Reliability Engineering and System Safety 134 (2015) 59–65

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

Unobserved heterogeneity in the power law nonhomogeneous Poisson process

Zeytu Gashaw Asfaw, Bo Henry Lindqvist *

Department of Mathematical Sciences, Norwegian University of Science and Technology, N-7491 Trondheim, Norway

ARTICLE INFO

Article history:
Received 29 March 2014
Received in revised form
2 September 2014
Accepted 2 October 2014
Available online 13 October 2014

Keywords:
Repairable system
Likelihood function
Frailty
Monte Carlo simulation

ABSTRACT

A study of possible consequences of heterogeneity in the failure intensity of repairable systems is presented. The basic model studied is the nonhomogeneous Poisson process with power law intensity function. When several similar systems are under observation, the assumption that the corresponding processes are independent and identically distributed is often questionable. In practice there may be an unobserved heterogeneity among the systems. The heterogeneity is modeled by introduction of unobserved gamma distributed frailties. The relevant likelihood function is derived, and maximum likelihood estimation is illustrated. In a simulation study we then compare results when using a power law model without taking into account heterogeneity, with the corresponding results obtained when the heterogeneity is accounted for. A motivating data example is also given.

1. Introduction

In the reliability literature, systems are generally classified as either non-repairable or repairable (see, e.g., Ascher and Feingold [1]). Non-repairable systems are those that do not get repaired when they fail. Thus, non-repairable system can fail only once, and a lifetime model such as the Weibull distribution provides the distribution of the time to failure of such systems.

On the other hand, repairable systems are those systems (machines, industrial plants, software, etc.) which, in the event of a failure, can be restored to satisfactory operation by any action, including part replacements or changes to adjustable settings. But, to what extent can the system perform after being returned back to its regular operation? We may have that the system's performance is in the same state that the system had at the start of the operation, which means an “as good as new” condition. Or, its performance may be returned to the same state as before the failure, which means an “as bad as old” condition.

The latter case is usually referred to as a “minimal repair”, modeled by a nonhomogeneous Poisson process (NHPP). Minimal repair thus means that a failed system is restored just back to a functioning state, and after repair the system continues as “if nothing had happened”. This implies that the likelihood of system failure, right after a failure and subsequent repair, is the same as it was immediately before the failure. Note that repair times in this kind of modeling are assumed to be negligible.

NHPP models, which are the main concern of this paper, are useful due to their flexible assumption that events are occurring randomly in time, with rates which may vary with time. This is in contrast to the more established homogeneous Poisson process (HPP), where the rate of events is constant in time.

The present paper is concerned with the problem of predicting the behavior of a system based on failure data from several similar systems. There is a well established theory for analysis of data for NHPPs. But as discussed for example in Lindqvist [9], there may be unobserved heterogeneity between the monitored systems which, if overlooked, may lead to non-optimal or possibly completely wrong decisions. An intuitive way of interpreting heterogeneity is to imagine an unknown covariate, with values that may vary between systems, and leading to an unexpected variation in the failure intensity of the different processes (see, e.g., Slimacek and Lindqvist [14]). Still it is believed that heterogeneity has been neglected in many reliability applications, and it is the purpose of the present paper, through a simulation study and analysis of a real data set, to point to some of the consequences that may result from not including heterogeneity in a model for repairable systems.

A striking example of heterogeneity is given by some data presented by Bhattacharjee et al. [2], presenting failure data for motor operated closing valves in safety systems at two boiling water reactor plants in Finland. Failures of the type “External Leakage” were considered for 104 valves with a follow-up time of 9 years. The data show an apparently unnormal variation in the number of failures per valve, suggesting a heterogeneity

* Corresponding author.
E-mail addresses: zeytugashaw@yahoo.com (Z.G. Asfaw), bo@math.ntnu.no (B.H. Lindqvist).

http://dx.doi.org/10.1016/j.ress.2014.10.005
0951-8320/

between valves. In their analysis, Bhattacharjee et al. [2] stressed the importance of taking heterogeneity into consideration and concluded that even very simple models may describe the heterogeneous behavior successfully. For illustration, their data are given in Section 5 together with a statistical analysis using the approach of the present paper.

The study of heterogeneity in reliability analyses of repairable systems is not new. An early reference is Engelhardt and Bain [5], and a detailed treatment is given in the monograph by Rigdon and Basu [13]. In the case of lifetime modelling, there is a rich literature on heterogeneity under the name of *frailty models*, a basic reference being Vaupel, Manton and Stallard [15]. In the reliability context, recent references are Cha and Finkelstein [3] and the monograph Finkelstein [6].

