CHAPTER 1

COMPETING RISK MODELING IN RELIABILITY

TIM BEDFORD

Department of Management Science
Strathclyde University
Glasgow, UK
E-mail: tim.bedford@strath.ac.uk

This paper gives a review of some work in the area of competing risk applied to reliability problems, focusing particularly on that of the author and co-workers. The results discussed cover a range of topics, starting with the identifiability problem, bounds and a characterization of marginal distributions with given competing risk information. We discuss the way in which the assumption of independence usually gives an optimistic view of failure behavior, possible models for maintenance, and generalizations of the competing risk problem to nonrenewal systems.

1.1. Introduction

The competing risk problem arises quite naturally in the reliability context. Maintenance logs often track the history of events occurring at a particular socket. The events can be failure mode specific, incipient failures, maintenance actions, etc. Where the cost of critical failure is large, the maintenance policy will ensure that the whole system is as good as new. Hence we can regard the data as arising from a renewal process in which we only see the "first" possible event occurring after renewal and we know what that event is. The different events can be regarded as competing risks. The "competing risk problem" is that we cannot identify the marginal distributions of the time to each event without making untestable distributional assumptions. Competing risk information is, at least implicitly, used in reliability databases such as the Center for Chemical Process Safety (CPSS) and European Industry Reliability Data (EIREDA) generic

MODERN STATISTICAL AND MATHEMATICAL METHODS IN RELIABILITY

databases amongst several others. These databases give data–failure rates or on-demand failure probabilities as appropriate–for each failure mode of the component. In order to derive this information it is necessary to have used a competing risk statistical model to interpret the underlying failure data. A discussion of the way in which such databases are built up and the role that the competing risk problem arises there is given in Cooke and Bedford.[1]

The paper gives an overview of competing risk work, specifically in the reliability context, with which I and co-workers have been associated. It does not attempt to give a general overview of competing risk. For a more general view we refer the reader to Crowder's recent book[2] and Deshpande's overview paper.[3] The issues we cover are:

- Independent and dependent competing risks
- Characterization of possible marginals
- Kolmogorov Smirnov Test
- The bias of independence
- Maintenance as a censoring mechanism
- Loosening the renewal assumption

In considering these issues we shall largely take a probabilistic modeling viewpoint, although one could equally well take a more statistical view or an operations research view.

The competing risk problem is a part of the more general issue of model identifiability. We build models in order to gain an understanding of system behavior. In general terms, the more tightly we specify the model class, the more likely we are to be able to identify the specific model within the class. If we define the model class too tightly though, the model class may not capture all the features contained in the data. However, our intuition about defining model classes more tightly does not always correspond to the functional constraints that imply identifiability or lack thereof. We shall give an example of this in Sec. 1.8.

In competing risk applications to reliability the model-identifiability issue means that we do need to specify "tight" families of models to apply in specific application situations. To do this we need a better understanding of the engineering context, particularly for applications to maintenance censoring.

1.2. Independent and Dependent Competing Risks

In general there may be several different processes going on that could remove a component from service. Hence the time to next removal, Y, is the minimum of a number of different potential event times $Y = \min(X_1, \ldots, X_n)$. For simplicity assume that the different events cannot occur together and, furthermore, just consider a single nonfailure event, for example unscheduled preventive maintenance. Hence there is a failure time X_1, which is the time (since the previous service removal) that the equipment would fail, and a single PM time X_2, which is the time at which the equipment would be preventively maintained. We only observe the smallest of the two variables but also observe which one it is, that is, we know whether we have observed a failure or a PM. Hence the observable data is of the form $Y = (\min(X_1, X_2), 1_{X_1<X_2})$. It would clearly be interesting to know about the distribution of X_1, that is, the behavior of the system with the maintenance effect removed. However, we cannot observe X_1 directly. The observations we have allow us only to estimate the *subdistribution function* $G_1(t) = P(X_1 \leq t, X_1 < X_2)$. The subdistribution function converges to the value $P(X_1 < X_2)$ as $t \to \infty$. We often talk about the *subsurvivor function* $S_1^*(t) = P(X_1 > t, X_1 < X_2)$, which is equal to $P(X_1 < X_2) - G_1(t)$. The normalized subsurvivor function is the quantity $S_{X_1}^*(t)/S_{X_1}^*(0)$ normalized to be equal to 1 at $t = 0$. The final important quantity that can be estimated directly from observable data is the probability of a censor after time t, $\Phi(t) = P(X_2 < X_1 | Y > t)$. The shapes of these functions can play a role in model selection.

