Published in final edited form as:

Scand Stat Theory Appl. 2017 March ; 44(1): 112–129. doi:10.1111/sjos.12244.

A Semi-parametric Transformation Frailty Model for Semi-competing Risks Survival Data

Fei Jiang and

Department of Statistics, University of South Carolina

Sebastien Haneuse

Department of Biostatistics, Harvard University

Abstract

In the analysis of semi-competing risks data interest lies in estimation and inference with respect to a so-called non-terminal event, the observation of which is subject to a terminal event. Multi-state models are commonly used to analyse such data, with covariate effects on the transition/intensity functions typically specified via the Cox model and dependence between the non-terminal and terminal events specified, in part, by a unit-specific shared frailty term. To ensure identifiability, the frailties are typically assumed to arise from a parametric distribution, specifically a Gamma distribution with mean 1.0 and variance, say, σ^2 . When the frailty distribution is misspecified, however, the resulting estimator is not guaranteed to be consistent, with the extent of asymptotic bias depending on the discrepancy between the assumed and true frailty distributions. In this paper, we propose a novel class of transformation models for semi-competing risks analysis that permit the non-parametric specification of the frailty distribution. To ensure identifiability, the class restricts to parametric specifications of the transformation and the error distribution; the latter are flexible, however, and cover a broad range of possible specifications. We also derive the semi-parametric efficient score under the complete data setting and propose a non-parametric score imputation method to handle right censoring; consistency and asymptotic normality of the resulting estimators is derived and small-sample operating characteristics evaluated via simulation. Although the proposed semi-parametric transformation model and non-parametric score imputation method are motivated by the analysis of semi-competing risks data, they are broadly applicable to any analysis of multivariate time-to-event outcomes in which a unit-specific shared frailty is used to account for correlation. Finally, the proposed model and estimation procedures are applied to a study of hospital readmission among patients diagnosed with pancreatic cancer.

Keywords

frailty; misspecification; multivariate survival analysis; semi-competing risks; semi-parametric models; transformation models

---

Fei Jiang, Department of Statistics, University of South Carolina, Columbia, SC, USA. homebovine@gmail.com.

Supporting information

Additional supporting information may be found in the online version of this article at the publishers web site.

1. Introduction

In biomedical studies, a patient may experience multiple failure types. One specific setting is where interest lies primarily with some so-called *non-terminal* event, the observation of which is subject to some *terminal* event. Consider, for example, the study of end-of-life quality of care for patients diagnosed with pancreatic cancer. In health policy literature, an important marker of quality of care is readmission to the emergency room following discharge from the hospital at which the diagnosis was given (Vest *et al.*, 2010; Warren *et al.*, 2011). In this setting, readmission is a non-terminal event in the sense that experiencing a readmission event does not preclude the patient from experiencing other relevant events including death. In contrast, death is terminal in the sense that experiencing a death event precludes the patient from subsequently experiencing the non-terminal readmission event, indeed any event. Data arising in this context are often referred to as *semi-competing risks data* (Fine *et al.*, 2001).

Notationally, let T_1 and T_2 be the non-terminal and terminal event times, respectively. Fine *et al.* (2001) showed that while the marginal distribution of T_2 is identifiable from semi-competing risks data, the marginal distribution of T_1 is not. Furthermore, the joint distribution of (T_1, T_2) is only identifiable on the upper wedge of the joint support. Towards analyzing semi-competing risks data, two general statistical frameworks have been proposed in the literature. The first imposes models directly on the marginal distributions of T_1 and T_2 , with the joint distribution specified via a copula (Fine *et al.*, 2001; Wang, 2003; Jiang *et al.*, 2005; Ghosh, 2006; Peng & Fine, 2007; Lakhal *et al.*, 2008; Hsieh *et al.*, 2008). The second framework views semi-competing risks data as arising from an underlying multi-state model, specifically the illness–death model (Liu *et al.*, 2004; Ye *et al.*, 2007; Zeng & Lin, 2009; Xu *et al.*, 2010; Zeng *et al.*, 2012; Lee *et al.*, 2015). Figure 1 provides a representation of this model.

In an illness–death multi-state model, the transitions between the three ‘states’ are characterized by three intensity or hazard functions, specifically:

$$\lambda_1(t_1) = \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \Pr(T_1 = t_1 + \Delta | T_1 \geq t_1, T_2 \geq t_1),$$

$$\lambda_2(t_2) = \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \Pr(T_2 = t_2 + \Delta | T_1 \geq t_2, T_2 \geq t_2), \quad (1)$$

$$\lambda_3(t_2 | t_1) = \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \Pr(T_2 = t_2 + \Delta | T_1 = t_1, T_2 \geq t_2).$$

In practice, models must be specified for each of these functions. One specification considered by Xu *et al.* (2010) was to model the dependence of these hazard functions on covariates, denoted $\mathbf{X}$, via Cox-type multiplicative models of the form:

$$\lambda_1(t_1|\mathbf{X}, \Gamma) = \Gamma \lambda_{01}(t_1) \exp(\boldsymbol{\beta}_1^T \mathbf{X}),$$

$$\lambda_2(t_2|\mathbf{X}, \Gamma) = \Gamma \lambda_{02}(t_2) \exp(\boldsymbol{\beta}_2^T \mathbf{X}), \quad (2)$$

$$\lambda_3(t_2|t_1, \mathbf{X}, \Gamma) = \Gamma \lambda_{03}(t_2) \exp(\boldsymbol{\beta}_3^T \mathbf{X}),$$

where $\lambda_{01}(t_1)$, $\lambda_{02}(t_2)$, and $\lambda_{03}(t_2)$ are baseline hazard functions and Γ is a patient-specific random effect or frailty. Intuitively, the inclusion of the Γ frailty terms in the model accounts for dependence between T_1 and T_2 that is induced marginally by covariates not included in $\mathbf{X}$. Because the baseline hazard functions are typically left unspecified, however, identifiability requires imposing a parametric assumption on the distribution of the frailties across the population (Hougaard, 1984; McGilchrist & Aisbett, 1991). If the frailties are assumed to follow a Gamma distribution, as Xu *et al.* (2010) do, the resulting (integrated) likelihood has the advantage of having a closed form expression and is therefore relatively simple to implement and maximize (Nielsen *et al.*, 1992; Kosorok *et al.*, 2004; Su & Wang, 2012). It is well-known, however, that if a parametric random effect/frailty distribution is misspecified, then estimation for other model components is no longer guaranteed to be consistent (White, 1982). This phenomenon has been extensively considered in the literature on related generalized linear mixed models (Neuhaus *et al.*, 1992; Heagerty & Kurland, 2001; Agresti *et al.*, 2004; Litière *et al.*, 2007; Neuhaus *et al.*, 2011) as well as on shared frailty models (Shih & Louis, 1995; Glidden, 1999; Agresti *et al.*, 2004; Cui & Sun, 2004). To our knowledge, however, potential misspecification of the frailty distribution has not been considered in the context of the analysis of semi-competing risks data.

Towards relaxing distributional assumptions regarding the shared frailty terms in an illness–death model for semi-competing risks data, we consider a novel analysis approach based on transformation models. In the literature, these models are motivated by a desire to move beyond the proportional hazards assumption. Two classes of transformation models for univariate and multivariate survival outcomes have been considered: the first considers transformations of the hazard function (Zeng & Lin, 2007; Zeng *et al.*, 2009) and the second considers transformations of the survival outcomes themselves, a class which includes the well-known accelerated failure time (AFT) model (Pan, 2001). In this paper, we build on the latter, focusing on developing models that consider the impact of covariates directly on the event times (rather than on the hazard function). The proposed framework is developed specifically to permit the frailty distribution to be non-parametric. To balance identifiability, we restrict to a class of flexible parametric transformation functions and error distributions. Towards performing estimation and inference, we derive the semi-parametric efficient score function under the complete data setting (i.e. in the absence of right censoring). To accommodate right censoring, we develop a score imputation-based approach for estimation and inference. Specifically, extending prior work in the univariate setting when covariates are discrete (Wang *et al.*, 2012), the proposed method is developed for the semi-competing

risks data where $\mathbf{X}$ can be a mixture of discrete and continuous covariates. For the latter, we derive the asymptotic properties of the resulting estimators and evaluate their small-sample operating characteristics via simulation. Finally, the data are applied to a motivating data set of readmission and death times for patients diagnosed with pancreatic cancer.

2. A semi-parametric transformation model

In statistical practice, the Cox model and AFT models are by far the most extensively used for investigating time-to-event outcomes. Both of these models can be shown to belong to a broader class of linear transformation models (Horowitz, 2010). To see this, let T denote a univariate time-to-event outcome and consider the model $G(T) = \beta^T \mathbf{X} + \xi$, where G is a monotone transformation function and ξ is a random error term. If one sets $G(\cdot) = \log(\cdot)$ and leaves the distribution of ξ arbitrary, then this model is equivalent to the AFT model; if one sets $G(\cdot) = \log \Lambda_0(\cdot)$, with $\Lambda_0(\cdot)$ an arbitrary cumulative hazard function, and assumes that ξ follows an extreme value distribution, then the induced model for the hazard is equivalent to the Cox model.

