

On the use of marginal posteriors in marginal likelihood estimation via importance sampling

Konstantinos Perrakis¹, Ioannis Ntzoufras¹ and Efthymios G. Tsionas²

¹*Department of Statistics, Athens University of Economics and Business, Greece*

²*Department of Economics, Athens University of Economics and Business, Greece*

Abstract

We investigate the efficiency of a marginal likelihood estimator where the product of the marginal posterior distributions is used as an importance sampling function. The approach is generally applicable to multi-block parameter vector settings, does not require additional Markov Chain Monte Carlo (MCMC) sampling and is not dependent on the type of MCMC scheme used to sample from the posterior. The proposed approach is applied to normal regression models, finite normal mixtures and longitudinal Poisson models, and leads to accurate marginal likelihood estimates.

Keywords: Finite normal mixtures, importance sampling, marginal posterior, marginal likelihood estimation, random effect models, Rao-Blackwellization

1 Introduction

The problem of estimating the marginal likelihood has received considerable attention during the last two decades. The topic is of importance in Bayesian statistics as it is associated with the evaluation of competing hypotheses or models via Bayes factors and posterior model odds. Consider, briefly, two competing models M_1 and M_2 with corresponding prior probabilities $\pi(M_1)$ and $\pi(M_2) = 1 - \pi(M_1)$. After observing a data vector \mathbf{y} , the evidence in favour of M_1 (or against M_2) is evaluated through the odds of the posterior model probabilities $p(M_1|\mathbf{y})$ and $p(M_2|\mathbf{y})$, that is,

$$\frac{p(M_1|\mathbf{y})}{p(M_2|\mathbf{y})} = \frac{m(\mathbf{y}|M_1)}{m(\mathbf{y}|M_2)} \times \frac{\pi(M_1)}{\pi(M_2)}.$$

The quantity $B_{12} = m(\mathbf{y}|M_1)/m(\mathbf{y}|M_2)$ is the ratio of the marginal likelihoods or prior predictive distributions of M_1 and M_2 and is called the Bayes factor of M_1 versus M_2 . The Bayes factor can also be interpreted as the ratio of the posterior odds to the prior odds. When M_1 and M_2 are assumed to be equally probable a-priori, the Bayes factor is equal to the posterior odds.

The marginal likelihood of a given model M_k associated with a parameter vector $\boldsymbol{\theta}_k$ is essentially the normalizing constant of the posterior $p(\boldsymbol{\theta}_k|\mathbf{y}, M_k)$, obtained by integrating the likelihood function $l(\mathbf{y}|\boldsymbol{\theta}_k, M_k)$ with respect to the prior density $\pi(\boldsymbol{\theta}_k|M_k)$, i.e.

$$m(\mathbf{y}|M_k) = \int l(\mathbf{y}|\boldsymbol{\theta}_k, M_k)\pi(\boldsymbol{\theta}_k|M_k)d\boldsymbol{\theta}_k. \quad (1)$$

The integration in (1) may be evaluated analytically for some elementary cases. Most often, it is intractable thus giving rise to the marginal likelihood estimation problem. Numerical integration methods can be used as an approach to the problem, but such techniques are of limited use when sample sizes are moderate to large or when vector $\boldsymbol{\theta}_k$ is of large dimensionality. In addition, the simplest Monte Carlo (MC) estimate, which is given by

$$\hat{m}_n(\mathbf{y}|M_k) = N^{-1} \sum_{n=1}^N l(\mathbf{y}|\boldsymbol{\theta}_k^{(n)}, M_k), \quad (2)$$

using draws $\{\boldsymbol{\theta}_k^{(n)} : n = 1, 2, \dots, N\}$ from the prior distribution is extremely unstable when the posterior is concentrated in relation to the prior. This scenario is frequently met in practice when flat, low-information prior distributions are used to express prior ignorance. A detailed discussion regarding Bayes factors and marginal likelihood estimation is provided by Kass and Raftery (1995).

It is worth noting that the problem of estimating (1) can be bypassed by considering model indicators as unknown parameters. This option has been investigated by several authors (e.g. Green, 1995, Carlin and Chib, 1995, Dellaportas et al., 2002) who introduce MCMC algorithms which sample simultaneously over parameter and model space and deliver directly posterior model probabilities. However, implementation of these methods can get quite complex since they require enumeration of all competing models and specification of tuning constants or “pseudopriors” (depending upon approach) in order to ensure successful mixing in model space. Moreover, since these methods focus on the estimation of posterior model probabilities, accurate estimation of the marginal likelihoods and/or Bayes factors will not be feasible in the cases where a dominating model exists in the set of models under consideration; tuning, following the lines of Ntzoufras et al. (2005), might be possible but is typically inefficient and time consuming.

In contrast, “direct” methods provide marginal likelihood estimates by utilizing the posterior samples of separate models. These methods are usually simpler to implement and are preferable in practice when the number of models under consideration is not large, namely when it is practically feasible to obtain a posterior sample for each of the competing models. Work along these lines includes the Laplace-Metropolis method (Lewis and Raftery, 1997), the harmonic-mean and the prior/posterior mixture importance sampling estimators (Newton and Raftery, 1994), bridge-sampling methods (Meng and Wong, 1994), candidate’s estimators for Gibbs sampling (Chib, 1995) and Metropolis-Hastings sampling (Chib and Jeliazkov, 2001), annealed importance sampling (Neal, 2001), importance-weighted marginal density estimators (Chen, 2005)

and nested sampling approaches (Skilling, 2006). More recently, Raftery et al. (2007) presented a stabilized version of the harmonic-mean estimator, while Friel and Pettitt (2008) and Weinberger (2012) proposed new approaches based on power posteriors and Lebesgue integration theory, respectively. It is worth mentioning that Bayesian evidence evaluation is also of particular interest in the astronomy literature where nested sampling is commonly used for marginal likelihood estimation (e.g. Feroz et al., 2009; Feroz et al., 2011). Recent reviews comparing popular methods based on MCMC sampling can be found in Friel and Wyse (2012) as well as in Ardia et al. (2012). Alternative approaches for marginal likelihood estimation include sequential Monte Carlo (Del Moral et al., 2006) and variational Bayes (Parise and Welling, 2007) methods.

In this paper we propose using the marginal posterior distributions on importance sampling estimators of the marginal likelihood. The proposed approach is particularly suited for the Gibbs sampler, but it is also feasible to use for other types of MCMC algorithms. The estimator can be implemented in a straightforward manner and it can be extended to multi-block parameter settings without requiring additional MCMC sampling apart from the one used to obtain the posterior sample.

The remainder of the paper is organized as follows. The proposed estimator and its variants are discussed in Section 2. In Section 3 the method is applied to normal regression models, to finite normal mixtures and also to hierarchical longitudinal Poisson models. Concluding remarks are provided in Section 4.

2 The proposed estimator

In the following we first introduce the proposed estimator in a two block setting. The more general multi-block case is considered next, explaining why the estimator will be useful in such cases. We further present details concerning the implementation of the proposed approach when the model formulation includes latent variables or nuisance parameters that are not of prime interest for model inference. The section continues with a description of the different estimation approaches of the posterior marginal distributions used as importance functions. We conclude with a note on a convenient implementation of the estimator for models where the posterior distribution becomes invariant under competing diffuse priors and brief remarks about the calculation of numerical standard errors. In the remaining of the paper, the dependence to the model indicator M_k (introduced in the previous section) is eliminated for notational simplicity.

2.1 Introducing the estimator in a two-block setting

Let us consider initially the 2-block setting where $l(\mathbf{y}|\boldsymbol{\theta}, \boldsymbol{\phi})$ is the likelihood of the data conditional on parameter vectors $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_p)^T$ and $\boldsymbol{\phi} = (\phi_1, \phi_2, \dots, \phi_q)^T$, which can be either independent, i.e. $\pi(\boldsymbol{\theta}, \boldsymbol{\phi}) = \pi(\boldsymbol{\theta})\pi(\boldsymbol{\phi})$, or dependent, e.g. $\pi(\boldsymbol{\theta}, \boldsymbol{\phi}) =$

$\pi(\boldsymbol{\theta}|\boldsymbol{\phi})\pi(\boldsymbol{\phi})$, a-priori. In general, one can improve the estimator in (2) by introducing a proper importance sampling density g and then calculate the marginal likelihood as an expectation with respect to g instead of the prior, i.e.

$$m(\mathbf{y}) = \int \frac{l(\mathbf{y}|\boldsymbol{\theta}, \boldsymbol{\phi})\pi(\boldsymbol{\theta}, \boldsymbol{\phi})}{g(\boldsymbol{\theta}, \boldsymbol{\phi})} g(\boldsymbol{\theta}, \boldsymbol{\phi}) d(\boldsymbol{\theta}, \boldsymbol{\phi}) = E_g \left[\frac{l(\mathbf{y}|\boldsymbol{\theta}, \boldsymbol{\phi})\pi(\boldsymbol{\theta}, \boldsymbol{\phi})}{g(\boldsymbol{\theta}, \boldsymbol{\phi})} \right].$$

This quantity can be easily estimated as

$$\hat{m}(\mathbf{y}) = N^{-1} \sum_{n=1}^N \frac{l(\mathbf{y}|\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n)})\pi(\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n)})}{g(\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n)})},$$

where $\boldsymbol{\theta}^{(n)}$ and $\boldsymbol{\phi}^{(n)}$, for $n = 1, 2, \dots, N$, are draws from g . Theoretically, an ideal importance sampling density is proportional to the posterior. In practice, we seek densities which are similar to the posterior and easy to sample from.

Given this consideration, we propose to use the product of the marginal posterior distributions as importance sampling density, i.e. $g(\boldsymbol{\theta}, \boldsymbol{\phi}) \equiv p(\boldsymbol{\theta}|\mathbf{y})p(\boldsymbol{\phi}|\mathbf{y})$. Under this approach

$$m(\mathbf{y}) = \int \int \frac{l(\mathbf{y}|\boldsymbol{\theta}, \boldsymbol{\phi})\pi(\boldsymbol{\theta}, \boldsymbol{\phi})}{p(\boldsymbol{\theta}|\mathbf{y})p(\boldsymbol{\phi}|\mathbf{y})} p(\boldsymbol{\theta}|\mathbf{y})p(\boldsymbol{\phi}|\mathbf{y}) d\boldsymbol{\theta} d\boldsymbol{\phi},$$

which yields the estimator

$$\hat{m}(\mathbf{y}) = N^{-1} \sum_{n=1}^N \frac{l(\mathbf{y}|\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n)})\pi(\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n)})}{p(\boldsymbol{\theta}^{(n)}|\mathbf{y})p(\boldsymbol{\phi}^{(n)}|\mathbf{y})}. \quad (3)$$

Note that the only twist in (3) is that the draws $\boldsymbol{\theta}^{(n)}$ and $\boldsymbol{\phi}^{(n)}$, for $n = 1, 2, \dots, N$, are draws from the marginal posteriors $p(\boldsymbol{\theta}|\mathbf{y})$ and $p(\boldsymbol{\phi}|\mathbf{y})$ and not from the joint posterior $p(\boldsymbol{\theta}, \boldsymbol{\phi}|\mathbf{y})$. In most cases the marginal posterior distributions will not be known, nevertheless, this does not constitute a major obstacle neither for sampling from the marginal posteriors nor for calculating the marginal probabilities which appear in the denominator of (3); the former issue is discussed here, the latter is handled in Section 2.4.

