Contents lists available at ScienceDirect



Computational Statistics and Data Analysis

www.elsevier.com/locate/csda



Iterative importance sampling with Markov chain Monte Carlo sampling in robust Bayesian analysis



Ivette Raices Cruz^{a,b,*}, Johan Lindström^c, Matthias C.M. Troffaes^d, Ullrika Sahlin^a

^a Centre for Environmental and Climate Science, Lund University, Sweden

^b Department of Biology, Lund University, Sweden

^c Centre for Mathematical Sciences, Lund University, Sweden

^d Department of Mathematical Sciences, Durham University, United Kingdom

ARTICLE INFO

Article history: Received 20 May 2021 Received in revised form 20 June 2022 Accepted 23 June 2022 Available online 2 July 2022

Dataset link: https:// github.com/Iraices/IIS_MCMC

Keywords: Bounds on probability Effective sample size Meta-analysis Random effects model Uncertainty quantification

ABSTRACT

Bayesian inference under a set of priors, called robust Bayesian analysis, allows for estimation of parameters within a model and quantification of epistemic uncertainty in quantities of interest by bounded (or imprecise) probability. Iterative importance sampling can be used to estimate bounds on the quantity of interest by optimizing over the set of priors. A method for iterative importance sampling when the robust Bayesian inference relies on Markov chain Monte Carlo (MCMC) sampling is proposed. To accommodate the MCMC sampling in iterative importance sampling, a new expression for the effective sample size of the importance sampling is derived, which accounts for the correlation in the MCMC samples. To illustrate the proposed method for robust Bayesian analysis, iterative importance sampling with MCMC sampling is applied to estimate the lower bound of the overall effect in a previously published meta-analysis with a random effects model. The performance of the method compared to a grid search method and under different degrees of prior-data conflict is also explored.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

Bayesian analysis quantifies uncertainty by precise probability derived from a prior (subjective) distribution for parameters and a likelihood for data given parameters (Gelman et al., 2013). Whereas statistical Bayesian inference usually uses non-informative priors as default, there are exceptions motivating the use of informative priors to reduce complexity (Simpson et al., 2017). In a decision context where one wants to use the best possible knowledge, informative priors are useful, or even needed, to integrate data with expert knowledge. Specifying a precise informative prior may be difficult, in particular for a model with many parameters or when experts disagree (Rinderknecht et al., 2012; Insua et al., 2000).

Robust Bayesian analysis is a way to consider the impact of the choice of prior on uncertainty in relevant quantities. The impact of different priors in Bayesian inference is important to evaluate for two reasons. First, it is common that more than one prior probability distribution could reasonably be chosen for the problem at hand. Second, when information in

https://doi.org/10.1016/j.csda.2022.107558 0167-9473/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

^{*} Corresponding author at: Centre for Environmental and Climate Science, Lund University, S ölvegatan 37, 22362 Lund, Sweden.

E-mail addresses: ivette.raices_cruz@cec.lu.se (I. Raices Cruz), johan.lindstrom@matstat.lu.se (J. Lindström), matthias.troffaes@durham.ac.uk (M.C.M. Troffaes), ullrika.sahlin@cec.lu.se (U. Sahlin).

data is weak (e.g. for small sample sizes), the choice of prior could matter a lot for the final outcome of an analysis. Robust Bayesian analysis has been used for sensitivity analysis towards the choice of prior (Berger, 1990).

A type of robust Bayesian analysis is to use sets of prior distributions or sets of likelihoods resulting in sets of posterior distributions. This can be seen as an extension of Bayesian inference which quantifies uncertainty by bounded (imprecise) probability instead of precise probability (Walley, 1991). In robust Bayesian analysis, one is often interested in estimating bounds on expectation. For instance, the lower bound on expectation of a function f with respect to a set of posterior distributions, \mathcal{M} , is expressed as

$$\underline{E}(f) := \inf_{p \in \mathcal{M}} \int f(x) p(x) dx.$$
(1)

Note that a set of posterior distributions is derived from a set of prior distributions.

Robust Bayesian analysis using sets of priors has been developed in a closed analytic form for conjugate models (Bernard, 2005; Quaeghebeur and de Cooman, 2005; Walley et al., 1996). In (Wei and Jiang, 2017), a range of posterior expectations are computed using a Monte Carlo method when considering uncertainty regarding the prior or likelihood.

Importance sampling has also been used to estimate bounds on expectations using independent samples drawn from arbitrary (e.g. not necessarily conjugate) models, as long as the posterior can be analytically evaluated up to a normalization constant (Fetz, 2017; Troffaes, 2017, 2018; Troffaes et al., 2018). In an iterative version of importance sampling, it has been suggested to iteratively change the sampling (also called proposal) distribution of importance sampling, in order to get an effective sample size (i.e. a measure of efficiency) as close as possible to the actual sample size (Troffaes, 2017, 2018). A small effective sample size means that the weights of importance sampling are too imbalanced and thus might be unreliable.

In (Troffaes, 2018), it is also suggested to use the posterior distribution directly as a sampling distribution where possible. The use of the posterior allows, in theory, for the effective sample size to be maximized across iterations. For this reason, the effective sample size of the importance sampling estimator is used as the stopping criterion of iterative importance sampling in (Troffaes, 2018). However, the effective sample size can be very poor if the sampling distribution is not carefully chosen, i.e. if the initial choice of posterior is far from the posterior that, say, minimizes the expectation of the quantity of interest. Further, the method can have issues with convergence across iterations as a result. The required number of samples for accurate estimation using importance sampling has also been discussed in (Agapiou et al., 2017; Chatterjee and Diaconis, 2018; Sanz-Alonso, 2018) by means of the Kullback Leibler divergence.

Markov chain Monte Carlo (MCMC) sampling is a method for Bayesian inference which does not require a closed form of the posterior (Brooks et al., 2011; Gelman et al., 2013). MCMC sampling allows for inference of simple as well as complex models. So far, few robust Bayesian analyses have used MCMC sampling (see (Vernon and Gosling, 2017) for an example). This is mainly due to limitations in existing methods, such as requirements on knowing the analytical form of the posteriors. The inability to use MCMC sampling severely restricts the use of robust Bayesian analysis on more complex models.

To use iterative importance sampling with MCMC samples there is a need to modify the effective sample size that is used in the stopping criterion in iterative importance sampling. In this paper, we combine iterative importance sampling with MCMC sampling by extending the method from (Troffaes, 2017, 2018) to a wider range of models, specifically those requiring MCMC sampling. To accomplish this, we derive an expression for the effective sample size which accounts for correlated MCMC samples.

In (Liesenfeld and Richard, 2008), an efficient importance sampling is proposed for improving a MCMC algorithm. The efficient importance sampling consists of selecting a proposal distribution (given a density kernel) using a least squares problem and then using the proposed distribution in an independent Metropolis Hasting sampling. Moreover, in (Llorente et al., 2021), a layered adaptive importance sampling algorithm is presented which combined MCMC algorithms with importance sampling and different strategies to re-use generated samples. The layered adaptive importance sampling algorithm generates samples using two layers. The upper layer generates samples using a MCMC algorithm which are later used in a multiple importance sampling scheme (lower layer) (Llorente et al., 2021). A recycling layered adaptive importance sampling scheme is presented which re-uses the samples from the upper layer in the lower layer (Llorente et al., 2021). This scheme is similar to the method proposed in this paper. However, a key difference between the (Liesenfeld and Richard, 2008) and (Llorente et al., 2021) papers and this paper is that we re-use the MCMC samples from one run by weighting with a different prior, and that we use an efficient sample size for determining when to re-run the MCMC sampler, which leads to fewer MCMC runs to cover the hyperparameter space.

A robust Bayesian analysis allows for a quantification of uncertainty in quantities of interest which is robust to the choice of prior. It can also be useful for Bayesian inference when priors are given as sets. To demonstrate the proposed method, iterative importance sampling with MCMC sampling is applied to estimate the lower bound of the overall effect of biomanipulation of freshwater lakes in an already published meta-analysis (Bernes et al., 2015).

