The relative frailty variance and shared frailty models

C. Paddy Farrington [1] and Steffen Unkel

The Open University, Milton Keynes, UK

Karim Anaya-Izquierdo

London School of Hygiene and Tropical Medicine, London, UK

[1] *Address for correspondence*: Paddy Farrington, Department of Mathematics and Statistics, The Open University, Milton Keynes MK7 6AA, United Kingdom. E-mail: c.p.farrington@open.ac.uk.

Summary. The relative frailty variance among survivors provides a readily interpretable measure of how the heterogeneity of a population, as represented by a frailty model, evolves over time. We discuss the properties of the relative frailty variance, show that it characterizes frailty distributions, and that, suitably rescaled, it may be used to compare patterns of dependence across models and data sets. In shared frailty models, the relative frailty variance is closely related to the cross-ratio function, which is estimable from bivariate survival data. We investigate the possible shapes of the relative frailty variance function for the purpose of model selection, and review available frailty distribution families in this context. We introduce several new families with contrasting properties, including simple but flexible time-varying frailty models. The benefits of the approach we propose are illustrated with two applications to bivariate current status data obtained from serological surveys.

Keywords: Cross-ratio function; Cure model; Current status data; Frailty; Heterogeneity; Relative frailty variance; Shared frailty model; Survival data; Time-varying frailty.

1 Introduction

Hougaard (1984) remarked that survival models with Gamma and inverse Gaussian frailties behave very differently, noting that the relative frailty distribution among survivors is independent of age for the Gamma, but becomes more homogeneous with time for the inverse Gaussian. Since this seminal paper other frailty distributions have been suggested, notably the power variance family (Hougaard 2000), which contains members whose relative frailty distribution in survivors becomes less homogeneous with time (Aalen *et al.* 2008).

The present paper takes as its focus the relative frailty distribution in survivors, concentrating on the variance of that distribution and how it evolves over time. Its varied behaviour suggests that it should have a role to play in identifying suitable frailty models, provided sufficient information to identify it is available. This is typically the case for bivariate survival data arising from a shared frailty model, studied by Oakes (1989). In such models, the cross-ratio function (Clayton 1978) may be estimated; it turns out to have a close relationship with the relative frailty variance.

In Section 2 the relative frailty variance is defined and related to the cross-ratio function from shared frailty models. In Section 3, the functional form of the relative frailty variance function is investigated and two families it characterizes are studied. In Section 4 further frailty models are presented, and some properties of the relative variance function, notably for a class of frailty cure models (Aalen *et al.* 2008), are obtained. Some models involving time-varying frailties are discussed in Section 5. Then in Section 6, two examples are presented of how the relative frailty variance can be used in practice. Technical proofs are given in the Appendix.

2 The relative frailty variance

2.1 The relative frailty distribution

Suppose that the random variable T denotes an individual's time to event from some time origin $T = 0$, and that U is a non-negative random variable. We shall assume for the

time being that U has finite mean and variance, with $E(U) = \mu$, and consider the frailty model for T specified by the hazard function

$$\lambda(t; U) = U\lambda(t),$$

where $\lambda(t)$ is the baseline hazard.

Consider now an individual with frailty U who survives to time t, that is, for whom $T > t$. This individual's hazard may be expressed as

$$\lambda(t; U) = U_t \times \lambda(t|T > t),$$

where

$$U_t = \frac{U}{E(U|T > t)}$$

and

$$\lambda(t|T > t) = E(U|T > t)\lambda(t).$$

This defines a frailty model in the selected population of survivors at time t, with frailty U_t and baseline hazard $\lambda(t|T > t)$. The distribution of the frailty in survivors was first investigated by Hougaard (1984). The mean of U_t conditional on $T > t$ is 1. Note that the frailty U_t is defined relative to the original frailty U by dividing it by its conditional mean; in particular, $U_0 = U/\mu$.

Definition U_t is the relative frailty at time t. The relative frailty variance at time t is

$$RFV^*(t) = \text{var}(U_t|T > t) = \frac{\text{var}(U|T > t)}{E(U|T > t)^2}. \quad (1)$$

■

The relative frailty variance thus describes the heterogeneity of survivors at time t. Equivalently, $RFV^*(t)$ is the square of the coefficient of variation of U conditional on $T > t$. Suppose now that U has distribution function $F(u)$ and Laplace transform

$$L(s) = E(e^{-sU}) = \int_0^\infty e^{-su} \mathrm{d}F(u).$$

It is sometimes more convenient to work with the cumulant generating function of U,

$$K(s) = \log L(-s).$$

Let

$$\Lambda(t) = \int_0^t \lambda(s)\mathrm{d}s$$

be the cumulative baseline hazard at time t. The baseline hazard $\lambda(t)$ and cumulative baseline hazard $\Lambda(t)$ are determined up to the multiplicative constant μ, the frailty mean. In what follows, we shall also require the absolute baseline hazard $\lambda_a(t) = \mu\lambda(t)$ and the absolute cumulative baseline hazard $\Lambda_a(t) = \mu\Lambda(t)$ (note that these are not the same as the population hazard and cumulative hazard).

The Laplace transform of the distribution of U in survivors at time t is $L(s+\Lambda(t))/L(\Lambda(t))$ (Aalen *et al.* 2008, page 243), from which it follows that $E(U|T > t) = K'(-\Lambda(t))$ and $\mathrm{var}(U|T > t) = K''(-\Lambda(t))$. Hence

$$RFV^*(t) = \frac{K''(-\Lambda(t))}{K'(-\Lambda(t))^2} = \frac{K''(-\Lambda_a(t)/\mu)}{K'(-\Lambda_a(t)/\mu)^2}. \tag{2}$$

Note that $RFV^*(t)$ depends on t only through $\Lambda_a(t)$, and thus may be characterized independently of the baseline hazard function by rescaling the time axis.

Definition Let U be as above. The (scaled) relative frailty variance is the function

$$RFV(s) = \frac{\mathrm{var}(U|T > \Lambda_a^{-1}(s))}{E(U|T > \Lambda_a^{-1}(s))^2},$$

where $s \geq 0$ belongs to the image set $\{\Lambda_a(t) : t \in (0, \infty)\}$. ■

From (2), it follows that the scaled relative variance function does not depend on the baseline hazard, and may be written

$$RFV(s) = \frac{K''(-s/\mu)}{K'(-s/\mu)^2}. \tag{3}$$

The scaled relative variance function is a feature solely of the frailty U. The absolute cumulative baseline hazard can be thought of as determining the application-specific time scale according to which the selection process takes place. From now on we refer to both $RFV(s)$ and $RFV^*(t)$ as 'the' relative frailty variance, stating explicitly whether it is rescaled only if required for clarity.

2.2 The relative frailty variance and the cross-ratio function

Frailty distributions are generally not identifiable from univariate data, unless strong parametric assumptions are made about the hazard function. In shared frailty models

for bivariate survival data, however, the frailty distribution is identifiable through the cross-ratio function.

Consider the shared frailty model with frailty U and baseline hazards $\lambda_j(t)$ for $j = 1, 2$ defined by

$$\lambda_j(t; U) = U\lambda_j(t).$$

The joint survivor function is

$$S(t_1, t_2) = E\Big(\exp\big(- U\{\Lambda_1(t_1) + \Lambda_2(t_2)\}\big)\Big),$$

where $\Lambda_j(t) = \int_0^t \lambda_j(s)\mathrm{d}s$ are the cumulative hazards, $j = 1, 2$. The cross-ratio function, introduced by Clayton (1978) and studied by Oakes (1989), is defined as follows.

Definition The cross-ratio function at (t_1, t_2) is

$$\theta^*(t_1, t_2) = \frac{S(t_1, t_2)S_{12}(t_1, t_2)}{S_1(t_1, t_2)S_2(t_1, t_2)}, \tag{4}$$

where

$$S_j(t_1, t_2) = \frac{\partial S(t_1, t_2)}{\partial t_j}, \quad j = 1, 2, \qquad S_{12}(t_1, t_2) = \frac{\partial^2 S(t_1, t_2)}{\partial t_1 \partial t_2}.$$

■

The cross-ratio function provides a convenient measure of local association for bivariate survival data. Since

$$S(t_1, t_2) = \exp\Big\{ K\big(- \Lambda_1(t_1) - \Lambda_2(t_2)\big)\Big\}$$

it follows that

$$\theta^*(t_1, t_2) = 1 + RFV(\Lambda_{a1}(t_1) + \Lambda_{a2}(t_2)), \tag{5}$$

where $\Lambda_{a1}(t) = \mu\Lambda_1(t_1)$ and $\Lambda_{a2}(t) = \mu\Lambda_2(t_2)$ are the absolute baseline hazards. The relationship (5) is implicit in Anderson *et al.* (1992). Oakes (1989) showed that in Archimedean copulas (which include shared frailty models) the cross-ratio function depends on (t_1, t_2) only through some function of the joint survivor function $S(t_1, t_2)$. Equation (5) shows that, in shared frailty models, it can also be expressed solely in terms of the sum of the absolute cumulative baseline hazards, $\Lambda_{a1}(t_1) + \Lambda_{a2}(t_2)$. Thus we can set $s = \Lambda_{a1}(t_1) + \Lambda_{a2}(t_2)$, and define

$$\theta(s) = \theta^*(t_1, t_2).$$

The connection between the relative frailty variance and the cross-ratio function is summarized in the following proposition.

