



University of
Salford
MANCHESTER

Event count distributions from renewal processes : fast computation of probabilities

Baker, RD and Kharrat, T

<http://dx.doi.org/10.1093/imaman/dpx008>

Title	Event count distributions from renewal processes : fast computation of probabilities
Authors	Baker, RD and Kharrat, T
Publication title	IMA Journal of Management Mathematics
Publisher	Oxford University Press
Type	Article
USIR URL	This version is available at: http://usir.salford.ac.uk/id/eprint/43500/
Published Date	2017

USIR is a digital collection of the research output of the University of Salford. Where copyright permits, full text material held in the repository is made freely available online and can be read, downloaded and copied for non-commercial private study or research purposes. Please check the manuscript for any further copyright restrictions.

For more information, including our policy and submission procedure, please contact the Repository Team at: library-research@salford.ac.uk.

Event count distributions from renewal processes: fast computation of probabilities

Rose Baker¹ and Tarak Kharrat²

¹Salford Business School, University of Salford, UK. r.d.baker@salford.ac.uk

²Salford Business School, University of Salford, UK. tarak.kharrat@gmail.com

June 28, 2017

Abstract

Discrete distributions derived from renewal processes, i.e. distributions of the number of events by some time t are beginning to be used in management science, econometrics and health sciences. A new fast method is presented for computation of the probabilities for these distributions. This will enable practitioners in management science to exploit this rich class of models.

We calculate the count probabilities by repeatedly convolving the discretized distribution, and then correct them using Richardson extrapolation. When just one probability is required, a second algorithm is described, an adaptation of De Pril's method, in which the computation time does not depend on the ordinality, so that even high-order probabilities can be rapidly found. Any survival distribution can be used to model the inter-arrival times, which gives models with great flexibility for modelling both underdispersed and overdispersed data. This work could pave the way for the routine use of these distributions as an additional tool for modelling event count data. An empirical example using fertility data illustrates the use of the method and has been fully implemented using an R (R Core Team, 2015) package `Countr` (Baker et al., 2016) developed by the authors and available from the Comprehensive R Archive Network (CRAN).

Keywords

Renewal process; duration dependence; count data; convolution; Richardson extrapolation; hurdle model.

1 Introduction

1.1 Modelling count variables

Modelling a count variable (the number of events occurring in a given time interval) is a common task in management science. The standard approach is to use the Poisson model, where $Y|x \sim \text{Poisson}(E(Y|x) = \exp(x'\gamma))$. Here Y is predicted given covariates with values x , using regression coefficients γ . There are of course many other models; see for example Winkelmann (2013) or Cameron and Trivedi (2013) for a review. The variance of the Poisson distribution equals its mean, and common departures from the Poisson distribution are overdispersion (variance greater than mean) or underdispersion (variance smaller than mean). Another common departure is an abnormal number of zero counts, dealt with by a hurdle model or a zero-inflated model.

The aim of this paper is to introduce and make available to the practitioner a rich class of models that should prove useful in modelling event counts. We have provided software to compute probabilities from these models, without which they would merely be of academic interest. We have produced fast algorithms and packaged them for easy use, so that the practitioner can focus on the problem at hand, and use these models without much worry about computation speed or accuracy. We do not claim that this method can outperform customised methods tailored to specific cases ¹, but rather that it is a fast general method.

An example is given later in the paper, where we fit a simple example of this class of models to a German fertility dataset. This shows off the ability of these models to cope with typical features of count data that cause it to depart from the Poisson distribution, underdispersion in this case.

The models we consider are derived from renewal processes. This is a class of stochastic processes that may be unfamiliar to many. Renewal processes find their main application in the field of maintenance and reliability, where the classic example of a renewal process is the series of replacements of a light bulb. Here the times to failure of the bulbs (the inter-arrival times) are independent and follow a common distribution, often taken as the Weibull distribution. In a modified or delayed renewal process, which we also need to consider, the time to first failure follows a different distribution from later inter-failure times.

1.2 Management Science implications

Renewal processes themselves find application in several areas of management science, e.g. in reliability and maintenance (Marquez et al., 2015), in warranty (Gonzalez-Prida, 2015) and in inventory (Larsen et al., 2008). However, count data (which can be derived from them) are ubiquitous in management science, and indeed in many other fields such as medicine. Some random examples are: in finance, numbers of takeover bids, numbers of defaults and numbers of unpaid instalments. In insurance, numbers of accident or other claims. In maintenance and reliability, number of failures of equipment. A firm may also record the number of patents awarded, the number of innovations, or the number of staff suggestions. In short, management requires the collection of data, and often interesting events are cumulated into intervals, giving rise to count data that must be analysed. A new class of models that can cope with the vagaries of count data is therefore very relevant.