The present paper is structured as follows. In Section 2 we give the formal definition of the NHPP and derive the likelihood function for the case of a power law intensity model, which will be the basic model considered in the paper. Section 3 introduces heterogeneity between systems, in the form of individual unobserved multiplicative *frailties* defined for each system, assumed independent and, for simplicity of exposition, gamma distributed with unit expectation. The likelihood function of data from several systems under heterogeneity is developed, leading to explicit expressions for the power law parameter estimates in the case where each process is observed on the same time frame. Section 4 is devoted to a simulation study; first giving an algorithm for simulation of data, and then performing a comprehensive simulation study with the aim of illustrating the main messages of the paper. The data from Bhattacharjee et al. [2] are analysed in Section 5, while some concluding remarks are given in the final Section 6.

2. The classical power law process

2.1. Characteristics of an NHPP model

An NHPP model is fully characterized by the intensity function, $w(t)$, commonly denoted ROCOF (Rate of occurrence of failures), see e.g. Rausand and Høyland [12]. It is furthermore convenient to introduce the cumulative rate function $W(t)=\int_0^t w(s)\,ds$, later called the CROCOF (cumulative ROCOF).

As is well known, the number of failures experienced in a time interval from 0 to t, $N(t)$, is Poisson-distributed with parameter $W(t)$, for any t, so that in particular $E[N(t)]=W(t)$ and $Var[N(t)]=W(t)$.

2.2. The power law NHPP

For illustration, we shall in this paper concentrate on the most commonly used parameterization of the NHPP, namely the power law model. One reason for its popularity is that the ROCOF as a function of t is of the same form as the hazard rate of a Weibull distribution. Hence the time to first failure of the power law NHPP is Weibull distributed. Because of this, the power law model is sometimes denoted the Weibull process.

The CROCOF of the power law is given by (see, e.g., Rausand and Høyland [12])

$$W(t)=\lambda t^{\beta} \quad \text{for } \lambda>0,\ \beta>0. \tag{1}$$

Thus, by differentiation, the ROCOF of the power law process is

$$w(t)=W'(t)=\lambda\beta t^{\beta-1}.$$

This intensity function was introduced in Crow [4] as a stochastic model for the Duane reliability growth postulate. The parameter β in the power law model gives information about the system as follows; if $0<\beta<1$, then the system is *improving* (happy); if $\beta>1$, then the system is *deteriorating* (sad); and if $\beta=1$ the model reduces to an HPP.

2.3. Maximum likelihood estimation in the power law NHPP

Suppose that data are available from m independent systems governed by NHPPs with the same intensity function $w(t)$, where system j is observed in the time interval $[S_j,T_j]$, $j=1,2,\ldots,m$, with events observed at times $t_{1j},t_{2j},\ldots,t_{n_j j}$.

The likelihood function of these data is given by (see, e.g., Meeker and Escobar [10])

$$L=\prod_{j=1}^{m}\left\{\prod_{i=1}^{n_j} w(t_{ij})\right\}e^{-[W(T_j)-W(S_j)]}, \tag{2}$$

which is the product of the individual likelihoods of each of the m systems. The log-likelihood function, which is usually easier to work with, is hence

$$l=\log L=\sum_{j=1}^{m}\left[\left\{\sum_{i=1}^{n_j}\log\ w(t_{ij})\right\}-[W(T_j)-W(S_j)]\right].$$

For the power law model, with the parameterization given in (1), the log-likelihood function is given by

$$l=n\log\ \lambda+n\log\ \beta+(\beta-1)\sum_{j=1}^{m}\sum_{i=1}^{n_j}\log\ t_{ij}-\lambda\sum_{j=1}^{m}\left[T_j^{\beta}-S_j^{\beta}\right] \tag{3}$$

where $n=\sum_{j=1}^{m} n_j$.

For simplicity we shall in the following assume that $S_j=0$ and $T_j=\tau$ for $j=1,\ldots,m$. Thus all the m processes are observed on the time interval from 0 to a fixed time τ. As we shall see, this simplifies several results, while the main ideas prevail. Using the standard method of finding the maximum likelihood estimators (MLEs) of $\hat{\lambda}$ of λ and $\hat{\beta}$ of β by setting the partial derivatives of the log-likelihood function with respect to each parameter equal to zero, we get from (3),

$$\frac{\partial l}{\partial\lambda}=\frac{n}{\lambda}-m\tau^{\beta}=0,$$

which implies

$$\hat{\lambda}=\frac{n}{m\tau^{\hat{\beta}}} \tag{4}$$

Similarly we get the equation

$$\frac{\partial l}{\partial\beta}=\frac{n}{\beta}+\sum_{j=1}^{m}\sum_{i=1}^{n_j}\log\ t_{ij}-\lambda m\tau^{\beta}\log\ \tau=0,$$

which by using (4) leads to

$$\hat{\beta}=\frac{n}{n\log\ \tau-\sum_{j=1}^{m}\sum_{i=1}^{n_j}\log\ t_{ij}}.$$

This gives an explicit solution for $\hat{\beta}$, which can afterwards be substituted in the expression (4) for $\hat{\lambda}$.

