



Universiteit
Leiden
The Netherlands

Statistical methods for frailty models: studies on old-age mortality and recurrent events

Böhnstedt, M.

Citation

Böhnstedt, M. (2021, November 30). *Statistical methods for frailty models: studies on old-age mortality and recurrent events*. Retrieved from <https://hdl.handle.net/1887/3243966>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3243966>

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

6

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

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, the samples are often selected at a later time, which results in delayed entry. Thus, only individuals who have not yet experienced the terminal event will be included in the study. We propose a method for estimating the joint frailty model from such left-truncated data. The frailty distribution among the selected survivors differs from the frailty distribution in the underlying population if the recurrence process and the terminal event are associated. The correctly adjusted marginal likelihood can be expressed as a ratio of two integrals over the frailty distribution, which may be approximated using Gaussian quadrature. The baseline rates are specified as piecewise constant functions, and the covariates are assumed to have multiplicative effects on the event rates. We assess the performance of the

This chapter has been submitted for publication as: M. Böhstedt, J. Gampe, M.A.A. Caljouw, and H. Putter. Incorporating delayed entry into the joint frailty model for recurrent events and a terminal event.

estimation procedure in a simulation study, and apply the method to estimate age-specific rates of recurrent urinary tract infections and mortality in an older population.

6.1 Introduction

Repeated occurrences of the same type of event in one individual arise in various applications. Examples of such recurrent event data include incidents of myocardial infarction, recurrent infections, fractures, or tumor relapses.

If the individual is additionally at risk of experiencing a terminal event such as death, which will stop the recurrent event process, this might induce dependent censoring of the recurrence process. Therefore, approaches for jointly modeling the two processes of the recurrent events and the terminal event have been developed. Moreover, studies might explicitly address the question of whether there is an association between the processes by asking, for instance, whether individuals who experience more recurrences also have a higher risk of experiencing the terminal event. Specific joint models can provide additional insights into the direction and the strength of the association between the event processes.

The choice of the time scale t , along which the recurrence process and the terminal event process are assumed to evolve, depends on the specific application. In clinical studies, the time since randomization is often used as the time scale; whereas in demographic studies of fertility, for instance, the most relevant time scale is age.

For a given time scale, the event processes can, in some cases, be observed from the time origin t_0 . Consider, for example, a medical study in which the time since the disease onset or diagnosis is used as the time scale. If each individual is observed from the disease onset or diagnosis onwards, then all individuals enter the study at time $t_0 = 0$.

However, studies are often initiated at a later point in time when the two processes have already started. If in the above setting patients do not participate in the clinical study until some period of time after their diagnosis, individual times of study entry will differ from 0, and might also vary between patients. Similarly, in studies of certain diseases in old-age populations, such as in register-based studies of cardiovascular disease and mortality, age can be considered the natural time scale (see, for instance, Modig et al., 2013). The individuals included in such analyses have already reached a certain advanced age by the beginning of the study period.

In these cases, the data are left-truncated in the sense that the individuals can enter the study only if they have not yet experienced the terminal event. If the recurrence process and the terminal event process are associated, this sample of survivors is not a random sample of the underlying population. Rather, the sample is comprised of individuals who tend to have a lower risk of experiencing the terminal event, and – if there is a positive association between the two processes – to also have a lower risk of experiencing recurrent events. Thus, to obtain valid inferences, correctly adjusting for the truncation is crucial.

The example we use in this study to illustrate our method focuses on recurrent urinary tract infections (UTIs) in older residents of long-term care facilities (LTCF). The original

study by Caljouw et al. (2014) investigated whether cranberry capsules are effective in preventing UTIs. Given that around 34% of the elderly study population died during the follow-up period, recurrent UTIs and mortality had to be modeled jointly. The original study used time since randomization as the time scale. By contrast, we have chosen to use age as the main time scale of the event processes, as mortality and, presumably, UTI recurrences naturally depend on age. Because the participants were between 64 and 102 years old when they entered the study, the observations are then left-truncated. In our analysis, we seek to assess the effects of the cranberry treatment, to estimate the age-specific rates of UTIs and death, and to determine whether there is an association between UTIs and death risks.

While different joint models for recurrent events and death have been proposed, we focus here on the joint frailty model introduced by Liu et al. (2004). 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 extended in several directions (e.g., to the setting of nested case-control studies in Jazić et al., 2019). The joint frailty model allows us to examine the shape of the rates 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 rate 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. In other models, the frailty affects both event rates 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). A common approach to estimating the joint frailty model that we consider here is using Gaussian quadrature to approximate the marginal likelihood (Liu and Huang, 2008; Rondeau et al., 2007). We will see that in this framework, the setting with left-truncated data can be handled in a straightforward manner by adapting the likelihood.

