The polysurvival model with long-term survivors

Josmar Mazucheli[a], Francisco Louzada[b] and Jorge A. Achcar[c]

[a]*Universidade Estadual de Maringá, DEs/UEM*
[b]*Universidade Federal de São Carlos, DEs/UFSCar*
[c]*Universidade de São Paulo, FMRP/USP*

Abstract. Long-term survival models have historically been considered for analyzing time-to-event data with long-term survivors fraction. However, situations in which a fraction $(1 - p)$ of systems is subject to failure from independent competing causes of failure, while the remaining proportion p is cured or has not presented the event of interest during the time period of the study, have not been fully considered in the literature. In order to accommodate such situations, we present in this paper a new long-term survival model. Maximum likelihood estimation procedure is discussed as well as interval estimation and hypothesis tests. A real dataset illustrates the methodology.

1 Introduction

Mixture models for long-term survivors have been widely used for fitting data where some individuals may never suffer the cause of failure under study (Maller and Zhou, 1996). In this type of modelling, it is assumed that, due to some unobserved prognostic factors, a certain fraction p of the population is immune to the cause of failure under study or a long-term survivor. The survivor function for the entire population can be written as

$$S(t) = p + (1 - p)S_0(t), \qquad (1.1)$$

where $S_0(t)$ denotes the survival function for the noncured group in the population. The long-term survivors cannot be identified but we can infer their presence in a data set if many of the largest times are censored. Common choices for $S_0(t)$ are the Gompertz, exponential and Weibull distributions. Yamaguchi (1992) considered the generalized log-gamma distribution for the cure rate in the context of accelerated failure-time regression models. Peng, Dear and Denham (1998) proposed a generalized F mixture model for the cure rate, which includes the most popular survival models as particular cases.

A problem however is that we can find in practice datasets where the fraction $(1 - p)$ of units are subject to failure from $k \geq 2$ competing causes. Besides, the exact cause of failure can be unknown, leading to the latent competing risk problem

(Louzada-Neto, 1999). According to Louzada-Neto (1998), in the classical competing risk scenarios the lifetime associated with a particular risk is not observable, rather we observe only the minimum (or maximum) lifetime value among all risks and which risk was responsible for the failure. A difficulty arises if the risks are latent in the sense that there is no information about which factor was responsible for the component failure (individual death). We call these latent competing risk data. For instance, in survival studies, patients are missing in follow-up, they die without autopsy, death is attributed to multiple causes, forms are not fully completed and only their death status is reported. In reliability, the components can be totally destroyed in the experiment. Further, the true cause of failure can be masked from our view. In modular systems, the need to keep a system running means that a module that contains many components can be replaced without the identification of the exact failing component.

As a simple illustration, consider data from a field trial of 4,992 circuit boards extracted from Chan and Meeker (1998, 1999). The data consists on the lifetimes of the circuit boards observed during a period of 10,000 hours of operation. There were 4,897 censored lifetimes. For the circuit boards that fail before the end of the experiment, engineering judgment indicates that failure can occur due to infant failure or wearout, but the exact cause of failure is unknown. There exists a large amount of censoring in the data, giving evidence of a possible presence of long-term survivors. Figure 1 shows the cumulative hazard plot for the data (Lawless, 1982). The curvature may indicate that another cause of failure becomes dominant as time progresses, corroborating the initial judgment. Two related models in the case of possible immunity for some systems were considered by Larson and Dinse (1985), which is based on a semi-Markov formulation of Lagakos, Sommer and


Figure 1   *Cumulative hazard plot.*

Zelen (1978), and by David and Moeschberger (1978). Both fully discussed in Maller and Zhou (1996).

In the present paper, we propose a general survival model for analyzing latent competing risk data with long-term survivals. The model is defined in Section 2 where we point out some of its particular cases and discuss the estimation procedure and computational aspects. Interval estimation is presented in Section 3. In Section 4, we discuss hypothesis tests. A brief discussion in Section 5 concludes the paper.

