

Semiparametric posterior limits

B. J. K. Kleijn

Korteweg-de Vries Institute for Mathematics, University of Amsterdam

May 2013

Abstract

We review the Bayesian theory of semiparametric inference following Bickel and Kleijn (2012) [5] and Kleijn and Knapik (2013) [47]. After an overview of efficiency in parametric and semiparametric estimation problems, we consider the Bernstein-von Mises theorem (see, *e.g.*, Le Cam and Yang (1990) [57]) and generalize it to (LAN) regular and (LAE) irregular semiparametric estimation problems. We formulate a version of the semiparametric Bernstein-von Mises theorem that does not depend on least-favourable submodels, thus bypassing the most restrictive condition in the presentation of [5]. The results are applied to the (regular) estimation of the linear coefficient in partial linear regression (with a Gaussian nuisance prior) and of the kernel bandwidth in a model of normal location mixtures (with a Dirichlet nuisance prior), as well as the (irregular) estimation of the boundary of the support of a monotone family of densities (with a Gaussian nuisance prior).

Contents

1	Introduction	2
2	Efficiency	4
2.1	Efficiency in parametric models	5
2.2	Le Cam's Bernstein-von Mises theorem	9
3	Semiparametric efficiency	12
3.1	Several semiparametric estimation problems	12
3.2	Efficiency in semiparametric models	15
4	Posterior consistency under perturbation	17
4.1	Least-favourable submodels	17
4.2	Posterior concentration	20
4.3	Application in examples	22
5	Local expansions for integrated likelihoods	24
5.1	Approximately least-favourable submodels	25
5.2	Integrated local asymptotic normality	26

5.3	Application in examples	27
6	Posterior asymptotic normality and exponentiality	29
6.1	Local limit shapes of marginal posteriors	29
6.2	Marginal posterior consistency and semiparametric bias	31
6.3	Application in examples	34
7	Main results	36
7.1	Main theorems	36
7.2	Partial linear regression	38
7.3	Normal location mixtures	40
7.4	Support boundary estimation	40
8	Proofs	40

1 Introduction

Consider estimation of a functional $\theta : \mathcal{P} \rightarrow \mathbb{R}^k$ on a dominated nonparametric model \mathcal{P} with metric g , based on a sample X_1, X_2, \dots , distributed *i.i.d.* according to $P_0 \in \mathcal{P}$. We introduce a prior Π on \mathcal{P} and consider the subsequent sequence of posteriors,

$$\Pi(A \mid X_1, \dots, X_n) = \int_A \prod_{i=1}^n p(X_i) d\Pi(P) \Big/ \int_{\mathcal{P}} \prod_{i=1}^n p(X_i) d\Pi(P), \quad (1)$$

where A is any measurable model subset. Typically, optimal (*e.g.* minimax) nonparametric posterior rates of convergence [31] are powers of n (possibly modified by a slowly varying function) that converge to zero more slowly than the parametric $n^{-1/2}$ -rate. Instances of inconsistency in nonparametric Bayesian statistics are numerous [29, 21, 16, 22, 30] but practical sufficient conditions for posterior consistency (Schwartz (1965) [66]) and rates of convergence (Ghosal, Ghosh and van der Vaart (2000) [31], Shen and Wasserman (2001) [68]) are well-known. Together, negative and positive results demonstrate that the choice of a nonparametric prior is a sensitive one that leaves room for unintended consequences unless due care is taken.

This lesson must also be taken seriously when one asks the question whether the marginal posterior for the parameter of interest in a semiparametric estimation problem displays Bernstein-Von Mises-type limiting behaviour. In this paper, our primary goal is efficient estimation of smooth, real-valued aspects of P_0 : parametrize the model in terms of a finite-dimensional *parameter of interest* $\theta \in \Theta$ and an infinite-dimensional *nuisance parameter* $\eta \in H$: $\mathcal{P} = \{P_{\theta, \eta} : \theta \in \Theta, \eta \in H\}$. We look for general sufficient conditions on model and prior such that the *marginal posterior for the parameter of interest* satisfies,

$$\sup_B \left| \Pi(\sqrt{n}(\theta - \theta_0) \in B \mid X_1, \dots, X_n) - N_{\tilde{\Delta}_n, \tilde{I}_{\theta_0, \eta_0}^{-1}}(B) \right| \rightarrow 0, \quad (2)$$

in P_{θ_0} -probability, where,

$$\tilde{\Delta}_n = \frac{1}{\sqrt{n}} \sum_{i=1}^n \tilde{I}_{\theta_0, \eta_0}^{-1} \tilde{\ell}_{\theta_0, \eta_0}(X_i), \quad (3)$$

$\tilde{\ell}_{\theta,\eta}$ is the efficient score function and $\tilde{I}_{\theta,\eta}$ the (non-singular) efficient Fisher information (for definitions, see subsection 3.2 below). Assertion (2) (roughly) implies efficiency of point-estimators like the posterior median, mode or mean and justifies asymptotic identification of credible regions with efficient confidence regions (see Section 2). From a practical point of view, the latter conclusion has an important implication: whereas, in many semiparametric estimation problems, it is hard to calculate optimal semiparametric confidence regions directly, simulation of a large sample from the marginal posterior (*e.g.* by MCMC techniques, see Robert (2001) [65] and many others) is sometimes comparatively straightforward.

Instances of the Bernstein-Von Mises limit have been studied in various semiparametric models; we mention references of a general nature and several model-specific discussions. The first general reference in this area is Shen (2002) [69] with application to partial linear regression, but his conditions appear hard to verify in other examples. Castillo (2012) [10] is inspired by and related to [69] and provides general conditions with two applications. Cheng and Kosorok (2008) [13] give a general perspective too, proving weak convergence of the posterior under sufficient conditions. Rivoirard and Rousseau (2009) [64] prove a version for linear functionals over the model, using a class of nonparametric priors based on infinite-dimensional exponential families on Sobolev and Besov spaces. Boucheron and Gassiat (2009) [8] consider the Bernstein-Von Mises theorem for families of discrete distributions with Dirichlet priors, motivated by information-theoretic questions. Johnstone (2010) [38] studies various marginal posteriors in the Gaussian sequence model, taking sieve-like limits of finite-dimensional posteriors. De Blasi and Hjort (2007,2009) [18, 19] analyse partial likelihood and Bayesian methods in Cox' proportional hazards model with a Beta process prior for the cumulative baseline hazard. In Kruijer and Rousseau (2013) [50], Gaussian time series with long-memory behaviour are analysed with an infinite-dimensional version of the FEXP model, using families of priors defined on approximating sieves. De Jonge and van Zanten (2013) [20] consider Gaussian regression problems with Gaussian priors to estimate the variance of the error and Knapik, van de Vaart and van Zanten (2011) [48] consider finite-dimensional marginals in Gaussian inverse problems with Gaussian priors.

The field of semiparametric Bayesian statistics is relatively new and the papers mentioned above explore a great variety of different methods to arrive at the Bernstein-von Mises limit. Many of those methods are model-specific and do not lend themselves to generalization (especially the Gaussian sequence model with a Gaussian prior has received a very large amount of attention). Questions remain and a coherent, unified point of view has not been established. For that reason the perspective of paper does not provide a comprehensive account of possible approaches to the Bayesian semiparametric problem; instead it is based primarily on the perspective of [5, 47]. We review the theory of efficient estimation in smooth parametric and semiparametric models and discuss the derivation of the semiparametric Bernstein-von Mises theorem in locally asymptotically normal [5] and locally asymptotically exponential [47] problems. To enhance applicability, a new version of the regular semiparametric Bernstein-von Mises theorem is formulated: where, previously, the construction depended on the existence of a smooth least-favourable submodel, the new version only requires that

the model permits a sequence of submodels that approximate least-favourable directions in a suitable way (see subsection 5.1). Where proofs change, full details are provided (see subsection 5.2 and section 8). Throughout, developments are related to the locally asymptotically exponential case (for which proofs run largely analogously).

Every major step in the development is illustrated with three running semiparametric examples: the first two, (regular) estimation of the linear coefficient in the partial linear regression model [5] and (irregular) estimation of support boundary points for a family of monotone densities [47], are analysed in full detail. The third is new and concerns (regular) estimation of kernel variance in a normal location mixture model, but the discussion is not as detailed and rigorous as that of the other two examples. Results are summarized in a general theorem and corollary (see subsection 7.1) and two model-specific Bernstein-von Mises theorems, for partial linear regression (see subsection 7.2), and support boundary estimation (see subsection 7.4). For lack of rigorous (aspects of) proofs, estimation of kernel variance in normal location mixtures is commented on in the form of a conjecture (see subsection 7.3).

Notation and conventions

The (frequentist) true distribution of the data is denoted P_0 and assumed to lie in \mathcal{P} , so that there exist $\theta_0 \in \Theta$, $\eta_0 \in H$ such that $P_0 = P_{\theta_0, \eta_0}$. In regular problems, θ is localized by introduction of $h = \sqrt{n}(\theta - \theta_0)$ with inverse $\theta_n(h) = \theta_0 + n^{-1/2}h$; in irregular problems we follow analogous definitions with rate n^{-1} . The (multivariate) normal distribution with mean μ and covariance Σ is denoted $N_{\mu, \Sigma}$. The location-scale family associated with the exponential distribution is denoted by $\text{Exp}_{\Delta, \lambda}^+$ and its negative version (supported on a half-line extending to $-\infty$) by $\text{Exp}_{\Delta, \lambda}^-$. The expectation of a random variable f with respect to a probability measure P is denoted Pf ; the sample average of $g(X)$ is denoted $\mathbb{P}_n g(X) = (1/n) \sum_{i=1}^n g(X_i)$ and $\mathbb{G}_n g(X) = n^{1/2}(\mathbb{P}_n g(X) - P g(X))$ (for other conventions and nomenclature customary in empirical process theory, see [76]). If f is an integrable random variable and h_n is stochastic, $P_{\theta_n(h_n), \eta}^n f$ denotes the integral $\int f(\omega) (dP_{\theta_n(h_n(\omega)), \eta}^n / dP_0^n)(\omega) dP_0^n(\omega)$. The Hellinger distance between P and P' is denoted $H(P, P')$ and induces a metric d_H on the space of nuisance parameters H by $d_H(\eta, \eta') = H(P_{\theta_0, \eta}, P_{\theta_0, \eta'})$, for all $\eta, \eta' \in H$. We endow the model with the Borel σ -algebra generated by the Hellinger topology and refer to [31] regarding issues of measurability.

2 Efficiency

Perhaps the most intuitive way to express statistical inference is formulation in terms of (frequentist) *confidence sets* or (Bayesian) *credible sets*. Typically, confidence sets are defined as neighbourhoods of an estimator with a certain coverage probability, based on the quantiles of its *sampling distribution*. Credible sets represent the same concept in Bayesian statistics and are defined with the posterior in the role of the sampling distribution. In what follows we shall not be too strict in Bayesian, subjectivist orthodoxy and interpret the posterior as a frequentist device, asking the natural question how its credible sets compare to confidence sets.

Since confidence sets and credible sets are conceptually so close, could it be that they are close also mathematically? To answer this question, we briefly review the modern theory of (point-)estimation of smooth parameters to arrive at a notion of asymptotic inferential optimality and we discuss the Bernstein-von Mises theorem (theorem 2.10 below) which demonstrates asymptotic equivalence of credible sets and optimal (or *efficient*) confidence sets.

2.1 Efficiency in parametric models

The concept of efficiency has its origin in Fisher's 1920's claim of asymptotic optimality of the maximum-likelihood estimator in differentiable parametric models (Fisher (1959) [28] and Cramèr (1946) [17]). Here, optimality of ML estimates means that they are consistent, achieve $n^{-1/2}$ rate of convergence and possess an asymptotic sampling distribution of minimal variance. To illustrate, consider the following classical result from M -estimation.

Theorem 2.1. *Let Θ be open in \mathbb{R}^k and assume that \mathcal{P} is characterized by densities $p_\theta : \mathcal{X} \rightarrow \mathbb{R}$ such that $\theta \mapsto \log p_\theta(x)$ is differentiable at θ_0 for all $x \in \mathcal{X}$, with derivative $\dot{\ell}_\theta(x)$. Assume that there exists a function $\dot{\ell} : \mathcal{X} \rightarrow \mathbb{R}$ such that $P_0 \dot{\ell}^2 < \infty$ and*

$$|\log p_{\theta_1}(x) - \log p_{\theta_2}(x)| \leq \dot{\ell}(x) \|\theta_1 - \theta_2\|,$$

for all θ_1, θ_2 in an open neighbourhood of θ_0 . Furthermore, assume that $\theta \mapsto P_{\theta_0} \log p_\theta$ has a second-order Taylor expansion around θ_0 of the form,

$$P_{\theta_0} \log p_\theta = P_{\theta_0} \log p_{\theta_0} + \frac{1}{2}(\theta - \theta_0)^T I_{\theta_0} (\theta - \theta_0) + o(\|\theta - \theta_0\|^2),$$

with non-singular I_{θ_0} . If $(\hat{\theta}_n)$ are (near-)maximizers of the likelihood such that $\hat{\theta}_n \xrightarrow{\theta_0} \theta_0$, then the estimator sequence is asymptotically linear,

$$n^{1/2}(\hat{\theta}_n - \theta_0) = n^{-1/2} \sum_{i=1}^n I_{\theta_0}^{-1} \dot{\ell}_{\theta_0}(X_i) + o_{P_{\theta_0}}(1),$$

in particular, $n^{1/2}(\hat{\theta}_n - \theta_0) \overset{\theta_0}{\rightsquigarrow} N(0, I_{\theta_0}^{-1})$.

For a proof, see theorem 5.23 in van der Vaart (1998) [77]. Associated *asymptotic* confidence sets are the approximate confidence sets one obtains upon approximation of sampling distributions by the limit distribution. Denoting quantiles of the χ^2 -distribution with k degrees of freedom by $\chi_{k,\alpha}^2$, we find that ellipsoids of the form,

$$C_\alpha(X_1, \dots, X_n) = \{\theta \in \Theta : n(\theta - \hat{\theta}_n)^T I_{\hat{\theta}_n} (\theta - \hat{\theta}_n) \leq \chi_{k,\alpha}^2\}, \quad (4)$$

have coverage probabilities converging to $1 - \alpha$ and are therefore asymptotic confidence sets.

Theorem 2.1 requires a rather large number of smoothness properties of the model which are there to guarantee that the ML estimator displays *regularity*. The prominence of regularity in the context of optimality questions was not fully appreciated until 1951, when Hodges revealed a phenomenon now known as *superefficiency* through formulation of shrinkage: the behaviour of estimators like $\hat{\theta}_n$ above can be adapted around certain points in the parameter

space to outperform the MLE and other estimators like it asymptotically, while doing equally well for all other points. Superefficiency indicated that Fisher's 1920's claim was false without further refinement and that a comprehensive understanding of optimality in differentiable estimation problems remained elusive.

To resolve the issue and arrive at a sound theory of asymptotic optimality of estimation in differentiable models, two concepts were introduced, the first being a concise notion of smoothness. (In the following we assume that the sample is *i.i.d.*, although usually the definition is extended to more general forms of data.)

Definition 2.2. (*Local asymptotic normality (LAN), Le Cam (1960) [53]*)

Let $\Theta \subset \mathbb{R}^k$ be open, parametrizing a model $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$ that is dominated by a σ -finite measure with densities p_θ . The model is said to be *locally asymptotically normal (LAN)* at θ_0 if, for any converging sequence $h_n \rightarrow h$:

$$\log \prod_{i=1}^n \frac{p_{\theta_0+n^{-1/2}h_n}(X_i)}{p_{\theta_0}} = h^T \Gamma_{n,\theta_0} - \frac{1}{2} h^T I_{\theta_0} h + o_{P_{\theta_0}}(1), \quad (5)$$

for random vectors Γ_{n,θ_0} such that $\Gamma_{n,\theta_0} \overset{\theta_0}{\rightsquigarrow} N_k(0, I_{\theta_0})$. \square

Differentiability of the log-density $\theta \mapsto \log p_\theta(x)$ at θ_0 for every x and continuity of the associated Fisher information (see, for instance, lemma 7.6 in [77]) imply that the model is LAN at θ_0 with $\Gamma_{n,\theta_0} = n^{-1/2} \sum_{i=1}^n \dot{\ell}_{\theta_0}(X_i)$. But local asymptotic normality can be achieved under a weaker condition.

Definition 2.3. (*Differentiability in quadratic mean (DQM)*)

Let Θ be an open subset of \mathbb{R}^k . A model $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$ that is dominated by a σ -finite measure μ with densities p_θ is said to be *differentiable in quadratic mean (DQM)* at $\theta_0 \in \Theta$, if there exists a score function $\dot{\ell}_{\theta_0} \in L_2(P_{\theta_0})$ such that:

$$\int \left(p_{\theta_0+h}^{1/2} - p_{\theta_0}^{1/2} - \frac{1}{2} h^T \dot{\ell}_{\theta_0} p_{\theta_0}^{1/2} \right)^2 d\mu = o(\|h\|^2),$$

as $h \rightarrow 0$. \square

Theorem 75.9 in Strasser (1985) [72] demonstrates equivalence of the DQM and LAN properties. In the proof of the semiparametric Bernstein-von Mises theorem below, we use a smoothness property that is slightly stronger.

Definition 2.4. (*Stochastic LAN (sLAN)*)

We say that a parametric model \mathcal{P} is *stochastically LAN* at θ_0 , if the LAN property of definition 2.1 is satisfied for every random sequence (h_n) that is bounded in probability, *i.e.* for all $h_n = O_{P_0}(1)$:

$$\log \prod_{i=1}^n \frac{p_{\theta_0+n^{1/2}h_n}(X_i)}{p_{\theta_0}} - h_n^T \Gamma_{n,\theta_0} - \frac{1}{2} h_n^T I_{\theta_0} h_n = o_{P_{\theta_0}}(1), \quad (6)$$

for random vectors Γ_{n,θ_0} such that $\Gamma_{n,\theta_0} \overset{\theta_0}{\rightsquigarrow} N_k(0, I_{\theta_0})$. \square

The second concept is a property that characterizes the class of estimators over which optimality is achieved (in particular excluding Hodges' shrinkage estimators and other examples of superefficiency, as becomes clear below). To prepare the definition heuristically, note that, given Hodges' counterexample, it is not enough to have estimators with pointwise convergence to limit laws; we must restrict the behaviour of estimators over $(n^{-1/2})$ -neighbourhoods rather than allow the type of wild variations that make superefficiency possible.

Definition 2.5. (*Regularity of estimation*)

Let $\Theta \subset \mathbb{R}^k$ be open. An estimator sequence (T_n) for the parameter θ is said to be *regular* at θ if there exists a L_θ such that for all $h \in \mathbb{R}^k$,

$$n^{1/2} \left(T_n - (\theta + n^{-1/2}h) \right) \rightsquigarrow L_\theta, \quad (\text{under } P_{\theta+n^{-1/2}h}), \quad (7)$$

i.e. with a limit law independent of h . □

So regularity describes the property that convergence of the estimator to a limit law is insensitive to *perturbation* of the parameter of size $n^{-1/2}h$. The LAN and regularity properties come together in the following theorem which forms the foundation for the convolution theorem that follows (see theorems 7.10, 8.3, 8.4 in van der Vaart (1998) [77]).

Theorem 2.6. (Gaussian limit experiment [54])

With $\Theta \subset \mathbb{R}^k$ open, let $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$ be LAN at θ_0 with non-singular Fisher information I_{θ_0} . Let (T_n) be regular estimators in the models $\{P_{\theta_0+n^{-1/2}h} : h \in \mathbb{R}^k\}$. Then there exists a (randomized) statistic T in the normal location model $\{N_k(h, I_{\theta_0}^{-1}) : h \in \mathbb{R}^k\}$ such that $T - h \sim L_{\theta_0}$ for all $h \in \mathbb{R}^k$.

Theorem 2.6 provides every regular estimator sequence with a limit in the form of a statistic in a very simple model in which the only parameter is the location of a normal distribution: the (weak) limit distribution that describes the local asymptotics of the sequence (T_n) under $P_{\theta_0+n^{-1/2}h}$ equals the distribution of T under h , for all $h \in \mathbb{R}^k$. Moreover regularity of the sequence (T_n) implies that under $N_k(h, I_{\theta_0}^{-1})$, the distribution of T relative to h is independent of h , an property known as *equivariance-in-law*. The class of equivariant-in-law estimators for location in the model $\{N_k(h, I_{\theta_0}^{-1}) : h \in \mathbb{R}^k\}$ is fully known: for any equivariant-in-law estimator T for h , there exists a distribution M such that T is distributed according to the convolution $N_k(h, I_{\theta_0}^{-1}) * M$. (The most straightforward example is $T = X$, for which $M = \delta_0$.) This argument gives rise to the following central result in the theory of efficiency.

Theorem 2.7. (Convolution theorem (Hájek (1970) [34]))

Let $\Theta \subset \mathbb{R}^k$ be open and let $\{P_\theta : \theta \in \Theta\}$ be LAN at θ_0 with non-singular Fisher information I_{θ_0} . Let (T_n) be a regular estimator sequence with limit distribution L_{θ_0} . Then there exists a probability distribution M_{θ_0} such that,

$$L_{\theta_0} = N_k(0, I_{\theta_0}^{-1}) * M_{\theta_0},$$

in particular, if L_{θ_0} has a covariance matrix Σ_{θ_0} , then $\Sigma_{\theta_0} \geq I_{\theta_0}^{-1}$.

