Modeling Interval-Censored, Clustered Cow Udder Quarter Infection Times Through the Shared Gamma Frailty Model

K. GOETHALS, B. AMPE, D. BERKVENS, H. LAEVENS, Paul JANSSEN, and Luc DUCHATEAU

Time to infection data are often simultaneously clustered and interval-censored. The time to infection is not known exactly; it is only known to have occurred within a certain interval. Moreover, observations often occur in clusters. Consequently, the independence assumption does not hold. Here we propose an extension of the shared gamma frailty model to handle the interval censoring and clustering simultaneously. We show that the frailties can be integrated out analytically, allowing maximization of the marginal likelihood to obtain parameter estimates. We apply our method to a longitudinal study with periodic follow-up on dairy cows to investigate the effect of parameters at the cow level (e.g., parity) and also parameters that can change within the cow (e.g., front or rear udder quarter) on time to infection. Dairy cows were assessed approximately monthly for the presence of a bacterial infection at the udder-quarter level, thus generating interval-censored data. Obviously, the four udder quarters are clustered within the cow. Based on simulations, we find that ignoring the interval-censored nature of the data can lead to biased parameter estimates.

Key Words: Dairy cow; Interval censoring; Mastitis; Multivariate survival data.

1. INTRODUCTION

Infectious disease data can be analyzed most efficiently by survival analysis techniques if infection times of individual units are available. But ordinary survival analysis techniques often must be extended due to the particular data structure. We investigate a data set with two characteristics that require extension of the currently available survival analysis techniques if they are to be dealt with simultaneously. First, infectious disease data are often hierarchically structured, with observation units grouped in clusters, so that the event

K. Goethals is a doctoral student (E-mail: *Klara.Goethals@UGent.be*), B. Ampe is a doctoral student, H. Laevens is a postdoctoral researcher, and Luc Duchateau is Associate Professor, Department of Physiology and Biometrics, Faculty of Veterinary Medicine, Ghent University, Salisburylaan 133, 9820 Merelbeke, Belgium. D. Berkvens is Associate Professor, Institute of Tropical Medicine, Nationalestraat 155, 2000 Antwerp, Belgium. Paul Janssen is Professor, Hasselt University, Center for Statistics, Agoralaan, 3590 Diepenbeek, Belgium.

Journal of Agricultural, Biological, and Environmental Statistics, Volume 14, Number 1, Pages 1–14
DOI: 10.1198/jabes.2009.0001

times within a cluster cannot be assumed to be independent. Second, the time to infection is not known exactly; it is only known that the infection occurred between the last negative test and the first positive test; therefore, the infection time is interval-censored (Finkelstein 1986). The infectious disease data set considered in this paper, the mastitis data (available as an online supplement to this article), corresponds to infection times of individual cow udder quarters with a bacterium (Laevens et al. 1997). Obviously, the four udder quarters are clustered within a cow (Adkinson et al. 1993), and udder quarters are sampled only approximately monthly, generating interval-censored data. We investigate the effect of covariates that change within a cow (e.g., front and rear udder quarters) as well as covariates that change between cows (e.g., parity [the number of previous calvings]). The correlation between udder infection times within a cow also is of interest, because this is a measure of the infectivity of the agent causing the disease (Barkema et al. 1997). Various models have been applied to address the problem of interdependence for right-censored observations; of these, the frailty model is the standard. Although numerous inferential techniques for analyzing independent interval-censored data have been described in the literature (Collet 1994; Radke 2003), analyzing data that are simultaneously correlated and interval-censored has received less attention. Bellamy et al. (2004) has proposed a method for fitting clustered interval-censored data assuming a normal distribution for the random effects and integrating out the random effects numerically using Gaussian quadrature.

In Section 2, we present an extension of the parametric shared gamma frailty model to interval-censored data. We show that a closed form of the marginal likelihood can be obtained by integrating out the gamma-distributed frailties. In Section 3, we demonstrate the method using an example. In Section 4, we evaluate the method's performance based on a simulation study. We present some conclusions in Section 5.

2. THE PARAMETRIC SHARED GAMMA FRAILTY MODEL WITH INTERVAL-CENSORED DATA

Consider the following proportional hazard frailty model:

$$h_{ij}(t) = h_0(t) w_i \exp(\underset{\sim}{\mathbf{x}}_{ij}^t \underset{\sim}{\boldsymbol{\beta}}), \qquad i = 1, \ldots, N, \ j = 1, \ldots, n_i, \tag{2.1}$$

where $h_{ij}(t)$ is the hazard at time t for udder quarter j of cow i, $h_0(t)$ is the baseline hazard at time t, $\underset{\sim}{\mathbf{x}}_{ij}$ is the vector of covariates for the corresponding udder quarter, and $\underset{\sim}{\boldsymbol{\beta}}$ is the vector of covariate effects. We further assume that the frailties $w_1, \ldots, w_N$ are independent realizations from a one-parameter gamma density with mean 1 and variance θ

$$f_W(w_i) = \frac{w_i^{1/\theta - 1} \exp(-w_i/\theta)}{\theta^{1/\theta} \Gamma(1/\theta)}. \tag{2.2}$$

In (2.1), the frailty w_i acts multiplicatively on the hazard rate and is gamma-distributed. The model formulation of Bellamy et al. (2004) expresses the frailty term as $\exp b_i$, with b_i the random effect working additively on the log hazard rate, and assumes a normal

distribution for the random effect (see Section 3 for a comparison of these two models). In what follows, we write θ for the variance of the frailty w_i and σ^2 for the variance of the random effect b_i.