Incorporating left truncation, which is also called delayed entry, has received varying levels of attention in studies based on different frailty models and 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). 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. It has been argued that for the joint frailty model for recurrent events and death, as specified by Liu et al. (2004), delayed entry can be easily incorporated (Rondeau et al., 2007). However, to our knowledge, no detailed account of this approach has previously been 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 3.3.2, date 2020-10-07, Rondeau et al., 2020). Extensions to left-truncated data have been considered

in some other joint models for recurrent events and death. 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. In the context of joint models for longitudinal data and death, estimation procedures based on left-truncated data have been derived for different models with shared random effects (see, for example, van den Hout and Muniz-Terrera, 2016; Crowther et al., 2016; Piulachs et al., 2021).

In this chapter, we propose a method for estimating the joint frailty model for recurrent events and a terminal event, as introduced by Liu et al. (2004), when data are left-truncated. Our approach adapts the method of Liu and Huang (2008), who use Gaussian quadrature to approximate the marginal likelihood of the joint frailty model.

The chapter is structured as follows. Section 6.2 first presents the joint frailty model and the corresponding likelihood in the usual setting without truncation, and then shows the adjustments for left-truncated data. The method of estimation is detailed in Section 6.3, and its performance is assessed via simulation studies in Section 6.4. We illustrate the approach using the data set on recurrent UTIs in Section 6.5, and conclude with a discussion in Section 6.6.

6.2 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.

6.2.1 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 or 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, $dN_i^{R^*}(t) = N_i^{R^*}((t + dt)^-) - N_i^{R^*}(t^-)$ is the increment of the recurrence process, equal to 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 the intensity $Y_i(t)\lambda_i(t|u_i)$ with

$$\begin{aligned} \mathbb{P}(dN_i^R(t) = 1 \mid \mathcal{F}_t^-, D_i \geq t) &= Y_i(t)\lambda_i(t|u_i)dt \\ \lambda_i(t|u_i)dt &= d\Lambda_i(t|u_i) = \mathbb{P}(dN_i^{R^*}(t) = 1 \mid \mathbf{z}_i, u_i, D_i \geq t). \end{aligned} \quad (6.1)$$

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

$$\begin{aligned} \mathbb{P}(dN_i^D(t) = 1 \mid \mathcal{F}_t^-) &= Y_i(t)h_i(t|u_i)dt \\ h_i(t|u_i)dt &= dH_i(t|u_i) = \mathbb{P}(dN_i^{D^*}(t) = 1 \mid \mathbf{z}_i, u_i, D_i \geq t). \end{aligned} \quad (6.2)$$

The first lines in (6.1) and (6.2) follow from the assumption that the censoring mechanism is conditionally independent of the two event processes given the process history.

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

$$\begin{aligned} \lambda_i(t|u_i) &= u_i e^{\boldsymbol{\beta}^\top \mathbf{z}_i} \lambda_0(t), \\ h_i(t|u_i) &= u_i^\gamma e^{\boldsymbol{\alpha}^\top \mathbf{z}_i} h_0(t). \end{aligned} \quad (6.3)$$

The baseline rates of recurrence and death, $\lambda_0(t)$ and $h_0(t)$, are affected by the covariates \mathbf{z}_i through a multiplicative model with effects $\boldsymbol{\beta}$ and $\boldsymbol{\alpha}$, respectively. The inclusion of the frailty u in the recurrence rate 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 rate of recurrence will also be subject to a higher hazard of death. For $\gamma < 0$, a higher rate of recurrence implies a lower hazard of death. If $\gamma = 0$, the intensities in (6.3) do not share any parameters, and the censoring of the recurrence process by death is non-informative.

A common choice for the distribution of the frailty U is a gamma distribution with a mean of one and a variance of θ . We will 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 (6.3) 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 or her frailty value u_i can be written as

$$\begin{aligned} L_i^{(c)}(u_i) &= \left[\prod_{j=1}^{J_i} \lambda_i(t_{ij}|u_i) \right] \exp \left\{ - \int_0^\infty Y_i(s) \lambda_i(s|u_i) ds \right\} \\ &\quad \cdot h_i(x_i|u_i)^{\delta_i} \exp \left\{ - \int_0^\infty Y_i(s) h_i(s|u_i) ds \right\} \\ &= \left[\prod_{j=1}^{J_i} u_i e^{\beta^\top z_i} \lambda_0(t_{ij}) \right] \exp \left\{ - \int_0^{x_i} u_i e^{\beta^\top z_i} \lambda_0(s) ds \right\} \\ &\quad \cdot \left[u_i^\gamma e^{\alpha^\top z_i} h_0(x_i) \right]^{\delta_i} \exp \left\{ - \int_0^{x_i} u_i^\gamma e^{\alpha^\top z_i} h_0(s) ds \right\}. \end{aligned} \quad (6.4)$$

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

$$L_i = \int_0^\infty L_i^{(c)}(u) dG_\theta(u) = \int_0^\infty L_i^{(c)}(u) g_\theta(u) du. \quad (6.5)$$