The occurrence of the inverse Fisher information is no coincidence: the estimator T is unbiased and satisfies the Cramér-Rao bound in the limiting model $\{N_k(h, I_{\theta_0}^{-1}) : h \in \mathbb{R}^k\}$. Hence, the last assertion of the convolution theorem says that, within the class of regular estimates, asymptotic variance is lower-bounded by the inverse Fisher information. A regular estimator that is optimal in this sense, is called *best-regular*. Anderson's lemma broadens this notion of optimality, in the sense that best-regular estimators outperform other regular estimators with respect to a large family of loss functions. Conversely, the asymptotic minimax theorem shows that best-regularity is *necessary* for optimality with respect to any such loss-function (Hájek (1972) [35]). Finally, we mention the following equivalence which characterizes efficiency concisely in terms of a weakly converging sequence.

Lemma 2.8. *In a LAN model, estimators (T_n) for θ are best-regular iff the (T_n) are asymptotically linear, i.e. for all θ in the model,*

$$n^{1/2}(T_n - \theta) = \frac{1}{\sqrt{n}} \sum_{i=1}^n I_{\theta}^{-1} \dot{\ell}_{\theta}(X_i) + o_{P_{\theta}}(1). \quad (8)$$

The random sequence of differences on the r.h.s. of (8) is denoted by Δ_{n,θ_0} below.

Coming back to theorem 2.1, we see that under stated conditions, a consistent sequence of MLE's ($\hat{\theta}_n$) is best-regular, finally giving substance to Fisher's 1920's claim. We now know that in a LAN model, confidence sets of the form (4) based on best-regular estimators ($\hat{\theta}_n$) share their optimality.

However, not all estimators are regular and not all model parameters are smooth. In the literature, situations in which regularity does not apply are collectively known as *irregular*. A prototypical irregular problem concerns the estimation a support boundary point for a density supported on a half-line. As a frequentist problem, it is well-understood (Ibragimov and Has'minskii (1981) [37]): assuming that the distribution P_{θ} of X is supported on the half-line $[\theta, \infty)$ and an *i.i.d.* sample X_1, X_2, \dots, X_n is given, we follow [37] and estimate θ with the ML estimator, the first order statistic $X_{(1)} = \min_i\{X_i\}$. If P_{θ} has a continuous Lebesgue density of the form $p_{\theta}(x) = \eta(x - \theta) 1\{x \geq \theta\}$, its rate of convergence is determined by the behaviour of the quantity $m(\epsilon) \mapsto \int_0^{\epsilon} \eta(x) dx$ for small values of ϵ . If $m(\epsilon) = \epsilon^{\alpha+1}(1 + o(1))$ as $\epsilon \downarrow 0$, for some $\alpha \in (-1, 1)$, then,

$$n^{1/(1+\alpha)}(X_{(1)} - \theta) = O_{P_{\theta}}(1). \quad (9)$$

For densities of this form, for any sequence θ_n that converges to θ at rate $n^{-1/(1+\alpha)}$, Hellinger distances obey (see Theorem VI.1.1 in [37]):

$$n^{1/2} H(P_{\theta_n}, P_{\theta}) = O(1). \quad (10)$$

If we substitute the estimators $\theta_n = \hat{\theta}_n(X_1, \dots, X_n) = X_{(1)}$, uniform tightness of the sequence in the above display signifies rate optimality of the estimator (*c.f.* Le Cam (1973, 1986) [55, 56]). Regarding asymptotic efficiency beyond rate-optimality, *e.g.* in the sense of minimal asymptotic variance (or other measures of dispersion of the limit distribution), one notices

that the (one-sided) limit distributions one obtains for the MLE $X_{(1)}$ can always be improved upon by de-biasing (see Section VI.6, examples 1–3 in [37] and Le Cam (1990) [58]).

In much of what follows we concentrate on the support boundary problem for a discontinuity ($\alpha = 0$) because in those cases the likelihood permits an expansion reminiscent of LAN: θ is represented in localised form, by centering on θ_0 and rescaling: $h = n(\theta - \theta_0) \in \mathbb{R}$. The following (irregular) local expansion of the likelihood is due to Ibragimov and Has'minskii (1981) [37].

Definition 2.9. (*Local asymptotic exponentiality (LAE)*)

Let $\Theta \subset \mathbb{R}$ be open; a model $\theta \mapsto P_\theta$ is said to be *locally asymptotically exponential (LAE)* at $\theta_0 \in \Theta$ if there exists a sequence of random variables (Δ_n) and a positive constant γ_{θ_0} such that for all (h_n) , $h_n \rightarrow h$,

$$\prod_{i=1}^n \frac{p_{\theta_0+n^{-1}h_n}(X_i)}{p_{\theta_0}} = \exp(h\gamma_{\theta_0} + o_{P_{\theta_0}}(1)) 1_{\{h \leq \Delta_n\}},$$

with Δ_n converging weakly to $\text{Exp}_{0,\gamma_{\theta_0}}^+$. □

2.2 Le Cam's Bernstein-von Mises theorem

To address the question of efficiency in smooth parametric models from a Bayesian perspective, we turn to the Bernstein-Von Mises theorem (Le Cam (1953) [52]). In the literature many different versions of the theorem exist, varying both in (stringency of) conditions and (strength or) form of the assertion. Following Le Cam and Yang (1990) [57] we state the theorem as follows. (For later reference define a parametric prior to be *thick* at θ_0 , if it has a Lebesgue density that is continuous and strictly positive at θ_0 .)

Theorem 2.10. (Bernstein-Von Mises theorem, Le Cam and Yang (1990) [57])

Assume that $\Theta \subset \mathbb{R}^k$ is open and that the model $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$ is identifiable and dominated. Suppose X_1, X_2, \dots forms an i.i.d. sample from P_{θ_0} for some $\theta_0 \in \Theta$. Assume that the model is LAN at θ_0 with non-singular Fisher information I_{θ_0} . Furthermore, suppose that, the prior Π_Θ is thick at θ_0 and that for every $\epsilon > 0$, there exists a test sequence (ϕ_n) such that,

$$P_{\theta_0}^n \phi_n \rightarrow 0, \quad \sup_{\|\theta - \theta_0\| > \epsilon} P_\theta^n (1 - \phi_n) \rightarrow 0.$$

Then the posterior distributions converge in total variation,

$$\sup_B \left| \Pi(\theta \in B \mid X_1, \dots, X_n) - N_{\hat{\theta}_n, (nI_{\theta_0})^{-1}}(B) \right| \rightarrow 0,$$

in P_{θ_0} -probability, where $(\hat{\theta}_n)$ denotes any best-regular estimator sequence.

For a proof, the reader is referred to [57, 77] (or to [46] for a proof under model misspecification that has a lot in common with the proof of theorem 6.1 below); see also Bickel and Yahav (1969) [2]. In figure 2.2, Bernstein-von Mises-type convergence of posterior densities is demonstrated through numerical simulation. Also displayed in figure 2.2 are the so-called

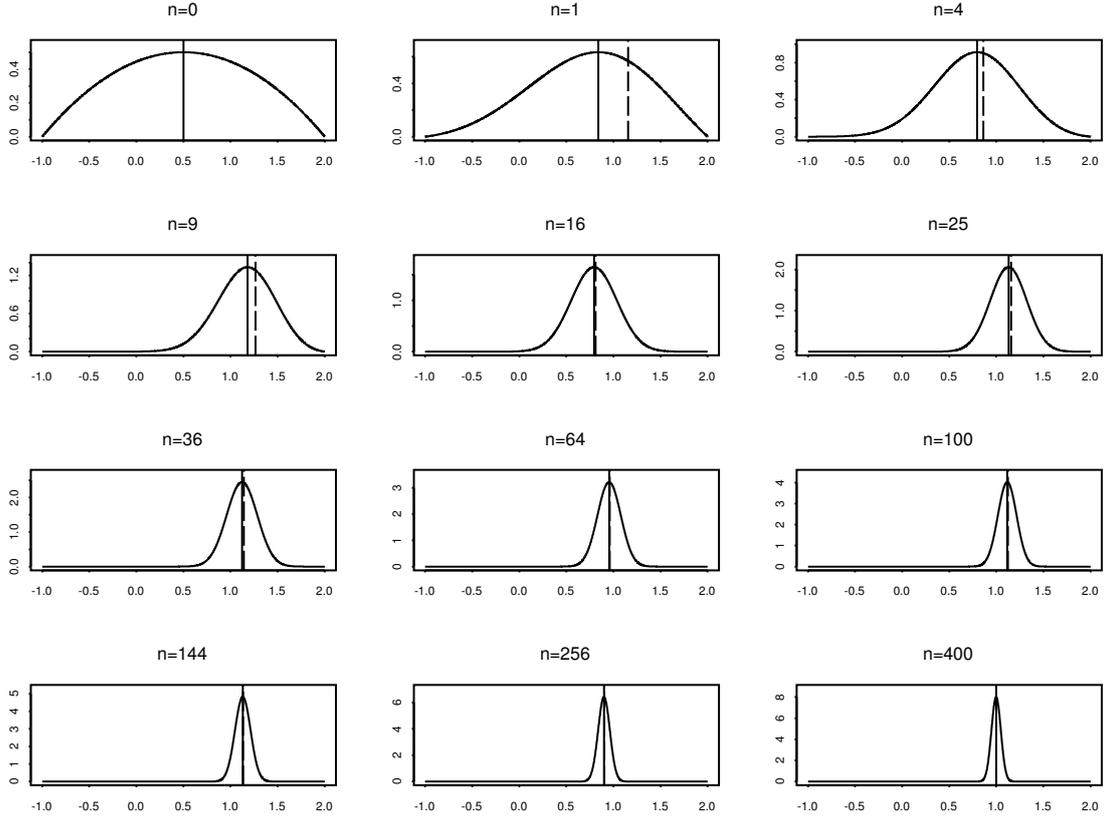


Figure 1: Convergence of the posterior density. The samples used for calculation of the posterior distributions consist of n observations; the model consists of all normal distributions with mean between -1 and 2 and variance 1 and has a polynomial prior, shown in the first ($n = 0$) graph. For all sample sizes, the *maximum a posteriori* and maximum likelihood estimators are indicated by a vertical line and a dashed vertical line respectively. (From Kleijn (2003))

MAP estimator and the ML estimator. It is noted that, here, the MLE is efficient so it forms a possible centring sequence for the limiting sequence of normal distributions in the assertion of the Bernstein-Von Mises theorem. Furthermore, it is noted that the posterior concentrates more and more sharply, reflecting the n^{-1} -proportionality of the variance of its limiting sequence of normals. It is perhaps a bit surprising in figure 2.2 to see limiting normality obtain already at such relatively low values of the sample size n . It cannot be excluded that, in this case, that is a manifestation the normality of the underlying model, but onset of normality of the posterior appears to happen at low values of n also in other smooth, parametric models. It suggests that asymptotic conclusions based on the Bernstein-Von Mises limit accrue validity fairly rapidly, for n in the order of several hundred to several thousand *i.i.d.* replications of the observation.

The uniformity in the assertion of the Bernstein-Von Mises theorem over model subsets

B implies that it holds also for model subsets that are random. (And because some authors content themselves with weaker statements they call “Bernstein-von Mises” assertions, it is noted that, crucially, pointwise convergence of the posterior distribution function does *not* constitute a sufficient condition.) In particular, given some $0 < \alpha < 1$, it is noted that the smallest sets $C_\alpha(X_1, \dots, X_n)$ such that,

$$N_{\hat{\theta}_n, (nI_{\theta_0})^{-1}}(C_\alpha(X_1, \dots, X_n)) \geq 1 - \alpha,$$

are ellipsoids of the form (4). According to the Bernstein-Von Mises limit, posterior coverage of C_α converges to the *l.h.s.* in the above display, so the C_α are asymptotic credible sets of posterior coverage $1 - \alpha$. Conversely, any sequence (C_n) of (data-dependent) credible sets of coverage $1 - \alpha$, is also a sequence of sets that have asymptotic confidence level $1 - \alpha$ (using the best-regularity of $\hat{\theta}_n$). So the Bernstein-von Mises theorem identifies inference based on frequentist best-regular point-estimators with inference based on Bayesian posteriors (in smooth, parametric models). From a practical perspective the Bernstein-Von Mises theorem offers an alternative way to arrive at asymptotic confidence sets, if we have an approximation of the posterior distribution of high enough quality (*e.g.* from MCMC simulation). In high dimensional parametric models, maximization of the likelihood may be much more costly computationally than generation of a sample from the posterior. As a consequence, the Bernstein-Von Mises theorem has an immediate practical implication of some significance. This practical point will continue to hold in semiparametric context where the comparative advantage is much greater.

The irregular example calls for estimation of a support boundary point of a density: consider an almost-everywhere differentiable Lebesgue density on \mathbb{R} that displays a jump at some point $\theta \in \mathbb{R}$; estimators for θ exist that converge at rate n^{-1} with exponential limit distributions [37]. To illustrate the form that this conclusion takes in Bayesian context, consider the following straightforward theorem with exponential densities.

Theorem 2.11. (Irregular posterior convergence)

For $\theta \in \mathbb{R}$, let $F_\theta(x) = (1 - e^{-(x-\theta)}) \vee 0$. Assume that X_1, X_2, \dots form an *i.i.d.* sample from F_{θ_0} , for some θ_0 . Let $\pi : \mathbb{R} \rightarrow (0, \infty)$ be a continuous Lebesgue probability density. Then the associated posterior distribution satisfies,

$$\sup_A \left| \Pi_n(\theta \in A \mid X_1, \dots, X_n) - \text{Exp}_{\hat{\theta}_n, n}^-(A) \right| \xrightarrow{\theta_0} 0,$$

where $\hat{\theta}_n = X_{(1)}$ is the maximum likelihood estimate for θ_0 .

Note that in this case the limiting posterior is a (negative) exponential distribution that can be identified as the distribution for which level sets define ML-based confidence sets. So here the asymptotic identification of credible sets and confidence intervals holds as well. What is missing in this case is the guarantee of optimality (for lack of an irregular analog of the convolution theorem). Indeed, the posterior follows the ML estimate and mean-square errors can be improved upon by simple de-biasing [37, 57] as a consequence.

3 Semiparametric efficiency

Semiparametric statistics asks parametric questions in nonparametric models. As such it combines the best of two worlds, diminishing the risk of misspecification by use of nonparametric models while maintaining much of the benefits of parametric inference, including the optimality theory for regular estimators in smooth models. Although the more general formulation calls for a nonparametric model \mathcal{P} with a finite-dimensional functional $\theta : \mathcal{P} \rightarrow \mathbb{R}^k$ of interest, we choose to parametrize model distributions in terms of a finite-dimensional *parameter of interest* $\theta \in \Theta$, for an open $\Theta \subset \mathbb{R}^k$, and an infinite-dimensional *nuisance parameter* $\eta \in H$; the nonparametric model is then represented as $\mathcal{P} = \{P_{\theta,\eta} : \theta \in \Theta, \eta \in H\}$. It is assumed that the model \mathcal{P} is identifiable and that the true distribution of the data P_0 is contained in the model, implying that there exist unique $\theta_0 \in \Theta$, $\eta_0 \in H$ such that $P_0 = P_{\theta_0,\eta_0}$. Furthermore it is assumed that the model is dominated by a σ -finite measure with densities $p_{\theta,\eta}$. Of course, we impose smoothness on the model in a suitable way and we intend to estimate θ with semiparametric efficiency.

3.1 Several semiparametric estimation problems

Before we discuss these matters in more detail, we mention several well-known semiparametric estimation problems (with references to Bayesian analyses in the literature).

Example (*Symmetric location problem, Stein (1956) [71]*)

Consider a distribution with Lebesgue density $\eta : \mathbb{R} \rightarrow [0, \infty)$ that is symmetric around 0, *i.e.* $\eta(x) = \eta(-x)$ for all $x \in \mathbb{R}$ with finite Fisher information for location $\int (\eta'/\eta)^2(x) \eta(x) dx < \infty$. Assume that X_1, X_2, \dots is an *i.i.d.* sample from a distribution P_0 with density $p_{\theta_0}(x) = \eta(x - \theta_0)$. We are interested in estimation of θ_0 without knowledge of the nuisance η . See Bickel (1982) [3]; for a Bayesian analysis with Bernstein-von Mises limits, see Shen (2002) [69] and Castillo (2012) [10]. \square

Example (*Semiparametric mixture models [63, 75, 4, 77]*)

Mixtures arise whenever a modelled random variable remains unobserved. In semiparametric mixture models we have a pair (Y, Z) of which only Y is observed and we consider the conditional distribution of Y given $Z = z$, assumed to be from a parametric family $\{\Psi_\theta(\cdot|z) : \theta \in \Theta\}$. We aim to estimate θ in the presence of the nuisance F , the unknown distribution of Z . As a simple example, consider the *normal location model*: a random variable X arises as $X = Z + e$, where the unobserved $Z \in [0, 1]$ has distribution $F \in \mathcal{D}[0, 1]$ and is independent of a normally distributed error $e \sim N(0, \sigma^2)$ with $\sigma \in \Sigma = [\sigma_-, \sigma_+] \subset (0, \infty)$. The model distributions $\{P_{\sigma,F} : \sigma \in \Sigma, F \in \mathcal{D}[0, 1]\}$ for X have densities of the form:

$$p_{\sigma,F}(x) = \int_0^1 \phi_\sigma(x - z) dF(z). \quad (11)$$

The semiparametric problem then consists of estimation of σ , in the presence of the nuisance parameter F . For later reference, we note that the model has an envelope $[U, L]$ that can be

described by:

$$\begin{aligned} U(x) &= \frac{\sigma_+}{\sigma_-} \left(\phi_{\sigma_+}(x) 1_{\{x < 0\}} + \phi_{\sigma_+}(x+1) 1_{\{x > 1\}} + \phi_{\sigma_+}(0) 1_{\{-M \leq x \leq M\}} \right), \\ L(x) &= \frac{\sigma_-}{\sigma_+} \left(\phi_{\sigma_-}(x+1) 1_{\{x < 1/2\}} + \phi_{\sigma_-}(x) 1_{\{x \geq 1/2\}} \right). \end{aligned} \quad (12)$$

The earliest Bayesian analyses of the normal location model based on the Dirichlet process prior (see, *e.g.* Ferguson (1973) [25]) can be found in Ferguson (1983) [27] and Lo (1984) [60]; a more modern perspective with Dirichlet priors is found in Ghosal and van der Vaart (2001,2007) [32, 33] and Kleijn and van der Vaart (2006) [45]. Numerical studies have been carried out in, for example, in Escobar and West (1995) [24]. Bernstein-von Mises-type Bayesian efficiency in semiparametric mixture models has not been considered in the literature yet [11]. Throughout this paper, the normal location model serves as an example and in subsection 7.3 a Bernstein-von Mises conjecture for this model is formulated.