Extending the univariate transformation model to the semi-competing risks data setting, consider the following semi-parametric linear transformation model for the i^{th} subject:

$$G_{\alpha_1}(T_{i1}) = \beta_1^T \mathbf{X}_i + \xi_{i1},$$

$$G_{\alpha_2}(T_{i2}) = \beta_2^T \mathbf{X}_i + \xi_{i2}, \text{ if the non-terminal event has not occurred,} \quad (3)$$

$$G_{\alpha_3}(T_{i2}) = \beta_3^T \mathbf{X}_i + \xi_{i3}, \text{ if the non-terminal event has occurred,}$$

where $\mathbf{X}_i = (X_{i0}, \dots, X_{ip})$ is a $(p+1)$ -dimensional covariate vector with the first entry at 1, $G_{\alpha_l}(\cdot)$ is a transition-specific transformation function, $l = 1, 2, 3$, and ξ_{il} follows a distribution for which the cumulative hazard is defined by $\Gamma_l A_{\kappa_l}(\cdot)$, where $\kappa_l(\cdot)$ is parametric function indexed by an unknown finite-dimensional parameter κ_l . As with the univariate transformation model, one can show that the class of models given by (3) includes models previously described in the literature. For example, the Xu *et al.* (2010) model given by (2) is a special case of (3) with $G_{\alpha_l}(\cdot)$ taken to be the log cumulative hazard function for the l^{th} transition and the $A_{\kappa_l}(\cdot) = \exp(\cdot)$ (i.e. each of the ξ_{il} error terms is taken to be distributed according to an extreme value distribution). Letting $\theta_l = (\alpha_l^T, \kappa_l^T, \beta_l^T)^T$, we note that the three transition-specific hazard functions corresponding to (3) have the following relatively simple representations:

$$\lambda_{\theta_1}(t_1 | \mathbf{X}_i, \Gamma_i) = \Gamma_i \lambda_{\theta_1}(t_1 | \mathbf{X}_i) = \Gamma_i \dot{A}_{\kappa_1} \{G_{\alpha_1}(t_1) - \beta_1^T \mathbf{X}_i\} \dot{G}_{\alpha_1}(t_1),$$

$$\lambda_{\theta_2}(t_2|\mathbf{X}_i, \Gamma_i) = \Gamma_i \lambda_{\theta_2}(t_2|\mathbf{X}_i) = \Gamma_i \dot{A}_{\kappa_2} \{G_{\alpha_2}(t_2) - \beta_2^T \mathbf{X}_i\} \dot{G}_{\alpha_2}(t_2), \quad (4)$$

$$\lambda_{\theta_3}(t_2|t_1, \mathbf{X}_i, \Gamma_i) = \Gamma_i \lambda_{\theta_3}(t_2|\mathbf{X}_i) = \Gamma_i \dot{A}_{\kappa_3} \{G_{\alpha_3}(t_2) - \beta_3^T \mathbf{X}_i\} \dot{G}_{\alpha_3}(t_2),$$

where $\dot{A}_{\kappa}(u) = dA_{\kappa}(u)/du$ and $\dot{G}_{\alpha}(t) = dG_{\alpha}(t)/dt$. Note that the final equality in the expression for $\lambda_{\theta_3}(\cdot)$ does not depend on t_1 . This corresponds to the so-called Markov assumption in which the hazard function for the transition from the non-terminal event to the terminal event is independent of when the non-terminal occurred. While this is a relatively standard assumption in the semi-competing risks literature (e.g. Xu *et al.*, 2010), we return to this point in the Discussion.

As we consider estimation and inference with respect to the unknown parameter vector $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3)$, we note that analysts will, in most settings, face two forms of missing data: (i) incomplete observance of the event times due to censoring and (ii) the unobserved latent subject-specific frailties, Γ_i . In the next two sections, we outline our proposed strategy for accommodating both types of missingness, with the specific goal of ensuring estimation and inference that is robust to possible misspecification of the (unknown) frailty distribution.

3. Robust estimation and inference in the absence of censoring

3.1. The full data likelihood

In the absence of censoring, outcome information in the semi-competing risks setting can take on one of two forms: a subject is either observed to experience both the non-terminal and terminal events or solely experiencing the terminal event. To formalize this distinction, let $\mathbf{D}_1 = \{Y_1, \mathbb{I}(T_1 \leq T_2), T_2\}^T \equiv \{\min(T_1, T_2), \mathbb{I}(T_1 \leq T_2), T_2\}^T$. Based on this notation, the observed outcome vector for a subject that experiences both events will be $\mathbf{d}_1 = \{t_1, 1, t_2\}$; the observed outcome vector for a subject that experiences the terminal event prior to the non-terminal event will be $\mathbf{d}_1 = \{t_2, 0, t_2\}$. Finally, let $\mathbf{F} = (\mathbf{D}_1^T, \mathbf{X}^T, \Gamma)^T$ denote the ‘full’ data vector (i.e. one in which there is no censoring and the frailty is observable).

Given an i.i.d sample of size n , the full data likelihood is a product of contributions each of the form:

$$\mathcal{L}_1(\mathbf{f}_i|\boldsymbol{\theta}, \eta_{\Gamma}, \eta_{\mathbf{X}}) = \mathcal{L}_1(\mathbf{d}_{i1}|\mathbf{x}_i, \gamma_i; \boldsymbol{\theta}) \eta_{\Gamma}(\gamma_i|\mathbf{x}_i) \eta_{\mathbf{X}}(\mathbf{x}_i), \quad (5)$$

where η_{Γ} and $\eta_{\mathbf{X}}$ are the densities of Γ_i and $\mathbf{X}_i$ respectively, and

$$\begin{aligned} \mathcal{L}_1(\mathbf{d}_{i1}|\mathbf{x}_i, \gamma_i; \boldsymbol{\theta}) &= \lambda_{\theta_1}(y_{it}|\mathbf{x}_i, \gamma_i)^{I(t_{i1} \leq t_{i2})} \exp\{-\Lambda_{\theta_1}(y_{it}|\mathbf{x}_i, \gamma_i)\} \\ &\times \lambda_{\theta_2}(y_{it}|\mathbf{x}_i, \gamma_i)^{\{1-I(t_{i1} \leq t_{i2})\}} \exp\{-\Lambda_{\theta_2}(y_{it}|\mathbf{x}_i, \gamma_i)\} \\ &\times \lambda_{\theta_3}(t_{i2}|y_{i1}, \mathbf{x}_i, \gamma_i)^{I(t_{i1} \leq t_{i2})} \exp\{-\Lambda_{\theta_3}(y_{it}, t_{i2}|\mathbf{x}_i, \gamma_i)\} \end{aligned}$$

with

$$\Lambda_{\theta_1}(t|\mathbf{x}_i, \gamma_i) = \int_0^t \lambda_{\theta_1}(s|\mathbf{x}_i, \gamma_i) ds = \gamma_i \int_0^t \lambda_{0\theta_1}(s|\mathbf{x}_i) ds,$$

$$\Lambda_{\theta_2}(t|\mathbf{x}_i, \gamma_i) = \int_0^t \lambda_{\theta_2}(s|\mathbf{x}_i, \gamma_i) ds = \gamma_i \int_0^t \lambda_{0\theta_2}(s|\mathbf{x}_i) ds,$$

$$\Lambda_{\theta_3}(u, t|\mathbf{x}_i, \gamma_i) = \int_0^t \lambda_{\theta_3}(s|u, \mathbf{x}_i, \gamma_i) ds = \gamma_i \int_u^t \lambda_{0\theta_3}(s|\mathbf{x}_i) ds.$$

A detailed derivation of this likelihood, based on counting processes, is provided in Section A of the Supporting Information Appendix. Note that the likelihood construction based on the counting process does not require the assumption of latent T_1 when it was truncated by T_2 , which may not be the true data-generating mechanism.

