UNIVERSITY OF Southampton
Southampton Statistical Sciences Research Institute

Working Paper M09/08
Methodology

Proportional Hazards Models With Discrete Frailty

Chrys Caroni, Martin Crowder, Alan Kimber

Abstract

We extend proportional hazards frailty models for lifetime data to allow a negative binomial, Poisson, Geometric or other discrete distribution of the frailty variable. This might represent, for example, the unknown number of flaws in an item under test. Zero frailty corresponds to a limited failure model containing a proportion of units that never fail (long-term survivors). Ways of modifying the model to avoid this are discussed. The models are applied to a previously published set of data on failures of printed circuit boards and to new data on breaking strengths of samples of cord.

Proportional hazards models with discrete frailty

Chrys Caroni[1], Martin Crowder[2] and Alan Kimber[3]

1. National Technical University of Athens, Greece
2. Imperial College London, England
3. University of Southampton, England

Address for correspondence:

C.Caroni, Dept of Mathematics

National Technical University of Athens

Tel: +30 210 7721707

Email: ccar@math.ntua.gr

Abstract

We extend proportional hazards frailty models for lifetime data to allow a negative binomial, Poisson, Geometric or other discrete distribution of the frailty variable. This might represent, for example, the unknown number of flaws in an item under test. Zero frailty corresponds to a limited failure model containing a proportion of units that never fail (long-term survivors). Ways of modifying the model to avoid this are discussed. The models are applied to a previously published set of data on failures of printed circuit boards and to new data on breaking strengths of samples of cord.

Keywords: reliability analysis; lifetime data; proportional hazards; frailty models; discrete distribution; limited failure model

NOTATION

c_i censoring indicator for unit i

D likelihood ratio statistic summed across groups

E s-expectation

G_z probability generating function of random variable Z

h_b baseline hazard function

ℓ, ℓ', ℓ'' log-likelihood function and its first and second derivatives, respectively

n_c number of censored observations in the sample

n_u number of uncensored observations in the sample

q_k probability that a unit has frailty (such as number of flaws) equal to k

S unconditional survival or reliability function

S_b baseline survival or reliability function

t_c fixed (Type I) censoring time

var s-variance

$\mathbf{x}, \mathbf{x}'$ vector of covariates and its transpose

Z frailty (such as the number of flaws in a unit)

β parameter vector of coefficients associated with covariate vector $\mathbf{x}$

λ parameter of Poisson distribution

$\hat{\lambda}$ maximum likelihood estimator of λ

ν, π parameters of the negative binomial distribution

ξ scale parameter of Weibull distribution

τ shape parameter of Weibull distribution

1 Introduction

A substantial part of the extensive literature on lifetime data analysis concerns so-called frailty models, which introduce heterogeneity between the lifetime distributions of individual units by means of an unobserved individual random effect, the frailty. In the standard proportional hazards framework that we shall adopt here, the effect of an individual frailty z is to change a baseline hazard function $h_b(t)$ to $zh_b(t)$ for that unit. The corresponding survivor or reliability function, conditional on z, becomes

$$S(t \mid z) := \mathrm{P}(T > t \mid z) = \exp\left\{-z\int_0^t h_b(s)ds\right\} = S_b(t)^z \tag{1.1}$$

where $S_b(t)$ is the baseline survivor function. The unconditional survivor function, $S(t)$, can be obtained by integrating (1.1) over the distribution of Z, once a frailty distribution has been specified. Published work on these models generally assumes that Z is a non-negative, continuous random variable. Frequently-used frailty distributions include the gamma [1] and positive stable [2].

In some circumstances, it is appropriate to consider discretely-distributed frailty, for example, when heterogeneity in lifetimes arises because of the presence of a random number of flaws in a unit or because of exposure to damage on a random number of occasions. Although the possibility of a discrete frailty distribution has been mentioned in the literature, it has not been investigated in detail. For example, Xue and Brookmeyer [3] stated that their main result also holds for discrete frailty distributions, as well as continuous ones, but did not pursue that line any further. Moreover, most such references turn out to consider finite mixtures, in which Z is a group or stratum indicator taking just a few values, rather than having a probability distribution over a wider range as we envisage here.

