

Bayes Extended Estimators for Curved Exponential Families

Michiko Okudo^{ID} and Fumiyasu Komaki^{ID}, *Member, IEEE*

Abstract—The Bayesian predictive density has complex representation and does not belong to any finite-dimensional statistical model except for in limited situations. In this paper, we introduce its simple approximate representation employing its projection onto a finite-dimensional exponential family. Its theoretical properties are established parallelly to those of the Bayesian predictive density when the model belongs to curved exponential families. It is also demonstrated that the projection asymptotically coincides with the plugin density with the posterior mean of the expectation parameter of the exponential family, which we refer to as the Bayes extended estimator. Information-geometric correspondence indicates that the Bayesian predictive density can be represented as the posterior mean of the infinite-dimensional exponential family. The Kullback–Leibler risk performance of the approximation is demonstrated by numerical simulations and it indicates that the posterior mean of the expectation parameter approaches the Bayesian predictive density as the dimension of the exponential family increases. It also suggests that approximation by projection onto an exponential family of reasonable size is practically advantageous with respect to risk performance and computational cost.

Index Terms—Bayesian prediction, curved exponential family, information geometry.

I. INTRODUCTION

CONSTRUCTING predictive densities is a fundamental problem in statistical analysis that aims at predicting the behavior of future samples using past observations. Let us suppose that we have observations $x^n = \{x(1), x(2), \dots, x(n)\}$ that are independently distributed according to a probability distribution with density function $p(x; \omega)$ that belongs to a statistical model

$$\mathcal{P} = \{p(x; \omega) \mid \omega \in \Omega\}.$$

The objective is to provide the predictive density of $y = x(n+1)$ that is independently distributed according to the same

density $p(y; \omega)$. We adopt the Kullback–Leibler divergence

$$D\{p(y; \omega); \hat{p}(y; \omega)\} = \int p(y; \omega) \log \frac{p(y; \omega)}{\hat{p}(y; \omega)} dy$$

as a loss function of a predictive density $\hat{p}(y; \omega)$. Then, the risk function and the Bayes risk with respect to a prior $\pi(\omega)$ can be written as

$$\begin{aligned} E[D\{p(y; \omega); \hat{p}(y; \omega)\}] \\ = \int p(x^n; \omega) D\{p(y; \omega); \hat{p}(y; \omega)\} dx^n, \end{aligned}$$

and

$$\int \pi(\omega) \int p(x^n; \omega) D\{p(y; \omega); \hat{p}(y; \omega)\} dx^n d\omega,$$

respectively.

Except for in limited cases, the Bayesian predictive distribution does not belong to any finite-dimensional model, which makes it intractable to obtain the full density although it is optimal with respect to the Bayes risk. The Bayesian predictive density is defined by

$$\hat{p}_\pi(y \mid x^n) = \int p(y; \omega) p_\pi(\omega \mid x^n) d\omega$$

where $p_\pi(\omega \mid x^n)$ is the posterior density

$$p_\pi(\omega \mid x^n) = \frac{p(x^n; \omega) \pi(\omega)}{\int p(x^n; \omega) \pi(\omega) d\omega}$$

of ω . It is shown in [1] that the Bayesian predictive density is optimal with respect to the Bayes risk in terms of the Kullback–Leibler divergence in the family of all probability densities, which we denote as \mathcal{F} . However, the full Bayesian predictive density is intractable in most problems due to the complex representation that involves averaging plugin densities about the model parameters. It is not included in the model \mathcal{P} or even in any finite-dimensional model in most problems, while plugin densities are always included in \mathcal{P} as they are constructed by plugging-in an estimator $\hat{\omega}(x^n)$ to the model.

In the present paper, we represent the Bayesian predictive density as the infinite-dimensional limit of a parameterized distribution of an exponential family. We demonstrate that the Bayesian predictive density can be considered as an infinite-dimensional extension of the plugin density with the posterior mean of the expectation parameter of an exponential family. It is shown that theoretical properties including optimality with respect to the Bayes risk are appropriately retained by this extension. The plugin density with the posterior mean

Manuscript received February 10, 2020; revised August 30, 2020; accepted October 15, 2020. Date of publication October 30, 2020; date of current version January 21, 2021. This work was supported in part by JSPS KAKENHI under Grant JP18J10499 (Japan Society for the Promotion of Science), in part by MEXT KAKENHI under Grant 16H06533 (Ministry of Education, Culture, Sports, Science and Technology), in part by JST CREST under Grant JPMJCR1763 (Japan Science and Technology Agency), and in part by AMED under Grant JP19dm0207001 (Japan Agency for Medical Research and Development). (*Corresponding author: Michiko Okudo.*)

Michiko Okudo is with the Department of Mathematical Informatics, The University of Tokyo, Tokyo 113-8656, Japan (e-mail: okudo@mist.i.u-tokyo.ac.jp).

Fumiyasu Komaki is with the Department of Mathematical Informatics, The University of Tokyo, Tokyo 113-8656, Japan, and also with the RIKEN Center for Brain Science, Wako 351-0198, Japan (e-mail: komaki@mist.i.u-tokyo.ac.jp).

Communicated by E. Gassiat, Associate Editor for Probability and Statistics. Digital Object Identifier 10.1109/TIT.2020.3035027

of the expectation parameter coincides with the projection of the Bayesian predictive density onto the exponential family with respect to the Fisher metric. It is shown that it approaches the Bayesian predictive density closer with respect to the risk as the projected exponential family increases. There is also information-geometrical correspondence between the Bayesian predictive density and the projected density that comes from the correspondence between \mathcal{F} and the exponential family. In practice, the Bayesian predictive density can be computationally approximated, for example, by taking the mean of plugin densities using Markov chain Monte Carlo simulations, or by performing an approximation of the posterior density using methods like the Laplace method, although the objective of this research is not to develop a complex approximation based on computational methods. Apart from computational approximations, a class of empirical Bayes predictive densities is proposed for multivariate normal models in [2] to avoid the intractable implementations of Bayesian predictive densities. Rather than constructing an approximation of the Bayesian predictive density that has good performance in terms of risk, we aim to formulate a simple interpretation of the Bayesian predictive density that maintains its theoretical properties on a finite-dimensional model.

The outline of the construction of the approximate predictive densities is explained below. We consider a model of a subspace of an exponential family as \mathcal{P} , namely we consider a statistical model of a curved exponential family

$$\mathcal{P} = \{p(x; \omega) = s(x) \exp(\theta^i(\omega)x_i - \Psi(\theta(\omega))) \mid \omega = (\omega^a) \in \Omega, a = 1, \dots, d, i = 1, \dots, m\},$$

where $\Omega \subset \mathbb{R}^d$ and $1 \leq d \leq m$. The model \mathcal{P} parametrized by ω is embedded in an exponential family parametrized by θ , thus here we represent θ as $\theta(\omega)$. Summation over a repeated index is automatically taken according to Einstein's summation convention: if an index occurs as an upper and lower index in one term, then the summation is implied. Curved exponential families embedded in exponential families can express a variety of models including network models (e.g., [3]) and time series models (e.g., [4]). They can also be applied to stochastic processes [5]. We consider predictive densities in a finite-dimensional full exponential family

$$\mathcal{E} = \{p(x; \theta) = s(x) \exp(\theta^i x_i - \Psi(\theta)) \mid \theta = (\theta^i) \in \Theta, i = 1, \dots, m\} \quad (\Theta \subseteq \mathbb{R}^m)$$

that includes the original curved exponential model \mathcal{P} . We refer to plugin densities in \mathcal{E} as extended plugin densities. The inclusion relation is $\mathcal{P} \subseteq \mathcal{E} \subseteq \mathcal{F}$ and we consider the middle layer of the three-layer structure. The coordinate system $\theta = (\theta^i)$ ($i = 1, \dots, m$) is called the natural parameter of exponential families. Another coordinate system $\eta = (\eta_i)$ defined by

$$\eta_i = E(x_i) = \frac{\partial}{\partial \theta^i} \Psi(\theta) \quad (i = 1, \dots, m) \quad (1)$$

is called the expectation parameter. The posterior mean of η is closely related to the Bayesian predictive density, and the extended plugin density with the posterior mean of η is considered in this paper. Based on the idea of covering \mathcal{F}

TABLE I
ALL PROBABILITY DENSITIES AND A FULL EXPONENTIAL FAMILY

\mathcal{F} : all probability densities	\mathcal{E} : a full exponential family
$p(x)$: m-representation	η : m-affine parameter
$\log p(x)$: e-representation	θ : e-affine parameter
infinite dimensional model	finite dimensional model

by extending exponential families, we specify the models that can be embedded in exponential families, and the theoretical properties described in the following sections are based on this embedding. The policy of expressing a probability density by extending exponential families has been investigated, for example, in [6], in which a log-density is approximated using series of polynomials and the rate of convergences is obtained. It should be noted that the practical advantage illustrated in the numerical experiments in Section IV can be attained for other models if it is possible to find an appropriate exponential family onto which we project the Bayesian predictive density.

From the viewpoint of information geometry, the posterior mean of η can be considered as the correspondence to the Bayesian predictive density in \mathcal{E} . Table I represents the infinite-finite correspondence of m and e representations. Here, "m" and "e" are short notations denoting "mixture" and "exponential," respectively. Exponential families and mixture families are important dual families in information geometry, and their typical representations are denoted as m-representation and e-representation, respectively. Concerning exponential families, the e-representation is $\log p(x) = \theta^i x_i - \Psi(\theta) + \log s(x)$, and θ is called the e-affine parameter, as the basis vector fields $(\partial/\partial \theta^i)p(x)$ ($i = 1, \dots, m$) are parallel vector fields with respect to the e-connection as defined in Section II of this paper. Concerning mixture families, the m-representation can be written as $p(x) = \eta_i q^i(x) + c(x)$, and η is the affine coordinate system about the m-connection (also defined in Section II). We can also set the m-affine coordinate system in exponential families, and it is defined by (1). Here, the m-affine (or e-affine) parameters are the finite-dimensional typical representations of m-representation (or e-representation, respectively). The Bayesian predictive density is the posterior mean about the m-representation (that is, density functions), and its finite-dimensional correspondence is the posterior mean about (η_i) . Since the Bayesian predictive density is optimal in the infinite-dimensional exponential family, we might expect that the posterior mean of (η_i) exhibit the same properties as the Bayesian predictive density, such as optimality with respect to the Bayes risk.

The properties of the posterior mean of η are investigated in following sections as follows. In Section III-A, we show that the extended plugin density with the posterior mean of η is optimal with respect to the Bayes risk in the finite-dimensional exponential family. We denote the posterior mean of η as the Bayes extended estimator. In Section III-B, the extended plugin density with the Bayes extended estimator is proved to be the projection of the Bayesian predictive density onto \mathcal{E} in terms of the Fisher metric. In Section III-C, its optimality with respect to the risk along orthogonal shift from the model is shown to be common with the property of optimality

in the case of the Bayesian predictive density. The relation between the projection angle under the Fisher metric and the risk difference between the Bayesian predictive density and the extended plugin density with the Bayes extended estimator is investigated. In Section IV, we compare the risk performance of the Bayes estimator, the Bayes extended estimator, and the Bayesian predictive density by conducting numerical simulations on the Gaussian spiked covariance models. We confirm that the projection angle converges to zero as the dimension of \mathcal{E} increases. The simulation results also suggest that the projection of the Bayesian predictive density is practically effective in approximating it with respect to the Kullback–Leibler risk and the computational cost.