2 Model formulation

Suppose that a system is subject to $k \geq 2$ different causes of failure (or that it is composed of k different components). Let the lifetime related to the jth cause (or the lifetime of the jth element), X_j, have a survival function $S_j(x)$. Only $T = \min(X_1, \ldots, X_k)$ is observed for each individual and the X_j's are assumed independent. Thus, the overall survival function for t is $S_0(t) = \prod_{j=1}^{k} S_j(t)$. Besides, assume that the system may never experience the type of failure under study. Then, following (1.1), T is said to have a polysurvival model with long-term survivor if its overall survival function is given by:

$$S(t) = p + (1 - p) \prod_{j=1}^{k} S_j(t), \qquad (2.1)$$

where $S_j(t)$ denotes the survivor function for the noncured group due to the jth cause of failure. Similarly to the standard long-term survivor model (1.1), we can choose $S_j(t)$ in (2.1) to be a product of usual survival functions, such as Weibull, log-normal and log-logistic ones, even of different families.

An advantage of model (2.1) is that for $p = 0$, implying in an absence of immunes in the population, model (2.1) results in a standard competing risk model (David and Moeschberger, 1978), while for $k = 1$ we have the standard long-term survival model (1.1).

Consider a sample of independent random variables $T_1, \ldots, T_n$ denoting the lifetimes of n units. Assume that T_i has associated an indicator variable defined by $\delta_i = 1$ if $T_i = t_i$ is an observed failure time and $\delta_i = 0$ if it is a right-censored observation. The maximum likelihood estimates (MLE) of the parameters are obtained by direct maximization of the likelihood function $L = \prod_{i=1}^{n} [pf(t_i)]^{\delta_i} [S(t_i)]^{1-\delta_i}$, where $f(t_i)$ is the probability density function obtained from (2.1), see, for example, Maller and Zhou (1996). The advantage of this procedure is that it runs immediately using existing statistical packages such as the R (Ihaka and Gentleman, 1996), or SAS/NLP (SAS, 2010). Interested readers can access the codes used by email to the authors. An important aspect of implementing the estimation procedure concerns parametrization. In our numerical examples and simulation

studies, we have not faced numerical problems, such as evidence of failure of convergence or end on multiple maximums, from parameters with unbounded ranges. The results however are all reported in the original scale.

A problem that appears is that the parameters in (2.1) are fundamentally not identifiable, since the likelihood function is unchanged by permutation of the component labels $1, \ldots, k$. To enforce parameter uniqueness, we have successfully used an order constraint on the shape parameters. For instance, consider the bi-Weibull long-term survival particular case, that is (2.1) with $j = 2$ and $S_0(t) = \prod_{j=1}^{2} S_j(t) = \prod_{j=1}^{2} \exp[-(t/\mu_j)]^{\beta_j}$. By considering $\beta_1 < \beta_2$ we avoid the identifiability problem. The rationality behind such enforcement relies on following idea. Considering the bi-Weibull long-term survival particular case, we may obtain two equivalent models renumbering the components 1 and 2 by a permutation. In this case, the likelihood function is bimodal, and its values do not change whatever the component order we consider. Then, letting us to assume that there is an order criterion that can be used when numbering the components. For instance, with $\beta_1 < \beta_2$, the parameters are obtained in a unique way. The idea of considering a restriction for solving the nonidentifiability problem is not new and have been considered for other authors. For instance, interested readers can refer to Jasra, Holmes and Stephens (2005) for a review of different approaches for solving the nonidentifiability problem.

To illustrate the methodology described above, we consider the Weibull model. This is however qualitatively similar to other particular cases of (2.1). Recall the circuit board data discussed in Section 1. For instance, we fitted a bi-Weibull long-term survival model to the data. The MLEs of μ_1, β_1, μ_2, β_2 and p are equal to 13208.20, 0.27, 8530.47, 8.39 and 0.981, respectively.

3 Interval estimation

Large sample inference for the parameters can be based, in principle, on the MLEs and their estimated standard errors. Considering the bi-Weibull long-term survival particular case fitted to the circuit board data discussed in Section 1 we obtain the usual asymptotic 90% confidence intervals for μ_1, β_1, μ_2, β_2 and p given, respectively, by (5140.34; 33935.83), (0.22; 0.34), (8195.95; 8877.97), (6.25; 11.25) and (0.977; 0.984).

A problem however is that, in lifetime studies, it is common to find datasets with a small or moderate quantity of observed lifetimes. Also, in the circuit board data, we can observe the presence of a large amount of censoring. In order to check the behavior of the asymptotic theory for small and moderate sized samples when censoring is observed, we performed a small-scale simulation study for examining the coverage probabilities of the confidence intervals for the parameters.