We now consider the Fisher information matrix for the computation of variances and covariances of the MLEs. The Fisher information matrix is used to measure the amount of information that the observed data carries about the unknown parameters. It is defined as

$$I(\lambda,\beta)=E\begin{bmatrix}-\frac{\partial^2 l(\lambda,\beta)}{\partial\lambda^2} & -\frac{\partial^2 l(\lambda,\beta)}{\partial\lambda\partial\beta}\\ -\frac{\partial^2 l(\lambda,\beta)}{\partial\lambda\partial\beta} & -\frac{\partial^2 l(\lambda,\beta)}{\partial\beta^2}\end{bmatrix},$$

and a straightforward computation in our case implies that

$$I(\lambda,\beta)=E\begin{bmatrix}\frac{n}{\lambda^2} & m\tau^\beta \log\,\tau\\ m\tau^\beta\log\,\tau & \frac{n}{\beta^2}+\lambda m\tau^\beta(\log\,\tau)^2\end{bmatrix}. \quad (5)$$

The only random element in the above matrix is n, so we need to substitute $E(n)=m\lambda\tau^\beta$ to get the final expression

$$I(\lambda,\beta)=m\tau^\beta\begin{bmatrix}\frac{1}{\lambda} & \log\,\tau\\ \log\,\tau & \lambda(\frac{1}{\beta^2}+(\log\,\tau)^2)\end{bmatrix}.$$

Standard theory for maximum likelihood tells us that the maximum likelihood estimates for λ and β are, for large samples, approximately normally distributed centered at the true parameter values and with variances given as the diagonal elements, respectively, of the inverse of the Fisher information matrix. It is hence possible to estimate those variances by inverting the estimated matrix $I\left(\hat{\lambda},\hat{\beta}\right)$.

Alternatively, one may invert the so called observed Fisher information matrix, which is the matrix (5) where one does not take the expectation, but instead uses the observed value of n, and substitute maximum likelihood estimates for the parameters.

3. Heterogeneity in the power law model

3.1. Heterogeneous NHPPs

Consider an NHPP with intensity function $w(t)$. With the inclusion of heterogeneity, this model is modified to assuming that the intensity is given by

$$w_a(t)=aw(t),$$

where $w(t)$ is the basic ("baseline") intensity function, and a is an unobserved positive constant, which may vary from system to system. More precisely, a is assumed to be a positive random variable with mean 1 and variance $\delta \ge 0$. The idea is that in the case of m systems, each system has its own value of a, i.e., $a_1,a_2,\ldots,a_m$, which are assumed to be independent draws from this distribution. Following the survival analysis literature, the a_j will in the following be denoted as *frailties*.

Although there are several potential distributions for the frailties a_j, we shall here apply the most commonly used one, namely the gamma distribution. The popularity of this distribution as a frailty distribution is due to both mathematical convenience and often good fit to actual data. There is, however, no physical justification to prefer gamma frailties instead of other models.

The density of the two-parameter gamma distribution is generally given as

$$f(a)=\frac{a^{k-1}e^{-a/\theta}}{\theta^k\Gamma(k)}$$

for $a>0$, where $k>0$ is the shape parameter and $\theta>0$ is the scale parameter. The corresponding expected value and variance are, respectively, $k\theta$ and $k\theta^2$. Since we require $E(a)=1$ and $Var(a)=\delta$, we use $k=1/\delta$ and $\theta=\delta$. The density of a hence becomes

$$h(a)=\frac{a^{1/\delta-1}e^{-a/\delta}}{\Gamma\left(\frac{1}{\delta}\right)\delta^{1/\delta}} \quad (6)$$

Fig. 1 shows several densities of gamma distributions with expected value 1.

The likelihood function for data from m systems modeled by NHPPs was given in (2). We now study the changes needed when including a frailty a. We use a similar argument as before, but now with $Var(a)=\delta$ as an additional parameter.

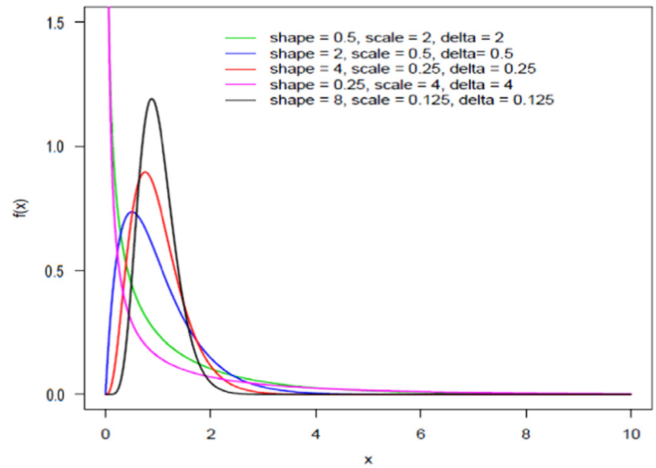

Fig. 1. Graphs of gamma densities with expected value 1.