II. PRELIMINARIES

In this section, we prepare some information-geometric notions. For details of the notions and notation concerning the differential geometry of curved exponential families, refer to [7].

Let a, b, \dots be indices for ω . Let $T_\omega \mathcal{P}$ be the tangent space of \mathcal{P} at a point ω . The tangent space $T_\omega \mathcal{P}$ is identified with the vector space spanned by $\partial_a p(x; \omega)$ ($a = 1, \dots, d$), where ∂_a denotes $\partial/\partial \omega^a$. We define inner products in the tangent space by

$$\langle \partial_a p(x; \omega), \partial_b p(x; \omega) \rangle = \int \frac{\partial_a p(x; \omega) \partial_b p(x; \omega)}{p(x; \omega)} dx. \quad (2)$$

In a statistical model \mathcal{P} , each component of the Fisher information matrix is defined by

$$g_{ab}(\omega) = \langle \partial_a p(x; \omega), \partial_b p(x; \omega) \rangle.$$

Let g^{ab} be a component of the inverse matrix of (g_{ab}) . Then, e-connection and the m-connection coefficients are defined as:

$$\overset{e}{\Gamma}_{abc}(\omega) = \int p(x; \omega) \{ \partial_a \partial_b \log p(x; \omega) \} \{ \partial_c \log p(x; \omega) \} dx$$

and

$$\overset{m}{\Gamma}_{abc}(\omega) = \int \frac{\partial_a \partial_b p(x; \omega) \partial_c p(x; \omega)}{p(x; \omega)} dx,$$

respectively. We define

$$\overset{e}{\Gamma}_{ab}^c = \overset{e}{\Gamma}_{abd} g^{dc}, \quad \overset{m}{\Gamma}_{ab}^c = \overset{m}{\Gamma}_{abd} g^{dc},$$

and

$$\begin{aligned} T_{abc} &= \overset{m}{\Gamma}_{abc} - \overset{e}{\Gamma}_{abc} \\ &= \int p(x; \omega) \{ \partial_a \log p(x; \omega) \} \{ \partial_b \log p(x; \omega) \} \\ &\quad \times \{ \partial_c \log p(x; \omega) \} dx. \end{aligned}$$

The Jeffreys prior density is given by

$$\pi_J(\omega) = \sqrt{|g(\omega)|},$$

where $|g(\omega)|$ is the determinant of the matrix $(g_{ab}(\omega))$.

The coordinate systems (θ^i) and (η_i) of the exponential family \mathcal{E} are dual to each other in the sense that

$$\left\langle \frac{\partial}{\partial \theta^i} p(x; \theta), \frac{\partial}{\partial \eta_j} p(x; \theta) \right\rangle = \delta_j^i \quad (3)$$

where δ_j^i is the Kronecker delta. In a curved exponential family, e-connection and m-connection coefficients are expressed as:

$$\overset{e}{\Gamma}_{abc} = (\partial_a \partial_b \theta^i)(\partial_c \eta_i) \quad \text{and} \quad \overset{m}{\Gamma}_{abc} = (\partial_a \partial_b \eta_i)(\partial_c \theta^i), \quad (4)$$

respectively.

In the rest of the paper, we assume regularity conditions to ensure that equalities such as

$$\int o_p(1) dy = o(1), \quad \int O_p(n^{-1}) dy = O(n^{-1}),$$

hold. For the details of the regularity conditions, see [8].

III. MAIN RESULTS

A. Optimality With Respect to Bayes Risk

The posterior mean of η (we denote as $\bar{\eta}_\pi$) of \mathcal{E} is evaluated as follows:

$$\bar{\eta}_\pi = \frac{\int \eta(\omega) p(x^n; \eta(\omega)) \pi(\omega) d\omega}{\int p(x^n; \eta(\omega)) \pi(\omega) d\omega}.$$

Note that $\bar{\eta}_\pi(x^n) \neq \eta(\bar{\omega}_\pi(x^n))$ in general where $\bar{\omega}_\pi(x^n)$ is the posterior mean of ω :

$$\bar{\omega}_\pi = \frac{\int \omega p(x^n; \omega) \pi(\omega) d\omega}{\int p(x^n; \omega) \pi(\omega) d\omega}.$$

We demonstrate that $p(y; \bar{\eta}_\pi)$ is optimal in \mathcal{E} with respect to the Bayes risk based on a prior π .

Proposition III.1: The Bayes risk of $p(y; \hat{\eta})$, where $\hat{\eta}$ is an estimator of η , is minimized when $\hat{\eta} = \bar{\eta}_\pi$.

Proof: Let $\hat{\theta}$ be an estimator of θ . Note that θ and η are functions of ω . The Kullback–Leibler loss of $p(y; \hat{\theta})$ is

$$\begin{aligned} D\{p(y; \theta(\omega)); p(y; \hat{\theta})\} \\ &= \int p(y; \theta) \log \left(\frac{\exp(\theta^i y_i - \Psi(\theta))}{\exp(\hat{\theta}^i y_i - \Psi(\hat{\theta}))} \right) dy \\ &= (\theta^i - \hat{\theta}^i) \eta_i - (\Psi(\theta) - \Psi(\hat{\theta})). \end{aligned}$$

Hence

$$\begin{aligned} &\int p_\pi(\omega | x^n) D\{p(y; \theta(\eta)); p(y; \hat{\theta}(\eta))\} d\omega \\ &= (\overline{\theta^i \eta_i} - \hat{\theta}^i \bar{\eta}_i) - (\overline{\Psi(\theta)} - \Psi(\hat{\theta})) \\ &= (\theta^i(\bar{\eta}) - \hat{\theta}^i) \bar{\eta}^i - (\Psi(\theta(\bar{\eta})) - \Psi(\hat{\theta})) \\ &\quad + (-\theta^i(\bar{\eta}) \bar{\eta}_i + \overline{\theta^i \eta_i} + \Psi(\theta(\bar{\eta})) - \overline{\Psi(\theta)}) \\ &= D\{p(y; \bar{\eta}); p(y; \hat{\theta})\} \\ &\quad + (-\theta^i(\bar{\eta}) \bar{\eta}^i + \overline{\theta^i \eta_i} + \Psi(\theta(\bar{\eta})) - \overline{\Psi(\theta)}), \quad (5) \end{aligned}$$

where, for a function $f(\eta)$,

$$\overline{f(\eta)} = \int p_\pi(\omega | x^n) f(\eta) d\omega.$$

It is minimized when $\hat{\theta} = \theta(\bar{\eta}) = \theta(\bar{\eta}_\pi)$. By multiplying (5) with $\int p(x^n; \omega) \pi(\omega) d\omega$ and then integrating with respect to x^n , it is shown that $p(y; \bar{\eta}_\pi)$ is optimal with respect to the Bayes risk in \mathcal{E} . \square

We refer to $\hat{\eta}_\pi$ as the Bayes extended estimator. Hereinafter, we denote the Bayes extended estimator and the Bayes estimator of ω about a prior π as $\hat{\eta}_\pi (= \bar{\eta}_\pi)$ and $\hat{\omega}_\pi$, respectively.

From Proposition III.1, the extended plugin density with $\hat{\eta}_\pi$ is the projection of the Bayesian predictive density onto \mathcal{E} about the Bayes risk. It is nearest to the Bayesian predictive density in \mathcal{E} regarding the Bayes risk, because the Bayesian predictive density is optimal about the Bayes risk in \mathcal{F} . In fact, $p(y, \hat{\eta}_\pi)$ coincides with the projection of the Bayesian predictive density onto \mathcal{E} regarding the Fisher metric asymptotically, as shown in Section III-B.

The choice of \mathcal{E} does not require to be fixed, and we can consider situations in which the size of the extended model \mathcal{E} can be increased, for example, by employing sequences of exponential families as in [6] and [9]. In those situations, the extended plugin density with the Bayes extended estimator $\hat{\eta}_\pi$ approaches the Bayesian predictive density as \mathcal{E} grows, as \mathcal{E} approaches the set of all probability distributions \mathcal{F} .

Here, we use a simple example to illustrate the difference of the plugin density with $\hat{\omega}_\pi$, the extended plugin density with $\hat{\eta}_\pi$, and the Bayesian predictive density.

Example (Fisher circle model) We consider two dimensional Gaussian distribution $N(\mu, I_2)$ with unknown mean vector μ and the identity covariance matrix I_2 . The density function is

$$\begin{aligned} p(x; \mu) &= \frac{1}{2\pi} \exp \left[-\frac{1}{2} \{ (x_1 - \mu_1)^2 + (x_2 - \mu_2)^2 \} \right] \\ &= \frac{1}{2\pi} \exp \left(-\frac{1}{2} (x_1^2 + x_2^2) \right) \\ &\quad \times \exp \left(x_1 \mu_1 + x_2 \mu_2 - \frac{1}{2} (\mu_1^2 + \mu_2^2) \right). \end{aligned}$$

When the mean vector μ is expressed as

$$\mu_1 = \cos \omega, \mu_2 = \sin \omega,$$

the one-dimensional submodel is called the Fisher circle model. Here, the following holds:

$$\theta^1 = \eta_1 = \mu_1, \quad \theta^2 = \eta_2 = \mu_2.$$

Then, we derive the Bayes estimator $\hat{\omega}_\pi$, the Bayes extended estimator $\hat{\eta}_\pi$, and the Bayesian predictive density. For $x^n = \{x(1), x(2), \dots, x(n)\}$,

$$\begin{aligned} p(x^n; \omega) &= \prod_{t=1}^n \frac{1}{2\pi} \exp \left(-\frac{\|x(t) - \mu(\omega)\|^2}{2} \right) \\ &= \frac{1}{(2\pi)^n} \exp \left(-\frac{\sum_{t=1}^n (x_1(t)^2 + x_2(t)^2)}{2} + \frac{n}{2} (\bar{x}_1^2 + \bar{x}_2^2) \right) \\ &\quad \times \exp \left(-\frac{n}{2} \|\bar{x} - \mu(\omega)\|^2 \right) \end{aligned}$$

where $\bar{x} = \sum_{t=1}^n x(t)/n$. Let $\bar{x} = (\|\bar{x}\| \cos \phi, \|\bar{x}\| \sin \phi)^\top$. Then by the law of cosines,

$$\begin{aligned} &-\frac{n}{2} \|\bar{x} - \mu(\omega)\|^2 \\ &= n \|\bar{x}\| \cos(\omega - \phi) + (\text{terms independent of } \omega) \end{aligned}$$




Fig. 1. The extended plugin density with $\hat{\eta}_\pi$ is the projection of the Bayesian predictive density onto \mathcal{E} . It is constructed by shifting from the plugin density with the maximum likelihood estimator (MLE).

and

$$\int_0^{2\pi} \exp(n \|\bar{x}\| \cos(\omega - \phi)) d\omega = 2\pi I_0(n \|\bar{x}\|)$$

where $I_0(\cdot)$ is the modified Bessel function of the first kind. See [10] (pp. 138–140) for the details. When the uniform prior

$$\pi(\omega) \propto 1$$

is adopted, the posterior density is

$$p_\pi(\omega | x^n) = \frac{1}{2\pi I_0(n \|\bar{x}\|)} \exp(n \|\bar{x}\| \cos(\omega - \phi)).$$

It follows that the plugin of the Bayes estimator is $N(\frac{\bar{x}}{\|\bar{x}\|}, I_2)$:

$$p(y; \hat{\omega}_\pi) = \frac{1}{2\pi} \exp \left(-\frac{1}{2} \left\| y - \frac{\bar{x}}{\|\bar{x}\|} \right\|^2 \right).$$

The extended plugin with the Bayes extended estimator is $N(\frac{I_1(n \|\bar{x}\|)}{I_0(n \|\bar{x}\|)} \frac{\bar{x}}{\|\bar{x}\|}, I_2)$:

$$p(y; \hat{\eta}_\pi) = \frac{1}{2\pi} \exp \left(-\frac{1}{2} \left\| y - \frac{I_1(n \|\bar{x}\|)}{I_0(n \|\bar{x}\|)} \frac{\bar{x}}{\|\bar{x}\|} \right\|^2 \right),$$

where $\hat{\eta}_\pi$ is not included in the circle parametrized by ω and $\hat{\eta}_\pi \neq \eta(\hat{\omega}_\pi)$. Here $I_1(\cdot)$ is the modified Bessel function of the first kind. On the other hand, the Bayesian predictive density is given by

$$p_\pi(y | x^n) = \frac{1}{2\pi} \frac{I_0(\|y + n\bar{x}\|)}{I_0(n \|\bar{x}\|)} \exp \left\{ -\frac{1}{2} (\|y\|^2 + 1) \right\}.$$

Therefore, $p_\pi(y | x^n)$ is not included in \mathcal{P} or \mathcal{E} because it is not a two-dimensional Gaussian with a covariance matrix I_2 .

