Title

Heterogeneous individual risk modelling of recurrent events

Journal

Biometrika, 108(1)

ISSN

0006-3444

Authors

Ma, Huijuan

Peng, Limin

Huang, Chiung-Yu

et al.

Publication Date

2021-03-02

DOI

10.1093/biomet/asaa053

Peer reviewed

HHS Public Access

Author manuscript

Biometrika. Author manuscript; available in PMC 2021 April 08.

Published in final edited form as:

Biometrika. 2021 March ; 108(1): 183–198. doi:10.1093/biomet/asaa053.

Heterogeneous Individual Risk Modeling of Recurrent Events

Huijuan Ma,

Academy of Statistics and Interdisciplinary Sciences, East China Normal University, Shanghai 200062, China

Limin Peng,

Department of Biostatistics and Bioinformatics, Emory University, Atlanta 30322, U.S.A.

Chiung-Yu Huang,

Department of Epidemiology and Biostatistics, University of California at San Francisco, San Francisco 94158, U.S.A.

Haoda Fu

Eli Lilly and Company, Indianapolis 46285, U.S.A.

Summary

Progression of chronic disease is often manifested by repeated occurrences of disease-related events over time. Delineating the heterogeneity in the risk of such recurrent events can provide valuable scientific insight for guiding customized disease management. In this paper, we propose a new sensible measure of individual risk of recurrent events and present a dynamic modeling framework thereof, which accounts for both observed covariates and unobservable frailty. The proposed modeling requires no distributional specification of the unobservable frailty, while permitting the exploration of dynamic effects of the observed covariates. We develop estimation and inference procedures for the proposed model through a novel adaptation of the principle of conditional score. The asymptotic properties of the proposed estimator, including the uniform consistency and weak convergence, are established. Extensive simulation studies demonstrate satisfactory finite-sample performance of the proposed method. We illustrate the practical utility of the new method via an application to a diabetes clinical trial that explores the risk patterns of hypoglycemia in Type 2 diabetes patients.

Keywords

Conditional score; Frailty; Quantile regression; Recurrent event

1. Introduction

In chronic disease follow-up studies, clinically important outcomes, such as opportunistic infections in HiV patients, repeated exacerbations in chronic obstructive pulmonary disease

hjma@fem.ecnu.edu.cn.

7. Supplementary Material

Supplementary material available at *Biometrika* online includes the technical proofs of the theorems and additional numerical results from simulation studies and real data analysis.

patients, and episodes of hypoglycemia in diabetic individuals, are often captured over time to track disease progression. understanding how the risk of such recurrent events varies across individuals pertaining to baseline characteristics can provide useful information for guiding customized disease management.

To characterize the individual risk of recurrent events, a popular approach is to model the counting process of recurrent events, $N_i^*(t)$, which depicts the trajectory of event occurrences in subject i. A well-known example is the classic proportional intensity model (Prentice et al., 1981; Andersen & Gill, 1982), where $N_i^*(t)$'s share a common baseline intensity function and baseline covariates $\widetilde{X}_i$ are assumed to have multiplicative effects on the intensity. To capture the within-subject dependency of recurrent events not explained by the observed covariates, an unobservable frailty ξ_i (Nielsen et al., 1992; Oakes, 1992) is often incorporated into the modeling of $N_i^*(t)$. For instance, given the subject-level traits including the observed covariates $\widetilde{X}_i$ and an unobservable frailty ξ_i, $N_i^*(t)$, may be specified as a nonstationary Poisson process with the intensity function $\xi_i \exp\left(\widetilde{X}_i^\tau b_0\right)\lambda_0(t)$, where $\lambda_0(t)$ is an unspecified baseline intensity function (Wang et al., 2001). Such a model sensibly accounts for both the observed covariates and the unobservable frailty for explicating the individual differences in recurrent event occurrences. However, as elaborated later, it is subject to a subtle but important limitation, which is, all covariates are assumed to have homogeneous effects across the whole study population. In practice, this assumption can be unreasonable, for example, when the efficacy of a testing drug varies across different risk groups. Such a limitation is commonly present in most existing models of recurrent event counting process (Pepe & Cai, 1993; Lawless & Nadeau, 1995; Lin et al., 2000; Wang et al., 2001; Schaubel et al., 2006, for example), posing an obstacle for probing some important heterogeneity in individual risk of recurrent events.

In this work, we develop a new modeling framework that is flexible and robust for delineating the heterogeneity in recurrent event risk. Taking a novel view of the multiplicative intensity model, we assume that $N_i^*(t)$, given a nonnegative random variable γ_i, is a nonstationary Poisson process with the intensity function,

$$\lambda(t|\gamma_i) = \gamma_i \cdot \lambda_0(t), \tag{1}$$

where the baseline intensity $\lambda_0(t)$ is an unknown, nonnegative, and continuous function. Here γ_i captures the scale shift of subject i's intensity process from the baseline intensity, and thus can serve as a sensible measure of subject-specific risk of recurrent events. Unlike in typical recurrent event frailty models, which usually address the influences of the observed covariates $\widetilde{X}_i$ and the unobservable frailty ξ_i separately, we shall utilize γ_i to simultaneously account for the individual variations in recurrent event occurrences explained by either $\widetilde{X}_i$ or ξ_i. The proportional intensity model (Andersen & Gill, 1982) is a special case of model (1) with $\gamma_i = \exp\left(\widetilde{X}_i^\tau b_0\right)$. When $\gamma_i = \xi_i \exp\left(\widetilde{X}_i^\tau b_0\right)$, model (1) reduces to Wang et al. (2001)'s semiparametric multiplicative intensity model. An important observation for these special cases is that a linear model is essentially exerted for $\log(\gamma_i)$; consequently, all

covariate effects are confined to be location-shifts. For example, Wang et al. (2001)'s model is equivalent to specifying $\log(\gamma_i) = \log\xi_i + \widetilde{X}_i^{\tau} b_0$. In a two-sample case comparing a new treatment versus placebo, the treatment effect is assumed to be a constant shift in $\log(\gamma_i)$ throughout its whole range, meaning a homogeneous treatment effect across different risk groups. This assumption may not be scientifically plausible, and automatically precludes inferences leading to more customized intervention strategies.

To enable exploring a potentially heterogeneous relationship between $\widetilde{X}_i$ and the latent individual risk measure γ_i, we propose to link $\log(\gamma_i)$ with $\widetilde{X}_i$ through a linear quantile regression model (Koenker & Bassett, 1978). The modeling strategy of quantile regression inherently permits dynamic effects of covariates that change across the range of γ_i. This effort leads to a broader class of multiplicative intensity models because existing ones usually assume a linear model for $\log(\gamma_i)$ over $\widetilde{X}_i$, which is a special case of a linear quantile regression model. Given this connection, inferential tools developed for the new general model can be used to test the goodness-of-fit of the existing multiplicative intensity models. Furthermore, as explained in Section 2, a linear quantile regression model can naturally account for sources of heterogeneity in γ_i that are not captured by the location-shift effects of the observed covariates; thus it is a more flexible venue to accommodate the unobservable frailty without explicitly specifying its distribution as well as the form of its interaction with $\widetilde{X}_i$.

To tackle the proposed quantile regression of γ_i, a main technical difficulty relates to the latent unobservable nature of γ_i. We address this challenge by novelly adapting the principle of conditional score (Stefanski & Carroll, 1987) in the settings of quantile regression and multiplicative intensity regression of recurrent events. Revising the conditional likelihood arguments in Wang et al. (2001), we can circumvent any parametric specification of the baseline intensity function $\lambda_0(t)$ in model (1). In addition, adopting the quantile regression modeling of γ_i allows us to avoid imposing distributional assumptions for any unobservable frailty. These nice properties indicate the flexibility and robustness of the proposed method for exploring the heterogeneity in individual recurrent event risk.

2. The Proposed Model

We begin with an introduction of the data and notation. Let $T_i^{(j)}$ denote the time to the jth recurrent event of subject i. The underlying counting process for the recurrent events is defined as $N_i^*(t) = \sum_{j=1}^{\infty} I\left(T_i^{(j)} \le t\right)$. Suppose the observation of recurrent events is terminated by a censoring time C_i. The observed counting process is given by $N_i(t) = N_i^*(t \wedge C_i) = \sum_{j=1}^{\infty} I\left(T_i^{(j)} \le t \wedge C_i\right)$, where $a \wedge b$ denotes the minimum of a and b. Let m_i be the total number of observed recurrent events of subject i, i.e. $m_i = N(C_i) = N^*(C_i)$. Define $X_i = \left(1, \widetilde{X}_i^{\tau}\right)^{\tau}$, where $\widetilde{X}_i$ is a $(p-1) \times 1$ vector capturing baseline covariates. The observed data include $\{N_i(t), C_i, X_i\}_{i=1}^n$. Hereafter the same notation without subscript i represent the corresponding population analogues.

