Supplementary materials for the manuscript "Fitting Double Hierarchical Models with the Integrated Nested Laplace Approximation"

Mabel Morales-Otero, Virgilio Gómez-Rubio and Vicente Núñez-Antón

1 Introduction

This supplementary material contains additional details and results not included in the main manuscript due to space constraints. The supplementary material is organized as follows. Section 2 includes additional results for the simulation studies, Section 3 contains additional results for the two examples presented in the main manuscript, and Section 4 includes a summary of the computational times required to fit the different models using MCMC and AMIS with INLA.

2 Simulation study

In this section we include additional results for the simulation study presented in the main manuscript. More specifically, we present graphical diagnostics to be able to assess the estimation of the posterior marginal for the model parameters that are fitted using importance sampling.

2.1 Poisson model with random effects with different precisions

Figure 1 presents the graphical diagnostics for the parameters in the models described in Section 5.1 in the main manuscript, which were fitted using importance sampling. Moreover, these parameters were used to model the log-precisions as $\log(\tau_i) = \gamma_0 + \gamma_1 z_i$. In this particular case, as can be seen in Figure 1, model fitting seems to be reasonably good.

Figure 1: Graphical diagnostics for the Poisson model used in the simulation study.

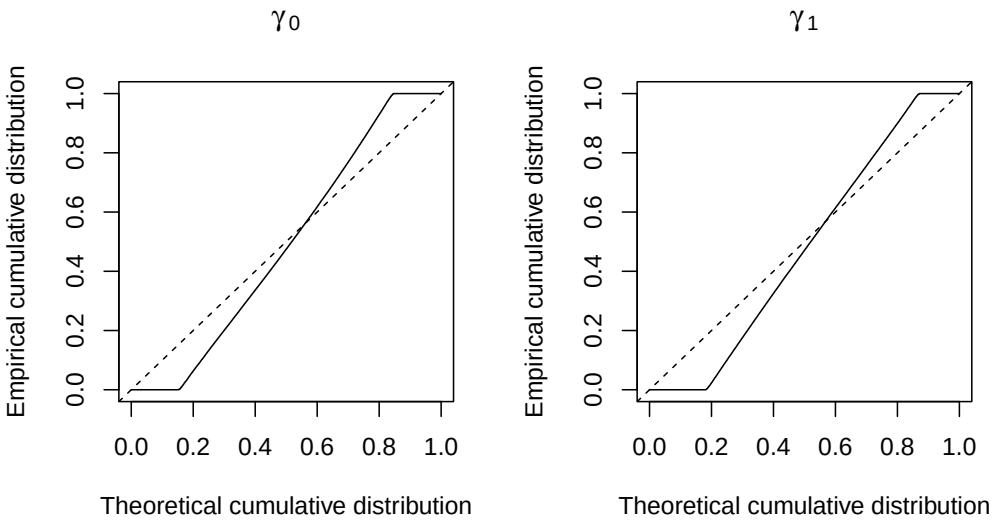

2.2 Negative binomial with different sample sizes

Figure 2 includes the graphical diagnostics for the parameters in the models described in Section 5.2 in the main manuscript, which were fitted using importance sampling. Moreover, these parameters were used to model the log-size parameter as $\log(k_i) = \gamma_0 + \gamma_1 z_i$. In this particular case, as can be seen in Figure 2, model fitting seems to be reasonably good for γ_0, whereas it is not as good for γ_1.

2.3 Gaussian model with different scale parameters

Figures 3, 4 and 5 present the graphical diagnostics for the parameters in the models described in Section 5.3 in the main manuscript, which were fitted using importance sampling. These models were included in the simulation study and used Gaussian data. With the exception of scenarios 1 and 3 (both with vague initial sampling distributions), model fitting seems to be very good.

Figure 2: Graphical diagnostics for the negative binomial model used in the simulation study.

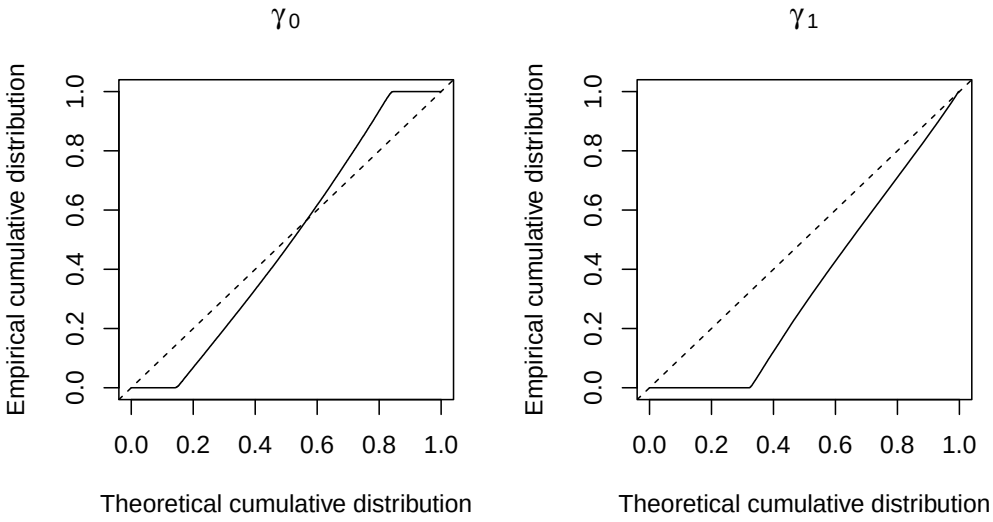

3 Examples

In this Section we include the graphical models, graphical diagnostics and some additional results for the examples included in Section 6 in the main manuscript.