The udder quarter infection times in the mastitis study are either right-censored or interval-censored. Cluster i consists of $n_i = 4$ observations (one observation per udder quarter), of which r_i are right-censored and d_i are interval-censored. We let R_{ij} denote the right-censored infection time for udder quarter j of cow i. If the information on the infection time is subject to interval censoring, then we denote the lower and upper bounds of the interval as L_{ij} and U_{ij}. For each cluster, we define two sets of indexes according to whether the infection time is right-censored or interval-censored:

$$R_i = \{j \in \{1,2,3,4\} : T_{ij} > R_{ij}\},$$
$$D_i = \{j \in \{1,2,3,4\} : L_{ij} < T_{ij} \le U_{ij}\},$$

where $R_i \cap D_i = \emptyset$, $R_i \cup D_i = \{1,2,3,4\}$, and T_{ij} is the unobservable infection time.

Assuming that the censoring process is not informative for the survival process, the conditional data likelihood contribution for cluster i consists of the product of differences of the conditional survival functions evaluated at the observed lower and upper time point for the interval-censored quarters and of the conditional survival function evaluated at the censoring time for the right-censored quarters,

$$L_i(\theta, \underset{\sim}{\boldsymbol{\xi}}, \underset{\sim}{\boldsymbol{\beta}} \mid w_i) = \prod_{j \in R_i} S_{ij}(R_{ij}) \prod_{j \in D_i} \{S_{ij}(L_{ij}) - S_{ij}(U_{ij})\}, \tag{2.3}$$

which results in

$$\begin{aligned} L_i(\theta, \underset{\sim}{\boldsymbol{\xi}}, \underset{\sim}{\boldsymbol{\beta}} \mid w_i) &= \exp\left\{-\sum_{j \in R_i} H_{ij}(R_{ij})\right\} \\ &\quad \times \prod_{j \in D_i} \left[\exp\{-H_{ij}(L_{ij})\} - \exp\{-H_{ij}(U_{ij})\}\right] \\ &= \exp(-w_i C_i) \times \prod_{j \in D_i} \{\exp(-w_i L_{ij}^*) - \exp(-w_i U_{ij}^*)\}, \end{aligned} \tag{2.4}$$

where $\underset{\sim}{\boldsymbol{\xi}}$ contains the parameters of the baseline hazard, $H_{ij}(\cdot) = H_0(\cdot) w_i \exp(\underset{\sim}{\mathbf{x}}_{ij}^t \underset{\sim}{\boldsymbol{\beta}})$, $C_i = \sum_{j \in R_i} H_0(R_{ij}) \exp(\underset{\sim}{\mathbf{x}}_{ij}^t \underset{\sim}{\boldsymbol{\beta}})$, $L_{ij}^* = H_0(L_{ij}) \exp(\underset{\sim}{\mathbf{x}}_{ij}^t \underset{\sim}{\boldsymbol{\beta}})$, and $U_{ij}^* = H_0(U_{ij}) \exp(\underset{\sim}{\mathbf{x}}_{ij}^t \underset{\sim}{\boldsymbol{\beta}})$, with $H_0(\cdot)$ the cumulative baseline hazard.

To write the product in the second factor of (2.4) in a general way, we define the following column vector $\underset{\sim}{\mathbf{a}}_i$ of length 2^{d_i}, where d_i is the number of elements in D_i:

$$\underset{\sim}{\mathbf{a}}_i = ({}_c a_{ik})_{k=1}^{2^{d_i}} = \bigotimes_{j \in D_i} \begin{pmatrix} \exp(-w_i L_{ij}^*) \\ -\exp(-w_i U_{ij}^*) \end{pmatrix},$$

where $\bigotimes_{j \in D_i}$ represents the Kronecker product of the vectors $(\exp(-w_i L_{ij}^*), -\exp(-w_i U_{ij}^*))^t$, $j \in D_i$. For example, the first element of this column vector is $\exp(-w_i \sum_{j \in D_i} L_{ij}^*)$. The last element is $\pm \exp(-w_i \sum_{j \in D_i} U_{ij}^*)$ with a positive sign if

the number of U_{ij}^*'s in the sum of the exponent is even and a negative sign if the number is odd. The number of U_{ij}^*'s in a_{ik} is denoted as n_{ik}.

Expression (2.4) then can be rewritten as

$$L_i(\theta, \underset{\sim}{\boldsymbol{\xi}}, \underset{\sim}{\boldsymbol{\beta}} \mid w_i) = \exp(-w_i C_i)\left(\sum_{k=1}^{2^{d_i}} a_{ik}\right).$$

This expression still contains the unobserved frailty term w_i. However, we can integrate out the frailty term, which is assumed to have the gamma density (2.2). We then obtain the marginal likelihood

$$\begin{aligned} L_i(\theta, \underset{\sim}{\boldsymbol{\xi}}, \underset{\sim}{\boldsymbol{\beta}}) &= \int_0^\infty \frac{w_i^{(1/\theta-1)} \exp(-w_i/\theta)}{\theta^{1/\theta}\Gamma(1/\theta)} \exp(-w_i C_i)\left(\sum_{k=1}^{2^{d_i}} a_{ik}\right) dw_i \\ &= \frac{1}{\theta^{1/\theta}\Gamma(1/\theta)} \sum_{k=1}^{2^{d_i}} \int_0^\infty w_i^{(1/\theta-1)} \exp\left\{-w_i\left(C_i + \frac{1}{\theta}\right)\right\} a_{ik}\, dw_i \\ &= \sum_{k=1}^{2^{d_i}} \frac{(-1)^{n_{ik}}}{(C_i + 1/\theta + \log p_{ik})^{1/\theta}\theta^{1/\theta}}, \end{aligned} \tag{2.5}$$

where $\underset{\sim}{\mathbf{p}}_i$ is the column vector,

$$\underset{\sim}{\mathbf{p}}_i = ({}_c p_{ik})_{k=1}^{2^{d_i}} = \bigotimes_{j \in D_i} \begin{pmatrix} \exp(L_{ij}^*) \\ \exp(U_{ij}^*) \end{pmatrix}.$$

To obtain the full marginal likelihood (L), we take the product of the N cluster-specific marginal likelihoods $\prod_{i=1}^N L_i(\theta, \underset{\sim}{\boldsymbol{\xi}}, \underset{\sim}{\boldsymbol{\beta}})$. We can then obtain maximum likelihood estimates by maximizing the full marginal likelihood using, for instance, the Newton–Raphson procedure. Because the second partial derivatives can be obtained for all parameters in the model (see the Appendix), an explicit expression for the information matrix is available, from which an estimate of the asymptotic variance–covariance matrix can be obtained. Various distributional assumptions for the baseline hazard are possible, as we discuss in the next section.

