



Universiteit
Leiden

The Netherlands

Incorporating delayed entry into the joint frailty model for recurrent events and a terminal event

Bohnstedt, M.; Gampe, J.; Caljouw, M.A.A.; Putter, H.

Citation

Bohnstedt, M., Gampe, J., Caljouw, M. A. A., & Putter, H. (2023). Incorporating delayed entry into the joint frailty model for recurrent events and a terminal event. *Lifetime Data Analysis*, 29, 585-607. doi:10.1007/s10985-022-09587-z

Version: Publisher's Version

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/3720586>

Note: To cite this publication please use the final published version (if applicable).



Incorporating delayed entry into the joint frailty model for recurrent events and a terminal event

Marie Böhnstedt^{1,2} · Jutta Gampe¹ · Monique A. A. Caljouw³ · Hein Putter²

Received: 30 August 2021 / Accepted: 12 December 2022 / Published online: 18 January 2023

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

Abstract

In studies of recurrent events, joint modeling approaches are often needed to allow for potential dependent censoring by a terminal event such as death. Joint frailty models for recurrent events and death with an additional dependence parameter have been studied for cases in which individuals are observed from the start of the event processes. However, samples are often selected at a later time, which results in delayed entry so that only individuals who have not yet experienced the terminal event will be included. In joint frailty models such left truncation has effects on the frailty distribution that need to be accounted for in both the recurrence process and the terminal event process, if the two are associated. We demonstrate, in a comprehensive simulation study, the effects that not adjusting for late entry can have and derive the correctly adjusted marginal likelihood, which can be expressed as a ratio of two integrals over the frailty distribution. We extend the estimation method of Liu and Huang (Stat Med 27:2665–2683, 2008. <https://doi.org/10.1002/sim.3077>) to include potential left truncation. Numerical integration is performed by Gaussian quadrature, the baseline intensities are specified as piecewise constant functions, potential covariates are assumed to have multiplicative effects on the intensities. We apply the method to estimate age-specific intensities of recurrent urinary tract infections and mortality in an older population.

Keywords Recurrent events · Joint modeling · Frailty · Left truncation · Gaussian quadrature

✉ Hein Putter
h.putter@lumc.nl

¹ Max Planck Institute for Demographic Research, Rostock, Germany

² Department of Biomedical Data Sciences, Leiden University Medical Center, Leiden, The Netherlands

³ Department of Public Health and Primary Care, Leiden University Medical Center, Leiden, The Netherlands

1 Introduction

Repeated occurrences of the same type of event, such as incidents of myocardial infarction, recurrent infections, fractures, or tumor relapses, arise in various applications. If individuals are additionally at risk of experiencing a terminal event such as death, which will stop the recurrent event process, then dependent censoring of the recurrence process is often induced. Therefore, approaches for jointly modeling the two processes—recurrent events and the terminal event—have been developed. Moreover, the association between the processes can be of direct interest, for instance, whether individuals who experience more recurrences also have a higher risk of experiencing the terminal event, and joint models can provide insights into the direction and strength of the association between the event processes.

When modeling event processes the time scale t , along which the recurrence and the terminal event process evolve, has to be chosen and the choice depends on the specific application. In many clinical studies time since randomization is used as the time scale, whereas in epidemiological context and in demographic studies age often is a more relevant time scale.

The choice of the time scale determines the time origin t_0 , and if individuals come under observation only at a time $t > t_0$ then this delayed entry has to be accounted for. In this case data are left-truncated since study entry is only possible if the terminal event has not been experienced yet. This leads to a length-biased sample (for the time to terminal event) and accounting for left truncation is standard in survival analysis (Klein and Moeschberger 2003).

If the recurrence process and the terminal event process are associated then left truncation has a more wide-ranging effect. Not only is the sample consisting of individuals who tend to have a lower risk of experiencing the terminal event but also—if there is a positive association between the two processes—to have a lower risk of experiencing recurrent events. The data hence are not a random sample of the underlying population regarding both processes studied. It is therefore crucial to correctly adjust for the truncation to obtain valid inferences.

In this paper we study the effect of left truncation in the joint frailty model for recurrent events and death introduced by Liu et al. (2004). Delayed entry into the sample implies that the frailty distribution for such individuals differs from the distribution at time t_0 , due to selective survival, and this change in frailty distribution has to be incorporated. We demonstrate the effect that neglecting this adjustment has on the estimation results in this setting, by presenting an extensive simulation study. And we show how the likelihood can be adapted to correctly incorporate the information that left-truncated observations provide in this joint model for the two processes. For practical estimation we extend the method of Liu and Huang (2008) to the current setting. We illustrate the method by analyzing data from a study of recurrent urinary tract infections (UTIs) in older residents of long-term care facilities (Caljouw et al. 2014).

The joint frailty model of Liu et al. (2004) allows to examine the shape of the intensity functions of recurrence and death, as well as the potential dependence between the two processes. The dependence is introduced by a shared individual random effect entering both the intensity function of recurrence and the hazard of death. An additional

parameter determines whether the processes are positively or negatively associated, and how strong this association is. This model has been applied repeatedly in medical studies (e.g., to study recurrent cancer events as in Rondeau et al., 2007, or recurrent heart failure hospitalizations as in Rogers et al., 2016) and it has been extended in several directions (e.g., to the setting of nested case-control studies in Jazić et al., 2019). In other models, the frailty affects the intensities of the recurrent events and death in the same way (Huang and Wang 2004), or the dependence between the processes is left completely unspecified (Cook and Lawless 1997; Ghosh and Lin 2003).

Left truncation has received varying levels of attention in different frailty models and in joint models. Several studies have discussed handling left truncation in shared frailty models for clustered survival data (e.g., Jensen et al. 2004; van den Berg and Drepper 2016; Rondeau et al. 2017). In a recurrent event setting, Balan et al. (2016) considered event dependent selection; i.e., individuals were included in the study only if they had experienced at least one recurrent event in a given time period. Recurrent event studies with selection dependent on survival were briefly discussed in Cook and Lawless (2007), but not specifically in the context of frailty models.

The option of delayed entry was mentioned in Rondeau et al. (2007) but no specific account was provided. At the time of writing, the R package `frailtypack`, which can be used for fitting a variety of frailty models, does not provide functionality for estimating the joint frailty model from left-truncated data (according to the manual of version 3.5.0, date 2021-12-20, Rondeau et al., 2021).

Extensions to left-truncated data have been considered in some different types of joint models. Emura et al. (2017) introduced a joint frailty-copula model for two event times that can be adapted to accommodate left truncation and recurrent event data. Outside of the class of shared frailty models, Cai et al. (2017) proposed a model for longitudinal measurements, recurrent events, and a terminal event with inferences based on estimation equations that can be generalized to allow for left-truncated data. Liu et al. (2012) presented an estimating equation procedure for a partial marginal model of recurrent events in the presence of a terminal event with left truncation. For joint models for longitudinal data and death, incorporation of left-truncated data was considered for different specifications of models with shared random effects, for example, van den Hout and Muniz-Terrera (2016, discrete longitudinal response following binomial or beta-binomial distribution), Crowther et al. (2016, linear mixed effects model for continuous longitudinal measurements), or Piulachs et al. (2021, zero-inflated negative binomial longitudinal model).

The rest of the paper is structured as follows. Section 2 first presents the joint frailty model and the corresponding likelihood in the setting without truncation, and then introduces the adjustments required for left-truncated data. We also discuss the method of estimation. The performance of the method is assessed via simulation studies in Sect. 3. We also study the effect of ignoring left truncation in various settings. Finally, we illustrate the approach using data on recurrent UTIs in elderly patients in Sect. 4, and conclude with a discussion in Sect. 5.

2 Joint frailty model and left truncation

The joint frailty model for recurrent events and a terminal event has been applied most frequently in situations in which the time of the terminal event is subject to independent right-censoring only. In the following, we will first present the model and the corresponding likelihood for such right-censored data. Then, we will lay out how certain assumptions and, in particular, the likelihood function are adjusted to the case of left-truncated data. Throughout, we will often refer to the terminal event as death for the sake of simplicity.

2.1 Joint frailty model

We consider independent individuals i , $i = 1, \dots, m$, who can experience recurrent events between time $t_0 = 0$ and the time D_i of the terminal event. Let C_i denote a censoring time, which is assumed to be independent of the recurrence and terminal event processes. An individual can then be observed only up to his/her follow-up time $X_i = \min(C_i, D_i)$, and $\delta_i = \mathbb{1}\{D_i \leq C_i\}$ will indicate whether the terminal event occurred before censoring, with the indicator function $\mathbb{1}\{\cdot\}$. The at-risk indicator at time $t \geq 0$ is given by $Y_i(t) = \mathbb{1}\{t \leq X_i\}$, if individuals enter the study at $t_0 = 0$.