B. Projection of Bayesian Predictive Densities in Terms of Fisher Metric

Here, we demonstrate that $p(y; \hat{\eta}_\pi)$ is the projection of the Bayesian predictive density regarding the Fisher metric. It is shown via asymptotic expansion of $p(y; \hat{\eta}_\pi)$, which is represented as a point in \mathcal{E} that is parallelly and orthogonally shifted from the plugin density of the maximum likelihood estimator of \mathcal{P} as shown in Figure 1.

Here, we proceed to obtain the asymptotic expansion of $\hat{\eta}_\pi$ around $\eta(\hat{\omega}_{\text{MLE}})$.

Theorem III.1: The Bayes extended estimator based on a prior $\pi(\omega)$ is expanded as

$$\begin{aligned}\hat{\eta}_\pi &= \eta(\hat{\omega}_{\text{MLE}}) \\ &+ \frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{2n} \left(\partial_a \partial_b \eta(\hat{\omega}_{\text{MLE}}) - \bar{\Gamma}_{ab}^c(\hat{\omega}_{\text{MLE}}) \partial_c \eta(\hat{\omega}_{\text{MLE}}) \right) \\ &+ \frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{n} \left(\partial_b \log \frac{\pi}{\pi_J}(\hat{\omega}_{\text{MLE}}) + \frac{T_b(\hat{\omega}_{\text{MLE}})}{2} \right) \\ &\times \partial_a \eta(\hat{\omega}_{\text{MLE}}) \\ &+ o_p(n^{-1}),\end{aligned}$$

where π_J is the density of the Jeffreys prior and $T_a = T_{abc}g^{bc}$.

Proof: See Appendix A. \square

We can obtain the asymptotic expansion of the extended plugin density with $\hat{\eta}_\pi$.

Theorem III.2: The extended plugin density with $\hat{\eta}_\pi$ is expanded as

$$\begin{aligned}p(y; \hat{\eta}_\pi) &= p(y; \eta(\hat{\omega}_{\text{MLE}})) \\ &+ \frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{2n} \left(\partial_a \partial_b \eta_i(\hat{\omega}_{\text{MLE}}) - \bar{\Gamma}_{ab}^c(\hat{\omega}_{\text{MLE}}) \partial_c \eta_i(\hat{\omega}_{\text{MLE}}) \right) \\ &\times \partial^i p(y; \hat{\omega}_{\text{MLE}}) \\ &+ \frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{n} \left(\partial_b \log \frac{\pi}{\pi_J}(\hat{\omega}_{\text{MLE}}) + \frac{T_b(\hat{\omega}_{\text{MLE}})}{2} \right) \\ &\times \partial_a p(y; \hat{\omega}_{\text{MLE}}) \\ &+ o_p(n^{-1}),\end{aligned}$$

where $\partial^i = \partial/\partial\eta_i$.

Proof: Symbols such as $\eta(\hat{\omega}_{\text{MLE}})$, $\partial_a \eta(\hat{\omega}_{\text{MLE}})$, and $\partial_a \partial_b \eta(\hat{\omega}_{\text{MLE}})$ are abbreviated to $\hat{\eta}$, $\partial_a \hat{\eta}$, and $\partial_{ab} \hat{\eta}$, respectively. Considering the asymptotic expansion introduced in Theorem III.1, we obtain the following:

$$\begin{aligned}p(y; \hat{\eta}_\pi) &= p(y; \hat{\eta}) + \left(\frac{\hat{g}^{ab}}{2n} \left(\partial_{ab} \hat{\eta}_i - \bar{\Gamma}_{ab}^c(\hat{\omega}_{\text{MLE}}) \partial_c \hat{\eta}_i \right) \right. \\ &\quad \left. + \frac{\hat{g}^{ab}}{n} \left(\partial_b \log \frac{\hat{\pi}}{\hat{\pi}_J} + \frac{T_b(\hat{\omega}_{\text{MLE}})}{2} \right) \partial_a \hat{\eta}_i \right) \partial^i p(y; \hat{\eta}) \\ &\quad + o_p(n^{-1}) \\ &= p(y; \hat{\eta}) + \frac{\hat{g}^{ab}}{2n} \left(\partial_{ab} \hat{\eta}_i - \bar{\Gamma}_{ab}^c(\hat{\omega}_{\text{MLE}}) \partial_c \hat{\eta}_i \right) \partial^i p(y; \hat{\eta}) \\ &\quad + \frac{\hat{g}^{ab}}{n} \left(\partial_b \log \frac{\hat{\pi}}{\hat{\pi}_J} + \frac{T_b(\hat{\omega}_{\text{MLE}})}{2} \right) \partial_a p(y; \hat{\eta}) + o_p(n^{-1}).\end{aligned}$$

\square

The shift from $p(y; \eta(\hat{\omega}_{\text{MLE}}))$ to $p(y; \hat{\eta}_\pi)$ in Theorem III.2 is composed of two components, one “parallel” and the other “orthogonal” to the model \mathcal{P} . That is, the term

$$\frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{n} \left(\partial_b \log \frac{\pi}{\pi_J}(\hat{\omega}_{\text{MLE}}) + \frac{T_b(\hat{\omega}_{\text{MLE}})}{2} \right) \partial_a p(y; \hat{\omega}_{\text{MLE}})$$

is included in the tangent space spanned by $\partial_a p(y; \eta)$ ($a = 1, \dots, d$) and the term

$$\begin{aligned}&\frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{2n} \left(\partial_a \partial_b \eta_i(\hat{\omega}_{\text{MLE}}) - \bar{\Gamma}_{ab}^c(\hat{\omega}_{\text{MLE}}) \partial_c \eta_i(\hat{\omega}_{\text{MLE}}) \right) \\ &\times \partial^i p(y; \hat{\omega}_{\text{MLE}})\end{aligned}\quad (6)$$

is orthogonal to $\partial_a p(x; \eta)$ ($a = 1, \dots, d$) with respect to the inner product (2), because

$$\begin{aligned}&\left\langle \left(\partial_a \partial_b \eta_i - \bar{\Gamma}_{ab}^c \partial_c \eta_i \right) \partial^i p(y; \eta), \partial_e p(y; \eta) \right\rangle \\ &= \int \partial_a \partial_b \eta_i \frac{\partial p(y; \eta)}{\partial \eta_i} \frac{\partial \theta^j}{\partial \omega^e} \frac{\partial p(y; \eta)}{\partial \theta^j} \frac{1}{p(y; \eta)} dy - \bar{\Gamma}_{ab}^c g_{ce} \\ &= \partial_a \partial_b \eta_i \partial_e \theta^i - \bar{\Gamma}_{abe} \\ &= 0\end{aligned}$$

by using (3) and (4).

We can compare the orthogonal shifts (6) to the orthogonal shifts from $p(y; \eta(\hat{\omega}_{\text{MLE}}))$ to the Bayesian predictive density, and we show (6) is the projection of the orthogonal shifts to the Bayesian predictive density onto \mathcal{E} . In [11], the Bayesian predictive density $p_\pi(y | x)$ is asymptotically expanded as

$$\begin{aligned}p_\pi(y | x) &= p(y; \eta(\hat{\omega}_{\text{MLE}})) + \frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{2n} \\ &\times \left(\partial_a \partial_b p(y; \hat{\omega}_{\text{MLE}}) - \bar{\Gamma}_{ab}^c(\hat{\omega}_{\text{MLE}}) \partial_c p(y; \hat{\omega}_{\text{MLE}}) \right) \\ &+ \frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{n} \\ &\times \left(\partial_b \log \frac{\pi}{\pi_J}(\hat{\omega}_{\text{MLE}}) + \frac{T_b(\hat{\omega}_{\text{MLE}})}{2} \right) \partial_a p(y; \eta(\hat{\omega}_{\text{MLE}})) \\ &+ o_p(n^{-1}).\end{aligned}$$

The parallel shift is identical to that of $p(y; \hat{\eta}_\pi)$. On the other hand, the orthogonal shift

$$\frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{2n} \left(\partial_a \partial_b p(y; \hat{\omega}_{\text{MLE}}) - \bar{\Gamma}_{ab}^c(\hat{\omega}_{\text{MLE}}) \partial_c p(y; \hat{\omega}_{\text{MLE}}) \right) \quad (7)$$

is different from that of $p(y; \hat{\eta}_\pi)$ and it is not included in the tangent space of \mathcal{E} . Therefore, the shifted density $p_\pi(y | x)$ is not included in \mathcal{E} , while $p(y; \hat{\eta}_\pi)$ is in \mathcal{E} .