3. EXAMPLE

Mastitis control is an important component of dairy herd health programs, because udder infections are closely associated with reduced milk yield and impaired milk quality (Seegers, Fourichon, and Beaudeau 2003). A total of 100 cows were screened monthly at the udder-quarter level for bacterial infections from the time of parturition, at which the cow was included in the cohort and assumed to be infection-free, until the end of the lactation period. Observations can be right-censored if no infection has occurred before the end of the lactation period, which is roughly 300–350 days (but different for every cow), or if the cow is lost to follow-up during the study, due to culling, for example. Because of the periodic follow-up, udder quarters that experience an event are interval-censored, with

the lower bound the last visit with a negative test and the upper bound the first visit with a positive test. Although farms are visited at regular time intervals, the time between two successive visits is not always exactly 1 month, and the recorded visiting times for each cow differ depending on when they were included in the cohort. Moreover, because of a lack of personnel in summer months, cows were screened only in July or August, meaning that at least one interval spans two months.

Two types of covariates are considered in the analysis. Cow-level covariates take the same value for every udder quarter of the cow (e.g., number of calvings or parity). Several studies have shown that both prevalence and incidence of intramammary infections increase with parity (Vecht, Wisselink, and Defize 1989; Weller, Saran, and Zeliger 1992). Several hypotheses have been suggested to explain these findings; for example, the teat end condition deteriorates with increasing parity (Neijenhuis et al. 2001). Because the teat end is a physical barrier that prevents organisms from invading the udder, impaired teat ends make the udder more vulnerable to intramammary infections. We consider the following categorical parity covariate: (a) primiparous cows (one calving, parity = 0), (b) cows with between two and four calvings (parity = 1), and (c) cows with more than four calvings (parity = 2). We have to categorize because for some of the parity levels, only a small number of cows is available. Udder quarter–level covariates change within the cow (e.g., position of the udder quarter [front or rear]). The difference in teats end condition between front and rear quarters also has been put forward to explain the difference in infection status (Adkinson et al. 1993; Barkema et al. 1997; Schepers et al. 1997). In total, 317 of 400 udder quarters were infected during the lactation period.

Various distributional assumptions for the baseline hazard are possible (Klein and Moeschberger 2003). Because of its mathematical simplicity, we first look at the exponential distribution with constant hazard function $h_0(t) = \lambda$. Because a constant hazard rate is probably not realistic for describing the time to intramammary infection, we consider the Weibull distribution with hazard function $h_0(t) = \lambda\gamma t^{\gamma-1}$. The hazard is monotone decreasing for $\gamma < 1$ and monotone increasing for $\gamma > 1$. We further consider another two-parameter distribution, the log-logistic distribution with hazard function $h_0(t) = \lambda\gamma t^{\gamma-1}/(1 + \lambda t^{\gamma})$. The numerator is the same as in the Weibull hazard, but the denominator ensures that the hazard is monotone decreasing for $\gamma \leq 1$. For $\gamma > 1$, the hazard initially increases to a maximum at time $\{(\gamma - 1)/\lambda\}^{1/\gamma}$, and then decreases to zero as time approaches infinity. The log-likelihood values are -786.637 for the model with an exponential distribution and -730.058 for the model with a Weibull distribution. Thus, the Weibull distribution is definitely preferable (likelihood ratio test; $p < 0.001$).

The comparison of the models with Weibull and log-logistic baseline hazards is based on the Akaike information criterion (AIC), because these models are not nested within each other. To calculate the AIC, we use the standard formula AIC $= -2\log L + 2 \times$ (number of parameters) (Izumi and Ohtaki 2004). The log-logistic and Weibull distributions have similar AIC values (1469.716 versus 1472.116). We proceed with the Weibull distribution, because this is the most standard choice. To assess the validity of the gamma frailty model with the Weibull baseline hazard, we compare the marginal survival functions obtained by the nonparametric estimator for interval-censored data proposed by Turnbull

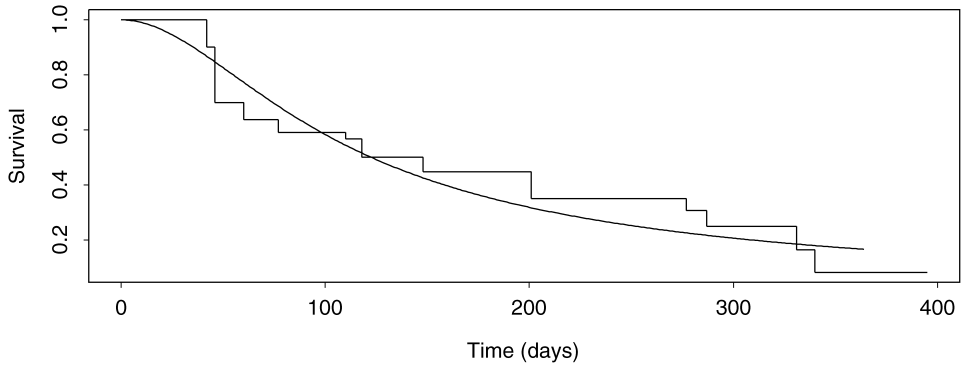

Figure 1. Estimated marginal survival functions for the front udder quarters of primiparous cows. The stepwise function corresponds to the nonparametric estimate for interval-censored data (Turnbull 1976), whereas the continuous curve is obtained from the gamma frailty model with a Weibull baseline hazard.