We have presented some initial work on discrete frailty distributions elsewhere [4, 5]. In this context we suppose that Z can take non-negative integer values, *i.e.* Z has a discrete distribution on $\{0, 1, 2, \ldots\}$ rather than a continuous distribution on $(0, \infty)$. The proportional hazards model then gives hazard function $kh_b(t)$ for a unit with $Z = k$. For example, for a unit subject to flaws, this supposes that the flaws act independently each with the same hazard, $h_b(t)$. The well-known Jelinski and Moranda [6] model for software reliability has this structure.

Let the probability distribution of Z be specified by $\mathrm{P}(Z = k) = q_k$ for $k = 0, 1, 2, \ldots$. Then, assuming proportional hazards, the unconditional survivor function of T is given by

$$S(t) = \mathrm{E}\{S_b(t)^Z\} = \sum_{k=0}^{\infty} q_k S_b(t)^k = G_Z\{S_b(t)\} \tag{1.2}$$

where G_Z is the probability generating function of Z. The case $k = 0$, which entails $\mathrm{P}(T > t \mid Z = 0) = 1$ for all t, will be addressed in detail below. As is usual, we assume here that the frailty of a unit does not change over time, *i.e.* Z is fixed 'at birth'. Also, we consider mostly parametric models for the q_k.

Standard discrete distributions such as the geometric, Poisson or negative binomial can be considered as models for the number of flaws in a unit. For the geometric distribution, with parameter $\pi \in (0, 1)$, $q_k = \pi^k(1 - \pi)$, which gives

$$S(t) = (1 - \pi)/\{1 - \pi S_b(t)\} \tag{1.3}$$

For the Poisson distribution, with parameter $\lambda > 0$, $q_k = \mathrm{e}^{-\lambda}\lambda^k/k!$ and then

$$S(t) = \exp\left[-\lambda\{1 - S_b(t)\}\right] \tag{1.4}$$

For the negative binomial distribution, with parameters $\nu > 0$ and $\pi \in (0, 1)$,

$$q_k = \binom{k + \nu - 1}{k} \pi^k (1 - \pi)^\nu$$

and then

$$S(t) = [(1 - \pi)/\{1 - \pi S_b(t)\}]^\nu \tag{1.5}$$

Taking $\nu = 1$ in (1.5) gives the geometric distribution (1.3), and the Poisson form (1.4) is recovered when $\nu \to \infty$ with $\pi = \lambda/\nu$. Thus, the negative binomial can be applied as an extended model in assessing goodness-of-fit of the geometric and Poisson distributions. Also worth noting when applying these models is that the geometric distribution has a heavier tail than the Poisson distribution: for the former, $q_{k+1}/q_k = \pi$ (fixed), whereas for the latter, $q_{k+1}/q_k = (k+1)^{-1}\lambda \to 0$ as $k \to \infty$.

The preliminary work [4, 5] drew attention to the difficulty of obtaining models in which the case of homogeneous frailty is a natural special case of the general model in which frailty varies randomly across units. Consider for a moment the continuous case. To achieve identifiability with the form $zh_b(t)$ it is often convenient to fix the mean of Z at 1 by a suitable constraint on the parameters of the frailty distribution. This is a reasonable choice because then the unit with mean frailty is the 'standard unit' whose hazard is h_b. Then, allowing $\mathrm{var}(Z)$ to tend to zero, with $\mathrm{E}(Z)$ fixed at 1, gives the no-frailty model. This is possible when the continuous distribution has a scale parameter that governs the variance. However, in general, discrete distributions on the integers do not have scale parameters and so the same approach cannot be applied.

The purpose of the present note is to set out some tractable discrete-frailty models. In Section 2, maximum likelihood estimation of discrete frailty models is outlined. In Section 3 various ways of accomodating zero frailty are considered. Some numerical results for two applications are presented in Sections 4 and 5. Concluding remarks are made in Section 6.

2 Estimation

Suppose that the data consists of a random sample $\{(t_i, c_i) : i = 1, \ldots, n\}$, where c_i is the censoring indicator, taking value 1 for an observed lifetime and 0 for a right-censored one; we assume here that the censoring is uninformative. Then the log-likelihood function for a model of specified form with parameter vector θ is

$$\ell(\theta) = \sum_{i=1}^{n} \{c_i \log f(t_i) + (1 - c_i) \log S(t_i)\}, \quad (2.1)$$

where $f(t) = -dS(t)/dt$ is the probability density function of T. Maximum likelihood estimates can be obtained by applying a standard function-optimisation routine to $\ell(\theta)$. In the applications described below a Matlab program to implement the BFGS algorithm [7] has been employed; derivatives were computed by differencing rather than relying on code to reflect their algebraic forms.

