

Exact Statistical Inference for Some Parametric Nonhomogeneous Poisson Processes

Bo Henry Lindqvist¹, Gunnar Taraldsen²

¹Department of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway.

²SINTEF ICT, Trondheim, Norway.

Abstract. Nonhomogeneous Poisson processes (NHPPs) are often used to model recurrent events, and there is thus a need to check model fit for such models. We study the problem of obtaining exact goodness-of-fit tests for certain parametric NHPPs, using a method based on Monte Carlo simulation conditional on sufficient statistics. A closely related way of obtaining exact confidence intervals in parametric models is also briefly considered.

Keywords. Conditional test, log linear NHPP, power low NHPP, sufficiency.

MSC: 60G55, 62B05, 62M02, 62N05, 65C05.

1 Introduction

Nonhomogenous Poisson processes (NHPP) are widely used as models for events occurring in time. In practice it may be of interest to check the NHPP property by statistical tests given observed data. In a recent paper, Lindqvist and Rannestad [7] presented a Monte Carlo approach for goodness-of-fit testing in NHPP models, where conditional tests are performed by simulating conditional samples given a sufficient statistic

Bo Henry Lindqvist(bo@math.ntnu.no), Gunnar Taraldsen(✉)(gunnar.taraldsen@sintef.no)

Received: June 10, 2011; Accepted: January 23, 2012

under the hypothesized NHPP model. In particular, they presented tests for the power law and the log-linear law NHPP models. It was noted that while exact algorithms for the power law case are well known in the literature, the availability of exact tests for the log-linear case are less known (Gaudoin [4] probably presented the first test of this kind).

The approach of the present paper is closely related to the one of [7], but differs in the way the conditional samples are obtained. More specifically, the present approach is based on the Monte Carlo simulation method of Lindqvist and Taraldsen [8].

It should be stressed that our emphasis is on how to obtain exact tests when a goodness-of-fit statistic W is given. The construction or choice of a test statistic W is hence beyond the primary scope of this article. A nice and informative discussion of how to perform goodness-of-fit testing in NHPP models is given in Baker [1]. Commonly used tests are based on the Cramer-von-Mises test or the Kolmogorov-Smirnoff test. We refer to [7] for more information on the choice of test statistics and relevant references.

2 Observations, Likelihood Functions and Sufficient Statistics

We assume that an NHPP is observed from time $t = 0$ and until n events have occurred. In reliability applications the events are usually failures of some repairable system, and the above given observation scheme is denoted as “failure truncation” or “failure censoring”. Most of our results may be modified to the case of “time truncation” or “time censoring”, where the process is observed until a given time t_0 , but for the sake of brevity this will not be done here.

Suppose we observe an NHPP with intensity function $\lambda(t)$ and let the n first event times be denoted $T = (T_1, \dots, T_n)$. The log likelihood function resulting from observed failure times $\{t_j; j = 1, 2, \dots, n\}$ is then given as (Crowder et al. [3]),

$$\sum_{j=1}^n \log \lambda(t_j) - \int_0^{t_n} \lambda(u) du. \quad (1)$$

Two popular parametrizations of NHPPs are the *power law*, with intensity function given by

$$\lambda(t) = abt^{b-1}$$

and the *log-linear law*, with intensity function

$$\lambda(t) = \exp(a + bt).$$

Substituting each of these functions into (1) and using the factorization criterion for sufficiency (Casella and Berger [2]), gives that the two-dimensional statistic

$$S = \left(\log T_n, \sum_{j=1}^{n-1} \log T_j \right) \quad (2)$$

is sufficient in the power law case, while

$$S = \left(T_n, \sum_{j=1}^{n-1} T_j \right) \quad (3)$$

is sufficient in the log-linear case (see also [7]).

3 Conditional Testing Given a Sufficient Statistic

Let T be the vector of failure times as described in the previous section, and consider testing of the null hypothesis H_0 that these data come from an NHPP of a particular parametric type. More precisely, consider null hypotheses of the form $H_0 : T_1, T_2, \dots$ are observations from an NHPP with intensity function $\lambda(t; \theta)$, where $\lambda(\cdot; \theta)$ is a parametric nonnegative function with θ varying in a specified finite-dimensional parameter set. Note that the parameter value θ is unspecified in the null hypothesis. Simple examples of such parametric functions are the power law and log-linear law intensity functions considered in the previous section, and which will be the objects of main study here.