Now the likelihood function for system j, for given value of the frailty, a_j, is

$$L_j(a_j)=\left\{\prod_{i=1}^{n_j} w(t_{ij})\right\}a_j e^{-a_j[W(T_j)-W(S_j)]}.$$

Since a_j is an unobservable random variable, the contribution to the full likelihood from this system is obtained by unconditioning with respect to a_j, which in practice will mean to compute the expected value of $L_j(a_j)$ with respect to the distribution of a_j. Since, furthermore, a_j is gamma distributed with expected value 1, and hence has probability density function (6), the expected value of $L_j(a_j)$ is

$$\begin{aligned}L_j&=E[L_j(a_j)]\\&=\int L_j(a_j)h(a_j)\,da_j\\&=\int\left\{\prod_{i=1}^{n_j}w(t_{ij})\right\}a_j e^{-a_j[W(T_j)-W(S_j)]}\frac{a_j^{1/\delta-1}e^{-a_j/\delta}}{\Gamma\left(\frac{1}{\delta}\right)\delta^{1/\delta}}\,da_j\\&=\frac{\prod_{i=1}^{n_j}w(t_{ij})}{\Gamma(\frac{1}{\delta})\delta^{1/\delta}}\int_0^\infty a_j^{r_j-1}e^{-a_j s_j}\,da_j\end{aligned}$$

where $r_j=n_j+\frac{1}{\delta}$ and $s_j=W(T_j)-W(S_j)+\frac{1}{\delta}$.

Now it is easy to show that $\int_0^\infty a^{r-1}e^{-sa}\,da=\Gamma(r)/s^r$ for all $r,s>0$, so we get

$$L_j=\frac{\prod_{i=1}^{n_j}w(t_{ij})}{\Gamma(\frac{1}{\delta})\delta^{1/\delta}}\frac{\Gamma\left(n_j+\frac{1}{\delta}\right)}{\left[W(T_j)-W(S_j)+\frac{1}{\delta}\right]^{n_j+1/\delta}}.$$

3.2. Maximum likelihood estimation for the heterogeneous power law NHPP

Specializing the above to the power law (1), we get

$$L_j=\frac{\lambda^{n_j}\beta^{n_j}\left(\prod_{i=1}^{n_j}t_{ij}\right)^{\beta-1}\Gamma\left(n_j+\frac{1}{\delta}\right)}{\Gamma\left(\frac{1}{\delta}\right)\delta^{1/\delta}\left[\lambda T_j^\beta-\lambda S_j^\beta+\frac{1}{\delta}\right]^{n_j+1/\delta}}.$$

Further, assuming $S_j = 0$, $T_j = \tau$ for all j, and then taking log and summing over all the m systems, we obtain the full log-likelihood

$$l(\lambda,\beta,\delta) = n \log \lambda + n \log \beta + (\beta - 1)\sum_{j=1}^{m}\sum_{i=1}^{n_j} \log t_{ij} + \sum_{j=1}^{m} \log \Gamma\left(n_j + \frac{1}{\delta}\right) - \left[m \log \Gamma\left(\frac{1}{\delta}\right) + m\frac{1}{\delta} \log \delta + \left[n + \frac{m}{\delta}\right] \log\left[\lambda\tau^{\beta} + \frac{1}{\delta}\right]\right]$$

In order to find the maximum likelihood estimators for λ, β, δ we first compute

$$\frac{\partial l(\lambda,\beta,\delta)}{\partial \lambda} = \frac{n}{\lambda} - \left[\frac{\tau^{\beta}}{\lambda\tau^{\beta} + \frac{1}{\delta}}\right]\left[n + \frac{m}{\delta}\right],$$

which when set to 0 implies

$$\hat{\lambda} = \frac{n}{m\tau^{\hat{\beta}}}. \tag{7}$$

Next we compute

$$\frac{\partial l(\lambda,\beta,\delta)}{\partial \beta} = \frac{n}{\beta} + \sum_{j=1}^{m}\sum_{i=1}^{n_j} \log t_{ij} - \left[\frac{\lambda\tau^{\beta} \log \tau}{\lambda\tau^{\beta} + \frac{1}{\delta}}\right]\left[n + \frac{m}{\delta}\right],$$

which when set to 0, using (7), leads to

$$\hat{\beta} = \frac{n}{n \log \tau - \sum_{j=1}^{m}\sum_{i=1}^{n_j} \log t_{ij}}. \tag{8}$$

Thus, $\hat{\lambda}$ and $\hat{\beta}$ are exactly the same functions of the data as for the power law case without frailities. (Note that this would not be the case if the observation time intervals were not all equal for all the m processes.)