Samples were generated according to a bi-Weibull survival model with long-term survivor with $p = 0.50$ and 0.90 and $\mu_1 = 30$, $\beta_1 = 0.5$, $\mu_2 = 100$ and

Table 1 *Coverage probabilities of nominal 90% asymptotic confidence intervals. In each cell the left result corresponds to $p = 0.50$ while the right result corresponds to $p = 0.90$*

n	p	μ_1	β_1	μ_2	β_2
75	0.62 0.56	0.59 0.55	0.60 0.57	0.77 0.72	0.68 0.61
100	0.69 0.65	0.82 0.79	0.65 0.61	0.81 0.76	0.72 0.68
200	0.71 0.67	0.88 0.84	0.82 0.77	0.91 0.86	0.87 0.83

$\beta_2 = 3$. The censoring times were generated independently of the survival times according to a uniform distribution given by $U(a, b)$, where a is equal to the biggest noncensored observation and $b = 10a$. We further fixed the amount of censoring at 50% and 90%, coinciding with the same amount of log-term survivals. A thousand samples were generated for each case for sample sizes of $n = 75$, 100 and 200. Table 1 shows the variation in coverage of nominal 90% asymptotic confidence intervals according to the sample size and the two percentages of long-term survivals. The 90% asymptotic confidence interval for the nominal coverage probability of 0.90 based on a sample of size equal to a thousand observations is given by $(0.884, 0.916)$. If a confidence interval has exact coverage of 0.90, roughly 90% of the observed coverages should be inside these bounds. There clearly is under coverage of the confidence intervals for small and moderate sized samples, particularly when the percentage of long-term survivals are heavy. Such findings are evidences for the need of more adequate procedures for small or moderate sized samples.

An alternative direct approach is the Bootstrap procedure, which is a simulation procedure that aims to obtain empirical interval estimations by re-sampling the original data set. There are two basic Bootstrap types: the parametric Bootstrap, where the simulating datasets are drawn by generating observations, in our case, from model (2.1) with the parameters replaced by their MLEs, and the nonparametric Bootstrap, where the simulating datasets are drawn with replacement directly from the original sample. More details about the Bootstrap technique may be seen in Davison and Hinkley (1997).

Consider μ_1 the parameter we have interest in and suppose that we are interested in constructing a confidence interval for it. For each resample, obtained by a parametric or a nonparametric way, we calculate the MLEs for μ_1 and we have at the end of R resamples $\widehat{\mu}_{1,1} < \cdots < \widehat{\mu}_{1,R}$ ordered MLEs values. Then, we use $\widehat{\mu}_{1,(R+1)(a/2)}$ and $\widehat{\mu}_{1,(R+1)(1-a/2)}$ as the lower and upper bounds of the Bootstrap percentile confidence interval $100(1 - a)\%$ for μ_1, respectively, and a is the significance level. The Bootstrap percentile intervals $100(1 - a)\%$ for the other model parameters can be analogously obtained.

In order to check the adequacy of the Bootstrap procedure for small and moderate sized samples when censoring is observed, we ran the simulation study described above again, with the same thousand samples generated for each case. We

Table 2 *Coverage probabilities of nominal* 90% *Bootstrap confidence intervals. In each cell the left result corresponds to* $p = 0.50$ *while the right result corresponds to* $p = 0.90$

n	Bootstrap	p	μ_1	β_1	μ_2	β_2
75	Parametric	0.86 0.83	0.86 0.84	0.87 0.83	0.90 0.87	0.87 0.85
	Nonparametric	0.84 0.82	0.85 0.82	0.85 0.84	0.89 0.86	0.86 0.86
100	Parametric	0.88 0.87	0.88 0.88	0.88 0.87	0.89 0.88	0.89 0.88
	Nonparametric	0.88 0.88	0.89 0.88	0.89 0.88	0.89 0.88	0.88 0.88
200	Parametric	0.90 0.89	0.91 0.89	0.91 0.88	0.90 0.90	0.91 0.89
	Nonparametric	0.92 0.90	0.90 0.89	0.91 0.91	0.90 0.90	0.92 0.90

further fixed R equals to 999. Table 2 shows the variation in coverage of nominal 90% Bootstrap confidence intervals according to the sample size and two censoring percentages. Based on the same criteria for indicating whether a confidence interval has exact coverage of 90%, that is, based on the 90% asymptotic confidence interval for the nominal coverage probability of 0.90 which is given by $(0.88, 0.92)$, there are evidences for the adequacy of the both Bootstrap procedures.