To give a more practically useful example of a semiparametric mixture model, consider the slightly more complicated but very similar *errors-in-variables regression model* (for an overview, see Anderson (1984) [1]), in which we observe an *i.i.d.* sample from (X, Y) related to an unobserved random variable Z through the regression equations,

$$X = Z + e, \quad \text{and} \quad Y = g_\theta(Z) + f,$$

where, usually, the errors (e, f) are assumed standard-normal and independent. (For a Bayesian analysis involving rates of convergence with non-parametric regression families, see chapter 4 in Kleijn (2003) [44].) The most popular formulation of the model involves a family of regression functions that is linear: $g_{\alpha, \beta}(z) = \alpha + \beta z$ and a completely unknown distribution F for the unobserved $Z \sim F$. Interest then goes to estimation of the parameter $\theta = (\alpha, \beta)$, while treating F as the nuisance parameter (see van der Vaart (1996) [75] and Taupin (2001) [73]). \square

Example (*Cox' proportional hazards model [15]*)

In medical studies (and many in other disciplines) one is interested in the relationship between the time of “survival” (which can mean anything from time until actual death, to onset of a symptom, or detection of a certain protein in a patients blood, etc.) and covariates believed to be of influence (like a regime of medication or specific patient habits). Observations consist of pairs (T, Z) associated with individual patients, where T is the survival time and Z is a vector of covariates. The probability of non-survival between t and $t + dt$, given survival up to time t is called the *hazard function* $\lambda(t)$,

$$\lambda(t) dt = P(t \leq T \leq t + dt \mid T \geq t).$$

The *Cox proportional hazards model* prescribes that the *conditional hazard function* given Z is of the form,

$$\lambda(t|Z) dt = e^{\theta^T Z} \lambda_0(t),$$

where λ_0 is the so-called *baseline hazard function*. The interpretation of the parameter of interest θ is easily established: if, for example, the component $Z_i \in \{0, 1\}$ describes presence

(or not) of certain characteristics in the patient (*e.g.* $Z = 0$ for a non-smoker and $Z = 1$ for a smoker), then e^{θ_i} is the ratio of hazard rates between two patients that differ only in that one characteristic Z_i . The parameter of interest is the vector of θ , while the baseline hazard rate is treated as an unknown nuisance parameter. Early Bayesian references in this model include Ferguson (1979) [26], Kalbfleisch (1978) [39] and Hjort (1990) [36]; see also De Blasi and Hjort (2007, 2009) [18, 19]. Kim and Lee (2004) [40] show that the posterior for the cumulative hazard function under right-censoring converges at rate $n^{-1/2}$ to a Gaussian centred at the Aalen-Nelson estimator for a class of neutral-to-the-right process priors. In Kim (2006) [41] the posterior for the baseline cumulative hazard function and regression coefficients in Cox' proportional hazard model are considered with similar priors. See Castillo (2012) [10] for a Bernstein-von Mises theorem for the proportional hazard rate. \square

Example (*Partial linear regression [12, 4, 61, 77]*)

Consider a situation in which one observes a vector $(Y; U, V)$ of random variables, assumed related through the regression equation,

$$Y = \theta U + \eta(V) + e, \quad (13)$$

with e independent of the pair (U, V) and such that $Ee = 0$, usually assumed normally distributed. The rationale behind this model would arise from situations where one is observing a linear relationship between two random variables Y and U , contaminated by an additive influence from V of largely unknown form. The parameter $\theta \in \mathbb{R}$ is of interest while the nuisance η is from some infinite-dimensional function space H . It is assumed that (U, V) has an unknown distribution P , Lebesgue absolutely continuous with density $p : \mathbb{R}^2 \rightarrow \mathbb{R}$. The distribution P is assumed to be such that $PU = 0$, $PU^2 = 1$ and $PU^4 < \infty$. At a later stage, we also impose $P(U - E[U|V])^2 > 0$ and a smoothness condition on the conditional expectation $v \mapsto E[U|V = v]$. Bayesian efficient estimation of the linear coefficient in the partial linear regression model was first discussed in Kimeldorf and Wahba (1970) [42]: the nuisance lies in the Sobolev space $H^k[0, 1]$ with prior defined through the zero-mean Gaussian process [78],

$$\eta(t) = \sum_{i=0}^k Z_i \frac{t^i}{i!} + (I_{0+}^k W)(t), \quad (14)$$

where $W = \{W_t : t \in [0, 1]\}$ is Brownian motion on $[0, 1]$, (Z_0, \dots, Z_k) form a W -independent, $N(0, 1)$ -*i.i.d.* sample and I_{0+}^k denotes $(I_{0+}^1 f)(t) = \int_0^t f(s) ds$, or $I_{0+}^{i+1} f = I_{0+}^1 I_{0+}^i f$ for all $i \geq 1$. Below, we summarize the analysis given in Bickel and Kleijn (2012) [5] with bounded Sobolev spaces and conditioned versions of the prior process (14) for the nuisance η . For another Bayesian analyses of the partial linear problem, see Shen (2002) [69]. MCMC simulations based on Gaussian priors have been carried out by Shively, Kohn and Wood (1999) [70]. \square

Example (*Semiparametric support boundary estimation [37, 54, 58]*)

Consider a model of densities with a discontinuity at θ . Observed is an *i.i.d.* sample X_1, X_2, \dots with marginal P_0 . The distribution P_0 is assumed to have a density with unknown location

θ for a nuisance density η in some space H . Model distributions $P_{\theta,\eta}$ are then described by densities,

$$p_{\theta,\eta} : [\theta, \infty) \rightarrow [0, \infty) : x \mapsto \eta(x - \theta),$$

for $\eta \in H$ and $\theta \in \Theta \subset \mathbb{R}$.

A Bayesian analysis of this irregular semiparametric estimation problem can be found in Kleijn and Knapik (2013) [47]. Our interest does not lie in modelling of the tail and we concentrate on specifying the behaviour at the discontinuity. For given $S > 0$, let \mathcal{L} denote the ball of radius S in the space $(C[0, \infty], \|\cdot\|_\infty)$ of continuous functions from the extended half-line to \mathbb{R} with uniform norm. Let $\alpha > S$ be fixed. We assume that $\eta : [0, \infty) \rightarrow [0, \infty)$ is differentiable and that $\dot{\ell}(t) = \eta'(t)/\eta(t)$ is a bounded continuous function with a limit at infinity. Define H as the image of \mathcal{L} under the map that takes $\dot{\ell} \in \mathcal{L}$ into densities $\eta_{\dot{\ell}}$ by an Esscher transform of the form,

$$\eta_{\dot{\ell}}(x) = \frac{e^{-\alpha x + \int_0^x \dot{\ell}(t) dt}}{\int_0^\infty e^{-\alpha y + \int_0^y \dot{\ell}(t) dt} dy}, \quad (15)$$

for $x \geq 0$. This map is uniform-to-Hellinger continuous (see Lemma 4.1 in [47]). To define a prior on this model, let $\{W_t : t \in [0, 1]\}$ be Brownian motion on $[0, 1]$ and let Z be independent and distributed $N(0, 1)$. We define the prior $\Pi_{\mathcal{L}}$ on \mathcal{L} as the distribution of the process,

$$\dot{\ell}(t) = S \Psi(Z + W_t), \quad (16)$$

where $\Psi : [-\infty, \infty] \rightarrow [-1, 1] : x \mapsto 2 \arctan(x)/\pi$. Then $\mathcal{L} \subset \text{supp}(\Pi_{\mathcal{L}})$. Then $\mathcal{L} \subset \text{supp}(\Pi_{\mathcal{L}})$. \square

3.2 Efficiency in semiparametric models

The strategy for finding efficient semiparametric estimators for θ_0 is based on the following argument: suppose that $\Theta \subset \mathbb{R}$ and that \mathcal{P}_0 is submodel of \mathcal{P} containing $P_0 = P_{\theta_0, \eta_0}$. Then estimation of θ_0 in the model \mathcal{P}_0 is no harder than it is in \mathcal{P} . When applied to smooth parametric models, this self-evident truth implies that estimation of the parameter θ is more accurate in \mathcal{P}_0 than in \mathcal{P} in the large sample limit. According to theorem 2.7, the Fisher information associated with θ in the larger model is smaller than or equal to that in the smaller model. Semiparametric information bounds are obtained as infima over the information bounds one obtains from collections of smooth, finite-dimensional submodels. That collection has to be somehow ‘‘rich enough’’ to capture the sharp information bound for (regular) semiparametric estimators (which represents the price one pays for a more general model).

Like in section 2, we introduce smoothness and regularity, here with respect to a collection of submodels. For simplicity assume that Θ is open in \mathbb{R} . Let U be an open neighbourhood of θ_0 and consider a map $\gamma : U \rightarrow \mathcal{P} : \theta \mapsto P_\theta$ such that $P_{\theta=\theta_0} = P_0$. To impose smoothness of γ , assume that there exists a P_0 -square-integrable score function $\dot{\ell}$ such that $P_0 \dot{\ell} = 0$ and

the LAN property is satisfied:

$$\log \prod_{i=1}^n \frac{p_{\theta_0+n^{-1/2}h_n}(X_i)}{p_0} = \frac{1}{\sqrt{n}} \sum_{i=1}^n h \dot{\ell}(X_i) - \frac{1}{2} h^2 P_0 \dot{\ell}^2 + o_{P_0}(1),$$

for $h_n \rightarrow h$. Let \mathcal{S} denote a collection of such smooth submodels. The corresponding collection of score functions $\{\dot{\ell} \in L_2(P_0) : \gamma \in \mathcal{S}\}$ may not be closed, but there exists an $\tilde{\ell}_{\mathcal{S}}$ in its $L_2(P_0)$ -closure such that:

$$\tilde{I}_{\mathcal{S}} := P_0 \tilde{\ell}_{\mathcal{S}}^2 = \inf_{\{\dot{\ell} : \gamma \in \mathcal{S}\}} P_0 \dot{\ell}^2. \quad (17)$$

In this text we refer to $\tilde{\ell}_{\mathcal{S}}$ and $\tilde{I}_{\mathcal{S}}$ as the *efficient score function* and *efficient Fisher information* for θ_0 at P_{θ_0, η_0} relative to \mathcal{S} . The efficient Fisher information $\tilde{I}_{\mathcal{S}}$ captures the notion of an “infimal Fisher information” (over \mathcal{S}) alluded to above. Clearly, $\tilde{I}_{\mathcal{S}}$ decreases if we enrich \mathcal{S} .

Call any estimator sequence (T_n) for θ_0 *regular* with respect to \mathcal{S} , if (T_n) is regular as an estimator for θ_0 in all $\gamma \in \mathcal{S}$ (c.f. definition 2.1). Theorem 2.7 applies in any $\gamma \in \mathcal{S}$ so we obtain a collection of Fisher information bounds, one for each $\gamma \in \mathcal{S}$. This implies that for any (T_n) regular with respect to \mathcal{S} , a convolution theorem can be formulated in which the inverse *efficient* Fisher information $\tilde{I}_{\mathcal{S}}$ represents a lower bound to estimation accuracy. For the following theorem, which can be found as theorem 25.20 in [77], define the *tangent set* to be $\{a\dot{\ell} : a \in [0, \infty), \gamma \in \mathcal{S}\} \subset L_2(P_0)$.

Theorem 3.1. (Semiparametric convolution, van der Vaart (1988) [74])

Let Θ be an open subset of \mathbb{R}^k and let H be an infinite-dimensional nuisance space and let \mathcal{P} be the corresponding semiparametric model. Let a collection of smooth submodels \mathcal{S} be given. Assume that the true distribution of the i.i.d. data is P_{θ_0, η_0} . For any estimator sequence (T_n) that is regular with respect to \mathcal{S} , the asymptotic covariance matrix is lower bounded by $\tilde{I}_{\mathcal{S}}$. Furthermore, if the tangent set is a convex cone, the limit distribution of (T_n) is of the form $N(0, \tilde{I}_{\mathcal{S}}^{-1}) * M$ for some probability distribution M .

If the collection \mathcal{S} of smooth submodels is too small, the efficient Fisher information relative to \mathcal{S} is too optimistic and, hence, $\tilde{I}_{\mathcal{S}}$ does not capture the semiparametric information bound, i.e. $\tilde{I}_{\mathcal{S}}$ does not give rise to a *sharp* bound on the asymptotic variance of regular estimators. To illustrate, consider an \mathcal{S} consisting only of $\gamma(\theta) = P_{\theta, \eta_0}$: in that case $\tilde{I}_{\mathcal{S}} = I_{\theta_0, \eta_0}$, the Fisher information associated with the score for θ . In such cases, optimal regular sequences (T_n) in theorem 3.1 (in the sense that $M = \delta_0$) do not exist in general (see, however, the definition of *adaptivity* in section 4). To continue the above illustration, generically semiparametric estimators for θ do not achieve the information bound I_{θ_0, η_0} ; that bound is associated with (parametric) estimation of θ in the presence of a *known* nuisance η_0 . To avoid this situation, one aims to reveal a class \mathcal{S} of smooth submodels that is sufficiently rich.

Definition 3.2. (*Efficient score and Fisher information*)

The efficient score function and efficient Fisher information relative to the *maximal* \mathcal{S} containing all LAN submodels are referred to as *the* efficient score function and *the* efficient Fisher information, denoted by $\tilde{\ell}_{\theta_0, \eta_0}$ and $\tilde{I}_{\theta_0, \eta_0}$ respectively. \square

The above implies a strategy for proving semiparametric efficiency: we make a clever proposal for a sequence of estimators (T_n) and for a collection \mathcal{S} of smooth submodels such that (T_n) is regular with respect to every $\gamma \in \mathcal{S}$ and attains the associated information bound, *i.e.* (T_n) is asymptotically normal with variance $\tilde{I}_{\mathcal{S}}$. By implication, the collection \mathcal{S} is “rich enough” and (T_n) is efficient. (Compare this with the manner in which we concluded that, under the conditions of theorem 2.1, the parametric ML estimator is efficient.)

Theorem 3.3. (Semiparametric efficiency)

Let \mathcal{S} denote a collection of smooth submodels of \mathcal{P} with corresponding efficient Fisher information $\tilde{I}_{\mathcal{S}}$. Let (T_n) be a regular estimator sequence for the parameter of interest. If,

$$n^{1/2}(T_n - \theta_0) \overset{\theta_0, \eta_0}{\rightsquigarrow} N(0, \tilde{I}_{\mathcal{S}}^{-1}),$$

then $\tilde{I}_{\theta_0, \eta_0} = \tilde{I}_{\mathcal{S}}$ and (T_n) is best-regular.

Like in the parametric case, semiparametric estimators (T_n) for θ_0 are best-regular *if and only if* the (T_n) are asymptotically linear, that is,

$$n^{1/2}(T_n - \theta_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \tilde{I}_{\theta_0, \eta_0}^{-1} \tilde{\ell}_{\theta_0, \eta_0}(X_i) + o_{P_{\theta}}(1). \quad (18)$$

(For a proof, see lemma 25.23 in [77].)

Efficiency in semiparametric statistics has yielded a rich literature [4], of which a large part concerns the so-called *calculus of scores*: in order to study sets of tangents \mathcal{S} and to control expansions of likelihoods, a theory has been built up around score functions and Fisher information coefficients in non-parametric models. The central fact of that theory is the orthogonality of the the efficient score to pure-nuisance scores. More precisely, $\tilde{\ell}_{\theta_0, \eta_0}$ is orthogonal in $L_2(P_{\theta_0, \eta_0})$ to all scores g associated with smooth submodels of the form $\gamma = \{P_{\theta_0, \eta_t} : t \in I\}$ (for some neighbourhood I of 0). (If not, one would be able to redefine all other smooth curves in such a way that the efficient Fisher information would be strictly smaller than before.) In fact, $\tilde{\ell}_{\theta_0, \eta_0}$ equals the $L_2(P_0)$ -projection of the ordinary score $\dot{\ell}_{\theta_0, \eta}$ onto the (closed) complement of the span of all scores g associated with variation of the nuisance.

4 Posterior consistency under perturbation

To discuss efficiency in Bayesian semiparametric context we introduce a form of posterior convergence that describes contraction around a curve rather than to a point. Alternatively one may think of this type of convergence as posterior consistency in the nuisance model with random perturbations of the θ -parameter. The curve in question is a so-called least-favourable submodel.

4.1 Least-favourable submodels

The conclusion of the previous section suggests one reasons as follows: in order to find a semiparametric efficient estimator, one would like to concentrate on a single smooth submodel

$\tilde{\gamma}$ for which the Fisher information equals the efficient Fisher information. If it exists, such a $\tilde{\gamma}$ is called a *least-favourable submodel* (see Stein (1956) [71] and [4, 77]). Any estimator sequence that is best-regular in $\tilde{\gamma}$ is automatically best-regular for θ_0 in \mathcal{P} , based on theorem 3.1.

To illustrate, consider the class of so-called *adaptive* problems (Bickel (1982) [3] and section 2.4 of Bickel *et al.* [4]): in adaptive estimation problems, the least-favourable direction *equals* the θ -direction and $\tilde{\gamma} = \{P_{\theta, \eta_0} : \theta \in \Theta\}$ is a least-favourable submodel. The efficient Fisher information equals the ordinary Fisher information for θ , signalling that estimation of θ is equally hard whether the true nuisance η_0 is known or not. An example of an adaptive semiparametric problem is Stein’s original *symmetric location problem* (see section 3.1): apparently, knowledge of the details of the symmetric density η_0 cannot be used to improve asymptotic performance of location estimators.

This useful perspective is not limited to the class of adaptive problems, as long as one is willing to re-define the nuisance parameter somewhat: if there exists an open neighbourhood U_0 of θ_0 and a map $\eta^* : U_0 \rightarrow H$ such that $\tilde{\gamma} = \{P_{\theta, \eta^*(\theta)} : \theta \in U_0\}$ forms a least-favourable submodel, it is possible to define an *adaptive* reparametrization: for all $\theta \in U_0$, $\eta \in H$:

$$(\theta, \eta(\theta, \zeta)) = (\theta, \eta^*(\theta) + \zeta), \quad (\theta, \zeta(\theta, \eta)) = (\theta, \eta - \eta^*(\theta)), \quad (19)$$

inviting the notation $Q_{\theta, \zeta}^* = P_{\theta, \eta^*(\theta) + \zeta}$. With $\zeta = 0$, $\theta \mapsto Q_{\theta, 0}$ describes the least-favourable submodel, implying that estimation of θ in the model for $Q_{\theta, 0}$ is adaptive. With a non-zero value of ζ , $\theta \mapsto Q_{\theta, \zeta}$ describes a version of the least-favourable submodel translated over a nuisance direction. Somewhat disappointingly, in many semiparametric problems least-favourable submodels like $\tilde{\gamma}$ do not exist. But even if *exactly least-favourable* submodels do not exist in a given problem, *approximately least-favourable* submodels may be substituted (see section 5).

At this stage, we leave the above remark for what it is and switch to the Bayesian perspective on semiparametric questions. Assuming measurability of the map $(\theta, \eta) \mapsto P_{\theta, \eta}$, we place a product prior $\Pi_{\Theta} \times \Pi_H$ on $\Theta \times H$ to define a prior on \mathcal{P} and calculate the posterior, in particular, the marginal posterior for the parameter of interest ($A = B \times H$ in (1), for measurable subsets B of Θ). Asymptotically the full posterior concentrates around least-favourable submodels. To see why, let us assume that for each θ in a neighbourhood U_0 of θ_0 , there exists a minimizer $\eta^*(\theta)$ of the Kullback-Leibler divergence,

$$-P_0 \log \frac{p_{\theta, \eta^*(\theta)}}{p_{\theta_0, \eta_0}} = \inf_{\eta \in H} \left(-P_0 \log \frac{p_{\theta, \eta}}{p_{\theta_0, \eta_0}} \right), \quad (20)$$

giving rise to a submodel $\mathcal{P}^* = \{P_{\theta}^* = P_{\theta, \eta^*(\theta)} : \theta \in U_0\}$. As is well-known [67], if \mathcal{P}^* is smooth it constitutes a least-favourable submodel and scores along \mathcal{P}^* are efficient. In the following, we refer to \mathcal{P}^* as a “least-favourable submodel” (whether it is smooth or not).