3.2. The efficient score function

In practice, because the subject-specific frailties are not observable, one cannot proceed directly with likelihood-based estimation and inference based on (5). A standard way forward is to assume that the Γ_i is i.i.d according to some distribution with the Gamma distribution being a common choice because the induced marginal likelihood has a closed form expression. If this assumption is incorrect, however, the induced marginal likelihood is misspecified and the maximum likelihood estimator is no longer guaranteed to be consistent (White, 1982). Towards ensuring estimation and inference that is robust to possible misspecification of the frailty distribution, we extend an approach developed by Tsiatis & Ma (2004) for measurement error models. Specifically, when Γ is observed, consider the

Hilbert space $\mathcal{H}^F$ consisting of all $d_{\theta} \equiv \sum_{l=1}^3 d_{\alpha_l} + d_{\kappa_l} + p + 1$ -dimensional mean-zero, finite variance functions of the full data. For elements $\phi_1(\cdot)$ and $\phi_2(\cdot) \in \mathcal{H}^F$, define the inner product as $\langle \phi_1, \phi_2 \rangle = E\{\phi_1^T(\mathbf{F})\phi_2(\mathbf{F})\}$. The nuisance tangent space Λ^F is the subspace of $\mathcal{H}^F$, which is the mean square closure of all the elements in the form of $\mathbf{B}\mathbf{S}$, where $\mathbf{S}$ is an arbitrary nuisance score function and $\mathbf{B}$ a matrix with d_{θ} rows (Ma & Zhu, 2012).

When Γ is observed, the nuisance tangent space with respect to the distributions of $\mathbf{X}$ and Γ can be written as

$$\Lambda^F = [\phi(\mathbf{X}, \Gamma) \in \mathcal{H}^F, \text{ such that } E\{\phi(\mathbf{X}, \Gamma)\} = \mathbf{0}].$$

Note the corresponding full data score function

$$\mathbf{U}_{\theta}^F(\mathbf{D}_1; \theta) = \frac{\partial}{\partial \theta} \log \mathcal{L}_1\{\mathbf{D}_1 | \mathbf{X}, \Gamma; \theta\},$$

is orthogonal to Λ^F because of the factorization in (5). Now, because the frailty Γ is not observable, the nuisance tangent space for the observed data, denoted $\mathbf{O}=(\mathbf{D}_1^T, \mathbf{X}^T)^T$, is

$$\Lambda=E(\Lambda^F|\mathbf{O})=[E\{\phi(\mathbf{X}, \Gamma)|\mathbf{O}\}, \phi(\mathbf{X}, \Gamma) \in \Lambda^F, \text{ such that } E\{\phi(\mathbf{X}, \Gamma)\}=\mathbf{0}], \quad (6)$$

with the corresponding score function:

$$\mathbf{U}_\theta(\mathbf{O};\theta)=E\{\mathbf{U}_\theta^F(\mathbf{D}_1;\theta)|\mathbf{O}\}=\frac{\partial}{\partial\theta}\log\int\mathcal{L}_1\{\mathbf{D}_1|\mathbf{X}, \gamma;\theta\}\eta_\Gamma(\gamma|\mathbf{X})d\gamma. \quad (7)$$

Because the expectation in expression (7) conditions on the observed data, $\mathbf{U}_\theta(\mathbf{O}; \theta)$ is no longer orthogonal to Λ . One can, however, show that the efficient score function is an element of the space $\Lambda^\perp$ defined as

$$\Lambda^\perp \equiv [\phi(\mathbf{O}) \text{ such that } E\{\phi(\mathbf{O})|\Gamma, \mathbf{X}\}=\mathbf{0}, a.s.].$$

Noting that there is an unique orthogonal project of $\mathbf{U}_\theta(\mathbf{O}; \theta)$ onto Λ of the form of $E\{\mathbf{a}_\theta(\mathbf{X}, \Gamma) | \mathbf{O}\}$, the efficient score is given by $\mathbf{U}_\theta(\mathbf{O}; \theta) - E\{\mathbf{a}_\theta(\mathbf{X}, \Gamma) | \mathbf{O}\}$, where $\mathbf{a}_\theta$ can be found by solving the integral equation

$$E[\mathbf{U}_\theta(\mathbf{O};\theta) - E\{\mathbf{a}_\theta(\mathbf{X}, \Gamma)|\mathbf{O}\}|\mathbf{X}, \Gamma]=\mathbf{0}$$

under the correct specification of the unknown η_Γ .

Now, consider the, arguably common, situation where η is misspecified. Let η_Γ^* denote the misspecified density function (corresponding to the analysts choice of distribution for Γ) and $E^*(\cdot)$ the expectation operator under the incorrect model. Let $\mathcal{H}^*$ denote the Hilbert space induced by the misspecified distribution η_Γ^* and define Λ^{*F} and Λ^* analogously to Λ^F and Λ but replacing the expectation $E(\cdot)$ with $E^*(\cdot)$. The following theorem shows that the space orthogonal to the nuisance tangent space is invariant to the specification of η_Γ^* .

Theorem 1—For the Hilbert space $\mathcal{H}^*$ induced by a possibly incorrect distribution η_Γ^* , the space orthogonal to the nuisance tangent space $\Lambda^* \in \mathcal{H}^*$ is given by

$$\Lambda^{*\perp}=[\phi(\mathbf{O}) \text{ such that } E\{\phi(\mathbf{O})|\Gamma, \mathbf{X}\}=\mathbf{0}, a.s.],$$

as long as η_Γ^* has the same support as the true η_Γ .

The proof of this theorem follows the same arguments as those presented in Theorem 1 in Tsiatis & Ma (2004).

Theorem 1 suggests that there exists a function $\mathbf{a}_\theta(\mathbf{X}, \Gamma) \in \Lambda^F$ such that $E^*\{\mathbf{a}_\theta(\mathbf{X}, \Gamma) | \mathbf{O}\}$ is in Λ^* , and $\mathbf{U}_\theta^*(\mathbf{O}; \theta) - E^*\{\mathbf{a}_\theta(\mathbf{X}, \Gamma) | \mathbf{O}\}$ is orthogonal to Λ^* . That is, there exists a function $\mathbf{a}_\theta(\mathbf{X}, \Gamma) \in \Lambda^F$ such that

$$E[\mathbf{U}_\theta^*(\mathbf{O}; \theta) - E^*\{\mathbf{a}_\theta(\mathbf{X}, \Gamma) | \mathbf{O}\} | \mathbf{X}, \Gamma] = \mathbf{0},$$

where $\mathbf{U}_\theta^*(\mathbf{O}; \theta)$ is defined in the same way as $\mathbf{U}_\theta(\mathbf{O}; \theta)$ in (7) but replacing η_Γ by η_Γ^* . Note, $E^*\{\mathbf{a}_\theta(\mathbf{X}, \Gamma) | \mathbf{O}\}$ is uniquely defined for any given η_Γ^* because it is the orthogonal projection of $\mathbf{U}_\theta^*(\mathbf{O}; \theta)$ onto Λ^* . From this construction and Theorem 1,

$$\mathbf{S}_{\text{eff}}(\mathbf{O}; \theta) = \mathbf{U}_\theta^*(\mathbf{O}; \theta) - E^*\{\mathbf{a}_\theta(\mathbf{X}, \Gamma) | \mathbf{O}\} \quad (8)$$

has mean $\mathbf{0}$ even if η_Γ^* is misspecified. Furthermore, the solution to $\sum_{i=1}^n \mathbf{S}_{\text{eff}}(\mathbf{O}_i; \theta) = \mathbf{0}$ is the semi-parametric efficient estimator of θ under the true specification for η_Γ and when there is no censoring (Tsiatis & Ma, 2004).

3.3. Procedures for obtaining $\mathbf{a}_\theta$

As is clear from expression (8), in order to evaluate $\mathbf{S}_{\text{eff}}$, one must be able to compute the function $\mathbf{a}_\theta$. Here, we describe an algorithm for doing so when the possibly misspecified η_Γ^* is taken to be a Gamma distribution, specifically one for which the density is

$\eta_\Gamma^*(\gamma) = \sigma^{-2/\sigma^2} \gamma^{1/\sigma^2 - 1} \exp(-\gamma/\sigma^2) / \Gamma(1/\sigma^2)$, where $\sigma^2 > 0$ is considered to be known. We restrict attention to this distribution, in part, because it is by far not only the most common assumption made in practice but also because, as mentioned earlier, the induced marginal likelihood has a closed form expression.