We assume a general multiplicative intensity model, model (1): given γ_i, which measures the individual risk of recurrent events for subject i, $N_i^*(t)$ is a nonstationary Poisson process with the intensity function, $\lambda(t|\gamma_i) = \gamma_i \cdot \lambda_0(t)$. We do not specify any parametric form for $\lambda_0(t)$; nonetheless, we impose a constraint,

$$\int_0^{\nu^*} \lambda_0(t)dt = 1, \tag{2}$$

for the purpose of model identifiability, where ν^* is a predetermined constant which can be chosen as the upper bound of C_i's support. It is clear that, without constraint (2), model (1) is not identifiable from an alternative specification of the intensity function, $(\gamma_i/c) \cdot \{c\lambda_0(t)\}$, where c is a positive constant. Thus, constraint (2) is an integral part of the assumed multiplicative intensity model (1). of note, model (1) implies

$$E\{N_i^*(t)|\gamma_i\} = \gamma_i \cdot \mu_0(t),$$

where $\mu_0(t) = \int_0^t \lambda_0(s)ds$, and $\mu_0(\nu^*) = 1$. This suggests an alternative interpretation of γ_i, which is the subject-specific scale shift in the mean function.

The core component of the proposed modeling is to use quantile regression to explore the heterogeneity in subject-specific risk of recurrent events, quantified by γ_i. Let γ denote the population analogue of γ_i and let $Q_\gamma(\tau|X_i)$ denote the τth conditional quantile of γ given $X = X_i$, namely, $Q_\gamma(\tau|X_i) = \inf\{u \geq 0 : \Pr(\gamma \leq u|X = X_i) \geq \tau\}$. We assume that

$$Q_\gamma(\tau|X_i) = \exp\{X_i^\tau \beta_0(\tau)\} \doteq \exp\{a_0(\tau) + \widetilde{X}_i^\tau b_0(\tau)\}, \tag{3}$$

where $\beta_0(\cdot) \doteq (a_0(\tau), b_0(\tau)^\tau)^\tau$ is a $p \times 1$ vector of unknown regression coefficient functions.

Our modeling strategy shares a similar spirit with the multilevel (or hierarchy) modeling in linear regression settings (Raudenbush & Bryk, 2002). That is, we first utilize γ_i to quantify the risk of recurrent events for subject i based on the individual multiplicative intensity model (1), and then probe the population-level heterogeneity in γ_i. The non-intercept coefficients in $b_0(\tau)$ are population quantities of key interest, representing the change in the τ-th quantile of $\log(\gamma)$ given one unit change in the corresponding covariate. For example, in a two-sample case where $\widetilde{X}_i$ indicates a new treatment versus placebo, $b_0(\tau)$ represents the treatment effect on the τ-th quantile of γ. Examining $b_0(\tau)$ across different τ's can provide a detailed picture of how γ_i's are distributed differently between subjects in the treatment group and those in the placebo group.

Under traditional multiplicative intensity modeling (Nielsen et al., 1992; Oakes, 1992; Wang et al., 2001, e.g.), γ_i is essentially specified by a log-linear model, $\log(\gamma_i) = \widetilde{X}_i^\tau b_0 + \log\xi_i$, where the exponentiated error term, ξ_i, corresponds to the so-called frailty. If $a_0(\tau)$ is a constant a_0, then $\log(\gamma_i)$ degenerates to $a_0 + \widetilde{X}_i^\tau b_0$ and the proposed model reduces to the proportional intensity model (Andersen & Gill, 1982). Such connections with the existing

models suggest that the non-constancy of $b_0(\tau)$, as permitted by model (3), accounts for sources of heterogeneity in γ_i not captured by the location-shift effects of the observed covariates (Koenker, 2005).

Based on the proposed models (1)–(3), we adopt the following censoring assumptions:

- **i.** C_i is independent of $N_i^*(\cdot)$ given γ_i;
- **ii.** C_i is independent of γ_i given $\tilde{X}_i$.

These assumptions allow C_i to depend $\tilde{X}_i$. In the special case of $\log(\gamma_i) = \tilde{X}_i^\tau b_0 + \log\xi_i$, the assumption (ii) is equivalent to assuming C_i is independent of the frailty ξ_i given $\tilde{X}_i$.

3. Estimation and Inference

3.1. Estimating equation

It is easy to see that model (3) is equivalent to $Q_{\log(\gamma)}(\tau|X_i) = X_i^\tau \beta_0(\tau)$, where $Q_{\log(\gamma)}(\tau|X_i)$ denotes the τth conditional quantile of $\log(\gamma)$ given $X = X_i$. If γ_i's were observed, we can easily estimate $\beta_0(\tau)$ through the score equation of the classic quantile loss function (Koenker & Bassett, 1978):

$$\sum_{i=1}^{n} X_i \cdot \psi_\tau\{\log(\gamma_i) - X_i^\tau b\} = 0, \tag{4}$$

where $\psi_\tau(v) = \tau - I(v < 0)$, $I(\cdot)$ denotes the indicator function, and $b \in \mathbb{R}^p$ is a p-dimensional unknown coefficients.

A key challenge for estimating $\beta_0(\tau)$ is that γ_i's are not observable. A naive approach to address this difficulty is to replace the γ_i in (4) by its observable proxy. Since model (1) implies $E\{N_i^*(t)|\gamma_i\} = \gamma_i \mu_0(t)$, an intuitive proxy of γ_i is given by $\hat{\gamma}_i = m_i/\hat{\mu}(C_i)$, where $\hat{\mu}(\cdot)$ is an estimator of $\mu_0(\cdot)$. An example of $\hat{\mu}(\cdot)$ is discussed at the end of this subsection. However, this naive approach, as evidenced by our simulation studies, can produce considerably biased estimation by ignoring the non-negligible deviation of $\hat{\gamma}_i$ from γ_i, on the subject-level, despite that $n^{-1}\sum_{i=1}^{n}\hat{\gamma}_i$ consistently estimate $E(\gamma)$.

Our strategy to deal with unobservable γ_i's is to apply the principle of conditional score (Stefanski & Carroll, 1987) to transform the score equation (4) that involves unobservable γ_i's to a valid estimating equation that only uses observable quantities. We utilize the fact that

$$\begin{aligned} 0 &= E\left[X \cdot \psi_\tau\{\log(\gamma) - X^\tau \beta_0(\tau)\}\right] = E\left[E\left\{X \cdot \psi_\tau\{\log(\gamma) - X^\tau \beta_0(\tau)\}\right| m, C, X\right] \\ &= E\left[\int_r X \cdot \psi_\tau\{\log(r) - X^\tau \beta_0(\tau)\} \cdot f\{r|m, C, X; \beta_0(\cdot), \mu_0(\cdot)\} dr\right], \end{aligned} \tag{5}$$

where $f\{\gamma|m, C, X; \beta_0(\cdot), \mu_0(\cdot)\}$ denotes the conditional density of γ given m, C and X, which depends on $\beta_0(\cdot)$ and $\mu_0(\cdot)$. Equation (5) reflects a critical idea that we choose (m, C)

as the surrogate data to recover the information on γ. As elaborated later, such a choice brings analytical convenience as well as computational feasibility.

By equation (5), we consider constructing an estimating equation based on

$$S_n(\beta, \mu, \tau) \doteq \frac{1}{n}\sum_{i=1}^{n}\int_r X_i \cdot \psi_\tau\{\log(r) - X_i^\tau \beta(\tau)\} f\{r|m_i, C_i, X_i; \beta(\cdot), \mu(\cdot)\} dr .$$

It is clearly seen from (5) that $E[S_n(\beta_0, \mu_0, \tau)] = 0$.

To utilize $S_n(\beta, \mu, \tau)$ to estimate $\beta_0(\tau)$, a crucial step is to derive the analytic form of $f\{r|m, C, X; \beta_0(\cdot), \mu_0(\cdot)\}$. To this end, we note that

$$\begin{aligned} f\{r|m, C, X; \beta_0(\cdot), \mu_0(\cdot)\} &= \frac{\rho\{m|r, C, X; \mu_0(\cdot)\} g\{r|C, X; \beta_0(\cdot)\}}{\int_r \rho\{m|r, C, X; \mu_0(\cdot)\} g\{r|C, X; \beta_0(\cdot)\} dr} \\ &= \frac{\rho\{m|r, C; \mu_0(\cdot)\} g\{r|X; \beta_0(\cdot)\}}{\int_r \rho\{m|r, C; \mu_0(\cdot)\} g\{r|X; \beta_0(\cdot)\} dr} . \end{aligned} \tag{6}$$

where $\rho\{m|r, C, X; \mu_0(\cdot)\}$ denotes the conditional probability mass function of m given ($\gamma = r, C, X$) and $g\{r|C, X; \beta_0(\cdot)\}$ denotes the conditional density of γ at $\gamma = r$ given (C, X). The censoring assumption (ii) implies that $g\{r|C, X; \beta_0(\cdot)\}$ is free of C and so we can simplify the notation $g\{r|C, X; \beta_0(\cdot)\}$ to $g\{r|X; \beta_0(\cdot)\}$. We can also omit X from $\rho\{m|r, C, X; \mu_0(\cdot)\}$ because $m = N^*(C)$ and thus its distribution is fully determined when γ and C are given. These justify the second equality in (6).