3.1 Infant mortality in Colombia

Figure 6 presents the graphical diagnostics for the parameters in the Poisson and negative binomial models described in Section 6.1 in the main manuscript, which were fitted using importance sampling. These parameters were used to model the log-precisions as $\log(\tau_i) = \gamma_0 + \gamma_1 \text{IBN}_i$. In general, model fitting is very similar for both models and, in addition, it is reasonably good.

3.2 Sleep deprivation study

Figure 7 includes the subject-specific reaction times together with their corresponding linear regression lines. It also illustrates the fact that variability of the reaction times among subjects is not uniform, with some subjects having a broader range of values than others. This is the reason that motivates

Figure 3: Graphical diagnostics for the Gaussian models in the simulation study under scenarios 1 (left panel) and 2 (right panel).

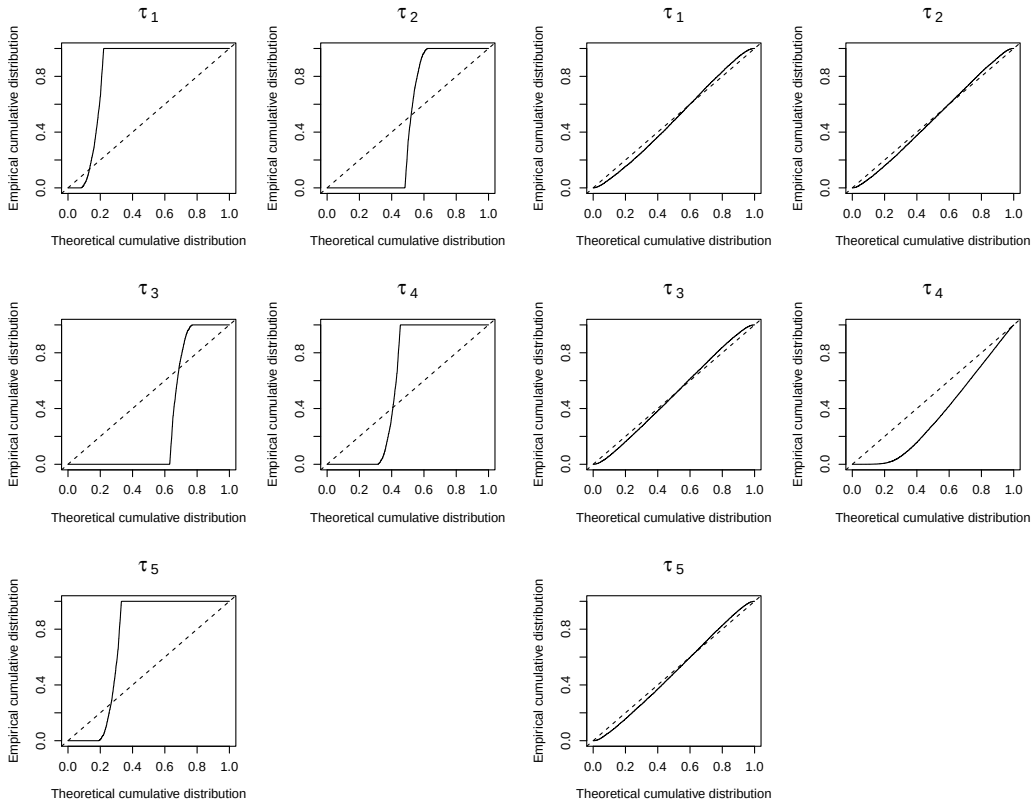

to propose the fitting of a model with random slopes per subject, where the within subject measurements precision is assumed to be different when using a DHGLM.

3.2.1 Model with shared slope

In addition to the model with random slopes presented in the main manuscript, we have also considered a model with a shared slope for all individuals (i.e., the slope for time is the same for all individuals):

Figure 4: Graphical diagnostics for the Gaussian models in the simulation study under scenarios 3 (left panel) and 4 (right panel).