1.3 The Poisson model

We now describe in more detail the Poisson model and its connection to renewal processes, the type of count model generalizing the Poisson model that renewal processes can provide, and previous work by others on computing count probabilities of this type.

The Poisson distribution for the number of events that have occurred by some time t follows from a Poisson process, which is a renewal process where the inter-arrival times are exponentially distributed. The Poisson model restricts the (conditional) variance to be equal to the (conditional) mean. This situation is rarely observed in real life data.

This limitation of the Poisson model results from the memorylessness property of the exponential distribution. This property states that the probability of having an arrival during the next $(t, t + \Delta t]$ time period (where $t > 0$ and $\Delta t > 0$) is independent of when the last arrival occurred. In many situations, this assumption is not realistic and the history of the process can be informative about future occurrences. For example, someone who consulted the doctor many times recently is more likely to have a higher number of doctor visits in the future (they are

¹i.e. specific inter-arrival time distributions

probably ill) than someone who did not. This is usually dealt with using the negative binomial model, where overdispersion is accommodated by making the hazard of a series of visits of an individual a random variable from a gamma distribution.

1.4 Generalizing the Poisson model

The distribution of $N(t)$, the number of renewal events by some time t offers an alternative to the Poisson model that preserves the connection between the count model and the timing process, but allows a more general event count distribution. Inter-arrival times between events are still assumed to be independent and identically distributed but the constant hazard function arising from an exponential distribution is replaced by a nonconstant hazard function. These type of models display *duration dependence* where negative duration dependence is obtained by a decreasing hazard function (of time) and positive duration dependence by an increasing hazard function. This gives a more flexible count distribution, and in particular, allows it to be overdispersed or underdispersed. A very good example in the marketing context was given in McShane et al. (2008), and these models have also been used in sport analytics (Boshnakov et al., 2017).

In a modified or delayed renewal process, the time to the first event has a different distribution from the subsequent inter-arrival periods. This can arise in the context of reliability, where for example a component may initially not be new; this example is given in Cox (1962). Another example is where replacement components are sourced differently than the one initially supplied. Modified renewal processes give rise to more flexible count models than ordinary renewal processes. If the hazard function of the initial distribution is small, the time to first failure will tend to be greater than subsequent inter-failure times. The count distribution will then have a higher probability of a zero count than before. With higher initial failure hazard, one can obtain count distributions with fewer zeros; thus this class of distributions is flexible enough to analyse data with an abnormal number of zero events. Hurdle models (see e.g. Mullahy (1986) for an account of hurdle models) are widely used for modelling when there is an excess of zero counts, and a modified renewal process can be thought of as a type of hurdle model

In the simplest hurdle model, we have a Bernoulli trial, followed by a zero-truncated Poisson distribution for the number of events. Greene (2011, chapter 25) comments apropos of hurdle models that it is difficult to test whether the hurdle is really there or not ('regime splitting' is occurring), as the hurdle model cannot reduce to the Poisson model and so give a nested model. However, modelling with a modified renewal process, we have to test only that the scale of the hazard function for the first event is equal to that for the later events, when the hurdle model reduces to a regular model. This can be done with a chi-squared test derived from the log-likelihood function. Also, tests for under or overdispersion are difficult with hurdle models, where the excess of zeros anyway induces overdispersion. With the modified Weibull process, a test for under or overdispersion even given a hurdle can be carried out by using a chi-squared test based on the log-likelihood to test whether the shape parameter β departs from unity. Renewal processes thus give rise to a rich and tractable class of models, but the slowness or unavailability of methods of computing the probabilities has so far largely prohibited their use.

Winkelmann (1995) was the first to comment on the usefulness of renewal process models and derived a count model based on gamma distributed inter-arrival times. The choice of the gamma distribution was justified by computational necessity. In fact, the reproductive property of the gamma distribution, that sums of independent gamma distributions are gamma distributed, leads to a simple form for the derived gamma count probability.