Standard errors for the parameters may be derived from the inverse Hessian matrix evaluated at the maximum likelihood estimate, $\ell''(\hat{\theta})^{-1}$. An alternative form, which is often more reliable numerically, is $\{\sum_{i=1}^{n} \ell_i'(\hat{\theta})\ell_i'(\hat{\theta})^{\mathrm{T}}\}^{-1}$, where $\ell_i(\hat{\theta})$ is the log-likelihood contribution from the ith case; this is guaranteed to be positive semi-definite in spite of rounding errors.

In some situations information might be more directly available on the frailty distribution. For example, suppose that frailty is the number of flaws in a unit. It might be possible to ascertain Z for a sample of units before or after failure. The contribution to the log-likelihood from a unit known to have k flaws is $\log\{q_k S_b(t)^k\}$ if still unfailed at time t, and $\log\{q_k f_k(t)\}$ if it failed at time t, where $f_k(t) = -dS_b(t)^k/dt$.

There is no difficulty in introducing covariates into the models. Thus, q_k and $S_b(t)$ can be modified to $q_k(x)$ and $S_b(t; x)$, where x is the vector of covariates. For example, in the geometric frailty distribution π may be expressed in logit-linear form: $\log\{\pi/(1-\pi)\} = x^{\mathrm{T}}\alpha$. Likewise, a log-linear model, $\log \xi = x^{\mathrm{T}}\beta$, may be used in a baseline exponential survival model of mean ξ.

3 Accommodating zero frailty

The unconditional survivor function (1.2) can be written as

$$S(t) = q_0 + \sum_{k=1}^{\infty} q_k S_b(t)^k. \tag{3.1}$$

Frailty distributions that allow $q_0 > 0$ can thus generate units with zero frailty. For such units, the proportional hazards model entails zero hazard, *i.e.* $S_b(t)^0 = 1$ for all t. This can be taken to describe long-term survivors, units that will never fail. In the medical context, such individuals are immune from or cured of the illness in question. Lifetime models with this feature have been used widely [8]. They are also known as limited failure models [9] and as split-population models [10].

Depending on the context, a model that allows zero risk of failure for some units might be unrealistic. We might take the pragmatic viewpoint that the model will be applied to data over a limited time span, so 'immortality' just means that such units have a negligible chance of failing within this period. However, with a parametric model for the q_k, such as the Poisson, the ratio of q_0 to other q_k is constrained and might then be inappropriate for the data. We now present some alternative strategies for dealing with this problem.

In certain circumstance it might be reasonable to modify the frailty distribution to exclude $Z = 0$, *i.e.* force $q_0 = 0$. One simple way of achieving this is to take $Z = 1 + W$, where W is distributed on $\{0, 1, 2, \ldots\}$, say with probabilities $\mathrm{P}(W = k) = r_k$; *e.g.* $r_k = \mathrm{e}^{-\lambda}\lambda^k/k!$. Then, with $q_k = \mathrm{P}(Z = k)$ as before,

$$q_0 = 0 \quad \text{and} \quad q_k = r_{k-1} \quad \text{for} \quad k = 1, 2, \ldots \tag{3.2}$$

Alternatively, one can simply truncate the distribution:

$$q_0 = 0 \quad \text{and} \quad q_k = r_k/(1 - r_0) \quad \text{for} \quad k = 1, 2, \ldots \tag{3.3}$$

However, forcing $q_0 = 0$ is sometimes too drastic since it means that all units, without exception, must have at least one flaw. This can be avoided by treating q_0 essentially as a separate parameter:

$$\mathrm{P}(Z=0)=q_0 \quad \text{and} \quad q_k=(1-q_0)r_{k-1} \quad \text{for} \quad k=1,2,\ldots \tag{3.4}$$

A different type of strategy for accommodating $Z=0$ as a separate case is to introduce a distinct hazard function, $h_0(t)$. It can be used in one of two ways, so that the hazard function for a unit with k flaws becomes either

$$h_0(t)+kh_b(t) \quad \text{for} \quad k=0,1,2,\ldots \tag{3.5}$$

or

$$h_0(t) \quad \text{for} \quad k=0, \quad kh_b(t) \quad \text{for} \quad k=1,2,\ldots \tag{3.6}$$