2. Importance sampling

Recall that the expectation of a function f with respect to a *target distribution*, p is given by

$$\mu := \mathsf{E}_p(f(x)) = \int f(x) p(x) dx,$$

which can be approximated by the standard Monte Carlo estimator of $E_p(f(x))$ as

$$\overline{\mu} \coloneqq \frac{1}{N} \sum_{i=1}^{N} f(X_i), \tag{3}$$

where $X_i \sim p$ are independent and identically distributed (i.i.d.) samples.

Sometimes, it is difficult to draw samples directly from p or there is a mismatch between f(x) and p(x) (Owen, 2013). In this case, importance sampling could be applied to estimate the expectation by weighting samples drawn from a sampling distribution, q, from which it is easier to generate samples (Egloff and Leippold, 2010). This technique has been applied in Bayesian inference (Gelman et al., 2013) and in numerical integration (Liu, 2008; Owen, 2013).

The sampling density function q must be such that q(x) > 0 whenever $p(x) f(x) \neq 0$. Sometimes p is only known up to a normalization constant, say only $p_u = cp$ is known, where c > 0 is an unknown constant. The expectation of f with respect to *p* can be written as

$$\mu \coloneqq \mathsf{E}_p\Big(f(x)\Big) = \frac{\int f(x)p_u(x)dx}{\int p_u(x)dx} = \frac{\int f(x)w_p(x)q(x)dx}{\int w_p(x)q(x)dx} = \frac{\mathsf{E}_q\Big(f(x)w_p(x)\Big)}{\mathsf{E}_q\Big(w_p(x)\Big)},\tag{4}$$

where $w_p(x) := \frac{p_u(x)}{q(x)}$. This expectation can be estimated by *self-normalized importance sampling*, which is defined as

$$\widetilde{\mu} \coloneqq \frac{\sum_{i=1}^{N} f(X_i) w_p(X_i)}{\sum_{i=1}^{N} w_p(X_i)} \quad \text{where } X_i \sim q \text{ (i.i.d)}.$$
(5)

In the following, we shall relax the assumption of independence. At this point, it suffices to point out that eq. (5) is still a valid estimator of $E_p(f(x))$ even if the X_i are dependent.

A measure to assess the quality of the importance sampling estimator is the effective sample size, ESS. It is defined by (Kong, 1992), as the ratio of the variances of $\overline{\mu}$ and $\widetilde{\mu}$ estimators, in eq. (3) and eq. (5), scaled to N:

$$\mathrm{ESS} \coloneqq \frac{N\mathrm{Var}_p(\overline{\mu})}{\mathrm{Var}_q(\widetilde{\mu})}.$$
(6)

The effective sample size represents the number of standard Monte Carlo samples that are needed for both estimators to have the same variance.

If the samples X_i from q are i.i.d. then ESS can be estimated by

$$ESS_{IS} := \frac{\left(\sum_{i=1}^{N} w_p(X_i)\right)^2}{\sum_{i=1}^{N} w_p^2(X_i)},$$
(7)

see (Owen, 2013). However, this formula is not applicable for correlated samples, as would be the case if the X_i are sampled from *q* using an MCMC algorithm.

3. Effective sample size of importance sampling using MCMC

Here, we want to derive an estimate of the effective sample size as defined in eq. (6) when the X_i are sampled from q through MCMC. First, we derive an approximation of $Var_q(\tilde{\mu})$.

Importance sampling with MCMC was introduced by (Hastings, 1970) for variance reduction (Bhattacharya, 2008). Hastings (Hastings, 1970) suggested the following approximate of the denominator in eq. (6):

$$\operatorname{Var}_{q}(\widetilde{\mu}) = \operatorname{Var}_{q}\left(\frac{\overline{Y}}{\overline{Z}}\right) \approx \frac{\operatorname{Var}_{q}(\overline{Y} - \mu \overline{Z})}{\left(\operatorname{E}_{q}(\overline{Z})\right)^{2}},\tag{8}$$

where

$$\overline{Y} := \frac{1}{N} \sum_{i=1}^{N} f(X_i) w_p(X_i),$$

$$\overline{Z} := \frac{1}{N} \sum_{i=1}^{N} w_p(X_i),$$
(10)

and $\mu := E_p(f(x))$ as defined earlier.

First, let us evaluate the denominator in eq. (8). Note that

$$E_q\left(w_p(X_i)\right) = \int w_p(x)q(x)dx = \int \frac{cp(x)}{q(x)}q(x)dx = c,$$
(11)

so $E_q(\overline{Z}) = c$, and we can approximate the denominator via

$$\left(\mathsf{E}_{q}(\overline{Z})\right)^{2} \approx \left(\frac{1}{N}\sum_{i=1}^{N}w_{p}(X_{i})\right)^{2}.$$
(12)

The numerator in eq. (8) is more tricky. Let

 $g(X_i) \coloneqq (f(X_i) - \mu) w_p(X_i).$ ⁽¹³⁾

With this notation, we get

$$\operatorname{Var}_{q}(\overline{Y} - \mu \overline{Z}) = \frac{1}{N^{2}} \left(\sum_{i=1}^{N} \operatorname{Var}_{q} \left(g(X_{i}) \right) + 2 \sum_{i < j} \operatorname{Cov}_{q} \left(g(X_{i}), g(X_{j}) \right) \right).$$
(14)

If the variables X_i form a stationary stochastic process, as in a converged MCMC algorithm, then $\operatorname{Var}_q(g(X_i))$ and $\operatorname{Cov}_q(g(X_i), g(X_{i+k}))$ depend only on k and not i. Hence, with $X \coloneqq X_1$,

$$\operatorname{Var}_{q}(\overline{Y} - \mu \overline{Z}) = \operatorname{Var}_{q}\left(g(X)\right)\left(\frac{1 + 2\sum_{k=1}^{N-k} \rho_{g}(k)}{N}\right),\tag{15}$$

where ρ_g is the autocorrelation function of the stationary process. According to (Geyer, 1992), for large *N*, the variance is approximately

$$\operatorname{Var}_{q}(\overline{Y} - \mu \overline{Z}) \approx \operatorname{Var}_{q}\left(g(X)\right) \left(\frac{1 + 2\sum_{k=1}^{\infty} \rho_{g}(k)}{N}\right).$$
(16)

If $\sum_{k=1}^{\infty} \rho_g(k)$ converges, then a standard approximation is (Geyer, 1992; Givens and Hoeting, 2012)

$$\sum_{k=1}^{\infty} \rho_g(k) \approx \sum_{k=1}^{\ell} \hat{\rho}_g(k), \tag{17}$$

where ℓ is the first index for which $\hat{\rho}_g(\ell + 1) < 0$, and where $\hat{\rho}_g$ is the empirical autocorrelation function of the sample $g(X_1), \ldots, g(X_N)$. Note that evaluating g requires knowledge of μ which is precisely the quantity we wish to estimate. So, instead we will use

$$\tilde{g}(X_i) \coloneqq (f(X_i) - \widetilde{\mu}) w_p(X_i), \tag{18}$$

as an approximation for $g(X_i)$, and therefore use

$$\sum_{k=1}^{\infty} \rho_g(k) \approx \sum_{k=1}^{\ell} \hat{\rho}_{\tilde{g}}(k).$$
(19)

We now estimate $\operatorname{Var}_q(g(X))$ in eq. (16). Using eq. (13), we also get that

$$\operatorname{Var}_{q}\left(g(X)\right) = \operatorname{Var}_{q}\left(f(X)w_{p}(X)\right) - 2\mu\operatorname{Cov}_{q}\left(f(X)w_{p}(X), w_{p}(X)\right) + \mu^{2}\operatorname{Var}_{q}\left(w_{p}(X)\right).$$

$$(20)$$

Now, following the same ideas as in (Elvira et al., 2018)[p. 5-6] (see Appendix A for details) we get

$$\operatorname{Var}_{q}\left(g(X)\right) \approx N\operatorname{Var}_{p}(\overline{\mu})\left(\frac{1}{N}\sum_{i=1}^{N}w_{p}^{2}(X_{i})\right).$$
(21)

