Journal of the Royal Statistical Society Series C: Applied Statistics, 2024, **73**, 598–620
https://doi.org/10.1093/jrsssc/qlae003
Advance access publication 1 February 2024
Original Article

ROYAL STATISTICAL SOCIETY
DATA | EVIDENCE | DECISIONS
Journal of the Statistics Society
Series C
Applied Statistics

Bayesian semi-parametric inference for clustered recurrent events with zero inflation and a terminal event

Xinyuan Tian, Maria Ciarleglio, Jiachen Cai, Erich J. Greene, Denise Esserman, Fan Li[†] and Yize Zhao[†]

Department of Biostatistics, Yale University, New Haven, CT, USA

Address for correspondence: Fan Li and Yize Zhao, Department of Biostatistics, Yale University, New Haven, CT 06511, USA. Emails: fan.f.li@yale.edu; yize.zhao@yale.edu

Abstract

Recurrent events are common in clinical studies and are often subject to terminal events. In pragmatic trials, participants are often nested in clinics and can be susceptible or structurally unsusceptible to the recurrent events. We develop a Bayesian shared random effects model to accommodate this complex data structure. To achieve robustness, we consider the Dirichlet processes to model the residual of the accelerated failure time model for the survival process as well as the cluster-specific shared frailty distribution, along with an efficient sampling algorithm for posterior inference. Our method is applied to a recent cluster randomized trial on fall injury prevention.

Keywords: accelerated failure time model, Bayesian survival analysis, Dirichlet process, pragmatic clinical trials, semi-competing risks, zero inflation

1 Introduction

Recurrent event data are common in clinical studies when participants are followed up longitudinally. Typically, each event occurrence can be subject to right censoring as well as a competing terminal event, such as death. In large pragmatic clinical trials, the event processes are often observed across a heterogeneous population, along with an informative competing event process subject to between-participant clustering. These features bring new challenges for the analysis of clustered recurrent events, due to the need for simultaneously characterizing the recurrent event process, non-terminal as well as terminal event survival processes as a function of covariates.

Falls are the leading cause of injury-related death among older Americans, and approximately one in four older adults experiences fall each year, resulting in numerous deaths and injury-related hospitalization and healthcare utilization annually (Choi et al., 2019; Verma et al., 2016). There has been a rising interest in implementing effective fall prevention strategies at a healthcare system level or provider level, to improve patient outcomes and reduce fall injury-related mortality (Hopewell et al., 2018). In 2014, the Patient-Centered Outcomes Research Institute and the National Institute on Aging in the United States funded a pragmatic trial, the Strategies to Reduce Injuries and Develop Confidence in Elders (STRIDE; Bhasin et al., 2020) study, to assess the effectiveness of a patient-centred intervention on fall injury prevention for older adults; our work is directly motivated by the STRIDE study. In STRIDE, more than 6,000 community-dwelling adults from 86 primary care practices were recruited, with 43 practices randomized to intervention and the remaining to usual care. Participants were followed up every 4 months via tele-interview (this is a relatively large number of clusters, as the upper quartile of

[†] F.L. and Y.Z. are co-senior authors.

Received: January 18, 2022. **Revised:** October 19, 2023. **Accepted:** January 5, 2024
© The Royal Statistical Society 2024. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

number of clusters in a past systematic review by Ivers et al., 2011 was only 52). All reported fall injuries were recorded, and a blinded adjudication committee confirmed serious fall injures via medical and claim records from the participating healthcare systems and Centers for Medicare and Medicaid Services data (Ganz et al., 2019). During the study, 89% of participants did not experience an adjudicated serious fall injury. As a simple illustration, we randomly select 50 patients from one random intervention practice and one random usual care practice, and present in Figure 1 the time trajectories for recurrent adjudicated serious fall injuries and an observed death event or censoring for each participant. Irrespective of the intervention, the recurrent event rate was relatively low, and there was an excessive number of participants without events, which signals potential zero inflation for the recurrent event process. In addition, online supplementary Figure 1 presents the descriptive Kaplan–Meier survival curves for the terminal event.

There is a growing body of literature on the analysis of recurrent events in the presence of a terminal event. For example, Lancaster and Intrator (1998) represented the first effort to develop a recurrent event model with patient-level frailty subject to non-informative terminal events. Sinha et al. (2008) provided a comprehensive review of methods for recurrent event analysis with dependent termination and developed the first Bayesian approach to analyse such data. More recent developments for recurrent event analysis with dependent termination include estimating equations approaches under a frequentist paradigm (Kalbfleisch et al., 2013) and parametric or semi-parametric models under a Bayesian paradigm (Li et al., 2019, 2020; Lin et al., 2017). A key feature of these methods is to characterize the dependence between non-terminal and terminal events under a semi-competing risk perspective (Fine et al., 2001), as ignoring this dependence can lead to a biased inference. To do so, one common strategy is to formulate a joint model with a shared participant-level frailty in the recurrent event and terminal event sub-models, where the sub-models can either be based on the intensity functions of the event processes (Lee et al., 2019; Liu et al., 2004) or the hazard rate of the gap time between two events (Paulon et al., 2020; Yu & Liu, 2011). Alternatively, Xu et al. (2021) developed a joint latent-class models to allow for class-specific risks for recurrence and termination. Their approach bypasses the distributional assumption of the shared random effect and can potentially lead to more interpretable covariate effects within and across latent classes. Despite this growing literature, few existing methods have simultaneously addressed the complications of cluster correlated data featured in the STRIDE study, whereas failure to account for clustering can result in an invalid inference (Lee et al., 2016). Jung et al. (2019) developed an approach that accounted for between-participant clustering in the presence of recurrent and terminal events. A similar joint model was also formulated in Rondeau et al. (2015) and implemented in the `R` package `frailtypack`. However, these existing approaches require strong parametric assumptions on the between-participant clustering effect and have not addressed population heterogeneity with respect to event susceptibility.

The contributions of our work are several-fold. First, we propose a new joint model to analyse recurrent event and survival processes in the presence of between-participant clustering and a competing terminal event. We introduce random effects at the participant level and the practice level, both of which contribute to connecting the recurrent event and survival processes. Second, we address potential zero inflation within our modelling framework by including a point mass at zero for the recurrent event intensity function. Using a latent indicator to define the status of unsusceptibility for each participant, we are able to directly inform population heterogeneity by separating the unobserved unsusceptible sub-population from the whole study population (Kim, 2021; Liu et al., 2016). Third, we consider separate non-parametric Dirichlet process priors (Ferguson, 1973) for the residual in the survival process as well as for the cluster-specific random effect, which, compared with conventional parametric formulations, alleviates potential bias due to model mis-specification. Finally, we apply the proposed Bayesian semi-parametric approach to analyse participant-level data from the STRIDE trial and generate new insights.

The rest of the article is organized as follows. In Sections 2 and 3, we introduce our Bayesian semi-parametric model including specifications of all sub-models, choice of priors, and posterior inference. We evaluate the model performance by comparing with other competing approaches using simulations in Section 4. We provide a comprehensive analysis of the STRIDE study in Section 5 using the proposed model and several other existing modelling techniques. We conclude with a discussion in Section 6.

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

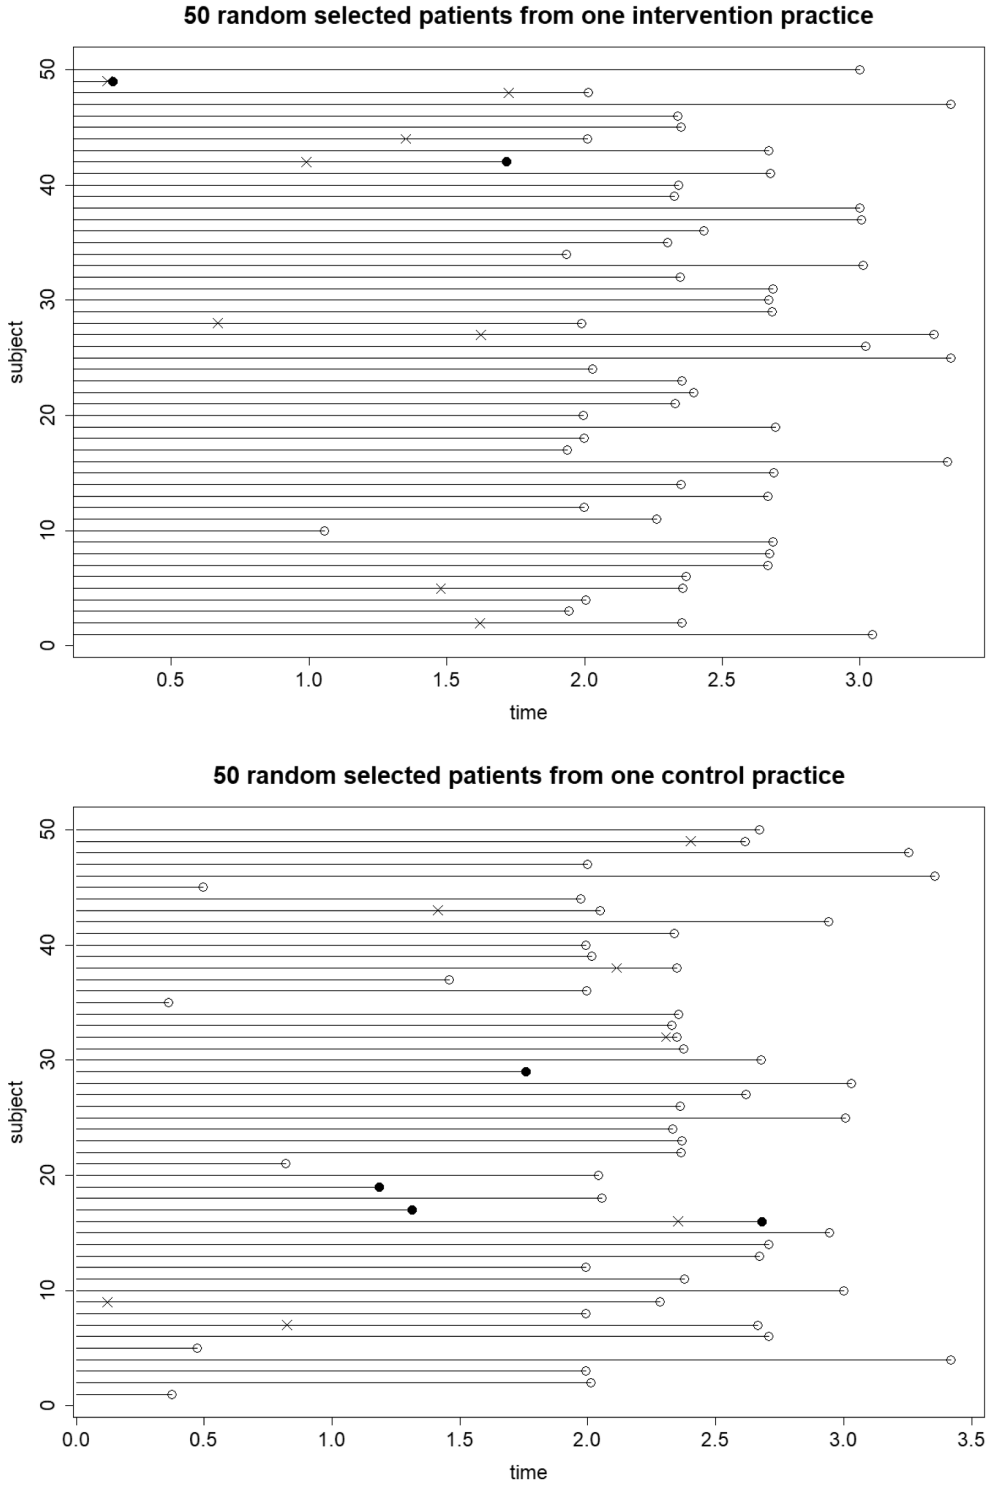

Figure 1. An illustration of (right-continuous) time trajectories for serious fall injury occurrence and terminal death event among randomly selected participants from both the intervention and control practices, where '∘' represents censoring, '•' represents occurrence of death, and '×' represents an occurrence of fall injury.