In (3.5) we assume that all units, those with flaws and those without, are susceptible to an additional cause of failure that operates independently of the flaws: the model is one of independent competing risks. In (3.6) the failure mechanisms for units with and without flaws are quite separate. When $h_0(t)=0$ both (3.5) and (3.6) reduce to the basic form $kh_b(t)$ for $k=0,1,2,\ldots$. When $h_0(t)=h_b(t)$ the 'zero-flaws hazard' is the same as the hazard per fault and then (3.5) gives $S_k^*(t)=\{S_b(t)\}^{k+1}$, which is equivalent to replacing Z by $1+W$, as in (3.2). The conditional survivor functions resulting from (3.5) and (3.6) are

$$S_k^*(t):=\mathrm{P}(T>t\mid k)=S_0(t)S_b(t)^k \quad \text{for} \quad k=0,1,2,\ldots$$

and

$$S_0^*(t)=S_0(t), \quad S_k^*(t)=S_b(t)^k \quad \text{for} \quad k=1,2,\ldots,$$

where $S_0(t)$ is the survivor function corresponding to $h_0(t)$, and the corresponding unconditional survivor functions are

$$S_0(t)G_Z\{S_b(t)\} \quad \text{and} \quad q_0\{S_0(t)-1\}+G_Z\{S_b(t)\}$$

4 Application 1

Meeker and LuValle [11] presented data on an accelerated test of circuit boards. They stated that these printed circuit boards fail because of the growth of conductive filaments through what should be insulating material. The data comprise four groups of circuit boards, tested at different relative humidities, with hours to failure (short-circuit) recorded under accelerated-life conditions. In Table 1 $rh\%$ is the relative humidity, n_u and n_c are the numbers of observed (uncensored) and right-censored failure times, and t_c is the fixed (Type I) censoring time in hours.

Table 1: Summary data on printed circuit board failures [11]

Group	rh %	n_u	n_c	t_c
1	49.5	22	48	4078
2	62.8	57	11	3067
3	75.4	70	0	–
4	82.4	70	0	–

One of the models fitted in [11] was a limited-failure population model, that is, a mixture of a proportion p of defective units with a Weibull distribution of lifetimes and a proportion $1-p$ of non-defective units that are not susceptible to failure. That paper went on to develop a more elaborate model of the processes involved, but when fitting the limited-failure model the authors suggested that the postulated defective units could be units with cracks. In that case a model in which the discrete frailty is the unknown number cracks in a circuit board appears to be worth considering. In particular, we will compare results between geometric, Poisson and negative binomial frailty distributions, in conjunction with a Weibull baseline survival distribution. Table 2 gives the resulting maximised log-likelihoods, for the three frailty distributions. Incidentally, the observed failure times given in [11] are intervalcensored in relatively narrow intervals. Our computations were performed both by treating the times as observed values falling at the mid-points of the intervals and, more precisely, by replacing the densities in the likelihood function by differences in the survivor functions at the end points. The two methods gave effectively the same results.

Table 2: Maximised log-likelihoods for three frailty models with Weibull baseline survival

Group	Frailty distribution		
	Geometric	Poisson	Neg. binomial
1	-177.525	-177.694	-176.070
2	-357.914	-360.253	-353.645
3	-309.790	-303.320	-303.320
4	-247.007	-240.857	-240.857

The results in Table 2 do not give an immediate indication of which model fits best. For groups 1 and 2 the geometric log-likelihood is slightly better than that for the Poisson, but for groups 3 and 4 the Poisson is substantially better; for groups 1 and 2 the negative binomial looks better than the Poisson, but loses this advantage for groups 3 and 4. A standard log-likelihood ratio test between the geometric and negative binomial models yields $\chi^2_4 = 36.69$ ($p < 0.001$) for the four groups combined. In the comparison between Poisson and negative binomial models, it must be taken into account that the parameter value giving the Poisson case ($\nu = \infty$) is on the boundary of the parameter space. For one group, minus twice the difference in maximized log-likelihoods should be assessed by reference to $\frac{1}{2}\chi^2_1$ because it takes the value zero with probability $\frac{1}{2}$ and has a χ^2_1 distribution with probability $\frac{1}{2}$ [12]. Adding these statistics over the four groups gives sum $D = 16.46$. The associated p-value can be found by observing that D takes value 0 with probability $(\frac{1}{2})^4$ and has a χ^2_j distribution with binomial probability $\binom{4}{j}\left(\frac{1}{2}\right)^4$. Consequently,