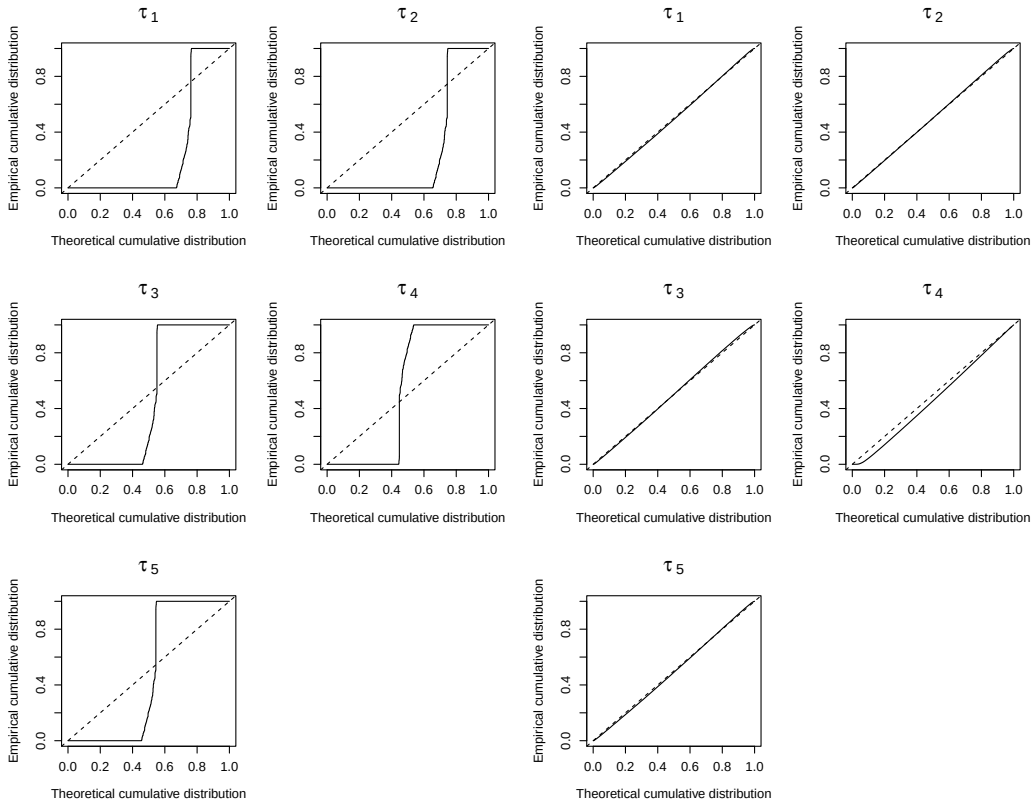

$$
\begin{aligned}
Y_{ij} &\sim N(\mu_{ij}, \tau_i); i = 1, \ldots, p;\ j = 1, \ldots, n_i \\
\mu_{ij} &= \beta_0 + \beta_1 \text{day}_{ij} \\
\log(\tau_i) &= \gamma + u_i \\
u_i &\sim N(0, \tau_u) \\
\tau_u &\sim Gamma(1, 0.00005) \\
\beta_0, \beta_1 &\sim N(0, 0.001) \\
\gamma &\sim N(0, 0.001)
\end{aligned} \tag{1}
$$

Table 1 includes estimations for the parameters in model (1) and Figure 8 presents their posterior densities estimations. As can be seen, in general, AMIS with INLA and MCMC produce similar estimates. In addition, for this specific case, AMIS with INLA results have an effective sample size of 71.906.

Figures 9 and 10 include the graphical diagnostic for the parameters in

Figure 5: Graphical diagnostics for the Gaussian models in the simulation study under scenarios 5 (left panel) and 6 (right panel).

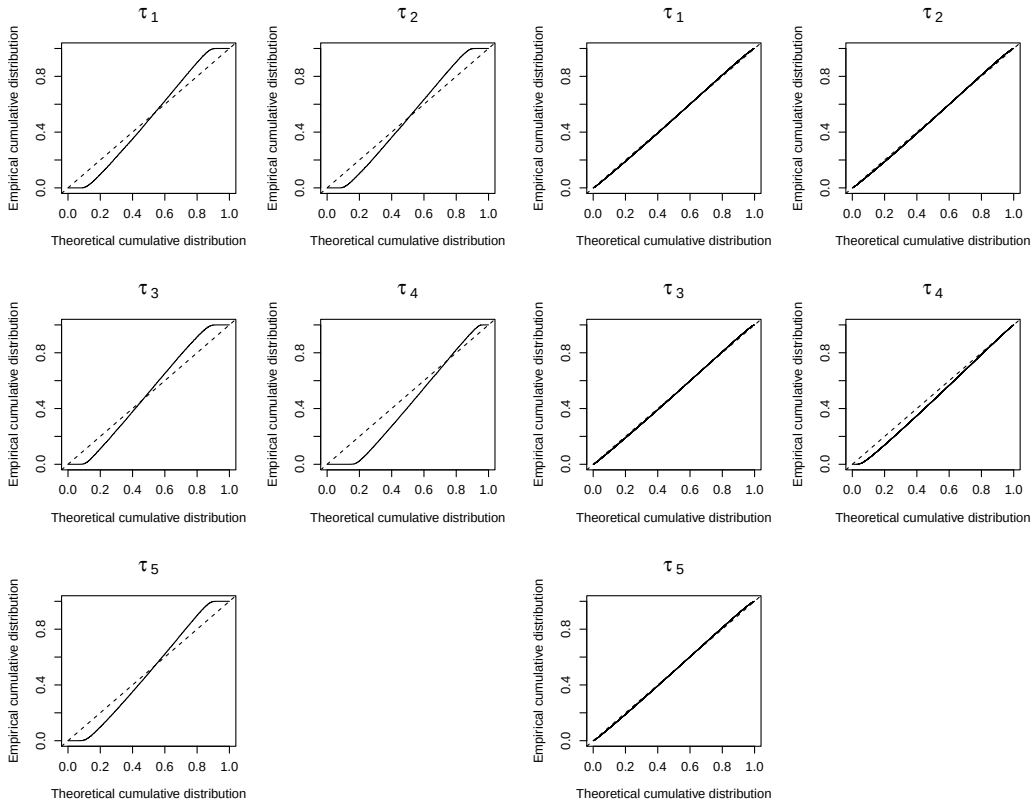

Table 1: Summary of the estimates for the Gaussian model fitted to the sleep study data.

	AMIS		MCMC	
Parameter	Mean	St. dev.	Mean	St. dev.
β_0	0.2576	0.0047	0.2572	0.0051
β_1	0.0109	0.0008	0.0105	0.0009
γ	6.6680	0.2496	6.5348	0.2889
τ_u	1.1676	0.4536	0.9794	0.4032

the model fitted using importance sampling. In general, model fitting is not as good as it was in the model with random slopes.