The partial derivative of the log likelihood with respect to δ involves the digamma function ψ defined by

$$\psi(x) = \frac{d}{dx} \log \Gamma(x) = \frac{\Gamma'(x)}{\Gamma(x)},$$

and we get

$$\begin{aligned}\frac{\partial l(\lambda,\beta,\delta)}{\partial \delta} &= -\frac{1^2}{\delta}\sum_{j=1}^{m} \psi\left(n_j + \frac{1}{\delta}\right) + \frac{m}{\delta^2}\psi\left(\frac{1}{\delta}\right) - m\left[-\frac{1}{\delta^2} \log \delta + \frac{1}{\delta^2}\right] \\ &\quad - \left[-\frac{m}{\delta^2} \log\left[\lambda\tau^{\beta} + \frac{1}{\delta}\right] - \frac{1}{\delta^2}\left[\frac{n + \frac{m}{\delta}}{\lambda\tau^{\beta} + \frac{1}{\delta}}\right]\right] \\ &= -\frac{1}{\delta^2}\sum_{j=1}^{m} \psi\left(n_j + \frac{1}{\delta}\right) + \frac{m}{\delta^2}\psi\left(\frac{1}{\delta}\right) + \frac{m}{\delta^2} \log \delta \\ &\quad - \frac{m}{\delta^2} + \frac{m}{\delta^2} \log\left[\lambda\tau^{\beta} + \frac{1}{\delta}\right] + \frac{1}{\delta^2}\left[\frac{n + \frac{m}{\delta}}{\lambda\tau^{\beta} + \frac{1}{\delta}}\right] \\ &= -\frac{1}{\delta^2}\left\{\sum_{j=1}^{m} \psi\left(n_j + \frac{1}{\delta}\right) - m\psi\left(\frac{1}{\delta}\right) - m \log \delta + m\right\} \\ &\quad + \frac{1}{\delta^2}\left\{m \log\left[\lambda\tau^{\beta} + \frac{1}{\delta}\right] - \frac{n + \frac{m}{\delta}}{\lambda\tau^{\beta} + \frac{1}{\delta}}\right\}\end{aligned}$$

The likelihood equation given by equating this to 0 is simplified by substituting the estimators for λ and β, which gives an equation of δ alone. No explicit expression for the maximum likelihood estimator $\hat{\delta}$ is available, however, so a numerical method like Newton–Raphson's method needs to be used.

4. Simulations

4.1. Simulate systems from the power law process with heterogeneity

The following probabilistic property of the NHPP can help us to simulate event times of an NHPP from that of an HPP. Namely, if $U_1, U_2, \ldots$ are the event times of an HPP with intensity 1, then it can be shown that $W^{-1}(U_1), W^{-1}(U_2), \ldots$ are the event times of an NHPP with CROCOF $W(t)$. Here the inverse function $W^{-1}(u)$ is uniquely determined from $W(t)$ if $w(t) > 0$ for all t.

We now apply this property to the power law NHPP with CROCOF defined by (1). Then it is seen that

$$W^{-1}(u) = (u/\lambda)^{1/\beta}.$$

Thus if $U_1, U_2, \ldots$ are the event times of an HPP with intensity 1, we obtain a simulated realization of the power law NHPP with given parameters λ and β as $(U_1/\lambda)^{1/\beta}$, $(U_2/\lambda)^{1/\beta}, \ldots$. Note here that the HPP, $U_1, U_2, \ldots$, can be simulated by first drawing U_1 from an exponential distribution with parameter 1; then letting $U_2 = U_1 + V_2$ where V_2 is an independent draw from the exponential distribution with parameter 1; and so on by adding new independent V_i from the exponential distribution with parameter 1, until the boundary time τ is reached for the transformed variables $(U_i/\lambda)^{1/\beta}$.

In order to simulate from a power law process with gamma distributed frailty, we first draw the value of a for each process, and then for the jth process replace λ by $a_j\lambda$ in the above simulation strategy.

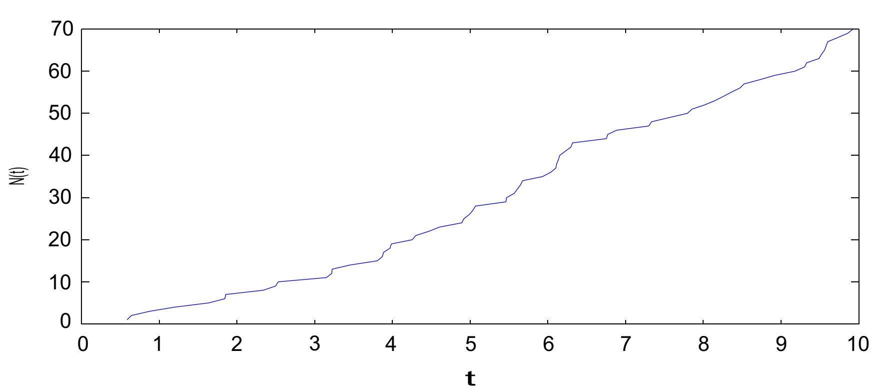

Fig. 2. Cumulative number of failures, $N(t)$, versus time for a power law process with $\lambda = 2$ and $\beta = 1.5$.

