Paper 1492-2014

Introduction to Frailty Models

John Amrhein, McDougall Scientific Ltd.

ABSTRACT

This paper introduces frailty models and their use in biostatistics to model time-to-event or survival data. Frailty models are survival models with at least one random effect. There are two main types of frailty models, univariate and multivariate. In a univariate model, the random effect pertains to an individual with one observation and accounts for unobserved heterogeneity at the level of the individual. The NLMIXED Procedure fits univariate models. In a multivariate model, the random effect models covariance associated with clustered observations. The clusters might be multiple individuals in a naturally occurring group or multiple measurements within an individual. These models are sometimes referred to as shared frailty models. The PHREG and NLMIXED Procedures fit multivariate frailty models.

INTRODUCTION

Frailty models are survival models with at least one random effect. For example, a proportional hazards model can be written as:

$$\lambda_j(t) = \lambda_0(t)e^{\beta' X_j(t)}$$

The subscript *j* indexes the individual. This model becomes a frailty model by adding a random effect as in:

$$\lambda_{ij}(t) = \lambda_0(t)e^{\beta' X_{ij}(t)+\omega_i} \text{ or } \lambda_{ij}(t) = \lambda_0(t)e^{\omega_i}e^{\beta' X_{ij}(t)}$$

The subscript *i* indexes a cluster of individuals and ω is the random effect.

If you have used a mixed model in a non-survival data situation and you have modeled survival data using, perhaps, a proportional hazards model, then you are familiar with concepts necessary to begin using frailty models.

This paper assumes that you have an awareness of random effects in a general linear model setting and that you are familiar with the basics of survival data; e.g. censoring. The purpose of this paper is to provide enough information to transition you to random effects models for time-to-event, or survival, data. The next section motivates the use of frailty models by describing situations in which a frailty model is a reasonable option. The next section describes two primary types of frailty models. One set of data is used in all examples to facilitate comparison of the models. The concluding section summarizes and mentions other approaches not covered in this paper.

This paper does not discuss different estimation theory and methods but instead focuses on application and interpretation. For more rigorous treatments of frailty models, the interested reader should consult citations under recommended reading.

MOTIVATION FOR FRAILTY MODELS

As already stated, frailty models are time-to-event, or survival, models with at least one random effect. One might say that they are “mixed” survival models in that they will (likely) contain both fixed and random effects. The name “frailty” was introduced by Vaupel (1979) and it is worth reviewing why he chose this name. Vaupel noted that demographic life tables published at the time ignored the fact that some individuals were more susceptible or prone to life-ending events than other individuals. These “frail” individuals tended to die at earlier ages than their more robust peers. Therefore, estimates of population-level hazards would under (over) estimate individual- or group- level hazards for more (less) frail individuals. Population mortality rates might be over-estimated with increasing age because the frail individuals drop out of the risk set.

To account for this heterogeneity in the population, Vaupel suggested the inclusion of a frailty factor in modeling hazards and survival. The frailty factor modifies the hazard multiplicatively. For example, an individual with a frailty of 2 has twice the hazard as the “standard” individual with frailty 1. Likewise, an individual with a frailty of ½ has one-half the hazard of the standard individual.

The frailty factor is, sometimes, a “catch-all” for unobserved covariates that operate on the individual (e.g. nutrition) or unobservable factors of the individual (e.g. hereditary factors).

TYPES OF FRAILTY MODELS

There are two main types of frailty models; univariate and multivariate. In univariate models, the frailty pertains to an individual with one observation. For example, a common response of interest in oncology is disease free survival

(DFS). A common scenario in clinical trials in oncology is one in which each patient's tumors are removed at the start of the trial (baseline), a treatment is administered, and then the patient's disease status is recorded at pre-specified intervals. The time to the first recurrence of the disease is the response of interest.