The remainder of this paper is laid out as follows. We start by reviewing the possible computation methods in Section 2. Section 3 discusses the situation when all probabilities

up to the m th are required. An alternative method is described in Section 4 when only the m th probability is of interest, in which case a faster computation can be done. Improvement by Richardson extrapolation is developed in Section 5. Section 6 contains a discussion on the generalisations to other survival distributions. In Section 7, we re-analyse the same data used in Winkelmann (1995) and compare a sequence of nested models starting with the basic Poisson regression. Using this approach allows us to highlight which features of the model are most critical to describe the data at hand. Future work and concluding remarks can be found in Section 8.

2 Possible computation methods for renewal processes

In this section, we review the possible methods for computing the count probabilities for other survival distributions besides the gamma. Lomnicki (1966) gave a method for computing a count model with Weibull inter-arrival times, based on an expansion of the exponential function into powers of t and also into Poissonian functions. McShane et al. (2008) used the expansion into powers of t to evaluate the discrete distribution probabilities and fit an underdispersed dataset (the one used in Winkelmann (1995) and fitted here). The same approach has been used in Jose and Abraham (2011) and Jose and Abraham (2013) to derive a counting process with Mittag-Leffler and Gumbel inter-arrival times respectively.

An expansion of the negative exponential is slow to converge. We found that this method can be improved by using techniques such as the Euler and van-Wijngaarden transformations (Press et al., 2007, Chapter 5), which are designed to speed up convergence of alternating-sign series. Nevertheless, convergence is not guaranteed for probabilities of large numbers of events and is not efficient if a high degree of accuracy is needed.

Throughout this paper we will use the Weibull distribution as our main example to illustrate the methodology, which can be applied more generally. The survival function $P_0(t)$, which is the probability of zero events by time t , is given by $P_0(t) = \exp(-(\alpha t)^\beta)$. This distribution allows both overdispersion ($\beta < 1$) and underdispersion ($\beta > 1$), and yields the Poisson distribution when $\beta = 1$. Before we develop our methodology to derive flexible count models based on renewal processes, we first summarise the obvious available computational techniques that can be used. They are:

- expand out the exponential, using series transformations to speed up convergence. This is specific to the Weibull renewal process, but could be developed for others;
- use (smart) Monte-Carlo simulation to generate renewal times up to time t and read off the number of events $N(t)$;
- use Laplace transforms, compute the survival distribution generating function, convert to the transform of the required probability, and invert the transform (e.g. Chaudhry et al. (2013));
- similarly, use the fast Fourier transform (FFT) which is often used for doing convolutions;
- evaluate the required probabilities directly as convolution integrals by discretizing the problem. This approach is the more attractive because De Pril (1985) presented a recursive algorithm for computing the probabilities for the sum of m discrete random variables, without computing the intermediate probabilities.

The Monte-Carlo method is very easy to program, and useful for checking results of other methods. However, it cannot deliver high accuracy. It can be made ‘smarter’ by methods such

as use of control variates, antithetic variation, or importance sampling, but one really needs to resort to Monte-Carlo simulation only for multidimensional integrals. For univariate integrals evaluation by conventional quadrature methods is quicker and more accurate.

Convolution can be done directly, or via taking the Laplace or Fourier transform of the probability density function (pdf) of the survival distribution and inverting the result. The drawback of directly doing convolutions is that the time goes as N^2 , where N is the number of points into which the probability is discretized. However, using Richardson extrapolation, N does not need to be very large, and so the advantage of transform methods largely disappears. The other advantage of transforms, that one can go straight to computation of the m th probability, is removed by the availability of the De Pril (1985) method. It is perhaps also worth noting that a quick look at transform methods throws up difficulties. For example, the non-periodicity of the survival pdf gives an error in the computed convolution. We have therefore used the direct method, for which the size of errors is most easily considered; transform methods undoubtedly have potential but are not explored further here.

This paper focuses on the use of the discretized convolution method. To increase accuracy, Richardson extrapolation is used. The use of the trapezoidal rule, together with Richardson extrapolation, is the basis of the well-known Romberg method of integration. Our approach is broadly similar. The methodology described here could be applied (at least in outline) to any survival distribution, and hence is more general. The first part of our methodology, the discretized convolution, can indeed be applied to any distribution. The details of the second (extrapolation) step depend on the order of the error, and so will be specific to a distribution, or to a class of distributions.