It is straightforward to see that the product marginal posterior is the optimal importance sampling density when $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are independent a-posteriori, since in this case $p(\boldsymbol{\theta}|\mathbf{y})p(\boldsymbol{\phi}|\mathbf{y}) = p(\boldsymbol{\theta}, \boldsymbol{\phi}|\mathbf{y})$ leading to the zero-variance estimator. Although posterior independence is not frequently met in practice, the product marginal posterior can serve as a good approximation to the joint posterior even if $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are not completely independent a-posteriori. First, it has exactly the same support as the joint posterior. Second, the blocking of the parameters can be such that the parameter blocks are close to orthogonal regardless whether the elements within $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are strongly correlated. Furthermore, appropriate reparameterizations can be used in order to form parameter blocks which are orthogonal or close to orthogonal (see e.g. Gilks and Roberts, 1996). Moreover, in generalized linear models the augmentation scheme of Ghosh and Clyde (2011) can be used to obtain orthogonal parameters.

It is worth noting that the estimator in (3) is similar to the Markov chain importance sampling approach described in Botev et al. (2012) and the marginal likelihood estimator proposed in Chan and Eisenstat (2013) based on the cross-entropy method. Botev et al. (2012) show that the product of the marginals is the best importance sampling density – in the sense of minimizing the Kullback-Leibler divergence with respect to the zero-variance importance sampling density – among all product form importance sampling densities, given that the zero-variance density is also decomposable in product form. Similarly, Chan and Eisenstat (2013) locate the importance sampling density minimizing Kullback-Leibler divergence with respect to block-independent factorizations of the joint posterior from distributions belonging to the same parametric families as the priors. The approach presented here differentiates from the above estimators since we consider directly the marginal posteriors and “manipulate” the joint MCMC sample in order to construct marginal samples, thus avoiding further importance sampling and leaving estimation of marginal densities as the main issue to deal with.

In general, marginal posterior samples can be obtained from any MCMC algorithm. The only problem is that a single MCMC chain corresponds to a sample from the joint distribution, with non-zero covariance between parameter blocks. One option is to use a different MCMC run for each block of parameters. In this case, one can calculate the estimator in (3) by using draws $\boldsymbol{\theta}^{(n)}$, $\boldsymbol{\phi}^{(n)}$ coming from two independent MCMC samples of equal size N . Nevertheless, this approach can considerably increase the number of MCMC iterations, especially for a large number of parameter blocks. In addition, the approach is not economical in the sense that only N posterior draws are used from a total sample of $2N$ draws.

A simple and more efficient solution is to re-order a single MCMC chain in such a way that it does not correspond to a sample from the joint posterior distribution. This can be easily implemented by systematically permuting either the sampled values of $\boldsymbol{\theta}$ or those of $\boldsymbol{\phi}$. For instance, consider one MCMC chain where the initial MCMC draws are indexed as $\{\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n)} : n = 1, 2, \dots, N_1, N_1 + 1, N_1 + 2, \dots, N\}$ where $N_1 = N/2$. Then, one can simply re-order the sample of $\boldsymbol{\phi}$ as $\{\boldsymbol{\phi}^{(n_1)} : n_1 = N_1 + 1, N_1 + 2, \dots, N, 1, 2, \dots, N_1\}$ and join the set of draws $\{\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n_1)}\}$, thus forming a sample of paired realizations from the distribution $g(\boldsymbol{\theta}, \boldsymbol{\phi}) = p(\boldsymbol{\theta}|\mathbf{y})p(\boldsymbol{\phi}|\mathbf{y})$. Obviously, when the paired sample has been formed, the distinction between the two sets of indices becomes irrelevant and a common index may be adopted as presented in equation (3). Reordering is trivial in implementation regardless of the number of parameter blocks; the only initial requirement is that the size of the final MCMC sample N must be dividable with the number of blocks, say B , so that B independent reorderings of the MCMC chain can be formed. This line of reasoning also holds for cases of multiple-chain MCMC sampling, which is a frequent MCMC strategy favoured mainly on the basis of MCMC convergence checks (e.g. Gelman and Rubin, 1992). For such cases, one can re-order the within-chain posterior samples in a similar manner to that previously described, and then form a joined sample from the multiple re-ordered chains.

Another, even simpler, alternative for removing correlations between marginal samples is to randomly permute each sample. Results using this approach will be similar to the above systematic re-ordering, except in extreme, unlikely cases where randomly permuted samples with non-zero sample correlation are generated by chance. Irrespective of the reordering scheme used, for the remainder of this paper we use a common index n across parameter blocks, referring to joined independent block samples.

2.2 Extension to multi-block settings

Generalization to multi-block hierarchical settings is straightforward. Consider B blocks $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_B$; in this case the product of the B marginal posteriors is used as importance sampling density and

$$m(\mathbf{y}) = \int \dots \int \frac{l(\mathbf{y}|\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_B)\pi(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_B)}{\prod_{i=1}^B p(\boldsymbol{\theta}_i|\mathbf{y})} \prod_{i=1}^B p(\boldsymbol{\theta}_i|\mathbf{y}) d\boldsymbol{\theta}_i,$$

which yields the following estimator

$$\hat{m}(\mathbf{y}) = N^{-1} \sum_{n=1}^N \frac{l(\mathbf{y}|\boldsymbol{\theta}_1^{(n)}, \dots, \boldsymbol{\theta}_B^{(n)})\pi(\boldsymbol{\theta}_1^{(n)}, \dots, \boldsymbol{\theta}_B^{(n)})}{\prod_{i=1}^B p(\boldsymbol{\theta}_i^{(n)}|\mathbf{y})}. \quad (4)$$

Note that the estimator in (4) may also refer to multiple unidimensional blocks where each parameter forms one block. The advantage of such an estimator raises from the fact that it is easy to construct good approximations of univariate marginal posterior distributions. On the other hand, any possible gain in the efficiency earned from the construction of good approximating densities for the marginal posteriors might be moderated by the use of an importance function which assumes overall independency. Therefore, the most efficient strategy is to choose blocks of minimal size constituted only by highly correlated parameters, which have at the same time weak between-block correlations. Of course, quite often the model design is such that this condition is already met, e.g. for reasons of efficient MCMC mixing. In addition, for cases of Gibbs sampling the natural blocking is the most convenient to use, since in such cases marginal posterior densities can be estimated accurately even for high-dimensional parameter blocks (see Section 2.4 for more details).

2.3 Handling latent variables and nuisance parameters

Many hierarchical models include a block component \mathbf{u} , usually not of main inferential interest, which is associated with a hyperparameter vector $\boldsymbol{\omega}$ through the relationship $\pi(\mathbf{u}, \boldsymbol{\omega}) = \pi(\mathbf{u}|\boldsymbol{\omega})\pi(\boldsymbol{\omega})$. For instance, \mathbf{u} may be a random effect vector or a latent vector used to facilitate posterior simulation through Gibbs sampling as in the data

augmentation setting introduced in Tanner and Wong (1987). In such cases inference usually focuses on the marginal sampling likelihood by integrating out \mathbf{u} . For example, when there is only one parameter block $\boldsymbol{\theta}$, the marginal sampling likelihood is $l(\mathbf{y}|\boldsymbol{\theta}, \boldsymbol{\omega}) = \int l(\mathbf{y}|\boldsymbol{\theta}, \mathbf{u})\pi(\mathbf{u}|\boldsymbol{\omega})d\mathbf{u}$. In this case, a marginal likelihood estimate is obtained through equation (3), where $\boldsymbol{\phi}$ is simply replaced by $\boldsymbol{\omega}$. The extension to the multi-block setting is essentially the same as in (4), only with the addition of $\boldsymbol{\omega}$, specifically

$$\hat{m}(\mathbf{y}) = N^{-1} \sum_{n=1}^N \frac{l(\mathbf{y}|\boldsymbol{\theta}_1^{(n)}, \dots, \boldsymbol{\theta}_B^{(n)}, \boldsymbol{\omega}^{(n)})\pi(\boldsymbol{\theta}_1^{(n)}, \dots, \boldsymbol{\theta}_B^{(n)})\pi(\boldsymbol{\omega}^{(n)})}{\prod_{i=1}^B p(\boldsymbol{\theta}_i^{(n)}|\mathbf{y})p(\boldsymbol{\omega}^{(n)}|\mathbf{y})}. \quad (5)$$

Alternatively, there is also the option of working with the hierarchical likelihood and including \mathbf{u} in the estimation process, i.e.

$$\hat{m}(\mathbf{y}) = N^{-1} \sum_{n=1}^N \frac{l(\mathbf{y}|\boldsymbol{\theta}_1^{(n)}, \dots, \boldsymbol{\theta}_B^{(n)}, \mathbf{u}^{(n)})\pi(\boldsymbol{\theta}_1^{(n)}, \dots, \boldsymbol{\theta}_B^{(n)})\pi(\mathbf{u}^{(n)}|\boldsymbol{\omega}^{(n)})\pi(\boldsymbol{\omega}^{(n)})}{\prod_{i=1}^B p(\boldsymbol{\theta}_i^{(n)}|\mathbf{y})p(\mathbf{u}^{(n)}|\mathbf{y})p(\boldsymbol{\omega}^{(n)}|\mathbf{y})}. \quad (6)$$

The latter approach is less practical to implement as it requires evaluation of $p(\mathbf{u}|\mathbf{y})$. In addition, marginalization over \mathbf{u} will in general lead to more precise marginal likelihood estimates due to scaling down the parameter space; see Vitoratou *et al.* (2013) for further details. Therefore, estimator (5) is overall preferable to estimator (6), except perhaps in cases where the likelihood in (5) is not available analytically and also estimation of $p(\mathbf{u}|\mathbf{y})$ is easy to handle based on the methods discussed next.

2.4 Estimating marginal posterior densities

As seen so far, the proposed approach is fairly simple to implement. The only remaining issue is the evaluation of the marginal posterior probabilities appearing in the denominators of estimators (3)–(6). Here we discuss some different approaches that can be adopted.

A first simple approach is to assume normality either directly or indirectly. Let us consider, for instance, the 2-block setting of Section 2.1 with parameter blocks $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. Suppose, for instance, that $\boldsymbol{\theta}$ relates to a vector of means or a vector of regression parameters. Then, for moderate to large sample sizes, a reasonable option is to assume that $\boldsymbol{\theta}|\mathbf{y} \sim \mathbf{N}(\bar{\boldsymbol{\theta}}, \boldsymbol{\Sigma}_{\boldsymbol{\theta}})$, where $\bar{\boldsymbol{\theta}}$ and $\boldsymbol{\Sigma}_{\boldsymbol{\theta}}$ are the estimated posterior mean vector and variance-covariance matrix from the MCMC output, respectively. Vector $\boldsymbol{\phi}$, on the hand, may refer to a vector of dispersion parameters, where the assumption of normality may not be suitable. One strategy, often sufficient in many cases, is to assume that a transformation $\mathbf{x} = t(\boldsymbol{\phi})$ is approximately normal, i.e. $\mathbf{x}|\mathbf{y} \sim \mathbf{N}(\bar{\mathbf{x}}, \boldsymbol{\Sigma}_{\mathbf{x}})$, and consequently $p(\boldsymbol{\phi}|\mathbf{y}) = p(\mathbf{x}|\mathbf{y})|\frac{dt^{-1}(\mathbf{x})}{d\mathbf{x}}|^{-1}$ for an appropriate invertible function $t(\cdot)$.

