I N S T I T U T D E S T A T I S T I Q U E
B I O S T A T I S T I Q U E E T
S C I E N C E S A C T U A R I E L L E S
(I S B A)

UNIVERSITÉ CATHOLIQUE DE LOUVAIN

D I S C U S S I O N
P A P E R

2013/37

A diagnostic plot for guiding the choice of the frailty distribution in clustered survival data

MUNDA, M. and C. LEGRAND

A diagnostic plot for guiding the choice of the frailty distribution in clustered survival data

Marco Munda* & Catherine Legrand

July 2013

Abstract

Clustered survival data are often analysed using frailty models. The frailty distribution provides a way to model the type of dependence between event times within a cluster. In this paper, we propose a new diagnostic plot to guide the choice of the frailty distribution. Our approach is based on the idea that although frailties are unobservable the dependence structure that they impose on the data can be observed. We estimate the observed association and we compare it with model-based structures obtained from different frailty distributions. To measure association, we use quantile dependence coefficients. The method easily accommodates any cluster sizes, non-ordered observations within clusters, and various censoring schemes.

Keywords: Clustered survival data; Frailty distribution; Diagnostic plot; Quantile dependence

1 Introduction

The Cox model is very popular in the field of survival data analysis, where interest lies in the time to the occurrence of an event such as, for example, disease progression after treatment. The Cox model requires independent (possibly censored) event times given the covariate information. Survival data, however, often display dependence due to clustering. That is, the population under study is divided into groups within which individuals are related in some way. Typical examples of clusters include families, geographical areas, or centres participating in a clinical trial. Shared frailty models (Duchateau & Janssen, 2008) have been proposed to deal with this type of

*marco.munda@uclouvain.be

data. Frailty models account for the within-cluster association via a random effect U, the so-called frailty, that modifies the hazard rate across clusters. Specifically, with $h_{ij}(t)$ the conditional (on the covariate information and on $U = u_i$) hazard rate at time t for member j of cluster i ($j = 1, \ldots, n_i$; $i = 1, \ldots, s$), the (shared) frailty model is an extension of the Cox model defined as

$$h_{ij}(t) = h_0(t) u_i \exp(\boldsymbol{x}'_{ij}\boldsymbol{\beta}) \tag{1}$$

with $h_0(\cdot)$ the baseline hazard function, $\boldsymbol{x}_{ij}$ the vector of covariates, $\boldsymbol{\beta}$ the associated vector of fixed effects parameters, and u_i the actual value taken by U in cluster i. The u_i's are independent and identically distributed with density $f(\cdot)$. To complete the model specification, one needs to postulate a parametric family for $f(\cdot)$.

Frailty models are based on a conditional independence assumption. If the frailties were known, event times within a cluster would be independent (given the covariate information). This means that the frailty term models the dependence in the data. Hougaard (1995) further showed that different choices for $f(\cdot)$ result in different types of dependence. The frailty distribution $f(\cdot)$ therefore needs to be carefully chosen to correctly model the dependence in the data. Because the u_i's are not observed, however, choosing the frailty distribution is a difficult issue.

Very often, a gamma frailty distribution is assumed. Shih & Louis (1995) showed that if this assumption is correct, then the average of the expected frailties at time t, conditional on the data information up to that time, takes constant value 1 for all t. On that basis, they proposed a graphical tool to assess the gamma frailty distribution based on the evolution of the average of the posterior expected frailties over time. Cui & Sun (2004) constructed a formal numerical test based on the largest deviation from 1. Alternatively, Geerdens et al. (2012) developed an order selection test for the null hypothesis that the frailty distribution is gamma against the alternative that it lies in an extended family of distributions constructed by means of certain polynomial expansions forming a nested sequence of new frailty distributions. All these methods are specific to gamma distributed frailties. Few diagnostic procedures accommodate other frailty distributions (like the positive stable). Economou & Caroni (2008) developed a method valid for frailty distributions having the "closure property"[1] (thus excluding, in particular, the positive stable distribution). In dimension 2 and without censoring, a diagnostic plot based on a measure of bivariate dependence, the cross ratio function, has been proposed by Oakes (1989), and later by Viswanathan & Manatunga (2001) with some refinements. This diagnostic plot, however, can only be applied when the clusters have a natural order (for example, in

[1] A frailty distribution is said to have the closure property if the probability density function of the frailty term among clusters with all their members event-free at some time t still belongs to the same family as the initial frailty distribution.

paired organs, observations can be sorted in such a way that observation 1 is taken from the left and observation 2 is taken from the right). Still based on the cross ratio function, Glidden (2007) defined a residual by taking the difference between a non-parametric estimator at some bivariate time point and a model-based estimator obtained via an estimation of the frailty distribution parameter(s). Cumulative (weighted) sums of these residuals over time define a model-checking process which is used to make a diagnostic plot and to construct a supremum-type test statistic whose critical value is determined via bootstrap resampling. This plot accommodates censoring and does not require any special ordering within clusters. An ad-hoc method is further suggested to handle higher-dimensional data by combining the estimates obtained from the bivariate subsamples.

The idea underlying the last three of the above papers is that although frailties are unobservable the dependence structure that they impose on the data can be observed. The problem of choosing the frailty distribution is thus closely related to the problem of having a good way of measuring dependence in survival data. The above papers rely on the cross ratio function. Estimating the dependence coefficient in a non-parametric way is the key ingredient to arrive at a diagnostic measure. Non-parametric estimation of the cross ratio function, however, is a non-trivial task, mainly because it is defined in terms of the hazard function. In this paper, we propose an alternative diagnostic measure. We use quantile dependence coefficients which are defined in terms of the survival function. Non-parametric estimation of survival functions is well developed. Quantile dependence coefficients provide a clear picture of the association pattern imposed by the frailty term. We plot model-based structures and we compare them with the observed (model-free) structure. The former are obtained by fitting the frailty model with different frailty distributions. For the observed structure, we derive a non-parametric estimate that is easily computed from clustered survival data, with any cluster sizes and possibly subject to right-censoring. The resulting plot reveals the best fitting frailty distribution.