To cope with the shifts that are orthogonal to \mathcal{P} , we introduce a coordinate system to the subspace of \mathcal{E} that is orthogonal to \mathcal{P} . We divide the tangent vectors of \mathcal{E} at η into two parts, namely, into those parallel to \mathcal{P} and those orthogonal to \mathcal{P} . For each point $\eta \in \mathcal{E}$, the tangent space $T_\eta \mathcal{E}$ is identified with the vector space spanned by

$$\frac{\partial}{\partial \eta_i} p(x; \eta) \quad (i = 1, \dots, m).$$

The tangent space $T_\omega \mathcal{P}$ is a subspace of $T_\eta \mathcal{E}$. Let $A(\omega)$ be an $(m-d)$ -dimensional smooth submanifold of \mathcal{E} attached to each point $\omega \in \mathcal{P}$ and assume that $A(\omega)$ orthogonally transverses \mathcal{P} at $\eta(\omega)$. Such a family of submanifolds $\{A(\omega) \mid \omega \in \Omega\}$ is called an ancillary family. We introduce an adequate

coordinate system $\xi = (\xi^\kappa)$ ($\kappa = d+1, \dots, m$) to $A(\omega)$ so that a pair (ω, ξ) uniquely specifies a point of \mathcal{E} in the neighborhood of $\eta(\omega)$. We adopt a coordinate system ξ on $A(\omega)$ so that $\eta(\omega, \xi) \in \mathcal{P}$ if $\xi = 0$. Then, we have

$$\text{span} \{ \partial^i p(x; \eta) \} = \text{span} \{ \partial_a p(x; \eta), \partial_\kappa p(x; \eta) \}$$

where $\partial_\kappa p(x; \eta) = (\partial / \partial \xi^\kappa) p(x; \eta)$. Since $A(\omega)$ orthogonally transverses \mathcal{E} , we have

$$\langle \partial_a p(x; \eta), \partial_\kappa p(x; \eta) \rangle = 0 \quad (a = 1, \dots, d, \kappa = d+1, \dots, m).$$

Now we show that the extended plugin density with $\hat{\eta}_\pi$ is the projection of the Bayesian predictive density onto \mathcal{E} asymptotically, as shown in Figure 1. The projection of (7) onto the tangent space of \mathcal{E} is

$$\left\langle \frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{2n} (\partial_a \partial_b p(y; \hat{\omega}_{\text{MLE}}) - \bar{\Gamma}_{ab}^c(\hat{\omega}_{\text{MLE}}) \partial_c p(y; \hat{\omega}_{\text{MLE}})), \partial_\lambda p(y; \hat{\omega}_{\text{MLE}}) \right\rangle g^{\lambda\kappa} \partial_\kappa p(y; \hat{\omega}_{\text{MLE}}) \quad (\lambda = d+1, \dots, m). \quad (8)$$

Because $\langle \partial_c p(x; \eta), \partial_\lambda p(x; \eta) \rangle = 0$ ($c = 1, \dots, d$), (8) is

$$\begin{aligned} & \frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{2n} \langle \partial_a \partial_b p(y; \hat{\omega}_{\text{MLE}}), \partial_\lambda p(y; \hat{\omega}_{\text{MLE}}) \rangle \\ & \times g^{\lambda\kappa} \partial_\kappa p(y; \hat{\omega}_{\text{MLE}}) \\ & = \frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{2n} \bar{H}_{ab}^\kappa \partial_\kappa p(y; \hat{\omega}_{\text{MLE}}) \end{aligned}$$

where

$$\begin{aligned} \bar{H}_{ab\kappa}^\kappa &= \bar{\Gamma}_{ab\kappa}^\kappa \\ &= \langle \partial_a \partial_b p(x; \omega), \partial_\kappa p(x; \eta) \rangle \\ &= (\partial_a \partial_b \eta_i)(\partial_\kappa \theta^i) \end{aligned}$$

($a, b = 1, \dots, d$, $\kappa = d+1, \dots, m$) is the mixture embedding curvature of \mathcal{P} in \mathcal{E} and $\bar{H}_{ab}^\kappa = \bar{H}_{ab\lambda} g^{\lambda\kappa}$. Here we represent $\bar{\Gamma}_{ab\kappa}^\kappa$ as $\bar{H}_{ab\kappa}^\kappa$ to clarify its geometrical interpretation as the mixture embedding curvature, distinguishing from $\bar{\Gamma}_{abc}^\kappa$, $\bar{\Gamma}_{a\kappa b}^\kappa$ and $\bar{\Gamma}_{\kappa ab}^\kappa$. We confirm that it coincides with the orthogonal shift to $p(y; \hat{\eta}_\pi)$ asymptotically as follows. Let

$$h_{ab} = (\partial_a \partial_b \eta_i(\omega)) \partial^i p(x; \eta) - \bar{\Gamma}_{ab}^c(\omega) \partial_c p(x; \omega),$$

then the orthogonal component in Theorem III.2 is

$$\frac{g^{ab}(\hat{\omega}_{\text{MLE}})}{2n} h_{ab}(\hat{\omega}_{\text{MLE}}). \quad (9)$$

Since h_{ab} ($a, b = 1, \dots, d$) are included in the space spanned by $\partial_\kappa p(x; \omega)$ ($\kappa = d+1, \dots, m$), the following holds:

$$\begin{aligned} h_{ab} &= \langle h_{ab}, \partial_\lambda p(x; \eta) \rangle g^{\lambda\kappa} \partial_\kappa p(x; \eta) \\ &= \left\langle \partial_a \partial_b \eta_i(\omega) \partial^i p(x; \eta) - \bar{\Gamma}_{ab}^c(\omega) \partial_c p(x; \omega), \partial_\lambda p(x; \eta) \right\rangle \\ &\quad \times g^{\lambda\kappa} \partial_\kappa p(x; \eta). \end{aligned}$$

As $\langle \partial_c p(x; \eta), \partial_\lambda p(x; \eta) \rangle = 0$ ($c = 1, \dots, d$, $\lambda = d+1, \dots, m$),

$$\begin{aligned} h_{ab} &= \langle \partial_a \partial_b \eta_i(\omega) \partial^i p(x; \eta), \partial_\lambda p(x; \omega) \rangle g^{\lambda\kappa} \partial_\kappa p(x; \eta) \\ &= \partial_a \partial_b \eta_i(\omega) \left\langle \partial^i p(x; \omega), \frac{\partial \theta^j}{\partial \xi^\lambda} \frac{\partial p(x; \eta)}{\partial \theta^j} \right\rangle g^{\lambda\kappa} \partial_\kappa p(x; \eta) \end{aligned}$$




Fig. 2. Angle of the projection.

$$\begin{aligned} &= \partial_a \partial_b \eta_i(\omega) \partial_\lambda \theta^i g^{\lambda\kappa} \partial_\kappa p(x; \eta) \\ &= \bar{H}_{ab}^\kappa \partial_\kappa p(x; \eta). \end{aligned} \quad (10)$$

C. Projection Angle and Risk Difference

In this section, we demonstrate that the extended plugin density with $\hat{\eta}_\pi$ is optimal with respect to the risk along the orthogonal shift from the model \mathcal{P} . This property is parallel to those of Bayesian predictive densities investigated in [11], as explained later in this section. Orthogonal shifts from \mathcal{P} can asymptotically improve the Kullback–Leibler risk from plugin densities of \mathcal{P} . By evaluating those risk improvements, the risk difference between $p(y; \hat{\eta}_\pi)$ and Bayesian predictive densities can be corresponded to the projection angle (as represented in Figure 2).

In the following discussion, we consider extended plugin densities $p(y; \hat{\eta})$ with estimators $\hat{\eta} = \eta(\hat{\omega}, \hat{\xi})$ where $\hat{\omega}$ and $\hat{\xi}$ can be expressed in the following forms, respectively:

$$\begin{aligned} \hat{\omega} &= \hat{\omega}_{\text{MLE}} + \frac{1}{n} \alpha(\hat{\omega}_{\text{MLE}}) + o_p(n^{-1}), \\ \hat{\xi} &= \frac{1}{n} \beta(\hat{\omega}_{\text{MLE}}) + o_p(n^{-1}). \end{aligned} \quad (11)$$

Here, $\alpha^a(\omega), \beta^\kappa(\omega)$ are smooth functions of $O_p(1)$. The density can be expanded as

$$\begin{aligned} p(y; \hat{\eta}) &= p_{\alpha, \beta}(y; \hat{\omega}_{\text{MLE}}) \\ &= p(y; \hat{\omega}_{\text{MLE}}) + \frac{1}{n} \alpha^a(\hat{\omega}_{\text{MLE}}) \partial_a p(y; \eta(\hat{\omega}_{\text{MLE}})) \\ &\quad + \frac{1}{n} \beta^\kappa(\hat{\omega}_{\text{MLE}}) \partial_\kappa p(y; \eta(\hat{\omega}_{\text{MLE}})) + o_p(n^{-1}). \end{aligned} \quad (12)$$

The following results hold also for other asymptotically efficient estimators of ω other than $\hat{\omega}_{\text{MLE}}$, although here we consider only $\hat{\omega}_{\text{MLE}}$ for simplicity. This class of extended plugin densities include $p(x; \hat{\omega}_{\text{MLE}})$ and $p(x; \hat{\eta}_\pi)$. For $\hat{\omega} = \hat{\omega}_{\text{MLE}}$ and $\hat{\xi} = 0$, the density is the plugin density with the maximum likelihood estimator $\hat{\omega}_{\text{MLE}}$. The extended plugin density $p(y; \hat{\eta}_\pi)$ in Theorem III.2 is given by, using (9) and (10),

$$\begin{aligned} \alpha^a(\hat{\omega}_{\text{MLE}}) &= g^{ab}(\hat{\omega}_{\text{MLE}}) \left(\partial_b \log \frac{\pi}{\pi_J} + \frac{T_b}{2} \right), \\ \beta^\kappa &= \frac{1}{2} \bar{H}_{ab}^\kappa(\hat{\omega}_{\text{MLE}}) g^{ab}(\hat{\omega}_{\text{MLE}}). \end{aligned}$$

We derive the Kullback–Leibler risk of extended plugin densities in this class.

Proposition III.2: The Kullback–Leibler risk of an extended plugin density $p(y; \hat{\eta})$ where $\hat{\eta} = (\hat{\omega}, \hat{\xi})$ is expressed as (11) is asymptotically expanded as follows:

$$\begin{aligned} E[D\{p(y; \omega), p(y; \hat{\eta})\}] &= \frac{d}{2n} + \frac{1}{2n^2} g_{ab}(\omega) \alpha^a(\omega) \alpha^b(\omega) + \frac{1}{n^2} \nabla_a^e \alpha^a(\omega) \\ &+ \frac{1}{2n^2} g_{\kappa\lambda}(\omega) \beta^\kappa(\omega) \beta^\lambda(\omega) - \frac{1}{2n^2} H_{ab\kappa}^m(\omega) g^{ab}(\omega) \beta^\kappa(\omega) \\ &+ (\text{terms independent of } \alpha, \beta) + o(n^{-2}), \end{aligned} \quad (13)$$

where $\nabla_a^e \alpha^b = \partial_a \alpha^b + \Gamma_{ac}^b \alpha^c$.

Proof: See Appendix B. \square

The risk expansion (13) confirms that the risk can be improved from a predictive density in \mathcal{P} by selecting an appropriate orthogonal shift β . We obtain the optimal orthogonal shift.