An alternative is to mimic the marginal posteriors by adopting appropriate distributional assumptions and matching parameters to posterior moments. This option

is more suitable when the assumption of normality is not particularly supported and appropriate transformation functions are hard to find. In such cases, one can also consider a wide range of options based on multivariate kernel methods (e.g. Scott, 1992) as an efficient alternative.

Moreover, when implementing Gibbs sampling where the normalizing constants of the full conditional distributions are known, marginal posterior densities can be estimated through an efficient, simulation-consistent technique referred as Rao–Blackwellization by Gelfand and Smith (1990). Consider, for instance, the B parameter block setting of Section 2.2; the Rao-Blackwell estimates in this case are

$$\begin{aligned}\hat{p}(\boldsymbol{\theta}_1|\mathbf{y}) &= L^{-1} \sum_{l=1}^L p(\boldsymbol{\theta}_1|\boldsymbol{\theta}_2^{(l)}, \dots, \boldsymbol{\theta}_B^{(l)}, \mathbf{y}), \\ \hat{p}(\boldsymbol{\theta}_b|\mathbf{y}) &= L^{-1} \sum_{l=1}^L p(\boldsymbol{\theta}_b|\boldsymbol{\theta}_1^{(l)}, \dots, \boldsymbol{\theta}_{b-1}^{(l)}, \boldsymbol{\theta}_{b+1}^{(l)}, \dots, \boldsymbol{\theta}_B^{(l)}, \mathbf{y}) \text{ for } b = 2, \dots, B-1 \\ \hat{p}(\boldsymbol{\theta}_B|\mathbf{y}) &= L^{-1} \sum_{l=1}^L p(\boldsymbol{\theta}_B|\boldsymbol{\theta}_1^{(l)}, \dots, \boldsymbol{\theta}_{B-1}^{(l)}, \mathbf{y}).\end{aligned}$$

Note that not all N posterior draws need to be used; usually a sufficiently large subsample of L posterior draws is adequate. For instance, in the examples presented next we find that samples between 200 to 500 draws are sufficient, which significantly reduces computational expense. It should also be noted that Rao-Blackwell estimates must be based on draws from the joint posterior distribution, that is draws from the initial non-permuted MCMC sample.

Finally, for cases of hybrid Gibbs sampling where only some full conditionals are known, one can use a combination of the methods discussed here. Rao-Blackwellization may be used for parameter blocks with known full conditional distributions, whereas for the remaining blocks one can choose among distributional approximations based on moment-fitting and kernel methods.

Estimators (3)–(6) will not be unbiased when approximating the marginal posterior densities using moment-matching strategies. Nevertheless, in practice such “proxies” can be very accurate for univariate as well as multivariate distributions. For instance, as illustrated in Section 3.3, high-dimensional marginal posteriors are approximated efficiently through multivariate normal distributions. In addition, the degree of bias can be empirically checked by comparing such estimates to the corresponding ones using importance samples from the moment-matched approximating distributions. The latter procedure yields an unbiased estimator of the marginal likelihood and, therefore, small observed differences will imply that the bias introduced is negligible.

2.5 Marginal likelihood estimation for diffuse priors

As known, the marginal likelihood is very sensitive to changes in the prior distribution, whereas the posterior distribution (after a point) is insensitive to the prior as the latter

becomes more and more diffuse. Therefore, a usual drawback of marginal likelihood estimators that are based solely on draws from the posterior distribution is that they are typically not reliable for evaluating the marginal likelihoods of different models when considering diffuse priors (see e.g. Friel and Wyse, 2012).

Nevertheless, this is not the case for the proposed estimator as it incorporates the prior in the estimation of the marginal likelihood. In fact, we can easily adopt estimator (3) in order to estimate the marginal likelihood under different diffuse priors (that have no essential effect on the posterior distribution) using a sample from a single MCMC run. To illustrate this, consider the 2-block setting of Section 2.1 and two diffuse priors π_0, π_1 under which the posterior distribution remains unchanged, i.e. $p_0 \equiv p_1$. Let us assume that draws $\{\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n)} : n = 1, 2, \dots, N\}$ are available from an initial MCMC run and that the marginal likelihood m_0 under π_0 has already been estimated through (3). Then, the marginal likelihood under π_1 can be accurately estimated by

$$\widehat{m}_1(\mathbf{y}) = N^{-1} \sum_{n=1}^N \frac{l(\mathbf{y}|\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n)})\pi_1(\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n)})}{p_0(\boldsymbol{\theta}^{(n)}|\mathbf{y})p_0(\boldsymbol{\phi}^{(n)}|\mathbf{y})}. \quad (7)$$

The estimator in (7) does not require additional MCMC sampling, likelihood evaluations or evaluations of the marginal posterior densities since the posterior distributions p_1 and p_0 are the same under π_1 and π_0 , respectively; the only extra effort involved is calculation of the prior probabilities $\pi_1(\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n)})$, for $n = 1, 2, \dots, N$.

2.6 Calculating the numerical standard error

The method of batch means provides a straightforward way for calculating the numerical or MC error of the estimator. Consider for instance the 2-block setting of Section 2.1; in this case the block-independent posterior sample $\boldsymbol{\zeta}^{(n)} \equiv (\boldsymbol{\theta}^{(n)}, \boldsymbol{\phi}^{(n)})$ is divided into K batches $\boldsymbol{\zeta}_1^{(n_k)}, \boldsymbol{\zeta}_2^{(n_k)}, \dots, \boldsymbol{\zeta}_K^{(n_k)}$ of size N_K , i.e. $n_k = 1, 2, \dots, N_K$ and $N = KN_K$, and one calculates

$$\widehat{m}_k(\mathbf{y}) = N_K^{-1} \sum_{n_k=1}^{N_K} \frac{l(\mathbf{y}|\boldsymbol{\theta}_k^{(n_k)}, \boldsymbol{\phi}_k^{(n_k)})\pi(\boldsymbol{\theta}_k^{(n_k)}, \boldsymbol{\phi}_k^{(n_k)})}{p(\boldsymbol{\theta}_k^{(n_k)}|\mathbf{y})p(\boldsymbol{\phi}_k^{(n_k)}|\mathbf{y})}, \quad (8)$$

for $k = 1, 2, \dots, K$. Then, an estimate of the standard error is given by

$$\widehat{\text{s.e.}}(\widehat{m}(\mathbf{y})) = \sqrt{\frac{1}{K(K-1)} \sum_{k=1}^K [\widehat{m}_k(\mathbf{y}) - \overline{m}(\mathbf{y})]^2}, \quad (9)$$

where $\overline{m}(\mathbf{y}) = K^{-1} \sum_{k=1}^K \widehat{m}_k(\mathbf{y})$ is the average batch mean estimate. Note that K must be large enough to ensure proper estimation of the variance (the usual choice is $30 \leq K \leq 50$) and N_K must also be sufficiently large so that the \widehat{m}_k 's are roughly independent (see e.g. Carlin and Louis, 1996).

Alternatively, we can consider the variance estimators of Newey and West (1987) and Geyer (1992) for dependent MCMC draws. Such estimators are suited when systematic re-ordering is used to form the block-independent posterior sample, since, in this case, the posterior dependency patterns will be the same as those of the initial MCMC sample. This is due to the fact that the number of parameter blocks B will be usually much smaller than the size of the posterior sample ($B \ll N$). Therefore, serial auto-correlations for lags greater than N/B are expected to be negligible (for converged MCMC runs), while auto-correlations of lower order are not affected by the re-ordering.

Finally, checking whether the variance is finite or not can be investigated empirically; if the variance is finite then one should expect that increasing the MCMC sample by a factor of d should lead to a decrease of the standard error estimate by a factor approximately equal to \sqrt{d} . As illustrated in Section 3.4, the variance of the proposed estimator is finite for the examples presented in Section 3.

3 Examples

In this section we apply our method to three common classes of models. First, we consider normal linear regression where the true marginal likelihood can be calculated analytically, and compare the proposed estimator to other estimators commonly used in practice. The second example concerns finite normal mixture models where marginal likelihood estimation has proven particularly problematic due to non-identifiability. In the third example, we apply the proposed methods to an hierarchical longitudinal Poisson model where the integrated sampling likelihood is analytically unavailable and, furthermore, standard Gibbs sampling cannot be implemented. The section closes with an empirical diagnostic for checking the assumption of finite variance by comparing the corresponding errors from samples of size N and $2N$.

In all illustrations, we denote the likelihood functions with $l(\cdot)$, prior densities with $\pi(\cdot)$ and posterior or full conditional distributions with $p(\cdot)$. Concerning specific distributional notation, the inverse-gamma density defined in terms of shape α and rate β is denoted by $\mathcal{IG}(\alpha, \beta)$, the Dirichlet distribution with k concentration parameters by $\text{Dir}(\alpha_1, \alpha_2, \dots, \alpha_k)$ and the p -dimensional inverse-Wishart distribution with ν degrees of freedom and scale matrix Ψ by $\mathcal{IW}_p(\nu, \Psi)$.