First, we examine $\rho\{m|r, C; \mu_0(\cdot)\}$ using the fact that under model (1), $N^*(t)$, given γ, is a nonhomogeneous Poisson process with mean function $\gamma\mu_0(t)$ (Lin et al., 2000). This implies that $\{\mu_0(T^{(1)}), \mu_0(T^{(2)}), \ldots\}$ can be viewed as random variates generated from a homogeneous Poisson process with mean function γt. Using standard probabilistic arguments, we show in Section 1 of the Supplementary Materials that, for both $m = 0$ and $m > 0$,

$$\rho\{m|r, C; \mu_0(\cdot)\} = \frac{\{r\mu_0(C)\}^m}{m!}\exp\{-r\mu_0(C)\} . \tag{7}$$

Next, we assess $g\{r|X; \beta_0(\cdot)\}$ using the relationship between the conditional density function and the conditional quantile function of γ (Wei & Carroll, 2009). Under model (3), we can write

$$g\{r|X; \beta_0(\cdot)\} = \lim_{\delta \to 0} \frac{\delta}{\exp\{X^\tau \beta_0(\tau_r + \delta)\} - \exp\{X^\tau \beta_0(\tau_r)\}}, \tag{8}$$

where $\tau_r = \{\tau \in (0, 1) : \exp\{X^\tau \beta_0(\tau)\} = r\}$. Using the results in (6), (7), and (8), we can express the $f\{r|m_i, C_i, X_i; \beta_0(\cdot), \mu_0(\cdot)\}$ explicitly in terms of $r, m_i, C_i, X_i, \mu_0(\cdot)$, and $\beta_0(\cdot)$. We assume the conditional density of γ given X, $g\{r|X; \beta_0(\cdot)\}$, is bounded away from 0 and ∞, and belong to a compact interval $[M_{g,l}, M_{g,u}]$.

To construct an estimating equation based on $S_n(\beta, \mu, \tau)$, there remains a major obstacle, which is the unknown infinitely-dimensional $\mu(\cdot)$. To address this difficulty, one may follow the conditional likelihood arguments in Wang et al. (2001) to obtain a nonparametric estimator of $\mu_0(t)$, which has a simple product-limit representation. Alternatively, we propose an asymptotic equivalent estimator of $\mu_0(\cdot)$, which takes the Nelson-Aalen form. Define the functions $S_C(t|\gamma) \doteq \Pr(C \geq t|\gamma)$ and $H_0(t) \doteq \log\{\mu_0(t)/\mu_0(v^*)\}$. Given the constraint (2), $\mu_0(v^*) = 1$, it is easy to see that $H_0(v^*) = 0$ and $\mu_0(t) = \exp\{H_0(t)\}$. Under the censoring assumption (i) that C_i is independent of $N_i^*(\cdot)$ given γ_i, the multiplicative intensity structure imposed by model (1) implies that

$$E\{dN_i(t)|\gamma_i\} = S_C(t|\gamma_i)E\{dN_i^*(t)|\gamma_i\} = S_C(t|\gamma_i)\gamma_i\lambda_0(t)dt,$$

and

$$E\{I(C_i \geq t)N_i(t)dH_0(t)|\gamma_i\} = S_C(t|\gamma_i)\gamma_i\mu_0(t)\frac{\lambda_0(t)}{\mu_0(t)}dt = S_C(t|\gamma_i)\gamma_i\lambda_0(t)dt\,.$$

It then follows that $E\{dM_i(t)\} = 0$, where $dM_i(t) \doteq dN_i(t) - I(C_i \geq t)N_i(t)dH_0(t)$. Solving $\sum_{i=1}^n dM_i(t) = 0$ yields an estimator of $\mu_0(t)$, which is given by $\hat{\mu}(t) = \exp\{\widehat{H}(t)\}$ with

$$\widehat{H}(t) = -\int_t^{v^*} \frac{\sum_{i=1}^n dN_i(s)}{\sum_{i=1}^n I(C_i \geq s)N_i(s)}\,.$$

Plug $\hat{\mu}(t)$ into the explicit expression of $f\{r|m_i, C_i, X_i; \beta(\cdot), \mu(\cdot)\}$ and denote the resulting $f\{r|m_i, C_i, X_i; \beta(\cdot), \mu(\cdot)\}$ and $S_n(\beta, \mu, \tau)$ by $f\{\gamma|m_i, C_i, X_i; \beta(\cdot), \hat{\mu}(\cdot)\}$ and $S_n(\beta, \hat{\mu}, \tau)$ respectively. Then the proposed estimating equation takes the form

$$n^{1/2}S_n(\beta, \hat{\mu}, \tau) = 0\,. \tag{9}$$

We shall derive the proposed estimator of $\beta_0(\tau)$, denoted by $\hat{\mu}(\tau)$, from this estimating equation.

3.2. Estimation algorithm

Solving estimating equation (9) is not straightforward because $g\{\gamma|X; \beta(\cdot)\}$ is expressed as a limit and $S_n(\beta, \hat{\mu}, \tau)$ involves integrals with respect to γ. To assess $g\{\gamma|X; \beta(\cdot)\}$, we adapt the strategy proposed by Wei & Carroll (2009) for quantile regression with covariate measurement errors, and approximate $\beta(\tau)$ by a cadlag piecewise-constant function that jumps only on an equally spaced grid on (0, 1), denoted by $S_{K_n} = \{0 = \tau_0 < \tau_1 < \tau_2 < \ldots < \tau_{K_n} < \tau_{K_n+1} = 1\}$ with $\tau_k = k/(K_n + 1)$. In the sequel, we may use notation K instead of K_n for notation simplicity.

Given (8), we propose to approximate $g\{\gamma|X; \beta(\cdot)\}$ by

$$\widetilde{g}_n\{r|X;\beta(\cdot)\} = \min\left\{M_{g,u}, \max\left(M_{g,l}, \sum_{k=1}^{K} \frac{\tau_k - \tau_{k-1}}{\exp\{X^\tau\beta(\tau_k)\} - \exp\{X^\tau\beta(\tau_{k-1})\}} \cdot I[\exp\{X^\tau\beta(\tau_{k-1})\} < r \le \exp\{X^\tau\beta(\tau_k)\}]\right)\right\},$$

with $\exp\{X^\tau\beta(0)\}$ fixed as 0, and $[M_{g,l}, M_{g,u}]$ assumed to bound $g\{\gamma|X;\beta(\cdot)\}$.

Let $\widetilde{f}_n\{r|m_i, C_i, X_i;\beta(\cdot),\hat{\mu}(\cdot)\}$ denote $f\{r|m_i, C_i, X_i;\beta(\cdot),\hat{\mu}(\cdot)\}$ with $\widetilde{g}_n\{r|X;\beta(\cdot)\}$ in place of $g\{r|X;\beta(\cdot)\}$. Note that $\widetilde{g}_n\{r|X;\beta(\cdot)\}$ only involves a $(K\cdot p)$-dimensional parameter, $\theta_n(\beta) = (\beta(\tau_1))^\tau, (\beta(\tau_2))^\tau, \ldots, (\beta(\tau_K)^\tau)^\tau$, and so does $\widetilde{f}_n\{r|m_i, C_i, X_i;\beta(\cdot),\hat{\mu}(\cdot)\}$. Write $\widetilde{g}_n\{r|X;\beta(\cdot)\}$ as $\widetilde{g}_n(r|X;\theta)$. Expressing $\widetilde{f}_n\{r|m_i, C_i, X_i;\beta(\cdot),\hat{\mu}(\cdot)\}$ as $\widetilde{f}_n(r|m_i, C_i, X_i;\theta)$ and using it in place of the $f\{r|m_i, C_i, X_i;\beta(\cdot),\hat{\mu}(\cdot)\}$ in $S_n(\beta,\mu,\tau)$, we transform equation (9) into an estimating equation, which can be written as

$$\widetilde{S}_n(\theta) \doteq \frac{1}{n}\sum_{i=1}^{n}\int_r \Psi\left\{\log(r) - X_i^\tau\theta\right\} \otimes X_i \cdot \widetilde{f}_n(r|m_i, C_i, X_i;\theta)dr = 0, \tag{10}$$

where $\Psi\left\{\log(r) - X_i^\tau\theta\right\} = \left(\psi_{\tau 1}\{\log(r) - X_i^\tau\beta(\tau_1)\}, \ldots, \psi_{\tau_K}\{\log(r) - X_i^\tau\beta(\tau_K)\}\right)^\tau$, and $\otimes$ denotes Kronecker product. Based on equation (10), we develop the following algorithm for estimating $\beta_0(\cdot)$:

Step 1. Set the initial value $\theta^{[0]} = (\hat{\beta}^{[0]}(\tau_1)^\tau, \ldots, \hat{\beta}^{[0]}(\tau_K)^\tau)^\tau$ as the naive estimates obtained from solving a standard quantile regression problem in equation (4) with $\hat{\gamma}_i$ replacing γ_i. Set $l = 1$.