Example (*Partial linear regression, cont.*)

The partial linear model has a well-defined least-favourable submodel \mathcal{P}^* : for any θ and η ,

$$-P_{\theta_0, \eta_0} \log(p_{\theta, \eta} / p_{\theta_0, \eta_0}) = \frac{1}{2} P_{\theta_0, \eta_0} ((\theta - \theta_0)U + (\eta - \eta_0)(V))^2,$$

so that for fixed θ , minimal KL-divergence over H obtains at $\eta^*(\theta) = \eta_0 - (\theta - \theta_0) E[U|V]$, P_0 -almost-surely. This defines a smooth least-favourable submodel $\mathcal{P}^* = \{P_{\theta, \eta^*(\theta)} : \theta \in \Theta\}$. The efficient score function equals $\tilde{\ell}_{\theta_0, \eta_0} = e(U - E[U|V])$ and the KL-divergence of P_θ^* with respect to P_0 is,

$$-P_{\theta_0, \eta_0} \log(p_\theta^*/p_{\theta_0, \eta_0}) = \frac{1}{2} \tilde{I}_{\theta_0, \eta_0} (\theta - \theta_0)^2,$$

with the efficient Fisher information as the coefficient of the leading, second order. The absence of a linear term characterizes least-favourable submodels. \square

Example (*Normal location mixtures, cont.*)

We view $\mathcal{D}[0, 1]$ as a convex, closed subset of the unit sphere in the dual of $C[0, 1]$, the space of all continuous functions on $[0, 1]$ with uniform norm. As a consequence of the Banach-Alaoglu theorem, $\mathcal{D}[0, 1]$ is weak-* compact [23]. Fix $\sigma \in \Sigma$ and consider the map $\mathcal{D}[0, 1] \rightarrow \mathcal{P} : F \mapsto P_{\sigma, F}$, *c.f.* (11). Let x be given; $z \mapsto \phi_\sigma(x - z)$ is bounded, so if $(F_\alpha)_{\alpha \in I}$ converges weak-* to F in $\mathcal{D}[0, 1]$ then $p_{\sigma, F_\alpha}(x) \rightarrow p_{\sigma, F}(x)$, which implies that $\log(p_{\sigma, F_\alpha}/p_{\sigma, F})(x) \rightarrow 0$. Using the bracket $[U, L]$ of (12), it is easily seen that $P_{\sigma', G} \log(U/L) < \infty$ for all $\sigma' \in \Sigma$ and $G \in \mathcal{D}[0, 1]$. Hence, by dominated convergence,

$$P_0 \log \frac{p_{\sigma, F_\alpha}}{p_{\sigma, F}} \rightarrow 0,$$

so that the map $F \mapsto -P_0 \log(p_{\sigma, F}/p_0)$ is weak-* continuous. Conclude that for every $\sigma \in \Sigma$, there exists an $F^*(\sigma)$ that minimizes the Kullback-Leibler divergence with respect to P_0 , *i.e.* there exists a “least-favourable submodel” $\mathcal{P}^* = \{P_{\sigma, F^*(\sigma)} : \sigma \in \Sigma\}$. To show that the \mathcal{P}^* is Hellinger continuous, note that for all $P, Q \in \mathcal{P}$ and all $a > 0$,

$$\int_{\{p/q > a\}} p(x) \left(\frac{p}{q}\right)^\delta dx \leq \int U(x) \left(\frac{U}{L}\right)^\delta(x) dx < \infty,$$

for $\delta \in (0, 1]$ such that $\delta \leq (\sigma_+/\sigma_-)^2 - 1$. According to theorem 5 of Wong and Shen (1995) [80], there exists a constant C and an $\epsilon > 0$ such that for all $P, Q \in \mathcal{P}$ with $H(P, Q) < \epsilon$,

$$-P \log \frac{q}{p} \leq C H^2(P, Q) \log \frac{1}{H(P, Q)}.$$

Let $\sigma \in \Sigma$ be given and consider,

$$\begin{aligned} H(P_\sigma^*, P_0) &\leq -P_0 \log \frac{p_\sigma^*}{p_0} = \inf_{F \in \mathcal{D}[0, 1]} -P_0 \log \frac{p_{\sigma, F}}{p_0} \\ &\leq C \inf_{F \in \mathcal{D}[0, 1]} H(P_{\sigma, F}, P_0) \log \frac{1}{H(P_{\sigma, F}, P_0)} \\ &\leq C H(P_{\sigma, F_0}, P_0) \log \frac{1}{H(P_{\sigma, F_0}, P_0)}. \end{aligned}$$

In the mixture model,

$$H(P_{\sigma, F_0}, P_0) \leq \|p_{\sigma, F_0} - p_{\sigma_0, F_0}\|_{1, \mu} \leq \|\phi_\sigma - \phi_{\sigma_0}\|_{1, \mu} = o(1),$$

as $\sigma \rightarrow \sigma_0$. It follows that the dependence $\sigma \mapsto P_\sigma^*$ is Hellinger continuous. Note that we have not shown \mathcal{P}^* to be smooth in any sense. \square

Example (*Support boundary estimation, cont.*)

In the support boundary model, notions of smoothness are lost and subsection 3.2 does not apply: although the central argument stands (estimation of the parameter of interest in submodels is easier), a suitable notion of optimality is lacking. The question remains on which subsets of the model the posterior concentrates. As it turns out [47], this role is taken over by the “adaptive” submodel $\{P_{\theta, \eta_0} : \theta \in \Theta\}$. \square

4.2 Posterior concentration

Neighbourhoods of \mathcal{P}^* are described with Hellinger balls in H of radius $\rho > 0$ around $\eta^*(\theta)$, for all $\theta \in U_0$:

$$D(\theta, \rho) = \{\eta \in H : d_H(\eta, \eta^*(\theta)) < \rho\}. \quad (21)$$

Furthermore, we define for all $\theta \in U_0$ the misspecified nonparametric models $\mathcal{P}_\theta = \{P_{\theta, \eta} : \eta \in H\}$. Kleijn and van der Vaart (2006) [45] show that the misspecified posterior on \mathcal{P}_θ concentrates asymptotically in any Hellinger neighbourhood of the point of minimal Kullback-Leibler divergence with respect to the true distribution of the data. Applied to \mathcal{P}_θ , we see that, under P_0 , Hellinger balls $D(\theta, \rho)$, ($\rho > 0$) receive posterior probability one asymptotically. We formulate this θ -dependent form of posterior convergence in terms of θ -conditional posteriors on H , given $\theta = \theta_n(h_n)$ with $h_n = O_{P_0}(1)$. We view posteriors on the models $\mathcal{P}_{\theta_n(h_n)}$ as random order- $n^{-1/2}$ perturbations of the posterior for \mathcal{P}_{θ_0} .

Definition 4.1. (*Consistency under perturbation*)

Given a rate sequence (ρ_n) , $\rho_n \downarrow 0$, we say that the conditioned nuisance posterior is *consistent under $n^{-1/2}$ -perturbation at rate ρ_n* , if,

$$\Pi_n(D^c(\theta, \rho_n) \mid \theta = \theta_0 + n^{-1/2}h_n; X_1, \dots, X_n) \xrightarrow{P_0} 0, \quad (22)$$

for all bounded, stochastic sequences (h_n) . \square

Note that consistency under perturbation is a property that expresses stability of an essential stage of the analysis (in this case, posterior consistency at rate ρ_n) against $n^{-1/2}$ -perturbation of the θ -component. Regularity of estimator sequences, *c.f.* (7), is another such property and prior stability (see Lemma 5.5) is yet another.

For posterior concentration to occur [66] sufficient prior mass must be present in Kullback-Leibler-type neighbourhoods [31, 45]. Presently these neighbourhoods take the form.

$$K_n(\rho, M) = \left\{ \eta \in H : P_0 \left(\sup_{\|h\| \leq M} -\log \frac{p_{\theta_n(h), \eta}}{p_{\theta_0, \eta_0}} \right) \leq \rho^2, \right. \\ \left. P_0 \left(\sup_{\|h\| \leq M} -\log \frac{p_{\theta_n(h), \eta}}{p_{\theta_0, \eta_0}} \right)^2 \leq \rho^2 \right\}, \quad (23)$$

for $\rho > 0$ and $M > 0$. The following theorem generalizes the main theorem in Ghosal, Ghosh and van der Vaart (2000) [31] to perturbed setting.

Theorem 4.2. (Posterior consistency under perturbation at a rate)

Assume that there exist a Hellinger-continuous \mathcal{P}^* and a sequence (ρ_n) with $\rho_n \downarrow 0$, $n\rho_n^2 \rightarrow \infty$ such that for all $M > 0$ and every bounded, stochastic (h_n) :

(i) There exists a constant $K > 0$ such that for large enough n ,

$$\Pi_H(K_n(\rho_n, M)) \geq e^{-Kn\rho_n^2}. \quad (24)$$

(ii) For all n large enough, $N(\rho_n, H, d_H) \leq e^{n\rho_n^2}$.

(iii) For all $L > 0$ and all bounded, stochastic (h_n) ,

$$\sup_{\{\eta \in H: d_H(\eta, \eta_0) \geq L\rho_n\}} \frac{H(P_{\theta_n(h_n), \eta}, P_{\theta_0, \eta})}{H(P_{\theta_0, \eta}, P_0)} = o(1), \quad (25)$$

and $d_H(\eta^*(\theta_n(h_n)), \eta_0) = o(\rho_n)$.

Then, for every bounded, stochastic (h_n) there exists an $L > 0$ such that the conditional nuisance posterior converges as,

$$\Pi(D^c(\theta, L\rho_n) \mid \theta = \theta_0 + n^{-1/2}h_n; X_1, \dots, X_n) = o_{P_0}(1), \quad (26)$$

under $n^{-1/2}$ -perturbation.

The proof of this theorem can be found in [5] and proceeds through the construction of tests based on the Hellinger geometry of the model, generalizing the approach of Birgé [6, 7] and Le Cam [56] to $n^{-1/2}$ -perturbed context. Consider the problem of testing/estimating η when θ_0 is known: we cover the nuisance model \mathcal{P}_{θ_0} by a minimal collection of Hellinger balls, all of radius (ρ_n) , each of which is testable against P_0 with power bounded by $\exp(-\frac{1}{4}nH^2(P_0, B))$ [56]. The tests for the covering Hellinger balls are combined into a single test for the alternative $\{P : H(P, P_0) \geq \rho_n\}$ against P_0 . The order of the cover controls the power of the combined test. Therefore the construction requires an upper bound to Hellinger metric entropy numbers [6, 7, 49, 76],

$$N(\rho_n, \mathcal{P}_{\theta_0}, H) \leq e^{n\rho_n^2}, \quad (27)$$

which is interpreted as indicative of the nuisance model's complexity in the sense that the lower bound to the collection of rates (ρ_n) solving (27), is the Hellinger minimax rate for estimation of η_0 . In the $n^{-1/2}$ -perturbed problem, the alternative does not just consist of the complement of a Hellinger-ball in the nuisance factor H , but also has an extent in the θ -direction shrinking at rate $n^{-1/2}$. Condition (25) guarantees that Hellinger covers of \mathcal{P}_{θ_0} are large enough to accommodate the θ -extent of the alternative, the implication being that the test sequence one constructs for the nuisance in case θ_0 is known, can also be used when θ_0 is known only up to $n^{-1/2}$ -perturbation. Geometrically, (25) requires that $n^{-1/2}$ -perturbed versions of the nuisance model are contained in a narrowing sequence of metric cones based at P_0 . In differentiable models, the Hellinger distance $H(P_{\theta_n(h_n), \eta}, P_{\theta_0, \eta})$ is typically of order $O(n^{-1/2})$ for all $\eta \in H$. So if, in addition, $n\rho_n^2 \rightarrow \infty$, limit (25) is expected to hold pointwise in η . Then only the uniform character of (25) truly forms a condition.

For corollary 7.3 we have a version of theorem 4.2 that only asserts consistency under $n^{-1/2}$ -perturbation at *some* rate while relaxing bounds for prior mass and entropy. In the statement of the corollary, we make use of the family of Kullback-Leibler neighbourhoods that would play a role for the posterior of the nuisance if θ_0 were known, $K(\rho) = K_{n=1}(\rho, M = 0)$.

Corollary 4.3. (Posterior consistency under perturbation)

Assume that there exists a Hellinger-continuous \mathcal{P}^* and that for all $\rho > 0$, $N(\rho, H, d_H) < \infty$, $\Pi_H(K(\rho)) > 0$ and,

(i) For all $M > 0$ there is an $L > 0$ such that for all $\rho > 0$ and large enough n , $K(\rho) \subset K_n(L\rho, M)$.

(ii) For every bounded random (h_n) , $\sup_{\eta \in H} H(P_{\theta_n(h_n), \eta}, P_{\theta_0, \eta}) = O(n^{-1/2})$.

Then there exists a sequence (ρ_n) , $\rho_n \downarrow 0$, $n\rho_n^2 \rightarrow \infty$, such that the conditional nuisance posterior converges under $n^{-1/2}$ -perturbation at rate (ρ_n) .

4.3 Application in examples

We apply corollary 4.3 in partial linear regression, normal location mixtures and support boundary estimation.

Example (*Partial linear regression, cont.*)

Note that for all $\eta_1, \eta_2 \in H$, $d_H(\eta_1, \eta_2) \leq -P_{\theta_0, \eta_2} \log(p_{\theta_0, \eta_1}/p_{\theta_0, \eta_2}) = \frac{1}{2} \|\eta_1 - \eta_2\|_{2, P}^2 \leq \frac{1}{2} \|\eta_1 - \eta_2\|_\infty^2$. Hence, for any $\rho > 0$, $N(\rho, \mathcal{P}_{\theta_0}, d_H) \leq N((2\rho)^{1/2}, H, \|\cdot\|_\infty) < \infty$. Similarly, one shows that for all η both $-P_0 \log(p_{\theta_0, \eta}/p_{\theta_0, \eta_0})$ and $P_0(\log(p_{\theta_0, \eta}/p_{\theta_0, \eta_0}))^2$ are bounded by $(\frac{1}{2} + D^2) \|\eta - \eta_0\|_\infty^2$. Hence, for any $\rho > 0$, $K(\rho)$ contains a $\|\cdot\|_\infty$ -ball. Assuming that $\eta_0 \in \text{supp}(\Pi_H)$, we see that the primary conditions of corollary 4.3 hold. Next, note that for $M > 0$, $n \geq 1$, $\eta \in H$,

$$\begin{aligned} \sup_{\|h\| \leq M} -\log \frac{p_{\theta_n(h), \eta}}{p_{\theta_0, \eta_0}} &= \frac{M^2}{2n} U^2 + \frac{M}{\sqrt{n}} |U(e - (\eta - \eta_0)(V))| \\ &\quad - e(\eta - \eta_0)(V) + \frac{1}{2}(\eta - \eta_0)^2(V), \end{aligned} \quad (28)$$

where $e \sim N(0, 1)$ under P_{θ_0, η_0} . Note that H is totally bounded in $C[0, 1]$, so that there exists a constant $D > 0$ such that $\|H\|_\infty \leq D$. Together with the help of the independence of e and (U, V) and the assumptions on the distribution of (U, V) , it is then verified that condition (i) of corollary 4.3 holds. Since $(p_{\theta_n(h), \eta}/p_{\theta_0, \eta}(X))^{1/2} = \exp((h/2\sqrt{n})eU - (h^2/4n)U^2)$, one derives the η -independent upper bound,

$$H^2(P_{\theta_n(h_n), \eta}, P_{\theta_0, \eta}) \leq \frac{M^2}{2n} P U^2 + \frac{M^3}{6n^2} P U^4 = O(n^{-1}),$$

for all bounded, stochastic (h_n) , so that condition (ii) of corollary 4.3 holds. \square

Example (*Normal location mixtures, cont.*)

The Hellinger-continuity of \mathcal{P}^* has been established earlier. The nuisance space is of finite

d_H -entropy (see theorem 3.1 in Ghosal and van der Vaart (2001,2007) [32, 33]). Let α be a finite Borel measure on $[0, 1]$ that dominates the Lebesgue measure and define $\Pi_H = D_\alpha$ to be the corresponding Dirichlet process prior for F . This prior satisfies $\Pi_H(K(\rho)) > 0$ for all $\rho > 0$ (as implied by the proof of theorem 5.1 in [32], or subsection 3.2 in Kleijn and van der Vaart (2006) [45]). A lengthy but elementary calculation shows that there exist constants $K_1, K_2 > 0$ such that for all $F \in \mathcal{D}[0, 1]$,

$$\begin{aligned} P_0 \left(\sup_{\|h\| \leq M} \log \frac{p_{\sigma_n(h), F}(X)}{p_0} \right) &\leq P_0 \log \frac{p_{\sigma_0, F}(X)}{p_0} + K_1 \frac{M^2}{n}, \\ P_0 \left(\sup_{\|h\| \leq M} \log \frac{p_{\sigma_n(h), F}(X)}{p_0} \right)^2 &\leq P_0 \left(\log \frac{p_{\sigma_0, F}(X)}{p_0} \right)^2 + K_2 \frac{M^2}{n}, \end{aligned}$$

for large enough n , so that condition (i) of corollary 4.3 holds. Let (h_n) be given. Lemma 17.3 in [72] says that if P, Q are distributions for (X, Z) and $Y = f(X, Z)$ has induced distributions P', Q' , then $H(P', Q') \leq H(P, Q)$. We apply this to (X, Z) and X to obtain,

$$\begin{aligned} &\sup_{F \in \mathcal{D}[0, 1]} n H^2(P_{\sigma_n(h_n), F}, P_{\sigma_0, F}) \\ &\leq \sup_{F \in \mathcal{D}[0, 1]} \iint n \left(\phi_{\sigma_n(h_n)}(x - z)^{1/2} - \phi_{\sigma_0}(x - z)^{1/2} \right)^2 dx dF(z) \\ &= \sup_{z \in [0, 1]} n H^2(\Phi_{\sigma_n(h_n)}, \Phi_{\sigma_0}) = O(1), \end{aligned}$$

i.e. condition (ii) of corollary 4.3 holds. □

Example (*Support boundary estimation, cont.*)

Given $0 < S < \alpha$, we define $\rho_0^2 = \alpha - S > 0$. Consider the distribution Q with Lebesgue density $q > 0$ given by $q(x) = \rho_0^2 e^{-\rho_0^2 x}$ for $x \geq 0$. Then the family $\mathcal{F} = \{x \mapsto \sqrt{\eta_{\dot{\ell}}/q(x)} : \dot{\ell} \in \mathcal{L}\}$ forms a subset of the collection of all monotone functions $\mathbb{R} \mapsto [0, C]$, where C is fixed and depends on α , and S . Referring to Theorem 2.7.5 in van der Vaart and Wellner (1996) [76], we conclude that the $L_2(Q)$ -bracketing entropy $N_{[\cdot]}(\epsilon, \mathcal{F}, L_2(Q))$ of \mathcal{F} is finite for all $\epsilon > 0$. Noting that,

$$d_H(\eta, \eta_0)^2 = d_H(\eta_{\dot{\ell}}, \eta_{\dot{\ell}_0})^2 = \int_{\mathbb{R}} \left(\sqrt{\frac{\eta_{\dot{\ell}}}{q}}(x) - \sqrt{\frac{\eta_{\dot{\ell}_0}}{q}}(x) \right)^2 dQ(x),$$

it follows that $N(\rho, H, d_H) = N(\rho, \mathcal{F}, L_2(Q)) \leq N_{[\cdot]}(2\rho, \mathcal{F}, L_2(Q)) < \infty$. Conclude that for all $\rho > 0$, $N(\rho, H, d_H) < \infty$. Since $\mathcal{L} \subset \text{supp}(\Pi_{\mathcal{F}})$, $\Pi(K(\rho)) > 0$ for all $\rho > 0$. Let $\rho > 0$ be given and let $\dot{\ell} \in \mathcal{L}$ be such that $\|\dot{\ell} - \dot{\ell}_0\|_\infty < \rho^2$. Without reproducing the derivation (see [47]), we state that

$$\begin{aligned} -P_0 \log \frac{p_{\theta_0, \eta}}{p_{\theta_0, \eta_0}} &\leq 2\rho^2 (P_0(X - \theta_0) + O(\rho^2)), \\ P_0 \left(\log \frac{p_{\theta_0, \eta}}{p_{\theta_0, \eta_0}} \right)^2 &\leq \rho^4 (P_0(X - \theta_0)^2 + 3P_0(X - \theta_0) + O(\rho^2)), \end{aligned}$$

which proves that there exists a constant L_1 such that $\{\eta_{\dot{\ell}} \in H : \|\dot{\ell} - \dot{\ell}_0\|_\infty \leq \rho^2\} \subset K(L_1\rho)$. Let $M > 0$ be given. With reasoning very similar to that which led to (28), one shows (see

[47]) that there exists an $L_2 > 0$ such that $K(L_1\rho) \subset K_n(L_2\rho, M)$. According to Lemma 4.4 in [47],

$$n H^2(P_{\theta_n(h_n), \eta}, P_{\theta_0, \eta}) \leq 2M \gamma_{\theta_0, \eta} + O(n^{-1}),$$

for all $\eta \in H$ and all bounded, stochastic (h_n) . According to Corollary 3.1 in [47] (which is completely analogous to Corollary 4.3 above), posterior consistency under n^{-1} -perturbation obtains at some rate (ρ_n) . \square

5 Local expansions for integrated likelihoods

Since the prior is of product form, the marginal posterior for the parameter $\theta \in \Theta$ depends on the nuisance factor only through the integrated likelihood ratio,

$$S_n : \Theta \rightarrow \mathbb{R} : \theta \mapsto \int_H \prod_{i=1}^n \frac{p_{\theta, \eta}}{p_{\theta_0, \eta_0}}(X_i) d\Pi_H(\eta). \quad (29)$$

(The localized version of S_n is denoted $h \mapsto s_n(h)$, $s_n(h) = S_n(\theta_0 + n^{-1/2}h)$.) The quantity S_n plays a central role in this section and the next, similar to that of the *profile likelihood* in semiparametric maximum-likelihood methods (see, *e.g.*, Severini and Wong (1992) [67] and Murphy and van der Vaart (2000) [62]), in the sense that S_n embodies the intermediate stage between nonparametric and semiparametric steps of the estimation procedure. Presently, we are interested in the local behaviour of S_n : the smoothness condition in the parametric Bernstein-Von Mises theorem is a LAN expansion of the likelihood, which is replaced in semiparametric context by a stochastic LAN expansion of the integrated likelihood (29). In this section, we consider sufficient conditions on model and prior for the following property.

Definition 5.1. (*Integral local asymptotic normality (ILAN)*)

The quantity S_n has the *integral LAN* property if s_n allows an expansion of the form,

$$\log \frac{s_n(h_n)}{s_n(0)} = \frac{1}{\sqrt{n}} \sum_{i=1}^{\infty} h_n^T \tilde{\ell}_{\theta_0, \eta_0} - \frac{1}{2} h_n^T \tilde{I}_{\theta_0, \eta_0} h_n + o_{P_0}(1), \quad (30)$$

for every random sequence $(h_n) \subset \mathbb{R}^k$ of order $O_{P_0}(1)$. \square

In Bickel and Kleijn (2012) [5] the LAN analysis departs from the assumption that the model possesses a *smooth* least-favourable submodel for which we can establish posterior consistency under perturbation. As we have seen above, the partial-linear regression model has such a smooth least-favourable submodel and corollary 5.2 of [5] applies. But in semiparametric mixture models (and this is generic in semiparametric models), no such guarantee can be given: although a Hellinger-continuous \mathcal{P}^* exists for which consistency under perturbation obtains, *smoothness* of this curve has not been shown and corollary 5.2 of [5] cannot be invoked. Below, we generalize the analysis to models that do not possess smooth least-favourable submodels.