Multivariate models are used to model responses from clustered observations. The clusters might be multiple individuals in a naturally occurring group (e.g. twins or patients within a clinic) or multiple measurements within an individual (e.g. examination of each eye or repeated measures over time). Shared frailty models are multivariate models in which the units within a cluster are assumed to experience (share) the same level of frailty.

"Multivariate" in the frailty sense refers to more than one observation per cluster or individual and does not refer to data set structure as in an ANOVA vs MANOVA sense. Data set structure for the multivariate frailty model follows that of a linear mixed model that is fit using PROC MIXED; i.e. multiple rows per individual or cluster.

A cluster-level frailty can be used to account for within-cluster correlation in a multivariate model. As in mixed models methods, this can be done using a population-averaged model (similar to modeling "R-side" effects) or subject-specific models (similar to modeling "G-side" effects). Examples of multivariate models in the next section will illustrate each of these methods.

FITTING FRAILTY MODELS USING SAS/STAT®

Table 1 summarizes SAS® procedures that you *might* consider for modeling time-to-event data. Each procedure's capabilities are listed with respect to censoring and type of frailty modeling.

SAS® Procedure	Model Description	Censoring	Univariate Frailty	Multivariate Models (clustered or repeated data)
PHREG	Cox Proportional Hazard Model	Right	No	Population-averaged (marginal) model via robust variance estimation Shared frailty model via RANDOM statement
LIFEREG	Parametric Accelerated Failure Time Model	Left, right, interval	No	No
MIXED	Linear Mixed Model	None	No	Population-averaged (marginal) model via R-side variance modeling Subject-specific (conditional) model via G-side variance modeling
GENMOD	Generalized Linear Model	Discrete-time survival model	No	Population-averaged (marginal) via GEEs and robust variance estimation
GLIMMIX	Generalized Linear Mixed Model	Discrete time survival model	Yes	Population-averaged (marginal) model via R-side variance modeling and robust variance estimation Subject-specific (conditional) model via G-side variance modeling
NLMIXED	Nonlinear Mixed Model	Program appropriate log-likelihood	Yes	Subject-specific (conditional) model via G-side variance modeling
MCMC	Bayesian Models	Program appropriate log-likelihood	Yes	Subject-specific (conditional) model via G-side variance modeling

Table 1. Frailty model capabilities of selected SAS/STAT® Procedures

You can eliminate PROCs LIFEREG and MIXED from consideration; LIFEREG cannot incorporate a frailty factor and MIXED cannot handle censored observations in an acceptable manner. PROCs GENMOD and GLIMMIX are limited to the discrete-time survival model for handling censored observations, but they can easily handle multiple random effects. See Allison (2010) for discussions on discrete time methods. Using robust variance estimation is a method to account for correlated observations and is an alternative to frailty models. However, models using robust variance estimation are not, themselves, frailty models. The remaining 3 procedures, PHREG, NLMIXED, and MCMC are flexible procedures able to handle one or both types of frailty models. The next section demonstrates how to fit a univariate frailty model using PROC NLMIXED. The succeeding section demonstrates fitting a shared frailty model using PHREG and NLMIXED.

PROC MCMC is beyond the scope of this introductory paper on frailty models. However, you are encouraged to read *Example 59.16 Piecewise Exponential Frailty Model* under The MCMC Procedure in SAS/STAT® 13.1 User's Guide. Example 59.16 uses the same data set as *Example 71.11 Analysis of Clustered Data* under The PHREG Procedure in SAS/STAT® 13.1 User's Guide. Between these two examples, you will learn several methods to fit a multivariate frailty model.

UNIVARIATE FRAILTY MODELS

The data for this example originate from Wei, Lin, and Weissfeld (1989), but were obtained from *Example 71.10 Analysis of Recurrent Events Data* under the PHREG Procedure in SAS/STAT® 13.1 User's Guide. Although these data are repeated measures survival data, this example illustrates modeling disease free survival (DFS) in a univariate frailty model using PROC NLMIXED.

