

Diagnostic tests for frailty

P. Economou and C. Caroni

Department of Mathematics
School of Applied Mathematics and Physical Sciences
National Technical University of Athens
9 Iroon Polytechniou, Zografou
155 80 Athens, Greece
(e-mail: polikon@math.ntua.gr, ccar@math.ntua.gr)

Abstract. A common way of allowing heterogeneity between individuals in models for lifetime data is to introduce an unobservable individual random effect Z . In a proportional hazards framework, the individual's hazard becomes $zh_b(t)$ where $h_b(t)$ is the baseline hazard. The random variable Z is often assumed to follow the Gamma or Inverse Gaussian distribution. We develop here diagnostic tests for these assumptions. One simple graphical diagnostic is based on the form of the unconditional survival function when $h_b(t)$ is assumed to be Weibull. Another plot uses a closure property of a family of frailty distributions, which implies that the frailty among survivors at time t has the same form as the original distribution of Z , with the same shape parameter but different scale parameter. In this method, we estimate the shape parameter at different times t and examine graphically whether it is constant. We give simulation results and examples to illustrate these methods.

Keywords: Lifetime data, frailty, proportional hazards, Burr distribution, Generalized Inverse Gaussian distribution.

1 Introduction

When modelling data obtained from time-to-event studies, it is often found that there is heterogeneity between individuals, over and above what can be accounted for by any available covariates. One common way of allowing for this heterogeneity is to introduce an unobservable individual random effect Z , the so-called frailty. This is usually assumed to operate in a proportional hazards framework, so that it acts multiplicatively on the baseline hazard function h_b which is common to all individuals. Thus the hazard function for an individual with frailty $Z = z$ is given by

$$h(t|z) = zh_b(t)$$

If there are also measured covariates, the model is usually extended to

$$h(t|z; X) = ze^{\beta' \mathbf{x}(t)} h_b(t)$$

where $\mathbf{x}(\cdot)$ is a q -dimensional vector of possibly time-dependent covariates.

Introducing heterogeneity in lifetime data by means of an unobserved quantity in this way was initiated by [Clayton, 1978], [Vaupel *et al.*, 1979] and [Hougaard, 1984]. The distribution of the random variable Z is often assumed to be Gamma [Vaupel *et al.*, 1979] or Inverse Gaussian [Hougaard, 1984]. As with any part of the process of statistical modelling, it is desirable to check that the assumed distribution is supported by the data. The purpose of the present paper is to develop diagnostic plots for the frailty distribution, with the emphasis on these two common choices, the Gamma and Inverse Gaussian. We will be assuming that the baseline hazard function h_b has been specified correctly and that the multiplicative proportional hazards frailty model is the proper one to describe the data.

2 Diagnostic plots for mixtures

From the proportional hazards assumption, it follows that the conditional survivor function for an individual with frailty z is

$$S(t|z) = [S_b(t)]^z$$

where S_b is the baseline survivor function. In particular, if the baseline model is taken to be Weibull (η, β) , then the survivor function conditional on frailty z is

$$S(t|z) = \exp(-zs)$$

where $s = (t/\eta)^\beta$. If Z has distribution function G on $(0, \infty)$, then the unconditional survivor function is given by

$$\int_0^\infty \exp(-zs) dG(z)$$

If G is taken to be Gamma with shape and scale parameters both equal to ν (so that the mean is one), then

$$S(t) = \left(1 + \frac{s}{\nu}\right)^{-\nu} \quad (1)$$

(the Burr distribution), while taking G to be Inverse Gaussian with scale 1 and shape λ yields the survivor function

$$S(t) = \exp\left(\lambda \left(1 - \sqrt{\frac{2s}{\lambda} + 1}\right)\right) \quad (2)$$

(In both cases, a constraint has been applied to the parameters of G to make the model identifiable. Other choices of constraint are possible but make no difference in principle.)

We now look for appropriate diagnostic plots to enable us to check that the assumption of one of these distributions is in fact correct. The idea is to

plot some function of the non-parametric Kaplan-Meier estimate of survival, $\hat{S}(t)$, against some function of t , to obtain the characteristic shape associated with the distributions (1) and (2).

Taking logarithms of (1), and supposing that s/ν is large enough so that $\log(1 + s/\nu) \approx \log(s/\nu)$, we see that $-\log\hat{S}(t)$ against $\log t$ should give a straight line.

