

Flexible objective Bayesian linear regression with applications in survival analysis

Francisco J. Rubio^{a*} and Keming Yu^b

^a*London School of Hygiene & Tropical Medicine, WC1E 7HT, UK;*

^b*Department of Mathematical Sciences, Brunel University, Uxbridge UB8 3PH, UK.*

(v4.1 released December 2013)

We study objective Bayesian inference for linear regression models with residual errors distributed according to the class of two-piece scale mixtures of normal distributions. These models allow for capturing departures from the usual assumption of normality of the errors in terms of heavy tails, asymmetry, and certain types of heteroscedasticity. We propose a general noninformative, scale-invariant, prior structure and provide sufficient conditions for the propriety of the posterior distribution of the model parameters, which cover cases when the response variables are censored. These results allow us to apply the proposed models in the context of survival analysis. This paper represents an extension to the Bayesian framework of the models proposed in [19]. We present a simulation study that shows good frequentist properties of the posterior credible intervals as well as point estimators associated to the proposed priors. We illustrate the performance of these models with real data in the context of survival analysis of cancer patients.

Keywords: Accelerated failure time model; residual life; noninformative prior; predictive; two-piece distributions.

Classification codes: 62E15; 62N01; 62N02; 62P10

1. Introduction

The use of normal residual errors in linear regression models (LRMs) is perhaps the most common distributional assumption. However, the normality assumption can be inappropriate in practice given that the inference about the regression parameters is affected when the true distribution of the errors is asymmetric or heavy tailed. In order to overcome this shortcoming, alternative distributional assumptions have been proposed. We refer the reader to [19] for an extensive review of the different distributional assumptions which include, for instance, the family of scale mixtures of normals (SMN) [9, 23, 24], skew-elliptical and skew-symmetric distributions [1, 4, 18, 21], semiparametric approaches such as quantile regression [13, 16, 28], among others.

In a Bayesian framework, it is often of interest to employ noninformative priors; for instance, when the prior knowledge about the model parameters is vague. These kinds of priors are functions of the parameters, not necessarily integrable, that induce a well-defined posterior distribution with good frequentist properties. In this direction, [9] proposed an improper prior structure for LRMs with residual errors distributed according to the family of SMN. In the context of survival regression models, [23] studied the use of Jeffreys-type priors for accelerated failure time (AFT) models (which are LRMs for the logarithm of a set of survival times) with SMN errors. However, the use of noninformative priors in LRMs with flexible errors that allow for capturing skewness has received little attention. In this line, [17] proposed an improper prior structure for AFT models with errors distributed according to the generalised extreme value distribution. They provided a list of sufficient conditions for the propriety of the corresponding posterior

*Corresponding author. Email: Francisco.Rubio@lshtm.ac.uk

distribution which involves truncating the parameter space. Recently, [18] proposed a general noninformative prior structure for LRMs with skew-symmetric errors. They provided conditions for the propriety of the posterior distribution that cover cases where the response variables are censored.

In this paper, we study the use of the class of two-piece scale mixtures of normal (TPSMN) distributions for modelling the residual errors in LRMs from a Bayesian perspective. These sorts of distributional assumptions enjoy several advantages. First, this family of error models contains the class of SMN distributions as a particular case, which has been used to account for the presence of outliers and certain types of heteroscedasticity [24]. In addition, TPSMN distributions can also be used to capture unobserved heterogeneity that induces asymmetry of the residual errors [19]. The implementation of these models is straightforward using the R package ‘twopiece’ (available under request). We propose a general improper prior structure for the models of interest that covers certain priors obtained by formal rules. We show that the corresponding posterior is proper under mild conditions that can be extended to cases where the response variables are censored, a common phenomenon in survival analysis. The contribution of this paper consists mainly of extending the LRMs in [19], who only consider likelihood-based inference and prediction, to the Bayesian framework. The Bayesian approach provides natural tools, namely, the posterior predictive distribution, for conducting prediction about right-censored responses. This paper also presents a tractable alternative strategy to that proposed in [18] to flexibly modelling the errors in LRMs. The rest of the paper is organised as follows. In Section 2, we present the family of distributions of interest (TPSMN) and briefly discuss some of their properties. In Section 3, we describe the LRMs with TPSMN errors and the proposed prior structure, and then provide sufficient conditions for the propriety of the corresponding posterior distribution. In Section 4, we discuss the propriety of the posterior distribution in cases when the response variables are censored. We link these results with survival regression models. In Section 5, we present a simulation study which shows good frequentist performance of the proposed models. In Section 6 we present two examples with real data in the context of survival times of cancer patients. Proofs of the results as well as tables associated to the simulation study in Section 5 are presented in the Supplemental Material.