Figure 6: Graphical diagnostics for the Poisson model (top panel) and the negative binomial model (bottom panel) fitted to the infant mortality data in Colombia.

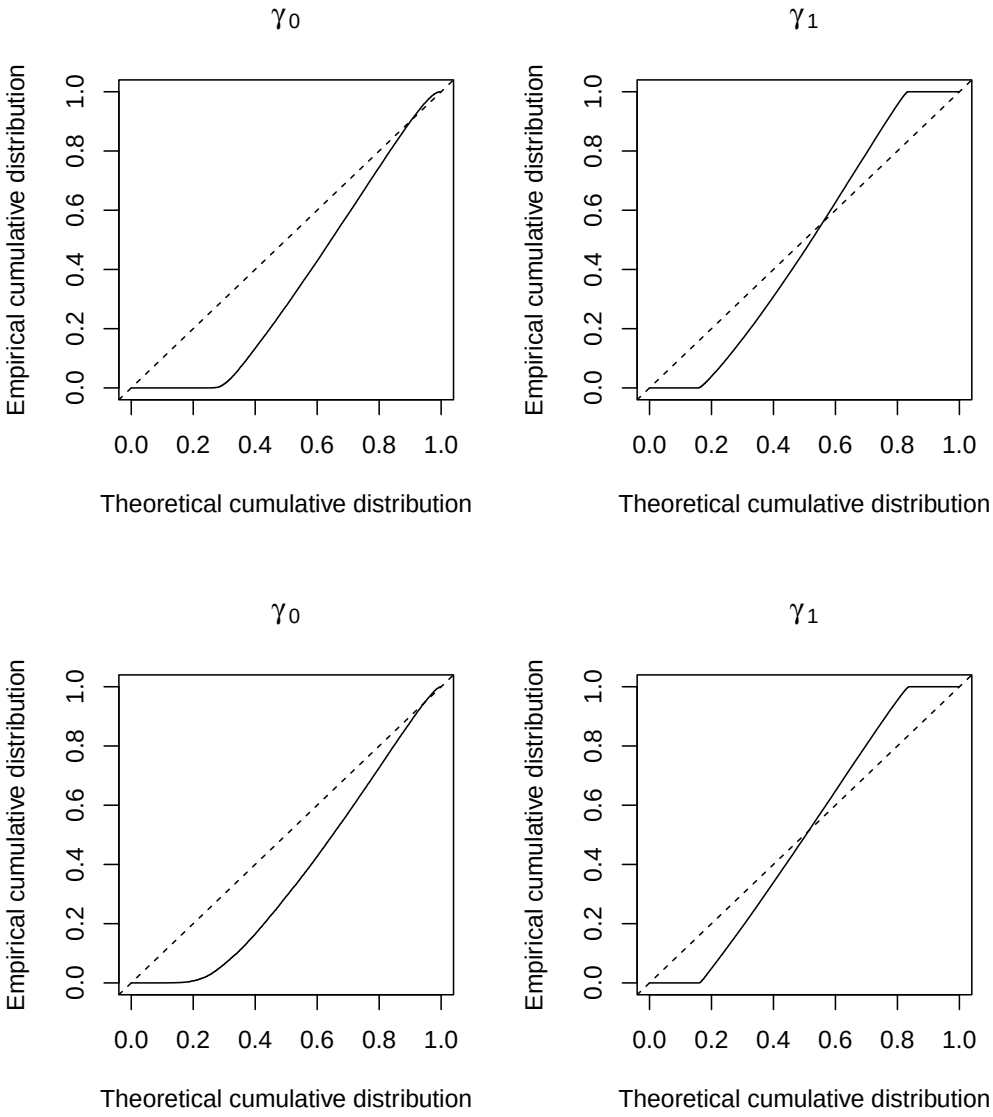

3.2.2 Model with random slopes

Figures 11 and 12 include the graphical diagnostic for the parameters in the model fitted using importance sampling. As can be seen, in general,

Figure 7: Effect of the number of days under sleep deprivation on the different subjects (based on code available in the `lme4` package).

