

RESEARCH ARTICLE

A tractable Bayesian joint model for longitudinal and survival data

Danilo Alvares*¹ | Francisco J. Rubio²

¹Department of Statistics, Pontificia
Universidad Católica de Chile, Macul, Chile

²Department of Statistical Science,
University College London, London, UK

Correspondence

*Danilo Alvares, Department of Statistics,
Pontificia Universidad Católica de Chile,
7820436, Chile. Email: dalvares@mat.uc.cl

Funding information

FONDECYT (Chile), Grant/Award Number:
11190018

Summary

We introduce a numerically tractable formulation of Bayesian joint models for longitudinal and survival data. The longitudinal process is modelled using generalised linear mixed models, while the survival process is modelled using a parametric general hazard structure. The two processes are linked by sharing fixed and random effects, separating the effects that play a role at the time scale from those that affect the hazard scale. This strategy allows for the inclusion of non-linear and time-dependent effects while avoiding the need for numerical integration, which facilitates the implementation of the proposed joint model. We explore the use of flexible parametric distributions for modelling the baseline hazard function which can capture the basic shapes of interest in practice. We discuss prior elicitation based on the interpretation of the parameters. We present an extensive simulation study, where we analyse the inferential properties of the proposed models, and illustrate the trade-off between flexibility, sample size, and censoring. We also apply our proposal to two real data applications in order to demonstrate the adaptability of our formulation both in univariate time-to-event data and in a competing risks framework. The methodology is implemented in `rstan`.

KEYWORDS:

Competing Risks, General Hazard Structure, Generalised Linear Mixed Models, Power Generalised Weibull

1 | INTRODUCTION

In medical statistics, it is common to come across scenarios where patients are followed-up for a period of time (typically, until death or a censoring time point), and some biomarkers, patient characteristics, or treatment information are recorded at different time points over this period. This produces a combination of longitudinal and survival information about each individual. Historically, both processes have been analysed separately. For example, modelling time-to-event data is typically done by using hazard-based regression models. These include the Cox Proportional Hazard (PH) model,¹ which assumes that the covariates have an effect at the hazard scale; Accelerated Failure Time (AFT) models,² which assume that the covariates have a direct effect on the survival time; Accelerated Hazard (AH) models, which assume that the effect of the covariates is only on the time scale of the hazard function; as well as other general hazard (GH) structures that generalise the PH, AFT, and AH assumptions.³ See Rubio et al⁴ for a general overview of such models. The longitudinal process is typically modelled using Generalised Linear Mixed Models (GLMMs), which allow for modelling repeated and correlated observations (see, e.g. McCulloch et al⁵ for a general overview). It has been shown that combining both the longitudinal and survival processes represents a powerful tool for incorporating the information in both processes. Joint modelling of longitudinal and survival processes has been extensively

discussed in recent literature. We refer the reader to literature^{6,7,8,9} for reviews on this sort of models. A common strategy in joint models consists of linking the survival and the longitudinal processes by means of including shared parameters on the models for the covariates. This allows for incorporating a number of statistical modelling tools already available in the literature, such as using flexible parametric models using splines for modelling the hazard or the cumulative hazard functions (see, Brilleman et al¹⁰ for a recent review on these methods), while the longitudinal process can be modelled using any techniques developed for GLMMs. Applications of joint models abound in a number of areas of medical statistics.^{11,12,13}

In this paper, we propose a numerically tractable and interpretable alternative formulation of joint models, where we allow the longitudinal process to be modelled using GLMMs, while the survival process is specified through a parametric general hazard structure. This formulation allows for a direct interpretation of the parameters, as they are formulated at the hazard scale, as well as a separation of the roles of the parameters that affect the time scale, from those that affect the hazard scale. Another appealing aspect of the proposed formulation of joint models is numerical tractability, as the evaluation of the hazard and cumulative hazard functions does not require numerical integration, allowing for a tractable implementation of the likelihood and posterior distribution functions. We discuss several choices for the baseline hazard that are able to capture a variety of shapes of the hazard function. We discuss prior elicitation, where the general idea is to use weakly information priors for shape parameters while, for regression parameters, we consider g-priors¹⁴ in order to ameliorate potential overfit of those variables modelled using splines. Thus, the proposed joint models can be coupled with a number of general-purpose MCMC samplers. We provide an implementation of these models in `rstan`¹⁵ and show a good performance of this sampler in our simulation study and applications. We provide an extensive simulation study that illustrates the performance of our joint specification as well as the trade-off between using flexible assumptions for modelling the baseline hazard and non-linear effects, with sample size and censoring. In addition, we use a data set on AIDS patients¹⁶ to illustrate our methodology in a standard joint model context. We also present another real data example, using the SANAD study,¹⁷ where the survival process contains competing risks, emphasising the flexibility of our formulation to be coupled with a variety of scenarios. The rest of the paper is organised as follows. In Section 2, we present the formulation of the joint model and discuss the interpretation of the parameters. In Section 3, we present the likelihood function in a general framework, and discuss prior elicitation for the case where the longitudinal process is modelled using a linear mixed model (LMM), which is the model used later in the real data applications. In Section 4, we discuss an extensive simulation study and indicate how to simulate from the proposed joint model. In Sections 5 and 6, we illustrate the proposed methodology with two real data applications in the contexts of univariate time-to-event and competing risks, respectively. Finally, in Section 7, we present a brief discussion of the proposal in this paper and conclude with some practical advice and potential directions for further research. Additional results, including summaries from the simulation study, alternative models in the application, as well as technical details are presented in the Supplementary Material. R code is available at: www.github.com/daniloalvares.

2 | THE JOINT MODEL

2.1 | Longitudinal model: generalised linear mixed model

The longitudinal component of the proposed joint model is specified through a GLMM.⁵ Let $y_{ij} = y_i(t_j)$ be the response variables associated to the i th individual, $i = 1, \dots, n$, measured at time t_j , $j = 1, \dots, n_i$. Let $\mathbf{x}_i \in \mathbb{R}^p$ be a vector of individual covariates corresponding to the i th individual. Define the conditional distribution of y_{ij} given Ψ_{1i} (parameters and random effects) as a member of the Exponential family:

$$\begin{aligned} y_{ij} | \Psi_{1i} &\stackrel{ind.}{\sim} f_L(y_{ij} | \Psi_{1i}), \\ f_L(y_{ij} | \Psi_{1i}) &= \exp \left\{ \frac{y_{ij} \xi_{ij} - \varphi(\xi_{ij})}{\tau^2} - c(y_{ij}, \tau) \right\}, \\ E[y_{ij} | \Psi_{1i}] &= \mu_{ij}, \\ g(\mu_{ij}) &= \tilde{\beta}_0 + \mathbf{s}_i^\top \boldsymbol{\beta} + \tilde{\mathbf{x}}_i^\top \boldsymbol{\gamma} P_1(t_{ij}) + b_{0i} + (\tilde{\beta}_1 + b_{1i}) P_2(t_{ij}), \end{aligned} \quad (1)$$

where φ is a known function, the conditional mean of y_{ij} given Ψ_{1i} is related to ξ_{ij} via the identity $\mu_{ij} = \frac{\partial \varphi(\xi_{ij})}{\partial \xi_{ij}}$, the conditional variance of y_{ij} given Ψ_{1i} is $\tau^2 \frac{\partial^2 \varphi(\xi_{ij})}{\partial \xi_{ij}^2}$, and g is the link function. Regarding the model on the mean μ_{ij} , $\tilde{\beta}_0$ is the intercept, $\tilde{\beta}_1$

is the time-dependent slope; $\mathbf{s}_i = (s_{i1}^\top, \dots, s_{ip}^\top)^\top \in \mathbb{R}^{pq}$, where $q = \sum_{i=1}^p q_i$, q_m , $m = 1, \dots, p$, is the dimension of s_{im} , and s_{im} is a spline expansion of x_{im} (for continuous variables, or simply x_{im} for categorical variables or variables with a linear effect); $\boldsymbol{\beta} \in \mathbb{R}^{pq}$ are the corresponding regression coefficients; $\tilde{\mathbf{x}}_i \subseteq \mathbf{x}_i \in \mathbb{R}^{\tilde{p}}$, $\tilde{p} \leq p$, is a vector of individual time-dependent covariates, and $\boldsymbol{\gamma} \in \mathbb{R}^{\tilde{p}}$ are the corresponding regression coefficients; b_{0i} and b_{1i} are the random effects, which represent a random intercept and a random slope. This formulation thus allows for the inclusion of linear and non-linear effects by using a spline expansion of the corresponding covariates. $P_1(t_{ij})$ and $P_2(t_{ij})$ represent polynomial expansions, which indicate the functional dependence of time of the time-dependent covariates $\tilde{\mathbf{x}}_i$, the slope $\tilde{\boldsymbol{\beta}}_1$, and the random slopes b_{1i} . These can be, for instance, a B -spline basis polynomial expansion, or simply the identity function.⁶ In practice, it is often assumed a linear relationship, unless the individual trajectories are suspected to be non-linear. We assume that the random effects, $\mathbf{b}_i = (b_{0i}, b_{1i})^\top$, given Σ , follow a joint bivariate normal distribution with zero mean and variance-covariance matrix Σ . This family of mixed models include linear mixed models, Poisson mixed models, Negative Binomial mixed models, binary mixed models, among others.