$$\mathrm{P}(D > d) = \sum_{j=1}^{4} \binom{4}{j} \mathrm{P}(\chi_j^2 > d);$$

in this application $\mathrm{P}(D > 16.46) = 0.0005$. On this basis the negative binomial has the most support. However, this result is almost solely determined by group 2. Further, a Poisson process for the occurrence of conductive filaments over time seems credible, in the sense that they occur in different places on the circuit board. Finally, the plots referred to below appear to support the simpler model. For these reasons we will adopt the Poisson/Weibull model.

The parameter vector for the Poisson/Weibull fit is $\theta = (\lambda, \xi, \tau)$, where $S_b(t) = \mathrm{e}^{-(t/\xi)^\tau}$ is the baseline Weibull survivor function. Maximum likelihood estimates are given in Table 3 with standard errors in parentheses. (Log-transformed parameters allow unconstrained optimisation and can improve asymptotic normal approximations for maximum likelihood estimators.) Some of the standard errors for groups 3 and 4 are large, reflecting a rather flat likelihood surface. Figure 4 shows the Kaplan-Meier survivor functions with a 95% pointwise confidence band along with the estimated reliability functions from the negative binomial/Weibull fits. The fits look quite good for all four groups: in fact, the fits shown here in Figure 1 for groups 1, 2 and 3 are somewhat better than those shown in Figure 5 of [11], where group 4 was omitted from the analyses.

For this set of data, the estimated proportions of units without flaws, the $\hat{q}_0$-values, are 0.20, 0.35, 0.00 and 0.09 for the four groups; the last two figures must be regarded as suspect because of the large standard errors associated with the corresponding $\hat{\lambda}$-values. In cases where a high proportion of units is not susceptible to failure, our model will be practically indistinguishable from a simple mixture model of defective and non-defective units. An extreme example is provided by the data used by Meeker [9] to introduce the limited failure model. In a life test of 4156 integrated circuits, only 28 failed up to 593 hours and the test

