Received 6 June 2015,

Accepted 17 January 2016



Published online 9 February 2016 in Wiley Online Library

(wileyonlinelibrary.com) DOI: 10.1002/sim.6897

# Bayesian linear regression with skew-symmetric error distributions with applications to survival analysis

Francisco J. Rubio<sup>a\*†</sup> and Marc G. Genton<sup>b</sup>

We study Bayesian linear regression models with skew-symmetric scale mixtures of normal error distributions. These kinds of models can be used to capture departures from the usual assumption of normality of the errors in terms of heavy tails and asymmetry. We propose a general noninformative prior structure for these regression models and show that the corresponding posterior distribution is proper under mild conditions. We extend these propriety results to cases where the response variables are censored. The latter scenario is of interest in the context of accelerated failure time models, which are relevant in survival analysis. We present a simulation study that demonstrates good frequentist properties of the posterior credible intervals associated with the proposed priors. This study also sheds some light on the trade-off between increased model flexibility and the risk of over-fitting. We illustrate the performance of the proposed models with real data. Although we focus on models with univariate response variables, we also present some extensions to the multivariate case in the Supporting Information. Copyright © 2016 John Wiley & Sons, Ltd.

# 1. Introduction

The classical assumption of normality of the residual errors in linear regression models (LRMs) can be restrictive in practice. In many cases, the use of more flexible parametric distributional assumptions is necessary to capture departures from normality. These departures often arise when the sample contains outliers or when the residual errors are asymmetric. Several approaches have been proposed for parametrically modelling these departures from normality such as the use of scale mixtures of normal (SMN) distributions [1], and skew-elliptical distributions [2–4], among others (see also [5] and [6]). In the absence of strong prior knowledge about the model parameters, a way for conducting Bayesian inference consists of using noninformative priors (often referred to as objective Bayesian inference). In a general sense, these kinds of priors are functions of the parameters that produce posterior distributions with good frequentist properties. The use of objective Bayesian inference in LRMs with symmetric errors has been widely studied. In particular, [1] studied the propriety of the posterior distribution associated with Bayesian LRMs with an improper prior structure and residual errors distributed according to the family of SMNs. The use of more general prior structures has been recently studied in [7] and [8] for particular members of the SMN family. In a related vein, [9] studied the use of Jeffreys-type priors in the context of accelerated failure time (AFT) models with SMN errors (which are LRMs for the logarithm of the survival times). However, there are fewer studies related to the use of noninformative priors with flexible error distributions that allow for capturing skewness. [10] proposed a class of multivariate twopiece distributions that allow for capturing skewness. They employed these distributions for modelling the errors in multivariate LRMs and also proposed an improper prior structure, which is similar to that in [11]. Another recent reference is [12], who studied AFT models with errors distributed according to the

Keywords: accelerated failure time model; flexible errors; model selection; multivariate; noninformative prior; skewness

<sup>&</sup>lt;sup>a</sup>Department of Statistics, University of Warwick, Coventry CV4 7AL, U.K.

<sup>&</sup>lt;sup>b</sup>CEMSE Division, King Abdullah University of Science and Technology, Thuwal 23955-6900, Saudi Arabia \*Correspondence to: Francisco J. Rubio, Department of Statistics, University of Warwick, Coventry, CV4 7AL, U.K. †E-mail: Francisco.Rubio@warwick.ac.uk

# Statistics in Medicine

generalised extreme value distribution. They proposed an improper prior for this LRM, but their conditions for the propriety of the posterior involve truncating the support of the prior distribution on the shape parameter, which in turns reduces the flexibility of the model, as well as a condition on the prior for the scale parameter.

Here, we study the use of the class of skew-symmetric scale mixtures of normal (SSSMN) distributions for modelling residual errors in LRMs. This class contains, for instance, the skew-normal distribution [13], the skew-*t* distribution [2] and the skew-logistic distribution [14], as well as the corresponding symmetric models as particular cases. We propose a general improper prior structure that preserves the propriety of the posterior distribution in the sense that the posterior exists under the same conditions as in the model with SMN errors. This result allows us to appeal to existent results in order to show the propriety of the posterior associated with LRMs with SSSMN errors and the proposed prior structure. We discuss the use of the proposed Bayesian models in survival analysis as well as the impact of using flexible distributions in the prediction of the remaining life of patients that survived beyond the end of the study.

The paper is organised as follows. In Section 2, we describe the LRMs of interest and introduce the proposed noninformative prior structure. We provide easy to check sufficient conditions for the propriety of the corresponding posterior distribution. These results represent an extension to previous results for LRMs with symmetric errors [1, 7, 8, 11]. In Section 3, we extend our results to cases where the sample contains censored response variables for application of the proposed models to survival analysis. In Section 4, we present a simulation study where we illustrate the good frequentist properties of the 95% credible intervals obtained with the proposed priors and discuss some issues associated with certain skew-symmetric models. In Section 5, we present two applications with publicly available data sets in the context of biometric measures and survival analysis. All proofs are presented in the Supporting Information.

#### 2. Linear regression with skew-symmetric errors

#### 2.1. Bayesian model

Recall first that a real random variable Z is said to be distributed according to a skew-symmetric distribution if its probability density function (PDF) can be written as [15]

$$s(z) = 2f(z)\varphi(z), \ z \in \mathbb{R},$$
(1)

where f is a symmetric PDF and  $\varphi$  :  $\mathbb{R} \to [0, 1]$  is a function that satisfies  $\varphi(z) = 1 - \varphi(-z)$ . We focus on the family of parametric skew-symmetric distributions of the type

$$s(z|\xi,\omega,\lambda,\delta) = \frac{2}{\omega} f\left(\left.\frac{z-\xi}{\omega}\right|\delta\right) G\left(\lambda\frac{z-\xi}{\omega}\right),\tag{2}$$

where *G* is a cumulative distribution function (CDF) with continuous symmetric PDF *g* with support on  $\mathbb{R}$ ,  $\xi \in \mathbb{R}$  is a location parameter,  $\omega > 0$  is a scale parameter,  $\lambda \in \mathbb{R}$ , and  $\delta \in \Delta \subset \mathbb{R}$  is a shape parameter. We will refer to *f* as the 'baseline density'. Note that this structure allows for the case where f = g. We do not consider the case where *g* contains an unknown shape parameter and  $g \neq f$  because this would produce a model with five parameters that may play redundant roles. This structure covers many cases of practical interest such as the skew-normal distribution, the skew-*t* distribution [2] with  $\delta$ degrees of freedom, the skew-logistic distribution [14] and the skew-slash distribution [16]. Density (2) is asymmetric for  $\lambda \neq 0$ , and s = f for  $\lambda = 0$ . If a random variable *Z* has distribution (2), we will denote it as  $Z \sim SS(\xi, \omega, \lambda, \delta; f, g)$ .