The data consist of 86 patients with tumors of the bladder. Each of the patients had all of his/her tumors removed at the initiation of the randomized clinical trial. They were treated with either a placebo or thiotepa in a randomized assignment, Thiotepa is a drug intended to prevent tumor recurrence. The patients were examined at various intervals for tumor recurrence. The data consist of the variables in Table 2.

Alphabetic List of Variables and Attributes					
#	**Variable**	**Type**	**Len**	**Format**	**Label**
1	ID	Num	8		Patient ID
4	NUMBER	Num	8		Baseline number of tumors
5	SIZE	Num	8		Baseline size of tumors
8	STATUS	Num	8		Recurrence status
3	TRT	Num	8	TREAT.	Treatment
2	TSTART	Num	8		Interval start time (months)
7	TSTOP	Num	8		Interval stop time (months)
6	VISIT	Num	8		Visit

Table 2. Variables in the Thiotepa data set

The original data, as will be shown in the next section, were multivariate (in the frailty sense) with 4 visits for each patient. For the DFS model, each patient appears only once; the first visit with tumor recurrence or the last follow-up visit if tumors did not recur. The research question is whether DFS, when treated with thiotepa, differs significantly from treatment with a placebo, while controlling for the initial number and size of tumors.

Before fitting a random effects model using PROC NLMIXED, it is a good idea to first fit the fixed-effects-only model. You should compare the results from the NLMIXED step with the results from a procedure designed to fit fixed-effect models; PROC LIFEREG in this case.

```
proc nlmixed data=thiodfs(where=(tstop ne 0));
   bounds gamma > 0;
   linp  = b0 + b1*trt + b2*number + b3*size;
   alpha = exp(-linp);
   G_t   = exp(-(alpha*tstop)**gamma);
   g     = gamma*alpha*((alpha*tstop)**(gamma-1))*G_t;
   ll    = (status=1)*log(g) + (status=0)*log(G_t);
   model tstop ~ general(ll);
run;

proc lifereg data=thiodfs;
   model tstop*status(0) = number size trt / dist=weibull;
run;
```

Example 68.5 Failure Time and Frailty Model under the NLMIXED Procedure in SAS/STAT® 13.1 User's Guide explains the decomposition of the likelihood in the above NLMIXED step. See Allison (2010, pp. 270-272) for an alternative formulation.

If censoring (right censoring only) is non-informative, then the log likelihood can be split into two terms, one for censored observations (U_c) and one for uncensored observations (U_u).

$$l(\beta, t) = \sum_{i \in U_u} \log g(t_i, \beta) + \sum_{i \in U_c} \log G(t_i, \beta)$$

In the NLMIXED step, each observation's contribution to the log likelihood is calculated on the statement defining the variable LL. The variable STATUS has value 1 when the observation is not censored (the event recurrence is observed), and 0 when it is censored. If you assume that DFS is distributed Weibull, then

$$G(t) = e^{-(\alpha t)^{\gamma}}$$

$$g(t) = \gamma \alpha t^{\alpha - 1} e^{-(\alpha t)^{\gamma}}$$

G(t) is the distribution of survival probabilities beyond time t. g(t) is the density of failure time. These quantities are programmed in the NLMIXED statements defining G_t and g respectively.

The WHERE= option on the PROC statements omits the first observation, which was censored at time 0. Zero-valued times cannot be evaluated for the log likelihood contribution. PROC LIFEREG omits this observation automatically.