Putting everything together, we get

I. Raices Cruz, J. Lindström, M.C.M. Troffaes et al.

=

$$\operatorname{Var}_{q}(\widetilde{\mu}) \approx \operatorname{NVar}_{p}(\overline{\mu}) \frac{\frac{1}{N} \sum_{i=1}^{N} w_{p}^{2}(X_{i})}{\left(\frac{1}{N} \sum_{i=1}^{N} w_{p}(X_{i})\right)^{2}} \left(\frac{1 + 2 \sum_{k=1}^{\ell} \hat{\rho}_{\tilde{g}}(k)}{N}\right), \tag{22}$$

$$N^{2} \operatorname{Var}_{n}(\overline{\mu})$$

$$= \frac{1}{\text{ESS}_{\text{IS}} \cdot \text{ESS}_{\text{MCMC}}},$$
(23)

where ESS_{IS} is the standard estimate of the effective sample size for importance sampling with independent samples, eq. (7), and ESS_{MCMC} is the MCMC effective sample size (for \tilde{g}),

$$\mathrm{ESS}_{\mathrm{MCMC}} \coloneqq \frac{N}{1 + 2\sum_{k=1}^{\ell} \hat{\rho}_{\tilde{g}}(k)}.$$

Substituting eq. (23) into eq. (6), we finally obtain the following estimate for the combined ESS

$$ESS \approx \frac{ESS_{MCMC}}{N} \cdot ESS_{IS}.$$
(24)

So, what is new in eq. (24) with respect to eq. (7) is the factor $\frac{\text{ESS}_{MCMC}}{N}$ which accounts for a reduction in effective sample size due to the correlation of the MCMC samples (a reduction since it is very unlikely that the correlation will be negative).

4. Importance sampling over a set of probability distributions

Let $\mathcal{M} \subset \mathbb{R}^d$ be a compact set and $p_t(\cdot) = \{p(\cdot|t) | t \in \mathcal{M}\}$ be a probability density function parameterized by t (i.e. hyperparameters). The lower expectation of a function f with respect to p_t for all $t \in \mathcal{M}$ is assumed to exist and is estimated by

$$\underline{E}(f) \coloneqq \min_{t \in \mathcal{M}} \int f(x) p_t(x) dx.$$
(25)

In practice, one can search for the minimum using numerical methods. For instance, an iterative version of standard importance sampling has been used to estimate bounds on expectations in robust Bayesian analysis (Fetz, 2017; Troffaes, 2017, 2018; Troffaes et al., 2018).

Using importance sampling, the lower expectation is estimated by

$$\widehat{\underline{E}}(f) \approx \min_{t \in \mathcal{M}} \frac{\sum_{i=1}^{N} f(X_i) w_t(X_i)}{\sum_{i=1}^{N} w_t(X_i)},\tag{26}$$

where $w_t(x) \coloneqq \frac{cp_t(x)}{q(x)}$.

Iterative importance sampling estimates the lower expectation of a function f by moving the sampling distribution, q, towards the optimal distribution (Fetz, 2017; Troffaes, 2017). The stopping criterion is that the effective sample size of importance sampling should be close enough to the desired independent sample size (denoted by ESS_{target}) that is fixed in advance. In this paper, we adapt iterative importance sampling introduced in (Fetz, 2017; Troffaes, 2017, 2018). Our contribution is the effective sample size of importance sampling with correlated MCMC samples which allows us to combine iterative importance sampling with MCMC sampling, thus allowing for robust analysis of more complex models. It is important to highlight that how large the sample size for MCMC samples should be and ESS_{target} in step 2 and step 4 respectively, are values fixed in advance (i.e. they are inputs to the procedure). The method goes as follows:

- Set $t = t^0$ where t^0 is an initial value in the feasible region. Step 1
- Generate samples from $q(x) = p_t(x)$ using MCMC sampling until the effective sample size for the MCMC sample is Step 2 large enough (i.e. exceed a specified threshold).
- Find $t_* = \arg \min_{t \in \mathcal{M}} \frac{\sum_{i=1}^{N} f(X_i) w_t(X_i)}{\sum_{i=1}^{N} w_t(X_i)}$ using an optimization algorithm. If ESS > ESS_{target}, or maximum number of iterations reached, then stop. Step 3
- Step 4
- Step 5 Set $t = t_*$ and go to Step 2.

The convergence of the method depends on the distributions and the parameter space. We have set a maximum number of iterations (i.e. in our case, 10 000) to stop the algorithm when it does not converge.

The condition of a large enough sample size is added in Step 2 to ensure that the optimization in Step 3 is based on a reliable sample. For example, in the application below, we first specify the ESS_{target} and then we require the effective sample size for the MCMC samples to be 20% greater than the ESStarget. The stopping criterion in Step 4 uses the effective sample size for importance sampling with correlated samples eq. (24), controlling for both convergence of the MCMC and quality of importance sampling. The reason for this stopping criterion is that MCMC sampling might require different numbers of iteration to produce a reasonable number of efficient samples.

Thinning chains in MCMC has been used in several papers; see for instance (Croll, 2006; Endo et al., 2019; Kopylev et al., 2009). Thinning consists of taking every k-th sample instead of all of them in order to reduce autocorrelation. In (Link and Eaton, 2012) it is shown that although thinning chains in MCMC reduces autocorrelation between MCMC samples, it also reduces the precision of the estimates (i.e. the average over a thinned sample set has greater variance than the average over the unthinned sample) (Geyer, 1992). Therefore, thinning is not advisable unless it is needed due to computer memory limitations (Link and Eaton, 2012).

5. An application

To illustrate the proposed method for robust Bayesian analysis, iterative importance sampling with MCMC sampling is applied on a previously published meta-analysis investigating the effect of biomanipulation (the intervention) on water quality in freshwater lakes (Bernes et al., 2015). We selected the random effects model (described below) for the meta-analysis of the change in the level of *Chlorophyll a* before and during biomanipulation (Bernes et al., 2015). Available data are estimated mean differences and estimation errors from 75 studies. The estimated effects range from -24.17 to $332.50 \ \mu g/l$ with a sample mean of 28.46 $\mu g/l$. The 5th and 95th percentile of the data are -11.10 and 76.29 $\mu g/l$ respectively. Here, a positive value corresponds to an improvement in water quality by biomanipulation (since we have turned the sign of the data). In addition, we investigate what happens with the performance of the suggested method when the set of priors is changed from a set with low to high prior-data conflict.

5.1. A Bayesian linear random effects model

The overall effect of the intervention μ is estimated by a linear random effects model (Fig. 1) according to

$$y_i|\delta_i \sim N(\delta_i, \sigma_i^2), \tag{27}$$

$$\delta_i | \mu, k, \tau_\mu \sim N(\mu, k^2 \tau_\mu^2), \tag{28}$$

where y_i is the observed intervention effect in study i (i = 1, ..., N) with known within-study variance σ_i^2 , δ_i is the specific intervention effect in study i, and $k^2 \tau_{\mu}^2$ is the between-study variance. This model can be expressed in its marginal form (Gelman et al., 2013; Röver, 2017) as

$$y_i|\mu, k, \tau_{\mu} \sim N(\mu, \sigma_i^2 + k^2 \tau_{\mu}^2).$$
 (29)

To implement this model in a Bayesian framework, we select the following prior distributions for the parameters

$$\mu | \tau_{\mu} \sim N(\mu_0, \tau_{\mu}^2), \tag{30}$$

$$\tau_{\mu} \sim U(\tau_l, \tau_0), \tag{31}$$

$$k \sim U(k_l, k_u), \tag{32}$$

where μ_0 is the mean of the overall intervention effect, τ_{μ} is the standard deviation of the overall intervention effect which ranges between τ_l and τ_0 , k is a proportionality constant which ranges between k_l and k_u . We let $\tau_l = 1$, $k_l = 1$ and $k_u = 5$.

The linear random effects model is implemented using MCMC sampling in Stan through the *rstan* package, the R interface to Stan (Stan Development Team, 2018) (see code availability for Stan code).

5.2. Selecting a set of priors

To expand the linear random effects model into a robust Bayesian framework, we consider a set of prior distributions for μ and τ_{μ} .