In Section 2 we review ways to measure dependence in survival data, and we discuss the association patterns that different frailty distributions put in the data. Quantile dependence coefficients are introduced in Section 3 where we explain how to obtain the model-based and the non-parametric estimates that are used to construct the diagnostic plot. The method is illustrated in Section 4 with two examples, and is assessed in Section 5 via simulations. Concluding remarks are given in Section 6.

2

Association measures in bivariate survival data

Association measures are mainly developed for bivariate data. In this section, we review a number of coefficients that evaluate association between two event times, (T_1, T_2).

2.1 An overall measure

The Kendall's τ is a standard measure of association between two random variables (see, e.g., Hougaard, 2000). In the case of frailty models, the Kendall's τ can be used to give a raw interpretation of the frailty distribution parameter(s) (Duchateau & Janssen, 2008, Chapter 4). To define the Kendall's τ, we need an independent copy of (T_1, T_2). We denote it by (T_1', T_2'). The Kendall's τ is the probability that these two pairs are concordant minus the probability that they are discordant,

$$\tau = \Pr\left[(T_1 - T_1')(T_2 - T_2') > 0\right] - \Pr\left[(T_1 - T_1')(T_2 - T_2') < 0\right]$$

It is a rank-based dependence measure which captures monotonic association. It can range from -1 to 1. Under independence, we have that $\tau = 0$ (but, in general, $\tau = 0$ does not imply independence). It is worth noting here that the position within a pair matters (comparing T_1 to T_1' is not the same as comparing T_1 to T_2'). The pairs must therefore have a natural order.

In Duchateau & Janssen (2008, Section 4.1.4), it is shown that the Kendall's τ can be written in terms of $S(\cdot, \cdot)$, the joint survival function of (T_1, T_2),

$$\tau = 4 \int_0^\infty \int_0^\infty S(t_1, t_2) \frac{\partial^2 S(t_1, t_2)}{\partial t_1 \partial t_2} \, \mathrm{d}t_1 \mathrm{d}t_2 - 1$$

It is further demonstrated that, in the case of frailty models, τ can be expressed in terms of the Laplace transform of U,

$$\tau = 4 \int_0^\infty x \mathcal{L}(x) \mathcal{L}^{(2)}(x) \, \mathrm{d}x - 1$$

with $\mathcal{L}(x) = \mathrm{E}(\exp(-Ux))$, $x \geq 0$, and $\mathcal{L}^{(2)}(\cdot)$ the second derivative of $\mathcal{L}(\cdot)$. In Appendix A we give, for a number of one-parameter frailty distributions, explicit formulas of τ (as a function of the frailty distribution parameter) obtained by substituting the corresponding Laplace transforms into the previous formula.

2.2 A local measure

The Kendall's τ is an overall measure of dependence. It cannot detect changes in the association structure over time. To address this question,

local measures are needed. Local measures of dependence are particularly relevant in the context of frailty models because the composition of the population changes in time. Individuals within "frail" clusters (those with a large value of U) tend to experience the event early while individuals within "strong" clusters are prone to be at risk for a longer time. This selection is to a large extent governed by the frailty distribution, and it will influence the evolution of the dependence. In particular, dependence may predominate at early or at late times (Hougaard, 2000, Section 3.2.1).

The cross ratio function $\zeta(t_1, t_2)$, introduced in Oakes (1989), is a local measure of dependence that compares the hazard of the first member to experience the event at time t_1 given that the partner has failed at time t_2 to the hazard of member 1 at t_1 given that member 2 has survived beyond t_2,

$$\zeta(t_1, t_2) = \frac{h_1(t_1 \,|\, T_2 = t_2)}{h_1(t_1 \,|\, T_2 > t_2)}$$

Unity indicates no association. Oakes (1989) showed that $\zeta(t_1, t_2)$ depends on t_1 and t_2 only through the joint survival function

$$\zeta(t_1, t_2) = \zeta^\star(S(t_1, t_2))$$

where $\zeta^\star(\cdot)$ determines the frailty distribution uniquely, up to a scale factor.

Explicit formulas for $\zeta^\star(\cdot)$ are given in Appendix A. For gamma frailties, $\zeta^\star(\cdot)$ is a constant, while $\zeta^\star(\cdot)$ decreases as $S(t_1, t_2)$ goes to zero both for inverse Gaussian and positive stable frailties.

Oakes (1989) developed a non-parametric estimator of the cross ratio function valid for orderable clusters of size 2 without censoring. This model-free estimator can be used as a diagnostic tool (Duchateau & Janssen, 2008, Section 4.2.6). In particular, substantial deviation from the constant line provides evidence against the gamma frailty distribution.

2.3 **Association patterns**

Another time-dependent association measure, with a simple conditional probability interpretation, is defined as (Anderson et al., 1992)

$$\psi(t_1, t_2) = \frac{S(t_1, t_2)}{S_1(t_1)\, S_2(t_2)} = \frac{\Pr\left[T_1 > t_1 \,|\, T_2 > t_2\right]}{\Pr\left[T_1 > t_1\right]}$$

with $S_1(\cdot)$ (resp. $S_2(\cdot)$) the survival function of T_1 (resp. T_2). Large values of $\psi(\cdot, \cdot)$ indicate positive association while unity indicates no association. In Figure 1, contour plots of $\psi(\cdot, \cdot)$ are depicted to display the evolution of the dependence in a graphical way (Duchateau & Janssen, 2008, Chapter 4). The gamma frailty distribution is characterised by late dependence (contour lines are very close together at late times) while the positive stable frailty distribution is characterised by early dependence (contour lines indicating

positive association appear early). The intuitive explanation relies on the behaviour of these frailty distributions in their tails: the positive stable has a fat right tail (leading to stronger dependence initially) while the gamma has a lot of probability mass skewed to the left (leading to stronger dependence at late times).