2.2 | Survival model: general hazard structure

In this section, we discuss the model for the survival process, in which we adopt a general hazard (GH) structure.^{3,4} Let $h_0(\cdot | \boldsymbol{\theta})$ be a parametric baseline hazard function, with parameter $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^d$. Define the hazard function:

$$h(t | \boldsymbol{\Psi}_{2i}) = h_0 \left(t \exp \left\{ \mathbf{w}_i^\top \boldsymbol{\kappa} + \alpha_1 (\tilde{\mathbf{x}}_i^\top \boldsymbol{\gamma} + b_{1i}) \right\} \mid \boldsymbol{\theta} \right) \exp \left\{ \tilde{\mathbf{w}}_i^\top \tilde{\boldsymbol{\kappa}} + \mathbf{s}_i^\top \boldsymbol{\lambda} + \alpha_0 b_{0i} \right\}, \quad (2)$$

where $t > 0$ represents the time; $\boldsymbol{\Psi}_{2i}^\top = (\boldsymbol{\theta}^\top, \boldsymbol{\gamma}^\top, \boldsymbol{\kappa}^\top, \tilde{\boldsymbol{\kappa}}^\top, \boldsymbol{\lambda}^\top, \mathbf{b}_i^\top, \alpha_0, \alpha_1)$ denotes the full vector of model parameters; \mathbf{w}_i and $\tilde{\mathbf{w}}_i$ are r - and \tilde{r} -dimensional vectors of additional covariates, affecting the time-scale and the hazard-scale, respectively, which may only be available for the survival process (*i.e.* this formulation allows for the inclusion of different variables in the longitudinal and survival models); $\boldsymbol{\kappa}$ and $\tilde{\boldsymbol{\kappa}}$ are the regression coefficients associated to \mathbf{w}_i and $\tilde{\mathbf{w}}_i$, respectively; $\boldsymbol{\lambda}$ are regression coefficients for the a expansion \mathbf{s}_i ; $\alpha_0 \in \mathbb{R}$ and $\alpha_1 \in \mathbb{R}$ are the parameters linking the longitudinal and the survival processes, often called *association parameters*.

The hazard structure (2) separates the roles of the time-dependent effects (which appear in the argument of the baseline hazard and, consequently, affect directly the time scale) from those effects on the hazard scale, which appear multiplying the baseline hazard.⁴ This is, the hazard structure (2) can be used to account for time-dependent effects as well as effects that either increase or decrease the hazard level while the link with the longitudinal process also explicitly separates these roles.⁴ We see this as an advantage of this formulation as it helps to identify the need for connecting the two processes via time-dependent and/or proportional hazard effects. This hazard model can be directly used in more complex scenarios such as competing risks models, which we illustrate in our real data application. Another appealing feature of this hazard structure is that the corresponding cumulative hazard can be written in closed-form, thus avoiding the need for numerical integration, as:

$$H(t | \boldsymbol{\Psi}_{2i}) = H_0 \left(t \exp \left\{ \mathbf{w}_i^\top \boldsymbol{\kappa} + \alpha_1 (\tilde{\mathbf{x}}_i^\top \boldsymbol{\gamma} + b_{1i}) \right\} \mid \boldsymbol{\theta} \right) \exp \left\{ \tilde{\mathbf{w}}_i^\top \tilde{\boldsymbol{\kappa}} + \mathbf{s}_i^\top \boldsymbol{\lambda} + \alpha_0 b_{0i} - [\mathbf{w}_i^\top \boldsymbol{\kappa} + \alpha_1 (\tilde{\mathbf{x}}_i^\top \boldsymbol{\gamma} + b_{1i})] \right\}, \quad (3)$$

where $H_0(\cdot | \boldsymbol{\theta})$ is the cumulative baseline hazard of $h_0(\cdot | \boldsymbol{\theta})$. This allows for a tractable implementation of the likelihood and posterior distributions, which in turns allows this joint model formulation to be coupled with any general-purpose MCMC sampler.

This model specification is identifiable provided that the baseline hazard is not the hazard associated to a Weibull distribution,^{3,4} since in this case the AFT, PH, and AH models coincide and it becomes impossible to distinguish the effects in the argument of the baseline hazard from those multiplying the baseline hazard. We do not consider this to be a big limitation as similar hazard shapes can be obtained with other distributions, as discussed next, and model selection tools can be used to identify simpler hazard structures.

A natural extension of the hazard structure (2) consists of using more than one link parameter α_1 , for instance, in cases where the vector $\tilde{\mathbf{x}}_i$ contains variables in very different scales. However, we argue that, in practice, this is not often required, and estimating more than one scaling parameter would require larger samples. This logic is also in line with the classical formulation of joint models,⁶ where only one link parameter is used. We also note that, under the formulation in (2), we are connecting the survival and longitudinal processes through the random effects and the time-dependent effects. It is, of course, possible to link them through the time-invariant effects $\mathbf{s}_i^\top \boldsymbol{\beta}$ as follows,

$$h(t | \boldsymbol{\Psi}_{2i}) = h_0 \left(t \exp \left\{ \mathbf{w}_i^\top \boldsymbol{\kappa} + \alpha_1 (\tilde{\mathbf{x}}_i^\top \boldsymbol{\gamma} + b_{1i}) \right\} \mid \boldsymbol{\theta} \right) \exp \left\{ \tilde{\mathbf{w}}_i^\top \tilde{\boldsymbol{\kappa}} + \alpha_0 (b_{0i} + \mathbf{s}_i^\top \boldsymbol{\beta}) \right\},$$

which reduces the number of parameters by omitting λ . However, one limitation of this approach is that it assumes that the scaling factor is the same for all covariates, which may be in different scales or may contain a combination of categorical and continuous variables. Nonetheless, it remains as an alternative formulation for modelling the survival process.

Baseline hazard function

The choice of the parametric baseline hazard function is crucial as this determines the hazard shapes the survival model (2) can capture. For instance, the log-normal hazard function is unimodal (up-then-down), while the Gamma hazard function can be increasing or decreasing. There exist other (three-parameter) distributions that can capture the basic shapes of the hazard (increasing, decreasing, unimodal, and bathtub), such as the Exponentiated Weibull, Generalised Gamma, and Power Generalised Weibull distributions. However, it is important to consider that an efficient estimation of the parameters of these distributions typically requires larger sample sizes, and that high censoring rates or early administrative censoring (short follow-up) may also be detrimental in estimating shape parameters (specially those that control the tail behaviour) of flexible parametric distributions.^{4,18} We consider four baseline hazard candidates: Log-normal, Gamma, Power Generalised Weibull (PGW), and Generalised Gamma (GG), based on their numerical tractability and flexibility. The PGW and GG distributions contain three parameters (a scale parameter, and two shape parameters). These distributions offer similar levels of tractability and flexibility¹⁹. Expressions for the PGW and GG probability density functions (pdf), survival functions, and hazard functions are presented in Sections A1-A2 in the Supplementary Material.

3 | BAYESIAN INFERENCE

3.1 | Likelihood function

The likelihood function of the full parameter vector and random effects of the joint model (1)–(2) is given by:

$$f(\text{Data} | \Psi) = \prod_{i=1}^n \prod_{j=1}^{n_i} f_L(y_{ij} | \Psi_{1i}) \prod_{i=1}^n f_S(t_i | \Psi_{2i}), \quad (4)$$

where $\Psi = (\Psi_1^\top, \Psi_2^\top)^\top = (\beta^\top, \tilde{\beta}^\top, \gamma^\top, \tau, \mathbf{b}_1^\top, \dots, \mathbf{b}_n^\top, \theta^\top, \kappa^\top, \tilde{\kappa}^\top, \lambda^\top, \alpha_0, \alpha_1)^\top$ denotes the full parameter vector and random effects; $f_L(y_{ij} | \Psi_{1i})$ denotes the conditional pdf of y_{ij} given Ψ_{1i} described in (1); and $f_S(t_i | \Psi_{2i})$ is the contribution of the i th time-to-event to the likelihood function. For example, for linear mixed models, which we describe in the next section, $f_L(y_{ij} | \Psi_{1i})$ can be the normal density with mean μ_{ij} and variance σ^2 . The contribution of the survival time t_i is described by:

$$f_S(t_i | \Psi_{2i}) = \begin{cases} h(t_i | \Psi_{2i})S(t_i | \Psi_{2i}), & \text{if exact lifetime,} \\ S(t_i | \Psi_{2i}), & \text{if right-censored observation,} \\ 1 - S(t_i | \Psi_{2i}), & \text{if left-censored observation,} \\ S(t_{i,L} | \Psi_{2i}) - S(t_{i,R} | \Psi_{2i}), & \text{if interval-censored observation,} \end{cases}$$

where $h(t_i | \Psi_{2i})$ is the hazard function (2) and $S(t_i | \Psi_{2i}) = \exp\{-H(t_i | \Psi_{2i})\}$ represents the survival function derived from the cumulative hazard introduced in (3).