2. Background on two-piece distributions

Let us first recall the definition of two-piece distributions. We refer the reader to [20] and [19] for a more extensive discussion on these models. A real random variable Z is said to be distributed according to a two-piece distribution, denoted $Z \sim \text{TP}(\mu, \sigma, \delta, \gamma; f)$, if its probability density function (PDF) can be written as:

$$g(z|\mu, \sigma, \delta, \gamma) = \frac{2}{\sigma[a(\gamma) + b(\gamma)]} \left[f\left(\frac{z - \mu}{\sigma b(\gamma)} \middle| \delta\right) I(z < \mu) + f\left(\frac{z - \mu}{\sigma a(\gamma)} \middle| \delta\right) I(z \geq \mu) \right], \quad z \in \mathbb{R} \quad (1)$$

where f is a symmetric PDF with support on \mathbb{R} and mode at 0, $\mu \in \mathbb{R}$ is a location parameter and the mode of the density, $\sigma \in \mathbb{R}_+$ is a scale parameter, $\delta \in \Delta \subset \mathbb{R}$ is a shape parameter, $\gamma \in \Gamma \subset \mathbb{R}$ is a skewness parameter, and $\{a(\gamma), b(\gamma)\}$ are positive functions of the parameter γ . Several parameterisations $\{a(\cdot), b(\cdot)\}$ of these models are studied in [2] and [20]. In our applications we will adopt the parameterisation proposed in [15]: $\{a(\gamma), b(\gamma)\} = \{1 - \gamma, 1 + \gamma\}$, $\gamma \in (-1, 1)$. Some properties of this family of distributions are presented below.

- (1) The tail behaviour of (1) is the same in each direction.
- (2) The moments of (1) exist whenever the moments of the baseline PDF f exist.
- (3) The Fisher information matrix associated to this sort of models is well defined [20], in contrast to some skew-symmetric models, such as the Azzalini’s skew-normal distribu-

tion [3].

- (4) Despite the fact that the PDF (1) is not twice differentiable at the mode, [2] showed that the maximum likelihood estimators of the parameters of (1) have good asymptotic properties.

Throughout, we focus on the case where f belongs to the family of SMN. Recall also that a symmetric PDF f is said to be a SMN if it can be written as:

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

where H is a mixing distribution with positive support, ϕ represents the standard normal PDF, and $\delta \in \Delta \subset \mathbb{R}$ is a shape parameter. This family contains distributions of great interest in practice such as the normal distribution, Logistic distribution, Laplace distribution, generalised hyperbolic distribution, and the Student- t distribution.

From the expressions in [2] we can obtain the cumulative distribution (CDF) associated to (1) as follows:

$$G(z|\mu, \sigma, \delta, \gamma) = \frac{2b(\gamma)}{a(\gamma) + b(\gamma)} F\left(\frac{z - \mu}{\sigma b(\gamma)} \middle| \delta\right) I(z < \mu) + \frac{b(\gamma) - a(\gamma)}{a(\gamma) + b(\gamma)} + \frac{2a(\gamma)}{a(\gamma) + b(\gamma)} F\left(\frac{z - \mu}{\sigma a(\gamma)} \middle| \delta\right) I(z \geq \mu), \quad z \in \mathbb{R}, \tag{3}$$

where F is the CDF associated to the PDF f . From the latter expression we can see that $\mathbb{P}[Z \leq \mu] = P_\gamma = \frac{b(\gamma)}{a(\gamma) + b(\gamma)}$. That is, the parameter μ is the P_γ -th quantile of Z , and the parameter γ controls the allocation of mass on either side of the mode μ . The PDF, CDF, quantile function, and random number generation of two-piece distributions are implemented in the R package ‘twopiece’, which is available under request. Some examples of the shape of the density (1), for some choices of the baseline density f , are presented in the Supplemental Material.

3. Linear regression with two-piece errors

Consider the linear regression model:

$$y_j = \mathbf{x}_j^\top \boldsymbol{\beta} + \varepsilon_j, \tag{4}$$

where $y_j \in \mathbb{R}$, $j = 1, \dots, n$, $\boldsymbol{\beta}$ is a p -dimensional vector of regression parameters, $\varepsilon_j \stackrel{i.i.d.}{\sim} \text{TP}(0, \sigma, \delta, \gamma; f)$, f is a SMN, and $\mathbf{X} = (\mathbf{x}_1^\top, \dots, \mathbf{x}_n^\top)^\top$ is a known $n \times p$ design matrix of full column rank. The resulting model is centred at the mode of the distribution of the errors (which is 0), which represents the P_γ -th quantile. We can re-centre the model at the mean, provided it exists, or any quantile of interest by adjusting the intercept after obtaining estimators for the corresponding parameters. The likelihood function associated to these assumptions is given by:

$$s(\mathbf{y}|\boldsymbol{\beta}, \sigma, \delta, \gamma) = \prod_{j=1}^n s(y_j - \mathbf{x}_j^\top \boldsymbol{\beta} | 0, \sigma, \delta, \gamma), \tag{5}$$

where s is the PDF given by (1). We adopt the prior structure:

$$\pi(\boldsymbol{\beta}, \sigma, \delta, \gamma) \propto \frac{\pi(\gamma)\pi(\delta)}{\sigma^q}, \tag{6}$$

where $q \geq 0$ and $\pi(\gamma)$ and $\pi(\delta)$ are proper priors. This prior structure covers the structure of some priors obtained by formal rules, for specific choices of the power hyperparameter q and the priors $\pi(\gamma)$ and $\pi(\delta)$. For instance, for the choices $q = 1$ and $\pi(\gamma) \propto (1 - \gamma)^{-\frac{1}{2}}$, the prior (6) corresponds to the independence Jeffreys prior (see [20] for a study of this prior in the context of location-scale TPSMN). Expressions for the reference prior and the Jeffreys prior have not been calculated, but we conjecture that they are particular cases of the prior structure (6). Their calculation represents a possible research direction. The following result provides general conditions for the propriety of the posterior distribution under the prior (6).