Let $W \equiv W(T)$ be any test statistic for revealing departure from the null model, assuming for simplicity that the null hypothesis is rejected for large values of W . Suppose that $S \equiv S(T)$ is a sufficient statistic for the unknown parameters under the null hypothesis, such as for example the ones given in (2) and (3).

An α -level conditional test of H_0 rejects, conditionally given $S = s$, when $W \geq k(s)$, where $k(s)$ is a critical value chosen such that

$$P_{H_0}(W \geq k(s) | S = s) = \alpha. \quad (4)$$

We here assume for convenience that exact equality to the right of (4) is always possible, which will be the case in our applications. It is readily seen by unconditioning that (4) implies $P_{H_0}(\text{reject } H_0) = \alpha$. Thus any α -level conditional test is an α -level (unconditional) test, while the opposite need not be true.

The clue of using a conditional test is that, by sufficiency of S , the conditional distribution of W given $S = s$ is independent of unknown parameters, and hence that it is, at least in principle, possible to calculate $k(s)$. Note also that we will need to compute $k(s)$ for the observed value of S only. Still this might, however, be difficult or even practically impossible. In the following we therefore describe an equivalent formulation of the problem which instead involves computation of (conditional) p -values.

Let the observed data be T_{obs} , and let $s_{obs} = S(T_{obs})$ and $w_{obs} = W(T_{obs})$ be the observed values of, respectively, the sufficient statistic and the test statistic. Instead of computing $k(s_{obs})$ we shall consider the conditional p -value,

$$p_{obs} = P_{H_0}(W \geq w_{obs} | S = s_{obs}). \quad (5)$$

Now (4) and (5) together imply that $p_{obs} \leq \alpha$ if and only if $w_{obs} \geq k(s_{obs})$, and hence the α -level conditional test is equivalent to the test which rejects the null hypothesis if $p_{obs} \leq \alpha$. The key to performing the conditional test is hence reduced to computation of the conditional p -value in (5).

In [7] this computation was made by simulating a large number, M say, of realizations T^* from the conditional distribution of T given $S = s_{obs}$, then computing $W^* = W(T^*)$ for each of them, and finally approximating the conditional p -value p_{obs} by the relative frequency

$$\hat{p}_{obs} = \#\{W^* \geq w_{obs}\} / M.$$

Basically, this is also what we will do here, but now using the method considered in Lindqvist and Taraldsen [8] which is a way of computing conditional expectations of the form $E(\phi(T) | S = s)$ when S is a sufficient statistic. Note that \hat{p}_{obs} given above indeed can be written in this form.

4 Monte Carlo Conditioning on a Sufficient Statistic

As described in the previous section, we let the model for the observation T under the null hypothesis be specified in terms of a finite-dimensional

parameter θ . The treatment of the present section is valid more generally than the case of parametric NHPP models.

Suppose that for a given value of θ we can simulate realizations of T by $T = \chi(U, \theta)$ for some function χ and a random vector U with known distribution. Further, suppose $S(T) \equiv S$ is a sufficient statistic for θ . Then of course S can be simulated by the function $\tau(U, \theta) \equiv S(\chi(U, \theta))$. How to choose U and χ is shown in the next section for, respectively, the power law and the log-linear NHPP models.

Consider now the computation of conditional expectations of the form $E(\phi(T)|S = s)$ for given functions ϕ , where s is the observed value of S . By the definition of sufficiency, this conditional expectation is independent of θ and the clue of the approach of the present paper is that it can be expressed in terms of ordinary expectations of functions of U (Theorems 1 and 2 below).