5.1 Approximately least-favourable submodels

Theorem 5.7 below proves the ILAN property under three conditions, consistency under $n^{-1/2}$ -perturbation for the nuisance posterior, sLAN expansions of model distributions and a domination condition. If \mathcal{P}^* can be approximated by sLAN models (in a suitable sense, see properties (33)–(35)) then one can lift the sLAN expansion of the integrand in (29) to an ILAN expansion of the form (30). Since the posterior concentrates in neighbourhoods of \mathcal{P}^* , only the least-favourable expansion at η_0 contributes to (30) asymptotically. For this reason, the integral LAN expansion is determined by the efficient score function (and not some other influence function). Ultimately, occurrence of the efficient score lends the marginal posterior (and statistics based upon it) properties of frequentist inferential optimality, in accordance with theorem 3.1.

In the derivation of theorem 5.7, the model is reparametrized *c.f.* (19) with approximately least-favourable models replacing η^* . To be more precise, we consider n -dependent model reparametrizations of the following form: for all $\theta \in U_0$, $\eta \in H$,

$$(\theta, \eta_n(\theta, \zeta)) = (\theta, \eta_n(\theta) + \zeta), \quad (\theta, \zeta_n(\theta, \eta)) = (\theta, \eta - \eta_n(\theta)), \quad (31)$$

depending on models $\mathcal{P}_n = \{P_{n,\theta} = P_{\theta, \eta_n(\theta)} : \theta \in U_0\}$ and we introduce the notation $Q_{n,\theta,\zeta} = P_{\theta, \eta_n(\theta) + \zeta}$.

Definition 5.2. (*Approximate least-favourability*)

Given a Hellinger-continuous \mathcal{P}^* , a sequence of submodels (\mathcal{P}_n) *c.f.* (31) is called *approximately least-favourable* at P_0 if it satisfies the following conditions. For all $n \geq 1$, $P_0 \in \mathcal{P}_n$ (*i.e.* $\eta_n(\theta_0) = \eta_0$) and the model is sLAN at θ_0 in the \mathcal{P}_n -direction(s) for all ζ in a Hellinger neighbourhood of $\zeta = 0$: noting that $Q_{n,\theta_0,\zeta} = P_{\theta_0, \eta_0 + \zeta}$ for all $n \geq 1$, we assume there exist $g_{n,\zeta} \in L_2(P_{\theta_0, \eta_0 + \zeta})$ such that for every random (h_n) that is bounded in $P_{\theta_0, \eta_0 + \zeta}$ -probability,

$$\log \prod_{i=1}^n \frac{q_{n,\theta_n,\zeta}(X_i)}{q_{n,\theta_0,\zeta}(X_i)} = \frac{1}{\sqrt{n}} \sum_{i=1}^n h_n^T g_{n,\zeta}(X_i) - \frac{1}{2} h_n^T I_{n,\zeta} h_n + R_n(h_n, \zeta), \quad (32)$$

where $\theta_n = \theta_0 + n^{-1/2} h_n$, $I_{n,\zeta} = Q_{n,\theta_0,\zeta}(g_{n,\zeta} g_{n,\zeta}^T)$ and $R_n(h_n, \zeta) = o_{P_{\theta_0, \eta_0 + \zeta}}(1)$. Furthermore, the models \mathcal{P}_n are assumed to approximate \mathcal{P}^* such that,

(i) the scores converge to the efficient score function in $L_2(P_0)$,

$$P_0 \|g_{n,0} - \tilde{\ell}_{\theta_0, \eta_0}\|^2 \rightarrow 0, \quad (33)$$

and there exists a rate (ρ_n) , $\rho_n \downarrow 0$ and $n\rho_n^2 \rightarrow \infty$ such that for all $M > 0$:

(ii) the models \mathcal{P}_n approximate \mathcal{P}^* ,

$$\sup_{\|h\| \leq M} d_H(\eta_n(\theta_n(h)), \eta^*(\theta_n(h))) = o(\rho_n), \quad (34)$$

(iii) and the quantities $U_n(\rho, h)$ defined by (see *Notation and conventions*),

$$U_n(\rho, h) = \sup_{\zeta \in B(\rho)} Q_{n,\theta_0,\zeta}^n \left(\prod_{i=1}^n \frac{q_{n,\theta_n(h),\zeta}(X_i)}{q_{n,\theta_0,\zeta}(X_i)} \right), \quad (35)$$

satisfies $U(\rho_n, h_n) = O(1)$ for all bounded, stochastic (h_n) .

□

The last requirement may be hard to interpret; however, for a single, fixed ζ , the condition says that the likelihood ratio remains integrable if we replace $\theta_n(h_n)$ by the maximum-likelihood estimator $\hat{\theta}_n(X_1, \dots, X_n)$ (see lemma 5.3); condition (35) imposes this uniformly over neighbourhoods of $\zeta = 0$. The following lemma shows that first-order Taylor expansions of likelihood ratios combined with a uniform limit for certain Fisher information coefficients suffices to satisfy $U(\rho_n, h_n) = O(1)$ for all bounded, stochastic (h_n) and every $\rho_n \downarrow 0$.

Lemma 5.3. *Let Θ be open in \mathbb{R} . Assume that there exists a $\rho > 0$ such that for all $\zeta \in B(\rho) = \{\zeta : d_H(\eta_0 + \zeta, \eta_0) < \rho\}$ and all x in the samplespace, the maps $\theta \mapsto \log(q_{n,\theta,\zeta}/q_{n,\theta_0,\zeta})(x)$ are continuously differentiable on $[\theta_0 - \rho, \theta_0 + \rho]$ with Lebesgue-integrable derivatives $g_{n,\theta,\zeta}(x)$ such that,*

$$\sup_{\zeta \in B(\rho)} \sup_{\{\theta: |\theta - \theta_0| < \rho\}} Q_{n,\theta,\zeta}(g_{n,\theta,\zeta})^2 = O(1). \quad (36)$$

Then, for every $\rho_n \downarrow 0$ and all bounded, stochastic (h_n) , $U_n(\rho_n, h_n) = O(1)$.

Proof See section 8. □

5.2 Integrated local asymptotic normality

Reparametrization leads to n - and θ -dependence in the prior for ζ . Below, it is shown that the prior mass of the relevant Hellinger neighbourhoods displays a type of stability, under a condition on the local behaviour of Hellinger distances in the least-favourable submodel and its approximations. Because these approximately least-favourable submodels are smooth, typically $d_H(\eta_n(\theta_n(h_n)), \eta_0) = O(n^{-1/2})$ for all bounded, stochastic (h_n) , which suffices for typical rates (ρ_n) .

Lemma 5.5. (Prior stability)

Assume that there exists a Hellinger-continuous \mathcal{P}^ and approximately least-favourable \mathcal{P}_n . Let (h_n) be a bounded, stochastic sequence of perturbations and let Π_H be any prior on H . For any rate (ρ_n) , $\rho_n \downarrow 0$ and $n\rho_n^2 \rightarrow \infty$ such that (34) is satisfied,*

$$\Pi_H(D(\theta_n(h_n), \rho_n)) = \Pi_H(D(\theta_0, \rho_n)) + o(1). \quad (37)$$

Proof See section 8. □

Prior stability is part of the construction underpinning the following theorem which roughly says that in models that allow approximately least-favourable submodels, consistency under $n^{-1/2}$ -perturbation is sufficient for the expansions of the form (30).

Theorem 5.7. (Integral local asymptotic normality)

Assume that there exists a Hellinger-continuous \mathcal{P}^ and approximately least-favourable \mathcal{P}_n . Furthermore assume that the posterior is consistent under $n^{-1/2}$ -perturbation at a rate (ρ_n) that is also valid in (34) and (35). Then the integral LAN-expansion (30) holds.*

Proof See section 8. □

With regard to the nuisance rate (ρ_n) , we first note that the proof of theorem 5.7 fails if the slowest rate required to satisfy (35) vanishes *faster* than the optimal rate for convergence under $n^{-1/2}$ -perturbation (as determined in (27) and (24)). However, the rate (ρ_n) does not appear in assertion (30), so if said contradiction between conditions (35) and (27)/(24) does not occur, the sequence (ρ_n) can remain entirely internal to the proof of theorem 5.7. More particularly, if condition (35) holds for *any* (ρ_n) such that $n\rho_n^2 \rightarrow \infty$ (as in lemma 5.3), integral LAN only requires consistency under $n^{-1/2}$ -perturbation at *some* such (ρ_n) . In that case, we may appeal to corollary 4.3 instead of theorem 4.2, thus relaxing conditions on model entropy and nuisance prior. Lemma 5.3 and this shortcut are used in all three examples of subsection 5.3.

5.3 Application in examples

In the partial linear example \mathcal{P}^* is a smooth least-favourable submodel. As a result, the formulation of [5] can be used (or choose $\mathcal{P}_n = \mathcal{P}^*$ for all $n \geq 1$ and apply the theorems of this paper).

Example (*Partial linear regression, cont.*)

For given (h_n) , $n \geq 1$ and fixed ζ , the submodel $\theta \mapsto Q_{\theta, \zeta}$ satisfies,

$$\begin{aligned} \log \prod_{i=1}^n \frac{p_{\theta_0 + n^{-1/2}h_n, \eta^*(\theta_0 + n^{-1/2}h_n) + \zeta}(X_i)}{p_{\theta_0, \eta_0 + \zeta}} & \\ &= \frac{h_n}{\sqrt{n}} \sum_{i=1}^n g_\zeta(X_i) - \frac{1}{2}h_n^2 P_{\theta_0, \eta_0 + \zeta} g_\zeta^2 + \frac{1}{2}h_n^2 (\mathbb{P}_n - P)(U - \mathbb{E}[U|V])^2, \end{aligned} \quad (38)$$

for all stochastic (h_n) , with $g_\zeta(X) = e(U - \mathbb{E}[U|V])$, $e = Y - \theta_0 U - (\eta_0 + \zeta)(V) \sim N(0, 1)$ under $P_{\theta_0, \eta_0 + \zeta}$. Since $PU^2 < \infty$, the last term on the right is $o_{P_{\theta_0, \eta_0 + \zeta}}(1)$ if (h_n) is bounded in probability. We conclude that $\theta \mapsto Q_{\theta, \zeta}$ is stochastically LAN for all ζ . For any $x \in \mathbb{R}^3$ and all ζ , the map $\theta \mapsto \log q_{\theta, \zeta}/q_{\theta_0, \zeta}(x)$ is continuously differentiable on all of Θ , with score $g_{\theta, \zeta}(X) = e(U - \mathbb{E}[U|V]) + (\theta - \theta_0)(U - \mathbb{E}[U|V])^2$. Since $Q_{\theta, \zeta} g_{\theta, \zeta}^2 = P(U - \mathbb{E}[U|V])^2 + (\theta - \theta_0)^2 P(U - \mathbb{E}[U|V])^4$ does not depend on ζ and is bounded over $\theta \in [\theta_0 - \rho, \theta_0 + \rho]$, lemma 5.3 says that $U(\rho_n, h_n) = O(1)$ for all $\rho_n \downarrow 0$ and all bounded, stochastic (h_n) . Posterior consistency under perturbation and the bound on Hellinger distances required for prior stability, *c.f.* Lemma 5.5, were shown to be valid in the previous section. We conclude that the integrated LAN expansion of (30) holds. □

Example (*Normal location mixtures, cont.*)

Let I be a open interval symmetric around 0. Let $\psi : I \rightarrow [0, \infty)$ be a probability density of suitable smoothness. Given a sequence (τ_n) of strictly positive τ_n that decrease to zero monotonously, let $\psi_n : I_n \rightarrow [0, \infty)$ denote the scaled kernel sequence $\psi_n(x) = \tau_n^{-1}\psi(x/\tau_n)$. Smooth the mixing distributions $F^*(\sigma)$ with the kernels ψ_n and shift the resulting curve to

compensate smoothing at F_0 ,

$$G_n : U_0 \rightarrow \mathcal{D}[0, 1] : \sigma \mapsto \int \psi_n(\sigma' - \sigma) F^*(\sigma') d\sigma,$$

and $F_n(\sigma) = (F_0 - G_n(\sigma_0)) + G_n(\sigma)$, to define submodels $\mathcal{P}_n = \{P_{\sigma, F_n(\sigma)} : \sigma \in U_0\}$. If the kernel ψ is smooth enough then for every $n \geq 1$, \mathcal{P}_n is differentiable; the score is a sum of the score for scaling in the normal model $\{\Phi_\sigma : \sigma \in \Sigma\}$ and the score along G_n . Smoothness is not influenced if we change the shift constant $(F_0 - G_n(\sigma_0))$, so smoothness holds for ζ -shifted submodels as well (for all ζ in a Hellinger neighbourhoods of 0). Because \mathcal{P}^* minimizes the Kullback-Leibler divergence, the scores $g_{n,0}$ converge to $\tilde{\ell}_{\sigma_0, F_0}$ c.f. (33). If \mathcal{P}^* is parametrized such that it is locally Lipschitz (of any order) around P_0 , the Hellinger degree of approximation between \mathcal{P}^* and \mathcal{P}_n can be controlled uniformly over U_0 and any rate (ρ_n) is achievable in (34) by letting (τ_n) decrease fast enough. Compare (36) with condition (2.7) in van der Vaart (1996) [75] and note that a considerable amount of control over properties of the functions $g_{n,\theta,\zeta}$ can be exercised through the choice of smoothing kernel ψ_n (e.g. of bounded Hölder norm: see examples 25.35–25.36 in [77], note that the σ -derivative of $\phi_\sigma(y)$ is bounded and conclude that the smoothing kernels ψ_n can be chosen such that the scores $g_{n, F - F_0}$ are bounded). These arguments make plausible (but do not prove) that a well-chosen sequence of kernels (ψ_n) smooths \mathcal{P}^* to \mathcal{P}_n 's that form a sequence of approximately least-favourable submodels. Since the posterior has already been shown to be consistent under perturbation, this claim implies that the ILAN expansion (30) holds. (The lack of a rigorous proof is one of two reasons why conjecture 7.8 is not a theorem.) \square

Conditions for integration of the LAE expansion are identical to those in the LAN case plus a requirement of one-sided contiguity. (In the LAN case, contiguity is implied by Le Cam's first lemma). For every $\eta \in D(\rho) = D(\theta_0, \rho)$, the sequence $(P_{\theta_n(h_n), \eta}^n)$ is required to be contiguous with respect to $(P_{\theta_0, \eta}^n)$. Lemma 5.9 below shows that such one-sided contiguity and domination as in (35) are closely related and both hold under a log-Lipschitz condition. Lemma 5.9 is a simple sufficiency statement that applies in the support boundary problem; various more general conditions for assertions (i) and (ii) of lemma 5.9 exist (see lemma 3.2 in Kleijn and Knapik (2013) [47]).

Lemma 5.9. *Suppose that there exists a constant m such that for all $\eta \in H$, all x and every θ in a neighbourhood of θ_0 ,*

$$|\log p_{\theta, \eta} - \log p_{\theta_0, \eta}|(x) 1_{A_{\theta_0}}(x) \leq m|\theta_0 - \theta|. \quad (39)$$

Then, for fixed $\rho > 0$ small enough,

(i) the model satisfies the domination condition

$$\sup_{\eta \in D(\rho)} P_{\theta_0, \eta}^n \left(\prod_{i=1}^n \frac{p_{\theta_n(h_n), \eta}(X_i)}{p_{\theta_0, \eta}} \right) = O(1),$$

(ii) and, for every $\eta \in D(\rho)$, $(P_{\theta_n(h_n), \eta}^n)$ is contiguous w.r.t. $(P_{\theta_0, \eta}^n)$.

We apply this lemma in the support boundary problem below.

Example (*Support boundary estimation, cont.*)

Since the space H consists of functions of bounded variation, Theorem V.2.2 in Ibragimov and Has'minskii (1981) [37] confirms that the model exhibits local asymptotic exponentiality in the θ -direction for every fixed η . In the notation of Definition 2.1, $\gamma_{\theta_0, \eta} = \eta(0)$, *i.e.* the size of the discontinuity at zero. According to Lemma 4.1 in [47], the map (15) is uniform-to-Hellinger continuous and the space H is a collection of probability densities that are (i) monotone decreasing with sub-exponential tails, (ii) continuously differentiable on $[0, \infty)$ and (iii) log-Lipschitz with constant $\alpha + S$. Hence (39) is satisfied with $m = \alpha + S$. We conclude that both the domination condition (35) (and the contiguity condition) are satisfied. Consistency under perturbation has been established in the previous section. According to Theorem 3.2 in [47] (the LAE analog of Theorem 5.7 above) the integral LAE-expansion holds, *i.e.*,

$$\begin{aligned} & \int_H \prod_{i=1}^n \frac{p_{\theta_n(h_n), \eta}(X_i)}{p_0} d\Pi_H(\eta) \\ &= \int_H \prod_{i=1}^n \frac{p_{\theta_0, \eta}(X_i)}{p_0} d\Pi_H(\eta) \exp(h_n \eta_0(0) + o_{P_0}(1)) 1_{\{h_n \leq \Delta_n\}}, \end{aligned}$$

for all stochastic (h_n) that are bounded in P_0 -probability. \square

6 Posterior asymptotic normality and exponentiality

In this section, it is shown that ILAN expansions of the form (30) induce asymptotic normality of marginal posteriors, *c.f.* (2), analogous to the way local asymptotic normality of parametric likelihoods induces the parametric Bernstein-Von Mises theorem. The underlying condition is marginal posterior consistency at rate $n^{-1/2}$ (which is also necessary for (2)). As it turns out, of all the conditions for the semiparametric Bernstein-von Mises limit, marginal consistency is the most stringent and hard to analyse in examples. The background of this issue is the possible occurrence of *semiparametric bias* (for an intriguing equivalence, see Klaassen (1987) [43] and relate to (18)).

6.1 Local limit shapes of marginal posteriors

The third major step in the proof of the semiparametric Bernstein-von Mises theorem is based on two observations: firstly, in a semiparametric problem the integrals S_n appear in the expression for the marginal posterior in exactly the same way as parametric likelihood ratios appear in the posterior for the parametric problem of theorem 2.10. Secondly, the parametric Bernstein-Von Mises proof depends on likelihood ratios *only* through the LAN property. As a consequence, local asymptotic normality for S_n offers the possibility to apply Le Cam and Yang's proof of posterior asymptotic normality in semiparametric context. We impose contraction at parametric rate for the marginal posterior to apply the LAN expansion of S_n and reach the conclusion that the marginal posterior satisfies the Bernstein-Von Mises assertion (2) (see theorem 6.1).

This shortcut is illustrated further by the following perspective. For given θ and n , $s_n(n^{1/2}(\theta - \theta_0))$ is a probability density for the stochastic vector (X_1, \dots, X_n) with respect to P_0^n , corresponding to the θ -conditioned (Π_H -prior predictive) distribution,

$$\tilde{P}_{n,\theta}(B) = P_0^n(1_B s_n(\sqrt{n}(\theta - \theta_0))),$$

(where B is measurable in the n -fold product of the samplespace). Indeed, keeping n fixed, we may view the map $\theta \mapsto \tilde{P}_{n,\theta}$ as a parametric model with a prior Π_Θ that is thick at θ_0 . Conditions then amount to stochastic local asymptotic normality and parametric posterior rate-optimality. This conceptual simplification comes at a price, though: firstly, this parametric model is misspecified, *i.e.* there is no $\theta \in \Theta$ such that $P_0^n = \tilde{P}_{n,\theta}$. Secondly, although we have assumed that the sample is distributed *i.i.d.*, in the parametric model above X_1, \dots, X_n are *not* independent, instead the sample (X_1, \dots, X_n) satisfies the weaker property of exchangeability under $\tilde{P}_{n,\theta}$ for every θ , in accordance with De Finetti's theorem. Although this enables application of methods put forth in Kleijn and van der Vaart [46], in the present context, results are sharper if we take into account the semiparametric background of the quantities $s_n(h)$.

Theorem 6.1. (Posterior asymptotic normality)

Let Θ be open in \mathbb{R}^k with a prior Π_Θ that is thick at θ_0 . Suppose that for large enough n , the map $h \mapsto s_n(h)$ is continuous P_0^n -almost-surely. Assume that there exists an $L_2(P_0)$ -function $\tilde{\ell}_{\theta_0, \eta_0}$ such that for every (h_n) bounded in probability, (30) holds, $P_0 \tilde{\ell}_{\theta_0, \eta_0} = 0$ and $\tilde{I}_{\theta_0, \eta_0}$ is non-singular. Furthermore suppose that for every (M_n) , $M_n \rightarrow \infty$,

$$\Pi_n(\|h\| \leq M_n \mid X_1, \dots, X_n) \xrightarrow{P_0} 1. \quad (40)$$

Then the sequence of marginal posteriors for θ converges to a normal distribution in total variation,

$$\sup_A \left| \Pi_n(h \in A \mid X_1, \dots, X_n) - N_{\tilde{\Delta}_n, \tilde{I}_{\theta_0, \eta_0}^{-1}}(A) \right| \xrightarrow{P_0} 0,$$

centred on $\tilde{\Delta}_n$ with covariance matrix $\tilde{I}_{\theta_0, \eta_0}^{-1}$.