THEOREM 1 Consider the model (4)–(6), where $\varepsilon_j \stackrel{i.i.d.}{\sim} TP(0, \sigma, \delta, \gamma; f)$ and f is a SMN. Consider the following conditions:

- (i) The posterior associated to the linear regression model (4), with errors distributed according to the symmetric baseline distribution f , together with the prior $\pi(\beta, \sigma, \delta) \propto \sigma^{-q} \pi(\delta)$ is proper.
- (ii) $\int_{\Gamma} \frac{h(\gamma)^{n+q-1}}{[a(\gamma) + b(\gamma)]^n} \pi(\gamma) d\gamma < \infty$, where $h(\gamma) = \min\{a(\gamma), b(\gamma)\}$,
- (iii) $\int_{\Gamma} \frac{H(\gamma)^{n+q-1}}{[a(\gamma) + b(\gamma)]^n} \pi(\gamma) d\gamma < \infty$, where $H(\gamma) = \max\{a(\gamma), b(\gamma)\}$.

Then, (i) and (ii) are necessary conditions, while (i) and (iii) are sufficient conditions for the propriety of the posterior distribution of $(\beta, \sigma, \delta, \gamma)$.

This result indicates that, in order to check the propriety of the LRMs with TPSMN errors, we only need to check the propriety of the posterior associated to the underlying model with residual errors distributed according to the symmetric baseline distribution f together with a condition on the parameterisation $\{a(\gamma), b(\gamma)\}$. In particular, for $q = 1$, conditions (ii) and (iii) are satisfied by any choice of $\{a(\gamma), b(\gamma)\}$. Moreover, if the functions $a(\cdot)$ and $b(\cdot)$ are bounded and $q > 1$, conditions (ii) and (iii) are automatically satisfied. The parameterisation $\{a(\gamma), b(\gamma)\} = \{1 - \gamma, 1 + \gamma\}$, proposed in [15], $\gamma \in (-1, 1)$, satisfies this boundedness condition. For unbounded parameterisations and $q > 1$ (such as the one proposed by [8]: $\{a(\gamma), b(\gamma)\} = \{\gamma, 1/\gamma\}$, $\gamma > 0$), the finiteness condition in (iii) depends on the choice of the prior $\pi(\gamma)$.

The following result presents conditions for the existence of the posterior for the case when the baseline density f belongs to the family of SMN and $q = 1$.

COROLLARY 1 Consider the model (4)–(6) with $q = 1$, where $\varepsilon_j \stackrel{i.i.d.}{\sim} TP(0, \sigma, \delta, \gamma; f)$ and f is a SMN. Then, the posterior distribution of (β, σ, γ) is proper provided that $\mathbf{y} \notin \mathcal{C}(\mathbf{X})$, where $\mathcal{C}(\mathbf{X})$ denotes the column space of \mathbf{X} , $n > p$, together with condition (iii) from Theorem 1.

This result is satisfied with probability one since the distribution of the residual errors is continuous. For the case when $q > 1$, conditions for the existence of the posterior associated to the model with symmetric errors become more restrictive. Next, we present some particular cases where the propriety of the posterior distribution can be easily checked.

COROLLARY 2 Consider the model (4)–(6) and suppose that the baseline density f in (1) is either a normal distribution, a Logistic distribution, a Laplace distribution, or a generalised hyperbolic distribution with fixed shape parameter. Suppose that $\mathbf{y} \notin \mathcal{C}(\mathbf{X})$, $n > p + 1 - q$, and condition (iii) from Theorem (1) are satisfied. Then, the posterior distribution (β, σ, γ) is proper.

Model (4)–(6) can be implemented by using the ‘twopiece’ R package. Moreover, several Markov Chain Monte Carlo (MCMC) samplers have been developed for this kind of models. For instance, [25] propose a blocked Metropolis-within-Gibbs algorithm that takes advantage of the representation of SMN distributions. For the case when f in (1) is a Laplace distribution,

[26] proposed an alternative MCMC algorithm based on a uniform mixture representation of the Laplace distribution.

We point out that Bayesian quantile LRMs [28] represent a family of models closely related to (4). These models can be interpreted as the LRM (4), with ϵ_j distributed according to a two-piece Laplace distribution, where the parameter γ is fixed according to the quantile of interest specified by the user. We emphasise that, in our context, we do not fix the parameter γ but, instead, obtain posterior inference about it using the prior structure (6).