Although we will focus on Bayesian inference for the parameters, we point out that the marginal likelihood function of the parameters $\Psi^* = \Psi_{-\{\mathbf{b}_1, \dots, \mathbf{b}_n\}}$ can be written as follows:

$$f(\text{Data} | \Psi^*, \Sigma) = \prod_{i=1}^n \int \left[\prod_{j=1}^{n_i} f_L(y_{ij} | \Psi_{1i}) \right] f_S(t_i | \Psi_{2i}) \phi(\mathbf{b}_i | \Sigma) d\mathbf{b}_i,$$

where $\phi(\mathbf{b}_i | \Sigma)$ is the bivariate normal density with zero mean and variance-covariance matrix Σ . Thus, the evaluation of the marginal likelihood function, under our joint model formulation, only requires numerical integration with respect to the distribution of the random effects.

3.2 | Prior distributions

In this section, we provide general guidelines about prior choice for the parameters of the proposed joint models. We identify weakly informative priors for the parameters of this new class of joint models, but we acknowledge the possibility of using alternative priors.

We define the prior distributions for the more particular case of linear mixed models (LMMs), which are the main interest in our applications. However, these can be extended to GLMMs by adapting the priors on the variance of the response variable, accordingly.²⁰ Consider the LMM:

$$y_i(t) = \tilde{\beta}_0 + \mathbf{s}_i^\top \boldsymbol{\beta} + \tilde{\mathbf{x}}_i^\top \boldsymbol{\gamma} P_1(t_{ij}) + b_{0i} + (\tilde{\beta}_1 + b_{1i}) P_2(t_{ij}) + \epsilon_i(t).$$

The residual errors are assumed conditionally independent and identically distributed as $(\epsilon_i(t) \mid \sigma^2) \sim N(0, \sigma^2)$. In order to avoid concerns about the propriety of the posterior distribution,²¹ we adopt a proper prior specification. For the parameters $(\tilde{\beta}_0, \tilde{\beta}_1, \boldsymbol{\gamma}, \Sigma)$, we adopt weakly informative priors:

$$\begin{aligned} \tilde{\beta}_j &\sim N(0, \phi_{\tilde{\beta}_j}^2), \quad j = 0, 1, \\ \boldsymbol{\gamma} &\sim N_{\tilde{p}}(\mathbf{0}, \boldsymbol{\Phi}_{\boldsymbol{\gamma}}), \\ \sigma^2 &\sim \text{Inv-Gamma}(0.01, 0.01), \\ \sigma_j^2 &\sim \text{Inv-Gamma}(0.01, 0.01), \quad j = 1, 2, \\ \frac{\rho + 1}{2} &\sim \text{Beta}(a_0, b_0), \end{aligned}$$

where $\rho \in (-1, 1)$ is the correlation between b_{0i} and b_{1i} . The variance hyperparameters, $\phi_{\tilde{\beta}_j}^2$ and the diagonal of $\boldsymbol{\Phi}_{\boldsymbol{\gamma}}$, are assumed to be large in order to reflect vague prior information. Alternatively, one could choose a half-Cauchy prior for the variance parameters.²¹ For the parameters $\boldsymbol{\beta}$ and $\boldsymbol{\lambda}$, which represent regression coefficients associated to covariates that may contain spline expansions, we consider the following prior specification that penalises overfit:

$$\begin{aligned} \pi(\boldsymbol{\beta} \mid \sigma^2) &= \prod_{I_O} N(\boldsymbol{\beta}_j; 0, \phi_{\boldsymbol{\beta}_j}^2) \prod_{I_S} N(\boldsymbol{\beta}_j; 0, g_\beta \mathbf{M}_j \sigma^2), \\ \pi(\boldsymbol{\lambda} \mid \eta^2) &= \prod_{I_O} N(\boldsymbol{\lambda}_j; 0, \phi_{\boldsymbol{\lambda}_j}^2) \prod_{I_S} N(\boldsymbol{\lambda}_j; 0, g_\lambda \mathbf{M}_j \eta^2), \end{aligned}$$

where $I_O = \{j : \mathbf{x}_{ij} = \mathbf{s}_{ij}, \text{ for all } i\}$ and $I_S = \{j : \mathbf{x}_{ij} \neq \mathbf{s}_{ij}, \text{ for all } i\}$ indicate the indexes of the variables expressed in the original scale and in a spline basis expansion, respectively, $\mathbf{M}_j = (\tilde{\mathcal{S}}_j^\top \tilde{\mathcal{S}}_j)^{-1}$, and $\tilde{\mathcal{S}}_j$ are the design matrices associated to spline basis expansions of the covariates \mathbf{x}_{ij} . This is, if the vector \mathbf{s}_i contains spline expansions of \mathbf{x}_i , we adopt g-priors;¹⁴ otherwise, we adopt weakly informative priors where the variance hyperparameters $\phi_{\boldsymbol{\beta}_j}^2$ and $\phi_{\boldsymbol{\lambda}_j}^2$ are assumed to be large. The hyperparameters $g_\beta = n/q$ and $g_\lambda = n/q$,¹⁸ assuming that all spline basis expansions have the same degree q (which can be easily relaxed, if necessary), induce a mild penalty that shrinks the parameters towards zero and help prevent overfitting (see¹⁸ for a discussion on this point and other choices of these hyperparameters). We keep a relatively simple prior choice in this paper, but we acknowledge the possibility of including other shrinkage priors that carry heavier penalties on model complexity.^{22,18} We emphasise that those priors can also be included in our approach and numerical implementation.

For the parameter $\boldsymbol{\theta}$ in baseline hazards, we consider the following priors:

- (i) *Log-normal* (LN). For the scale parameter, say $\eta > 0$, we adopt a weakly informative prior $\eta \sim \text{half-Cauchy}(0, s_\eta)$, in the sense that it is a heavy tailed prior.²¹ For the log-location parameter, we assume $\mu \sim N(0, \phi_\mu^2)$, where the hyperparameter ϕ_μ^2 is assumed to be large.
- (ii) *Gamma*. For the scale and shape parameters, we adopt a weakly informative prior $\eta \sim \text{half-Cauchy}(0, s_\eta)$ and $\nu \sim \text{half-Cauchy}(0, s_\nu)$, respectively.
- (iii) *Power Generalised Weibull* (PGW). For the scale, shape, and power parameters, (η, ν, δ) , we adopt weakly informative priors specified as $\eta \sim \text{half-Cauchy}(0, s_\eta)$, $\nu \sim \text{half-Cauchy}(0, s_\nu)$, and $\delta \sim \text{Gamma}(1.83, 0.65)$. The prior on the parameter δ represents an approximation to the weakly informative prior *BTV*(1, 1) obtained with the method proposed in Dette et al²³. A full description of this prior is presented in Section A2 in the Supplementary Material.
- (iv) *Generalised Gamma* (GG). Similar to PGW specification, $\eta \sim \text{half-Cauchy}(0, s_\eta)$, $\nu \sim \text{half-Cauchy}(0, s_\nu)$, and $\delta \sim \text{Gamma}(1.83, 0.65)$.

For the association parameters, we adopt the weakly informative priors $\alpha_k \sim N(0, \phi_{\alpha_k}^2)$, $k = 0, 1$, where the variance hyperparameters are assumed to be large. We point out that shrinkage priors²⁴ could also be considered for these parameters, which might be useful to enforce parsimony in the link between the longitudinal and survival processes.

4 | SIMULATIONS

4.1 | Simulating from the joint model

We now describe a method to simulate from the proposed joint model (1)–(2). This method basically requires simulating the random effects first, followed by the simulation of the survival process using the GH structure (2), and finally simulating the longitudinal process (1). The steps for the simulation procedure are described in Algorithm 1. We can see that simulating from the survival model GH is relatively simple, in contrast to other joint models discussed in the literature, provided that one can simulate times-to-event from the baseline model. Thus, the choice of tractable baseline hazards model also facilitates the simulation of the joint model. Regarding the distribution of the distance between repeated observations (DDBRO), we have several scenarios of practical importance. For instance, in medical scenarios with periodic consultations, the DDBRO would be equidistant. In more complex scenarios, this distance might be random, for instance visits to the hospital due to some treatment or illness complication, which are also recorded and monitored; or even a combination of periodic and random visits to the hospital. Our formulation allows for the inclusion of all of these types of DDBRO. Censored survival times can be induced in the standard way, by either inducing administrative censoring or simulating random censoring points.

Algorithm 1 Simulation from the proposed joint model

For each individual $i = 1, \dots, n$, and for given values of the parameters and the design matrix:

Random Effects. Simulate $\mathbf{b}_i \sim N_2(0, \Sigma)$.

Survival Process. In order to simplify notation, let us denote:

$$A = \exp \left\{ \mathbf{w}_i^\top \boldsymbol{\kappa} + \alpha_1 \left(\tilde{\mathbf{x}}_i^\top \boldsymbol{\gamma} + b_{1i} \right) \right\},$$

$$B = \exp \left\{ \tilde{\mathbf{w}}_i^\top \tilde{\boldsymbol{\kappa}} + \mathbf{s}_i^\top \boldsymbol{\lambda} + \alpha_0 b_{0i} - \left[\mathbf{w}_i^\top \boldsymbol{\kappa} + \alpha_1 \left(\tilde{\mathbf{x}}_i^\top \boldsymbol{\gamma} + b_{1i} \right) \right] \right\}.$$

The individual survival function is $S(t | \boldsymbol{\Psi}_{2i}) = \exp[-H(t | \boldsymbol{\Psi}_{2i})]$, we can apply the probability integral transform directly to obtain:

$$t_i = \frac{F_0^{-1} \left[1 - \exp \left\{ \frac{\log(1 - u_i)}{B} \right\} \mid \boldsymbol{\theta} \right]}{A},$$

where F_0 is the cumulative distribution function associated to the baseline hazard h_0 , and $u_i \sim U(0, 1)$.