Consider now the linear regression model:

$$y_j = \mathbf{x}_j^{\mathsf{T}} \boldsymbol{\beta} + \varepsilon_j, \ j = 1, \dots, n,$$
(3)

where  $y_j \in \mathbb{R}$ ,  $\boldsymbol{\beta}$  is a *p*-dimensional vector of regression parameters,  $\boldsymbol{\varepsilon}_j \stackrel{i.i.d.}{\sim} SS(0, \omega, \lambda, \delta; f, g)$ , and  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^{\mathsf{T}}$  is a known  $n \times p$  design matrix of full column rank. In some cases, we might be interested on centring the regression model on the mean or some quantile (such as the median); in which case, we have to properly centre the error distribution. This centring strategy is typically carried out after the estimation



We adopt the general prior structure:

$$\pi(\boldsymbol{\beta}, \boldsymbol{\omega}, \boldsymbol{\lambda}, \boldsymbol{\delta}) = \frac{p(\boldsymbol{\lambda}, \boldsymbol{\delta})}{\boldsymbol{\omega}^a},\tag{4}$$

**Statistics** 

where  $p(\lambda, \delta)$  is a proper prior. This prior structure covers the structure of several priors obtained by formal rules such as the independence Jeffreys prior, the Jeffreys-rule prior and the reference prior [17]. Given that this prior is improper, we need to investigate conditions for the existence of the corresponding posterior distribution.

In order to provide a general propriety result, we restrict our study to the case when the baseline density f in (2) belongs to the family of SMN distributions. Recall that a symmetric PDF f is said to be an SMN if it can be written as

$$f(z|\delta) = \int_{\mathbb{R}_+} \tau^{1/2} \phi\bigl(\tau^{1/2} z\bigr) \, dH(\tau|\delta), \ z \in \mathbb{R},$$

where *H* is a mixing distribution with positive support and  $\phi$  represents the standard normal PDF. The family of SMNs contains important models such as the normal distribution, the logistic distribution, the Student-*t* distribution, the Laplace distribution, the symmetric generalised hyperbolic distribution and the slash distribution. Under this setting, we have the following results (see the Supporting Information for a formal proof).

- (1) Consider the model (3)–(4). Assume that f in (2) is a scale mixture of normals and that rank(**X**) = p. If a = 1, a sufficient condition for the propriety of the posterior of ( $\beta, \omega, \delta, \lambda$ ) is n > p.
- (2) If a > 1, a sufficient condition for the propriety of the posterior is n > p + 1 a and

$$\int \tau^{-\frac{a-1}{2}} \tilde{p}(\delta) dH(\tau|\delta) d\delta < \infty,$$
(5)

where  $\tilde{p}(\delta)$  represents the marginal prior on  $\delta$ .

These results indicate that the introduction of skewness in the distribution of the residual errors, by means of the skew-symmetric construction (2), does not affect the existence of the posterior distribution, as long as we use a proper prior on the skewness parameter  $\lambda$ . For a = 1, the propriety of the posterior is guaranteed by having more observations than covariates for the class of LRMs with SSSMN errors. The results presented here are satisfied with probability one (for almost all sets of observations). We refer the reader to [11, 18], [1] and [9] for a discussion on zero-probability events that induce improper posteriors. This includes, for instance, samples with a certain number of responses that can be represented as an exact combination of their covariates which, in principle, have probability zero of occurrence under a continuous model but that may appear in practice due to rounding of the measurements. For a > 1, the conditions for the existence of the posterior distribution become more restrictive, and condition (5) has to be checked case by case. For the cases when the residual error distribution is a skew-normal or a skewlogistic (normal or logistic baseline f), the condition (5) is automatically satisfied [9], and therefore, the condition n > p + 1 - a is sufficient for the propriety of the posterior. Similarly, for cases when the baseline PDF f is a generalised hyperbolic distribution (with either fixed tail parameter  $\delta > 0$  or with a compactly supported marginal prior  $\tilde{p}(\delta)$ ) or a Laplace distribution, the condition n > p + 1 - a is also sufficient for the existence of the posterior [7,8].

2.1.1. The role of  $\lambda$  in skew-symmetric models. An aspect that has been little discussed in the literature is the overall influence of the parameter  $\lambda$  on the shape of a skew-symmetric PDF (1). We already mentioned that the PDF (1) is left/right skewed for  $\lambda \leq 0$ . However, the parameter  $\lambda$  controls other features as well such as spread and location of the mode. Moreover, for some combinations of f and G, the parameter  $\lambda$ has little influence of the shape of the PDF for a certain range of values around  $\lambda = 0$ . For instance, in the skew-normal PDF, the shape parameter  $\lambda$  has little influence on the asymmetry when  $|\lambda| \leq 1$ . In this region,  $\lambda$  mainly controls the location and spread of the PDF. Figure 1 shows that a skew-normal PDF with parameter  $\lambda = 1$  can be reasonably well approximated with a normal PDF. This phenomenon is also observed in other models where f and G have lighter tails than normal, while models with heavier tails



**Figure 1.** Normal probability density function with parameters (0.505, 0.817) (dashed line) *versus* skew-normal probability density function with parameters (0, 1, 1) (continuous line).

seem to be exempt from this problem. Intuitively, this suggests that it is hard to estimate the parameters  $(\xi, \omega, \lambda)$  (both from a Classical and a Bayesian perspective) in some light-tailed skew-symmetric models when the true value lies in a certain interval centred at  $\lambda = 0$  and the sample size is small or moderate. Consequently, if one uses a skew-symmetric distribution, with lighter or equal tails than normal, to model the errors in a LRM, it is expected to observe high correlation between the estimator of the intercept regression parameter, the scale parameter and the skewness parameter (Section 4). Given that these issues only appear when the errors are nearly symmetric, a simple solution consists of testing for  $\lambda = 0$  in order to identify the best and most parsimonious model for the data. Moreover, in the context of our paper, this issue has little relevance because we are only considering the cases when *f* is an SMN, leaving the normal as the only problematic case.

#### 2.2. Discussion on the choice of the priors for the shape parameters

The choice of the priors for the shape parameters  $(\lambda, \delta)$  deserves further discussion. We present a summary of some informative a noninformative priors used for the shape parameters in skew-symmetric models.

#### 2.2.1. Priors for $\lambda$ .

(i). Jeffreys priors. For the skew-normal regression model, we can employ the reference prior of the skewness parameter  $\lambda$ . This prior was obtained by [19], who showed that this is proper and well defined. [20] showed that this prior can be reasonably well approximated using a Student-*t* distribution with 1/2 degrees of freedom and scale parameter  $\sigma_0 = \pi/2$ . In [21], a similar result was also established for the skew-*t* model with degrees of freedom  $\delta > 2$ . [17] characterised the marginal reference prior of the parameter  $\lambda$  in a large class of skew-symmetric models and showed that this is symmetric, with tails of order  $O(|\lambda|^{-3/2})$ , and proper. These kinds of priors have been shown to induce posteriors with good frequentist properties [17, 19, 21]. However, the reference priors of  $\lambda$  are well defined at  $\lambda = 0$  only if the second moment of *f* is finite. As shown in [17], the Jeffreys/reference prior of  $\lambda$  can be calculated as follows:

$$p(\lambda) \propto \sqrt{\int_0^\infty x^2 f(x) \frac{g(\lambda x)^2}{G(\lambda x)[1 - G(\lambda x)]}} dx.$$
 (6)

- (ii). **Matching priors**. [22] calculated the matching prior for the skewness parameter  $\lambda$  in the skewnormal model. They showed that this prior is proper, data-independent, bimodal, with tails of order  $O(|\lambda|^{-3/2})$ , and that this prior induces a posterior with good frequentist properties. However, the bimodality of this prior might complicate sampling from the posterior for small and moderate samples. To the best of our knowledge, matching priors have not been calculated for other skew-symmetric models.
- (iii). **Heuristic priors**. An alternative noninformative prior can be constructed by using the popular alternative parameterisation of the skew-normal distribution in terms of the parameter  $\rho = \lambda/\sqrt{1 + \lambda^2} \in (-1, 1)$ . If we assign a uniform prior on the parameter  $\rho$ , this induces the following

proper prior on  $\lambda$ :

$$p(\lambda) \propto \frac{1}{(1+\lambda^2)^{3/2}}.$$
(7)

This prior has been used recently in [23].