Proposition 1 Suppose that the frailty U has mean μ and relative variance function $RFV(s)$. In a shared frailty model with frailty U and baseline cumulative hazards $\Lambda_j(t)$, $j = 1, 2$, the cross-ratio function $\theta^*(t_1, t_2)$ depends only on t_1 and t_2 through the sum of the absolute cumulative baseline hazards $\Lambda_{aj}(t) = \mu\Lambda_j(t)$ with

$$\theta(s) = 1 + RFV(s), \tag{6}$$

where $s = \Lambda_{a1}(t_1) + \Lambda_{a2}(t_2)$ and $\theta(s) = \theta^*(t_1, t_2)$. ■

When bivariate data are thought to arise from a shared frailty model, diagnostic plots based on the cross-ratio function may be used to suggest an appropriate frailty distribution (Viswanathan and Manatunga 2001, Duchateau and Janssen 2008). However, hitherto interest has focused largely on whether the cross-ratio function (and hence the relative frailty variance) is constant, suggesting a Gamma frailty, or decreases, perhaps suggesting an inverse Gaussian or other Hougaard frailties. In fact, the relative frailty variance function can take much more diverse profiles.

3 Functional form of the relative frailty variance

3.1 A characterization of relative frailty variance functions

Equation (3) expresses the relative frailty variance in terms of the cumulant generating function of the frailty distribution. Suppose now that $a(s)$ is an arbitrary non-negative function on $[0, \infty)$. When is $a(s)$ the relative variance function of some frailty distribution, and what is this distribution?

We begin with the following inversion theorem, obtained by solving the differential equation (3). The result can be verified directly by differentiation.

Proposition 2 Suppose that U is a frailty with mean $E(U) = \mu$ and Laplace transform $L(s)$ defined and twice differentiable on $(0, \infty)$. Let $a(s)$ be a non-negative integrable function defined on $[0, \infty)$. Then the relative frailty variance function of U, $RFV(s)$,

equals $a(s)$ if and only if

$$L(s) = \exp\Big(- \int_0^{\mu s} \frac{1}{1 + A(t)} \mathrm{d}t \Big), \tag{7}$$

where $A(t) = \int_0^t a(s)\mathrm{d}s$. ■

The relative frailty variance function only determines the frailty distribution, via its Laplace transform, up to its mean μ. Relative frailty variance functions thus partition the space of frailty distributions into equivalence classes.

Let $\eta(t)$ denote the population hazard. Note that the population survivor function

$$S(t) = \exp\Big(- \int_0^t \eta(s)\mathrm{d}s \Big)$$

is equal to $L(\Lambda(t))$. Thus, changing variables in (7) from s to $\Lambda(t)$ yields the following corollary.

Proposition 3 Suppose that U is a frailty with mean μ and relative frailty variance $a(s)$, and that the absolute baseline hazard is $\lambda_a(t)$. Then the population hazard is

$$\eta(t) = \frac{\lambda_a(t)}{1 + A(\Lambda_a(t))}, \tag{8}$$

where $A(t) = \int_0^t a(s)\mathrm{d}s$ and $\Lambda_a(t) = \int_0^t \lambda_a(s)\mathrm{d}s$. ■

Expression (8) generalizes to arbitrary frailties the expression for the population hazard of a Gamma frailty of unit mean (Aalen *et al.* (2008), page 237).

Proposition 2 yields only a partial answer to the question 'which functions $a(s)$ are relative frailty variances', in the sense that some functions $a(s)$ yield an $L(s)$ in (7) that is not a Laplace transform. One such example, to be described in Subsection 3.2, is $a(s) = \gamma e^{\alpha s}$, where $\gamma > 0$ and $\alpha > \gamma + 1$.

The next result provides necessary and sufficient conditions on $a(s)$ to be the relative frailty variance of a frailty distribution.

Proposition 4 A non-negative, integrable function $a(s)$ is the relative variance function of a frailty if and only if it is infinitely differentiable and

$$\forall n \geq 1 \quad \forall s > 0 \qquad (-1)^n D^{(n)}(s) \geq 0, \tag{9}$$

where the $D^{(n)}(s)$ are defined recursively by

$$\begin{aligned} D^{(1)}(s) &= -(1 + A(s))^{-1}, \\ D^{(n+1)}(s) &= D^{(1)}(s) D^{(n)}(s) + \frac{\mathrm{d}D^{(n)}(s)}{\mathrm{d}s}, \end{aligned} \tag{10}$$

with $A(s) = \int_0^s a(t)\mathrm{d}t$. ■

This characterization theorem follows immediately from Proposition 2 (with $\mu = 1$) and Bernstein's Theorem (Feller (1966), page 415) applied to $L(s)$ in expression (7) and the fact that $L(0) = 1$. Note that many variants of the recursive relations (10) are possible; we have not found any to be very useful in practical applications so have given one which is concisely stated.

3.2 Families of frailties characterized by their relative frailty variance functions

A natural way to characterize families of frailties is by choosing a simple yet flexible parametric family of functions $a(s) \geq 0$ and studying the properties of the resulting frailty distributions.

Example 1: The power variance family

Frailty distributions in the power variance family (Hougaard 2000, Aalen *et al.* 2008) have relative frailty variances of the form

$$a(s) = \frac{\alpha+1}{\nu}\Big(1+\frac{s}{\nu}\Big)^{\alpha}, \qquad \nu > 0, \alpha > -1. \tag{11}$$

For these functions, $1 + A(s) = (1 + s/\nu)^{\alpha+1}$ and the resulting Laplace transform (for a frailty of unit mean) is

$$L(s) = \exp\Big(-\frac{\nu}{\alpha}\Big\{1-\Big(\frac{\nu}{\nu+s}\Big)^{\alpha}\Big\}\Big), \qquad \nu > 0, \alpha > -1. \tag{12}$$

When $\alpha = 0$ the frailty is Gamma distributed, and when $-1 < \alpha < 0$ it is a Hougaard distribution (Hougaard 1986); a special case is the inverse Gaussian, for $\alpha = -0.5$. When $\alpha > 0$ the distribution of the frailty U is compound Poisson with a non-zero probability that $U = 0$. The limiting distribution when $\alpha \to \infty$ with $\alpha/\nu \to 1$ is the Poisson. The values $\alpha < -1$ do not correspond to densities on $[0, \infty)$. Figure 1 shows the relative variance functions of some members of this family.

Example 2: The Addams family

The power variance family stems from a relative variance function with a simple rational form. An obvious alternative is the family of exponential relative variance functions

$$a(s) = \gamma e^{\alpha s}, \qquad \gamma > 0, \quad \alpha \text{ unrestricted.} \tag{13}$$

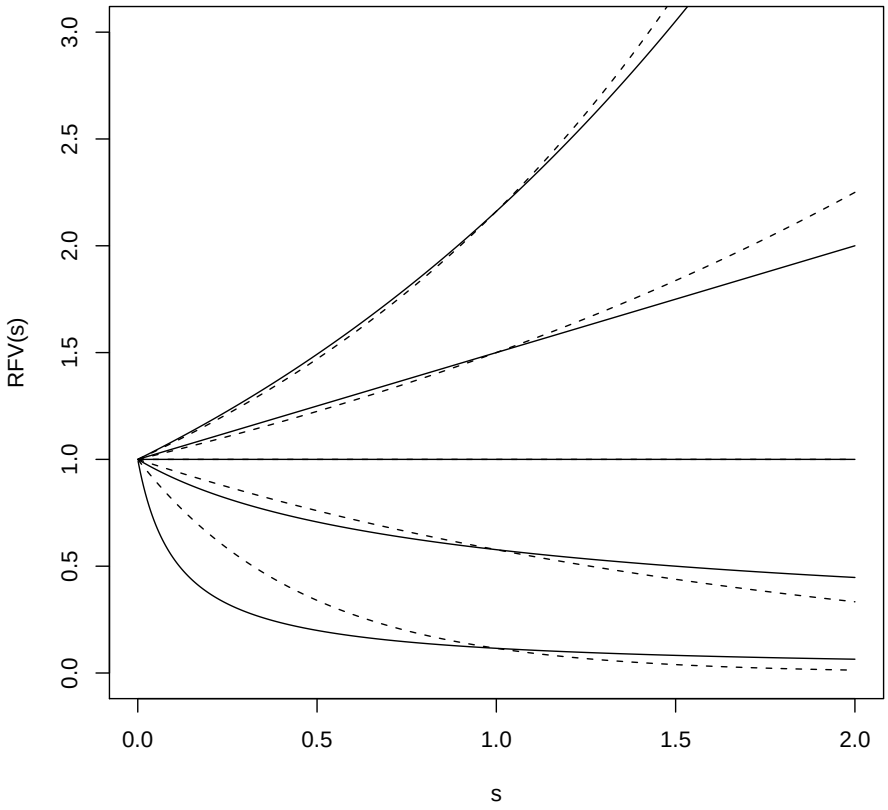

Figure 1: Relative frailty variance functions from the PVF (full lines) and Addams (dashed lines) families. The PVF lines are, from top to bottom, $\alpha = 5, 1, 0, -0.5, -0.9$ and $\nu = \alpha + 1$. The Addams lines are chosen to intersect the PVF lines at $s = 0$ and $s = 1$.

Then, $1 + A(s) = \alpha^{-1}\{\alpha - \gamma + \gamma e^{\alpha s}\}$ and so, by Proposition 2 with $\mu = 1$,

$$\begin{aligned} L(s) &= \Big\{\big(1 - \frac{\gamma}{\alpha}\big)e^{-\alpha s} + \frac{\gamma}{\alpha}\Big\}^{\frac{1}{\alpha-\gamma}}, && \text{for } \alpha \neq \gamma, \alpha \neq 0, \\ &= \exp\Big\{\frac{1}{\gamma}(e^{-\gamma s} - 1)\Big\}, && \text{for } \alpha = \gamma, \\ &= (1 + \gamma s)^{-\frac{1}{\gamma}}, && \text{for } \alpha = 0. \end{aligned} \tag{14}$$

We now describe the distributions (with unit mean) corresponding to different values of α, for fixed $\gamma > 0$.