Proof The proof is identical to that of theorem 2.1 in [46] upon replacement of parametric likelihoods with integrated likelihoods. \square

In the irregular LAE case a completely analogous statement can be made, leading to the assertion that the posterior is asymptotically exponential:

$$\sup_A \left| \Pi_n(h \in A \mid X_1, \dots, X_n) - \text{Exp}_{\tilde{\Delta}_n, \gamma_{\theta_0, \eta_0}}^-(A) \right| \xrightarrow{P_0} 0, \quad (41)$$

with local parameter h such that $\theta_n(h) = \theta_0 + n^{-1}h$. Here, the rate condition on the marginal posterior must enable the LAE expansion, *i.e.* it must imply one-sided, rate- n^{-1} consistency.

6.2 Marginal posterior consistency and semiparametric bias

Condition (40) in theorem 6.1 requires that the posterior measures of a sequence of model subsets of the form,

$$\Theta_n \times H = \{(\theta, \eta) \in \Theta \times H : \sqrt{n}\|\theta - \theta_0\| \leq M_n\}, \quad (42)$$

converge to one in P_0 -probability, for every sequence (M_n) such that $M_n \rightarrow \infty$. Essentially, this condition enables us to restrict the proof of theorem 6.1 to the shrinking domain in which (30) applies. Marginal posteriors in nonparametric models have not received much specific attention in the literature on posterior asymptotics thus far. Questions concerning testing in the presence of nuisance parameters (see [14] and many others) lie at the centre of this problem.

To fix a perspective to frame the question, consider the following lemma, which is a variation on lemma 6.1 of Bickel and Kleijn (2012) [5] and appears easier to satisfy in models that are everywhere smooth (see also condition (B3) of theorem 8.2 in Lehmann and Casella (1998) [59]).

Lemma 6.3. (Marginal parametric rate (I))

Given some P_0 , assume that the model possesses globally defined approximately least-favourable submodels $\Theta \mapsto \mathcal{P} : \theta \mapsto Q_{n,\theta,\zeta}$ for all ζ . Let the sequence of maps $\theta \mapsto S_n(\theta)$ be P_0 -almost-surely continuous and such that (30) is satisfied. Furthermore, assume that there exists a constant $C > 0$ such that for any (M_n) , $M_n \rightarrow \infty$,

$$P_0^n \left(\sup_{\zeta \in H} \sup_{\theta \in \Theta_n^c} \mathbb{P}_n \log \frac{q_{n,\theta,\zeta}}{q_{n,\theta_0,\zeta}} \leq -\frac{C M_n^2}{n} \right) \rightarrow 1. \quad (43)$$

Then, for any nuisance prior Π_H and parametric prior Π_Θ , thick at θ_0 ,

$$\Pi(n^{1/2}\|\theta - \theta_0\| > M_n \mid X_1, \dots, X_n) \xrightarrow{P_0} 0, \quad (44)$$

for any (M_n) , $M_n \rightarrow \infty$.

Proof Let (M_n) , $M_n \rightarrow \infty$ be given. Define (A_n) to be the events in (43) so that $P_0^n(A_n^c) = o(1)$ by assumption. In addition, let,

$$B_n = \left\{ \int_{\Theta} S_n(\theta) d\Pi_{\Theta}(\theta) \geq e^{-\frac{1}{2} C M_n^2} S_n(\theta_0) \right\}.$$

By (30) and lemma 6.3 in [5], $P_0^n(B_n^c) = o(1)$ as well. Then,

$$\begin{aligned} P_0^n \Pi(\theta \in \Theta_n^c \mid X_1, \dots, X_n) &\leq P_0^n \Pi(\theta \in \Theta_n^c \mid X_1, \dots, X_n) 1_{A_n \cap B_n} + o(1) \\ &\leq e^{\frac{1}{2} C M_n^2} P_0^n \left(S_n(\theta_0)^{-1} \right. \\ &\quad \left. \times \int_H \int_{\Theta_n^c} \prod_{i=1}^n \frac{q_{n,\theta,\zeta}}{q_{n,\theta_0,\zeta}}(X_i) \prod_{i=1}^n \frac{q_{n,\theta_0,\zeta}}{q_{n,\theta_0,\zeta_0}}(X_i) d\Pi_{\Theta} d\Pi_H 1_{A_n} \right) + o(1) = o(1), \end{aligned}$$

which proves (44). □

Essentially the proof of lemma 6.3 revolves around suppressing the diverging inverse prior masses of $\Theta_n \times H$ by the uniform bound on likelihood ratios implied by (43). The condition says that the *semiparametric likelihood ratio statistic* associated with the marginal estimation of θ must have testing power. Requirements of this type also play a prominent role in frequentist semiparametric theory, typically to assure that the centred and rescaled limit-distribution of the estimator is tight.

For example, in general formulations of *profile likelihood* methods (Severini and Wong (1992) [67], Murphy and van der Vaart (2000) [62]), conditions are formulated to exclude the possibility that θ -estimators (which arises as ML estimators in models for $\theta|\eta$ in which η has been replaced by a likelihood-maximizer $\hat{\eta}_n$, compare with (43) above)) develops what is known as *semiparametric bias*: local variations of the nuisance parameter distort the model for $\theta|\eta$ to such an extent, that the “plug-in” $\eta = \hat{\eta}_n$ does not give rise to a *tight limit law* for $n^{1/2}(\hat{\theta}_n - \theta_0)|\hat{\eta}_n$. To exclude this possibility, Murphy and van der Vaart (2000) [62] (see also [77]) impose so-called *no-bias conditions* of the the following form: for the maximizers $\hat{\eta}_n, \hat{\theta}_n$ and all $\theta \in \Theta, \eta \in H$,

$$\begin{aligned} P_{\hat{\theta}_n, \eta} \tilde{\ell}_{\hat{\theta}_n, \hat{\eta}_n} &= o(n^{-1/2} + \|\hat{\theta}_n - \theta_0\|), \\ P_{\theta, \eta} \|\tilde{\ell}_{\hat{\theta}_n, \hat{\eta}_n} - \tilde{\ell}_{\theta, \eta}\|^2 &= o_P(1), \quad P_{\hat{\theta}_n, \eta} \|\tilde{\ell}_{\hat{\theta}_n, \hat{\eta}_n}\|^2 = O_P(1). \end{aligned} \tag{45}$$

The requirements that are second order in scores control the local behaviour of Fisher information coefficients and play a (dominating) role comparable to that of (35). The essential condition is the first one, linear in the efficient score: if, when varying the nuisance η , the expectation of the efficient score cannot be controlled to be (roughly) $o(n^{-1/2})$ uniformly, profile ML estimators $\hat{\theta}_n$ tend to drift off with a bias of order $O(n^{-1/2})$ or worse: according to theorem 25.59 in [77], under certain general conditions, solutions $\hat{\theta}_n$ to efficient score equations satisfy,

$$n^{1/2}(\hat{\theta}_n - \theta_0) = \tilde{\Delta}_n + P_{\hat{\theta}_n, \eta} \tilde{\ell}_{\hat{\theta}_n, \hat{\eta}_n} + o_{P_0}(1), \tag{46}$$

(compare with (18)). Conditions (45) also determine bias in the score tests associated with the (likelihood ratio) tests of (43). In quite some generality [62, 77, 4] one can say that a (sufficient) no-bias condition is that the “plug-in” $\hat{\eta}_n$ is consistent and that the family \mathcal{F}_0 of score functions $\tilde{\ell}_{\theta_0, \eta}$ (with η in neighbourhoods of η_0) forms a P_0 -Donsker class.

Of course, lemma 6.3 formulates mere sufficient conditions, so one could suspect that bias issues occur as a result of our chosen (ML) methods rather than being intrinsic to the problem. However, the following straightforward lemma shows that inconvenient uniformities are also part of a strictly Bayesian approach.

Lemma 6.5. (Marginal parametric rate (II))

Let Π_Θ and Π_H be given. Assume that there exists a sequence (H_n) of subsets of H , such that the following two conditions hold:

- (i) The nuisance posterior concentrates on H_n asymptotically,

$$\Pi(\eta \in H \setminus H_n \mid X_1, \dots, X_n) \xrightarrow{P_0} 0. \tag{47}$$

(ii) For every (M_n) , $M_n \rightarrow \infty$,

$$\sup_{\eta \in H_n} \Pi(n^{1/2}\|\theta - \theta_0\| > M_n \mid \eta, X_1, \dots, X_n) \xrightarrow{P_0} 0. \quad (48)$$

Then the marginal posterior for θ concentrates at parametric rate, i.e.,

$$\Pi(n^{1/2}\|\theta - \theta_0\| > M_n \mid X_1, \dots, X_n) \xrightarrow{P_0} 0,$$

for every sequence (M_n) , $M_n \rightarrow \infty$,

Proof Let (M_n) , $M_n \rightarrow \infty$ be given and consider the posterior for the complement of (42). By assumption (i) and Fubini's theorem,

$$\begin{aligned} & P_0^n \Pi(\theta \in \Theta_n^c \mid X_1, \dots, X_n) \\ & \leq P_0^n \int_{H_n} \Pi(\theta \in \Theta_n^c \mid \eta, X_1, \dots, X_n) d\Pi(\eta \mid X_1, \dots, X_n) + o(1) \\ & \leq P_0^n \sup_{\eta \in H_n} \Pi(n^{1/2}\|\theta - \theta_0\| > M_n \mid \eta, X_1, \dots, X_n) + o(1), \end{aligned}$$

the first term of which is $o(1)$ by assumption (ii). \square

Note that condition (48) requires *uniform* posterior convergence to θ_0 at rate $n^{-1/2}$ under P_0 , in all misspecified parametric models $\mathcal{P}_\eta = \{P_{\theta, \eta} : \theta \in \Theta\}$ for $\eta \in H_n$. From this perspective, it is clear how semiparametric bias manifests itself in Bayesian context: according to Kleijn and van der Vaart [46] and Kleijn (2003) [44], the misspecified posterior on \mathcal{P}_η concentrates at parametric rate *around the minimizer* $\theta^*(\eta)$ of $\theta \mapsto -P_0 \log(p_{\theta, \eta}/p_0)$ rather than around θ_0 . So one hopes for the eventuality that marginal posteriors for η concentrate on subsets H_n such that the corresponding KL-minimizers θ^* satisfy,

$$\sup_{\eta \in H_n} \|\theta^*(\eta) - \theta_0\| = O(n^{-1/2}), \quad (49)$$

in order for posterior concentration to occur on the strips (42). (Recall that marginal $n^{-1/2}$ -consistency is *necessary* for (2).)

The difference between bias terms of orders $o(n^{-1/2})$ or $O(n^{-1/2})$ is the absence or presence of bias in the limit experiment [54]. More precisely, a bias of order $O(n^{-1/2})$ does not ruin marginal posterior consistency at $n^{-1/2}$ -rate but biases the centring sequence $\tilde{\Delta}_n$ of (3), in congruence with (46) and theorem 25.59 in [77]. In theorem 6.1, biased centring cannot occur because the form of $\tilde{\Delta}_n$ follows exclusively from the ILAN expansion of the integrated likelihood. Nevertheless this form of semiparametric bias occurs with surprisingly high frequency in semiparametric Bernstein-von Mises analyses [64, 9]: certain priors (Gaussian in [64, 9], but the problem cannot be expected to be not limited to this class) combine with the likelihood and distort LAN expansions by undesirable order- $O(n^{-1/2})$ bias of the form $P_{\theta_0, \eta} \tilde{\ell}_{\theta_0, \eta_0}$. The interested reader is referred to Castillo (2011) [9].

To conclude we mention a lemma that proves the intuitively reasonable assertion that convergence at rate $n^{-1/2}$ of the posterior measures for a sequence of (misspecified) parametric

submodels to their individual Kullback-Leibler minima implies their convergence to the true value of the parameter, if the sequence of minima itself converges at rate $1/\sqrt{n}$. The sequence of submodels may be chosen stochastically, for example, we may “model-select” from $\{\mathcal{P}_\eta : \eta \in H\}$ by means of a point-estimator $\hat{\eta}_n$.

Lemma 6.7. *Let \mathcal{P}_n be a (possibly stochastic) sequence of parametric models. Assume that the sequence of Kullback-Leibler minima θ_n^* satisfies:*

$$\sqrt{n}(\theta_n^* - \theta_0) = O_{P_0}(1). \quad (50)$$

Furthermore, assume that for each of the (misspecified) models \mathcal{P}_n , the posterior concentrates around θ_n^ at rate $n^{-1/2}$ in P_0 -expectation. Then, for every sequence M_n such that $M_n \rightarrow \infty$*

$$\Pi_n(\sqrt{n}\|\theta - \theta_0\| > M_n | X_1, \dots, X_n) \rightarrow 0,$$

in P_0 -probability.

Proof See lemma 4.18 in Kleijn (2003) [44]. □

To resolve semiparametric bias, alternative point-estimation methods (*e.g.* regularized likelihood maximization by inclusion of suitable penalties, or, replacement of score equations by general estimating equations) are applied. This suggests that to prevent semiparametric bias in Bayesian context, one should expect either the model to be small enough for \mathcal{F}_0 to satisfy the Donsker property, or, the prior to be concentrated (enough) on submodels for which \mathcal{F}_0 is Donsker.

6.3 Application in examples

In the three examples below (including the irregular boundary support problem), the model is kept such that \mathcal{F}_0 is a Donsker class.

Example (*Partial linear regression, cont.*)

Concerning marginal consistency at parametric rate in the partial linear model, let (M_n) , $M_n \rightarrow \infty$ be given and define Θ_n as in (42). Here, $\mathcal{P}_n = \mathcal{P}^*$ for all $n \geq 1$, all ζ and all $\theta \in \Theta_n^c$,

$$\begin{aligned} \mathbb{P}_n \log \frac{q_{n,\theta,\zeta}}{q_{n,\theta_0,\zeta}} &= (\theta - \theta_0) \sum_{i=1}^n g_\zeta(X_i) \\ &\quad - \frac{n}{2}(\theta - \theta_0)^2 P_{\theta_0, \eta_0 + \zeta}(g_\zeta)^2 - \frac{n}{2}(\theta - \theta_0)^2 (\mathbb{P}_n - P_{\theta_0, \eta_0 + \zeta})(g_\zeta)^2. \end{aligned}$$

Because $P_{\theta_0, \eta_0 + \zeta}(g_\zeta)^2 = \tilde{I}_{\theta_0, \eta_0}$ for all ζ and $PU^4 < \infty$, the last term converges to zero P_0 -almost-surely, uniform in ζ . Note that for all ζ , the no-bias condition is satisfied *exactly*,

$$P_{\theta_0, \eta_0 + \zeta} g_\zeta = P_{\theta_0, \eta_0 + \zeta}(e(U - E[U|V])) = 0,$$

so that $n^{1/2}\mathbb{P}_n g_\zeta = \mathbb{G}_n g_\zeta$ under all $P_{\theta_0, \eta_0 + \zeta}$. We conclude that,

$$\sup_{\theta \in \Theta_n^c} \sup_{\zeta \in H - \eta_0} \mathbb{P}_n \log \frac{q_{n,\theta,\zeta}}{q_{n,\theta_0,\zeta}} \leq M_n \sup_{\zeta \in H - \eta_0} |\mathbb{G}_n g_\zeta| - \frac{1}{2} M_n^2 \tilde{I}_{\theta_0, \eta_0} + o_{P_0}(1).$$

Assume that H is a P_0 -Donsker class and that the efficient Fisher information is non-singular at P_0 . The Donsker property guarantees asymptotic tightness of $\sup_{\zeta \in H - \eta_0} |\mathbb{G}_n g_\zeta|$ so that lemma 6.3 holds and (40) is valid. We note that $h \mapsto s_n(h)$ is continuous for every $n \geq 1$ (see (38)) and ILAN according to the previous section. Applying theorem 6.1, we conclude that (2) holds. \square

Example (*Normal location mixtures, cont.*)

Based on examples 25.35, 25.36 and 25.61 in [77], we see that the bias term in the first condition in (45) vanishes *exactly* and *globally* in the normal location mixture model:

$$P_{\sigma_0, F} \tilde{\ell}_{\sigma_0, F_0} = 0,$$

for all $F \in \mathcal{D}[0, 1]$. Referring to the previous example, this suggests strongly that semiparametric bias does not play a role and problems regarding marginal posterior convergence at parametric rate are not anticipated. Indeed, preliminary calculations indicate that for all $\sigma \in \Theta_n^c$, the empirical KL divergences $-\mathbb{P}_n \log(q_{n, \sigma, F - F_0} / q_{n, \sigma_0, F - F_0})$ stay above $D(\sigma - \sigma_0)^2$ (for some $D > 0$) up to $o_{P_0}(1)$ uniformly in $F \in \mathcal{D}[0, 1]$, so that (43) would be satisfied. (The lack of a rigorous proof is the second reason why conjecture 7.8 is not a theorem.) We note that $h \mapsto s_n(h)$ is continuous for every $n \geq 1$ and ILAN according to the previous section. Applying theorem 6.1, we conclude that (2) holds. \square

Example (*Support boundary estimation, cont.*)

Integral LAE was verified in the previous section and continuity of $h \mapsto s_n(h)$ (on $(-\infty, \Delta_n]$) is implied by the integral LAE expansion. In the irregular case, marginal consistency at rate n^{-1} follows from lemma 3.3 in [47], which is completely analogous to lemma 6.1 in [5]. To show that the condition is satisfied, note that for fixed x and η , the map $\theta \mapsto p_{\theta, \eta}(x)$ is monotone increasing. Therefore

$$\sup_{\theta \in \Theta_n^c} \frac{1}{n} \log \prod_{i=1}^n \frac{p_{\theta, \eta}(X_i)}{p_{\theta_0, \eta}(X_i)} \leq \frac{1}{n} \log \prod_{i=1}^n \frac{\eta(X_i - \theta^*)}{\eta(X_i - \theta_0)} 1_{\{X_{(1)} \geq \theta^*\}}(\underline{X}_n),$$

where $\theta^* = X_{(1)}$ if $X_{(1)} \geq \theta_0 + M_n/n$, or $\theta_0 - M_n/n$ otherwise. We first note that $X_{(1)} < \theta_0 + M_n/n$ with probability tending to one. Indeed, shifting the distribution to $\theta = 0$, we calculate,

$$P_{0, \eta_0}^n \left(X_{(1)} \geq \frac{M_n}{n} \right) = \left(1 - \int_0^{\frac{M_n}{n}} \eta_0(x) dx \right)^n \leq \exp \left(-n \int_0^{\frac{M_n}{n}} \eta_0(x) dx \right).$$

By lemma 5.1 in [47], the right-hand side of the above display is bounded further as follows,

$$\exp \left(-\gamma_{\theta_0, \eta_0} M_n + M_n \int_0^{\frac{M_n}{n}} |\eta_0'(x)| dx \right) \leq \exp \left(-\frac{\gamma_{\theta_0, \eta_0}}{2} M_n \right),$$

for large enough n . We continue with $\theta^* = \theta_0 - M_n/n$. By absolute continuity of η we have

$$\eta(X_i - \theta^*) = \eta(X_i - \theta_0) + \int_{X_i - \theta_0}^{X_i - \theta^*} \eta'(y) dy,$$

and the conditions on the nuisance η yield the following bound,

$$\int_{X_i - \theta_0}^{X_i - \theta^*} \eta'(y) dy \leq (\theta_0 - \theta^*)(S - \alpha)\eta(X_i - \theta_0).$$

Therefore

$$\frac{1}{n} \log \prod_{i=1}^n \frac{\eta(X_i - \theta^*)}{\eta(X_i - \theta_0)} 1_{\{X_{(1)} \geq \theta^*\}}(\underline{X}_n) \leq \frac{1}{n} \log \left(1 - \frac{(\alpha - S)M_n}{n} \right)^n \leq -\frac{(\alpha - S)M_n}{n}.$$

With $C < \alpha - S$, the condition of lemma 3.3 in [47] is satisfied. We conclude that (41) holds. \square

7 Main results

Before we state the main results of this paper, general conditions imposed on models and priors are formulated.

(i) Model assumptions

The model \mathcal{P} is assumed to be well-specified and dominated by a σ -finite measure on the samplespace and parametrized identifiably on $\Theta \times H$, with $\Theta \subset \mathbb{R}^k$ open and H a subset of a metric vector-space with metric d_H . It is assumed that $(\theta, \eta) \mapsto P_{\theta, \eta}$ is continuous. We also assume that there exists an open neighbourhood $U_0 \subset \Theta$ of θ_0 on which approximately least-favourable submodels $\eta_n : U_0 \rightarrow H$ are defined.

(ii) Prior assumptions

For the prior Π on (the Borel σ -algebra of) \mathcal{P} we endow $\Theta \times H$ with a Borel product-prior $\Pi_\Theta \times \Pi_H$. Also it is assumed that the prior Π_Θ is *thick* (that is, Lebesgue absolutely continuous with continuous and strictly positive density).

7.1 Main theorems

With the above general considerations for model and prior in mind, we formulate the main result of this paper.