(iv). **Informative priors**. In cases where there is reliable prior information, one can appeal to any proper prior that can capture the features of our prior beliefs. In particular, [23] proposed the use of the skew-normal distribution as a prior distribution, with hyperparameters ( $\xi_0, \omega_0, \lambda_0$ ), for  $\lambda$ . This prior allows the user to control the mass allocated on values of  $\lambda \leq 0$ , which are related to negative and positive skewness.

2.2.2. Priors for  $\delta$ . Regarding the tail parameter  $\delta$ , [24] proposed a general method for constructing weakly informative priors for kurtosis parameters. The idea consists of assigning a uniform prior to a bounded measure of kurtosis applied to the symmetric baseline density  $f(\cdot|\delta)$ . In order for this method to be applicable, the chosen measure of kurtosis must be a one-to-one function of the parameter  $\delta$ . This strategy induces a proper prior on the parameter  $\delta$  which can be interpreted as a weakly informative prior, in the sense that it assigns a flat prior on a function that represents the influence of the parameter  $\delta$  on the shape of the density. In the case of the degrees of freedom of the Student-*t* distribution, [24] showed that this strategy produces a prior with a behaviour similar to that of the approximation to the Jeffreys prior for this parameter [25] proposed in [26].

In addition, these priors on  $\delta$  can be coupled with the Jeffreys priors in order to produce a joint prior on  $(\lambda, \delta)$  by using the decomposition  $p(\lambda, \delta) = p(\lambda|\delta)p(\delta)$ , where

$$p(\lambda|\delta) \propto \sqrt{\int_0^\infty x^2 f(x|\delta) \frac{g(\lambda x)^2}{G(\lambda x)[1 - G(\lambda x)]}} dx.$$

This prior also has tails of order  $O(|\lambda|^{-3/2})$ , for each value of  $\delta$ . If the prior  $p(\lambda|\delta)$  is not too variable for different values of  $\delta$ , we may opt for an independent structure  $p(\lambda, \delta) = p(\lambda)p(\delta)$  for practical purposes. For instance, in the skew-*t* distribution, the prior  $p(\lambda|\delta)$  can be reasonably well approximated with a Student-*t* distribution with 1/2 degrees of freedom and scale parameter  $\sigma_0$ . A reasonable value of  $\sigma_0$ depends on the degrees of freedom  $\delta$ . Figure 2 shows the values of  $\sigma_0$  that produce a reasonable approximation (obtained by matching the mode) for  $\delta = 2.1, 2.2, ..., 50$ . We observe that  $\sigma_0$  varies between 0.5 and  $\pi/2$ . In this case, for the sake of simplicity, we adopt the prior  $p(\lambda, \delta) = p(\lambda)p(\delta)$ , where  $p(\lambda)$  is a Student-*t* distribution with 1/2 degrees of freedom and scale parameter  $\sigma_0 = \pi/2$ , which corresponds to the reference prior of  $\lambda$ . In Section 4, we show, through a simulation study, that this prior structure induces a posterior with good frequentist properties.



**Figure 2.** Value of the scale parameter  $\omega_0$  used in the Student-*t* approximation to the Jeffreys prior of  $\lambda$  in the skew-*t* model with  $\delta$  degrees of freedom.

# 3. Accelerated failure time models

A closely related family of LRMs that are of great interest in survival analysis is AFT models. The basic idea behind this kind of LRM consists of modelling a set of survival times  $\mathbf{T} = (T_1, \dots, T_n)^{\mathsf{T}}$  in terms of a set of covariates  $\boldsymbol{\beta}$  through the model equation:

$$\log T_j = \mathbf{x}_j^{\mathsf{T}} \boldsymbol{\beta} + \boldsymbol{\varepsilon}_j, \ j = 1, \dots, n.$$
(8)

In this context, the distribution of the errors  $\varepsilon_j$  is typically assumed to be normal or that  $\exp(\varepsilon_j)$  follows a Weibull distribution. Given that the normality assumption can be restrictive in practice, several alternative models have been proposed. For instance, from a classical inferential framework, [27] proposed the use of the Birnbaum–Saunders distribution, which corresponds to modelling the survival times  $T_j$  as a log-Birnbaum–Saunders distribution, and [28] employed the class of log-two-piece distributions. From a Bayesian perspective, [9] employed the class of log scale mixture of normals for modelling survival times; meanwhile, [12] employed the log generalised extreme value distribution.

We consider the class of log skew-symmetric (LSS) distributions [29],

$$s_l(t|\xi,\omega,\lambda,\delta) = \frac{2}{\omega t} f\left(\left.\frac{\log t - \xi}{\omega}\right|\delta\right) G\left(\lambda\frac{\log t - \xi}{\omega}\right), \quad t > 0, \tag{9}$$

for modelling the survival times  $T_j$  in (8), where f and g are as described in Section 2. It includes, for example, the log-skew-normal and log-skew-t distributions [30, 31]. This class of distributions can also be seen as an extension of the class of log-symmetric distributions (which is obtained with  $\lambda = 0$ ). If a positive random variable T has distribution (9), we denote it as  $T \sim \text{LSS}(\xi, \omega, \lambda, \delta; f, g)$ . More specifically, we assume that  $T_j | \mathbf{x}_j, \boldsymbol{\beta}, \omega, \lambda, \delta \sim \text{LSS}(\mathbf{x}_j^\top \boldsymbol{\beta}, \omega, \lambda, \delta; f, g)$ , which f in (9) is a scale mixture of normals, and that rank( $\mathbf{X}$ ) = p. If we adopt the prior structure (4) for this model, then the corresponding posterior is proper under the conditions in Section 2. However, in the context of AFT models, the presence of (right–, left–, and interval–) censored responses is quite common. Suppose that  $T_j | \mathbf{x}_j, \boldsymbol{\beta}, \omega, \lambda, \delta \sim \text{LSS}(\mathbf{x}_j^\top \boldsymbol{\beta}, \omega, \lambda, \delta; f, g)$ , which f in (9) is a scale mixture of normals, and that rank( $\mathbf{X}$ ) = p. Consider the prior structure (4) for this model. Suppose also that  $n_c \leq n$  observations are censored, and  $n_o = n - n_c$  are observed. Then,

- (1). If a = 1, a sufficient condition for the propriety of the corresponding posterior is  $n_0 > p$ .
- (2). If a > 1, a sufficient condition for the propriety of the posterior is  $n_o > p + 1 a$  together with (5).

