Vol. 16 (2022) 5182-5225
ISSN: 1935-7524
https://doi.org/10.1214/22-EJS2063

# Sparse linear mixed model selection via streamlined variational Bayes* 

Emanuele Degani<br>Dipartimento di Scienze Statistiche Università degli Studi di Padova, Padova, Italy<br>$e$-mail: degani@stat.unipd.it<br>Luca Maestrini<br>Research School of Finance, Actuarial Studies and Statistics The Australian National University, Canberra, Australia<br>e-mail: luca.maestrini@anu.edu.au

Dorota Toczydłowska and Matt P. Wand<br>School of Mathematical and Physical Sciences<br>University of Technology Sydney, Sydney, Australia<br>e-mail: dtoczydlowska@gmail.com, e-mail: matt.wand@uts.edu.au


#### Abstract

Linear mixed models are a versatile statistical tool to study data by accounting for fixed effects and random effects from multiple sources of variability. In many situations, a large number of candidate fixed effects is available and it is of interest to select a parsimonious subset of those being effectively relevant for predicting the response variable. Variational approximations facilitate fast approximate Bayesian inference for the parameters of a variety of statistical models, including linear mixed models. However, for models having a high number of fixed or random effects, simple application of standard variational inference principles does not lead to fast approximate inference algorithms, due to the size of model design matrices and inefficient treatment of sparse matrix problems arising from the required approximating density parameters updates.

We illustrate how recently developed streamlined variational inference procedures can be generalized to make fast and accurate inference for the parameters of linear mixed models with nested random effects and globallocal priors for Bayesian fixed effects selection. Our variational inference algorithms achieve convergence to the same optima of their standard implementations, although with significantly lower computational effort, memory usage and time, especially for large numbers of random effects. Using simulated and real data examples, we assess the quality of automated procedures for fixed effects selection that are free from hyperparameters tuning and only rely upon variational posterior approximations. Moreover, we show high accuracy of variational approximations against model fitting via Markov Chain Monte Carlo sampling.


MSC2020 subject classifications: Primary 62F15; secondary 62H12, 62J05, 62J07.

[^0]Keywords and phrases: Mean field variational Bayes, multilevel models, longitudinal data analysis, fixed effects selection, global-local shrinkage priors.

Received October 2021.

## 1. Introduction

A variety of statistical models can be formulated as linear regression models incorporating both fixed and random effects in the linear predictor. The former are effects associated with the entire population or repeatable levels of experimental factors; the latter arise from individual experimental units drawn at random. In the statistical literature, models admitting both fixed and random effects are known as mixed-effects models [68]. These models are employed in an assortment of regression-type studies, including the analysis of classical longitudinal data [e.g. 26], repeated measurements [e.g. 84], blocked designs [e.g. 49], multilevel data [e.g. 30], as well as semi-parametric regression models [e.g. 72] such as those including spatial or spline-type components.

The focus of this work is on Bayesian fitting of linear mixed-effects models with nested random effects structures, which are commonly used for the analysis of longitudinal, multilevel and panel data [e.g. 83, 3], or small area estimation [e.g. 70]. These data are typically collected from experimental units that can be grouped into different levels of nesting, and the interest is in modeling within-group correlations. In areas of application such as genome-wide association studies [e.g. 43, 75, 47]) and medical research [e.g. 14], datasets typically possess a large number of group-invariant predictors of which only a few are effectively relevant. A common misleading strategy is that of including all the predictors as fixed effects in the model specification. This may compromise the parsimony of the model specification and validity of inferential conclusions, especially in sparse covariate settings. Therefore, a proper fixed effects selection procedure is recommended to identify the effectively relevant effects.

Although many frequentist procedures have been developed to tackle this problem [e.g. 74, 25, 33, 38, 48]), little exists in the Bayesian literature. Bayesian approaches are mostly focused on random effects selection induced by the decomposition of their covariance matrix [e.g. 20, 89], or joint fixed and random effects selection [e.g. 41, 90].

The current work focuses on fixed effects selection procedures from a Bayesian perspective. This may be advantageous over frequentist approaches especially in high-dimensional settings when likelihood-based inference is computationally intractable and allow for prior knowledge about the parameters to be incorporated in the model specification. Markov Chain Monte Carlo (MCMC) sampling still represents the reference toolkit for exact Bayesian inference and all the aforementioned references on Bayesian approaches for effects selection perform model fitting via MCMC. Although not accounting for selection procedures, the brms package [15] allows to fit Bayesian multilevel models in R [79] making use of the popular probabilistic programming language Stan [17]. Using this package,
practitioners only have to specify the appropriate model structure, the model is automatically fitted and convergence can be assessed. However, automatic sampling procedures usually generate higher computational times and, in general, proper MCMC procedures necessitate convergence assessment for all the model parameters, which may arise problematics such as poor mixing connected to the model parameterization. These and other drawbacks have supported the development of variational approximations for linear mixed models to improve the speed of convergence, at the cost of employing an approximation to the true posterior distribution for carrying out inferential conclusions.
[88] provide some insights on how to implement variational approximations for approximate Bayesian inference in hierarchical models through Infer.NET [57]. Although this computational framework is suitable for longitudinal and multilevel models, its computational advantage quickly decreases for high numbers of groups and sub-groups, limiting the usefulness of variational inference. Algorithm 3 of [65], and Algorithms 3 and 5 of [50] allow to implement variational inference for fitting longitudinal and multilevel data; however, they do not perform efficiently for large dimensions, as they include naïve updates based on inefficient matrix inversions.
[44] developed a streamlined updating scheme for variational inference making efficient use of sub-matrix inversion operations whose number is linear in the size of groups at each level. The streamlined scheme represents an improvement of two orders of magnitude over naïve implementations of variational approximations. These results have also been extended to the class of generalized linear mixed-effects models and applied, for instance, to models for multiple longitudinal markers [37]. [60] took advantage of the sparse matrix results developed in [62] for deriving streamlined algorithms and performing efficient Bayesian variational approximations for linear mixed models with two and three-level nested random effects structures. This framework, named streamlined variational inference, allows to dramatically reduce computational times when compared to naïve implementations of variational inference, although achieving the same approximation. Furthermore, streamlined variational inference allows to efficiently store the matrices needed to perform algorithm updates, hence providing significant memory savings.

These developments have recently inspired streamlined algorithms for linear mixed-effects models with crossed random effects [55] and group-specific curves [56]. Many extensions can be envisaged and are motivated by the high demand for fast and accurate processing methods for big amounts of data from clinical studies, psychological experiments or surveys in social sciences. The current streamlined variational algorithms have been developed and tested using generic uninformative priors over the fixed effects vector. In this work, we introduce streamlined variational inference for models with priors inducing Bayesian posterior shrinkage and study an efficient selection procedure for fixed effects.

### 1.1. Contribution and article organization

To the best of our knowledge, scalable variational approximation methods for fixed effects selection in linear mixed models are rarely present in literature. [1] propose a sparse variational Bayes analysis of linear mixed models which focuses on random effects shrinkage via decomposition of the random effects vector covariance matrix. A more recent contribution is [81], where the suggested approach performs simultaneous fixed-effect selection and parameter estimation via variational Bayes and Bayesian adaptive lasso. However, the approach is limited to high-dimensional two-level generalized linear mixed models and does not account for any streamlined updating improvements.

The current work extends the results and algorithms of [60] by developing streamlined Bayesian variational approximations for multilevel linear mixed models with two or three-level random effects where a subset of fixed effects is subject to selection. The selection is performed by first placing global-local priors over the fixed effects being subject to selection, which ensures good shrinkage properties towards the origin for irrelevant fixed effects marginal posteriors, and then identifying those being relevant via an automated selection procedure free from hyperparameters tuning.

The article is organized as follows. Section 2 provides an overview of linear mixed models from a Bayesian perspective, with a specific focus on twoand three-level random effects specifications. Section 3 explains variational approximations for this class of models, with a particular focus on issues arising from naïve implementations of variational algorithms and the benefits of streamlined variational inference. Section 4 discusses automated approximate Bayesian methods for performing variable selection when global-local shrinkage priors are introduced in a linear regression model. Section 5 connects the previous two sections and provides streamlined variational Bayes algorithms for mixed-effects models with global-local priors placed over a subset of fixed effects which are subject to selection. Section 6 provides a detailed simulation study that demonstrates the advantages provided by the methodology proposed in this work. A real data illustration is included in Section 7. The article is supported by additional supplementary material containing details on distributions, complementary algorithms and derivations [21].

### 1.2. Notation

The notation of this article matches the one of [60]. Here we briefly recall some essential notation. Data vectors and design matrices can be combined using stack and blockdiag operators defined as

$$
\left.\underset{1 \leq i \leq d}{\operatorname{stack}}\left(\boldsymbol{M}_{i}\right) \equiv\left[\begin{array}{c}
\boldsymbol{M}_{1} \\
\vdots \\
\boldsymbol{M}_{d}
\end{array}\right] \quad \text { and } \quad \begin{array}{c}
\text { blockdiag } \\
1 \leq i \leq d
\end{array} \boldsymbol{M}_{i}\right) \equiv\left[\begin{array}{cccc}
\boldsymbol{M}_{1} & \boldsymbol{O} & \cdots & \boldsymbol{O} \\
\boldsymbol{O} & \boldsymbol{M}_{2} & \cdots & \boldsymbol{O} \\
\vdots & \vdots & \ddots & \vdots \\
\boldsymbol{O} & \boldsymbol{O} & \cdots & \boldsymbol{M}_{d}
\end{array}\right]
$$

for a sequence of matrices $\boldsymbol{M}_{1}, \ldots, \boldsymbol{M}_{\boldsymbol{d}}$. The stack operator requires $\boldsymbol{M}_{i}$ matrices with the same number of columns. If $\boldsymbol{M}$ is a square matrix, $\operatorname{diagonal}(\boldsymbol{M})$ is the main diagonal of $\boldsymbol{M}$ and $\operatorname{tr}(\boldsymbol{M})$ is the trace of $\boldsymbol{M}$. For a vector $\boldsymbol{v}$ of length $d, \operatorname{diag}(\boldsymbol{v})$ produces a diagonal matrix having the elements of $\boldsymbol{v}$ as main diagonal. Unless specified otherwise, given two vectors $\boldsymbol{v}_{1}$ and $\boldsymbol{v}_{2}$ of same length, $\boldsymbol{v}_{1} / \boldsymbol{v}_{2}$ indicates their element-wise division. Element-wise addition, subtraction and multiplication are similarly defined.

We use $\mathfrak{p}$ and $\mathfrak{q}$ for density functions. In particular, $\mathfrak{q}$ is used for densities arising from variational approximations. The letters $p$ and $q$ are used for the dimensions of model vectors and matrices.

We use $\mu_{\mathfrak{q}(\theta)} \equiv E_{\mathfrak{q}}(\theta)$ for a generic parameter $\theta, \boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\theta})} \equiv E_{\mathfrak{q}}(\boldsymbol{\theta})$ for a generic vector of parameters $\boldsymbol{\theta}$ and $\boldsymbol{M}_{\mathfrak{q}(\boldsymbol{\Theta})} \equiv E_{\mathfrak{q}}(\boldsymbol{\Theta})$ for a generic matrix of parameters $\boldsymbol{\Theta}$, with $E_{\mathfrak{q}}(\cdot)$ denoting the expectation with respect to the probability density function $\mathfrak{q}$.

## 2. Linear mixed models

This article treats linear mixed models with Gaussian responses and homoskedastic independent errors from a Bayesian inference perspective. A general formulation for these models is

$$
\begin{gather*}
\boldsymbol{y}\left|\boldsymbol{\beta}, \boldsymbol{u}, \sigma^{2} \sim \mathrm{~N}\left(\boldsymbol{X} \boldsymbol{\beta}+\boldsymbol{Z} \boldsymbol{u}, \sigma^{2} \boldsymbol{I}\right), \quad \boldsymbol{u}\right| \boldsymbol{G} \sim \mathrm{N}(\mathbf{0}, \boldsymbol{G}), \quad \boldsymbol{\beta} \sim \mathrm{N}\left(\boldsymbol{\mu}_{\boldsymbol{\beta}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}}\right) \\
\sigma^{2} \mid a_{\sigma^{2}} \sim \text { Inverse- } \chi^{2}\left(\nu_{\sigma^{2}}, 1 / a_{\sigma^{2}}\right), \quad a_{\sigma^{2}} \sim \text { Inverse- } \chi^{2}\left(1,1 /\left(\nu_{\sigma^{2}} s_{\sigma^{2}}^{2}\right)\right)  \tag{1}\\
\boldsymbol{G} \mid \boldsymbol{A}_{\boldsymbol{G}} \sim \mathfrak{p}\left(\boldsymbol{G} \mid \boldsymbol{A}_{G}\right), \quad \boldsymbol{A}_{\boldsymbol{G}} \sim \mathfrak{p}\left(\boldsymbol{A}_{\boldsymbol{G}}\right)
\end{gather*}
$$

where $\boldsymbol{y}$ is a vector of observed data, $\boldsymbol{\beta}$ and $\boldsymbol{u}$ are respectively the vectors of fixed and random effects, $\boldsymbol{X}$ and $\boldsymbol{Z}$ are the associated fixed and random effects design matrices, $\sigma^{2}$ is the variance of the unit-specific error term and $\boldsymbol{G}$ is the random effects covariance matrix.

A very general prior specification for the parameters of model (1) is considered. The fixed effects vector $\boldsymbol{\beta}$ has a multivariate Normal prior with hyperparameters $\boldsymbol{\mu}_{\boldsymbol{\beta}}$ and $\boldsymbol{\Sigma}_{\boldsymbol{\beta}}$. Following [28], the hierarchical prior specification on $\sigma^{2}$ generates a Half- $t$ distribution on $\sigma$ with $\nu_{\sigma^{2}}$ degrees of freedom and scale parameter $s_{\sigma^{2}}$, where larger values of $s_{\sigma^{2}}$ correspond to weaker informativity. A similar hierarchical prior is imposed on the random effects vector covariance matrix $G$ : if, for instance, $\mathfrak{p}\left(\boldsymbol{G} \mid \boldsymbol{A}_{\boldsymbol{G}}\right)$ is an Inverse-G-Wishart $\left(G_{\text {full }}, \nu_{\boldsymbol{G}}+2 q-2, \boldsymbol{A}_{\boldsymbol{G}}^{-1}\right)$ density function and $\mathfrak{p}\left(\boldsymbol{A}_{G}\right)$ is an Inverse-G-Wishart $\left(G_{\text {diag }}, 1, \boldsymbol{\Lambda}_{\boldsymbol{A}_{G}}\right)$ density function
with $\boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{G}}} \equiv\left\{\nu_{\boldsymbol{G}} \operatorname{diag}\left(s_{\boldsymbol{G}, 1}^{2}, \ldots, s_{\boldsymbol{G}, q}^{2}\right)\right\}^{-1}$, then according to [36] such a prior imposition may induce arbitrarily noninformative priors on the standard deviation parameters for large values of $s_{G, 1}, \ldots, s_{G, q}$ and a Uniform $(-1,1)$ distribution over the correlation parameters. The notations $G_{\text {full }}$ and $G_{\text {diag }}$ symbolize fully connected and disconnected graphs arising from the structure of $\boldsymbol{G}^{-1}$ and $\boldsymbol{A}_{\boldsymbol{G}}^{-1}$, as explained in [53].

The structures of $\boldsymbol{Z}, \boldsymbol{u}$ and $\boldsymbol{G}$ embed a rich ensemble of mixed model specifications [92]. Hereafter, we will focus on multilevel models having two-level and three-level random effects specifications.

### 2.1. Two-level linear mixed models

Multilevel models with two-level random effects arise from applications where observations from different units belonging to separate groups are available, and the interest is in capturing the within-group variability. Let $m$ be the number of groups, each composed by $o_{i}$ units, $1 \leq i \leq m$. A two-level linear mixed model can be expressed in terms of the observations from the $i$ th group as follows:

$$
\begin{gather*}
\boldsymbol{y}_{i}\left|\boldsymbol{\beta}, \boldsymbol{u}_{i}, \sigma^{2} \stackrel{\text { ind. }}{\sim} \mathrm{N}\left(\boldsymbol{X}_{i} \boldsymbol{\beta}+\boldsymbol{Z}_{i} \boldsymbol{u}_{i}, \sigma^{2} \boldsymbol{I}\right), \quad \boldsymbol{u}_{i}\right| \boldsymbol{\Sigma} \stackrel{\text { ind. }}{\sim} \mathrm{N}(\mathbf{0}, \boldsymbol{\Sigma}), \quad 1 \leq i \leq m, \\
\boldsymbol{\beta} \sim \mathrm{~N}\left(\boldsymbol{\mu}_{\boldsymbol{\beta}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}}\right), \quad \sigma^{2} \mid a_{\sigma^{2}} \sim \text { Inverse- } \chi^{2}\left(\nu_{\sigma^{2}}, 1 / a_{\boldsymbol{\sigma}^{2}}\right), \\
a_{\sigma^{2}} \sim \text { Inverse- } \chi^{2}\left(1,1 /\left(\nu_{\sigma^{2}} s_{\sigma^{2}}^{2}\right)\right),  \tag{2}\\
\boldsymbol{\Sigma} \mid \boldsymbol{A}_{\boldsymbol{\Sigma}} \sim \text { Inverse-G-Wishart }\left(G_{\text {full }}, \nu_{\boldsymbol{\Sigma}}+2 q-2, \boldsymbol{A}_{\boldsymbol{\Sigma}}^{-1}\right), \\
\boldsymbol{A}_{\boldsymbol{\Sigma}} \sim \text { Inverse-G-Wishart }\left(G_{\text {diag }}, 1, \boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}}}\right), \\
\boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}}} \equiv\left\{\nu_{\boldsymbol{\Sigma}} \operatorname{diag}\left(s_{\boldsymbol{\Sigma}, 1}^{2}, \ldots, s_{\boldsymbol{\Sigma}, q}^{2}\right)\right\}^{-1} .
\end{gather*}
$$

Here $\boldsymbol{\Sigma}$ is the covariance matrix for the group-specific random effects vector $\boldsymbol{u}_{i}$ of length $q$. Notice model (2) is a particular case of model (1), with

$$
\begin{array}{rrr}
\boldsymbol{y}=\underset{1 \leq i \leq m}{\operatorname{stack}}\left(\boldsymbol{y}_{i}\right), & \boldsymbol{X}=\underset{1 \leq i \leq m}{\operatorname{stack}}\left(\boldsymbol{X}_{i}\right), & \boldsymbol{Z}=\underset{1 \leq i \leq m}{\operatorname{blockdiag}}\left(\boldsymbol{Z}_{i}\right),  \tag{3}\\
\boldsymbol{u}=\operatorname{stack}_{1 \leq i \leq m}\left(\boldsymbol{u}_{i}\right), & \boldsymbol{G}=\boldsymbol{I}_{m} \otimes \boldsymbol{\Sigma}, & \boldsymbol{A}_{\boldsymbol{G}}=\boldsymbol{I}_{m} \otimes \boldsymbol{A}_{\boldsymbol{\Sigma}} .
\end{array}
$$

The structure of $\boldsymbol{Z}$ is such that $\boldsymbol{Z} \boldsymbol{u}=\sum_{i=1}^{m} \boldsymbol{Z}_{i} \boldsymbol{u}_{i}$ and notice that as $m$ increases, $\boldsymbol{Z}$ becomes sparser with only the $(100 / m) \%$ of its cells being non-zero.

### 2.2. Three-level linear mixed models

Multilevel models with three-level random effects extend two-level models by adding a further hierarchy level. Such structures are employed when there is interest in capturing both the variability within groups and that within their subgroups.