There are several methods to elicit priors from experts (O'Hagan et al., 2006; Hanea et al., 2021; Hartmann et al., 2020; Gosling, 2018), but there is no obvious alternative to specify a set of priors. The approach that we use to specify a set of priors for multiple parameters is chosen to illustrate robust Bayesian analysis with IIS and MCMC sampling, and other approaches could have been used as well. The prior distributions for each parameter are assumed to belong to the same family of probability distributions (eq. (30) and eq. (31)), but with different hyperparameters. In order to consider interaction between parameters, the specification of priors is made using an approach similar to prior predictive check (Daimon, 2008; Schad et al., 2019). A compact set of hyperparameters is selected by comparing the distribution of the intervention effect for a random study conditional on the hyperparameters $\delta^* | \mu, \tau_{\mu}$ to an elicited range *R*.

The selection of a set of priors can be summarized as:

- 1. Specify a range *R* where the effect size of a randomly selected study is expected to fall with a probability of at least h% (i.e. target coverage).
- 2. Specify a regular grid of hyperparameters (μ_0 , τ_0).



Fig. 1. Linear random effects graphical model. Unknown quantities (parameters) are represented by ellipses, known quantities (priors) by circles and observed data by filled squares. The plate indicates repeated cases.



Fig. 2. Prior check. Selection of priors, a curve fitted to the area where more than 90% of the generated specific intervention effect fell inside of the elicited range.

- 3. For each combination of hyperparameters (μ_0 , τ_0), generate a random sample δ_1^* , ..., δ_M^* from eq. (28).
- 4. Identify hyperparameters where the proportion of generated samples falling inside *R* exceeds the target coverage *h*.
- 5. Find a function that discriminates hyperparameters complying with the target coverage and use the function to select a compact set of hyperparameters.

Here, we select two sets of priors that represent situations with a low and high prior-data conflict, respectively. In actual applications, priors should be elicited by structured expert judgement (O'Hagan et al., 2006).

A set of priors that represents a situation with low prior-data conflict is derived using an elicited range of specific intervention effect of R = [-20, 80], a target coverage of 90%, and a regular grid of 200×200 of hyperparameters ($-100 \le \mu_0 \le 100$ and $5 \le \tau_0 \le 50$). Following the procedure described above yields to the set

$$\mathcal{M} = \left\{ \begin{array}{l} -8 \le \mu_0 \le 68\\ 5 \le \tau_0 \le 16\\ r(\mu_0, \tau_0) - 0.9 \ge 0 \end{array} \right\}$$

where $r(\mu_0, \tau_0) = \tau_0 - (-0.01\mu_0^2 + 0.46\mu_0 + 9.56)$ is a function used to capture hyperparameters fulfilling the target coverage (Fig. 2).

5.3. Setting up the iterative importance sampling

The quantity of interest is the expected value of the overall intervention effect (μ) on the linear random effects model given in section 5.1. Iterative importance sampling is applied to estimate the lower bound of the quantity of interest over



Fig. 3. From left, effective sample sizes ESS, ESS_{IS} and ESS_{MCMC} given MCMC samples using hyperparameters $\mu_0 = 5$ and $\tau_0 = 7$ (the dot). Here, ESS_{target} is 5 000 and if the optimization in Step 3 moves the hyperparameters into a region with a lower ESS we need to re-run the MCMC in Step 2, using the 'new' hyperparameters.

the compact domain of priors \mathcal{M} (see Appendix C for the proof of the existence of the minimum in our example). The lower bound of μ is

$$\widehat{\underline{E}}(\mu|\mathbf{y},\mu_{0},\tau_{0}) \approx \min_{(\mu_{0},\tau_{0})\in\mathcal{M}} \frac{\sum_{i=1}^{N} \mu^{(i)} w(X^{(i)};\mu_{0},\tau_{0})}{\sum_{i=1}^{N} w(X^{(i)};\mu_{0},\tau_{0})},$$
(33)

where $X^{(i)} = \{\mu^{(i)}, k^{(i)}, \tau^{(i)}_{\mu}\}$ and $\mu^{(i)}$ are samples drawn from a sampling distribution using MCMC based on hyperparameters μ'_0 and τ'_0 . The weights are given by

$$w(X^{(i)};\mu_0,\tau_0) = \frac{p(\mu^{(i)},k^{(i)},\tau_{\mu}^{(1)}|\mathbf{y},\mu_0,\tau_0)}{p(\mu^{(i)},k^{(i)},\tau_{\mu}^{(i)}|\mathbf{y},\mu_0',\tau_0')},$$
(34)

where $p(\mu, k, \tau_{\mu} | \mathbf{y}, \mu_0, \tau_0)$ is the posterior distribution of the target distribution corresponding to hyperparameters μ_0 and τ_0 .

Using eq. (34) as weights in eq. (24) gives the effective sample size of importance sampling with MCMC samples. The search for the lower bound in Step 3 will stop when the optimal prior is close to the prior used in the MCMC (Fig. 3). How close is determined by the assigned ESS_{target} in Step 4, which in our example is set to 5 000 and 10 000 samples. The effective sample size for the MCMC samples in Step 2 is set to exceed ESS_{target} by at least 20%, i.e. MCMC sampling in Step 2 will be run until an effective sample size of 6 000 and 12 000 samples has been reached.

The optimization in Step 3 is performed by simulated annealing, (Du and Swamy, 2016; Henderson et al., 2003), using the *optim* function from the *optimx* package (Nash and Varadhan, 2011) in R.

5.4. The lower bound of the expected overall effect

We illustrate the method using two different $\text{ESS}_{\text{target}}$ and two choices of initial values ($\mu'_0 = -7$, $\tau'_0 = 6$) and ($\mu'_0 = 10$, $\tau'_0 = 10$). The number of iterations of iterative importance sampling ranges between 2 and 3 with an effective sample size between 11 640 and 14 305, see (Table 1, 2 and 3). Note that the effective samples size for the MCMC sample is derived based on the hyperparameters in the previous iteration.

The estimated lower bound on the expected overall effect over \mathcal{M} is 12.1 µg/l (Table 1, 2 and 3). We also run the method using different random seeds which gives an estimated lower bound on the expected overall effect between 12.07 and 12.24 µg/l. The bound is lower than 19.9 which is the estimated effect in a standard Bayesian analysis using a flat prior centered at zero ($\mu_0 = 0$ and $\tau_0 = 1$ 000). The lower bound is obtained at the corner of the prior region (i.e. $\mu_0 = -8$ and

Table 1

Summary of iterative importance sampling (IIS) of the lower expected value over the set M, ESS_{target} = 5 000 and initial values (μ_0 ; τ_0) = (-7; 6).

| | Iter | Samples | μ_{0_*} | $	au_{0_*}$ | ESS | ESS _{MCMC} | ESSIS | $\hat{\mu}_{(\mu_{0_*},\tau_{0_*})}$ | Time |
|-------------|-------------|-----------------------|------------------------|---------------------|-------------------|-----------------------|-------------------|--------------------------------------|----------------------|
| Grid Search | - | 20 000 | -8 | 5 | - | - | - | 12.091 | 11.45 h |
| IIS | 0 1 2 | - 20 000 20 000 | -7 -7.706 -7.999 | 6 5.051 5.045 | - 27 11 640 | - 18 005 12 307 | - 30 18 916 | - 10.988 12.166 | – – 11.78 mins |

Table 2

Summary of iterative importance sampling (IIS) of the lower expected value over the set \mathcal{M} and ESS_{target} = 10 000 and initial values (μ_0 ; τ_0) = (-7; 6).