(1976) with the marginal survival function obtained from the gamma frailty model with a Weibull baseline hazard. For example, for the front udder quarters of primiparous cows, the marginal survival function from the gamma frailty model closely follows the nonparametric estimate (Figure 1). Similar findings are obtained for other covariate-level combinations.

The parameter estimates obtained from the method proposed in the previous section with the Weibull baseline hazard are given in the first column of Table 1. These parameter estimates are obtained using infection times in terms of quarters of the year rather than days, to ensure sufficiently large values for estimating λ; however, for the figures and the interpretation in the text, we rescale to days to make the interpretation easier.

The parameter estimate $\hat{\gamma}$ is >1, indicating that the hazard is increasing with time. The hazard ratio of cows with more than four calvings versus primiparous cows is 4.06 ($p = 0.004$), with 95% confidence interval [1.56; 10.51]. The hazard ratio of cows with two to four calvings versus primiparous cows is 0.82, with 95% confidence interval [0.42; 1.58]. This confirms the findings of Neijenhuis et al. (2001) for the older cows with parity

Table 1. Parameter estimates (Est) and standard errors (SE) for a gamma frailty model with parity (with $\hat{\beta}_{p_1}$ the effect of parity category 2 and $\hat{\beta}_{p_2}$ the effect of parity category 3) and udder location (with $\hat{\beta}_r$ the effect of the rear udder quarter) as the covariate and Weibull baseline hazard.

	Exact Est (SE)	Midpoint Est (SE)	Upper bound Est (SE)	Gaussian quadrature Est (SE)
$\hat{\theta}$	1.600 (0.279)	1.606 (0.276)	1.708 (0.289)	–
$\hat{\sigma}^2$	–	–	–	2.517 (0.534)
$\hat{\lambda}$	0.721 (0.185)	0.749 (0.191)	0.481 (0.125)	0.244 (0.070)
$\hat{\gamma}$	1.936 (0.109)	2.003 (0.107)	2.417 (0.129)	2.015 (0.115)
$\hat{\beta}_r$	0.180 (0.122)	0.176 (0.119)	0.184 (0.119)	0.174 (0.123)
$\hat{\beta}_{p1}$	−0.201 (0.336)	−0.238 (0.335)	−0.252 (0.345)	0.037 (0.373)
$\hat{\beta}_{p2}$	1.400 (0.486)	1.391 (0.464)	1.181 (0.464)	1.878 (0.540)

NOTE: Results are shown for the proposed method (exact), midpoint, and upper bound of the interval as exact event times and Gaussian quadrature.

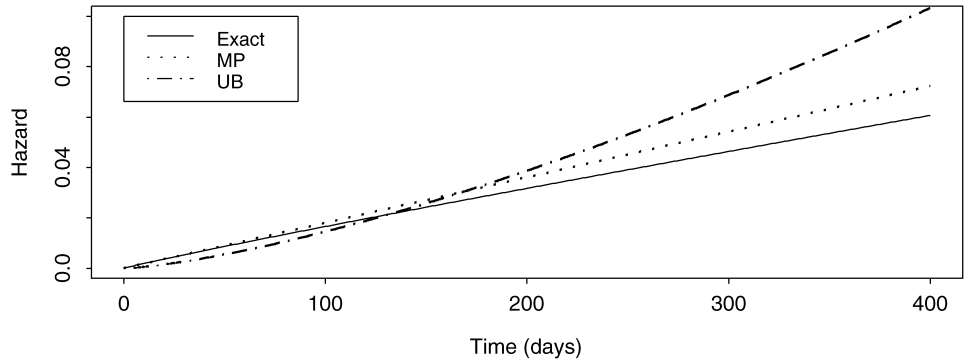

Figure 2. Estimated hazard functions from the gamma frailty model for a front udder quarter of a primiparous cow with frailty equal to 1, either taking into account the interval censoring (Exact) or imputing the midpoint (MP) or the upper bound (UB).

exceeding four, but not for the cows of parity between two and four. The hazard ratio of rear versus front udder quarters is 1.20, with 95% confidence interval [0.94; 1.52]. The estimated θ is 1.6. Kendall's tau can be obtained from this estimate. Kendall's tau corresponds to $\mathrm{E}[\mathrm{sign}\{(T_{ik} - T_{jk})(T_{ik'} - T_{jk'})\}]$, where $k \neq k'$ and T_{ik} is the infection time of the kth udder quarter of cow i (Oakes 1989). Thus, infection times within the cow are highly correlated.

We now compare our proposed method with the naive method of imputing the midpoint or the upper bound of the interval as an exact event time. Although using the midpoint gives similar results (with only slightly higher parameter estimates of θ, λ, and γ; see Table 1), imputation of the upper bound has a significant effect on the parameter estimates $\hat{\theta}$, $\hat{\lambda}$, and $\hat{\gamma}$. The overestimation of γ is particularly eye-catching and leads to a more rapidly increasing hazard compared with that obtained using imputation of the midpoint or the exact method based on interval-censored data. This is clearly shown in Figure 2, which depicts the estimated hazard functions for the three models for a front udder quarter of a primiparous cow with frailty equal to 1. The choice of the upper bound as exact event time ensures that no events occur within the first 30 days after calving; therefore, the model based on imputation of the upper bound leads to a more rapidly increasing hazard function, to account for the fact that the hazard rate should be as low as possible during the first 30 days.

It is interesting to compare our results with the estimates obtained from the method proposed by Bellamy et al. (2004). We consider the following model:

$$h_{ij}(t) = h_0(t)\exp(b_i)\exp(\underset{\sim}{\mathbf{x}}_{ij}^{t}\underset{\sim}{\boldsymbol{\beta}}), \qquad i = 1, \ldots, N,\ j = 1, \ldots, n_i, \tag{3.1}$$

where $h_0(t) = \lambda\gamma t^{\gamma-1}$ and $b_i \sim \mathrm{N}(0, \sigma^2)$.