[Wolstenholme, 1999] suggests this plot for the Pareto distribution. This is the special case of the Burr distribution when the baseline hazard is exponential, which is a special case of the Weibull ($\beta = 1$). However, we observe that the above approximation is poor for the early failures. These give the plot a characteristic horizontal section, whose length depends on ν , disappearing as ν becomes small (high degree of heterogeneity) (Figure 1). When the frailty




Fig. 1. Diagnostic plot for Burr distribution for various values of the parameter ν . 1000 simulations of samples of size 1000, baseline Weibull parameters 1000 (scale) and 2 (shape). Type I censoring at $t=3000$.

distribution is the Inverse Gaussian, taking logarithms of (2) suggests the plot of $\log(-\log\hat{S}(t))$ against $\log t$ if λ is large. Under these circumstances, there is only a small degree of heterogeneity and the distribution will not be very different from the baseline Weibull distribution. This is the standard diagnostic plot used to check for the Weibull distribution. It gives a straight line with slope equal to the shape parameter β . On the other hand, suppose that λ is large. It is easy to show that in this case the Weibull(η, β) - Inverse Gaussian($1, \lambda$) mixture tends to the Weibull with scale parameter $\eta/(2\lambda)^{1/\beta}$ and shape $\beta/2$. Thus the same plot gives a straight line with slope $\beta/2$. For intermediate values of λ , the plot should be curved with slope falling from β to $\beta/2$ as time increases. Examples are shown in Figure 2.

3 Closure property of the frailty distributions

To develop another kind of diagnostic plot, we start with a closure property of Gamma frailty in the proportional hazards model ([Vaupel *et al.*, 1979]).

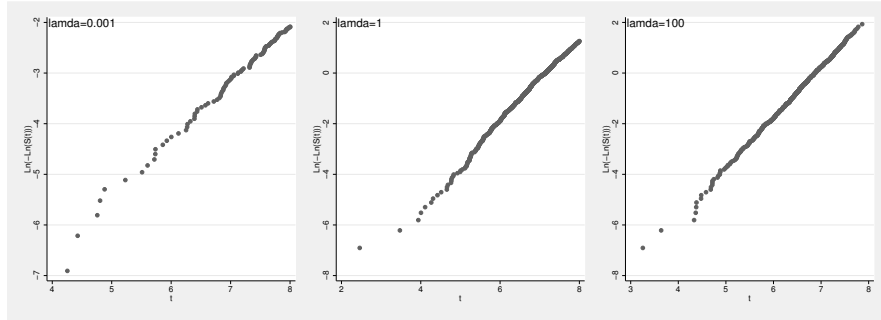


Fig. 2. Diagnostic plot for Weibull-Inverse Gaussian mixture for various values of the parameter λ . Details of simulations as in Figure 1.

Given that the frailty distribution among all individuals is Gamma with scale parameter κ and shape parameter λ , then the frailty distribution among the population of survivors at time t is again Gamma with the same shape parameter λ but different scale parameter given by $\kappa + H_b(t)$, where $H_b(t)$ is the cumulative baseline hazard function. This property can be generalized to a whole family of distributions.

3.1 Generalization of the Gamma frailty property

The Gamma frailty property given by [Vaupel *et al.*, 1979] can be generalized first to the case of available covariates. More specifically, given that the frailty distribution is $\text{Gamma}(\kappa, \lambda)$, then the frailty distribution among survivors at time t , conditional on the value of the covariates, is $\text{Gamma}(\kappa + H_b^x(t), \lambda)$, where $H_b^x(t)$ is given by

$$H_b^x(t) = \int_0^t e^{\beta' \mathbf{x}(u)} h_b(u) du.$$

The closure property is not a characterization of the Gamma distribution only. It is quite easy to show that a similar property holds also for the Inverse Gaussian and the Generalized Inverse Gaussian (GIG). Furthermore, a similar property holds for a whole class of distributions that belong to the exponential family [Hougaard, 1984]. Let frailty Z be a random variable with distribution $F(\alpha)$ on $(0, \infty)$, where α is the parameter vector, with p.d.f. of the form

$$f_z(z) = \frac{e^{-[z, g(z)][\eta_1(\alpha), \eta_2(\alpha)]'}}{\Phi(\alpha)} \xi(z)$$

which is an exponential family distribution with canonical statistics z and $g(z)$ [Shao, 1998]. For this frailty distribution the following theorem holds, which extends Hougaard's result by including covariates.

Theorem 1 *Given the frailty distribution $F(\alpha)$ with p.d.f. as above, then under the proportional hazards frailty model the frailty distribution among survivors at time t is again $F(\cdot)$. The value of $\eta_1(\alpha)$, the element of the parameter vector corresponding to z , changes, but the components of $\eta_2(\alpha)$ do not. More specifically, the p.d.f. of frailty among survivors at time t is given by*

$$f_{Z|T>t}(z) = \frac{e^{-[z,g(z)][\eta_1^*(\alpha),\eta_2(\alpha)]'}}{\Phi^*(\alpha)} \xi(z)$$

where $\eta_1^*(\alpha) = \eta_1(\alpha) + H_b^x(t)$ and $\Phi^*(\alpha) = \Phi(\alpha)S_T(t)$.