Now, it is well-known that any function can be approximated by a step function arbitrarily well as the number of component intervals increases. With this in mind, we assume that for each $\mathbf{X}$, $\mathbf{a}_\theta(\mathbf{X}, \Gamma)$ can be approximated by

$$\mathbf{a}_\theta(\mathbf{X}, \Gamma = \gamma) = \sum_{j=1}^m \mathbf{r}_j I(\gamma_j \leq \gamma < \gamma_{j+1}), \quad (9)$$

where $(\gamma_1, \dots, \gamma_{m+1})$ is a partition of the support of Γ and $\mathbf{r}_j$ is the interval-specific value of an unspecified d_θ -dimensional function. In practice, we can find the intervals by first fitting a working model for Γ and then use the quantiles from the estimated frailty distribution as the partition endpoints. Under this specification,

$$E[E^*\{\mathbf{a}_\theta(\mathbf{X}, \Gamma) | \mathbf{O}\} | \mathbf{X}, \Gamma = \gamma_i] = \sum_{j=1}^m \int \mathbf{r}_j \{F_g(\gamma_{j+1} | \mathbf{O}) - F_g(\gamma_j | \mathbf{O})\} \mathcal{L}_1(d_1 | \mathbf{X}, \Gamma = \gamma_j; \theta) dd_1,$$

where F_g the posterior distribution of Γ given $\mathbf{O}$ under the possibly misspecified η_{Γ}^* . The latter can be shown to correspond to a Gamma

$[\Delta_1+1/\sigma^2+1, 1/\sigma^2+A_{\kappa_1}\{G_{\alpha_1}(Y_t)-\beta_1^T\mathbf{X}\}+A_{\kappa_2}\{G_{\alpha_2}(Y_t)-\beta_2^T\mathbf{X}\}+A_{\kappa_3}\{G_{\alpha_3}(T_2)-\beta_3^T\mathbf{X}\}-A_{\kappa_3}\{G_{\alpha_3}(Y_t)-\beta_3^T\mathbf{X}\}]$ distribution. Define the matrix $\mathbf{R}$ to be the $m \times m$ matrix with $[i, j]^{\text{th}}$ entry given by

$$R_{ij}(\mathbf{X}) = \int \{F_g(\gamma_{j+1}|\mathbf{O}) - F_g(\gamma_j|\mathbf{O})\} \mathcal{L}_1(\mathbf{d}_1|\mathbf{X}, \Gamma=\gamma_i; \boldsymbol{\theta}) d\mathbf{d}_1.$$

Furthermore, let $\mathbf{b}(\mathbf{X}) = [E^*\{\mathbf{U}_{\boldsymbol{\theta}}^*(\mathbf{O}; \boldsymbol{\theta})|\Gamma=\gamma_1, \mathbf{X}\}^T, \dots, E^*\{\mathbf{U}_{\boldsymbol{\theta}}^*(\mathbf{O}; \boldsymbol{\theta})|\Gamma=\gamma_m, \mathbf{X}\}^T]^T$, where

$$E^*\{\mathbf{U}_{\boldsymbol{\theta}}^*(\mathbf{O}; \boldsymbol{\theta})|\Gamma=\gamma_i, \mathbf{X}\} = \int \frac{\partial}{\partial \boldsymbol{\theta}} \{\log f \mathcal{L}_1(\mathbf{d}_1|\mathbf{X}, \gamma; \boldsymbol{\theta}) \eta_{\Gamma}^*(\gamma|\mathbf{X}) d\gamma\} \mathcal{L}_1(\mathbf{d}_1|\mathbf{X}, \Gamma=\gamma_i; \boldsymbol{\theta}) d\mathbf{d}_1.$$

One can then obtain the values of $\mathbf{r} = (r_1, \dots, r_m)^T$ by solving:

$$\mathbf{R}(\mathbf{X})\mathbf{r}^T(\mathbf{X}) = \mathbf{b}^T(\mathbf{X}).$$

Finally, after obtaining $\mathbf{r}_j$, the efficient score function $\mathbf{S}_{\text{eff}}$ can be approximated by

$$\frac{\partial}{\partial \boldsymbol{\theta}} \log \int \mathcal{L}_1(\mathbf{d}_1|\mathbf{X}, \gamma; \boldsymbol{\theta}) \eta_{\Gamma}^*(\gamma|\mathbf{X}) d\gamma - \sum_{j=1}^m \mathbf{r}_j \{F_g(\gamma_{j+1}|\mathbf{O}) - F_g(\gamma_j|\mathbf{O})\}.$$

4. Robust estimation and inference in the presence of censoring

4.1. The censored data likelihood

In the presence of (right) censoring, $\mathbf{D}_1$ will not be observed for some subjects. To accommodate this, let C denote the censoring time and $\mathbf{D}_2 = \{Y_c, \mathbb{I}(T_1 \leq C), C\}^T \equiv \{\min(T_1, C), \mathbb{I}(T_1 \leq C), C\}^T$. Based on this notation, the outcome vector for a subject that is censored prior to either of the events being observed would be $\mathbf{d}_2 = \{c, 0, c\}$; the outcome vector for a subject that is observed to experience a non-terminal event but is subsequently censored would have $\mathbf{d}_2 = \{t_1, 1, c\}$. For either of these subjects, the likelihood contribution is of the form:

$$\mathcal{L}_2(\mathbf{d}_{i2}, \mathbf{x}_i, \gamma_i | \boldsymbol{\theta}, \eta_{\Gamma}, \eta_C, \eta_{\mathbf{X}}) = \mathcal{L}_2(\mathbf{d}_{i2} | \mathbf{x}_i, \gamma_i; \boldsymbol{\theta}) \eta_{\Gamma}(\gamma_i | \mathbf{x}_i) \eta_C(c_i | \mathbf{x}_i) \eta_{\mathbf{X}}(\mathbf{x}_i), \quad (10)$$

where $\eta_C(c_i | \mathbf{x}_i)$ is the density for C and

$$\mathcal{L}_2(\mathbf{d}_{i2} | \mathbf{x}_i, \gamma_i; \boldsymbol{\theta}) = \lambda_{\theta_1}(y_{ic} | \mathbf{x}_i, \gamma_i)^{\mathbb{I}(t_{i1} \leq c_i)} \exp\{-\Lambda_{\theta_1}(y_{ic} | \mathbf{x}_i, \gamma_i)\} \times \exp\{-\Lambda_{\theta_2}(y_{ic} | \mathbf{x}_i, \gamma_i)\} \times \exp\{-\Lambda_{\theta_3}(y_{ic}, c_i | \mathbf{x}_i, \gamma_i)\}.$$

4.2. Estimation via non-parametric score imputation

Building on the derivation of the efficient score function for the non-censored observations in Section 3.2, we derive the score function for censored observations based on a score

imputation method. Towards this, let $Y_{j2} = \min(T_{j2}, C_j)$, $Y_{j1} = \min(T_{j1}, Y_{j2})$, $\Delta_{j1} = I(T_{j1} > Y_{j2})$ and $\Delta_{j2} = I(T_{j2} > C_j)$. Following Wang *et al.* (2012), the overarching strategy is to replace $S_{\text{eff}}(\mathbf{O}; \boldsymbol{\theta})$ by its conditional expectation giving the observed data. Specifically, for an individual that experienced a non-terminal event but was censored prior to observation of a terminal event (i.e. $\Delta_{j1}=1$ and $\Delta_{j2}=0$), the unobserved efficient score function is replaced by

$$Q_1(C, \mathbf{X}; \boldsymbol{\theta}) = E\{S_{\text{eff}}(\mathbf{O}; \boldsymbol{\theta}) | T_{j2} > C > T_{j1}, C, \mathbf{X}\}.$$

For an individual that is censored prior to observation of either event (i.e. $\Delta_{j1}=0$ and $\Delta_{j2}=0$), the unobserved efficient score function is replaced by

$$Q_2(C, \mathbf{X}; \boldsymbol{\theta}) = E\{S_{\text{eff}}(\mathbf{O}; \boldsymbol{\theta}) | \min(T_{j1}, T_{j2}) > C, C, \mathbf{X}\}.$$

Based on these expressions, the resulting censored data score function across all four observed data scenarios is

$$\psi(\boldsymbol{\theta}) = \Delta_2 S_{\text{eff}}(\mathbf{O}; \boldsymbol{\theta}) + \Delta_1 (1 - \Delta_2) E\{S_{\text{eff}}(\mathbf{O}; \boldsymbol{\theta}) | T_{j2} > C > T_{j1}, C, \mathbf{X}\} + (1 - \Delta_1)(1 - \Delta_2) E\{S_{\text{eff}}(\mathbf{O}; \boldsymbol{\theta}) | \min(T_{j2}, T_{j1}) > C, C, \mathbf{X}\}. \quad (11)$$