Let $m$ denote the number of level 1 (L1) groups, $n_{i}$ be the number of level 2 (L2) subgroups belonging to the $i$ th group, $1 \leq i \leq m$, and $o_{i j}$ be the number of
units belonging to the $j$ th subgroup, $1 \leq j \leq n_{i}$, of the $i$ th group. A three-level linear mixed model can be defined in terms of the observations from the $j$ th subgroup belonging to the $i$ th group as follows:

$$
\begin{gather*}
\boldsymbol{y}_{i j} \mid \boldsymbol{\beta}, \boldsymbol{u}_{i}^{\mathrm{L} 1}, \boldsymbol{u}_{i j}^{\mathrm{L} 2}, \sigma^{2} \stackrel{\text { ind. }}{\sim} \mathrm{N}\left(\boldsymbol{X}_{i j} \boldsymbol{\beta}+\boldsymbol{Z}_{i j}^{\mathrm{L} 1} \boldsymbol{u}_{i}^{\mathrm{L} 1}+\boldsymbol{Z}_{i j}^{\mathrm{L} 2} \boldsymbol{u}_{i j}^{\mathrm{L} 2}, \sigma^{2} \boldsymbol{I}\right), \\
{\left.\left[\begin{array}{c}
\boldsymbol{u}_{i}^{\mathrm{L} 1} \\
\boldsymbol{u}_{i j}^{\mathrm{L} 2}
\end{array}\right] \right\rvert\, \boldsymbol{\Sigma}^{\mathrm{L} 1}, \boldsymbol{\Sigma}^{\mathrm{L} 2} \stackrel{\text { ind. }}{\sim} \mathrm{N}\left(\left[\begin{array}{c}
\mathbf{0} \\
\mathbf{0}
\end{array}\right],\left[\begin{array}{cc}
\boldsymbol{\Sigma}^{\mathrm{L} 1} & \boldsymbol{O} \\
\boldsymbol{O} & \boldsymbol{\Sigma}^{\mathrm{L} 2}
\end{array}\right]\right), \quad 1 \leq i \leq m, 1 \leq j \leq n_{i},} \\
\boldsymbol{\beta} \sim \mathrm{~N}\left(\boldsymbol{\mu}_{\boldsymbol{\beta}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}}\right), \quad \sigma^{2} \mid a_{\boldsymbol{\sigma}^{2}} \sim \operatorname{Inverse}-\chi^{2}\left(\nu_{\sigma^{2}}, 1 / a_{\sigma^{2}}\right), \\
a_{\sigma^{2}} \sim \operatorname{Inverse}-\chi^{2}\left(1,1 /\left(\nu_{\sigma^{2}} s_{\sigma^{2}}^{2}\right)\right), \\
\boldsymbol{\Sigma}^{\mathrm{L} 1} \mid \boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}} \sim \operatorname{Inverse-G-Wishart}\left(G_{\text {ful1 }}, \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}+2 q_{1}-2, \boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}^{-1}\right),  \tag{4}\\
\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}} \sim \operatorname{Inverse-G-Wishart}\left(G_{\mathrm{diag}}, 1, \boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}}\right), \\
\boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}} \equiv\left\{\nu_{\boldsymbol{\Sigma}^{\mathrm{L} 1}} \operatorname{diag}\left(s_{\boldsymbol{\Sigma}^{\mathrm{L} 1}, 1}^{2}, \ldots, s_{\boldsymbol{\Sigma}^{\mathrm{L} 1}, q_{1}}^{2}\right)\right\}^{-1}, \\
\boldsymbol{\Sigma}^{\mathrm{L} 2} \mid \boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}} \sim \operatorname{Inverse-G-Wishart}\left(G_{\text {ful1 }}, \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}+2 q_{2}-2, \boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}^{-1}\right), \\
\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}} \sim \operatorname{Inverse-G-Wishart}\left(G_{\mathrm{diag}}, 1, \boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}}\right), \\
\boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}} \equiv\left\{\nu_{\boldsymbol{\Sigma}^{\mathrm{L} 2}} \operatorname{diag}\left(s_{\boldsymbol{\Sigma}^{\mathrm{L} 2}, 1}^{2}, \ldots, s_{\boldsymbol{\Sigma}^{\mathrm{L} 2}, q_{2}}^{2}\right)\right\}^{-1} .
\end{gather*}
$$

Here $\boldsymbol{\Sigma}^{\mathrm{L} 1}$ is the covariance matrix for the group-specific random effects vector $\boldsymbol{u}_{i}^{\mathrm{L} 1}$ of length $q_{1}$ and $\boldsymbol{\Sigma}^{\mathrm{L} 2}$ is that for the subgroup-specific random effects vector $\boldsymbol{u}_{i j}^{\mathrm{L} 2}$ having length $q_{2}$. Notice model (4) is a particular case of model (1), with

$$
\begin{gather*}
\boldsymbol{y}=\underset{1 \leq i \leq m}{\operatorname{stack}}\left(\underset{1 \leq j \leq n_{i}}{\operatorname{stack}}\left(\boldsymbol{y}_{i j}\right)\right), \quad \boldsymbol{X}=\underset{1 \leq i \leq m}{\operatorname{stack}}\left(\underset{1 \leq j \leq n_{i}}{\operatorname{stack}}\left(\boldsymbol{X}_{i j}\right)\right), \\
\boldsymbol{Z}=\underset{1 \leq i \leq m}{\operatorname{blockdiag}}\left(\left[\left.\underset{1 \leq j \leq n_{i}}{\operatorname{stack}}\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 1}\right)\right|_{1 \leq j \leq n_{i}} ^{\operatorname{blockdiag}}\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 2}\right)\right]\right), \\
\boldsymbol{u}=\underset{1 \leq i \leq m}{\operatorname{stack}}\left(\left[\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)^{T} \mid\left(\underset{1 \leq j \leq n_{i}}{\operatorname{stack}}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)\right)^{T}\right]^{T}\right),  \tag{5}\\
\boldsymbol{G}=\underset{1 \leq i \leq m}{\operatorname{blockdiag}}\left(\left[\begin{array}{cc}
\boldsymbol{\Sigma}^{\mathrm{L} 1} & \boldsymbol{O} \\
\boldsymbol{O} & \boldsymbol{I}_{n_{i}} \otimes \boldsymbol{\Sigma}^{\mathrm{L} 2}
\end{array}\right]\right) \\
\boldsymbol{A}_{\boldsymbol{G}}=\underset{1 \leq i \leq m}{\operatorname{blockdiag}}\left(\left[\begin{array}{cc}
\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}} & \boldsymbol{O} \\
\boldsymbol{O} & \boldsymbol{I}_{n_{i}} \otimes \boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}
\end{array}\right]\right)
\end{gather*}
$$

The structure of $\boldsymbol{Z}$ is more involved than the one of two-level random effects models and is such that $\boldsymbol{Z} \boldsymbol{u}=\sum_{i=1}^{m} \sum_{j=1}^{n_{i}}\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 1} \boldsymbol{u}_{i}^{\mathrm{L} 1}+\boldsymbol{Z}_{i j}^{\mathrm{L} 2} \boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)$. As $m$ and the $n_{i}$ 's increase, $\boldsymbol{Z}$ becomes sparser with only the $\left\{\left(q_{1}+q_{2}\right) /\left(q_{1} m+q_{2} \sum_{i=1}^{m} n_{i}\right) \times 100\right\} \%$ of its cells being non-zero. Notice that in the particular case where $n_{i}=1$ for all $1 \leq i \leq m$ and $q_{1}=q_{2}=q$, the three-level specification corresponds to the two-level one.

## 3. Variational Bayesian inference

In this section, we provide a brief overview of mean field variational Bayes, shortly MFVB [see e.g. 9, 10], which is the variational approximation technique we employ in this work for fitting two- and three-level linear mixed models.

### 3.1. Overview

Consider a generic model with an observed data vector $\boldsymbol{y}$ and a parameter vector $\boldsymbol{\theta} \in \boldsymbol{\Theta}$. Let $\mathfrak{q}(\boldsymbol{\theta})$ be an arbitrary density function over the parameter space $\boldsymbol{\Theta}$. Then the logarithm of the marginal likelihood $\mathfrak{p}(\boldsymbol{y})$ satisfies

$$
\begin{equation*}
\log \mathfrak{p}(\boldsymbol{y})=\int_{\boldsymbol{\Theta}} \mathfrak{q}(\boldsymbol{\theta}) \log \left\{\frac{\mathfrak{p}(\boldsymbol{y}, \boldsymbol{\theta})}{\mathfrak{q}(\boldsymbol{\theta})}\right\} \mathrm{d} \boldsymbol{\theta}+\int_{\boldsymbol{\Theta}} \mathfrak{q}(\boldsymbol{\theta}) \log \left\{\frac{\mathfrak{q}(\boldsymbol{\theta})}{\mathfrak{p}(\boldsymbol{\theta} \mid \boldsymbol{y})}\right\} \mathrm{d} \boldsymbol{\theta} \tag{6}
\end{equation*}
$$

where $\mathfrak{p}(\boldsymbol{\theta} \mid \boldsymbol{y}) \equiv \mathfrak{p}(\boldsymbol{y}, \boldsymbol{\theta}) / \mathfrak{p}(\boldsymbol{y})$ is the posterior density function and $\mathfrak{p}(\boldsymbol{y}, \boldsymbol{\theta})$ is the model joint density function. The first addend of (6) is a lower bound to the marginal log-likelihood and it is maximized when the second addend, that is the Kullback-Leibler divergence $\operatorname{KL}\{\mathfrak{q}(\boldsymbol{\theta}) \| \mathfrak{p}(\boldsymbol{\theta} \mid \boldsymbol{y})\}$, is minimized. The lower bound corresponds to $\log \mathfrak{p}(\boldsymbol{y})$ when $\mathfrak{q}(\boldsymbol{\theta})=\mathfrak{p}(\boldsymbol{\theta} \mid \boldsymbol{y})$, although in practical situations exact computation of the posterior density function is infeasible.

The central idea of MFVB is to approximate $\mathfrak{p}(\boldsymbol{\theta} \mid \boldsymbol{y})$ with an approximating density function $\mathfrak{q}(\boldsymbol{\theta})$ that solves the following optimization problem:

$$
\begin{equation*}
\mathfrak{q}^{*}(\boldsymbol{\theta})=\underset{\mathfrak{q}(\boldsymbol{\theta}) \in \mathcal{Q}}{\arg \min } \operatorname{KL}\{\mathfrak{q}(\boldsymbol{\theta}) \| \mathfrak{p}(\boldsymbol{\theta} \mid \boldsymbol{y})\} \tag{7}
\end{equation*}
$$

Tractable solutions arise when $\mathfrak{q}(\boldsymbol{\theta})$ is restricted to some convenient product of densities such that

$$
\mathcal{Q}=\left\{\mathfrak{q}(\boldsymbol{\theta}): \mathfrak{q}(\boldsymbol{\theta})=\prod_{i=1}^{M} \mathfrak{q}_{i}\left(\boldsymbol{\theta}_{i}\right), \text { for some partition }\left\{\boldsymbol{\theta}_{1}, \ldots, \boldsymbol{\theta}_{M}\right\} \text { of } \boldsymbol{\theta}\right\}
$$

It is possible to show [e.g. 65] that under this restriction the optimal $\mathfrak{q}$-density functions satisfy

$$
\begin{equation*}
\mathfrak{q}_{i}^{*}\left(\boldsymbol{\theta}_{i}\right) \propto \exp \left\{E_{\mathfrak{q}\left(-\boldsymbol{\theta}_{i}\right)}\left\{\log \mathfrak{p}\left(\boldsymbol{\theta}_{i} \mid \text { rest }\right)\right\}\right\}, \quad 1 \leq i \leq M \tag{8}
\end{equation*}
$$

where $E_{\mathfrak{q}\left(-\boldsymbol{\theta}_{i}\right)}$ denotes expectation with respect to $\prod_{j \neq i} \mathfrak{q}_{j}\left(\boldsymbol{\theta}_{j}\right)=\mathfrak{q}(\boldsymbol{\theta}) / \mathfrak{q}_{i}\left(\boldsymbol{\theta}_{i}\right)$ and $\mathfrak{p}\left(\boldsymbol{\theta}_{i} \mid\right.$ rest $)$ is the full-conditional density function of $\boldsymbol{\theta}_{i}$. An iterative coordinate ascent procedure can be used to solve (7) and maximize the first addend on the right-hand side of (6) through iterative updates for the optimal approximating densities arising from (8). Tractability is achieved when the updating steps reduce to updates of the approximating densities parameters and this typically occurs when all the $\mathfrak{q}_{i}^{*}\left(\boldsymbol{\theta}_{i}\right)$ 's belong to known families of parametric distributions. Convergence to at least local optima is guaranteed from convexity properties of the lower bound [13] and once it is reached inference can be performed employing the $\mathfrak{q}_{i}^{*}\left(\boldsymbol{\theta}_{i}\right)$ densities in place of the corresponding $\mathfrak{p}\left(\boldsymbol{\theta}_{i} \mid \boldsymbol{y}\right)$ marginal posterior densities.

### 3.2. Nä̈ve variational inference

Assume the posterior density function $\mathfrak{p}\left(\boldsymbol{\beta}, \boldsymbol{u}, \sigma^{2}, a_{\sigma^{2}}, \boldsymbol{G}, \boldsymbol{A}_{\boldsymbol{G}} \mid \boldsymbol{y}\right)$ of the generic linear mixed model (1) is approximated by a $\mathfrak{q}$-density function factorized as follows:

$$
\mathfrak{q}\left(\boldsymbol{\beta}, \boldsymbol{u}, \sigma^{2}, a_{\sigma^{2}}, \boldsymbol{G}, \boldsymbol{A}_{\boldsymbol{G}}\right)=\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u}) \mathfrak{q}\left(\sigma^{2}\right) \mathfrak{q}\left(a_{\sigma^{2}}\right) \mathfrak{q}(\boldsymbol{G}) \mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{G}}\right)
$$

Using arguments from Section 3.1, it is possible to show that the optimal approximating densities that are function of $(\boldsymbol{\beta}, \boldsymbol{u}), \sigma^{2}$ and $a_{\sigma^{2}}$ are the following:

$$
\begin{align*}
& \quad \mathfrak{q}^{*}(\boldsymbol{\beta}, \boldsymbol{u}) \text { is a } \mathrm{N}\left(\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}, \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}\right) \text { density function, } \\
& \mathfrak{q}^{*}\left(\sigma^{2}\right) \text { is an Inverse- } \chi^{2}\left(\xi_{\mathfrak{q}\left(\sigma^{2}\right)}, \lambda_{\mathfrak{q}\left(\sigma^{2}\right)}\right) \text { density function }  \tag{9}\\
& \text { and } \quad \mathfrak{q}^{*}\left(a_{\sigma^{2}}\right) \text { is an Inverse- } \chi^{2}\left(\xi_{\mathfrak{q}\left(a_{\sigma^{2}}\right)}, \lambda_{\mathfrak{q}\left(a_{\sigma^{2}}\right)}\right) \text { density function. }
\end{align*}
$$

Let $\boldsymbol{C} \equiv[\boldsymbol{X} \mid \boldsymbol{Z}]$ and denote with $n$ the number of its rows. The parameters of these densities can be obtained by iteratively performing the updates

$$
\begin{gather*}
\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})} \longleftarrow\left(\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \boldsymbol{C}^{T} \boldsymbol{C}+\left[\begin{array}{cc}
\boldsymbol{\Sigma}_{\boldsymbol{\beta}}^{-1} & \boldsymbol{O} \\
\boldsymbol{O} & E_{\mathfrak{q}}\left(\boldsymbol{G}^{-1}\right)
\end{array}\right]\right)^{-1}, \\
\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})} \longleftarrow \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}\left(\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \boldsymbol{C}^{T} \boldsymbol{y}+\left[\begin{array}{c}
\boldsymbol{\Sigma}_{\boldsymbol{\beta}}^{-1} \boldsymbol{\mu}_{\boldsymbol{\beta}} \\
\mathbf{0}
\end{array}\right]\right),  \tag{10}\\
\xi_{\mathfrak{q}\left(\sigma^{2}\right)} \longleftarrow \nu_{\sigma^{2}}+n, \\
\lambda_{\mathfrak{q}\left(\sigma^{2}\right)} \longleftarrow \mu_{\mathfrak{q}\left(1 / a_{\sigma^{2}}\right)}+\left\|\boldsymbol{y}-\boldsymbol{C} \boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}\right\|^{2}+\operatorname{tr}\left\{\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})} \boldsymbol{C}^{T} \boldsymbol{C}\right\}, \\
\xi_{\mathfrak{q}\left(a_{\sigma^{2}}\right)} \longleftarrow \nu_{\sigma^{2}}+1 \quad \text { and } \quad \lambda_{\mathfrak{q}\left(a_{\sigma^{2}}\right)} \longleftarrow \mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)}+1 /\left(\nu_{\sigma^{2}} s_{\sigma^{2}}^{2}\right)
\end{gather*}
$$

until convergence, where

$$
\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \longleftarrow \xi_{\mathfrak{q}\left(\sigma^{2}\right)} / \lambda_{\mathfrak{q}\left(\sigma^{2}\right)} \quad \text { and } \quad \mu_{\mathfrak{q}\left(1 / a_{\sigma^{2}}\right)} \longleftarrow \xi_{\mathfrak{q}\left(a_{\sigma^{2}}\right)} / \lambda_{\mathfrak{q}\left(a_{\sigma^{2}}\right)}
$$

The approximating densities for the matrices $\boldsymbol{G}$ and $\boldsymbol{A}_{\boldsymbol{G}}$ vary according to the random effects structure considered and are related to the structure of the random effects design matrix. For the two-level random effects specification, the structures of $\boldsymbol{G}$ and $\boldsymbol{A}_{\boldsymbol{G}}$ are given by the second row of (3), which involves matrices $\boldsymbol{\Sigma}$ and $\boldsymbol{A}_{\boldsymbol{\Sigma}}$. Their approximating densities are the following:

$$
\mathfrak{q}^{*}(\boldsymbol{\Sigma}) \text { is an Inverse-G-Wishart }\left(G_{\text {full }}, \xi_{\mathfrak{q}(\boldsymbol{\Sigma})}, \boldsymbol{\Lambda}_{\mathfrak{q}(\boldsymbol{\Sigma})}\right) \text { density function }
$$ and $\quad \mathfrak{q}^{*}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)$ is an Inverse-G-Wishart $\left(G_{\text {diag }}, \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)}, \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)}\right)$ density function.

The parameters of these $\mathfrak{q}$-densities are updated according to:

$$
\begin{aligned}
& \xi_{\mathfrak{q}(\boldsymbol{\Sigma})} \longleftarrow \nu_{\boldsymbol{\Sigma}}+m+2 q-2, \boldsymbol{\Lambda}_{\mathfrak{q}(\boldsymbol{\Sigma})} \longleftarrow \boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}^{-1}\right)}+\sum_{i=1}^{m}\left\{\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)} \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}^{T}+\boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}\right\}, \\
& \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)} \longleftarrow \nu_{\boldsymbol{\Sigma}}+q \quad \text { and } \quad \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)} \longleftarrow \boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}}}+\operatorname{diag}\left\{\operatorname{diagonal}\left(\boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{-1}\right)}\right)\right\},
\end{aligned}
$$

where $\boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}^{-1}\right)} \longleftarrow \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)} \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)}^{-1}$, and $\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}$ and $\boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}$ respectively correspond to the sub-vector of $\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ and sub-matrix of $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ associated with the $i$ th group random effects vector $\boldsymbol{u}_{i}$, for $1 \leq i \leq m$. Also, the update for $E_{\mathfrak{q}}\left(\boldsymbol{G}^{-1}\right)$ appearing in the first line of (10) is

$$
\begin{equation*}
E_{\mathfrak{q}}\left(\boldsymbol{G}^{-1}\right) \longleftarrow \boldsymbol{I}_{m} \otimes \boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{-1}\right)} \tag{11}
\end{equation*}
$$

with $\boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{-1}\right)} \longleftarrow\left(\xi_{\mathfrak{q}(\boldsymbol{\Sigma})}-q+1\right) \boldsymbol{\Lambda}_{\mathfrak{q}(\boldsymbol{\Sigma})}^{-1}$.
For the three-level random effects specification, the structures of $\boldsymbol{G}$ and $\boldsymbol{A}_{\boldsymbol{G}}$ are given by the last row of (5), which involves matrices $\boldsymbol{\Sigma}^{\mathrm{L} 1}, \boldsymbol{\Sigma}^{\mathrm{L} 2}, \boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}$ and $\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}$. Their approximating densities are the following: $\mathfrak{q}^{*}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)$ is an Inverse-G-Wishart $\left(G_{\text {full }}, \xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)}, \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)}\right)$ density function, $\mathfrak{q}^{*}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}\right)$ is an Inverse-G-Wishart $\left(G_{\text {diag }}, \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}\right)}, \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}\right)}\right)$ density function, $\mathfrak{q}^{*}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)$ is an Inverse-G-Wishart $\left(G_{f u l l}, \xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)}, \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)}\right)$ density function and $\mathfrak{q}^{*}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\right)$ is an Inverse-G-Wishart $\left(G_{\text {diag }}, \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\right)}, \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\right)}\right)$ density function.
The parameters of these $\mathfrak{q}$-densities are updated according to:

$$
\begin{gathered}
\xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)} \longleftarrow \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}+m+2 q_{1}-2, \\
\boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)} \longleftarrow \boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}^{-1}\right)}+\sum_{i=1}^{m}\left\{\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)} \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}^{T}+\boldsymbol{\Sigma}_{\left.\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)\right\},},\right. \\
\xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}\right)}^{\leftarrow} \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}+q_{1}, \\
\boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\left.\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)}\right.} \longleftarrow \boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}}+\operatorname{diag}\left\{\operatorname{diagonal}\left(\boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)^{-1}\right)}\right)\right\}, \\
\xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)} \longleftarrow \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}+\sum_{i=1}^{m} n_{i}+2 q_{2}-2, \\
\boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)} \longleftarrow \boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}^{-1}\right)}+\sum_{i=1}^{m} \sum_{j=1}^{n_{i}}\left\{\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)} \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}^{T}+\boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}\right\}, \\
\xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\right)} \longleftarrow \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}+q_{2} \\
\text { and } \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\right)} \longleftarrow \boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}}+\operatorname{diag}\left\{\operatorname{diagonal}\left(\boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)^{-1}\right)}\right)\right\},
\end{gathered}
$$

where

$$
\boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}^{-1}\right)} \longleftarrow \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}\right)} \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}\right)}^{-1}, \quad \boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}^{-1}\right)} \longleftarrow \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\right)} \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\right)}^{-1}
$$

$\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}$ and $\boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}$ respectively correspond to the sub-vector of $\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ and sub-matrix of $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ associated with the $i$ th group random effects vector at level $1 \boldsymbol{u}_{i}^{\mathrm{L} 1}$, and $\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}$ and $\boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}$ respectively correspond to the sub-vector of $\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ and sub-matrix of $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ associated with the $j$ th subgroup of the $i$ th group random effects vector at level $2 \boldsymbol{u}_{i j}^{\mathrm{L} 2}$, for $1 \leq i \leq m$ and $1 \leq j \leq n_{i}$. Furthermore, the update for $E_{\mathfrak{q}}\left(\boldsymbol{G}^{-1}\right)$ appearing in the first line of (10) is

$$
E_{\mathfrak{q}}\left(\boldsymbol{G}^{-1}\right) \longleftarrow \underset{1 \leq i \leq m}{\text { blockdiag }}\left(\left[\begin{array}{cc}
\boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)^{-1}\right)} & \boldsymbol{O}  \tag{12}\\
\boldsymbol{O} & \boldsymbol{I}_{n_{i}} \otimes \boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)^{-1}\right)}
\end{array}\right]\right)
$$

with

$$
\begin{gathered}
\boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)^{-1}\right)} \longleftarrow\left(\xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)}-q_{1}+1\right) \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)}^{-1} \\
\text { and } \quad \boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)^{-1}\right)} \longleftarrow\left(\xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)}-q_{2}+1\right) \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)}^{-1} .
\end{gathered}
$$

The term naïve is used in this work when the MFVB updates described in this section are implemented without exploiting sparse matrix structures. Note, for example, that the update for $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ in (10) involves the inversion of a potentially massive matrix whose sparse structure is induced by those of $\boldsymbol{Z}$ and $\boldsymbol{G}$. As explained in Sections 2.1 and 2.2, when the model dimensions increase such matrices may become extremely sparse and the inversion operation can face many complications, both in terms of memory storage and computational efficiency. By taking advantage of the specific random effects structure it is possible to perform efficient streamlined variational updates.

### 3.3. Streamlined variational inference

The concept of streamlined variational inference for linear mixed models first appears in [44], where the sparse structure of $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ is exploited for efficiently fitting a particular version of model (2) via MFVB. [62] define sparse matrix classes arising from two-level and three-level random effects specifications and provide efficient mathematical solutions to the associated matrix inversion problems in their Theorems 2.2, 2.3, 3.2 and 3.3. [60] implement such results and develop streamlined MFVB algorithms for linear mixed models having both twolevel and three-level random effects specifications. These results are presented as solutions of two- and three-level sparse matrix problems.