Longitudinal Process. Once a simulated time-to-event t_i is obtained from the previous step, specify the distribution of the distance between the repeated observations (e.g. equidistant or random). This produces the time points t_{ij} , $j = 1, \dots, n_i$, at which the repeated observations are recorded. The longitudinal process simulation is thus obtained by plugging-in the corresponding values of the parameters and covariates in μ_{ij} , and simulating from the corresponding GLMM based on (1).

4.2 | Simulation study

In this section, we conduct an extensive simulation study where we present the performance of the proposed joint model and estimation methods. More specifically, we illustrate the parameter estimation, ability to recover the baseline hazard shapes, as well as the effect of sample size and censoring rates on inference. For the survival process, we consider a scenario where the available variables are age at diagnosis, sex, and the presence of comorbidities. This setting is common in population studies in cancer epidemiology.²⁵ We analyse the following simulation scenarios, in increasing order of complexity.

Scenario 1. The longitudinal model:

$$y_i(t) = \tilde{\beta}_0 + \tilde{\beta}_1 t + \beta_1 \text{sex}_i + \beta_2 \text{age}_i + b_{0i} + b_{1i} t + c_i(t).$$

The survival process:

$$h(t \mid \Psi_{2i}) = h_0 \left(t \exp \{ \alpha_1 b_{1i} \} \mid \theta \right) \exp \{ \tilde{\kappa}_1 \text{comorb}_i + \lambda_1 \text{sex}_i + \lambda_2 \text{age}_i + \alpha_0 b_{0i} \}.$$

In addition, we consider **Scenario 0** where we simulate the model described in **Scenario 1** but we fit the following joint model. The longitudinal model:

$$y_i(t) = \tilde{\beta}_0 + \tilde{\beta}_1 t + \beta_1 \text{sex}_i + \beta_2 \text{age}_i + b_{0i} + b_{1i} t + \epsilon_i(t).$$

The survival process:

$$h(t \mid \Psi_{2i}) = h_0(t \mid \theta) \exp \{ \tilde{\kappa}_1 \text{comorb}_i + \lambda_1 \text{sex}_i + \lambda_2 \text{age}_i + \alpha_0 b_{0i} \}.$$

Scenario 2. The longitudinal model:

$$y_i(t) = \tilde{\beta}_0 + \tilde{\beta}_1 t + \gamma_1 \{ t \text{age}_i \} + \beta_1 \text{sex}_i + \beta_2 \text{age}_i + b_{0i} + b_{1i} t + \epsilon_i(t).$$

The survival process:

$$h(t \mid \Psi_{2i}) = h_0 \left(t \exp \{ \alpha_1 (\gamma_1 \text{age}_i + b_{1i}) \} \mid \theta \right) \exp \{ \tilde{\kappa}_1 \text{comorb}_i + \lambda_1 \text{sex}_i + \lambda_2 \text{age}_i + \alpha_0 b_{0i} \}.$$

Scenario 3. The longitudinal model:

$$y_i(t) = \tilde{\beta}_0 + \tilde{\beta}_1 t + \gamma_1 \{ t \text{age}_i \} + \beta_1 \text{sex}_i + g_1(\text{age}_i) + b_{0i} + b_{1i} t + \epsilon_i(t).$$

The survival process:

$$h(t \mid \Psi_{2i}) = h_0 \left(t \exp \{ \alpha_1 (\gamma_1 \text{age}_i + b_{1i}) \} \mid \theta \right) \exp \{ \tilde{\kappa}_1 \text{comorb}_i + \lambda_1 \text{sex}_i + g_2(\text{age}_i) + \alpha_0 b_{0i} \},$$

where g_1 and g_2 are B-spline expansions of the variable age_i of degree $q = 3$.

Thus, in Scenarios 2-3, the variable ‘‘age’’ represents a time-dependent effect both on the longitudinal and the survival process. This can be seen as it appears multiplying t in the longitudinal process, while it affects the time-scale directly on the survival process. In Scenario 3, the variable ‘‘age’’ also includes non-linear effects in both processes. The variable ‘‘age’’ is simulated from a mixture of uniform variables with probabilities 0.25 for the age group (30, 65), 0.35 for the age group (65, 75), and 0.4 for the age group (75, 85). This variable is then centered at 70 and scaled by a factor of 10. The variables ‘‘comorbidity’’ and ‘‘sex’’ are simulated from a Binomial distribution with probability parameter 0.5. The parameter values in each of these scenarios are presented in Section A3 in the Supplementary Material. These values are selected, together with the administrative censoring points, in order to obtain 5% and 35% censoring rates, which will allow us to assess the effect of censoring. We also consider two sample sizes in each scenario, $n = 200$ and $n = 500$, in order to evaluate the effect of sample size. The residual errors are assumed conditionally independent and identically distributed as $(\epsilon_i(t) \mid \sigma^2) \sim N(0, \sigma^2)$.

For each simulation scenario, we simulate $N = 100$ data sets and obtain 2,000 posterior samples of the parameters of the corresponding joint models using `rstan`. We apply a burn-in period of 1,000 iterations as well as a thinning period of 5 iterations to these posterior samples, for a total of 200 posterior samples. Under this configuration, we have observed convergence of the posterior samples. The number of Monte Carlo iterations is based on a trade-off between CPU time while trying to minimise the Monte Carlo error. Simulations were performed on an iMac with 3.3 GHz Quad-Core Intel Core i7, 16 GB RAM, macOS Catalina.

The results from this simulation study are presented in Sections A4 and A5 in the Supplementary Material. Tables A7–A38 display summaries of the posterior samples. These tables present the averages of the posterior means, posterior medians, 2.5% posterior quantiles, and 97.5% posterior quantiles for each of the posterior samples in each scenario. We observe a good performance of the Bayesian point estimators, as these means are close to the true values of the parameters. There is a clear effect of the sample size and censoring rates on the accuracy of the estimates. Unsurprisingly, the larger the sample, the more accurate the estimates. A similar conclusion is obtained for the censoring rates, as we can see that lower censoring rates produce more accurate estimates. Model complexity is an important factor (for instance, log-normal vs. PGW baseline, or linear model vs. splines model), as we can see that the more complex the model is, the wider the credible intervals are, which also interacts with the effects of sample size and censoring rates. A more interesting conclusion is observed for the length of follow-up, as reducing the length of follow-up has a marked effect on the accuracy of the estimates of the parameters of the survival model. In particular, the estimation of the shape parameters in the PGW distribution suffers from early administrative censoring as these parameters control the tail of the distribution, while early administrative censoring removes information about the tails of the distribution. This phenomenon remains even after increasing the sample size, indicating that a longer follow-up might

be as important as increasing the sample size if the aim is to improve the accuracy of the estimates. In addition, Figures A3–A22 in the Supplementary Material show the baseline hazards associated to the plugging posterior median estimators as well as the posterior predictive baseline hazards (which are defined as the ratio of the posterior predictive probability density function and the posterior predictive survival function). We observe a similar situation about the effect of the censoring rate, sample size, and model complexity on the ability to recover the shape of the baseline hazard. Another interesting result, obtained from comparing Scenarios 0 and 1, is that not including $\alpha_1 b_{1i}$ (i.e. not sharing the random slopes) induces a bias in the estimation of α_0 . However, in order to be able to estimate the parameter α_1 accurately, a longer follow-up or a larger sample size is necessary. Finally, comparing Scenarios 1 and 2, we can see that it is easier to estimate the link parameters when fixed and random time-dependent effects are combined (for 3-parameter baseline hazard), in contrast to the case when only random time-dependent effects are considered. This applies to all sample sizes and censoring rates. In this case, the use of an appropriate simpler model (in the sense that it can capture the true hazard shape) improves the estimation of the link parameters.

We conduct additional simulation studies (Section A6) where we assess the effect of higher censoring rates (60% censoring) and for binary longitudinal outcomes (which is modelled using a logistic mixed effects model). Results are shown in Tables A39–A58 in the Appendix. The conclusions are the same as those in the previous paragraph in terms of the interplay between sample size, censoring, and the flexibility of the baseline hazard and the functional form of the regression model. This section also illustrates the fact that results are the same for negative correlations and negative regression coefficients (as the role of the parameters remains the same). Finally, Section A6 also presents a simulation study using a Generalised Gamma baseline hazard instead of PGW. The results are comparable to those obtained in the PGW case, however, the need for using special functions for evaluating the GG hazard and cumulative hazard functions has a large cost in terms of computing times as the posterior sampling is slowed down by a factor of 7, compared to the PGW model, despite the efficient implementation of these functions discussed in the appendix. Thus, although equally flexible, the GG model also carries a larger computational cost, which is more apparent in the Bayesian framework where the evaluation of the posterior needs to be done thousands of times to obtain a posterior MCMC sample.