Note that because $\boldsymbol{\psi}(\boldsymbol{\theta})$ has mean $\mathbf{0}$ under the true model, estimation based on (11) is consistent. An important practical challenge, however, is that evaluation of the conditional expectation components of $\boldsymbol{\psi}(\boldsymbol{\theta})$ requires specification of the frailty distribution. To resolve this, we propose to use a kernel estimator for $E\{S_{\text{eff}}(\mathbf{O}; \boldsymbol{\theta}) | T_{j2} > C > T_{j1}, C, \mathbf{X}\}$ given by

$$\frac{\sum_{j=1}^n \Delta_{j2} I(T_{j2} > C > T_{j1}) S_{\text{eff}}(\mathbf{O}_j; \boldsymbol{\theta}) \widehat{F}_C(Y_{j2} | \mathbf{X}_j)^{-1} K_{\mathbf{h}}(\mathbf{X} - \mathbf{X}_j)}{\sum_{j=1}^n \Delta_{j2} I(T_{j2} > C > T_{j1}) \widehat{F}_C(Y_{j2} | \mathbf{X}_j)^{-1} K_{\mathbf{h}}(\mathbf{X} - \mathbf{X}_j)}, \quad (12)$$

where $\widehat{F}_C$ is a root- n consistent estimator for $1 - F_C$, the survival function for the censoring time, and $K_{\mathbf{h}}(\mathbf{X} - \mathbf{X}_j) = \prod_{k=1}^p K_{h_k}(X_k - X_{jk})$ is a multivariate kernel with a p dimensional bandwidth $\mathbf{h} = (h_k, k = 1, \dots, p)^T$. Note that we choose the product kernel for $\mathbf{X}$ because it permits the analyst to monitor the bandwidth for each covariate individually which, in turn, facilitates the theoretical derivations and numerical evaluations. For the discrete covariates, we can simply take $K_{h_k}(X_k - X_{jk}) = I(X_k = X_{jk})$, although in practice any sensible multivariate kernel could be used. Similarly, we estimate $E\{S_{\text{eff}}(\mathbf{O}; \boldsymbol{\theta}) | \min(T_{j1}, T_{j2}) > C, C, \mathbf{X}\}$ by

$$\frac{\sum_{j=1}^n \Delta_{j2} I\{\min(T_{j1}, T_{j2}) > C\} S_{\text{eff}}(\mathbf{O}_j; \boldsymbol{\theta}) \widehat{F}_C(Y_{j2} | \mathbf{X}_j)^{-1} K_{\mathbf{h}}(\mathbf{X} - \mathbf{X}_j)}{\sum_{j=1}^n \Delta_{j2} I\{\min(T_{j1}, T_{j2}) > C\} \widehat{F}_C(Y_{j2} | \mathbf{X}_j)^{-1} K_{\mathbf{h}}(\mathbf{X} - \mathbf{X}_j)}. \quad (13)$$

Finally, combining expressions (11), (12) and (13), the proposed estimator is obtained as the

solution to $\hat{\psi}(\boldsymbol{\theta}) = \sum_{i=1}^n \hat{\psi}_i(\boldsymbol{\theta}) = \mathbf{0}$, where

$$\begin{aligned} \hat{\psi}_i(\boldsymbol{\theta}) = & \Delta_{i2} \mathbf{S}_{\text{eff}}(\mathbf{O}_i; \boldsymbol{\theta}) + \Delta_{i1}(1 - \Delta_{i2}) \hat{E}\{\mathbf{S}_{\text{eff}}(\mathbf{O}_i; \boldsymbol{\theta}) | T_{i2} > C_i > T_{i1}, C_i, \mathbf{X}_i\} \\ & + (1 - \Delta_{i1})(1 - \Delta_{i2}) \hat{E}\{\mathbf{S}_{\text{eff}}(\mathbf{O}_i; \boldsymbol{\theta}) | \min(T_{i2}, T_{i1}) > C_i, C_i, \mathbf{X}_i\}. \end{aligned}$$

5. Asymptotic results

To establish asymptotic properties of the estimator proposed on Section 4, we first state the necessary conditions:

- (A1) $E\{\mathbf{S}_{\text{eff}}(\mathbf{O}; \boldsymbol{\theta})\} = \mathbf{0}$ has unique solution at $\boldsymbol{\theta} = \boldsymbol{\theta}_0$ and $\boldsymbol{\Psi}(\boldsymbol{\theta})$ has a continuous and bounded derivative.
- (A2) In the kernel function $K_{h_k}(x) = h_k K(h_k^{-1}x)$, h_k are bandwidths for covariate X_k , $k = 1, \dots, p$ and K is an ℓ^{th} -order symmetric kernel function that satisfies $\ell \geq 1$, $\int K(x)dx = 1$, $\int K(x)^2 dx < \infty$ and $\int x^2 K(x)dx < \infty$. The bandwidths satisfy $nh_k^2 \rightarrow \infty$ for each $k = 1, \dots, p$ and $n \sum_{k=1}^p h_k^4 \rightarrow 0$.
- (A3) $\mathbf{X}$ has finite second moment.
- (A4) The joint density function for $\mathbf{X}$ is bounded away from zero and infinity on its support.
- (A5) The censoring time is conditionally independent of (T_1, T_2) given $\mathbf{X}$, the effects of which do not depend on time.

For the most part, these conditions are standard. Specifically, condition (A1) is the usual condition for identifiability of model parameters while condition (A2) guarantees convergence of the kernel estimators. Conditions (A3) and (A4) ensure that the asymptotic variance of the estimator is bounded. The first part of condition (A5) is analogous to the usual non-informative censoring assumption in univariate time-to-event analyses. The second part of condition (A5) is an assumption that permits the use of a Cox proportional hazards model for the censoring time:

$$\bar{F}_C(t|\mathbf{X}) = \exp\{-\Lambda_{c0}(t) \exp(\boldsymbol{\beta}_{c0}^T \mathbf{X})\}.$$

Under this model, the partial likelihood estimator $\hat{\boldsymbol{\beta}}_c$, the Breslow baseline hazard estimator $\hat{\Lambda}_c(t, \hat{\boldsymbol{\beta}})$ and the corresponding censoring distribution function estimator

$$\hat{\bar{F}}_C(t|\mathbf{X}) = \exp\left\{-\hat{\Lambda}_c(t, \hat{\boldsymbol{\beta}}) \exp\left(\hat{\boldsymbol{\beta}}_c^T \mathbf{X}\right)\right\}$$

are all root n consistent to $\boldsymbol{\beta}_{c0}$, $\Lambda_{c0}(t)$ and $\bar{F}_C(t|\mathbf{X})$ (Lemma 1). The form of the censoring distribution covers the case when C is assumed to be independent of both $\mathbf{X}$ and $\mathbf{T}$, which is

used in Ma & Yin (2010) for an inverse probability weighting-based method. This independent censoring assumption would not affect the theoretic properties of the survival function estimators except that $\beta_{c0} = \mathbf{0}$ is assumed to be known in advanced. It is worthwhile to note that any consistent method can be used here to estimate the censoring distribution.

With the aforementioned conditions/assumptions, the following theorem describes the asymptotic properties of the proposed estimator:

Theorem 2

Let $\mathbf{W}_i = (\mathbf{O}_i^T, C_i)^T$ and $\hat{\boldsymbol{\theta}}$ the solution to

$$\hat{\boldsymbol{\psi}}(\boldsymbol{\theta}) = \sum_{i=1}^n \hat{\boldsymbol{\psi}}_i(\boldsymbol{\theta}) = \mathbf{0}.$$

Under conditions (A1)–(A5), $\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0 = o_p(1)$, where $\boldsymbol{\theta}_0$ is the true value of $\boldsymbol{\theta}$. Furthermore,

$$\begin{aligned} n^{1/2}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) = & - \left[E \left\{ \frac{\partial \boldsymbol{\psi}(\boldsymbol{\theta}_0)}{\partial \boldsymbol{\theta}^T} \right\} \right]^{-1} \left\{ n^{-1/2} \sum_{i=1}^n \mathbf{J}_1(\mathbf{W}_i; \boldsymbol{\theta}_0) \right. \\ & + n^{-1/2} \sum_{i=1}^n \int_0^\tau \mathbf{J}_2(s, \mathbf{W}_i; \boldsymbol{\theta}_0) dM_{ic}(s) \\ & \left. + n^{1/2} \int_0^\tau \mathbf{J}_3(s; \boldsymbol{\theta}_0) ds (\hat{\beta}_c - \beta_{c0}) \right\} + o_p(1) \end{aligned}$$

has an asymptotic mean-zero normal distribution with variance–covariance matrix, Σ , given by

$$\begin{aligned} & \left[E \left\{ \frac{\partial \boldsymbol{\psi}(\boldsymbol{\theta}_0)}{\partial \boldsymbol{\theta}^T} \right\} \right]^{-1} \left[E \{ \mathbf{J}_1(\mathbf{W}_i; \boldsymbol{\theta}_0)^{\otimes 2} \} + E \left\{ \int_0^\tau \mathbf{J}_2(s, \mathbf{W}_i; \boldsymbol{\theta}_0)^{\otimes 2} R_i(s) \lambda_{0c}(s) \exp(\boldsymbol{\beta}_{c0}^T \mathbf{X}_i) ds \right\} \right. \\ & \left. + \left\{ \int_0^\tau \mathbf{J}_3(s; \boldsymbol{\theta}_0) ds \right\} \Sigma_c^{-1} \left\{ \int_0^\tau \mathbf{J}_3(s; \boldsymbol{\theta}_0) ds \right\}^T \right] \left(\left[E \left\{ \frac{\partial \boldsymbol{\psi}(\boldsymbol{\theta}_0)}{\partial \boldsymbol{\theta}^T} \right\} \right]^{-1} \right)^T, \end{aligned}$$

where $\mathbf{J}_1, \mathbf{J}_2, \mathbf{J}_3$ are defined in Section D of the Supporting Information Appendix.