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

2 Modelling clustered recurrent events in the presence of a terminal event

We consider a clustered data structure with recurrent events that are subject to a terminal event, such as death. We assume J clusters (primary care practices) are recruited, with N_j participants in cluster j and $N = \sum_{j=1}^{J} N_j$ participants in total. Define $Q_{ij}(t)$ as the number of recurrent events prior to or at time t for participant i ($i = 1, \ldots, N_j$) within cluster j ($j = 1, \ldots, J$). In our motivating STRIDE study, adjudicated serious fall injuries for patients are considered as recurrent events, subject to the risk of death as a terminal event. After a terminal event, recurrent events are no longer observable. In this case, we define time to the terminal event for each participant as R_{ij} and the usual right censoring time (such as administrative censoring) as C_{ij}. The observed follow-up time is $\widetilde{R}_{ij} = R_{ij} \wedge C_{ij}$ with a censoring indicator $\Delta_{ij} = 1$ if the terminal event is observed and 0 if censored. Throughout, we assume that the censoring time can at most depend on the observed baseline covariates, and therefore do not consider dependent-censoring based on the unmeasured or time-varying information. Equivalently, we observe a total of $Q_{ij}(\widetilde{R}_{ij})$ recurrent events for participant i in cluster j. We also write $T_{ijk} \leq \widetilde{R}_{ij}$ as the time when the kth ($1 \leq k \leq Q_{ij}(\widetilde{R}_{ij})$) recurrent event is observed. For notation purposes, we define the collection of recurrent event times for each participant with at least one event as $\mathbf{T}_{ij} = \{T_{ij1}, \ldots, T_{ij,Q_{ij}(\widetilde{R}_{ij})}\}$, and for those with zero events as $\mathbf{T}_{ij} = \emptyset$.

As shown in Figure 1, a substantial proportion of participants in our motivating study have not experienced recurrent events, suggesting that some patients may be structurally unsusceptible to fall injuries during the study period, and could have distinctive characteristics from the remaining population. This requires us to separately consider this sub-group for plausibly uncovering the actual event mechanisms. To model zero inflation, we introduce a latent indicator D_{ij} with $D_{ij} = 1$ if participant i in cluster j belongs to the sub-group that is unsusceptible to recurrent event during the study period and 0 otherwise. We consider a point mass mixture of non-homogeneous Poisson process (NHPP) to model the recurrent event hazard (or intensity) function for each participant as

$$\lambda_{ij}(t) = \begin{cases} \gamma_{ij}\lambda_0(t)\exp\left(\boldsymbol{\beta}^T\mathbf{X}_{ij} + \mu_j\right) & \text{if } D_{ij} = 0; \\ 0 & \text{if } D_{ij} = 1. \end{cases} \tag{1}$$

In the hazard function (1), $\mathbf{X}_{ij}$ represents the set of covariates including the treatment arm and additional baseline characteristics potentially related to the recurrent process, $\boldsymbol{\beta}$ are the coefficients representing the relationship between $\mathbf{X}_{ij}$ and recurrent event process among the susceptible sub-group with $D_{ij} = 0$, and $\lambda_0(t)$ is the associated baseline hazard. By definition, a participant belongs to the susceptible sub-group if $Q_{ij}(\widetilde{R}_{ij}) > 0$; otherwise, the participant can belong to either the susceptible or unsusceptible sub-group. In addition, γ_{ij} is the subject-specific frailty accounting for the correlation between recurrent events for the same participant, and μ_j is the cluster-specific random effect that captures between-participant correlation within the same practice.

There are different options to specify $\lambda_0(t)$ under a Bayesian paradigm. For instance, we could assume a power-law model (Lee et al., 2019) with $\lambda_0(t) = \psi t^{\psi-1}$, which corresponds to a Weibull baseline hazard with scale parameter 1 and shape parameter ψ controlling the rate of event occurrences. Alternatively, we could also consider a non-parametric specification for $\lambda_0(t)$ with a piece-wise constant function (Jung et al., 2019; McKeague & Tighiouart, 2000)

$$\lambda_0(t) = \sum_{g=1}^{G} \mathbf{I}\left[s_{g-1} < t \leq s_g\right] \cdot \lambda_{0g}, \tag{2}$$

where $\mathbf{I}[\cdot]$ is the indicator function, $s_0 = 0$, s_G represents the largest recurrent event time, and $\{s_1, \ldots, s_{G-1}\}$ are $G - 1$ grid points that partition the time interval such that baseline hazard is a constant λ_{0g} over $(s_{g-1}, s_g]$. While the power-law model assumes a monotone baseline hazard, the piece-wise constant model can be more flexible and more robust to model assumptions. In what follows, we will primarily focus on the piece-wise constant model (2); additional details and numerical results under the power-law baseline hazard are provided in the online supplementary Appendix S2.

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

For the survival process of the terminal event, we consider an accelerated failure time (AFT) model incorporating the hierarchical random effects shared with the recurrent event model

$$\log(R_{ij}) = \alpha_0 + \boldsymbol{\alpha}^T \mathbf{Z}_{ij} + \xi_1 \log(\gamma_{ij}) + \xi_2 \mu_j + \kappa_{ij}^{-1} \epsilon_{ij}, \tag{3}$$

where κ_{ij}^{-1} is the participant-specific shape parameter and ϵ_{ij} is the independent and identically distributed residual error for the log survival time. In equation (3), α_0 is the intercept that captures the common factor across subjects, $\mathbf{Z}_{ij}$ is the set of covariates associated with the terminal event time with coefficients $\boldsymbol{\alpha}$ and can differ from $\mathbf{X}_{ij}$ in the recurrent event model (1), and coefficients ξ_1 and ξ_2 control the degree of unobserved associations between the recurrent and terminal event processes at the participant level and cluster level, respectively. This above model representation indicates that the participant-level frailty γ_{ij} and the cluster-level random effect μ_j jointly affect the relative change in survival time for the terminal event to account for the variation beyond that captured by the observed covariates. Meanwhile, model (3) and the recurrent event intensity model (1) share the hierarchical random effects to induce an informative terminal event process. To interpret this in the STRIDE study, an elderly participant who is more susceptible to repeated occurrences of falls may be either more likely or unlikely to survive until the end of the study, as captured by the participant-level frailty γ_{ij} and its coefficient ξ_1. Similar interpretation also applies to the practice-level frailty μ_j and its coefficient ξ_2 in the terminal process sub-model. Several prior studies (Mitchell et al., 2013; Rietdyk et al., 2022) have shown that individuals who are more likely to falls might also face a heightened risk of severe injuries, which may impact on their survival until the end of a study. Consequently, employing shared random effects provides a mechanism to capture the potential interconnections among different outcomes at both the participant and practice levels. However, we acknowledge that while the shared random effects are justified in our application, they might not be suitable for other studies. We refer to a discussion on this point in Section 6.

For AFT model (3), a canonical parametric specification is to assume that residual error ϵ_{ij} follows a standard extreme value distribution and $\kappa_{ij} = \kappa$, $\forall i, j$. Under this parameterization, the AFT model implies a Weibull hazard function for the terminal event time with $h_{ij}(t \mid \kappa_{ij} = \kappa) = \gamma_{ij}^{-\kappa\xi_1} t^{\kappa-1} \kappa \exp\{-\kappa(\alpha_0 + \boldsymbol{\alpha}^T \mathbf{Z}_{ij} + \xi_2 \mu_j)\}$. Accordingly, the survival function becomes $H_{ij}(t \mid \kappa_{ij} = \kappa) = \exp[-\gamma_{ij}^{-\kappa\xi_1} t^{\kappa} \exp\{-\kappa(\alpha_0 + \boldsymbol{\alpha}^T \mathbf{Z}_{ij} + \xi_2 \mu_j)\}]$. Such an AFT model with a homogeneous error distribution, although easy to implement, may be less robust to between-participant heterogeneity in their baseline risk to the terminal event. To enhance model robustness, we consider a non-parametric Dirichlet process ($\mathcal{DP}$) to model the error distribution. Specifically, we assume the participant-specific shape parameter

$$\kappa_{ij} \mid F \overset{\text{i.i.d}}{\sim} F, \quad j = 1, \ldots, J; \; i = 1, \ldots, N_j; \quad F \sim \mathcal{DP}(\phi_0, F_0). \tag{4}$$

Here, F_0 is called a base measure that defines the expectation of the random probability $F \in \mathbb{R}$ from which κ_{ij} is sampled, and ϕ_0 is the scale parameter describing the overall sampling concentration or the variance of the random probability measure. We specify F_0 as a Gamma distribution $\mathcal{G}(a_\kappa, b_\kappa)$, and assign a weakly informative Gamma distribution for scale parameter $\phi_0 \sim \mathcal{G}(1, 1)$ to ensure adequate flexibility. Essentially, model (4) induces a non-parametric realization for the shape parameters, which then corresponds to a more flexible form of the hazard and survival functions. To elaborate on this point, we can represent the $\mathcal{DP}$ model in equation (4) by an infinite mixture of point masses (Sethuraman, 1994)

$$F = \sum_{k=1}^{\infty} \pi_k \mu_{\theta_k}, \quad \text{with} \quad \pi_k = \pi'_k \prod_{h=1}^{k-1} (1 - \pi'_h), \tag{5}$$

where μ_θ is a probability measure concentrated at θ, and the two sets of independent and identically distributed random variables $\{\pi'_k\}_{k=1}^{\infty}$ and $\{\theta_k\}_{k=1}^{\infty}$ follow

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

$$\pi'_k \mid F_0, \phi \sim \text{Beta}(1, \phi_0); \quad \theta_k \mid F_0, \phi \sim F_0; \quad k = 1, \ldots, \infty. \tag{6}$$

Here, $\{\theta_k\}_{k=1}^{\infty}$ is a sequence of independent draws from the base measure F_0, and $\{\pi_k\}_{k=1}^{\infty}$ are the weight parameters constructed via a stick-breaking representation. With probability one, F is a discrete distribution as a combination of infinite number of point masses. Under weights $\{\pi_k\}_{k=1}^{\infty}$, realization of each κ_{ij} will be obtained directly from F consisting of components $\{\theta_k\}_{k=1}^{\infty}$. The induced survival function for participant i in cluster j then becomes an infinite mixture of Weibull survival functions given by

$$\sum_{k=1}^{\infty} \pi_k H_{ij}(t \mid \kappa_{ij} = \theta_k) = \sum_{k=1}^{\infty} \pi_k \exp\Big[-\gamma_{ij}^{-\theta_k \xi_1} t^{\theta_k} \exp\big\{-\theta_k\big(\alpha_0 + \boldsymbol{\alpha}^T \mathbf{Z}_{ij} + \xi_2 \mu_j\big)\big\}\Big],$$

and the associated hazard function corresponds to a similar infinite mixture of Weibull hazards

$$\sum_{k=1}^{\infty} \left\{ \frac{\pi_k H_{ij}(t \mid \kappa_{ij} = \theta_k)}{\sum_{l=1}^{\infty} \pi_l H_{ij}(t \mid \kappa_{ij} = \theta_l)} \right\} \gamma_{ij}^{-\theta_k \xi_1} t^{\theta_k - 1} \theta_k \exp\Big\{-\theta_k\big(\alpha_0 + \boldsymbol{\alpha}^T \mathbf{Z}_{ij} + \xi_2 \mu_j\big)\Big\},$$

both of which are arguably much more flexible than their canonical, fully parametric counterparts. Meanwhile, as shown in equation (5), with k increased, π_k decreases exponentially and concentrates the sampling on a number of initial components. This allows the residual error distributions to group based on their identical shape parameter values, and in turn, induces a clustering effect to dissect sub-group of individuals sharing a similar shape of the survival function. Finally, the canonical AFT specification can be considered as a special case of equation (5) with a degenerate Dirac measure.

3 Bayesian inference

3.1 Prior specification

To jointly characterize the zero-inflated recurrent events and terminal event process, the proposed joint modelling framework involves the following unknown parameters: regression coefficients $\boldsymbol{\beta}$, $\boldsymbol{\alpha}$, ξ_1, and ξ_2, participant-level and cluster-level random effects $\boldsymbol{\gamma} = \{\gamma_{ij}\}$ and $\boldsymbol{\mu} = \{\mu_j\}$, latent indicator $\mathbf{D} = \{D_{ij}\}$, shape parameter $\boldsymbol{\kappa} = \{\kappa_{ij}\}$ for the terminal event sub-model, grid points $\mathbf{s}$, and piece-wise constants $\boldsymbol{\lambda} = (\lambda_{01}, \ldots, \lambda_{0G})$ for the recurrent event sub-model.