6.2.2 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 notations 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 or 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 Section 6.2.1 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 Section 6.2.1, 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 ${}_vY_i(t)\lambda_i(t|u_i)$ and (6.1) is adapted as

$$P(d {}_vN_i^R(t) = 1 \mid {}_v\mathcal{F}_{t-}, D_i \geq t) = {}_vY_i(t)\lambda_i(t|u_i). \quad (6.1')$$

The intensity of the observed terminal event process is, correspondingly, ${}_vY_i(t)h_i(t|u_i)$, such that (6.2) is modified as

$$P(d {}_vN_i^D(t) = 1 \mid {}_v\mathcal{F}_{t-}) = {}_vY_i(t)h_i(t|u_i). \quad (6.2')$$

Based on this, we can develop the likelihood of the joint frailty model (6.3) for left-truncated data. The conditional likelihood contribution of individual i given u_i is constructed in analogy to (6.4), with $Y_i(s)$ replaced by ${}_vY_i(s)$, and appropriately restricting to the observed recurrence times $t_{ij} \geq v_i$; that is,

$$\begin{aligned} {}_vL_i^{(c)}(u_i) &= \left[\prod_{t_{ij} \geq v_i} \lambda_i(t_{ij}|u_i) \right] \exp \left\{ - \int_0^\infty {}_vY_i(s)\lambda_i(s|u_i) ds \right\} \\ &\quad \cdot h_i(x_i|u_i)^{\delta_i} \exp \left\{ - \int_0^\infty {}_vY_i(s)h_i(s|u_i) ds \right\} \\ &= \left[\prod_{t_{ij} \geq v_i} u_i e^{\beta^\top z_i} \lambda_0(t_{ij}) \right] \exp \left\{ - \int_{v_i}^{x_i} u_i e^{\beta^\top z_i} \lambda_0(s) ds \right\} \\ &\quad \cdot \left[u_i^\gamma e^{\alpha^\top z_i} h_0(x_i) \right]^{\delta_i} \exp \left\{ - \int_{v_i}^{x_i} u_i^\gamma e^{\alpha^\top z_i} h_0(s) ds \right\}. \end{aligned} \quad (6.6)$$

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 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

$${}_vL_i = \int_0^\infty {}_vL_i^{(c)}(u) dG_\theta(u \mid D_i > v_i, V_i = v_i, z_i) = \int_0^\infty {}_vL_i^{(c)}(u) dG_\theta(u \mid D_i > v_i, z_i), \quad (6.7)$$

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

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 , $G_\theta(u \mid D > v)$, 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

$${}_vL_i^{\text{naive}} = \int_0^\infty {}_vL_i^{(c)}(u) dG_\theta(u) = \int_0^\infty {}_vL_i^{(c)}(u) g_\theta(u) du, \quad (6.8)$$

would lead to an invalid inference if $\gamma \neq 0$. We will illustrate the resulting biases in the parameter estimates in a simulation study in Section 6.4.

For computing the correct marginal likelihood (6.7), we first apply Bayes' theorem to find that

$$g_\theta(u \mid D_i > v_i, \mathbf{z}_i) = \frac{P(D_i > v_i \mid u, \mathbf{z}_i) g_\theta(u)}{P(D_i > v_i \mid \mathbf{z}_i)} = \frac{\exp\left\{-\int_0^{v_i} h_i(s \mid u) ds\right\} g_\theta(u)}{\int_0^\infty P(D_i > v_i \mid u) g_\theta(u) du}, \quad (6.9)$$

where we suppress the dependence on the covariates \mathbf{z}_i in the last expression for notational convenience. Combining equations (6.6), (6.7), and (6.9), we can express the marginal likelihood contribution ${}_vL_i$ of individual i with left-truncated data as

$$\frac{\int_0^\infty \left[\prod_{t_{ij} \geq v_i} \lambda_i(t_{ij} \mid u) \right] \exp\left\{-\int_{v_i}^{x_i} \lambda_i(s \mid u) ds\right\} h_i(x_i \mid u)^{\delta_i} \exp\left\{-\int_0^{x_i} h_i(s \mid u) ds\right\} g_\theta(u) du}{\int_0^\infty P(D_i > v_i \mid u) g_\theta(u) du}. \quad (6.10)$$

Interestingly, the formula for ${}_vL_i$ in (6.10) 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

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

As the integrals over the frailty distribution in (6.10) will not, in general, have closed-form expressions, we will use numerical integration in the following.

6.3 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 (6.3). In combination with a piecewise constant specification of the baseline rates, this approach allows for the direct maximization of the approximated likelihood.