Theorem 4.1. (Lindqvist and Taraldsen [8]) *Suppose that θ and S take values in R^k and suppose that the equation $\tau(u, \theta) = s$ has the unique solution $\theta = \hat{\theta}(u, s)$ for each fixed u and s . Let $f(\theta)$ be a nonnegative function defined on the parameter space, and let $\det \partial_\theta \tau(u, \theta)$ be the determinant of the matrix of partial derivatives of $\tau(u, \theta)$ for fixed u . Then*

$$E(\phi(T)|S = s) = \frac{E_U \left[\phi(\chi(U, \hat{\theta}(U, s))) \left| \frac{f(\theta)}{\det \partial_\theta \tau(U, \theta)} \right|_{\theta=\hat{\theta}(U, s)} \right]}{E_U \left[\left| \frac{f(\theta)}{\det \partial_\theta \tau(U, \theta)} \right|_{\theta=\hat{\theta}(U, s)} \right]} \tag{6}$$

It is tacitly assumed above that f is such that the expectations exist and such that the denominator is positive, but f may otherwise be arbitrarily chosen. It should be noted that f can be thought of as a (possibly improper) density of a distribution for θ , similar to a prior distribution in Bayesian statistics. For such connections to Bayesian and also fiducial statistics we refer to [8]. Lindqvist and Taraldsen [9] argue that traditional noninformative priors such as the Jeffreys' prior for f will generally work well. The problem of choosing f will be further discussed in connection with the application in Section 5.2.

The practical consequence of Theorem 4.1 is that the expectations can be computed by simulation by drawing independent versions of the random vector U (which has a known distribution) and averaging the expressions inside the expectations on the right hand side of (6). It follows that if a function $f(\cdot)$ can be chosen so that $|f(\theta)/\det \partial_\theta \tau(U, \theta)|_{\theta=\hat{\theta}(U, s)}$ does not depend on U , then we will have $E(\phi(T)|S = s) = E_U [\phi(\chi(U, \hat{\theta}(U, s)))]$ which means that the function $\chi(U, \hat{\theta}(U, s))$ can be used to

sample directly from the conditional distribution of T given $S = s$. Unfortunately, it is not always possible to find such an f , but the following sufficient condition can be given.

Theorem 4.2. (Lindqvist and Taraldsen [8]) *Let the situation be as in Theorem 4.1. Assume that there exist functions r and $\tilde{\tau}$ with $\tau(u, \theta) = \tilde{\tau}(r(u), \theta)$, such that the equation $\tilde{\tau}(v, \theta) = s$ has a unique solution $v = \tilde{v}(\theta, s)$ for all (θ, s) . Then $\chi(U, \hat{\theta}(U, s))$ is distributed as the conditional distribution of T given $S = s$.*

The new assumption of Theorem 4.2 means that $\tau(u, \theta)$ depends on u only through $r(u)$, which usually has a much lower dimension than u , and has the property that for given θ , $r(u)$ is uniquely determined by s . Note that $\tilde{v}(\theta, S)$ is in this case a pivotal quantity in the classical meaning in statistical theory. The new condition of the theorem is therefore called the *pivotal* condition.

As we shall see in the next section, the pivotal condition turns out to hold for an NHPP with the power law intensity function, while it does not hold for the log-linear intensity in which case Theorem 4.1 should be used.

5 Conditional Simulation for Parametric NHPP Models

It is well known (e.g. Ross [14]) that the first n events of an NHPP with intensity function $\lambda(\cdot)$ can be simulated by letting $U = (U_1, U_2, \dots, U_n)$ be the first n events of a homogeneous Poisson process with unit intensity, and then letting

$$T_j = \Lambda^{-1}(U_j) ; j = 1, \dots, n, \tag{7}$$

where Λ^{-1} is the inverse function of the cumulative intensity function $\Lambda(t) \equiv \int_0^t \lambda(u) du$.

5.1 Power Law Intensity

From Section 2 follows that $T_j = (U_j/a)^{1/b}$. With notation from the previous section we simulate T by

$$\chi(u; a, b) = ((u_1/a)^{1/b}, \dots, (u_n/a)^{1/b}),$$

while the sufficient statistic $S = (\log T_n, \sum_{j=1}^{n-1} \log T_j)$ is simulated by

$$\tau(u; a, b) = ((\log u_n - \log a)/b, (\sum_{j=1}^{n-1} \log u_j - (n-1) \log a)/b).$$