3 Computation of probabilities by convolution

Before discussing the convolution method and how it can be used to compute count probabilities, we recall the general framework used to build up the connection between the count model and inter-arrival timing process. Let $X_k, k \in \mathbf{N}$ be a sequence of *waiting times* between the $(k-1)$ th and the k th event. The arrival time of the m th event is :

$$a_m = \sum_{k=1}^m X_k, \quad m = 1, 2, \dots$$

Denote by N_t the total number of events in $[0, t)$. If t is fixed, $N_t = N(t)$ is the count variable we wish to model. It follows that:

$$N_t < m \iff a_m > t$$

Thus, if F_m is the distribution function of a_m , we have

$$P(N_t < m) = P(a_m > t) = 1 - F_m(t),$$

Furthermore,

$$\begin{aligned} P(N_t = m) &= P(N_t < m + 1) - P(N_t < m) \\ &= F_m(t) - F_{m+1}(t) \\ &= P_m(t) \end{aligned} \tag{1}$$

Equation (1) is the fundamental relationship between the count variable and the timing process. If the X_k are iid with common density $f(x)$, the process is called a *renewal process* (See Feller

(1970) for a formal definition). In this case, Equation (1) can be extended to obtain the following recursive relationship:

$$\begin{aligned} P_{m+1}(t) &= \int_0^t F_m(t-u) dF(u) - \int_0^t F_{m+1}(t-u) dF(u) \\ &= \int_0^t P_m(t-u) dF(u), \end{aligned} \quad (2)$$

where we have that $P_0(u) = S(u) = 1 - F(u)$, sometimes denoted the survival function. Equation (2) can be understood intuitively: the probability of exactly $m + 1$ events occurring by time t is the probability that the first event occurs at time $0 \leq u < t$, and that exactly m events occur in the remaining time interval, integrated over all times u . Evaluating this integral, $P_1(t) \cdots P_m(t)$ can be generated in turn.

This is an attractive method of generating the required probabilities, because the integrand is positive, so there are no subtractions to increase rounding error. To compute the integral, we use a method similar to the extended or composite midpoint rule (e.g. Press et al. (2007, section 4.1.4)). We have:

$$\int_0^{Nh} f(x) dx = h \sum_{j=1}^N f\{(j-1/2)h\} + O(h^2),$$

where there are N steps with stepsize h , and $Nh = t$. This is an open rule, i.e. it does not require evaluating f at the limits of the integral. Thus

$$\int_{(j-1)h}^{jh} g(u) dF(u) = \int_{(j-1)h}^{jh} g(u)f(u) du \simeq g\{(j-1/2)h\}(F\{jh\} - F\{(j-1)h\}),$$

where $g(u) = P_m(t-u)$ for some m , and f is the pdf of the survival distribution. We make the choice of doing the integral of the pdf $f(u)$ analytically, so that

$$f((j-1/2)h) \simeq (F\{jh\} - F\{(j-1)h\})/h, \quad (3)$$

because this is simple for the Weibull distribution (and eventually other distributions) and increases accuracy to $O(h^{1+\beta})$.

The basic procedure is implemented in `getAllProbsUtil.cpp()` function in the `Counter` package (Baker et al., 2016). It generates probabilities $P_0 \dots P_m$. On exit, the P array (local) contains the probabilities $P_0 \cdots P_m$. This code sets up q (local) to contain P_0 at the midpoints $h/2 \cdots (N-1/2)h$, sets up the $F\{jh\} - F\{(j-1)h\}$ array, and carries out the convolutions. The array $q[\]$ initially contains P_0 , and this is overwritten to contain P_1 etc.

A crucial step is the shifting of the probabilities $q[k]$ left by $h/2$. This is necessary because g must be used at the midpoint of each step, and the integral computes g at the end of the step. With this correction, the result is $O(h^2)$ when $\beta \geq 1$, and $O(h^{\beta+1})$ for $\beta < 1$, as shown in appendix B. The algorithm uses $2N$ evaluations of the (Weibull) survival function (which is expensive) and then does $(m-1)N(N+3)/2 + N$ multiplications. Clearly, computing time increases as N^2 for large N .

4 Computing one probability: adaptation of De Pril's method

The method presented above computes all probabilities up to the m th, which is slow if we need only the m th probability. It can be improved so that computing time is $O(\ln(m)N^2)$ instead

of $O(mN^2)$, using the addition chain method. This is essentially an adaptation of a method that is used by compilers for fast computation of integer powers of a variable with the minimum number of multiplications. The details are in Appendix A. This method, which we also call the ‘naïve method’ is useful for timing comparisons, but our main interest is in the De Pril method, which can compute the m th probability in $O(N^2)$ operations.