The aim of Gaussian quadrature is to replace the integral of a function with a weighted sum of function values. The Gauss-Hermite quadrature rule gives an approximation for a specific integral of a function $f(x)$,

$$\int_{-\infty}^{\infty} f(x)e^{-x^2} dx \approx \sum_{q=1}^Q w_q f(x_q).$$

The quadrature points x_q need to be determined as the roots of the Q^{th} -order Hermite polynomial, while w_q specify corresponding weights. As the marginal likelihood of a model with normal random effects is easily rewritten in the above form, it follows that such a likelihood can be approximated as

$$\int_{-\infty}^{\infty} L^{(c)}(b) \phi(b) db \approx \sum_{q=1}^Q \tilde{w}_q L^{(c)}(\tilde{x}_q) \phi(\tilde{x}_q), \quad (6.11)$$

where $\phi(\cdot)$ denotes the standard normal density, while $\tilde{x}_q = \sqrt{2}x_q$ and $\tilde{w}_q = \sqrt{2}w_q e^{x_q^2}$ are modified quadrature points and weights, respectively. Marginal likelihoods integrated over non-normal random effects can be expressed as an integral of the form in (6.11) by applying the probability integral transformation (see Nelson et al., 2006; Liu and Huang, 2008). We show in Section 6.7.1 how this quadrature approach can be used to derive an approximation of the marginal likelihood of the joint frailty model with left truncation.

The evaluation of the approximate marginal likelihood depends on the specific form of the baseline rates $\lambda_0(t)$ and $h_0(t)$. As suggested by Liu and Huang (2008), 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 rates were assumed to follow a simple parametric model, such as the Weibull model. Nonetheless, we recommend the use of the more flexible piecewise constant rate 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. 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 6.7.2.

6.4 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.

6.4.1 Settings

In the base scenario, we generated data from a joint frailty model (6.3). The time scale t is the age of the individual. The hazard of death and the rate of recurrence 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 parameter 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 rate does not affect the mortality risk ($\gamma = 0$).

The baseline rates $h_0(t)$ and $\lambda_0(t)$ were designed to mimic a study in an older population among whom the death rates as well as the recurrent event rates increase exponentially with age. Hence, we chose for both baseline rates 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 rates were comparable to the estimated rates for the high risk group in the data example in Section 6.5.

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 or 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 (6.3) 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 Figure 6.1 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 Section 6.5. 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 Figure 6.1 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 individuals at risk across ages is considerably smaller than the total sample size of about 500.

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 Figure 6.1). 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 rates were expected

to agree with the rates 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 6.7.3.




Figure 6.1: 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.

6.4.2 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 rates $\lambda_0(t)$ and $h_0(t)$. For both rates, 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 6.7.2.

Figures 6.2 and 6.3 illustrate the results of the base scenario with different underlying associations, $\gamma \in \{-0.5, 0, 0.5\}$. The top panels of Figure 6.2 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 Figure 6.2 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. If the recurrence process and the mortality process are positively associated, this implies that the frailty values and the recurrence rates 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 rate, 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 rate 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.

The estimates of the baseline rates, displayed for the base scenario with positive dependence $\gamma = 0.5$ in Figure 6.3, 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 Figures 6.4 and 6.5 show, if the inference is based on the naively constructed marginal likelihood (6.8), 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. The baseline rates of recurrence and death are increasingly underestimated for advancing age in the base scenario with positive dependence ($\gamma = 0.5$), as depicted in Figure 6.5. 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 rate 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 (6.7), introduces marked biases in the estimates and the standard errors. Only if the event rates 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 (6.7) 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 6.7.3. 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 6.8 and 6.9).




Figure 6.2: 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.




Figure 6.3: Estimates (gray) of the baseline rate 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 rate.

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 rates.

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, occasionally with negative estimates (see Figure 6.10). In addition, the first piece of each of the baseline rates shows an upward bias (cf. top panels of Figure 6.11) because few individuals entered the study at the younger ages. Although the intervals for the rate pieces 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 6.10).




Figure 6.4: 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.




Figure 6.5: Estimates (gray) of the baseline rate 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 rate.

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.

6.5 Recurrent infections and mortality in an older population

We use the proposed method to analyze recurrent urinary tract infections and mortality in an institutionalized elderly population. The data come from a double-blind, randomized, placebo-controlled trial in long-term care facilities that aimed to assess the effect of cranberry capsules on the occurrence of UTIs in vulnerable older persons (Caljouw et al., 2014). At baseline, 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. The participants took cranberry or placebo capsules twice a day over a period of one year. Occurrences of UTIs were recorded by

the treating physicians according to a clinical definition based on international practice guidelines for LTCF residents, and a strict definition based on scientific criteria. We focus here on the occurrence of the more broadly defined clinical UTIs.

The final study population 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 considered to be at low baseline UTI risk. Individuals entered the study between ages 64 and 102, as shown in the left panel of Figure 6.6, and were followed on average for about a year (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.

Unlike in the original study, we modeled recurrent UTIs and mortality to evolve with age, where $t_0 = 0$ corresponded to age 64. Because of the specific distribution of the ages at study entry in conjunction with the short individual follow-up times, relatively few individuals were under observation at a given age, in particular at the youngest and oldest ages, as the right panel of Figure 6.6 shows.




Figure 6.6: Distribution of ages at study entry (left) and 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).