The number of recurrent events experienced by individual i up to time t is recorded by the actual recurrent event process $N_i^{R*}(t)$. Similarly, we define the actual counting process of the terminal event as $N_i^{D*}(t) = \mathbb{1}\{D_i \leq t\}$. However, due to right-censoring, we can only observe the processes $N_i^D(t) = \mathbb{1}\{X_i \leq t, \delta_i = 1\}$ and $N_i^R(t) = \int_0^t Y_i(s) dN_i^{R*}(s) = N_i^{R*}(\min(t, X_i))$. Here, the increment of the recurrence process $dN_i^{R*}(t) = N_i^{R*}((t+dt)^-) - N_i^{R*}(t^-)$ equals the number of events in the small interval $[t, t+dt)$, with t^- as the left-hand limit.

The observed data of individual i up to time t are given by $\mathcal{O}_i(t) = \{Y_i(s), N_i^R(s), N_i^D(s), 0 \leq s \leq t; \mathbf{z}_i\}$, including the observed time-fixed covariate vector \mathbf{z}_i . Individual risks will depend on the covariates as well as on the unobservable frailty value u_i , where the frailties u_i are independent realizations of a positive random variable U .

In the joint frailty model introduced by Liu et al. (2004), the observed recurrence process is assumed to have, conditional on u_i , the intensity $Y_i(t)\lambda_i(t|u_i)$ with

$$\begin{aligned} P(dN_i^R(t) = 1 \mid \mathcal{F}_t^-, D_i \geq t) &= Y_i(t)\lambda_i(t|u_i)dt \\ \lambda_i(t|u_i)dt &= d\Lambda_i(t|u_i) = P(dN_i^{R*}(t) = 1 \mid \mathbf{z}_i, u_i, D_i \geq t). \end{aligned} \quad (1)$$

Here, $\mathcal{F}_t = \sigma\{\mathcal{O}_i(s), 0 \leq s \leq t, u_i; i = 1, \dots, m\}$ denotes the σ -algebra generated by the frailty and the observed data. The terminal event process is characterized by the intensity $Y_i(t)h_i(t|u_i)$ with

$$\begin{aligned} P(dN_i^D(t) = 1 \mid \mathcal{F}_t^-) &= Y_i(t)h_i(t|u_i)dt \\ h_i(t|u_i)dt &= dH_i(t|u_i) = P(dN_i^{D*}(t) = 1 \mid \mathbf{z}_i, u_i, D_i \geq t). \end{aligned} \quad (2)$$

The first lines in (1) and (2) follow from the assumption that the censoring mechanism is conditionally independent of the two event processes given the process history. Additionally, censoring is assumed to be non-informative.

Following Liu et al. (2004), we specify the intensities as

$$\begin{aligned}\lambda_i(t|u_i) &= u_i e^{\beta' z_i} \lambda_0(t), \\ h_i(t|u_i) &= u_i^\gamma e^{\alpha' z_i} h_0(t),\end{aligned}\tag{3}$$

with baseline intensities $\lambda_0(t)$ and $h_0(t)$. The covariates z_i affect the intensities of the recurrent event process and death through a multiplicative model with effects β and α , respectively. The inclusion of the frailty u in the recurrence intensity accommodates heterogeneity across individuals and dependence between the recurrences within one individual. The association between the recurrent events and death results from the fact that the shared frailty u also enters the hazard of death. Due to the additional parameter γ , the model can capture associations of variable magnitudes and in different directions. For positive $\gamma > 0$, individuals with a higher intensity of recurrent events will also be subject to a higher hazard of death. For $\gamma < 0$, a higher intensity of recurrent events implies a lower hazard of death. If $\gamma = 0$, the intensities in (3) do not share any parameters, and the censoring of the recurrence process by death is non-informative.

The frailty U is a positive random variable. Among the distributions with positive support the gamma and the log-normal distribution received most attention as frailty distributions. The inverse-Gaussian and positive stable distribution are less frequently used alternatives. For a comprehensive discussion of potential frailty distributions see Wienke (2011) and Duchateau and Janssen (2008). The use of the log-normal distribution is most prominent in models with multivariate frailties, for example, in twin or family studies where different levels of relatedness should be represented. Undoubtedly, the most common choice for the distribution of the frailty U is a gamma distribution. Besides mathematical convenience, Abbring and van den Berg (2007) showed that in a large class of proportional hazards models the frailty distribution converges over time to a gamma, thus lending strong support to this choice. In the practical applications we will also choose a gamma distribution for the frailty U , but our derivations in this section do not depend on a specific distribution.

Hence we more generally consider that the frailties u_i follow a distribution with a density of $g_\theta(u)$ and a corresponding distribution function $G_\theta(u)$ with parameter θ . This assumption refers to the initial distribution of frailties in the population at time $t_0 = 0$. However, if $\gamma \neq 0$, the distribution of frailties in the population will change over time due to selection effects, which will cause the population at time t to be composed of survivors with lower mortality risks.

We now formulate the likelihood of the joint frailty model (3) when individuals are observed from time $t_0 = 0$. Let t_{ij} , $j = 1, \dots, J_i$, be the observed recurrence times of individual i . Based on the arguments stated in Liu et al. (2004), the conditional likelihood contribution of individual i given his/her frailty value u_i can be written as

$$\begin{aligned}
L_i^{(c)}(u_i) &= \left[\prod_{j=1}^{J_i} \lambda_i(t_{ij}|u_i) \right] \exp \left\{ - \int_0^\infty Y_i(s) \lambda_i(s|u_i) ds \right\} \\
&\quad \cdot h_i(x_i|u_i)^{\delta_i} \exp \left\{ - \int_0^\infty Y_i(s) h_i(s|u_i) ds \right\} \\
&= \left[\prod_{j=1}^{J_i} u_i e^{\beta' z_i} \lambda_0(t_{ij}) \right] \exp \left\{ - \int_0^{x_i} u_i e^{\beta' z_i} \lambda_0(s) ds \right\} \\
&\quad \cdot \left[u_i^\gamma e^{\alpha' z_i} h_0(x_i) \right]^{\delta_i} \exp \left\{ - \int_0^{x_i} u_i^\gamma e^{\alpha' z_i} h_0(s) ds \right\}.
\end{aligned} \tag{4}$$

The marginal likelihood L_i of the observed data of individual i is obtained by integrating the above expression over the frailty distribution,

$$L_i = \int_0^\infty L_i^{(c)}(u) dG_\theta(u) = \int_0^\infty L_i^{(c)}(u) g_\theta(u) du. \tag{5}$$

2.2 Adjusting for left truncation

We will now extend the above framework to allow for left truncation, that is, individuals entering the study at times that may be later than $t_0 = 0$. Before deriving the likelihood for the left-truncated data, we introduce some additional notation and assumptions.

A sample of m_V independent individuals i , $i = 1, \dots, m_V$, is left-truncated if the individuals i enter the study only at times $V_i \geq t_0$, with strict inequality for at least some individuals. Then, the observation of individual i is conditional on his/her survival up to the entry time, $D_i > V_i$, and events can only be observed in the interval $[V_i, X_i]$. Hence, the at-risk indicator $Y_i(t)$ of Sect. 2.1 is replaced by ${}_V Y_i(t) = \mathbb{1}\{V_i \leq t \leq X_i\}$.

As a consequence, the observed recurrent event process ${}_V N_i^R(t) = \int_0^t {}_V Y_i(s) dN_i^{R*}(s) = [N_i^{R*}(\min(t, X_i)) - N_i^{R*}(V_i)] \mathbb{1}\{t > V_i\}$ in this setting records only the recurrences after study entry at V_i . Analogously, the left-truncated counting process of the terminal event is given by ${}_V N_i^D(t) = \mathbb{1}\{V_i \leq X_i \leq t, \delta_i = 1\}$. The observed data for individual i are then ${}_V \mathcal{O}_i(t) = \{{}_V Y_i(s), {}_V N_i^R(s), {}_V N_i^D(s), V_i \leq s \leq t; z_i; V_i\}$, and the σ -algebra is modified as ${}_V \mathcal{F}_t = \sigma\{{}_V \mathcal{O}_i(s), 0 \leq s \leq t, u_i; i = 1, \dots, m_V\}$.