The hierarchical random effects play an important role in connecting the recurrent and survival processes, since they represent shared unmeasured factors in addition to those captured by the baseline covariates. The frailty $\boldsymbol{\gamma}$ is directly grouped by different practices the participants belong to and provides quantification of between-participant heterogeneity, while the practice-specific random effects $\boldsymbol{\mu}$ account for between-practice heterogeneity. We assume independence between elements of $\boldsymbol{\gamma}$ and assign $\gamma_{ij} \sim \mathcal{LN}(0, \tau_j^2)$, where $\mathcal{LN}$ represents a Log-Normal distribution and τ_j^2 represents a participant-specific variance parameter; and we adopt an Inverse Gamma ($\mathcal{IG}$) hyper-prior such that $\tau_j^2 \sim \mathcal{IG}(a_0, b_0)$. For $\boldsymbol{\mu}$, instead of using parametric conjugate priors, we consider a non-parametric $\mathcal{DP}$ prior by assuming

$$\mu_j \mid G \overset{\text{i.i.d}}{\sim} G, \quad j = 1, \ldots, J; \quad G \sim \mathcal{DP}(\phi, G_0). \tag{7}$$

We specify base measure G_0 as a Normal distribution, $\mathcal{N}(0, \sigma^2)$, and assign $\phi \sim \mathcal{G}(1, 1)$ to ensure adequate flexibility. Prior (7) induces a non-parametric representation for the random effects over practices. Since the inference of model parameters may be sensitive to parametric assumptions of the practice-level random effects (Gasparini et al., 2019), this non-parametric prior can induce more robust characterization of the quality of care in each practice. To facilitate posterior inference under equation (7), following equations (5) and (6), we also resort to an infinite mixtures

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

of point masses representation under the point mass random set $\{\eta_l\}_{l=1}^{\infty}$, where each μ_j sampled from under the weights $\{\tilde{\pi}_l\}_{l=1}^{\infty}$ [η and $\tilde{\pi}$ are analogues to those introduced in model (5)]. This also groups realizations of each element within $\boldsymbol{\mu}$ together by their identical values, indicating, for example, similar quality of care across the included practices. In STRIDE, the primary care practices are nested within different healthcare systems, which could induce inter-practice similarity. Although we do not directly account for heterogeneity across health systems beyond that across practices, the implicit clustering effect due to the $\mathcal{DP}$ prior automatically identifies more similar practices according to values of μ_j, either within or across different health systems, and provides additional flexibility beyond a single random effect at the health system level. Alternatively, in the absence of a clear grouping pattern between practices in terms of quality of care, we can still rely on equation (7) to potentially reduce the number of unknown practice-level random effects. Of note, in the analysis of the STRIDE trial, we have specified the above Log-Normal parametric prior for the patient-level frailty because the recurrent event rate was relatively low; however, a relatively large number of practices in STRIDE supports a non-parametric $\mathcal{DP}$ prior for the practice-level random effects. In addition, we have considered a $\mathcal{DP}$ shape-mixture of errors in the AFT terminal event model as well as a $\mathcal{DP}$ prior for the practice-level random effects in the same terminal event model. This double non-parametric prior specification does not lead to non-identifiability because the practice-level random effects are shared between the recurrent event and terminal event models and posterior inference for the practice-level random effects will be based on additional information beyond the terminal event process.

We further assume the latent indicator D_{ij} follows $D_{ij} \sim \text{Bern}(p_{ij})$, $i = 1, \ldots, N_j$, $j = 1, \ldots, J$, with p_{ij} being the participant-specific probability to be classified into the unsusceptible sub-group. When there is prior knowledge on potential risk factors that are associated with an individual's susceptibility status for recurrent events, we can adopt a logistic model

$$\text{logit}(p_{ij}) = \boldsymbol{\zeta}^T \mathbf{U}_{ij}, \tag{8}$$

where $\mathbf{U}_{ij}$ includes an intercept as well as risk factors for susceptibility and $\boldsymbol{\zeta}$ represents the regression coefficients (Cooner et al., 2007; Joseph & Robert, 1952). In other cases without strong prior information on such covariates, we could instead assume that p_{ij} takes a constant value, say 0.5, which leads to a non-informative prior for the latent indicator D_{ij}, and can be regarded as a special case of equation (8). For generality, we will discuss posterior inference under a general logistic formulation (8). In terms of the baseline hazard in the recurrent event process, following Jung et al. (2019), we pre-specify $\mathbf{s}$ as quantiles based on the minimum to the maximum recurrent event time, and adopt a uniform prior $(0, \infty)$ for each element within $\boldsymbol{\lambda}$, i.e. $p(\lambda_{0g}) \propto 1$. This improper uniform prior is a convenient choice and still leads to well-defined posterior distribution that integrates to 1, and results under a proper uniform prior over $(0, 100)$ are no different for the analysis of STRIDE (omitted for brevity). To complete prior specification, we assign priors for the remaining model parameters such that $\boldsymbol{\beta} \sim \mathcal{N}(\mathbf{0}, \sigma_\beta^2 \mathbf{I})$, $\boldsymbol{\alpha} \sim \mathcal{N}(\mathbf{0}, \sigma_\alpha^2 \mathbf{I})$, $\boldsymbol{\zeta} \sim \mathcal{N}(\mathbf{0}, \sigma_\zeta^2 \mathbf{I})$, $\xi_1 \sim \mathcal{N}(0, \sigma_{\xi_1}^2)$, $\xi_2 \sim \mathcal{N}(0, \sigma_{\xi_2}^2)$; and further assign conjugate $\mathcal{IG}$ hyper-priors for σ_β^2 and σ_α^2 and pre-specify the remaining hyper-parameters with reasonable values without strong prior impact on the posterior inference; as our model includes a substantial amount of parameters, sensitivity analyses to choice of hyper-parameters are also recommended. For an overview of our method, Figure 2 provides a graphical illustration of the data structure along with key modelling assumptions.

3.2 Likelihood and posterior inference

Given the observed data $\mathcal{O}_{ij} = \{\widetilde{R}_{ij}, \Delta_{ij}, \mathbf{T}_{ij}, \underline{Q}_{ij}(\widetilde{R}_{ij}), \mathbf{X}_{ij}, \mathbf{Z}_{ij}, \mathbf{U}_{ij}\}$ for each subject $i (i = 1, \ldots, N_j)$ within practice $j (j = 1, \ldots, J)$, we need to carefully distinguish between different events and survival states, as well as the sub-group each subject belongs to, in order to carry out inference for all model parameters. For example, while each participant may be or may not be susceptible to recurrent events, all participants are susceptible to the terminal events in the STRIDE application. With the unknown parameters $\Theta = \{\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\zeta}, \boldsymbol{\gamma}, \boldsymbol{\mu}, \mathbf{D}, \xi_1, \xi_2, \boldsymbol{\kappa}; \phi \text{ or } \boldsymbol{\lambda}\}$, the observed data likelihood

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

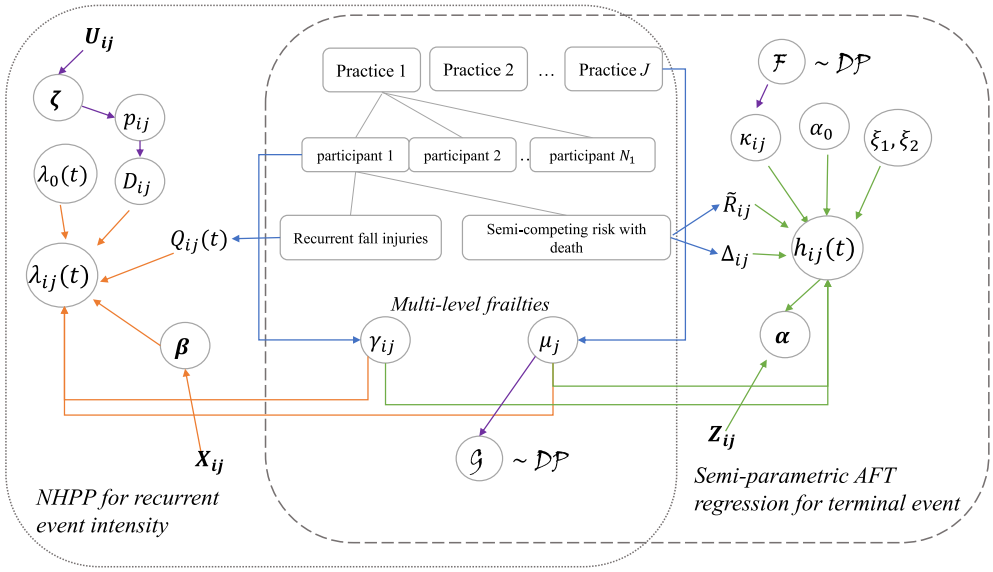

Figure 2. A graphical demonstration of the multi-level data structure and our proposed Bayesian joint model for the zero-inflated recurrent events and semi-competing survival process.

involves a combination of probabilities for structural zeros among the unsusceptible sub-group, recurrent events, and terminal events, and is given by

$$\mathcal{L}(\{\mathcal{O}_{ij}\} \mid \Theta) = \prod_{j=1}^{J}\prod_{i=1}^{N_j}\Big\{D_{ij} + (1-D_{ij})(1-\Delta_{ij})S_{ij}(\widetilde{R}_{ij})H_{ij}(\widetilde{R}_{ij}) + (1-D_{ij})\Delta_{ij}S_{ij}(\widetilde{R}_{ij})$$
$$\times f_{ij}(\widetilde{R}_{ij})\Big\}^{\mathbf{I}[Q_{ij}(\tilde{R}_{ij})=0]} \times \left\{\Delta_{ij}\prod_{k=1}^{Q_{ij}(\tilde{R}_{ij})}\lambda_{ij}(T_{ijk})S_{ij}(\widetilde{R}_{ij})f_{ij}(\widetilde{R}_{ij})\right.$$
$$\left. + (1-\Delta_{ij})\prod_{k=1}^{Q_{ij}(\tilde{R}_{ij})}\lambda_{ij}(T_{ijk})S_{ij}(\widetilde{R}_{ij})H_{ij}(\widetilde{R}_{ij})\right\}^{\mathbf{I}[Q_{ij}(\tilde{R}_{ij})>0]},$$

where $f_{ij}(\widetilde{R}_{ij}) = h_{ij}(\widetilde{R}_{ij})H_{ij}(\widetilde{R}_{ij})$ is the density function for the terminal event process of participant i in practice j evaluated at the observed survival time $\widetilde{R}_{ij}$, and the indicator function $\mathbf{I}[\,\cdot\,]$ separating the likelihood for those with and without recurrent events. By combining the observed data likelihood with our prior specification, we obtain the joint posterior distribution of Θ, from which we perform estimation and inference for each of the unknown parameters.

To achieve posterior inference, we develop a Markov chain Monte Carlo (MCMC) algorithm based on a combination of Gibbs sampler and Metropolis–Hastings (MH) updates. The full computational details of our MCMC are provided in the online supplementary material. In brief, under random initials, the algorithm cycles through the following steps:

- Sample each element of $\mathbf{D}$ from its posterior Bernoulli distribution.
- For the recurrent event sub-model, update each element of $\boldsymbol{\beta}$ via its MH step; and update each element of $\boldsymbol{\lambda}$ in the baseline hazard from its MH step.
- For the terminal event sub-model, update α_0 and each element of $\boldsymbol{\alpha}$ via the corresponding MH steps. For the individual shape parameter $\boldsymbol{\kappa}$, we implement an approximate sampling procedure under the truncated stick-breaking process (Ishwaran & James, 2001; Li et al., 2015), where a conservative upper bound K larger than the possible number of latent groups for the mixture of κ_{ij}'s is assigned. By introducing a mapping indicator set $\mathbf{v} = (v_{11}, \ldots, v_{N_J J})$

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

with $v_{ij} \in \{1, \ldots, K\}$ following a Multinomial distribution with probabilities $\{\pi_1, \ldots, \pi_K\}$, we align each κ_{ij} to its latent membership label v_{ij}. Within the same group membership label, the κ_{ij}'s are considered identical. Therefore, we update each v_{ij} from the posterior Multinomial distribution, sample κ_{ij} within each of the K clusters via an MH step, and update $\pi_k = \pi'_k \prod_{h<k}(1-\pi'_h)$ with π'_h sampled from the Beta distribution.

- For the participant-specific frailty, update each γ_{ij} via the MH step and sample the frailty variance τ_j^2 from its posterior $\mathcal{IG}$ distribution for $i = 1, \ldots, N_j$, $j = 1, \ldots, J$.
- For the practice-specific random effect $\boldsymbol{\mu}$, we implement a similar sampling procedure as that for $\boldsymbol{\kappa}$ by assigning a conservative upper bound L and introducing a mapping indicator set $\mathbf{m} = (m_1, \ldots, m_J)$ with each element following a Multinomial distribution with probabilities $\{\tilde{\pi}_1, \ldots, \tilde{\pi}_L\}$. The update for each μ_j, m_j, and $\tilde{\pi}_l$ follow a similar procedure to that used in updating the shape parameter of the terminal event sub-model.
- For the shared random effects in the terminal event sub-model, update ξ_1 and ξ_2 via the corresponding MH steps.
- For the logistic model, update each element of $\boldsymbol{\zeta}$ via the corresponding MH step.