We estimated the joint frailty model for UTIs and overall mortality with age as the time scale 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. The baseline rate 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 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 6.1. In the group with a high baseline UTI risk, the infection rates varied between participants with an estimated frailty variance of $\hat{\theta} = 0.380$ (SE: 0.086). In particular, individuals with a higher rate 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 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 infection rates, while males had higher mortality levels than females in both groups.

Table 6.1: 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)

The estimated baseline rates displayed in Figure 6.7 demonstrate nicely the age dependence of the recurrence rate and the hazard of death. For the individuals with a high baseline UTI risk, both the rate 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 individuals with a high baseline UTI risk tended to experience higher rates of recurrent infection and death than the individuals with a low baseline UTI risk.

The original study, which used a different time 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.




Figure 6.7: Estimated baseline rates (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).

6.6 Discussion

We have proposed a method for estimating the joint frailty model for recurrent events and a terminal event based on left-truncated data. The marginal likelihood of the model can be expressed as a ratio of two integrals over the frailty distribution, each of which is approximated using Gauss-Hermite quadrature. The direct maximization of the approximate marginal likelihood is possible if the baseline rates are specified as piecewise constant functions.

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 rates 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, we need to choose a frailty distribution as well as the number of quadrature points and the intervals for the baseline rates. 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 number of quadrature points then 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. Regarding the baseline rate functions, the number of intervals corresponds to the number of parameters for the rates, and should thus be selected to allow for sufficient flexibility of the shape of the rates, while retaining numerical stability and the computational feasibility of the method. Adequate results can often be obtained with moderate numbers of up to 10 intervals.

Nevertheless, in some applications, it may seem appealing to aim for smooth rate estimates, such as through the use of penalized splines. However, automatic smoothing parameter selection in joint frailty models is an issue that needs further investigation.

Moreover, the current approach is limited to applications in which there is heterogeneity in recurrence rates. 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.

6.7 Supplementary material

6.7.1 Approximation of the marginal likelihood using Gaussian quadrature

In this section, we will elaborate on the use of Gauss-Hermite quadrature for approximating the marginal likelihood of the joint frailty model. We will first recap the quadrature approach as proposed by Liu and Huang (2008) in the setting with right-censoring only, and then show how to adapt the method to the setting with left truncation.

An approximation to marginal likelihoods integrated over normal random effects based on Gauss-Hermite quadrature was already presented in Section 6.3. But the marginal likelihoods of the joint frailty model given in Section 6.2 involve integrals over non-normal random effects. Thus, we use the probability integral transformation (Nelson et al., 2006; Liu and Huang, 2008) to rewrite the integrals over the random effect u with distribution function $G_\theta(u)$ as integrals over standard normal random effects. This relies on the fact that the $G_\theta(u)$ have a standard uniform distribution, such that their transformations $a = \Phi^{-1}[G_\theta(u)]$ follow a standard normal distribution, if $\Phi(\cdot)$ denotes the standard normal distribution function.

For the marginal likelihood contribution (6.5) of individual i in the joint frailty model without truncation, the substitution $u = G_\theta^{-1}[\Phi(a)]$ yields

$$L_i = \int_0^\infty L_i^{(c)}(u) g_\theta(u) du = \int_{-\infty}^\infty L_i^{(c)}(G_\theta^{-1}[\Phi(a)]) \phi(a) da.$$

We can then employ Gauss-Hermite quadrature as in (6.11) to arrive at the approximate marginal likelihood contribution of the form

$$L_i \approx \sum_{q=1}^Q L_i^{(c)}(G_\theta^{-1}[\Phi(\tilde{x}_q)]) \phi(\tilde{x}_q) \tilde{w}_q.$$

In the setting with left-truncated data, the marginal likelihood of the joint frailty model has a slightly more complex structure. In (6.10), the marginal likelihood contribution ${}_vL_i$ of individual i is expressed as a ratio of two integrals over the density $g_\theta(u)$. Hence, we will approximate the likelihood by applying the above approach separately to the two integrals,

$${}_vL_i = \frac{\int_0^\infty {}_vL_i^{(c)}(u) \exp\{-H_i(v_i|u)\} g_\theta(u) du}{\int_0^\infty \exp\{-H_i(v_i|u)\} g_\theta(u) du}$$

$$\approx \frac{\sum_{q=1}^Q vL_i^{(c)}(G_\theta^{-1}[\Phi(\tilde{x}_q)]) \exp\{-H_i(v_i|G_\theta^{-1}[\Phi(\tilde{x}_q)])\} \phi(\tilde{x}_q) \tilde{w}_q}{\sum_{q=1}^Q \exp\{-H_i(v_i|G_\theta^{-1}[\Phi(\tilde{x}_q)])\} \phi(\tilde{x}_q) \tilde{w}_q},$$

with $vL_i^{(c)}(u)$ given in (6.6). The approximate marginal likelihood of the joint frailty model with left truncation is then given by

$$\prod_{i=1}^{m_V} \frac{\sum_{q=1}^Q vL_i^{(c)}(G_\theta^{-1}[\Phi(\tilde{x}_q)]) \exp\{-H_i(v_i|G_\theta^{-1}[\Phi(\tilde{x}_q)])\} \phi(\tilde{x}_q) \tilde{w}_q}{\sum_{q=1}^Q \exp\{-H_i(v_i|G_\theta^{-1}[\Phi(\tilde{x}_q)])\} \phi(\tilde{x}_q) \tilde{w}_q}.$$