Step 2. Based on $\theta^{[l-1]}$, evaluate

$$f^{[l]}\left(r|m_i, C_i, X_i;\theta^{[l-1]}\right) = \frac{\rho\{m_i|r, C_i;\hat{\mu}(\cdot)\}\widetilde{g}\left(r|X_i;\theta^{[l-1]}\right)}{\int_r \rho\{m_i|r, C_i;\hat{\mu}(\cdot)\}\widetilde{g}\left(r|X_i;\theta^{[l-1]}\right)dr},$$

where

$$\rho\{m_i|r, C_i;\hat{\mu}(\cdot)\} = \frac{\{r\hat{\mu}(C_i)\}^{m_i}}{m_i!}\exp\{-r\hat{\mu}(C_i)\}.$$

Step 3. Update $\theta^{[l]} = (\hat{\beta}^{[l]}(\tau_1)^\tau, \ldots, \hat{\beta}^{[l]}(\tau_K)^\tau)^\tau$ by the solution to (10) with $\widetilde{f}(r|m_i, C_i, X_i;\theta)$ evaluated at $f^{[l]}\left(r|m_i, C_i, X_i;\theta^{[l-1]}\right)$ Increase l by 1.

Step 4. Repeat Steps 2 and 3 until the algorithm converges.

Step 5. The proposed estimator is given by

$$\hat{\beta}(\tau) = \sum_{k=1}^{K+1} \hat{\beta}(\tau_{k-1}) I(\tau_{k-1} \le \tau < \tau_k).$$

To implement the presented algorithm, we adopt numerical integration to assess the integrals with respect to r. In Step 3, finding the solution to (10) can be transformed to a weighted quantile regression problem. Let $\tilde{\gamma}_i^{[l]} = (\tilde{r}_{i,1}^{[l]}, \tilde{r}_{i,2}^{[l]}, \ldots, \tilde{r}_{i,J}^{[l]})^\tau$ be a fine grid of possible r_i values in the lth step. The estimating equations (10) can be approximated by

$$\sum_{i=1}^{n} \sum_{j=1}^{J-1} X_i \cdot \psi_{\tau_k}\left\{\log(\tilde{r}_{i,j}^{[l]}) - X_i^\tau \beta(\tau_k)\right\} f^{[l]}(\tilde{r}_{i,j}^{[l]} \Big| m_i, C_i, X_i)(\tilde{r}_{i,j+1}^{[l]} - \tilde{r}_{i,j}^{[l]}) = 0, \quad (11)$$

for $k = 1, \ldots, K$. Equation (11) can be viewed as a weighted quantile regression problem with responses, $\log(\tilde{r}_{i,j}^{[l]})$, and covariates, X_i, along with weights $f^{[l]}(\tilde{r}_{i,j}^{[l]} \Big| m_i, C_i, X_i)(\tilde{r}_{i,j+1}^{[l]} - \tilde{r}_{i,j}^{[l]})$. Then estimating equation (11) can be solved by standard statistical software, such as the *rq*() function in R package *quantreg*.

The presented estimation algorithm involves the choice of the τ-grid S_K and the γ-grid $\tilde{\gamma}_i^{[l]}$. By our asymptotic studies, we require the grid size of S_K, defined as $\|S_K\| \doteq \max\{\tau_{k+1} - \tau_k, k = 1, \ldots, K-1\}$, which equals $(K_n + 1)^{-1}$, is of asymptotic order $o(n^{-1/2})$. We also suggest choosing $J = K$ and setting $\tilde{\gamma}_i^{[l]}$ as $(\exp\{X_i^\tau \hat{\beta}^{[l-1]}(\tau_1)\}, \ldots, \exp\{X_i^\tau \hat{\beta}^{[l-1]}(\tau_K)\})^\tau$ for computational simplicity.

3.3. Large sample studies

We first introduce the regularity conditions and necessary notation. For a vector u, let $\|u\|$ denote its Euclidean form, $u^{(j)}$ denote the jth component of u, and $u^{(s:t)}$ denote the subvector consisting of the s-th to t-th component of u. Let $\bar{R} = R \cup \{-\infty\}$. For $y \doteq (y_1, \ldots y_p)^\tau \in \bar{R}^p$, define

$$\tilde{D}(y) = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ -1 & e^{-y_2} & & 0 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & 0 & \cdots & e^{-y_p} \end{pmatrix}.$$

Let χ denotes the support of X, $\beta(\bar{S}) = \{f : [0,1) \to \bar{R}^p; f(\cdot)$ is cadlag; for any $\tau \in [0,1)$, $f(\tau)^{(1)} \le \bar{S}$, and $\|f(\tau)^{(2:p)}\| \le \bar{S}\}$, and $\mathcal{U} = \{f : [0,\infty) \to R; f(\cdot)$ is cadlag, nonnegative, and non-decreasing with $f(0) = 0\}$. Let $C^p[0,1)$ denote the set of p-dimensional differentiable functions on $[0,1)$. Define $h_x(\tau) = \{d \exp\{x^\tau \beta_0(\tau)\}/d\tau\}^{-1}$ and $\dot{h}_x(\tau) = dh_x(\tau)/d\tau$. The following are the regularity conditions:

C1. (a) χ is compact; (b) $\beta_0 \in \beta(\bar{S})$ for some $\bar{S} < \infty$;; (c) $N^*(\nu^*)$ is bounded, a.s..

C2. (a) $\Pr(C > \nu^*) = 0$ and $\Pr(C = \nu^*) > 0$; (b) $\Pr(C < \nu^*) = 0$ for some $\nu^* \in (0, \nu^*)$.

C3. $\inf_{t \in [v*, v^*]} E\{S_C(t|\gamma)\gamma\}\mu_0(t) > 0$.

C4. (a) $\beta_0(\tau)$ is continuously differentiable in $\tau \in (0, 1)$; (b) $0 < M_{g,l} < \sup_{x \in \chi} h_x(\tau) < M_{g,u} < \infty$; (c) there exist constants M_h and v_1, $v_2 > -1$ such that $\sup_{x \in \chi} \dot{h}_x(\tau)$ is bounded above by $M_h \tau^{v_1} (1 - \tau)^{v_2}$.

C5. For $\beta \in \beta(\bar{S})$ and $\mu \in \mathcal{U}$ in a neighborhood of μ_0,

$$\text{eigmin}_{\eta j \in \mathcal{R}(\bar{S}), 1 \le j \le p} w_n(\eta_1, \ldots, \eta_p; \beta, \mu)^{-1} > 0$$

for $n >$ some N_0, where eigmin$(\cdot)$ denotes the minimum eigenvalue of a matrix and

$$w_n(\eta_1, \ldots, \eta_p; \beta, \mu) = \mathrm{E}\{XX^{\tau}\mathrm{diag}(\tilde{f}_n\{X^{\tau}\eta_j | m, C, X; \beta(\cdot), \mu(\cdot)\}e^{\tilde{X}^{\tau}\eta^{(2:p)}}\tilde{D}(\eta), 1 \le j \le p)\}.$$

C6. $\beta_0(\cdot)$ is the unique solution to $s(\beta, \mu_0, \tau) = 0$ for $\beta \in C^p[0, 1)$ and $\tau \in (0, 1)$, where $s(\beta, \mu, \tau) = \mathrm{E}\{\mathrm{S_n}(\beta, \mu, \tau)\}$.

Condition C1 implies realistic boundedness assumptions for X, γ and m. Condition C2 (a) is always satisfied when a truncated censoring time $C^* = \min(C, \nu^*)$ is adopted in place of C, given ν^* is a constant less than the upper bound of C's support. Under Condition C2, we essentially only utilize recurrent events occurred up to time ν^*, and thus we need to set ν^* as large as possible to minimize information loss. Conditions C3 and C2(b) play an important role to achieve desirable asymptotic behaviors of $\hat{\mu}(C)$. Condition C4(a) assumes the smoothness of the true coefficient function $\beta_0(\tau)$, which has been commonly adopted in quantile regression literature. The assumptions in conditions C4(b) and C4(c) were similarly adopted by Wei & Carroll (2009) and can help justify the approximation of $g\{\gamma|X; \beta(\cdot)\}$ by $\tilde{g}_n(\gamma|X; \theta)$. Conditions C5 and C6 are the key assumptions to ensure the identifiability of $\beta_0(\tau)$, and consequently the uniform consistency of $\hat{\beta}(\tau)$.