3.1 Normal regression models

Here we consider the data set presented in Montgomery et al. (2001, p.128) concerning 25 direct current (DC) electric charge measurements (volts) and wind velocity measurements (miles/hour). The goal is to infer about the effect of wind velocity on the production of electricity from a water mill. The models under consideration are

- i) M_0 : the null model with the intercept,
- ii) M_1 : intercept + $(x_1 - \bar{x}_1)$,

- iii) M_2 : intercept + $(x_2 - \bar{x}_2)$ and
- iv) M_3 : intercept + $(x_1 - \bar{x}_1) + x_1^2$,

where x_1 is wind velocity and x_2 is the logarithm of wind velocity. Let j denote the model indicator, i.e. $j = 0, 1, 2, 3$. The likelihood and prior assumptions are the following

$$\begin{aligned} \mathbf{y} | \boldsymbol{\beta}_j, \sigma_j^2 &\sim \mathcal{N}(\mathbf{X}_j \boldsymbol{\beta}_j, \mathbf{I} \sigma_j^2) \\ \boldsymbol{\beta}_j | \sigma_j^2 &\sim \mathcal{N}(\mathbf{0}, \mathbf{V}_j \sigma_j^2) \\ \sigma_j^2 &\sim \mathcal{IG}(10^{-3}, 10^{-3}) \end{aligned}$$

where $\boldsymbol{\beta}_j$ and \mathbf{X}_j correspond to the regression vector and design matrix of model j , respectively, and $\mathbf{V}_j = n^2 (\mathbf{X}_j^T \mathbf{X}_j)^{-1}$ with $n = 25$. In relation to the context of Section 2 this is a 2-block setting where $\boldsymbol{\beta} \equiv \boldsymbol{\theta}$ and $\sigma^2 \equiv \phi$. Under this conjugate prior design the distributions $p(\boldsymbol{\beta}_j | \sigma_j^2, \mathbf{y})$, $p(\sigma_j^2 | \boldsymbol{\beta}_j, \mathbf{y})$, $p(\boldsymbol{\beta}_j | \mathbf{y})$ and $p(\sigma_j^2 | \mathbf{y})$ are all of known form. We treat the posterior distribution as unknown and implement a Gibbs sampler in R. Specifically, one Gibbs chain is iterated 10,000 times and the first 1000 iterations are discarded as burn-in, resulting in a final posterior sample of 9,000 draws for each model. We calculate two variations of estimator (3), considering: i) the true marginals $p(\boldsymbol{\beta}_j | \mathbf{y})$ and $p(\sigma_j^2 | \mathbf{y})$, and ii) Rao-Blackwell estimates of $p(\boldsymbol{\beta}_j | \mathbf{y})$ and $p(\sigma_j^2 | \mathbf{y})$ based on reduced samples of 200 posterior draws. The two variants are denoted by $\hat{m}(\mathbf{y})_{\text{mp}}$ and $\hat{m}(\mathbf{y})_{\text{RB}}$, respectively.

For comparison reasons, we also consider the following commonly used marginal likelihood estimators: the Laplace-Metropolis estimator (Lewis and Raftery, 1997), the importance-weighted marginal density estimator of Chen (2005), the candidate's estimator from Gibbs sampling (Chib, 1995) and the optimal bridge-sampling estimator (Meng and Wong, 1996). For the Laplace-Metropolis we require only the MCMC estimated posterior mean vector and posterior covariance matrix. For the second estimator, which requires specification of approximating densities, we use normal distributions for the $\boldsymbol{\beta}_j$'s and inverse gamma distributions for the σ_j^2 's which mimic the respective component-wise marginal posteriors through moment-fitting. In addition, the points which maximize the unnormalized posterior density of each model are used as posterior ordinates. In order to apply Chib's estimator in a realistic context (using reduced Gibbs sampling) the posterior ordinates (the points maximizing the unnormalized posterior) are decomposed according to the univariate densities. The reduced posterior ordinates are calculated via Rao-Blackwellization based on 9,000 draws from further Gibbs updating (additional sampling is not needed for the simple intercept-model). Finally, for the optimal bridge-sampling estimator, which is calculated iteratively, we utilize the same approximating densities as in the implementation of Chen's estimator and iterate 1000 times using the geometric bridge-sampling estimates, also presented in Meng and Wong (1996), as starting values. The three additional estimators are denoted by $\hat{m}(\mathbf{y})_{\text{LM}}$, $\hat{m}(\mathbf{y})_{\text{Chen}}$, $\hat{m}(\mathbf{y})_{\text{Chib}}$ and $\hat{m}(\mathbf{y})_{\text{obs}}$, respectively.

In order to calculate MC errors the posterior samples are divided into 30 batches of 300 draws. Batch mean estimates, MC errors and the true marginal log-likelihoods

Estimator		Model			
		M_0	M_1	M_2	M_3
Laplace-Metropolis	$\log \hat{m}(\mathbf{y})_{\text{LM}}$	-35.1381 (0.0092)	-12.3676 (0.0124)	-0.4044 (0.0092)	-0.9044 (0.0112)
Importance-weighted	$\log \hat{m}(\mathbf{y})_{\text{Chen}}$	-34.8815 (0.0029)	-13.1407 (0.0039)	-1.5979 (0.0031)	-2.2277 (0.0068)
Candidate's	$\log \hat{m}(\mathbf{y})_{\text{Chib}}$	-34.8789 (0.0020)	-13.1420 (0.0028)	-1.5962 (0.0023)	-2.2337 (0.0067)
Optimal bridge-sampling	$\log \hat{m}(\mathbf{y})_{\text{obs}}$	-34.8807 (0.0011)	-13.1412 (0.0019)	-1.5979 (0.0022)	-2.2294 (0.0030)
Proposed method					
Exact marginals	$\log \hat{m}(\mathbf{y})_{\text{mp}}$	-34.8786 (0.0023)	-13.1420 (0.0035)	-1.5932 (0.0030)	-2.2302 (0.0030)
Rao-Blackwellization	$\log \hat{m}(\mathbf{y})_{\text{RB}}$	-34.8782 (0.0023)	-13.1405 (0.0030)	-1.5919 (0.0030)	-2.2280 (0.0033)
Target value	$\log m(\mathbf{y})$	-34.8797	-13.1429	-1.5953	-2.2270

Table 1: Estimated marginal log-likelihood values compared with true values for Example 1; average batch mean estimates (MC errors in parentheses) are presented using 30 batches of size 300.

are presented in Table 1. In practical terms, we found $\hat{m}(\mathbf{y})_{\text{LM}}$ being the easiest to compute. On the other hand, this estimator performs poorly in comparison to the others, as seen in Table 1. Variations of $\hat{m}(\mathbf{y})_{\text{LM}}$ based on multivariate medians (L_1 centers) and maximum density points of the unnormalized posteriors (not presented here) did not yield substantially different estimates. In contrast, estimators $\hat{m}(\mathbf{y})_{\text{Chen}}$, $\hat{m}(\mathbf{y})_{\text{Chib}}$ and $\hat{m}(\mathbf{y})_{\text{obs}}$ perform substantially better. Implementation for $\hat{m}(\mathbf{y})_{\text{obs}}$ is in general somewhat more complicated in comparison to $\hat{m}(\mathbf{y})_{\text{Chen}}$ as it requires an iterative solution in addition to specification of approximating densities, whereas $\hat{m}(\mathbf{y})_{\text{Chib}}$ requires additional Gibbs sampling for models M_1 , M_2 and M_3 . The estimators proposed here, $\hat{m}(\mathbf{y})_{\text{mp}}$ and $\hat{m}(\mathbf{y})_{\text{RB}}$, only require as input the posterior marginal samples and yield comparable batched mean estimates, with MC errors lower than those of $\hat{m}(\mathbf{y})_{\text{Chen}}$ and just slightly higher than those of estimator $\hat{m}(\mathbf{y})_{\text{obs}}$. Also, note that $\hat{m}(\mathbf{y})_{\text{Chen}}$ and $\hat{m}(\mathbf{y})_{\text{Chib}}$ yield higher MC errors for M_3 . In addition, the estimates derived through Rao-Blackwellization are very similar to the estimates obtained from the true marginal posteriors, while the MC errors are similar across models.

We proceed by testing the ability of the proposed estimator to capture the sensitivity of the marginal likelihood over different diffuse prior distributions which have minimal effect on the posterior distributions of the regression coefficients β_j . The prior used, with $\mathbf{V}_j = n^2(\mathbf{X}_j^T \mathbf{X}_j)^{-1}$, corresponds to a Zellner g -prior (Zellner, 1986) with g set equal to n^2 . For this particular data set the posterior distribution of β_j is

sensitive to the prior when setting g equal to \sqrt{n} and n , which are among the commonly used options (see Fernández et al., 2001). Therefore, we assume more diffuse priors and use the values of 1000, 1500 and 2000 for g ; for these choices the posterior distributions are essentially equivalent with posterior expectations (means, standard deviations etc.) being exact up to the 3rd decimal place. We adopt the approach discussed in Section 2.5 using estimator (7) and the model with $g = 1000$ as the base model from which we sample from the posterior 10,000 draws discarding the first 1000 as burn-in. Batch mean estimates from the Rao-Blackwell estimator and MC errors, based on 30 batches 300 draws, along with the true marginal log-likelihoods are presented in Table 2.

g -prior	Target vs. Estimate	Model			
		M_0	M_1	M_2	M_3
$g = n^2 = 625$	$\log m(\mathbf{y})$	-34.8797	-13.1429	-1.5953	-2.2270
	$\log \hat{m}(\mathbf{y})_{\text{RB}}$	-34.8782 (0.0023)	-13.1405 (0.0030)	-1.5919 (0.0030)	-2.2280 (0.0033)
$g = 1000$	$\log m(\mathbf{y})$	-35.0673	-13.2125	-1.0198	-1.6312
	$\log \hat{m}(\mathbf{y})_{\text{RB}}$	-35.0696 (0.0022)	-13.2063 (0.0043)	-1.0189 (0.0024)	-1.6368 (0.0036)
$g = 1500$	$\log m(\mathbf{y})$	-35.2437	-13.3897	-0.8038	-1.4529
	$\log \hat{m}(\mathbf{y})_{\text{RB}}$	-35.2461 (0.0022)	-13.3836 (0.0043)	-0.8040 (0.0032)	-1.4569 (0.0051)
$g = 2000$	$\log m(\mathbf{y})$	-35.3743	-13.5616	-0.7686	-1.4716
	$\log \hat{m}(\mathbf{y})_{\text{RB}}$	-35.3767 (0.0022)	-13.5556 (0.0044)	-0.7694 (0.0040)	-1.4754 (0.0067)

Table 2: Estimated marginal log-likelihood Rao-Blackwell estimates compared with the true values for Example 1 for the initial g -prior and three diffuse g -priors for the regression vector; average batch mean estimates (MC errors in parentheses) are presented using 30 batches of size 300. The estimates for $g = 1500$ and $g = 2000$ are based on posterior samples from the models with $g = 1000$.

As seen in Table 2, the estimates are accurate despite the fact that the posterior distributions remain the same. In addition, using estimator (7) based on draws from the model with $g = 1000$ required only calculation of prior probabilities for the models with g equal to 1500 and 2000, and led to consistent marginal likelihood estimates. The estimates based on the exact marginal posteriors (not presented here) are equivalent.

3.2 Finite normal mixture models

In this example we consider the well-known galaxy data which where initially presented by Postman et al. (1986). The data are velocities (km's per second) of 82 galaxies from six separated conic sections of the Corona Borealis region. The data set is taken from MASS library in R which contains a "typo"; the value of the 78th observation was

corrected to 26960. The goal is to investigate whether the galaxies can be classified into different clusters according to their velocities, as suggested in astronomical theories. Gaussian finite mixture models are used in the related literature with the purpose of finding the most plausible number of clusters or components. Under this modeling assumption, the likelihood of the velocity data $\mathbf{y} = (y_1, y_2, \dots, y_n)^T$ for a model with k components $w_j \in (0, 1)$, such that $\sum_j w_j = 1$ for $j = 1, 2, \dots, k$, is given by

$$l(\mathbf{y}|\boldsymbol{\mu}, \boldsymbol{\sigma}^2, \mathbf{w}) = \prod_{i=1}^n \sum_{j=1}^k w_j \phi(y_i|\mu_j, \sigma_j^2), \quad (10)$$

where $\mathbf{w} = (w_1, w_2, \dots, w_k)^T$, $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_k)^T$, $\boldsymbol{\sigma}^2 = (\sigma_1^2, \sigma_2^2, \dots, \sigma_k^2)^T$ and $\phi(\cdot)$ is the p.d.f. of the normal distribution. Vectors $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$ consist of the component-specific means and variances, respectively. As originally shown in Dempster et al. (1977), any mixture model can be expressed in terms of missing or latent data; if $z_i \in \{1, 2, \dots, k\}$ represents a latent indicator variable associated with observation y_i , so that $\Pr(z_i = j) = w_j$ and $l(y_i|z_i = j, \mu_j, \sigma_j^2) = \phi(y_i|\mu_j, \sigma_j^2)$, then we have that

$$l(y_i, z_i = j|\mu_j, \sigma_j^2) = l(y_i|z_i = j, \mu_j, \sigma_j^2) \Pr(z_i = j) = \phi(y_i|\mu_j, \sigma_j^2) w_j.$$

Summation over the components w_j results in the complete marginalized data likelihood presented in (10).