In the naive likelihood (6.8), there is only one integral over the frailty distribution, and hence only one approximation is required,

$$vL_i^{\text{naive}} = \int_0^\infty vL_i^{(c)}(u) g_\theta(u) du \approx \sum_{q=1}^Q vL_i^{(c)}(G_\theta^{-1}[\Phi(\tilde{x}_q)]) \phi(\tilde{x}_q) \tilde{w}_q.$$

6.7.2 Computational details

We used R (R Core Team, 2020) to implement the estimation procedure. The quadrature points and weights were calculated using function `gauss.quad()` from package `statmod` (Smyth, 1998). For numerical optimization of the approximate marginal log-likelihood, we applied function `nlm()`, which performs minimization based on a Newton-type algorithm, to the negative log-likelihood. The Hessian of the marginal log-likelihood was approximated numerically using function `hessian()` from package `numDeriv` (Gilbert and Varadhan, 2019).

As the frailty variance and the parameters of the piecewise constant baseline rates are restricted to be non-negative, the log-likelihood was maximized with respect to the log-transform of these parameters, which guaranteed non-negative estimates. The delta-method was then applied to derive the respective standard errors.

However, when specifying the baseline rates as piecewise constant functions, the numerically computed Hessian of the log-likelihood may in some cases not be invertible. To overcome this issue, one can add to the log-likelihood small, fixed ridge penalties (e.g., with penalty parameter 10^{-6}) on the logarithm of the rate parameters and derive the Hessian matrix from this penalized log-likelihood.

6.7.3 Supplement to the simulation studies

Generation of different truncation patterns

A brief description of how the truncation patterns A and B were generated was already given in Section 6.4. For both patterns, truncation times were simulated from a truncated normal distribution defined on the corresponding age range [75, 95] or [64, 105]. However, the parameters μ and σ^2 of the underlying normal distribution were chosen such that

the density of the resulting truncated normal distribution was either increasing over the whole age range with mode equal to the maximum age or unimodal with a mode within the age range. These two distinct shapes yielded the desired truncation patterns A (TrA) or B (TrB) in the final samples. Table 6.2 reports, for the different settings, the parameter values for the distribution of the truncation times as well as the initial sample size M needed to obtain truncated samples with an average size of $m_V = 500$.

When implementing the simulation study in R, we used function `rttruncnorm()` from package `truncnorm` (Mersmann et al., 2018) for drawing random numbers from a truncated normal distribution.

Table 6.2: Initial sample size M and parameter values for the distribution of the truncation times.

γ	age range	pattern	M	μ	σ^2
0.5	[75, 95]	TrA	$1.07 \cdot 10^4$	109	124
0.5	[75, 95]	TrB	$1.20 \cdot 10^4$	90	18
0.5	[64, 105]	TrA	$7.65 \cdot 10^4$	120	225
0.5	[64, 105]	TrB	$4.75 \cdot 10^4$	93	52
0	[75, 95]	TrA	$3.50 \cdot 10^4$	115	110
-0.5	[75, 95]	TrA	$3.44 \cdot 10^4$	115	109

Additional figures for the simulation results

The following figures illustrate the impact of different censoring and truncation patterns on the performance of the estimation procedure. All settings are modifications of the base scenario with positive dependence $\gamma = 0.5$, that was presented in Section 6.4. In particular, samples with a target size of $m_V = 500$ were generated from a joint frailty model with covariate effects $\alpha = \beta = 0.5$, frailty variance $\theta = 0.5$, and Gompertz-Makeham baseline rates.

- In Figures 6.8 and 6.9, the effect of different censoring mechanisms is studied by comparing the estimation results for different lengths of the planned individual follow-up $t_C \in \{1, 4, 8\}$ years.
- Changes in the distribution of truncation times are examined in Figures 6.10 and 6.11. The displayed settings assume different distributions of the study entry times in the final sample, both in terms of the shape (TrA: truncation pattern A, or TrB: truncation pattern B) and the support of the distribution (75+: ages 75 to 95, or 64+: ages 64 to 105).