Thus the pivotal condition in Theorem 4.2 holds with $r(u) = (\log u_n, \sum_{j=1}^{n-1} \log u_j)$. Letting the observed times be $t = (t_1, \dots, t_n)$ and letting $s = (\log t_n, \sum_{j=1}^{n-1} \log t_j)$, it is straightforward to obtain $\hat{\theta}(u, s) \equiv (\hat{a}(u, s), \hat{b}(u, s))$ by solving $\tau(u; a, b) = s$ for a and b . The solutions can in fact be written in the simple form

$$\begin{aligned} \hat{b}(u, s) &= \frac{\sum_{j=1}^{n-1} \log(u_j/u_n)}{\sum_{j=1}^{n-1} \log(t_j/t_n)}, \\ \hat{a}(u, s) &= u_n/t_n^{\hat{b}(u,s)}. \end{aligned}$$

Samples $\tilde{t} = (\tilde{t}_1, \dots, \tilde{t}_n)$ from the conditional distribution of T given $S = s$ can now be obtained by first sampling $u = (u_1, \dots, u_n)$ from a unit intensity homogeneous Poisson process and then computing $\tilde{t} = \chi(u; \hat{\theta}(u, s))$. We get in this way

$$\tilde{t}_j = (u_j/u_n)^{1/\hat{b}(u,s)} t_n; \quad j = 1, \dots, n,$$

which can be easily simulated since the u_j/u_n are distributed as the order statistics of a sample of $n - 1$ independent uniforms on $[0, 1]$. In fact this is the same formula for \tilde{t}_j that one would get from Lemma 1 of [7] by a straightforward modification to the failure truncation case.

5.2 Log Linear Intensity

In this case (7) becomes $T_j = \log(1 + be^{-a}U_j)/b$, so T is simulated by

$$\chi(u; a, b) = (\log(1 + be^{-a}u_1)/b, \dots, \log(1 + be^{-a}u_n)/b),$$

while $S = (T_n, \sum_{j=1}^{n-1} T_j)$ is obtained as

$$\tau(u; a, b) = (\log(1 + be^{-a}u_n)/b, \sum_{j=1}^{n-1} \log(1 + be^{-a}u_j)/b).$$

It is clear that the pivotal condition of Theorem 4.2 is not satisfied here. This should be no surprise since it is well known that the log-linear NHPP model has no interesting pivotal statistic. In order to compute

p -values of conditional tests, we thus have to use Theorem 4.1 with some properly chosen function $f(a, b)$ (see below for a discussion of possible choices).

First we find $\hat{\theta}(u, s) \equiv (\hat{a}(u, s), \hat{b}(u, s))$ by solving for a, b the equations

$$t_n = \frac{1}{b} \log(1 + be^{-a}u_n), \tag{8}$$

$$\sum_{j=1}^{n-1} t_j = \frac{1}{b} \sum_{j=1}^{n-1} \log(1 + be^{-a}u_j). \tag{9}$$

From (8) we readily get

$$e^{-a} = \frac{e^{bt_n} - 1}{bu_n} \tag{10}$$

which by substitution in (9) gives

$$b \sum_{j=1}^{n-1} t_j = \sum_{j=1}^{n-1} \log \left[1 + \frac{u_j}{u_n} (e^{t_n b} - 1) \right], \tag{11}$$

which is an equation in b only. By differentiating twice it is seen that the right hand side of (11) is convex in b . Further consideration leads to the conclusion that equation (11), in addition to the trivial solution $b = 0$, has a unique additional solution, which is the one that solves our problem, and which is easily found by numerical methods. The solution for a is then finally obtained from (10).

Now the entries of $\tilde{t} = \chi(u; \hat{a}(u, s), \hat{b}(u, s))$ are given by

$$\tilde{t}_j = \log \left(1 + (e^{t_n \hat{b}(u, s)} - 1)(u_j/u_n) \right) / \hat{b}(u, s),$$

and a straightforward computation shows that the determinant $\det \partial_{a,b} \tau(u; a, b)$ with $(\hat{a}(u, s), \hat{b}(u, s))$ substituted for (a, b) is given by

$$\left(t_n \sum_{j=1}^{n-1} h_j - h_n \sum_{j=1}^{n-1} t_j \right) / \hat{b}(u, s)^2,$$

where $h_j = ((e^{t_n \hat{b}} - 1)(u_j/u_n)) / (1 + (e^{t_n \hat{b}} - 1)(u_j/u_n))$ and $\hat{b} \equiv \hat{b}(u, s)$. As for the power law case we can simulate the u_j/u_n for $j = 1, \dots, n - 1$ as the order statistics from a set of $n - 1$ i.i.d. uniforms on $[0, 1]$.