Another confidence interval comparison criteria is based on the confidence interval average amplitudes, which are only reported here. Although, in comparing average amplitudes of confidence intervals it is preferred to compare intervals with nearly the same coverage probability (Jeng and Meeker, 1999), which is not our case, we notice that the average amplitudes of the confidence intervals increase when the sample size decreases. We also observe that the asymptotic confidence intervals have the smallest average amplitudes compared with the confidence intervals obtained by resampling. This is in full agreement with Davison and Hinkley (1997), which pointed out that first order likelihood procedures generally under estimate the variance. The nonparametric percentile confidence intervals have similar average amplitudes when compared with the parametric percentile ones.

Then, we prefer the nonparametric Bootstrap procedure since it has smaller computational cost in comparison with the parametric Bootstrap version. To illustrate the methodology described above, we consider bi-Weibull long-term survival fitted to the circuit board data discussed in Section 1. The 90% nonparametric percentile confidence intervals for μ_1, β_1, μ_2, β_2 and p are given by $(6697.918, 29130.992)$, $(0.242, 0.310)$, $(8275.010, 8798.596)$, $(6.892, 10.571)$ and $(0.978, 0.983)$.

4 Hypothesis tests

There are two major problems that should be addressed from the hypothesis tests point of view related to the polysurvival model with long term survivor (2.1). The first problem is related to the test of a particular polysurvival function with a

number k of components. The second problem is related to the test for presence of long-term survivors, that is, to test the hypothesis $H_0 : p = 0$.

4.1 Testing a particular polysurvival function

For testing the adequacy of a particular polysurvival function, we can consider the likelihood ratio statistics (LRS), $w = 2(l_2 - l_1)$, where l_1 and l_2 are the log-likelihood functions for models 1 and 2, respectively. Large positive values of w give favourable evidence to model 2. A difficulty with the present setting is that the tests can be non nested, since the components of $S_j(t)$ in (2.1) can be a product of usual survival functions, but even of different families and even with different number k of components. So, the regularity conditions, on which the standard asymptotic theory is based (Cox and Hinkley, 1974), will not hold. Then we cannot assume that w is asymptoticaly distributed like a chi-square distribution with one degree of freedom. For instance, we present the results of a small scale simulation study ran in order to check the asymptotic behavior of w for testing a single-Weibull long-term survival model, that is (2.1) with $j = 1$ and $S_0(t) = \exp{-(t/\mu)^\beta}$, against a bi-Weibull long-term survival model, that is (2.1) with $j = 2$ and $S_0(t) = \prod_{j=1}^{2} S_j(t) = \prod_{j=1}^{2} \exp(-(t/\mu_j)^{\beta_j})$. Samples were generated under the null hypothesis according to a single-Weibull survival model with long-term survivor with $p = 0.50$ and 0.90 and $\mu = 30$ and $\beta = 0.5$ and 3. The censoring times were generated independently of the survival times according to a uniform distribution given by $U(a, b)$, where a is equal to the biggest noncensored observation and $b = 10a$. We further assumed that the amount of censoring are fixed at 50% and 90%, coinciding with the amount of long-term survivals. A thousand samples were generated for each case for sample sizes of $n = 75$, 100 and 200. We then verified whether the null hypothesis was rejected at the 5% level or not. Table 3 shows the empirical significance levels (in percentage) according to the sample size, β value and two percentages of long-term survivals. As discussed before, the empirical significance levels indicate that the LRS, w, is not distributed as a chi-squared distribution for small and moderate size samples.