3 Quantile dependence

In this section we introduce quantile dependence, an additional association measure that captures the dependence pattern in the tails, and we propose both a model-based estimator (assuming, for example, the gamma or the positive stable frailty distribution) and a non-parametric (model-free) estimator that easily accommodates both censoring and clusters of any sizes with no natural ordering. By comparing the model-free estimates with the model-based estimates, we obtain a way to choose the best fitting frailty distribution.

3.1 Definitions

Consider two event times (T_1, T_2), with joint survival function $S(\cdot,\cdot)$ and marginal distribution (resp. quantile) functions $F_1(\cdot) = 1 - S_1(\cdot)$ (resp. $F_1^{-1}(\cdot)$) and $F_2(\cdot) = 1 - S_2(\cdot)$ (resp. $F_2^{-1}(\cdot)$). Quantile dependence measures local association between T_1 and T_2. Let $q \in (0,1)$. The lower quantile dependence coefficient $\lambda_\ell(q)$ is the conditional probability of one member to have the event before $F_1^{-1}(q)$ given that the event occurred before $F_2^{-1}(q)$ for the partner,

$$\lambda_\ell(q) = \Pr\left[T_1 \leq F_1^{-1}(q) \;\middle|\; T_2 \leq F_2^{-1}(q)\right]$$

When T_1 and T_2 are independent, $\lambda_\ell(\cdot)$ is the identity function. If we let q go to zero, then we obtain the lower tail dependence coefficient (see, e.g., Schmidt & Stadtmüller, 2006),

$$\Lambda_\ell = \lim_{q \to 0} \lambda_\ell(q)$$

The upper counterparts, $\lambda_u(\cdot)$ and Λ_u, are similarly defined,

$$\lambda_u(q) = \Pr\left[T_1 > F_1^{-1}(q) \;\middle|\; T_2 > F_2^{-1}(q)\right] \quad \text{and} \quad \Lambda_u = \lim_{q \to 1} \lambda_u(q)$$

A distribution is said to display lower (resp. upper) tail dependence if $\Lambda_\ell > 0$ (resp. $\Lambda_u > 0$). In contrast, tail independence means that extreme events occur independently in both tails ($\Lambda_\ell = 0$ and $\Lambda_u = 0$).

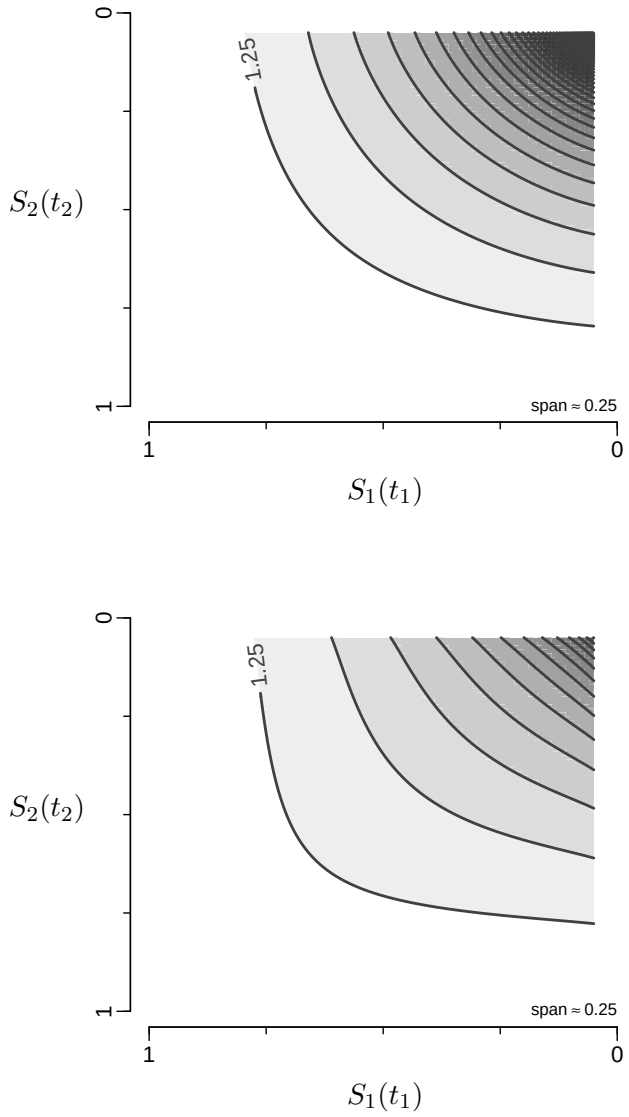

Figure 1 – Contour plots of $\psi(t_1, t_2)$ for gamma frailties (upper panel) and for positive stable frailties (lower panel) with $\tau = 0.4$