What remains in order to estimate the expectations in (6) for the current application is then to choose $f(a, b)$. It is tempting to try first

a constant function $f(a, b) \equiv 1$. It turns out, however, that in this particular case this leads to non-convergence of the method. This was in fact pointed out for a modified version of the current problem in personal communication with Federico O'Reilly, who found satisfactory results by using Jeffreys' prior as the choice for f . The computation of Jeffreys' prior seems, however, to be unnecessarily complicated in the present case. It turns out on the other hand that the problem of using a constant function f is due to the behaviour of f around $b = 0$. In order to avoid this problem, the following function f was used:

$$f(a, b) = \begin{cases} 0 & \text{if } b \in (-\epsilon, \epsilon) \\ 1 & \text{otherwise} \end{cases}$$

for some $\epsilon > 0$, where suitable values can be found by trial and error. It should furthermore be noted that in the present case there is no function f for which the weight $|f(\theta)/\det \partial_{\theta}\tau(U, \theta)|_{\theta=\hat{\theta}(U, s)}$ does not depend on U . This is because, as already noted, the pivotal condition (see end of Section 4) is not satisfied in the log-linear case.

6 Statistical Inference in NHPP Models

6.1 Goodness-of-Fit Testing

The identity (7) is equivalent to $U_j = \Lambda(T_j)$. It follows that if $\Lambda(\cdot)$ is the cumulative intensity function of the NHPP T_1, T_2, \dots , then $\Lambda(T_1), \Lambda(T_2), \dots$ is a homogeneous Poisson process with unit intensity. Thus by a standard result on Poisson processes ([14]), the transformed times $V_j = \Lambda(T_j)/\Lambda(T_n)$ for $j = 1, \dots, n-1$ are distributed as the order statistic of $n-1$ i.i.d. uniform variables on $[0, 1]$. If $\Lambda^*(\cdot)$ is an estimate of $\Lambda(\cdot)$ based on data $t = (t_1, \dots, t_n)$, then we shall define estimated transformed times v_1^*, \dots, v_{n-1}^* by

$$v_j^* = \Lambda^*(t_j)/\Lambda^*(t_n).$$

One then anticipates these to behave much similar to uniform variables, and goodness-of-fit testing may thus be based on comparing the behaviour of the estimated transformed times to that of uniform variates. The quantities v_j^* are in fact the basis of a large class of goodness-of-fit tests for NHPPs.

Baker [1], see also Rigdon [13], showed for the power law process that when Λ^* is based on the maximum likelihood estimates for the parameters, then the estimated transformed times are pivots, i.e. have

distributions which do not depend on the unknown parameters. This follows in fact from the representation $t_j = (u_j/a)^{1/b}$ in Section 5.1, noting that $b^* = -n/\sum_{j=1}^{n-1} \log(t_j/t_n)$ is the maximum likelihood estimate of b based on the data (Crowder et al. 1991). Then we have

$$v_j^* = (t_j/t_n)^{b^*} = (u_j/u_n)^{-n/\sum_{j=1}^{n-1} \log(u_j/u_n)}$$

which are independent of the parameters.

Baker [1] derived a class of score tests based on the estimated transformed times. A special case, which we shall use for illustration, is

$$W = \sum_{j=1}^n (v_j^* - v_{j-1}^*)^2,$$

where $v_0^* = 0, v_n^* = 1$. This statistic was called the Greenwood statistic in [7]. The null hypothesis of an NHPP of a particular parametric type is rejected for either too small or too large values of this statistic. Note that since the v_j^* are pivots, we can in the power law case compute (by simulation) the unconditional p -values which in this two-sided case becomes

$$2(\min\{P_{H_0}(W \leq w_{\text{obs}}), P_{H_0}(W \geq w_{\text{obs}})\}).$$

In the case of log-linear intensity (Section 5.2) we get

$$v_j^* = (e^{b^*t_j} - 1)/(e^{b^*t_n} - 1),$$

where b^* is the maximum likelihood estimate of b . In Crowder et al. [3] it is shown that b^* is the solution for b of the equation

$$\bar{t} + \frac{1}{b} - \frac{t_n e^{bt_n}}{e^{bt_n} - 1} = 0, \tag{12}$$

where $\bar{t} = \sum_{j=1}^n t_j/n$.