In the Supporting Information, we provide a study of the extreme case in which all of the observations are censored (Corollary 2). This scenario requires a more careful analysis given that, intuitively, samples containing only censored observations contain little information about the parameters, and thus, one has to be careful about using noninformative priors in such extreme scenarios.

# 4. Simulation study

In this section, we present a simulation study that illustrates the performance of the proposed models. We study the regression model:

$$y_j = \beta_1 + \beta_2 x_{1j} + \beta_3 x_{2j} + \varepsilon_j, \ j = 1, \dots, n,$$

where we simulate the variables  $x_{1j}$  and  $x_{2j}$  from a standard normal distribution and consider different combinations of the distribution of the residual errors and the sample size *n*.

In the first scenario, we simulate the residual errors from a skew-normal distribution with scale parameter  $\omega = 0.5$  and skewness parameter  $\lambda = 0, 1, 2, 3, 4, 5, (\beta_1, \beta_2, \beta_3) = (1, 2, 3)$  and n = 100, 250, 500. Negative values of  $\lambda$  produce similar results and are, therefore, omitted. We adopt the product prior structure (4) with a = 1 and  $p(\lambda)$  given by the reference prior of this parameter [19]. We use the Student-*t* approximation from [20] to facilitate its implementation. For each of these scenarios, we obtain N = 1000samples of size 2000 from the posterior distribution using the R [32] t-walk MCMC sampler [33] after a burn-in period of 5000 iterations and thinned to every 25th iteration (55 000 iterations in total). Then, we calculate the proportion of 95% credible intervals that include the true value of the parameters, the 5%, 50% and 95% quantiles of the maximum likelihood estimators, the maximum *a posteriori* estimators, and posterior median estimators, as well as the median Bayes factor associated with the hypothesis  $H_0$ :  $\lambda = 0$ . The Bayes factor is approximated using the Savage–Dickey density ratio [34]. In the second scenario, we simulate the residual error from a skew-logistic distribution with scale parameter  $\omega = 0.5$  and skewness parameter  $\lambda = 0, 1, 2, 3, 4, 5$ . We adopt the product prior structure (4) with a = 1 and  $p(\lambda)$  given by the reference prior of this parameter [17]. We use the Student-*t* approximation (1/2 degrees of freedom and scale 4/3) from [17] to facilitate its implementation. In the third scenario, we simulate the residual errors from a skew-*t* distribution with degrees of freedom  $\delta = 3$  and scale parameter  $\omega = 0.5$ . We adopt the prior on the skewness parameter described in the first simulation scenario. For the degrees of freedom  $\delta$ , we use the approximation to the Jeffreys prior proposed in [26]:

$$\pi(\delta) = \frac{2d\delta}{(\delta+d)^3}.$$
(10)

**Statistics** 

Medicine

We fix the hyperparameter d = 2, which induces a prior with mode at  $\delta = 1$ .

Results are reported in Tables I–III in succeeding text and Section 3, and Tables 1S–6S of the Supporting Information. For the first scenario (skew-normal errors, Tables I–III), we can observe that the

<b>Table I.</b> Point estimators, coverage proportions and Bayes factors: linear regression model with residual errors simulated from a skew-normal distribution, $n = 100$ .											
Method		MLE			MAP			Median		Coverage	BF
	5%	50%	95%	5%	50%	95%	5%	50%	95%		
$\lambda = 0$											
$\beta_1$	4.456	4.917	5.521	4.463	4.996	5.519	4.559	4.999	5.380	0.987	—
$\beta_2$	1.913	2.002	2.089	1.911	2.001	2.091	1.916	2.003	2.085	0.956	_
$\beta_3$	2.908	3.001	3.087	2.905	2.999	3.086	2.906	3.000	3.087	0.942	—
ω	0.492	0.613	0.761	0.475	0.542	0.714	0.504	0.581	0.699	0.849	—
λ	-1.993	0.285	2.224	-1.452	-0.003	1.613	-1.196	0.009	1.520	0.989	2.141
$\lambda = 1$											
$\beta_1$	4.788	5.009	5.317	4.796	5.057	5.649	4.856	5.193	5.507	0.947	—
$\beta_2$	1.931	2.001	2.069	1.927	2.001	2.069	1.928	2.002	2.067	0.950	—
$\beta_3$	2.930	3.001	3.066	2.929	3.001	3.071	2.931	3.002	3.066	0.963	—
ω	0.373	0.485	0.635	0.388	0.450	0.613	0.417	0.483	0.603	0.994	—
λ	0.007	0.992	2.550	-0.988	0.190	2.063	-0.713	0.248	1.974	0.964	2.029
$\lambda = 2$											
$\beta_1$	4.886	5.003	5.252	4.885	5.009	5.435	4.900	5.059	5.359	0.911	—
$\beta_2$	1.941	2.001	2.058	1.938	2.003	2.060	1.942	2.001	2.058	0.955	—
$\beta_3$	2.944	3.000	3.060	2.941	3.001	3.061	2.943	3.001	3.059	0.952	—
ω	0.354	0.496	0.600	0.336	0.447	0.593	0.365	0.462	0.589	0.947	—
λ	0.451	2.083	4.942	-0.315	1.585	3.389	0.035	1.441	3.973	0.901	1.093
$\lambda = 3$											
$\beta_1$	4.909	5.002	5.115	4.904	5.003	5.147	4.915	5.020	5.257	0.912	—
$\beta_2$	1.946	1.998	2.047	1.944	1.998	2.050	1.947	1.999	2.047	0.955	—
$\beta_3$	2.949	3.000	3.054	2.948	3.000	3.054	2.950	3.001	3.053	0.950	—
ω	0.395	0.491	0.590	0.319	0.484	0.585	0.359	0.482	0.587	0.929	_
λ	1.488	3.225	10.269	-0.154	2.463	5.112	0.577	2.742	6.939	0.904	0.223
$\lambda = 4$	4.020	5 002	5 00 4	4 0 2 2	5 000	5 001	4.020	5 000	5 107	0.046	
$p_1$	4.928	5.002	5.084	4.922	5.000	5.091	4.929	5.008	5.137	0.946	
$\beta_2$	1.954	2.000	2.046	1.953	2.001	2.047	1.954	2.001	2.045	0.964	_
$p_3$	2.954	3.002	3.047	2.953	3.003	3.048	2.955	3.002	3.048	0.964	_
	0.407	0.495	0.5/6	0.369	0.488	0.574	0.380	0.490	0.574	0.950	510-2
л 1	2.298	4.505	108.781	-1.076	3.227	8.417	1.502	3.807	12.553	0.941	5×10 -
$\lambda = 5$	4.022	4 000	5.0(0	4.020	4 007	5 074	4.020	5 000	5.000	0.054	
$p_1$	4.933	4.998	5.068	4.930	4.997	5.074	4.938	5.000	5.096	0.954	_
$p_2$	1.954	1.998	2.043	1.953	1.999	2.043	1.956	1.999	2.044	0.955	_
$p_3$	2.960	3.003	3.046	2.959	3.003	3.04/	2.962	3.003	3.045	0.969	_
ω	0.418	0.499	0.575	0.397	0.494	0.574	0.402	0.498	0.578	0.958	
λ	2.997	5.883	138.560	-1/.640	3.930	12.810	2.192	4.929	21.107	0.943	$2 \times 10^{-2}$

MAP, maximum a posteriori; MLE, maximum likelihood estimator.