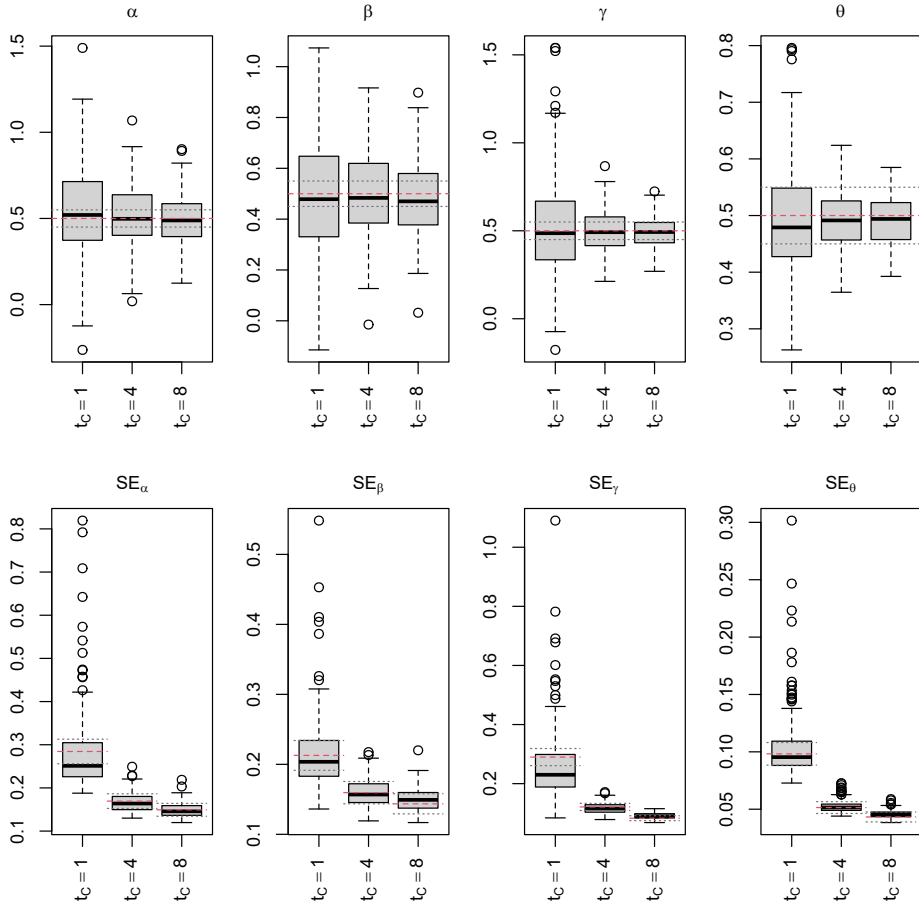


Figure 6.8: Box plots of the parameter estimates (top) and estimated standard errors (bottom) in the joint frailty model with positive dependence ($\gamma = 0.5$) as in the base scenario with truncation pattern A for ages 75+, but varying the planned individual follow-up $t_C \in \{1, 4, 8\}$. Left to 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.




Figure 6.9: Estimates (gray) of the baseline rate 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 ($\gamma = 0.5$). As in the base scenario truncation follows pattern A for ages 75+, but planned individual follow-up is $t_C = 1$ (top) or $t_C = 8$ (bottom) year(s). The red bold line gives the true baseline rate. (Note the different scales of the vertical axes.)




Figure 6.10: Box plots of the parameter estimates (top) and estimated standard errors (bottom) in the joint frailty model with positive dependence ($\gamma = 0.5$) as in the base scenario with $t_C = 4$, but truncation times distributed with different shapes (TrA: truncation pattern A, TrB: truncation pattern B) and across different age ranges (75+: [75, 95], 64+: [64, 105]). 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.




Figure 6.11: Estimates (gray) of the baseline rate 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 ($\gamma = 0.5$) and planned individual follow-up of $t_C = 4$ years as in the base scenario. Truncation times are distributed according to pattern B across ages [75, 95] (top; TrB, 75+) or according to pattern A across ages [64, 105] (bottom; TrA, 64+). The red bold line gives the true baseline rate. (Note the different scales of the horizontal and vertical axes.)