In addition to the assumption of conditionally independent censoring already made in Sect. 2.1, we further assume that the truncation times V_i are conditionally independent of the recurrence and terminal event processes given the process history. Hence, the intensity of the observed recurrence process is given by ${}_V Y_i(t) \lambda_i(t|u_i)$ and (1) is adapted as

$$P(d {}_V N_i^R(t) = 1 \mid {}_V \mathcal{F}_{t-}, D_i \geq t) = {}_V Y_i(t) \lambda_i(t|u_i). \tag{1'}$$

The intensity of the observed terminal event process is, correspondingly, ${}_V Y_i(t) h_i(t|u_i)$, such that (2) is modified as

$$P(d {}_V N_i^D(t) = 1 \mid {}_V \mathcal{F}_{t-}) = {}_V Y_i(t) h_i(t|u_i). \tag{2'}$$

Based on this, we can develop the likelihood of the joint frailty model (3) for left-truncated data. The conditional likelihood contribution of individual i given u_i is constructed in analogy to (4), with $Y_i(s)$ replaced by $vY_i(s)$, and noting that only the recurrence times $t_{ij} \geq v_i$ are observed; that is,

$$\begin{aligned} vL_i^{(c)}(u_i) &= \left[\prod_{t_{ij} \geq v_i} \lambda_i(t_{ij}|u_i) \right] \exp \left\{ - \int_0^\infty vY_i(s) \lambda_i(s|u_i) ds \right\} \\ &\quad \cdot h_i(x_i|u_i)^{\delta_i} \exp \left\{ - \int_0^\infty vY_i(s) h_i(s|u_i) ds \right\} \\ &= \left[\prod_{t_{ij} \geq v_i} u_i e^{\beta' z_i} \lambda_0(t_{ij}) \right] \exp \left\{ - \int_{v_i}^{x_i} u_i e^{\beta' z_i} \lambda_0(s) ds \right\} \\ &\quad \cdot \left[u_i^\gamma e^{\alpha' z_i} h_0(x_i) \right]^{\delta_i} \exp \left\{ - \int_{v_i}^{x_i} u_i^\gamma e^{\alpha' z_i} h_0(s) ds \right\}. \end{aligned} \quad (6)$$

In contrast to (4), the likelihood contribution (6) gives the probability of the observed data conditional not only on the frailty value u_i , but also on $D_i > V_i$, as the data are now observed conditionally on the individual's survival to study entry.

The marginal likelihood contribution is again obtained by integrating out the frailty. However, as the frailty distribution in the sample of survivors differs from the frailty distribution at time t_0 , we need to integrate over the conditional frailty distribution given survival to the time of entry into the study. This point has previously been discussed in the context of clustered survival data, among others, by van den Berg and Drepper (2016) and Eriksson et al. (2015) and for general state duration models by Lawless and Fong (1999). More formally, the marginal likelihood contribution of individual i is thus

$$\begin{aligned} vL_i &= \int_0^\infty vL_i^{(c)}(u) dG_\theta(u | D_i > v_i, V_i = v_i, z_i) \\ &= \int_0^\infty vL_i^{(c)}(u) dG_\theta(u | D_i > v_i, z_i), \end{aligned} \quad (7)$$

under the assumption that the truncation time V_i is independent of U_i .

In particular, if $\gamma > 0$ such that the recurrence process and the mortality process are positively associated, individuals who survived up to time v will tend to have lower frailty values than individuals who died before time v , for given z_i . Hence, the frailty distribution among survivors beyond time v , $G_\theta(u | D > v)$, will tend to have more probability mass at lower values u than the frailty distribution $G_\theta(u)$ in the underlying population at time t_0 . Consequently, neglecting the effect of the survivor selection on the frailty distribution in the sample and constructing a marginal likelihood as

$$vL_i^{\text{naive}} = \int_0^\infty vL_i^{(c)}(u) dG_\theta(u) = \int_0^\infty vL_i^{(c)}(u) g_\theta(u) du, \quad (8)$$

would lead to invalid inference if $\gamma \neq 0$. We will illustrate the resulting biases in the parameter estimates in a simulation study in Sect. 3.

For computing the correct marginal likelihood (7), we first apply Bayes' theorem to find that

$$g_\theta(u | D_i > v_i, \mathbf{z}_i) = \frac{P(D_i > v_i | u, \mathbf{z}_i) g_\theta(u)}{P(D_i > v_i | \mathbf{z}_i)} = \frac{\exp\left\{-\int_0^{v_i} h_i(s|u) ds\right\} g_\theta(u)}{\int_0^\infty P(D_i > v_i | u) g_\theta(u) du}, \quad (9)$$

where we suppress the dependence on the covariates \mathbf{z}_i in the last expression for notational convenience. Combining equations (6), (7), and (9), we can express the marginal likelihood contribution ${}_{\nu}L_i$ of individual i with left-truncated data as

$$\frac{\int_0^\infty \left[\prod_{t_{ij} \geq v_i} \lambda_i(t_{ij}|u) \right] \exp\left\{-\int_{v_i}^{x_i} \lambda_i(s|u) ds\right\} h_i(x_i|u)^{\delta_i} \exp\left\{-\int_0^{x_i} h_i(s|u) ds\right\} g_\theta(u) du}{\int_0^\infty P(D_i > v_i | u) g_\theta(u) du}. \quad (10)$$

Interestingly, the formula for ${}_{\nu}L_i$ in (10) could have been equivalently derived as the marginal (with respect to the frailty) probability of the recurrence and follow-up data on individual i , conditional on individual i being included in the study, $D_i > v_i$. To see this, let us denote by E_i the event that individual i has follow-up time x_i with indicator δ_i and the observed recurrence times t_{ij} over $[v_i, x_i]$, and consider

$$P(E_i | D_i > v_i) = \frac{P(E_i \cap \{D_i > v_i\})}{P(D_i > v_i)} = \frac{P(E_i)}{P(D_i > v_i)} = \frac{\int_0^\infty P(E_i | u) g_\theta(u) du}{\int_0^\infty P(D_i > v_i | u) g_\theta(u) du}.$$

As the integrals over the frailty distribution in (10) will not, in general, have closed form expressions, we will use numerical integration in the following.

2.3 Estimation of the joint frailty model

Liu and Huang (2008) proposed using Gaussian quadrature to approximate the marginal likelihood of frailty proportional hazards models, including the joint frailty model (3). In combination with a suitable specification of the baseline intensities, this approach allows for the direct maximization of the approximated likelihood which results in shorter computation times as compared to approaches like the EM algorithm. We also use Gaussian quadrature to estimate the parameters in the joint frailty model for left-truncated data. The details are given in Section S.1 of the supplementary material.

The evaluation of the approximate marginal likelihood depends on the specific form of the baseline intensities $\lambda_0(t)$ and $h_0(t)$. We adopt a piecewise constant model for these functions,

$$\lambda_0(t) = \sum_{k=1}^{K^R} \lambda_{0k} \mathbb{1}\{t \in I_k^R\} \quad \text{and} \quad h_0(t) = \sum_{k=1}^{K^D} h_{0k} \mathbb{1}\{t \in I_k^D\}.$$

with intervals $I_k^R = (t_{k-1}^R, t_k^R]$, $k = 1, \dots, K^R$, and $I_k^D = (t_{k-1}^D, t_k^D]$, $k = 1, \dots, K^D$. Specifications with a moderate number of up to 10 intervals and the cut-points t_k , $k \geq 1$, which have been determined based on the quantiles of the observed event times, are generally expected to lead to good results in practice (see Cook and Lawless 2007; Liu and Huang 2008). In the setting with left truncation, appropriate choices have to be made for the starting points of the first intervals, t_0^R and t_0^D . Depending on the study design, they might be set equal to the lowest study entry time, $\min_i v_i$, or a lower value $t^* \geq t_0$.

The direct maximization of the marginal likelihood would also be possible if the baseline intensities were assumed to follow a simple parametric model, such as the Weibull model. Nonetheless, we recommend the use of the more flexible piecewise constant intensity models, unless prior knowledge allows for an informed choice of a specific parametric model.

Finally, the parameter estimates in the joint frailty model with left-truncated data are obtained by maximizing the approximate marginal log-likelihood. Given the piecewise constant specification of the baseline intensities, this corresponds to standard parametric maximum likelihood estimation with the usual asymptotic properties of the estimators. The calculation of the standard errors is based on the inverse of the negative Hessian matrix of the approximate marginal log-likelihood. We give additional details on the implementation in Section S.2 of the supplementary material.

3 Simulation study

To evaluate the performance of the proposed method for estimating the parameters of the joint frailty model in case of left-truncated data, we conducted a simulation study. We will also demonstrate which biases can arise if the likelihood is not correctly adjusted to the survivor selection, in particular, to the selection effects on the frailty distribution.