As illustrated in West (1992) and Diebolt and Robert (1994), data-augmentation facilitates posterior simulation via Gibbs sampling from the full conditional densities of $\mathbf{w}, \boldsymbol{\mu}, \boldsymbol{\sigma}^2$ and \mathbf{z} . The conjugate priors are $\mu_j \sim \mathcal{N}(\mu_0, \sigma_0^2)$, $\sigma_j^2 \sim IG(\nu_0/2, \delta_0/2)$ and $\mathbf{w} \sim \text{Dir}(\alpha_1, \alpha_2, \dots, \alpha_k)$. The prior for \mathbf{z} is fixed by model design, since $\Pr(z_i = j) = w_j$. Gibbs sampling is straightforward to implement, given these prior assumptions; let $T_j = \{i : z_i = j\}$ be the set of observation indices for those y_i classified into the j -th cluster and let n_j denote the number of observations falling into the j -th cluster. Then, we sample sequentially

$$\begin{aligned} z_i|\mathbf{y}, \mu_j, \sigma_j^2, w_j &\sim \Pr(z_i = j|\mathbf{y}, w_j, \mu_j, \sigma_j^2) \propto w_j \phi(y_i|\mu_j, \sigma_j^2), \\ \mu_j|\mathbf{y}, \sigma_j^2, \mathbf{z} &\sim \mathcal{N}(\hat{\mu}_j, \hat{\sigma}_j^2), \\ \sigma_j^2|\mathbf{y}, \mu_j, \mathbf{z} &\sim \mathcal{IG}\left(\frac{\nu_0 + n_j}{2}, \frac{\delta_0 + \delta_j}{2}\right), \\ \mathbf{w}|\mathbf{y}, \mathbf{z} &\sim \text{Dir}(\alpha_1 + n_1, \alpha_2 + n_2, \dots, \alpha_k + n_k), \end{aligned}$$

where $\hat{\mu}_j = \hat{\sigma}_j^2(\sigma_0^{-2}\mu_0 + \sigma_j^{-2}\sum_{i \in T_j} y_i)$, $\hat{\sigma}_j^2 = (\sigma_0^{-2} + \sigma_j^{-2}n_j)^{-1}$ and $\delta_j = \sum_{i \in T_j} (y_i - \mu_j)^2$. We are also interested in models which have a common variance term in (10); in this case the full conditional of σ^2 is $\mathcal{IG}(\frac{\nu_0 + n}{2}, \frac{\delta_0 + \delta}{2})$ with $\delta = \sum_{j=1}^k \sum_{i \in T_j} (y_i - \mu_j)^2$.

A central point in the discussion that follows is the identifiability problem which is present in mixture models, known as “label-switching”. Non-identifiability arises from the fact that relabelling the mixture components w_j does not change the likelihood in (10). Therefore, when the priors are also invariant to label permutations, the posterior

distribution has $k!$ symmetrical modes. In terms of posterior sampling this implies that common MCMC samplers will most probably fail to explore adequately all $k!$ modes as it is very likely that an MCMC chain will get “trapped” in one particular mode thus leaving the remaining $k! - 1$ modes unvisited. A first suggestion proposed in the early literature is to impose prior ordering constraints, e.g. $w_1 < w_2 < \dots < w_k$ or $\mu_1 < \mu_2 < \dots < \mu_k$, which translate to truncated priors that restrict inference to constrained unimodal posteriors. Robert and Mengersen (1999) further extended this strategy to the use of improper priors through reparameterization. Nevertheless, other authors object to the use of prior identifiability constraints and recommend sampling from the unconstrained posterior. Among them, Celeux et al. (2000) propose tempered transition algorithms and appropriate loss functions for permutation invariant posteriors, while Marin et al. (2005) suggest ex-post reordering schemes.

Chib (1995) was the first who estimated directly the marginal likelihoods of these data for two and three component models via the candidate’s formula and Gibbs updating for the estimation of reduced posterior ordinates. Nevertheless, as pointed out in Neal (1998), Chib’s use of the Gibbs sampler for mixture models results in biased marginal likelihood estimates due to lack of label-switching within the Gibbs sampler. A simple approach to correct for bias is to multiply the marginal likelihood estimates with a factor of $k!$, but as Neal remarked the bias correction will only be valid when the symmetrical modes are well-separated (i.e. when label-switching is not likely to occur). Therefore, Neal (1998) suggests either to introduce special relabelling transitions into the Gibbs sampler or to enforce constrained priors during Gibbs updating which will be $k!$ times larger than the unconstrained priors, as general but computationally demanding solutions. Motivated by the practical bias-correction approach, Berkhof et al. (2003) present simulation consistent marginal likelihood estimators based on a stratification principle and ex-post randomly permuted samples. Frühwirth-Schnatter (2004), on the other hand, recommends to use MCMC samplers which adequately explore all $k!$ labeling posterior subspaces and presents bridge-sampling estimators based on draws from the unconstrained random permutation sampler introduced in Frühwirth-Schnatter (2001).

We consider the same models as Chib (1995) and show that the estimator proposed here can accurately estimate the marginal likelihoods either by taking into account the bias-correction of Neal (1998) or through the use of MCMC samplers which explore effectively the unconstrained posterior space. Specifically, interest lies in the 2-component equal-variance model and 3-component models with equal and unequal variances (i.e. $k = 2, 3$), under the prior assumptions $\mu_0 = 20$, $\sigma_0^2 = 100$, $\nu_0 = 6$, $\delta_0 = 40$ and $\alpha_j = 1$ for $j = 1, 2, \dots, k$. We further take into account a 4-component equal variance model ($k = 4$) with the same prior assumptions. Models with more than four clusters are not considered due to the fact that there is not enough information in the data to support $k > 3$, which gives rise to serious convergence problems due to non-identifiability of parameters for more than four clusters; see Carlin and Chib (1995).

The Gibbs sampler for these models can be easily implemented through package

`bayesmix` (Grün, 2011) in R, which also allows for ex-post reordering and random permutation sampling. We iterate the Gibbs sampler 13,000 times and discard the first 1000 iterations as burn-in. In the context of Section 2 this is a multi-block problem ($\theta_1 \equiv \mu, \theta_2 \equiv \sigma^2, \theta_3 \equiv w$), including a latent vector ($u \equiv z$) which is integrated out. We divide the reordered product marginal posterior sample into $K = 30$ batches of $N_K = 400$ draws and calculate the marginal likelihood for each batch as

$$\hat{m}(\mathbf{y}) = N_K^{-1} \sum_{n=1}^{N_K} \frac{l(\mathbf{y} | \mu^{(n)}, \sigma^{2(n)}, w^{(n)}) \pi(\mu^{(n)}) \pi(\sigma^{2(n)}) \pi(w^{(n)})}{p(\mu^{(n)} | \mathbf{y}) p(\sigma^{2(n)} | \mathbf{y}) p(w^{(n)} | \mathbf{y})}.$$

Marginal posterior densities are estimated through Rao-Blackwellization based on reduced samples of size $L = 500$, i.e. $p(\mu | \mathbf{y}) = \prod_j [\frac{1}{L} \sum_{l=1}^L p(\mu_j | \mathbf{y}, \sigma_j^{2(l)}, z^{(l)})]$, $p(\sigma^2 | \mathbf{y}) = \prod_j [\frac{1}{L} \sum_{l=1}^L p(\sigma_j^2 | \mathbf{y}, \mu_j^{(l)}, z^{(l)})]$ and $p(w | \mathbf{y}) = \frac{1}{L} \sum_{l=1}^L p(w | \mathbf{y}, z^{(l)})$, for $j = 2, 3, 4$. For the equal-variance models we have that $p(\sigma^2 | \mathbf{y}) = \frac{1}{L} \sum_{l=1}^L p(\sigma^2 | \mathbf{y}, \mu_j^{(l)}, z^{(l)})$. Table 3 shows batch mean estimates on log scale and the corresponding MC errors for the simple estimator $\hat{m}(\mathbf{y})$, the bias-corrected estimator $\hat{m}_{bc}(\mathbf{y})$ (obtained by adding the constant $\log k!$ to $\log \hat{m}(\mathbf{y})$) and the estimator based on ex-post random permutation sampling $\hat{m}_{rp}(\mathbf{y})$.

Estimator	Model			
	2 clusters	3 clusters	3 clusters	4 clusters
	equal variance	equal variance	unequal variance	equal variance
$\log \hat{m}(\mathbf{y})^*$	-240.458 (0.002)	-228.597 (0.003)	-228.595 (0.029)	-229.027 (0.045)
$\log \hat{m}_{bc}(\mathbf{y})^\dagger$	-239.765 (0.002)	-226.805 (0.003)	-226.803 (0.029)	-225.849 (0.045)
$\log \hat{m}_{rp}(\mathbf{y})^\ddagger$	-239.762 (0.010)	-226.778 (0.018)	-226.771 (0.051)	-225.922 (0.060)
Neal (1998) estimates	-239.764	-226.803	-226.791	–

* $\hat{m}(\mathbf{y})$: Simple (biased) estimator; $^\dagger \hat{m}_{bc}(\mathbf{y})$: Bias-corrected estimator; $^\ddagger \hat{m}_{rp}(\mathbf{y})$: Random permutation sampling estimator.

Table 3: Estimated marginal log-likelihood values for Example 2; average batch mean estimates (MC errors in parentheses) are presented using 30 batches of size 400.

The benchmark results reported by Neal (1998), based on 10^8 draws from the prior distributions, are also included in Table 3; the corresponding standard errors are 0.005 for the 2 component equal-variance model, 0.040 for the three component equal-variance model and 0.089 for the three component unequal-variance model. It is obvious, that the simple estimator $\hat{m}(\mathbf{y})$ results in biased estimates, which are very similar to the ones presented in Chib (1995); see Neal (1998) for the “typo-corrected” estimate of Chib for the 3rd model. On the other hand, as reflected in the bias-corrected estimator $\hat{m}_{bc}(\mathbf{y})$, simply adding the term $\log k!$ results in accurate marginal likelihood

estimates which are in agreement with the estimates of Neal. Interestingly, the MC errors of $\hat{m}_{bc}(\mathbf{y})$ are similar to the “coefficients of variation” in Steele et al. (2006) who handle marginal likelihood estimation through an incremental mixture importance sampling approach based on marginalization. Nevertheless, Steele et al. (2006) adopt different prior assumptions and, therefore, their marginal likelihood estimates are not comparable to the ones in Table 3.