Case 1. $\gamma > 0 > \alpha$. In this case

$$L(s) = \Big(\frac{-\alpha/(\gamma-\alpha)}{1-\big(1+\frac{\alpha}{\gamma-\alpha}\big)e^{\alpha s}}\Big)^{\frac{1}{\gamma-\alpha}} \times e^{\frac{\alpha s}{\gamma-\alpha}},$$

which is the Laplace transform of the sum of two independent random variables: $-\alpha$ times the negative binomial $NB\big(\nu = \frac{1}{\gamma-\alpha}, p = \frac{-\alpha}{\gamma-\alpha}\big)$, and the degenerate (i.e. point mass) distribution at $-\alpha/(\gamma-\alpha)$. Thus, this distribution is $-\alpha$ times a negative binomial shifted by $(\gamma-\alpha)^{-1}$. In this case the frailty distribution has lower terminal > 0, which increases towards 1 as α tends to $-\infty$.

Case 2. $\gamma > \alpha = 0$. The frailty distribution is Gamma with mean 1 and variance γ^{-1}.

Case 3. $\gamma > \alpha > 0$. In this case

$$L(s) = \Big(\frac{\alpha/\gamma}{1-\big(1-\frac{\alpha}{\gamma}\big)e^{-\alpha s}}\Big)^{\frac{1}{\gamma-\alpha}},$$

which is the Laplace transform of α times the negative binomial variable $NB\big(\nu = \frac{1}{\gamma-\alpha}, p = \frac{\alpha}{\gamma}\big)$. The relative frailty variance function $RFV(s) = \gamma\exp(\alpha s)$ is increasing, as with all cases where $\alpha > 0$.

Case 4. $\alpha = \gamma > 0$. $L(s)$ is the Laplace transform of α times the Poisson $P(\alpha^{-1})$. Thus when $\alpha = \gamma = 1$, the distribution is Poisson.

Case 5. $\alpha > \gamma > 0$. Not all values of α result in a Laplace transform. The characterization of Proposition 4 is of little help in establishing which values of α correspond to probability distributions. Instead we use the fact that the function

$$L(s) = (1-p+pe^{-as})^b,$$

where $a > 0$, $b > 0$ and $0 < p < 1$, is the Laplace transform of a random variable if and only if b is a positive integer, in which case the random variable is the scaled binomial $a \times B(b,p)$ on $\{0, a, 2a, 3a, ..., ba\}$. This follows from an adaptation of the argument in Kemp (1979).

In our case, $p = (\alpha-\gamma)/\alpha$, $a = \alpha$ and $b = (\alpha-\gamma)^{-1}$. Thus, $L(s)$ is a Laplace transform if and only if $\alpha - \gamma = n^{-1}$ for integer n, in which case $L(s)$ is the Laplace transform of α times the binomial $B\big(n, (n\alpha)^{-1}\big)$. When $n = 1$, the frailty reduces to a two-point distribution. Figure 1 shows some examples of relative variance functions from the Addams family, superimposed on those of the power variance function family.

The distributions within this family are somewhat eccentric, whence the name we chose for it. Nevertheless, in some areas of application, discrete frailty distributions may be appropriate. One such area is infectious diseases transmitted by sexual contact, in which heterogeneity could be represented by number of sexual partners. The Gamma distribution is the only continuous density within this family, and in the present context can be regarded as the limit of a sequence of discrete distributions with increasing number of distinct frailty classes. It is convenient to include it in any family, as it is the only distribution yielding a constant relative frailty variance, and thus plays the role of reference distribution.

4 Homogeneous families of frailty distributions

For modelling purposes, the decision to use a continuous or discrete distribution, or a mixed distribution with a single atom at zero, should arguably be based on subject-matter knowledge. For example, Aalen and Tretli (1999) make a case for using the compound Poisson distribution based on an understanding of the underlying biology of testicular cancer. Moreover, Figure 1 suggests that the shapes of the relative frailty variance functions can be similar for distinct families, which are thus unlikely to be identifiable from data.

The decision as to whether the relative frailty variance should increase or decrease, however, can in some circumstances be based on data, for example from shared frailty models. In this situation it would be convenient to have at one's disposal families of densities which are all discrete (with the Gamma as a limiting case), or all continuous, or all with an atom at zero, but allow for a range of different increasing and decreasing relative variance functions.

The Addams family is homogeneous in the sense described. However, the power variance family shows that it is not possible to guarantee such homogeneity when starting from a parametric form for the relative variance function. A different approach to obtaining homogeneous families of continuous frailties is required.

4.1 Tilting the Gamma density

The derivative of the relative frailty variance function, obtained from (3), is

$$RFV'(s) = \frac{1}{\mu}\Big\{ - \frac{K'''(-s/\mu)}{K'(-s/\mu)^2} + 2\frac{K''(-s/\mu)^2}{K'(-s/\mu)^3}\Big\},$$

where μ is the frailty mean. Thus $RFV(s)$ is monotone increasing if and only if

$$\forall s > 0 \qquad K'''(-s/\mu) < 2\frac{K''(-s/\mu)^2}{K'(-s/\mu)}. \tag{15}$$

It is monotone decreasing if and only if the above inequality is reversed. Provided that the required moments exist, a necessary condition for monotonicity may be obtained by substituting $s = 0$ in (15). Thus if $RFV(s)$ is increasing, then

$$\rho_3 < 2\Big(\frac{\sigma}{\mu}\Big)^2,$$

where σ^2 is the variance of the frailty and ρ_3 is its third central moment. The Gamma distribution satisfies $\rho_3 = 2(\sigma/\mu)^2$. Thus, if a frailty distribution is to have a monotone increasing (decreasing) relative frailty variance, it must be less (more) skewed than the Gamma with the same mean and variance.

This suggests that one way to create a family of frailty densities which includes members with increasing and decreasing relative frailty variances is to introduce a tilting factor into a Gamma density to alter its skewness. We therefore consider densities of the form

$$f(u;\alpha,\zeta,\nu) = \frac{u^{\alpha-1}e^{-\nu u}t(u;\zeta)}{I(\alpha,\zeta,\nu)}, \qquad u,\nu > 0, \tag{16}$$

where $t(u;\zeta) > 0$ is continuous and monotone and

$$I(\alpha,\zeta,\nu) = \int_0^\infty u^{\alpha-1}e^{-\nu u}t(u;\zeta)\mathrm{d}u$$

is the normalizing constant. The density of the relative frailty, conditional on $T > t$, belongs to the same family, with

$$f(u|T > t;\alpha,\zeta,\nu) = \frac{u^{\alpha-1}e^{-(\nu+\Lambda(t))u}t(u;\zeta)}{I(\alpha,\zeta,\nu+\Lambda(t))}. \tag{17}$$

Densities of the form (16), also considered by Hougaard (1984), have both s and $\log(s)$ as canonical statistics. It follows that the Laplace transform, the positive moments

and the relative frailty variance can be expressed in terms of the normalizing function $I(\alpha, \zeta, \nu)$. Specifically,

$$\begin{aligned} L(s) &= \frac{I(\alpha, \zeta, \nu + s)}{I(\alpha, \zeta, \nu)}, \\ E(U^k) &= \frac{I(\alpha + k, \zeta, \nu)}{I(\alpha, \zeta, \nu)}, \\ RFV(s) &= \frac{I(\alpha + 2, \zeta, \nu + s/\mu) I(\alpha, \zeta, \nu + s/\mu)}{I(\alpha + 1, \zeta, \nu + s/\mu)^2} - 1, \end{aligned} \tag{18}$$

where $\mu = E(U)$. This feature greatly simplifies computations. An interesting property of such families of densities is that, in some circumstances, their relative variance functions converge to a positive asymptote.

Proposition 5 Suppose that a frailty U with density as described in equation (16) which is obtained by tilting a Gamma density with a monotone function $t(u; \zeta)$ such that

$$\lim_{u \to 0} t(u; \zeta) = c(\zeta),$$

where $c(\zeta) > 0$ is a (finite) constant. Then the relative frailty variance $RFV(s)$ tends to α^{-1} as $s \to \infty$. ■

This asymptotic behaviour distinguishes these frailty distributions from all so far mentioned (other than the Gamma).

4.2 Two tilted Gamma densities

We describe two families obtained by tilting the Gamma density using a two-parameter tilting function $t(u; \zeta)$ with $\zeta = (\beta, \delta)$.

Example 3: The Kummer family

This family of densities was introduced into the frailty literature by Aalen *et al.* (2008), page 286, who specifically recommend it for use in shared frailty models. The density is

$$f(u; \alpha, \beta, \delta, \nu) = \frac{\delta^{\beta - \alpha}}{\Gamma(\alpha) U(\alpha, \alpha - \beta + 1, \nu\delta)} u^{\alpha - 1} \exp(-\nu u)(\delta + u)^{-\beta}, \qquad u > 0, \tag{19}$$

where $\alpha, \nu, \delta > 0$ and β is unrestricted. The function $U(a, b, z)$ is the confluent hypergeometric function of the second kind, or Kummer function, with the integral representation

$$\Gamma(a) U(a, b, z) = \int_0^\infty e^{-zt} t^{a-1} (1 + t)^{b-a-1} \mathrm{d}t$$

for a, z positive (Abramovitz and Stegun (1972), section 13.2.5). The Laplace transform of the Kummer density is

$$L(s) = \frac{U(\alpha, \alpha - \beta + 1, (s + \nu)\delta)}{U(\alpha, \alpha - \beta + 1, \nu\delta)}.$$

Limiting cases of this distribution include the Beta distribution of the second kind and the Pareto distribution; see Aalen *et al.* (2008) for details.