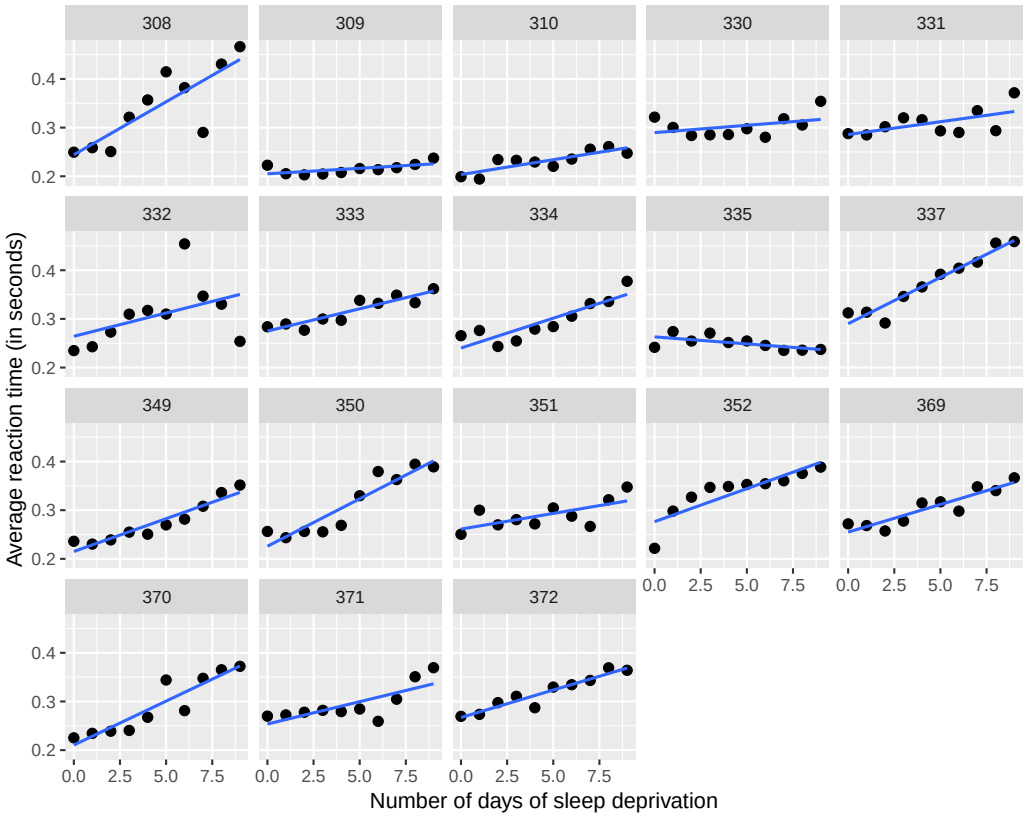

model fitting is reasonably good. It is relevant to mention that, in this case, $\tau_i,\ i = 1, \ldots, 18$ represent the within-individual precisions.

4 Computation times

In order to provide an overview of the computation times (in seconds) required to fit the different models discussed in the paper, Table 2 includes some approximate model fitting times. We would also like to mention that establishing a direct comparison between AMIS with INLA and MCMC is somehow difficult, mainly because of the different implementations used for these methods. More specifically, MCMC results have been obtained with JAGS (which is based on C++), whereas INLA with AMIS relies of calling the `inla()` function repeatedly from `R`, which introduces a considerable overhead.

Figure 8: Posterior marginal distributions of the estimated parameters obtained by fitting the Gaussian model with shared slope fitted to the sleep study data, using both the MCMC and AMIS-INLA approaches.

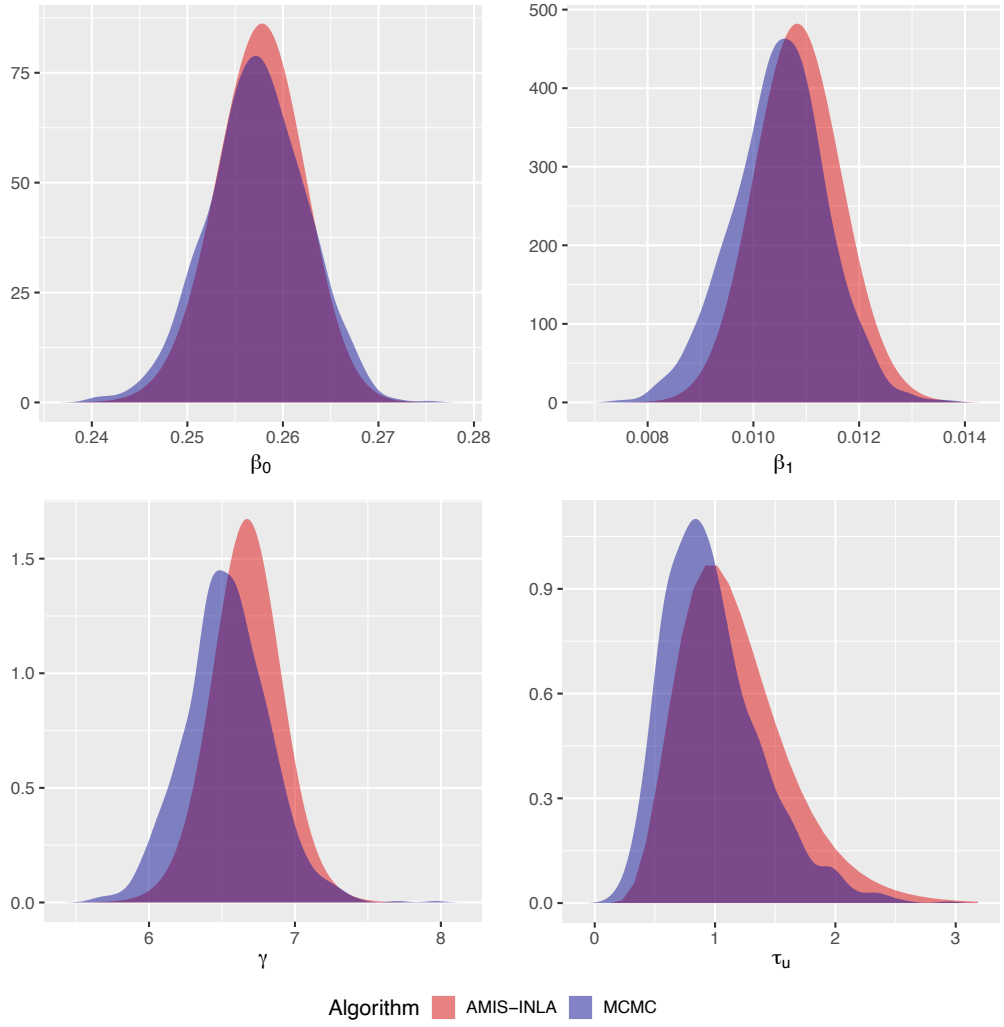

In Table 2, elapsed times refer to the time required since starting the AMIS with INLA or the MCMC sampling and until the end of the sampling process. Time per sample is the elapsed time divided by the total number of samples carried out. As stated in the main manuscript, simulations have been carried out on a Linux Ubuntu 18.2 cluster using 60 cores Intel(R) Xeon(R) CPU E5-2683 v4 @ 2.10GHz for INLA with AMIS, and using a single core for MCMC. Therefore, it is clear that INLA with MCMC times can still be reduced by increasing the number of cores used in the simulations.