Analysis of Maximum Likelihood Parameter Estimates							
Parameter	**DF**	**Estimate**	**Standard Error**	**95% Confidence Limits**		**Chi-Square**	**Pr > ChiSq**
Intercept	1	4.0862	0.4473	3.2095	4.9628	83.46	<.0001
NUMBER	1	-0.3256	0.0986	-0.5188	-0.1325	10.92	0.0010
SIZE	1	-0.0751	0.1345	-0.3388	0.1886	0.31	0.5767
TRT	1	0.7859	0.4023	-0.0025	1.5744	3.82	0.0507
Scale	1	1.2839	0.1575	1.0095	1.6329		
Weibull Shape	1	0.7789	0.0956	0.6124	0.9906		

Table 3. Parameter estimates from PROC LIFEREG for fixed-effects-only model

Parameter Estimates									
Parameter	**Estimate**	**Standard Error**	**DF**	**t Value**	**Pr > |t|**	**Alpha**	**Lower**	**Upper**	**Gradient**
gamma	0.7789	0.09556	85	8.15	<.0001	0.05	0.5889	0.9689	-0.0001
b0	4.0862	0.4473	85	9.14	<.0001	0.05	3.1969	4.9755	-0.00005
b1	0.7859	0.4023	85	1.95	0.0540	0.05	-0.01389	1.5858	-0.00006
b2	-0.3256	0.09856	85	-3.30	0.0014	0.05	-0.5216	-0.1297	-0.00019
b3	-0.07509	0.1345	85	-0.56	0.5782	0.05	-0.3426	0.1924	0.000029

Table 4. Parameter estimates from PROC NLMIXED for fixed-effects-only model

Table 3 and Table 4 show that the parameter estimates match, validating that the log likelihood has been correctly programmed in the NLMIXED step. B0 corresponds to the intercept, B1 to TRT, B2 to NUMBER, B3 to SIZE and GAMMA is the Weibull shape parameter.

The addition of a random effect transforms the model into a frailty model.

```
proc nlmixed data=thiodfs(where=(tstop ne 0));
   parms b0=4.0862 b1=0.7859 b2=-0.3256 b3=-0.07509 gamma=0.7789;
   bounds gamma > 0;
   linp  = b0 + b1*trt + b2*number + b3*size + z;
   alpha = exp(-linp);
   G_t   = exp(-(alpha*tstop)**gamma);
   g     = gamma*alpha*((alpha*tstop)**(gamma-1))*G_t;
   ll    = (status=1)*log(g) + (status=0)*log(G_t);
   model tstop ~ general(ll);
   random z ~ normal(0,exp(2*logsig)) subject=id;
run;
```

The PARMS statement sets the initial values to the fixed-effect-only estimates, which is likely to improve convergence over the default starting values. The random effect is entered as Z at the end of the linear predictor. Its distribution is specified as a normal distribution on the RANDOM statement. A normal distribution is the only distribution currently available in PROC NLMIXED. Table 5 displays the parameter estimates.

Parameter Estimates									
Parameter	**Estimate**	**Standard Error**	**DF**	**t Value**	**Pr > |t|**	**Alpha**	**Lower**	**Upper**	**Gradient**
b0	3.9131	0.5003	84	7.82	<.0001	0.05	2.9181	4.9081	-0.0002
b1	0.6245	0.4080	84	1.53	0.1296	0.05	-0.1869	1.4359	0.000294
b2	-0.3293	0.1110	84	-2.97	0.0039	0.05	-0.5501	-0.1085	-0.00018
b3	-0.09219	0.1348	84	-0.68	0.4959	0.05	-0.3602	0.1759	0.000149
gamma	1.7492	0.6333	84	2.76	0.0071	0.05	0.4899	3.0085	-0.00018
logsig	0.3748	0.1421	84	2.64	0.0100	0.05	0.09211	0.6574	-0.00022

Table 5. Parameter estimates from PROC NLMIXED for a frailty model

Table 6 displays estimate for both the fixed-effects and frailty models for comparison.