The GIG distribution (and hence the Gamma and Inverse Gaussian distributions which it includes) belongs to the above class of the exponential family. Unfortunately, some other distributions, like the lognormal, which are also widely used as frailty distributions, do not belong to this class because they do not have z as a canonical statistic. This obstacle can be overcome by considering a generalized distribution, adding one more parameter [Hougaard, 1986] which will be zero initially. So, the Theorem can be applied to all distributions $F(\alpha)$ with p.d.f given by

$$f_z(z) = \frac{e^{-T(z)[\eta(\alpha)]'}}{\Phi(\alpha)} \xi(z)$$

where $T(z)$ does not contain z as a component, since the above distribution can be seen as a special case of $GF(\alpha, \beta)$ with p.d.f. given by

$$f_z(z) = \frac{e^{-[z,T(z)][\beta,\eta(\alpha)]'}}{\Phi_G(\alpha, \beta)} \xi(z)$$

for $\beta = 0$. $\Phi_G(\alpha, \beta)$ is the integral over the range of z , $R(z)$, of the numerator of the previous relationship. Applying the Theorem to the distribution $GF(\alpha, 0)$ shows that the frailty distribution among the survivors at time t will be again GF but with parameter vector given by $(\alpha, H_b^x(t))$.

3.2 Plots

Given that the frailty distribution has been chosen correctly, then our above Theorem shows that the vector $\eta_2(\alpha)$ of the initial parameters does not change when we restrict our attention to the frailty distribution among those units that have survived until time t . Let $\hat{\eta}_{2i}(\alpha)_{|T>t}$ denote component i of the maximum likelihood estimate of this vector among the survivors at time t . If our assumption of the frailty distribution is correct, then $\hat{\eta}_{2i}(\alpha)_{|T>t}$ for any time t is an asymptotically unbiased estimator of the same quantity $\eta_{2i}(\alpha)$. Therefore, a plot of $\hat{\eta}_{2i}(\alpha)_{|T>t}$ against time should give a straight line parallel to the time axis.

For the Gamma(κ, λ) and Inverse Gaussian(κ, λ) distributions of frailty this plot reduces to $\hat{\lambda}_{|T>t}$ versus t since $\eta_2(\kappa, \lambda) = \lambda$. For the GIG(λ, δ, γ), the proposed plots are of $\hat{\lambda}_{|T>t}$ and $\hat{\delta}_{|T>t}$ against t , since $\eta_2(\lambda, \delta, \gamma) = (\lambda - 1, \delta^2)$ for this distribution.

In all cases, the maximum likelihood estimates of the model's parameters that are required for the plots are obtained by maximising the logarithm of the usual likelihood function for lifetime data ($t_i : i = 1, 2 \dots n$)

$$L = \prod_{i=1}^n \{h(t_i)^{\delta_i} S(t_i)\}$$

where δ_i is the censoring indicator which takes the value 1 if t_i is an observed lifetime and zero if it represents a right censored observation. The expressions for $S(t)$ are given above in (1) and (2) for the Gamma and Inverse Gaussian frailty distributions, respectively, and the hazard function $h(t)$ can be obtained as minus the derivative of $\log S$. We first carry out this estimation using all the data. Then we select a sequence of convenient time points τ_j ($j = 1, 2 \dots k$) and repeat the estimation k times, using in the i th estimation only those data points t_i satisfying $t_i \geq \tau_j$.

3.3 Simulations

To illustrate the method, we simulated a set of 1000 uncensored data points from the Burr distribution (Weibull-Gamma mixture) and produced the above plot based on repeated estimates of ν . Then we fitted the incorrect Weibull-Inverse Gaussian mixture to the same data and produced the corresponding plot (Figure 3). Next we repeated the exercise with the roles of the two distributions reversed. Thus we generated a set of data from the Weibull-Inverse Gaussian mixture and produced the plots for both the correct model and for the incorrect Burr distribution (Figure 4). In both cases, the plots




Fig. 3. Plot defined in Section 3.2, fitting (left) correct Burr distribution, (right) incorrect Weibull-Inverse Gaussian mixture to data generated from Burr ($\nu = 1$).