**Table II.** Point estimators, coverage proportions and Bayes factors: linear regression model with residual errors simulated from a skew-normal distribution, n = 250.

Method		MLE			MAP			Median		Coverage	BF
	5%	50%	95%	5%	50%	95%	5%	50%	95%		
$\lambda = 0$											
$\beta_1$	4.574	4.965	5.424	4.568	5.008	5.429	4.665	4.994	5.335	0.988	_
$\beta_2$	1.948	2.000	2.057	1.945	1.999	2.057	1.947	2.000	2.056	0.945	_
$\beta_3$	2.947	2.998	3.052	2.945	2.996	3.055	2.946	2.997	3.053	0.952	
ω	0.504	0.584	0.683	0.484	0.525	0.656	0.508	0.555	0.640	0.859	
λ	-1.406	0.079	1.400	-1.240	0.030	1.234	-1.035	0.015	1.017	0.988	2.547
$\lambda = 1$											
$\beta_1$	4.869	5.000	5.290	4.872	5.024	5.548	4.897	5.157	5.430	0.903	_
$\beta_2$	1.957	2.000	2.047	1.957	2.000	2.049	1.959	2.000	2.047	0.943	_
$\beta_3$	2.957	3.001	3.044	2.956	3.001	3.044	2.958	3.001	3.043	0.949	_
ω	0.393	0.498	0.589	0.401	0.440	0.581	0.420	0.465	0.564	0.986	_
λ	0.014	1.018	1.782	-0.809	0.548	1.690	-0.428	0.390	1.553	0.903	2.282
$\lambda = 2$											
$\beta_1$	4.922	5.002	5.100	4.922	5.002	5.117	4.928	5.013	5.217	0.919	_
$\beta_2$	1.967	2.001	2.035	1.966	2.001	2.038	1.966	2.001	2.035	0.955	_
$\beta_4$	2.966	3.000	3.033	2.965	3.001	3.034	2.966	3.001	3.033	0.967	_
ω	0.420	0.498	0.565	0.358	0.492	0.564	0.386	0.489	0.564	0.919	
λ	1.231	2.046	3.159	0.718	1.871	2.886	0.592	1.897	3.009	0.917	0.080
$\lambda = 3$											
$\beta_1$	4.948	5.001	5.064	4.946	5.001	5.067	4.949	5.005	5.074	0.944	_
$\beta_2$	1.971	2.000	2.028	1.970	2.001	2.029	1.971	2.000	2.028	0.966	—
$\beta_3$	2.970	3.001	3.029	2.969	3.000	3.029	2.970	3.000	3.029	0.963	—
ω	0.437	0.497	0.552	0.431	0.494	0.550	0.430	0.495	0.551	0.950	
λ	2.034	3.093	4.676	1.851	2.845	4.117	1.826	2.950	4.330	0.947	$1 \times 10^{-11}$
$\lambda = 4$											
$\beta_1$	4.958	4.998	5.050	4.958	4.998	5.051	4.960	5.000	5.054	0.962	_
$\beta_2$	1.973	2.000	2.028	1.972	2.000	2.027	1.973	2.000	2.027	0.965	_
$\beta_3$	2.973	2.999	3.026	2.973	2.999	3.027	2.973	2.999	3.026	0.970	—
ω	0.449	0.497	0.549	0.445	0.495	0.545	0.446	0.496	0.546	0.966	—
λ	2.820	4.180	6.669	2.539	3.781	5.629	2.687	3.979	6.161	0.947	$1 \times 10^{-13}$
$\lambda = 5$											
$\beta_1$	4.962	4.997	5.042	4.960	4.997	5.044	4.962	4.999	5.044	0.953	—
$\beta_2$	1.976	2.001	2.025	1.975	2.001	2.026	1.976	2.001	2.026	0.967	—
$\beta_3$	2.975	3.000	3.026	2.974	2.999	3.026	2.974	2.999	3.026	0.959	—
ω	0.447	0.499	0.548	0.447	0.497	0.545	0.448	0.499	0.547	0.956	—
λ	3.550	5.289	8.647	3.196	4.692	7.200	3.356	4.965	7.821	0.959	$6 \times 10^{-12}$

MAP, maximum a posteriori; MLE, maximum likelihood estimator.

coverage of the credible intervals of the parameter  $\omega$  is greatly affected when the true value of  $\lambda$  is zero or one. Moreover, the Bayes factors associated with  $H_0$ :  $\lambda = 0$  favour the normal model in both cases, for all sample sizes, when  $\lambda = 0, 1$ , which emphasises the difficulty in identifying models with  $|\lambda| \leq 1$ . These scenarios correspond to the cases described in Section 2.1.1 where the skew-normal distribution is nearly or exactly symmetric, which in turns complicates the identification of the parameters ( $\beta_1, \omega, \lambda$ ), and induces a strong correlation between these parameters. Figures 3 and 4 show the scatter plots associated with a posterior sample of ( $\beta_1, \beta_2, \beta_3, \omega, \lambda$ ) for n = 250 and  $\lambda = 1, 5$ , respectively. These figures show the strong correlation between the parameters ( $\beta_1, \omega, \lambda$ ) for  $\lambda = 1$ , while the correlation between these parameters seems to be lower for  $\lambda = 5$ . We can also observe from these tables that the coverage improves as  $\lambda \ge 2$  and that the coverage proportions converge to the nominal value as the sample size increases. This study also indicates that we need at least 100 observations in order to obtain a good coverage. For the second and third scenarios (Tables 1S–6S), we observe a good coverage for all the sample sizes as well as a good behaviour of the Bayesian estimators compared with that of the corresponding maximum likelihood estimators.

In the Supporting Information, Tables 7S–9S, we present an additional simulation study using the configuration in the first scenario with the prior (7) for the parameter  $\lambda$ . This study shows that the prior

St	atisti	CS	
in	Med	licine	2

**Table III.** Point estimators, coverage proportions and Bayes factors: linear regression model with residual errors simulated from a skew-normal distribution, n = 500.