Figure 1: Histograms of posterior means for the three models from Gibbs sampling.

Histograms of posterior means from Gibbs sampling are presented in Figure 1. As seen, the Gibbs sampler remains in one particular mode for the 2-component and 3-component equal variance models. This is not the case for the 3-component unequal variance model, where label-switching does actually occur for parameters μ_1 and μ_2 . For the 4-component model label-switching is noticeable for all posterior means. Despite that fact, the bias-corrected estimator still performs well for these models. Nevertheless, we would not warrant to guarantee that the bias-corrected estimator will always perform well, especially as the number of clusters gets larger and the posterior modes are not well separated.

Alternatively, one can simply use random permutation sampling and estimate the marginal likelihoods without the need to account for bias-correction. In addition, random permutation sampling will probably prove to be a more reliable solution for models with many components, since the marginal posteriors from random permutation sampling capture all possible modes. The estimates from $\hat{m}_{rp}(\mathbf{y})$ are indeed very similar to Neal’s estimates and to the bias-corrected estimates. The MC errors are slightly higher for the random permutation estimates, nevertheless, this is understandable since ex-post random permutation artificially increases MCMC variability. Histograms of posterior means from random permutation sampling are presented in




Figure 2: Histograms of posterior means for the three models from random permutation sampling.

Figure 2; the symmetries in the posterior distributions due to non-identifiability are now apparent. In accordance to the discussion in Carlin and Chib (1995), the histograms for the 4-component model show that only three modes are estimated efficiently as there is a significant overlap between the 2nd and 3rd mode.

In conclusion, both $\hat{m}_{bc}(\mathbf{y})$ and $\hat{m}_{rp}(\mathbf{y})$ yield satisfactory results. For models with a small number of components (i.e. when label-switching is not likely to occur) the bias-corrected estimator will most probably be sufficient. For more complicated models and when the two estimators result in estimates which are in disagreement, we would recommend to use either the correction for the candidate estimator proposed by Marin and Robert (2008) or the estimator based on alternative MCMC strategies (e.g. Frühwirth-Schnatter, 2001; Geweke, 2007).

3.3 Longitudinal Poisson models

As a last example, we consider a data set taken from Diggle et al. (1995), consisting of seizure counts y_{it} from a group of epileptics ($i = 1, 2, \dots, 59$) which is monitored initially over an 8-week baseline period ($t = 0$) and then over four subsequent 2-week periods ($t = 1, 2, 3, 4$). Each patient is randomly assigned either a placebo or the drug progabide after the baseline period. This example is chosen mainly because standard Gibbs sampling is not possible to implement for the model presented next. In addition, the epilepsy data is also considered by Chib et al. (1998) and Chib and Jeliazkov (2001) who present marginal likelihood estimates based on the candidate's formula and Metropolis-Hastings sampling. Reduced posterior ordinates are calculated through

kernel density estimation in Chib et al. (1998), whereas Chib and Jeliazkov (2001) employ Metropolis-Hastings updating. For the sake of comparison, we adopt exactly the same modeling assumptions.

The main model under consideration is

$$\begin{aligned} y_{it} | \boldsymbol{\beta}, \mathbf{b}_i &\sim \text{Poisson}(\lambda_{it}), \\ \log \lambda_{it} &= \log \tau_{it} + \beta_1 x_{it1} + \beta_2 x_{it2} + b_{i1} + b_{i2} x_{it2}, \\ \mathbf{b}_i &\sim \mathcal{N}_2(\boldsymbol{\eta}, \mathbf{D}), \end{aligned}$$

where τ_{it} is the offset which equals 8 when $t = 0$ and 2 otherwise, x_{it1} is an indicator of treatment (0 for placebo, 1 for progabide treatment), x_{it2} is an indicator of time period (0 for baseline, 1 otherwise) and $\mathbf{b}_i = (b_{i1}, b_{i2})$ are latent random effects for $i = 1, 2, \dots, 58$ (subject 49 is removed from the analysis due to unusually high pre-and post-randomization seizure counts). The prior assumptions are bivariate normal distributions for $\boldsymbol{\beta}$ and $\boldsymbol{\eta}$ and a bivariate inverse-Wishart for \mathbf{D} , namely $\boldsymbol{\beta} \sim \mathcal{N}_2(\mathbf{0}, 100\mathbf{I}_2)$, $\boldsymbol{\eta} \sim \mathcal{N}_2(\mathbf{0}, 100\mathbf{I}_2)$ and $\mathbf{D} \sim \mathcal{IW}_2(4, \mathbf{I}_2)$, where \mathbf{I}_2 is the 2×2 identity matrix. The full conditionals of $\boldsymbol{\eta}$ and \mathbf{D} are known, specifically we have that

$$\begin{aligned} \boldsymbol{\eta} | \mathbf{D}, \mathbf{b}_1, \dots, \mathbf{b}_{58} &\sim \mathcal{N}_2(\hat{\boldsymbol{\eta}}, \mathbf{V}), \\ \mathbf{D} | \boldsymbol{\eta}, \mathbf{b}_1, \dots, \mathbf{b}_{58} &\sim \mathcal{IW}_2\left(58 + 4, \mathbf{I}_2 + \sum_{i=1}^{58} (\mathbf{b}_i - \boldsymbol{\eta})(\mathbf{b}_i - \boldsymbol{\eta})^T\right), \end{aligned}$$

where $\hat{\boldsymbol{\eta}} = \mathbf{V} \sum_{i=1}^{58} \mathbf{D}^{-1} \mathbf{b}_i$ and $\mathbf{V} = (100^{-1} \mathbf{I}_2 + 58 \mathbf{D}^{-1})^{-1}$. The full conditionals for $\boldsymbol{\beta}$ and the \mathbf{b}_i 's are not known distributions and thus standard Gibbs sampling is not feasible. Another complication is that the integrated sampling likelihood $l(\mathbf{y} | \boldsymbol{\beta}, \boldsymbol{\eta}, \mathbf{D}) = \prod_{i=1}^{58} \int l(\mathbf{y}_i | \boldsymbol{\beta}, \mathbf{b}_i) \pi(\mathbf{b}_i) d\mathbf{b}_i$, with $l(\mathbf{y}_i | \boldsymbol{\beta}, \mathbf{b}_i) = \prod_{t=0}^4 l(y_{it} | \boldsymbol{\beta}, \mathbf{b}_i)$, is also not available analytically. Therefore, evaluating $l(\mathbf{y} | \boldsymbol{\beta}, \boldsymbol{\eta}, \mathbf{D})$ requires either numerical integration or some other efficient technique, such as importance sampling for instance.

We utilize WinBUGS software (Spiegelhalter et al., 2003) to sample from the posterior. Specifically, one chain is iterated 31,000 times and the first 1000 iterations are discarded as burn-in, resulting in a final sample of 30,000 draws. Posterior means and standard deviations, for the parameters of scientific interest, are presented in Table 4. The estimates for the main model are comparable to the Metropolis-Hastings estimates presented in Chib et al. (1998). Table 4 also includes the estimates for the simpler model without the random effects related to time. For this model we assume a-priori that $b_i \sim \mathcal{N}(\eta_1, D_{11})$, for $i = 1, 2, \dots, 58$. In order to keep equivalent prior assumptions to the main model, we define the priors as $\eta_1 \sim \mathcal{N}(0, 100)$ and $D_{11} \sim \mathcal{IG}(2, 1/2)$.

As discussed in Section 2.3, there are two approaches for estimating the marginal likelihood of this model. The first is to treat it as a 3-block setting, i.e. consider the product of the marginal posteriors of $\boldsymbol{\beta}, \boldsymbol{\eta}, \mathbf{D}$ as importance sampling density. In this case one needs to estimate the integrated likelihood which is unknown. The

second approach is to treat the problem as a 4-block setting, i.e. also include the joint marginal posterior of the \mathbf{b}_i 's in the importance sampling density. The advantage with this approach is that we can work directly with the hierarchical Poisson likelihood.

Initially, let us consider the first approach which corresponds to estimator (5) of Section 2.3; in this case the parameters of scientific interest are $\boldsymbol{\theta}_1 \equiv \boldsymbol{\beta}, \boldsymbol{\theta}_2 \equiv \boldsymbol{\eta}, \boldsymbol{\theta}_3 \equiv \mathbf{D}$, while $\mathbf{u} \equiv \{\mathbf{b}_1, \dots, \mathbf{b}_{58}\}$ is used only for Rao-Blackwellization. First, we appropriately re-order the posterior sample in order to correspond to a sample from the product marginal posterior and then we split the sample into $K = 30$ batches of $N_K = 1000$ draws. The marginal likelihood estimate for each batch is calculated as

$$\hat{m}_{b_3}(\mathbf{y}) = N_K^{-1} \sum_{n=1}^{N_K} \frac{l(\mathbf{y}|\boldsymbol{\beta}^{(n)}, \boldsymbol{\eta}^{(n)}, \mathbf{D}^{(n)})\pi(\boldsymbol{\beta}^{(n)})\pi(\boldsymbol{\eta}^{(n)})\pi(\mathbf{D}^{(n)})}{p(\boldsymbol{\beta}^{(n)}|\mathbf{y})p(\boldsymbol{\eta}^{(n)}|\mathbf{y})p(\mathbf{D}^{(n)}|\mathbf{y})}.$$

Marginal posterior probabilities for $\boldsymbol{\eta}$ and \mathbf{D} are estimated via Rao-Blackwellization based on reduced posterior samples of $L = 200$ draws which are randomly re-sampled from the initial MCMC sample, i.e. $p(\boldsymbol{\eta}|\mathbf{y}) = \frac{1}{L} \sum_{l=1}^L p(\boldsymbol{\eta}|\mathbf{D}^{(l)}, \mathbf{b}_1^{(l)}, \dots, \mathbf{b}_{58}^{(l)})$ and $p(\mathbf{D}|\mathbf{y}) = \frac{1}{L} \sum_{l=1}^L p(\mathbf{D}|\boldsymbol{\eta}^{(l)}, \mathbf{b}_1^{(l)}, \dots, \mathbf{b}_{58}^{(l)})$. For the marginal posterior of $\boldsymbol{\beta}$ we assume that $p(\boldsymbol{\beta}|\mathbf{y}) \approx \mathcal{N}_2(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}})$, where $\tilde{\boldsymbol{\beta}}$ and $\tilde{\boldsymbol{\Sigma}}$ are estimated from the MCMC output. Similarly to Chib et al. (1998) and Chib and Jeliazkov (2001), we employ further importance sampling to evaluate the likelihood $l(\mathbf{y}|\boldsymbol{\beta}, \boldsymbol{\eta}, \mathbf{D})$. We employ multivariate normals as importance sampling functions, namely $p_{IS}(\mathbf{b}_1, \dots, \mathbf{b}_{58}|\mathbf{y}) \approx \mathcal{N}_{116}(\tilde{\mathbf{b}}, \tilde{\mathbf{B}})$ for the main model and $p_{IS}(b_1, \dots, b_{58}|\mathbf{y}) \approx \mathcal{N}_{58}(\tilde{\mathbf{b}}, \tilde{\mathbf{B}})$ for the reduced model, where $\tilde{\mathbf{b}}, \tilde{\mathbf{B}}$ are the vector of means and the complete covariance matrix of the random effects estimated from the MCMC output. Likelihood estimation is based on 100 importance sampling draws.