In terms of the joint survival function $S(\cdot, \cdot)$, $\lambda_\ell(q)$ and $\lambda_u(q)$ are written as

$$\lambda_\ell(q) = \frac{S\left(F_1^{-1}(q), F_2^{-1}(q)\right) + 2q - 1}{q} \tag{2}$$

and

$$\lambda_u(q) = \frac{S\left(F_1^{-1}(q), F_2^{-1}(q)\right)}{1 - q} \tag{3}$$

3.2 Model-based structures

Within the framework of frailty models, the joint survival function has a Laplace transform representation (Duchateau & Janssen, 2008, Section 4.1.2)

$$S(t_1, t_2) = \mathcal{L}\left\{\mathcal{L}^{-1}\left(S_1(t_1)\right) + \mathcal{L}^{-1}\left(S_2(t_2)\right)\right\}$$

with $\mathcal{L}(\cdot)$ the Laplace transform of the frailty term. By plugging this representation of $S(t_1, t_2)$ into (2) and (3), it follows that, in the case of frailty models, the quantile dependence coefficients can be reexpressed as

$$\lambda_\ell(q) = \frac{\mathcal{L}\left\{2\mathcal{L}^{-1}(1 - q)\right\} + 2q - 1}{q} \tag{4}$$

and

$$\lambda_u(q) = \frac{\mathcal{L}\left\{2\mathcal{L}^{-1}(1 - q)\right\}}{1 - q} \tag{5}$$

In Figure 2, we display several model-based association structures by plotting $\lambda_\ell(q)$ and $\lambda_u(q)$ versus q for gamma, positive stable, and inverse Gaussian frailties. The explicit formulas, obtained by substituting the corresponding Laplace transform into (4) and (5), are given in Appendix A. As expected (cf. Section 2.3), the gamma frailty distribution displays upper tail dependence ($\Lambda_\ell = 0$ and $\Lambda_u > 0$) while the positive stable displays lower tail dependence ($\Lambda_\ell > 0$ and $\Lambda_u = 0$). On the other hand, inverse Gaussian frailties yield independence in both tails ($\Lambda_\ell = 0$ and $\Lambda_u = 0$).

To obtain a model-based estimate $\hat{\lambda}_\ell(q)$ of $\lambda_\ell(q)$, we need to plug in an estimate for the frailty distribution parameter into the Laplace transform. This estimate is obtained by fitting the frailty model. To the best of our knowledge, there is no publicly available software for the semi-parametric (i.e. with the baseline hazard function $h_0(\cdot)$ left unspecified) frailty model, except with the gamma frailty distribution[2]. We therefore follow a parametric approach, using the `parfm` package in R (Munda et al., 2012).

[2]The log-normal frailty distribution is also available in the semi-parametric setting, but we do not consider it here since its Laplace transform does not exist in a closed form.

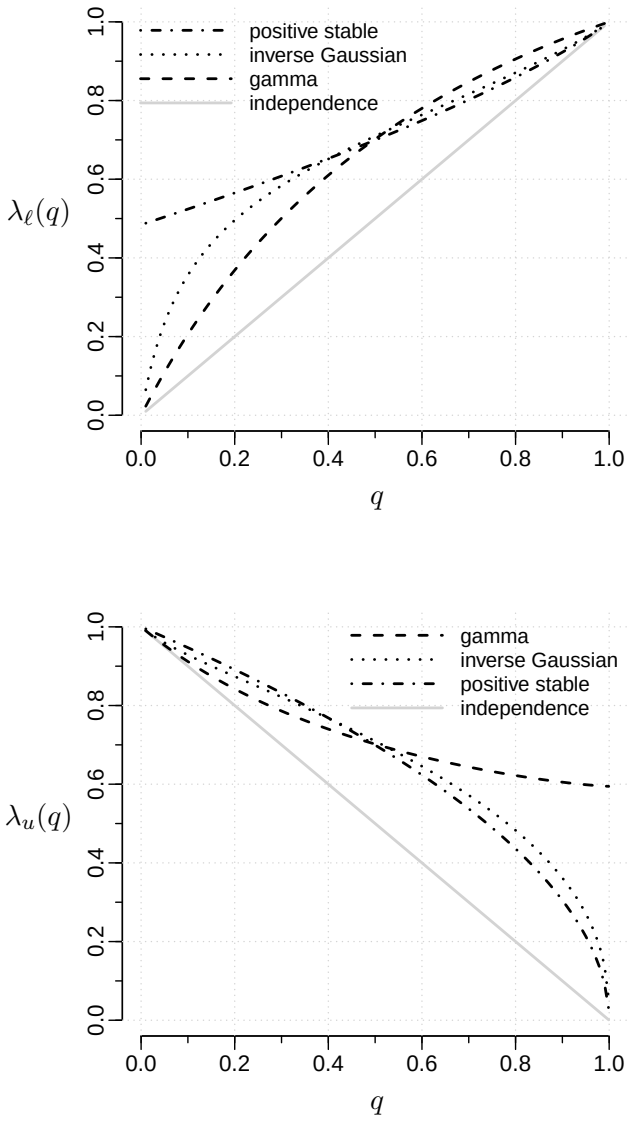

Figure 2 – Model-based association structures obtained from Formula (4) and Formula (5). The frailty distribution parameter is chosen to yield a Kendall's τ of 0.4

3.3 Observed structure

To compare the model-based association structures with the observed (model-free) structure, we now derive a non-parametric estimate of quantile dependence. Because censoring often prevents from going far enough into the upper tail, we focus on the lower coefficient.

3.3.1 In dimension 2 without covariate

We start with pairs of event times (T_1, T_2), possibly right-censored by two non-negative random variables (C_1, C_2). Then, for $j = 1, 2$, the random variables that we observe are $Y_j = \min(T_j, C_j)$ and $\delta_j = I(T_j \leq C_j)$. We denote by $S(\cdot, \cdot)$ and $G(\cdot, \cdot)$ the joint survival functions of (T_1, T_2) and (C_1, C_2), respectively.