Method		MLE			MAP			Median		Coverage	BF
	5%	50%	95%	5%	50%	95%	5%	50%	95%		
$\lambda = 0$											
$\beta_1$	4.612	5.022	5.376	4.608	4.991	5.379	4.680	5.000	5.287	0.981	_
$\beta_2$	1.964	2.001	2.035	1.963	2.000	2.037	1.965	2.001	2.036	0.957	
$\beta_3$	2.962	3.000	3.039	2.959	3.000	3.040	2.962	3.000	3.040	0.951	
ω	0.502	0.567	0.649	0.487	0.520	0.630	0.507	0.541	0.615	0.848	
λ	-1.134	-0.045	1.181	-1.066	0.009	1.140	-0.807	-0.007	0.898	0.982	2.953
$\lambda = 1$											
$\beta_1$	4.897	5.000	5.263	4.896	5.012	5.450	4.910	5.112	5.333	0.892	_
$\beta_2$	1.969	1.999	2.029	1.967	1.999	2.031	1.969	1.999	2.031	0.948	_
$\beta_3$	2.970	3.000	3.031	2.968	3.000	3.033	2.970	3.000	3.032	0.950	_
ω	0.411	0.501	0.569	0.405	0.437	0.567	0.424	0.463	0.560	0.965	_
λ	0.055	1.010	1.602	-0.430	0.833	1.547	-0.130	0.544	1.505	0.885	2.279
$\lambda = 2$											
$\beta_1$	4.950	4.999	5.062	4.950	5.000	5.064	4.954	5.004	5.081	0.950	—
$\beta_2$	1.975	2.001	2.025	1.974	2.000	2.025	1.975	2.000	2.024	0.953	
$\beta_3$	2.976	3.001	3.026	2.976	3.001	3.027	2.976	3.001	3.026	0.949	
ω	0.447	0.499	0.545	0.440	0.497	0.544	0.436	0.496	0.543	0.942	
λ	1.446	2.022	2.706	1.362	1.943	2.585	1.317	1.961	2.633	0.937	$2 \times 10^{-26}$
$\lambda = 3$											
$\beta_1$	4.963	4.999	5.040	4.962	5.000	5.041	4.965	5.000	5.043	0.958	_
$\beta_2$	1.979	1.999	2.020	1.978	1.999	2.022	1.979	1.999	2.021	0.967	_
$\beta_3$	2.977	3.000	3.022	2.977	3.000	3.022	2.978	3.000	3.022	0.955	—
ω	0.458	0.498	0.538	0.456	0.497	0.537	0.458	0.498	0.538	0.957	_
λ	2.351	3.048	4.019	2.246	2.928	3.874	2.295	2.984	3.925	0.954	$4 \times 10^{-53}$
$\lambda = 4$											
$\beta_1$	4.968	5.001	5.033	4.968	5.001	5.034	4.969	5.001	5.034	0.950	
$\beta_2$	1.981	2.000	2.017	1.981	2.000	2.019	1.981	2.000	2.018	0.961	—
$\beta_3$	2.981	2.999	3.020	2.981	2.999	3.020	2.981	2.999	3.019	0.964	
ω	0.460	0.499	0.536	0.458	0.498	0.537	0.459	0.499	0.536	0.941	
λ	3.142	4.073	5.415	2.996	3.895	5.119	3.073	3.976	5.216	0.954	$3 \times 10^{-55}$
$\lambda = 5$											
$\beta_1$	4.974	5.000	5.027	4.974	5.000	5.029	4.974	5.000	5.028	0.962	
$\beta_2$	1.983	2.000	2.017	1.982	2.000	2.017	1.982	2.000	2.017	0.963	
$\beta_3$	2.983	2.999	3.017	2.983	2.999	3.017	2.983	2.999	3.017	0.966	
ω	0.465	0.499	0.532	0.463	0.497	0.532	0.465	0.498	0.532	0.953	—
λ	3.913	5.199	7.083	3.669	4.895	6.654	3.795	5.043	6.800	0.956	$5 \times 10^{-46}$



**Figure 3.** Scatter plot for a single posterior simulation of  $(\beta_1, \beta_2, \beta_3, \omega, \lambda)$  with n = 250 and  $\lambda = 1$ .





**Figure 4.** Scatter plot for a single posterior simulation of  $(\beta_1, \beta_2, \beta_3, \omega, \lambda)$  with n = 250 and  $\lambda = 5$ .

(7), which has been considered as a noninformative prior in the literature, induces a posterior with poor frequentist properties.

# 5. Applications

In this section, we illustrate the performance of the proposed Bayesian models with two real data sets. In the first example, we revisit the popular 'Australian Athletes' data set. We compare the performance of four types of distributional assumptions on the residual errors in an LRM studied in [35]. In the second example, we present an application in the context of survival analysis of cancer patients. Simulations from the corresponding posterior distributions are obtained using the t–walk [33]. Model comparison is conducted via Bayes factors which are calculated using an importance sampling (and corroborated using Laplace's method), the Bayesian information criterion (BIC) and log-predictive marginal likelihood (LPML). LPML is a measure that ranks the models of interest in terms of their predictive performance (see [36] for an overview that includes the case with interval censored observations). Larger values of LPML indicate a better predictive performance. The use of the Bayes factors with improper priors is justified in this context because we use the same prior structure (with improper priors only on the common parameters) for the different models in order to avoid the presence of arbitrary constants (see [9] for a discussion on this point). R codes are available upon request.

#### 5.1. Australian athletes data

Our first application concerns the study of the regression model [35]:

$$Lbm_i = \beta_1 Ht_i + \beta_2 Wt_i + \varepsilon_i, \quad j = 1, \dots, 102,$$
(11)

where  $Lbm_j$ ,  $Ht_j$  and  $Wt_j$  denote the lean body mass, the height and the weight of 102 Australian male athletes from the Australian Institute of Sport data set. We compare the models obtained with four distributional assumptions on the residual errors  $\varepsilon_j$ : a skew-normal distribution SN(0,  $\omega$ ,  $\lambda$ ) [13], a skew Student-*t* distribution SSt(0,  $\omega$ ,  $\lambda$ ,  $\delta$ ) with unknown degrees of freedom and the corresponding symmetric submodels (normal and Student-*t*). We adopt the prior structure (4) with a = 1. For the skewness parameter, in both cases, we use the Student-*t* approximation (with 1/2 degrees of freedom and scale  $\pi/2$ ) to the reference prior proposed in [20]. For the degrees of freedom, we use the prior (10) with hyperparameter d = 10. This prior structure allows us to match the priors associated with the different distributional assumptions, which in turn helps us to justify the use of Bayes factors for model comparison. The propriety of the corresponding posterior distributions is guaranteed by Proposition 1 and Corollary 1. We obtain a posterior sample of size 10 000 after a burn-in period of 25 000 iterations and thinned to every 25th iteration for each of these models (275 000 iterations in total). Table IV presents a summary of the posterior samples as well as the BIC, Bayes factors and LPML. All of the model selection criteria favour the model with skew-*t* errors, which suggest the presence of heavy tails and skewness.

In order to assess the goodness of fit of the SSt model, we propose a Bayesian residual analysis as follows. First, for each posterior sample of the parameters  $(\beta_1^{(k)}, \beta_2^{(k)})$ , k = 1, ..., 10000, we obtain the median of the residual errors  $\varepsilon_j^{(k)} = Lbm_j - \beta_1^{(k)}Ht_j - \beta_2^{(k)}Wt_j$ . By using the posterior samples of  $(\omega, \lambda, \delta)$ , we approximate the posterior predictive distribution of the errors and obtain a sample of size 10 000 from this distribution. Then, we obtain M = 10000 sub-samples of size 102 from the sample of the predictive distribution and construct a QQ-plot for each of these samples. Finally, using these M QQ-plots, we generate a predictive envelope by taking the 5% and 95% quantiles of the QQ-plots at each sample quantile point. This envelope can be used to visually assess the fit of the residuals and to detect shortcomings of a model in specific regions. Figure 5 shows the predictive quantile envelope associated with both the SSt and normal models. We notice that the envelope associated with the SSt model covers all the data points as well as the straight line, which represents perfect fit. On the other hand, the envelope produced with the normal model does not contain the straight line in several regions.