Theorem 7.1. (Semiparametric Bernstein-Von Mises)

Let X_1, X_2, \dots be distributed i.i.d.- P_0 , with $P_0 \in \mathcal{P}$ and let Π_Θ be thick at θ_0 . Suppose that for large enough n , $h \mapsto s_n(h)$ is continuous P_0^n -almost-surely. Also assume that the $\theta \mapsto Q_{n, \theta, \zeta}$ are stochastically LAN at θ_0 in the θ -direction for all ζ in a neighbourhood of $\zeta = 0$ and that the efficient Fisher information $\tilde{I}_{\theta_0, \eta_0}$ is non-singular. Furthermore, assume that there exists a sequence (ρ_n) with $\rho_n \downarrow 0$, $n\rho_n^2 \rightarrow \infty$ such that (35) holds and:

(i) For all $M > 0$, there exists a $K > 0$ such that, for large enough n ,

$$\Pi_H(K_n(\rho_n, M)) \geq e^{-Kn\rho_n^2}.$$

(ii) For all n large enough, the Hellinger metric entropy satisfies,

$$N(\rho_n, H, d_H) \leq e^{n\rho_n^2}.$$

(iii) For every bounded, stochastic (h_n) and all $L > 0$, Hellinger distances satisfy the uniform bound,

$$\sup_{\{\eta \in H: d_H(\eta, \eta_0) \geq L\rho_n\}} \frac{H(P_{\theta_n(h_n), \eta}, P_{\theta_0, \eta})}{H(P_{\theta_0, \eta}, P_0)} = o(1).$$

(iv) For every (M_n) , $M_n \rightarrow \infty$, the posterior satisfies,

$$\Pi_n(\|h\| \leq M_n \mid X_1, \dots, X_n) \xrightarrow{P_0} 1.$$

Then the sequence of marginal posteriors for θ converges in total variation to a normal distribution,

$$\sup_A \left| \Pi_n(h \in A \mid X_1, \dots, X_n) - N_{\tilde{\Delta}_n, \tilde{I}_{\theta_0, \eta_0}^{-1}}(A) \right| \xrightarrow{P_0} 0, \quad (51)$$

centred on $\tilde{\Delta}_n$ with covariance matrix $\tilde{I}_{\theta_0, \eta_0}^{-1}$.

Proof The assertion follows from combination of theorem 4.2, corollary 4.3, theorem 5.7 and theorem 6.1. \square

Conditions (i) and (ii) also arise when considering Hellinger rates for nonparametric posterior convergence, and the methods of Ghosal *et al.* (2000) [31] can be applied in the present context with minor modifications. Typically, the numerator in condition (iii) is of order $O(n^{-1/2})$ so it is easily satisfied for nonparametric rates (ρ_n) . Condition (iv) of theorem 7.1 is the most significant one: note, first of all, that (iv) is necessary. Subsection 6.2 shows that formulation of straightforward sufficient conditions is hard in generality. Condition (iv) involves the nuisance prior and, as such, poses a condition for the nuisance prior. In the examples, the ‘hard work’ stems from condition (iv): for example, $\alpha > 1/2$ Hölder smoothness and boundedness of the family of regression functions in corollary 7.6 are imposed in order to satisfy this condition. Since conditions (i) and (ii) appear quite reasonable and conditions (35) and (iii) are satisfied relatively easily, condition (iv) should be viewed as the most complicated in an essential way.

Consider the rate (ρ_n) : on the one hand, (ρ_n) must converge to zero fast enough to satisfy the second-order approximation condition (35), on the other hand, (ρ_n) is fixed at or above the minimax Hellinger rate for estimation of the nuisance (with known θ_0) by condition (ii) and must converge to zero slowly enough to satisfy conditions (i) and (iii). Lemma 5.3 shows that in many semiparametric models approximately least-favourable reparametrizations exist that satisfy (35) for *any* (ρ_n) . In that case, conditions (i), (ii) and (iii) can be weakened and the rate (ρ_n) does not need to be mentioned explicitly in the formulation of the theorem. This enables a rate-free corollary in which conditions (i) and (ii) above are weakened to become comparable to those of Schwartz (1965) [66] for nonparametric posterior consistency, rather than those for posterior rates of convergence following Ghosal, Ghosh and van der Vaart [31].

Corollary 7.3. (Semiparametric Bernstein-Von Mises, rate-free)

Let X_1, X_2, \dots be distributed i.i.d.- P_0 , with $P_0 \in \mathcal{P}$ and let Π_Θ be thick at θ_0 . Suppose that for large enough n , $h \mapsto s_n(h)$ is continuous P_0^n -almost-surely. Assume that the $\theta \mapsto Q_{n, \theta, \zeta}$ are stochastically LAN at θ_0 in the θ -direction, for all ζ in a neighbourhood of $\zeta = 0$ and that

the efficient Fisher information $\tilde{I}_{\theta_0, \eta_0}$ is non-singular. Also assume that there exists a $\rho > 0$ such that (35) holds with $\rho_n = \rho$. If,

(i) for all $\rho > 0$, the Hellinger metric entropy satisfies, $N(\rho, H, d_H) < \infty$ and the nuisance prior satisfies $\Pi_H(K(\rho)) > 0$,

(ii) for every $M > 0$, there exists an $L > 0$ such that for all $\rho > 0$ and large enough n , $K(\rho) \subset K_n(L\rho, M)$,

(iii) Hellinger distances satisfy, $\sup_{\eta \in H} H(P_{\theta_n(h_n), \eta}, P_{\theta_0, \eta}) = O(n^{-1/2})$,

(iv) for every $M_n \rightarrow \infty$, we have $\Pi_n(\|h\| \leq M_n \mid X_1, \dots, X_n) \xrightarrow{P_0} 1$,

then the sequence of marginal posteriors for θ converges in total variation to a normal distribution,

$$\sup_A \left| \Pi_n(h \in A \mid X_1, \dots, X_n) - N_{\tilde{\Delta}_n, \tilde{I}_{\theta_0, \eta_0}^{-1}}(A) \right| \xrightarrow{P_0} 0,$$

centred on $\tilde{\Delta}_n$ with covariance matrix $\tilde{I}_{\theta_0, \eta_0}^{-1}$.

Proof Due to the fact that (35) holds for *any* rate, under conditions (i), (ii), (iii) and the stochastic LAN assumption, the assertion of corollary 4.3 holds. Condition (iv) then suffices for the assertion of theorem 6.1. \square

7.2 Partial linear regression

For the following theorem we think of the regression function and the process (14) as elements of the Banach space $(C[0, 1], \|\cdot\|_\infty)$. In the corollary that follows, we relate to Banach subspaces with stronger norms to complete the argument.

Theorem 7.5. *Let X_1, X_2, \dots be an i.i.d. sample from the partial linear model (13) with $P_0 = P_{\theta_0, \eta_0}$ for some $\theta_0 \in \Theta$, $\eta_0 \in H$. Assume that H is a subset of $C[0, 1]$ of finite metric entropy with respect to the uniform norm and that H forms a P_0 -Donsker class. Regarding the distribution of (U, V) , suppose that $PU = 0$, $PU^2 = 1$ and $PU^4 < \infty$, as well as $P(U - E[U|V])^2 > 0$, $P(U - E[U|V])^4 < \infty$ and $v \mapsto E[U|V = v] \in H$. Endow Θ with a prior that is thick at θ_0 and $C[0, 1]$ with a prior Π_H such that $H \subset \text{supp}(\Pi_H)$. Then the marginal posterior for θ satisfies the Bernstein-Von Mises limit,*

$$\sup_{B \in \mathcal{B}} \left| \Pi(\sqrt{n}(\theta - \theta_0) \in B \mid X_1, \dots, X_n) - N_{\tilde{\Delta}_n, \tilde{I}_{\theta_0, f_0}^{-1}}(B) \right| \xrightarrow{P_0} 0, \quad (52)$$

where $\tilde{\ell}_{\theta_0, \eta_0}(X) = e(U - E[U|V])$ and $\tilde{I}_{\theta_0, \eta_0} = P(U - E[U|V])^2$.

In the following we choose a prior by picking a suitable k in (14) and conditioning on $\|\eta\|_\alpha < M$. The resulting prior (denoted $\Pi_{\alpha, M}^k$) is shown to be well-defined and have full support.

Corollary 7.6. *Let $\alpha > 1/2$ and $M > 0$ be given; choose $H = \{\eta \in C^\alpha[0, 1] : \|\eta\|_\alpha < M\}$ and assume that $\eta_0 \in C^\alpha[0, 1]$. Suppose the distribution of the covariates (U, V) is as in theorem 7.5. Then, for any integer $k > \alpha - 1/2$, the conditioned prior $\Pi_{\alpha, M}^k$ is well-defined and gives rise to a marginal posterior for θ satisfying (52).*

Proof Choose k as indicated; the Gaussian distribution of η over $C[0, 1]$ is based on the RKHS $H^{k+1}[0, 1]$ and denoted Π^k . Since η in (14) has smoothness $k + 1/2 > \alpha$, $\Pi^k(\eta \in C^\alpha[0, 1]) = 1$. Hence, one may also view η as a Gaussian element in the Hölder class $C^\alpha[0, 1]$, which forms a separable Banach space even with strengthened norm $\|\cdot\| = \|\eta\|_\infty + \|\cdot\|_\alpha$, without changing the RKHS. The trivial embedding of $C^\alpha[0, 1]$ into $C[0, 1]$ is one-to-one and continuous, enabling identification of the prior induced by η on $C^\alpha[0, 1]$ with the prior Π^k on $C[0, 1]$. Given $\eta_0 \in C^\alpha[0, 1]$ and a sufficiently smooth kernel ϕ_σ with bandwidth $\sigma > 0$, consider $\phi_\sigma * \eta_0 \in H^{k+1}[0, 1]$. Since $\|\eta_0 - \phi_\sigma * \eta_0\|_\infty$ is of order σ^α and a similar bound exists for the α -norm of the difference [78], η_0 lies in the closure of the RKHS both with respect to $\|\cdot\|_\infty$ and to $\|\cdot\|$. Particularly, η_0 lies in the support of Π^k , in $C^\alpha[0, 1]$ with norm $\|\cdot\|$. Hence, $\|\cdot\|$ -balls centred on η_0 receive non-zero prior mass, *i.e.* $\Pi^k(\|\eta - \eta_0\| < \rho) > 0$ for all $\rho > 0$. Therefore, $\Pi^k(\|\eta - \eta_0\|_\infty < \rho, \|\eta\|_\alpha < \|\eta_0\|_\alpha + \rho) > 0$, which guarantees that $\Pi^k(\|\eta - \eta_0\|_\infty < \rho, \|\eta\|_\alpha < M) > 0$, for small enough $\rho > 0$. This implies that $\Pi^k(\|\eta\|_\alpha < M) > 0$ and,

$$\Pi_{\alpha, M}^k(B) = \Pi^k(B \mid \|\eta\|_\alpha < M),$$

is well-defined for all Borel-measurable $B \subset C[0, 1]$. Moreover, it follows that $\Pi_{\alpha, M}^k(\|\eta - \eta_0\|_\infty < \rho) > 0$ for all $\rho > 0$. We conclude that k times integrated Brownian motion started at random, conditioned to be bounded by M in α -norm, gives rise to a prior that satisfies $\text{supp}(\Pi_{\alpha, M}^k) = H$. As is well-known [76], the entropy numbers of H with respect to the uniform norm satisfy, for every $\rho > 0$, $N(\rho, H, \|\cdot\|_\infty) \leq K\rho^{-1/\alpha}$, for some constant $K > 0$ that depends only on α and M . The associated bound on the bracketing entropy gives rise to finite bracketing integrals, so that H is universally Donsker. Then, if the distribution of the covariates (U, V) is as assumed in theorem 7.5, the Bernstein-Von Mises limit (52) holds. \square

Comparing the above result with sufficient conditions from the frequentist literature on this model, one notices that the restriction $\alpha > 1/2$ is in line with earlier analyses but *boundedness* of the α -norm is more restrictive than expected. However, there are good reasons to suspect that the restriction on the regression class can be avoided here as well. To see this, note that the Bernstein-Von Mises limit (52) holds for any value of the constant $M > 0$ that lies above the α -norm of η_0 , as in corollary 7.6. Therefore there exists a sequence (M_n) , $M_n \rightarrow \infty$, such that the corresponding sequence of priors (Π_{α, M_n}^k) gives rise to marginal posteriors for the parameter θ that still satisfy,

$$\sup_{B \in \mathcal{B}} \left| \Pi_{\alpha, M_n}^k(\sqrt{n}(\theta - \theta_0) \in B \mid X_1, \dots, X_n) - N_{\tilde{\Delta}_n, \tilde{I}_{\theta_0, f_0}^{-1}}(B) \right| \xrightarrow{P_0} 0.$$

Then, one constructs an infinite convex combination of the priors (Π_{α, M_n}^k) to obtain a prior that does not depend on the bound M any longer. However, since we do not know in advance which sequences of bounds (M_n) diverge slowly enough to maintain Bernstein-Von Mises convergence, this proposal does not possess great practical advantage.

7.3 Normal location mixtures

Based on the discussion of the problem of variance estimation in normal location mixtures as presented above, we conjecture the following.

Conjecture 7.8. *Let X_1, X_2, \dots be an i.i.d. sample from $P_0 = P_{\sigma_0, F_0}$ in the semiparametric normal location mixture model parametrized by (11). Let $\Sigma = [\sigma_-, \sigma_+] \subset (0, \infty)$ have a thick prior and $\mathcal{D}[0, 1]$ a Dirichlet prior D_α with finite base measure α that dominates the Lebesgue measure on $[0, 1]$. Assume that the efficient Fisher information at P_0 is non-singular. Then the marginal posterior for the kernel variance σ has a Bernstein-von Mises limit of the form:*

$$\sup_{B \in \mathcal{B}(\Sigma)} \left| \Pi(\sqrt{n}(\sigma - \sigma_0) \in B \mid X_1, \dots, X_n) - N_{\tilde{\Delta}_n, \tilde{I}_{\sigma_0, F_0}^{-1}}(B) \right| \xrightarrow{P_0} 0.$$

In this case we do not have a closed-form expression for the efficient score function and, as a result, it could prove difficult to point-estimate the efficient Fisher information at θ_0 numerically. As a result, computationally, there is no easy way to find confidence ellipsoids of the form (4). This circumstance clearly demonstrates the practical value of the semiparametric Bernstein-von Mises theorem: when simulating the marginal posterior for σ (with methods very similar to those discussed in Escobar and West (1995) [24]), we do not need estimates for the efficient Fisher information to approximate credible sets, and hence, confidence sets.

7.4 Support boundary estimation

Locally asymptotically exponential semiparametric problems are covered by theorem 2.2 in Kleijn and Knapik (2013) [47]: besides requiring LAE instead of LAN smoothness, conditions are identical to those of theorem 7.1 with one addition: we require one-sided contiguity (condition (iv) of theorem 2.2 in [47]) which is implicit in LAN context.

Based on the results of previous sections regarding the support boundary problem, we now have the following irregular Bernstein-von Mises limit for the marginal posterior.

Theorem 7.9. *Let X_1, X_2, \dots be an i.i.d. sample from the location model of definition (15) with $P_0 = P_{\theta_0, \eta_0}$ for some $\theta_0 \in \Theta$, $\eta_0 \in H$. Endow Θ with a prior that is thick at θ_0 and \mathcal{L} with the prior $\Pi_{\mathcal{L}}$ of definition (16) (or any other prior such that $\mathcal{L} \subset \text{supp}(\Pi_{\mathcal{L}})$). Then the marginal posterior for θ satisfies,*

$$\sup_A \left| \Pi(n(\theta - \theta_0) \in A \mid X_1, \dots, X_n) - \text{Exp}_{\Delta_n, \gamma_{\theta_0, \eta_0}}^{-}(A) \right| \xrightarrow{P_0} 0, \quad (53)$$

where Δ_n is exponentially distributed with scale $\gamma_{\theta_0, \eta_0} = \eta_0(0)$.

In an example concerning a scaling parameter for which the model satisfies the LAE property, a similar result is derived in Kleijn and Knapik (2013) [47].

8 Proofs

In this section we collect several proofs from earlier sections. Following definition (21), Hellinger neighbourhoods of \mathcal{P}_n are given by

$$D_n(\theta, \rho) = \{ \eta \in H : d_H(\eta, \eta_n(\theta)) < \rho \}, \quad (54)$$

for all $\theta \in U_0$.

Proof of lemma 5.3 Let (h_n) be stochastic and upper-bounded by $M > 0$. For every ζ and all $n \geq 1$,

$$\begin{aligned} & Q_{n,\theta_0,\zeta}^n \left| \prod_{i=1}^n \frac{q_{n,\theta_n(h_n),\zeta}(X_i)}{q_{n,\theta_0,\zeta}} - 1 \right| \\ &= Q_{n,\theta_0,\zeta}^n \left| \int_{\theta_0}^{\theta_n(h_n)} \sum_{i=1}^n g_{n,\theta',\zeta}(X_i) \prod_{j=1}^n \frac{q_{n,\theta',\zeta}(X_j)}{q_{n,\theta_0,\zeta}} d\theta' \right| \\ &\leq \int_{\theta_0 - \frac{M}{\sqrt{n}}}^{\theta_0 + \frac{M}{\sqrt{n}}} Q_{n,\theta',\zeta}^n \left| \sum_{i=1}^n g_{n,\theta',\zeta}(X_i) \right| d\theta' \\ &\leq \sqrt{n} \int_{\theta_0 - \frac{M}{\sqrt{n}}}^{\theta_0 + \frac{M}{\sqrt{n}}} \sqrt{Q_{n,\theta',\zeta}(g_{n,\theta',\zeta})^2} d\theta', \end{aligned}$$

with use of the Cauchy-Schwartz inequality. For large enough n , $\rho_n < \rho$ and the square-root of (36) dominates the difference between $U(\rho, h_n)$ and 1. \square

Proof of lemma 5.5 Let (h_n) and (ρ_n) be given. Denote $\theta_n = \theta_n(h_n)$, $E_n = D(\theta_n, \rho_n)$ and $F_n = D(\theta_0, \rho_n)$ for all $n \geq 1$. Since,

$$\left| \Pi_H(E_n) - \Pi_H(F_n) \right| \leq \Pi_H((E_n \cup F_n) \setminus (E_n \cap F_n)),$$

we consider the sequence of symmetric differences. Note that, since the submodels \mathcal{P}_n are sLAN, we have $d_H(\eta_n(\theta_n(h_n)), \eta_0) = O(n^{-1/2})$. Fix some $0 < \alpha < 1$; for all $\eta \in E_n$,

$$d_H(\eta, \eta_0) \leq d_H(\eta, \eta^*(\theta_n)) + d_H(\eta^*(\theta_n), \eta_n(\theta_n)) + d_H(\eta_n(\theta_n), \eta_0),$$

which is dominated by $(1 + \alpha)\rho_n$ for large enough n , in accordance with (34). As a result, $E_n \cup F_n \subset D(\theta_0, (1 + \alpha)\rho_n)$. Furthermore, for any $\eta \in D(\theta_0, (1 - \alpha)\rho_n)$,

$$\begin{aligned} d_H(\eta, \eta^*(\theta_n)) &\leq d_H(\eta, \eta_0) + d_H(\eta_0, \eta_n(\theta_n)) + d_H(\eta_n(\theta_n), \eta^*(\theta_n)) \\ &\leq (1 - \alpha)\rho_n + o(\rho_n), \end{aligned}$$

so that $D(\theta_0, (1 - \alpha)\rho_n) \subset E_n \cap F_n$ for large enough n . Therefore,

$$(E_n \cup F_n) \setminus (E_n \cap F_n) \subset D(\theta_0, (1 + \alpha)\rho_n) \setminus D(\theta_0, (1 - \alpha)\rho_n) \rightarrow \emptyset,$$

for large enough n , which implies (37). \square

Proof of theorem 5.7 Throughout this proof $G_n(h, \zeta) = \sqrt{n} h^T \mathbb{P}_n g_{n,\zeta} - \frac{1}{2} h^T I_{n,\zeta} h$, for all h and all ζ . Furthermore, we abbreviate $\theta_n(h_n)$ to θ_n , $D(\theta_n, \rho_n)$ to D_n and omit explicit notation for (X_1, \dots, X_n) -dependence in several places. Let $\delta, \epsilon > 0$ be given and let $\theta_n = \theta_0 + n^{-1/2} h_n$ with (h_n) bounded in P_0 -probability. Then there exists a constant $M > 0$ such that $P_0^n(\|h_n\| > M) < \frac{1}{2}\delta$ for all $n \geq 1$. With (h_n) bounded, the assumption of consistency under $n^{-1/2}$ -perturbation says that,

$$P_0^n \left(\log \Pi(D_n \mid \theta = \theta_n; X_1, \dots, X_n) \geq -\epsilon \right) > 1 - \frac{1}{2}\delta,$$

for large enough n . This implies that the posterior's numerator and denominator are related through,

$$\begin{aligned} P_0^n \left(\int_H \prod_{i=1}^n \frac{p_{\theta_n, \eta}}{p_{\theta_0, \eta_0}}(X_i) d\Pi_H(\eta) \right. \\ \left. \leq e^\epsilon 1_{\{\|h_n\| \leq M\}} \int_{D_n} \prod_{i=1}^n \frac{p_{\theta_n, \eta}}{p_{\theta_0, \eta_0}}(X_i) d\Pi_H(\eta) \right) > 1 - \delta. \end{aligned} \quad (55)$$