Overall, this simulation study provides some guidelines (and a warning) about the use of flexible models with many parameters when there are high censoring rates or small samples. Thus, in practice, it is recommended to compare simple models against more complex models using formal model selection tools. We will illustrate this idea in the following section.

5 | THE AIDS STUDY: UNIVARIATE TIME-TO-EVENT JOINT MODEL

The `aids` data consists of 467 patients with advanced human immunodeficiency virus infection during antiretroviral treatment who had failed or were intolerant to zidovudine therapy.¹⁶ This data set is publicly available in the R package `JM`²⁶ and the main objective is to identify associations between the time to death and the following covariates: `CD4`: CD4 cells count (longitudinal biomarker); `prevOI`: a factor with levels `AIDS` denoting previous opportunistic infection (AIDS diagnosis) at study entry, and `noAIDS` denoting no previous infection; `drug`: a factor with levels `ddC` denoting zalcitabine and `ddI` denoting didanosine; `gender`: a factor with levels `female` and `male`; `AZT`: a factor with levels `intolerance` and `failure` denoting AZT intolerance and AZT failure, respectively.

For simplicity, we used only the baseline variable `prevOI`, coded as 0 (`noAIDS`) and 1 (`AIDS`), to illustrate our methodological approach. In addition, the longitudinal CD4 variable was transformed by applying the square root. This selection is also based on a preliminary study on the importance of these variables on the longitudinal outcome of interest.

5.1 | Model 1 (M_1)

Our first proposal specifies the longitudinal model with fixed effects for the intercept ($\tilde{\beta}_0$), slope ($\tilde{\beta}_1$), and `prevOI` (β); random effects for the intercept (b_0) and slope (b_1); and a time-dependent effect for the `prevOI` covariate (γ). Mathematically, we can represent $y_i(t)$ as the $\sqrt{\text{CD4}}$ at time t for patient $i = 1, \dots, n = 467$ and therefore the longitudinal model is defined by:

$$y_i(t) = \tilde{\beta}_0 + \tilde{\beta}_1 t + \gamma \text{prevOI}_i t + \beta \text{prevOI}_i + b_{0i} + b_{1i} t + \epsilon_i(t), \quad (5)$$

where $\epsilon_i(t) \stackrel{\text{ind.}}{\sim} N(0, \sigma^2)$ is the random error at time t and we assume that the random effects, $\mathbf{b}_i = (b_{0i}, b_{1i})^\top$, given Σ , follow a joint bivariate normal distribution with zero mean and variance-covariance matrix Σ .

The hazard function at time t is defined as:

$$h(t | \Psi_{2i}) = h_0 \left(t \exp \{ \alpha_1 (\gamma \text{prevOI}_i + b_{1i}) \} \mid \theta \right) \exp \{ \lambda \text{prevOI}_i + \alpha_0 b_{0i} \}, \quad (6)$$

where h_0 is a baseline hazard function that will be specified here as Log-normal, Gamma, Generalised Weibull, and Generalised Gamma (see Section 3.2 for more details of these specifications); α_0 and α_1 denote the association parameters; and λ is the regression coefficient for prevOI .

5.2 | Model 2 (M_2)

The longitudinal model of our second proposal is slightly different from the specification of (5), as now we do not include the time-dependent effect for the prevOI covariate. Then, the linear mixed model (LMM) is expressed by:

$$y_i(t) = \tilde{\beta}_0 + \tilde{\beta}_1 t + \beta \text{prevOI}_i + b_{0i} + b_{1i} t + \epsilon_i(t). \quad (7)$$

In this joint approach specification, the survival model shares only the random effects:

$$h(t | \Psi_{2i}) = h_0 \left(t \exp \{ \alpha_1 b_{1i} \} \mid \theta \right) \exp \{ \lambda \text{prevOI}_i + \alpha_0 b_{0i} \}. \quad (8)$$

5.3 | Model 3 (M_3)

Our third proposal models the longitudinal process as in (7), but shares only the random intercept:

$$h(t | \Psi_{2i}) = h_0(t | \theta) \exp \{ \lambda \text{prevOI}_i + \alpha_0 b_{0i} \}. \quad (9)$$

The prior distributions for the parameters and hyperparameters of models M_1 , M_2 and M_3 are specified as in Section 3.2.

5.4 | Bayesian model comparison

Suppose we have m Bayesian models, say M_1, \dots, M_m , to be compared. So, the relative plausibility of a particular model M_v given its prior probability and the evidence from the data is quantified by the so-called *posterior model probability* (PMP),²⁷ defined as follows:

$$\text{PMP}_v = P(M_v | \text{Data}) = \frac{P(\text{Data} | M_v)P(M_v)}{\sum_{j=1}^m P(\text{Data} | M_j)P(M_j)}, \quad v = 1, \dots, m, \quad (10)$$

where again we assumed that the models are equally probable *a priori*.

5.5 | Results

We start the analysis by comparing the joint models M_1 , M_2 and M_3 using the following baseline hazards: Log-normal, Gamma, Power Generalised Weibull, and Generalised Gamma. Table 1 shows the approximate calculation of posterior model probabilities, obtained with the `post_prob` function, available in the R package `bridgesampling`²⁸.

TABLE 1 Baseline hazard model comparison based on posterior model probability (PMP).

Model	Log-normal	Gamma	Power Generalised Weibull	Generalised Gamma
1	0	0	0	0
2	0	0.9944	0	0.0001
3	0	0.0011	0.0044	0

The results indicate M_2 with the Gamma baseline hazard as the best model. Table 2 shows a posterior summary for this model. The last column of this table contains the posterior probability that the corresponding parameter is positive. A probability equal to 0.5 indicates that a positive value of the parameter is equally likely than a negative one.

TABLE 2 Posterior summary for model M_2 with the Gamma baseline hazard specification.

Interpretation	Parameter	Mean	Median	2.5%	97.5%	$P(\cdot > 0 \mid \text{Data})$
intercept	$\tilde{\beta}_0$	3.111	3.111	2.986	3.242	1.000
slope	$\tilde{\beta}_1$	-0.042	-0.042	-0.050	-0.034	0.000
prevOI (AIDS)	β	-0.910	-0.909	-1.067	-0.754	0.000
intercept RE variance	$\sigma_{b_0}^2$	0.582	0.582	0.500	0.678	-
slope RE variance	$\sigma_{b_1}^2$	0.002	0.002	0.001	0.002	-
RE correlation	ρ	0.040	0.039	-0.143	0.231	0.655
error variance	σ^2	0.134	0.134	0.120	0.150	-
prevOI (AIDS)	λ	1.622	1.613	1.181	2.108	1.000
intercept RE association	α_0	-0.935	-0.935	-1.207	-0.666	0.000
time association	α_1	-25.681	-25.587	-44.110	-8.053	0.002
Gamma scale	η	33.081	32.007	20.172	52.348	-
Gamma shape	ν	1.760	1.740	1.440	2.195	-

prevOI: Previous Opportunistic Infection at study entry. RE: Random Effect.

The first seven parameters in Table 2 refer to the longitudinal model for the $\sqrt{\text{CD4}}$. In particular, it is important to note that the posterior mean of the amount of CD4 (in square root scale) to patients with previous opportunistic infection (AIDS diagnosis) at study entry was less than that of patients with no previous infection by $E(\beta \mid \text{Data}) = -0.91$ units.

As expected, the group of patients with previous opportunistic infection at study entry has a higher risk of death, $E(\lambda \mid \text{Data}) = 1.615$. In addition, the association parameters (α_0 and α_1) indicate that by having a low CD4 amount at study entry (intercept) or even decreasing this amount throughout the study (slope), the risk of death increases.

Considering the following parameterisation of the Gamma probability distribution function $f_0(t) = \zeta^\nu t^{\nu-1} e^{-\zeta t} / \Gamma(\nu)$, Figure 1 shows the plots of the predictive baseline hazard, $h_0(t) = f_0(t) / S_0(t)$, and survival, $S_0(t)$, and their respective 95% credible interval using the posterior samples of ζ and ν from model M_2 .

6 | THE SANAD STUDY: COMPETING RISKS JOINT MODEL

The SANAD (Standard and New Anti-epileptic Drugs) study, designed and analysed by Marson et al¹⁷, is an unblinded randomised controlled trial in hospital-based outpatient clinics conducted between 1998 and 2006 in the UK. Partial data from this study is publicly available in the R package `joiner`²⁹, where the main objective is to investigate the time to treatment failure (here defined as the withdrawal of a randomised drug or addition of another) based on a standard anti-epileptic drug (carbamazepine, CBZ) and a new drug (lamotrigine, LTG). The time to treatment failure can occur due to two competing events: inadequate seizure control (ISC) or unacceptable adverse effects (UAE). Table 3 shows a brief summary of the baseline covariates and time to events for each competing event.