Inspection of the influence function given in Theorem 2 highlights that the uncertainty comes from three sources: uncertainty associated with the kernel estimator, uncertainty in estimate of the censoring distribution and uncertainty in the estimate $\hat{\beta}_c$. A proof of this result, along with those of two necessary lemmas, are presented in Appendices C and D. Because calculation of the empirical version of Σ is rather involved, in practice, one could use the bootstrap to obtain standard error estimates and/or confidence intervals (Wang *et al.*, 2012).

6. Simulation studies

6.1. Simulation set-up

To evaluate the small-sample operating characteristics of the proposed methods, we conducted a simulation study. As emphasized in the Introduction, the primary purpose of the patient-specific frailty terms is to account for dependence that is not taken into account by measured covariates included in the model. Our simulation studies are therefore structured to mimic this scenario. Specifically, we considered a data-generating mechanism based on model (3) setting $G(\cdot) = \log(\cdot)$, $l = 1, 2, 3$ and $A(\cdot) = \exp(\cdot/\kappa_l)$ so that the error terms follow a scaled extreme value distribution with $\kappa_l = 0.5$. For covariates, we considered the two-dimensional vector $\mathbf{X}_i = (1, X_{1i})^T$ and an additional latent variable Z_i associated indirectly with the survival time through X_{1i} and the frailty. More specifically, Z_i is drawn from a standard normal distribution. Given Z_i , a binary X_{1i} was generated as a random Bernoulli draw with probability $\Pr(X_{1i} = 1 | Z_i) = \Phi(Z_i)$, with $\Phi(\cdot)$ the cumulative distribution function for the standard normal distribution. Furthermore, a patient's frailty was generated from a mixture distribution depending on the value of another variable I_i with probability $\Pr(I_i = 1 | Z_i) = \Phi(Z_i)$; if $I_i = 1$, Γ_i was drawn from a Gamma distribution with mean 1 and variance 0.5, denoted here by $\text{Gamma}(1, 0.5)$; if $I_i = 0$, Γ_i was drawn from $2 + \text{Gamma}(1, 6)$, a location shifted Gamma distribution. By mixing the two Gamma distributions, the induced marginal distribution has a longer tail than either one of the component distributions; Fig. 2 provides a graphical representation. Note that, although I_i and X_{1i} have the same distribution, they are independent, given the Z_i , which suggests that Γ_i does not account for the dependence explained by $\mathbf{X}_i$. After obtaining multiple samples of $\mathbf{X}_i$, Γ_i , survival times based on model (3) were generated using methods described in Jiang & Haneuse (2014) with the 'true' values of $\boldsymbol{\theta}$, denoted $\boldsymbol{\theta}_0$, given in the first column of Table 1. Finally, independent right censoring was introduced by generating a censoring time based on an exponential distribution with the rate parameter chosen to achieve a certain censoring rate.

Based on the previous structure, we generated four sets of 1,000 simulated semi-competing risks data sets, which varied in sample size ($n = 100$ and 500) and overall censoring rates (0 per cent and 20 per cent). For each of the data sets, we applied the proposed semi-parametric approach described in Section 4. Throughout, to focus the simulations on the misspecification of the distribution of the frailties, we assumed the true transformation and cumulative hazard functions. To approximate the efficient score function, we choose $m = 25$ equally spaced grid points on the support of the frailty distribution (Section 4.2). Based on preliminary analyses not shown, we found this choice to achieve a reasonable balance between computational tractability and accuracy. Finally, for comparison, we also fit model (3) assuming the frailties arise from a common parametric distribution, specifically a Gamma distribution with mean 1 and variance σ^2 .

Finally, we note that the methods described in Sections 2–5 have been implemented in R (R Core Team, 2015). The code, together with the scripts for the simulation studies, are freely available via GitHub at <https://github.com/homebovine/semifrailty.git>.

6.2. Results

Results from the simulation study are provided in Table 1. Comparing the results between the proposed semi-parametric transformation (upper half) and the parametric analysis (lower half), we see that both analyses estimate the κ_j parameters (that index the distribution of the ε_{ij} terms) with low bias. In contrast, estimates of the regression parameters exhibit substantial bias under the parametric analysis and virtually none under the proposed approach; because both analyses assume the correct functional form for the transformations and the error terms, this difference is solely because of the differences in treatment of the frailty distribution. Comparing the empirical estimates of the standard deviations of the sampling distributions, we see that the estimates based on the proposed approach generally exhibit greater uncertainty. However, the empirical mean squared error indicates that the overall performance of the proposed estimator is substantially better. Finally, across the board, the general conclusions we draw regarding the differences between the proposed approach and a fully parametric analysis do not change as the sample sizes changes nor as the censoring rate changes.

7. Analysis of Medicare data

The methods proposed in this manuscript are applied to a motivating data set consisting of patients diagnosed with pancreatic cancer. Specifically, we consider $n = 621$ Medicare enrollees aged 87 years and older who received a diagnosis for pancreatic cancer during a hospitalization between 01/2005 and 11/2008 and were subsequently discharged (i.e. they did not die during the index hospitalization). For these analyses, the non-terminal event of interest is whether or not the patient was readmitted to the emergency room after having being discharged from the index hospitalization; the terminal event is death. Furthermore, observation time was administratively censored at 90 days because interest typically lies with readmission events soon after a patient is discharged (Lee *et al.*, 2015). Among the $n = 621$ patients in the study sample, 56 patients were observed to experience both events, 34 patients were observed to experience a readmission event but were then censored prior to death, 401 patients were observed to experience a death event prior to readmission and 130 patients were censored prior to experiencing either event.

7.1. Analyses

To illustrate the approach, we follow the simulations and consider a single covariate, specifically gender, where $X_{1j} = 0/1$ if the j^{th} patient is male/female. Four sets of analyses are considered. Specifically, we consider both the proposed semi-parametric approach as well as a fully parametric model (i.e. the patient-specific frailties are assumed to arise from a Gamma distribution). For both of these approaches, we consider the transformation functions to be (i) fixed as the log-transformation and (ii) taken from the Box–Cox family of transformations, specifically with $G_{\alpha,l}(\eta) = (\eta^l - 1)/\alpha_l$, $l = 1, 2, 3$. Throughout, as in the simulations, we set $A_{\kappa}(\xi) = \exp(\xi/\kappa)$ and approximate $\mathbf{a}_{\theta}(\mathbf{X}, \Gamma)$ based expression (9) with $m = 25$ equally spaced grid points. Finally, as is clear from the expression in Theorem 2, the functional form of asymptotic variance–covariance matrix is complex requiring the empirical evaluation of numerous integrals. We, therefore, report bootstrap-based standard error estimates and 95% confidence intervals, based on 1,000 replicates.

7.2. Results

Results from the four sets of analyses are shown in Tables 2 and 3 and in Fig. 3, and Figs B. 1 and B.2 in the Supporting Information Appendix. Focusing on the regression parameter for gender (i.e. β_{11}), all analyses indicate that females diagnosed with pancreatic cancer have lower risk for readmission and death, both prior to and post readmission, than males. In both tables, the point estimates for the impact of gender on risk of readmission are different between the proposed semi-parametric and parametric analyses; in Table 2, the estimates are $\hat{\beta}_{11} = -0.35$ and $\hat{\beta}_{11} = -0.45$ under the two analyses, respectively; in Table 3 the difference is even greater. To aid in the interpretation of these differences, Fig. 3 provides estimates of the conditional survival function (i.e. setting the patient-specific frailty to $\Gamma = 1$) for readmission stratified by gender, from the four analyses; estimates for the corresponding survival functions for death are provided in Section B in the Supporting Information Appendix.

Comparing panel (a) to panel (b) and panel (c) to panel (d), we again see a clear indication of differences in the results between the proposed semi-parametric analysis and one that is fully parametric. Collectively, these results indicate that there is some sensitivity to the specification of the frailty distribution, with the proposed semi-parametric analysis likely providing more reliable results. Comparing the results for the semi-parametric analysis between Tables 2 and 3, we see that while the point estimates for the effect of gender are similar, the standard error estimates are substantially smaller under the model in which a Box–Cox power transformation was used. From Table 3, the results for the power transformation parameters provide a clear indication of a better fit of this model than one in which the transformation is taken as the log function. Further evidence of this is given by the differences in the results for σ^2 between the two tables; under the log-transformation, there is substantial residual between-patient variation that is absorbed by the frailties ($\hat{\sigma}^2=1.23$; 95% CI 0.89, 1.71); under the Cox–Box specification, there is virtually no residual between-patient variation ($\hat{\sigma}^2=0.04$; 95% CI 2.87e-06, 0.72).