When $\beta = 0$, the Kummer density reduces to the Gamma, with constant relative variance function $RFV(s) = \alpha^{-1}$. When $\beta < 0$, $RFV(s)$ is monotone increasing while when $\beta > 0$ $RFV(s)$ is monotone decreasing. By Proposition 5, in each case $RFV(s)$ tends to the asymptote α^{-1}. For given β, changing δ alters the shape of the relative variance function but not its direction. Figure 2 shows some examples of relative frailty variance functions from this family.

Example 4: The extended generalized Gamma and inverse Gaussian (Egg) family

The Kummer family is obtained by tilting the Gamma density using a rational function. An alternative is to use an exponential tilting function. This gives rise to the density

$$f(u; \alpha, \beta, \delta, \nu) = \frac{u^{\alpha-1} e^{-\nu u} e^{-\delta u^{\beta}}}{I(\alpha, \beta, \delta, \nu)}, \qquad u > 0, \tag{20}$$

where $\alpha, \beta, \delta, \nu > 0$. When $\beta = 1$ the density reduces to the Gamma. For $0 < \beta < 1$, the relative frailty variance $RFV(s)$ decreases, whereas for $\beta > 1$ it increases. By Proposition 5, in each case $RFV(s)$ tends towards the asymptote α^{-1}. The densities (20) include the generalized Gamma, corresponding to the case $\nu = 0$ (Johnson *et al.* (1994), page 388, Balakrishnan and Peng (2006), Aalen *et al.* (2008), page 285).

This family has a second branch, with α unrestricted, $\beta < 0$, and $\delta, \nu > 0$. When $\beta = 0$ the density again reduces to the Gamma provided that $\alpha > 0$. This branch of the family includes the generalized inverse Gaussian, corresponding to the case $\beta = -1$ (Johnson *et al.* (1994), page 284) and is thus an extended generalized inverse Gaussian family.

The family as a whole is referred to as the Egg family, the extended generalized Gamma being its first branch and the extended generalized inverse Gaussian its second branch.

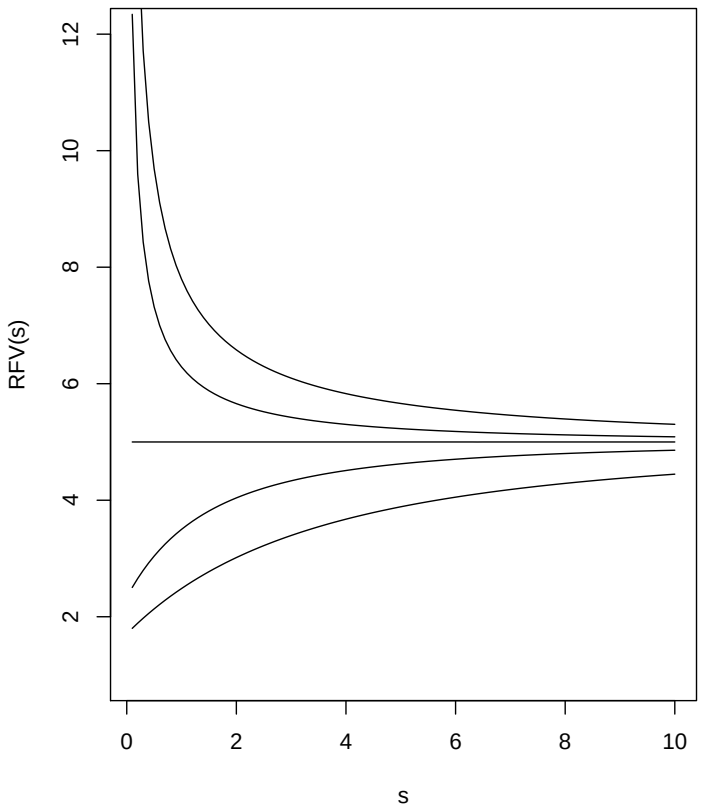

Figure 2: Relative frailty variance functions from the Kummer family. All lines have $\alpha = 0.2$. From top to bottom, (β, δ, ν) equal: $(1, 0.01, 0.01)$, $(1, 0.3, 0.01)$, $(0, 0.05, 1)$, $(-1, 0.3, 1)$, and $(-1, 0.15, 1)$.

The Egg family includes a wide range of distributions, notably the Gamma, the inverse Gaussian, the reciprocal Gamma, the truncated Normal, the Weibull and, in the limit as $\alpha \to \infty$, the Lognormal distributions. The following proposition establishes that the family is well-defined.

Proposition 6 Both branches of the Egg family are well-defined with finite positive moments. ■

The next proposition describes the behaviour of $RFV(s)$ for the second branch as $s \to \infty$.

Proposition 7 The relative frailty variance $RFV(s)$ for the second branch of the Egg

family declines to zero as $s \to \infty$. ■

Figure 3 shows some examples of relative frailty variance functions from this family.

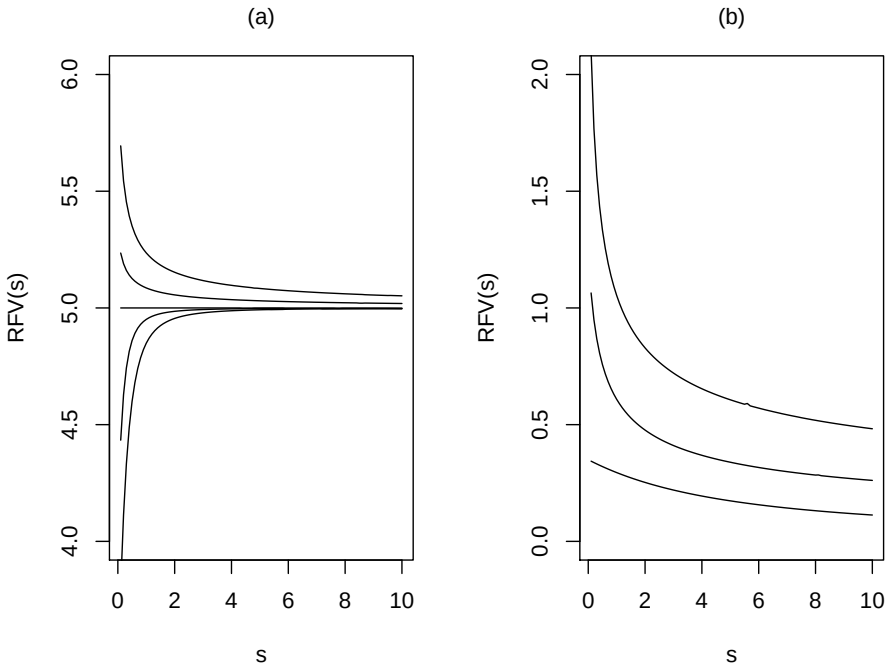

Figure 3: Relative frailty variance functions from the Egg family. (a): First branch with $\alpha = 0.2$ and $\nu = 0.6$. From top to bottom, (β, δ) equal: $(0.7, 0.8)$, $(0.7, 0.2)$, $(1, 2)$, $(2, 0.1)$ and $(2, 2)$. (b): Second branch with $\nu = 0.6$. From top to bottom, (α, β, δ) equal: $(-2, -0.3, 3)$, $(-2, -0.5, 2)$ and $(2, -3, 1)$.

4.3 Scaling the Kummer and Egg families

In practical applications, it is convenient to work with frailties of unit mean, so as to separate the frailty model from the model for the baseline hazards. The Kummer and Egg families may be reparameterized to include a scale parameter which can then be used to constrain the mean to be one.

We begin with the Kummer density. Setting $\theta = \delta$ and $\lambda = \nu\delta$, the density (19) may be expressed as

$$f(u;\alpha,\beta,\theta,\lambda) = \frac{\theta^{-1}}{\Gamma(\alpha)U(\alpha,\alpha-\beta+1,\lambda)}(u/\theta)^{\alpha-1}e^{-\lambda(u/\theta)}(1+(u/\theta))^{-\beta}.$$

It follows that

$$E(U) = \theta \times \frac{\alpha U(\alpha+1,\alpha-\beta+2,\lambda)}{U(\alpha,\alpha-\beta+1,\lambda)}.$$

Hence, restricting θ to equal

$$\theta = \frac{U(\alpha,\alpha-\beta+1,\lambda)}{\alpha U(\alpha+1,\alpha-\beta+2,\lambda)}$$

will ensure that the frailty has unit mean.