We establish the uniform consistency and weak convergence of the proposed estimator $\hat{\beta}(\tau)$ in the following theorems:

THEOREM 1. *Under regularity conditions C1–C6, if* $K_n \to \infty$ *and* $K_n/n^{a0} \to 0$ *for some* $a_0 > 0$, *then* $\lim_{n \to \infty} \sup_{\tau \in [\zeta_1, \zeta_c]} \|\hat{\beta}(\tau) - \beta_0(\tau)\| \to_p 0$, *where* ζ_1 *and* ζ_2 *are constants satisfying* $0 < \zeta_1 < \zeta_2 < 1$.

THEOREM 2. *Under regularity conditions C1–C6, if* $n^{-1/2}K_n \to \infty$ *and* $K_n/n^{a0} \to 0$ *for some* $a_0 > 1/2$, *then* $n^{1/2}\{\hat{\beta}(\tau) - \beta_0(\tau)\}$ *converges weakly to a Gaussian process for* $\tau \in [\zeta_1, \zeta_2]$, *where* ζ_1 *and* ζ_2 *are constants satisfying* $0 < \zeta_1 < \zeta_2 < 1$.

Note that Theorems 1-2 are focused on the asymptotic properties of $\hat{\beta}(\tau)$ *with* $\tau \in [\zeta_1, \zeta_2]$, a closed subset of (0, 1) away from 0. This is necessary because model (3) implies that $\exp\{X^{\tau}\beta_0(0)\} = 0$ and hence $\|\beta_0(0)\}\| = \infty$. Following the strategy in Peng & Huang (2007),

we circumvent this difficulty by considering reparameterizing $\beta_0(\tau)$ by $\alpha_0(\tau) = \kappa^{-1}\{\beta_0(\tau)\}$, where $\kappa^{-1}(y) = \left(e^{y_1}, e^{y_1+y_2}, \ldots, e^{y_1+y_p}\right)^\tau$ for $y = (y_1, \ldots, y_p)^\tau \in R^p$. A key advantage of working on $\alpha_0(\tau)$ is that the derivative matrix of the estimating function with respect to $\alpha_0(\tau)$ can have eigenvalues bounded away zero uniformly across $\tau \in (0, 1)$ (see condition C5). Thus, we are able to establish the uniform consistency of $\hat{\alpha}(\tau) \doteq \kappa\left\{\hat{\beta}(\tau)\right\}$ over $\tau \in (0, 1)$. Through sophisticated derivations, we also uncover the link between $n^{1/2}\{\hat{\alpha}_n(\tau) - \alpha_0(\tau)\}$ and a tight Gaussian process via a Fredholm integral equation of the second kind, which entails the weak convergence of $n^{1/2}\{\hat{\alpha}_n(\tau) - \alpha_0(\tau)\}$. Given the one-to-one and smooth transformation between $\hat{\alpha}$ and $\hat{\beta}$ and that between α_0 and β_0, it follows the uniform consistency of $\hat{\beta}(\tau)$ and weak convergence of $n^{1/2}\left\{\hat{\beta}(\tau) - \beta_0(\tau)\right\}$ for $\tau \in [\zeta_1, \zeta_2]$ where the derivative of $\kappa(\cdot)$ around $\beta_0(\tau)$ is uniformly bounded. Detailed proofs of Theorem 1 and Theorem 2 are provided in Sections 1–4 of Supplementary Materials.

As suggested by Theorems 1 and 2, K_n needs to be chosen properly. The requirement of $K_n \to \infty$ is well expected to make $\tilde{g}_n\{r|X; \beta(\cdot)\}$ closely approximate $g\{r|X; \beta(\cdot)\}$. The assumption, $\lim_{n\to\infty} K_n/n^{a_0} = 0$ for some $a_0 > 0$, is to ensure $\tilde{S}_n(\theta)$ is uniformly close to its smooth counterpart $\tilde{s}_n(\theta)$ despite the diverging dimension of θ. By this assumption, K_n should be at most the polynomial order of n, and is not allowed to increase with n too fast, say at the exponential rate. The assumption, $\lim_{n\to\infty} n^{-1/2}K_n = \infty$, is to control the estimation errors from the grid approximation of $\beta_0(\cdot)$ by $o(n^{-1/2})$. Note that a larger K_n requires more computation efforts. In practice, we recommend setting $K_n = O(n^r)$ with $1/2 < r \leq 1$ for a good balance between estimation performance and computational intensity.

3.4. Bootstrapping-based inference

To make inference about $\beta_0(\tau)$, a bootstrapping procedure may be preferred provided the complexity of the asymptotic covariance matrix derived in the proof of Theorem 2. Specifically, we may resample the observed data with replacement and obtain an estimator of $\beta_0(\tau)$ based on the resampled sample, denoted by $\beta^*(\tau)$. Repeating this procedure many times can generate a large number of realizations of $\beta^*(\tau)$. For a fixed $\tau^* \in [\tau_1, \tau_K]$, the variance of $\hat{\beta}(\tau*)$ can be estimated by the empirical variance of $\beta^*(\tau^*)$. The confidence intervals for $\beta_0(\tau^*)$ can be constructed using a normal approximation or by referring to the empirical percentiles of $\beta^*(\tau^*)$.

In addition, one may be interested in testing whether some components of $b_0(\tau)$ are constant over τ or not. Rejecting the constancy hypothesis would indicate the lack-of-fit of an existing model that imposes location-shift effects for all covariates, such as the proportional intensity model and Wang et al. (2001)'s semiparametric multiplicative intensity model. The practical implication is that the influence of the corresponding covariate on γ may not be homogeneous across all subjects. Such a finding, coupled with an examination of the heterogeneous pattern of the coefficient estimates over τ, can often lead to useful scientific insight. The second-stage inference procedures can follow similar lines of other work on quantile regression (Peng & Huang, 2008; Peng & Fine, 2009, e.g.); details are omitted here.

4. Numerical studies

4.1. Monte-Carlo simulations

In this subsection, we conduct simulation studies to evaluate the finite sample performance of the proposed methods. Specifically, for subject i, we generate recurrent event times $\left\{T_i^{(j)}, j = 1, 2, \ldots\right\}$ from a Poisson process with rate γ_i. In this case, model (1) is met with $\mu_0(t) = t$.

We first consider the situation where γ_i satisfies a log-linear model with homogeneous errors:

$$\log(\gamma_i) = X_i^{\tau} b + 0.5\epsilon_i, \tag{12}$$

where $b = (b_0, b_1, b_2)^{\tau}$ and $X_i = (1, X_{i,1}, X_{i,2})^{\tau}$. We let $X_{i,1} \sim Unif(0, 1)$, and $X_{i,2} \sim Bernoulli(0.5)$. We let ϵ_i follow the standard normal distribution, $N(0, 1)$, or the Student's t-distribution, t_3. In this set up, model (3) holds with $\beta^{(1)}(\tau) = b_0 + 0.5Q_{\in}(\tau)$, $\beta^{(2)}(\tau) = b_1$,, $\beta^{(3)}(\tau) = b_2$, where the superscript $^{(k)}$ indicates the kth component of a vector, and $Q_{\epsilon}(\tau)$ represents the τth quantile of ϵ. We generate the censoring time C_i from Unif(2/3, 1), independent of $\tau_i^{(j)}$ and $\widetilde{X}_i$. We set $b_0 = \log(3) + 1$, $b_1 = b_2 = 1$, yielding the average number of observed recurrent events per subject is about 24 or 25.8 corresponding to $N(0, 1)$ or t_3 error respectively. Under each configuration, we generate 500 simulated datasets with sample size $n = 500$. For each simulated dataset, 100 bootstrapping samples are drawn to calculate the estimated standard error and coverage probability. To implement the proposed method, S_K is set as an equally spaced grid between 0.02 and 0.98 with the step size 0.01. The naive estimate for $\beta_0(\tau)$ is calculated as the solution to a standard quantile regression problem, $n^{-1/2}\sum_{i=1}^{n} X_i \cdot \psi_{\tau}\{\log(\hat{\gamma}_i) - X_i^{\tau} b\} = 0$ where $\hat{\gamma}_i = \max(1, m_i)/\hat{\mu}(C_i)$. For the iterative algorithm, the maximum iteration number is set to be 100, and the stop criterion is $\sum_{k=1}^{K} \left\|\beta^{[l-1]}(\tau_k) - \beta^{[l]}(\tau_k)\right\|^2 < 0.01$..