8. Discussion

In this paper, we have proposed a novel semi-parametric transformation model for the analysis of semi-competing risks data. The primary goal of this work was to develop a framework within which the assumption of a Gamma distribution for patient-specific frailties, standard in multi-state models for semi-competing risks data, can be relaxed. To our knowledge, this is the first paper in the semi-competing risks literature that attempts to do so. While the proposed methodology requires the analyst to adopt a *working frailty distribution*, the estimation procedures developed in Sections 3 and 4 yield a consistent estimator that is robust to the specific choice. This robustness comes at a price; however, to ensure identifiability, we restrict attention to specifications for the three hazard functions in the illness–death model that are fully parametric. To mitigate this, we propose a new class of flexible transformation models within which the transformation functions and error distributions can be made as simple or complex as analysts see fit. To accommodate right censoring, we use non-parametric score imputation and show that the resulting estimator is $\sqrt{n}$ -consistent and asymptotically normally distributed.

The complexity of semi-competing risks data, coupled with finite sample sizes, typically requires analysts to balance structure and flexibility in their model specifications. Towards this, some parametric structure is invariably imposed either in the way in which covariates influence the underlying risk (i.e. regression parameters) or in terms of distributional assumptions. In principle, analysts can vary these assumptions and examine sensitivity in the results. This will be particularly important in ‘data poor’ settings where either the event rates are low either due to the overall sample size or if the censoring rate is high. The modelling and estimation framework we propose will be useful in this context, giving analysts a complementary analysis approach in which a standard parametric assumption can be relaxed. While the proposed transformation models are flexible, the proposed estimation procedures can also be used if analysts are interested in restricting to Cox-type or AFT model in order to facilitate interpretation of regression parameters. As with all transformation models, if the $G_{\alpha}(\cdot)$ functions are permitted to be more flexible, one can compare risk across sub-populations graphically as we do in Section 7.

Although this paper is motivated by the analysis of semi-competing risks data, the semiparametric transformation model proposed in Section 2 and the estimation procedures derived in Sections 3 and 4 are much more broadly applicable. In particular, the methods could be applied to any analysis of multivariate survival data in which a shared frailty term is used to account for correlation. These include, but are not limited to, the modelling of natural disease history (page 271–300 in Aalen *et al.* (2008) & Yen *et al.* (2010)) and multi-state models for competing risks data (Gorfine & Hsu, 2011).

Finally, a number of extensions of the proposed methods are possible, which should give analysts even greater flexibility. Specifically, we are investigating relaxing the fully parametric specification for the hazard functions by considering estimation/inference for a Cox model for which the baseline hazard is left unspecified. This requires characterizing identifiability when the overall model has two non-parametric components, specifically the frailty distribution and the baseline hazard functions; this is an active area of our own research. Related to this is that while we adopt a Markov assumption for the transition from the non-terminal event to the terminal event (i.e. the hazard $\lambda_{\theta_3}(t_2 | t_1)$ in expression (4)), analysts may want to consider other formulations such as the semi-Markov assumption in which one models the so-called sojourn time between the two events (e.g. Xu *et al.* (2010) & Lee *et al.* (2015)). Finally, an additional important extension will be to permit the observed data to be subject to left truncation and/or interval censoring, which will broaden the scope of applications to which the methods can be applied.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

Aalen, O., Borgan, O., Gjessing, H. Survival and event history analysis: A process point of view. New York: Springer Science & Business Media; 2008.

- Agresti A, Caffo B, Ohman-Strickland P. Examples in which misspecification of a random effects distribution reduces efficiency, and possible remedies. *Comput. Statist. Data Anal.* 2004; 47(3):639–653.
- Cui S, Sun Y. Checking for the gamma frailty distribution under the marginal proportional hazards frailty model. *Statist. Sinica.* 2004; 14(1):249–267.
- Fine JP, Jiang H, Chappell R. On semi-competing risks data. *Biometrika.* 2001; 88(4):907–919.
- Ghosh D. Semiparametric inferences for association with semi-competing risks data. *Statist. Med.* 2006; 25(12):2059–2070.
- Glidden DV. Checking the adequacy of the gamma frailty model for multivariate failure times. *Biometrika.* 1999; 86(2):381–393.
- Gorfine M, Hsu L. Frailty-based competing risks model for multivariate survival data. *Biometrics.* 2011; 67(2):415–426. [PubMed: 20707868]
- Heagerty PJ, Kurland BF. Misspecified maximum likelihood estimates and generalised linear mixed models. *Biometrika.* 2001; 88(4):973–985.
- Horowitz, JL. Semiparametric and nonparametric methods in econometrics. Vol. 692. New York: Springer; 2010.
- Hougaard P. Life table methods for heterogeneous populations: Distributions describing the heterogeneity. *Biometrika.* 1984; 71(1):75–83.
- Hsieh JJ, Wang W, Adam Ding A. Regression analysis based on semicompeting risks data. *J. R. Statist. Soc. B.* 2008; 70(1):3–20.
- Jiang F, Haneuse S. Simulation of semicompeting risk survival data and estimation based on multistate frailty model. Harvard University Biostatistics Working Paper Series. 2014
- Jiang H, Fine JP, Kosorok MR, Chappell R. Pseudo self-consistent estimation of a copula model with informative censoring. *Scand. J. Statist.* 2005; 32(1):1–20.
- Kosorok MR, Lee BL, Fine JP. Robust inference for univariate proportional hazards frailty regression models. *Ann. Statist.* 2004; 32(4):1448–1491.
- Lakhal L, Rivest LP, Abdous B. Estimating survival and association in a semicompeting risks model. *Biometrics.* 2008; 64(1):180–188. [PubMed: 17645782]
- Lee KH, Haneuse S, Schrag D, Dominici F. Bayesian semiparametric analysis of semicompeting risks data: Investigating hospital readmission after a pancreatic cancer diagnosis. *Journal of the Royal Statistical Society: Series C (Applied Statistics).* 2015; 64(2):253–273.
- Litière S, Alonso A, Molenberghs G. Type I and type II error under random-effects misspecification in generalized linear mixed models. *Biometrics.* 2007; 63(4):1038–1044. [PubMed: 17425642]
- Liu L, Wolfe RA, Huang X. Shared frailty models for recurrent events and a terminal event. *Biometrics.* 2004; 60(3):747–756. [PubMed: 15339298]
- Ma Y, Yin G. Semiparametric median residual life model and inference. *Canad. J. Statist.* 2010; 38(4):665–679.
- Ma Y, Zhu L. A semiparametric approach to dimension reduction. *J. Amer. Statist. Assoc.* 2012; 107(497):168–179.
- McGilchrist CA, Aisbett CW. Regression with frailty in survival analysis. *Biometrics.* 1991; 47(1):461–466. [PubMed: 1912255]
- Neuhaus JM, Hauck WW, Kalbfleisch JD. The effects of mixture distribution misspecification when fitting mixed-effects logistic models. *Biometrika.* 1992; 79(4):755–762.
- Neuhaus JM, McCulloch CE, Boylan R. A note on type II error under random effects misspecification in generalized linear mixed models. *Biometrics.* 2011; 67(2):654–656. [PubMed: 21689077]
- Nielsen GG, Gill RD, Andersen PK, Sørensen TI. A counting process approach to maximum likelihood estimation in frailty models. *Scand. J. Statist.* 1992; 19(1):25–43.
- Pan W. Using frailties in the accelerated failure time model. *Lifetime Data Anal.* 2001; 7(1):55–64. [PubMed: 11280848]
- Peng L, Fine JP. Regression modeling of semicompeting risks data. *Biometrics.* 2007; 63(1):96–108. [PubMed: 17447934]
- R Core Team. R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing; 2015.