Conditional p -values for the test based on W can then be computed from Theorem 4.1 with $\phi(T) = I(W \leq w_{\text{obs}})$ and $\phi(T) = I(W \geq w_{\text{obs}})$, and using the further specifications in Section 5.2. The resulting test was considered also in [7], where a Gibbs sampling algorithm for computation of conditional p -values was used.

6.2 Exact Confidence Intervals

From the general approach of Lillegård and Engen [6] it follows how the $\hat{b}(u, s)$ of Section 5 can be used to obtain exact confidence intervals for

the parameter b , both for the power law and the log-linear law cases. Draw independent realizations u^1, \dots, u^m of the vector U , for a specified positive integer m . Let s be the observed value of the sufficient statistic and let $\tilde{b}_{(1)} < \dots < \tilde{b}_{(m)}$ be the ordered values of the $\hat{b}(u^j, s)$. Then $(\tilde{b}_{(k)}, \tilde{b}_{(m-k+1)})$ is an exact $1 - 2k/(m + 1)$ confidence interval for b for any m . Note that the intervals will depend on the actually drawn u^j .

For the power law case it can be seen that the above constructed interval, as $m \rightarrow \infty$, is the same as the classical one based on the pivotal statistic $2nb/b^*$, which is known to be chi-square distributed with $2(n - 1)$ degrees of freedom (e.g. Rausand and Høyland [12]).

The intervals obtained by this method for the log-linear case, however, appear to be new. Note that they are not given by a closed expression, even when $m \rightarrow \infty$, but need to be computed by simulation.

6.3 Application to Data Set

We apply the above results to data from a reliability growth program, taken from Leitch [5]. There are $n = 10$ failures, at times

$$103, 315, 801, 1183, 1345, 2957, 3909, 5702, 7261, 8245.$$

Suppose first that it is of interest to know whether the data are consistent with a power law or a log-linear NHPP.

In the power law case we get for the Greenwood statistic W of Section 6.1,

$$w_{\text{obs}} = 0.1263.$$

Simulating the distribution of W as described in Section 6.1, using 10000 repetitions, we find

$$P_{H_0}(W \leq 0.1263) \approx 0.024,$$

which implies some evidence against the power law NHPP. More precisely, as the test based on W is two-sided, we get a p -value of 0.048.

In the log-linear case we compute for the same W ,

$$w_{\text{obs}} = 0.1466, \quad s_{\text{obs}} = (8245, 23576)$$

and

$$P_{H_0}(W \leq w_{\text{obs}} | S = s_{\text{obs}}) \approx 0.217,$$

following the recipe described in Section 5.2 and using the function f given there. This result does not lead to rejection of the log-linear NHPP assumption.

An exact 90% confidence interval for b in the power law model, using the chi-square distributed pivotal statistic as explained in Section 6.2, follows by first computing the maximum likelihood estimate $b^* = 0.6249$. The resulting interval is then $(0.2933, 0.9020)$. Since the interval does not contain 1, we can conclude that we have reliability growth. It should be noted, however, that we have already rejected at 5% level the null hypothesis of power law NHPP. Thus the confidence interval may have lost its meaning and is included mostly for illustration of the method.

For the log-linear case, where the model has not been rejected, an exact 90% confidence interval for b was computed as

$$(-5.60 \cdot 10^{-4}, -3.83 \cdot 10^{-5}),$$

using 10000 simulations of $\hat{b}(u, s)$ as described in Section 6.2. Since the interval is completely on the left side of 0, there is an indication of reliability growth. It may also be of interest to compute the maximum likelihood estimate of b , which by solving the equation (12) is obtained as $b^* = -1.715 \cdot 10^{-4}$.