Because a normal distribution is assumed for the random effects b_i, the frailties follow a lognormal distribution, and it is no longer possible to obtain a closed-form expression for the marginal likelihood by integrating out the frailties exactly. Consequently, Bellamy et al. (2004) used Gaussian quadrature to integrate out the frailties, and then maximized the marginal likelihood. Unlike. We use the proportional hazards (PH) model representa-

tion instead of the accelerated failure time (AFT) model representation used by Bellamy et al. (2004). With β_r representing the effect of the rear udder quarter, β_{p_1} the effect of the second parity category, and β_{p_2} the effect of the third parity category, the parameter estimates correspond to $\hat{\sigma}^2 = 2.517$, $\hat{\lambda} = 0.244$, $\hat{\gamma} = 2.015$, $\hat{\beta}_r = 0.174$, $\hat{\beta}_{p_1} = 0.037$, and $\hat{\beta}_{p_2} = 1.878$.

Comparing the gamma frailty model (2.1) with model (3.1) with normally distributed random effects is not straightforward. The frailties corresponding to the random effects with mean 0 in the last model do not have mean 1. In this particular case, the mean is estimated by $\exp(0.5\hat{\sigma}^2) = 3.52$. Therefore, it is good practice to compare the two models in terms of a medically relevant quantity, such as the median time to infection (M). Both the random-effects model and the frailty models induce heterogeneity in the median time to infection between cows. The density function for the median time to infection in the frailty model is given by

$$f_M(m) = \gamma \left(\frac{\log 2}{\theta \lambda \exp(\underset{\sim}{\mathbf{x}}_{ij}^t \underset{\sim}{\boldsymbol{\beta}})} \right)^{1/\theta} \frac{1}{\Gamma(1/\theta)} \left(\frac{1}{m} \right)^{1+\gamma/\theta} \exp\left(-\frac{\log 2}{\theta \lambda m^\gamma \exp(\underset{\sim}{\mathbf{x}}_{ij}^t \underset{\sim}{\boldsymbol{\beta}})} \right). \quad (3.2)$$

In the case of normally distributed random effects, the density function for M is (Legrand et al. 2005)

$$f_M(m) = \frac{\alpha}{m\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2\sigma^2} \left\{ \log\left(\frac{\log 2}{m\lambda \exp(\underset{\sim}{\mathbf{x}}_{ij}^t \underset{\sim}{\boldsymbol{\beta}})} \right) \right\}^2 \right]. \quad (3.3)$$

Although the two density functions for a front udder quarter of a primiparous cow appear rather similar (Figure 3), compared with the model with normally distributed random effects, the gamma frailty model assumes that the median times to infection are less skewed to the right and thus have a somewhat higher peak at the more central median time to infection values.

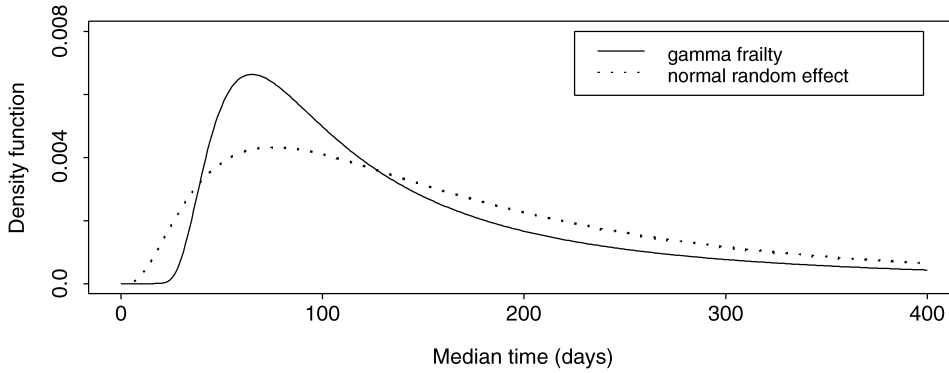

Figure 3. Density functions for the median time to infection from the frailty model with gamma- and lognormal-distributed frailty for the front udder quarter of a primiparous cow.

4. SIMULATION STUDY

We conducted a simulation study to evaluate the performance of our proposed methodology. Using the same data structure as in the example data set, we generated 1000 data sets consisting of 100 clusters each, with four observations per cluster. We generated the frailties (w_i) from the one-parameter gamma density (2.2), with $E(w_i) = 1$ and $\text{var}(w_i) = \theta = 1.8$. The data were simulated from the frailty model (2.1) assuming Weibull-distributed event times with scale parameter $\lambda = 0.9$, shape parameter $\gamma = 1.9$, and $\beta = 0.2$. The true values for the simulation study correspond to the parameter estimates obtained by fitting the proposed model to the mastitis data with a single covariate, the udder location parameter. A binary covariate x takes the value 1 for the first two observations in a cluster and 0 for the other two observations in the cluster.

Asynchronous intervals of 30 days are generated around the simulated infection time as follows. The number of days in which a particular cow was in lactation before the first visit is simulated from the uniform distribution with a minimum of 1 and a maximum of 29. Visits are conducted at fixed time points (e.g., 0–30–60–90), but because each cow entered the study at a different moment (i.e., its first day of lactation), the endpoints of the intervals are adjusted to the number of days in lactation, so they can take any arbitrary value. All cows are assumed to be infection-free at the start of their lactation period, so if an udder quarter is already infected at the first visit, then the infection is assumed to occur between the start of the lactation period and the first visit. The study endpoint was set at 1 year; thus udder quarters with simulated infection times exceeding 1 year are right-censored. The upper bound of the last interval is used as the censoring time.