The simulation results when ϵ follows $N(0, 1)$ distribution are provided in Table 1. It is shown that the naive estimator can produce large biases, especially for large τ's. In contrast, the empirical biases of the proposed estimator are quite small. Table 1 compares the empirical standard deviations with the estimated standard errors of the proposed estimator. We observe that they match with each other very well. For 95% confidence intervals constructed by normal approximations that use bootstrapping standard errors, the empirical coverage probabilities are reasonably close to the nominal level 95%. It is seen from Table 1 that the square-root mean square error of the proposed estimator is generally smaller than that of the naive estimators. We have similar observations when ϵ follows t_3 distribution, and the detailed results are relegated to Section 5 of the Supplementary Materials.

We also consider the situation where the non-intercept coefficients in $\beta_0(\tau)$ are not constant. To simulate such data, we let γ_i follow a log-linear model with heteroscedastic errors:

$$\log(\gamma_i) = X_i^{\tau} b + \left(X_i^{\tau} d\right) \epsilon_i \ . \tag{13}$$

We generate X_i and C_i in the same way as in the first set-up. We set $b = (b_0, b_1, b_2)^{\tau}$ as before, and set $d = (d_0, d_1, d_2)^{\tau} = (0.1, 0.1, 0.1)^{\tau}$. We only consider ϵ_i that follows the standard normal distribution $N(0, 1)$ in this heteroscedastic case. The average number of observed recurrent events per subject is about 22.8. Under model (13), $\beta^{(i)}(\tau) = b_{i-1} + d_{i-1} Q_{\epsilon}(\tau)$, $i = 1, 2, 3$, which are changing with τ. In Table 2, the simulation results are displayed in the same manner as those in Table 1. It is shown that the proposed estimator $\hat{\beta}^{(1)}(\tau)$ has small biases for τ's ranging from 0.1 to 0.9. Meanwhile, the biases of $\hat{\beta}^{(2)}(\tau)$ and $\hat{\beta}^{(3)}(\tau)$ are negligible except for that corresponding to extremely small and large τ's. The estimated standard errors agree well with the empirical standard deviations, and the confidence intervals yield quite accurate empirical coverages probabilities. Overall, our simulation results suggest satisfactory finite-sample performance of the proposed methods.

4.2. The DURABLE Data Example

The DURABLE study (Buse et al., 2009) is an open-label randomized clinical trial in Type 2 diabetes patients. It was designed to compare the efficacy and safety of two starter insulin regimens, twice-daily lispro mix 75/25, which is 75% lispro protamine suspension plus 25% lispro, or once-daily glargine, in addition to oral antihyperglycemic drugs. This study enrolled 2,187 insulin-naive patients with type 2 diabetes from 11 countries, aged 30 to 80 years, with HbA1c > 7.0%, and on at least two oral antihyperglycemic agents.

Hypoglycemia, as an important safety endpoint, was closely monitored during this study. The number of hypoglycemia episodes observed for each subject ranges from 0 to 137, with mean 10.8 and median 5. These descriptive statistics suggest a high degree of heterogeneity in the individual risk of hypoglycemia presented in the DURABLE trial. Exploring the risk factors for hypoglycemia and, moreover, potentially different risk mechanisms between high risk versus low risk patients are of great clinical interest. The proposed quantile regression framework for recurrent event data is tailored to address these interests, in particular the latter one, which cannot be addressed by routine recurrent event data analyses.

We apply the proposed method to the DURABLE data. The recurrent event time $T^{(j)}$ corresponds to the time from study enrollment to the jth episode of hypoglycemia. We consider baseline covariates including *therapy*, which is 1 if the patient had twice-daily lispro mix 75/25 and 0 otherwise, *basfglu*, which represents baseline fasting blood glucose, *basfins*, which represents baseline fasting insulin, *bmibase*, which represents baseline body mass index, *durdiab*, which represents duration of type 2 diabetes, *tzduse*, which is 1 if the patient used thiazolidinedione and 0 otherwise, and *sulfouse*, which is 1 if the patient used sulfonylurea and 0 otherwise. These covariates are summarized in Table S1 of the Supplementary Material. In our analysis, we standardize continuous covariates, and exclude subjects with missing covariates or those falling outside the reference range. The final sample size is $n = 2{,}003$. We choose S_K as an equally space grid between 0.02 and 0.98 with

step size 0.02. Inferences are carried out based on 200 bootstrapping samples. Other set-ups are the same as those in the simulation studies.

The analysis results for $\tau \in [0.1, 0.9]$ are displayed In Fig. 1. Under the proposed models (1)–(3), positive coefficients indicate covariate effects associated with higher risk of hypoglycemia, which is measured as subject-specific positive scale shift of the intensity function of hypoglycemia recurrence. We can see from Fig. 1 that patients receiving lispro mix 75/25 demonstrate higher risk of hypoglycemia than patients in the glargine group. The results in Fig. 1 also suggest that lower baseline glucose, lower baseline insulin, lower baseline body mass index, longer diabetes duration, or using sulfonylurea, are associated with higher risk of hypoglycemia. We note that the naive estimates sometimes show significant departures from the proposed estimates. For example, the naive estimates for *therapy*'s coefficients are beyond on the upper bound of the confidence intervals when $\tau \in [0.15, 0.25]$. This may be a sign of large estimation bias resulted from using the naive approach.

We also conduct second-stage inference to summarize the estimated covariate effects by average covariate effects, defined as $\int_{\tau_L}^{\tau_U} \beta_0^{(j)}(\tau) d\tau / (\tau_U - \tau_L)$, where $\tau_L = 0.1$, $\tau_U = 0.9$ and $j = 2, \ldots, p$. The inferences on the average effects are conducted by following the lines of Peng & Fine (2009). We also fit the data with the standard proportional intensity model (Andersen & Gill, 1982), which is a special case of the proposed models with all coefficients in $\beta_0(\tau)$ being constant over τ. In Table 1, we present the estimated average effects and the corresponding standard errors and Wald-test p-values, along with the coefficient estimates and the corresponding standard errors and p-values based on the proportional intensity model. It is seen that the proposed method generates quite consistent findings regarding the impact of covariates on the risk of hypoglycemia. The standard errors based on the proposed method are larger than those based on the proportional intensity model. This reflects a tradeoff between greater model flexibility and reduced estimation efficiency.

We next employ second-stage inference to test the constancy of each coefficient function in $\beta_0(\tau)$. We apply this test to the coefficient for each covariate. The results indicate that *durdiab* and *sulfouse* have non-constant effects over τ, while constant effects are adequate for other covariates. Combined with the observation from Fig. 1, this suggests that the elevated risk of hypoglycemia associated with the use of sulfonylurea may be amplified in subjects who are susceptible to frequent hypoglycemia (corresponding to large τ's). A clinical implication is that caution may be needed for using sulfonylurea in patients who are known or projected to have a high risk of hypoglycemia based on patient history and clinical judgement. The non-constancy of the coefficients for *durdiab* and *sulfouse* also provide an evidence for the lack of fit of the proportional intensity model to the DURABLE data.

5. Discussion

There are other applications or generalizations of quantile regression to recurrent event data. For example, Luo et al. (2013) studied the quantile regression modeling of gap times between recurrent events. Huang & Peng (2009) and Sun et al. (2016) proposed the accelerated recurrence time model, which can reduce to a quantile regression model when

the event of interest is not recurrent. The modeling perspective proposed in this work is fundamentally different from these existing approaches by its unique focus on a sensible latent measure of subject-specific recurrent event risk. This new strategy is expected to yield more straightforward interpretations regarding the heterogeneity in individual recurrent event risk.

The proposed modeling allows for the prediction of median or other quantiles of the latent risk measure γ given the observed covariates based on the estimation results for model (3). For example, for subject i with covariates in X_i, $Q_\gamma(0.5|X_i)$ can be predicted by $\exp\left\{X_i^T\hat{\beta}(0.5)\right\}$, and this prediction may be used as a proxy to reflect the recurrent event risk of this subject.

It is easy to show that a different choice of v^* in the assumed models would only induce a constant scale shift to $\lambda_0(t)$ and a constant location shift to the intercept coefficient $a_0(\tau)$, while the covariate coefficients, $b_0(\tau)$, the estimand of key interest, would remain the same. The discrepancies in estimates for $b_0(\tau)$ are thus expected to be asymptotically negligible.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgement

We thank the editor, associate editor, and two referees for their valuable comments. We are grateful to Dr. Yonggang Yao for helping improve the implementation of the proposed method. This research was primarily supported by the U.S. NIH grant R01HL113548 (Dr. Peng). We acknowledge the support of other grants, R01CA193888 (Dr. Huang), and National NSF of China grants 11901200 and 71931004 and Shanghai Pujiang Program 19PJ1403400 (Dr. Ma).

References

Andersen PK & Gill RD (1982). Cox's regression model for counting processes: a large sample study. The Annals of Statistics 10, 1100–1120.