References

- Balan, T. A., M. A. Jonker, P. C. Johannesma, and H. Putter (2016). Ascertainment correction in frailty models for recurrent events data. *Statistics in Medicine* 35, 4183–4201.
- Cai, Q., M.-C. Wang, and K. C. G. Chan (2017). Joint modeling of longitudinal, recurrent events and failure time data for survivor's population. *Biometrics* 73, 1150–1160.
- Caljouw, M. A. A., W. B. van den Hout, H. Putter, W. P. Achterberg, H. J. M. Cools, and J. Gussekloo (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. *Journal of the American Geriatrics Society* 62, 103–110.
- Cook, R. J. and J. F. Lawless (1997). Marginal analysis of recurrent events and a terminating event. *Statistics in Medicine* 16, 911–924.
- Cook, R. J. and J. F. Lawless (2007). *The Statistical Analysis of Recurrent Events*. Statistics for Biology and Health. New York: Springer.
- Crowther, M. J., T. M.-L. Andersson, P. C. Lambert, K. R. Abrams, and K. Humphreys (2016). Joint modelling of longitudinal and survival data: incorporating delayed entry and an assessment of model misspecification. *Statistics in Medicine* 35, 1193–1209.
- Emura, T., M. Nakatochi, K. Murotani, and V. Rondeau (2017). A joint frailty-copula model between tumour progression and death for meta-analysis. *Statistical Methods in Medical Research* 26, 2649–2666.
- Eriksson, F., T. Martinussen, and T. H. Scheike (2015). Clustered survival data with left-truncation. *Scandinavian Journal of Statistics* 42, 1149–1166.
- Ghosh, D. and D. Y. Lin (2003). Semiparametric analysis of recurrent events data in the presence of dependent censoring. *Biometrics* 59, 877–885.
- Gilbert, P. and R. Varadhan (2019). *numDeriv: Accurate Numerical Derivatives*. R package version 2016.8-1.1. <https://CRAN.R-project.org/package=numDeriv>.
- Huang, C.-Y. and M.-C. Wang (2004). Joint modeling and estimation for recurrent event processes and failure time data. *Journal of the American Statistical Association* 99, 1153–1165.
- Jazić, I., S. Haneuse, B. French, G. MacGrogan, and V. Rondeau (2019). Design and analysis of nested case-control studies for recurrent events subject to a terminal event. *Statistics in Medicine* 38(22), 4348–4362.

- Jensen, H., R. Brookmeyer, P. Aaby, and P. K. Andersen (2004). *Shared Frailty Model for Left-Truncated Multivariate Survival Data*. Biostatistisk Afdeling: Museum Tusulanum.
- Lawless, J. F. and D. Y. T. Fong (1999). State duration models in clinical and observational studies. *Statistics in Medicine* 18, 2365–2376.
- Liu, D., D. E. Schaubel, and J. D. Kalbfleisch (2012). Computationally efficient marginal models for clustered recurrent event data. *Biometrics* 68, 637–647.
- Liu, L. and X. Huang (2008). The use of Gaussian quadrature for estimation in frailty proportional hazards models. *Statistics in Medicine* 27, 2665–2683.
- Liu, L., R. A. Wolfe, and X. Huang (2004). Shared frailty models for recurrent events and a terminal event. *Biometrics* 60, 747–756.
- Mersmann, O., H. Trautmann, D. Steuer, and B. Bornkamp (2018). *truncnorm: Truncated Normal Distribution*. R package version 1.0-8. <https://CRAN.R-project.org/package=truncnorm>.
- Modig, K., T. Andersson, S. Drefahl, and A. Ahlbom (2013). Age-specific trends in morbidity, mortality and case-fatality from cardiovascular disease, myocardial infarction and stroke in advanced age: Evaluation in the Swedish population. *PLOS ONE* 8. doi: 10.1371/journal.pone.0064928.
- Nelson, K. P., S. R. Lipsitz, G. M. Fitzmaurice, J. Ibrahim, M. Parzen, and R. Strawderman (2006). Use of the probability integral transformation to fit nonlinear mixed-effects models with nonnormal random effects. *Journal of Computational and Graphical Statistics* 15(1), 39–57.
- Piulachs, X., E.-R. Andrinopoulou, M. Guillén, and D. Rizopoulos (2021). A Bayesian joint model for zero-inflated integers and left-truncated event times with a time-varying association: Applications to senior health care. *Statistics in Medicine* 40(1), 147–166.
- R Core Team (2020). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Rogers, J. K., A. Yaroshinsky, S. J. Pocock, D. Stokar, and J. Pogoda (2016). Analysis of recurrent events with an associated informative dropout time: Application of the joint frailty model. *Statistics in Medicine* 35, 2195–2205.
- Rondeau, V., J. R. Gonzalez, Y. Mazroui, A. Mauguen, A. Diakite, A. Laurent, M. Lopez, A. Król, C. L. Sofeu, J. Dumerc, and D. Rustand (2020). *frailtypack: General Frailty Models: Shared, Joint and Nested Frailty Models with Prediction; Evaluation of Failure-Time Surrogate Endpoints*. R package version 3.3.2. <https://CRAN.R-project.org/package=frailtypack>.

-
- Rondeau, V., S. Mathoulin-Pelissier, H. Jacqmin-Gadda, V. Brouste, and P. Soubeyran (2007). Joint frailty models for recurring events and death using maximum penalized likelihood estimation: application on cancer events. *Biostatistics* 8(4), 708–721.
- Smyth, G. K. (1998). Numerical integration. In P. Armitage and T. Colton (Eds.), *Encyclopedia of Biostatistics*, pp. 3088–3095. London: Wiley.
- van den Berg, G. J. and B. Drepper (2016). Inference for shared-frailty survival models with left-truncated data. *Econometric Reviews* 35, 1075–1098.
- van den Hout, A. and G. Muniz-Terrera (2016). Joint models for discrete longitudinal outcomes in aging research. *Journal of the Royal Statistical Society, Series C* 65, 167–186.