Table 3 *Empirical significance levels for testing a single-Weibull long term survival model against a bi-Weibull long-term survival model according to the sample size, β value and two censoring percentages. In each cell the left result corresponds to $p = 0.50$ while the right result corresponds to $p = 0.90$*

n	$\beta = 0.5$	$\beta = 1$	$\beta = 3$
75	15.6 17.7	11.7 12.9	13.2 15.5
100	13.3 15.8	9.2 10.1	11.6 13.4
200	8.1 11.3	5.8 7.4	6.9 9.5

An alternative direct approach is to Bootstrap the LRS w in order to obtain its empirical distribution (Davison and Hinkley, 1997). This can be done parametrically or nonparametrically. The parametric Bootstrap technique consists of generating R datasets from the model under the null hypothesis (model 1) with the parameters substituted by their MLEs obtained by using the procedure discussed in Section 2, record $w_1^* < \cdots < w_R^*$, and use $w_{(R+1)(1-a)}^*$ as the critical point to test the null hypothesis with size a. We consider here for safe R equals to 999, according to Hall (1986) however this number is bigger than the the number of replications required to get a critical level of 0.10 from the 0.90 percentile of the empirical distribution of the LRS. The nonparametric Bootstrap technique operates in much the same way, but instead of generating datasets from the model under the null hypothesis (model 1), with the parameters substituted by their MLEs, we draw R samples with replacement of n observations each from the original dataset $t_1, \ldots, t_n$. The same thousand samples considered above were considered here for each case for sample sizes of $n = 75$, 100 and 200. We then applied the suggested parametric and nonparametric Bootstrap procedures and verified whether the null hypothesis was rejected at the 5% level or not. Table 4 shows the bootstrap empirical significance levels (in percentage) according to the sample size, β value and two percentages of long-term survivals. The suggested bootstrap procedures (parametric and nonparametric) can properly address the anti-conservative behavior of the empirical significance levels based on the asymptotic theory.

For instance, recall the circuit board data discussed in Section 1. The LRS w for testing a single-Weibull long-term survival model, that is (2.1) with $j = 1$ and $S_0(t) = \exp -(t/\mu)^\beta$, against a bi-Weibull long-term survival model, that is (2.1) with $j = 2$ and $S_0(t) = \prod_{j=1}^2 S_j(t) = \prod_{j=1}^2 \exp(-(t/\mu_j)^{\beta_j})$, is equal to 21.86. The empirical p-values equal to 0.003 and 0.014 according to the parametric and non parametric Bootstrap scheme described above, respectively, which gives a strong evidence in favor of the full model. This is corroborated by the comparison of the Kaplan–Meier estimate and the fitted models in Figure 2.

Table 4 *Bootstrap empirical significance levels for testing a single-Weibull long term survival model against a bi-Weibull long-term survival model according to the sample size, β value and two censoring percentages. In each cell the left result corresponds to $p = 0.50$ while the right result corresponds to $p = 0.90$*

| n	Bootstrap	$\beta = 0.5$	$\beta = 1$	$\beta = 3$
75	Parametric	10.0 10.1	8.9 9.7	9.2 10.0
	Nonparametric	10.9 10.7	8.2 9.1	9.2 10.5
100	Parametric	8.2 8.3	6.8 7.1	7.6 8.2
	Nonparametric	8.6 8.4	7.1 7.9	7.7 8.9
200	Parametric	5.3 6.2	4.9 5.6	5.1 6.0
	Nonparametric	5.3 6.1	5.3 5.9	5.3 6.2

Figure 2 *Kaplan–Meier estimate and the fitted models of the circuit boards data.* (—): *bi-Weibull long-term survival model fit*, ($\cdots$): *single-Weibull long-term survival model fit.*

4.2 Testing the presence of long-term survivals

For testing for the presence of long-term survivors, we should perform the hypothesis tests $H_0 : p = 0$ versus $H_1 : p > 0$. Let $u = 2(l_{\text{unres}} - l_{\text{res}})$ be the LRS for testing for the presence of long-term survivors, where l_{res} and l_{unres} are the log-likelihoods for the model under the restricted hypothesis H_0 and under the unrestricted hypothesis H_1. Large positive values of u give favourable evidence to the full model. Following Ghitany and Maller (1992), it can be shown that, under the polysurvival model with long term survivor (2.1), u is asymptotically distributed like a symmetric mixture of a chi-squared distribution with one degree of freedom and a point-mass at zero. Then, $\lim_{n \to \infty} P(u_n \leq c) = 1/2 + 1/2 \ P(\chi_1^2 \leq c)$, where $P(\chi_1^2 \leq c)$ denotes a random variable with a chi-square distribution with one degree of freedom.