	Frailty Model			**Fixed Effects Model**		
Parameter	**Estimate**	**Standard Error**	**Pr > |t|**	**Estimate**	**Standard Error**	**Pr > |t|**
b0	3.9131	0.5003	<.0001	4.0862	0.4473	<.0001
b1	0.6245	0.4080	0.1296	0.7859	0.4023	0.0540
b2	-0.3293	0.1110	0.0039	-0.3256	0.09856	0.0014
b3	-0.09219	0.1348	0.4959	-0.07509	0.1345	0.5782
gamma	1.7492	0.6333	0.0071	0.7789	0.09556	<.0001
logsig	0.3748	0.1421	0.0100			

Table 6. Comparison of parameter estimates from a frailty and fixed effects model

The treatment effect is rendered far less significant than in the fixed effects model indicating that patient to patient variability is an important factor in understanding DFS. The Weibull shape parameter has more than doubled indicating a much different distribution of DFS times than estimated by the fixed effects model.

You can use the UNIVARIATE Procedure to estimate Weibull parameters for the raw survival times as shown in Figure 1. The kernel density estimate of the shape parameter indicates a distribution between an exponential (shape parameter < 1) and a highly skewed lognormal-type distribution (shape parameter > 1). Search “Weibull Distribution” on Wikipedia.org for illustrations of various Weibull shapes.

MULTIVARIATE FRAILTY MODELS

All patient visits (4 each) in the thiotepa data are used to illustrate multivariate (shared) frailty models in this section. Three models are demonstrated; a population-averaged model using PROC PHREG and the robust sandwich estimator of variance, a frailty model using the RANDOM statement in PROC PHREG, and a frailty model using PROC NLMIXED.

The COVS(AGGREGATE) option on the PROC PHREG statement requests robust variance estimation.

```
proc phreg data=thiotepa covs(aggregate);
   model (TStart, TStop) * Status(0) = Trt Number Size;
   id id;
   where TStart < TStop;
run;
```