Figure 9: Graphical diagnostics for the model with shared slope in the sleep deprivation study (1/2).

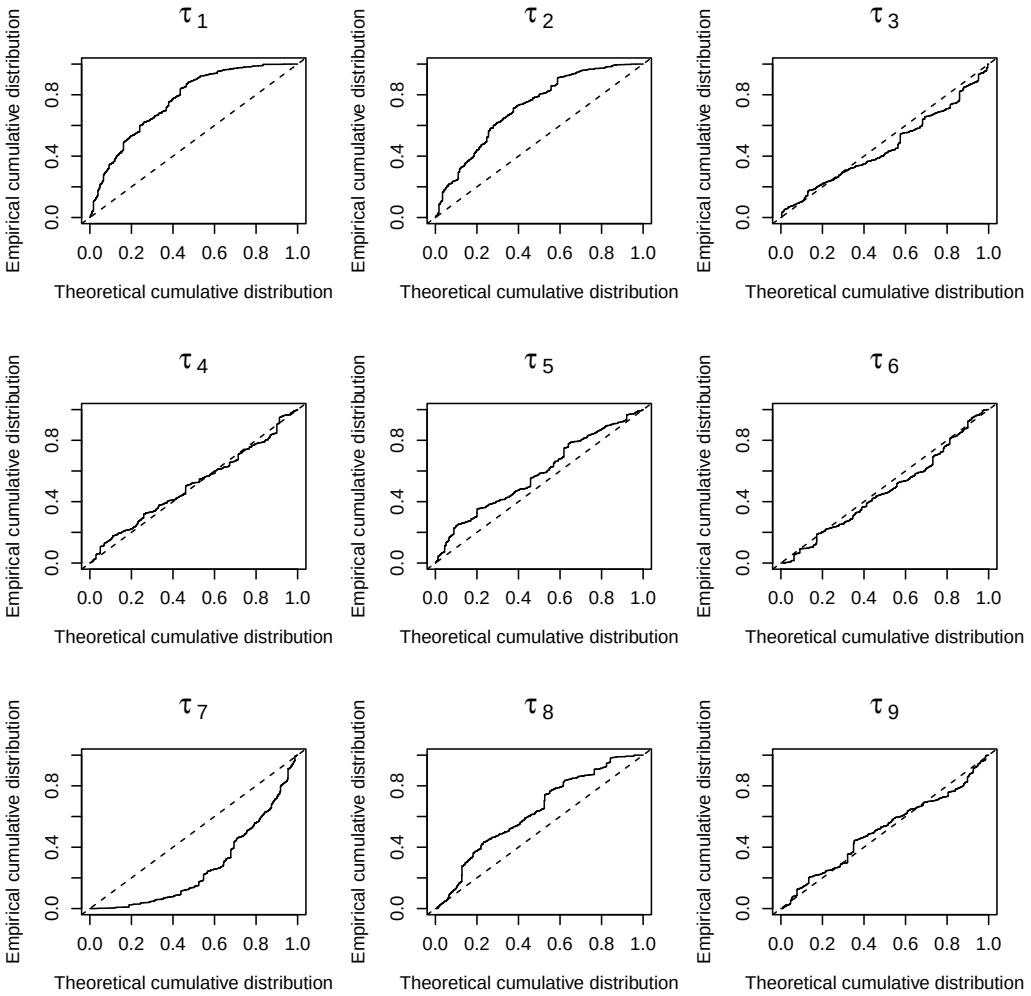

Finally, we would like to mention that, although times per iterations are considerably smaller for MCMC, it is worth indicating that, while MCMC provides a sample from the joint posterior distribution of the parameters in the model, INLA with IS provides a sample from the parameters estimated with IS and the (conditional) posterior marginal distributions for the remaining parameters in the model. Hence, in our view, IS with INLA provides a more informative output per iteration than MCMC. This issue clearly illustrates the fact that time per iteration times cannot be really directly compared between the MCMC and IS with INLA approaches.

Figure 10: Graphical diagnostics for the model with shared slope in the sleep deprivation study (2/2).