Two-level sparse matrix problems are described in Section 2 of [62]. These problems are related to finding the vector $\boldsymbol{x}$ such that $\boldsymbol{A} \boldsymbol{x}=\boldsymbol{a}$, where

$$
\boldsymbol{A} \equiv\left[\begin{array}{ccccc}
\boldsymbol{A}_{11} & \boldsymbol{A}_{12,1} & \boldsymbol{A}_{12,2} & \cdots & \boldsymbol{A}_{12, m} \\
\boldsymbol{A}_{12,1}^{T} & \boldsymbol{A}_{22,1} & \boldsymbol{O} & \cdots & \boldsymbol{O} \\
\boldsymbol{A}_{12,2}^{T} & \boldsymbol{O} & \boldsymbol{A}_{22,2} & \cdots & \boldsymbol{O} \\
\vdots & \vdots & \vdots & \ddots & \vdots \\
\boldsymbol{A}_{12, m}^{T} & \boldsymbol{O} & \boldsymbol{O} & \cdots & \boldsymbol{A}_{22, m}
\end{array}\right], \boldsymbol{a} \equiv\left[\begin{array}{c}
\boldsymbol{a}_{1} \\
\boldsymbol{a}_{2,1} \\
\boldsymbol{a}_{2,2} \\
\vdots \\
\boldsymbol{a}_{2, m}
\end{array}\right] \text { and } \boldsymbol{x} \equiv\left[\begin{array}{c}
\boldsymbol{x}_{1} \\
\boldsymbol{x}_{2,1} \\
\boldsymbol{x}_{2,2} \\
\vdots \\
\boldsymbol{x}_{2, m}
\end{array}\right]
$$

and obtaining the sub-blocks of $\boldsymbol{A}^{-1}$ corresponding to the non-zero blocks of $\boldsymbol{A}$. The structure of $\boldsymbol{A}^{-1}$ is

$$
\boldsymbol{A}^{-1} \equiv\left[\begin{array}{ccccc}
\boldsymbol{A}^{11} & \boldsymbol{A}^{12,1} & \boldsymbol{A}^{12,2} & \cdots & \boldsymbol{A}^{12, m} \\
\boldsymbol{A}^{12,1 T} & \boldsymbol{A}^{22,1} & \times & \cdots & \times \\
\boldsymbol{A}^{12,2 T} & \times & \boldsymbol{A}^{22,2} & \cdots & \times \\
\vdots & \vdots & \vdots & \ddots & \vdots \\
\boldsymbol{A}^{12, m} & \times & \times & \cdots & \boldsymbol{A}^{22, m}
\end{array}\right]
$$

The blocks represented by the $X$ symbol are not of interest. The relevant blocks of $\boldsymbol{A}^{-1}$ can be efficiently computed applying Theorem 1 of [60] and the formulas therein can be used to derive streamlined MFVB updates and achieve fast computation.

Three-level sparse matrix problems are useful for the treatment of three-level random effects models and details about this class of problems are provided in Section S.2.2 of the supplementary material.

Algorithms using streamlined updates achieve the same MFVB approximations obtained with naïve updates, yet reducing memory usage and performing algebraic steps more efficiently. The former is obtained by circumventing the need of storing the zero sub-blocks of $\boldsymbol{C}$ and the sub-blocks of $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ which are not needed for performing the updates. The latter is achieved by computing the useful sub-blocks of $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ with faster lower-dimensional matrix inversions, and the updates of $\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ and $\lambda_{\mathfrak{q}\left(\sigma^{2}\right)}$ solely relying on the non-zero sub-blocks of $\boldsymbol{C}$ and $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$.

Excellent performances both in terms of approximation accuracy, computational time and memory saving when compared to naïve MFVB or efficient MCMC samplers are shown in [60], especially for large values of $m$. Nevertheless, this reference only treats the generic $\boldsymbol{\beta} \sim \mathrm{N}\left(\boldsymbol{\mu}_{\boldsymbol{\beta}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}}\right)$ prior specification for the fixed effects vector given in (1). In this article, we develop streamlined variational inference procedures allowing for more general prior specifications on $\boldsymbol{\beta}$ aiding selection of fixed effects.

## 4. Approximate variable selection with global-local priors

Regression modeling is often concerned with the problem of selecting an optimal subset of plausible regressors with a significant impact on explaining the variability of the response variable. This is of particular interest in sparse covariate settings, where a large set of regressors is considered but only a small proportion of them is effectively relevant. We refer to [63] and references therein for an exhaustive introductory review on variable selection procedures from a Bayesian perspective.

### 4.1. Bayesian methods for variable selection

Most common Bayesian approaches involve placing suitable prior distributions over the parameters subject to selection. Approaches of this type can be essentially subdivided into two main families, based on the so-called spike-andslab priors [e.g. 58, 29, 40, 39, 23, 11] and global-local shrinkage priors [e.g. 18, 19, 31, 69, 2].

Spike-and-slab priors are two-component mixture priors. The first prior component, the spike, is a point mass function at zero characterizing the noise, usually given by a Dirac delta function or a Gaussian density function having mean zero and very small variance. The second component, the slab, is an absolutely continuous density function representing the signal density of nonzero
coefficients associated with relevant covariates. The slab is usually given by Laplace or Gaussian density functions and is typically centered around zero. A weight parameter taking values in the unit interval is used to balance the contribution of the two components. Although being highly appealing and allowing for separate control of the level of sparsity and the size of the signal coefficients, these priors may suffer from computational hurdles in high-dimensions.

Global-local shrinkage priors are absolutely continuous shrinkage priors that are placed on each coefficient $\beta_{h}, 1 \leq h \leq H$, which is subject to selection. These priors admit the following convenient scale mixture representation [69], for proper choices of $\mathfrak{p}(\tau)$ and $\mathfrak{p}\left(\zeta_{h}\right)$ :

$$
\begin{equation*}
\beta_{h} \mid \tau, \zeta_{h} \stackrel{\text { ind. }}{\sim} \mathrm{N}\left(0, \tau^{2} / \zeta_{h}\right), \quad \tau \sim \mathfrak{p}(\tau), \quad \zeta_{h} \sim \mathfrak{p}\left(\zeta_{h}\right), \quad 1 \leq h \leq H . \tag{13}
\end{equation*}
$$

The global variance parameter $\tau^{2}$ is common to all the coefficients and induces shrinkage towards the origin in the associated posterior density; the local variance parameter $\zeta_{h}$ is coefficient-specific. A more general specification including the model response error variance parameter is proposed in [7]. Depending on the distributional specifications for $\tau$ and $\zeta_{h}$ in (13), many well-known shrinkage priors arise. Examples are the Horseshoe prior of [18], [19], the Bayesian lasso of [67], the Normal-Gamma prior of [31], the Normal-Exponential-Gamma prior of [32], the generalized double Pareto prior of [2], the Dirichlet-Laplace prior of [8] and the horseshoe+ prior of [5]. Longer lists are given in Table 1 of [77] and Table 2 of [6].

For spike-and-slab priors, the posterior distributions of negligible effects present a higher weight for the spike: this provides a direct way to detect relevant effects, and therefore to perform the selection. For global-local priors, there is no posterior spike. The posterior density function, instead, is continuous with probability mass highly concentrated around zero, and a direct way for determining relevant effects is usually unavailable.

We employ global-local priors as they may offer substantial computational advantages over spike-and-slab priors due to their convenient representation as Gaussian scale mixtures, which give rise to convenient conjugate updates for all the $\beta_{h}$ 's and $\zeta_{h}$ 's. [8] also show that such priors exhibit improved posterior concentrations. Furthermore, the estimates of frequentist regularization procedures such as ridge [35], lasso [80], bridge [27] and elastic net [94] can be recasted as posterior mode estimates from models with global-local priors.

### 4.2. Variational inference with global-local priors

Before considering their use within linear mixed model specifications, we illustrate the essential elements of variational inference for the simpler linear regression model. Without loss of generality, we will treat three of the most commonly adopted global-local prior specifications, namely:

$$
\begin{equation*}
\beta_{h}\left|\tau \stackrel{\text { ind. }}{\sim} \operatorname{Laplace}(0, \tau), \beta_{h}\right| \tau \stackrel{\text { ind. }}{\sim} \operatorname{Horseshoe}(0, \tau) \text { or } \beta_{h} \mid \tau \stackrel{\text { ind. }}{\sim} \operatorname{NEG}(0, \tau, \lambda) \tag{14}
\end{equation*}
$$



FIG 1. Visual comparison of the probability density functions for the Laplace, Horseshoe and Normal-Exponential-Gamma (NEG) distributions with zero mean and unit standard deviation. For easiness of comparison, the standard Gaussian probability density function is also displayed.

Table 1
Hierarchical formulation of the Laplace, Horseshoe and Negative-Exponential-Gamma (NEG) priors (14) following the general global-local representation (13), for known $\tau$.

| Prior specification $\mathfrak{p}\left(\beta_{h}\right)$ | $\mathfrak{p}\left(\beta_{h} \mid \zeta_{h}\right)$ | $\mathfrak{p}\left(\zeta_{h} \mid a_{\zeta_{h}}\right)$ | $\mathfrak{p}\left(a_{\zeta_{h}}\right)$ |
| :--- | :---: | :---: | :---: |
| Laplace $(0, \tau)$ | $\mathrm{N}\left(0, \tau^{2} / \zeta_{h}\right)$ | Inverse- $\chi^{2}(2,1)$ | - |
| Horseshoe $(0, \tau)$ | $\mathrm{N}\left(0, \tau^{2} / \zeta_{h}\right)$ | $\operatorname{Gamma}\left(1 / 2, a_{\zeta_{j}}\right)$ | $\operatorname{Gamma}(1 / 2,1)$ |
| $\mathrm{NEG}(0, \tau, \lambda)$ | $\mathrm{N}\left(0, \tau^{2} / \zeta_{h}\right)$ | Inverse- $\chi^{2}\left(2,2 a_{\zeta_{j}}\right)$ | $\operatorname{Gamma}(\lambda, 1)$ |

for each model coefficient $\beta_{h}, 1 \leq h \leq H$. Hereafter $\lambda>0$ is an additional shape parameter that we always assume being user-specified.

The Laplace, Horseshoe and Normal-Exponential-Gamma (NEG) distributions account for different degrees of prior shrinkage towards zero and have different tail behaviors, as shown in Figure 1. Each prior specification in (14) can be recasted into the scale mixture framework (13), as summarized in Table 1. For the Horseshoe and NEG priors, we use convenient hierarchical representations of $\mathfrak{p}\left(\zeta_{h}\right)$ based on auxiliary variables $a_{\zeta_{h}}, 1 \leq h \leq H$, involving tractable Gamma and Inverse- $\chi^{2}$ distributions. Details about the involved density functions and auxiliary variable representations are summarized in Section S. 1 of the supplementary material.

A linear regression model with one of the three global-local priors in (14) is then expressible as:

$$
\begin{gather*}
\boldsymbol{y} \mid \beta_{0}, \boldsymbol{\beta}, \sigma^{2} \sim \mathrm{~N}\left(\beta_{0} \mathbf{1}+\boldsymbol{X} \boldsymbol{\beta}, \sigma^{2} \boldsymbol{I}\right), \\
\beta_{0} \sim \mathrm{~N}\left(\mu_{\beta_{0}}, \sigma_{\beta_{0}}^{2}\right), \quad \boldsymbol{\beta} \mid \boldsymbol{\zeta}, \tau^{2} \sim \mathrm{~N}\left(\mathbf{0}, \tau^{2} \operatorname{diag}(\boldsymbol{\zeta})^{-1}\right), \\
\sigma^{2} \mid a_{\sigma^{2}} \sim \text { Inverse- } \chi^{2}\left(\nu_{\sigma^{2}}, 1 / a_{\sigma^{2}}\right), \quad a_{\sigma^{2}} \sim \text { Inverse- } \chi^{2}\left(1,1 /\left(\nu_{\sigma^{2}} s_{\sigma^{2}}^{2}\right)\right),  \tag{15}\\
\tau^{2} \mid a_{\tau^{2}} \sim \text { Inverse- } \chi^{2}\left(1,1 / a_{\tau^{2}}\right), \quad a_{\tau^{2}} \sim \text { Inverse- } \chi^{2}\left(1,1 / s_{\tau^{2}}^{2}\right), \\
\zeta_{h} \mid a_{\zeta_{h}} \stackrel{\text { ind. }}{\sim} \mathfrak{p}\left(\zeta_{h} \mid a_{\zeta_{h}}\right), \quad a_{\zeta h} \stackrel{\text { ind. }}{\sim} \mathfrak{p}\left(a_{\zeta_{h}}\right), \quad 1 \leq h \leq H,
\end{gather*}
$$

where $\beta_{0}$ is the model intercept (which is usually excluded from the selection procedure), $\mathbf{1}$ is a vector full of ones, $\boldsymbol{\beta}=\left[\beta_{1}, \ldots, \beta_{H}\right]^{T}$ is the vector of coefficients having global-local shrinkage priors, $\boldsymbol{X}=\left[\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{p}\right]$ is the associated design matrix, $\boldsymbol{\zeta}=\left[\zeta_{1}, \ldots, \zeta_{H}\right]^{T}$ and $\boldsymbol{a}_{\boldsymbol{\zeta}}=\left[a_{\zeta_{1}}, \ldots, a_{\zeta_{H}}\right]^{T}$. The common Gaussian and Half-t prior distributions are considered for $\beta_{0}$ and $\sigma$, respectively. Importantly, we specify a Half- $t\left(s_{\tau^{2}}, 1\right)$ distribution with $s_{\tau^{2}}>0$ for the global scale parameter $\tau$, allowing for weak prior informativeness about the global degree of sparseness when a large scale parameter $s_{\tau^{2}}$ is used. Hence, we let the variable-selection procedure be free from hyperparameters that have to be manually tuned by the user. The densities $\mathfrak{p}\left(\zeta_{h} \mid a_{\zeta_{h}}\right)$ and $\mathfrak{p}\left(a_{\zeta_{h}}\right)$ vary according to the global-local prior specification adopted, as shown in Table 1.

The model posterior density function $\mathfrak{p}\left(\beta_{0}, \boldsymbol{\beta}, \sigma^{2}, a_{\sigma^{2}}, \tau^{2}, a_{\tau^{2}}, \boldsymbol{\zeta}, \boldsymbol{a}_{\boldsymbol{\zeta}} \mid \boldsymbol{y}\right)$ admits a tractable MFVB approximation when the following mean-field restriction is used:
$\mathfrak{q}\left(\beta_{0}, \boldsymbol{\beta}, \sigma^{2}, a_{\sigma^{2}}, \tau^{2}, a_{\tau^{2}}, \boldsymbol{\zeta}, \boldsymbol{a}_{\boldsymbol{\zeta}}\right)=\mathfrak{q}\left(\beta_{0}, \boldsymbol{\beta}\right) \mathfrak{q}\left(\sigma^{2}\right) \mathfrak{q}\left(\tau^{2}\right) \mathfrak{q}\left(a_{\sigma^{2}}\right) \mathfrak{q}\left(a_{\tau^{2}}\right) \prod_{h=1}^{H}\left\{\mathfrak{q}\left(\zeta_{h}\right) \mathfrak{q}\left(a_{\zeta_{h}}\right)\right\}$.
The optimal $\mathfrak{q}$-density functions then results as follows:

$$
\begin{align*}
& \mathfrak{q}^{*}\left(\beta_{0}, \boldsymbol{\beta}\right) \text { is a } \mathrm{N}\left(\boldsymbol{\mu}_{\mathfrak{q}\left(\beta_{0}, \boldsymbol{\beta}\right)}, \boldsymbol{\Sigma}_{\mathfrak{q}\left(\beta_{0}, \boldsymbol{\beta}\right)}\right) \text { density function, } \\
& \mathfrak{q}^{*}\left(\sigma^{2}\right) \text { is an Inverse- } \chi^{2}\left(\xi_{\mathfrak{q}\left(\sigma^{2}\right)}, \lambda_{\mathfrak{q}\left(\sigma^{2}\right)}\right) \text { density function, } \\
& \mathfrak{q}^{*}\left(a_{\sigma^{2}}\right) \text { is an Inverse- } \chi^{2}\left(\xi_{\mathfrak{q}\left(a_{\sigma^{2}}\right)}, \lambda_{\mathfrak{q}\left(a_{\sigma^{2}}\right)}\right) \text { density function, } \\
& \mathfrak{q}^{*}\left(\tau^{2}\right) \text { is an Inverse- } \chi^{2}\left(\xi_{\mathfrak{q}\left(\tau^{2}\right)}, \lambda_{\mathfrak{q}\left(\tau^{2}\right)}\right) \text { density function, } \\
& \mathfrak{q}^{*}\left(a_{\tau^{2}}\right) \text { is an Inverse- } \chi^{2}\left(\xi_{\mathfrak{q}\left(a_{\tau^{2}}\right)}, \lambda_{\mathfrak{q}\left(a_{\tau^{2}}\right)}\right) \text { density function, } \\
& \mathfrak{q}^{*}\left(\zeta_{h}\right) \text { is }\left\{\begin{array}{l}
\text { an Inverse-Gaussian }\left(\mu_{\mathfrak{q}\left(\zeta_{h}\right)}, 1\right) \text { d. f. for a Laplace prior } \\
\text { a } \operatorname{Gamma}\left(1, \lambda_{\mathfrak{q}\left(\zeta_{h}\right)}\right) \text { d. f. for a Horseshoe prior } \\
\text { an Inverse-Gaussian }\left(\mu_{\mathfrak{q}\left(\zeta_{h}\right)}, \lambda_{\mathfrak{q}\left(\zeta_{h}\right)}\right) \text { d. f. for a NEG prior }
\end{array}\right.  \tag{16}\\
& \text { and } \mathfrak{q}^{*}\left(a_{\zeta_{h}}\right) \text { is }\left\{\begin{array}{c}
- \\
\text { for a Laplace prior } \\
\operatorname{a~} \operatorname{Gamma}\left(1, \lambda_{\mathfrak{q}\left(a_{\zeta_{h}}\right)}\right) \text { d. f. for a Horseshoe prior } \\
\operatorname{a~} \operatorname{Gamma}\left(\lambda+1, \lambda_{\mathfrak{q}\left(a_{\zeta_{h}}\right)}\right) \text { d. f. for a NEG prior, }
\end{array}\right.
\end{align*}
$$

for $1 \leq h \leq H$. When a Laplace prior is specified, the model does not include the $a_{\zeta_{h}}$ auxiliary variables and therefore $\mathfrak{q}\left(a_{\zeta_{h}}\right)$ is not included. The expressions for the parameter updates of these approximating densities, together with their full derivations, are reported in Section S. 3 of the supplementary material.

### 4.3. From shrinkage to selection: The signal adaptive variable selector

One limitation of continuous global-local shrinkage priors is the unavailability of direct information from the posterior of each $\beta_{h}$ for selecting relevant effects.

Typically, the posterior distributions of less relevant coefficients arising from such priors are highly concentrated around zero, with marked peaks and negligible tails, although not having full mass at zero. Therefore, global-local shrinkage priors do not provide any sparse posterior solution. This issue becomes even more relevant when variational inference is used to fit the model because the approximate marginal posterior densities of $\boldsymbol{\beta}$ are Gaussian, and so the peaks of the true marginal posterior densities are approximated by bell-shaped curves.

Several heuristic methods have been developed for post-processing posterior distributions arising from global-local priors and determining whether the associated covariates have to be selected or not. A simple but possibly misleading solution is to select as relevant the covariates associated with coefficients whose posterior credible intervals do not contain the zero. Nonetheless, this approach usually exhibits poor performances due to the difficulty of accurately estimating the uncertainty in high dimensional problems, and depends on the chosen credible level. [19] define a local shrinkage factor which can take values in the unit interval and help determine whether each variable is suggested to be selected or not according to a pre-specified threshold, analogously to the classical posterior inclusion probability of [4]. [12] propose a method based on posterior credible regions, although its implementation and results rely upon the use of conjugate Normal priors. [91] extended the method to global-local priors and propose an intuitive approach to tune the prior hyperparameters based on minimizing a discrepancy measure between the induced distribution of $R^{2}$ from the prior and the desired distribution.

All these methods and many others have a common issue, that is the dependence on the choice of one or more thresholds. [8] propose grouping the entries of posterior medians into null and non-null groups using 2-means clustering. While this approach does not require any tuning parameters, issues emerge when there are signals of varying strengths. [46] propose a similar approach which is based on first obtaining a posterior distribution of the number of signals by clustering the signal and the noise coefficients and then estimating the signals from the posterior median.

In this work, we opt for the signal adaptive variable selector (SAVS) partially motivated by [34] and accurately developed by [71]. The SAVS approach post-processes a point estimate from the posterior distribution of a coefficient having global-local prior distribution via soft-thresholding to determine whether the associated covariate is assumed to be relevant or not. We adapt this procedure for usage in variational inference and propose Algorithm 1 as an implementation of the SAVS approach based on the optimal approximate posterior densities $\mathfrak{q}^{*}\left(\beta_{h}\right), 1 \leq h \leq H$. The procedure takes the approximate posterior mean parameter of a generic coefficient subject to selection and the associated unstandardized covariate as inputs. It then returns a sparsified approximate posterior summary estimate $\mu_{\mathfrak{q}^{*}\left(\beta_{h}\right)}^{*}$, together with a binary variable $\gamma_{h}$ indicating whether the $h$ th covariate is suggested to be selected or not. The attractiveness of this approach comes from the fact that it is completely automated and does not require any tuning parameters. [71] provide a theoretical justification for the SAVS approach, noticing that its output can be obtained by solving an

```
Algorithm 1 Signal Adaptive Variable Selector (SAVS) algorithm for perform-
ing variable selection using the optimal approximate density function \(\mathfrak{q}^{*}\left(\beta_{h}\right)\) of
a generic coefficient with global-local prior.
Inputs: \(\mu_{\mathfrak{q}^{*}\left(\beta_{h}\right)} \equiv E_{\mathfrak{q}^{*}}\left(\beta_{h}\right)\) and \(\mathbf{x}_{h}\) being the covariate vector corresponding to \(\beta_{h}, 1 \leq h \leq H\).
If \(\left\|\mathbf{x}_{h}\right\|^{2} \leq\left|\mu_{\mathfrak{q}^{*}\left(\beta_{h}\right)}\right|^{-3}\) :
    \(\mu_{\mathfrak{q}^{*}\left(\beta_{h}\right)}^{*}=0 \quad\) and \(\quad \gamma_{h}=0 ;\)
else:
    \(\mu_{\mathfrak{q}^{*}\left(\beta_{h}\right)}^{*}=\operatorname{sign}\left(\mu_{\mathfrak{q}^{*}\left(\beta_{h}\right)}\right)\left\|\mathbf{x}_{h}\right\|^{-2}\left(\left|\mu_{\mathfrak{q}^{*}\left(\beta_{h}\right)}\right|\left\|\mathbf{x}_{h}\right\|^{2}-\mu_{\mathfrak{q}^{*}\left(\beta_{h}\right)}^{-2}\right) \quad\) and \(\quad \gamma_{h}=1\).
Output: A sparse estimate \(\mu_{\mathfrak{q}^{*}\left(\beta_{h}\right)}^{*}\) for \(\mathfrak{q}^{*}\left(\beta_{h}\right)\) and the associated binary selector \(\gamma_{h}\).
```

optimization problem closely related to the adaptive lasso of [93], and show it is highly competitive among alternative Bayesian selection procedures.