Estimator performance will depend on various aspects of the observation scheme. One aspect is the distribution of the study entry times V_i , in which both the range of the distribution and its shape matter. Furthermore, the censoring mechanism—that is, the length of the individual follow-up periods and the number of additional drop-outs—will influence the performance of the method. To study these issues, we will first present a base scenario, and will then assess how different observational settings affect the results.

3.1 Settings

In the base scenario, we generated data from a joint frailty model (3). The time scale t is the age of the individual. The hazard of death and the intensity of recurrent events are each affected by a single binary covariate, which is drawn from a Bernoulli distribution with parameter 0.5. The regression coefficients are $\alpha = 0.5$ (death) and $\beta = 0.5$ (recurrence), respectively. The frailty values are realizations of a gamma distribution with a mean of one and a variance of $\theta = 0.5$. The values of the dependence param-

ter γ were chosen to cover a positive ($\gamma = 0.5$) and a negative ($\gamma = -0.5$) association between the recurrence process and death, as well as the case in which the recurrence intensity does not affect the mortality risk ($\gamma = 0$).

The baseline functions $h_0(t)$ and $\lambda_0(t)$ were designed to mimic a study in an older population among whom the risks of death as well as of experiencing the recurrent event increase exponentially with age. Hence, we chose for both baseline functions the Gompertz-Makeham form, $ae^{bt} + c$, where $t = 0$ corresponds to age 75. By setting $a = 0.984$, $b = 0.045$, and $c = 0$ for the recurrence process ($\lambda_0(t)$), and $a = 0.108$, $b = 0.07$, and $c = 0.12$ for the survival process ($h_0(t)$), the baseline functions were comparable to the estimated intensities for the high risk group in the data example in Sect. 4.

To arrive at the left-truncated samples, the following steps were combined. For each individual, a survival time D (i.e., age > 75) and an entry time into study V were simulated. Only those individuals who survived beyond his/her entry time—that is, for whom $D > V$ —were included in the final sample (i.e., were 'observed'). Therefore, the distribution of entry times V that are observed in the final sample depends on both the mortality model in (3) and the initial distribution of the truncation times before selection.

In the base scenario, our aim was to have entry times in the final sample that were distributed across the total age range—here, ages 75 to 95—with higher numbers of study entries at the younger ages than at the older ages. This scheme will be referred to as truncation pattern A in the following.

To obtain a final observed sample with such characteristics, the entry times V were drawn from a truncated normal distribution defined on the age range 75 to 95. More specifically, the truncated normal distribution was specified to have a mode equal to the maximum age of 95 with parameter values chosen so that the distribution of the observed study entry times in the truncated sample had the desired shape (the left panel of Fig. 1 illustrates this procedure). The initial number of generated survival times was chosen such that the final truncated samples had an average size of about $m_V = 500$ individuals.

An independent censoring mechanism was imposed in the following way. For most individuals, the censoring times were the end of a planned individual follow-up period of $t_C = 4$ years. However, some of the follow-up times were longer than four years, and some premature random drop-outs occurred. Again, this was done in response to the situation that we observed in the data application of Sect. 4. Accordingly, we generated random durations from a mixture distribution with an 85% point mass at t_C , a 10% uniform distribution on $[0, t_C]$, and a 5% uniform distribution on $[t_C, t_C + 0.5]$, with the latter two covering the drop-outs before t_C and the longer follow-up periods, respectively. These random durations were added to the individual V_i , and the individual censoring time C_i was the minimum of this sum and age 95.

The right panel of Fig. 1 illustrates how the mechanisms of truncation and censoring jointly determine the number of individuals at risk at any time t across the age range $[75, 95]$. Truncation pattern A causes the number of individuals at risk to increase steeply at the early ages, and then to decrease only gradually across the age range. However, due to the relatively short individual follow-up times, the number of

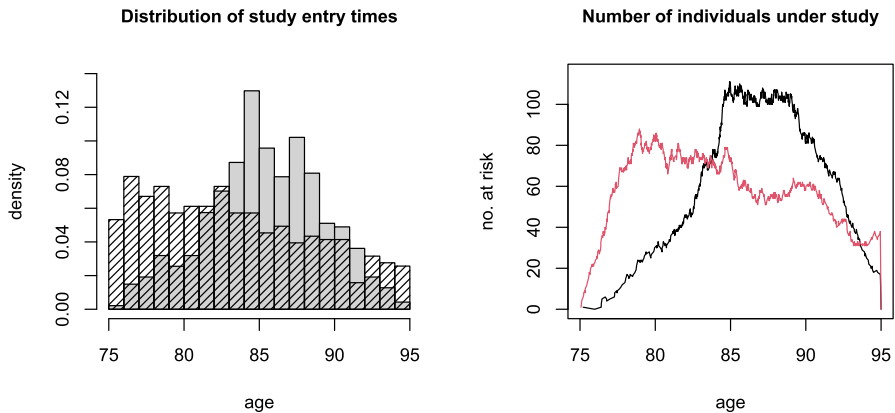


Fig. 1 Distribution of the ages at study entry (left) and the number of individuals at risk across the age range [75, 95] (right) for one simulated sample from the base scenario with truncation pattern A (shaded bars, red line) or truncation pattern B (gray bars, black line), both with planned individual follow-ups of four years (Color figure online)

individuals at risk across ages is considerably smaller than the total sample size of about 500.

Under this observation scheme and with the Gompertz-Makeham baseline intensities, the average number of observed recurrent events per individual ranges from 1.62 to 2.40 (with a mean of 2.04) across 200 final truncated samples generated from the base scenario with a positive association. If there is no association, the range is [2.88, 3.93] (with a mean of 3.33). In the base scenario with a negative association, the average number of recurrences per individual is comparatively high, ranging from 5.33 to 6.93 (with a mean of 6.09). This is due to the fact that the sample of survivors tends to have lower death risks, which in this scenario with $\gamma < 0$ are associated with relatively higher recurrence intensities.

In the setting with a positive association between the recurrence and mortality process ($\gamma = 0.5$), additional simulation scenarios were set up by varying the censoring and truncation patterns.

First, we considered the effect of changing the planned individual follow-up times to $t_C = 1$ year or $t_C = 8$ years, respectively. Longer individual follow-up times increase the number of individuals who are under observation at a certain time t , and are therefore expected to improve the estimator performance.

Second, we explored a scenario with a more unimodal distribution of the study entry times in which relatively few individuals entered the study at the youngest and the oldest ages (see Fig. 1). This is truncation pattern B. To obtain a final sample with these characteristics, we simulated the initial truncation times again from a truncated normal distribution on the age range 75 to 95. However, in this scenario, the distribution had a mode equal to 90, that is, within the above age range.

Finally, we examined a setting with a wider age range of [64, 105]. If $t = 0$ was now expected to correspond to age 64, but the Gompertz-Makeham baseline intensities were expected to agree with the intensities of the base scenario over [75, 95], the parameters needed to be adapted. This was achieved by maintaining the values of b and c , but

setting $a = 0.6$ or $a = 0.05$ for the recurrence and death processes, respectively. The initial distributions of the study entry times were adapted to produce truncation patterns A or B on the wider age range [64, 105]. In all of the additional scenarios, the truncated samples again had a target size of $m_V = 500$ individuals. The parameter values for the distributions of the study entry times and the initial sample sizes for the different scenarios can be found in Section S.3 of the supplementary material.

3.2 Estimation and results

The estimation of the joint frailty model was carried out under the assumptions of gamma-distributed frailties with a mean of one and piecewise constant models for the two baseline intensities $\lambda_0(t)$ and $h_0(t)$. For both intensities, 10 intervals were used that were denoted by I_k^R (recurrence process) and I_k^D (mortality), $k = 1, \dots, 10$. The intervals were determined by the deciles of the observed recurrence and survival times, respectively. We set $t_0^R = t_0^D = 75$ (or $t_0^R = t_0^D = 64$) equal to the starting point of the respective age range and $t_{10}^R = t_{10}^D$ equal to the maximum follow-up time in the sample. The marginal likelihood was approximated using non-adaptive Gauss-Hermite quadrature with $Q = 30$ quadrature points. We ran 200 replications in each setting. All computations were performed in R (R Core Team 2020). Further details on the implementation are provided in Section S.2 of the supplementary material.