#### 5.2. North Central Cancer Treatment Group data

We now analyse the popular North Central Cancer Treatment Group data lung cancer data. This data set contains the survival times of n = 227 patients (the total number of patients is 228, but for the sake of simplicity, we have removed one patient with a missing covariate) with advanced lung cancer from the North Central Cancer Treatment Group. The sample contains 63 censored observations. The goal of this study was to compare the descriptive information from a questionnaire applied to a group of patients against the information obtained by each patient's physician, in terms of prognostic power [37]. We fit an AFT model (8) with three covariates 'age' (in years), 'sex' (Male = 1, Female = 2), 'ph.ecog' (ECOG performance score, 0 = good - 5 = dead) as well as an intercept, together with four distributional assumptions on the residual errors. We compare the inference obtained with skew-logistic errors (SS Logistic), skew-normal errors (SS Normal) and the corresponding symmetric sub-models (logistic and normal). We adopt the prior structure (4) with a = 1 and the reference prior for  $\lambda$  proposed in [17] (which

<b>Table IV.</b> AIS Data: posterior median, 95% credible intervals, BIC and Bayes factors againstthe SSt model.								
Model	SSt	SN	t	Ν				
$\beta_1$	0.05 (0.03,0.08)	0.05 (0.04,0.08)	0.06 (0.04,0.08)	0.08 (0.06,0.10)				
$\beta_2$	0.81 (0.75,0.87)	0.81 (0.76,0.85)	0.77 (0.74,0.81)	0.73 (0.69,0.77)				
ω	1.85 (1.18,2.77)	3.51 (2.94,4.22)	1.24 (0.96,1.62)	2.29 (2.01,2.64)				
λ	-2.43 (-7.04,-0.57)	-4.08 (-7.89, -2.18)	—	—				
δ	2.69 (1.54,6.03)	—	2.33 (1.44,4.28)	—				
BIC	231.07	238.21	233.00	241.83				
Bayes factor	_	$4 \times 10^{-4}$	$8 \times 10^{-3}$	$3 \times 10^{-7}$				
LPML	-217.39	-224.48	-218.84	-232.18				

AIS, Australian Institute of Sport; LMPL, log-predictive marginal likelihood;



Figure 5. Bayesian residual analysis: (a) skew-normal errors; (b) skew-t errors.

the SS logistic	c model.			
Model	SS logistic	SS normal	Logistic	Normal
Intercept	6.79 (5.82, 7.74)	7.19 (6.20, 8.28)	5.96 ( 4.98, 6.96)	6.48 (5.30, 7.62)
Age	-0.01 (-0.02,0.00)	-0.01 (-0.02,0.00)	-0.01 (-0.02, 0.00)	-0.02 (-0.03, 0.00)
Sex	0.44 (0.19,0.72)	0.45 (0.17,0.72)	0.49 (0.22, 0.76)	0.53 (0.23, 0.84)
ph.ecog	-0.37 (-0.55,-0.19)	-0.34 (-0.53, -0.15)	-0.41 (-0.60, -0.22)	-0.36 (-0.57,-0.16)
ω	0.71 (0.57, 0.86)	1.46 (1.28, 1.67)	0.55 (0.48,0.63)	1.04 (0.93, 1.17)
λ	-2.03 (-4.20,-0.64)	-43.26 (-5.90, -1.85)	_	_
BIC	555.93	563.67	562.17	580.96
Bayes factor	—	0.008	0.108	$1 \times 10^{-5}$
LPML	-268.53	-272.48	-272.91	-283.23

**Table V.** NCCTG lung cancer data: posterior median, 95% credible intervals, BIC and Bayes factors against the SS logistic model.

LPML, log-predictive marginal likelihood; NCCTG, North Central Cancer Treatment Group.

<b>Table VI.</b> NCCTG lung cancer data: quantiles of the predictive residual           life distribution for the median SS logistic and logistic models.								
Quantile	5%	25%	50%	75%	95%			
SS Logistic model								
Patient 1	1054.1	1163.4	1576.7	2120.3	3539.5			
Patient 2	1045.8	1155.6	1347.6	1693.3	2693.3			
Patient 3	1008.1	1200.1	1516.2	2048.7	3436.6			
Patient 4	896.0	1217.3	1706.0	2478.9	4363.4			
Patient 5	882.6	1071.5	1374.5	1881.8	3174.8			
Logistic model								
Patient 1	1043.9	1211.7	1549.4	2318.6	5700.5			
Patient 2	1052.9	1206.9	1520.3	2241.5	5453.9			
Patient 3	997.8	1160.0	1485.5	2225.7	5481.2			
Patient 4	857.6	1034.8	1378.8	2136.7	5387.6			
Patient 5	869.3	1013.5	1302.4	1956.2	4829.0			

NCCTG, North Central Cancer Treatment Group.

is approximated using a Student-*t* distribution with 1/2 degrees of freedom and scale 3/4). The propriety of the corresponding posterior distributions is guaranteed by Theorem 2 in the Supporting Information. We obtain a posterior sample of size 10 000 after a burn-in of 50 000 and thinned to every 25th iteration (300 000 iterations in total). Table V shows a summary of the posterior samples as well as the model comparison. The Bayes factors, BIC and LPML favour the model with SS logistic errors, which suggests the presence of skewness and slightly heavier tails than normal. In this case, we cannot obtain a residual analysis given that the data set contains censored observations. Although there have been some attempts to produce Bayesian residual analyses in the presence of censored observations [38], there does not seem to be an agreement on which of these methods is more appropriate.

Now, using the AFT models with SS Logistic and Logistic errors, we study the remaining life of five patients which survived beyond the end of the study. A proper prediction of the individual remaining life is of great importance in medicine because this information can be used for planning health care. Table VI presents a summary of the quantiles of these distributions for the first five censored patients. These quantiles were obtained by using the posterior AFT model centred at the median (averaged over all the posterior samples). We can observe that the quantiles higher than 50% are much larger in the model with Logistic errors. This difference seems to be produced by the larger values in the posterior sample of  $\omega$  in the model with Logistic errors, which are overestimated by the lack of flexibility of this error distribution.

#### 6. Discussion

We proposed a flexible class of LRMs with SSSMN errors that can capture a variety of tail behaviours and skewness. The proposed models represent an extension to LRM with SMN errors. The latter has been widely used to capture heteroscedasticity and the presence of outliers, but they cannot capture departures from symmetry. The use of SSSMN error distributions can capture additional unobserved heterogeneity which has the effect of inducing asymmetry in the residual errors. We introduced a noninformative prior structure that induces proper posteriors under rather mild conditions. We presented propriety results in a unified framework that covers the usual LRMs and AFT models. We provided tangible conditions for checking the propriety of the posterior distribution in cases when the response variables are censored. We have illustrated the need for this sort of extension with simulated and real data. Our simulation studies also indicate that the proposed prior structure induces posterior distributions with appealing frequentist properties. We have emphasised the usefulness of the proposed models in survival analysis, which are relevant in medicine. However, they can also be applied to other contexts such as finance, biology or industrial applications.

We implemented the proposed models using a general-purpose MCMC sampler. This was possible given that we were using distributions with numerically tractable PDF and CDF. In a more general framework, the Metropolis within Gibbs sampling strategy (which takes advantage of the stochastic representation of SMN) proposed in [9] can be extended to models with SSSMN errors by using the stochastic representation of this sort of distributions [15]. This only implies an additional step in the Gibbs sampler from [9]. In those cases when the PDF and CDF of the SSSMN of interest are tractable, it is possible to implement a Metropolis-within-Gibbs sampler with already available R packages such as 'spBayes' [39].