## 5. Linear mixed models with global-local priors on fixed effects

Variational approximations for linear mixed models with two- and three-level random effects are described in Section 3, using a generic $\boldsymbol{\beta} \sim N\left(\boldsymbol{\mu}_{\boldsymbol{\beta}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}}\right)$ prior distribution for the fixed effects parameter vector. In this work, our interest is in developing variational approximations for generalizations of model (1) embedding prior specifications for fixed effects selection such as those discussed in Section 4.

In order to do so, we subdivide the $p$-dimensional fixed effects vector $\boldsymbol{\beta}$ as follows:

$$
\boldsymbol{\beta}=\left[\begin{array}{l}
\boldsymbol{\beta}^{\mathrm{R}} \\
\boldsymbol{\beta}^{\mathrm{A}} \\
\boldsymbol{\beta}^{\mathrm{S}}
\end{array}\right]
$$

where $\boldsymbol{\beta}^{\mathrm{R}}$ is a $p_{\mathrm{R}}$-dimensional vector of fixed effects associated to the random $(R)$ effects component of the model, $\boldsymbol{\beta}^{\mathrm{A}}$ is a $p_{\mathrm{A}}$-dimensional vector of additional $(A)$ fixed effects and $\boldsymbol{\beta}^{\mathrm{S}}$ is a $p_{\mathrm{S}}$-dimensional vector of fixed effects which are subject to selection $(S)$. Here $p=p_{\mathrm{R}}+p_{\mathrm{A}}+p_{\mathrm{S}}$, with $p_{\mathrm{R}}, p_{\mathrm{A}}$ and $p_{\mathrm{S}}$ varying according to the application of interest. For the two- and three-level mixed models considered in this work, $p_{\mathrm{R}}=q$ and $p_{\mathrm{R}}=\max \left(q_{1}, q_{2}\right)$ respectively. Typically, $q, q_{1}$ and $q_{2}$ are relatively small, while $p_{\mathrm{A}}$ and $p_{\mathrm{S}}$ could take moderate to large values.

Similarly, we subdivide the fixed effects design matrix as follows:

$$
\boldsymbol{X}=\left[\begin{array}{l|l|l}
\boldsymbol{X}^{\mathrm{R}} & \boldsymbol{X}^{\mathrm{A}} & \boldsymbol{X}^{\mathrm{S}}
\end{array}\right]
$$

with $\boldsymbol{X}^{\mathrm{S}}$ assumed to have columns with zero mean and unit variance, unless differently specified. The fixed effects linear contribution of model (1) then factorizes into $\boldsymbol{X} \boldsymbol{\beta}=\boldsymbol{X}^{\mathrm{R}} \boldsymbol{\beta}^{\mathrm{R}}+\boldsymbol{X}^{\mathrm{A}} \boldsymbol{\beta}^{\mathrm{A}}+\boldsymbol{X}^{\mathrm{S}} \boldsymbol{\beta}^{\mathrm{S}}$, and the same applies to the two-level mixed model (2) and the three-level mixed model (4) specifications. Notice that for the former specification $\boldsymbol{Z}_{i}=\boldsymbol{X}_{i}^{\mathrm{R}}$, for all $1 \leq i \leq m$.

We assume without loss of generality that $\boldsymbol{\beta}^{\mathrm{R}}, \boldsymbol{\beta}^{\mathrm{A}}$ and $\boldsymbol{\beta}^{\mathrm{S}}$ are a-priori independent from each other, and specify the following prior distributions:

$$
\boldsymbol{\beta}^{\mathrm{R}} \sim \mathrm{~N}\left(\boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{R}}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{R}}}\right), \quad \boldsymbol{\beta}^{\mathrm{A}} \sim \mathrm{~N}\left(\boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{A}}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{A}}}\right), \quad \boldsymbol{\beta}^{\mathrm{S}} \sim \mathfrak{p}\left(\boldsymbol{\beta}^{\mathrm{S}}\right)=\prod_{h=1}^{p_{\mathrm{S}}} \mathfrak{p}\left(\beta_{h}^{\mathrm{S}}\right)
$$

with hyperparameters $\boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{R}}} \in \mathbb{R}^{p_{\mathrm{R}}}, \boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{A}}} \in \mathbb{R}^{p_{\mathrm{A}}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{R}}}$ and $\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{A}}}$ symmetric positive definite matrices. The prior specification for $\boldsymbol{\beta}^{\mathrm{S}}$ assumes a-priori independence among all the coefficients subject to selection, with $\mathfrak{p}\left(\beta_{h}^{S}\right)$ taking one of the three different global-local prior distributions treated in Section 4 for $1 \leq h \leq p_{\mathrm{S}}$, specifically:

$$
\beta_{h}^{\mathrm{S}} \mid \tau \stackrel{\text { ind. }}{\sim} \text { Laplace }(0, \tau), \quad \beta_{h}^{\mathrm{S}} \mid \tau \stackrel{\text { ind. }}{\sim} \operatorname{Horseshoe}(0, \tau) \quad \text { or } \quad \beta_{h}^{\mathrm{S}} \mid \tau \stackrel{\text { ind. }}{\sim} \operatorname{NEG}(0, \tau, \lambda)
$$

The resulting linear mixed model is a generalization of (1) that accounts for global-local prior specification over a subset of the fixed effects, and can be expressed as:

$$
\begin{align*}
& \boldsymbol{y} \mid \boldsymbol{\beta}^{\mathrm{R}}, \boldsymbol{\beta}^{\mathrm{A}}, \boldsymbol{\beta}^{\mathrm{S}}, \boldsymbol{u}, \sigma^{2} \sim \mathrm{~N}\left(\boldsymbol{X}^{\mathrm{R}} \boldsymbol{\beta}^{\mathrm{R}}+\boldsymbol{X}^{\mathrm{A}} \boldsymbol{\beta}^{\mathrm{A}}+\boldsymbol{X}^{\mathrm{S}} \boldsymbol{\beta}^{\mathrm{S}}+\boldsymbol{Z} \boldsymbol{u}, \sigma^{2} \boldsymbol{I}\right), \\
& \boldsymbol{u} \mid \boldsymbol{G} \sim \mathrm{N}(\mathbf{0}, \boldsymbol{G}), \\
& {\left.\left[\begin{array}{c}
\boldsymbol{\beta}^{\mathrm{R}} \\
\boldsymbol{\beta}^{\mathrm{A}} \\
\boldsymbol{\beta}^{\mathrm{S}}
\end{array}\right] \right\rvert\, \boldsymbol{\zeta}, \tau^{2} \sim \mathrm{~N}\left(\left[\begin{array}{c}
\boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{R}}} \\
\boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{A}}} \\
\boldsymbol{0}
\end{array}\right],\left[\begin{array}{ccc}
\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{R}}} & \boldsymbol{O} & \boldsymbol{O} \\
\boldsymbol{O} & \boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{A}}} & \boldsymbol{O} \\
\boldsymbol{O} & \boldsymbol{O} & \tau^{2} \operatorname{diag}(\boldsymbol{\zeta})^{-1}
\end{array}\right]\right),} \\
& \sigma^{2} \mid a_{\sigma^{2}} \sim \text { Inverse- } \chi^{2}\left(\nu_{\sigma^{2}}, 1 / a_{\sigma^{2}}\right), \quad a_{\sigma^{2}} \sim \text { Inverse- } \chi^{2}\left(1,1 /\left(\nu_{\sigma^{2}} S_{\sigma^{2}}^{2}\right)\right), \\
& \zeta_{h} \left\lvert\, a_{\zeta_{h}} \stackrel{\text { ind. }}{\sim}\left\{\begin{array}{ll}
\text { Inverse- } \chi^{2}(2,1) & \text { for a Laplace prior } \\
\text { Gamma }\left(1 / 2, a_{\zeta_{h}}\right) & \text { for a Horseshoe prior } \\
\text { Inverse- } \chi^{2}\left(2,2 a_{\zeta_{h}}\right)
\end{array} \quad\right. \text { for a NEG prior, }\right.  \tag{17}\\
& a_{\zeta_{h}} \stackrel{\text { ind. }}{\sim}\left\{\begin{array}{cl}
- & \text { for a Laplace prior } \\
\operatorname{Gamma}(1 / 2,1) & \text { for a Horseshoe prior } \\
\operatorname{Gamma}(\lambda, 1) & \text { for a NEG prior },
\end{array}\right. \\
& \text { for } 1 \leq h \leq p_{\mathrm{S}} \text {, } \\
& \tau^{2} \mid a_{\tau^{2}} \sim \text { Inverse- } \chi^{2}\left(1,1 / a_{\tau^{2}}\right), \quad a_{\tau^{2}} \sim \text { Inverse- } \chi^{2}\left(1,1 / s_{\tau^{2}}^{2}\right), \\
& \boldsymbol{G} \mid \boldsymbol{A}_{\boldsymbol{G}} \sim \mathfrak{p}\left(\boldsymbol{G} \mid \boldsymbol{A}_{G}\right), \quad \boldsymbol{A}_{\boldsymbol{G}} \sim \mathfrak{p}\left(\boldsymbol{A}_{\boldsymbol{G}}\right) .
\end{align*}
$$

Here $\boldsymbol{\zeta} \equiv\left[\zeta_{1}, \ldots, \zeta_{p_{\mathrm{S}}}\right]^{T}$ and $\boldsymbol{a}_{\boldsymbol{\zeta}} \equiv\left[a_{\zeta_{1}}, \ldots, a_{\zeta_{p_{\mathrm{S}}}}\right]^{T}$. This model can be fitted via MFVB assuming that the full posterior density function is approximated as

$$
\begin{align*}
& \mathfrak{p}\left(\boldsymbol{\beta}, \boldsymbol{u}, \sigma^{2}, a_{\sigma^{2}}, \boldsymbol{\zeta}, \boldsymbol{a}_{\boldsymbol{\zeta}}, \tau^{2}, a_{\tau^{2}}, \boldsymbol{G}, \boldsymbol{A}_{\boldsymbol{G}} \mid \boldsymbol{y}\right) \approx  \tag{18}\\
& \mathfrak{q}\left(\boldsymbol{\beta}, \boldsymbol{u}, \sigma^{2}, a_{\sigma^{2}}, \boldsymbol{\zeta}, \boldsymbol{a}_{\boldsymbol{\zeta}}, \tau^{2}, a_{\tau^{2}}, \boldsymbol{G}, \boldsymbol{A}_{G}\right)
\end{align*}
$$

and a tractable solution arises with the following mean-field restriction:

$$
\begin{gather*}
\mathfrak{q}\left(\boldsymbol{\beta}, \boldsymbol{u}, \sigma^{2}, a_{\sigma^{2}}, \boldsymbol{\zeta}, \boldsymbol{a}_{\boldsymbol{\zeta}}, \tau^{2}, a_{\tau^{2}}, \boldsymbol{G}, \boldsymbol{A}_{\boldsymbol{G}}\right)= \\
\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u}) \mathfrak{q}\left(\sigma^{2}\right) \mathfrak{q}\left(a_{\sigma^{2}}\right)\left\{\prod_{h=1}^{p_{\mathrm{S}}} \mathfrak{q}\left(\zeta_{h}\right) \mathfrak{q}\left(a_{\zeta_{h}}\right)\right\} \mathfrak{q}\left(\tau^{2}\right) \mathfrak{q}\left(a_{\tau^{2}}\right) \mathfrak{q}(\boldsymbol{G}) \mathfrak{q}\left(\boldsymbol{A}_{G}\right) . \tag{19}
\end{gather*}
$$

Arguments similar to those given in Section 3.2 and 4.2 lead to the optimal approximating densities being:

$$
\begin{align*}
& \mathfrak{q}^{*}(\boldsymbol{\beta}, \boldsymbol{u}) \text { is a } \mathrm{N}\left(\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}, \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}\right) \text { density function, } \\
& \mathfrak{q}^{*}\left(\sigma^{2}\right) \text { is an Inverse- } \chi^{2}\left(\xi_{\mathfrak{q}\left(\sigma^{2}\right)}, \lambda_{\mathfrak{q}\left(\sigma^{2}\right)}\right) \text { density function, } \\
& \mathfrak{q}^{*}\left(a_{\sigma^{2}}\right) \text { is an Inverse- } \chi^{2}\left(\xi_{\mathfrak{q}\left(a_{\sigma^{2}}\right)}, \lambda_{\mathfrak{q}\left(a_{\sigma^{2}}\right)}\right) \text { density function, } \\
& \mathfrak{q}^{*}\left(\zeta_{h}\right) \text { is }\left\{\begin{array}{l}
\text { an } \operatorname{Inverse-Gaussian}\left(\mu_{\mathfrak{q}\left(\zeta_{h}\right)}, 1\right) \text { d. f. for a Laplace prior } \\
\text { a } \operatorname{Gamma}\left(1, \lambda_{\mathfrak{q}\left(\zeta_{h}\right)}\right) \text { d. f. for a Horseshoe prior } \\
\text { an Inverse-Gaussian }\left(\mu_{\mathfrak{q}\left(\zeta_{h}\right)}, \lambda_{\mathfrak{q}\left(\zeta_{h}\right)}\right) \text { d. f. for a NEG prior, }
\end{array}\right.  \tag{20}\\
& \mathfrak{q}^{*}\left(a_{\zeta_{h}}\right) \text { is }\left\{\begin{array}{c}
- \\
\text { for a Laplace prior } \\
\operatorname{a~} \operatorname{Gamma}\left(1, \lambda_{\mathfrak{q}\left(a_{\zeta_{h}}\right)}\right) \text { d. f. for a Horseshoe prior } \\
\text { a } \operatorname{Gamma}\left(\lambda+1, \lambda_{\mathfrak{q}\left(a_{\zeta_{h}}\right)}\right) \text { d.f. for a NEG prior, }
\end{array}\right. \\
& \text { for } 1 \leq h \leq p_{\mathrm{S}}, \\
& \mathfrak{q}^{*}\left(\tau^{2}\right) \text { is an Inverse- } \chi^{2}\left(\xi_{\mathfrak{q}\left(\tau^{2}\right)}, \lambda_{\mathfrak{q}\left(\tau^{2}\right)}\right) \text { density function } \\
& \text { and } \mathfrak{q}^{*}\left(a_{\tau^{2}}\right) \text { is an Inverse- } \chi^{2}\left(\xi_{\mathfrak{q}\left(a_{\tau^{2}}\right)}, \lambda_{\mathfrak{q}\left(a_{\tau^{2}}\right)}\right) \text { density function. }
\end{align*}
$$

The optimal approximating densities for the matrices $\boldsymbol{G}$ and $\boldsymbol{A}_{\boldsymbol{G}}$ vary according to the adopted random effect structure, as explained in Section 3.2. Notice that (19) jointly approximates $\boldsymbol{\beta}^{\mathrm{R}}, \boldsymbol{\beta}^{\mathrm{A}}, \boldsymbol{\beta}^{\mathrm{S}}$ and $\boldsymbol{u}$, allowing all the fixed effects to share posterior dependence with the random effects. Also, the $\mathfrak{q}^{*}\left(\beta_{h}^{S}\right)$ approximating densities are Gaussian, although different global-local prior specifications may lead to marginal posterior density functions $\mathfrak{p}\left(\beta_{h}^{\mathrm{S}} \mid \boldsymbol{y}\right)$ having shapes around zero that are different from the typical bell-shaped behavior, especially those of fixed effects associated to irrelevant covariates. Using MFVB we sacrifice some degree of accuracy, yet obtaining a substantial computational time advantage over standard MCMC sampling procedures. Section 6 investigates the quality of these approximations and the fixed effects selection performances.

### 5.1. Naïve updates

The updates of the parameters of the $\mathfrak{q}$-densities in (20) can be derived combining and adapting the results discussed in Sections 3.2 and 4.2 , and references therein. In particular, those related to the $\mathfrak{q}^{*}(\boldsymbol{\beta}, \boldsymbol{u})$ optimal density function are:

$$
\begin{align*}
& \Sigma_{q(\beta, u)} \longleftarrow \\
& \left(\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \boldsymbol{C}^{T} \boldsymbol{C}+\left[\begin{array}{cccc}
\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathfrak{R}}}^{-1} & O & O & O \\
O & \boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{A}}}^{-1} & O & O \\
O & O & \mu_{\mathfrak{q}\left(1 / \tau^{2}\right)} \operatorname{diag}\left(\boldsymbol{\mu}_{\mathfrak{q}(\zeta)}\right) & \boldsymbol{O} \\
\boldsymbol{O} & O & O & E_{\mathfrak{q}}\left(\boldsymbol{G}^{-1}\right)
\end{array}\right]\right)^{-1}  \tag{21}\\
& \text { and } \quad \boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})} \longleftarrow \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}\left(\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \boldsymbol{C}^{T} \boldsymbol{y}+\left[\begin{array}{c}
\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{R}}}^{-1} \boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{R}}} \\
\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{A}}}^{-1} \boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{A}}} \\
\mathbf{0} \\
\mathbf{0}
\end{array}\right]\right) \text {. }
\end{align*}
$$

The update for $E_{\mathfrak{q}}\left(\boldsymbol{G}^{-1}\right)$ is given by (11) or (12) depending on whether we are considering a two-level or a three-level mixed model specification, respectively. The main difference with the expressions given in (10) is that a globallocal prior specification on each element of $\boldsymbol{\beta}^{\mathrm{S}}$ introduces the diagonal matrix $\mu_{\mathfrak{q}\left(1 / \tau^{2}\right)} \operatorname{diag}\left(\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\zeta})}\right)$ of dimension $p_{\mathrm{S}} \times p_{\mathrm{S}}$ inside the update expression for $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$. This diagonal matrix is updated at each iteration of the MFVB algorithm employing the updated values of $\mu_{\mathfrak{q}\left(1 / \tau^{2}\right)}$ and $\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\zeta})}$, accordingly to the global-local prior adopted.

The updates for the parameters of $\mathfrak{q}^{*}\left(\sigma^{2}\right)$ and $\mathfrak{q}^{*}\left(a_{\sigma^{2}}\right)$ are identical to those in (10). The updates for the parameters of the optimal approximating $\mathfrak{q}$-densities of $\boldsymbol{G}$ and $\boldsymbol{A}_{\boldsymbol{G}}$ are identical to those described in Section 3.2. The updates for the parameters of $\mathfrak{q}^{*}\left(\zeta_{h}\right)$ and $\mathfrak{q}^{*}\left(a_{\zeta_{h}}\right), 1 \leq h \leq p_{\mathrm{S}}, \mathfrak{q}^{*}\left(\tau^{2}\right)$ and $\mathfrak{q}^{*}\left(a_{\tau^{2}}\right)$ follow with minor modifications from Section 4.2 , replacing $H$ with $p_{\mathrm{S}}$.

Nevertheless, updates (21) suffer from the same problematics elucidated in Section 3.2 for large $m$ values. Therefore, an appropriate streamlined enhancement for efficient implementations is required.

### 5.2. Streamlined updates

Results 1 and 4 from [60] can be extended to derive a streamlined MFVB algorithm for efficiently updating the parameters of the densities in (20) and, in particular, for exploiting the sparse matrix structures presented in updates (21). Such results are based on the two- and three-level sparse matrix least squares problems defined in [60], and make use of the associated

SolveTwoLevelSparsematrix and SolveThreeLevelSparseMatrix
routines, which are recalled in Section S. 2 of the supplementary material. The following results explain how to efficiently compute $\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ and the relevant
sub-blocks of $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ which are necessary for finding the optimal $\mathfrak{q}$-density listed in (20).