Figures 2 and 3 illustrate the results of the base scenario with different underlying associations, $\gamma \in \{-0.5, 0, 0.5\}$. The top panels of Fig. 2 show that the covariate effects α and β , the dependence parameter γ , and the frailty variance θ are estimated without significant bias. The estimated standard errors of these parameters in the bottom panels of Fig. 2 are largely in line with the empirical standard deviations of the respective parameter estimates across the replications. Nevertheless, we notice that the estimator performance varies for different true values of the dependence parameter.

This pattern can be explained to some extent by different survivor selection effects. The truncated sample consists of survivors, who tend to have lower mortality risks. For $\gamma \neq 0$, these lower death risks are associated with relatively lower or higher recurrence intensities, which are also reflected in the differences in the average numbers of recurrences per individual given in Sect. 3.1. In particular, if the recurrence process and the mortality process are positively associated, this implies that the frailty values and the recurrence intensities are lower in the sample of survivors. In the current setting, this lower frailty variance in the sample is favorable for the estimation of θ ; whereas the low recurrence intensity, which is associated with higher probabilities of having no observed recurrent event, increases the variability in the corresponding estimated covariate effect $\hat{\beta}$. The opposite effects are observed if the event processes are negatively associated. If the recurrence intensity has no effect on survival ($\gamma = 0$), the method still yields reliable results, and the parameters exclusively affecting survival are estimated with higher levels of precision.

We also calculated the coverage probabilities of 95% Wald confidence intervals for the covariate effects, the dependence parameter, and the frailty variance across the 200 replications of the base scenario for each choice of $\gamma \in \{-0.5, 0, 0.5\}$. The values, reported in Table 1, are close to the nominal level, indicating adequate coverage.

Table 1 Coverage probabilities of 95% confidence intervals under the base scenario with varying dependence parameter $\gamma \in \{-0.5, 0, 0.5\}$ based on the correct or the naive likelihood, across 200 replications each

Dependence γ	Correct likelihood (7)				Naive likelihood (8)			
	α	β	γ	θ	α	β	γ	θ
0.5	0.945	0.955	0.920	0.960	0.855	0.420	0.855	0.810
0	0.945	0.955	0.955	0.945	0.960	0.945	0.960	0.920
-0.5	0.940	0.960	0.955	0.955	0.925	0.750	0.915	0.000

The estimates of the baseline intensities, displayed for the base scenario with positive dependence $\gamma = 0.5$ in Fig. 3, also perform satisfactorily.

It is instructive to look at how the results change if the effects of survivor selection on the frailty distribution in the sample are not taken into account correctly. As Figs. 4 and 5 show, if the inference is based on the naively constructed marginal likelihood (8), biases can be seen in all parameter estimates in case the recurrence process and the mortality process are associated. Moreover, as the estimated standard errors for the covariate effects are substantially smaller than those obtained using the correct likelihood, they do not adequately reflect the uncertainty in the parameter estimates. As a result, the coverage of the corresponding 95% Wald confidence intervals is mostly much lower than the nominal level (see Table 1). The baseline intensities of recurrence and death are increasingly underestimated for advancing age in the base scenario with positive dependence ($\gamma = 0.5$), as depicted in Fig. 5. This is because in a setting with a positive association, the distribution of frailty among the survivors tends to be concentrated at lower values. Accordingly, for negative associations, the recurrence intensity will be overestimated at the older ages, while the hazard of death will again be underestimated at the older ages. Hence, failing to construct the marginal likelihood based on the correct distribution of the frailty, see (7), introduces marked biases in the estimates and the standard errors. Only if the event intensities are not associated ($\gamma = 0$) is the distribution of frailty among survivors equal to the initial frailty distribution G_θ , such that the naive marginal likelihood coincides with the correct marginal likelihood (7) and yields valid inferences.

Lastly, we want to examine the results for the additional simulation scenarios with modified censoring and truncation patterns. The figures illustrating these results can be found in Section S.3 of the supplementary material. In the scenario with a planned individual follow-up of only $t_C = 1$ year, we find increased variability in all parameter estimates (see Figures S.1 and S.2). This is expected, because with shorter individual follow-up times, fewer individuals are observed at a given age t than in the base scenario. Further extending the planned individual follow-up times of the base scenario from $t_C = 4$ to $t_C = 8$ years does not lead to considerable improvements, apart from some reduced variability in the estimates of the frailty variance and the baseline intensities.

A change in the distribution of the study entry times can markedly influence the estimation results. In the modified base scenario with truncation pattern B, the estimated covariate effects $\hat{\alpha}$ and $\hat{\beta}$ are more variable than under truncation pattern A,

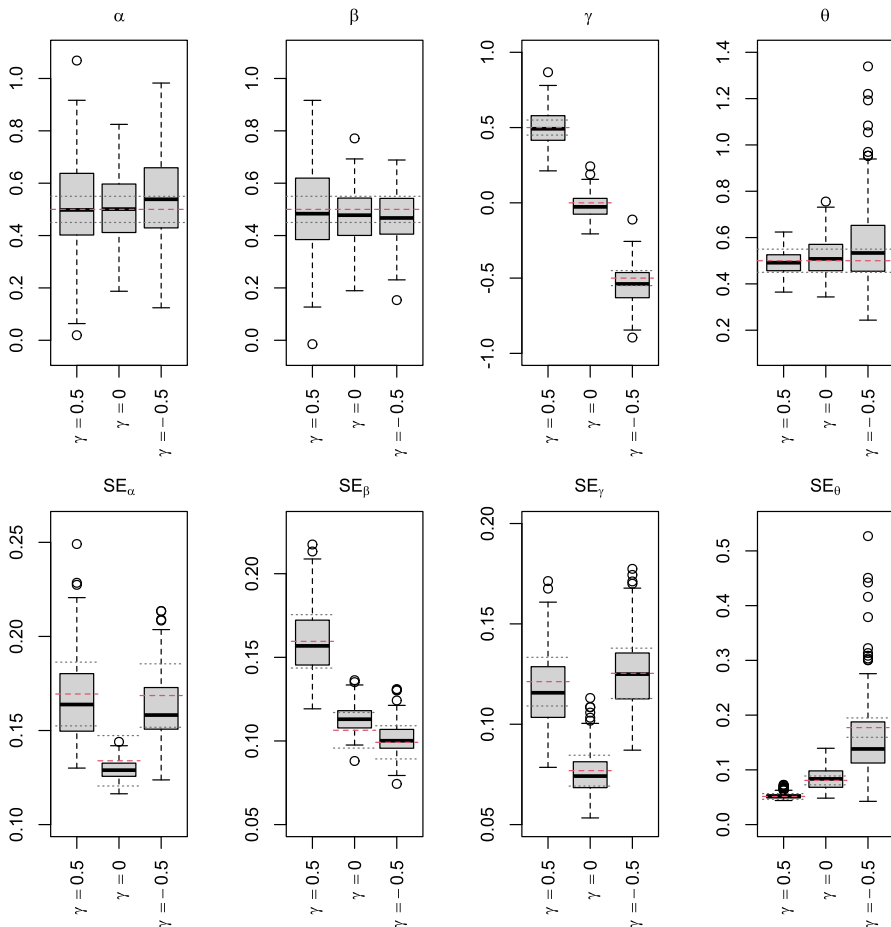


Fig. 2 Box plots of the parameter estimates (top) and estimated standard errors (bottom) in the joint frailty model for positive ($\gamma = 0.5$), no ($\gamma = 0$), or negative ($\gamma = -0.5$) dependence under the base scenario. Left to right: covariate effect on mortality (α) and on recurrences (β), dependence parameter (γ), and frailty variance (θ) based on 200 truncated samples with a target size of 500. The red dashed line marks the true parameter value (top) or empirical standard deviation (bottom); the gray dotted lines mark 10% deviations from the respective value (Color figure online)

occasionally with negative estimates (see Figure S.3). In addition, the first piece of each of the baseline intensities shows an upward bias (cf. top panels of Figure S.4) because few individuals entered the study at the younger ages. Although the intervals for the pieces of the baseline intensities were constructed to contain roughly equal numbers of observed events, the first intervals cover a relative large age range with few individuals under study at a given age t due to the delayed entry.

The last scenario combines a wider age range [64, 105] and truncation times spread across the whole age range according to pattern A or B, with individual follow-ups planned for $t_C = 4$ years. This setting is more demanding because the amount of information available at a given age t is considerably smaller than it is in the scenarios

with age range [75, 95]. Therefore, the variability in the estimates tends to increase, and the estimates of the dependence parameter and the frailty variance exhibit a small downward bias (see Figure S.3).