Additionally, at each clinical visit, the drug dose of each patient is adjusted if necessary. So, the dose at each visit is a longitudinal marker potentially associated with the time until the events of interest. This time-dependent endogenous covariate, typically modelled through a linear mixed-effects specification, is linked to the competing risks model by means of a joint modelling.³⁰ Williamson et al^{31,32} were the first to analyse this dataset using a competing risk model without longitudinal

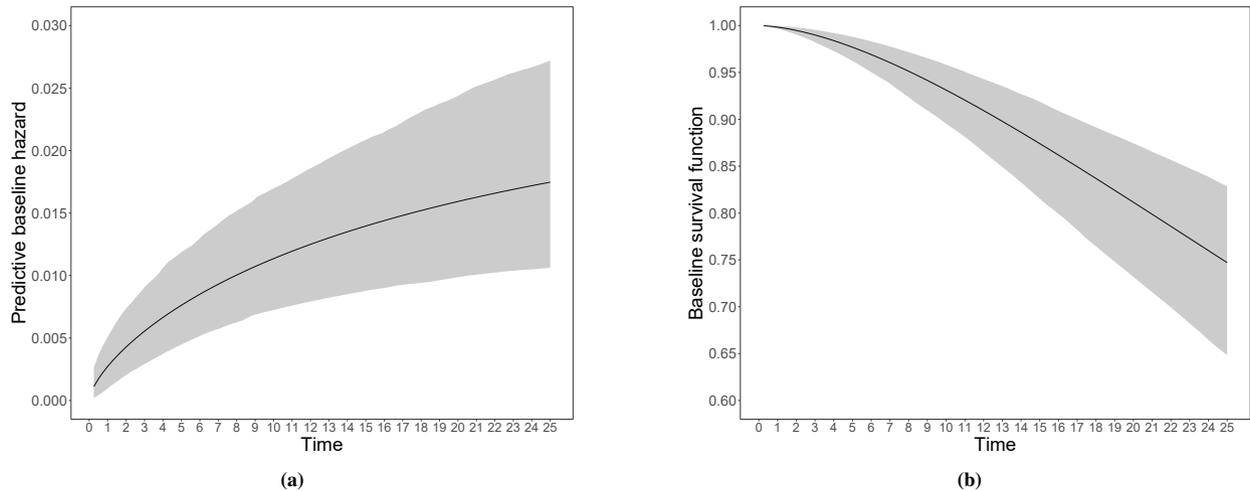


FIGURE 1 Gamma predictive (a) baseline hazard and (b) survival functions, and their respective 95% credible intervals for model M_2 .

TABLE 3 Competing event status, baseline covariates and time to events.

	Censored	ISC	UAE
n	391	120	94
gender: Female Male	173 218	52 68	37 57
treat: CBZ LTG	179 212	55 65	58 36
age: Mean (SD) [in years]	38.4 (19.1)	33.6 (16.7)	38.9 (19.0)
time: Median (SD) [in years]	2.2 (1.7)	1.3 (1.3)	0.5 (0.9)

information. Later, Williamson et al³³ proposed a joint modelling approach and more recently Hickey et al¹² compared different specifications of competing risks joint models for these data.

To model this problem, we propose three flexible specifications for joint models for longitudinal and competing risks data. All proposals model the longitudinal dose variable as a linear mixed model (LMM) and the competing risks data as a cause-specific hazards model³⁴ using the log-normal baseline specification. The details of each model are described below.

6.1 | Model 1 (M_1)

Our first proposal specifies the longitudinal model with fixed effects for the intercept ($\tilde{\beta}_0$), slope ($\tilde{\beta}_1$), gender (β_1), treat (β_2), and age (β_3); random effects for the intercept (b_0) and slope (b_1); and a time-dependent effect for the age covariate (γ). Mathematically, we can represent $y_i(t)$ as the drug dose at time t for patient $i = 1, \dots, n = 605$ and therefore the longitudinal model is defined by:

$$y_i(t) = \tilde{\beta}_0 + \tilde{\beta}_1 t + \gamma \text{age}_i t + \beta_1 \text{gender}_i + \beta_2 \text{treat}_i + \beta_3 \text{age}_i + b_{0i} + b_{1i} t + \epsilon_i(t), \quad (11)$$

where $\epsilon_i(t) \stackrel{\text{ind.}}{\sim} N(0, \sigma^2)$ is the random error at time t and we assume that the random effects, $\mathbf{b}_i = (b_{0i}, b_{1i})^\top$, given Σ , follow a joint bivariate normal distribution with zero mean and variance-covariance matrix Σ .

The cause-specific hazard function of the k th treatment failure at time t is defined as:

$$h_k(t | \Psi_{2i}) = h_{k0}(t \exp \{ \alpha_{k1}(\gamma \text{age}_i + b_{1i}) \} | \mu_k, \eta_k) \exp \{ \lambda_{k1} \text{gender}_i + \lambda_{k2} \text{treat}_i + \lambda_{k3} \text{age}_i + \alpha_{k0} b_{0i} \}, \quad (12)$$

where h_{k0} is a log-normal baseline hazard function with log-location μ_k and scale η_k parameters; α_{k0} and α_{k1} denote the association parameters; λ_{k1} , λ_{k2} and λ_{k3} are the regression coefficients for gender, treat and age; and $k = I, U$ represent ISC and UAE events, respectively.

6.2 | Model 2 (M_2)

The longitudinal model of our second proposal is slightly different from the specification of (11), as now we do not include the time-dependent effect for the age covariate. Then, the LMM is expressed by:

$$y_i(t) = \tilde{\beta}_0 + \tilde{\beta}_1 t + \beta_1 \text{gender}_i + \beta_2 \text{treat}_i + \beta_3 \text{age}_i + b_{0i} + b_{1i} t + \epsilon_i(t). \quad (13)$$

In this joint approach specification, the competing risks model shares only the random effects:

$$h_k(t | \Psi_{2i}) = h_{k0}(t \exp \{ \alpha_{k1} b_{1i} \} | \mu_k, \eta_k) \exp \{ \lambda_{k1} \text{gender}_i + \lambda_{k2} \text{treat}_i + \lambda_{k3} \text{age}_i + \alpha_{k0} b_{0i} \}. \quad (14)$$

6.3 | Model 3 (M_3)

Our third proposal models the longitudinal process as in (13), but shares only the random intercept:

$$h_k(t | \Psi_{2i}) = h_{k0}(t | \mu_k, \eta_k) \exp \{ \lambda_{k1} \text{gender}_i + \lambda_{k2} \text{treat}_i + \lambda_{k3} \text{age}_i + \alpha_{k0} b_{0i} \}. \quad (15)$$

The prior distributions for the parameters and hyperparameters of models M_1 , M_2 and M_3 are specified as in Section 3.2.

For these analyses, we coded gender as 0 (Female) and 1 (Male), treat as 0 (CBZ) and 1 (LTG), age was standardised, and the longitudinal dose variable for both groups of drugs were rescaled to have the same range of values. From now on we will refer to the dose as the calibrated dose due to this scale transformation.

6.4 | Bayesian model comparison

In addition to the posterior model probability (see Section 5.4), we also used the Bayes factor. Let M_v and M_j be two Bayesian models competing with each other, then the Bayes factor in favour of M_v against M_j is defined by:

$$\text{BF}_{vj} = \frac{P(\text{Data} | M_v)}{P(\text{Data} | M_j)} = \frac{P(M_v | \text{Data}) P(M_v)}{P(M_j | \text{Data}) P(M_j)}, \quad (16)$$

where we assumed that M_v and M_j are equally probable *a priori*, so that $P(M_v) = P(M_j)$ and therefore $\text{BF}_{vj} = P(M_v | \text{Data}) / P(M_j | \text{Data})$. In order to show the comparative results on a more friendly scale, we used \log_{10} -Bayes factor (LBF) with the interpretations proposed by Kass and Raftery³⁵.

6.5 | Results

We started the analysis by comparing the joint models M_1 , M_2 and M_3 introduced in previous sections. Table 4 shows the approximate calculation of Bayes factors and posterior model probabilities, obtained with the `bf` and `post_prob` functions, respectively, available in the R package `bridgesampling`²⁸.

TABLE 4 Model comparison based on posterior model probability (PMP) and \log_{10} -Bayes factor (LBF).

Posterior model probability			Log ₁₀ -Bayes factor		
PMP ₁	PMP ₂	PMP ₃	LBF ₁₂	LBF ₁₃	LBF ₂₃
0	1	0	-7.20	19.56	26.76

The results are decisively favorable to model M_2 and indicates the model M_3 as the worst option. Table 5 shows a posterior summary for model M_2 and the results for models M_1 and M_3 are presented in Section A5 in the Supplementary Material.

TABLE 5 Posterior summary for model M_2 with a log-normal baseline hazard specification.