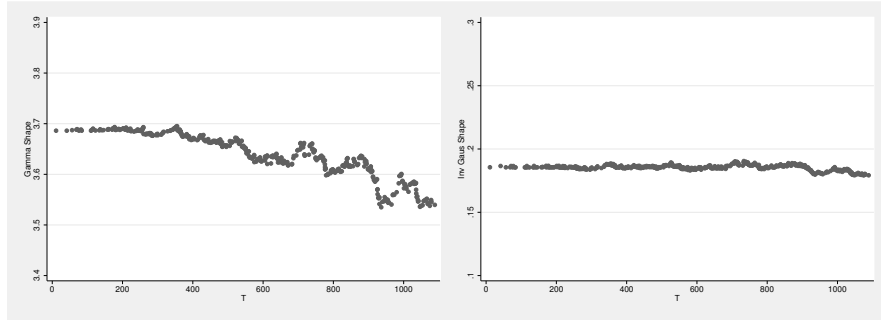


Fig. 4. Plot defined in Section 3.2, fitting (left) incorrect Burr distribution, (right) correct Weibull-Inverse Gaussian mixture to data generated from Weibull-Inverse Gaussian mixture ($\lambda = 0.5$).

discriminate extremely well between the two frailty distribution; the plot for the correct distribution is a horizontal straight line as predicted, but the plot for the incorrect distribution departs clearly from a horizontal line.

4 Example

For a real-data illustration of our methods, we used data on the duration of a treadmill test undertaken by 978 successive patients at a cardiac clinic in Athens. Figure 5 shows the simple diagnostic plots that were developed in Section 2. The plot for the Burr distribution, on the right, has the expected shape of a straight line preceded by a horizontal section. The plot for the Weibull-Inverse Gaussian mixture, on the left, is curved as expected, but the curvature is greater than it should be if this is the correct model. Figure 6 shows the diagnostic plots that were developed in Section 3. These indicate




Fig. 5. Plots defined in Section 2 for data on 978 cardiac patients. Left: Weibull-Inverse Gaussian mixture; right: Burr.




Fig. 6. Plots defined in Section 3 for data on 978 cardiac patients. Left: Weibull-Inverse Gaussian mixture; right: Burr.

clearly that the assumption of a Gamma distribution for frailty is acceptable, because the estimates of its shape parameter at different times are scattered about a horizontal line, but the Inverse Gaussian assumption is not.

5 Further research

Although graphical diagnostics have proved very useful in statistical modelling, it can also be valuable to have formal statistical tests for the presence of any frailty, thus showing whether or not it is necessary to use the models considered here. In these models, one parameter of the distribution controls both the presence and the degree of frailty. For example, when the frailty distribution is $\text{Gamma}(\nu, \nu)$, the parameter ν controls the amount of frailty since $V(Z) = 1/\nu \rightarrow 0$ ($\nu \rightarrow \infty$). If the basic distribution is Weibull and the unconditional distribution is therefore Burr, the presence of any frailty can be examined by testing the null hypothesis $\nu = \infty$ (or $1/\nu = 0$) by likelihood-based methods applied to the Burr distribution. The theory for one of these methods, the score test, was given by [Crowder and Kimber, 1997] for multivariate lifetime data and the details for the univariate case which we are interested in by [Kimber, 1996]. (In fact, the test also holds for other Weibull mixtures, not just for Weibull-Gamma = Burr.) Other likelihood-based tests that can be applied include a Wald test and a likelihood ratio for this parameter. A difficulty that arises is that the null value of the parameter being tested falls on the boundary of the parameter space. In such cases, the distribution of minus twice the log likelihood ratio is not given by the usual chi-squared approximation. Instead, a mixture of chi-squared distributions usually applies. We intend to complete a study of the likelihood ratio test for this model and then carry out a simulation study to compare the properties of the different tests in order to recommend the best one for use in practice.

References

- [Clayton, 1978]D.G. Clayton. 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, 1978.
- [Crowder and Kimber, 1997]M. Crowder and A. Kimber. A score test for the multivariate burr and other weibull mixture distributions. *Scandinavian Journal of Statistics*, 24:419–432, 1997.
- [Hougaard, 1984]P. Hougaard. Life table methods for heterogeneous populations: Distributions describing the heterogeneity. *Biometrika*, 71:75–83, 1984.
- [Hougaard, 1986]P. Hougaard. Survival models for heterogeneous populations derived from stable distributions. *Biometrika*, 73:387–396, 1986.
- [Kimber, 1996]A. Kimber. A weibull-based score test for heterogeneity. *Lifetime Data Analysis*, 2:63–71, 1996.
- [Shao, 1998]J. Shao. *Mathematical Statistics*. Springer-Verlag, New York, 1998.
- [Vaupel *et al.*, 1979]J.A. Vaupel, K.G. Manton, and E. Stallard. The impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography*, 16:439–454, 1979.
- [Wolstenholme, 1999]L.C. Wolstenholme. *Reliability Modelling: A Statistical Approach*. Chapman and Hall/CRC, Boca Raton, 1999.