Choice of the prior for γ

Rubio and Steel [20] proposed a prior elicitation strategy for the parameter γ , based on the interpretation of this parameter, that can be used to construct a weakly informative proper prior. They propose assigning a $\text{Beta}(a_0, b_0)$ prior on a measure of skewness which is an injective function of the parameter γ . This strategy induces a proper prior on γ which can be used to construct informative and noninformative priors on γ . For the case when $a_0 = b_0 = 1/2$, coupled with the parameterisation in [15], this strategy simply leads to the Jeffreys prior of γ . If $a_0 = b_0 = 1$, this strategy leads to a uniform prior on $\gamma \in (-1, 1)$. Throughout this paper we adopt this prior with $a_0 = b_0 = 1/2$, this is $\pi(\gamma) \propto (1 - \gamma^2)^{-\frac{1}{2}}$. This prior has been shown to induce a posterior distribution with good frequentist properties in the context of location-scale models [20].

4. Accelerated failure time models

4.1 Propriety results

AFT models are of great interest in survival analysis given that they can be used for modelling a set of survival times $\mathbf{T} = (T_1, \dots, T_n)$ in terms of a set of covariates β through the model equation:

$$y_j = \log(T_j) = \mathbf{x}_j^\top \beta + \epsilon_j, \quad j = 1, \dots, n, \tag{7}$$

where β is a p -dimensional vector of regression parameters, and $\mathbf{X} = (\mathbf{x}_1^\top, \dots, \mathbf{x}_n^\top)^\top$ is a known $n \times p$ design matrix of full column rank. The use of normal and Logistic residual errors represent the most common distributional assumptions. Other distributional assumptions were discussed Section 1 and in [19].

We assume that $\epsilon_j \stackrel{i.i.d.}{\sim} \text{TP}(0, \sigma, \delta, \gamma; f)$, where the baseline density f is a SMN. If we adopt the prior structure (6) for this model, then the corresponding posterior is proper under the conditions in Corollaries 1 and 2. However, a common challenge that arises in the context of the analysis of time-to-event data is the presence of censored observations (see [19] for a discussion on this). The following result provides sufficient conditions for the propriety of the posterior distribution for the case when the sample contains both censored and uncensored observations.

THEOREM 2 *Consider the linear regression model (7) with prior (6). Suppose that $n_c \leq n$ survival times are censored and $n_o = n - n_c$ are observed. Let \mathbf{y}_o be the set of uncensored observations and \mathbf{X}_o be the corresponding design matrix. Then, the posterior distribution of $(\beta, \sigma, \delta, \gamma)$ is proper provided that the posterior distribution associated to the n_o uncensored observations is proper.*

Since this result relies only on the sub-sample of uncensored observations, we can use the results in the previous section to check the propriety of the posterior. Corollary 1 provides conditions for the case when f is a SMN and $q = 1$, while Corollary 2 provides conditions for the

case when $q \geq 1$ and certain particular choices of f .

An extreme case that arises in practice is when the sample contains only censored observations. The next result presents sufficient conditions for the existence of the posterior distribution in this scenario.

COROLLARY 3 Consider the model (7) with prior (6). Suppose that f is a scale mixture of normals, $q = 1$, $n_I \leq n$ observations are interval censored, where the length of these intervals is finite, and that the other $n - n_I$ observations are censored of any other type. Denote the n_I interval-censored observations as (I_1, \dots, I_{n_I}) , and let \mathbf{X}_{n_I} be the corresponding design submatrix. Then, the corresponding posterior is proper if $\mathcal{E} = I_1 \times \dots \times I_{n_I}$ and the column space of \mathbf{X}_{n_I} are disjoint, together with the condition $n_I > p$, and condition (iii) from Theorem 1.

Similarly, for $q > 1$ we have the following results.

COROLLARY 4 Consider the model (7) with prior (6). Suppose that f is either a normal, Logistic, Laplace or generalised hyperbolic distribution; $n_I \leq n$ observations are interval censored, where the length of these intervals is finite, and that the other $n - n_I$ observations are censored of any other type. Then, the corresponding posterior is proper if $\mathcal{E} = I_1 \times \dots \times I_{n_I}$ and the column space of \mathbf{X}_{n_I} are disjoint, together with the condition $n_I > p + 1 - q$, and condition (iii) from Theorem 1.

Checking that \mathcal{E} and the column space of \mathbf{X}_{n_I} are disjoint can be formulated as a linear programming problem (LP). Denote $\boldsymbol{\eta} \in \mathbb{R}^p$, $\boldsymbol{\xi} = (\xi_1, \dots, \xi_{n_I}) \in \mathcal{E}$, and $I_j = [l_j, u_j]$, $j = 1, \dots, n_I$. Define the LP problem:

$$\begin{aligned} & \text{Find } \max_{\boldsymbol{\eta}, \boldsymbol{\xi}} 1, \\ & \text{Subject to } \mathbf{X}_{n_I} \boldsymbol{\eta} = \boldsymbol{\xi}, \\ & \text{and } \log(l_j) \leq \xi_j \leq \log(u_j), \quad j = 1, \dots, n_I. \end{aligned} \tag{8}$$

Thus, the disjointness condition is equivalent to verifying the infeasibility of the LP problem (8), for which there are several theoretical and numerical tools (LP solvers) available [7]. It is *important* to notice that the maximisation step in (8) (which is trivial) is just a tool to connect the propriety conditions in Corollaries 3 and 4 with the feasibility of the restrictions in a LP problem.