The classical competing risks problem is to identify the marginal distributions of the competing risk variables from the competing risk data. It is well known[4] that the marginal and joint distributions of (X_1, X_2) are in general "nonidentifiable," that is, there are many different joint distributions which share the same subdistribution functions. Is it also well known[5,6,7] that if X_1 and X_2 are independent, nonatomic and share essential suprema, their marginal distributions are identifiable. Given a pair of subsurvivor functions we can assume an underlying independent model, but may have to accept that one of the random variables has a degenerate distribution, that is, an atom at infinity.[8] In any case, if we are prepared to assume independence then–subject to some technical conditions–we can identify marginals. For the moment we shall concentrate on general bounds that make as few assumptions as possible, and explain the source of nonidentifiability.

Figure 1.1 shows the (X_1, X_2) plane and the events whose probabilities can be estimated from observable data. For example, given times $t_1 < t_2$ we can estimate the probability of the event $t_1 < X_1 \leq t_2, X_1 < X_2$ which corresponds to the vertical hatched region on the figure, while given times $t_3 < t_4$ we can estimate the probability of the event $t_3 < X_2 \leq t_4, X_2 < X_1$, which corresponds to the horizontal hatched region on the figure.

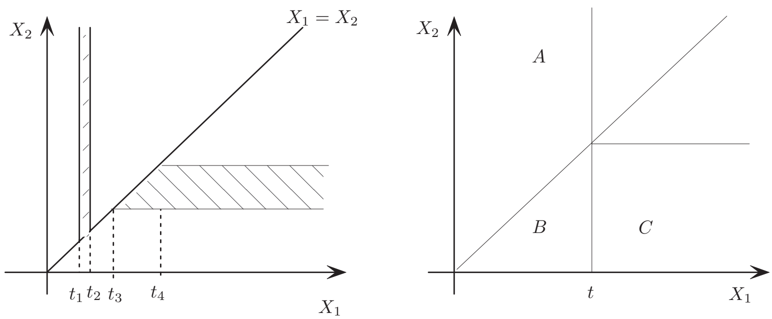

Fig. 1.1. (L) Events whose probabilities can be estimated by competing risk data; (R)Geometry of events determining the upper and lower bounds.

We are able to estimate the probability of any such region, but we cannot estimate how the probability is distributed *along* such a region. Now we can see why the distribution of X_1 is not identifiable; varying the mass within the horizontal region changes the distribution of X_1 without changing the distribution of the observable quantities. Considering the "extreme" ways in which that probability mass could be distributed leads to upper and lower bounds on the marginal distribution function of X_1. The Peterson bounds[9] are pointwise upper and lower bounds on the value of the marginal distribution function. They say that for any $t \geq 0$ we have $G_1(t) \leq F_1(t) \leq F_Y(t)$. A functional bound was found by Crowder[10] who showed that the distance between the distribution function F_1 and the Peterson lower bound, $F_1(t) - G_1(t)$ is nondecreasing.

A simple proof of these bounds is possible through considering the geometry of the events in the (X_1, X_2)-plane. Figure 1.1(R) shows, for a given t, three events marked A, B and C. The probability of event A is the lower bound probability in the Peterson bound, $G_1(t)$. This is clearly less than or equal to the probability of event $A \cup B$, which is just $F_1(t)$. This in turn is less than or equal to the probability of event $A \cup B \cup C$, which is just $F_Y(t)$.

This shows that the lower and upper Peterson bounds hold. To get the functional lower bound of Crowder just note that the difference $F_1(t) - G_1(t)$ is the probability of the event $\{X_1 \leq t, X_2 < X_1\}$, that is, the region marked B on the figure. Clearly as t increases, we get an increasing sequence of events whose probabilities must therefore also be nondecreasing. This demonstrates the functional bound. Both Peterson and Crowder make constructions to show when there is a joint distribution satisfying the bounds. A result in Bedford and Meilijson[11] however improved these results slightly while giving a very simple geometric construction, which we now consider.