Without any covariate, T_1 and T_2 have the same marginal distribution function ($F_1(\cdot) = F_2(\cdot) = F(\cdot)$) which can be estimated via the Kaplan-Meier estimator

$$\hat{F}(t) = 1 - \prod_{y_{(k)} \leq t} \left(1 - \frac{d_k}{r_k}\right)$$

with $y_{(1)} < \cdots < y_{(K)}$ the ordered distinct event times, d_k the number of events at $y_{(k)}$, and r_k the number of observations still in the risk set just prior to $y_{(k)}$.

Let $t_1 = t_2 = \hat{F}^{-1}(q)$ for some $q \in (0, 1)$. Based on s independent realisations of $(Y_1, \delta_1, Y_2, \delta_2)$, $\{(y_{i1}, \delta_{i1}, y_{i2}, \delta_{i2}) \,|\, i = 1, \ldots, s\}$, we obtain a non-parametric estimate $\tilde{\lambda}_\ell(q)$ of $\lambda_\ell(q)$ by substituting a non-parametric estimate $\tilde{S}(t_1, t_2)$ of $S(t_1, t_2)$ into Equation (2).

Under the hypothesis of independence between the event times and the censoring times, we have

$$S(t_1, t_2) = \frac{\Pr[T_1 > t_1, Y_2 > t_2]}{G(t_1, t_2)}$$

The numerator is the joint survival function of the observables (Y_1, Y_2) at (t_1, t_2), which we can estimate by its natural empirical counterpart

$$\frac{1}{s} \sum_{i=1}^{s} I(y_{i1} > t_1, y_{i2} > t_2)$$

The estimation of the denominator, on the other hand, depends on the censoring mechanism (Wang & Wells, 1997). In numerous applications, C_1 and C_2 are independent and identically distributed so that

$$G(t_1, t_2) = \Pr[C_1 > t_1] \Pr[C_2 > t_2] = \Pr[C_1 > t_1] \Pr[C_1 > t_2]$$

Univariate censoring provides another common framework for which $C_1 = C_2 = C$, leading to

$$G(t_1, t_2) = \Pr[C > \max(t_1, t_2)]$$

In these two cases, $G(t_1, t_2)$ can be estimated using the Kaplan-Meier estimator by interchanging the role of the event times and the censoring times. More cases are covered in Wang & Wells (1997).

3.3.2 With covariates

In the presence of covariates, the overall Kaplan-Meier estimator $\hat{F}$ is no longer valid. For a single binary covariate, we can use two different Kaplan-Meier estimators (and we let $t_1 = \hat{F}_1^{-1}(q)$ and $t_2 = \hat{F}_2^{-1}(q)$). For more than two groups, or in the presence of a continuous covariate, the observations can be partitioned into two groups based on a prognostic index that combines information from the available covariates. Alternatively, the Kaplan-Meier estimator can be superseded by a kernel-based estimator (Beran, 1981).

3.3.3 In higher dimensions

With cluster sizes larger than 2 (possibly different across clusters), we use a similar ad-hoc method as in Glidden (2007). We propose to repeat step (i) and step (ii) a large number of times, with

(i) select one pair of observations per cluster, and

(ii) obtain $\tilde{\lambda}_\ell(q)$ based on these s clusters of size 2,

and to combine the bivariate estimates by taking the average over the repetitions.

3.3.4 Remark

In Appendix B, we show how one can use $\tilde{\lambda}(q)$ to derive a new non-parametric estimator of the Kendall's τ under censoring.

4 Examples

We illustrate the proposed method with two examples from the book by Duchateau & Janssen (2008, Example 1.4 and Example 1.2). For the first example, we use the real data. For the second one, the data cannot be made publicly available but a data set with a similar structure has been downloaded from the book's website (`www.vetstat.ugent.be/research/frailty/`).

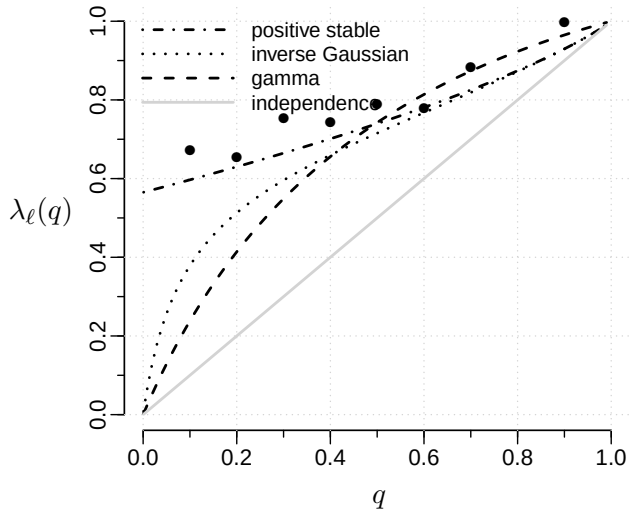

Figure 3 – Diagnostic plot for the data on times to infection in dairy cows (cf. Section 4.1)

4.1 Time to infection in dairy cows

The sample consists of 100 dairy cows which were followed-up for infection in one or more udder quarters. We distinguish between multiparous and primiparous cows. Observations are at the udder level. Each cluster is thus of size 4, and the covariate information does not change within a cluster. About 20% of the observations are censored (no infection occurred during follow-up), with censoring at the cow level. The diagnostic plot is given in Figure 3. The model-based association structures are obtained using a Weibull baseline hazard function. The diagnostic plot displays evidence against the gamma distribution. This is in line with the result in Geerdens et al. (2012) where the same data were analysed using an order selection test procedure. In addition, our diagnostic plot suggests the positive stable frailty distribution to model dependence in these data. Using the positive stable frailty distribution, we find a hazard ratio of $\exp(\hat{\beta}) = 1.762$ (95%CI : $[0.871, 3.562]$) and a Kendall's τ of $\hat{\tau} = 0.479$. The gamma frailty distribution would have lead to a hazard ratio of $\exp(\hat{\beta}) = 1.404$ (95%CI : $[0.737, 2.674]$) and a Kendall's τ of $\hat{\tau} = 0.475$.