5. Simulation study

In this section we present a simulation study that illustrates the performance of the proposed prior structure. We adopt the simulation scenarios used in [19] in order to allow for qualitative comparisons. We study the LRM:

$$y_j = \beta_1 + \beta_2 x_{1j} + \beta_3 x_{2j} + \varepsilon_j, \quad j = 1, \dots, n, \tag{9}$$

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 two-piece normal distribution with unit scale parameter and skewness parameter $\gamma = 0, 0.25, 0.5, 0.75$, $(\beta_1, \beta_2, \beta_3) = (1, 2, 3)$, and $n = 100, 250, 500$. We fit the LRM (9) with $\varepsilon_j \stackrel{i.i.d.}{\sim} \text{TP}(0, \sigma, \gamma; f)$, where f is the standard normal PDF. We adopt the product prior structure (6) with $q = 1$ and the Jeffreys prior on the

parameter γ . For each of these scenarios, we obtain $N = 1,000$ samples of size 2,000 from the posterior distribution using the R t-walk sampler [6] after a burn-in period of 5,000 iterations and thinned to every 25th iteration (this is, a chain of length 55,000). Then, we calculate the proportion of 95% credible intervals that include the true value of the parameter, the median posterior estimators, the maximum *a posteriori* (MAP) estimators, the maximum likelihood estimators (MLEs) for comparison, as well as the median of the Bayes factors associated to the hypothesis $H_0 : \gamma = 0$ (approximated using the Savage-Dickey density ratio).

In the second and third scenarios, we simulate the residual errors from a two-piece Student-*t* distribution with degrees of freedom $\delta = 2, 5$. For these scenarios we fit the LRM (9) with $\varepsilon_j \stackrel{i.i.d.}{\sim} \text{TP}(0, \sigma, \delta, \gamma; f)$, where f is the Student-*t* PDF with $\delta > 0$ degrees of freedom. We adopt the prior structure (6) with $q = 1$ and the Jeffreys prior on the parameter γ . For the degrees of freedom δ , we use the approximation to the Jeffreys prior for this parameter proposed in [11]:

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

We choose the hyperparameter $d = 10$, which induces a prior with mode at $\delta = 5$. In the fourth scenario, we simulate from the linear regression model:

$$\log(y_j) = \mathbf{x}_j^\top \boldsymbol{\beta} + \varepsilon_j, \quad j = 1, \dots, n,$$

with $n = 100, 250, 500$, $\boldsymbol{\beta} = (1, 2, 3)^\top$, and $\mathbf{x}_j = (1, x_{j1}, x_{j2})^\top$. The second and third entries of the covariates \mathbf{x}_j are simulated from a right-half-normal with scale parameter 1/3. The errors ε_j are simulated from a two-piece normal distribution with parameters $\mu = 0$, $\sigma = 0.25$, and $\gamma = 0, 0.25, 0.5, 0.75$. We truncate the observations y_j that are greater than 17.5, producing samples with 15%–35% censored observations. Results are reported in Tables 1–12 of the Supplemental Material. In the first scenario we can observe a good coverage as well as good frequentist properties of the estimators associated to the proposed model overall. We can also observe that the Bayes factors clearly identify the case when the errors are symmetric. In the second and third scenarios we observe a good coverage of the credible intervals associated to the regression parameters $(\beta_1, \beta_2, \beta_3)$ as well as an accurate point estimation. However, in order to get a decent coverage of the credible intervals associated to the scale and tail parameters (σ, δ) , we need at least 250 observations. This is a well known phenomenon about the estimation of degrees of freedom of the Student-*t* distribution. Interestingly, the level of skewness does not seem to affect the performance of the credible intervals, even though for the case when $\gamma = 0$ we are fitting an overparameterised model. Although the proposed model performs well even when the true distribution of the residual errors is symmetric, in practice, we recommend conducting a formal model selection between the models with symmetric and asymmetric errors in order to avoid overparameterisation, which has a other unpleasant effects such as increasing the length of the credible intervals. The presence of mild levels of censored observations does not greatly affect the performance of the Bayes estimators and the coverage proportions as we can see from Tables 10–12 in the Supplemental Material.

6. Applications

In this section we present two examples with publicly available real data to illustrate the usefulness and performance of the proposed Bayesian LRMs. Both examples concern the study of the survival times of cancer patients. In the second example, we discuss the impact of using flexible errors in terms of prediction. For the models with two-piece residual errors, we employ the parameterisation in [15]. Posterior samples are obtained using the R `twalk` sampler [6]. Since the implementation of the log-posterior associated to the models of interest is very tractable, other samplers (such as Metropolis-Hastings or Metropolis-within-Gibbs samplers) can also be easily implemented. R codes and data used for these examples are available upon request.

Model comparison is conducted in terms of three formal model selection tools: Bayesian information criterion (BIC), Bayes factors, and log-predictive marginal likelihood (LPML, [22]). Bayes factors are calculated using an importance sampling technique. The use of the Bayes factors with the proposed improper prior structure is justified since we are employing improper priors only on the common parameters of the different models, while the priors on the shape parameters have the same interpretation across the different models (see [23] for a discussion on this point). Bayes factors and BIC are useful to identify the model that provides the best fit. On the other hand, LPML is a measure that ranks the models of interest in terms of their predictive performance [22]. Therefore, these two variables provide complementary information. Their combination is particularly relevant in survival analysis given that we are interested on selecting the best model for the data but, since this model is often used for prediction of the residual life of patients that survived beyond the end of the study (see [19]), it is important to check that the model also has a better predictive performance than the competitor models.