Theorem III.3: The optimal β^κ in (12) is given by

$$\beta_{\text{opt}}^\kappa(\omega) = \frac{1}{2} \frac{m}{H_{ab}^\kappa}(\omega) g^{ab}(\omega). \quad (14)$$

Proof: The risk in Proposition III.2 is

$$\begin{aligned} E[D\{p(y; \omega), p(y; \hat{\eta})\}] &= \frac{1}{2n^2} g_{ab} \alpha^a \alpha^b + \frac{1}{n^2} \nabla_a^e \alpha^a(\omega) \\ &+ \frac{1}{2n^2} g_{\kappa\lambda} \left(\beta^\lambda - \frac{1}{2} \frac{m}{H_{ab}^\lambda} g^{ab} \right) \left(\beta^\kappa - \frac{1}{2} \frac{m}{H_{cd}^\kappa} g^{cd} \right) \\ &- \frac{1}{8n^2} H_{ab}^\lambda H_{cd}^\kappa g^{ab} g^{cd} g_{\kappa\lambda} + (\text{terms independent of } \alpha, \beta) \\ &+ o(n^{-2}). \end{aligned}$$

Therefore, β is optimal when

$$\beta^\kappa(\omega) = \frac{1}{2} \frac{m}{H_{ab}^\kappa}(\omega) g^{ab}(\omega). \quad \square$$

Accordingly, the orthogonal component of the shift in Theorem III.2 is optimal, and the extended plugin density with $\hat{\eta}_\pi$ has the optimal shift (as illustrated in Figure 3). The risk improvement by the optimal shift is evaluated by the inner product of the optimal shifts $\beta_{\text{opt}}^\kappa$ and given by

$$\frac{1}{8n^2} H_{ab}^\lambda H_{cd}^\kappa g^{ab} g^{cd} g_{\kappa\lambda} + o(n^{-2}),$$

which does not depend on the parallel shift α . Here, $\frac{m}{H_{ab}^\lambda} H_{cd}^\kappa g^{ab} g^{cd} g_{\kappa\lambda}$ is the mixture mean curvature of \mathcal{P} embedded in \mathcal{E} at ω , thus this risk improvement has a geometrical interpretation.

These results are related to the properties of Bayesian predictive densities. In [11], it is shown that Bayesian predictive densities are optimal along orthogonal shift from \mathcal{P} . The orthogonal shift is not included in the tangent space of \mathcal{E} , as explained in Section III-B. The risk improvement achieved by the orthogonal shift is the mixture mean curvature of \mathcal{P} embedded in \mathcal{F} , while the risk improvement of $p(y; \hat{\eta}_\pi)$ is the mixture mean curvature of \mathcal{P} embedded in \mathcal{E} . The risk improvement is evaluated by inner products of the optimal shifts, and as a result the cosine of the angle between the two




Fig. 3. Extended plugin density with $\hat{\eta}_\pi$ has the optimal orthogonal shift.

orthogonal shifts, as shown in Figure 2, can be found as the square root of the ratio of the two risk improvements.

Example (Fisher circle model, continued) We have $g_{\omega\omega} = 1$ and $\Gamma_{\omega\omega}^m = 0$. Thus, the optimal orthogonal shift is

$$\begin{aligned} h_{\omega\omega}(x; \eta) &= \left(\partial_{\omega\omega} \eta_i - \Gamma_{\omega\omega}^m \partial_{\omega} \eta_i \right) \partial_i p(x; \eta) \\ &= p(x; \eta) (-(x_1 - \eta_1) \cos \omega - (x_2 - \eta_2) \sin \omega) \end{aligned}$$

and the risk improvement obtained by the optimal orthogonal shift is

$$\begin{aligned} &\frac{1}{8n^2} H_{ab}^\lambda H_{cd}^\kappa g^{ab} g^{cd} g_{\kappa\lambda} + o(n^{-2}) \\ &= \frac{1}{8n^2} E[(-(y_1 - \eta_1) \cos \omega - (y_2 - \eta_2) \sin \omega)^2] + o(n^{-2}) \\ &= \frac{1}{8n^2} + o(n^{-2}). \end{aligned}$$

The risk improvement corresponding to the optimal shift (7) is $3/(8n^2) + o(n^{-2})$.

If the variance of x_1, x_2 is σ^2 , the risk improvements corresponding to the optimal orthogonal shift and to the shift (7) can be obtained as follows, respectively:

$$\frac{\sigma^2}{8n^2}, \quad \frac{\sigma^2 + 2}{8n^2}.$$

Therefore, when σ^2 is large, the risk improvement obtained by $p(y; \hat{\eta}_\pi)$ becomes relatively significant as well, and the performance of the Bayes extended estimator is close to that of the Bayesian predictive density. The cosine of the angle between the two shift vectors is

$$\sqrt{\frac{\sigma^2}{8n^2}} / \sqrt{\frac{\sigma^2 + 2}{8n^2}} = \sqrt{\frac{\sigma^2}{\sigma^2 + 2}},$$

which approaches 0 as σ^2 increases. In this way, the Bayesian predictive density and $p(y; \hat{\eta}_\pi)$ approach each other.

IV. NUMERICAL STUDIES

The numerical simulations of the Kullback–Leibler risk of $p(y; \hat{\eta}_\pi)$ and Bayesian predictive densities are shown in a curved Gaussian model. It confirmed the theoretical results so far and also illustrates the practical importance of the projection of Bayesian predictive densities.

The model \mathcal{P} is the spiked covariance model (for more information about related models, see e.g. [12]), that is

l -dimensional Gaussian $N_l(0, \Sigma)$ with the covariance matrix Σ expressed as

$$\Sigma(\lambda, \mathbf{u}) = \lambda \mathbf{u} \mathbf{u}^\top + I_l, \quad (15)$$

where the vector $\mathbf{u} \in \mathbb{R}^l$ satisfies $\mathbf{u}^\top \mathbf{u} = 1$ and $\lambda > 0$. The eigenvalues of the matrix Σ are $\lambda + 1, 1, \dots, 1$, and \mathbf{u} is the first eigenvector.

The model \mathcal{P} parametrized by $\omega = (\mathbf{u}, \lambda)$ is embedded in the larger full exponential family $\mathcal{E} = \{N_l(0, \Sigma) \mid \Sigma\}$, and the expectation parameter η comprises the components of Σ . The extended plugin distribution $p(y; \hat{\eta}_\pi)$ is $N_l(0, \hat{\Sigma}_\pi)$, where $\hat{\Sigma}_\pi$ is the posterior mean of Σ . The Bayes estimator $(\hat{\lambda}_\pi, \hat{\mathbf{u}}_\pi)$ of \mathcal{P} is composed of $\hat{\lambda}_\pi$, the posterior mean of λ , and $\hat{\mathbf{u}}_\pi$, the first eigen vector of $\hat{\Sigma}_\pi$. The plugin distribution of the Bayes estimator is $N_l(0, \hat{\lambda}_\pi \hat{\mathbf{u}}_\pi \hat{\mathbf{u}}_\pi^\top + I_l)$. The settings of l and n are $(l, n) = (5, 20)$ and $(80, 320)$. The posterior mean of λ and Σ is computed by the 1000 MCMC samples for $l = 5$ and 2000 MCMC samples for $l = 80$ produced by Gibbs sampling with 250 and 500 burn-in samples, respectively. The Bayesian predictive density is computed by taking the mean of the plugin densities of those MCMC samples of (λ, \mathbf{u}) . The Kullback–Leibler risk is derived as the mean of 2000 trials.

The results are illustrated in Figure 4, and it is confirmed that $p(y; \hat{\eta}_\pi)$ approaches the Bayesian predictive density as the size of \mathcal{E} increases. In Figure 4a, the three-layer structure of the plugin density, the extended plugin density, and the Bayesian predictive density is seen in the risk comparison. On the other hand, in Figure 4b, it could be seen that the risk plots of $p(y; \hat{\eta}_\pi)$ and the Bayesian predictive density are quite close, which means the projection angle between them is close to zero. The dimension of the parameter space of \mathcal{P} is l and that of \mathcal{E} is $l(l+1)/2$; thus, \mathcal{P} is embedded in relatively larger full exponential families when l increases. Therefore, it is natural that the two risk performances approach as l increases because the extended model \mathcal{E} approaches the set of all probability densities \mathcal{F} , and $p(y; \hat{\eta}_\pi)$ approaches the Bayesian predictive density.

Figure 4b also illustrates the practical advantage of $p(y; \hat{\eta}_\pi)$, as it shows that projection of Bayesian predictive densities onto finite-dimensional models of reasonable size is an effective way of approximation. Bayesian predictive densities are typically approximated by the mean of plugin densities, because obtaining the full density function is intractable. In these experiments, the Bayesian predictive density is the mean of 2000 plugin densities. The problem of this approximation is that it requires large space and time to compute density, because storing all MCMC samples is necessary, and we have to take the mean of these plugin densities for each y . Figure 4b demonstrates that the approximation by $p(y; \hat{\eta}_\pi)$, a single point in \mathcal{E} is comparable to the mean of 2000 points in \mathcal{P} . Approximation by $p(y; \hat{\eta}_\pi)$ does not require storing MCMC samples, and the full density function is available avoiding taking mean of plugin densities every time. Table II provides the computation time required to evaluate the density of 1000 new samples y and the memory size of the Bayesian predictive density and $p(y; \hat{\eta}_\pi)$. It is shown that both computational time and memory size are effectively saved when we utilize $p(y; \hat{\eta}_\pi)$. Therefore, the extended plugin density $p(y; \hat{\eta}_\pi)$ is an

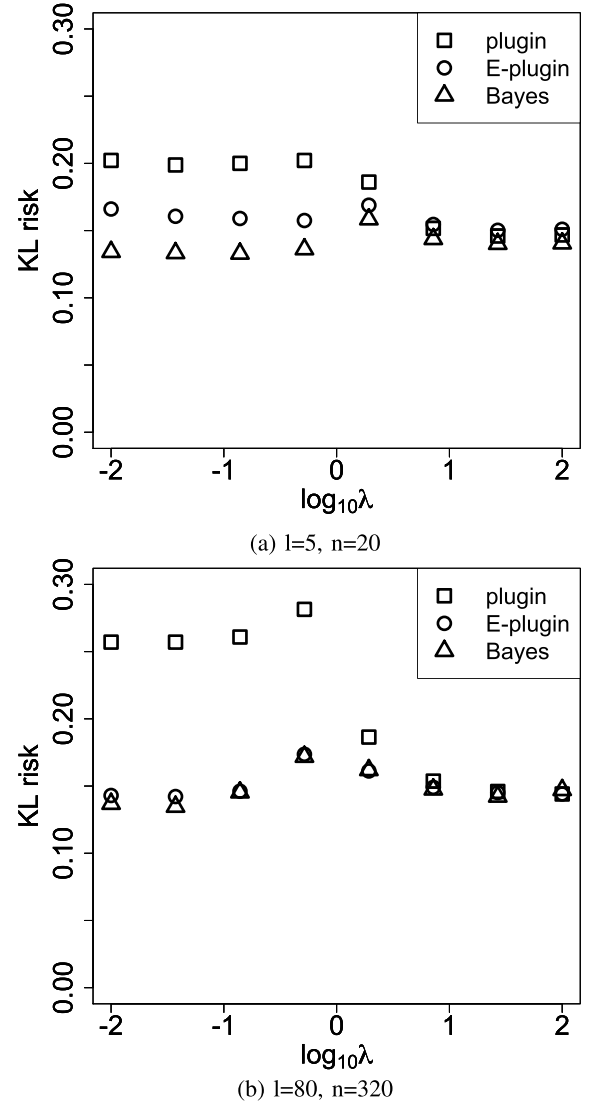


Fig. 4. Kullback–Leibler (KL) risk of plugin density with Bayes estimator, extended plugin density (E-plugin) with Bayes extended estimator, and Bayesian predictive density.

TABLE II
COMPUTATIONAL TIME REQUIRED TO EVALUATE PREDICTIVE DENSITY FOR 1000 NEW SAMPLES AND MEMORY SIZE OF EACH PREDICTIVE DENSITY ($l = 80, \lambda = 1$)

	Bayes	E-plugin
computational time	1.24×10^2 s	1.6×10^{-2} s
predictive density size	1.29×10^6 bytes	5.14×10^4 bytes

effective approximation of Bayesian predictive densities, with a large advantage in terms of computational costs, which has a natural theoretical interpretation as a projection onto a finite dimensional model.