Overall, the simulation studies suggest that the proposed method for the estimation of the joint frailty model based on left-truncated data performs satisfactorily. The parameter estimates are largely unbiased if the study design ensures that a reasonable number of individuals are under observation across time t . Including a relatively large number of individuals early on and a preferably stable number of study entries across the remaining time range benefits the estimation. In addition, the individual follow-up times should be sufficiently long given the total time window and the sample size. As expected, the patterns of censoring and truncation that cause more information to be lost negatively affect the estimator performance.

4 Recurrent infections and mortality in an older population

We apply the proposed method to data from a study on recurrent urinary tract infections (UTIs) and mortality in an institutionalized elderly population (Caljouw et al. 2014). The original study was set up to investigate whether cranberry capsules are effective in preventing UTIs in elderly individuals (vs. placebo). Study participants were residents of long-term care facilities and were between 64 and 102 years old when they entered the study. In the present analysis the aim is to study the association between recurrent UTIs and mortality, so both processes shall be modeled jointly. We consider age as the main time scale of the event processes, as mortality and presumably also susceptibility to infections naturally depend on age. Around 34% of the participants died during the follow-up period which was set to one year. Ages at entry stretch over almost forty years, so observations are left-truncated.




Fig. 3 Estimates (gray) of the baseline intensity of recurrence (left) and of death (right) based on 200 truncated samples with a target size of 500 generated from a joint frailty model with positive dependence ($\gamma = 0.5$) under the base scenario. The red bold line gives the true baseline intensity (Color figure online)

Apart from determining whether there is an association between the rate of UTIs and mortality, we also seek to estimate the age-specific risks of UTIs and death as well as the effects of the cranberry treatment. The participants were stratified into two groups of high or low UTI risk depending on whether they had diabetes mellitus, a urinary catheter, or at least one treated UTI in the preceding year. Within these two strata, the participants were randomly assigned to the treatment or the control group. For details on drug administration and diagnosis of UTIs see Caljouw et al. (2014).

The sample consisted of 928 individuals, most of whom were women (703; 75.8%). Of these individuals, 516 were considered to be at high baseline UTI risk, while 412 were in the low risk group. The distribution of entry ages is shown in the left panel of Fig. 6. Follow-up time was about one year on average (mean: 332 days, median: 372 days). A total of 317 participants (34.2%) died during the study period. The number of observed UTIs per individual ranged from zero to 10, with 62.2% of the individuals having no UTIs, 20.8% having one UTI, and 17.0% experiencing two or more UTIs during the follow-up period.

The origin of the time scale $t_0 = 0$ corresponds to age 64, the youngest age of entry into the study. Because of the specific distribution of the ages at study entry in conjunction with the comparatively short individual follow-up times, relatively few individuals were under observation at any given age, in particular at the youngest and oldest ages, as the right panel of Fig. 6 shows.

We estimated the joint frailty model for UTIs and overall mortality separately for the groups with high and low baseline UTI risk. Two binary covariates for treatment and gender were included, and frailties were assumed to follow a gamma distribution with a mean of one (at t_0). The baseline intensity of UTI recurrence and the hazard of death were specified as piecewise constant functions with 10 intervals over the age range 64 to 103 in the high risk group and 64 to 104 in the low risk group. Separately for the two risk groups, the cut-points for the intervals were determined based on the

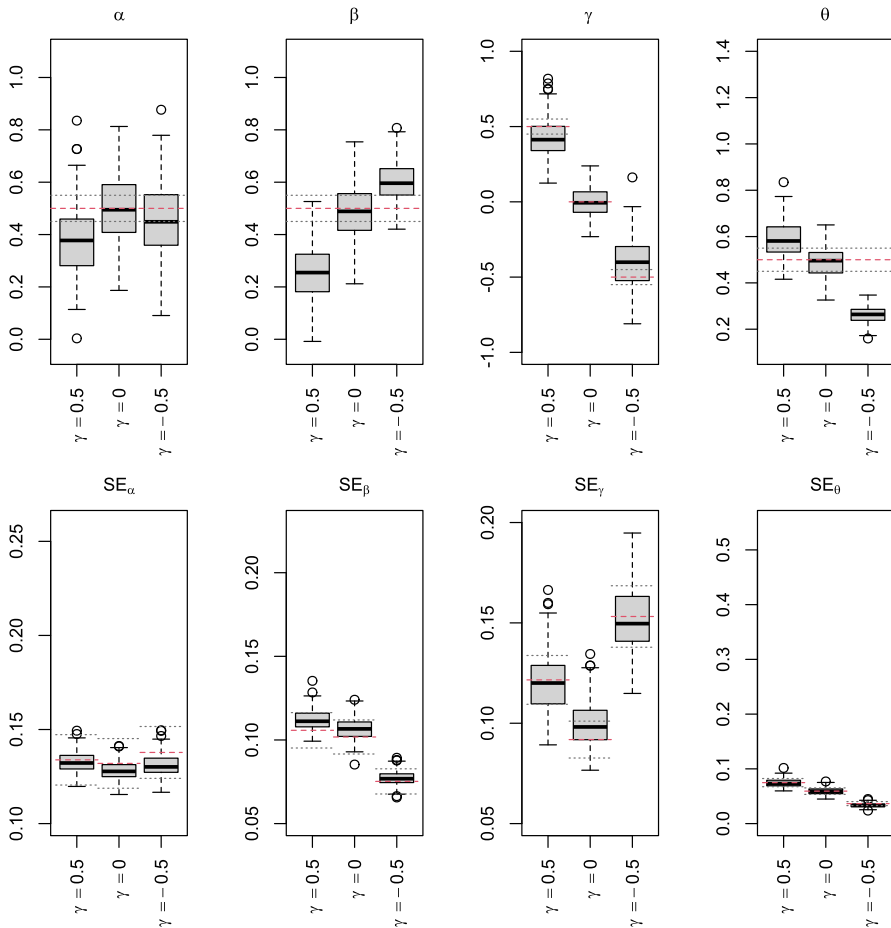


Fig. 4 Box plots of the parameter estimates (top) and estimated standard errors (bottom) based on the naive likelihood of the joint frailty model for positive ($\gamma = 0.5$), no ($\gamma = 0$), or negative ($\gamma = -0.5$) dependence under the base scenario. Left to right: covariate effect on mortality (α) and on recurrences (β), dependence parameter (γ), and frailty variance (θ) based on 200 truncated samples with a target size of 500. The red dashed line marks the true parameter value (top) or empirical standard deviation (bottom); the gray dotted lines mark 10% deviations from the respective value (Color figure online)

deciles of the observed recurrence or death times, respectively. The likelihood was approximated using non-adaptive Gaussian quadrature with 30 nodes.

The parameter estimates for both risk groups are reported in Table 2. In the group with a high baseline UTI risk, the intensities of recurrent infection varied between participants with an estimated frailty variance of $\hat{\theta} = 0.380$ (SE: 0.086). In particular, individuals with a higher intensity of recurrent infections tended to also experience higher mortality risks, as indicated by the positive estimate of the dependence parameter, $\hat{\gamma} = 0.181$ (SE: 0.084). The participants in the low risk group seemed to be more heterogeneous ($\hat{\theta} = 1.122$, SE: 0.316), but the analysis did not detect an association between the occurrence of UTIs and survival ($\hat{\gamma} = 0.058$, SE: 0.044). The results

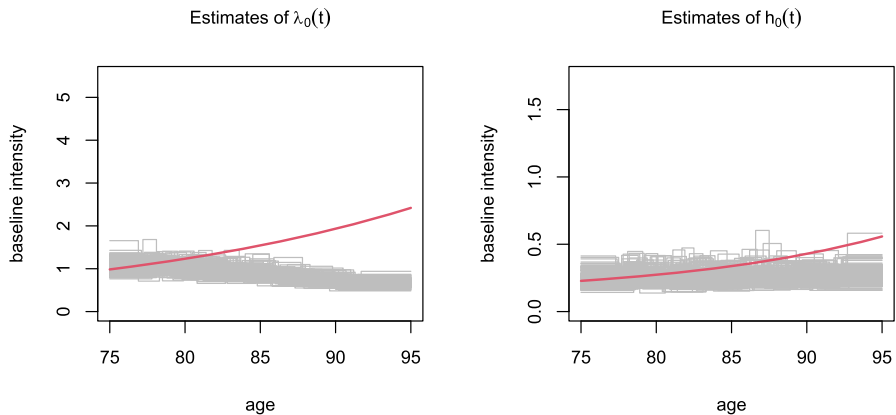


Fig. 5 Estimates (gray) of the baseline intensity of recurrence (left) and of death (right) based on the naive likelihood for 200 truncated samples with a target size of 500 generated from a joint frailty model with positive dependence ($\gamma = 0.5$) under the base scenario. The red bold line gives the true baseline intensity (Color figure online)