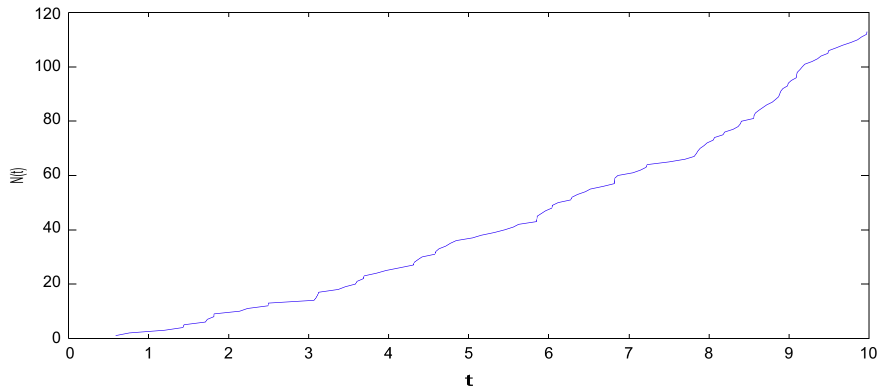

Fig. 3. Cumulative number of failures, $N(t)$, versus time for a single power law process with $\lambda=2$, $\beta=1.5$ and $\delta=0.2$.

Table 1
Data simulated from the heterogeneity model with varying δ, and estimation done as discussed in Section 3.2.

Data	m	True value			n		Estimates					
10 000	20	λ	β	δ	Ave.	St.D	$\hat{\lambda}$		$\hat{\beta}$		$\hat{\delta}$	
							Ave.	St.D	Ave.	St.D	Ave.	St.D
		2	1.5	0	63.1451	8.0541	2.0059	0.2057	1.5008	0.0430	0.0021	0.0036
				0.1	63.0648	21.4421	2.0075	0.2456	1.5005	0.0421	0.0942	0.0352
				0.2	63.2006	29.6606	2.0042	0.2853	1.5013	0.0425	0.1897	0.0651
				0.4	63.0469	40.1756	2.0033	0.3450	1.5009	0.0426	0.3795	0.1208
				0.6	63.3880	49.3983	2.0063	0.4033	1.5014	0.0431	0.5718	0.1750
				0.8	63.0149	56.1749	2.0050	0.4452	1.5013	0.0430	0.7511	0.2223
				1	63.4131	64.1655	2.0033	0.5023	1.5021	0.0432	0.9309	0.2610
			1	0	20.0288	4.4658	2.0051	0.2526	1.0019	0.0509	0.0064	0.0119
				0.2	19.9885	9.9238	1.9987	0.3231	1.0023	0.0509	0.1879	0.0774
				0.4	20.2335	13.4745	2.0052	0.3802	1.0026	0.0508	0.3742	0.1305
				0.6	19.8679	15.8587	1.9968	0.4285	1.0043	0.0514	0.5520	0.1769
				0.8	20.1814	18.2127	1.9923	0.4691	1.0051	0.0518	0.7158	0.2107
				1	20.4585	20.8251	1.9843	0.5149	1.0083	0.0524	0.8679	0.2421
			0.75	0	11.2920	3.3512	1.9992	0.2653	0.7535	0.0507	0.0116	0.0215
				0.2	11.8200	6.8717	2.0535	0.3142	0.7502	0.0493	0.1888	0.0924
				0.4	11.2779	7.8143	1.9923	0.3934	0.7557	0.0519	0.3620	0.1381
				0.6	11.1454	9.2048	1.9846	0.4330	0.7572	0.0522	0.5234	0.1763
				0.8	11.1906	10.4320	1.9702	0.4873	0.7624	0.0539	0.6685	0.2054
				1	11.1918	11.4700	1.9590	0.5285	0.7663	0.0552	0.7917	0.2274

4.2. Simulated single processes

Throughout the simulation study we assume that there are $m=20$ systems, each observed on the fixed time interval from 0 to $\tau=10$. The failure processes will be power law NHPPs with basic ROCOF $\lambda\beta t^{\beta-1}$, for varying values of λ and β, but possibly with heterogeneity obtained by multiplying the system intensities by independent random variables a from the gamma distribution with expected value 1 and variance δ.

Fig. 2 shows a simulation of a single power law process observed on the time interval [0,10], where parameter values are $\lambda=2$ and $\beta=1.5$. Fig. 3 shows similarly a single process from a power law process with the same λ and β, but with a frailty parameter δ $(=Var(a))=0.2$.

We also illustrate, for the data in Fig. 2, the estimated parameters, and the inverse of the observed Fisher information matrix. The ML estimates are $\hat{\lambda}=1.9926$ and $\hat{\beta}=1.4999$, while the inverse of the observed information matrix is

$$I^{-1}\left(\hat{\lambda},\hat{\beta}\right)=\begin{bmatrix} 0.0407 & -0.0082 \\ -0.0082 & 0.0018 \end{bmatrix}.$$

By taking the square roots of the diagonal elements of this matrix we obtain the estimated standard errors of $\hat{\lambda}$ and $\hat{\beta}$, respectively, 0.2018 and 0.0423.