APPENDIX A PROOF OF THEOREM III.1

Let $L(\eta) = \frac{1}{n} \sum_{t=1}^n \log p(x(t); \eta)$. Then, $\hat{\eta}_\pi(x^n) = \{(\hat{\eta}_\pi)_i(x^n)\}$ is given by

$$\begin{aligned} \hat{\eta}_\pi(x^n) &= \frac{\int \eta(\omega) p(x^n; \eta(\omega)) \pi(\omega) d\omega}{\int p(x^n; \eta(\omega)) \pi(\omega) d\omega} \\ &= \frac{\int \eta(\omega) \exp(nL(\eta(\omega))) \pi(\omega) d\omega}{\int \exp(nL(\eta(\omega))) \pi(\omega) d\omega}. \end{aligned}$$

We approximate $\hat{\eta}_\pi$ using the Laplace method (e.g., [13], Sec. 3.6 and Sec. 4.6, and [14]). The proposed approximation follows Theorem 4.6.1 in [13]. First we expand $\pi(\omega) \exp(nL(\omega))$ around $\omega = \hat{\omega}_{\text{MLE}}$. In the following, symbols such as $\eta(\hat{\omega}_{\text{MLE}})$, $\partial_a \eta(\hat{\omega}_{\text{MLE}})$, and $\partial_a \partial_b \eta(\hat{\omega}_{\text{MLE}})$ are abbreviated to $\hat{\eta}$, $\partial_a \hat{\eta}$, and $\partial_{ab} \hat{\eta}$ respectively. Let $\omega = \hat{\omega}_{\text{MLE}} + \phi/\sqrt{n}$. Then,

$$\begin{aligned} & \pi(\omega) \exp(nL(\omega)) \\ &= \left(\hat{\pi} + \frac{(\partial_a \hat{\pi}) \phi^a}{\sqrt{n}} + \frac{(\partial_{ab} \hat{\pi}) \phi^a \phi^b}{2n} + \frac{(\partial_{abc} \hat{\pi}) \phi^a \phi^b \phi^c}{6n\sqrt{n}} \right. \\ & \quad + O_p(n^{-2}) \Big) \exp \left(n\hat{L} + \frac{(\partial_{ab} \hat{L}) \phi^a \phi^b}{2} + \frac{(\partial_{abc} \hat{L}) \phi^a \phi^b \phi^c}{6\sqrt{n}} \right. \\ & \quad + \frac{(\partial_{abcd} \hat{L}) \phi^a \phi^b \phi^c \phi^d}{24n} + \frac{(\partial_{abcde} \hat{L}) \phi^a \phi^b \phi^c \phi^d \phi^e}{120n\sqrt{n}} \\ & \quad + O_p(n^{-2}) \Big) \\ &= \hat{\pi} e^{n\hat{L}} e^{(\partial_{ab} \hat{L}) \phi^a \phi^b / 2} \left(1 + \frac{(\partial_a \hat{\pi}) \phi^a}{\hat{\pi} \sqrt{n}} + \frac{(\partial_{ab} \hat{\pi}) \phi^a \phi^b}{2\hat{\pi} n} \right. \\ & \quad + \frac{(\partial_{abc} \hat{\pi}) \phi^a \phi^b \phi^c}{6\hat{\pi} n \sqrt{n}} + O_p(n^{-2}) \Big) \left(1 + \frac{(\partial_{abc} \hat{L}) \phi^a \phi^b \phi^c}{6\sqrt{n}} \right. \\ & \quad + \frac{(\partial_{abcd} \hat{L}) \phi^a \phi^b \phi^c \phi^d}{24n} \\ & \quad + \frac{(\partial_{abc} \hat{L}) (\partial_{a'b'c'} \hat{L}) \phi^a \phi^b \phi^c \phi^{a'} \phi^{b'} \phi^{c'}}{72n} + O_p(n^{-3/2}) \Big) \\ &= \hat{\pi} e^{n\hat{L}} e^{-\hat{J}_{ab} \phi^a \phi^b / 2} \left(1 + \frac{(\partial_a \hat{\pi}) \phi^a}{\hat{\pi} \sqrt{n}} + \frac{(\partial_{abc} \hat{L}) \phi^a \phi^b \phi^c}{6\sqrt{n}} \right. \\ & \quad + \frac{(\partial_{ab} \hat{\pi}) \phi^a \phi^b}{2\hat{\pi} n} + \frac{(\partial_{abc} \hat{\pi}) (\partial_{bcd} \hat{L}) \phi^a \phi^b \phi^c \phi^d}{6\hat{\pi} n} \\ & \quad + \frac{(\partial_{abc} \hat{\pi}) \phi^a \phi^b \phi^c}{6\hat{\pi} n \sqrt{n}} + \frac{(\partial_{ab} \hat{\pi}) (\partial_{cde} \hat{L}) \phi^a \phi^b \phi^c \phi^d \phi^e}{12\hat{\pi} n \sqrt{n}} \\ & \quad + \frac{(\partial_a \hat{\pi}) (\partial_{bcde} \hat{L}) \phi^a \phi^b \phi^c \phi^d \phi^e}{24\hat{\pi} n \sqrt{n}} \\ & \quad + \frac{(\partial_a \hat{\pi}) (\partial_{bcd} \hat{L}) (\partial_{efg} \hat{L}) \phi^a \phi^b \phi^c \phi^d \phi^e \phi^f \phi^g}{72\hat{\pi} n \sqrt{n}} \\ & \quad + \frac{C_1}{n} + O_p(n^{-2}) \Big). \end{aligned}$$

Here, $\hat{J}_{ab} = -\partial_{ab} \hat{L}$ and we denote the terms of $O_p(n^{-1})$ which do not depend on π as C_1/n , which is $O_p(n^{-1})$. Let (\hat{J}^{ab}) be the inverse matrix of (\hat{J}_{ab}) . Next, we integrate both sides of the above equation with respect to ω . By changing the variables from ω to ϕ , and by using the formula of moments of multivariate Gaussian distributions, we obtain

$$\begin{aligned} & \int \pi(\omega) \exp(nL(\omega)) d\omega \\ &= C_2 \hat{\pi} \left(1 + \frac{(\partial_{ab} \hat{\pi})}{2\hat{\pi} n} \int \phi^a \phi^b e^{-\hat{J}_{cd} \phi^c \phi^d / 2} d\phi \right. \\ & \quad + \frac{(\partial_{abc} \hat{\pi}) (\partial_{bcd} \hat{L})}{6\hat{\pi} n} \int \phi^a \phi^b \phi^c \phi^d e^{-\hat{J}_{ef} \phi^e \phi^f / 2} d\phi + \frac{C_1}{n} \\ & \quad + O_p(n^{-2}) \Big) \end{aligned}$$

$$\begin{aligned} &= C_2 \hat{\pi} \left(1 + \frac{(\partial_{ab} \hat{\pi}) \hat{J}^{ab}}{2\hat{\pi} n} \right. \\ & \quad + \frac{(\partial_{abc} \hat{\pi}) (\partial_{bcd} \hat{L}) (\hat{J}^{ab} \hat{J}^{cd} + \hat{J}^{ac} \hat{J}^{bd} + \hat{J}^{ad} \hat{J}^{bc})}{6\hat{\pi} n} + \frac{C_1}{n} \\ & \quad + O_p(n^{-2}) \Big) \\ &= C_2 \hat{\pi} \left(1 + \frac{(\partial_{ab} \hat{\pi}) \hat{J}^{ab}}{2\hat{\pi} n} + \frac{(\partial_{abc} \hat{\pi}) (\partial_{bcd} \hat{L}) \hat{J}^{ab} \hat{J}^{cd}}{2\hat{\pi} n} + \frac{C_1}{n} \right. \\ & \quad + O_p(n^{-2}) \Big). \end{aligned}$$

Here C_2 does not depend on π . Replace $\pi(\omega)$ by $\eta_i(\omega)\pi(\omega)$ and we have

$$\begin{aligned} & \int \eta_i(\omega) \pi(\omega) \exp(nL(\omega)) d\omega \\ &= C_2 \hat{\eta}_i \hat{\pi} \left(1 + \frac{\{\partial_{ab}(\hat{\eta}_i \hat{\pi})\} \hat{J}^{ab}}{2\hat{\eta}_i \hat{\pi} n} + \frac{\{\partial_a(\hat{\eta}_i \hat{\pi})\} (\partial_{bcd} \hat{L}) \hat{J}^{ab} \hat{J}^{cd}}{2\hat{\eta}_i \hat{\pi} n} \right. \\ & \quad + \frac{C_1}{n} + O_p(n^{-2}) \Big). \end{aligned}$$

Therefore, the posterior mean of η_i is expanded as

$$\begin{aligned} & (\hat{\eta}_\pi)_i \\ &= \frac{\int \eta_i \exp(nL(\omega)) \pi(\omega) d\omega}{\int \exp(nL(\omega)) \pi(\omega) d\omega} \\ &= C_2 \hat{\eta}_i \hat{\pi} \left(1 + \frac{\partial_{ab}(\hat{\eta}_i \hat{\pi}) \hat{J}^{ab}}{2\hat{\eta}_i \hat{\pi} n} \right. \\ & \quad + \frac{\{\partial_a(\hat{\eta}_i \hat{\pi})\} (\partial_{bcd} \hat{L}) \hat{J}^{ab} \hat{J}^{cd}}{2\hat{\eta}_i \hat{\pi} n} + \frac{C_1}{n} + O_p(n^{-2}) \Big) \\ & \quad \Big/ C_2 \hat{\pi} \left(1 + \frac{(\partial_{ab} \hat{\pi}) \hat{J}^{ab}}{2\hat{\pi} n} + \frac{(\partial_{abc} \hat{\pi}) (\partial_{bcd} \hat{L}) \hat{J}^{ab} \hat{J}^{cd}}{2\hat{\pi} n} \right. \\ & \quad + \frac{C_1}{n} + O_p(n^{-2}) \Big) \\ &= \hat{\eta}_i \left(1 + \frac{\hat{J}^{ab}}{2n} \left(\frac{\partial_{ab}(\hat{\eta}_i \hat{\pi})}{\hat{\eta}_i \hat{\pi}} - \frac{\partial_{ab} \hat{\pi}}{\hat{\pi}} \right) \right. \\ & \quad + \frac{\hat{J}^{ab} \hat{J}^{cd} \partial_{bcd} \hat{L}}{2n} \left(\frac{\partial_a(\hat{\eta}_i \hat{\pi})}{\hat{\eta}_i \hat{\pi}} - \frac{\partial_a \hat{\pi}}{\hat{\pi}} \right) + O_p(n^{-2}) \Big) \\ &= \hat{\eta}_i + \frac{\hat{J}^{ab}}{2n} \left(\partial_{ab} \hat{\eta}_i + \frac{2(\partial_a \hat{\eta}_i)(\partial_b \hat{\pi})}{\hat{\pi}} \right) + \frac{\hat{J}^{ab} \hat{J}^{cd} \partial_{bcd} \hat{L}}{2n} \partial_a \hat{\eta}_i \\ & \quad + O_p(n^{-2}). \end{aligned}$$