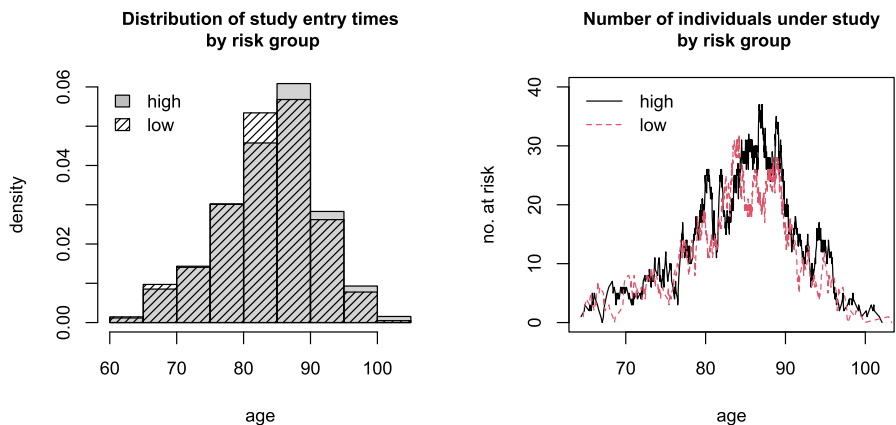


Fig. 6 Distribution of the ages at study entry (left) and the number of individuals under observation across the age range (right) in the cranberry data set, separately for the groups with high baseline UTI risk (gray bars, black solid line) and low baseline UTI risk (shaded bars, red dashed line) (Color figure online)

suggest that the cranberry capsules did not have a noticeable effect on the occurrence of UTIs irrespective of the baseline UTI risk. When we look at gender differences, we see that males and females experienced similar intensities of infection, while males had higher mortality levels than females in both groups.

The estimated baseline intensities displayed in Fig. 7 demonstrate nicely the age dependence of the recurrence intensity and the hazard of death. For the individuals with a high baseline UTI risk, both the intensity of recurrent infection and the mortality risk showed a general tendency to increase with age, although the small number of observations leads to considerable uncertainty at the highest ages. In addition, the

Table 2 Parameter estimates (with standard errors) for the joint frailty model fitted to the cranberry data set, separately by risk group

	High baseline UTI risk	Low baseline UTI risk
Recurrent UTIs		
Treatment (cranberry)	0.000 (0.161)	0.189 (0.217)
Gender (male)	0.061 (0.218)	− 0.384 (0.381)
Mortality		
Treatment (cranberry)	0.107 (0.152)	− 0.001 (0.197)
Gender (male)	0.396 (0.178)	0.787 (0.210)
Association		
Dependence γ	0.181 (0.084)	0.058 (0.044)
Frailty variance θ	0.380 (0.086)	1.122 (0.316)

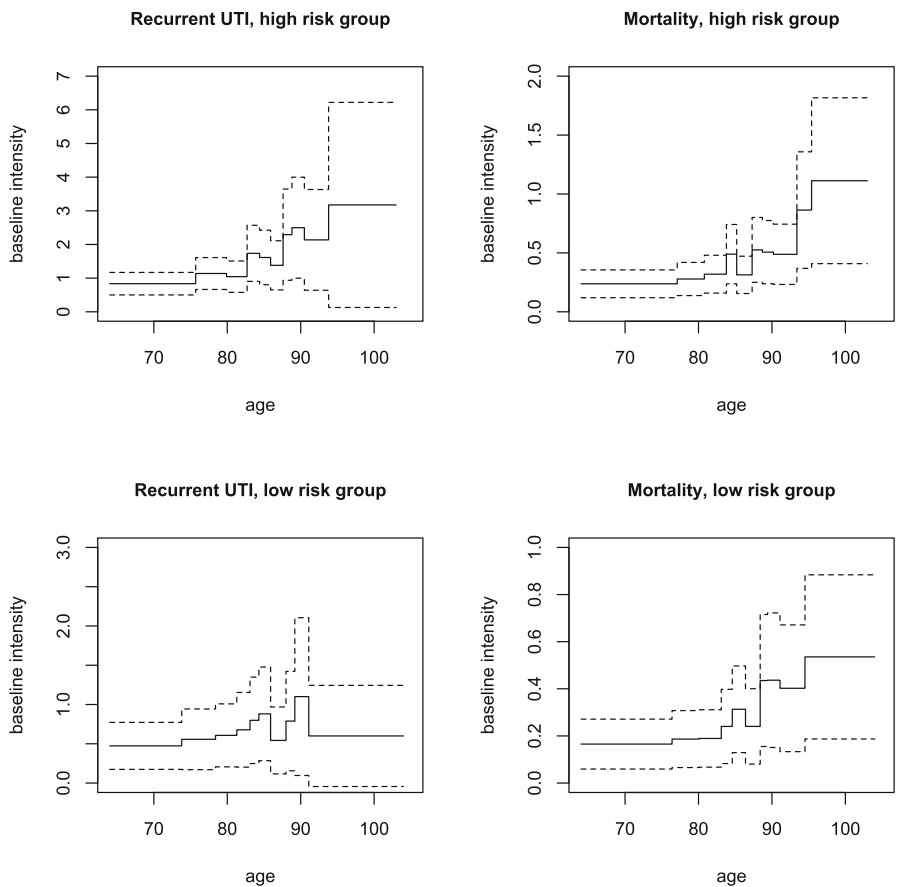


Fig. 7 Estimated baseline intensities (solid) of recurrence (left) and mortality (right) with ± 2 SE-confidence bounds (dashed) for the cranberry data, separately for the groups with high baseline UTI risk (top) and low baseline UTI risk (bottom)

individuals with a high baseline UTI risk tended to experience higher intensities of recurrent infection and death than the individuals with a low baseline UTI risk.

The original study, which used time since randomization as scale, reported a positive treatment effect of the intake of cranberry capsules only in the group with a high baseline UTI risk and only for the outcome of UTI incidence (first infection during follow-up). When all recurrent UTIs were analyzed in a gamma-frailty model, no treatment effect was detected, which is in line with the findings presented here.

5 Discussion

We have studied the extension of a method for estimating the joint frailty model for recurrent events and a terminal event to the setting with left-truncated data. The marginal likelihood of the model can be expressed as a ratio of two integrals over the frailty distribution. For computing the marginal likelihood we used Gauss-Hermite quadrature.

The simulation studies presented here have shown that the estimation procedure performs satisfactorily in general, and have demonstrated how different observation schemes affect the estimator performance. While any pattern of truncation or censoring results in incomplete information, study designs should still aim to provide enough information to meet the needs of a model as complex as the joint frailty model. Having a sufficient number of individuals under observation across most of the time range, and especially at the start of the process, seems to be crucial for the method to yield reliable results.

Allowing for left truncation in frailty models requires us to consider carefully how the frailty distribution in the sample of survivors may differ from the frailty distribution in the underlying population due to selection effects. We illustrated through simulations the biases that can arise in the parameter estimates of the joint frailty model if this difference in the frailty distributions is ignored.

Extending the framework of the joint frailty model to incorporate delayed entry allowed us to study age-specific risks of recurrent urinary tract infections and death in an older population. Similarly, the proposed approach enables researchers to use the joint frailty model in a wider variety of contexts in which subjects are included in a study only if they have not yet experienced the terminal event. Apart from clinical studies with delayed entry, these contexts may include register-based studies of event processes evolving with age as the main time scale, with individuals entering at different ages.

For a complete specification of the model and the approximate likelihood function, the frailty distribution and the baseline intensities have to be chosen. The number of quadrature points also has to be fixed.

Although the simulation study and the application covered only the common choice of a gamma distribution for the frailties, the quadrature approach can be employed with other frailty distributions that have a closed form inverse distribution function, or a log-normal distribution. The use of Gaussian quadrature, in the way it was presented here, is not the only option. However, comparisons with some readily available

‘all-purpose’ numerical integration procedure (the R-function *integrate()*) showed considerable advantages regarding computation times.

If the baseline intensities are specified by parsimonious parametric models, such as, for example, Weibull hazards, then the marginal likelihood will be a function of a moderate number of parameters that can be maximized directly. However, global parametric specifications require problem-specific knowledge as they always bear the risk of misspecification and therefore of biases. Piecewise constant specifications of the baseline intensities are a flexible alternative, still preserving the computational advantages of parametric models, and have therefore repeatedly been suggested in the literature. Naturally, sample size will guide the flexibility, that is, the number of intervals (parameters), that can be afforded, but specifications with a moderate number of up to 10 intervals generally lead to good results in practice (see Cook and Lawless 2007; Liu and Huang 2008). The positions of the cut points commonly are determined based on the quantiles of the observed event times and capture the required flexibility well. If flexible smooth estimates of the intensities are required then approaches like penalized splines are an alternative, however, automatic smoothing parameter selection in joint frailty models needs further investigation.