Based on the alternative 4-block approach, corresponding to estimator (6) of Section 2.3, the batched marginal likelihood estimates are calculated as

$$\hat{m}_{b_4}(\mathbf{y}) = N_K^{-1} \sum_{n=1}^{N_K} \left[\frac{l(\mathbf{y}|\boldsymbol{\beta}^{(n)}, \mathbf{b}_1^{(n)}, \dots, \mathbf{b}_{58}^{(n)}) \times \pi(\boldsymbol{\beta}^{(n)})\pi(\mathbf{b}_1^{(n)}, \dots, \mathbf{b}_{58}^{(n)}|\boldsymbol{\eta}^{(n)}, \mathbf{D}^{(n)})\pi(\boldsymbol{\eta}^{(n)})\pi(\mathbf{D}^{(n)})}{\left\{ p(\boldsymbol{\beta}^{(n)}|\mathbf{y})p(\mathbf{b}_1^{(n)}, \dots, \mathbf{b}_{58}^{(n)}|\mathbf{y})p(\boldsymbol{\eta}^{(n)}|\mathbf{y})p(\mathbf{D}^{(n)}|\mathbf{y}) \right\}^{-1}} \right].$$

With this approach, there is no need to implement further importance sampling for evaluating the likelihood function since the data conditional on the random effects parameters follow the Poisson distribution. Despite this convenient aspect, the downside of the 4-block estimator is that it requires estimation of the high-dimensional joint marginal $p(\mathbf{b}_1, \dots, \mathbf{b}_{58}|\mathbf{y})$. Given that the Rao-Blackwell device cannot be used, we adopt a simple assumption and namely use the importance sampling functions used for likelihood evaluation in estimator $\hat{m}_{b_3}(\mathbf{y})$, i.e. we assume that $p(\mathbf{b}_1, \dots, \mathbf{b}_{58}|\mathbf{y}) \approx \mathcal{N}_{116}(\tilde{\mathbf{b}}, \tilde{\mathbf{B}})$ for the model with time effects and $p(b_1, \dots, b_{58}|\mathbf{y}) \approx \mathcal{N}_{58}(\tilde{\mathbf{b}}, \tilde{\mathbf{B}})$ for the model not including time effects.

Average batch mean marginal likelihood estimates and MC errors for the two models in question are presented in Table 4. The 3-block approach provides accurate

Parameter	Model with time effect		Model without time effect	
	Mean	St.Dev.	Mean	St.Dev.
Constant η_1	1.065	0.146	1.095	0.138
Treatment β_1	-0.0003	0.209	-0.071	0.190
Time η_2	0.005	0.111	-	-
Interaction β_2	-0.349	0.156	-0.191	0.052
D_{11}	0.474	0.100	0.531	0.105
D_{12}	0.017	0.057	-	-
D_{22}	0.243	0.063	-	-
Log-marginal likelihood estimates				
$\log \hat{m}_{b_3}(\mathbf{y})$	-914.992 (0.035)		-966.971 (0.018)	
$\log \hat{m}_{b_4}(\mathbf{y})$	-914.485 (0.137)		-966.814 (0.064)	
Chib et al. (1998)	-915.404		-969.824	
Chib & Jeliazkov (2001)	-915.230			

Table 4: Posterior means and standard deviations from 30,000 posterior draws for the parameters of two epilepsy models and average batch mean marginal likelihood estimates (MC errors in brackets) from 30 batches of size 1000 for the 3-block and 4-block estimators for Example 3.

estimates with MC errors being very low in comparison to the magnitude of the batched means. The 4-block estimates are in agreement with the 3-block estimates, but have higher Monte Carlo errors; approximately four times higher than the MC errors of $\hat{m}_{b_3}(\mathbf{y})$. Nevertheless, this is expected due to the much larger augmented parameter space and the use of the normal approximation of the high-dimensional joint posterior of the random effects. Overall, estimator $\hat{m}_{b_4}(\mathbf{y})$ is computationally less demanding than $\hat{m}_{b_3}(\mathbf{y})$ and the resulting MC errors, although higher than those of $\hat{m}_{b_3}(\mathbf{y})$, are still relatively low in comparison to the batched mean marginal likelihood values.

Table 4 also includes the estimates presented in Chib et al. (1998) and Chib and Jeliazkov (2001) from 10,000 posterior draws. Both latter studies report standard errors, based on the variance estimator of Newey and West (1987), of approximately 0.1 for the main model. The reduced model is briefly considered in Chib et al. (1998) without reporting a standard error. Our 3-block and 4-block estimates are comparable to those of Chib et al. (1998) and Chib and Jeliazkov (2001) but not in strict agreement. Our experience from this particular example is that a long MCMC chain is needed in order to obtain accurate posterior estimates; initial 3-block estimates based on 10,000 posterior draws were actually closer to the estimates of Chib et al. (1998) and Chib and Jeliazkov (2001), namely -915.566 for the main model and -969.580 for the reduced model, but with considerably higher MC errors.

3.4 MC errors for different number of iterations

Here we briefly investigate the issue of finite variance for the illustrated examples presented in this section. Following the discussion in Section 2.6, an informal empirical

Examples and models	MCMC length	MC error
<u>Example 1: Normal regression models[*] ($N = 9000$)</u>		
intercept	N	0.0023
	$2N$	0.0018 (0.0016)
intercept+wind speed	N	0.0030
	$2N$	0.0021 (0.0021)
intercept+log(wind speed)	N	0.0030
	$2N$	0.0023 (0.0021)
intercept+wind speed+(centered wind speed) ²	N	0.0033
	$2N$	0.0026 (0.0023)
<u>Example 2: Gaussian mixture models[†] ($N = 12000$)</u>		
2 clusters equal variance	N	0.010
	$2N$	0.007 (0.007)
3 clusters equal variance	N	0.018
	$2N$	0.015 (0.013)
3 clusters unequal variance	N	0.051
	$2N$	0.039 (0.036)
4 clusters equal variance	N	0.060
	$2N$	0.047 (0.042)
<u>Example 3: Longitudinal Poisson models[‡] ($N = 30000$)</u>		
Model with time effect	N	0.137
	$2N$	0.105 (0.097)
Model without time effect	N	0.063
	$2N$	0.040 (0.045)

^{*}MC errors for the Rao-Blackwell estimator for $g = n^2$; [†]MC errors for the random permutation estimator; [‡]MC errors for the 4-block estimator

Table 5: MC errors of marginal log-likelihood estimates for MCMC runs of length N and $2N$; the number of batches is 30 for all models. The expected MC errors for samples of size $2N$ under the assumption of finite variance are presented in parentheses.

diagnostic for checking whether the variance is finite can be performed by comparing the MC errors from MCMC samples of different sizes. In general, increasing the MCMC sample by a factor of d should lead to a decrease of MC errors by a factor of \sqrt{d} for estimators which have finite variance. Table 5 depicts the estimated MC errors for the Rao-Blackwell estimator for Example 1, the estimator based on random-permutation sampling for Example 2 and the 4-block estimator for Example 3. In all cases, the MC errors from the original posterior samples of size N are compared with the corresponding errors from samples of size $2N$ and with the errors from the original samples scaled down by a factor of $\sqrt{2}$, which are the expected MC errors under the assumption of finite variance. From this table, it is evident that the MC errors from the MCMC runs with length equal to $2N$ are roughly equal to the MC errors from the chains with length N divided by $\sqrt{2}$, for all models, indicating that the variance of the corresponding estimators is finite.

4 Concluding remarks

In this paper we have presented a method of marginal likelihood estimation based on utilizing the product marginal posterior as importance sampling density. The approach is in general straightforward to implement even for multi-block parameter settings as it is non-iterative and does not require adaptations in MCMC sampling. As illustrated, the estimator is accurate in capturing changes in the marginal likelihood due to different diffuse prior setups that do not affect the posterior distribution. For such cases, the computational demands for estimating marginal likelihoods of competing models under diffuse priors are reduced significantly, since only one MCMC run is required. In general, the overall performance of the estimator depends on; i) the efficiency of approximating the joint posterior through independent univariate or multivariate marginals and ii) the accuracy in estimating marginal posterior densities.

Arguably, the method can fail when the product of marginal posteriors is a poor approximation to the joint posterior. Nevertheless, appropriate parameter blocking and reparameterizations can always improve the performance of the method, so that it will be feasible to work with a few parameter blocks that are close to orthogonal regardless whether the elements within the blocks are highly correlated. In the three, relatively diverse, examples handled in this paper the natural blocking of the parameters proved to be sufficient in delivering accurate estimates. It is worth noting that similar estimators based on importance sampling from independent posterior factorizations have shown to perform well (Botev et al., 2012; Chan and Eisenstat, 2013). Moreover, independent posterior factorization is also extensively used for the variational Bayes (Bishop, 2006) and expectation-propagation (Minka, 2001) approaches in the machine-learning literature. As a last remark concerning this topic, Ghosh and Clyde (2011) present a methodology for linear and binary regression models that augments non-orthogonal designs to obtain orthogonal designs based on Gibbs sampling for the “missing” response variables. With some additional effort one could consider

this orthogonalization approach which would guarantee an optimal importance sampling density, thus leaving estimation of univariate marginal posterior densities as the only remaining source of error.

With respect to estimating marginal probabilities, the approach proposed here is particularly suited for Gibbs sampling settings where Rao-Blackwellization can be used to obtain simulation-consistent marginal posterior density estimates. Practically, the proposed estimator can get computationally demanding when using Rao-Blackwellization for the entire posterior sample. Nevertheless, the related coding work basically requires averaging and is straightforward, without requiring any special effort in implementation or fine-tuning of parameters in trial and error runs. In addition, as illustrated in the examples, the sample needed for Rao-Blackwellization is substantially smaller than the total MCMC sample and is obtained once as a random sub-sample of the MCMC chain.

The approach can also be applied under other types of MCMC schemes by adopting other strategies for estimating the marginal posterior densities such as normal approximations, fitting posterior moments, kernel methods and so forth. In strict theory, the method will not yield unbiased estimates when using such approximating strategies, nevertheless, in practical terms such approaches can often be sufficient and can lead to accurate estimates even for high dimensional multivariate approximations, as demonstrated in Section 3.3. In addition, the degree of bias can be checked indirectly by using the approximating densities as importance sampling densities. It is worth noting, that more elaborated strategies can also be considered, for instance the methods discussed in Oh (1999) based on importance-weighted marginal density estimation (Chen, 1994) or the integrated nested Laplace approximations (INLA's) presented in Rue et al. (2009).