Recall the circuit board data discussed in Section 1. The LRS u for testing $H_0 : p = 0$ versus $H_1 : p > 0$ is equal to 17.59 with p-value equal to 0.000027 according to the limiting distribution of u, which is a strong evidence in favor of the model with long-term survivors.

As stated before, a problem however is that in survival and reliability studies we can have samples with small or moderate amount of observed lifetimes and with a huge amount of censoring. In these settings, the asymptotic approach may not be adequate. For instance, we present the results of a small scale simulation study ran in order to check the asymptotic behavior of u for testing $H_0 : p = 0$. Samples were generated according to a bi-Weibull survival model with long-term survivor with $p = 0.50$ and 0.90 and $\mu_1 = 30$, $\beta_1 = 0.5$, $\mu_2 = 100$ and $\beta_2 = 3$.

Table 5　*Empirical significance levels for testing $H_0 : p = 0$ according to the sample size and two censoring percentages. In each cell the left result corresponds to the asymptotic approach while the middle and right results corresponds to the parametric and nonparametric approaches, respectively*

n	$p = 0.50$	$p = 0.90$
75	22.4 12.3 13.0	24.3 12.7 14.6
100	18.6 9.1 10.3	21.1 9.9 10.7
200	10.9 6.3 6.9	14.2 6.8 7.2

The censoring times were generated independently of the survival times according to a uniform distribution given by $U(a, b)$, where a is equal to the biggest non-censored observation and $b = 10a$. We also assumed that the amount of censoring are fixed at 50% and 90%. A thousand samples were generated for each case for sample sizes of $n = 75$, 100 and 200. We then verified whether the null hypothesis was rejected at the 5% level. Table 5 shows the empirical significance levels (in percentage) according to the sample size, β value and two percentages of long-term survivals. The left results in each cell motivate to a Bootstrap approach for this kind of hypothesis tests. As stated in Section 4.1, the idea is to obtain the empirical distribution of u. The parametric Bootstrap version consists of generating R datasets from the model under the null hypothesis (model without long-term survivors) with the parameters substituted by their MLEs obtained by using the procedure discussed in Section 2, record $u_1^* < \cdots < u_R^*$, and use $u_{(R+1)(1-a)}^*$ as the critical point to test the null hypothesis with size a. We consider here R equal to 999. The nonparametric Bootstrap version of the procedure operates in the same way, but instead of generating datasets from the model under the null hypothesis, with the parameters substituted by their MLEs, we draw R samples with replacement of n observations each from the original dataset $t_1, \ldots, t_n$. The same thousand samples considered above were considered here for each case for sample sizes of $n = 75$, 100 and 200. We then applied the suggested parametric and nonparametric Bootstrap procedures and verified whether the null hypothesis was rejected at the 5% level or not. According to the middle and right results in each cell of Table 5, the suggested bootstrap procedures (parametric and nonparametric) can properly address the anti-conservative behavior of the empirical significance levels based on the asymptotic theory.

Recall again the circuit board data discussed in Section 1. The LRS u for testing $H_0 : p = 0$ versus $H_1 : p > 0$ is equal to 17.59 with empirical p-values equal to 0.042 and 0.075, which are evidences in favor of the model with long-term survivors, at least at a 10% significance level.

5 Discussion

The poly-survival model with long-term survivors proposed in this paper allows a broad class of survival models. The model provides a reasonable physical interpretation of the phenomenon underlying latent competing risk data with long-term survivors. Maximum likelihood inference can be implemented straightforwardly and parametric and nonparametric simulation can be successfully used for hypothesis testing and generating precise confidence intervals for the parameters.

In the paper, following Louzada-Neto (1999), we assume independent competing causes of failure. However, as pointed out by a referee, assuming independent competing causes of failure is a rather strong assumption, which should be more investigated in the context of our modeling in future. For instance, Maller and Zhou (1996) mention that the David and Moeschberger formulation (David and Moeschberger, 1978), can be generalized to cases when the competing causes are not independent. Besides, the mixture model approach proposed Larson and Dinse (1985) avoids the independence assumption. The finite mixture model approach presented by Chao (1998) do not assumes independent cause-specific failure times. And, finally, a common way to model dependent competing causes is to consider the so called frailty models (Gordon, 2002 and Box-Steffensmeier and Jones, 2004).