Buse J, Wolffenbuttel B, Herman W, Shemonsky N, Jiang H, Fahrbach J, Scism-Bacon J & Martin S (2009). Durability of basal versus lispro mix 75/25 insulin efficacy (durable) trial 24-week results: safety and efficacy of insulin lispro mix 75/25 versus insulin glargine added to oral antihyperglycemic drugs in patients with type 2 diabetes. Diabetes Care 32, 1007–1013. [PubMed: 19336625]

Huang Y & Peng L (2009). Accelerated recurrence time models. Scandinavian Journal of Statistics 36, 636–648.

Koenker R (2005). Quantile regression, vol. 38. Cambridge Univ Pr.

Koenker R & Bassett G (1978). Regression quantiles. Econometrica 46, 33–50.

Lawless J & Nadeau C (1995). Some simple robust methods for the analysis of recurrent events. Technometrics 37, 158–168.

Lin DY, Wei LJ, Yang I & Ying Z (2000). Semiparametric regression for the mean and rate functions of recurrent events. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 62, 711–730.

Luo X, Huang C-Y & Wang L (2013). Quantile regression for recurrent gap time data. Biometrics 69, 375–385. [PubMed: 23489055]

Nielsen GG, Gill RD, Andersen PK & Sørensen TI (1992). A counting process approach to maximum likelihood estimation in frailty models. Scandinavian journal of Statistics , 25–43.

Oakes D (1992). Frailty models for multiple event times. In Survival Analysis: State of the Art. Springer, pp. 371–379.

Peng L & Fine J (2009). Competing risks quantile regression. Journal of the American Statistical Association 104, 1440–1453.

Peng L & Huang Y (2007). Survival analysis with temporal covariate effects. Biometrika 94, 719–733.

Peng L & Huang Y (2008). Survival analysis with quantile regression models. Journal of the American Statistical Association 103, 637–649.

Pepe MS & Cai J (1993). Some graphical displays and marginal regression analyses for recurrent failure times and time dependent covariates. Journal of the American Statistical Association 88, 811–820.

Prentice R, Williams B & Peterson A (1981). On the regression analysis of multivariate failure time data. Biometrika 68, 373–379.

Raudenbush SW & Bryk AS (2002). Hierarchical linear models : applications and data analysis methods. Thousand Oaks, CA: Sage Publications, 2nd Ed.

Schaubel DE, Zeng D & Cai J (2006). A semiparametric additive rates model for recurrent event data. Lifetime Data Analysis 12, 389–406. [PubMed: 17031499]

Stefanski L & Carroll R (1987). Conditional scores and optimal scores for generalized linear measurement-error models. Biometrika 74, 703–716.

Sun X, Peng L, Huang Y & Lai HJ (2016). Generalizing quantile regression for counting processes with applications to recurrent events. Journal of the American Statistical Association 111, 145–156. [PubMed: 27212738]

Wang M-C, Qin J & Chiang CT (2001). Analyzing recurrent event data with informative censoring. Journal of the American Statistical Association 96, 1057–1065.

Wei Y & Carroll RJ (2009). Quantile regression with measurement error. Journal of the American Statistical Association 104, 1129–1143. [PubMed: 20305802]

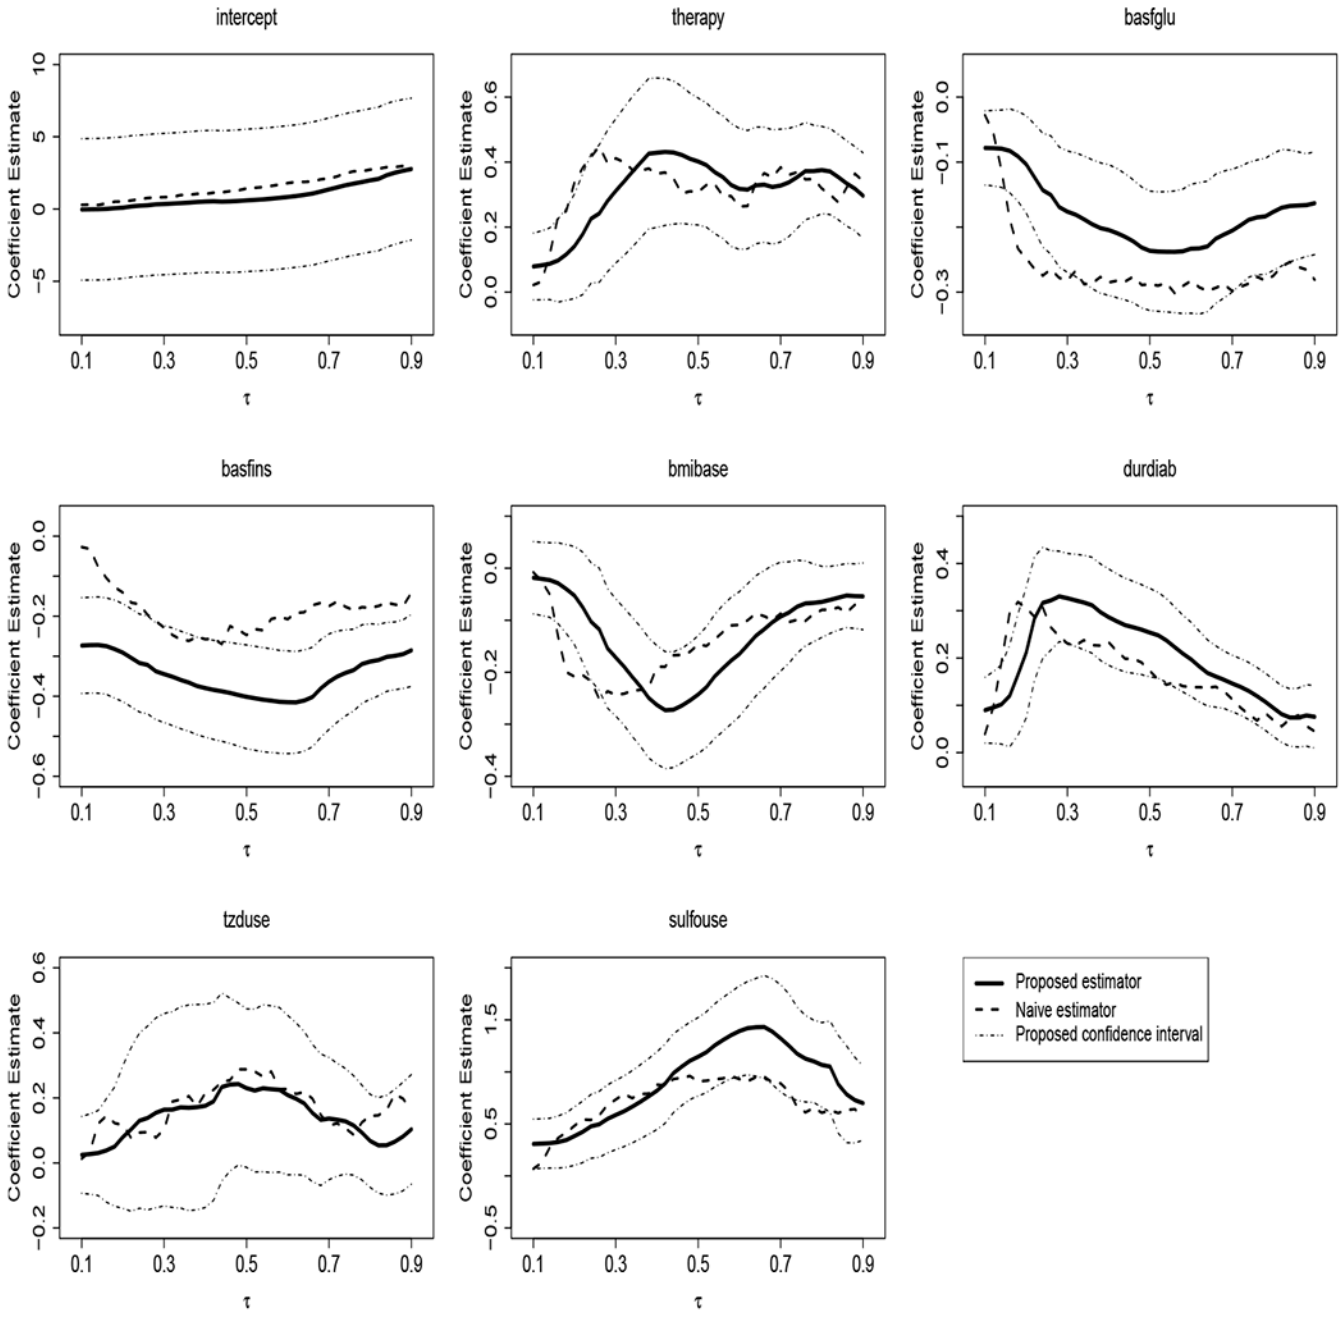

Fig. 1.

The DURABLE data example: the proposed coefficient estimates (solid lines) and the corresponding point-wise confidence intervals (dash dotted lines), along with the naive estimates (dashed lines).

Table 1.