The number of quadrature points determines the accuracy of the integral approximations in the marginal likelihood, as well as the computation time. In line with previous recommendations for gamma frailty models (see Liu and Huang 2008), we used $Q = 30$ quadrature points, which produced good results in a reasonable period of time in our settings.

The joint frailty model generally can also incorporate time-dependent covariates (Liu et al. 2004). Yet, in the situation with delayed entry, including a time-dependent covariate in the survival model would require us to know the covariate values at times $t \geq t_0$ before study entry, which may not be available depending on the application.

Moreover, the current approach is limited to applications in which there is heterogeneity in recurrence intensities. Due to the specific dependence structure in the joint frailty model considered here, the association between the recurrence process and the terminal event process cannot be assessed if the frailty variance is close to zero.

Finally, in some applications, it might be of interest to extend the proposed method to use information on recurrences before entry into the study. These additional event times can be included in the marginal likelihood, and are expected to lead to increased precision in the estimation of the model for the recurrence process. However, when using such an approach, researchers should reflect critically on the quality of the retrospectively collected data, as recollections by study participants may be less reliable than data drawn from other sources, such as registries.

We have focused on adapting the joint frailty model as introduced by Liu et al. (2004), which specifies the complete intensities for both event processes, to left-truncated observations. An alternative approach to the analysis of recurrent events in the presence of a terminal event was presented by Kalbfleisch et al. (2013). They consider a partial marginal rate function of the recurrent event process that is conditional only on the frailty and the covariates. Inference is based on a set of estimating equations with Breslow-type estimators for the baseline rate and hazard function. Adapting this approach to take left truncation into account seems feasible, but is beyond the scope of

this paper. A thorough investigation of the method's performance would be required, in particular, to assess the effects of intervals with few individuals at risk. This is a topic for future research.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10985-022-09587-z>.

Funding The authors received no specific funding for this work.

Data availability The data set on recurrent urinary tract infections analyzed in the present manuscript cannot be made available for reasons of confidentiality.

Code availability

R-code to reproduce the results of Section 3 is available from the corresponding author upon request.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

- Abbring JH, van den Berg GJ (2007) The unobserved heterogeneity distribution in duration analysis. *Biometrika* 94(1):87–99
- Balan TA, Jonker MA, Johannesma PC, Putter H (2016) Ascertainment correction in frailty models for recurrent events data. *Stat Med* 35:4183–4201. <https://doi.org/10.1002/sim.6968>
- Cai Q, Wang M-C, Chan KCG (2017) Joint modeling of longitudinal, recurrent events and failure time data for survivor's population. *Biometrics* 73:1150–1160. <https://doi.org/10.1111/biom.12693>
- Caljouw MAA, van den Hout WB, Putter H, Achterberg WP, Cools HJM, Gussekloo J (2014) Effectiveness of cranberry capsules to prevent urinary tract infections in vulnerable older persons: a double-blind randomized placebo-controlled trial in long-term care facilities. *J Am Geriatr Soc* 62:103–110. <https://doi.org/10.1111/jgs.12593>
- Cook RJ, Lawless JF (1997) Marginal analysis of recurrent events and a terminating event. *Stat Med* 16:911–924
- Cook RJ, Lawless JF (2007) The statistical analysis of recurrent events. *Statistics for biology and health*. Springer, New York
- Crowther MJ, Andersson TM-L, Lambert PC, Abrams KR, Humphreys K (2016) Joint modelling of longitudinal and survival data: incorporating delayed entry and an assessment of model misspecification. *Stat Med* 35:1193–1209. <https://doi.org/10.1002/sim.6779>
- Duchateau L, Janssen P (2008) The frailty model. Springer, Berlin
- Emura T, Nakatochi M, Murotani K, Rondeau V (2017) A joint frailty-copula model between tumour progression and death for meta-analysis. *Stat Methods Med Res* 26:2649–2666. <https://doi.org/10.1177/0962280215604510>
- Eriksson F, Martinussen T, Scheike TH (2015) Clustered survival data with left-truncation. *Scand J Stat* 42:1149–1166. <https://doi.org/10.1111/sjost.12157>
- Ghosh D, Lin DY (2003) Semiparametric analysis of recurrent events data in the presence of dependent censoring. *Biometrics* 59:877–885. <https://doi.org/10.1111/j.0006-341X.2003.00102.x>
- Huang C-Y, Wang M-C (2004) Joint modeling and estimation for recurrent event processes and failure time data. *J Am Stat Assoc* 99:1153–1165. <https://doi.org/10.1198/016214504000001033>
- Jazić I, Haneuse S, French B, MacGrogan G, Rondeau V (2019) Design and analysis of nested case-control studies for recurrent events subject to a terminal event. *Stat Med* 38(22):4348–4362. <https://doi.org/10.1002/sim.8302>
- Jensen H, Brookmeyer R, Aaby P, Andersen PK (2004) Shared frailty model for left-truncated multivariate survival data. *Museum Tusculanum, Biostatistisk Afdeling*

- Kalbfleisch JD, Schaubel DE, Ye Y, Gong Q (2013) An estimating function approach to the analysis of recurrent and terminal events. *Biometrics* 69:366–374. <https://doi.org/10.1111/biom.12025>
- Klein JP, Moeschberger ML (2003) *Survival analysis. Techniques for censored and truncated data*, 2nd edn. Springer, Berlin
- Lawless JF, Fong DYT (1999) State duration models in clinical and observational studies. *Stat Med* 18:2365–2376
- Liu D, Schaubel DE, Kalbfleisch JD (2012) Computationally efficient marginal models for clustered recurrent event data. *Biometrics* 68:637–647. <https://doi.org/10.1111/j.1541-0420.2011.01676.x>
- Liu L, Huang X (2008) The use of Gaussian quadrature for estimation in frailty proportional hazards models. *Stat Med* 27:2665–2683. <https://doi.org/10.1002/sim.3077>
- Liu L, Wolfe RA, Huang X (2004) Shared frailty models for recurrent events and a terminal event. *Biometrics* 60:747–756. <https://doi.org/10.1111/j.0006-341X.2004.00225.x>
- Piulachs X, Andrinopoulou E-R, Guillén M, Rizopoulos D (2021) A Bayesian joint model for zero-inflated integers and left-truncated event times with a time-varying association: applications to senior health care. *Stat Med* 40(1):147–166. <https://doi.org/10.1002/sim.8767>
- R Core Team (2020) *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria
- Rogers JK, Yaroshinsky A, Pocock SJ, Stokar D, Pogoda J (2016) Analysis of recurrent events with an associated informative dropout time: application of the joint frailty model. *Stat Med* 35:2195–2205. <https://doi.org/10.1002/sim.6853>
- Rondeau V, Mathoulin-Pelissier S, Jacqmin-Gadda H, Brouste V, Soubeyran P (2007) Joint frailty models for recurring events and death using maximum penalized likelihood estimation: application on cancer events. *Biostatistics* 8(4):708–721. <https://doi.org/10.1093/biostatistics/kxl043>
- Rondeau V, Mauguén A, Laurent A, Berr C, Helmer C (2017) Dynamic prediction models for clustered and interval-censored outcomes: investigating the intra-couple correlation in the risk of dementia. *Stat Methods Med Res* 26(5):2168–2183
- Rondeau V, Gonzalez JR, Mazroui Y, Mauguén A, Diakite A, Laurent A, Lopez M, Król A, Sofeu CL, Dumerc J, Rustand D, Chauvet J, Coent, Le Q (2021) frailtypack: general frailty models: shared, joint and nested frailty models with prediction; evaluation of failure-time surrogate endpoints. R package version 3.5.0. <https://CRAN.R-project.org/package=frailtypack>
- van den Berg GJ, Drepper B (2016) Inference for shared-frailty survival models with left-truncated data. *Econ Rev* 35:1075–1098. <https://doi.org/10.1080/07474938.2014.975640>
- van den Hout A, Muniz-Terrera G (2016) Joint models for discrete longitudinal outcomes in aging research. *J R Stat Soc Ser C* 65:167–186. <https://doi.org/10.1111/rssc.12114>
- Wienke A (2011) *Frailty models in survival analysis*. Chapman & Hall/CRC Biostatistics Series, Boca Raton

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.