4.2 Time to fracture healing in dogs

Each of 106 dogs was evaluated for fracture healing by means of two methods, and times to healing were recorded. For comparison's sake, we also

apply the diagnostic plot introduced in Oakes (1989). This can be done because (i) each cluster is of size 2, (ii) there is no censoring, and (iii) for each dog, observation 1 always refers to method 1 while observation 2 always refers to method 2. The two diagnostic plots are given in Figure 4. For the model-based structures, we used a log-logistic baseline hazard which provides a better fit than the Weibull. Both diagnostic plots clearly suggest that frailties follow the positive stable distribution. In that case, the hazard ratio equals $\exp(\hat{\beta}) = 3.124$ (95%CI : $[2.193, 4.450]$) and the Kendall's τ is $\hat{\tau} = 0.356$. The gamma frailty distribution would have lead to a hazard ratio of $\exp(\hat{\beta}) = 2.132$ (95%CI : $[1.548, 2.935]$) and a Kendall's τ of $\hat{\tau} = 0.072$.

5 Simulations

In this section we report the results of a simulation study that we conducted to evaluate the performance of the proposed diagnostic plot. Parameters that might influence the performance are the number of clusters/observations per cluster ($(s, n) = (125, 2)$, $(50, 5)$, $(25, 10)$, $(10, 25)$, $(250, 2)$, $(50, 10)$, $(25, 20)$, and $(10, 50)$), the degree of global dependence ($\tau = 0.25$ and $\tau = 0.50$), and the censoring fraction (0, 0.2, and 0.5). We generate data according to the frailty model (1), with gamma or positive stable frailties, and we report the proportion of times, out of 500 simulations, that the correct frailty distribution (either gamma or positive stable) is selected. Decision is made by choosing the frailty distribution that minimises the sum of squared distances between the model-based estimates $\hat{\lambda}_\ell(q)$ and the non-parametric estimates $\tilde{\lambda}_\ell(q)$ (calculated for q between 0.1 and 0.9 with span 0.1).

5.1 Data generation

We consider a binary covariate, whose values are drawn from a Bernoulli distribution, which divides each cluster into two roughly balanced groups. We set the associated coefficient to -0.3. Then, the conditional hazard ratio equals $\exp(-0.3) \approx 0.75$. A constant baseline hazard function is used, $h_0(t) = \lambda$ for all $t > 0$, resulting in exponentially distributed event times (given the frailties). We set $\lambda = 0.2$. For each event time, a censoring time is also generated from a common exponential distribution whose parameter is chosen to achieve the targeted percentage of censoring. We record the minimum between the event times and the censoring times, together with 0-1 indicators.

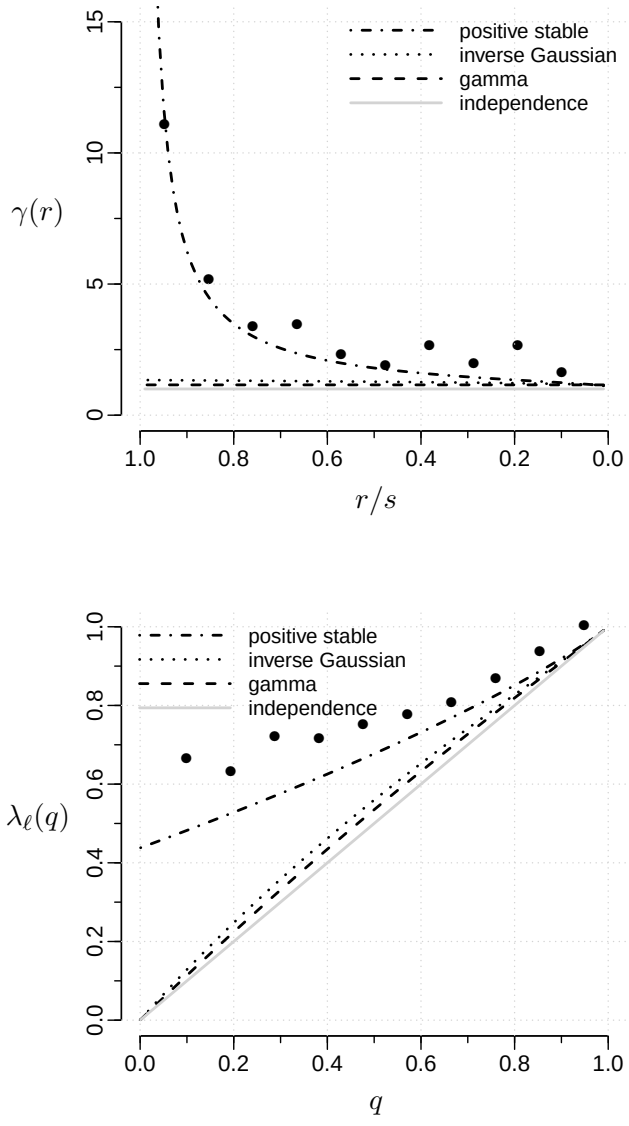

Figure 4 – Diagnostic plot for the data on times to fracture healing in dogs (cf. Section 4.2). The top panel is based on the cross ratio function (see Duchateau & Janssen, 2008, Section 4.2.6). The bottom panel is the new diagnostic plot based on quantile dependence

5.2 Results

Table 1 records the results. Better performances are observed under the high level of global dependence ($\tau = 0.50$). For a fixed value of τ, increasing the number of clusters greatly improves the proportion of correct decisions. On the other hand, increasing the number of observations per cluster does not seem to make a tremendous difference, except if the level of censoring is high. Indeed, high levels of censoring raise concerns with limited cluster sizes. This is made manifest in bivariate data where heavy censoring, which produces many clusters that only contain censored observations, deteriorates the quality of the diagnostic plot. In summary, this simulation study indicates that the proposed method performs well as long as the number of clusters is large enough, with possibly the need for moderately large cluster sizes if censoring is high.