Simulation results with the homogeneous error $\epsilon \sim N(0, 1)$

		Proposed					Naive	
τ		Bias	SD	ESE	CP	sMSE	Bias	sMSE
0.1	$\hat{\beta}^{(1)}$	−0.0246	0.1235	0.1347	95.0	0.1259	−0.3022	0.3307
	$\hat{\beta}^{(2)}$	0.0178	0.1702	0.1811	94.2	0.1712	0.1398	0.2304
	$\hat{\beta}^{(3)}$	0.0161	0.1061	0.1071	94.6	0.1073	0.1513	0.1871
0.2	$\hat{\beta}^{(1)}$	−0.0070	0.0974	0.1062	94.6	0.0977	−0.1621	0.1894
	$\hat{\beta}^{(2)}$	0.0041	0.1374	0.1435	94.0	0.1374	0.0774	0.1587
	$\hat{\beta}^{(3)}$	0.0082	0.0842	0.0843	95.6	0.0846	0.0834	0.1169
0.3	$\hat{\beta}^{(1)}$	−0.0054	0.0880	0.0957	94.8	0.0881	−0.0970	0.1302
	$\hat{\beta}^{(2)}$	0.0051	0.1264	0.1289	94.8	0.1265	0.0490	0.1333
	$\hat{\beta}^{(3)}$	0.0050	0.0764	0.0754	93.4	0.0765	0.0518	0.0904
0.4	$\hat{\beta}^{(1)}$	−0.0004	0.0834	0.0913	94.4	0.0834	−0.0496	0.0940
	$\hat{\beta}^{(2)}$	−0.0023	0.1242	0.1234	93.8	0.1243	0.0238	0.1224
	$\hat{\beta}^{(3)}$	0.0031	0.0698	0.0721	95.2	0.0699	0.0290	0.0734
0.5	$\hat{\beta}^{(1)}$	0.0035	0.0820	0.0891	92.6	0.0821	−0.0118	0.0777
	$\hat{\beta}^{(2)}$	−0.0048	0.1181	0.1197	94.2	0.1182	0.0072	0.1119
	$\hat{\beta}^{(3)}$	0.0016	0.0679	0.0705	96.2	0.0679	0.0103	0.0654
0.6	$\hat{\beta}^{(1)}$	0.0099	0.0809	0.0885	94.2	0.0815	0.0155	0.0763
	$\hat{\beta}^{(2)}$	−0.0084	0.1162	0.1181	94.6	0.1165	−0.0060	0.1084
	$\hat{\beta}^{(3)}$	−0.0020	0.0655	0.0700	96.8	0.0655	−0.0024	0.0603
0.7	$\hat{\beta}^{(1)}$	0.0082	0.0804	0.0905	95.4	0.0808	0.0368	0.0812
	$\hat{\beta}^{(2)}$	−0.0088	0.1150	0.1211	95.0	0.1153	−0.0145	0.1056
	$\hat{\beta}^{(3)}$	−0.0050	0.0670	0.0717	94.6	0.0672	−0.0155	0.0629
0.8	$\hat{\beta}^{(1)}$	0.0111	0.0865	0.0961	94.2	0.0872	0.0594	0.0970
	$\hat{\beta}^{(2)}$	−0.0169	0.1203	0.1290	96.2	0.1215	−0.0304	0.1157
	$\hat{\beta}^{(3)}$	−0.0069	0.0749	0.0756	95.0	0.0752	−0.0268	0.0744
0.9	$\hat{\beta}^{(1)}$	0.0111	0.0936	0.1073	95.8	0.0943	0.0677	0.1091
	$\hat{\beta}^{(2)}$	−0.0218	0.1336	0.1479	95.4	0.1354	−0.0302	0.1330

		Proposed					Naive	
τ		Bias	SD	ESE	CP	sMSE	Bias	sMSE
	$\hat{\beta}^{(3)}$	−0.0085	0.0875	0.0876	93.6	0.0879	−0.0301	0.0859

SD, empirical standard deviation; ESE, estimated standard error; CP, empirical coverage probability; sMSE, square-root mean square error.

Table 2.

Simulation results with the heterogeneous error $\epsilon \sim N(0, 1)$

		Proposed					Naive	
τ		Bias	SD	ESE	CP	sMSE	Bias	sMSE
0.1	$\hat{\beta}^{(1)}$	−0.0193	0.1259	0.1310	95.4	0.1274	−0.3117	0.3410
	$\hat{\beta}^{(2)}$	0.0087	0.1801	0.1835	94.2	0.1803	0.1519	0.2487
	$\hat{\beta}^{(3)}$	0.0173	0.1045	0.1089	95.0	0.1059	0.1600	0.1917
0.2	$\hat{\beta}^{(1)}$	−0.0013	0.0984	0.1043	95.4	0.0984	−0.1650	0.1927
	$\hat{\beta}^{(2)}$	−0.0086	0.1451	0.1479	94.8	0.1453	0.0745	0.1641
	$\hat{\beta}^{(3)}$	0.0082	0.0847	0.0860	94.2	0.0851	0.0931	0.1247
0.3	$\hat{\beta}^{(1)}$	0.0057	0.0887	0.0943	93.6	0.0889	−0.0979	0.1297
	$\hat{\beta}^{(2)}$	−0.0147	0.1341	0.1333	93.6	0.1349	0.0463	0.1368
	$\hat{\beta}^{(3)}$	0.0049	0.0760	0.0770	95.4	0.0762	0.0581	0.0923
0.4	$\hat{\beta}^{(1)}$	0.0099	0.0821	0.0885	94.4	0.0826	−0.0451	0.0905
	$\hat{\beta}^{(2)}$	−0.0195	0.1256	0.1254	94.2	0.1271	0.0169	0.1192
	$\hat{\beta}^{(3)}$	0.0016	0.0710	0.0722	94.6	0.0711	0.0323	0.0748
0.5	$\hat{\beta}^{(1)}$	0.0135	0.0764	0.0872	95.0	0.0776	−0.0068	0.0719
	$\hat{\beta}^{(2)}$	−0.0221	0.1187	0.1218	95.0	0.1208	−0.0016	0.1106
	$\hat{\beta}^{(3)}$	−0.0010	0.0692	0.0710	94.0	0.0692	0.0128	0.0654
0.6	$\hat{\beta}^{(1)}$	0.0209	0.0780	0.0869	93.2	0.0808	0.0242	0.0749
	$\hat{\beta}^{(2)}$	−0.0260	0.1229	0.1216	93.2	0.1256	−0.0162	0.1123
	$\hat{\beta}^{(3)}$	−0.0086	0.0678	0.0707	95.0	0.0683	−0.0048	0.0632
0.7	$\hat{\beta}^{(1)}$	0.0194	0.0787	0.0894	96.4	0.0810	0.0521	0.0873
	$\hat{\beta}^{(2)}$	−0.0262	0.1235	0.1246	94.8	0.1263	−0.0307	0.1161
	$\hat{\beta}^{(3)}$	−0.0105	0.0724	0.0725	94.2	0.0732	−0.0202	0.0687
0.8	$\hat{\beta}^{(1)}$	0.0282	0.0791	0.0932	94.2	0.0840	0.0752	0.1035
	$\hat{\beta}^{(2)}$	−0.0393	0.1252	0.1315	94.6	0.1312	−0.0439	0.1236
	$\hat{\beta}^{(3)}$	−0.0168	0.0770	0.0771	93.6	0.0788	−0.0333	0.0779
0.9	$\hat{\beta}^{(1)}$	0.0344	0.0900	0.1026	94.0	0.0963	0.0965	0.1275
	$\hat{\beta}^{(2)}$	−0.0581	0.1372	0.1456	93.6	0.1490	−0.0545	0.1451

		Proposed					Naive	
τ		Bias	SD	ESE	CP	sMSE	Bias	sMSE
	$\hat{\beta}^{(3)}$	−0.0218	0.0871	0.0868	93.6	0.0898	−0.0453	0.0936

SD, empirical standard deviation; ESE, estimated standard error; CP, empirical coverage probability; sMSE, square-root mean square error.

Table 3.

Comparison of the Proposed Method with the Proportional Intensity Model

	Average Effect			Proportional Intensity Model		
	Estimate	**Standard Error**	***p*-value**	**Estimate**	**Standard Error**	***p*-value**
therapy	0.315	0.072	< 0.001	0.274	0.015	< 0.001
basfglu	−0.182	0.036	< 0.001	−0.173	0.008	< 0.001
basfins	−0.349	0.052	< 0.001	−0.384	0.013	< 0.001
bmibase	−0.135	0.042	0.001	−0.045	0.008	< 0.001
durdiab	0.201	0.033	< 0.001	0.118	0.007	< 0.001
tzduse	0.142	0.092	0.125	0.113	0.016	< 0.001
sulfouse	0.903	0.152	< 0.001	0.755	0.037	< 0.001