A similar trick is possible with both branches of the Egg family. Setting $\theta = \delta^{-1/\beta}$ and $\lambda = \nu\delta^{-1/\beta}$, the density (20) may be expressed as

$$f(u;\alpha,\beta,\theta,\lambda) = \frac{(u/\theta)^{\alpha-1}e^{-\lambda(u/\theta)}e^{-(u/\theta)^{\beta}}}{I^*(\alpha,\beta,\theta,\lambda)},$$

where

$$I^*(\alpha,\beta,\theta,\lambda) = \int_0^\infty (u/\theta)^{\alpha-1}e^{-\lambda(u/\theta)}e^{-(u/\theta)^{\beta}}\mathrm{d}u$$

is the normalizing constant. It is easily seen that

$$I^*(\alpha,\beta,\theta,\lambda) = \theta I^*(\alpha,\beta,1,\lambda)$$

and that

$$E(U) = \theta \times \frac{I^*(\alpha+1,\beta,1,\lambda)}{I^*(\alpha,\beta,1,\lambda)},$$

from which it follows that the frailty has unit mean provided that

$$\theta = \frac{I^*(\alpha,\beta,1,\lambda)}{I^*(\alpha+1,\beta,1,\lambda)}.$$

Note that the reparameterization of the Kummer family can be expressed in the same way, with, in this case,

$$I^*(\alpha,\beta,\theta,\lambda) = \theta\Gamma(\alpha)U(\alpha,\alpha-\beta+1,\lambda).$$

For both the Kummer and the Egg families, the relative variance function then has the form

$$RFV(s) = \frac{I^*(\alpha+2,\beta,1,\lambda+s\theta)I^*(\alpha,\beta,1,\lambda+s\theta)}{I^*(\alpha+1,\beta,1,\lambda+s\theta)^2} - 1,$$

where θ is the function of (α,β,λ) described above.

4.4 Mixed frailty models

The power variance family contains mixed distributions, which have an atom at zero and are continuous elsewhere. Models of this type, which we shall call mixed frailty models, are important as they represent a situation in which a proportion of individuals are not susceptible, and thus comprise a subset of the wider class of cure models (Aalen *et al.* (2008), page 244). However, the compound Poisson models all have increasing relative frailty variance. Is it possible to find a family of cure models with decreasing as well as increasing relative frailty variances?

We show that no such family can exist. Suppose that the frailty U, with mean μ, has a mixed distribution with $P(U = 0) = p$ and density $f(u)$ for $u > 0$, with Laplace transform $L^*(s)$ and mean μ^*. Thus, $\mu = (1-p)\mu^*$. The probability p can depend on the parameters of $f(u)$, as in the compound Poisson models (this is required, for example, if $\mu = 1$), or not, as in the distribution proposed by Longini and Halloran (1996). The Laplace transform of U is

$$L(s) = p + (1-p)L^*(s)$$

and so

$$K(-s/\mu) = \log\{p + (1-p)\exp(K^*(-s/\mu))\}.$$

Using expression (3), it follows that the relative frailty variance is

$$a(s) = \{1 + a^*(s/(1-p))\}\{1 + \frac{p}{1-p}L^*(s/\mu)^{-1}\} - 1,$$

where $a^*(s)$ is the relative frailty variance corresponding to the frailty with density $f(u)$. Thus

$$a(s) \geq \frac{p}{1-p}L^*(s/\mu)^{-1},$$

which tends to ∞ as $s \to \infty$. This proves the following proposition.

Proposition 8 Suppose that U is a frailty with non-zero probability $p = P(U = 0)$. Then the relative frailty variance of U cannot be monotone decreasing on the interval $[a, \infty)$ for any $a \geq 0$. ■

The relative frailty variance functions for mixed frailty models need not be monotone increasing: an example with a bathtub shape (thus answering a query in Paik *et al.*

(1994) about the existence of such profiles) is shown in Figure 4, in which the continuous component is inverse Gaussian. Proposition 8 implies that, if the relative frailty variance is known to be monotone decreasing, a mixed frailty model is inappropriate. Note that such models are not the only ones giving rise to frailties with bathtub-shaped relative frailty variances: these shapes can also arise from continuous frailty distributions, for example the truncated Cauchy and other Student t distributions.

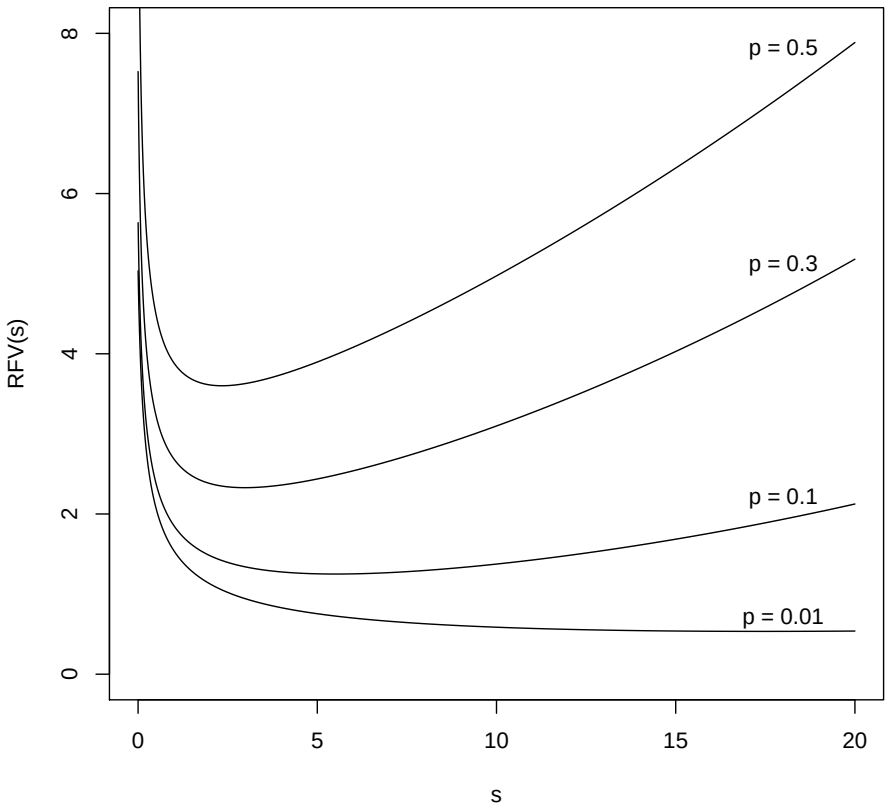

Figure 4: Relative frailty variance functions for mixed frailty models with atom $p = P(U = 0)$ and inverse Gaussian continuous component with $\nu = 0.1$.

5 Time-varying frailties

So far it has been assumed that the frailty U does not vary over time. Suppose now that the hazard at time t is

$$\lambda(t; U(t)) = U(t)\lambda(t)$$

where $U(t)$ is a non-negative stochastic process and $\lambda(t)$ is a deterministic function. As in the time-invariant case, one may define the relative frailty U_t in survivors at time t by

$$U_t = \frac{U(t)}{E(U(t)|T > t)}.$$

Thus the relative frailty variance can be defined as

$$RFV^*(t) = \frac{\mathrm{var}(U(t)|T > t)}{E(U(t)|T > t)^2}.$$

The key difference with the time-invariant case is that there is no longer a canonical time scale determined solely by the cumulative hazard: calendar time t plays a role as well. This is illustrated in the following example.

Example 5: Lévy processes

Lévy processes provide a very rich framework in which to model a stochastic hazard function. Consider the following simple example of such a process with hazard $\lambda(t; U(t)) = U(t)\lambda(t)$ where $U(t)$ is a subordinator with Laplace exponent $\Phi(z)$. Using results from Aalen *et al.* (2008), pages 440-441, we obtain

$$RFV^*(t) = \frac{-\int_0^t \Phi''(\Lambda(t) - \Lambda(u))\mathrm{d}u}{\left[\int_0^t \Phi'(\Lambda(t) - \Lambda(u))\mathrm{d}u\right]^2}.$$

Unlike the $RFV^*(t)$ for time-invariant frailties in (2), it is not possible to rescale the time axis so as to remove dependence on the baseline hazard $\lambda(t)$.

In consequence, there is no natural equivalent to the rescaled $RFV(s)$. Nevertheless, the relative frailty variance $RFV^*(t)$ remains a useful index of how the heterogeneity of the population of survivors varies over time, even though selection effects are no longer the sole cause of this variation.

A simple alternative to modelling the hazard with a stochastic process, while retaining much flexibility, is to assume that $U(t) = g(U_1, U_2, ..., U_k, t)$ for some function g which

may involve parameters to be estimated, where U_1,...,U_k are time-invariant frailties of unit mean. In this model, in which the interpretation of the U_j should preferably be motivated by subject-matter considerations, the frailties are modulated over time in a deterministic fashion. For example, consider the additive frailty model

$$U(t) = \sum_{j=1}^{k} h_j(t) U_j, \qquad \sum_{j=1}^{k} h_j(t) = 1, \quad 0 \leq h_j(t) \leq 1, \quad j = 1, \ldots, k. \tag{21}$$

The restrictions on the functions $h_j(t)$ ensure that $U(t)$ is non-negative with unit mean. Suppose that the U_j are independent with Laplace transforms $L_j(s)$. For a baseline hazard function $\lambda(t)$, the population survivor function is

$$S(t) = \prod_{j=1}^{k} L_j(H_j(t)),$$

where $H_j(t) = \int_0^t h_j(s)\lambda(s)$.

In bivariate shared frailty models, the relative frailty variance can be extended to the relative frailty covariance,

$$RFV^*(t_1, t_2) = \frac{\mathrm{cov}(U(t_1), U(t_2) | T_1 > t_1, T_2 > t_2)}{E(U(t_1)|T_1 > t_1) E(U(t_2)|T_2 > t_2)},$$

which reduces to $RFV^*(t)$ when $t_1 = t_2 = t$. However, time-dependent frailty models are not in general Archimedean, so the cross-ratio function $\theta^*(t_1, t_2)$ cannot generally be expressed as a function of $S(t_1, t_2)$ (Oakes 1989). In consequence, there is no general relationship between $\theta^*(t_1, t_2)$ and $RFV^*(t_1, t_2)$. In one important special case of the additive model, however, such a relationship exists.