1.3. Characterization of Possible Marginals

The characterization tells us exactly which marginal distributions are possible for given subdistribution functions. To simplify things in this presentation we assume that we are only going to deal with continuous (sub)distributions and the reader is referred to Bedford and Meilijson[11] for the details of the general case of more than 2 variables, which may have atoms and could have ties.

The key idea here is that of a co-monotone representation. We have already seen that the Crowder functional bound writes the distribution function as a sum of two monotone functions: the subdistribution function and the nondecreasing "gap." More generally we define a *co-monotone representation* of a continuous real-valued function f as a pair of monotone nondecreasing continuous functions f_1 and f_2 such that $f = f_1 + f_2$.

The characterization result will show when we can find a pair of random variables compatible with the observable subdistribution functions. It is therefore necessary to define abstractly what a pair of subdistribution functions is without reference to a pair of random variables. We define a pair of functions G_1, G_2 to be a lifetime subdistribution pair if

(1) $G_i : [0, \infty) \to \mathbb{R}$, $i = 1, 2$.
(2) They are nondecreasing continuous real-valued functions with $G_1(0) = G_2(0) = 0$.
(3) $\lim_{t\to\infty} G_1(t) = p_1$ and $\lim_{t\to\infty} G_2(t) = p_2$ with $p_1 + p_2 = 1$.

Suppose we are given such a pair of functions. (As stated above, the subdistribution functions can be estimated from competing risk data.) Consider the unit strip $[0, \infty) \times [0, 1]$ and subdivide it into two strips of heights p_1 and p_2, respectively as shown in Fig. 1.2. In the lower and upper strips we plot the functions $G_1(t)$ and $G_2(t)$, respectively. In the same strips we choose, arbitrarily, two nondecreasing right continuous functions whose graphs increase from the bottom to the top of the strip, while lying under the graphs of the functions we have already drawn. For reasons that will become clear we call these new functions $F_2 - G_2$ and $F_1 - G_1$, respectively. See Fig. 1.3. The notion of "lying under" will be made clear in the statement of the theorem below.

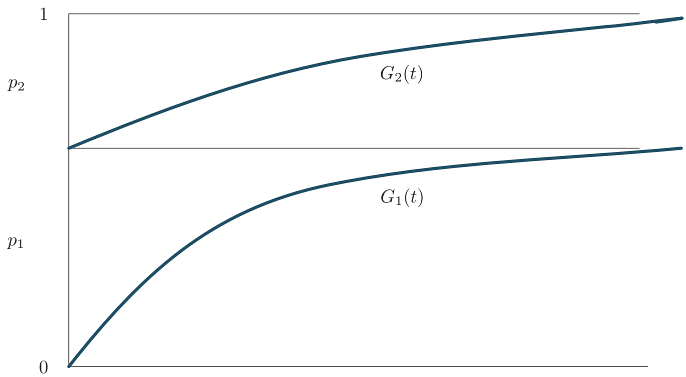

Fig. 1.2. Co-monotone construction - 1.

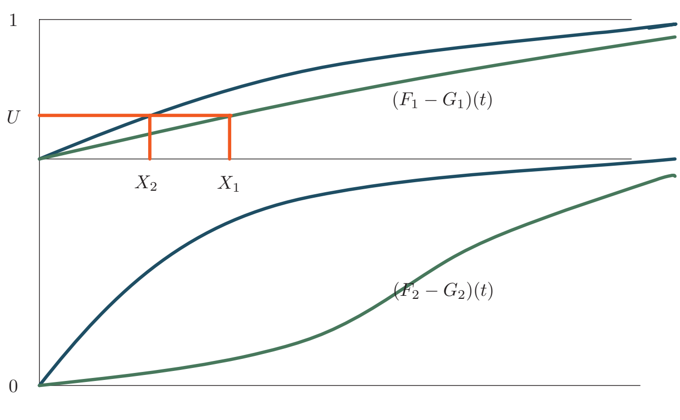

Fig. 1.3. Co-monotone construction - 2.