| Samples | μ_{0_*} | τ_{0_*} | ESS | ESS _{MCMC} | ESSIS | $\hat{\mu}_{(\mu_{0_*}, \tau_{0_*})}$ | Time |
|---------|---------------------------------|--|---|---|---|---|--|
| 20 000 | -8 | 5 | - | - | - | 12.091 | 11.45 h |
| | | | | | | | |
| - | -7 | 6 | - | - | - | - | - |
| 40 000 | -7.900 | 5.035 | 59 | 31 579 | 75 | 12.006 | - |
| 20 000 | -7.998 | 5.136 | 13 899 | 13 901 | 19 997 | 12.135 | 18.73 mins |
| | 20 000 - 40 000 20 000 | Samples μ_{0_*} 20 000 -8 - -7 40 000 -7.900 20 000 -7.998 | Samples μ_{0_*} τ_{0_*} 20 000 -8 5 - -7 6 40 000 -7.900 5.035 20 000 -7.998 5.136 | Samples μ_{0_*} τ_{0_*} ESS 20 000 -8 5 - - -7 6 - 40 000 -7.900 5.035 59 20 000 -7.998 5.136 13 899 | Samples μ_{0_*} τ_{0_*} ESS ESS _{MCMC} 20 000 -8 5 - - - -7 6 - - 40 000 -7.900 5.035 59 31 579 20 000 -7.998 5.136 13 899 13 901 | Samples μ_{0_*} τ_{0_*} ESS ESS _{MCMC} ESS _{IS} 20 000 -8 5 - - - - -7 6 - - - 40 000 -7.900 5.035 59 31 579 75 20 000 -7.998 5.136 13 899 13 901 19 997 | Samples μ_{0_*} τ_{0_*} ESS ESS _{MCMC} ESS $\mu_{(\mu_{0_*}, \tau_{0_*})}$ 20 000 -8 5 - - - 12.091 - -7 6 - - - - 40.000 - 7.5 12.006 20 000 -7.998 5.136 13.899 13.901 19.997 12.135 |

Table 3

Summary of iterative importance sampling (IIS) of the lower expected value over the set \mathcal{M} and ESS_{target} = 10 000 and initial values (μ_0 ; τ_0) = (10; 10).

| | Iter | Samples | μ_{0_*} | τ_{0_*} | ESS | ESS _{MCMC} | ESS _{IS} | $\hat{\mu}_{(\mu_{0_*},\tau_{0_*})}$ | Time |
|-------------|------|---------|-------------|--------------|--------|---------------------|-------------------|--------------------------------------|------------|
| Grid Search | - | 20 000 | -8 | 5 | - | - | - | 12.091 | 11.45 h |
| | | | | | | | | | |
| IIS | 0 | - | 10 | 10 | - | - | - | - | - |
| | 1 | 40 000 | -6.734 | 5.148 | 5 | 40 000 | 5 | 12.762 | - |
| | 2 | 20 000 | -7.993 | 5.008 | 3 995 | 13 907 | 5 746 | 12.145 | - |
| | 3 | 20 000 | -7.995 | 5.221 | 14 305 | 14 305 | 20 000 | 12.103 | 23.72 mins |
| | | | | | | | | | |

 $\tau_0 = 5$). The bound obtained by iterative importance sampling is compared to one estimated by a grid search method across a regular grid of hyperparameters (which yields a total of 4 773 hyperparameter values). The grid search method gives a lower expected bound of 12.09 µg/l which is close to the one found using iterative importance sampling. It takes much longer time to run a grid search method, over 10 hours, compared to the iterative importance sampling taking roughly 15 to 25 minutes (Table 1, 2 and 3).

Moreover, we run the optimization algorithm (simulated annealing) over the set \mathcal{M} , but without the resampling step, (e.g. just re-run the MCMC for each parameter that is evaluated). The method gives a lower bound of 12.13 µg/l and it takes much longer time to run, over 20 hours.

The iterative importance sampling converges with fewer iterations when the initial values are relative close to the hyperparameters corresponding to the lower bound.

5.5. The influence of prior data conflict

In order to illustrate what happens when the set of prior has a high conflict with data, we specify a new set M' using the procedure described in subsection 5.2, but with R = [30, 100]. This range does not include the median (14.6 µg/l) of the observed effects in the 75 studies or the posterior mean from a standard Bayesian analysis with a flat prior. This gives the set

$$\mathcal{M}' = \left\{ \begin{array}{c} 42 \le \mu_0 \le 88 \\ 5 \le \tau_0 \le 11 \\ r'(\mu_0, \tau_0) - 0.90 \ge 0 \end{array} \right\},$$

where $r'(\mu_0, \tau_0) = \tau_0 - (-0.011\mu_0^2 + 1.427\mu_0 - 35.104)$.

Shifting the prior region from \mathcal{M} to \mathcal{M}' , (i.e. towards higher values of μ), increases the lower bound of the expected overall effect from 12.1 to 26.8 µg/l, (Table 4). The lower bound is in this case obtained for a less precise prior $\tau_0 = 10.5$.

We also compare the bound obtained by iterative importance sampling to the one estimated by a grid search method across a regular grid of hyperparameters (which yields a total of 1 649 hyperparameter values). The grid search method gives a slightly lower expected bound of 26.7 μ g/l, but takes much longer time, 4 hours, compared to 13 minutes for the iterative importance sampling (Table 4).

Table 4

Summary of iterative importance sampling of the lower expected values over the set \mathcal{M}' and ESS_{target} = 10 000 and initial values (μ_0 ; τ_0) = (60; 11).

| | Iter | Samples | μ_{0_*} | $	au_{0_*}$ | ESS | ESS _{MCMC} | ESSIS | $\hat{\mu}_{(\mu_{0_*},\tau_{0_*})}$ | Time |
|-------------|------|-------------|-------------|-------------|-------------|---------------------|-------------|--------------------------------------|----------------|
| Grid Search | - | 20 000 | 57 | 10.5 | - | - | - | 26.754 | 3.70 h |
| lic | 0 | | 60 | 11 | | | | | |
| 115 | 1 | - 40 000 | 59.064 | 10.859 | - 15 371 | - 17 533 | - 35 067 | 26.811 | - 12.5 mins |

6. Conclusions

Robust Bayesian analysis (i.e. Bayesian inference under a set of priors) offers a way to quantify epistemic uncertainty by bounded instead of precise probability. This type of analysis is useful for evaluating sensitivity to the choice of prior. It is also a solution for Bayesian inference when prior information is upfront given as a set of distributions.

We have derived an expression for the effective sample size of importance sampling with correlated MCMC samples, which ensures reliable samples from both MCMC and importance sampling procedures when combining them. The combination of iterative importance sampling with MCMC sampling was used for robust Bayesian analysis on an existing meta-analysis based on a random effects model (Bernes et al., 2015). The estimated lower bound on the expected overall effect of the intervention in the meta-analysis is a conservative estimate compared to what would have been the result from selecting one prior in the set and using a standard Bayesian analysis.

Iterative importance sampling with MCMC sampling allows for robust Bayesian analysis on a wider range of models not limited to conjugate models. The flexibility in the choice of model may allow for more applications of robust Bayesian analysis. The method was demonstrated on a relatively simple model with two different choices of prior sets. It would be useful to evaluate the proposed method on more complex models to further explore the theoretical and practical challenges associated with robust Bayesian analysis; as well as to investigate how other discrepancy measures such as those proposed by (Martino et al., 2017) and the Kullback-Leibler (KL) divergence measure could be used.

Funding

This work was supported by the Swedish Research Council Formas through the project "Scaling up uncertain environmental evidence" (219-2013-1271) and the strategic research areas BECC (Biodiversity and Ecosystem Services in a Changing Climate) and MERGE (Modelling the Regional and Global Climate/Earth system).

Code availability

The Stan and R codes to run the analysis, except the data, are available through the following link: https://github.com/ Iraices/IIS_MCMC.

Acknowledgements

We thank Claes Bernes for supporting this work with data on the meta-analysis in the evidence synthesis example.