Result 1 (Result 1 of [60], revisited) The $M F V B$ updates of $\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ and each of the sub-blocks of $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ that are relevant for variational inference concerning model (20) with a two-level random effects specification are expressible as a two-level sparse matrix problem of the form $\boldsymbol{A} \boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}=\boldsymbol{a}$, where $\boldsymbol{a}$ and the non-zero sub-blocks of $\boldsymbol{A}$ are, according to the notation in Section 3.3:

$$
\begin{gathered}
\boldsymbol{A}_{11}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \sum_{i=1}^{m}\left(\boldsymbol{X}_{i}^{T} \boldsymbol{X}_{i}\right)+\left[\begin{array}{ccc}
\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{R}}^{-1} & \boldsymbol{O} & \boldsymbol{O} \\
\boldsymbol{O} & \boldsymbol{\Sigma}_{\boldsymbol{\beta}^{A}}^{-1} & \boldsymbol{O} \\
\boldsymbol{O} & \boldsymbol{O} & \mu_{\mathfrak{q}\left(1 / \tau^{2}\right)} \operatorname{diag}\left(\boldsymbol{\mu}_{\mathfrak{q}(\zeta)}\right)
\end{array}\right] \\
\boldsymbol{a}_{1}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \sum_{i=1}^{m}\left(\boldsymbol{X}_{i}^{T} \boldsymbol{y}_{i}\right)+\left[\begin{array}{c}
\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{R}}^{-1} \boldsymbol{\mu}_{\boldsymbol{\beta}^{R}} \\
\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{A}}^{-1} \boldsymbol{\mu}_{\boldsymbol{\beta}^{A}} \\
\mathbf{0}
\end{array}\right] \\
\boldsymbol{A}_{22, i}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \boldsymbol{Z}_{i}^{T} \boldsymbol{Z}_{i}+\boldsymbol{M}_{q\left(\boldsymbol{\Sigma}^{-1}\right)}, \quad \boldsymbol{A}_{12, i}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \boldsymbol{X}_{i}^{T} \boldsymbol{Z}_{i} \\
\boldsymbol{a}_{2, i}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \boldsymbol{Z}_{i}^{T} \boldsymbol{y}_{i}
\end{gathered}
$$

for $1 \leq i \leq m$. Moreover, $\boldsymbol{X}_{i}=\left[\boldsymbol{X}_{i}^{R}\left|\boldsymbol{X}_{i}^{A}\right| \boldsymbol{X}_{i}^{S}\right]$. The SolveTwoLevelSparsematrix routine efficiently solves the associated linear system and provides the solutions:

$$
\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}=\boldsymbol{x}_{1}, \quad \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta})}=\boldsymbol{A}^{11}
$$

and, for $1 \leq i \leq m$,

$$
\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}=\boldsymbol{x}_{2, i}, \quad \boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}=\boldsymbol{A}^{22, i}, \quad E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}\right)^{T}\right\}=\boldsymbol{A}^{12, i}
$$

Result 2 (Result 4 of [60], revisited) The $M F V B$ updates of $\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ and each of the sub-blocks of $\boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}$ that are relevant for variational inference concerning model (20) with a three-level random effects specification are expressible as a three-level sparse matrix problem of the form $\boldsymbol{A} \boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}=\boldsymbol{a}$, where $\boldsymbol{a}$ and the non-zero sub-blocks of $\boldsymbol{A}$ are, according to the notation in Section S.2.2 of the supplementary material:

$$
\begin{gathered}
\boldsymbol{A}_{11}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \sum_{i=1}^{m} \sum_{j=1}^{n_{i}}\left(\boldsymbol{X}_{i j}^{T} \boldsymbol{X}_{i j}\right)+\left[\begin{array}{ccc}
\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{R}}^{-1} & \boldsymbol{O} & \boldsymbol{O} \\
\boldsymbol{O} & \boldsymbol{\Sigma}_{\boldsymbol{\beta}^{A}}^{-1} & \boldsymbol{O} \\
\boldsymbol{O} & \boldsymbol{O} & \mu_{\mathfrak{q}\left(1 / \tau^{2}\right)} \operatorname{diag}\left(\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\zeta})}\right)
\end{array}\right], \\
\boldsymbol{a}_{1}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \sum_{i=1}^{m} \sum_{j=1}^{n_{i}}\left(\boldsymbol{X}_{i j}^{T} \boldsymbol{y}_{i j}\right)+\left[\begin{array}{c}
\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{R}}^{-1} \boldsymbol{\mu}_{\boldsymbol{\beta}^{R}} \\
\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{A}}^{-1} \boldsymbol{\mu}_{\boldsymbol{\beta}^{A}} \\
\mathbf{0}
\end{array}\right]
\end{gathered}
$$

$$
\begin{aligned}
& \boldsymbol{A}_{22, i}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \sum_{j=1}^{n_{i}}\left(\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 1}\right)^{T} \boldsymbol{Z}_{i j}^{\mathrm{L} 1}\right)+\boldsymbol{M}_{q\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)^{-1}\right)}, \\
& \boldsymbol{A}_{12, i}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \sum_{j=1}^{n_{i}}\left(\boldsymbol{X}_{i j}\right)^{T} \boldsymbol{Z}_{i j}^{\mathrm{L} 1}, \\
& \boldsymbol{a}_{2, i}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \sum_{j=1}^{n_{i}}\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 1}\right)^{T} \boldsymbol{y}_{i j}, \\
& \boldsymbol{A}_{22, i j}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)}\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 2}\right)^{T} \boldsymbol{Z}_{i j}^{\mathrm{L} 2}+\boldsymbol{M}_{q\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)^{-1}\right)}, \\
& \boldsymbol{A}_{12, i, j}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)}\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 1}\right)^{T} \boldsymbol{Z}_{i j}^{\mathrm{L} 2}, \quad \boldsymbol{A}_{12, i j}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \boldsymbol{X}_{i j}^{T} \boldsymbol{Z}_{i j}^{\mathrm{L} 2}, \\
& \boldsymbol{a}_{2, i j}=\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)}\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 2}\right)^{T} \boldsymbol{y}_{i j},
\end{aligned}
$$

for $1 \leq i \leq m$ and $1 \leq j \leq n_{i}$. Moreover, $\boldsymbol{X}_{i j}=\left[\boldsymbol{X}_{i j}^{R}\left|\boldsymbol{X}_{i j}^{A}\right| \boldsymbol{X}_{i j}^{S}\right]$. The SolveThreelevelSparsematrix routine efficiently solves the associated linear system, and provides the solutions:

$$
\begin{gathered}
\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}=\boldsymbol{x}_{1}, \quad \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta})}=\boldsymbol{A}^{11}, \\
\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}=\boldsymbol{x}_{2, i}, \quad \boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}=\boldsymbol{A}^{22, i}, \\
E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}\right)^{T}\right\}=\boldsymbol{A}^{12, i}, \quad 1 \leq i \leq m, \quad \text { and } \\
\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}=\boldsymbol{x}_{2, i j}, \quad \boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}=\boldsymbol{A}^{22, i j}, \quad E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}^{T}\right\}=\boldsymbol{A}^{12, i j},\right. \\
E_{\mathfrak{q}}\left\{\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}\right)\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}\right)^{T}\right\}=\boldsymbol{A}^{12, i, j}, \quad 1 \leq i \leq m, \quad 1 \leq j \leq n_{i} .
\end{gathered}
$$

We employ these two results to derive streamlined MFVB algorithms for determining the optimal parameters of the densities in (20) for multilevel models having two-level and three-level random effects. The former specification is accommodated by Algorithm 2, the latter by Algorithm 3. Notice all the sub-blocks of $\boldsymbol{A}$ and components of $\boldsymbol{a}$ described in Results 1 and 2 can be updated simply performing linear transformations of matrices involving multiplications of sub-vectors of $\boldsymbol{y}$ and sub-matrices of $\boldsymbol{X}$ and $\boldsymbol{Z}$ which need to be computed only once instead of at each iteration of the associated algorithms.

Similar results are provided in [60] in terms of sparse least squares problems of the type $\left\|\boldsymbol{b}-\boldsymbol{B} \boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}\right\|^{2}$, after exploiting the equalities $\boldsymbol{A}=\boldsymbol{B}^{T} \boldsymbol{B}$ and $\boldsymbol{a}=\boldsymbol{B}^{T} \boldsymbol{b}$. This class of problems can be solved via efficient QR-decompositions, instead of sparse matrix problems of the type $\boldsymbol{A} \boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})}=\boldsymbol{a}$ which rely upon matrix inversion routines. Section 2.1 of [62] claims that the former class of problems and associated

## SolveTwoLevelSparseLeastSquares and SolveThreeLevelSparseLeastSquares

routines proposed in Appendix A of [60] are numerically preferred to the latter, since QR-decomposition methods are more computationally stable. However, streamlined MFVB approximations based on the QR-decomposition require an
additional set of matrices to be used at each algorithm iteration. Efficient QRdecomposition routines and functions for performing matrix multiplications with the associated $\boldsymbol{Q}$ and $\boldsymbol{R}$ matrices are not available in all standard computing environments and programming languages. Furthermore, the sub-blocks of $\boldsymbol{B}$ and components of $\boldsymbol{b}$ become particularly sparse for large $p$, as for the cases considered in this work, entailing onerous memory consumption and compromising efficiency of the QR-decomposition. For these reasons, we opt for routines based on sparse matrix linear system problems such as those of Results 1 and 2.

The convergence of Algorithms 2 and 3 can be assessed in several ways. One way is to monitor the relative increment of the variational lower bound to the marginal log-likelihood having general expression given by the first addend of (6). The algorithm can be stopped when the increment falls below a pre-specified threshold. The drawback is that much tedious algebra is required to derive an explicit lower bound expression. An alternative way is to measure the absolute relative increments among all the parameters updated at each algorithm iteration and stopping when the maximum increment falls below a pre-specified threshold. However, both the approaches can be computationally expensive, especially for large $m$ values, and these convergence checks may significantly slow down the overall streamlined iterations. Therefore, we suggest letting Algorithms 2 and 3 run for a reasonable number of iterations fixed in advance and increase the number of iterations if convergence checks suggest that the desired level of convergence has not been achieved.

A code bundle with R and $\mathrm{C}++$ implementations of our algorithms is publicly available at github.com/lucamaestrini/DMTWcode, in the GitHub repository DMTWcode.

```
Algorithm 2 Streamlined algorithm for obtaining the mean field variational Bayes approximate posterior density functions (20) for the parameters of the linear mixed model (17) with the two-level random effects specification (2). The approximation is based on the mean-field density restriction (19). The algorithm description requires more than one page and is continued on a subsequent page.
Data Inputs: \(\boldsymbol{y}_{i}\left(o_{i} \times 1\right), \boldsymbol{X}_{i}^{\mathrm{R}}\left(o_{i} \times p_{\mathrm{R}}\right), \boldsymbol{X}_{i}^{\mathrm{A}}\left(o_{i} \times p_{\mathrm{A}}\right), \boldsymbol{X}_{i}^{\mathrm{S}}\left(o_{i} \times p_{\mathrm{S}}\right), \boldsymbol{Z}_{i}\left(o_{i} \times q\right)\),
\(1 \leq i \leq m\). Build \(\boldsymbol{X}_{i}=\left[\boldsymbol{X}_{i}^{\mathrm{R}}\left|\boldsymbol{X}_{i}^{\mathrm{A}}\right| \boldsymbol{X}_{i}^{\mathrm{S}}\right]\left(o_{i} \times p\right)\).
Global-local prior type choice: Laplace, Horseshoe or NEG.
Hyperparameter Inputs: \(\boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{R}}}\left(p_{\mathrm{R}} \times 1\right), \boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{A}}}\left(p_{\mathrm{A}} \times 1\right)\),
\[
\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{R}}}\left(p_{\mathrm{R}} \times p_{\mathrm{R}}\right), \boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{A}}}\left(p_{\mathrm{A}} \times p_{\mathrm{A}}\right) \text { both symmetric and positive }
\]
\[
\text { definite, } \nu_{\sigma^{2}}, s_{\sigma^{2}}, s_{\tau^{2}}, \nu_{\boldsymbol{\Sigma}}, s_{\boldsymbol{\Sigma}, 1}, \ldots, s_{\boldsymbol{\Sigma}, q}>0 . \text { If NEG: } \lambda>0
\]
Initialize: \(\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)}>0, \mu_{\mathfrak{q}\left(1 / a_{\sigma^{2}}\right)}>0, \mu_{\mathfrak{q}\left(1 / \tau^{2}\right)}>0, \mu_{\mathfrak{q}\left(1 / a_{\tau^{2}}\right)}>0\), \(\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\zeta})}\left(p_{\mathrm{S}} \times 1\right), \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{a}_{\boldsymbol{\zeta}}\right)}\left(p_{\mathrm{S}} \times 1\right)\), \(\boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{-1}\right)}(q \times q), \boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}^{-1}\right)}(q \times q)\) both symmetric and positive definite.
\(\xi_{\mathfrak{q}\left(\sigma^{2}\right)} \longleftarrow \nu_{\sigma^{2}}+\sum_{i=1}^{m} o_{i} ; \quad \xi_{\mathfrak{q}(\boldsymbol{\Sigma})} \longleftarrow \nu_{\boldsymbol{\Sigma}}+m+2 q-2 ;\)
\(\xi_{\mathfrak{q}\left(a_{\sigma^{2}}\right)} \longleftarrow \nu_{\sigma^{2}}+1 ; \quad \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)} \longleftarrow \nu_{\boldsymbol{\Sigma}}+q ;\)
\(\xi_{\mathfrak{q}\left(\tau^{2}\right)} \longleftarrow p_{\mathrm{S}}+1 ; \quad \xi_{\mathfrak{q}\left(a_{\tau^{2}}\right)} \longleftarrow 2 ; \quad \boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}}} \longleftarrow\left\{\nu_{\boldsymbol{\Sigma}} \operatorname{diag}\left(s_{\boldsymbol{\Sigma}, 1}^{2}, \ldots, s_{\boldsymbol{\Sigma}, q}^{2}\right)\right\}^{-1}\).
Cycle until convergence:
continued on a subsequent page . . .
```

```
\(\overline{\text { Algorithm } 2 \text { continued. This is a continuation of the description of this }}\)
algorithm that commences on a preceding page.
    Compute \(\boldsymbol{a}_{1}, \boldsymbol{A}_{11},\left\{\boldsymbol{a}_{2, i}, \boldsymbol{A}_{22, i}, \boldsymbol{A}_{12, i}: 1 \leq i \leq m\right\}\) with expressions from Result 1.
    \(\mathcal{S}_{1} \longleftarrow \operatorname{SolveTwoLevelSparseMatrix}\left(\boldsymbol{a}_{1}, \boldsymbol{A}_{11},\left\{\boldsymbol{a}_{2, i}, \boldsymbol{A}_{22, i}, \boldsymbol{A}_{12, i}: 1 \leq i \leq m\right\}\right)\).
    \(\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})} \longleftarrow \boldsymbol{x}_{1}\) component of \(\mathcal{S}_{1} ; \quad \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta})} \longleftarrow \boldsymbol{A}^{11}\) component of \(\mathcal{S}_{1} ;\)
    \(\boldsymbol{\mu}_{\mathfrak{q}\left(\left(\beta^{\mathrm{S}}\right)^{2}\right)} \longleftarrow \operatorname{diagonal}\left(\boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{\beta}^{\mathrm{S}}\right)}+\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{\beta}^{\mathrm{S}}\right)} \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{\beta}^{\mathrm{S}}\right)}^{T}\right)\);
    \(\lambda_{\mathfrak{q}\left(\sigma^{2}\right)} \longleftarrow \mu_{\mathfrak{q}\left(1 / a_{\sigma^{2}}\right)} ; \quad \boldsymbol{\Lambda}_{\mathfrak{q}(\boldsymbol{\Sigma})} \longleftarrow M_{\mathfrak{q}\left(A_{\boldsymbol{\Sigma}}^{-1}\right)}\).
    For \(i=1, \ldots, m\) :
        \(\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)} \longleftarrow \boldsymbol{x}_{2, i}\) component of \(\mathcal{S}_{1} ; \quad \boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)} \longleftarrow \boldsymbol{A}^{22, i}\) component of \(\mathcal{S}_{1} ;\)
        \(E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}\right)^{T}\right\} \longleftarrow \boldsymbol{A}^{12, i}\) component of \(\mathcal{S}_{1} ;\)
        \(\lambda_{\mathfrak{q}\left(\sigma^{2}\right)} \longleftarrow \lambda_{\mathfrak{q}\left(\sigma^{2}\right)}+\left\|\boldsymbol{y}_{i}-\boldsymbol{X}_{i} \boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}-\boldsymbol{Z}_{i} \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}\right\|^{2}+\operatorname{tr}\left(\boldsymbol{X}_{i}^{T} \boldsymbol{X}_{i} \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\)
        \(+\operatorname{tr}\left(\boldsymbol{Z}_{i}^{T} \boldsymbol{Z}_{i} \boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}\right)+2 \operatorname{tr}\left[\boldsymbol{Z}_{i}^{T} \boldsymbol{X}_{i} E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}\right)^{T}\right\}\right] ;\)
        \(\boldsymbol{\Lambda}_{\mathfrak{q}(\boldsymbol{\Sigma})} \longleftarrow \boldsymbol{\Lambda}_{\mathfrak{q}(\boldsymbol{\Sigma})}+\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)} \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}^{T}+\boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)} ;\)
    \(\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \longleftarrow \xi_{\mathfrak{q}\left(\sigma^{2}\right)} / \lambda_{\mathfrak{q}\left(\sigma^{2}\right)} ; \quad \boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{-1}\right)} \longleftarrow\left(\xi_{\mathfrak{q}(\boldsymbol{\Sigma})}-q+1\right) \boldsymbol{\Lambda}_{\mathfrak{q}(\boldsymbol{\Sigma})}^{-1} ;\)
    \(\lambda_{\mathfrak{q}\left(a_{\sigma^{2}}\right)} \longleftarrow \mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)}+1 /\left(\nu_{\sigma^{2}} s_{\sigma^{2}}^{2}\right) ; \quad \mu_{\mathfrak{q}\left(1 / a_{\sigma^{2}}\right)} \longleftarrow \xi_{\mathfrak{q}\left(a_{\sigma^{2}}\right)} / \lambda_{\mathfrak{q}\left(a_{\sigma^{2}}\right)} ;\)
    \(\lambda_{\mathfrak{q}\left(\tau^{2}\right)} \longleftarrow \mu_{\mathfrak{q}\left(1 / a^{2}\right)}+\boldsymbol{\mu}_{\mathfrak{q}(\zeta)}^{T} \boldsymbol{\mu}_{\mathfrak{q}\left(\left(\beta^{S}\right)^{2}\right)} ; \quad \mu_{\mathfrak{q}\left(1 / \tau^{2}\right)} \longleftarrow \xi_{\mathfrak{q}\left(\tau^{2}\right)} / \lambda_{\mathfrak{q}\left(\tau^{2}\right)} ;\)
    \(\lambda_{\mathfrak{q}\left(a_{\tau^{2}}\right)} \longleftarrow \mu_{\mathfrak{q}\left(1 / \tau^{2}\right)}+1 / s_{\tau^{2}}^{2} ; \quad \mu_{\mathfrak{q}\left(1 / a_{\tau^{2}}\right)} \longleftarrow \xi_{\mathfrak{q}\left(a_{\tau^{2}}\right)} / \lambda_{\mathfrak{q}\left(a_{\tau^{2}}\right)} ;\)
    \(g \longleftarrow \frac{1}{2} \mu_{\mathfrak{q}\left(1 / \tau^{2}\right)} \mu_{\mathfrak{q}\left(\left(\beta^{\mathrm{S}}\right)^{2}\right)} ;\)
    If Laplace: \(\boldsymbol{\mu}_{\mathfrak{q}(\zeta)} \longleftarrow \sqrt{\mathbf{1 / ( 2 g )}}\);
    If Horseshoe: \(\lambda_{\mathfrak{q}(\zeta)} \longleftarrow \mu_{\mathfrak{q}\left(a_{\zeta}\right)}+g ; \quad \mu_{\mathfrak{q}(\zeta)} \longleftarrow \mathbf{1} / \lambda_{\mathfrak{q}(\zeta)} ;\)
        \(\lambda_{\mathfrak{q}\left(a_{\zeta}\right)} \longleftarrow \mu_{\mathfrak{q}(\zeta)}+1 ; \quad \mu_{\mathfrak{q}\left(a_{\zeta}\right)} \longleftarrow 1 / \lambda_{\mathfrak{q}\left(a_{\zeta}\right)} ;\)
    If NEG: \(\boldsymbol{\lambda}_{\mathfrak{q}(\zeta)} \longleftarrow 2 \boldsymbol{\mu}_{\mathfrak{q}\left(a_{\zeta}\right)} ; \quad \boldsymbol{\mu}_{\mathfrak{q}(\zeta)} \longleftarrow \sqrt{\lambda_{\mathfrak{q}(\zeta)} /(2 \boldsymbol{g})}\);
        \(\mu_{\mathfrak{q}(1 / \zeta)} \longleftarrow 1 / \mu_{\mathfrak{q}(\zeta)}+1 /\left(2 \mu_{\mathfrak{q}\left(a_{\zeta}\right)}\right) ;\)
        \(\lambda_{\mathfrak{q}\left(a_{\zeta}\right)} \longleftarrow \mu_{\mathfrak{q}(1 / \zeta)}+\mathbf{1} ; \quad \mu_{\mathfrak{q}\left(a_{\zeta}\right)} \longleftarrow(\lambda+1)\left(\mathbf{1} / \lambda_{\mathfrak{q}\left(a_{\zeta}\right)}\right) ;\)
    \(\boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)} \longleftarrow \operatorname{diag}\left\{\operatorname{diagonal}\left(\boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{-1}\right)}\right)\right\}+\boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}}} ; \quad \boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}^{-1}\right)} \longleftarrow \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)} \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)}^{-1}\).
Outputs: \(\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}, \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta})},\left\{\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}, \boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}, E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}\right)}\right)^{T}\right\}: 1 \leq i \leq m\right\}\),
    \(\xi_{\mathfrak{q}\left(\sigma^{2}\right)}, \lambda_{\mathfrak{q}\left(\sigma^{2}\right)}, \xi_{\mathfrak{q}\left(a_{\sigma^{2}}\right)}, \lambda_{\mathfrak{q}\left(a_{\sigma^{2}}\right)}, \xi_{\mathfrak{q}\left(\tau^{2}\right)}, \lambda_{\mathfrak{q}\left(\tau^{2}\right)}, \xi_{\mathfrak{q}\left(a_{\tau^{2}}\right)}, \lambda_{\mathfrak{q}\left(a_{\tau^{2}}\right)}\),
    if Laplace: \(\boldsymbol{\mu}_{\mathfrak{q}(\zeta)}\), if Horseshoe: \(\left\{\boldsymbol{\lambda}_{\mathfrak{q}(\zeta)}, \boldsymbol{\lambda}_{\mathfrak{q}\left(a_{\zeta}\right)}\right\}\), if NEG: \(\left\{\boldsymbol{\mu}_{\mathfrak{q}(\zeta)}, \boldsymbol{\lambda}_{\mathfrak{q}(\zeta)}, \boldsymbol{\lambda}_{\mathfrak{q}\left(a_{\zeta}\right)}\right\}\),
    \(\xi_{\mathfrak{q}(\boldsymbol{\Sigma})}, \boldsymbol{\Lambda}_{\mathfrak{q}(\boldsymbol{\Sigma})}, \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)}, \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right)}\).
```

```
Algorithm 3 Streamlined algorithm for obtaining the mean field variational
Bayes approximate posterior density functions (20) for the parameters of the
linear mixed model (17) with the three-level random effects specification (4). The
approximation is based on the mean-field density restriction (19). The algorithm
description requires more than one page and is continued on a subsequent page.
Data Inputs: \(\boldsymbol{y}_{i j}\left(o_{i j} \times 1\right), \boldsymbol{X}_{i j}^{\mathrm{R}}\left(o_{i j} \times p_{\mathrm{R}}\right), \boldsymbol{X}_{i j}^{\mathrm{A}}\left(o_{i j} \times p_{\mathrm{A}}\right), \boldsymbol{X}_{i j}^{\mathrm{S}}\left(o_{i j} \times p_{\mathrm{S}}\right)\),
    \(\boldsymbol{Z}_{i j}^{\mathrm{L} 1}\left(o_{i j} \times q_{1}\right), \boldsymbol{Z}_{i j}^{\mathrm{L} 2}\left(o_{i j} \times q_{2}\right), 1 \leq i \leq m, 1 \leq j \leq n_{i}\).
    Build \(\boldsymbol{X}_{i j}=\left[\boldsymbol{X}_{i j}^{\mathrm{R}}\left|\boldsymbol{X}_{i j}^{\mathrm{A}}\right| \boldsymbol{X}_{i j}^{\mathrm{S}}\right]\left(n_{i} \times p\right)\).
Global-local prior type choice: Laplace, Horseshoe or NEG.
Hyperparameter Inputs: \(\boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{R}}}\left(p_{\mathrm{R}} \times 1\right), \boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{A}}}\left(p_{\mathrm{A}} \times 1\right)\),
                    \(\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{R}}}\left(p_{\mathrm{R}} \times p_{\mathrm{R}}\right)\) and \(\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{A}}}\left(p_{\mathrm{A}} \times p_{\mathrm{A}}\right)\) both symmetric
                    and positive definite, \(\nu_{\sigma^{2}}, s_{\sigma^{2}}, s_{\tau^{2}}, \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}, \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\),
                    \(s_{\boldsymbol{\Sigma}^{\mathrm{L} 1}, 1}, \ldots, s_{\boldsymbol{\Sigma}^{\mathrm{L} 1}, q_{1}}, s_{\boldsymbol{\Sigma}^{\mathrm{L} 2}, 1}, \ldots, s_{\boldsymbol{\Sigma}^{\mathrm{L} 2}, q_{2}}>0\). If NEG: \(\lambda>0\).
Initialize: \(\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)}>0, \mu_{\mathfrak{q}\left(1 / a_{\sigma^{2}}\right)}>0, \mu_{\mathfrak{q}\left(1 / \tau^{2}\right)}>0, \mu_{\mathfrak{q}\left(1 / a_{\tau^{2}}\right)}>0\),
    \(\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\zeta})}\left(p_{\mathrm{S}} \times 1\right), \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{a}_{\zeta}\right)}\left(p_{\mathrm{S}} \times 1\right)\),
    \(\boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)^{-1}\right)}\left(q_{1} \times q_{1}\right), \boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)^{-1}\right)}\left(q_{2} \times q_{2}\right), \boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{A}_{\left.\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)}^{-1}\right)}\left(q_{1} \times q_{1}\right)\),
    \(\boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}^{-1}\right)}\left(q_{2} \times q_{2}\right)\) all symmetric and positive definite.
\(\xi_{\mathfrak{q}\left(\sigma^{2}\right)} \longleftarrow \nu_{\sigma^{2}}+\sum_{i=1}^{m} \sum_{j=1}^{n_{i}} o_{i j} ; \quad \xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)} \longleftarrow \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}+m+2 q_{1}-2 ;\)
\(\xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)} \longleftarrow \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}+\sum_{i=1}^{m} n_{i}+2 q_{2}-2 ;\)
\(\xi_{\mathfrak{q}\left(a_{\sigma^{2}}\right)} \longleftarrow \nu_{\sigma^{2}}+1 ; \quad \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}\right)} \longleftarrow \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}+q_{1} ; \quad \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\right)} \longleftarrow \nu_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}+q_{2} ;\)
\(\xi_{\mathfrak{q}\left(\tau^{2}\right)} \longleftarrow p_{\mathrm{S}}+1 ; \quad \xi_{\mathfrak{q}\left(a_{\tau^{2}}\right)} \stackrel{\Sigma}{ } \longleftarrow 2 ;\)
\(\boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}} \longleftarrow\left\{\nu_{\boldsymbol{\Sigma}^{\mathrm{L} 1}} \operatorname{diag}\left(s_{\boldsymbol{\Sigma}^{\mathrm{L} 1}, 1}^{2}, \ldots, s_{\boldsymbol{\Sigma}^{\mathrm{L} 1}, q_{1}}^{2}\right)\right\}^{-1} ;\)
\(\boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}} \longleftarrow\left\{\nu_{\boldsymbol{\Sigma}^{\mathrm{L} 2}} \operatorname{diag}\left(s_{\boldsymbol{\Sigma}^{\mathrm{L} 2}, 1}^{2}, \ldots, s_{\boldsymbol{\Sigma}^{\mathrm{L} 2}, q_{2}}^{2}\right)\right\}^{-1}\).
Cycle until convergence:
    Compute \(\boldsymbol{a}_{1}, \boldsymbol{A}_{11},\left\{\boldsymbol{a}_{2, i}, \boldsymbol{A}_{22, i}, \boldsymbol{A}_{12, i}: 1 \leq i \leq m\right\},\left\{\boldsymbol{a}_{2, i j}, \boldsymbol{A}_{22, i j}, \boldsymbol{A}_{12, i j}, \boldsymbol{A}_{12, i, j}:\right.\)
    \(\left.1 \leq i \leq m, 1 \leq j \leq n_{i}\right\}\) with expressions from Result 2 .
    \(\mathcal{S}_{2} \longleftarrow \operatorname{SolveThreeLevelSparseMatrix}\left(\boldsymbol{a}_{1}, \boldsymbol{A}_{11},\left\{\boldsymbol{a}_{2, i}, \boldsymbol{A}_{22, i}, \boldsymbol{A}_{12, i}: 1 \leq i \leq m\right\}\right.\),
    \(\left.\left\{\boldsymbol{a}_{2, i j}, \boldsymbol{A}_{22, i j}, \boldsymbol{A}_{12, i j}, \boldsymbol{A}_{12, i, j}: 1 \leq i \leq m, 1 \leq j \leq n_{i}\right\}\right)\).
    \(\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})} \longleftarrow \boldsymbol{x}_{1}\) component of \(\mathcal{S}_{2} ; \quad \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta})} \longleftarrow \boldsymbol{A}^{11}\) component of \(\mathcal{S}_{2} ;\)
    \(\boldsymbol{\mu}_{\mathfrak{q}\left(\left(\boldsymbol{\beta}^{\mathrm{S}}\right)^{2}\right)} \longleftarrow\) diagonal \(\left(\boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{\beta}^{\mathrm{S}}\right)}+\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{\beta}^{\mathrm{S}}\right)} \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{\beta}^{\mathrm{S}}\right)}^{T}\right) ;\)
    \(\lambda_{\mathfrak{q}\left(\sigma^{2}\right)} \longleftarrow \mu_{\mathfrak{q}\left(1 / a_{\sigma^{2}}\right)} ; \quad \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)} \longleftarrow \boldsymbol{M}_{\mathfrak{q}\left(A_{\Sigma^{\mathrm{L} 1}}^{-1}\right)} ; \quad \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)} \longleftarrow \boldsymbol{M}_{\mathfrak{q}\left(\boldsymbol{A}_{\Sigma^{\mathrm{L} 2}}^{-1}\right)}\).
```