Interpretation	Parameter	Mean	Median	2.5%	97.5%	$P(\cdot > 0 \mid \text{Data})$
intercept	$\tilde{\beta}_0$	1.812	1.812	1.674	1.950	1.000
slope	$\tilde{\beta}_1$	0.347	0.346	0.299	0.395	1.000
gender (Male)	β_1	0.089	0.089	-0.062	0.240	0.877
treat (LTG)	β_2	-0.030	-0.029	-0.180	0.120	0.350
age	β_3	0.005	0.005	-0.068	0.080	0.553
intercept RE variance	$\sigma_{b_0}^2$	0.745	0.743	0.649	0.851	-
slope RE variance	$\sigma_{b_1}^2$	0.162	0.161	0.127	0.204	-
RE correlation	ρ	0.043	0.042	-0.075	0.163	0.754
error variance	σ^2	0.199	0.198	0.186	0.212	-
ISC						
gender (Male)	λ_{I1}	0.039	0.038	-0.332	0.415	0.579
treat (LTG)	λ_{I2}	-0.248	-0.248	-0.627	0.132	0.100
age	λ_{I3}	-0.254	-0.253	-0.464	-0.046	0.009
intercept RE association	α_{I0}	0.181	0.182	-0.053	0.409	0.937
time association	α_{I1}	-7.139	-7.065	-9.795	-4.896	0.000
log-normal log-location	μ_I	2.959	2.927	2.084	4.014	1.000
log-normal scale	η_I	2.716	2.708	2.301	3.167	-
UAE						
gender (Male)	λ_{U1}	0.196	0.195	-0.308	0.708	0.775
treat (LTG)	λ_{U2}	-1.009	-1.005	-1.522	-0.508	0.000
age	λ_{U3}	0.139	0.138	-0.102	0.384	0.869
intercept RE association	α_{U0}	-1.278	-1.275	-1.617	-0.957	0.000
time association	α_{U1}	-12.060	-11.427	-21.827	-5.635	0.000
log-normal log-location	μ_U	3.339	3.295	2.003	4.935	1.000
log-normal scale	η_U	2.887	2.872	2.383	3.463	-

LTG: Lamotrigine. RE: Random Effect. ISC: Inadequate Seizure Control. UAE: Unacceptable Adverse Effects.

The first nine parameters in Table 5 refer to the longitudinal model for the calibrated drug dose. $E(\tilde{\beta}_0 \mid \text{Data}) = 1.812$ represents the posterior mean of the average value of dose at baseline with its respective underlying posterior variance (intercept random effect) of 0.745 among patients. The posterior mean dose increment each year was 0.347 and its posterior mean inter-individual variation (slope random effect) was 0.162. The mean posterior correlation between the random effects was positive but small (0.043). The posterior mean of the amount of dose delivered to men was marginally higher than that of women by $E(\beta_1 \mid \text{Data}) = 0.089$ units. The posterior mean of the amount of dose administered to the LTG-treated patient group was insignificantly less than in those CBZ-treated ($E(\beta_2 \mid \text{Data}) = -0.030$). The age of patients was irrelevant in terms of the amount of dose delivered to them. The posterior mean of the error variance was 0.199.

Regarding the risk of ISC, the patient's gender had virtually no influence, whereas treatment LTG and age produced a decrease in risk of ISC. As the posterior mean of the association parameter for the random intercept was positive, also confirmed by $P(\alpha_{I0} > 0 \mid \text{Data}) = 0.937$, and so a high baseline dose leads to a higher risk of ISC. On the other hand, the posterior mean of the association parameter for the random slope was negative, it implies that an increase in dose is associated with a decrease in risk of ISC.

Regarding the risk of UAE, male and/or older patients had a higher risk of UAE, whereas LTG-treated patient group had a reduced risk of UAE compared to those CBZ-treated. For this competing event, a high baseline dose and its increase over time reduced the risk of the patient experiencing UAE, since the posterior mean of the association parameters for the random intercept and slope were distinctly negative.

In order to visually compare the competing risks, we have used the *predictive baseline hazards* and the *cumulative incidence functions* based on the posterior sample of the log-normal log-location (μ) and scale (η) parameters for ISC and UAE treatment failures. The predictive baseline hazard function for the k th risk is described as:

$$h_{k0}(t | \text{Data}) = \frac{f_{k0}(t | \text{Data})}{S_0(t | \text{Data})}, \quad k = \text{I, U}, \quad (17)$$

where

$$f_{k0}(t | \text{Data}) = \int_{\mathbb{R}_+} \int_{\mathbb{R}} h_{k0}(t | \mu_k, \eta_k) S_0(t | \mu_k, \eta_k) \pi(\mu_k, \eta_k | \text{Data}) d\mu_k d\eta_k,$$

$$S_0(t | \text{Data}) = \prod_{k=\text{I, U}} \int_{\mathbb{R}_+} \int_{\mathbb{R}} \exp \left\{ - \int_0^t h_{k0}(u | \mu_k, \eta_k) du \right\} \pi(\mu_k, \eta_k | \text{Data}) d\mu_k d\eta_k,$$

are the baseline posterior predictive sub-density function for the k th risk and the overall posterior predictive survival function, respectively. The (baseline) posterior predictive cumulative incidence function represents the probability of failure from cause k before time t in the presence of all other possible causes,³⁶ and is defined as:

$$F_{k0}(t | \text{Data}) = \int_0^t f_{k0}(u | \text{Data}) du, \quad k = \text{I, U}. \quad (18)$$

These quantities can be approximated using Monte Carlo integration based on the posterior samples. Figure 2 shows the plots of (17) and (18) according to ISC and UAE risks, and their respective 95% credible interval for model M_2 .

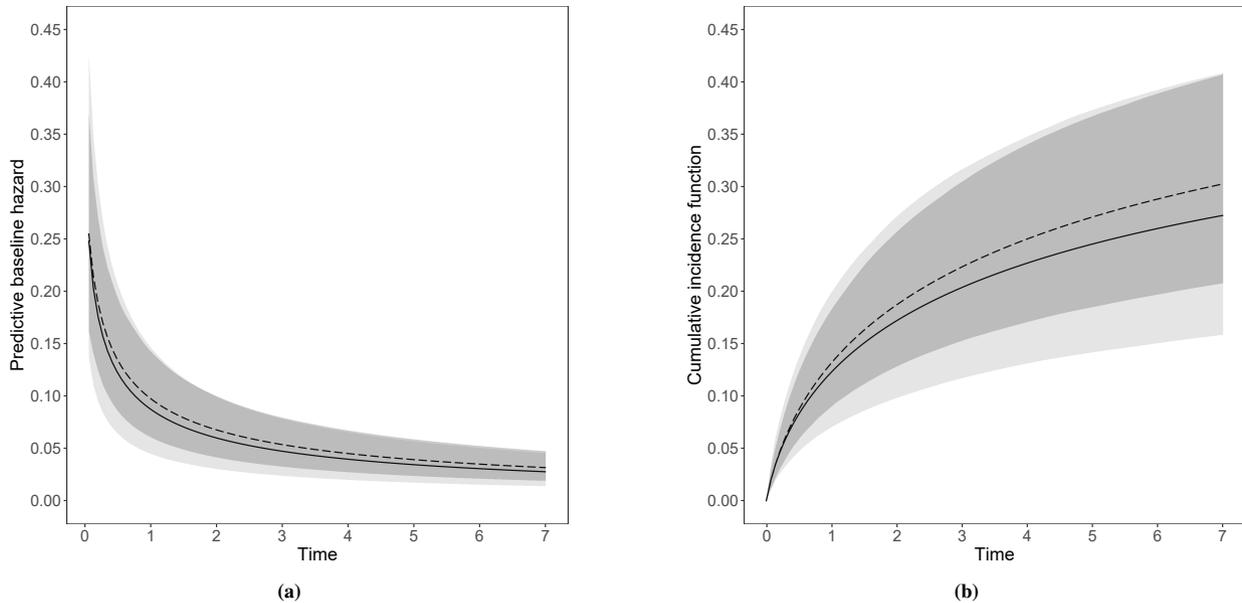


FIGURE 2 Log-normal predictive (a) baseline hazard and (b) cumulative incidence function of ISC (dashed line) and UAE (solid line), and their respective 95% credible intervals (dark and light grey) for model M_2 .

The interpretations presented here are consistent with previous work that analysed this data.^{31,32,33,12}

7 | DISCUSSION

We have proposed a formulation of Bayesian joint models for longitudinal and survival data which allows for a relatively simple interpretation of the parameters and a tractable implementation. The idea is to model the survival process using a general hazard structure that separates the roles of the variables acting on the time scale from those that affect the hazard scale. This formulation can be coupled with the use of flexible parametric baseline hazards (e.g. PGW or GG), which can capture a variety of hazard shapes, avoiding the need for numerical integration. We connect the survival process with the longitudinal process by sharing parameters with a similar interpretation. The longitudinal process can be modelled using GLMMs, allowing for the inclusion of a variety of response variables including continuous and categorical, within the Exponential family. This formulation facilitates the implementation of the proposed joint models in a Bayesian framework using MCMC methods. In this paper, we have focused on the use of `rstan`, but other methods that allow for efficiently sampling from models with random effects can be used as well. We have presented a honest characterisation of the limitations of the proposed joint specification, which include guidelines on cases with high censoring rates, or with early administrative censoring. In such cases, the use of flexible parametric baseline hazards has to be taken with some care as, intuitively and as shown in our simulation study, there is not enough information to estimate the parameters controlling the tails. This is reflected on the resulting wide posterior distributions, compared to those associated to simpler choices of the baseline hazard (e.g. log-normal). Model selection tools, such as Bayes factors or posterior model probabilities, are thus useful to identify the best model. In fact, the study of the performance of Bayesian model selection tools in the context of joint modelling represents a potential future research direction. The real data applications presented here illustrate the flexibility of our formulation to be adapted to settings with competing or semi-competing risks in the survival process, adding another option to the toolbox for modelling these challenging scenarios^{37,38}. There are several natural extensions of the proposed joint models. For instance, other flexible parametric baseline hazards and flexible distributions on the longitudinal models can be employed instead of the ones presented here.²¹ Although we have presented a careful prior elicitation step, combining g-priors and weakly informative priors, we do not claim to have the last word on this point. Thus, other priors can be used as well (see Rossell and Rubio¹⁸ for a discussion on different priors for survival models), and our R codes can be easily adapted for that purpose.