- Shih JH, Louis TA. Assessing gamma frailty models for clustered failure time data. *Lifetime Data Anal.* 1995; 1(2):205–220. [PubMed: 9385102]
- Su YR, Wang JL. Modeling left-truncated and right-censored survival data with longitudinal covariates. *Ann. Statist.* 2012; 40(3):1465–1488.
- Tsiatis AA, Ma Y. Locally efficient semiparametric estimators for functional measurement error models. *Biometrika.* 2004; 91(4):835–848.
- Vest JR, Gamm LD, Oxford BA, Gonzalez MI, Slawson KM. Determinants of preventable readmissions in the United States: A systematic review. *Implement. Sci.* 2010; 5(1):88. [PubMed: 21083908]
- Wang W. Estimating the association parameter for copula models under dependent censoring. *J. R. Stat. Soc. Ser. B.* 2003; 65(1):257–273.
- Wang Y, Garcia TP, Ma Y. Nonparametric estimation for censored mixture data with application to the Cooperative Huntington's Observational Research Trial. *J. Amer. Statist. Assoc.* 2012; 107(500): 1324–1338.
- Warren JL, Barbera L, Bremner KE, Yabroff KR, Hoch JS, Barrett MJ, Luo J, Krahn MD. End-of-life care for lung cancer patients in the United States and Ontario. *J. Natl. Cancer Inst.* 2011; 103(11): 853–862. [PubMed: 21593012]
- White H. Maximum likelihood estimation of misspecified models. *Econometrica.* 1982; 50(1):1–25.
- Xu J, Kalbfleisch JD, Tai B. Statistical analysis of illness–death processes and semicompeting risks data. *Biometrics.* 2010; 66(3):716–725. [PubMed: 19912171]
- Ye Y, Kalbfleisch JD, Schaubel DE. Semiparametric analysis of correlated recurrent and terminal events. *Biometrics.* 2007; 63(1):78–87. [PubMed: 17447932]
- Yen AM, Chen TH, Duffy SW, Chen CD. Incorporating frailty in a multi-state model: Application to disease natural history modelling of adenoma-carcinoma in the large bowel. *Statist. Methods in Med.* 2010; 19(5):529–546.
- Zeng D, Chen Q, Chen MH, Ibrahim JG. Amgen Research Group. Estimating treatment effects with treatment switching via semicompeting risks models: An application to a colorectal cancer study. *Biometrika.* 2012; 99(1):167–184. [PubMed: 23049136]
- Zeng D, Chen Q, Ibrahim JG. Gamma frailty transformation models for multivariate survival times. *Biometrika.* 2009; 96(2):277–291. [PubMed: 24948838]
- Zeng D, Lin D. Maximum likelihood estimation in semiparametric regression models with censored data. *J. R. Statist. Soc. B.* 2007; 69(4):507–564.
- Zeng D, Lin D. Semiparametric transformation models with random effects for joint analysis of recurrent and terminal events. *Biometrics.* 2009; 65(3):746–752. [PubMed: 18945267]

Fig. 1.
Multi-state model representation of the semi-competing risks data setting.

Fig. 2. Graphical representation of the mixture Gamma frailty distribution used in the simulations of Section 6.

Fig. 3.

Estimated conditional survival functions for readmission from four analyses that vary in terms of the transformation functions, $G(\cdot)$, and treatment of the patient-specific frailties. In each panel, black lines correspond to males and grey lines to females. Dashed lines indicate pointwise bootstrap-based 95% confidence bands. Throughout, the patient-specific frailty was set to $\Gamma = 1$.

Table 1

Comparison of operating characteristics, based on 1,000 simulated samples, between the proposed semi-parametric transformation model for the analysis of semi-competing risks data to a parametric analysis in which the frailty random effects are assumed to arise from a Gamma distribution. Values are as follows: θ_0 is the true value of the parameter; Bias($\hat{\theta}$) is the absolute bias for θ ; sd($\hat{\theta}$) is empirical standard deviation; MSE is the empirical mean square error of the estimators

Proposed semi-parametric transformation model												
Censoring rate = 0%												
n 100 θ_0 Bias($\hat{\theta}$) sd($\hat{\theta}$) MSE θ κ_1 κ_2 κ_3 β_{10} β_{11} β_{20} β_{21} β_{30} β_{31} n 500 $\hat{\theta}$ Bias($\hat{\theta}$) sd($\hat{\theta}$) MSE n 100 $\hat{\theta}$ Bias($\hat{\theta}$) sd($\hat{\theta}$) MSE   Bias($\hat{\theta}</$												

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

β_{11}	-1.1	0.21	0.167	0.072	0.21	0.071	0.048	0.21	0.174	0.074	0.21	0.073	0.050
β_{20}	-0.2	0.42	0.177	0.210	0.42	0.077	0.182	0.41	0.195	0.206	0.42	0.083	0.183
β_{21}	-1.1	0.19	0.189	0.073	0.21	0.091	0.052	0.20	0.197	0.077	0.21	0.092	0.053
β_{30}	-0.2	0.42	0.142	0.196	0.42	0.063	0.180	0.42	0.150	0.200	0.42	0.066	0.181
β_{31}	-1.1	0.21	0.174	0.075	0.20	0.075	0.045	0.21	0.175	0.074	0.20	0.084	0.048

Table 2

Results from the analysis of the Medicare data based on the proposed framework with $G_{\text{ol}}(t)$ fixed to be the log-transformation, $l = 1, 2, 3$. The estimators are the average over 1,000 bootstrap replicates. Standard errors were estimated using the bootstrap, based on 1,000 replicates. 95% confidence intervals are based on the quantiles of the bootstrap estimate of the sampling distribution

Parameter, θ^a	Proposed semi-parametric model			Fully parametric model		
	$\hat{\theta}$	sd($\hat{\theta}$)	95% CI	$\hat{\theta}$	sd($\hat{\theta}$)	95% CI
Readmission						
κ_1	1.14	0.511	(0.60, 2.92)	0.85	0.086	(0.72, 1.02)
β_{10}	1.41	0.177	(1.08, 2.11)	1.19	0.222	(0.87, 1.63)
β_{11}	-0.35	0.296	(-0.91, 0.38)	-0.45	0.236	(-0.81, 0.09)
Death prior to readmission						
κ_2	0.80	0.146	(0.56, 1.37)	0.84	0.049	(0.76, 0.93)
β_{20}	-0.31	0.232	(-0.68, 0.09)	-0.14	0.121	(-0.33, 0.10)
β_{21}	-0.28	0.186	(-0.71, 0.21)	-0.30	0.168	(-0.59, 0.01)
Death after readmission						
κ_3	0.61	0.176	(0.39, 1.52)	0.91	0.154	(0.66, 1.32)
β_{30}	0.59	0.196	(-0.19, 0.91)	0.69	0.216	(0.31, 1.07)
β_{31}	-1.09	0.421	(-2.36, -0.42)	-0.95	0.465	(-1.69, -0.08)
Frailty variance						
σ^2	-	-	-	1.23	0.223	(0.89, 1.71)

^a κ_l indexes the scaled exponential distribution for the ξ_{il} error terms; (β_l, β_{ll}) are the regression parameters, $l = 1, 2, 3$; σ^2 is variance component indexing the Gamma distribution in the fully parametric analysis.

Table 3

Results from the analysis of the Medicare data based on the proposed framework with $G_{\alpha_l}(t) = (t^{\alpha_l} - 1)/\alpha_l$, the Box–Cox transformation, $l=1,2,3$. The estimators are the average over 1,000 bootstrap replicates. Standard errors were estimated using the bootstrap, based on 1,000 replicates. 95% confidence intervals are based on the quantiles of the bootstrap estimate of the sampling distribution

Parameter, θ^a	Proposed semi-parametric model			Fully parametric model		
	$\hat{\theta}$	sd($\hat{\theta}$)	95% CI	$\hat{\theta}$	sd($\hat{\theta}$)	95% CI
Readmission						
κ_1	1.93	0.631	(1.09, 3.31)	1.32	0.178	(1.01, 1.77)
β_{10}	2.79	0.137	(2.51, 3.18)	2.58	0.383	(1.79, 3.43)
β_{11}	-0.38	0.173	(-0.93, -0.01)	-0.55	0.352	(-1.09, 0.26)
α_1	-0.35	0.126	(-0.64, -0.15)	-0.16	0.071	(-0.29, -0.02)
Death prior to readmission						
κ_2	1.40	0.149	(1.15, 1.76)	1.40	0.118	(1.12, 1.59)
β_{20}	0.11	0.204	(-0.45, 0.49)	0.47	0.121	(0.14, 0.69)
β_{21}	-0.30	0.177	(-0.70, 0.15)	-0.34	0.190	(-0.64, 0.06)
α_2	-0.37	0.044	(-0.46, -0.27)	-0.28	0.044	(-0.34, -0.16)
Death after readmission						
κ_3	1.75	0.094	(1.34, 1.91)	1.73	0.317	(1.18, 3.36)
β_{30}	0.28	0.167	(-0.23, 0.55)	0.26	0.399	(-0.51, 1.02)
β_{31}	-1.23	0.156	(-1.74, -0.93)	-1.08	0.728	(-2.90, 0.30)
α_3	0.20	0.089	(0.08, 0.52)	0.33	0.264	(-0.40, 0.74)
Frailty variance component						
σ^2	-	-	-	0.04	0.057	(2.87e-06, 0.72)

^a κ_l /indexes the scaled exponential distribution for the ξ_l /error terms; (β_0, β_l) are the regression parameters, $l=1,2,3$; α_l is the Box–Cox power parameter, $l=1,2,3$; σ^2 is variance component indexing the Gamma distribution in the fully parametric analysis.