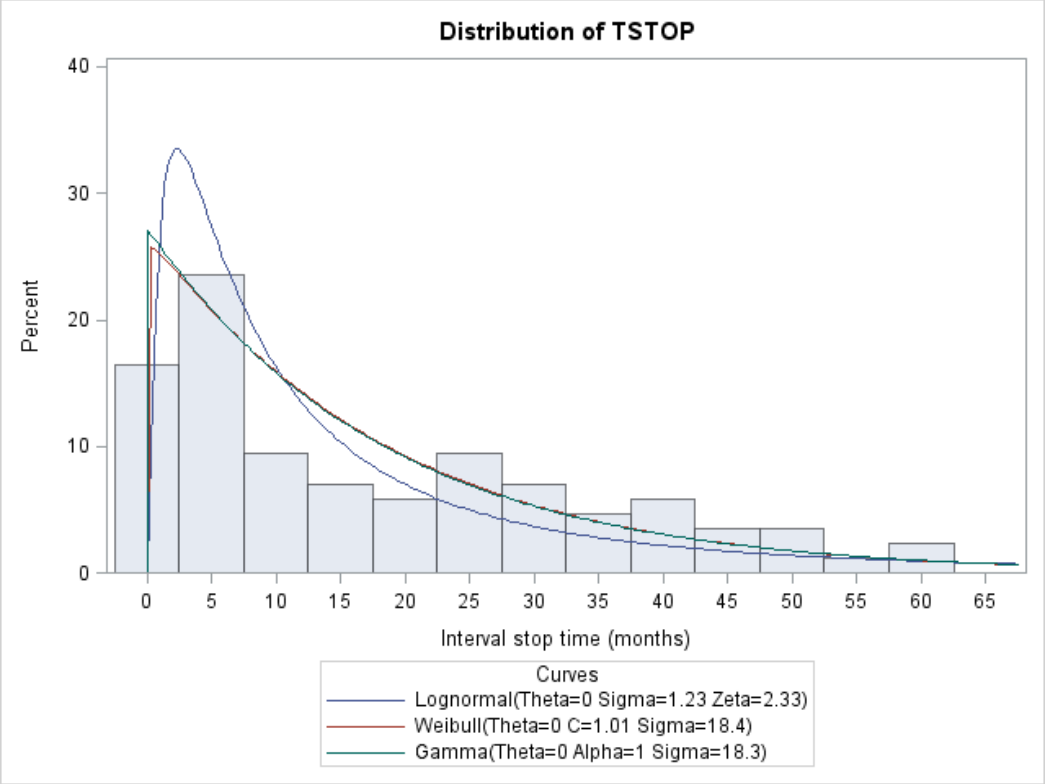

Figure 1. Distribution of DFS months for the univariate frailty model

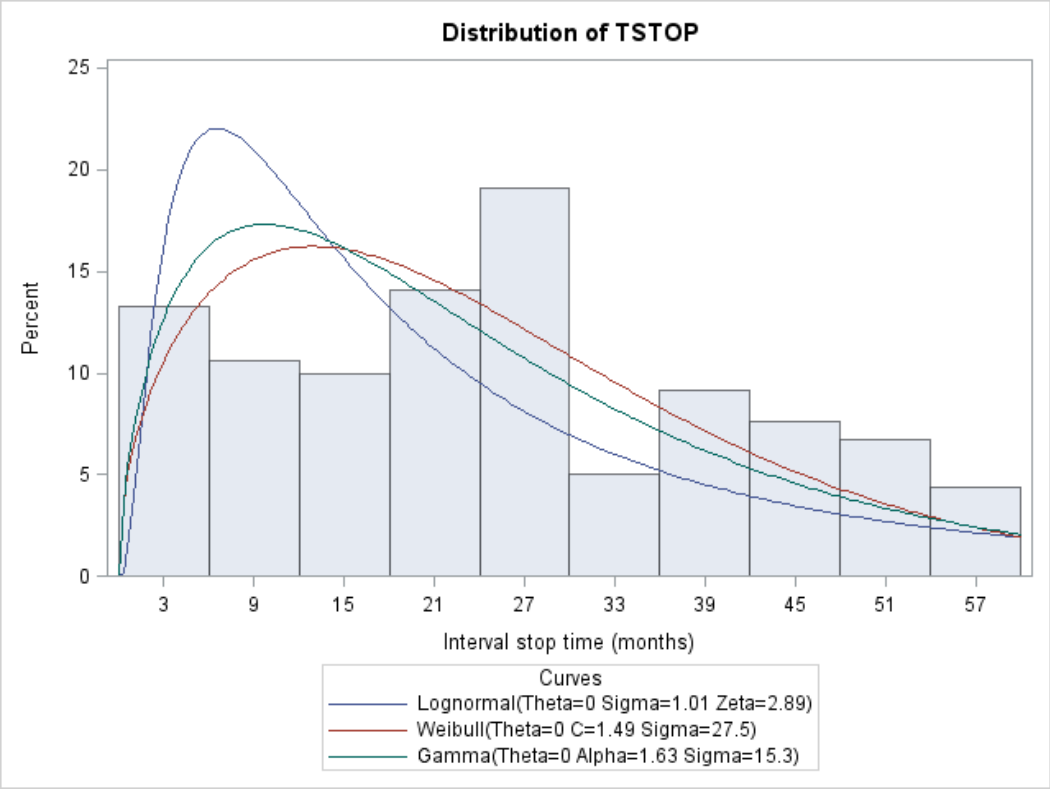

Figure 2. Distribution of recurrence months for the multivariate frailty model

Table 7 displays the parameter estimates for the population-averaged model.

Analysis of Maximum Likelihood Estimates with Sandwich Variance Estimate								
Parameter	**DF**	**Parameter Estimate**	**Standard Error**	**StdErr Ratio**	**Chi-Square**	**Pr > ChiSq**	**Hazard Ratio**	**Label**
TRT	1	-0.45979	0.25801	1.290	3.1757	0.0747	0.631	Treatment
NUMBER	1	0.17165	0.06131	1.296	7.8373	0.0051	1.187	Baseline number of tumors
SIZE	1	-0.04256	0.07555	1.094	0.3174	0.5732	0.958	Baseline size of tumors

Table 7. Parameter estimates from a marginal model using the robust sandwich estimator of variance

This model is referred to as the Proportional Means model. This model is useful if you consider the correlation between repeated measures of DFS to be a mere nuisance. That is, you focus your attention on the test for treatment effect, but want to properly estimate variance for the test by taking into account correlation within patient.