6 Concluding remarks

Many of the diagnostic checks mentioned in Section 1 are either very specific (like the plot in Shih & Louis (1995), only valid for gamma distributed frailties, or the one derived from Oakes (1989), only valid for clusters of size 2 with a natural ordering and no censoring), or quite theoretically and mathematically involved (like in Geerdens et al. (2012) or in Glidden (2007)). Further, no implementation has appeared in standard statistical packages. From a practical-oriented point of view, these are serious limitations.

The main advantage of the diagnostic plot proposed here, above its simplicity, is that it guides the choice of the frailty distribution rather than only (not) rejecting one particular candidate. Provided that the number of clusters is large enough, the proposed method is reliable for data involving different cluster sizes and various censoring schemes. The proposed diagnostic plot can further be used to evaluate the need to adjust for clustering, hence supplementing the test $\mathrm{H}_0 : \theta = 0$ versus $\mathrm{H}_1 : \theta > 0$, with θ the frailty variance, for which guidance is still discussed in the literature (Wienke, 2010, Section 7.5). An R code implementing the method is available as supplementary material.

Diagnostic plots have proven to be useful to examine many aspects of Cox models, like the assumption of proportional hazards (see, e.g., Collett, 2003, Chapter 4). Diagnostic plots are generally not intended to replace formal numerical tests, but they are certainly useful companions to guide the model construction.

Table 1 – Simulation results: proportion of times, out of 500 simulations, that the correct frailty distribution—either gamma or positive stable—is selected

τ	cens.	$s \times n = 250$				$s \times n = 500$			
	frac.	125×2	50×5	25×10	10×25	250×2	50×10	25×20	10×50
gamma									
0.25	0	0.97	0.99	0.99	0.98	1.00	1.00	1.00	0.99
	0.2	0.86	0.97	0.99	0.98	0.94	1.00	1.00	0.98
	0.5	0.60	0.78	0.86	0.89	0.70	0.95	0.97	0.94
0.50	0	1.00	1.00	1.00	0.97	1.00	1.00	1.00	0.99
	0.2	0.98	0.99	1.00	0.98	1.00	1.00	1.00	0.98
	0.5	0.81	0.93	0.95	0.94	0.88	0.99	0.99	0.97
pos. stable									
0.25	0	0.94	0.81	0.67	0.38	0.99	0.84	0.67	0.57
	0.2	0.90	0.81	0.69	0.45	0.98	0.82	0.66	0.59
	0.5	0.77	0.70	0.65	0.49	0.87	0.80	0.66	0.45
0.50	0	0.99	0.89	0.80	0.58	1.00	0.96	0.87	0.59
	0.2	0.99	0.89	0.82	0.58	1.00	0.92	0.84	0.58
	0.5	0.82	0.75	0.73	0.55	0.94	0.80	0.75	0.48

Appendices

A

Explicit formulas

	$\underset{\theta>0}{\mathrm{Gam}(\theta)}$	$\underset{\theta>0}{\mathrm{IG}(\theta)}$	$\underset{0<\theta<1}{\mathrm{PS}(\theta)}$
$\mathcal{L}(x)$	$(1+\theta x)^{-1/\theta}$	$\exp\left\{\frac{1}{\theta}\left(1-\sqrt{1+2\theta x}\right)\right\}$	$\exp(-x^{1-\theta})$
$\mathcal{L}^{-1}(x)$	$\frac{1}{\theta}\left(x^{-\theta}-1\right)$	$\frac{\log(x)}{2}\left(\theta\log(x)-2\right)$	$\left(-\log(x)\right)^{1/(1-\theta)}$
$\tau^{\ddagger}$	$\frac{\theta}{\theta+2}$	$\frac{1}{2}-\frac{1}{\theta}+\frac{2}{\theta^2}\exp\left(\frac{2}{\theta}\right)\int_{2/\theta}^{\infty}u^{-1}\exp(-u)\,\mathrm{d}u$	θ
$\zeta^{\star}(v)$	$\theta+1$	$1+\frac{1}{(1/\theta)-\log(v)}$	$1-\frac{\theta}{(1-\theta)\log(v)}$
$\lambda_{\ell}(q)$	$\frac{\left(2(1-q)^{-\theta}-1\right)^{-1/\theta}+2q-1}{q}$	$\frac{\mathcal{L}\left\{\log(1-q)\left(\theta\log(1-q)-2\right)\right\}+2q-1}{q}$	$\frac{\exp\left\{2^{1-\theta}\log(1-q)\right\}+2q-1}{q}$
$\lambda_{u}(q)$	$\frac{\left(2(1-q)^{-\theta}-1\right)^{-1/\theta}}{1-q}$	$\frac{\mathcal{L}\left\{\log(1-q)\left(\theta\log(1-q)-2\right)\right\}}{1-q}$	$\frac{\exp\left\{2^{1-\theta}\log(1-q)\right\}}{1-q}$
Λ_{ℓ}	0	0	$2(1-2^{-\theta})$
Λ_{u}	$2^{-1/\theta}$	0	0

‡ The Kendall's τ for an IG frailty term is always lower than $1/2$.