Appendix A. Calculations

Explicit calculations for eq. (20) are given here. Recall eq. (20)

$$\operatorname{Var}_{q}(g(X)) = \operatorname{Var}_{q}((f(X) - \mu)w_{p}(X))$$

We will use that, for any random variable h(X),

$$E_q(h(X)w_p(X)) = \int h(x)w_p(x)q(x)dx = \int h(x)\frac{cp(x)}{q(x)}q(x)dx$$
$$= c\int h(x)p(x)dx = cE_p(h(X)),$$

which gives

$$\operatorname{Var}_{q}\left(h(X)w_{p}(X)\right) = c\operatorname{E}_{p}\left(h^{2}(X)w_{p}(X)\right) - c^{2}\operatorname{E}_{p}\left(h(X)\right)^{2}$$
(A.1)

Taking $h(X) = f(X) - \mu$ and noting that $E_p(h(X)) = 0$, the expression in eq. (20) expands as follows

I. Raices Cruz, J. Lindström, M.C.M. Troffaes et al.

$$\operatorname{Var}_{q}\left(g(X)\right) = \operatorname{Var}_{q}\left((f(X) - \mu)w_{p}(X)\right) = c\operatorname{E}_{p}\left((f(X) - \mu)^{2}w_{p}(X)\right)$$

Applying a second order delta method to the last expectation at $E_p(w_p(X))$ and $E_p(h(X))$, gives (see appendix B for details)

$$\operatorname{Var}_{q}\left(g(X)\right) \approx c\left(\operatorname{E}_{p}\left(h(X)\right)^{2}\operatorname{E}_{p}\left(w_{p}(X)\right) + \operatorname{E}_{p}\left(h(X)\right)\operatorname{Cov}_{p}\left(w_{p}(X), h(X)\right) + \operatorname{E}_{p}\left(w_{p}(X)\right)\operatorname{Var}_{p}\left(h(X)\right)\right)$$
(A.2)

$$= c \mathbf{E}_p \left(w_p(X) \right) \mathbf{E}_p \left(h^2(X) \right)$$
(A.3)

$$= c E_p \Big((f(X) - \mu)^2 \Big) E_p \Big(w_p(X) \Big)$$
(A.4)

$$= c \operatorname{Var}_{p} \left(f(X) \right) \operatorname{E}_{p} \left(w_{p}(X) \right).$$
(A.5)

Using eq. (A.5) and the equality $\operatorname{Var}_p(f(X)) = N\operatorname{Var}_p(\overline{\mu})$ we obtain

$$\operatorname{Var}_{q}(g(X)) \approx N\operatorname{Var}_{p}(\overline{\mu})\operatorname{cE}_{p}\left(w_{p}(X)\right)$$
(A.6)

$$= N \operatorname{Var}_{p}(\overline{\mu}) \operatorname{E}_{q}\left(w_{p}^{2}(X)\right)$$
(A.7)

$$= N \operatorname{Var}_{p}(\overline{\mu}) \left(\frac{1}{N} \sum_{i=1}^{N} w_{p}^{2}(X_{i}) \right).$$
(A.8)

Here we have used that eq. (A.1) with $h(X) = w_p(X)$ gives the equality $cE_p(w_p(X)) = E_q(w_p^2(X))$.

Appendix B. Delta method

Let g(u, v) be a function twice differentiable at $(u, v) = (a_1, a_2)$. Then, the second order Taylor polynomial for g(u, v) near the point $(u, v) = (a_1, a_2)$ is:

$$g(u, v)_{(a_1, a_2)} = g(a_1, a_2) + \frac{\partial g}{\partial u}(a_1, a_2)(u - a_1) + \frac{\partial g}{\partial v}(a_1, a_2)(v - a_2) + \frac{1}{2}\frac{\partial^2 g}{\partial u^2}(a_1, a_2)(u - a_1)^2 + \frac{1}{2}\frac{\partial^2 g}{\partial u \partial v}(a_1, a_2)(u - a_1)(v - a_2) + \frac{1}{2}\frac{\partial^2 g}{\partial v^2}(a_1, a_2)(v - a_2)^2.$$
(B.1)

Now, we can use the second order Taylor polynomial approximation to estimate the mean (this is also known as second order delta method).

Let *U* and *V* be random variables with mean $\theta_1 = E_p(U)$ and $\theta_2 = E_p(V)$ respectively.

$$\begin{split} \mathbf{E}_{p}\Big(g(U,V)\Big)_{(\theta_{1},\theta_{2})} &\approx g(\theta_{1},\theta_{2}) + \frac{\partial g}{\partial u}(\theta_{1},\theta_{2})\mathbf{E}_{p}(U-\theta_{1}) + \frac{\partial g}{\partial v}(\theta_{1},\theta_{2})\mathbf{E}_{p}(V-\theta_{2}) \\ &+ \frac{1}{2}\frac{\partial^{2} g}{\partial u^{2}}(\theta_{1},\theta_{2})\mathbf{E}_{p}\Big((U-\theta_{1})^{2}\Big) \\ &+ \frac{1}{2}\frac{\partial^{2} g}{\partial u \partial v}(\theta_{1},\theta_{2})\mathbf{E}_{p}\Big((U-\theta_{1})(V-\theta_{2})\Big) \\ &+ \frac{1}{2}\frac{\partial^{2} g}{\partial v^{2}}(\theta_{1},\theta_{2})\mathbf{E}_{p}\Big((V-\theta_{2})^{2}\Big). \end{split}$$
(B.2)

Note that $E_p(U - \theta_1) = 0$ and $E_p(V - \theta_2) = 0$. Taking $g(u, v) := v^2 u$ where V = h(X) and $U = w_p(X)$ (to match notation in eq. (A.2))

$$E_{p}\left(g(U,V)\right)_{(\theta_{1},\theta_{2})} \approx g(\theta_{1},\theta_{2}) + \theta_{2}\operatorname{Cov}_{p}(U,V) + \theta_{1}\operatorname{Var}_{p}(V)$$

$$\approx \left(E_{p}(V)\right)^{2}E_{p}(U) + E_{p}(V)\operatorname{Cov}_{p}(U,V) + E_{p}(U)\operatorname{Var}_{p}(V)$$

$$\approx \left(E_{p}\left(h(X)\right)\right)^{2}E_{p}\left(w_{p}(X)\right) + E_{p}\left(h(X)\right)\operatorname{Cov}_{p}\left(w_{p}(X),h(X)\right)$$
(B.3)

I. Raices Cruz, J. Lindström, M.C.M. Troffaes et al.

$$+ E_p(w_p(X)) \operatorname{Var}_p(h(X)). \tag{B.4}$$

Appendix C. Proof of existence of minimum

To guarantee the minimum exist in our example, it is enough to prove that the expectation is continuous as a function of (μ_0, τ_0) on $[-8, 68] \times [5, 16]$.

By Fubini's theorem in our example, we have

$$E_{(\mu_0,\tau_0)}(\mu) \coloneqq \int_{1}^{\tau_0} \int_{-\infty}^{5} \int_{-\infty}^{+\infty} \mu \cdot p_{(\mu_0,\tau_0)}(\mu,\tau_{\mu},k) \, d\mu \, dk \, d\tau_{\mu},$$
(C.1)

where

$$p_{(\mu_0,\tau_0)}(\mu,\tau_{\mu},k) = \frac{1}{c(\mu_0,\tau_0)} \cdot \left(\prod_{i=1}^N \frac{1}{\sqrt{\sigma_i^2 + k^2 \tau_{\mu}^2}}\right) \exp\left\{-\frac{1}{2} \sum_{i=1}^N \frac{(y_i - \mu)^2}{\sigma_i^2 + k^2 \tau_{\mu}^2}\right\} \cdot \frac{1}{\tau_{\mu}} \exp\left\{-\frac{1}{2} \frac{(\mu - \mu_0)^2}{\tau_{\mu}^2}\right\}$$

is the posterior probability distribution. The proportionality constant $c(\mu_0, \tau_0)$ is given by

$$c(\mu_{0},\tau_{0}) = \int_{1}^{\tau_{0}} \int_{1}^{5} \left(\prod_{i=1}^{N} \frac{1}{\sqrt{\sigma_{i}^{2} + k^{2} \tau_{\mu}^{2}}} \right) \frac{1}{\tau_{\mu}}$$
(C.2)
$$\left[\int_{-\infty}^{+\infty} \exp\left\{ -\frac{1}{2} \sum_{i=1}^{N} \frac{(y_{i} - \mu)^{2}}{\sigma_{i}^{2} + k^{2} \tau_{\mu}^{2}} \right\} \cdot \exp\left\{ -\frac{1}{2} \frac{(\mu - \mu_{0})^{2}}{\tau_{\mu}^{2}} \right\} d\mu \right] dk \, d\tau_{\mu}$$
$$- \int_{0}^{\tau_{0}} \int_{1}^{5} \left(\prod_{i=1}^{N} \frac{1}{-1} \right) \frac{1}{2}$$
(C.3)