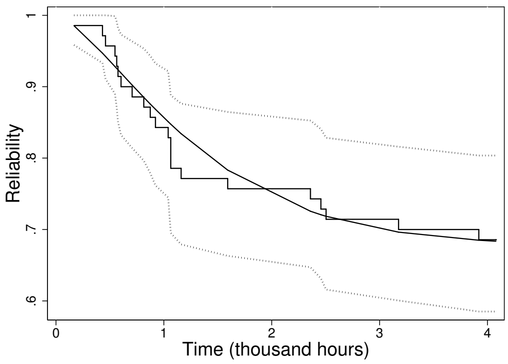

(a) Relative humidity 49.5 %

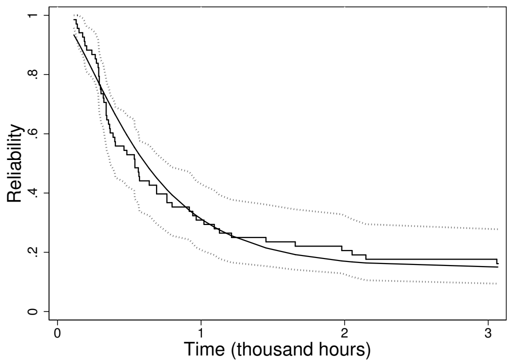

(b) Relative humidity 62.8 %

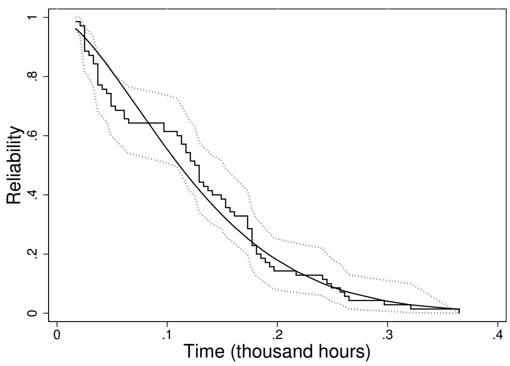

(c) Relative humidity 75.4 %

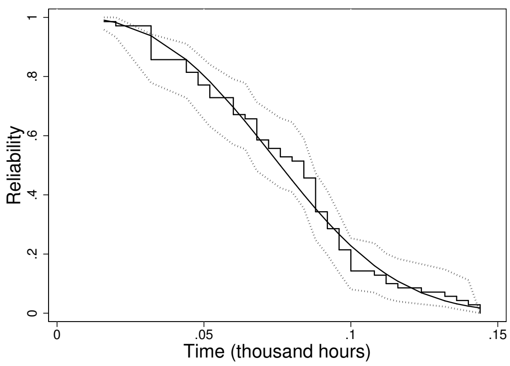

(d) Relative humidity 82.4 %

Figure 1: Estimated reliability function (continuous curve) compared to Kaplan-Meier estimate (step function) with 95 % pointwise confidence band (broken curves), in each experimental group of circuit boards.

was stopped at 1370 hours with all the remaining 4128 times right-censored at this point. The proportion of non-defectives must be close to 4128/4156 (99.3%), giving an estimated Poisson parameter of 0.00676 in our model and hence the expected number of units with two or more flaws is only 0.09 even in this very large sample.

Table 3: Parameter estimates for Poisson/Weibull model

group	$\log\hat{\xi}$	$\log\hat{\tau}$	$\log\hat{\lambda}$
1	0.472 (0.17)	0.365 (0.26)	-0.945 (0.22)
2	0.041 (0.11)	0.402 (0.12)	0.646 (0.15)
3	3.721 (3.97)	0.433 (0.11)	8.762 (6.08)
4	0.887 (2.68)	1.013 (0.10)	9.171 (7.34)

5 Application 2

The data here come from a consultancy problem and full details may not be given for reasons of confidentiality. Briefly, the observations are of strengths of a type of braided cord used in safety netting. In practice, the cord spends much of its working life in the open air. The purpose of the study was to investigate the effects of weathering, or ageing, on the strengths of three different types of cord denoted here as white, red and yellow. A large number of pieces of cord were involved; some were left to weather naturally and others were kept in store throughout the trial. At the end of a specified period the cords were strength-tested to destruction. The data used here correspond to the strengths (in coded units) of cords kept in store. An understanding of the properties of these control data is an important preliminary to a full analysis of the cords weathered in the open air. A simple weakest-link argument, together with previous experience of such materials, suggests that a Weibull baseline distribution is a reasonable assumption. However, two types of heterogeneity were envisaged. In one the occurrence of flaws in a cord will weaken its capacity for load-bearing. In the other, there is variability in quality between cords arising from variations in thickness, poor braiding and so on. It was expected that the yellow cord would not suffer from these problems because it was manufactured to high standards, but the white and red cords might suffer from one or more of these problems. Exploratory analysis of the data showed mild curvature of the log-cumulative hazazd plots for white and red cord but not for yellow.

The first type of heterogeneity may be modelled by a discrete frailty model such as (3.5). The second type of heterogeneity may be modelled using a standard continuous frailty model. A general test for frailty where the frailty distribution has finite variance [13] yielded statistically significant results ($p < 0.05$) for white and red cord but not for yellow. Table 4 shows maximised log-likelihoods resulting from fitting standard Weibull models without frailty and models of type (3.5) with Poisson frailty and Weibull forms for S_0 and S_b. Evidently, model (3.5) only achieves an improvement in fit for the red cord, though even in this case twice the log-likelihood difference is a modest 2.34.

Thus, on the basis of these initial analyses, as expected there is no evidence of a departure from a Weibull model for the yellow cord. There is no evidence of flaws in the white cord but there appears to be variability in quality. However, for the red cord there is a weak indication of the occurrence of flaws along with more general variability in quality.

Table 4: Maximised log-likelihoods for cord data

Group	Model	
	Weibull	(3.5)
white	254.964	254.964
red	253.983	255.154
yellow	245.008	245.008

6 Conclusion

The models considered here provide a basis for frailty models when heterogeneity can be attributed partly to unmeasured discrete-valued factors. In certain circumstances sensible modelling considerations lead to discrete frailty rather than continuous frailty or finite mixtures, the latter having a fixed upper bound on the number of flaws. As suggested in the Introduction, likely situations include cases where a variable number of flaws are present in a unit, or where variable numbers of exposures have caused damage. Another situation is when a variable number of contacts takes place; this might be applicable to times to infection after a disease enters a closed population.