Theorem 1.1: [11] (i) Let X_1 and X_2 be lifetime random variables. Then, using the notation established above,

(1) $F_i = G_i + (F_i - G_i)$ is a nonnegative co-monotone representation of a nondecreasing continuous function, for $i = 1, 2$.
(2) $F_i(t) \leq G_1(t) + G_2(t)$ for all t, and the Lebesgue measure of the range set $\{(F_i - G_i)(t) | F_i(t) = G_1(t) + G_2(t)\}$ is zero, for $i = 1, 2$.
(3)(a) $F_i(0) = 0$ and $F_i(\infty) = 1$, for $i = 1, 2$.
  (b) $G_1(\infty) + G_2(\infty) = 1$.

(ii) If nondecreasing right continuous functions F_1, F_2, G_1, and G_2 satisfy the conditions (1)-(3) of (i) then there is a pair of random variables (X_1, X_2) for which F_i and G_i are the distribution and subdistribution functions respectively ($i = 1, 2$).

The construction of the random variables (X_1, X_2) can be done quite simply and is shown in Fig. 1.3. Draw a uniform random variable U. If it is in the lower strip (that is, $U < p_1$) then we can invert U through the functions drawn in the lower strip. Define $X_1 = G_1^{-1}(U)$ and $X_2 = (F_2 - G_2)^{-1}(U)$. The ordering of these two functions implies that $X_1 \leq X_2$, and furthermore that $X_1 = X_2$ with probability 0. If U is in the upper strip then a similar construction applies with the roles of X_1 and X_2 reversed. See Fig. 1.3.

This geometrical construction shows that the conditions of the theorem are sufficient for the existence of competing risk variables. Condition 1 is necessary as it is Crowders functional bound. Condition 3 is an obvious necessary condition. The first part of Condition 2 is the Peterson upper bound, while the second part is a rather subtle "light touch" condition whose proof is fairly technical and for which we refer the reader to Bedford and Meilijson.[11]

1.4. Kolmogorov-Smirnov Test

The complete characterization described above was used to produce a statistical test based on the Kolmogorov-Smirnov statistic in which a hypothesized marginal distribution can be tested against available data. The functional bound tells us that, given a dataset, the difference between the empirical subdistribution and the unobserved empirical marginal distribution function should be nondecreasing. A little thought shows that the distances can be computed at the "jumps" of these functions, which of course occur at the times recorded in the dataset. Recall that the Kolmogorov-Smirnov test takes a hypothesized distribution and uses the asymptotic relation that the functional difference between true population distribution and empirical distribution converges, when suitably normalized, to a Brownian bridge. Extremes of the Brownian bridge can then be used to establish classical confidence intervals. Although in our competing risk situation we cannot estimate the maximal difference between empirical distribution function and hypothesized distribution function (as we are not able to observe the empirical distribution function), the functional bound can be used to give lower bound estimates on that maximal difference, thus enabling a conservative Kolmogorov-Smirnov test to be developed. A dynamic programming algorithm can be used to determine the maximum difference. Theoretical details of the test are in Bedford and Meilijson[11] with more implementation details about the dynamic programming and application examples given in Bedford and Meilijson.[12]

1.5. Conservatism of Independence

As we noted at the beginning, commercial reliability databases make use of competing risk models in interpreting and presenting data. Common assumptions are that underlying times to failure from different failure modes are exponential and that the censoring is independent. Clearly, assumptions need to be made, but one can ask whether or not these assumptions are going to bias the numerical results in any consistent way. It turns out that the functional bounds can be used to show that the assumption of independence tends to give an optimistic assessment of the marginal of X_1. It was shown in Bedford and Meilijson[13] that any other dependence structure would have given a *higher* estimate of the failure rate; the "independent" failure rate is the lower endpoint of the interval of constant failure rates compatible with the competing risk data. This result, which is generalized to a wider class of parametric families (those ordered by monotone likelihood ratio), is essentially based on a constraint implied by differentiating the functional bound at the origin. Consider the following simple example from Bedford and Meilijson.[13]

Suppose that the lifetime Y of a machine is exponentially distributed, there are two failure modes with failure times X_1, X_2, and its cause of failure $I = 1, 2$ is independent of Y. Suppose also that the failure rate of Y is θ and let $P(I = 1) = p_1$. The unique independent model (X_1, X_2) for this observed data joint distribution makes X_1 and X_2 exponentially distributed with respective failure rates $p_1\theta$ and $(1 - p_1)\theta$.