    For \(i=1, \ldots, m\) :
    \(\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)} \longleftarrow \boldsymbol{x}_{2, i}\) component of \(\mathcal{S}_{2} ; \quad \boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)} \longleftarrow \boldsymbol{A}^{22, i}\) component of \(\mathcal{S}_{2} ;\)
    \(E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}\right)^{T}\right\} \longleftarrow \boldsymbol{A}^{12, i}\) component of \(\mathcal{S}_{2} ;\)
    \(\boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)} \longleftarrow \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)}+\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)} \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}^{T}+\boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}\).
        For \(j=1, \ldots, n_{i}\) :
        \(\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)} \longleftarrow \boldsymbol{x}_{2, i j}\) component of \(\mathcal{S}_{2}\);
        \(\boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)} \longleftarrow \boldsymbol{A}^{22, i j}\) component of \(\mathcal{S}_{2} ;\)
        continued on a subsequent page ...
    ```
\(\overline{\text { Algorithm } 3 \text { continued. This is a continuation of the description of this }}\)
algorithm that commences on a preceding page.
    \(E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}\right)^{T}\right\} \longleftarrow \boldsymbol{A}^{12, i j}\) component of \(\mathcal{S}_{2} ;\)
    \(E_{\mathfrak{q}}\left\{\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}\right)\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}\right)^{T}\right\} \longleftarrow \boldsymbol{A}^{12, i, j}\) component of \(\mathcal{S}_{2} ;\)
    \(\lambda_{\mathfrak{q}\left(\sigma^{2}\right)} \longleftarrow \lambda_{\mathfrak{q}\left(\sigma^{2}\right)}+\left\|\boldsymbol{y}_{i j}-\boldsymbol{X}_{i j} \boldsymbol{\mu}_{\mathfrak{q}(\beta)}-\boldsymbol{Z}_{i j}^{\mathrm{L} 1} \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}-\boldsymbol{Z}_{i j}^{\mathrm{L} 2} \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}\right\|^{2}\)
        \(+\operatorname{tr}\left(\boldsymbol{X}_{i j}^{T} \boldsymbol{X}_{i j} \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta})}\right)+\operatorname{tr}\left(\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 1}\right)^{T} \boldsymbol{Z}_{i j}^{\mathrm{L} 1} \boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}\right)\)
        \(+\operatorname{tr}\left(\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 2}\right)^{T} \boldsymbol{Z}_{i j}^{\mathrm{L} 2} \boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}\right)\)
        \(+2 \operatorname{tr}\left[\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 1}\right)^{T} \boldsymbol{X}_{i j} E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}\right)^{T}\right\}\right]\)
        \(+2 \operatorname{tr}\left[\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 2}\right)^{T} \boldsymbol{X}_{i j} E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}\right)^{T}\right\}\right]\)
        \(+2 \operatorname{tr}\left[\left(\boldsymbol{Z}_{i j}^{\mathrm{L} 1}\right)^{T} \boldsymbol{Z}_{i j}^{\mathrm{L} 2} E_{\mathfrak{q}}\left\{\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}\right)\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}\right)^{T}\right\}\right] ;\)
        \(\boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)} \longleftarrow \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)}+\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)} \boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}^{T}+\boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}\).
    \(\mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)} \longleftarrow \xi_{\mathfrak{q}\left(\sigma^{2}\right)} / \lambda_{\mathfrak{q}\left(\sigma^{2}\right)} ;\)
    \(\boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)^{-1}\right)} \longleftarrow\left(\xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)}-q_{1}+1\right) \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)}^{-1} ;\)
    \(\boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)^{-1}\right)} \longleftarrow\left(\xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)}-q_{2}+1\right) \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)}^{-1} ;\)
    \(\lambda_{\mathfrak{q}\left(a_{\sigma^{2}}\right)} \longleftarrow \mu_{\mathfrak{q}\left(1 / \sigma^{2}\right)}+1 /\left(\nu_{\sigma^{2}} s_{\sigma^{2}}^{2}\right) ; \quad \mu_{\mathfrak{q}\left(1 / a_{\sigma^{2}}\right)} \longleftarrow \xi_{\mathfrak{q}\left(a_{\sigma^{2}}\right)} / \lambda_{\mathfrak{q}\left(a_{\sigma^{2}}\right)} ;\)
    \(\lambda_{\mathfrak{q}\left(\tau^{2}\right)} \longleftarrow \mu_{\mathfrak{q}\left(1 / a^{2}\right)}+\boldsymbol{\mu}_{\mathfrak{q}(\zeta)}^{T} \boldsymbol{\mu}_{\mathfrak{q}\left(\left(\beta^{S}\right)^{2}\right)} ; \quad \mu_{\mathfrak{q}\left(1 / \tau^{2}\right)} \longleftarrow \xi_{\mathfrak{q}\left(\tau^{2}\right)} / \lambda_{\mathfrak{q}\left(\tau^{2}\right)} ;\)
    \(\lambda_{\mathfrak{q}\left(a_{\tau^{2}}\right)} \longleftarrow \mu_{\mathfrak{q}\left(1 / \tau^{2}\right)}+1 / s_{\tau^{2}}^{2} ; \quad \mu_{\mathfrak{q}\left(1 / a_{\tau^{2}}\right)} \longleftarrow \xi_{\mathfrak{q}\left(a_{\tau^{2}}\right)} / \lambda_{\mathfrak{q}\left(a_{\tau^{2}}\right)} ;\)
    \(g \longleftarrow \frac{1}{2} \mu_{\mathfrak{q}\left(1 / \tau^{2}\right)} \boldsymbol{\mu}_{\mathfrak{q}\left(\left(\beta^{\mathrm{S}}\right)^{2}\right)} ;\)
    If Laplace: \(\boldsymbol{\mu}_{\mathrm{q}(\zeta)} \longleftarrow \sqrt{\mathbf{1 / ( 2 g )}}\);
    If Horseshoe: \(\lambda_{\mathfrak{q}(\zeta)} \longleftarrow \mu_{\mathfrak{q}\left(a_{\zeta}\right)}+g ; \quad \mu_{\mathfrak{q}(\zeta)} \longleftarrow 1 / \lambda_{\mathfrak{q}(\zeta)}\);
        \(\lambda_{\mathfrak{q}\left(a_{\zeta}\right)} \longleftarrow \mu_{\mathfrak{q}(\zeta)}+1 ; \quad \mu_{\mathfrak{q}\left(a_{\zeta}\right)} \longleftarrow 1 / \lambda_{\mathfrak{q}\left(a_{\zeta}\right)} ;\)
    If NEG: \(\lambda_{\mathfrak{q}(\zeta)} \longleftarrow 2 \mu_{\mathfrak{q}\left(a_{\zeta}\right)} ; \quad \mu_{\mathfrak{q}(\zeta)} \longleftarrow \sqrt{\lambda_{\mathfrak{q}(\zeta)} /(2 g)}\);
        \(\mu_{\mathfrak{q}(1 / \zeta)} \longleftarrow 1 / \mu_{\mathfrak{q}(\zeta)}+1 /\left(2 \mu_{\mathfrak{q}\left(a_{\zeta}\right)}\right) ;\)
        \(\lambda_{\mathfrak{q}\left(a_{\zeta}\right)} \longleftarrow \mu_{\mathfrak{q}(1 / \zeta)}+\mathbf{1} ; \quad \mu_{\mathfrak{q}\left(a_{\zeta}\right)} \longleftarrow(\lambda+1)\left(\mathbf{1} / \lambda_{\mathfrak{q}\left(a_{\zeta}\right)}\right) ;\)
    \(\boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}\right)} \longleftarrow \operatorname{diag}\left\{\operatorname{diagonal}\left(\boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)^{-1}\right)}\right)\right\}+\boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}} ;\)
    \(\boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L}}}\right)^{-1}\right)} \longleftarrow \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}\right)} \boldsymbol{\Lambda}_{\boldsymbol{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L}}}\right)}^{-1}\).
    \(\boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\right)} \longleftarrow \operatorname{diag}\left\{\operatorname{diagonal}\left(\boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)^{-1}\right)}\right)\right\}+\boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}} ;\)
    \(\boldsymbol{M}_{\mathfrak{q}\left(\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\right)^{-1}\right)} \longleftarrow \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}\right)} \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}{ }^{\mathrm{L} 2}\right)}^{-1}\).
Outputs: \(\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}, \boldsymbol{\Sigma}_{\mathfrak{q}(\boldsymbol{\beta})},\left\{\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}, \boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}, E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i}^{\mathrm{L} 1}\right)}\right)^{T}\right\}: 1 \leq i \leq m\right\}\),
    \(\left\{\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}, \boldsymbol{\Sigma}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}, E_{\mathfrak{q}}\left\{\left(\boldsymbol{\beta}-\boldsymbol{\mu}_{\mathfrak{q}(\boldsymbol{\beta})}\right)\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}-\boldsymbol{\mu}_{\mathfrak{q}\left(\boldsymbol{u}_{i j}^{\mathrm{L} 2}\right)}\right)^{T}\right\}: 1 \leq i \leq m, 1 \leq j \leq n_{i}\right\}\),
    \(\xi_{\mathfrak{q}\left(\sigma^{2}\right)}, \lambda_{\mathfrak{q}\left(\sigma^{2}\right)}, \xi_{\mathfrak{q}\left(a_{\sigma^{2}}\right)}, \lambda_{\mathfrak{q}\left(a_{\sigma^{2}}\right)}, \xi_{\mathfrak{q}\left(\tau^{2}\right)}, \lambda_{\mathfrak{q}\left(\tau^{2}\right)}, \xi_{\mathfrak{q}\left(a_{\tau^{2}}\right)}, \lambda_{\mathfrak{q}\left(a_{\tau^{2}}\right)}\),
    if Laplace: \(\boldsymbol{\mu}_{\mathfrak{q}(\zeta)}\), if Horseshoe: \(\left\{\boldsymbol{\lambda}_{\mathfrak{q}(\zeta)}, \boldsymbol{\lambda}_{\mathfrak{q}\left(a_{\zeta}\right)}\right\}\), if NEG: \(\left\{\boldsymbol{\mu}_{\mathfrak{q}(\zeta)}, \boldsymbol{\lambda}_{\mathfrak{q}(\zeta)}, \boldsymbol{\lambda}_{\mathfrak{q}\left(a_{\zeta}\right)}\right\}\),
    \(\left.\xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)}, \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 1}\right)}, \xi_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)}, \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{\Sigma}^{\mathrm{L} 2}\right)}, \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}{ }^{\mathrm{L} 1}\right)}, \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}\right)}, \xi_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}{ }^{\mathrm{L} 2}\right)}, \boldsymbol{\Lambda}_{\mathfrak{q}\left(\boldsymbol{A}_{\boldsymbol{\Sigma}}\right.} \boldsymbol{L}^{\mathrm{L} 2}\right)\).
```


## 6. Simulation study investigations

We discuss the results of a simulation study conducted to assess:

1. the accuracy of the optimal $\mathfrak{q}^{*}$-density approximations compared to the marginal posterior density functions obtained via MCMC, when globallocal priors are specified;
2. fixed effects selection performances via the SAVS procedure for effectively discriminating the relevant fixed effects from those being irrelevant;
3. computational timings and memory storage requirements for both naïve and streamlined algorithm implementations, especially when the number of parameters increases.
Notice that the accuracy and variable selection performances are not affected by the use of streamlined updates in place of naïve counterparts, given that both the implementations converge to the same solution.

The simulation study was performed on a MacBook Pro laptop with a 1.4 gigahertz processor and 8 gigabytes of random access memory. To allow for maximal speed, both the MFVB algorithms and corresponding MCMC sampling schemes were implemented in C++ employing the Armadillo library [73] and executed into an R environment using the RcppArmadillo package [22].

The simulation study focused on three-level random effect models, which give rise to more complex three-level sparse matrix structures. We simulated 50 datasets according to model specification (17) with $m=100$ groups, each with $n_{i}=15$ sub-groups, each having $o_{i j}=20$ units, for $1 \leq i \leq 100,1 \leq j \leq 15$. We included a random intercept and a random slope for both the group and sub-group levels $\left(q_{1}=q_{2}=p_{\mathrm{R}}=2\right)$ with true parameter values being $\boldsymbol{\beta}^{\mathrm{R}}=$ [0.58 1.98] ${ }^{T}$, and $p_{\mathrm{A}}=3$ additional fixed effects having true parameter values $\boldsymbol{\beta}^{\mathrm{A}}=\left[\begin{array}{lll}0.7 & -0.9 & 1.8\end{array}\right]^{T}$. We also considered a sparse design setting with $p_{\mathrm{S}}=50$ fixed effects such that the first 10 of them were $\beta_{1}^{\mathrm{S}}=1.91, \beta_{2}^{\mathrm{S}}=1.96, \beta_{3}^{\mathrm{S}}=$ $-0.10, \beta_{4}^{\mathrm{S}}=1.62, \beta_{5}^{\mathrm{S}}=-1.45, \beta_{6}^{\mathrm{S}}=-1.53, \beta_{7}^{\mathrm{S}}=0.24, \beta_{8}^{\mathrm{S}}=1.76, \beta_{9}^{\mathrm{S}}=1.79$ and $\beta_{10}^{\mathrm{S}}=-0.15$, while the remaining 40 were assumed to be irrelevant and hence the corresponding parameters had true value equal to zero $\left(\beta_{h}^{S}=0\right.$ for each $11 \leq$ $h \leq 50$ ). The true variance parameter $\sigma^{2}$ was fixed to 0.7 . The random effects vectors $\boldsymbol{u}_{i}^{\mathrm{L} 1}$ and $\boldsymbol{u}_{i j}^{\mathrm{L} 2}$ were respectively generated independently from $\mathrm{N}\left(\mathbf{0}, \boldsymbol{\Sigma}^{\mathrm{L} 1}\right)$ and $\mathrm{N}\left(\mathbf{0}, \boldsymbol{\Sigma}^{\mathrm{L} 2}\right)$ distributions, having

$$
\boldsymbol{\Sigma}^{\mathrm{L} 1}=\left[\begin{array}{cc}
0.42 & -0.09 \\
-0.09 & 0.52
\end{array}\right] \quad \text { and } \quad \boldsymbol{\Sigma}^{\mathrm{L} 2}=\left[\begin{array}{cc}
0.80 & -0.24 \\
-0.24 & 0.75
\end{array}\right]
$$

For each data replication, the slope-associated column of $\boldsymbol{X}^{\mathrm{R}}$ was generated from a standard Gaussian distribution, while the rows of $\boldsymbol{X}^{\mathrm{A}}$ and $\boldsymbol{X}^{\mathrm{S}}$ were generated from two multivariate Gaussian distributions having zero mean vector and covariance matrices generated from $\mathrm{Wishart}\left(p_{\mathrm{A}}, \boldsymbol{I}\right)$ and $\mathrm{Wishart}\left(p_{\mathrm{S}}, \boldsymbol{I}\right)$ distributions, respectively. This strategy produces covariates with different variability and non-zero correlations that better mimic a real-data scenario. The setting
embeds a fixed effects design matrix $\boldsymbol{X}$ of size $30,000 \times 55$, and a sparse random effects design matrix $\boldsymbol{Z}$ of size $30,000 \times 3,200$ with 96 millions cells, of which $99.875 \%$ are zeros.

We proceed fitting model (17) for each data replication using diffuse priors with hyperparameters $\boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{R}}}=\boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{A}}}=\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{R}}}=\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{A}}}=10^{10} \boldsymbol{I}, \nu_{\sigma^{2}}=1$, $\nu_{\boldsymbol{\Sigma}^{\mathrm{L} 1}}=\nu_{\boldsymbol{\Sigma}^{\mathrm{L} 2}}=2, s_{\sigma^{2}}=s_{\boldsymbol{\Sigma}^{\mathrm{L} 1}, 1}=s_{\boldsymbol{\Sigma}^{\mathrm{L} 1,2}}=s_{\boldsymbol{\Sigma}^{\mathrm{L} 2}, 1}=s_{\boldsymbol{\Sigma}^{\mathrm{L} 2}, 2}=10^{5}$. Without loss of generality, a Horseshoe prior for all the elements of $\boldsymbol{\beta}^{\mathrm{S}}$ have been specified, with $s_{\tau^{2}}=10^{5}$ to limit prior information about the degree of sparsity. Along all the simulation study, the MFVB approximations were obtained by running 200 iterations of the streamlined or naïve algorithms, i.e. four times the number of iterations used in the numerical experiments of [60]. This number of MFVB iterations was chosen to guarantee appropriate convergence of the variational algorithms, as well as a fair computational time comparison between naïve and streamlined MFVB. [53] suggest running variational algorithms until the relative change in the variational parameters reaches a certain threshold. In our numerical experiments, the 200 MFVB iterations guaranteed that the relative change in the variational parameter estimates was smaller that $10^{-3}$ and for this or smaller values the accuracy and fixed effect selection performance of MFVB was not impacted. Variational algorithm convergence could also be assessed by monitoring the lower bound at every iteration or after a certain number of iterations; however, deriving the variational lower bound is subject to calculation errors and its computation can be time consuming.

### 6.1. Accuracy assessment

We measured the quality of the variational approximations using the accuracy index proposed in Section 3.4 of [24]. For a generic univariate parameter $\theta \in \Theta$, this is defined as

$$
\begin{equation*}
\operatorname{Accuracy}(\theta) \equiv\left(1-\frac{1}{2} \int_{\Theta}\left|\mathfrak{q}^{*}(\theta)-\mathfrak{p}(\theta \mid \boldsymbol{y})\right| \mathrm{d} \theta\right) \% \tag{22}
\end{equation*}
$$

where $\mathfrak{p}(\theta \mid \boldsymbol{y})$ is the marginal posterior density of $\theta$ and $\mathfrak{q}^{*}(\theta)$ is the associated optimal approximating density function obtained via MFVB. This index takes values between $0 \%$ and $100 \%$, with a score of $100 \%$ indicating perfect matching between $\mathfrak{q}^{*}(\theta)$ and $\mathfrak{p}(\theta \mid \boldsymbol{y})$, and $0 \%$ if the densities have no overlapping mass. In practice, computation of the true $\mathfrak{p}(\theta \mid \boldsymbol{y})$ is numerically challenging, so we employed binned kernel density estimation with direct plug-in bandwidth selection [see Section 3.6.1 of 86] and applied it to the MCMC samples using the R package KernSmooth [85]. The MCMC samplings were performed using a warmup of length 5,000 followed by 100,000 iterations, to which we applied a thinning value of 20 . The integration in (22) was then accurately performed via trapezoidal numerical quadrature. Figure 2 shows the accuracy results through boxplots for the 50 data replicates. The boxplots refer to all the entries of $\boldsymbol{\beta}^{\mathrm{R}}$ and $\boldsymbol{\beta}^{\mathrm{A}}$, to 10 elements of $\beta^{\mathrm{S}}$ chosen such that 5 of them have non-zero values $\left(\beta_{1}^{\mathrm{S}}, \beta_{3}^{\mathrm{S}}\right.$, $\beta_{4}^{\mathrm{S}}, \beta_{6}^{\mathrm{S}}$ and $\left.\beta_{8}^{\mathrm{S}}\right)$ and 5 are null $\left(\beta_{11}^{\mathrm{S}}, \beta_{23}^{\mathrm{S}}, \beta_{24}^{\mathrm{S}}, \beta_{36}^{\mathrm{S}}\right.$ and $\left.\beta_{48}^{\mathrm{S}}\right)$, the intercept- and


FIG 2. Side-by-side boxplots of the accuracy scores from the simulation study for a selection of model parameters and random effects. Outliers are displayed as solid points.
slope-associated elements of $\boldsymbol{u}_{i}^{\mathrm{L} 1}$ and $\boldsymbol{u}_{i j}^{\mathrm{L} 2}$, for $1 \leq i \leq 3$ and $j=1$, the entries of $\boldsymbol{\Sigma}^{\mathrm{L} 1}$ and $\boldsymbol{\Sigma}^{\mathrm{L} 2}, \sigma^{2}$, and $\tau^{2}$. The intercept- and slope-associated parameters are identified by the subscripts "int" and "slp", respectively.

Variational approximations showed high accuracy scores for all the model parameters considered and across the different data replications. All the fixed effects subject to selection and having non-zero true values exhibited accuracy scores greater than $90 \%$. In contrast, those having true values equal to zero had lower accuracy scores between $75 \%$ and $85 \%$ due to the spiky marginal posterior densities being approximated by Gaussian variational densities. All the other fixed effects parameters, random effects and variance parameters showed accuracy scores greater than $90 \%$. The accuracy scores of the global variance parameter $\tau^{2}$ were under $50 \%$. Despite not directly shown, all the $\mathfrak{q}^{*}\left(\zeta_{h}\right)$ 's had accuracy scores between $75 \%$ and $80 \%$. Similar results were obtained for the alternative global-local prior specifications.

Figure 3 displays the MFVB and MCMC approximate posterior densities, together with the associated accuracy scores, obtained for the first data replication of the simulation study. These plots allow to visually assess the quality of the approximation and show that the approximate posterior density functions are generally concentrated around the true parameter values. Notice that the global-local prior shrinks the MCMC marginal posterior densities of the $\beta^{\mathrm{S}}$ parameters having null true value towards zero. The corresponding $\mathfrak{q}^{*}$-densities are Gaussian and, although they are not able to capture peak and tail behaviors, they provide satisfactory approximations. Notice also that the $\tau^{2}$ approximate posterior densities are concentrated around a very small value, as expected for the sparse design setting we considered in the simulation study.


FIG 3. Approximate posterior density functions of some of the three-lever random effects model parameters obtained from the first replication of the simulation study. Each plot shows the optimal approximate posterior density function $\mathfrak{q}^{*}(\theta)$ obtained via MFVB (black dashed curves) and the MCMC-based $\mathfrak{p}(\theta \mid \boldsymbol{y})$ densities (grey curves). A vertical line indicates the parameter true value. The percentages of accuracy are also provided.