In practice, we can assess the posterior convergence by both trace plots and the Gelman-Rubin method (Gelman & Rubin, 1992). Based on the posterior samples (after burn-in), we can directly obtain the point and credible interval estimators for each parameter using the posterior mean and associated quantiles.

4 Simulation studies

We carry out simulation studies to assess the finite-sample performance of the proposed Bayesian semi-parametric joint model and compare with alternative modelling approaches. Although our motivating STRIDE study recruited 86 practices, we simulate 60, 40, 20 practices, representing more challenging scenarios with fewer clusters. We assume equal numbers of participants per practice and consider $N = 1{,}800$, $N = 1{,}200$, and $N = 600$ as three levels of total sample sizes. For each participant, we specify the covariates $\mathbf{Z}_{ij}$ for the terminal event as a three-dimensional vector with each element generated from $\mathcal{N}(0, 0.1^2)$ and set $\boldsymbol{\alpha} = (0.2, 0.3, 0.4)^T$. We then generate the frailty $\gamma_{ij} \sim \mathcal{LN}(0, 0.25)$ with a common variance component across all practices. In reality, there may be unobserved heterogeneity at the practice level and the participant level. Relaxing the typical single-component Normal assumption for the practice-level random effect, we draw the practice-level random effect from a five-component mixture of Normals, given by,

$$\mu_j \overset{\text{ind}}{\sim} 0.2\mathcal{N}(-0.4, 0.1^2) + 0.2\mathcal{N}(-0.2, 0.1^2) + 0.2\mathcal{N}(0, 0.1^2) + 0.2\mathcal{N}(0.2, 0.1^2) + 0.2\mathcal{N}(0.4, 0.1^2).$$

This is equivalent to assuming that the practices can be roughly divided into five equal-sized strata, within which the practice-level random effect follows a single-component Normal distribution. Similarly, we accommodate unobserved participant-level heterogeneity by simulating the terminal event time R_{ij} from a mixture of Weibull distributions with the shape parameter κ_{ij} randomly sampled from the set $\{0.7, 2.2, 5.2, 8.2\}$ with replacement. Intuitively, this assumes that participants with the same covariate values may be additionally stratified into four equal-sized sub-groups, depending on the shape of their underlying terminal event distributions. These data-generating assumptions are designed to reflect unobserved heterogeneity at different levels and to assess the robustness of our models in relatively challenging scenarios. We also fix $\alpha_0 = 0.15$, $\xi_1 = 0.1$, $\xi_2 = -0.5$ for illustration. We then generate the censoring indicator independently from $\Delta_{ij} \sim \text{Bern}(0.5)$; that is, roughly 50% of participants exhibiting censored status for the terminal event process. When $\Delta_{ij} = 0$, we generate the observed survival time $\widetilde{R}_{ij}$ from a Uniform distribution under $(0, R_{ij})$; otherwise, we directly set $\widetilde{R}_{ij} = R_{ij}$. For the recurrent event process, we first specify the covariates $\mathbf{X}_{ij}$ for the recurrent events including the first two elements of $\mathbf{Z}_{ij}$ and a third element generated from $\mathcal{N}(0, 0.1^2)$, then we set $\boldsymbol{\beta} = (0.4, 0.3, 0.2)^T$. We generate the latent indicator D_{ij} from a

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

Bernoulli distribution with participant-specific probability (p_{ij}) to be classified into the unsusceptible sub-group. We specify p_{ij} based on the logistic model $\text{logit}(p_{ij}) = \boldsymbol{\zeta}^T \mathbf{U}_{ij}$, where U_{ij} includes an intercept as well as all elements of $\mathbf{Z}_{ij}$, and set $\zeta = (1, 0.8, 1, 1.2)^T$. Under this specification, approximately 70% of participants belong to unsusceptible sub-group. As a sensitivity check, we also generate D_{ij} from a Bernoulli distribution with $p_{ij} = 0.5$ to be classified into the unsusceptible sub-group, and the results are similar (online supplementary Table 7). To generate the recurrent event process $Q_{ij}(t)$, we consider a piece-wise constant baseline hazard specified by quintile grids and $(\lambda_{01}, \lambda_{02}, \lambda_{03}, \lambda_{04}, \lambda_{05})^T = (2, 2.3, 2.1, 2.4, 1.7)^T$ and censor the recurrent events at time $\widetilde{R}_{ij}$. For the unsusceptible sub-group, we set $Q_{ij}(\widetilde{R}_{ij}) = 0$ but without affecting the terminal event time.

Besides the above data-generating process (referred to as DGP1), we also consider additional scenarios to assess the robustness of our method. We first repeat the above data generation process except that we simulate the practice-level random effects from a single-component Normal distribution $\mathcal{N}(0, 0.1^2)$ (DGP2), which represents a simpler case. Second, we consider a scenario where stronger participant-level heterogeneity exists for the terminal event process such that the shape parameter κ_{ij} in the AFT model is randomly generated from a Gamma distribution $\mathcal{G}(1, 1)$ (DGP3). To further assess the model comparison results under a covariate-dependent censoring scheme, we generate the censoring time $C_{ij} \sim \text{Exp}(\text{rate} = \eta^T \mathbf{Z}_{ij})$, where $\eta = (0.5, 0.2, 0.4)^T$; we then set $\widetilde{R}_{ij} = R_{ij} \wedge C_{ij}$, and $\Delta_{ij} = \mathbf{I}[T_{ij} < C_{ij}]$. Under this specification, approximately 30% of the participants exhibit censoring status for the terminal event process (DGP4). In addition, we vary the value of the variance components within the set $\{0.5, 0.75, 1\}$, for both the participant-level and practice-level random effects to determine the sensitivity of inference to different degrees of unobserved heterogeneity (DGP5). For each setting, the results are based on 250 simulated data replicates.

To implement our method, we set $\sigma^2_\zeta = \sigma^2_{\xi_1} = \sigma^2_{\xi_2} = 10$ to set weakly informative Normal priors, $a_\kappa = b_\kappa = 1$ for the Gamma base measure, $\sigma^2 = 1$ for the Normal base measure, and $\mathcal{IG}(1/2, 1/2)$ for the conjugate priors of σ^2_β and σ^2_α. We also consider $G = 5$ to specify the quantile grids in $\mathbf{s}$. In each implementation, multiple chains with randomly generated initial values are run for 10,000 iterations with the first 5,000 as burn-in. It takes roughly 8 hr to complete 10,000 iterations on the high-performance computing cluster where we implemented our suggested model. Our results show that the posterior inference is insensitive to the initial values with a proper mixing for each parameter. In addition to implementing our proposed model (abbreviated as BMZ-$\mathcal{DP}$ for the Bayesian multi-level zero-inflated $\mathcal{DP}$ model), we also consider three variations of BMZ-$\mathcal{DP}$ by simplifying certain model components: (a) BM-$\mathcal{DP}$, which ignores the structural zeros by modelling recurrent event hazard with a single mode Poisson process; (b) BZ-$\mathcal{DP}$, which ignores the multi-level data structure by omitting the practice-level random effects; (c) BMZ, which replaces the non-parametric $\mathcal{DP}$ prior for μ_j with a fully parametric Normal prior and the non-parametric $\mathcal{DP}$ prior for κ_{ij} with a fully parametric Gamma prior, as well as (d) the joint frailty model under a frequentist paradigm implemented in the `R` package `frailtypack` (Rondeau et al., 2012, 2015), which accounts for the multi-level data structure but ignores structural zeros. Of note, BZ-$\mathcal{DP}$ is a variation of the approach developed in Lee et al. (2019) with the addition of the susceptible sub-group, and BMZ is a pure Bayesian parametric implementation. The priors and hyper-parameters for the three Bayesian model variations largely follow those for BMZ-$\mathcal{DP}$; for the frequentist joint frailty model (denoted as frailty), we use the `frailtypack` function which is designed to fit a joint frailty model for clustered data and closest to our setting (with Gamma-distributed participant-level frailty and practice-level frailty). For each method, we summarize the mean or posterior mean, percentage bias (%) relative to the true value, and the 95% confidence or credible intervals (CIs) for the primary parameters of interest, $\boldsymbol{\beta} = (\beta_1, \beta_2, \beta_3)^T$ and $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \alpha_3)^T$ in Table 1 under the primary data-generating process (DGP1). The results under DGP2-DGP5 are summarized in online supplementary Tables 2–5, respectively.

Based on the results in Table 1, the proposed BMZ-$\mathcal{DP}$ model achieves the overall best performance under all levels of sample size, with the smallest or among the smallest percentage bias and highest coverage for all parameters. Under the proposed BMZ-$\mathcal{DP}$ model, the coverage probabilities for the survival process parameters are generally at the 95% level, whereas the coverage probabilities for the recurrent event process parameters are slightly lower than 95% (though still the

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

Table 1. Simulation results under sample sizes 600, 1,200, 1,800 for all the methods summarized by mean or posterior mean (mean), percentage bias [bias (%)] and coverage probability of the 95% credible interval [coverage (%)]

			Recurrent process			Survival process		
N		Method	β_1	β_2	β_3	α_1	α_2	α_3
		BMZ-$\mathcal{DP}$	0.41	0.31	0.25	0.20	0.32	0.40
		BM- $\mathcal{DP}$	−0.04	−0.15	−0.22	0.21	0.31	0.41
	Mean	BZ-$\mathcal{DP}$	0.52	0.45	0.37	0.21	0.32	0.40
		BMZ	0.43	0.35	0.29	0.22	0.33	0.42
		Frailty	−0.08	−0.16	−0.28	0.02	0.06	−0.05
		BMZ-$\mathcal{DP}$	2.20	2.93	26.22	2.21	6.25	0.67
		BM-$\mathcal{DP}$	−111.15	−148.38	−210.02	4.63	2.85	2.57
600	Bias (%)	BZ-$\mathcal{DP}$	29.03	48.87	84.75	6.92	8.15	0.58
		BMZ	9.40	17.41	44.09	10.72	9.69	3.83
		Frailty	−118.88	−153.00	−237.96	−90.21	−80.77	−113.39
		BMZ-$\mathcal{DP}$	88.80	84.80	78.80	96.00	92.80	96.80
		BM-$\mathcal{DP}$	1.60	1.60	1.60	93.60	90.80	75.20
	Coverage (%)	BZ-$\mathcal{DP}$	66.80	54.40	51.60	93.60	89.20	95.80
		BMZ	82.40	79.20	71.00	91.20	85.60	89.60
		Frailty	1.60	1.60	0.00	79.60	74.80	63.20
		BMZ-$\mathcal{DP}$	0.38	0.31	0.23	0.20	0.31	0.40
		BM-$\mathcal{DP}$	−0.05	−0.13	−0.24	0.21	0.31	0.43
	Mean	BZ-$\mathcal{DP}$	0.50	0.44	0.37	0.21	0.32	0.41
		BMZ	0.42	0.35	0.27	0.21	0.33	0.41
		Frailty	−0.08	−0.17	−0.28	0.05	−0.03	−0.05
		BMZ-$\mathcal{DP}$	−4.66	2.93	16.37	0.27	3.52	0.96
		BM-$\mathcal{DP}$	−111.81	−144.49	−220.16	4.59	3.43	6.38
1,200	Bias (%)	BZ-$\mathcal{DP}$	25.30	46.01	84.45	4.03	6.89	1.77
		BMZ	−121.00	16.70	36.13	2.64	8.34	3.57
		Frailty	6.18	−158.67	−238.52	−77.35	−108.81	−113.37
		BMZ-$\mathcal{DP}$	84.80	84.80	83.60	94.80	90.40	95.20
		BM-$\mathcal{DP}$	1.20	1.20	0.00	84.40	90.80	52.40
	Coverage (%)	BZ-$\mathcal{DP}$	58.80	40.80	25.20	87.00	85.20	95.60
		BMZ	80.80	72.40	71.60	82.40	85.60	85.60
		Frailty	1.20	1.20	1.20	71.20	62.80	56.00
		BMZ-$\mathcal{DP}$	0.37	0.31	0.22	0.20	0.30	0.40
		BM-$\mathcal{DP}$	−0.05	−0.13	−0.24	0.21	0.31	0.42
	Mean	BZ-$\mathcal{DP}$	0.51	0.44	0.35	0.22	0.32	0.40
		BMZ	0.43	0.37	0.27	0.21	0.32	0.42
		Frailty	−0.07	−0.16	−0.27	0.05	−0.01	−0.04
		BMZ-$\mathcal{DP}$	−7.87	4.98	13.44	1.35	1.60	1.38
		BM-$\mathcal{DP}$	−114.50	−142.90	−220.50	4.04	2.47	6.03
1,800	Bias (%)	BZ-$\mathcal{DP}$	26.98	46.85	76.99	9.16	5.54	1.15
		BMZ	6.14	21.86	34.30	3.56	7.52	4.66
		Frailty	−117.34	−154.61	−237.30	−77.21	−101.53	−110.87
		BMZ-$\mathcal{DP}$	79.20	85.60	81.40	94.80	93.20	96.80