Since $\hat{J}_{ab} = \hat{g}_{ab} + o_p(1)$, $\partial_{bcd} \hat{L} = E(\partial_{bcd} \hat{L}) + o_p(1)$ by the law of large numbers, and

$$E(\partial_{bcd} L) = -\partial_b g_{cd} - \overset{\text{e}}{\Gamma}_{cdb} = -\partial_b g_{cd} - \overset{\text{m}}{\Gamma}_{cdb} + T_{bcd},$$

and

$$g^{cd} \partial_b g_{cd} = \partial_b \log(|g|) = 2\partial_b \log \pi_J$$

hold, we obtain

$$\begin{aligned} & (\hat{\eta}_\pi)_i \\ &= \hat{\eta}_i + \frac{\hat{g}^{ab}}{2n} (\partial_{ab} \hat{\eta}_i + 2\partial_a \hat{\eta}_i \partial_b \log \hat{\pi}) \end{aligned}$$

$$\begin{aligned}
& + \frac{\hat{g}^{ab}}{2n} \left(-2\partial_b \log \hat{\pi}_J - \hat{g}^{cd} \Gamma_{cdb}^m(\hat{\omega}_{\text{MLE}}) + T_b(\hat{\omega}_{\text{MLE}}) \right) \partial_a \hat{\eta}_i \\
& + o_p(n^{-1}) \\
= & \hat{\eta}_i + \frac{\hat{g}^{ab}}{2n} \left(\partial_{ab} \hat{\eta}_i - \Gamma_{ab}^c(\hat{\omega}_{\text{MLE}}) \partial_c \hat{\eta}_i \right) \\
& + \frac{\hat{g}^{ab}}{n} \left(\partial_b \log \frac{\hat{\pi}}{\hat{\pi}_J} + \frac{T_b(\hat{\omega}_{\text{MLE}})}{2} \right) \partial_a \hat{\eta}_i + o_p(n^{-1}).
\end{aligned}$$

APPENDIX B PROOF OF PROPOSITION III.2

We abbreviate symbols such as $g_{ab}(\omega)$, $\bar{H}_{ab\kappa}^m(\omega)$, and $p(y; \omega)$ to g_{ab} , $\bar{H}_{ab\kappa}^m$, and p , respectively.

The Kullback–Leibler divergence from $p(y; \omega)$ to $p_{\alpha, \beta}(y; \hat{\omega}_{\text{MLE}})$ is expanded as

$$\begin{aligned}
& D(p(y; \omega), p_{\alpha, \beta}(y; \hat{\omega}_{\text{MLE}})) \\
& = D(p, p(y; \hat{\omega}, \hat{\xi})) \\
& = - \int p \{ \log p(y; \hat{\omega}, \hat{\xi}) - \log p \} dy,
\end{aligned}$$

and we can expand $\log p(y; \hat{\omega}, \hat{\xi}) - \log p$ as follows:

$$\begin{aligned}
& \log p(y; \hat{\omega}, \hat{\xi}) - \log p \\
& = \tilde{\omega}^a \partial_a \log p + \hat{\xi}^\kappa \partial_\kappa \log p + \frac{1}{2} \tilde{\omega}^a \tilde{\omega}^b \partial_{ab} \log p \\
& + \tilde{\omega}^a \hat{\xi}^\kappa \partial_{a\kappa} \log p + \frac{1}{2} \hat{\xi}^\kappa \hat{\xi}^\lambda \partial_{\kappa\lambda} \log p + \frac{1}{6} \tilde{\omega}^a \tilde{\omega}^b \tilde{\omega}^c \partial_{abc} \log p \\
& + \frac{1}{2} \tilde{\omega}^a \tilde{\omega}^b \hat{\xi}^\kappa \partial_{ab\kappa} \log p + \frac{1}{24} \tilde{\omega}^a \tilde{\omega}^b \tilde{\omega}^c \tilde{\omega}^d \partial_{abcd} \log p \\
& + o_p(n^{-2})
\end{aligned}$$

where $\tilde{\omega} = \hat{\omega} - \omega$. Because $E[\partial_a \log p] = 0$, $E[-\partial_{ab} \log p] = g_{ab}$, and $g_{a\kappa} = 0$ for $a = 1, \dots, d$, $\kappa = d+1, \dots, m$, we have

$$\begin{aligned}
& D(p(y; \omega), p_{\alpha, \beta}(y; \hat{\omega}_{\text{MLE}})) \\
& = \frac{1}{2} g_{ab} \tilde{\omega}^a \tilde{\omega}^b + \frac{1}{2} g_{\kappa\lambda} \hat{\xi}^\kappa \hat{\xi}^\lambda - \frac{1}{6} \tilde{\omega}^a \tilde{\omega}^b \tilde{\omega}^c E[\partial_{abc} \log p] \\
& - \frac{1}{2} \tilde{\omega}^a \tilde{\omega}^b \hat{\xi}^\kappa E[\partial_{ab\kappa} \log p] - \frac{1}{24} \tilde{\omega}^a \tilde{\omega}^b \tilde{\omega}^c \tilde{\omega}^d E[\partial_{abcd} \log p] \\
& + o_p(n^{-2}) \\
& = \frac{1}{2} g_{ab} \tilde{\omega}^a \tilde{\omega}^b + \frac{1}{2} g_{\kappa\lambda} \hat{\xi}^\kappa \hat{\xi}^\lambda + \left(\frac{1}{2} \Gamma_{abc}^m - \frac{1}{3} T_{abc} \right) \tilde{\omega}^a \tilde{\omega}^b \tilde{\omega}^c \\
& + \left\{ \frac{1}{2} (\Gamma_{ab\kappa}^m + \Gamma_{a\kappa b}^m + \Gamma_{\kappa ab}^m) - T_{ab\kappa} \right\} \tilde{\omega}_{\text{MLE}}^a \tilde{\omega}_{\text{MLE}}^b \hat{\xi}^\kappa \\
& + K_{abcd} \tilde{\omega}_{\text{MLE}}^a \tilde{\omega}_{\text{MLE}}^b \tilde{\omega}_{\text{MLE}}^c \tilde{\omega}_{\text{MLE}}^d + o_p(n^{-2}),
\end{aligned}$$

where $\tilde{\omega}_{\text{MLE}} = \hat{\omega}_{\text{MLE}} - \omega$ and

$$\begin{aligned}
& K_{abcd} \\
& = \frac{1}{24} \int \left\{ 6 \left(\frac{\partial_a p}{p} \frac{\partial_b p}{p} \frac{\partial_c p}{p} \frac{\partial_d p}{p} \right) - 12 \left(\frac{\partial_a p}{p} \frac{\partial_b p}{p} \frac{\partial_c \partial_d p}{p} \right) \right. \\
& \quad \left. + 3 \left(\frac{\partial_a \partial_b p}{p} \frac{\partial_c \partial_d p}{p} \right) + 4 \left(\frac{\partial_a p}{p} \frac{\partial_b \partial_c \partial_d p}{p} \right) \right\} p dy.
\end{aligned}$$

Therefore, the Kullback–Leibler risk from $p(y; \omega)$ to $p_{\alpha, \beta}(y; \hat{\omega}_{\text{MLE}})$ is expanded as

$$E[D\{p(y; \omega), p_{\alpha, \beta}(y; \hat{\omega}_{\text{MLE}})\}]$$

$$\begin{aligned}
& = \frac{1}{2} g_{ab} E(\tilde{\omega}^a \tilde{\omega}^b) + \frac{1}{2n^2} g_{\kappa\lambda} E(\hat{\beta}^\kappa \hat{\beta}^\lambda) \\
& + \left(\frac{1}{2} \Gamma_{abc}^m - \frac{1}{3} T_{abc} \right) E(\tilde{\omega}^a \tilde{\omega}^b \tilde{\omega}^c) \\
& + \frac{1}{n} \left\{ \frac{1}{2} (\Gamma_{ab\kappa}^m + \Gamma_{a\kappa b}^m + \Gamma_{\kappa ab}^m) - T_{ab\kappa} \right\} E(\tilde{\omega}_{\text{MLE}}^a \tilde{\omega}_{\text{MLE}}^b \hat{\beta}^\kappa) \\
& + (\text{terms independent of } \alpha, \beta) + o(n^{-2}).
\end{aligned}$$

Because $\hat{\beta}$ is a smooth function of $O_p(1)$ and $\hat{\beta} = \beta + o_p(1)$,

$$E(\hat{\beta}^\kappa \hat{\beta}^\lambda) = \beta^\kappa \beta^\lambda + o(1),$$

$$E(\tilde{\omega}_{\text{MLE}}^a \tilde{\omega}_{\text{MLE}}^b \hat{\beta}^\kappa) = \frac{1}{n} g^{ab} \beta^\kappa + o(n^{-1})$$

hold and

$$\begin{aligned}
& E(\tilde{\omega}^a \tilde{\omega}^b \tilde{\omega}^c) \\
& = (\alpha^a E(\tilde{\omega}_{\text{MLE}}^b \tilde{\omega}_{\text{MLE}}^c) + \alpha^b E(\tilde{\omega}_{\text{MLE}}^c \tilde{\omega}_{\text{MLE}}^a) \\
& + \alpha^c E(\tilde{\omega}_{\text{MLE}}^a \tilde{\omega}_{\text{MLE}}^b)) / n + (\text{terms independent of } \alpha, \beta) \\
& + o(n^{-2}) \\
& = (\alpha^a g^{bc} + \alpha^b g^{ca} + \alpha^c g^{ab}) / n^2 \\
& + (\text{terms independent of } \alpha, \beta) + o(n^{-2}).
\end{aligned}$$

Lastly we derive $E(\tilde{\omega}^a \tilde{\omega}^b)$. It can be expanded as (see e.g. (10.19) in [15], as its multivariate version is presented here)

$$\begin{aligned}
& E(\tilde{\omega}^a \tilde{\omega}^b) = \frac{1}{n} g^{ab} + \frac{1}{n} (g^{ca} \partial_c E(\tilde{\omega}^b) + g^{cb} \partial_c E(\tilde{\omega}^a)) \\
& + E\{(\tilde{\omega}^a - g^{ac} \partial_c L)(\tilde{\omega}^b - g^{bd} \partial_d L)\}.
\end{aligned}$$

From the likelihood equation,

$$\begin{aligned}
& \partial_a L(\hat{\omega}_{\text{MLE}}) = 0, \\
& \partial_a L + \tilde{\omega}_{\text{MLE}}^b \partial_{ab} L + \frac{1}{2} \tilde{\omega}_{\text{MLE}}^c \tilde{\omega}_{\text{MLE}}^d \partial_{acd} L + o_p(n^{-1}) = 0, \\
& g_{ab} \tilde{\omega}_{\text{MLE}}^b = \partial_a L + \tilde{\omega}_{\text{MLE}}^b (\partial_{ab} L + g_{ab}) \\
& + \frac{1}{2} \tilde{\omega}_{\text{MLE}}^c \tilde{\omega}_{\text{MLE}}^d \partial_{acd} L + o_p(n^{-1}).
\end{aligned}$$