Next, time to recurrence is modeled using a shared frailty.

```
proc phreg data=thiotepa;
   class id;
   model (TStart, TStop) * Status(0) = Trt Number Size;
   id id;
   random id;
   where TStart < TStop;
run;
```

The RANDOM statement adds a normally distributed random effect for the specified variable. In SAS/STAT® 13.1, you can specify the gamma or lognormal distribution using the DIST= option on the RANDOM statement. Table 8 displays the parameter estimates for the shared frailty PHREG model.

Analysis of Maximum Likelihood Estimates							
Parameter	**DF**	**Parameter Estimate**	**Standard Error**	**Chi-Square**	**Pr > ChiSq**	**Hazard Ratio**	**Label**
TRT	1	-0.54590	0.30240	3.2588	0.0710	0.579	Treatment
NUMBER	1	0.21575	0.07987	7.2973	0.0069	1.241	Baseline number of tumors
SIZE	1	-0.02163	0.10157	0.0453	0.8314	0.979	Baseline size of tumors

Table 8. Parameter estimates from a shared frailty model using PROC PHREG

The shared frailty model fit with PROC PHREG, or PROC NLMIXED, is a subject-specific model. By introducing the frailty, you explicitly model unobserved heterogeneity of the individuals. The predicted frailties (e.g. using the SOLUTION option on the RANDOM statement in PROC PHREG), give insight into the magnitude of the heterogeneity. The parameter estimates for the fixed effects have changed from the proportional means model because the likelihood has changed; it now contains the frailty term, which corrects "… some or all of the bias in the coefficients caused by unobserved heterogeneity" (Allison, 2010, p. 270).

Lastly, a shared frailty model is fit using PROC NLMIXED.

```
proc nlmixed data=john.thiotepa(where=(tstop ne 0));
   parms  b0=3.9 b1=0.62 b2=-0.33 b3=-0.09 gamma=1.75 logsig=0.37;
   bounds gamma > 0;
   linp  = b0 + b1*trt + b2*number + b3*size + z;
   alpha = exp(-linp);
   G_t   = exp(-(alpha*tstop)**gamma);
   g     = gamma*alpha*((alpha*tstop)**(gamma-1))*G_t;
   ll    = (status=1)*log(g) + (status=0)*log(G_t);
   model tstop ~ general(ll);
   random z ~ normal(0,exp(2*logsig)) subject=id out=EB;
   predict 1-G_t out=cdf;
run;
```

The PARMS statement specifies the parameter estimates from the DFS model, again, in an attempt to help convergence by providing starting values believed to be close to the final values. The OUT=EB option on the RANDOM statement provides a temporary data set named EB containing the empirical Bayes estimates for the random effect Z. You can review the estimates to gain insight into the variability (unobserved) across patients and also check their distribution against the assumed normal distribution.

The PREDICT statement provides predicted failure times for each patient in a data set named CDF. See *Example 68.5 Failure Time and Frailty Model* under the NLMIXED Procedure in SAS/STAT® 13.1 User's Guide to learn how to use this data set to plot failure times for the sample of patients.

Table 9 displays the parameter estimates for the shared frailty model using PROC NLMIXED.