Suppose now that we believe that X_1 has a marginal exponential distribution, but we are not sure about possible dependence. What is then the range of possible values of its failure rate λ? Since the sub-distribution function of X_1 is $G_1(t) = (1 - e^{-\theta t})p_1$ and λ must satisfy $F_1'(t) = (1 - e^{-\lambda t})' \geq G_1'(t)$ for all $t > 0$, we have that $p_1\theta \leq \lambda \leq \theta$. The upper Peterson bound tells us that $\lambda \leq \theta$. However, the "light touch" condition discussed above shows that equality is not possible thus giving a range of feasible λ values as

$$p_1\theta \leq \lambda < \theta \ . \qquad (1)$$

This shows that the lowest, most optimistic, failure rate compatible with the general competing risk bounds is that obtained from the independence assumption.

1.6. The Bias of Independence

As a further illustration of the potential bias created by assumption of independence when it is not clearly appropriate we consider an example from reliability prediction discussed in Bedford and Cooke.[14] This gives the theoretical background to work carried out for a European Space Agency project. Four satellites were due to be launched to carry out a scientific project, which required the functioning of all satellites throughout the mission period of 2 years post launch. A preliminary assessment of the satellite system reliability using a standard model (independent exponentially distributed lifetimes of subsystems) indicated a rather low probability of mission success. The study included discussions with mission engineers about the mechanisms of possible mission failure which suggested that, in contrast to the assumptions made in the standard reliability modeling, mission risk was highly associated to events such as vibration damage during launch and satellite rocket firing, and to thermal shocks caused by moving in and out of eclipses. A new model was built in which the total individual satellite failure probability over the mission profile was kept at that predicted in the original reliability model. Now, however, a mission phase-dependent failure rate was used to capture the engineering judgement associating higher failure rates to the above-mentioned mission phases, and a minimally informative copula was used to couple failure rates to take account of residual couplings not captured by the main effects.

The effect of positive correlation between satellite lifetimes is to improve system lifetime. This may seem counter intuitive–especially to those with an engineering risk background used to thinking of common cause effects as being a "bad thing." However, it is intuitively easy to understand. For if we sample 4 positively correlated lifetime variables all with marginal distribution F then the realizations will tend to be more tightly bunched as compared to the 4 independent realizations from the same marginal distribution F. Hence the series system lifetime, which is the minimum of the satellite lifetimes, will tend to be larger when the lifetimes are more highly correlated. This is illustrated in Fig. 1.4, taken from Bedford and Cooke,[14] which is based on simulation results.

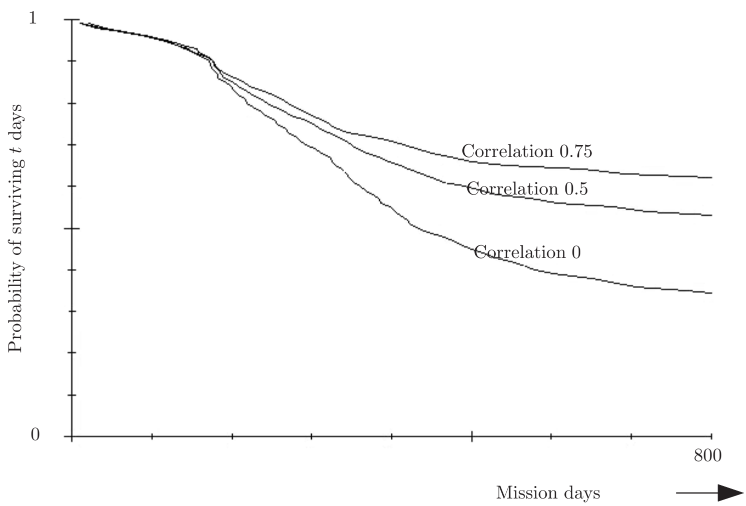

Fig. 1.4. Cluster system survival probability depending on correlation.

1.7. Maintenance as a Censoring Mechanism

One very interesting and practical area of application of competing risk ideas is in understanding the impact of maintenance on failure data. There is a lot of anecdotal evidence to suggest that there can be major differences in performance between different plant (for example nuclear plant) that are not caused by differences in design or different usage patterns, and that are therefore likely to be related to different maintenance practices and/or policies.