7 Concluding Remarks

The present paper complements the computational methods of [7] for exact goodness-of-fit testing in parametric NHPP models. The common idea of the two papers is that such tests can be derived by conditioning appropriate test statistics on a sufficient statistic under the hypothesized model. As regards the power law case, although the derivations of the methods of [7] and the method presented here are apparently different, it has been demonstrated that the resulting methods are in fact identical. On the other hand, for the log-linear case, the methods turn out to be quite different in nature. Here [7] uses Gibbs-sampling, following ideas of Lockhart et al. [10], while the present paper uses the weighted sampling suggested in Lindqvist and Taraldsen [8]. One particular difference is that Gibbs sampling produces dependent samples while the weighted sampling is based on independent samples. The final results should, however, still be independent of the method used to compute them. A relevant comparison between the methods might therefore be based on computing times and accuracy. For example, one may record for each of the two methods the necessary computing time to obtain a given precision; or, for a given computing time one may compare the precision of the results of the two methods. Rannestad [11] did such a

study with a conclusion that the Gibbs sampling methods of [7] are more computationally efficient than the methods described in the present paper. The reason for this is most probably the need for solving equations and computing weights in the latter method.

As pointed out in [7], an alternative to the exact testing approach based on conditional simulations considered here is the use of parametric bootstrapping. Indeed, if one relaxes the requirement of *exactly* computed p -values, then bootstrapping is both a powerful and intuitive tool. The basic idea is then to simulate samples from the null hypothesis model by substituting parameter values estimated under the null model.

Acknowledgements

The authors would like to thank the Special Issue Editor and two referees for constructive comments that improved the presentation of this article.

References

- [1] Baker, R. D. (1996), Some new tests of the power law process. *Technometrics*, **38**, 256-265.
- [2] Casella, G. and Berger, R. L. (2002), *Statistical Inference*. (2nd ed.), Pacific Grove, Calif.: Duxbury.
- [3] Crowder, M. J., Kimber, A. C., Smith, R. L., and Sweeting, T. J. (1991), *Statistical Analysis of Reliability Data*. London: Chapman & Hall.
- [4] Gaudoin, O. (1999), CPIT Goodness-of-fit tests for reliability growth models. In: *Statistical and Probabilistic Models in Reliability*, D.C. Ionescu and N. Limnios eds, Boston: Birkhäuser, 27-37.
- [5] Leitch, R. D. (1995), *Reliability Analysis for Engineers*. New York: Oxford University Press.
- [6] Lillegård, M. and Engen, S. (1999), Exact confidence intervals generated by conditional parametric bootstrapping. *Journal of Applied Statistics*, **26**, 447-459.
- [7] Lindqvist, B. H. and Rannestad, B. (2011), Monte Carlo exact goodness-of-fit tests for nonhomogeneous Poisson processes. *Applied Stochastic Models in Business and Industry*, **27**(3), 329-341.

- [8] Lindqvist, B. H. and Taraldsen G. (2005), Monte Carlo conditioning on a sufficient statistic. *Biometrika*, **92**, 451-464.
- [9] Lindqvist, B. H. and Taraldsen, G. (2007), Conditional Monte Carlo based on sufficient statistics with applications. In: *Advances in Statistical Modeling and Inference. Essays in Honor of Kjell A. Doksum*, V. Nair ed, Singapore: World Scientific, 545-562.
- [10] Lockhart, R. A., O'Reilly, F. J., and Stephens, M. A. (2007), Use of the Gibbs sampler to obtain conditional tests, with applications. *Biometrika*, **94**, 992-998.
- [11] Rannestad, B. (2007), Exact statistical inference in nonhomogeneous Poisson processes, based on simulation. Master thesis, Department of Mathematical Sciences, Norwegian University of Science and Technology.
- [12] Rausand, M. and Høyland, A. (2004), *System Reliability Theory: Models, Statistical Methods, and Applications*. (2nd ed.), Hoboken, NJ: Wiley.
- [13] Rigdon, S. E. (1989), Testing goodness-of-fit for the power law process. *Communications in Statistics: Theory and Methods*, **18**, 4665-4676.
- [14] Ross, S. M. (2007), *Introduction to Probability Models*. (9th ed.), Amsterdam: Elsevier.