The advantage of not depending on the type of MCMC scheme used to sample from the posterior becomes obvious for classes of models like the finite normal mixtures, considered here, where conventional Gibbs sampling fails to explore multi-modal posterior surfaces. This implies that the proposed method will also work well for models with similar posterior symmetries, based on alternative samplers (e.g. Frühwirth-Schnatter, 2001; Geweke, 2007), without increased complexity in estimation.

A possibly interesting extension of the idea presented here is to incorporate it within bridge-sampling estimation by using the product marginal posterior as approximating density.

Acknowledgements

The authors would like to thank two anonymous referees for their interesting comments and suggestions.

References

- Ardia, D., Baştürk, N., Hoogerheide, L. and van Dijk, H.K. (2012). A comparative study of Monte Carlo methods for efficient evaluation of marginal likelihood. *Computational Statistics and Data Analysis*, **56**, 3398–3414.
- Berkhof, J., van Mechelen, I. and Gelman, A. (2003). A Bayesian approach to the selection and testing of mixture models. *Statistica Sinica*, **13**, 423–442.
- Bishop, C.M. (2006). *Pattern Recognition and Machine Learning*. New York: Springer.
- Botev, Z., L’Ecuyer, P. and Tuffin, B. (2013). Markov chain importance sampling with applications to rare event probability estimation. *Statistics and Computing*, **23**, 271–285.
- Carlin, B.P. and Chib, S. (1995). Bayesian model choice via Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society B*, **57**, 473–484.
- Carlin, B.P. and Louis, T. (1996). *Bayes and Empirical Bayes Methods for Data Analysis*. London: Chapman & Hall\CRC.
- Celeux, G., Hurn, M. and Robert, C. (2000). Computational and inferential difficulties with mixtures posterior distribution. *Journal of the American Statistical Association*, **95**, 957–979.
- Chan, J. and Eisenstat, E. (2013). Marginal likelihood estimation with the cross-entropy method. *Econometric Reviews*, *forthcoming*.
- Chen, M.-H. (1994). Importance-weighted marginal Bayesian posterior density estimation. *Journal of the American Statistical Association*, **89**, 818–824.
- Chen, M.-H. (2005). Computing marginal likelihoods from a single MCMC output. *Statistica Neerlandica*, **59**, 16–29.
- Chib, S. (1995). Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association*, **90**, 1313–132.
- Chib, S., Greenberg, E. and Winkelmann, R. (1998). Posterior simulation and Bayes factors in panel count data models. *Journal of Econometrics*, **86**, 33–54.
- Chib, S. and Jeliazkov, I. (2001). Marginal likelihood from the Metropolis-Hastings output. *Journal of the American Statistical Association*, **96**, 270–281.
- Dellaportas, P., Forster, J.J. and Ntzoufras, I. (2002). On Bayesian model and variable selection using MCMC. *Statistics and Computing*, **12**, 27–36.
- Del Moral, P., Doucet, A. and Jasra, A. (2006). Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society B*, **68**, 411–436.

- Dempster, A. P., Laird, N. M. and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society B*, **39**, 1–38.
- Diebolt, J. and Robert, C.P. (1994). Estimation of finite mixture distributions through Bayesian sampling. *Journal of the Royal Statistical Society B*, **56**, 363–375.
- Diggle, P., Liang, K.-Y. and Zeger, S.L. (1995). *Analysis of Longitudinal Data*. Oxford: Oxford University Press.
- Fernández, C., Ley, E. and Steel, M.F.J. (2001). Benchmark priors for Bayesian model averaging. *Journal of Econometrics*, **100**, 381–427.
- Feroz, F., Balan, S.T. and Hobson, M.P. (2011). Bayesian evidence for two companions orbiting HIP 5158. *Monthly Notices of the Royal Astronomical Society*, **416**, L104–L108.
- Feroz, F., Hobson, M.P. and Bridges, M. (2009). MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics. *Monthly Notices of the Royal Astronomical Society*, **398**, 1601–1614.
- Friel, N. and Pettitt, A.N. (2008). Marginal likelihood estimation via power posteriors. *Journal of the Royal Statistical Society B*, **70**, 589–607.
- Friel, N. and Wyse, J. (2012). Estimating the evidence – a review. *Statistica Neerlandica*, **66**, 288–308.
- Frühwirth-Schnatter, S. (2001). Markov chain Monte Carlo estimation of classical and dynamic switching and mixture models. *Journal of the American Statistical Association*, **96**, 194–209.
- Frühwirth-Schnatter, S. (2004). Estimating marginal likelihoods for mixture and Markov switching models using bridge sampling techniques. *The Econometrics Journal*, **7**, 143–167.
- Gelfand, A.E. and Smith, A.F.M. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, **85**, 398–409.
- Gelman, A. and Rubin, D.B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, **7**, 457–511.
- Geweke, J. (2007). Interpretation and inference in mixture models: simple MCMC works. *Computational Statistics and Data Analysis*, **51**, 3529–3550.
- Geyer, C.J. (1992). Practical Markov chain Monte Carlo. *Statistical Science*, **4**, 473–483.

- Ghosh, J. and Clyde, M.A. (2011). Rao-Blackwellization for Bayesian variable selection and model averaging in linear and binary regression: a novel data augmentation approach. *Journal of the American Statistical Association*, **106**, 1041–1052.
- Gilks, W.R. and Roberts, G.O. (1996). Strategies for improving MCMC. In *Markov Chain Monte Carlo in Practice*, 6, eds. W.R. Gilks, S. Richardson and D.J. Spiegelhalter, London: Chapman & Hall\CRC, pp. 89–114.
- Green, P.J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, **82**, 711–732.
- Grün, B. (2011). *bayesmix: Bayesian Mixture Models with JAGS*. Available at <http://cran.r-project.org/web/packages/bayesmix/index.html>.
- Kass, R.E. and Raftery, A.E. (1995). Bayes factors and model uncertainty. *Journal of the American Statistical Association*, **90**, 773–795.
- Lewis, S.M. and Raftery, A.E. (1997). Estimating Bayes factors via posterior simulation with the Laplace-Metropolis estimator. *Journal of the American Statistical Association*, **92**, 648–655.
- Marin, J.-M., Mengersen, K. and Robert, C. (2005). Bayesian modelling and inference on mixtures of distributions. In *Handbook of Statistics*, vol. 25, eds. C. Rao and D. Dey, New York: Elsevier, pp. 459–507.
- Marin, J.-M and Robert, C.P. (2008). Approximating the marginal likelihood in mixture models. *Bulletin of the Indian Chapter of ISBA*, **V**(1), 2–7; also available as [arXiv0804.2414](https://arxiv.org/abs/0804.2414).
- Meng, X.-L. and Wong, W.H. (1996). Simulating ratios of normalizing constants via a simple identity: a theoretical exploration. *Statistica Sinica*, **6**, 831–860.
- Minka, T.P. (2001). Expectation propagation for approximate Bayesian inference. *Uncertainty in Artificial Intelligence*, **17**, 362–369.
- Montgomery, D.C., Peck, E.A. and Vining, G.G. (2001). *Introduction to Linear Regression Analysis*. New York: John Wiley.
- Neal, R.M. (1998). Erroneous results in ‘Marginal likelihood from the Gibbs output’. Available at <http://www.cs.utoronto.ca/radford/radford@stat.utoronto.ca>.
- Neal, R.M. (2001). Annealed importance sampling. *Statistics and Computing*, **11**, 125–139.
- Newey, W. K. and West, K. D. (1987). A simple positive semi-definite heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, **55**, 703–708.

- Newton, M.A. and Raftery, A.E. (1994). Approximate Bayesian inference with the weighted likelihood bootstrap. *Journal of the Royal Statistical Society B*, **56**, 3–48.
- Ntzoufras, I., Katsis, A. and Karlis, D. (2005). Bayesian assessment of the distribution of insurance claim counts using reversible jump MCMC. *North American Actuarial Journal*, **9**, 90–108.
- Oh, M.-S. (1999). Estimation of posterior density functions from a posterior sample. *Computational Statistics and Data Analysis*, **29**, 411–427.
- Parise, S. and Welling, M. (2007). Bayesian model scoring in Markov random fields. In *Advances in Neural Information Processing Systems*, 19, eds. B. Schölkopf, J. Platt and T. Hoffman, Cambridge, MA: MIT Press, pp. 1073–1080.
- Postman, M., Huchra, J.P. and Geller, M.J. (1986). Probes of large-scale structures in the Corona Borealis region. *The Astronomical Journal*, **92**, 1238–1247.
- Raftery, A.E., Newton, M.A., Satagopan, J.M. and Krivitsky, P.N. (2007). Estimating the integrated likelihood via posterior simulation using the harmonic mean identity. In *Bayesian Statistics*, 8, eds. J.M. Bernardo, M.J. Bayarri, J.O. Berger, A.P. Dawid, D. Heckerman, A.F.M. Smith and M. West, Oxford, U.K.: Oxford University Press, pp. 1–45.
- Robert, C. P. and Mengersen, K. L. (1999). Reparameterization issues in mixture modelling and their bearing on MCMC algorithms. *Computational Statistics and Data Analysis*, **29**, 325–343.
- Rue, H., Martino, S. and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations (with discussion). *Journal of the Royal Statistical Society B*, **71**, 319–92.
- Scott, D.W. (1992). *Multivariate Density Estimation*. New York: John Wiley.
- Skilling, J. (2006). Nested sampling for general Bayesian computation. *Bayesian Analysis*, **1**, 833–860.
- Spiegelhalter, D., Thomas, A., Best, N. and Lunn, D. (2003). *WinBUGS User Manual, Version 1.4*. UK: MRC Biostatistics Unit, Institute of Public Health and Department of Epidemiology and Public Health, Imperial College School of Medicine. Available at <http://www.mrc-bsu.cam.ac.uk/bugs/winbugs/manual14.pdf>
- Steele, R.J., Raftery, A.E. and Emond, M.J. (2006). Computing normalizing constants for finite mixture models via incremental mixture importance sampling. *Journal of Computational and Graphical Statistics*, **15**, 712–734.
- Tanner, M.A. and Wong, W. (1987). The calculation of posterior distributions by data augmentation (with discussion). *Journal of the American Statistical Association*, **82**, 528–550.

- Vitoratou, S., Ntzoufras, I. and Moustaki I. (2013). Explaining the behavior of joint and marginal Monte Carlo estimators in latent variable models with independence assumptions. [arXiv:1311.0656\[stat.CO\]](#) (submitted).
- Weinberg, M.D. (2012). Computing the Bayes factor from a Markov chain Monte Carlo simulation of the posterior distribution. *Bayesian Analysis*, **7**, 737–770.
- West, M. (1992). Modelling with mixtures (with discussion). In *Bayesian Statistics*, 4, eds. J.M. Bernardo, J.O. Berger, A.P. Dawid and A.F.M. Smith, Oxford, U.K.: Oxford University Press, pp. 503–524.
- Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g-prior distributions. In *Bayesian Inference and Decision Techniques: Essays in Honour of Bruno de Finetti*, eds. P. Goel and A. Zellner, Amsterdam: North-Holland, pp. 233–243.