Example 6: Piecewise Gamma frailty

Paik *et al.* (1994) proposed a piecewise frailty model with a nested structure. Ignoring the nesting, the model is piecewise Gamma on disjoint intervals $I_j = (t_{j-1}, t_j]$. Thus

$$U(t) = \sum_{j=1}^{k} U_j I_j(t),$$

where the U_j are independent Gamma variables with unit mean and variance γ_j^{-1}, and $I_j(t) = 1$ if and only if $t \in I_j$. In this model,

$$\theta^*(t_1, t_2) = 1 + RFV^*(t_1, t_2),$$

and

$$\begin{aligned} RFV^*(t_1, t_2) &= \gamma_j^{-1} \quad \text{if } t_1 \text{ and } t_2 \text{ belong to the same interval } I_j, \\ &= 0 \quad \text{otherwise.} \end{aligned}$$

Thus, in this specific example the cross-ratio function and the relative frailty (co)variance remain closely related.

An alternative to the additive model (21) is the multiplicative model

$$U(t) = \prod_{j=1}^{k} (1 + (U_j - 1)h_j(t)), \qquad 0 \leq h_j(t) \leq 1, \quad j = 1, \ldots, k. \tag{22}$$

This model compounds the frailties in a manner which may in some settings be more natural than its additive counterpart. However, neither the additive model (21) nor the multiplicative model (22) have simple expressions for the relative frailty variance, which involve an interplay between selection effects and temporal effects. This relationship is elucidated in the following simple case.

Example 7: A simple time-varying frailty model

Suppose that

$$U(t) = 1 + (U - 1)h(t),$$

where U is a time invariant frailty of unit mean and $0 \leq h(t) \leq 1$. This model is a special case of (21) with $k = 2$ and $\text{var}(U_2) = 0$, and also a special case of (22). The support of $U(t)$ is contained in $(1 - h(t), \infty)$.

Let $RFV^*(t)$ denote the relative frailty variance of $U(t)$, and $RFV_0^*(t)$ the relative frailty variance of U. Also, let

$$\mu(t) = E(U|T > t)$$

be the mean of U in survivors at time t. Direct calculation yields

$$RFV^*(t) = RFV_0^*(t) \times \left\{ \frac{h(t)}{1 + \mu(t)^{-1}(1 - h(t))} \right\}^2. \tag{23}$$

The first term in (23) describes the impact of selection effects induced by U on the relative frailty variance. The second term depends largely on the modulating effect of $h(t)$. If $RFV^*(t)$ is observed to vary with t, then this could be due to the fact that the relative

frailty variance of U, $RFV_0^*(t)$, varies with t, or, alternatively, to the fact that $h(t)$ is not identically 1, or both. Note that if U is Gamma (and hence $RFV_0^*(t)$ is constant) and $h(t)$ is non-decreasing, then $RFV^*(t)$ may still decrease, since $\mu(t)$ is a decreasing function.

Time-varying frailty models of the type proposed in (21) and (22) allow great flexibility in the relative frailty variance. However, as shown in Example 7, the distinct contributions of selection effects and temporal variation in the frailty are not identifiable from the relative frailty variance. Subject-matter knowledge, or additional data, are required to make inferences about their relative contributions.

6 Applications

Serological data are obtained by testing blood serum residues for the presence of antibodies to one or more infections. For each infection, a positive (negative) test result indicates prior infection (susceptibility to infection), giving rise to current status data. We illustrate the methods developed in this paper with two datasets, collected as part of two large surveys undertaken in the United Kingdom (Osborne *et al.* 2000).

6.1 Description of the data

Toxoplasma and *Helicobacter pylori* infections

This survey was undertaken in 1996. A total of 3,632 individuals aged 1-84 years were tested for antibodies to Toxoplasma and *H. pylori* infection. Both infections are transmitted by oral ingestion of contaminated matter, and heterogeneity in hygiene levels is likely to result in association between the two infections. Such heterogeneity is most appropriately represented by a continuous frailty distribution.

Unkel and Farrington (2010) introduced a new association measure ϕ for use with current status data that tracks the relative frailty variance, and which therefore can serve as a diagnostic tool for suggesting suitable families of frailty distributions. With the present datasets, the measure ϕ gives results not dissimilar to the log odds ratio, but is more readily interpretable.

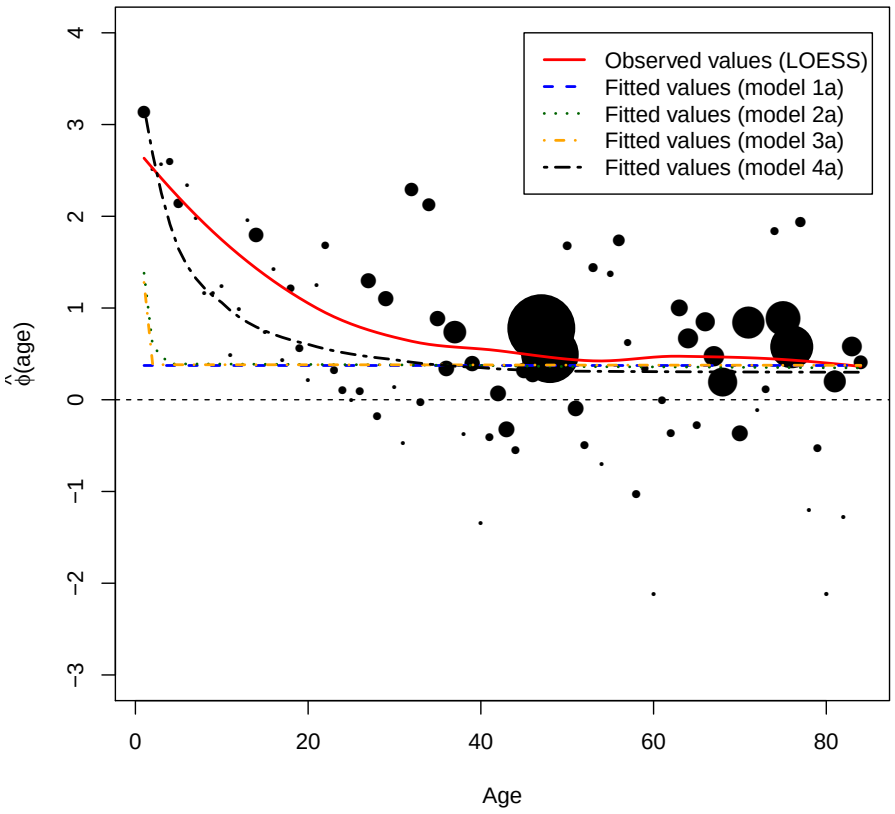

Figure 5: Toxoplasma and *H. pylori* infection data: observed and fitted values of ϕ by age (horizontal dashed line at zero: no association).

The LOESS smoother superimposed on the plot of $\hat{\phi}$ in Figure 5 (with plotting symbols proportional to the precisions) suggests that the heterogeneity in the survivor population may be decreasing towards a positive asymptote. This could be due to a selection effect caused by a time-invariant frailty model, or to temporal variation of the frailty itself. With respect to the former, the data could be represented by a shared frailty model with the frailty either being Kummer distributed with $\beta > 0$, or Egg (branch 1) distributed with $\alpha > 0$ and $0 < \beta < 1$.

Parvovirus B19 and Cytomegalovirus infections

This survey was undertaken in 1991. A total of 1331 individuals aged 1-96 years were tested for antibodies to B19 and CMV infection. As the data are sparse at older ages, only the 1268 cases aged 1-44 are included in the analysis. B19 is transmitted by the respiratory route, whereas CMV is transmitted by ingestion of bodily secretions. In childhood, different transmission routes are confounded to some degree, so correlations are to be expected at younger ages owing to heterogeneity of individual social contact intensities. Such heterogeneity is again most appropriately represented by a continuous frailty distribution.

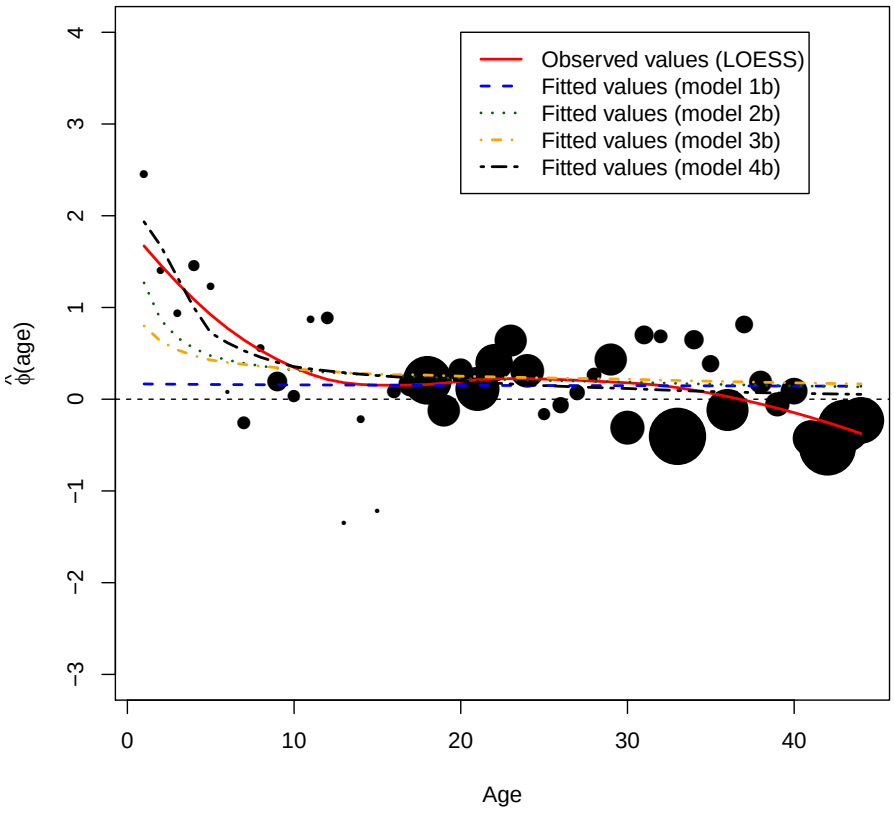

Figure 6: B19 and CMV infection data: observed and fitted values of ϕ by age (horizontal dashed line at zero: no association).