4.3. Simulation study

For each setup of parameters we consider 10,000 simulations, each consisting of $m=20$ systems. The results are shown in Table 1. For each simulation we estimate parameters and their standard errors by maximum likelihood, and report averages of these numbers based on the 10,000 simulations. These numbers can hence be viewed as approximations of expected values of the parameters, which enables consideration of possible bias in the estimators. Further, the columns named by “St.D” give empirical standard errors of the corresponding 10,000 computed estimates, obtained as the square roots of empirical variances.

The table also gives the averages (approximation of expected values) of number of failures in the time interval [0, 10] for each parameter combination, as well as the corresponding standard deviations, computed as square roots of the empirical variances.

Note that each such number is based on $m \times 10,000$ simulated processes, i.e. 200,000 simulations.

In the discussion below we use the property that the maximum likelihood estimators $\hat{\lambda}$ and $\hat{\beta}$ are given by the same function of the data whatever be the value of δ (see Section 3.2). Thus, the estimates are the same whether we assume the power law model without frailty or the power law model with frailty. (It should be noted that this would not be the case if the observation intervals of the m systems differed). Here it enables us to reach at many interesting conclusions regarding heterogeneity.

The main conclusions to be drawn from the simulation study are given below.

4.3.1. An ordinary power law is anticipated, while there may be an unrecognized heterogeneity

Consider first the situation where one thinks that the ordinary power law model is the true model, but that in reality there is heterogeneity between the $m=20$ systems.

For each combination of λ and β is seen that, as δ increases, the average number of failures per system is approximately constant, which is in fact exactly true by a theoretical computation of expected values. However, the standard deviations (St.D) of the number of failures per system increase with δ. This is caused by the heterogeneity, which in practice means that some system will have a higher failure intensity, while others will have a lower intensity than for the base case. (On average, the intensity will be the same as for the no heterogeneity case, however, since the variables a_i have expected value 1.)

Now let us consider the estimated parameters. It is remarkable that neither the expected value nor the standard error of $\hat{\beta}$ are much influenced by the heterogeneity. As regards $\hat{\lambda}$, this is close to the true value, 2, but is increasingly biased downwards as δ increases. This is most clearly seen when $\beta \neq 1$. Its standard error, on the other hand, clearly increases with δ for all cases. Practical implications of these results are: (i) For predictions of number of failures of new systems, by an erroneous assumption of *no* heterogeneity, one gets too short predicted intervals for the number of failures in a given time period. In fact, the expected value is correctly estimated, but the variation could be much bigger than expected if heterogeneity is not accounted for. For example, suppose the true values are $\lambda=2$ and $\beta=0.75$. In a model not taking heterogeneity into account, the number of failures will be predicted to be in the interval (expected value $\pm$ 2 standard deviations) approximately from 4.5 to 18. If the true heterogeneity variance is, e.g., $\delta=0.4$, this interval would be much wider, namely from 0 to 27. (ii) A similar problem is seen for the estimation of λ. Here one would get a too optimistic estimate for the standard error by assuming no heterogeneity, and also a downward bias will be present in the estimate.

Table 2
Failure times for 104 closing valves, with follow-up time of 3286 h, at two boiling water reactor plants in Finland. Failure type is "External Leakage".

System #	Failure times							
1	610	614	943	2024	2087	2104	2399	2525
2	126	323	943	1132	2087	2399	2426	
3	860	915	1606	3181				
4	10	19	104	2352				
5	293	2567						
6	2434	2676						
7	1963							
8	1262							
9	2501							
10	1963							
11	132							
12	1623							
13	3127							
14	3211							
15	1225							
16	1222							
17–104	–							

4.3.2. The correct model, a power law model with heterogeneity, is used for statistical inference

In this case, the table contains information on all the maximum likelihood estimators of the model, $\hat{\lambda}, \hat{\beta}, \hat{\delta}$.

It is seen that the maximum likelihood estimator $\hat{\delta}$ slightly underestimates the true value of δ. Further, its standard error increases with δ as should be expected. The conclusion is that the estimator $\hat{\delta}$ seems to behave quite satisfactorily. The properties of the estimators $\hat{\lambda}$ and $\hat{\beta}$ are in fact already discussed, since the formulas for their maximum likelihood estimators are the same with and without including δ in the model.

5. The data example from Bhattacharjee et al. (2003)

Recall the closing valve failure example (Bhattacharjee et al. [2]), which was considered in the introduction. There are $m=104$ systems, each observed on the time interval [0,3286] (h). The failure times are given in Table 2.

Let the model be as given in Section 3.2. Using the approach in that subsection we get $\hat{\lambda}=4.594 \times 10^{-4}$, $\hat{\beta}=0.8215$, $\hat{\delta}=8.340$. The estimate of δ thus reveals a considerable heterogeneity between the systems. This heterogeneity is also clearly visible from computation of the standard deviation of the number of failures per system. In fact, the empirical standard deviation of number of observed failures per system is 1.206. On the other hand, using the estimates for λ and β and assuming an ordinary power law model without heterogeneity, the standard deviation would be estimated to 0.3558. The latter number is obtained as the square root of the estimated cumulative intensity of the process at time 3286, using the fact that expected value equals variance for a Poisson distributed random variable.