Theoretical models for maintenance optimization require us to know the lifetime distribution of the components in question. Therefore it is major significance to understand the impact that competing risk has in masking our knowledge of that distribution due to current maintenance (or other) practices that censor the lifetime variable of interest.

1.7.1. *Dependent copula model*

The sensitivity of predicted lifetime to the assumptions made about dependency between PM and failure time was investigated in Bunea and Bedford[15] using a family of dependent copulae. This is based on the results in Zheng and Klein[16] where a generalization of the Kaplan-Meier estimator is defined that gives a consistent estimator based on an assumption about the underlying copula of (X_1, X_2).

To illustrate this in an optimization context, the interpretation given was that of choosing an age-replacement maintenance policy. Existing data, corresponding to failure and/or unscheduled PM events, is taken as input. In order to apply the age replacement maintenance model we need the lifetime distribution of the equipment. Hence it is necessary to "remove" the effect of the unscheduled PM from the lifetime data.

The objective was to show what the costs of assuming the wrong model would be, if one was trying to optimize an age replacement policy. The conclusion of this paper is that the costs of applying the wrong model can indeed be very substantial. The costs arise because the age replacement interval is incorrectly set for the actual lifetime distribution when the lifetime distribution has been incorrectly estimated using false assumptions about the form of censoring. See Fig. 1.5, taken from Bunea and Bedford.[15]

1.7.2. *Random clipping*

This is not really a competing risk model, but is sufficiently close to be included here. The idea, due to Cooke, is that the component life time is exponential, and that the equipment emits a warning at some time before the end of life. That warning period is independent of the lifetime of the equipment, although we only see those warnings that occur while the equipment is in use. The observable data here is the time at which the warning occurs. An application of the memoryless property of the exponential distribution shows that the observable data (the warning times) has the same distribution as the underlying failure time. Hence we can estimate the MTBF just by the mean of the the warning times data. This model, and the following one, is discussed further in Bedford and Cooke.[17]

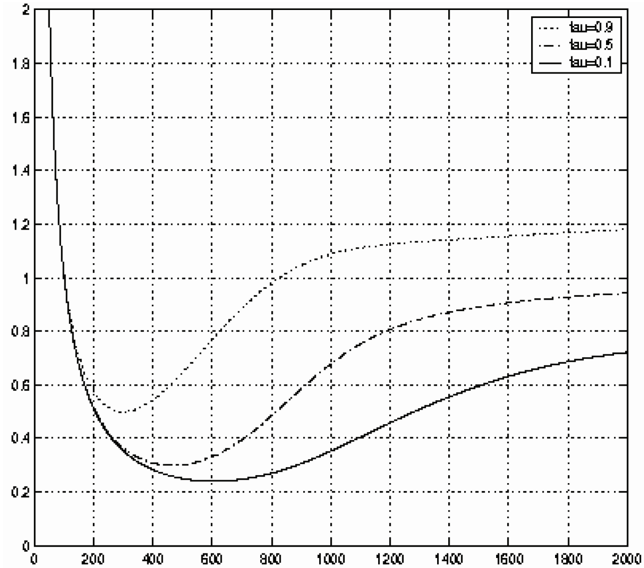

Fig. 1.5. Optimizing to the wrong model: effect of assuming different correlations.

1.7.3. *Random signs*

This model[18] uses the idea that the time at which PM might occur is related to the time of failure. The PM (censoring time) X_2 is equal to the failure time X_1 plus an a random quantity ξ, $X_2 = X_1 - \xi$. Now, while ξ might not be statistically independent of X_1, its sign is. In other words PM is trying to be effective and to occur round about the time of the failure, but might miss the failure and occur too late. The chance of failure or PM is independent of the time at which the failure would occur. This is quite a plausible model, but is not always compatible with the data. Indeed, Cooke has shown that the model is consistent with the distribution of observable data if and only if the normalized subsurvivor functions are ordered. In other words, this model can be applied if and only if the normalized subsurvivor function for X_1 always lies above that for X_2.

1.7.4. *LBL model*

This model, proposed in Langseth and Lindqvist,[19] develops a variant of the random signs model in which the likelihood of an early (that is, before failure) intervention by the maintainer is proportional to the unconditional failure intensity for the component. Maintenance is possibly imperfect in this model. The model is identifiable.