(*continued*)

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

Table 1. Continued

			Recurrent process			Survival process		
N		Method	β_1	β_2	β_3	α_1	α_2	α_3
	Coverage (%)	BM-$\mathcal{DP}$	1.60	1.20	1.20	85.20	79.60	28.80
		BZ-$\mathcal{DP}$	49.60	30.80	21.60	82.40	86.80	94.80
		BMZ	76.40	68.00	60.80	79.20	86.40	84.60
		Frailty	1.60	1.60	1.60	84.40	72.80	40.40

Note. BMZ-$\mathcal{DP}$ refers to the proposed Bayesian multi-level zero-inflated Dirichlet process model.

closest to 95% among all competing methods). With this level of sample size, the under-coverage for recurrent event model parameters is not unexpected, due to the complicated event structure in the recurrent process including zero inflation as well as nested random effects. When we remove the structural zeros in the data-generating process and fit the BM-$\mathcal{DP}$ model without zero inflation, we observe that the coverage probabilities for the recurrent event process parameters were elevated to be closer to 95% (online supplementary Table 6). In the presence of structural zeros, BMZ-$\mathcal{DP}$ outperforms BM-$\mathcal{DP}$ particularly in uncovering the recurrent process by tracking zero inflation among the population. Meanwhile, we observe substantial estimation bias from both BZ-$\mathcal{DP}$ and BMZ, as they either ignore the between-participant clustering or assume a fully parametric specification of the practice-level random effect. The performance of the frequentist joint frailty model is unsatisfactory in estimating both recurrent and survival model parameters, likely due to its omission for structural zeros and mis-specification of the distributions of the practice-level random effect and residual error of the survival process. For completeness, we have also included results for the association parameters, ξ_1 and ξ_2, as well as the variance component τ^2, under the BMZ-$\mathcal{DP}$ model in online supplementary Table 8. While the percent bias appears reasonably small, the frequentist coverage for these parameters is below 95% given the sample sizes we considered. Finally, the model comparison results for the proposed BMZ-$\mathcal{DP}$ and its variations under alternative data-generating processes (online supplementary Tables 2–5) are qualitatively similar to those in Table 1, suggesting that the proposed model remains to have the best performance even under covariate-dependent censoring schemes or with larger variance components. Notably, the coverage probabilities for recurrent event model parameters under BMZ-$\mathcal{DP}$ increase with a lower censoring rate, a lower zero-inflation rate, or smaller values of the variance components.

Overall, the difference in the performance of our method and the competing approaches across different scenarios we have investigated helps reinforce the necessity of the key components of our proposed model. In the presence of zero inflation, BM-$\mathcal{DP}$ and the frequentist joint frailty model carry the largest estimation bias and lowest coverage for the recurrent process parameters, indicating that accurate inference for recurrent event processes depends critically on adjusting for population heterogeneity regarding the structural zeros. Similarly, BZ-$\mathcal{DP}$ also suffers with large bias and low coverage, suggesting the necessity for explicitly accounting for the multi-level data structure with our complex survival outcomes. Finally, the performance of BMZ demonstrates that it may be critical to consider a flexible non-parametric prior for the practice-level random effect, as the estimation of model parameters can otherwise biased. We have also conducted a similar set of simulations when the baseline hazard for the recurrent process is generated from a Weibull distribution and the conclusions are essentially no different. The details of the simulation setting, model implementation and results are provided in online supplementary Appendix S2.

5 Application to the STRIDE cluster randomized trial

5.1 Strategies to reduce injuries and develop confidence in elders (STRIDE) trial

As stated in Section 1, STRIDE was a pragmatic, parallel-arm CRT aimed at reducing serious falls among community-dwelling older adults. A total of 5,419 participants aged 70 years and older from 86 primary care practices are included in our final analysis. Primary care practices range

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

in size from 10 to 158 participants, with a mean cluster size 63 and coefficient of variation 0.52. The participants were followed for a maximum of 44 months, at which point the survival outcomes were right censored due to study termination. During the follow-up, each occurrence of fall injury and its severity level, along with adverse events including hospitalization and death, were recorded periodically. The descriptive statistics summarizing the number of recurrent adjudicated serious fall injuries, observed death events, and key baseline covariates are presented in online supplementary Table 9 by arm. In the intervention and control primary care practices, the event rate (first serious fall-related injury) was 5.2 and 4.9 per 100 person-years of follow-up, respectively, while in both treatment arms the death rate was lower, at 3.4 per 100 person-years of follow-up. Under cluster randomization, the baseline characteristics were generally balanced between arms, although slightly more white elderly patients appeared in the intervention practices, and patients from the intervention practices tended to have slightly more chronic disease conditions at baseline.

5.2 Model specification and implementation

We implement our Bayesian semi-parametric joint model and the competing approaches described in Section 4 to analyse the STRIDE cluster randomized trial, investigating how the intervention and covariates are associated with serious fall injuries and death among the elderly participants who are 70 years of age or older. We are interested in the recurrent adjudicated serious fall injuries (falls resulting in a fracture, joint dislocation, cut requiring closure, or overnight hospitalization, reported by participants and confirmed by medical records or claims data) and deaths. Besides the intervention, we adjust for several risk factors, including age, sex, race, and number of chronic coexisting conditions (NCD), to study their effect on fall prevention and survival. We implement the proposed BMZ-$\mathcal{DP}$ model and the competing approaches, i.e. BZ-$\mathcal{DP}$, BM-$\mathcal{DP}$, BMZ as well as the joint latent-class model developed by Xu et al. (2021), assuming all the risk factors can potentially impact the recurrent events and survival. In particular, the joint latent-class model does not account for zero inflation nor clustering by practice; however, it allows for latent-class-specific effect on recurrent event and survival. To implement the first four approaches, we set $\sigma^2_\zeta = \sigma^2_{\zeta_1} = \sigma^2_{\zeta_2} = 10$ as prior variances for the regression parameters. For the $\mathcal{DP}$ prior of the practice-specific random effects, we set $\sigma^2 = 1$ for the Normal base measure. We further set $a_\kappa = b_\kappa = 1$ for the Gamma base measure for $\mathcal{DP}$ prior of the participant-specific shape parameter κ_{ij}. We assign $\mathcal{IG}(1/2,\ 1/2)$ as conjugate hyper-priors for the prior variances σ^2_β, σ^2_α and τ_l. For the recurrent event baseline hazard function, we consider $G = 5$ quantile grids $\mathbf{s}$ and an improper prior for each element of $\boldsymbol{\lambda}$ as indicated in Section 3.1. Other hyper-parameter specifications closely follow those in Section 4. We consider a logistic model to represent p_{ij} to allow for dependence on treatment, sex and race, which accommodates the potential effect of intervention and baseline risk factors on the individual susceptibility status. For the joint latent class model, we consider the default implementation and mixture of finite mixtures hierarchical prior for latent-class probability explained in Xu et al. (2021). For each model, under random initials, we run MCMC for 20,000 iterations with the first 10,000 as burn-in; it takes approximately 24 hr to complete 20,000 iterations for the proposed BMZ-$\mathcal{DP}$ model. The trace plots for several key parameters are provided in the online supplementary Figure 2. Finally, the posterior results for each method are summarized in Table 2. The joint frailty model under a frequentist paradigm (based on the `frailtypack`) is not discussed further because the model did not converge after running for 48 hr.

5.3 Results from the proposed model

We first investigate the impact of the intervention and risk factors on recurrent fall injuries. Under the proposed BMZ-$\mathcal{DP}$ model, NCD and age appear to be associated with fall injury intensity, and their corresponding 95% credible intervals exclude zero. Exponentiating the posterior means of the model parameters, we find that one additional chronic condition multiplies the serious fall rate by 1.15 and a 1-year increase in age reduces injury intensity by around 5%. The latter result provides seemingly counter-intuitive evidence since it was originally believed that older age increases the risk for fall injury. However, because the population recruited in our study are 70 years of age or older, it is also likely that a further increase in age could start to prevent them from potential triggers for serious fall such as exercise or intensive movement. In addition, female

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

Table 2. Posterior inference results for parameters in the recurrent and survival processes under different methods for the analysis of the STRIDE study

	BMZ-$\mathcal{DP}$	BM-$\mathcal{DP}$	BZ-$\mathcal{DP}$	BMZ	Xu et al. (2021)
Recurrent process					
Intervention	−0.08 (−0.34, 0.16)	−0.25 (−0.48, 0.01)	−0.02 (−0.26, 0.22)	−0.12 (−0.28, 0.04)	−0.34 (−0.57, −0.11)
NCD	0.14 (0.06, 0.22)	0.14 (0.08, 0.18)	0.14 (0.07, 0.22)	0.10 (−0.04, 0.24)	−0.22 (−0.29, −0.15)
Age	−0.05 (−0.06, −0.04)	−0.05 (−0.06, −0.04)	0.01 (−0.01, 0.01)	−0.04 (−0.09, −0.04)	−0.02 (−0.24, 0.20)
Sex (Female)	0.08 (−0.12, 0.30)	0.03 (−0.12, 0.20)	0.18 (−0.01, 0.36)	−0.02 (−0.18, 0.14)	0.08 (−0.02, 0.18)
Race (White)	0.18 (−0.18, 0.47)	0.17 (−0.13, 0.44)	0.81 (0.38, 1.29)	0.89 (0.48, 1.30)	0.11 (−0.43, 0.65)
Survival process					
Intervention	0.02 (−0.23, 0.24)	−0.05 (−0.20, 0.11)	−0.04 (−0.38, 0.28)	0.11 (−0.16, 0.38)	−0.10 (−0.34, 0.14)
NCD	−0.16 (−0.26, −0.10)	−0.15 (−0.21, −0.12)	−0.19 (−0.31, −0.09)	−0.18 (−0.30, −0.05)	−0.26 (−0.33, −0.19)
Age	−0.03 (−0.06, −0.01)	−0.04 (−0.05, −0.03)	−0.04 (−0.07, −0.01)	0.02 (0.00, 0.04)	0.10 (−0.03, 0.23)
Sex (Female)	0.50 (0.22, 0.81)	0.45 (0.32, 0.61)	0.65 (0.29, 1.16)	0.97 (0.50, 1.50)	1.17 (0.85, 1.49)
Race (White)	0.06 (−0.45, 0.52)	0.09 (−0.08, 0.30)	0.04 (−0.69, 0.62)	0.12 (−0.33, 0.67)	0.19 (0.07, 0.31)
LPML	−1534.12	−3585.48	−1624.85	−2092.51	−3466.78

Note. BMZ-$\mathcal{DP}$ refers to the proposed Bayesian multi-level zero-inflated Dirichlet process model; NCD is the abbreviation for number of chronic coexisting conditions; LPML is the abbreviation for log pseudo-marginal likelihood.

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

and white patients are more likely to experience falls than other sub-groups. On the hazard scale, the intervention appears to have a small effect on reducing the risks for recurrent fall injuries. For the terminal event—death, NCD, age, and sex remain significant predictors, with their 95% credible intervals excluding zero. An increase in NCD or age leads to shorter survival times, as expected, but female and white patients appear to have longer survival. Consistent with the previous analysis focusing on the first occurrence of fall injury (Bhasin et al., 2020; Chen & Li, 2022; Li et al., 2022), the intervention reduces the rate of recurrent falls as well as benefits survival experiences, but 95% credible intervals of the intervention effect parameter in both processes include zero.