6 Acknowledgments

References

Box-Steffensmeier, J. M. and Jones, B. S. (2004). *Event History Modeling: A Guide for Social Scientists*. Cambridge/New York: Cambridge Univ. Press.

Chan, V. V. and Meeker, W. Q. (1998). A competing-risk limited failure population model for product failure times. Technical report, Dept. Statistics, Iowa State Univ., Ames, IA.

Chan, V. V. and Meeker, W. Q. (1999). A Failure-time model for infant-mortality and wearout failure modes. *IEEE Transactions on Reliability* **48**, 377–387.

Chao, E. C. (1998). Gibbs sampling for long-term survival data with competing risks. *Biometrics* **54**, 350–366.

Cox, D. R. and Hinkley, D. V. (1974). *Theoretical Statistics*. London: Chapman & Hall. MR0370837

David, H. A. and Moeschberger, M. L. (1978). *The Theory of Competing Risks. Griffin's Statistical Monograph Series* **39**. New York: Macmillan. MR0592960

Davison, A. C. and Hinkley, D. V. (1997). *Bootstrap Methods and Their Applications*. Cambridge/New York: Cambridge Univ. Press. MR1478673

Ghitany, M. E. and Maller, R. A. (1992). Asymptotic results for exponential mixture models with long-term survivors. *Statistics* **23**, 321–336. MR1238254

Gordon, S. C. (2002). Stochastic dependence in competing risks. *American Journal of Political Science* **46**, 200–217.

Hall, P. P. (1986). On the bootstrap and confidence intervals. *The Annals of Statistics* **14**, 1431–1452. MR0868310

Ihaka, R. R. and Gentleman, R. R. (1996). R: A language for data analysis and graphics. *Journal of Computational and Graphical Statistics* **5**, 299–314.

Jasra, A. A., Holmes, C. C. and Stephens, D. A. (2005). Markov chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling. *Statistical Science* **20**, 50–67. MR2182987

Jeng, S. L. and Meeker, W. Q. (1999). Comparisons of approximate confidence interval procedures for type I censored data. *Scandinavian Journal of Statistics* **20**, 1–15.

Lagakos, S. W., Sommer, C. J. and Zelen, M. M. (1978). Semi-Markov models for partially censored data. *Biometrika* **65**, 311–317. MR0513932

Larson, M. G. and Dinse, G. E. (1985). A mixture model for the regression analysis of competing risks data. *Journal of the Royal Statistical Society, Ser. C: Applied Statistics* **34**, 201–211. MR0827668

Lawless, J. F. (1982). *Statistical Models and Methods for Lifetime Data*. New York: Wiley. MR0640866

Louzada-Neto, F. F. (1998). Hazard models for lifetime data. Ph.D. thesis, Oxford Univ.

Louzada-Neto, F. F. (1999). Polyhazard regression models for lifetime data. *Biometrics* **55**, 1281–1285.

Maller, R. A. and Zhou, X. X. (1996). *Survival Analysis With Long-Term Survivors*. Chichester: Wiley. MR1453117

Peng, Y. Y., Dear, K. B. G. and Denham, J. W. (1998). A generalized F mixture model for cure rate estimation. *Statistics in Medicine* **17**, 813–830.

SAS (2010). *The NLP Procedure, SAS/OR® 9.22 User's Guide: Mathematical Programming*. Cary, NC: SAS Institute Inc.

Yamaguchi, K. K. (1992). Accelerated failure-time regression models with a regression model of surviving fraction: An application to the analysis of permanent employment in Japan. *Journal of the American Statistical Association* **87**, 284–292.

J. Mazucheli
Departamento de Estatística, DEs/UEM
Universidade Estadual de Maringá
Maringá, PR
Brazil
E-mail: jmazucheli@uem.br

F. Louzada
Departamento de Matemática Aplicada e Estatística
Universidade de São Paulo
São Carlos - SP
Brazil
E-mail: louzada@icmc.usp.br

J. A. Achcar
Departamento de Medicina Social, FMRP/USP
Universidade de São Paulo
Ribeirão Preto, SP
Brazil
E-mail: achcar@fmrp.usp.br