$$= \int_{1}^{+} \int_{1}^{+} \left(\prod_{i=1}^{1} \frac{1}{\sqrt{\sigma_{i}^{2} + k^{2} \tau_{\mu}^{2}}} \right)^{\frac{1}{\tau_{\mu}}}$$
(C.3)
$$\left[\int_{-\infty}^{+\infty} \exp\left\{ -\frac{1}{2} \left(a(\tau_{\mu}, k, \mu_{0}) \mu^{2} - 2b(\tau_{\mu}, k, \mu_{0}) \mu + c(\tau_{\mu}, k, \mu_{0}) \right) \right\} d\mu \right] dk \, d\tau_{\mu}$$

where

$$a(\tau_{\mu}, k, \mu_{0}) = \sum_{i=1}^{N} \frac{1}{\sigma_{i}^{2} + k^{2} \tau_{\mu}^{2}} + \frac{1}{\tau_{\mu}^{2}},$$
(C.4)

$$b(\tau_{\mu}, k, \mu_{0}) = \sum_{i=1}^{N} \frac{y_{i}}{\sigma_{i}^{2} + k^{2} \tau_{\mu}^{2}} + \frac{\mu_{0}}{\tau_{\mu}^{2}},$$
(C.5)

$$c(\tau_{\mu}, k, \mu_{0}) = \sum_{i=1}^{N} \frac{y_{i}^{2}}{\sigma_{i}^{2} + k^{2} \tau_{\mu}^{2}} + \frac{\mu_{0}^{2}}{\tau_{\mu}^{2}}.$$
(C.6)

Since $a(\tau_{\mu}, k, \mu_0) > 0$, then by Lemma Appendix D.1 in (C.2), we get that the integral with respect to μ is

$$\sqrt{2\pi} \cdot a(\tau_{\mu}, k, \mu_{0})^{\frac{-1}{2}} \cdot \exp\left\{-\frac{1}{2}\left(c(\tau_{\mu}, k, \mu_{0}) - \frac{b(\tau_{\mu}, k, \mu_{0})^{2}}{a(\tau_{\mu}, k, \mu_{0})}\right)\right\}.$$
(C.7)

Note that $a(\tau_{\mu}, k, \mu_0)$, $b(\tau_{\mu}, k, \mu_0)$ and $c(\tau_{\mu}, k, \mu_0)$ are continuously differentiable and $\tau_{\mu} \in [1, \tau_0]$ and $k \in [1, 5]$. Thus, by Leibniz integral rule it follows that $c(\mu_0, \tau_0)$ is continuously differentiable on $[-8, 68] \times [5, 16]$.

Analogously, applying Lemma Appendix D.2 to eq. (C.1), we get that the integral with respect to μ is

$$\sqrt{2\pi} \cdot a(\tau_{\mu}, k, \mu_{0})^{\frac{-3}{2}} \cdot b(\tau_{\mu}, k, \mu_{0}) \cdot \exp\left\{-\frac{1}{2}\left(c(\tau_{\mu}, k, \mu_{0}) - \frac{b(\tau_{\mu}, k, \mu_{0})^{2}}{a(\tau_{\mu}, k, \mu_{0})}\right)\right\}.$$
(C.8)

Furthermore, by Leibniz integral rule it yields that $E_{(\mu_0,\tau_0)}(\mu)$ is continuously differentiable on $[-8, 68] \times [5, 16]$.

(D.2)

Appendix D. Lemmas

Lemma Appendix D.1. Let $a > 0, b \in \mathbb{R}$ and $c \in \mathbb{R}$. It holds that:

$$\int_{-\infty}^{+\infty} \exp\left\{-\frac{1}{2}(ax^2 - 2bx + c)\right\} dx \tag{D.1}$$

$$= \sqrt{2\pi} e^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}\left(e^{-\frac{b^2}{2}}\right)\right\} \tag{D.2}$$

$$=\sqrt{2\pi} \cdot a^{-2} \cdot \exp\left\{-\frac{1}{2}\left(c - \frac{1}{a}\right)\right\}.$$

Proof. By completing the square, we obtain

$$\int_{-\infty}^{+\infty} \exp\left\{-\frac{1}{2}\left(\frac{x-\frac{b}{a}}{a^{-\frac{1}{2}}}\right)^2\right\} \cdot \exp\left\{-\frac{1}{2}\left(c-\frac{b^2}{a}\right)\right\} dx \tag{D.3}$$

$$=\sqrt{2\pi} \cdot a^{-\frac{1}{2}} \cdot \exp\left\{-\frac{1}{2}\left(c-\frac{b^{2}}{a}\right)\right\} \cdot \int_{-\infty}^{+\infty} \frac{1}{\sqrt{2\pi} \cdot a^{-\frac{1}{2}}} \exp\left\{-\frac{1}{2}\left(\frac{x-\frac{b}{a}}{a^{-\frac{1}{2}}}\right)^{2}\right\} dx.$$
 (D.4)

Note that the integrand is the probability density function of a normally distributed random variable with mean $\frac{b}{a}$ and variance $\frac{1}{a}$, (i.e. $N\left(\frac{b}{a}, \frac{1}{a}\right)$). Thus, the desired result immediately follows. \Box

Lemma Appendix D.2. Let $a > 0, b \in \mathbb{R}$ and $c \in \mathbb{R}$. It holds that:

$$\int_{-\infty}^{+\infty} x \exp\left\{-\frac{1}{2}(ax^2 - 2bx + c)\right\} dx \tag{D.5}$$

$$=\sqrt{2\pi} \cdot a^{-\frac{3}{2}} \cdot b \cdot \exp\left\{-\frac{1}{2}\left(c - \frac{b^2}{a}\right)\right\}.$$
 (D.6)

Proof. By completing the square, we obtain

$$\int_{-\infty}^{+\infty} x \cdot \exp\left\{-\frac{1}{2}\left(\frac{x-\frac{b}{a}}{a^{-\frac{1}{2}}}\right)^2\right\} \cdot \exp\left\{-\frac{1}{2}\left(c-\frac{b^2}{a}\right)\right\} dx \tag{D.7}$$

$$=\sqrt{2\pi} \cdot a^{-\frac{1}{2}} \cdot \exp\left\{-\frac{1}{2}\left(c - \frac{b^2}{a}\right)\right\} \cdot \int_{-\infty}^{+\infty} \frac{1}{\sqrt{2\pi} \cdot a^{-\frac{1}{2}}} x \exp\left\{-\frac{1}{2}\left(\frac{x - \frac{b}{a}}{a^{-\frac{1}{2}}}\right)^2\right\} dx.$$
 (D.8)

Note that the previous integral is the expected value of a normally distributed random variable with mean $\frac{b}{a}$ and variance $\frac{1}{a}$, (i.e. $N\left(\frac{b}{a}, \frac{1}{a}\right)$). Thus, the desired result immediately follows. \Box

References

Bernard, J.-M., 2005. An introduction to the imprecise Dirichlet model for multinomial data. Int. J. Approx. Reason. 39, 123-150.

Brooks, S., Gelman, A., Jones, G.L., Meng, X.-L. (Eds.), 2011. Handbook of Markov Chain Monte Carlo. Chapman and Hall/CRC, New York.

Agapiou, S., Papaspiliopoulos, O., Sanz-Alonso, D., Stuart, A.M.S., 2017. Importance sampling: intrinsic dimension and computational cost. Stat. Sci. 32, 405-431. https://doi.org/10.1214/17-STS611.

Berger, J.O., 1990. Robust Bayesian analysis: sensitivity to the prior. J. Stat. Plan. Inference 25, 303-328.

Bernes, C., Carpenter, S.R., Gårdmark, A., Larsson, P, Persson, L., Skov, C., Speed, J.D., Donk, E.V., 2015. What is the influence of a reduction of planktivorous and benthivorous fish on water quality in temperate eutrophic lakes? A systematic review. Environmen. Evid. 4, 7. https://doi.org/10.1186/s13750-015-0032-9.