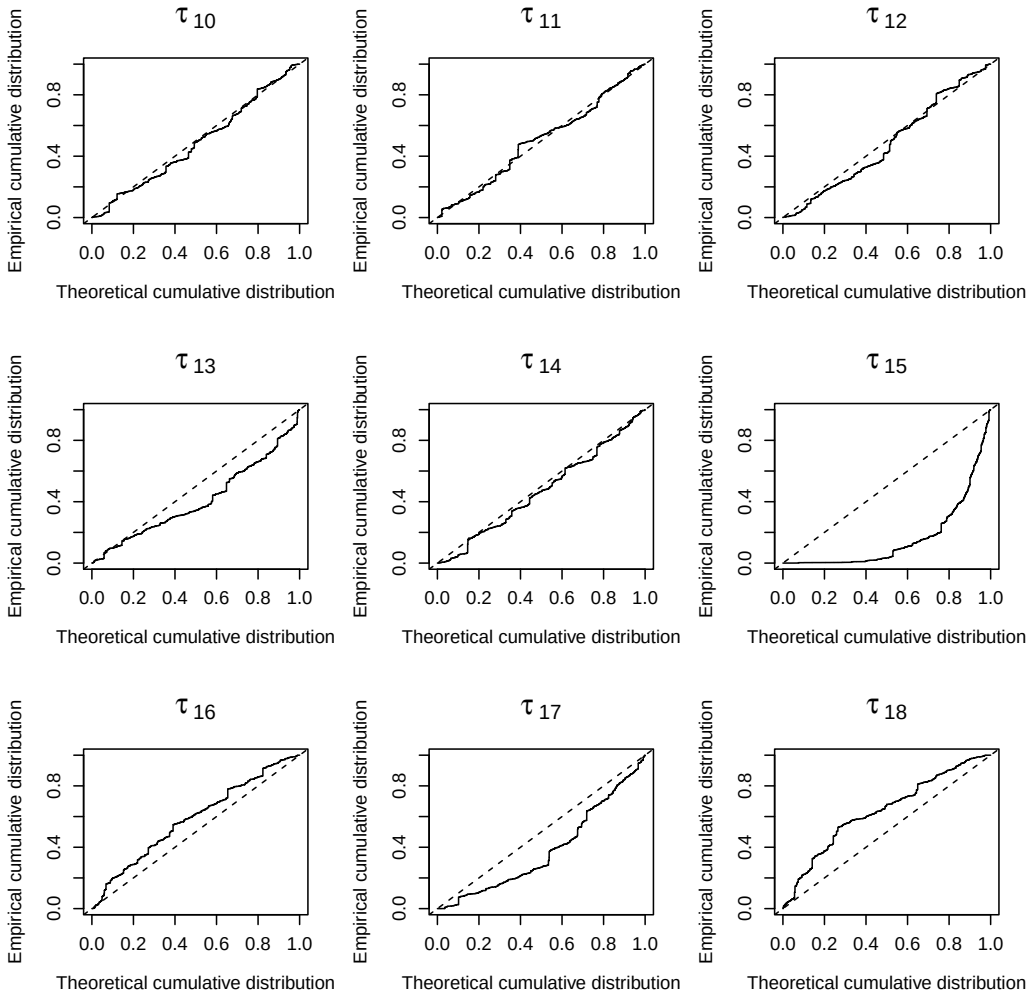

Figure 11: Graphical diagnostics for the model with random slopes in the sleep deprivation study (1/2).