There may be another reason for considering discrete frailty models. For example, frailty is often revealed by the presence of upper outliers relative to a fitted no-frailty survival model such as the Weibull [14]. The full impact of observations with zero flaws may be hidden by right censoring but observations with a relatively large number of flaws will tend to have surpisingly short lifetimes relative to the baseline distribution. In effect, the discrete frailty model can explain apparent lower outliers in the data.

It has to be said that fitting models like (3.5) is not entirely trouble-free. We have run some simulations and found that the likelihood surface can tend to be rather flat, giving rise to numerical problems such as near-singularity of the Hessian matrix. Overall, Nelder-Mead optimisation seems to be more suited to this type of likelihood than quasi-Newton methods such as BFGS [7]: the former seems to be less likely to give up the search prematurely because the gradients are small.

References

[1] J.W. Vaupel, K.G. Manton and E. Stallard, "The impact of heterogeneity in individual frailty on the dynamics of mortality," *Demography*, vol. 16, pp. 439-454, 1979.

[2] P. Hougaard, "Life table methods for heterogeneous populations: distributions describing the heterogeneity," *Biometrika*, vol. 71, pp. 75-83, 1984.

[3] X. Xue and R. Brookmeyer, "Regression analysis of discrete time survival data under heterogeneity," *Statist. Medicine*, vol. 16, pp. 1983-1993, 1997.

[4] C. Caroni, M.J. Crowder and A.C. Kimber, "Discrete frailty models for heterogeneity in lifetime distributions," in *Proceedings of the 17th International Workshop on Statistical Modelling*, M. Stasinopoulos and G. Touloumi, Eds., 2002, pp. 157-162.

[5] C. Caroni, M.J. Crowder and A.C. Kimber, "Discrete frailty," in *MMR 2002: 3rd International Conference on Mathematical Methods in Reliability*, H. Langseth and B. Lindqvist, Eds. Trondheim: NTNU, 2002, pp. 135-138.

[6] Z. Jelinski and P.B. Moranda, "Software reliability research," in *Statistical Computer Performance Evaluations*, W. Freiberger, Ed. New York: Academic Press, 1972, pp. 465-484.

[7] W.H. Press, B.P. Flannery, S.A. Teukolsky and W.T. Vetterling, *Numerical Recipes in C*. Cambridge: Cambridge University Press, 1988.

[8] R. Maller and X. Zhou, *Survival Analysis with Long-term Survivors*. Chichester: John Wiley, 1999.

[9] W.Q. Meeker, "Limited failure population life tests: application to integrated circuit reliability," *Technometrics*, vol. 29, pp. 51-65, 1987.

[10] P. Schmidt and A.D. Witte, "Predicting criminal recidivism using 'split population' survival time models," *J. Econometrics*, vol. 40, pp. 141-159, 1989.

[11] W.Q. Meeker and M.J. LuValle, "An accelerated life test model based on reliability kinetics," *Technometrics*, vol. 37, pp. 133-146, 1995.

[12] S.G. Self and K.Y. Liang, "Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under non-standard conditions," *J. Amer. Statist. Assoc.*, vol. 82, pp. 605-610, 1987.

[13] A.C. Kimber, "A Weibull-based score test for heterogeneity," *Lifetime Data Analysis*, vol. 2, pp. 63-71, 1996.

[14] C. Caroni and A.C. Kimber, "Detection of frailty in Weibull lifetime data using outlier tests," *J. Statist. Comp. Simul.*, vol. 74, pp. 15-23, 2004.

Chrys Caroni is an Associate Professor in the Department of Mathematics, School of Applied Mathematical and Physical Sciences of the National Technical University, Athens, Greece. She obtained her PhD in Statistics from the University of Southampton, England. Her research interests focus on reliability and survival analysis and on methods for detecting outliers.

Martin Crowder recently retired as Professor of Stochastic Modelling in the Department of Mathematics, Imperial College London. He now has an honorary research position in the Institute for Mathematical Sciences at the College. His main areas of research are survival analysis, estimating equations and asymptotic theory of statistics.

Alan Kimber is a Reader in Statistics in the School of Mathematics and the Southampton Statistical Sciences Research Institute at the University of Southampton, UK, having previously been a member of faculty at Sussex, Surrey and Reading Universities. He received his PhD in Mathematical Statistics from the University of Hull, UK in 1981 and became a Chartered Statistician in 1993. His current major research interest is in survival analysis methodology with applications in reliability and medicine.