Financial disclosure

None reported.

Conflict of interest

The authors declare no potential conflict of interests.

References

1. Cox D. Regression Models and Life-Tables. *Journal of the Royal Statistical Society. Series B (Methodological)* 1972; 34(2): 187–202. doi: 10.1111/j.2517-6161.1972.tb00899.x
2. Kalbfleisch J, Prentice R. *The Statistical Analysis of Failure Time Data*. 360. John Wiley & Sons . 2011.
3. Chen Y, Jewell N. On a General Class of Semiparametric Hazards Regression Models. *Biometrika* 2001; 88(3): 687–702. doi: 10.1093/biomet/88.3.687
4. Rubio F, Remontet L, Jewell N, Belot A. On a General Structure for Hazard-Based Regression Models: An Application to Population-Based Cancer Research. *Statistical Methods in Medical Research* 2019; 28(8): 2404–2417. doi: 10.1177/0962280218782293
5. McCulloch C, Searle S, Neuhaus J. *Generalized, Linear, and Mixed Models*. Wiley . 2008.
6. Rizopoulos D. *Joint Models for Longitudinal and Time-to-Event Data: With Applications in R*. Chapman and Hall/CRC . 2012.

7. Elashoff R, Li G, Li N. *Joint Modeling of Longitudinal and Time-to-Event Data*. CRC Press . 2016.
8. Furgal A, Sen A, Taylor J. Review and Comparison of Computational Approaches for Joint Longitudinal and Time-to-Event Models. *International Statistical Review* 2019; 87(2): 393–418. doi: 10.1111/insr.12322
9. Alsefri M, Sudell M, García-Fiñana M, Kolamunnage-Dona R. Bayesian Joint Modelling of Longitudinal and Time to Event Data: A Methodological Review. *BMC Medical Research Methodology* 2020; 20(94): 1–17. doi: 10.1186/s12874-020-00976-2
10. Brilleman S, Crowther M, Moreno-Betancur M, et al. Joint Longitudinal and Time-to-Event Models for Multilevel Hierarchical Data. *Statistical Methods in Medical Research* 2019; 28(12): 3502–3515. doi: 10.1177/0962280218808821
11. Rué M, Andrinopoulou E, Alvares D, Armero C, Forte A, Blanch L. Bayesian Joint Modeling of Bivariate Longitudinal and Competing Risks Data: An Application to Study Patient-Ventilator Asynchronies in Critical Care Patients. *Biometrical Journal* 2017; 59(6): 1184–1203. doi: 10.1002/bimj.201600221
12. Hickey G, Philipson P, Jorgensen A, Kolamunnage-Dona R. A Comparison of Joint Models for Longitudinal and Competing Risks Data, With Application to an Epilepsy Drug Randomized Controlled Trial. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 2018; 181(4): 1105–1123. doi: 10.1111/rssa.12348
13. Mauff K, Steyerberg E, Kardys I, Boersma E, Rizopoulos D. Joint Models With Multiple Longitudinal Outcomes and a Time-to-Event Outcome: A Corrected Two-Stage Approach. *Statistics and Computing* 2020; 30: 999–1014. doi: 10.1007/s11222-020-09927-9
14. Zellner A. *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno De Finetti*. On Assessing Prior Distributions and Bayesian Regression Analysis with g-Prior Distributions: 233–243; Elsevier/North-Holland . 1986.
15. Stan Development Team . *RStan: the R Interface to Stan*. Stan; <http://mc-stan.org/>: 2020.
16. Goldman A, Carlin B, Crane L, et al. Response of CD4 Lymphocytes and Clinical Consequences of Treatment Using ddI or ddC in Patients with Advanced HIV Infection. *Journal of Acquired Immune Deficiency Syndromes and Human Retrovirology* 1996; 11(2): 161–169. doi: 10.1097/00042560-199602010-00007
17. Marson A, Al-Kharusi A, Alwaidh M, et al. The SANAD Study of Effectiveness of Carbamazepine, Gabapentin, Lamotrigine, Oxcarbazepine, or Topiramate for Treatment of Partial Epilepsy: An Unblinded Randomised Controlled Trial. *The Lancet* 2007; 369(9566): 1000–1015. doi: 10.1016/S0140-6736(07)60460-7
18. Rossell D, Rubio F. Additive Bayesian Variable Selection Under Censoring and Misspecification. *arXiv Preprint* 2019; na(na): na–na. doi: 1907.13563
19. Jones M, Noufaily A. Log-Location-Scale-Log-Concave Distributions for Survival and Reliability Analysis. *Electronic Journal of Statistics* 2015; 9(2): 2732–2750. doi: 10.1214/15-EJS1089
20. Li Y, Clyde M. Mixtures of g-Priors in Generalized Linear Models. *Journal of the American Statistical Association* 2018; 113(524): 1828–1845. doi: 10.1080/01621459.2018.1469992
21. Rubio F, Steel M. Flexible Linear Mixed Models With Improper Priors for Longitudinal and Survival Data. *Electronic Journal of Statistics* 2018; 12(1): 572–598. doi: 10.1214/18-EJS1401
22. Simpson D, Rue H, Riebler A, Martins T, Sørbye S. Penalising Model Component Complexity: A Principled, Practical Approach to Constructing Priors. *Statistical Science* 2017; 32(1): 1–28. doi: 10.1214/16-STS576
23. Dette H, Ley C, Rubio F. Natural (Non-) Informative Priors for Skew-Symmetric Distributions. *Scandinavian Journal of Statistics* 2018; 45(2): 405–420. doi: 10.1111/sjos.12306
24. Andrinopoulou E, Rizopoulos D. Bayesian shrinkage approach for a joint model of longitudinal and survival outcomes assuming different association structures. *Statistics in Medicine* 2016; 35(26): 4813–4823.

25. Rubio F, Rachet B, Giorgi R, Maringe C, Belot A. On Models for the Estimation of the Excess Mortality Hazard in Case of Insufficiently Stratified Life Tables. *Biostatistics* 2021; 22(1): 51–67.
26. Rizopoulos D. *JM: Joint Modeling of Longitudinal and Survival Data*. R package version 1.4-8; <https://CRAN.R-project.org/package=JM>: 2018.
27. Berger J, Molina G. Posterior Model Probabilities via Path-Based Pairwise Priors. *Statistica Neerlandica* 2005; 59(1): 3–15. doi: 10.1111/j.1467-9574.2005.00275.x
28. Gronau Q, Singmann H, Wagenmakers E. bridgesampling: An R Package for Estimating Normalizing Constants. *Journal of Statistical Software* 2020; 92(10): 1–29. doi: 10.18637/jss.v092.i10
29. Philipson P, Sousa I, Diggle P, et al. *joineR: Joint Modelling of Repeated Measurements and Time-to-Event Data*. R package version 1.2.5; <https://CRAN.R-project.org/package=joineR>: 2020.
30. Elashoff RM, Li G, Li N. An Approach to Joint Analysis of Longitudinal Measurements and Competing Risks Failure Time Data. *Statistics in Medicine* 2007; 26(14): 2813–2835. doi: 10.1002/sim.2749
31. Williamson P, Kolamunnage-Dona R, Smith C. The Influence of Competing-Risks Setting on the Choice of Hypothesis Test for Treatment Effect. *Biostatistics* 2007; 8(4): 689–694. doi: 10.1093/biostatistics/kx1040
32. Williamson P, Smith C, Sander J, Marson A. Importance of Competing Risks in the Analysis of Anti-Epileptic Drug Failure. *Trials* 2007; 8(12): 1–10. doi: 10.1186/1745-6215-8-12
33. Williamson P, Kolamunnage-Dona R, Philipson P, Marson A. Joint Modelling of Longitudinal and Competing Risks Data. *Statistics in Medicine* 2008; 27(30): 6426–6438. doi: 10.1002/sim.3451
34. Putter H, Fiocco M, Geskus R. Tutorial in Biostatistics: Competing Risks and Multi-State Models. *Statistics in Medicine* 2007; 26(11): 2389–2430. doi: 10.1002/sim.2712
35. Kass R, Raftery A. Bayes Factors. *Journal of the American Statistical Association* 1995; 90(430): 773–795. doi: 10.2307/2291091
36. Bakoyannis G, Touloumi G. Practical Methods for Competing Risks Data: A Review. *Statistical Methods in Medical Research* 2012; 21(3): 257–272. doi: 10.1177/0962280210394479
37. Andrinopoulou E, Rizopoulos D, Takkenberg J, Lesaffre E. Joint modeling of two longitudinal outcomes and competing risk data. *Statistics in Medicine* 2014; 33(18): 3167–3178.
38. Andrinopoulou E, Rizopoulos D, Takkenberg J, Lesaffre E. Combined dynamic predictions using joint models of two longitudinal outcomes and competing risk data. *Statistical Methods in Medical Research* 2017; 26(4): 1787–1801.