For each of the 1000 data sets, we fit three models: the model proposed in Section 2 using the interval-censored data and the two naive models that ignore interval censoring and impute the midpoint or upper bound as the exact event time. We compared the mean of the 1000 obtained estimates for the parameters θ, λ, γ, and β with the true value, and investigate the differences among the three models. Standard errors are obtained by taking the inverse of the Hessian matrix at the end of the optimization procedure. We also calculate the empirical standard error obtained from the 1000 data sets. Finally, we determine the coverage, defined as the percentage of the 1000 data sets containing the true population parameter within their 95% confidence interval. The results of our simulation study suggest that the estimates obtained with our proposed model and by imputation of the midpoint are close to the true population parameters of interest (see Table 2). But for the upper bound imputation, the estimate of λ is biased downward, and the estimate of γ is biased upward. The coverage is good if the interval-censored nature of the data is taken into account or if imputation of the midpoint is used. As can be expected, given the large bias for λ and γ, coverages for these parameters are unacceptable when the upper bound imputation is used.

Based on these simulations, it might seem that our new technique provides no advantage over imputation of the midpoint. This is not always the case, however. Consider, for instance, the same simulation setting as before ($\lambda = 0.9$, $\beta = 0.2$, $\theta = 1.8$, 30-day intervals), but with the value of the parameter γ changed to 0.5. Changing the value of the parameter γ from 1.9 to 0.5 means that the hazard is no longer increasing, but is decreasing over time.

Table 2. The averages of estimated model parameters (Est), the empirical standard error (Emp SE), and coverage (Cov) from 1000 simulated data sets using the proposed method (Exact), midpoint, and upper bound of the interval as exact event times.

	Exact Est (Emp SE, Cov)	Midpoint Est (Emp SE, Cov)	Upper bound Est (Emp SE, Cov)
		$\gamma = 1.9$	
$\hat{\theta}$	1.804 (0.285, 94.3)	1.760 (0.276, 93.1)	1.882 (0.286, 96.3)
$\hat{\lambda}$	0.896 (0.156, 93.6)	0.881 (0.151, 92.4)	0.596 (0.101, 21.5)
$\hat{\gamma}$	1.900 (0.104, 93.7)	1.887 (0.095, 93.9)	2.267 (0.115, 6.3)
$\hat{\beta}$	0.202 (0.124, 93.6)	0.198 (0.121, 94.1)	0.206 (0.126, 93.4)
		$\gamma = 0.5$	
$\hat{\theta}$	1.842 (0.359, 92.7)	2.141 (0.397, 89.7)	2.156 (0.398, 89.3)
$\hat{\lambda}$	0.917 (0.170, 94.5)	0.878 (0.181, 90.0)	0.709 (0.139, 62.1)
$\hat{\gamma}$	0.492 (0.039, 94.2)	0.735 (0.045, 0.0)	0.881 (0.056, 0.0)
$\hat{\beta}$	0.199 (0.145, 95.3)	0.212 (0.156, 93.4)	0.211 (0.156, 93.5)

NOTE: True values for the parameters are given by $\lambda = 0.9$, $\beta = 0.2$, $\theta = 1.8$, and $\gamma = 1.9$ in the first part of the table and 0.5 in the second part.

At the start of the study, all udder quarters are at risk; in the case of an increasing hazard, few events occur in the beginning, and many udder quarters remain at risk toward the end of the study, when more events occur. Therefore, much information is available throughout the study period for obtaining parameter estimates. For a decreasing hazard ($\gamma < 1$), many events occur in the beginning, leaving only few udder quarters at risk near the end of the study. Thus when the interval-censored nature of the data is ignored in this setting, insufficient information remains for obtaining adequate parameter estimates. As can be seen in Table 2, the exact method performs well for estimating all parameters, including γ, but the techniques of imputing the midpoint or upper bound both fail to estimate γ, and imputation also performs worse than the exact method in estimating the other parameters.

We also investigated situations with broader censoring intervals. We generated intervals of width 60, 90, and 120 days around the simulated infection times in the same manner as before. Figure 4 shows what happens with the estimates of the four model parameters when intervals become broader for the exact technique (empty box) and the midpoint imputation method (shaded box). The dashed horizontal line represents the true value of the parameter. The box represents the interquartile range, and the solid line represents the median of the 1000 simulated data sets. The whiskers are drawn to the nearest value not beyond 1.5 times the interquartile range. It can be seen that the exact method performs better than the midpoint imputation approach in estimating the parameters θ, λ, and β. The latter method tends to underestimate these parameters. The bias is larger when intervals are broader. When the exact method is used, coverages are good, but with the imputation technique, coverages decrease as intervals become broader. For example, when the interval spans 120 days, the coverage for λ is 59.8% with midpoint imputation, versus 95.2% with the exact technique. However, the parameter γ is underestimated by the exact technique and overestimated by midpoint imputation. For both techniques, coverages decrease as intervals become broader. The results of our simulation studies clearly show that the exact method outperforms the two imputation-based methods.