1.7.5. *Mixed exponential model*

A new model capturing a class of competing risk data not previously covered by the above was presented in Bunea et al.[20] The underlying model is that X_1 is drawn from a mixture of two exponential distributions, while X_2 is also exponential and independent of X_1. This is therefore a special case of the independent competing risks model, but in a very specific parametric setting. Important features of this model that differ from the previous models are: (1) The normalized subdistribution functions are mixtures of exponential distribution functions, (2) The function $\Phi(t)$ increases continuously as a function of t. This model was developed for an application to OREDA data in which these phenomena were observed.

1.7.6. *Delay time model*

The Delay time model[21] is well known within the maintenance community. Here the two times X_1 and X_2 are expressed in terms of a warning variable and supplementary times, $X_1 = W + X_1{}'$, $X_2 = W + X_2{}'$, where $W, X_1{}', X_2{}'$ are mutually independent life variables. As above we observe the minimum of X_1 and X_2.

In the case that these variables are all exponential it can be shown (see Hokstadt and Jensen[22]) that (a) The normalized subdistribution functions are equal and are exponential distribution functions, (b) The function $\Phi(t)$ is constant as a function of t.

1.8. Loosening the Renewal Assumption

The main focus of the paper is on competing risks, and this implies that when we consider a reliability setting we are assuming that the data can be interpreted as though generated by a renewal process. In practice this is not always a good assumption. In general, the higher the risk is that is potentially caused by failure of the equipment and the lower the cost of repair, the more likely it is that maintenance crew will try to get the system back to a "good as new state."

As discussed in the introduction, the competing risk identifiability question is part of the general issue of model identifiability. It is natural therefore to consider loosening the renewal assumption. In the context of a series system it is quite natural to think about three types of maintenance programs. The first is "good as new," that is, when one component fails, all components are restored to as good as new, thus enabling us to make an assumption of a renewal process. The second is partial renewal: only the component that fails is restored to as good as new. The fact that the other component(s) are not new may lead to extra stress being placed on all components. The third possibility is "as bad as old" where a minimal repair is applied to the failed component that restores it to the functioning state but leaves the failure intensity for the whole system in the same state as just before the failure. This situation is discussed by Bedford and Lindqvist[23] where it is assumed that each component has a failure intensity depending on the components own lifetime plus another term that depends on the age of each component.

In this context, identifiability means that we can estimate the failure intensities of the components using data from a single socket. This means that we start off with a single unit and replace the components according to the maintenance policy that was determined, recording failure times as we go.

As stated above, the "good as new" policy essentially means that we are in the classical competing risk situation. Our model class is sufficiently general that it is not identifiable. The least "intensive" maintenance policy, the "bad as old" policy is also not identifiable (indeed, from a single socket data we never revisit any times, so cannot possibly make estimates of failure probabilities). The partial repair policy, however, is rather different. Here, at least under some quite reasonable technical conditions, we are able to identify the model. The reason for this is enlightening. We can consider the vector valued stochastic process, which tells us the current ages of the components at each time point. This process has no renewal properties whatsoever in the minimal repair case. In the full and partial repair cases however it can be considered as a continuous time continuous state Markov process. However, in the full repair case there is no mixing, whereas the partial repair case (under suitable technical, but weak, conditions) the process is ergodic. This ergodicity implies that a single sample path will (with probability 1) visit the whole sample space, thus enabling us to estimate the complete intensity functions.

1.9. Conclusion

(Dependent) competing risk models are increasingly being developed to support the analysis of reliability data. Because of the competing risk problem we cannot identify the joint distribution or marginal distributions without making nontestable assumptions. The validation of such models on a statistical basis is therefore impossible, and validation must therefore be of a "softer" nature, relying on assessment of the engineering and organizational context. This is particularly so in the area of maintenance policy. Clearly, there is a whole area of modeling that can be developed.

Acknowledgments

I would like to thank my various collaborators over the last few years including Isaco Meilijson (Tel Aviv University), Bo Lindqvist (NTNU Trondheim), Cornel Bunea (George Washington University), Hans van der Weide (TU Delft), Helge Langseth (NTNU Trondheim), Sangita Karia Kulathinal (National Public Health Institute, Helsinki), Isha Dewan (ISI New Delhi), Jayant Deshpande (Pune University), Catalina Mesina (Free University, Amsterdam) and especially Roger Cooke (TU Delft) who introduced me to the area.