Because $\tilde{\omega}_{\text{MLE}}^b = g^{ab} \partial_a L + O_p(n^{-1})$ and $\tilde{\omega}_{\text{MLE}} = \hat{\omega}_{\text{MLE}} - \omega = \hat{\omega} - \alpha/n - \omega + o_p(n^{-1})$,

$$\begin{aligned}
& g_{ab} \tilde{\omega}_{\text{MLE}}^b = \partial_a L + g^{cd} \partial_d L (\partial_{ac} L + g_{ac}) \\
& + \frac{1}{2} \tilde{\omega}_{\text{MLE}}^c \tilde{\omega}_{\text{MLE}}^d \partial_{acd} L + o_p(n^{-1}), \\
& \tilde{\omega}^b = \frac{\alpha^b}{n} + g^{ab} \partial_a L + g^{ab} g^{cd} \partial_d L (\partial_{ac} L + g_{ac}) \\
& + \frac{1}{2} g^{ab} \tilde{\omega}_{\text{MLE}}^c \tilde{\omega}_{\text{MLE}}^d \partial_{acd} L + o_p(n^{-1}).
\end{aligned}$$

Thus

$$E(\tilde{\omega}^b) = \frac{\alpha^b}{n} + (\text{terms independent of } \alpha, \beta) + o(n^{-1})$$

and

$$\begin{aligned}
& E\{(\tilde{\omega}^a - g^{ac} \partial_c L)(\tilde{\omega}^b - g^{bd} \partial_d L)\} \\
& = \frac{\alpha^a \alpha^b}{n^2} + \frac{\alpha^a}{n} \times E[g^{be} g^{cd} \partial_d L (\partial_{ce} L + g_{ce}) \\
& + \frac{1}{2} g^{be} \tilde{\omega}_{\text{MLE}}^c \tilde{\omega}_{\text{MLE}}^d \partial_{cde} L] \\
& + \frac{\alpha^b}{n} E[g^{ae} g^{cd} \partial_d L (\partial_{ce} L + g_{ce})
\end{aligned}$$

$$+ \frac{1}{2} g^{ae} \tilde{\omega}_{MLE}^c \tilde{\omega}_{MLE}^d \partial_{cde} L \Big] \\ + (\text{terms independent of } \alpha, \beta) + o(n^{-2}).$$

Using $E[\partial_d L \partial_{ce} L] = \tilde{\Gamma}_{ecd}^e/n$ and $E[\partial_{cde} L] = -\tilde{\Gamma}_{cde}^e - \tilde{\Gamma}_{dec}^e - \tilde{\Gamma}_{ecd}^e - T_{ecd}$, we have

$$\begin{aligned} & E\{(\tilde{\omega}^a - g^{ac} \partial_c L)(\tilde{\omega}^b - g^{bd} \partial_d L)\} \\ &= \frac{\alpha^a \alpha^b}{n^2} + \frac{\alpha^a}{n^2} E \left[g^{be} g^{cd} \tilde{\Gamma}_{ecd}^e \right. \\ & \quad \left. + \frac{1}{2} g^{be} g^{cd} (-\tilde{\Gamma}_{cde}^e - \tilde{\Gamma}_{dec}^e - \tilde{\Gamma}_{ecd}^e - T_{ecd}) \right] \\ & \quad + \frac{\alpha^b}{n^2} E \left[g^{ae} g^{cd} \tilde{\Gamma}_{ecd}^e \right. \\ & \quad \left. + \frac{1}{2} g^{ae} g^{cd} (-\tilde{\Gamma}_{cde}^e - \tilde{\Gamma}_{dec}^e - \tilde{\Gamma}_{ecd}^e - T_{ecd}) \right] \\ & \quad + (\text{terms independent of } \alpha, \beta) + o(n^{-2}) \\ &= \frac{\alpha^a \alpha^b}{n^2} - \frac{\alpha^a g^{be} g^{cd} \tilde{\Gamma}_{cde}^m}{2n^2} - \frac{\alpha^b g^{ae} g^{cd} \tilde{\Gamma}_{cde}^m}{2n^2} \\ & \quad + (\text{terms independent of } \alpha, \beta) + o(n^{-2}) \end{aligned}$$

(in the last equation we use $g^{be} g^{cd} \tilde{\Gamma}_{ecd}^e = g^{be} g^{cd} \tilde{\Gamma}_{dec}^e$ and $g^{ae} g^{cd} \tilde{\Gamma}_{ecd}^e = g^{ae} g^{cd} \tilde{\Gamma}_{dec}^e$.) Hence we obtain

$$\begin{aligned} & E(\tilde{\omega}^a \tilde{\omega}^b) \\ &= \frac{1}{n} g^{ab} + \frac{1}{n^2} (g^{ca} \partial_c \alpha^b + g^{cb} \partial_c \alpha^a) + \frac{\alpha^a \alpha^b}{n^2} \\ & \quad - \frac{\alpha^a g^{be} g^{cd} \tilde{\Gamma}_{cde}^m}{2n^2} - \frac{\alpha^b g^{ae} g^{cd} \tilde{\Gamma}_{cde}^m}{2n^2} \\ & \quad + (\text{terms independent of } \alpha, \beta) + o(n^{-2}). \end{aligned}$$

Here,

$$\tilde{\Gamma}_{ab\kappa}^m = \tilde{H}_{ab\kappa}^m$$

and

$$-\tilde{\Gamma}_{b\kappa a}^m + T_{ab\kappa} = \tilde{H}_{ab\kappa}^m$$

hold. The first one follows from the definition, and the second one is from the relation

$$\begin{aligned} & \partial_b g_{a\kappa} \\ &= \int \partial_a \partial_b p \frac{\partial_\kappa p}{p} dy + \int \partial_b \partial_\kappa p \frac{\partial_a p}{p} dy - \int \frac{\partial_a p \partial_\kappa p \partial_b p}{p^2} dy \\ &= \tilde{\Gamma}_{ab\kappa}^m + \tilde{\Gamma}_{b\kappa a}^m - T_{ab\kappa} \end{aligned}$$

and $g_{a\kappa} = 0$ ($a = 1, \dots, d$, $\kappa = d+1, \dots, m$).

Therefore, the risk is expanded as follows:

$$\begin{aligned} & E[D\{p(y; \omega), p_{\alpha, \beta}(y; \tilde{\omega}_{MLE})\}] \\ &= \frac{d}{2n} + \frac{g_{ab}}{2n^2} (g^{ca} \partial_c \alpha^b + g^{cb} \partial_c \alpha^a) + \frac{g_{ab}}{2n^2} \alpha^a \alpha^b \\ & \quad + \alpha^a \frac{g^{bc}}{2n^2} \left(\tilde{\Gamma}_{abc}^m + \tilde{\Gamma}_{bca}^m + \tilde{\Gamma}_{cab}^m - 2T_{abc} \right) \\ & \quad - \frac{g_{ab}}{2n^2} \alpha^a g^{be} g^{cd} \tilde{\Gamma}_{cde}^m + \frac{g_{\kappa\lambda}}{2n^2} \beta^\kappa \beta^\lambda \\ & \quad + \beta^\kappa \frac{g^{ab}}{2n^2} \left(\tilde{\Gamma}_{\kappa ab}^m + \tilde{\Gamma}_{b\kappa a}^m + \tilde{\Gamma}_{ab\kappa}^m - 2T_{ab\kappa} \right) \\ & \quad + (\text{terms independent of } \alpha, \beta) + o(n^{-2}) \end{aligned}$$

$$\begin{aligned} &= \frac{d}{2n} + \frac{g_{ab}}{2n^2} \alpha^a \alpha^b + \frac{1}{n^2} \left(\alpha^a \tilde{\Gamma}_{abc}^e g^{bc} + \partial_a \alpha^a \right) \\ & \quad + \frac{g_{\kappa\lambda}}{2n^2} \beta^\kappa \beta^\lambda - \frac{1}{2n^2} \tilde{H}_{ab\kappa}^m g^{ab} \beta^\kappa \\ & \quad + (\text{terms independent of } \alpha, \beta) + o(n^{-2}). \end{aligned}$$

ACKNOWLEDGMENT

The authors would like to thank the referees' comments on an earlier version. They would also like to thank T. Sei for helpful comments. This article is based on a part of the first author's Ph.D. thesis done at The University of Tokyo.

REFERENCES

- [1] J. Aitchison, "Goodness of prediction fit," *Biometrika*, vol. 62, no. 3, pp. 547–554, Dec. 1975.
- [2] X. Xu and D. Zhou, "Empirical Bayes predictive densities for high-dimensional normal models," *J. Multivariate Anal.*, vol. 102, no. 10, pp. 1417–1428, Nov. 2011.
- [3] D. R. Hunter and M. S. Handcock, "Inference in curved exponential family models for networks," *J. Comput. Graph. Statist.*, vol. 15, no. 3, pp. 565–583, Sep. 2006.
- [4] N. Ravishanker, E. L. Melnick, and C.-L. Tsai, "Differential geometry of ARMA models," *J. Time Ser. Anal.*, vol. 11, no. 3, pp. 259–274, May 1990.
- [5] U. K  chler and M. S  rensen, *Exponential Families of Stochastic Processes*. New York, NY, USA: Springer, 1997.
- [6] A. R. Barron and C.-H. Sheu, "Approximation of density functions by sequences of exponential families," *Ann. Statist.*, vol. 19, no. 3, pp. 1347–1369, Sep. 1991.
- [7] S. Amari, *Differential-Geometrical Methods in Statistics*. New York, NY, USA: Springer-Verlag, 1985.
- [8] J. A. Hartigan, "The maximum likelihood prior," *Ann. Statist.*, vol. 26, no. 6, pp. 2083–2103, Dec. 1998.
- [9] A. R. Barron, L. Gy  rfi, and E. C. van der Meulen, "Distribution estimation consistent in total variation and in two types of information divergence," *IEEE Trans. Inf. Theory*, vol. 38, no. 5, pp. 1437–1454, Sep. 1992.
- [10] R. A. Fisher, *Statistical Methods and Scientific Inference*, 3rd ed. New York, NY, USA: Hafner, 1973.
- [11] F. Komaki, "On asymptotic properties of predictive distributions," *Biometrika*, vol. 83, no. 2, pp. 299–313, Jun. 1996.
- [12] I. M. Johnstone, "On the distribution of the largest eigenvalue in principal components analysis," *Ann. Statist.*, vol. 29, no. 2, pp. 295–327, 2001.
- [13] R. E. Kass and P. W. Vos, *Geometrical Foundations of Asymptotic Inference*. New York, NY, USA: Wiley, 1997.
- [14] L. Tierney and J. B. Kadane, "Accurate approximations for posterior moments and marginal densities," *J. Amer. Stat. Assoc.*, vol. 81, no. 393, pp. 82–86, Mar. 1986.
- [15] B. Efron, "Defining curvature of a statistical problem (with applications second order efficiency)," *Ann. Statist.*, vol. 3, no. 6, pp. 1189–1242, Nov. 1975.

Michiko Okudo received the B.E., M.S., and Ph.D. degrees from The University of Tokyo in 2015, 2017, and 2020, respectively. She is currently an Assistant Professor with the Department of Mathematical Informatics, The University of Tokyo.

Fumiyasu Komaki (Member, IEEE) received the B.Eng. and M.Eng. degrees in mathematical engineering from The University of Tokyo, Japan, in 1987 and 1989, respectively, and the Ph.D. degree in statistical science from the Graduate University for Advanced Studies, Japan, in 1992. He is currently a Professor with the Department of Mathematical Informatics, The University of Tokyo, and a Unit Leader with the RIKEN Center for Brain Science. His research interests include prediction theory, Bayesian theory, information geometry, and statistical modeling.