The LOESS smoother superimposed on the plot of $\hat{\phi}$ in Figure 6 suggests that the

heterogeneity in the survivor population may be decreasing towards zero. An appropriate choice for the distribution of the frailty could therefore be a member of the PVF family with $-1 < \alpha < 0$, that is, a Hougaard distribution, or the second branch of the Egg family with α unrestricted and $\beta < 0$. Alternatively, a time-varying frailty (with age as the time line) may be appropriate.

6.2 Fitting frailty models to paired serological survey data

The paired data at each year of age t is multinomial of order four, with index equal to the number tested for both infections at age t, the four counts being cross-classified according to infection status at t for the two infections. For some individuals, data on only one infection were available: in this case the data are binomial. The overall likelihood is product multinomial, and involves the parameters for the baseline hazards $\lambda_j(t)$ $(j = 1, 2)$, together with those required to specify the frailty. The function `nlm` in R version 2.12.1 (R Development Core Team 2010) is used to minimise the deviance. The `integrate` function is used to evaluate the normalizing constants and hence the Laplace transforms for members of the Kummer and Egg families. The baseline hazards $\lambda_j(t)$ $(j = 1, 2)$ were parameterized as piecewise constant on age classes chosen on epidemiological grounds. We investigated sensitivity to the parameterisation of the baseline hazards. Other choices (including continuous baseline hazards) resulted in worse fits but did not alter the optimal choices of frailty distributions. All computer code is available upon request.

Toxoplasma and *Helicobacter pylori* infections

We fitted shared frailty models with the following frailty distributions: Gamma with unit mean and shape parameter γ, Kummer, and Egg (first branch), both with the reparameterisation described in Subsection 4.3, and the time-varying multiplicative model presented in equation (22) with two components U_j $(j = 1, 2)$. For this last model, U_1 and U_2 are independent Gamma variables with unit mean and shape γ_1 and γ_2, respectively, $h_1(t) = \exp(-(t/\rho)^2)$, where ρ is an unknown constant, and $h_2(t) = 1$. The rationale for this model is that the first component represents transient heterogeneity in childhood and the second represents persistent heterogeneity due to differences in hygiene in adulthood.

The baseline hazards were modelled with eight age classes. Results of the analyses are given in Table 1.

Table 1: Results for Toxoplasma and *H. pylori* infection data.

Model	Frailty distribution	Frailty density parameter estimates	deviance	df	p-value	AIC
1a	Gamma	$\hat{\gamma} = 2.2091$	412.0591	366	.0484	8504.62
2a	Kummer	$\hat{\alpha} = 5.4375$ $\hat{\beta} = 6.5920$ $\hat{\lambda} = 0.3578$	408.1992	364	.0547	8504.76
3a	Egg 1st branch	$\hat{\alpha} = 2.5865$ $\hat{\beta} = 0.5936$ $\hat{\lambda} = 0.5124$	408.5725	364	.0533	8505.13
4a	2-component multiplicative double Gamma	$\hat{\gamma}_1 = 0.0093$ $\hat{\rho} = 2.1540$ $\hat{\gamma}_2 = 2.9230$	393.2316	364	.1400	8489.79

The lowest deviance and lowest AIC are obtained with the time-varying 2-component multiplicative double Gamma model (model 4a). At any given age t, let $(s00, s01, s10, s11)$ denote the observed vector of proportions for the paired data, where s00 denote the proportion of individuals susceptible to both infections, $s01$ denote the proportion of individuals uninfected by infection 1 but previously infected by infection 2, and so on. Figure 7 shows the observed proportions together with the fitted proportions $(f00, f01, f10, f11)$ superimposed. Only the lines corresponding to the highest and lowest deviances are presented; these lines are virtually indistinguishable.

The plot provides no guidance on how well the models match the observed association pattern. This is best achieved by comparing the observed and fitted values of the association measure ϕ in Figure 5. This shows that none of the time-invariant frailty models with positive asymptotes are sufficiently flexible to give a good fit to the observed association pattern. The fitted curve that most closely resembles the observed pattern corresponds to the time-varying frailty model.

The observed pattern may also be consistent with heterogeneity declining slowly to

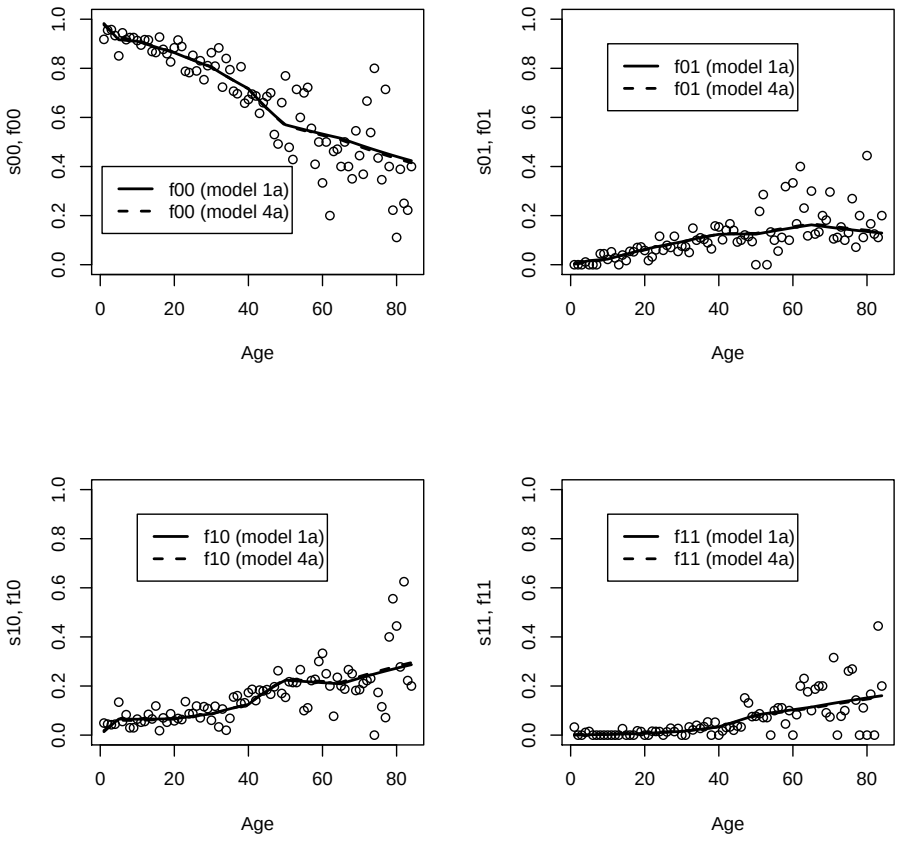

Figure 7: Toxoplasma and *H. pylori* infection data: observed (circles) and fitted values by age.

zero, beyond the range of the data. To investigate this possibility we also fitted PVF and Egg branch 2 models. The PVF model provided an acceptable fit (deviance 402.4101 on 365 d.f., AIC 8496.97) but the Egg branch 2 model did not (deviance 411.1689 on 364 d.f., AIC 8507.73). Overall, the best model is the 2-component time-varying model.

Parvovirus B19 and Cytomegalovirus infections

As with the previous example, the baseline hazards are parameterized by means of a piecewise constant model. Since the data only comprise ages 1-44, five age classes are used. Both the Hougaard and Egg branch 2 families include the inverse Gaussian, so the data are fitted first by a shared inverse Gaussian frailty model with unit mean and

variance γ^{-1}. We also fitted PVF and Egg branch 2 models, and the 1-component time-varying frailty model described in Example 7, in which the single frailty U is Gamma distributed with unit mean and variance γ^{-1}, and $h(t) = \exp\{-(t/\rho)^2\}$. In this setting, the single latent variable U represents transient heterogeneities in childhood. Table 2 summarizes the results.

Table 2: Results for B19 and CMV infection data.

Model	Frailty distribution	Frailty density parameter estimates	deviance	df	p-value	AIC
1b	inverse Gaussian	$\hat{\gamma} = 5.4696$	244.7048	209	.0457	8739.09
2b	PVF	$\hat{\alpha} = -0.8938$ $\hat{\nu} = 1.1644 \times 10^{-12}$	232.9842	208	.1129	8729.37
3b	Egg 2nd branch	$\hat{\alpha} = -1.6981$ $\hat{\beta} = -161.3055$ $\hat{\lambda} = 0.0025$	238.6739	207	.0648	8737.06
4b	1-component Gamma	$\hat{\gamma} = 0.1434$ $\hat{\rho} = 4.1778$	230.5422	208	.1357	8726.93

The best model in terms both of the deviance and the AIC is the time-varying 1-component Gamma model (model 4b). The results for model 2b indicate that the PVF model is being optimized close to the boundary of the parameter space, where $\alpha = -1$ and $\nu = 0$. There is an identifiability issue in this region: it is difficult to obtain reliable estimates of ν and the baseline hazards separately. The results for model 3b indicate that the fitted density is close to the limiting distribution obtained when $\alpha < -1$ and $\beta \to \infty$ and $\lambda \to 0$, namely the Pareto distribution with shape parameter $-\alpha$ on the interval $[1 + \alpha^{-1}, \infty)$.

