929. PMLR.
- [Syring and Martin, 2020] Syring, N. and Martin, R. (2020). Gibbs posterior concentration rates under sub-exponential type losses. *arXiv:2012.04505*.
- [Van der Vaart, 2000] Van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.
- [Varin et al., 2011] Varin, C., Reid, N., and Firth, D. (2011). An overview of composite likelihood methods. *Statistica Sinica*, 21:5 – 42.
- [Waddell et al., 2002] Waddell, P. J., Kishino, H., and Ota, R. (2002). Very fast algorithms for evaluating the stability of ml and Bayesian phylogenetic trees from sequence data. In *Genome informatics*, volume 13, pages 82 – 92.
- [Wang and Blei, 2019] Wang, Y. and Blei, D. M. (2019). Variational Bayes under model misspecification. In *In Advances in Neural Information Processing Systems*.
- [Wang et al., 2017] Wang, Y., Kucukelbir, A., and Blei, D. M. (2017). Robust probabilistic modeling with Bayesian data reweighting. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3646 – 3655. PMLR.
- [White, 1982] White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica*, 50:1 – 25.
- [Wu and Martin, 2020] Wu, P.-S. and Martin, R. (2020). A comparison of learning rate selection methods in generalized Bayesian inference. *arxiv:2012.11349*.
- [Zhang and Nalisnick, 2021] Zhang, Y. and Nalisnick, E. (2021). On the inconsistency of Bayesian inference for misspecified neural networks. In *Third Symposium on Advances in Approximate Bayesian Inference*.