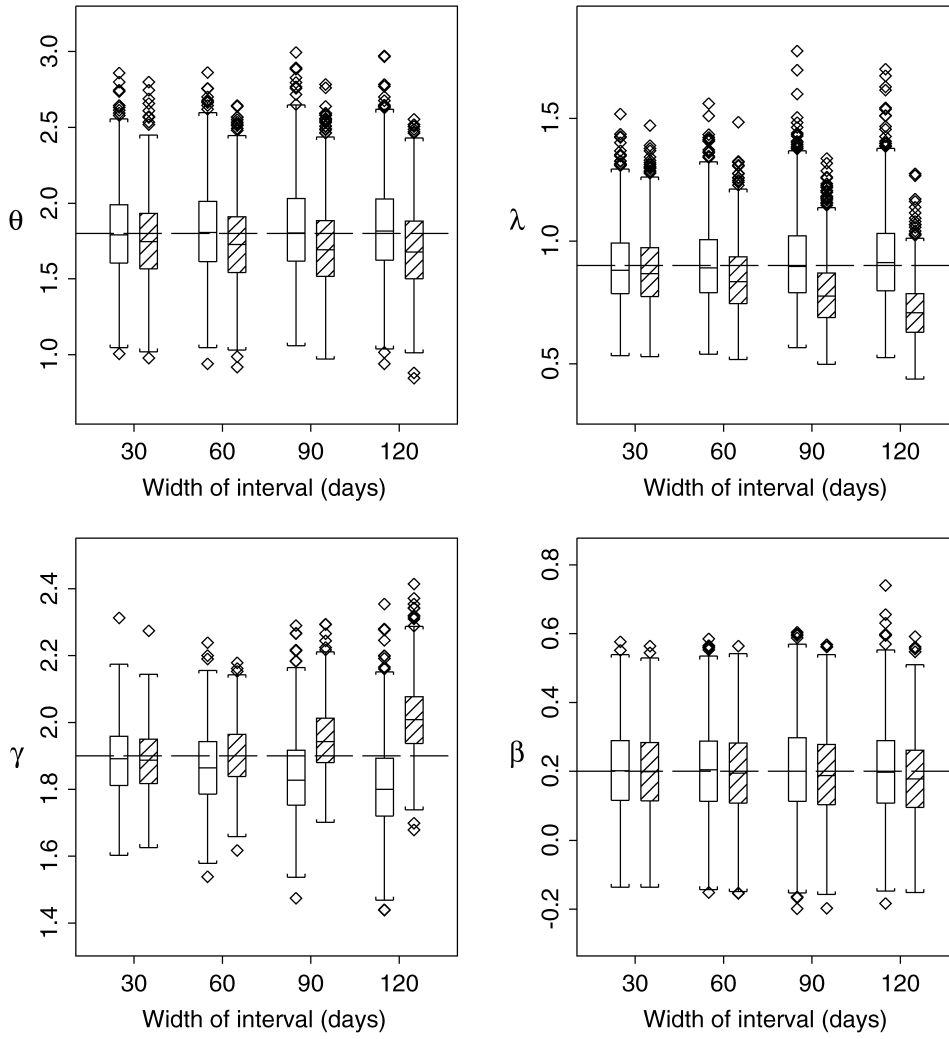

Figure 4. Boxplots for the estimated model parameters from 1000 simulated data sets using the proposed method (empty box) and the midpoint of the interval as exact event time (shaded box) for different interval widths (30–60–90–120). True values for the parameters are given by $\lambda = 0.9$, $\beta = 0.2$, $\theta = 1.8$, and $\gamma = 1.9$.

5. CONCLUSIONS

In this article we have proposed a shared gamma frailty model for clustered, interval-censored data with varying baseline hazard functions. The model could be made even more flexible by using penalized splines for the baseline hazard using the same techniques (Rondeau, Commenges, and Joly 2003). Assuming a gamma distribution for the frailty allows us to integrate out the frailties analytically and obtain a closed-form expression for the marginal likelihood, which can then be maximized with an optimization procedure, such as the nlm function in R to obtain parameter estimates (R code is available as an online supplement to this article). Furthermore, exact expressions for the second derivatives of

the likelihood, and thus estimates of the parameter variances, can be obtained by inverting the matrix of second derivatives.

In the example, we compare our technique with that of Bellamy et al. (2004), who proposed normally distributed random effects. Under this assumption, a closed form of the likelihood cannot be obtained, and frailties are integrated out using Gaussian quadrature. Some of the parameters appearing in the two models have the same meaning (γ and β) and thus can be compared. Linking the λ parameter in the proposed model to parameters in the model of Bellamy et al. (2004) is more difficult. This is due to the specification of the cluster effects in terms of normally distributed random effects with mean 0 (i.e., a lognormal distribution at the frailty level, but with a mean different from 1). Thus comparing these two hazard functions is not straightforward, and we can expect to find significant differences between them for large values of the variance of the lognormal and gamma distributions (Therneau and Grambsch 2000). Indeed, the larger the value of the variance parameter, the more different the densities. But the two models can be meaningfully compared when the models are translated in terms of the density function of the median infection time. The two models result in comparable density functions.

We applied our proposed method to a data set consisting of 100 cows screened approximately monthly for the presence of a bacterial infection at the udder-quarter level. This method also can be used in a wide variety of other applications, because it involves few or no data constraints. The data set must consist of a number of clusters from which the members cannot be observed continuously. Intervals can be of variable length, although the parameter γ tends to be increasingly biased as intervals become broader. In our first simulation setting, accurate estimates were obtained using the proposed technique and imputation of the midpoint. But in the second simulation setting, where the parameter γ was <1, imputation of the midpoint failed. Using the upper bound as an exact event time led to biased estimates especially for λ and γ, in both simulation settings. Also, as censoring intervals grew broader, using the midpoint imputation technique led to biased estimates for all parameters. Overall, our simulation studies demonstrate that our technique outperforms the imputation techniques.

APPENDIX: INFORMATION MATRIX

Here we use the second partial derivative for β in case of one covariate of the log-likelihood for clustered, interval-censored data as an example. The other partial derivatives can be obtained similarly.

$$\frac{\partial^2 l_i}{\partial \beta^2} = (-1)\left\{\sum_{k=1}^{2^{d_i}} \frac{1}{(C_i + \log p_k + 1/\theta)^{1/\theta}} \cdot s_k\right\}$$
$$- 2 \cdot \left\{\frac{1}{\theta}\sum_{k=1}^{2^{d_i}} \frac{1}{(C_i + \log p_k + 1/\theta)^{(1/\theta+1)}} \cdot (C_{i,\beta} + \log p_{k,\beta}) \cdot (-s_k)\right\}^2$$

$$+\left\{\sum_{k=1}^{2^{d_i}}\frac{1}{(C_i+\log p_k+1/\theta)^{1/\theta}}\cdot s_k\right\}^{-1}$$
$$\cdot\left[\frac{1}{\theta}\sum_{k=1}^{2^{d_i}}\left\{\left(\frac{-1}{\theta}-1\right)\frac{1}{(C_i+\log p_k+1/\theta)^{(1/\theta+2)}}\cdot(C_{i,\beta}+\log p_{k,\beta})^2\right.\right.$$
$$\left.\left.+\frac{1}{(C_i+\log p_k+1/\theta)^{(1/\theta+1)}}\cdot(C_{i,\beta\beta}+\log p_{k,\beta\beta})\right\}(-s_k)\right],$$