6. Discussion and concluding remarks

The motivation for the present paper is the fact that an unobserved heterogeneity between observed systems of the same kind may, if ignored, lead to wrong conclusions or bad predictions of system behavior.

We show, in a partly tutorial manner, how a possible heterogeneity between systems may be included in a statistical investigation of repairable systems. In order to simplify the exposition we consider the fairly standard case of a power law nonhomogeneous Poisson process with heterogeneity of the gamma type. The advantage of the approach is that it leads to relatively simple formulas and procedures, while main ideas, pitfalls and possible remedies are clearly demonstrated.

The model that we consider can be viewed as a hierarchical model where the failure process for each system is conditional on the value of the gamma-distributed frailty a. While our approach can be viewed as an empirical Bayes approach, a fully Bayesian approach is of course also feasible. One example of such a study is George et al. [8], who did a fully Bayesian approach for the corresponding HPP model with heterogeneity.

As already noted, the main reason for assuming gamma-distributed frailties is mathematical tractability. In certain cases, however, the assumption may be questioned, and there are examples where the use of other distributions would lead to different conclusions, see, e.g., the note by Nelder et al. [11]. This problem is, however, beyond the scope of this paper.

We have, further, assumed the observation lengths, τ, for each system to be equal. This had the nice effect of giving simple explicit expressions (7) and (8) for the parameter estimates of λ and β, which are not influenced by the assumption of heterogeneity. It is in principle straightforward to generalize the likelihood functions to the case of different observation lengths, which however would destroy the simple computation of estimates. Thus, in order to simplify the presentation, we decided to keep the observation lengths equal. For some applications, this assumption would even be natural (see, e.g., the example of Section 5).

As a final remark on our assumptions, while we have assumed the heterogeneity of the systems to be connected to the scale-parameter λ, it might in applications be natural to expect that also the trend parameter β varies between systems. In principle we can derive the likelihood function and do the analyses also for this case, but of course the analyses would be more involved.

An apparently different extension of the power law process has recently been considered by Le Gat [7], essentially assuming that the rate of occurrence of failures of a system at any time depends on the perviously experienced number of failures. In a forthcoming paper we will explore the relation between this dynamic extension of the power law process and the heterogeneity extension considered in the present paper.

References

[1] Ascher H, Feingold H. Repairable systems reliability: modeling, inference, misconceptions and their causes. New York: M. Dekker; 1984.
[2] Bhattacharjee M, Arjas E, Pulkkinen U. Modeling heterogeneity in nuclear power plant valve failure data. In: Lindqvist BH, Doksum KA, editors. Mathematical and statistical methods in reliability. Singapore: World Scientific Publishing; 2003. p. 341–53.
[3] Cha JH, Finkelstein M. Some notes on unobserved parameters (frailties) in reliability modeling. Reliab Eng Sys Saf 2014;123:99–103.
[4] Crow LH. Reliability analysis for complex, repairable systems. In: Proschan F, Serfling RJ, editors. Reliability and biometry. Philadelphia, Pennsylvania: SIAM; 1974. p. 379–410.
[5] Engelhardt M, Bain LJ. Statistical analysis of a compound power-law model for repairable systems. IEEE Trans Reliab 1987;R-36:392–6.
[6] Finkelstein M. Failure rate modelling for reliability and risk. Springer, London; 2010.
[7] LeGat Y. Extending the Yule process to model recurrent pipe failures in water supply networks. Urban Water J 2013. http://dx.doi.org/10.1080/1573062X.2013.783088.
[8] George EI, Makov UE, Smith AFM. Conjugate likelihood distributions. Scand J Stat 1993;20:147–56.
[9] Lindqvist BH. On the statistical modelling and analysis of repairable systems. Stat Sci 2013;21:532–51.
[10] Meeker WQ, Escobar LA. Statistical methods for reliability data. New York: Wiley; 1998.
[11] Nelder JA, George EI, Makov UE, Smith AFM. Reader reaction: a re-analysis of the pump-failure data [with response]. Scand J Stat 1994;21:187–91.
[12] Rausand M, Høyland A. System reliability theory: models, statistical methods, and applications. New York: Wiley; 2004.
[13] Rigdon SE, Basu AP. Statistical methods for the reliability of repairable systems. New York: Wiley; 2000.
[14] Slimacek V, Lindqvist BH. Failure rate of wind turbines modeled by homogeneous Poisson process with covariates and unobserved heterogeneity. In: Safety, reliability and risk analysis: beyond the horizon. Proceedings of the European safety an reliability conference, ESREL 2013. CRC Press; 2014.
[15] Vaupel JW, Manton KG, Stallard E. The impact of heterogeneity in individual frailty on the dynamics of mortality. Demography 1979;16:439–54.