Figure 6 shows the observed and fitted association patterns for models 1b-4b. As before, the marginal plots of observed and fitted values (not shown) are far less sensitive to the choice of frailty model than the plot of the association measure. For completeness, we also fitted Kummer and Egg branch 1 models; neither provided good fits (deviances in excess of 244.8 both on 207 d.f., AIC in excess of 8743). Overall, the best fit is provided by the 1-component time-varying model.

7 Final comments

This paper has focused on the relative frailty variance as a tool for identifying, classifying, and generating new frailty models. While many of the results apply to univariate frailty models, the methods come into their own as tools for data analysis in the context of shared frailty models, for which the frailty is, in principle, identifiable through its relationship with the cross-ratio function. While we have only considered bivariate survival data, the methods could be applied to pairwise investigations in higher dimensional settings, as a prelude to fitting more ambitious models.

A major theme of the paper has been to identify what shapes of relative frailty variances are possible. While we have suggested some new families of frailty distributions, the applications demonstrate that more work is needed to obtain further frailty distributions with more flexible relative variance functions, particularly with non-zero asymptotes.

For inference, we propose using homogeneous families of frailties, the choice between families of continuous, mixed and discrete models best being left to subject-matter considerations, if possible. The identification of frailty models, nevertheless, remains a tricky problem. Our methods provide no guidance on whether patterns exhibited by the relative frailty variance result from selection processes underpinned by a time-invariant frailty, or from temporal variation of the frailty variance itself. Nevertheless, exploring the possible patterns consistent with the first explanation may, we believe, help to throw light on the plausibility of the second.

Acknowledgements

This research was supported by a project grant from the Medical Research Council. We thank Dr Richard Pebody of the Health Protection Agency, London, who supplied the serological data. We also thanks two anonymous referees and an Associate Editor, as well as the Joint Editor, for constructive comments.

Appendix

Proof of Proposition 5. Using the transformation $y = (\nu + s)u$ gives, for $k \geq 0$,

$$\begin{aligned} I(\alpha + k, \zeta, \nu + s) &= \frac{1}{(\nu + s)^{\alpha+k}} \int_0^\infty y^{\alpha+k-1} e^{-y} t\big(\frac{y}{\nu + s}; \zeta\big) \mathrm{d}y \\ &\to \frac{c(\zeta)\Gamma(\alpha + k)}{(\nu + s)^{\alpha+k}} \qquad \text{as } s \to \infty, \end{aligned}$$

by the monotone convergence theorem and the fact the t is monotone. Hence

$$\begin{aligned} RFV(s) &= \frac{I(\alpha + 2, \zeta, \nu + s/\mu) I(\alpha, \zeta, \nu + s/\mu)}{I(\alpha + 1, \zeta, \nu + s/\mu)^2} - 1 \\ &\to \frac{\Gamma(\alpha + 2)\Gamma(\alpha)}{\Gamma(\alpha + 1)^2} - 1 \qquad \text{as } s \to \infty \\ &= \frac{1}{\alpha}. \end{aligned}$$

Proof of Proposition 6. Let $g(u) = u^{\alpha-1} e^{-\nu u} \exp -\delta u^\beta$. We show that $I(\alpha, \beta, \delta, \nu) = \int_0^\infty g(u) \mathrm{d}u$ is finite for the stated parameter ranges.

(i) For $\alpha > 0$, $g(u)$ is dominated by $u^{\alpha-1} e^{-\nu u}$ and so is integrable.

(ii) For $\alpha < 0$ and $\beta < 0$,

$$\begin{aligned} \int_0^\infty g(u) \mathrm{d}u &\leq \int_0^\infty u^{\alpha-1} \exp -\delta u^\beta \mathrm{d}u \\ &= \frac{1}{-\beta} \int_0^\infty v^{\frac{\alpha}{\beta}-1} e^{-\delta v} \mathrm{d}v, \qquad \text{substituting } v = u^\beta, \end{aligned}$$

and this is finite since $\alpha/\beta > 0$.

(iii) For $\alpha = 0$ and $\beta < 0$, the integral of $g(u)$ is dominated by those in (ii) for $0 < u < 1$ and by those in (i) for $1 \leq u < \infty$.

Since $I(\alpha, \beta, \delta, \nu) < \infty$ for all α, $I(\alpha + k, \beta, \delta, \nu) < \infty$ for all $k \geq 0$, so all positive moments are finite.

Proof of Proposition 7. Suppose that $\beta < 0$. Let

$$f(u) = \log g(u) = (\alpha - 1) \log(u) - \nu u - \delta u^\beta$$

and let u_0 be its maximum. This is the solution of

$$\nu u^{1-\beta} - (\alpha - 1) u^{-\beta} + \delta\beta = 0.$$

The solution has the form

$$u_0 = \big(-\frac{\delta\beta}{\nu}\big)^{\frac{1}{1-\beta}}\{1+o(1)\},$$

where $o(1)$ denotes a quantity tending to zero as $\nu \to \infty$. The Laplace approximation, for large ν, to the integral $I(\alpha,\beta,\delta,\nu) = \int_0^\infty e^{f(u)}\mathrm{d}u$ is obtained by expanding $f(u)$ about u_0 to second order:

$$f(u) \simeq f(u_0) - \frac{(u-u_0)^2}{2\sigma_0^2},$$

where

$$\begin{aligned}\sigma_0^2 &= -\frac{1}{f''(u_0)} \\ &= \frac{u_0^2}{(\alpha-1)+\delta\beta(\beta-1)u_0^\beta}.\end{aligned}$$

The Laplace approximation works in spite of the fact that $u_0 \to 0$ because $(u_0/\sigma_0)^2 \to \infty$ as $\nu \to \infty$, yielding

$$\begin{aligned}I(\alpha,\beta,\delta,\nu) &\simeq \sqrt{2\pi\sigma_0^2}\quad \exp(f(u_0)) \\ &= u_0^\alpha \times \sqrt{2\pi(\sigma_0/u_0)^2} \times e^{-\nu u_0 - \delta u_0^\beta}.\end{aligned}$$

As $\nu \to \infty$,

$$u_0 \simeq \big(-\frac{\delta\beta}{\nu}\big)^{\frac{1}{1-\beta}}$$

and

$$\sigma_0^2 \simeq \frac{u_0^{2-\beta}}{\delta\beta(\beta-1)}.$$

It follows that, for large ν, the integral $I(\alpha,\beta,\delta,\nu)$ is approximated to first order in ν^{-1} by

$$I(\alpha,\beta,\delta,\nu) \simeq \Big(-\frac{\delta\beta}{\nu}\Big)^{\frac{\alpha}{1-\beta}} R(\beta,\delta,\nu),$$

where the function R does not involve α. Thus as $s \to \infty$,

$$I(\alpha+k,\beta,\delta,\nu+s) \to \Big(-\frac{\delta\beta}{\nu+s}\Big)^{\frac{\alpha+k}{1-\beta}} R(\beta,\delta,\nu+s),$$

and so, from (18), $RFV(s) \to 0$ as $s \to \infty$.

References

Aalen, O. O. and Tretli, S. (1999). Analyzing incidence of testis cancer by means of a frailty model. *Cancer causes and control* **10**, 285 - 292.

Aalen, O. O., Borgan, Ø., and Gjessing, H. K. (2008). *Survival and Event History Analysis*. New York: Springer.

Abramovitz, M. and Stegun, I. A. (1972). *Handbook of Mathematical Functions*. New York: Dover.

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

Balakrishnan, N. and Peng, Y. (2006). Generalized gamma frailty model. *Statistics in Medicine*, **25**, 2797 - 2816.

Clayton, D. G. (1978). A model for association in bivariate life tables and its application in epidemiological studies of familial tendency in chronic disease incidence. *Biometrika*, **65**, 141 - 151.

Duchateau, L. and Janssen, P. (2008). *The Frailty Model*. New York: Springer.

Feller, W. (1966). *An Introduction to Probability Theory and Its Applications*, Volume II. New York: Wiley.

Hougaard, P. (1984). Life table methods for heterogeneous populations: Distributions describing the heterogeneity. *Biometrika*, **71**, 75-83.

Hougaard, P. (1986). Survival models for heterogeneous populations derived from stable distributions. *Biometrika*, **73**, 387 - 396.

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

Johnson, N. L., Kotz, S. and Balakrishnan, N. (1994). *Continuous Univariate Distributions*, Volume I. Second edition. New York: Wiley.

Kemp, A. W. (1979). Convolutions involving binomial pseudo-variables. *Sankhyā, Series A*, **41**, 242 - 243.

Longini, I. M. and Halloran, M. E. (1996). A frailty mixture model for estimating vaccine efficacy. *Journal of the Royal Statistical Society, Series C*, **45**, 165 - 173.

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

Osborne, K., Gay, N. J., Hesketh, L. M., Morgan-Capner, P., and Miller, E. (2000). Ten years of serological surveillance in England and Wales: methods, results, implications and action. *International Journal of Epidemiology*, 29, 362 - 368.

Paik, M. C., Tsai, W.-Y. and Ottman, R. (1994). Multivariate survival analysis using piecewise Gamma frailty. *Biometrics*, **50**, 975 - 988.

R Development Core Team (2010). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Unkel, S. and Farrington, C. P. (2010). A new measure of time-varying association for shared frailty models with bivariate current status data, Open University Statistics Group Technical Report 10/14, available from http://statistics.open.ac.uk/technical_reports.

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