Bhattacharya, S., 2008. Consistent estimation of the accuracy of importance sampling using regenerative simulation. Stat. Probab. Lett. 78, 2522-2527. https://doi.org/10.1016/j.spl.2008.02.030.

Chatterjee, S., Diaconis, P., 2018. The sample size required in importance sampling. Ann. Appl. Probab. 28, 1099–1135. https://doi.org/10.1214/17-AAP1326. Croll, B., 2006. Markov chain Monte Carlo methods applied to photometric spot modeling. Publ. Astron. Soc. Pac. 118, 1351-1359.

Daimon, T., 2008. Predictive checking for Bayesian interim analyses in clinical trials. Contemp. Clin. Trials 29, 740–750. https://doi.org/10.1016/j.cct.2008.05. 005.

Du, K.-L., Swamy, M., 2016. Search and Optimization by Metaheuristics. Techniques and Algorithms Inspired by Nature. Chapman & Hall/CRC Texts in Statistical Science. Springer International Publishing, Switzerland.

Egloff, D., Leippold, M., 2010. Quantile estimation with adaptive importance sampling. Ann. Stat. 38, 1244–1278.

Elvira, V., Martino, L., Robert, C.P., 2018. Rethinking the effective sample size. arXiv:1809.04129 [stat.CO].

Endo, A., van Leeuwen, E., Baguelin, M., 2019. Introduction to particle Markov chain Monte Carlo for disease dynamics modellers. Epidemics 29, 100363. https://doi.org/10.1016/j.epidem.2019.100363.

Fetz, T., 2017. Efficient computation of upper probabilities of failure. In: Bucher, C., Ellingwood, B.R., Frangopol, D.M. (Eds.), 12th International Conference on Structural Safety & Reliability, pp. 493–502.

Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A., Rubin, D.B., 2013. Bayesian Data Analysis, third edition. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis.

Geyer, C.J., 1992. Practical Markov chain Monte Carlo. Stat. Sci. 7, 473-483.

Givens, G.H., Hoeting, J.A., 2012. Computational Statistics, second edition. John Wiley & Sons.

Gosling, J.P., 2018. SHELF: the Sheffield elicitation framework. In: Elicitation: The Science and Art of Structuring Judgement. In: International Series in Operations Research & Management Science, vol. 261. Springer International Publishing, pp. 61–93.

Hanea, A.M., Hemming, V., Nane, G.F., 2021. Uncertainty quantification with experts: present status and research needs. Risk Anal. https://doi.org/10.1111/ risa.13718.

Hartmann, M., Agiashvili, G., Bürkner, P., Klami, A., 2020. Flexible prior elicitation via the prior predictive distribution. In: Peters, J., Sontag, D. (Eds.), Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI). In: PMLR Proceedings of Machine Learning Research, vol. 124, pp. 1129–1138. Hastings, W., 1970. Monte Carlo sampling methods using Markov chains and their applications. Biometrika 57, 97–109.

Henderson, D., Jacobson, S.H., Johnson, A.W., 2003. Handbook of Metaheuristics. The Theory and Practice of Simulated Annealing. Springer US.

Insua, D.R., Ruggeri, F., Martín, J., 2000. Bayesian sensitivity analysis. In: Saltelli, A., Chan, K., Scott, E.M. (Eds.), Sensitivity Analysis. Wiley, pp. 225–244.

Kong, A., 1992. A note on Importance Sampling using standardized weights. Technical Report 348. University of Chicago, Dept. of Statistics. Kopylev, L., Fox, J., Chen, C., 2009. Combining risks from several tumors using Markov chain Monte Carlo. In: Uncertainty Modeling in Dose Response: Bench Testing Environmental Toxicity. John Wiley & Sons, pp. 197–205.

Liesenfeld, R., Richard, J.-F., 2008. Improving MCMC, using efficient importance sampling. Comput. Stat. Data Anal. 53, 272–288. https://doi.org/10.1016/j. csda.2008.07.028.

Link, W.A., Eaton, M.J., 2012. On thinning of chains in MCMC. Methods Ecol. Evol. 3, 112–115. https://doi.org/10.1111/j.2041-210X.2011.00131.x.

Liu, J.S., 2008. Monte Carlo Strategies in Scientific Computing. Springer, New York.

Llorente, F., Curbelo, E., Martino, L., Elvira, V., Delgado, D., 2021. MCMC-driven importance samplers. https://doi.org/10.48550/ARXIV.2105.02579.

Martino, L., Elvira, V., Louzada, F., 2017. Effective sample size for importance sampling based on discrepancy measures. Signal Process. 131, 386–401. https:// doi.org/10.1016/j.sigpro.2016.08.025.

Nash, J.C., Varadhan, R., 2011. Unifying optimization algorithms to aid software system users: optimx for R. J. Stat. Softw. 43, 1-14.

O'Hagan, A., Buck, C.E., Daneshkhah, A., Eiser, J.R., Garthwaite, P.H., Jenkinson, D.J., Oakley, J.E., Rakow, T., 2006. Uncertain Judgements: Eliciting Experts' Probabilities. Wiley, Chichester.

Owen, A.B., 2013. Monte Carlo theory, methods and examples. http://statweb.stanford.edu/~owen/mc/.

Quaeghebeur, E., de Cooman, G., 2005. Imprecise probability models for inference in exponential families. In: Cozman, F.G., Nau, R., Seidenfeld, T. (Eds.), ISIPTA'05: Proceedings of the Fourth International Symposium on Imprecise Probabilities and Their Applications. Pittsburgh, USA, pp. 287–296. http:// www.sipta.org/isipta05/proceedings/019.html.

Rinderknecht, S.L., Borsuk, M.E., Reichert, P., 2012. Bridging uncertain and ambiguous knowledge with imprecise probabilities. Environ. Model. Softw. 36, 122–130. https://doi.org/10.1016/j.envsoft.2011.07.022.

Röver, C., 2017. Bayesian random-effects meta-analysis using the bayesmeta R package. arXiv:1711.08683 [stat.CO].

Sanz-Alonso, D., 2018. Importance sampling and necessary sample size: an information theory approach. SIAM/ASA J. Uncertain. Quantificat. 6, 867–879. https://doi.org/10.1137/16M1093549.

Schad, D.J., Betancourt, M., Vasishth, S., 2019. Toward a principled Bayesian workflow in cognitive science. arXiv:1904.12765 [stat.ME].

Simpson, D.P., Rue, H., Martins, T.G., Riebler, A., Sørbye, S.H., 2017. Penalising model component complexity: a principled, practical approach to constructing priors. Stat. Sci. 32, 1–28.

Stan Development Team. RStan: the R interface to Stan. R package version 2.17.4.

Troffaes, M.C.M., 2017. A note on imprecise Monte Carlo over credal sets via importance sampling. In: Antonucci, A., Corani, G., Couso, I., Destercke, S. (Eds.), Proceedings of the Tenth International Symposium on Imprecise Probability: Theories and Applications. In: PMLR Proceedings of Machine Learning Research, vol. 62, pp. 325–332.

Troffaes, M.C.M., 2018. Imprecise Monte Carlo simulation and iterative importance sampling for the estimation of lower previsions. Int. J. Approx. Reason. 101, 31–48. https://doi.org/10.1016/j.ijar.2018.06.009. arXiv:1806.10404.

Troffaes, M.C.M., Fetz, T., Oberguggenberger, M., 2018. Iterative importance sampling for estimating expectation bounds under partial probability specifications. In: The 8th International Workshop on Reliable Engineering Computing (REC2018).

Vernon, I., Gosling, J.P., 2017. A Bayesian computer model analysis of robust Bayesian analyses. arXiv:1703.01234.

Walley, P., 1991. Statistical Reasoning with Imprecise Probabilities. Chapman & Hall, London.

Walley, P., Gurrin, L., Burton, P., 1996. Analysis of clinical data using imprecise prior probabilities. J. R. Stat. Soc. 45, 457-485.

Wei, W., Jiang, W., 2017. On Monte Carlo computation of posterior expectations with uncertainty. J. Stat. Comput. Simul. 87, 2038–2049. https://doi.org/10. 1080/00949655.2017.1311895.