To further interpret the treatment effect on the risk of fall injury, we borrow the counterfactual outcome framework to investigate the participant-average treatment effect (Kahan et al., 2023). Specifically, suppose an individual remains alive at time t, the recurrence rate for that individual at time t is given by $\mu_{ij}(t) = E[N_{ij}(t)] = p_{ij} \int_0^t \lambda_{ij}(u) H_{ij}(u) du$, where we recall p_{ij} is the participant-specific probability to be classified into the unsusceptible sub-group (therefore not a structural zero), $\lambda_{ij}(t)$ is the participant-specific hazard function for the recurrent event, and $H_{ij}(t)$ is the survival function for the terminal event; similar definition has also been discussed in Xu et al. (2021) in the absence of zero inflation. Then the counterfactual recurrence rate had an individual received the intervention (possibly contrary to fact) can be expressed by $\mu_{ij}(1, t) = p_{ij}(do(\text{Treat} = 1)) \int_0^t \lambda_{ij}(u \mid do(\text{Treat} = 1)) H_{ij}(u \mid do(\text{Treat} = 1)) du$, where we use the do-calculus notation to indicate the critical step of setting the treatment variable to be 1 when computing the probability to be in the unsusceptible sub-group, recurrent event hazard, and terminal event survival functions (Pearl, 2000). Analogously we define $\mu_{ij}(0, t) = p_{ij}(do(\text{Treat} = 0)) \int_0^t \lambda_{ij}(u \mid do(\text{Treat} = 0)) H_{ij}(u \mid do(\text{Treat} = 0)) du$. Therefore, the participant-average treatment effect for fall injury at time t, on the rate difference scale and rate ratio scale, can be expressed as

$$\text{rate difference}(t) = \frac{\sum_{j=1}^{J} \sum_{i=1}^{N_j} \{\mu_{ij}(1, t) - \mu_{ij}(0, t)\}}{\sum_{j=1}^{J} N_j} \quad \text{rate ratio}(t) = \frac{\sum_{j=1}^{J} \sum_{i=1}^{N_j} \mu_{ij}(1, t)}{\sum_{j=1}^{J} \sum_{i=1}^{N_j} \mu_{ij}(0, t)}.$$

Figure 3 plots the posterior mean and 95% credible intervals on the counterfactual recurrent rates, rate difference, and rate ratio as a function of follow-up time in years. Panel (a) suggests a very mild treatment effect since intervention leads to a slightly lower counterfactual recurrence rate. Indeed, from panels (b) and (c) the recurrence rate difference at year 3 is around −0.029, suggesting around 29 falls prevented per 1,000 patients; the recurrence rate ratio at year 3 is around 0.80. Interestingly, the 95% point-wise credible bands just exclude null for both the rate difference and rate ratio effect measures since the start of follow-up until approximately year 2, but include the null from that time onward. It is important to note that these counterfactual treatment effect quantities are not identical to the treatment effect parameter due to non-collapsibility, and may be more interpretable when the interest lies in measuring the population impact of intervention due to switching from usual care to the fall injury prevention programme.

Finally, we examine the structural zero probabilities and cluster-specific random effects. We summarize the marginal posterior inclusion probability for $D_{ij} = 1$ over subjects who may be considered naturally unsusceptible for the recurrent fall events in panel (a) of Figure 4. With a 0.5 cut-off (Barbieri & Berger, 2004), there could be a substantial fraction of participants who may not be susceptible to serious fall injuries during the study period. We present the distribution of those inclusion probabilities within different practices in panel (b) of Figure 4. Almost all practices include a large amount of patients from the unsusceptible sub-group, and the inclusion probability varies both within and between practices. This visualization can help identify practices with substantially more unsusceptible patients for falls, although the exact scientific mechanism for unsusceptibility remains to be further studied. We also provide the posterior mean along with the 95% credible intervals obtained from the inference of each μ_j in panel (c) of Figure 4, where the practices are ordered by their point estimates. Clearly, the cluster-level frailties show substantial heterogeneity across practices, are all negative, and are all significantly different from zero. Online supplementary Figure 3a additionally presents the posterior mean of the density estimates of μ_j, overlaid by a Normal density with matching moments. This comparison indicates that the density of μ_j may deviate from a Normal density and there can be some benefits in leveraging

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

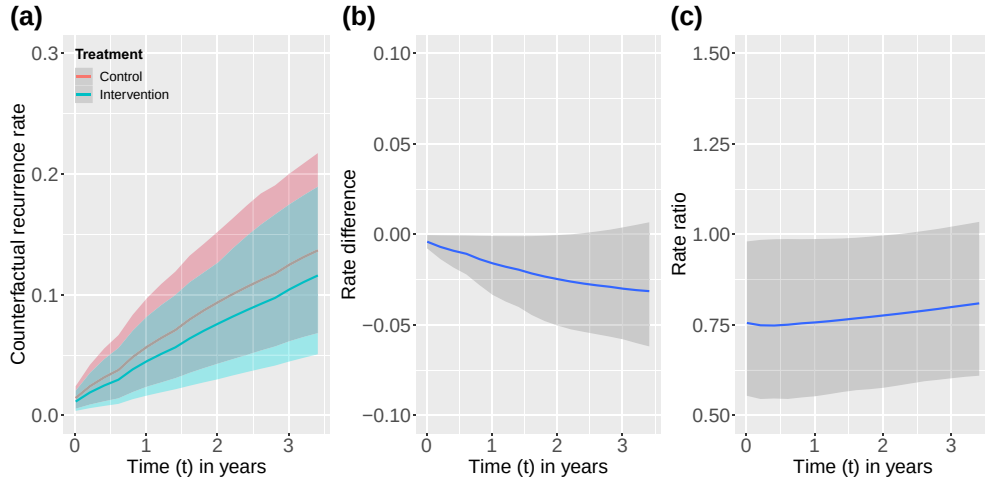

Figure 3. (a) Posterior mean and 95% point-wise credible bands for counterfactual recurrence rates over time. (b) Posterior mean and 95% point-wise credible bands for the counterfactual recurrence rate difference over time. (c) Posterior mean and 95% point-wise credible bands for the counterfactual recurrence rate ratio over time.

the non-parametric prior to capture the potential non-normality. In addition, online supplementary Figure 3b presents the posterior mean of the variance components τ_j^2, grouped by each practice. It is evident that there exists substantial variability within and across practices on participant-level frailty variance. Such information on between-practice heterogeneity and between-participant heterogeneity, however, would not be available if the analysis model ignores the multi-level feature of the STRIDE data.

5.4 Model comparison and sensitivity to priors

Table 2 additionally summarizes results from the competing models. Although results from the models without zero inflation (BM-$\mathcal{DP}$), multi-level data structure (BZ-$\mathcal{DP}$), or both DP priors (BMZ) are generally consistent with those from BMZ-$\mathcal{DP}$, we notice differences in the effect estimates and credible intervals. For example, the main conclusions based on BM-$\mathcal{DP}$ align with those from BMZ-$\mathcal{DP}$, but failure to account for the structural zeros tends to exaggerate the intervention effects for the serious fall intensity and the survival. The results from BZ-$\mathcal{DP}$ did not identify age as an important predictor for recurrent fall injury, and the BMZ model results also reveal larger intervention effect estimates than the proposed model. Finally, the joint latent-class model ignores both the zero inflation and clustering, and tends to overestimate the covariate and intervention effects for both recurrence and survival processes.

To evaluate model fit, we consider model validation diagnostics for recurrent events data using the conditional predictive ordinate (CPO); the CPO refers to the conditional predictive ordinate used for detecting surprising observations over subjects (Pettit, 1990) and has similarly been operationalized in Sinha et al. (2008) for non-clustered recurrent event data. By plotting CPO_{ij} against the observed follow-up time $\widetilde{R}_{ij}$, we can visually assess whether the posterior analysis has any unusually low prediction capabilities for certain values of the observed survival time. Online supplementary Figure 4 indicates lack of systematic association between CPO_{ij} and $\widetilde{R}_{ij}$ under the proposed model, therefore suggesting no strong evidence against the model adequacy. In addition, we compare the proposed model with the competing models by calculating the log pseudo-marginal likelihood (LPML) based on leave-one-out-cross-validation as $\text{LPML}=\sum_{j=1}^{J}\sum_{i=1}^{N_j}\log(\text{CPO}_{ij})$. This is a commonly used metric in Bayesian survival analyses to compare model performance, with a larger value suggesting a better fit to the data. As shown in Table 2, our proposed model has the largest LPML and therefore demonstrates the best fit to the analysis of STRIDE data among the competing models.

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

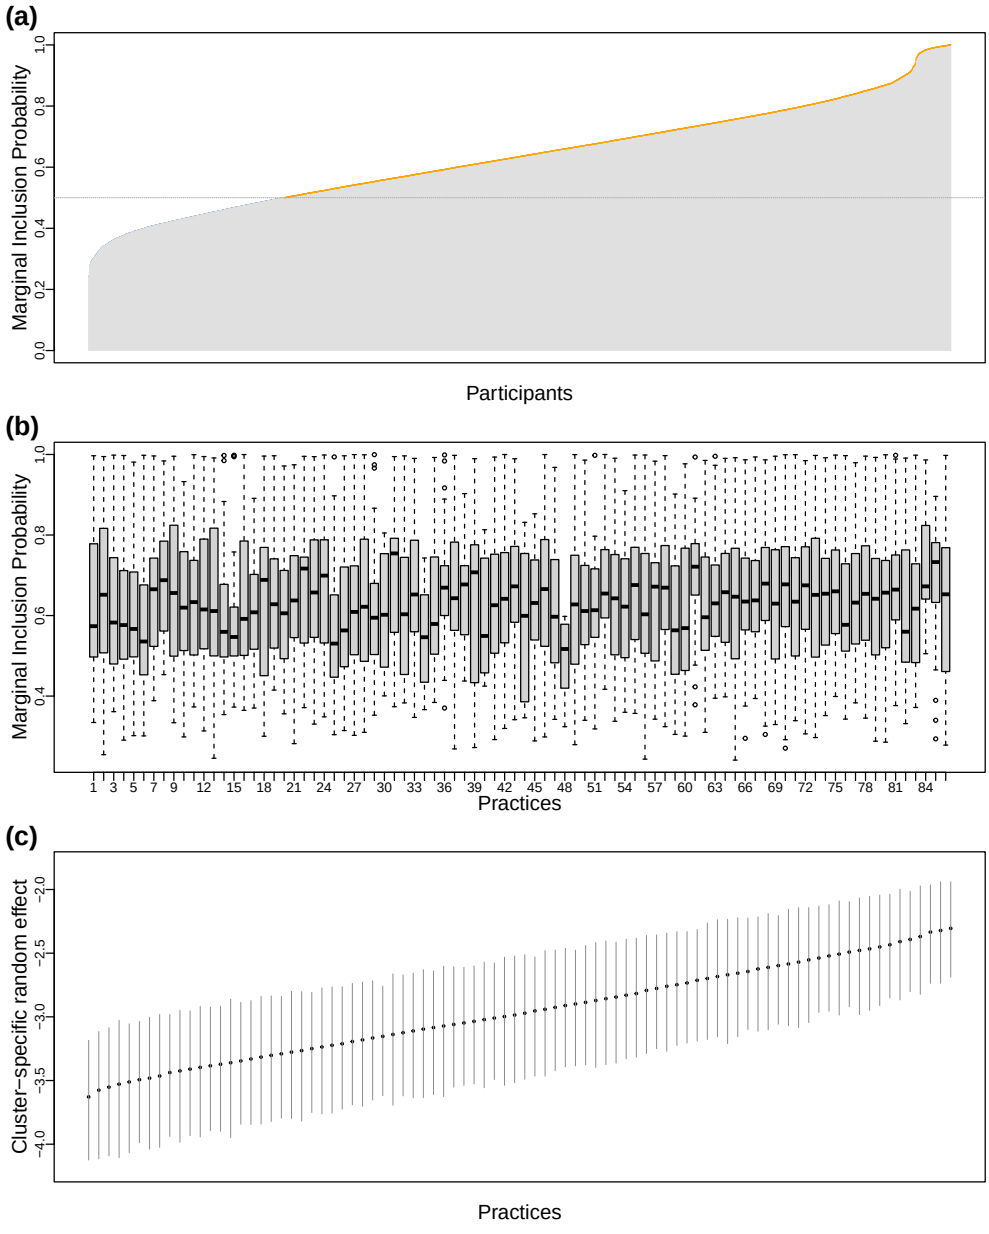

Figure 4. Analysis results on structural zero probabilities and cluster-specific random effects for STRIDE. (a) Marginal posterior inclusion probability of $D_{ij} = 1$ ordered from smallest to largest. (b) Distribution of marginal posterior inclusion probability of $D_{ij} = 1$ within different practices. (c) Posterior inference of μ_j ordered by posterior mean.