6.1 *Small Cell Cancer Data*

We analyse the data set from [27] about a lung cancer study with two different types of treatment. The data set contains $n = 121$ survival times (in days) of patients with small cell lung cancer (SCLC) that were administrated two types of therapies. For patients with SCLC the standard treatment consists of a combination of etoposide (E) and cisplatin (P); however the optimal order for the administration of these two treatments has not been established [27]. The group of patients was splitted into two groups: Arm A (62 patients), whose therapy consisted of P followed by E, and Arm B (59 patients), whose therapy consisted of E followed by P. The covariates used for this study are the “Entry age” (in years) and the type of treatment (Arm A and Arm B). The sample contains $n_c = 23$ right-censored observations. We fit an AFT model (7) with 4 distributional assumptions for the residual errors: two-piece Laplace (TP Laplace), two-piece Normal (TP Normal), as well as corresponding symmetric submodels (Laplace and Normal). The propriety of the corresponding posterior distributions is guaranteed by Theorem 1 and Corollary 2. We adopt the prior structure (6) with $q = 1$ and the Jeffreys prior on the skewness parameter $\gamma \in (-1, 1)$. For each of these models, a sample of size 10,000 was obtained from the posterior distribution after a burn-in period of 50,000 iterations and thinned to every 25 iterations (this is, 300,000 MCMC iterations in total). Table 1 presents a summary of the posterior samples as well as the model comparison tools. The TP Laplace model performs better overall (closely followed by the TP Normal model) in terms of BIC, LPML and Bayes factors, which suggests the need for a model with heavier tails than normal and asymmetry.

Model	TP Laplace	TP Normal	Laplace	Normal
Intercept	6.690 (5.964,7.445)	7.150 (6.282,8.150)	7.114 (6.168,8.294)	7.633 (6.560,8.653)
Entry age	-0.009 (-0.021,0.004)	-0.016 (-0.031,-0.002)	-0.011 (-0.029,0.004)	-0.017 (0.033, -0.0003)
Treatment	-0.446 (-0.682,-0.197)	-0.387 (-0.660,-0.107)	-0.403 (-0.683,-0.137)	-0.408 (-0.703,-0.139)
σ	0.650 (0.517,0.792)	0.785 (0.664,0.915)	0.648 (0.532,0.795)	0.759 (0.660,0.890)
γ	-0.395 (-0.599,-0.1872)	-0.383 (-0.639,-0.116)	–	–
BIC	283.09	284.75	292.49	287.59
Bayes factor	–	0.721	0.016	0.174
LPML	-134.893	-136.416	-141.0203	-138.1511

Table 1. SCLC Lung Cancer data: Posterior median and 95% credible intervals and model comparison tools. The Bayes factors are calculated against the model with TP Laplace errors.

6.2 North Central Cancer Treatment Group (NCCTG) Lung Cancer Data

In this application we analyse the NCCTG Lung Cancer data set, which is available in the ‘survival’ R package. The data set with complete cases (removing missing covariates) contains the survival times (in days) of $n = 227$ patients with advanced lung cancer from the NCCTG. The sample contains $n_c = 63$ right-censored observations. The aim of this study was to compare the information from a questionnaire applied to a group of patients against the information obtained by the patient’s physician in terms of prognostic power [14]. We fit an AFT model 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 with 4 residual error distributions: two-piece Logistic errors (TP Logistic), two-piece normal errors (TP Normal), and the corresponding symmetric sub-models (Logistic and normal). We adopt the prior structure (6) with $q = 1$ and the Jeffreys prior on γ . The propriety of the corresponding posterior distributions is guaranteed by Theorem 1 and Corollary 2. We obtain a posterior sample of size 10,000 after a burn-in of 50,000 and thinned to every 25th iteration (300,000 MCMC iterations in total). Table 2 shows a summary of the posterior samples as well as the model comparison tools. The model with TP Logistic errors performs better overall, which suggests the presence of skewness and slightly heavier tails than normal.

Model	TP Logistic	TP Normal	Logistic	Normal
Intercept	6.531 (5.514, 7.565)	6.940 (5.840,7.979)	5.965 (4.985, 6.962)	6.477 (5.309, 7.628)
Age	-0.010 (-0.025, 0.004)	-0.015 (-0.029, 0.001)	-0.008 (-0.023, 0.006)	-0.018 (-0.034, -0.002)
Sex	0.435 (0.188, 0.720)	0.446 (0.197,0.726)	0.496 (0.222, 0.761)	0.529 (0.231, 0.842)
ph.ecog	-0.363 (-0.533, -0.167)	-0.326 (0.507,-0.119)	-0.407 (-0.601, -0.221)	-0.359 (-0.571,-0.157)
σ	0.495 (0.429, 0.569)	0.906 (0.806,1.018)	0.548 (0.479,0.628)	1.043 (0.929, 1.170)
γ	0.384 (0.129, 0.600)	0.481 (0.270,0.669)	–	–
BIC	556.60	566.50	562.17	580.96
Bayes factor	–	0.006	0.019	2×10^{-6}
LPML	-268.966	-274.415	-272.91	-283.23

Table 2. NCCTG Lung Cancer data: Posterior median and 95% credible intervals and model comparison tools. The Bayes factors are calculated against the model with TP Logistic errors.