### 6.2. Fixed effects selection assessment

We assessed fixed effects selection performances by running Algorithm 3 on the same 50 data replications, also experimenting the other two global-local priors considered in this work: the Laplace and Normal-Exponential-Gamma with $\lambda=0.25$. We also considered the Gaussian prior specification described in [60] with hyperparameters $\boldsymbol{\mu}_{\boldsymbol{\beta}^{\mathrm{s}}}=\mathbf{0}$ and $\boldsymbol{\Sigma}_{\boldsymbol{\beta}^{\mathrm{s}}}=10^{10} \boldsymbol{I}$. Then, for each data replication and each of the different four prior specifications considered we used the approximating densities corresponding to fixed effects subject to selection to perform the SAVS procedure of Algorithm 1. We employed this procedure
also for the approximate posteriors obtained from the Gaussian prior to provide a comparison with a prior that does not belong to the global-local family.

Let TP (true positives) denote the number of selected fixed effects having true value different from zero and TN (true negatives) denote the number of unselected irrelevant fixed effects having true value equal to zero. We measured the fixed effects selection performance for each prior choice using the $F_{1}$-score [82], which is defined as:

$$
\begin{equation*}
F_{1} \equiv\left(\frac{2 \times \text { precision } \times \text { recall }}{\text { precision }+ \text { recall }}\right) \times 100 \% \tag{23}
\end{equation*}
$$

with precision $\equiv \mathrm{TP} /(\mathrm{TP}+\mathrm{FP})$ and recall $\equiv \mathrm{TP} /(\mathrm{TP}+\mathrm{FN})$, where FP and FN denote the number of false positives (incorrectly selected fixed effects) and false negatives (relevant fixed effects that have not been selected), respectively. This index takes values between $0 \%$ and $100 \%$, with higher values to be preferred.

The median $F_{1}$-score was $63.5 \%$ (1st quartile: $43.5 \%$; 3rd quartile: $94.2 \%$ ) for the Gaussian prior and $95.24 \%$ (1st quartile: $80 \%$; 3rd quartile: $100 \%$ ) for the Laplace prior. Both the Horseshoe and Negative-Exponential-Gamma priors exhibited a $F_{1}$-score equal to $100 \%$ for all the 50 data replications of the simulation study, meaning that perfect selection ( $\mathrm{TP}=10$ and $\mathrm{TN}=40$ ) was always achieved. From this simplified simulated sparse data scenario it is apparent that global-local priors effectively provide better variable selection performances than the Gaussian prior. The lower performances of the Gaussian and Laplace priors are due to the SAVS procedure being applied to optimal approximate $\mathfrak{q}^{*}$-densities that have not been properly shrunk towards zero and so irrelevant fixed effects tend to be selected. Nonetheless, all the four different priors gave $\mathrm{FP}=0$, meaning that they did not incorrectly select irrelevant fixed effects.

### 6.3. Speed and memory saving assessment

Streamlined variational inference has been conceived to obtain efficient implementations of variational algorithms. Hence, the assessment of speed and memory savings is another important aspect. We assessed speed and memory performances of the streamlined variational Algorithm 3 and its naïve counterpart, whose updates are described in Section 5.1. Four different group numbers and three different lengths for the vector of fixed effects to be selected were considered, namely $m \in\{10,50,100,200\}$ and $p_{\mathrm{S}} \in\{25,100,200\}$, aiming to explore scalability of the streamlined methodology to high dimensions. For each combination of $m$ and $p_{\mathrm{S}}$, we simulated 10 data replications from model (17) with a random intercept and one slope for a single continuous predictor, choosing the sub-group dimensions $n_{i}$ uniformly on the discrete set $\{10, \ldots, 20\}$ and the sub-group specific unit dimensions $o_{i j}$ uniformly on the discrete set $\{20, \ldots, 30\}$. This setting allows to test models with heterogeneous dimensions but the same number of groups $m$. We considered a Horseshoe prior for the fixed effects subject to selection, and all the other model dimensions and hyperparameters specifications were the same as those described before.

Table 2
Average (standard deviation of) elapsed computing times in seconds and average (standard deviation of) total size of required data inputs in megabytes for fitting model (17) with three-level random effects specification and $p_{S}$ fixed effects subject to penalization via Horseshoe prior. Results are shown for different group sizes $m$ and different values for $p_{S}$.

| Total runtime of the algorithm (seconds) |  |  |  |  | Total size of required data inputs (Mbytes) |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $m \quad p_{\mathrm{S}}$ | Stream. MFVB | Naïve MFVB | $\frac{\text { Naive } M F V B}{\text { Str. } M F V B}$ | MCMC | Stream. MFVB | Nä̈ve MFVB | $\frac{\text { Naive } M F V B}{\text { Str. } M F V B}$ |
| $10 \quad 25$ | $0.70{ }_{(0.06)}$ | $3.86{ }_{(0.50)}$ | 5.54 | $13.03_{(0.96)}$ | $1.39_{(0.07)}$ | $11.06_{(1.00)}$ | 7.97 |
| 100 | $3.90_{(0.23)}$ | $6.51{ }_{(0.82)}$ | 1.67 | $46.99_{(1.47)}$ | $3.68{ }_{(0.24)}$ | $12.73_{(1.34)}$ | 3.46 |
| 200 | $9.71_{(1.22)}$ | $12.09_{(2.04)}$ | 1.24 | $176.22_{(9.37)}$ | $6.60{ }_{(0.52)}$ | $15.09_{(1.89)}$ | 2.29 |
| $50 \quad 25$ | $3.29_{(0.11)}$ | $365.07_{(37.10)}$ | 111.01 | $63.28_{(3.76)}$ | $6.69{ }_{(0.20)}$ | $240.40_{(13.92)}$ | 35.91 |
| 100 | $19.66_{(0.87)}$ | $415.29_{(51.66)}$ | 21.12 | $138.42_{(5.72)}$ | $18.37{ }_{(0.81)}$ | $254.04_{(20.84)}$ | 13.83 |
| 200 | $49.25_{(1.47)}$ | $501.33_{(39.76)}$ | 10.18 | $305.25_{(6.51)}$ | $34.00_{(1.04)}$ | $269.34_{(14.90)}$ | 7.92 |
| 10025 | $6.77{ }_{(0.32)}$ | $2877.33_{(177.77)}$ | 424.72 | $151.05_{(10.48)}$ | $13.56_{(0.34)}$ | $967.34_{(42.15)}$ | 71.33 |
| 100 | $40.42_{(1.66)}$ | $3172.66_{(116.83)}$ | 78.50 | $266.71_{(10.06)}$ | $36.77_{(1.29)}$ | $1010.25_{(26.94)}$ | 27.47 |
| 200 | $97.81_{(2.07)}$ | $3403.55_{(204.38)}$ | 34.80 | $494.15_{(9.34)}$ | $67.53_{(1.31)}$ | $1013.90_{(48.67)}$ | 15.01 |
| 20025 | $13.30_{(0.30)}$ | $>5$ hours | > 1355 | $328.87_{(20.04)}$ | $26.90_{(0.56)}$ | $3817.06_{(113.85)}$ | 141.88 |
| 100 | $79.97{ }_{(1.18)}$ | $>5$ hours | $>225$ | $578.88_{(13.45)}$ | $73.74_{(1.01)}$ | $3941.15{ }_{(97.90)}$ | 53.45 |
| 200 | $197.20_{(4.13)}$ | $>5$ hours | > 95 | $904.64_{(15.77)}$ | $135.84_{(2.73)}$ | $3991.66_{(102.48)}$ | 29.38 |

For each simulated dataset, we collected the computational timings and the total size of the input data required for performing both the streamlined and naïve algorithms. We also ran MCMC (warmup of length 5,000 followed by 25,000 iterations, to which we applied a thinning value of 5) and recorded its computational timings, although MFVB and MCMC do not admit a genuine comparison due to the fact that both depend upon different convergence requirements, as explained in Section 5.1 of [66]. Moreover, our efficient implementation of MCMC is based on independent resampling from the full conditional densities of $\boldsymbol{\beta}, \boldsymbol{u}_{i}^{\mathrm{L} 1}$ and $\boldsymbol{u}_{i j}^{\mathrm{L} 2}$, whereas MFVB uses a joint approximation $\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})$ for all these vectors. Nonetheless, MCMC timings are also reported to provide an intuition on the computational effort required for sampling from the posterior distribution when $m$ and $p_{\mathrm{S}}$ increase. The tabulated results are shown in Table 2. The "ratio" columns help understand the gain obtained employing the streamlined MFVB methodology over its naïve counterpart. For increasing $m$, streamlined MFVB reached convergence faster than naïve MFVB. Notice that in the biggest scenario under examination (last row of the table), the streamlined MFVB algorithm ran in less then 4 minutes on average, while naïve MFVB required more than 5 hours. Bigger scenarios are computationally demanding for the naïve implementation and may fail to run due to excessive storage demand, whilst we did not experience these issues with Algorithm 3. Similar comments apply to the huge saving of memory allocation for the required input data provided by the streamlined MFVB implementation.

We conclude the discussion by noticing that our methodology takes the approximate joint posterior dependence between the fixed effects and random effects parameters into account through the multivariate Gaussian approximating density $\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})$. This choice allows to better capture the a posteriori covariance structure between $\boldsymbol{\beta}$ and $\boldsymbol{u}$ and ensures better results in terms of approximation accuracy. Nevertheless, alternative and less restrictive factorizations can
be considered, especially if the quality of the approximation can be sacrificed and faster algorithms are desired. For instance, the additional factorization $\mathfrak{q}(\boldsymbol{\beta}, \boldsymbol{u})=\mathfrak{q}\left(\boldsymbol{\beta}^{\mathrm{R}}, \boldsymbol{u}\right) \mathfrak{q}\left(\boldsymbol{\beta}^{\mathrm{A}}, \boldsymbol{\beta}^{\mathrm{S}}\right)$ could remarkably speed up the computations if $p_{\mathrm{R}}$ is small, regardless of the size of $p_{\mathrm{A}}$ and $p_{\mathrm{S}}$. For implementing this and any other additional independence constraints, mean field restrictions such as (19) and the algorithms proposed in Section 5.2 need to be modified accordingly.

## 7. Application to data from a perinatal study

We present an application of the methodology and algorithms proposed in Section 5 to the National Collaborative Perinatal Project data ([42]), a multisite prospective cohort study which took place in the United States of America between 1959 and 1974. This study was designed to identify the effects of complications during pregnancy or the perinatal period on birth and child outcomes. The data are publicly available from the U.S. National Archives with identifier 606622. Many online resources already employed this dataset or subset of it for several analysis and over the years the dataset has become a high-quality reference for biomedical and behavioral research in many areas such as obstetrics, perinatology, pediatrics, and developmental psychology.

The same data were examined in [60] and we expressly account for the same model specification to experiment a suitable fixed effects selection for a moderately large set of regressors that have been excluded from their analysis and that may have a relevant impact onto explaining the response variable. A full-blown analysis goes beyond the scope of this paper and we focused on predicting the height-for-age z-score for 37,257 infants followed longitudinally over their first year of life, following indications from [78]. The height-for-age z-score is a standardized measure of the World Health Organization for the height of children after accounting for age; see [64] for insights on how notable discrepancies from this index standard reference values constitutes an alarm signal for malnutrition symptoms.

We performed a Bayesian analysis that accounts for the heterogeneity of the evolution of such index across infants. Our analysis was performed through a two-level random effects linear model having a random intercept, and linear and quadratic slopes $(q=3)$ to account for the quadratic evolution of that score over time. All the fixed effects regressors, excluding the intercept, age of the infant (in days since birth) and its square, were subject to selection; these include characteristics of the infant at birth (e.g. weight, length, head circumference, sex and Apgar scores), and characteristics of the mother, father and family. We also accounted for possible interactions between some infant characteristics and sex, for a total of 38 candidate predictors subject to selection.

The model we fitted respects the general specification (17) and can be expressed for the generic $i$ th infant as follows:

$$
\boldsymbol{y}_{i} \mid \boldsymbol{\beta}, \boldsymbol{u}_{i}, \sigma^{2} \stackrel{\text { ind. }}{\sim} \mathrm{N}\left(\boldsymbol{X}_{i}^{\mathrm{R}} \boldsymbol{\beta}^{\mathrm{R}}+\boldsymbol{X}_{i}^{\mathrm{S}} \boldsymbol{\beta}^{\mathrm{S}}+\boldsymbol{Z}_{i} \boldsymbol{u}_{i}, \sigma^{2} \boldsymbol{I}\right)
$$

$$
\begin{align*}
& \boldsymbol{u}_{i} \mid \Sigma \stackrel{\text { ind. }}{\sim} \mathrm{N}_{3}(\mathbf{0}, \boldsymbol{\Sigma}), \quad 1 \leq i \leq 37257, \\
& {\left.\left[\begin{array}{c}
\boldsymbol{\beta}^{\mathrm{R}} \\
\boldsymbol{\beta}^{\mathrm{S}}
\end{array}\right] \right\rvert\, \boldsymbol{\zeta}, \tau^{2} \sim \mathrm{~N}\left(\left[\begin{array}{c}
\mathbf{0}_{3} \\
\mathbf{0}_{38}
\end{array}\right],\left[\begin{array}{cc}
10^{10} \boldsymbol{I}_{3} & \boldsymbol{O} \\
\boldsymbol{O} & \tau^{2} \operatorname{diag}(\boldsymbol{\zeta})^{-1}
\end{array}\right]\right),} \\
& \sigma^{2} \mid a_{\sigma^{2}} \sim \text { Inverse- } \chi^{2}\left(1,1 / a_{\sigma^{2}}\right), \quad a_{\sigma^{2}} \sim \text { Inverse- } \chi^{2}\left(1,10^{-10}\right), \\
& \boldsymbol{\Sigma} \mid \boldsymbol{A}_{\boldsymbol{\Sigma}} \sim \text { Inverse-G-Wishart }\left(G_{\text {full }}, 6, \boldsymbol{A}_{\boldsymbol{\Sigma}}^{-1}\right), \\
& \boldsymbol{A}_{\boldsymbol{\Sigma}} \sim \operatorname{Inverse-G-Wishart}\left(G_{\text {diag }}, 1,2 \times 10^{-10} \boldsymbol{I}_{3}\right) \text {, } \\
& \zeta_{h} \left\lvert\, a_{\zeta_{h}} \stackrel{\text { ind. }}{\sim} \begin{cases}\text { Inverse- } \chi^{2}(2,1) & \text { for a Laplace prior } \\
\operatorname{Gamma}\left(1 / 2, a_{\zeta_{h}}\right) & \text { for a Horseshoe prior } \\
\text { Inverse- } \chi^{2}\left(2,2 a_{\zeta_{h}}\right) & \text { for a NEG prior, }\end{cases} \right.  \tag{24}\\
& a_{\zeta_{h}} \stackrel{\text { ind. }}{\sim}\left\{\begin{array}{cl}
- & \text { for a Laplace prior } \\
\operatorname{Gamma}(1 / 2,1) & \text { for a Horseshoe prior } \\
\operatorname{Gamma}(\lambda, 1) & \text { for a NEG prior, }
\end{array}\right. \\
& \text { for } 1 \leq h \leq 38 \text {, } \\
& \tau^{2} \mid a_{\tau^{2}} \sim \text { Inverse- } \chi^{2}\left(1,1 / a_{\tau^{2}}\right), \quad a_{\tau^{2}} \sim \text { Inverse- } \chi^{2}\left(1,10^{-10}\right) .
\end{align*}
$$

For the $i$ th infant, $o_{i}$ time-point measurements were recorded, ranging from one to four in number. The $\boldsymbol{X}_{i}^{\mathrm{R}}$ matrix has size $n_{i} \times 3$ with the first column being a vector of ones, the second one consists of the time-point measurements for the $i$ th infant and the third column containing the square of the elements of the second one. We set $\boldsymbol{X}_{i}^{\mathrm{A}}=\boldsymbol{O}$, while the $\boldsymbol{X}_{i}^{\mathrm{S}}$ matrix of size $n_{i} \times 38$ consisted of all the considered predictors subject to selection. Moreover, $\boldsymbol{Z}_{i}=\boldsymbol{X}_{i}^{\mathrm{R}}$ by definition. Uninformative priors were placed over all the model parameters. We fitted the model using the three global-local priors treated in this work, and the Gaussian prior for $\beta^{\mathrm{S}}$ treated in [60].

Streamlined MFVB and MCMC were used for model fitting. The former was performed running Algorithm 2 and stopping it after 200 iterations, while the latter was performed running 25,000 iterations to which a thinning factor of 5 was applied after discarding 5,000 burnin iterations. The whole input data required approximately 200 megabytes of memory storage, while a naïve MFVB procedure would necessitate several gigabytes of memory to entirely store the $\boldsymbol{Z}$ matrix, that is composed by $11,693,705,562$ cells of which the $99.997 \%$ are zeros. All the covariates excluding the binary ones were standardized and the estimates were rescaled back to the original scale before presenting the results. The streamlined MFVB algorithms took 2 to 3 minutes to run for each prior specification, while the associated MCMC samplers took more than 35 minutes.

We omit the presentation of model interpretation, goodness-of-fit analysis, accuracy of the approximations, convergence of the MCMC chains and visualization of the fitted height-for-age z-score trajectories over time. Instead, in Figure 4 we present $90 \%$ high posterior density credible intervals for all the fixed effects subject to selection. The thicker lines represent the intervals ob-

## E. Degani et al.



Fig 4. The $90 \%$ high posterior density credible intervals for the fixed effects subject to selection in the real data application. The four different priors for $\boldsymbol{\beta}^{S}$ are represented with different colors. For each fixed effect, the thicker lines correspond to the intervals obtained from the MCMC approximate marginal posterior densities, while the thinner lines represent those obtained from the streamlined MFVB approximating densities.
tained from the MCMC posterior samples through the emp.hpd function of the TeachingDemos package [76], while the thinner lines represent those calculated from the determined optimal MFVB approximate posteriors. The two lines are superimposed to facilitate immediate comparison, for each different prior specification.

Overall, MFVB provided very accurate high posterior density credible intervals when compared to MCMC. The MCMC chain associated to the birth length fixed effect showed some convergence problems, probably due to its moderate correlation with the response variable and this reduced the overlap with MFVB. Importantly, the Laplace global-local prior produced results very sim-
ilar to those provided by the Gaussian prior specification. On the other hand, both the Horseshoe and Normal Exponential Gamma priors seem to have effectively shrunk most of the fixed effects towards zero, especially those associated to dummy variables.

The SAVS procedure applied to the approximate posterior densities classifies the following fixed effects as relevant: indicator that the infant is male, infant's year of birth, birth length, birth head circumference, the Apgar score assessed 5 minutes after the infant's birth and the mother height. Some effects such as the ethnicity of the mother, the place where she conceived, her marital status and informations on previous pregnancies were not indicated as relevant by the SAVS procedure, albeit some of their associated high posterior density credible intervals are far from zero. Notice also that, for some fixed effects, global-local priors drag the intervals towards the origin, with different intensities probably depending on the variability and correlation of the covariates. More ad-hoc analyses are firmly suggested, although these go beyond the scope of the current work.

## 8. Conclusions

In this work, we developed streamlined mean-field variational Bayes procedures for Gaussian response linear mixed models having nested random effects structures and admitting fixed effects prior specifications alternative to the classical Gaussian one. The priors we considered are amenable to automated and hyperparameter-free fixed effects selection procedures. Simulated and real data examples showed how streamlined variational inference can provide impressive benefits in terms of computational time and memory saving when compared to inefficient implementations of variational approximations. Albeit the marginal posterior densities of fixed effects subject to selection are approximated by bellshaped curves, our studies showed high performances of the automated selection procedure. It is also worth mentioning that the more restrictive mean-field approximations discussed at the end of Section 6 can sensibly improve the benefit of streamlined variational inference for large $p$ scenarios and, in general, when speed is more important than approximation accuracy.

Numerous ramifications of the proposed methodology can be envisaged. These include the treatment of models with unit-specific errors, heteroskedastic covariance structures for groups and sub-groups, higher levels of nesting or crossed random effects and the extension to generalized linear mixed models. [65, 87, $61,52,54]$ provide variational inference algorithms for models with a variety of response distributions such as Bernoulli with probit or logit link, Poisson, Negative Binomial, $t$, Asymmetric Laplace, Skew Normal, Skew $t$ and Finite Normal Mixtures. The variational algorithms presented in these references have been derived using data-augmentation representations of the response distributions. The same strategy can be adopted to derive MFVB algorithms for non-Gaussian response mixed models and allow for the implementation of streamlined variational inference through a relatively straightforward adaptation of the algorithms presented in this work. [51] also provides explicit streamlined variational