Parameter Estimates									
Parameter	**Estimate**	**Standard Error**	**DF**	**t Value**	**Pr > |t|**	**Alpha**	**Lower**	**Upper**	**Gradient**
b0	4.6684	0.3925	84	11.89	<.0001	0.05	3.8878	5.4490	-0.00023
b1	0.6229	0.3195	84	1.95	0.0546	0.05	-0.01250	1.2583	-0.00015
b2	-0.2434	0.08662	84	-2.81	0.0062	0.05	-0.4156	-0.07113	-0.00052
b3	0.01244	0.1061	84	0.12	0.9070	0.05	-0.1986	0.2234	-0.00033
gamma	1.3505	0.1177	84	11.47	<.0001	0.05	1.1164	1.5846	-0.00038
logsig	0.07305	0.1535	84	0.48	0.6353	0.05	-0.2321	0.3782	0.000302

Table 9. Parameter estimates from a shared frailty model using PROC NLMIXED

The NLMIXED step models survival whereas PROC PHREG models hazards. Therefore, you cannot directly compare parameter estimates. In general, the signs of the estimates will be opposite since survival increases when hazard decreases. The NLMIXED results are consistent with all other results; thiotepa does not appear to improve a patient's tumor recurrence status over the use of a placebo.

CONCLUSION

A frailty model is a subject-specific, random effects, survival model that corrects for bias in fixed effect parameter estimates and estimates of standard errors (and therefore test statistics) due to unobserved heterogeneity (or correlation between observations). Statisticians analyzing survival or time-to-event data can use PROC PHREG or NLMIXED to fit frailty models. PROC NLMIXED can fit both univariate and multivariate type frailty models. PROC PHREG can fit shared frailty models.

The examples shown in this paper are not the only, and perhaps not the best, choice for analyzing the thiotepa data.

- The thiotepa data are a special case of multivariate survival data because the events of interest are recurrent; each visit is an examination of tumor recurrence. If one or more tumors are found, they are removed. Consequently, a patient cannot have a second recurrence unless a first recurrence was observed. One may wish to fit a conditional model that accounts for this phenomenon. See *Example 71.10 Analysis of Recurrent Events Data* under the PHREG Procedure in SAS/STAT® 13.1 User's Guide for a demonstration on fitting total time and gap time models for recurrent data. Hosmer and Lemeshow (1999) also discuss several possible models suitable for recurrent event data, the choice of which depends on your research question and the method of data collection.
- Censoring has been assumed to be non-informative. This is often violated in oncology trials when high-risk patients censor due to events related to their condition. See Rondeau (2010) and Lu and Liu (2008) for discussions on modeling survival data with terminal events. Competing risk multi-state models may be another alternative.

REFERENCES

Allison, Paul D. 2010. *Survival Analysis Using SAS®: A Practical Guide, Second Edition*. Cary, NC. SAS Institute Inc.

Hosmer, David W and Stanley Lemeshow. 1999. *Applied Survival Analysis: Regression Modeling of Time to Event Data*. New York. John Wiley & Sons, Inc. 308-317.

Lu, Li and Chenwei Liu. 2008. "Analysis of Correlated Recurrent and Terminal Events Data in SAS®." *Proceedings of Northeastern SAS Users' Group 2008 Conference.*Cary, NC. SAS Institute Inc.

Rondeau, V. 2010. “Statistical models for recurrent events and death: Application to cancer events.” *Mathematical and Computer Modeling.* 52:949-955.

Vaupel, James W., Kenneth G. Manton and Eric Stallard. 1979. “The Impact of Heterogeneity in Individual Frailty on the Dynamics of Mortality.” *Demography*. 16:3:439-454.

ACKNOWLEDGMENT

The author thanks Lorne Rothman of SAS Institute (Canada) for his review and helpful suggestions that greatly improved this paper.

RECOMMENDED READING

- *Survival Analysis Using SAS®: A Practical Guide, Second Edition* by Paul Allison. Cary, NC. SAS Institute Inc.
- *SAS/STAT® 13.1 User’s Guide, the NLMIXED, PHREG, and MCMC Procedure examples.*

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

John Amrhein
McDougall Scientific Ltd.
789 Don Mills Rd, Suite 802
Toronto, Ontario, Canada M3C 1T5
jamrhein@mcdougallscientific.com
www.mcdougallscientific.com