Finally, we perform a series of sensitivity analyses to evaluate how much the posterior inference results would change according to alternative prior specifications. We independently check the estimation results under the following scenarios. (a) The hyper-prior for the marginal inclusion indicator is uniformly set to $p_{ij} = 0.5$ without dependence on covariates. (b) The hyper-priors for the prior variances associated with the regression coefficients are set to be $\sigma_\beta^2 \sim \mathcal{IG}(0.01, 0.01)$ and $\sigma_\alpha^2 \sim \mathcal{IG}(0.01, 0.01)$ instead of $\mathcal{IG}(1/2, 1/2)$, so they are further less informative. (c) The number of quantile grids s associated with the baseline hazard in the recurrent event process is expanded to $G = 8$ (rather than $G = 5$ in our original implementation), and we specify a proper uniform prior

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

(0, 100) for each element within $\boldsymbol{\lambda}$ (rather than the improper prior in our original implementation). (d) We set $\sigma^2_{\zeta} = \sigma^2_{\xi_1} = \sigma^2_{\xi_2} = 3$ to give slightly more informative priors for the frailty association parameters and the logistic structural-zero regression parameters. (e) We set the hyper-prior for the practice-specific variance parameter to be $\tau_j \sim \mathcal{IG}(0.01, 0.01)$ to make it less informative. The posterior inference results for each specification are summarized in Table 3, and we observe that the results are generally stable across different specifications. We further obtain the LPML for each analysis and found that only specification (e) produces slightly larger LPML than (but very close to) our primary implementation in Section 5.2. The posterior inference results for the regression coefficients are not substantially different between these two specifications.

6 Discussion

In this paper, we propose a new Bayesian semi-parametric joint model framework to simultaneously characterize the recurrent event process and survival in the presence of clustering and potential zero inflations. To accommodate the between-participant clustering commonly seen in pragmatic clinical trials, we introduce hierarchical random effects at the participant and practice levels, both of which bridge the recurrent and survival processes. To further relax the parametric assumptions, we specify separate non-parametric realizations for the baseline hazard, practice-level random effect and terminal event survival function, to enhance the model robustness compared with the existing alternative modelling strategies. We also demonstrate the necessity of each component within our joint modelling framework and develop MCMC algorithms to enable posterior inference for all model parameters. Through simulations and our empirical application to a recent pragmatic cluster randomized trial, we demonstrate the advantage of our proposed method.

To account for zero inflation, we define the structural zeros as those contributed by participants who are unsusceptible to the recurrent event during the study period, and we impose a two-component mixture to represent the recurrent event hazard. This definition is similar to the existing cure model literature (Rondeau et al., 2013; Yu et al., 2004) with straightforward interpretation of unsusceptibility. When recurrent events are subject to a terminal event, alternative definitions of unsusceptibility exist. For example, one can either define the unsusceptibility status based on the unsusceptibility status of recurrent events or further differentiate the structural zero from the random zero based on the terminal event. The former definition has been adopted in Xu et al. (2018) and Han et al. (2020) as well as our current model specification, given study participants should always be susceptible to typical terminal events such as death, particularly in studies with an elderly population and a longer follow-up. Liu et al. (2016) provided a brief explanation on the difference of these two definitions of unsusceptibility driven by a specific application. In certain applications, our modelling framework can be adapted under an alternative definition of unsusceptibility incorporating the terminal event.

A notable feature of the proposed Bayesian joint model is that the random effects are shared between the recurrent event and terminal event sub-models, through some latent association parameters. This type of specification has also been considered in the context of joint longitudinal and survival modelling for multi-level data (Brilleman et al., 2019) and is expected to provide more efficient inference for the model parameters when, in reality, the random effects between sub-models do overlap. However, when the sources of additional variation are distinct for the recurrent event and terminal event processes, the validity of the proposed model requires additional investigation, and it is possible that ignoring distinct random-effect structures in different sub-models will lead to bias and hence compromise the validity of the model estimates. As a potential improvement, one may expand the joint model to additionally incorporate independent random slopes at the practice level and the participant level, thus accounting for sources of variation that may be unique to each sub-model. However, given the complexity of this more elaborate model structure, it is yet to be studied whether the proposed MCMC procedure can return numerically stable estimates in finite samples within a reasonable computation time. Alternatively, a promising approach to alleviate the computational intensity is by employing variational inference algorithm (Attias, 1999). Rather than sampling from the posterior distribution, this approach utilizes a variational distribution to approximate the full posterior and is often computationally much faster. However, the variational inference provides only an approximation to the true posterior and can introduce bias depending on the chosen variational family. A comparison between variational

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

Table 3. Posterior inference for parameters under the proposed model but with different prior specifications

	Section 5.2	(a)	(b)	(c)	(d)	(e)
Recurrent process						
Intervention	−0.08 (−0.34, 0.16)	−0.12 (−0.28, 0.04)	−0.03 (−0.22, 0.17)	−0.01 (−0.25, 0.23)	−0.05 (−0.25, 0.17)	−0.06 (−0.38, 0.20)
NCD	0.14 (0.06, 0.22)	0.10 (−0.04, 0.24)	0.11 (0.05, 0.19)	0.12 (0.04, 0.19)	0.11 (0.04, 0.19)	0.13 (0.05, 0.23)
Age	−0.05 (−0.06, −0.04)	−0.04 (−0.09, −0.04)	−0.05 (−0.06, −0.04)	−0.05 (−0.06, −0.04)	−0.05 (−0.07, −0.04)	−0.04 (−0.05, −0.04)
Sex (Female)	0.08 (−0.12, 0.30)	−0.02 (−0.18, 0.14)	0.06 (−0.11, 0.25)	0.02 (−0.17, 0.24)	0.05 (−0.18, 0.26)	0.07 (−0.13, 0.26)
Race (White)	0.18 (−0.18, 0.47)	0.89 (0.48, 1.30)	0.21 (−0.21, 0.56)	0.31 (−0.11, 0.69)	0.16 (−0.27, 0.50)	0.26 (−0.13, 0.82)
Survival process						
Intervention	0.02 (−0.23, 0.24)	0.11 (−0.16, 0.38)	−0.03 (−0.34, 0.29)	0.04 (−0.29, 0.33)	−0.01 (−0.29, 0.23)	0.05 (−0.22, 0.34)
NCD	−0.16 (−0.26, −0.10)	−0.18 (−0.30, −0.05)	−0.18 (−0.28, −0.07)	−0.17 (−0.28, −0.08)	−0.15 (−0.24, −0.08)	−0.14 (−0.23, −0.07)
Age	−0.03 (−0.06, −0.01)	0.02 (0.00, 0.04)	−0.03 (−0.05, −0.01)	−0.03 (−0.05, −0.02)	−0.01 (−0.02, 0.01)	−0.01 (−0.02, −0.00)
Sex (Female)	0.50 (0.22, 0.81)	0.97 (0.50, 1.50)	0.62 (0.30, 0.98)	0.58 (0.27, 0.95)	0.52 (0.25, 0.81)	0.44 (0.20, 0.72)
Race (White)	0.06 (−0.45, 0.52)	0.12 (−0.33, 0.67)	0.00 (−0.54, 0.49)	0.31 (−0.11, 0.69)	0.09 (−0.36, 0.53)	−0.07 (−0.61, 0.38)
LPML	−1534.12	−2092.51	−1580.62	−1554.66	−1689.87	−1527.97

Note. The second column represents the primary implementation in Section 5.2, and the remaining columns correspond to sensitivity analyses under the following scenarios: (a) prior for marginal inclusion independent of covariates; (b) alternative hyper-priors for the prior variances of the regression coefficients; (c) expanding quantile grids for the recurrent event baseline hazard; (d) slightly more informative priors for frailty coefficients and logistic structural-zero regression coefficients; (e) less informative hyper-prior for practice-specific variance parameters. LPML = log pseudo-marginal likelihood; NCD = number of chronic coexisting conditions.

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

inference and MCMC inference for complex joint survival models with multi-level data will be left for future work.

One potential limitation of our modelling strategy is that we have not distinguished the survival hazard between subjects who experience recurrent event(s) and those who directly move to the terminal event. The impact upon the survival function from the previous recurrent event occurrence is primarily controlled by the shared frailty terms, γ_{ij} and μ_j. An alternative modelling strategy can be based on the multi-state model (Lee et al., 2016; Li et al., 2022), where one can more explicitly split the event occurrence paths into (a) recurrent events only, (b) recurrent events followed by the terminal event, and (c) terminal event only, and characterize state-specific hazard functions for each of them in conjunction with different random effects. Under multi-state modelling, we may be able to capture the influence from both fixed and random effects on different hazard functions, potentially allowing for richer information extraction. However, a practical issue of multi-state modelling in our case comes from the growing number of unknown parameters, since we would need an additional set of coefficients as well as more random effects. This could induce computational and inferential challenges, especially with a small sample size, and merits additional investigation. Another potential limitation is that we have not distinguished the terminal event survival processes for the susceptible and unsusceptible patient sub-populations. In STRIDE, It is possible that patients who are unsusceptible to recurrent fall injuries are more likely to survive until the end of the study that those who are susceptible. Generally, this assumption is challenging to test empirically from data because in our model formulation these two sub-populations are not fully observed from the study sample alone. A possible improvement of the proposed model is to modify the likelihood in Section 3.2 to allow the terminal event survival function and density function to further depend on D_{ij}. However, the associated posterior inference can become substantially more challenging with our non-parametric prior specifications at multiple levels. Similarly, a primary care practice including more than usual unsusceptible patients may be expected to have lower risk of death, and it would be interesting to further control for $N_j^{-1}\sum_{i=1}^{N_j} D_{ij}$ in the terminal event model. Lastly, we have defined the latent indicator D_{ij} based on the initial recurrence without accounting for the potential switch in susceptibility status after this initial occurrence. As a result, after a participant experiences their first fall injury, he or she will continue to be categorized as susceptible, regardless of their subsequent health status or the length of follow-up. In the context of fall injury for the elderly, participants who have experienced an initial fall generally have a higher risk for subsequent falls in comparison to others (O'Loughlin et al., 1993; Tinetti et al., 1988), and fall prevention programmes are generally not considered as a cure for fall injury (although it holds promises to reduce the rates of fall injury); therefore, it is unlikely for participants to switch susceptibility status after the first fall injury in our context. However, in other applications, one may expect the susceptibility status to vary over time, which was not addressed in the proposed model. Alternatively, the multi-state model allows for changes in susceptibility status over time, reflecting the natural progression of recurrent fall injuries and cumulative intervention effects that can eventually prevent recurrences. By accounting for the transitions between susceptibility status over time, one may enhance our understanding of the dynamics of recurrences and gain insights into the optimal timing of future interventions. On the other hand, modelling time-varying transition between susceptibility status for each participant can come with a substantial computational demand. Although these relevant extensions are promising for understanding the complex event processes in STRIDE, the identifiability and efficient posterior inference based on these extensions are beyond the scope of this article and will be pursued in future research.

Conflict of interest: None declared.

Funding

This work was supported by National Institutes of Health (NIH) grants RF1AG081413, R01EB034720, RF1AG068191, R01HL168202, P30AG066508, UL1 TR0001863, as well as a Patient-Centered Outcomes Research Institute Award (PCORI Award ME-2020C3-21072). The STRIDE study was funded by the Patient-Centered Outcomes Research Institute (PCORI), with additional support from the National Institute on Aging at NIH (U01AG048270) and the Claude D. Pepper Older Americans Independence Center at Yale School of Medicine

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

(P30AG021342). The statements presented in this article are solely the responsibility of the authors and do not necessarily represent the views of NIH, PCORI, or its Board of Governors or Methodology Committee. The authors are grateful to the Associate Editor and two anonymous referees for their constructive suggestions, which have led to substantial improvement in our work.

Data availability

The de-identified data that went into the analyses, tables, and figures in this manuscript were not publicly available at the time of writing. The data will be made available through the NIA data repository (NIA AgingResearchBiobank) per agreement with the National Institute on Aging, and may be requested from the NIA Biorepository. Until the data are available through the NIA Biorepository, the requests for data access can be sent to Dr. Erich Greene at erich.greene@yale.edu and will be evaluated on a case-by-case basis. Sample code to implement the proposed Bayesian models is available at https://github.com/xt83/Bayesian_semi_parametric_inference_for_clustered_recurrent_event.

Supplementary material

Supplementary material is available online at *Journal of the Royal Statistical Society: Series C*.

References

Attias H. (1999). A variational Bayesian framework for graphical models. In S. A. Solla, T. K. Leen, & K.-R. Müller (Eds.), *Advances in neural information processing systems* (pp. 209–215). Cambridge: MIT Press.