where $\mathbf{s}$ is the following column vector:

$$\mathbf{s}=({}_c s_k)_{k=1}^{2^{d_i}}=\bigotimes_{j=1}^{d_i}\begin{pmatrix}1\\-1\end{pmatrix},$$

$$C_{i,\beta}=\sum_{j\in R_i}H_{w_i}(R_{ij})x_{ij}\exp(x_{ij}\beta),$$

$$\mathbf{p}_{k,\beta}=\bigotimes_{j=1}^{d_i}\begin{pmatrix}x_{ij}\exp(L_{ij}^*)\\x_{ij}\exp(U_{ij}^*)\end{pmatrix},$$

$$C_{i,\beta\beta}=\sum_{j\in R_i}H_{w_i}(R_{ij})(x_{ij})^2\exp(x_{ij}\beta),$$

$$\mathbf{p}_{k,\beta\beta}=\bigotimes_{j=1}^{d_i}\begin{pmatrix}(x_{ij})^2\exp(L_{ij}^*)\\(x_{ij})^2\exp(U_{ij}^*)\end{pmatrix}.$$

To obtain the second partial derivative for the full marginal log-likelihood, we need to sum over the N cluster-specific second partial derivatives.

ACKNOWLEDGMENTS

The work of Goethals was supported by a grant from the Research Fund of Ghent University (B/04953/01). The work of Duchateau and Janssen was supported by the Interuniversity Attraction Poles research network P6/03 of the Belgian Government (Belgian Science Policy).

[Received January 2007. Revised February 2008.]

REFERENCES

Adkinson, R. W., Ingawa, K. H., Blouin, D. C., and Nickerson, S. C. (1993), "Distribution of Clinical Mastitis Among Quarters of the Bovine Udder," *Journal of Dairy Science*, 76, 3453–3459.

Barkema, H. W., Schukken, Y. H., Lam, T. J. G. M., Galligan, D. T., Beiboer, M. L., and Brand, A. (1997), "Estimation of Interdependence Among Quarters of the Bovine Udder With Subclinical Mastitis and Implications for Analysis," *Journal of Dairy Science*, 80, 1592–1599.

Bellamy, S. L., Li, Y., Ryan, L. M., Lipsitz, S., Canner, M. J., and Wright, R. (2004), "Analysis of Clustered and Interval Censored Data From a Community-Based Study in Asthma," *Statistics in Medicine*, 23, 3607–3621.

Collet, D. (1994), *Modelling Survival Data in Medical Research*, London: Chapman & Hall.

Finkelstein, D. M. (1986), "A Proportional Hazards Model for Interval-Censored Failure Time Data," *Biometrics*, 42, 845–854.

Izumi, S., and Ohtaki, N. (2004), "Aspects of the Armitage–Doll Gamma Frailty Model for Cancer Incidence Data," *Environmetrics*, 15, 209–218.

Klein, J. P., and Moeschberger, M. L. (2003), *Survival Analysis: Techniques for Censored and Truncated Data*, New York: Springer.

Laevens, H., Deluyker, H., Schukken, Y. H., De Meulemeester, L., Vandermeersch, R., De Muêlenaere, E., and De Kruif, A. (1997), "Influence of Parity and Stage of Lactation on the Somatic Cell Count in Bacteriologically Negative Dairy Cows," *Journal of Dairy Science*, 80, 3219–3226.

Legrand, C., Ducrocq, V., Janssen, P., Sylvester, R., and Duchateau, L. (2005), "A Bayesian Approach to Jointly Estimate Centre and Treatment by Centre Heterogeneity in a Proportional Hazards Model," *Statistics in Medicine*, 24, 3789–3804.

Neijenhuis, F., Barkema, H. W., Hogeveen, H., and Noordhuizen, J. P. T. M. (2001), "Relationship Between Teat-End Callosity and Occurrence of Clinical Mastitis," *Journal of Dairy Science*, 84, 2664–2672.

Oakes, D. (1989), "Bivariate Survival Models Induced by Frailties," *Journal of the American Statistical Association*, 84, 487–493.

Radke, B. R. (2003), "A Demonstration of Interval-Censored Survival Analysis," *Preventive Veterinary Medicine*, 59, 241–256.

Rondeau, V., Commenges, D., and Joly, P. (2003), "Maximum Penalized Likelihood Estimation in a Gamma-Frailty Model," *Lifetime Data Analysis*, 9, 139–153.

Schepers, A. J., Lam, T. J. G. M., Schukken, Y. H., Wilmink, J. B. M., and Hanekamp, W. J. A. (1997), "Estimation of Variance Components for Somatic Cell Counts to Determine Thresholds for Uninfected Quarters," *Journal of Dairy Science*, 80, 1833–1840.

Seegers, H., Fourichon, C., and Beaudeau, F. (2003), "Production Effects Related to Mastitis and Mastitis Economics in Dairy Cattle Herds," *Veterinary Research*, 34, 475–491.

Therneau, T. M., and Grambsch, P. M. (2000), *Modeling Survival Data: Extending the Cox Model*, New York: Springer.

Turnbull, B. W. (1976), "Empirical Distribution Function With Arbitrarily Grouped, Censored and Truncated Data," *Journal of the Royal Statistical Society*, Ser. B, 38, 290–295.

Vecht, U., Wisselink, H. J., and Defize, P. R. (1989), "Dutch National Mastitis Survey—the Effect of Herd and Animal Factors on Somatic-Cell Count," *Netherlands Milk and Dairy Journal*, 43, 425–435.

Weller, J. I., Saran, A., and Zeliger, Y. (1992), "Genetic and Environmental Relationships Among Somatic-Cell Count, Bacterial Infection, and Clinical Mastitis," *Journal of Dairy Science*, 75, 2532–2540.