We now analyse the impact of using more flexible errors in terms of prediction. As discussed in [19] and [18], it is often of interest to study the distribution of the residual life of patients that survived beyond the end of the study. In order to obtain these predictions, consider the AFT model (7) with a general residual error distribution $\varepsilon_j \stackrel{i.i.d.}{\sim} TP(c(\sigma, \gamma), \sigma, \gamma; f)$, where $c(\sigma, \gamma)$ denotes the point at which the AFT model is centred (e.g. the mode, mean, or median). Denote by $\theta = (\beta, \sigma, \gamma)$ the model parameters, and let $\pi(\theta)$ be the corresponding prior. Suppose that the j th subject survived beyond time T_j , and therefore the corresponding observation is right-censored. Then, the posterior predictive CDF of the residual life for this subject is given by:

$$\Pi_R(t|\mathbf{T}, j) = \frac{\Pi(t|\mathbf{T}, j) - \Pi(T_j|\mathbf{T}, j)}{1 - \Pi(T_j|\mathbf{T}, j)}, \quad t > T_j, \tag{11}$$

where

$$\Pi(t|\mathbf{T}, j) = \int_0^t \pi(r|\mathbf{T}, j)dr,$$

is the posterior predictive CDF associated to this model, and

$$\pi(r|\mathbf{T}, j) = \int \frac{1}{r} s[\log r|\boldsymbol{\theta}, \mathbf{x}_j] \pi(\boldsymbol{\theta}|\mathbf{T}) d\boldsymbol{\theta}, \quad r > 0, \tag{12}$$

is the posterior predictive PDF associated to subject j , and $\pi(\boldsymbol{\theta}|\mathbf{T})$ represents the posterior distribution of $\boldsymbol{\theta}$. We recommend centring the model (after sampling from the posterior distribution) around the median rather than the mean since the latter may not exist for certain combinations of the distribution of the residual errors and priors. Moreover, median estimators are robust to the presence of outliers. The posterior predictive survival function of the residual life of subject j is given by $\mathcal{S}_R(t|\mathbf{T}, j) = 1 - \Pi_R(t|\mathbf{T}, j)$. This estimator takes into account the uncertainty on the model parameters given that they are integrated out with respect to the corresponding posterior distribution in (12). If we have a sample from the posterior distribution $\pi(\boldsymbol{\theta}|\mathbf{T})$, then we can approximate (11) by using a Monte Carlo approximation of (12). One advantage of the predictive estimator (11) over the plug-in estimator proposed in [19] is that this incorporates the posterior uncertainty about the model parameters.

Table 3 presents a summary of the quantiles of the residual life distributions, for the first 5 censored patients, using the AFT models with TP Logistic and Logistic errors centred at the median. As we discussed before, the model with TP Logistic errors produces a better fit and a better predictive performance. The quantiles higher than 50% associated to the model with Logistic errors are much larger than those obtained with the model with TP Logistic errors. Therefore, although the inference on the regression parameters is very similar for these two models (see Table 2), the corresponding prediction intervals are very different. These differences can be explained using the expressions (11)–(12) which indicate the dependence of the predictions on the residual error distribution and the posterior distribution (which are different in this case).

Quantile	5%	25%	50%	75%	95%
TP Logistic model					
Patient 1	1037.6	1163.9	1382.0	1774.0	2976.8
Patient 2	1040.2	1126.7	1287.3	1605.0	2641.4
Patient 3	992.4	1117.3	1332.4	1719.3	2886.6
Patient 4	868.8	1078.9	1406.8	1927.5	3358.7
Patient 5	866.5	986.3	1187.5	1545.6	2604.5
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

Table 3. NCCTG Lung Cancer data: Quantiles of the predictive residual life distribution for the Median TP Logistic and Logistic models.

7. Discussion

We introduced a flexible class of LRM that can account for departures from normality of the residual errors in terms of heavy tails, asymmetry and certain kinds of heteroscedasticity. We proposed a general noninformative prior structure and provided easy to check conditions for the propriety of the corresponding posterior distribution. A simulation study suggests a good frequentist performance of the proposed Bayesian models. The propriety results cover cases when

the response variables are censored, which allows for implementing the proposed Bayesian models in survival analysis. The implementation of these models is tractable using already available R packages. For instance, the R package ‘twopiece’ provides commands for the implementation of the PDF and CDF associated to TPSMN distributions.

In the real data applications, we have compared the proposed models against appropriate competitors using different sorts of model comparison tools. In the context of survival analysis, we advocated for the use of model selection tools that provide information about the model that better fits the data, as well as tools that provide information about the predictive performance of the models. This is particularly important in cases when the selected model is used for predicting the remaining life of a censored subject.