# Acknowledgements

We thank an Associate Editor and two referees for their very helpful comments. FJR gratefully acknowledges research support from EPSRC grant EP/K007521/1. MGG's research is supported by King Abdullah University of Science and Technology (KAUST).

# References

- 1. Fernández C, Steel M. Bayesian regression analysis with scale mixtures of normals. Econometric Theory 2000; 80:80-101.
- 2. Azzalini A, Capitanio A. Distributions generated by perturbation of symmetry with emphasis on a multivariate skew-t distribution. *Journal of Royal Statistical Society, Series B* 2003; **65**:367–389.
- Sahu S, Dey D, Branco M. A new class of multivariate skew distributions with applications to Bayesian regression models. *Canadian Journal of Statistics* 2003; 31:129–150.
- 4. Azzalini A, Genton M. Robust likelihood methods based on the skew-*t* and related distributions. *International Statistical Review* 2008; **76**:106–129.
- 5. Azzalini A, Capitanio A. The Skew-normal and Related Families. Cambridge University Press: Cambridge, United Kingdom, 2014.
- 6. Genton M. Skew-elliptical Distributions and Their Applications: A Journey Beyond Normality. Chapman & Hall/CRC: Boca Raton, Florida, 2004.
- 7. Choi H, Hobert J. Analysis of MCMC algorithms for Bayesian linear regression with Laplace errors. *Journal of Multivariate Analysis* 2013; **117**:32–40.
- Jung Y, Hobert J. Spectral properties of MCMC algorithms for Bayesian linear regression with generalized hyperbolic errors. Statistics & Probability Letters 2014; 95:92–100.
- 9. Vallejos C, Steel M. Objective Bayesian survival analysis using shape mixtures of log-normal distributions. *Journal of the American Statistical Association* 2015; **110**:697–710.
- 10. Ferreira J, Steel M. A new class of skewed multivariate distributions with applications to regression analysis. *Statistica Sinica* 2007; **17**:505–529.
- 11. Fernández C, Steel MFJ. Multivariate Student-t regression models: pitfalls and inference. Biometrika 1999; 86:153-167.
- 12. Roy V, Dey D. Propriety of posterior distributions arising in categorical and survival models under generalized extreme value distribution. *Statistica Sinica* 2014; **24**:699–722.
- 13. Azzalini A. A class of distributions which includes the normal ones. Scandinavian Journal of Statistics 1985; 12:171–178.
- 14. Wahed A, Ali M. The skew logistic distribution. *Journal of Statistical Research* 2001; **35**:71–80.
- 15. Wang J, Boyer J, Genton M. A skew symmetric representation of multivariate distributions. *Statistica Sinica* 2004; **14**: 1259–1270.
- 16. Wang J, Genton M. The multivariate skew-slash distribution. *Journal of Statistical Planning and Inference* 2006; **136**: 209–220.
- 17. Rubio FJ, Liseo B. On the independence Jeffreys prior for skew-symmetric models. *Statistics & Probability Letters* 2014; **85**:91–97.
- 18. Fernández C, Steel M. On the dangers of modelling through continuous distributions: a Bayesian perspective. In *Bayesian statistics 6*, Bernardo JM, Berger JO, Dawid AP, Smith AFM (eds). Oxford University Press: Oxford, 1999.
- 19. Liseo B, Loperfido N. A note on reference priors for the scalar skew-normal distribution. *Journal of Statistical Planning and Inference* 2006; **136**:373–389.
- 20. Bayes C, Branco M. Bayesian inference for the skewness parameter of the scalar skew-normal distribution. *Brazilian Journal of Probability and Statistics* 2007; **21**:141–163.
- 21. Branco M, Genton M, Liseo B. Objective Bayesian analysis of skew-t distributions. *Scandinavian Journal of Statistics* 2012; **40**:63–85.

# Statistics in Medicine

- 22. Cabras S, Racugno W, Castellanos M, Ventura L. A matching prior for the shape parameter of the skew-normal distribution. *Scandinavian Journal of Statistics* 2012; **39**(2):236–247.
- 23. Canale A, Scarpa B. Informative Bayesian inference for the skew-normal distribution. *arXiv preprint arXiv:1305.3080* 2013.
- Rubio F, Steel M. Bayesian modelling of skewness and kurtosis with two-piece scale and shape distributions. *Electronic Journal of Statistics* 2015; 9(2):1884–1912.
- 25. Fonseca T, Ferreira M, Migon H. Objective Bayesian analysis for the Student-*t* regression model. *Biometrika* 2008; **95**: 325–333.
- Juárez M, Steel M. Model-based clustering of non-Gaussian panel data based on skew-t distributions. Journal of Business and Economic Statistics 2010; 28:52–66.
- 27. Barros M, Paula G, Leiva V. A new class of survival regression models with heavy-tailed errors: robustness and diagnostics. *Lifetime Data Analysis* 2008; 14:316–332.
- Rubio F, Hong Y. Survival and lifetime data analysis with a flexible class of distributions. *Journal of Applied Statistics*. in press. DOI:10.1080/02664763.2015.1120710.
- Marchenko Y, Genton M. Multivariate log-skew-elliptical distributions with applications to precipitation data. *Environmetrics* 2010; 21:318–340.
- Azzalini A, Dal Cappello T, Kotz S. Log-skew-normal and log-skew-t distributions as models for family income data. Journal of Income Distribution 2003; 11:12–20.
- 31. Lin G, Stoyanov J. The logarithmic skew-normal distributions are moment indeterminate. *Journal of Applied Probability* 2009; **46**:909–916.
- 32. R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing: Vienna, Austria, 2015. http://www.R-project.org/.
- Christen JA, Fox C. A general purpose sampling algorithm for continuous distributions (the twalk). *Bayesian Analysis* 2010; 5:263–282.
- Dickey J. The weighted likelihood ratio, linear hypotheses on normal location parameters. *The Annals of Mathematical Statistics* 1971; 42:204–223.
- Arellano-Valle R, Castro L, Genton M, Gómez H. Bayesian inference for shape mixtures of skewed distributions, with application to regression analysis. *Bayesian Analysis* 2008; 3:513–540.
- Sinha D, Chen MH, Ghosh S. Bayesian analysis and model selection for interval-censored survival data. *Biometrics* 1999; 55:585–590.
- Loprinzi CL, Laurie JA, Wieand HS, Krook JE, Novotny PJ, Kugler JW, Bartel J, Law M, Bateman M, Klatt NE. et al.. Prospective evaluation of prognostic variables from patient-completed questionnaires. Journal of Clinical Oncology 1994; 23:601–607.
- 38. Chaloner K. Bayesian residual analysis in the presence of censoring. Biometrika 1991; 78:637-644.
- 39. Finley AO, Banerjee S, Carlin BP. SPBayes: an R package for univariate and multivariate hierarchical point-referenced spatial models. *Journal of Statistical Software* 2007; **19**:1–24.

# **Supporting information**

Additional supporting information may be found in the online version of this article at the publisher's web site.