Table 2 – Explicit formulas for the gamma (Gam), the inverse gaussian (IG), and the positive stable (PS) frailty distributions

B

A new non-parametric estimator of the Kendall's τ under censoring

Based on complete data (no censoring), the Kendall's τ is easily estimated in a non-parametric way by enumerating the concordant and discordant pairs (see, e.g., Hougaard, 2000, Section 4.2). The presence of censoring makes it more difficult because there are cases where the concordant/discordant status cannot be ascertained. We propose a new non-parametric estimate of the Kendall's τ under censoring as a direct by-product of our methodology.

An alternative to the Kendall's τ, which requires one pair only, is given by the median concordance coefficient (Hougaard, 2000, Section 4.4)

$$\beta = \Pr\left[(T_1 - \tilde{t}_1)(T_2 - \tilde{t}_2) > 0\right] \Pr\left[(T_1 - \tilde{t}_1)(T_2 - \tilde{t}_2) < 0\right]$$

with $\tilde{t}_1$ and $\tilde{t}_2$ the medians of T_1 and T_2, respectively. The median concordance coefficient is also known as the Blomqvist's β. Like the Kendall's τ, the Blomqvist's β lies between -1 and 1, equals 0 under independence, and, in the world of frailty models, can be written in terms of the Laplace transform as

$$\beta = 4\mathcal{L}\left\{\mathcal{L}^{-1}(1/2)\right\} - 1$$

so that (cf. Equation (4))

$$\beta = 2\lambda_\ell(1/2) - 1$$

The Blomqvist's β often provides an accurate approximation of the Kendall's τ (Nelsen, 2006, Section 5.4.1). As an illustration, the relationship between β and τ is depicted in Figure 5 for the three frailty distributions that we have considered in this paper.

Thus

$$\hat{\beta} := 2\tilde{\lambda}_\ell(1/2) - 1$$

with $\tilde{\lambda}_\ell(1/2)$ given in Section 3.3, provides a non-parametric estimate of the Blomqvist's β that can be used as an accurate approximation of the Kendall's τ.

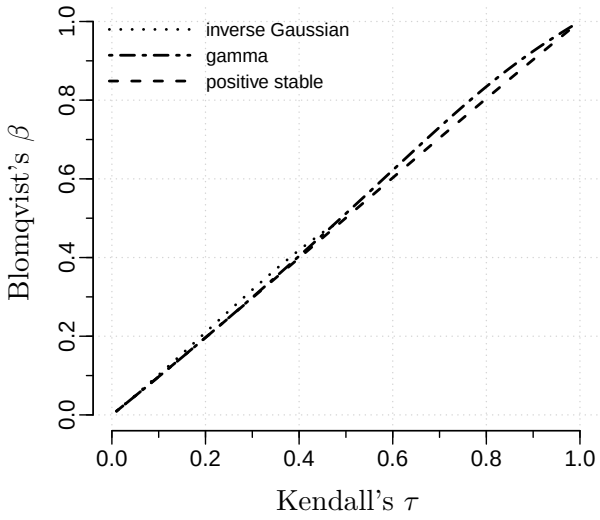

Figure 5 – The Blomqvist's β versus the Kendall's τ

References

Anderson, J. E., Louis, T. A., Holm, N. V. & Harvald, B. (1992). Time-dependent association measures for bivariate survival distributions. *Journal of the American Statistical Association* **87**, 641–650.

Beran, R. (1981). Nonparametric regression with randomly censored data. Tech. rep., University California, Berkeley.

Collett, D. (2003). *Modelling Survival Data in Medical Research*. Chapman and Hall/CRC, 2nd ed.

Cui, S. & Sun, Y. (2004). Checking for the gamma frailty distribution under the marginal proportional hazards frailty model. *Statistica Sinica* **14**, 249–267.

Duchateau, L. & Janssen, P. (2008). *The Frailty Model*. Springer.

Economou, P. & Caroni, C. (2008). Graphical tests for the frailty distribution in the shared frailty model. *Communications in Statistics - Simulation and Computation* **37**, 978–992.

Geerdens, C., Claeskens, G. & Janssen, P. (2012). Goodness-of-fit tests for the frailty distribution in proportional hazards models with shared frailty. *Biostatistics* , 1–14.

Glidden, D. V. (2007). Pairwise dependence diagnostics for clustered failure-time data. *Biometrika* **94**, 371–385.

Hougaard, P. (1995). Frailty models for survival data. *Lifetime Data Analysis* **1**, 255–273.

Hougaard, P. (2000). *Analysis of Multivariate Survival Data*. Springer.

Munda, M., Rotolo, F. & Legrand, L. (2012). parfm: Parametric frailty models in R. *Journal of Statistical Software* **51**, ??–??

Nelsen, R. B. (2006). *An Introduction to Copulas*. Springer, 2nd ed.

Oakes, D. (1989). Bivariate survival models induced by frailties. *Journal of the American Statistical Association* **84**, 487–493.

Schmidt, R. & Stadtmüller, U. (2006). Non-parametric estimation of tail dependence. *Scandinavian Journal of Statistics* **33**, 307–335.

Shih, J. H. & Louis, T. A. (1995). Assessing gamma frailty models for clustered failure time data. *Lifetime Data Analysis* **1**, 205–220.

Viswanathan, B. & Manatunga, A. K. (2001). Diagnostic plots for assessing the frailty distribution in multivariate survival data. *Lifetime Data Analysis* **7**, 143–155.

Wang, W. & Wells, M. T. (1997). Nonparametric estimators of the bivariate survival function under simplified censoring conditions. *Biometrika* **84**, 863–880.

Wienke, A. (2010). *Frailty Models in Survival Analysis.* Chapman and Hall/CRC.