References

- [1] Arellano-Valle, R. B., Castro, L. M., Genton, M. G. and Gómez, H. W. (2008). Bayesian inference for shape mixtures of skewed distributions, with application to regression analysis. *Bayesian Analysis* 3: 513–540.
- [2] Arellano-Valle, R. B., Gómez, H. W. and Quintana, F. A. (2005). Statistical inference for a general class of asymmetric distributions. *Journal of Statistical Planning and Inference* 128: 427–443.
- [3] Azzalini, A. (1985). A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics* 12: 171–178.
- [4] Azzalini, A. and Genton, M. G. (2008). Robust likelihood methods based on the skew- t and related distributions. *International Statistical Review* 76: 106–129.
- [5] Choi, H. M. and Hobert, J. P. (2013). Analysis of MCMC algorithms for Bayesian linear regression with Laplace errors. *Journal of Multivariate Analysis* 117: 32–40.
- [6] Christen, J. A. and Fox, C. (2010). A general purpose sampling algorithm for continuous distributions (the t -walk). *Bayesian Analysis* 5: 263–282.
- [7] Dantzig, G. B., and Thapa, M. N. (1997). *Linear programming 1: Introduction*. Springer-Verlag, New York.
- [8] Fernández, C. and Steel, M. F. J. (1998). On Bayesian modeling of fat tails and skewness. *Journal of the American Statistical Association* 93: 359–371.
- [9] Fernández, C. and Steel, M. F. J. (2000). Bayesian Regression Analysis with Scale Mixtures of Normals. *Econometric Theory* 16: 80–101.
- [10] Geweke, J. (1993). Bayesian treatment of the independent Student- t model. *Journal of Applied Econometrics* 8: 19–40.
- [11] Juárez, M. A. and Steel, M. F. J. (2010). Model-based clustering of non-Gaussian panel data based on skew- t distributions. *Journal of Business and Economic Statistics* 28: 52–66.
- [12] Jung, Y. J. and Hobert, J. P. (2014). Spectral properties of MCMC algorithms for Bayesian linear regression with generalized hyperbolic errors. *Statistics and Probability Letters* 95: 92–100.
- [13] Lin, J., Sinha, D., Lipsitz, S., and Polpo, A. (2012). Semiparametric Bayesian Survival Analysis using Models with Log-linear Median. *Biometrics* 68: 1136–1145.
- [14] Loprinzi C. L., Laurie, J. A., Wieand, H. S., Krook, J. E., Novotny, P. J., Kugler, J. W., Bartel, J., Law, M., Bateman, M., Klatt, N. E. et al. (1994). Prospective evaluation of prognostic variables from patient-completed questionnaires. *Journal of Clinical Oncology* 12: 601–607.
- [15] Mudholkar, G. S., and Hutson, A. D. (2000). The epsilon-skew-normal distribution for analyzing near-normal data. *Journal of Statistical Planning and Inference* 83: 291–309.
- [16] Reich, B. J. and Smith, L. B. (2013). Bayesian Quantile Regression for Censored Data. *Biometrics* 69: 651–660.
- [17] Roy, V., and Dey, D. K. (2014). Propriety of posterior distributions arising in categorical and survival models under generalized extreme value distribution. *Statistica Sinica* 24: 699–722.
- [18] Rubio, F. J. and Genton, M. G. (2016). Bayesian linear regression with skew-symmetric error distributions with applications to survival analysis. *Statistics in Medicine*, in press.
- [19] Rubio, F. J. and Hong, Y. (2015). Survival and lifetime data analysis with a flexible class of distributions. *Journal of Applied Statistics*, in press.
- [20] Rubio, F. J. and Steel, M. F. J. (2014). Inference in Two-Piece Location-Scale models with Jeffreys Priors (with discussion). *Bayesian Analysis* 9: 1–22.
- [21] Sahu, S. J., Dey, D. K. and Branco, M. D. (2003). A new class of multivariate skew distributions with applications to Bayesian regression models. *Canadian Journal of Statistics* 2: 129–150.
- [22] Sinha, D., Chen, M. H. and Ghosh, S. K. (1999). Bayesian Analysis and Model Selection for Interval-Censored Survival Data. *Biometrics* 55: 585–590.
- [23] Vallejos, C. and Steel, M. F. J. (2015). Objective Bayesian survival analysis using shape mixtures of log-normal

- distributions. *Journal of the American Statistical Association* 110: 697–710.
- [24] West, M. (1984). Outlier models and prior distributions in Bayesian linear regression. *Journal of the Royal Statistical Society B* 46: 431–439.
- [25] Wichitaksorn, N., Choy, S. T. B. and Gerlach, R. (2014). A generalized class of skew distributions and associated robust quantile regression models. *Canadian Journal of Statistics* 42: 579–596.
- [26] Wichitaksorn, N., Wang, J. J. J., Choy, S. T. B., and Gerlach, R. (2015). Analyzing return asymmetry and quantiles through stochastic volatility models using asymmetric Laplace error via uniform scale mixtures. *Applied Stochastic Models in Business and Industry* 31: 584–608.
- [27] Ying, Z., Jung, S. H., and Wei, L. J. (1995). Survival analysis with median regression models. *Journal of the American Statistical Association* 90: 178–184.
- [28] Yu, K. and Moyeed, R. A. (2001). Bayesian quantile regression. *Statistics and Probability Letters* 54: 437–447.