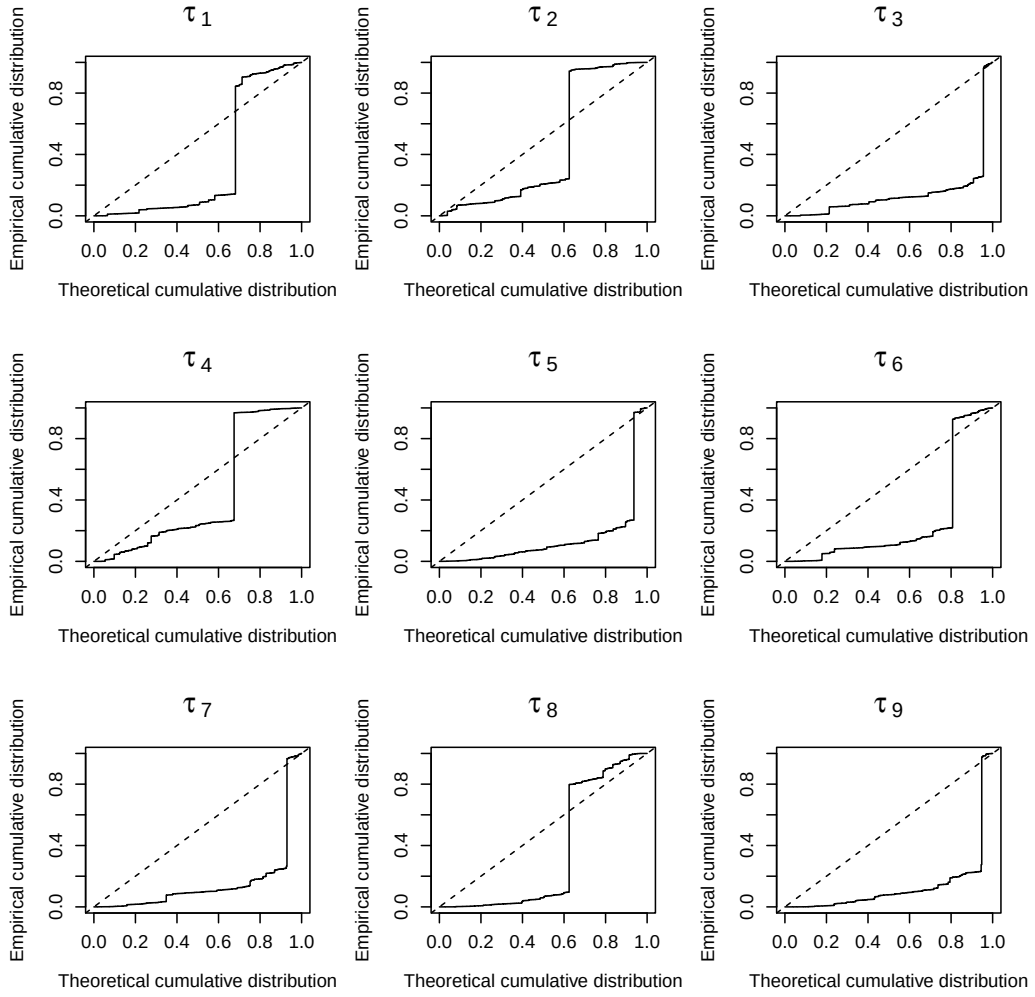

Figure 12: Graphical diagnostics for the model with random slopes in the sleep deprivation study (2/2).

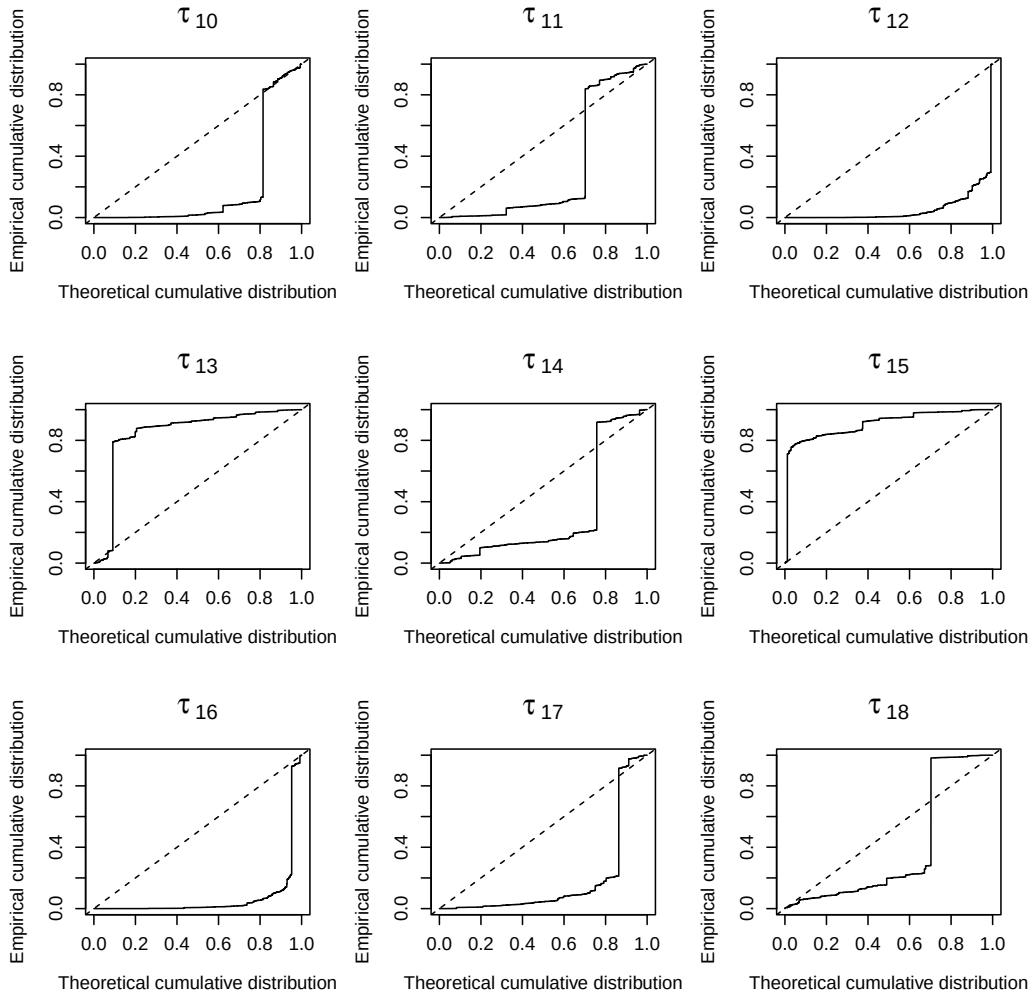

Table 2: Summary of the computation times (in seconds) required to fit the different models using INLA with AMIS and MCMC.

Simulation study				
	AMIS		MCMC	
Model	Total Elapsed	Time per sample	Total Elapsed	Time per sample
Poisson	1673.41	0.11	801.24	0.0073
Negative binomial	37814.47	2.52	593.14	0.0054
Gaussian (scenario 1)	3382.17	0.22	589.52	0.0054
Gaussian (scenario 2)	3303.17	0.22	589.52	0.0054
Gaussian (scenario 3)	2209.47	0.20	589.52	0.0054
Gaussian (scenario 4)	2366.28	0.22	589.52	0.0054
Gaussian (scenario 5)	4083.92	0.27	589.52	0.0054
Gaussian (scenario 6)	23854.80	0.43	589.52	0.0054
Examples				
	AMIS		MCMC	
Model	Total Elapsed	Time per sample	Total Elapsed	Time per sample
Infant mortality (Poisson)	2323.63	0.15	29.84	0.0003
Infant mortality (neg. binomial)	4865.71	0.32	15.97	0.0001
Sleep study (shared slope)	8198.79	0.39	18.76	0.0002
Sleep study (random slopes)	8561.52	0.41	20.73	0.0002