inference algorithms for two-level mixed models with binary-logistic and Poisson responses, and generic Gaussian priors for the fixed effects coefficients.

Regarding the selection of fixed effects, alternative shrinkage priors belonging to the global-local family can be accounted for with minor modifications to the proposed algorithms. Following the prescriptions of [59], it is possible to study global-local prior specifications that are not based on auxiliary variables and investigate whether this could lead to better approximation accuracies, although at the cost of more computationally intensive variational updates. We also mention the possibility of admitting spike-and-slab prior formulations following, for example, a variational inference approach similar to that of [16]. The drawback is that spike-and-slab priors require more involved algebra for implementing streamlined variational inference. Using the Bayesian adaptive lasso (BaLasso) prior proposed by [45] it is also possible to extend our individual fixed-effects selection approach to account for (ordered) group selection by imposing different shrinkage levels to different coefficients. A variational Bayes approach for generalized linear models with priors of this type has been explored by [81], where the variational parameter updates of their VBGLMM algorithm could be streamlined using the framework studied in our work.

One last possible direction to explore concerns streamlined variational inference for models with priors for selecting random effects. Consider for simplicity the two-level random effects case in which $\boldsymbol{u}_{i} \stackrel{\text { ind. }}{\sim} \mathrm{N}(\mathbf{0}, \boldsymbol{\Sigma})$ for $1 \leq i \leq m$. If the covariance matrix $\boldsymbol{\Sigma}$ is supposed to be diagonal, i.e. $\operatorname{Cov}\left(u_{i h}, u_{i k}\right)=0$ for each $1 \leq h, k \leq q$ with $h \neq k$ and $1 \leq i \leq m$, and a global-local prior distribution is imposed to each random effect, then explicit streamlined MFVB updates can be obtained with straightforward manipulation of the results presented in the supplementary material and references therein. This applies to three-level models in a similar way. However, assuming that $\boldsymbol{\Sigma}$ is diagonal may sometimes be quite restrictive and the highest computational advantage of streamlined MFVB over its naïve counterpart is achieved when the random effects covariance matrix is non-diagonal. Furthermore, it is typically easier to provide an interpretation to the selection of fixed effects than explaining why certain random effects are relevant only for a subset of clusters or sub-clusters. For these reasons the present work provides a first extension of the streamlined MFVB approach of [60] focusing on fixed-effects selection. Another relevant extension is the derivation of (streamlined) MFVB algorithms with random effects selection procedures admitting non-diagonal structures for $\boldsymbol{\Sigma}$. The covariance matrix decomposition approach of [20] and related methods discussed in Section 1 can be exploited for this extension.

## Supplementary Material

## Supplementary Material for: Sparse Linear Mixed Model Selection via Streamlined Variational Bayes <br> (doi: 10.1214/22-EJS2063SUPP; .pdf).

S.1: Distributions and Associated Useful Results. We provide details on the probability density functions associated to the probability distributions are used throughout the manuscript, together with additional useful results.
S.2: Multilevel Sparse Matrix Problem Algorithms. Algorithms 2 and 3 described in Section 5 rely upon two matrix algebraic routines for efficiently solving the two-level and three-level versions of the multilevel sparse matrix problems defined in [62]. We briefly describe two-level and three-level sparse matrix structures, and give explicit definitions of the routines for efficiently solving the associated linear system problems, following Appendix A of [60].
S.3: Derivations. We derive explicit updates for the parameters of the optimal density functions in (16) using arguments similar to those provided in [59] with some adjustments. Such updates are then combined with the derivations in Appendix B of [60] for deriving the streamlined MFVB algorithms for our two- and three- level linear mixed-effects models.

## References

[1] Armagan, A. and Dunson, D. B. (2011). Sparse variational analysis of linear mixed models for large data sets. Statistics \& Probability Letters 81 1056-1062. MR2803744
[2] Armagan, A., Dunson, D. B. and Lee, J. (2013). Generalized double Pareto shrinkage. Statistica Sinica 23 119-143. MR3076161
[3] Baltagi, B. H. (2021). Econometric Analysis of Panel Data, Sixth ed. Springer Cham. MR4393585
[4] Barbieri, M. M. and Berger, J. O. (2004). Optimal predictive model selection. The Annals of Statistics 32 870-897. MR2065192
[5] Bhadra, A., Datta, J., Polson, N. G. and Willard, B. (2017). The horseshoe+ estimator of ultra-sparse signals. Bayesian Analysis 12 1105-1131. MR3724980
[6] Bhadra, A., Datta, J., Polson, N. G. and Willard, B. (2019). Lasso meets horseshoe: a survey. Statistical Science 34 405-427. MR4017521
[7] Bhattacharya, A., Chakraborty, A. and Mallick, B. K. (2016). Fast sampling with Gaussian scale mixture priors in high-dimensional regression. Biometrika 103 985-991. MR3620452
[8] Bhattacharya, A., Pati, D., Pillai, N. S. and Dunson, D. B. (2015). Dirichlet-Laplace priors for optimal shrinkage. Journal of the American Statistical Association 110 1479-1490. MR3449048
[9] Bishop, C. M. (2006). Pattern Recognition and Machine Learning. Information Science and Statistics. Springer, New York. MR2247587
[10] Blei, D. M., Kucukelbir, A. and McAuliffe, J. D. (2017). Variational inference: a review for statisticians. Journal of the American Statistical Association 112 859-877. MR3671776
[11] Bogdan, M. g., Chakrabarti, A., Frommlet, F. and Ghosh, J. K. (2011). Asymptotic Bayes-optimality under sparsity of some multiple testing procedures. The Annals of Statistics 39 1551-1579. MR2850212
[12] Bondell, H. D. and Reich, B. J. (2012). Consistent high-dimensional Bayesian variable selection via penalized credible regions. Journal of the American Statistical Association 107 1610-1624. MR3036420
[13] Boyd, S. and Vandenberghe, L. (2004). Convex Optimization. Cambridge University Press, Cambridge. MR2061575
[14] Brown, H. and Prescott, R. (2015). Applied Mixed Models in Medicine, Third ed. John Wiley \& Sons.
[15] Bürkner, P.-C. (2018). Advanced Bayesian Multilevel Modeling with the R Package brms. The R Journal 10 395-411.
[16] Carbonetto, P. and Stephens, M. (2012). Scalable variational inference for Bayesian variable selection in regression, and its accuracy in genetic association studies. Bayesian Analysis 7 73-107. MR2896713
[17] Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P. and Riddell, A. (2017). Stan: A probabilistic programming language. Journal of Statistical Software 76 1-32.
[18] Carvalho, C. M., Polson, N. G. and Scott, J. G. (2009). Handling Sparsity via the Horseshoe. In Proceedings of the Twelth International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research 5 73-80. PMLR, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA.
[19] Carvalho, C. M., Polson, N. G. and Scott, J. G. (2010). The horseshoe estimator for sparse signals. Biometrika 97 465-480. MR2650751
[20] Chen, Z. and Dunson, D. B. (2003). Random effects selection in linear mixed models. Biometrics. Journal of the International Biometric Society 59 762-769. MR2025100
[21] Degani, E., Maestrini, L., ToczydŁowska, D. and Wand, M. P. (2022). Supplement to "Sparse linear mixed model selection via streamlined variational Bayes". DOI: 10.1214/22-EJS2063SUPP.
[22] Eddelbuettel, D. and Sanderson, C. (2014). RcppArmadillo: Accelerating R with high-performance $\mathrm{C}++$ linear algebra. Computational Statistics and Data Analysis 71 1054-1063. MR3132026
[23] Efron, B. (2008). Microarrays, empirical Bayes and the two-groups model. Statistical Science 23 1-22. MR2431866
[24] Faes, C., Ormerod, J. T. and Wand, M. P. (2011). Variational Bayesian inference for parametric and nonparametric regression with missing data. Journal of the American Statistical Association 106 959-971. MR2894756
[25] Fan, Y. and Li, R. (2012). Variable selection in linear mixed effects models. The Annals of Statistics 40 2043-2068. MR3059076
[26] Fitzmaurice, G., Davidian, M., Verbeke, G. and Molenberghs, G. (2008). Longitudinal Data Analysis. CRC Press. MR1500110
[27] Frank, I. E. and Friedman, J. H. (1993). A Statistical View of Some Chemometrics Regression Tools. Technometrics 35 109-135.
[28] Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). Bayesian Analysis 1 515-533. MR2221284
[29] George, E. I. and Mcculloch, R. E. (1997). Approaches for Bayesian variable selection. Statistica Sinica 7 339-374.
[30] Goldstein, H. (2010). Multilevel Statistical Models, Fourth ed. John Wiley \& Sons Inc.
[31] Griffin, J. E. and Brown, P. J. (2010). Inference with normal-gamma prior distributions in regression problems. Bayesian Analysis 5 171-188. MR2596440
[32] Griffin, J. E. and Brown, P. J. (2011). Bayesian hyper-lassos with non-convex penalization. Australian $\mathcal{E}^{2}$ New Zealand Journal of Statistics 53 423-442. MR2910027
[33] Groll, A. and Tutz, G. (2014). Variable selection for generalized linear mixed models by $L_{1}$-penalized estimation. Statistics and Computing 24 137-154. MR3165544
[34] Hahn, P. R. and Carvalho, C. M. (2015). Decoupling shrinkage and selection in Bayesian linear models: a posterior summary perspective. Journal of the American Statistical Association 110 435-448. MR3338514
[35] Hoerl, A. E. and Kennard, R. W. (1970). Ridge Regression: Biased Estimation for Nonorthogonal Problems. Technometrics 12 55-67.
[36] Huang, A. and Wand, M. P. (2013). Simple marginally noninformative prior distributions for covariance matrices. Bayesian Analysis 8 439-451. MR3066948
[37] Hughes, D. M., García-Fiñana, M. and Wand, M. P. (2021). Fast approximate inference for multivariate longitudinal data. Biostatistics. (Year and page numbers pending.).
[38] Hui, F. K. C., Müller, S. and Welsh, A. H. (2017). Joint selection in mixed models using regularized PQL. Journal of the American Statistical Association 112 1323-1333. MR3735380
[39] Ishwaran, H. and Rao, J. S. (2005). Spike and slab variable selection: frequentist and Bayesian strategies. The Annals of Statistics 33 730-773. MR2163158
[40] Johnstone, I. M. and Silverman, B. W. (2005). Empirical Bayes selection of wavelet thresholds. The Annals of Statistics 33 1700-1752. MR2166560
[41] Kinney, S. K. and Dunson, D. B. (2007). Fixed and random effects selection in linear and logistic models. Biometrics. Journal of the International Biometric Society 63 690-698. MR2395705
[42] Klebanoff, M. A. (2009). The Collaborative Perinatal Project: a 50 -year retrospective. Paediatric and Perinatal Epidemiology 23 2-8.
[43] Korte, A., Vilhjálmsson, B. J., Segura, V., Platt, A., Long, Q. and Nordborg, M. (2012). A mixed-model approach for genome-wide association studies of correlated traits in structured populations. Nature Genetics 44 1066-1071.
[44] Lee, C. Y. Y. and Wand, M. P. (2016). Streamlined mean field vari-
ational Bayes for longitudinal and multilevel data analysis. Biometrical Journal 58 868-895. MR3527420
[45] Leng, C., Tran, M.-N. and Nott, D. (2014). Bayesian adaptive Lasso. Annals of the Institute of Statistical Mathematics 66 221-244. MR3171404
[46] Li, H. and Pati, D. (2017). Variable selection using shrinkage priors. Computational Statistics \& Data Analysis 107 107-119. MR3575062
[47] Li, J., Wang, Z., Li, R. and Wu, R. (2015). Bayesian group Lasso for nonparametric varying-coefficient models with application to functional genome-wide association studies. The Annals of Applied Statistics 9 640-664. MR3371329
[48] Li, Y., Wang, S., Song, P. X. K., Wang, N., Zhou, L. and Zhu, J. (2018). Doubly regularized estimation and selection in linear mixed-effects models for high-dimensional longitudinal data. Statistics and its Interface 11 721-737. MR3858527
[49] Lindner, C. C. and Rodger, C. A. (2008). Design Theory, second ed. Discrete Mathematics and its Applications. Chapman and Hall/CRC, Boca Raton, FL. MR2469212
[50] Luts, J., Broderick, T. and Wand, M. P. (2014). Real-time semiparametric regression. Journal of Computational and Graphical Statistics 23 589-615. MR3224647
[51] Maestrini, L. (2018). On Variational Approximations for Frequentist and Bayesian Inference, PhD thesis, Università degli Studi di Padova, Italy.
[52] Maestrini, L. and Wand, M. P. (2018). Variational message passing for skew $t$ regression. Stat 7 e196, 11. MR3905852
[53] Maestrini, L. and Wand, M. P. (2021). The Inverse G-Wishart distribution and variational message passing. Australian \& New Zealand Journal of Statistics 63 517-541. MR4374511
[54] McLean, M. W. and Wand, M. P. (2019). Variational message passing for elaborate response regression models. Bayesian Analysis 14 371-398. MR3934090
[55] Menictas, M., Credico, G. D. and Wand, M. P. (2022). Streamlined variational inference for linear mixed models with crossed random effects. Journal of Computational and Graphical Statistics. (volume and page numbers pending).
[56] Menictas, M., Nolan, T. H., Simpson, D. G. and Wand, M. P. (2021). Streamlined variational inference for higher level group-specific curve models. Statistical Modelling 21 479-519. MR4334020
[57] Minka, T. P., Winn, J. M., Guiver, J. P., Zaykov, Y., Fabian, D. and Bronskill, J. (2018). Infer.NET 0.3. Microsoft Research Cambridge. http://dotnet.github.io/infer.
[58] Mitchell, T. J. and Beauchamp, J. J. (1988). Bayesian variable selection in linear regression. Journal of the American Statistical Association 83 1023-1036. MR0997578
[59] Neville, S. E., Ormerod, J. T. and Wand, M. P. (2014). Mean field variational Bayes for continuous sparse signal shrinkage: pitfalls and remedies. Electronic Journal of Statistics 8 1113-1151. MR3263115
[60] Nolan, T. H., Menictas, M. and Wand, M. P. (2020). Streamlined computing for variational inference with higher level random effects. Journal of Machine Learning Research (JMLR) 21 Paper No. 157, 62. MR4209443
[61] Nolan, T. H. and Wand, M. P. (2017). Accurate logistic variational message passing: algebraic and numerical details. Stat 6 102-112. MR3633590
[62] Nolan, T. H. and Wand, M. P. (2020). Streamlined solutions to multilevel sparse matrix problems. The Australian $\mathcal{E}$ New Zealand Industrial and Applied Mathematics Journal 62 18-41. MR4130820
[63] O'Hara, R. B. and Sillanpää, M. J. (2009). A review of Bayesian variable selection methods: what, how and which. Bayesian Analysis 485-117. MR2486240
[64] World Health Organization (2006). WHO child growth standards: length/height-for-age, weight-for-age, weight-for-length, weight-for-height and body mass index-for-age: methods and development. Available at https://apps.who.int/iris/handle/10665/43413.
[65] Ormerod, J. T. and Wand, M. P. (2010). Explaining variational approximations. The American Statistician 64 140-153. MR2757005
[66] Ormerod, J. T., You, C. and Müller, S. (2017). A variational Bayes approach to variable selection. Electronic Journal of Statistics 11 3549-3594. MR3709863
[67] Park, T. and Casella, G. (2008). The Bayesian lasso. Journal of the American Statistical Association 103 681-686. MR2524001
[68] Pinheiro, J. C. and Bates, D. M. (2006). Mixed-Effects Models in S and S-Plus. Springer, New York.
[69] Polson, N. G. and Scott, J. G. (2011). Shrink globally, act locally: sparse Bayesian regularization and prediction. In Bayesian Statistics, 9 501-538. Oxford Univ. Press, Oxford. MR3204017
[70] Rao, J. N. K. and Molina, I. (2015). Small Area Estimation, second ed. Wiley Series in Survey Methodology. John Wiley \& Sons, Inc., Hoboken, NJ. MR3380626
[71] Ray, P. and Bhattacharya, A. (2018). Signal Adaptive Variable Selector for the Horseshoe Prior. arXiv preprint arXiv:1810.09004.
[72] Ruppert, D., Wand, M. P. and Carroll, R. J. (2003). Semiparametric Regression. Cambridge Series in Statistical and Probabilistic Mathematics 12. Cambridge University Press, Cambridge. MR1998720
[73] Sanderson, C. and Curtin, R. (2016). Armadillo: a template-based C++ library for linear algebra. Journal of Open Source Software 1.
[74] Schelldorfer, J., Bühlmann, P. and van de Geer, S. (2011). Estimation for high-dimensional linear mixed-effects models using $\ell_{1}$-penalization. Scandinavian Journal of Statistics. Theory and Applications 38 197-214. MR2829596
[75] Sikorska, K., Rivadeneira, F., Groenen, P. J., Hofman, A., Uitterlinden, A. G., Eilers, P. H. and Lesaffre, E. (2013). Fast linear mixed model computations for genome-wide association studies with longitudinal data. Statistics in Medicine 32 165-180. MR3017892
[76] Snow, G. (2020). TeachingDemos: Demonstrations for Teaching and Learning $R$ package version 2.12. https://CRAN.R-project.org/package= TeachingDemos.
[77] Tang, X., Ghosh, M., Xu, X. and Ghosh, P. (2018). Bayesian variable selection and estimation based on global-local shrinkage priors. Sankhya A. The Indian Journal of Statistics 80 215-246. MR3850065
[78] Taylor, P. M. (1980). The First Year of Life: The Collaborative Perinatal Project of the National Institute of Neurological and Communicative Disorders and Stroke. Journal of the American Medical Association 244 1503-1503.
[79] R Core Team (2020). R: A Language and Environment for Statistical Computing R Foundation for Statistical Computing, Vienna, Austria https://www.R-project.org/.
[80] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society. Series B (Statistical Methodology) 58 267-288. MR1379242
[81] Tung, D. T., Tran, M.-N. and Cuong, T. M. (2019). Bayesian adaptive lasso with variational Bayes for variable selection in high-dimensional generalized linear mixed models. Communications in Statistics. Simulation and Computation 48 530-543. MR3933861
[82] van Rijsbergen, C. J. (2004). The Geometry of Information Retrieval. Cambridge University Press, Cambridge. MR2136161
[83] Verbeke, G. and Molenberghs, G. (2000). Linear Mixed Models for Longitudinal Data. Springer Series in Statistics. Springer-Verlag, New York. MR1880596
[84] Vonesh, E. F. and Chinchilli, V. M. (1997). Linear and Nonlinear Models for the Analysis of Repeated Measurements. Statistics: Textbooks and Monographs 154. Marcel Dekker, Inc., New York. MR1633822
[85] Wand, M. P. (2020). KernSmooth: Functions for Kernel Smoothing Supporting Wand and Jones (1995) R package version 2.23-18. https://CRAN.R-project.org/package=KernSmooth.
[86] Wand, M. P. and Jones, M. C. (1995). Kernel Smoothing. Monographs on Statistics and Applied Probability 60. Chapman and Hall, Ltd., London. MR1319818
[87] Wand, M. P., Ormerod, J. T., Padoan, S. A. and Frührwirth, R. (2011). Mean field variational Bayes for elaborate distributions. Bayesian Analysis 6 847-900. MR2869967
[88] WANG, S. S. J. and Wand, M. P. (2011). Statistical computing and graphics using Infer.NET for statistical analyses. The American Statistician 65 115-126. MR2768622
[89] Yang, M. (2013). Bayesian nonparametric centered random effects models with variable selection. Biometrical Journal 55 217-230. MR3045842
[90] Yang, M., Wang, M. and Dong, G. (2020). Bayesian variable selection for mixed effects model with shrinkage prior. Computational Statistics 35 227-243. MR4066283
[91] Zhang, Y. and Bondell, H. D. (2018). Variable selection via penal-
ized credible regions with Dirichlet-Laplace global-local shrinkage priors. Bayesian Analysis 13 823-844. MR3807868
[92] Zhao, Y., Staudenmayer, J., Coull, B. A. and Wand, M. P. (2006). General design Bayesian generalized linear mixed models. Statistical Science 21 35-51. MR2275966
[93] Zou, H. (2006). The adaptive lasso and its oracle properties. Journal of the American Statistical Association 101 1418-1429. MR2279469
[94] Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society. Series B (Statistical Methodology) 67 301-320. MR2137327


[^0]:    *The research presented in this article was supported by the Australian Research Council Discovery Project DP180100597 and the Australian Research Council Center of Excellence for Mathematical and Statistical Frontiers (grant CE140100049).