References

1. R. Cooke and T. Bedford, Reliability Databases in Perspective, *IEEE Transactions on Reliability* **51**, 294-310 (2002).
2. M. Crowder, Classical competing risks, Chapman and Hall/CRC (2001).
3. J. V. Deshpande, Some recent advances in the theory of competing risks, Presidential Address, Section of Statistics, Indian Science Congress 84th Session (1997).
4. A. Tsiatis, A nonidentifiablility aspect in the problem of competing risks, Proceedings of National Academy of Science, USA **72**, 20-22 (1975).
5. E. L. Kaplan and P. Meier, On the identifiability crisis in competing risks analysis, *Jourmal of American Statistical Association* **53**, 457-481 (1958).
6. A. Nádas, On estimating the distribution of a random vector when only the smallest coordinate is observable, *Technometrics* **12**, 923-924 (1970).
7. D. R. Miller, A note on independence of multivariate lifetimes in competing risk models, *Ann. Statist.* **5**, 576-579 (1976).
8. J. A. M. Van der Weide and T. Bedford, Competing risks and eternal life, *Safety and Reliability* (Proceedings of ESREL'98), S. Lydersen, G.K. Hansen, H.A. Sandtorv (eds), Vol. **2**, 1359-1364, Balkema, Rotterdam (1998).
9. A. Peterson, Bounds for a joint distribution function with fixed subdistribution functions: Application to competing risks, *Proc. Nat. Acad. Sci. USA* **73**, 11-13 (1976).

10. M. Crowder, On the identifiability crisis in competing risks analysis, *Scand. J. Statist.* **18**, 223-233 (1991).
11. T. Bedford and I. Meilijson, A characterization of marginal distributions of (possibly dependent) lifetime variables which right censor each other, *Annals of Statistics* **25**, 1622-1645 (1997).
12. T. Bedford and I. Meilijson, A new approach to censored lifetime variables, *Reliability Engineering and System Safety* **51**, 181-187 (1996).
13. T. Bedford and I. Meilijson, The marginal distributions of lifetime variables which right censor each other, in H. Koul and J. Deshpande (Eds.), *IMS Lecture Notes Monograph Series* **27** (1995).
14. T. Bedford and R. M. Cooke, Reliability Methods as management tools: dependence modelling and partial mission success, *Reliability Engineering and System Safety* **58**, 173-180 (1997).
15. C. Bunea and T. Bedford, The effect of model uncertainty on maintenance optimization, *IEEE Transactions in Reliability* **51**, 486-493 (2002).
16. M. Zheng and J. P. Klein, Estimates of marginal survival for dependent competing risks based on an assumed copula, *Biometrika* **82**, 127-138 (1995).
17. T. Bedford and R. Cooke, Probabilistic Risk Analysis: Foundations and Methods, Cambridge University Press (2001).
18. R. Cooke, The total time on test statistic and age-dependent censoring, *Stat. and Prob. Let.* **18** (1993).
19. H. Langseth and B. Lindqvist, A maintenance model for components exposed to several failure mechanisms and imperfect repair, in B. L. K. Doksum (Ed.), *Mathematical and Statistical Methods in Reliability*, pp. 415-430. World Scientific Publishing (2003).
20. C. Bunea, R. Cooke, and B. Lindqvist, Competing risk perspective over reliability databases, in H. Langseth and B. Lindqvist (Eds.), Proceedings of Mathematical Methods in Reliability, NTNU Press (2002).
21. A. Christer, Stochastic Models in Reliability and Maintenance, Chapter A review of delay time analysis for modelling plant maintenance, Springer (2002).
22. P. Hokstadt and U. Jensen, "Predicting the failure rate for components that go through a degradation state," *Safety and Reliability*, Lydersen, Hansen and Sandtorv(eds) Balkema, Rotterdam, pp. 389-396, (1998).
23. T. Bedford and B. H. Lindqvist, The Identifiability Problem for Repairable Systems Subject to Competing Risks, *Advances in Applied Probability* **36**, 774-790 (2004).