Barbieri M. M., & Berger J. O. (2004). Optimal predictive model selection. *The Annals of Statistics*, *32*(3), 870–897. https://doi.org/10.1214/009053604000000238

Bhasin S., Gill T. M., Reuben D. B., Latham N. K., Ganz D. A., Greene E. J., Dziura J., Basaria S., Gurwitz J. H., Dykes P. C., McMahon S., Storer T.W., Gazarian P., Miller M.E., Travison T.G., Esserman D., Carnie M.B., Goehring L., Fagan M., …Alexander N. (2020). A randomized trial of a multifactorial strategy to prevent serious fall injuries. *New England Journal of Medicine*, *383*, 129–140. https://doi.org/10.1056/NEJMoa2002183

Brilleman S. L., Crowther M. J., Moreno-Betancur M., Novik J. B., Dunyak J., Al-Huniti N., Fox R., Hammerbacher J., & Wolfe R. (2019). Joint longitudinal and time-to-event models for multilevel hierarchical data. *Statistical Methods in Medical Research*, *28*(12), 3502–3515. https://doi.org/10.1177/0962280218808821

Chen X., & Li F. (2022). Finite-sample adjustments in variance estimators for clustered competing risks regression. *Statistics in Medicine*, *41*(14), 2645–2664. https://doi.org/10.1002/sim.v41.14

Choi N. G., Choi B. Y., DiNitto D. M., Marti C. N., & Kunik M. E. (2019). Fall-related emergency department visits and hospitalizations among community-dwelling older adults: Examination of health problems and injury characteristics. *BMC Geriatrics*, *19*(1), 1–10. https://doi.org/10.1186/s12877-019-1329-2

Cooner F., Banerjee S., Carlin B. P., & Sinha D. (2007). Flexible cure rate modeling under latent activation schemes. *Journal of the American Statistical Association*, *102*(478), 560–572. https://doi.org/10.1198/016214507000000112

Ferguson T. S. (1973). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, *1*(2), 209–230. https://doi.org/10.1214/aos/1176342360

Fine J. P., Jiang H., & Chappell R. (2001). On semi-competing risks data. *Biometrika*, *88*(4), 907–919. https://doi.org/10.1093/biomet/88.4.907

Ganz D. A., Siu A. L., Magaziner J., Latham N. K., Travison T. G., Lorenze N. P., Lu C., Wang R., Greene E. J., Stowe C. L., Harvin L. N., Araujo K. L. B., Gurwitz J. H., Agrawal Y., Correo-De-Araujo R., Peduzzi P., & Gill T. M. (2019). Protocol for serious fall injury adjudication in the strategies to reduce injuries and develop confidence in elders (STRIDE) study. *Injury Epidemiology*, *6*(1), 1–8. https://doi.org/10.1186/s40621-019-0190-2

Gasparini A., Clements M. S., Abrams K. R., & Crowther M. J. (2019). Impact of model misspecification in shared frailty survival models. *Statistics in Medicine*, *38*(23), 4477–4502. https://doi.org/10.1002/sim.v38.23

Gelman A., & Rubin D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, *7*, 457–472. https://doi.org/10.1214/ss/1177011136

Han D., Su X., Sun L., Zhang Z., & Liu L. (2020). Variable selection in joint frailty models of recurrent and terminal events. *Biometrics*, *76*(4), 1330–1339. https://doi.org/10.1111/biom.v76.4

Hopewell S., Adedire O., Copsey B. J., Boniface G. J., Sherrington C., Clemson L., Close J. C., & Lamb S. E. (2018). Multifactorial and multiple component interventions for preventing falls in older people living in the community. *Cochrane Database of Systematic Reviews*, 1–261. https://doi.org/10.1002/14651858.CD012221

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

Ishwaran H., & James L. F. (2001). Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association*, *96*(453), 161–173. https://doi.org/10.1198/016214501750332758
Ivers N., Taljaard M., Dixon S., Bennett C., McRae A., Taleban J., Skea Z., Brehaut J., Boruch R., & Eccles M. (2011). Impact of CONSORT extension for cluster randomised trials on quality of reporting and study methodology: Review of random sample of 300 trials, 2000–8. *British Medical Journal*, *343*, d5886. https://doi.org/10.1136/bmj.d5886
Joseph B., & Robert P. G. (1952). Survival curve for cancer patients following treatment. *Journal of the American Statistical Association*, *47*(259), 501–515. https://doi.org/10.1080/01621459.1952.10501187
Jung T. H., Peduzzi P., Allore H., Kyriakides T. C., & Esserman D. (2019). A joint model for recurrent events and a semi-competing risk in the presence of multi-level clustering. *Statistical Methods in Medical Research*, *28*(10–11), 2897–2911. https://doi.org/10.1177/0962280218790107
Kahan B. C., Li F., Copas A. J., & Harhay M. O. (2023). Estimands in cluster-randomized trials: Choosing analyses that answer the right question. *International Journal of Epidemiology*, *52*(1), 107–118. https://doi.org/10.1093/ije/dyac131
Kalbfleisch J. D., Schaubel D. E., Ye Y., & Gong Q. (2013). An estimating function approach to the analysis of recurrent and terminal events. *Biometrics*, *69*(2), 366–374. https://doi.org/10.1111/biom.v69.2
Kim Y.-J. (2021). Joint model for bivariate zero-inflated recurrent event data with terminal events. *Journal of Applied Statistics*, *48*(4), 738–749. https://doi.org/10.1080/02664763.2020.1744539
Lancaster T., & Intrator O. (1998). Panel data with survival: Hospitalization of HIV-positive patients. *Journal of the American Statistical Association*, *93*(441), 46–53. https://doi.org/10.1080/01621459.1998.10474086
Lee J., Thall P. F., & Lin S. H. (2019). Bayesian semiparametric joint regression analysis of recurrent adverse events and survival in esophageal cancer patients. *The Annals of Applied Statistics*, *13*(1), 221. https://doi.org/10.1214/18-AOAS1182
Lee K. H., Dominici F., Schrag D., & Haneuse S. (2016). Hierarchical models for semicompeting risks data with application to quality of end-of-life care for pancreatic cancer. *Journal of the American Statistical Association*, *111*(515), 1075–1095. https://doi.org/10.1080/01621459.2016.1164052
Li F., Lu W., Wang Y., Pan Z., Greene E. J., Meng G., Meng C., Blaha O. B., Zhao Y., Peduzzi P. N., et al. (2022). A comparison of analytical strategies for cluster randomized trials with survival outcomes in the presence of competing risks. *Statistical Methods in Medical Research*, *31*(7), 1224–1241. https://doi.org/10.1177/09622802221085080
Li F., Zhang T., Wang Q., Gonzalez M. Z., Maresh E. L., & Coan J. A. (2015). Spatial Bayesian variable selection and grouping for high-dimensional scalar-on-image regression. *The Annals of Applied Statistics*, *9*, 687–713. https://doi.org/10.1214/15-AOAS818
Li Z., Chinchilli V. M., & Wang M. (2019). A Bayesian joint model of recurrent events and a terminal event. *Biometrical Journal*, *61*(1), 187–202. https://doi.org/10.1002/bimj.v61.1
Li Z., Chinchilli V. M., & Wang M. (2020). A time-varying Bayesian joint hierarchical copula model for analysing recurrent events and a terminal event: An application to the cardiovascular health study. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, *69*(1), 151–166. https://doi.org/10.1111/rssc.12382
Lin L.-A., Luo S., Chen B. E., & Davis B. R. (2017). Bayesian analysis of multi-type recurrent events and dependent termination with nonparametric covariate functions. *Statistical Methods in Medical Research*, *26*(6), 2869–2884. https://doi.org/10.1177/0962280215613378
Liu L., Huang X., Yaroshinsky A., & Cormier J. N. (2016). Joint frailty models for zero-inflated recurrent events in the presence of a terminal event. *Biometrics*, *72*(1), 204–214. https://doi.org/10.1111/biom.12376
Liu L., Wolfe R. A., & Huang X. (2004). Shared frailty models for recurrent events and a terminal event. *Biometrics*, *60*(3), 747–756. https://doi.org/10.1111/biom.2004.60.issue-3
McKeague I. W., & Tighiouart M. (2000). Bayesian estimators for conditional hazard functions. *Biometrics*, *56*(4), 1007–1015. https://doi.org/10.1111/biom.2000.56.issue-4
Mitchell R. J., Watson W. L., Milat A., Chung A. Z., & Lord S. (2013). Health and lifestyle risk factors for falls in a large population-based sample of older people in Australia. *Journal of Safety Research*, *45*, 7–13. https://doi.org/10.1016/j.jsr.2012.11.005
O'Loughlin J., Robitaille Y., Boivin J.-F., & Suissa S. (1993). Incidence of and risk factors for falls and injurious falls among the community-dwelling elderly. *American Journal of Epidemiology*, *137*(3), 342–354. https://doi.org/10.1093/oxfordjournals.aje.a116681
Paulon G., De Iorio M., Guglielmi A., & Ieva F. (2020). Joint modeling of recurrent events and survival: A Bayesian non-parametric approach. *Biostatistics*, *21*, 1–14. https://doi.org/10.1093/biostatistics/kxy026
Pearl J. (2000). *Causality: Models, reasoning, and inference*. Cambridge University Press.
Pettit L. (1990). The conditional predictive ordinate for the normal distribution. *Journal of the Royal Statistical Society: Series B (Methodological)*, *52*(1), 175–184. https://doi.org/10.1111/j.2517-6161.1990.tb01780.x
Rietdyk S., Ambike S., Amireault S., Haddad J. M., Lin G., Newton D., & Richards E. A. (2022). Co-occurrences of fall-related factors in adults aged 60 to 85 years in the United States National Health and nutrition examination survey. *PLoS One*, *17*(11), 1–17. https://doi.org/10.1371/journal.pone.0277406

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026

Rondeau V., Mazroui Y., & Gonzalez J. R. (2012). frailtypack: An R package for the analysis of correlated survival data with frailty models using penalized likelihood estimation or parametrical estimation. *Journal Statistics Software*, *47*(4), 1–28. https://doi.org/10.18637/jss.v047.i04

Rondeau V., Pignon J.-P., Michiels S., & Group M.-N. C. (2015). A joint model for the dependence between clustered times to tumour progression and deaths: A meta-analysis of chemotherapy in head and neck cancer. *Statistical Methods in Medical Research*, *24*(6), 711–729. https://doi.org/10.1177/0962280211425578

Rondeau V., Schaffner E., Corbiere F., Gonzalez J. R., & Mathoulin-Pélissier S. (2013). Cure frailty models for survival data: Application to recurrences for breast cancer and to hospital readmissions for colorectal cancer. *Statistical Methods in Medical Research*, *22*(3), 243–260. https://doi.org/10.1177/0962280210395521

Sethuraman J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica*, *4*, 639–650. https://www.jstor.org/stable/24305538

Sinha D., Maiti T., Ibrahim J. G., & Ouyang B. (2008). Current methods for recurrent events data with dependent termination: A Bayesian perspective. *Journal of the American Statistical Association*, *103*(482), 866–878. https://doi.org/10.1198/016214508000000201

Tinetti M. E., Speechley M., & Ginter S. F. (1988). Risk factors for falls among elderly persons living in the community. *New England Journal of Medicine*, *319*(26), 1701–1707. https://doi.org/10.1056/NEJM198812293192604

Verma S. K., Willetts J. L., Corns H. L., Marucci-Wellman H. R., Lombardi D. A., & Courtney T. K. (2016). Falls and fall-related injuries among community-dwelling adults in the United States. *PLoS One*, *11*(3), e0150939. https://doi.org/10.1371/journal.pone.0150939

Xu C., Chinchilli V. M., & Wang M. (2018). Joint modeling of recurrent events and a terminal event adjusted for zero inflation and a matched design. *Statistics in Medicine*, *37*(18), 2771–2786. https://doi.org/10.1002/sim.v37.18

Xu Z., Sinha D., & Bradley J. R. (2021). Joint analysis of recurrence and termination: A Bayesian latent class approach. *Statistical Methods in Medical Research*, *30*(2), 508–522. https://doi.org/10.1177/0962280220962522

Yu M., Law N. J., Taylor J. M., & Sandler H. M. (2004). Joint longitudinal-survival-cure models and their application to prostate cancer. *Statistica Sinica*, *14*, 835–862. https://www.jstor.org/stable/24307418

Yu Z., & Liu L. (2011). A joint model of recurrent events and a terminal event with a nonparametric covariate function. *Statistics in Medicine*, *30*(22), 2683–2695. https://doi.org/10.1002/sim.4297

Downloaded from https://academic.oup.com/jrsssc/article/73/3/598/7596387 by guest on 29 January 2026