We continue with the integral over D_n under the restriction $\|h_n\| \leq M$. Next, parametrize the model locally in terms of (θ, ζ_n) , *c.f.* (31). Define B_n as the image of D_n under reparametrization (31) and C_n by $D(\theta_0, \rho_n) = \eta_0 + C_n$ (*i.e.* the image of $D(\theta, \rho_n)$ with $\theta = \theta_0$ under any of the reparametrizations (31) or (19)).

$$\int_{D_n} \prod_{i=1}^n \frac{p_{\theta_n, \eta}}{p_{\theta_0, \eta_0}}(X_i) d\Pi_H(\eta) = \int_{B_n} \prod_{i=1}^n \frac{q_{n, \theta_n, \zeta}}{q_{n, \theta_0, 0}}(X_i) d\Pi_n(\zeta | \theta = \theta_n), \quad (56)$$

where $\Pi_n(\cdot | \theta)$ denotes the prior for ζ_n given θ , *i.e.* Π_H translated over $\eta_n(\theta)$. Next we note that by Fubini's theorem and the domination condition (35), there exists a constant $L > 0$ such that,

$$\begin{aligned} \left| P_0^n \int_{B_n} \prod_{i=1}^n \frac{q_{n, \theta_n, \zeta}}{q_{n, \theta_0, 0}}(X_i) d\Pi_n(\zeta | \theta = \theta_n) \right. \\ \left. - P_0^n \int_{C_n} \prod_{i=1}^n \frac{q_{n, \theta_n, \zeta}}{q_{n, \theta_0, 0}}(X_i) d\Pi_n(\zeta | \theta = \theta_0) \right| \\ \leq L \left| \Pi_n(B_n | \theta = \theta_n) - \Pi_n(C_n | \theta = \theta_0) \right| \\ = L \left| \Pi_H(D(\theta_n, \rho_n)) - \Pi_H(D(\theta_0, \rho_n)) \right|, \end{aligned}$$

for large enough n . Lemma 5.5 asserts that the difference on the *r.h.s.* of the above display is $o(1)$, so that,

$$\int_{B_n} \prod_{i=1}^n \frac{q_{n, \theta_n, \zeta}}{q_{n, \theta_0, 0}}(X_i) d\Pi_n(\zeta | \theta = \theta_n) = \int_{C_n} \prod_{i=1}^n \frac{q_{n, \theta_n, \zeta}}{q_{n, \theta_0, 0}}(X_i) d\Pi(\zeta) + o_{P_0}(1), \quad (57)$$

where we use the notation $\Pi(A) = \Pi_n(\zeta \in A | \theta = \theta_0)$ for brevity. We define for all $\zeta, \epsilon > 0, n \geq 1$ the events $F_n(\zeta, \epsilon) = \{\sup_h |G_n(h, \zeta) - G_n(h, 0)| \leq \epsilon\}$. With (35) as a domination condition, Fatou's lemma and the fact that $F_n^c(0, \epsilon) = \emptyset$ lead to,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \int_{C_n} Q_{n, \theta_n, \zeta}^n(F_n^c(\zeta, \epsilon)) d\Pi(\zeta) \\ \leq \int \limsup_{n \rightarrow \infty} 1_{C_n \setminus \{0\}}(\zeta) Q_{n, \theta_n, \zeta}^n(F_n^c(\zeta, \epsilon)) d\Pi(\zeta) = 0, \end{aligned} \quad (58)$$

(again using (35) in the last step). Combined with Fubini's theorem, this suffices to conclude that,

$$\int_{C_n} \prod_{i=1}^n \frac{q_{n, \theta_n, \zeta}}{q_{n, \theta_0, 0}}(X_i) d\Pi(\zeta) = \int_{C_n} \prod_{i=1}^n \frac{q_{n, \theta_n, \zeta}}{q_{n, \theta_0, 0}}(X_i) 1_{F_n(\zeta, \epsilon)} d\Pi(\zeta) + o_{P_0}(1), \quad (59)$$

and we continue with the first term on the *r.h.s.*. By stochastic local asymptotic normality for every ζ , expansion (32) of the log-likelihood implies that,

$$\prod_{i=1}^n \frac{q_{n,\theta_n,\zeta}}{q_{n,\theta_0,0}}(X_i) = \prod_{i=1}^n \frac{q_{n,\theta_0,\zeta}}{q_{n,\theta_0,0}}(X_i) e^{G_n(h_n,\zeta)+R_n(h_n,\zeta)}, \quad (60)$$

where the rest term is of order $o_{Q_{n,\theta_0,\zeta}}(1)$. Accordingly, we define, for every ζ , the events $A_n(\zeta, \epsilon) = \{|R_n(h_n, \zeta)| \leq \frac{1}{2}\epsilon\}$, so that $Q_{\theta_0,\zeta}^n(A_n^c(\zeta, \epsilon)) \rightarrow 0$. Contiguity then implies that $Q_{n,\theta_n,\zeta}^n(A_n^c(\zeta, \epsilon)) \rightarrow 0$ as well. Reasoning as in (59) we see that,

$$\begin{aligned} & \int_{C_n} \prod_{i=1}^n \frac{q_{n,\theta_n,\zeta}}{q_{n,\theta_0,0}}(X_i) 1_{F_n(\zeta,\epsilon)} d\Pi(\zeta) \\ &= \int_{C_n} \prod_{i=1}^n \frac{q_{n,\theta_n,\zeta}}{q_{n,\theta_0,0}}(X_i) 1_{A_n(\zeta,\epsilon) \cap F_n(\zeta,\epsilon)} d\Pi(\zeta) + o_{P_0}(1). \end{aligned} \quad (61)$$

For fixed n and ζ and for all $(X_1, \dots, X_n) \in A_n(\zeta, \epsilon) \cap F_n(\zeta, \epsilon)$:

$$\left| \log \prod_{i=1}^n \frac{q_{n,\theta_n,\zeta}}{q_{n,\theta_0,0}}(X_i) - G_n(h_n, 0) \right| \leq 2\epsilon,$$

so that the first term on the *r.h.s.* of (61) satisfies the bounds,

$$\begin{aligned} & e^{G_n(h_n,0)-2\epsilon} \int_{C_n} \prod_{i=1}^n \frac{q_{n,\theta_0,\zeta}}{q_{n,\theta_0,0}}(X_i) 1_{A_n(\zeta,\epsilon) \cap F_n(\zeta,\epsilon)} d\Pi(\zeta) \\ & \leq \int_{C_n} \prod_{i=1}^n \frac{q_{n,\theta_n,\zeta}}{q_{n,\theta_0,0}}(X_i) 1_{A_n(\zeta,\epsilon) \cap F_n(\zeta,\epsilon)} d\Pi(\zeta) \\ & \leq e^{G_n(h_n,0)+2\epsilon} \int_{C_n} \prod_{i=1}^n \frac{q_{n,\theta_0,\zeta}}{q_{n,\theta_0,0}}(X_i) 1_{A_n(\zeta,\epsilon) \cap F_n(\zeta,\epsilon)} d\Pi(\zeta). \end{aligned} \quad (62)$$

The integral factored into lower and upper bounds can be relieved of the indicator for $A_n \cap F_n$ by reversing the argument that led to (59) and (61) (with θ_0 replacing θ_n), at the expense of an $e^{o_{P_0}(1)}$ -factor. Substituting in (62) and using, consecutively, (61), (59), (57) and (55) for the bounded integral, we find,

$$e^{G_n(h_n,0)-3\epsilon+o_{P_0}(1)} s_n(0) \leq s_n(h_n) \leq e^{G_n(h_n,0)+3\epsilon+o_{P_0}(1)} s_n(0).$$

According to (33) $g_{n,\zeta=0}$ converges to $\tilde{\ell}_{\theta_0,\eta_0}$ in $L_2(P_0)$. As a result, $I_{n,\zeta=0} = \|g_{n,0}\|_{P_0,2}^2$ converges to $\|\tilde{\ell}_{\theta_0,\eta_0}\|_{P_0,2}^2 = \tilde{I}_{\theta_0,\eta_0}$ and, by Markov's inequality and the boundedness of h_n ,

$$n^{1/2} \mathbb{P}_n h_n^T (g_{n,0} - \tilde{\ell}_{\theta_0,\eta_0}) = \mathbb{G}_n h_n^T (g_{n,0} - \tilde{\ell}_{\theta_0,\eta_0}) = o_{P_0}(1).$$

So $G_n(h_n, 0)$ differs from the *r.h.s.* of (30) only by an $o_{P_0}(1)$ -term and we conclude that (30) holds. \square

Acknowledgements

The author thanks P. Bickel, Y. D. Kim and B. Knapik and A. van der Vaart for valuable discussions and suggestions.

References

- [1] T. ANDERSON, *Estimating linear statistical relationships*, Ann. Statist. **12** (1984), 1–45.
- [2] P. BICKEL and J. YAHAV, *Some contributions to the asymptotic theory of Bayes solutions*, Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete **11** (1969), 257–276.
- [3] P. BICKEL, *On adaptive estimation*, Ann. Statist. **10** (1982), 647–671.
- [4] P. BICKEL, C. KLAASSEN Y. RITOV, and J. WELLNER, *Efficient and adaptive estimation for semiparametric models (2nd edition)*, Springer, New York (1998).
- [5] P. BICKEL and B. KLEIJN, *The semi-parametric Bernstein-Von Mises theorem*, Ann. Statist. **40** (2012), 206–237.
- [6] L. BIRGÉ, *Approximation dans les espaces métriques et théorie de l'estimation*, Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete **65** (1983), 181–238.
- [7] L. BIRGÉ, *Sur un théorème de minimax et son application aux tests*, Probability and Mathematical Statistics **3** (1984), 259–282.
- [8] S. BOUCHERON and E. GASSIAT, *A Bernstein-Von Mises theorem for discrete probability distributions*, Electron. J. Statist. **3** (2009), 114–148.
- [9] I. CASTILLO, *Semiparametric Bernstein-Von Mises theorem and bias, illustrated with Gaussian process priors*, CNRS preprint (2011).
- [10] I. CASTILLO, *A semiparametric Bernstein-Von Mises theorem for Gaussian process priors*, Prob. Theory Rel. Fields **152** (2012), 53–99.
- [11] M. W. CHAE, Y. D. KIM and B. KLEIJN, *Bayesian efficiency for semiparametric Dirichlet mixtures*, (in preparation).
- [12] H. CHEN and J. SHIAU, *A two-stage spline-smoothing method for partially linear models*, Journal of Statistical Planning and Inference **27** (1991), 187–201.
- [13] G. CHENG and M. KOSOROK, *General frequentist properties of the posterior profile distribution*, Ann. Statist. **36** (2008), 1819–1853.
- [14] S. CHOI, W. HALL and A. SCHICK, *Asymptotically uniformly most powerful tests in parametric and semiparametric models*, Ann. Statist. **24** (1996), pp. 841–861.
- [15] D. R. COX, *Regression models and life-tables*, J. Roy. Statist. Soc. **B34** (1972), 187–220.
- [16] D. D. COX, *An analysis of Bayesian inference for nonparametric regression*, Ann. Statist. **21** (1993), 903–924.
- [17] H. CRAMÉR, *Mathematical methods of statistics*, Princeton University Press, Princeton (1946).
- [18] P. DE BLASI and N. HJORT, *Bayesian Survival Analysis in Proportional Hazard Models with Logistic Relative Risk*, Scand. J. Statist. **34** (2007), 229–257.
- [19] P. DE BLASI and N. HJORT, *The Bernstein-von Mises theorem in semiparametric competing risks models*, J. Statist. Plann. Inference **139** (2009), 23162328.
- [20] R. DE JONGE and J. VAN ZANTEN, *Semiparametric Bernstein-von Mises for the error standard deviation*, Electron. J. Statist. **7** (2013), 217–243.

- [21] P. DIACONIS and D. FREEDMAN, *On the consistency of Bayes estimates*, Ann. Statist. **14** (1986), 1–26.
- [22] P. DIACONIS and D. FREEDMAN, *Consistency of Bayes estimates for nonparameteric regression: Normal theory*, Bernoulli **4** (1998), 411–444.
- [23] N. DUNFORD and J. SCHWARTZ, *Linear operators, Part I, General Theory* Wiley-Interscience, New York (1958).
- [24] M. ESCOBAR and M. WEST, *Bayesian density estimation and inference with mixtures*, J. Amer. Statist. Assoc. **90** (1995), 577–588.
- [25] T. FERGUSON, *A Bayesian Analysis of Some Nonparametric Problems*, Ann. Statist. **1** (1973), 209–230.
- [26] T. FERGUSON and E. PHADIA, *Bayesian nonparametric estimation based on censored data*, Ann. Statist. **7** (1979), 163–186.
- [27] T. FERGUSON, *Bayesian density estimation by mixtures of Normal distributions*, In *Recent Advances in Statistics*, (eds. M. Rizvi, J. Rustagi and D. Siegmund, eds.), 287–302. Academic Press, New York (1983).
- [28] R. FISHER, *Statistical methods and scientific inference* (2nd edition), Oliver and Boyd, London (1959).
- [29] D. FREEDMAN, *On the asymptotic behavior of Bayes estimates in the discrete case I*, Ann. Math. Statist. **34** (1963), 1386–1403.
- [30] D. FREEDMAN, *On the Bernstein-von Mises theorem with infinite dimensional parameters*, Ann. Statist. **27** (1999), 1119–1140.
- [31] S. GHOSAL, J. GHOSH and A. VAN DER VAART, *Convergence rates of posterior distributions*, Ann. Statist. **28** (2000), 500–531.
- [32] S. GHOSAL and A. VAN DER VAART, *Entropies and rates of convergence for maximum likelihood and Bayes estimation for mixtures of normal densities*, Ann. Statist. **29** (2001), 1233–1263.
- [33] S. GHOSAL and A. VAN DER VAART, *Posterior convergence rates of Dirichlet mixtures at smooth densities*, Ann. Statist. **35** (2007), 697–723.
- [34] J. HÁJEK, *A characterization of limiting distributions of regular estimates*, Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete **14** (1970), 323–330.
- [35] J. HÁJEK, *Local asymptotic minimax and admissibility in estimation*, Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability **1**, 175–194. University of California Press, Berkeley (1972).
- [36] N. HJORT, *Nonparametric Bayes estimators based on beta processes in models for life history data*, Ann. Statist. **18** (1990), 1259–1294.
- [37] I. IBRAGIMOV and R. HAS’MINSKI, *Statistical estimation: asymptotic theory*, Springer, New York (1981).
- [38] I. JOHNSTONE, *High dimensional Bernstein-von Mises: simple examples*, in *Borrowing Strength: Theory Powering Applications – A Festschrift for Lawrence D. Brown*, (eds. J. Berger, T. Cai and I. Johnstone) IMS, Beachwood (2010), 87–98.
- [39] J. KALBFLEISCH, *Non-parametric Bayesian analysis of survival time data*, J. Roy. Statist. Soc. **B40** (1978), 214–221.
- [40] YONGDAI KIM and JAEYONG LEE, *A Bernstein Von Mises theorem in the nonparametric right-censoring model*, Ann. Statist. **4** (2004), 1492–1512.

- [41] YONGDAI KIM, *The Bernstein Von Mises theorem for the proportional hazard model*, Ann. Statist. **4** (2006), 1678–1700.
- [42] G. KIMELDORF and G. WAHBA, *A correspondence between Bayesian estimation on stochastic processes and smoothing by splines*, Ann. Math. Statist. **41** (1970), 495–502.
- [43] C. KLAASSEN, *Consistent estimation of the influence function of locally asymptotically linear estimators*, Ann. Statist. **15** (1987), 1548–1562.
- [44] B. KLEIJN, *Bayesian asymptotics under misspecification*. PhD. Thesis, Free University Amsterdam (2003).
- [45] B. KLEIJN and A. VAN DER VAART, *Misspecification in ∞ -dimensional Bayesian statistics*. Ann. Statist. **34** (2006), 837–877.
- [46] B. KLEIJN and A. VAN DER VAART, *The Bernstein-Von-Mises theorem under misspecification*. Electron. J. Statist. **6** (2012), 354–381.
- [47] B. KLEIJN and B. KNAPIK, *Semiparametric posterior limits under local asymptotic exponentiality*, (submitted for publication).
- [48] B. KNAPIK, A. VAN DER VAART, and J. VAN ZANTEN, *Bayesian inverse problems with Gaussian priors*, Ann. Statist. **39** (2011), 2626–2657.
- [49] A. KOLMOGOROV and V. TIKHOMIROV, *Epsilon-entropy and epsilon-capacity of sets in function spaces*, American Mathematical Society Translations (series 2), **17** (1961), 277–364.
- [50] W. KRUIJER and J. ROUSSEAU, *Bayesian semi-parametric estimation of the long-memory parameter under FEXP-priors*, preprint (arXiv:1202.4863 [math.ST])
- [51] H. LEAHU, *On the Bernsteinvon Mises phenomenon in the Gaussian white noise model*, Electron. J. Stat. **5** (2011) 373404.
- [52] L. LE CAM, *On some asymptotic properties of maximum-likelihood estimates and related Bayes estimates*, University of California Publications in Statistics, **1** (1953), 277–330.
- [53] L. LE CAM, *Locally asymptotically normal families of distributions*, University of California Publications in Statistics, **3** (1953), 37–98.
- [54] L. LE CAM, *Limits of Experiments*, Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability **1**, 245–261. University of California Press, Berkeley (1972).
- [55] L. LE CAM, *Convergence of Estimates Under Dimensionality Restrictions*, Ann. Statist. Volume **1** (1973), 38–53.
- [56] L. LE CAM, *Asymptotic methods in statistical decision theory*, Springer, New York (1986).
- [57] L. LE CAM and G. YANG, *Asymptotics in Statistics: some basic concepts*, Springer, New York (1990).
- [58] L. LE CAM, *Maximum likelihood: An introduction*, Internat. Statist. Rev. **58** (1990), 153–171 (Previously published as *Lecture Notes No. 18*, Statistics Branch of the Dept. of Mathematics, Univ. of Maryland (1979)).
- [59] E. LEHMANN and G. CASELLA, *Theory of point estimation*, Springer, New York (1998).
- [60] A. Y. LO, *On a class of Bayesian nonparametric estimates I: Density estimates*, Ann. Statist. **12** (1984), 351357.

- [61] E. MAMMEN and S. VAN DE GEER, *Penalized quasi-likelihood estimation in partial linear models*, Ann. Statist. **25** (1997), 1014–1035.
- [62] S. MURPHY and A. VAN DER VAART, *On profile likelihood*, J. Amer. Statist. Assoc. **95** (2000), 449–485.
- [63] J. PFANZAGL, *Consistency of maximum likelihood estimators for certain nonparametric families, in particular: mixtures*, J. Statist. Plann. Inference **19** (1988), 137–158.
- [64] V. RIVOIRARD and J. ROUSSEAU, *Bernstein-Von Mises theorem for linear functionals of the density*, preprint 2009, arXiv:math.ST/0908.4167v1
- [65] C. ROBERT, *The Bayesian choice: from decision-theoretic foundations to computational implementation*, Springer, New York (2001).
- [66] L. SCHWARTZ, *On Bayes procedures*, Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete **4** (1965), 10–26.
- [67] T. SEVERINI and W. WONG, *Profile likelihood and conditionally parametric models*, Ann. Statist. **20** (1992), 1768–1802.
- [68] X. SHEN and L. WASSERMAN, *Rates of convergence of posterior distributions*, Ann. Statist. **29** (2001), 687–714.
- [69] X. SHEN, *Asymptotic normality of semiparametric and nonparametric posterior distributions*, Journal of the American Statistical Association **97** (2002), 222–235.
- [70] T. SHIVELY, R. KOHN and S. WOOD, *Variable selection and function estimation in additive nonparametric regression using a data-based prior*, Journal of the American Statistical Association **94** (1999), 777–804.
- [71] C. STEIN, *Efficient nonparametric testing and estimation*, Proc. Third Berkeley Symp. Math. Statist. Prob. **1** (1956), 187–196.
- [72] H. STRASSER, *The mathematical theory of Statistics*, De Gruyter, Berlin (1985).
- [73] M. TAUPIN, *Semi-parametric estimation in the nonlinear structural errors-in-variables model*, Ann. Statist. **29** (2001), 66–93.
- [74] A. VAN DER VAART, *Statistical Estimation in Large Parameter Spaces*, CWI Tract **44**, CWI Amsterdam (1988).
- [75] A. VAN DER VAART, *Efficient maximum likelihood estimation in semiparametric mixture models*, Ann. Statist. **24** (1996), 862–878.
- [76] A. VAN DER VAART and J. WELLNER, *Weak Convergence and Empirical Processes*, Springer Verlag, New York (1996).
- [77] A. VAN DER VAART, *Asymptotic Statistics*, Cambridge University Press, Cambridge (1998).
- [78] A. VAN DER VAART and J. VAN ZANTEN, *Rates of contraction of posterior distributions based on Gaussian process priors*, Ann. Statist. **36** (2008), 1435–1463.
- [79] G. WAHBA, *Improper priors, spline smoothing and the problem of guarding against model error in regression*, J. Roy. Statist. Soc. **B40** (1978), 364–372.
- [80] W. H. WONG and X. SHEN, *Probability Inequalities for Likelihood Ratios and Convergence Rates of Sieve MLE's*, Ann. Statist. **23** (1995), 339–362.