De Pril (1985) gave a method for computing the m -fold convolution of a discrete distribution. He found that the idea dated back a long way, being first used in other applications than probability before 1956. We refer the reader to De Pril’s paper for two derivations of this amazing algorithm and its history, and simply present it here: let q_i be the value of probability density function of the survival distribution evaluated at points $t_i \geq 0$ where $q_0 > 0$. Then the probability of m events is $f_N^{(m)}$, the m -fold convolution of q , given by

$$f_0^{(m)} = q_0^m,$$

and for $N > 0$ by the recursion

$$f_N^{(m)} = q_0^{-1} \sum_{j=1}^N \left(\frac{(m+1)j}{N} - 1 \right) f_{N-j}^{(m)} q_j. \quad (4)$$

This algorithm when applied to our case requires three arrays: one to hold the survival function, one for the probability mass q , and one work array to hold f .

To apply this method to continuous distributions like the Weibull, we first discretized the distribution, so that $q_j = F((j+1)h) - F(jh)$. The probability mass $f_0^{(m)}$ has contributions from the m random variables all taking the value zero, up to them all taking the value $h - \epsilon$. We should therefore estimate the mean as $mh/2$ rather than zero, so an approximation to the continuous case is that all probability masses such as the N th should be taken as pertaining to time $(N + m/2)h$. To apply this continuity correction, we do not need to copy the $f_N^{(m)}$ into different array locations, but simply to reduce the time interval in the survival function in (8). Finally, for even m , the latest probability mass occurs exactly at time t , and so we take only half of this probability mass. With these two crucial modifications, the method yields the same accuracy as the earlier methods, and Richardson extrapolation can be applied as before. The results are very similar to the addition-chain method, but are usually slightly more accurate, and computation is of course faster. An unexpected additional gain is that for even m , the survival function is not required at half-integer values of h , so saving time on these computations. It had been feared that the presence of the minus sign in the recursion (4) would degrade accuracy, but running the program in quadruple precision gave identical results, so that is not a problem.

Sometimes data are censored, and we only know that at least m events have occurred. This probability $P_{\geq m}$ is then needed for likelihood-based inference. For the direct method (Section 3), one would compute $P_{\geq m} = 1 - \sum_{i=0}^{m-1} P_i(t)$, but for this method, which delivers $f_m(u)$, we compute $P_{\geq m} = \int_0^t f_m(u) du$; the routine supplied in the R package `Countr` returns this. This is an advantage of this and the addition chain method, because small probabilities obtained by differencing are subject to large errors.

The next section describes how Richardson extrapolation can be used to improve the accuracy, without necessitating a large value of N and consequent slow computation.

5 Improvement by Richardson extrapolation

In Romberg integration, the trapezoidal rule is used to generate approximations of error $O(h^2)$, and Richardson extrapolation is used to progressively remove errors of order h^2 , h^4 etc. Clearly,

if an estimate $S_1 = S + \gamma h^\delta$ and $S_2 = S + \gamma(h/2)^\delta$, where S_1 and S_2 are the approximations with N and $2N$ steps respectively and S is the true value, we can remove the error and estimate S as

$$S_3 = (2^\delta S_2 - S_1)/(2^\delta - 1). \quad (5)$$

Subsequently, higher-order errors can be removed in the same way until the required accuracy is attained. Romberg integration can also be done with the extended-midpoint rule (e.g. Press et al. (2007)).

The situation for convolutions is less straightforward, but a satisfactory solution can be found, and the details are given in Appendix B. We now study the proportional errors of probabilities, because these are what determine the error in the in the log-likelihood. Figure 1 shows absolute proportional errors $\delta p/p$ for the first 15 probabilities with $\beta = 1.1$, for the naïve computation, after applying a Richardson extrapolation for error $h^{1+\beta}$, and after applying the second transformation to remove error $O(h^2)$. It can be seen that the errors reduce substantially. Figure 2 shows the estimated power of h of the error, derived by applying (13), with $\beta = 1.2$. It can be seen that this is initially around 2 (because $1 + \beta > 2$), and increases to 2.2, then to 3-4 after the second extrapolation. Figure 3 shows the 3 errors for $\beta = 0.6$.

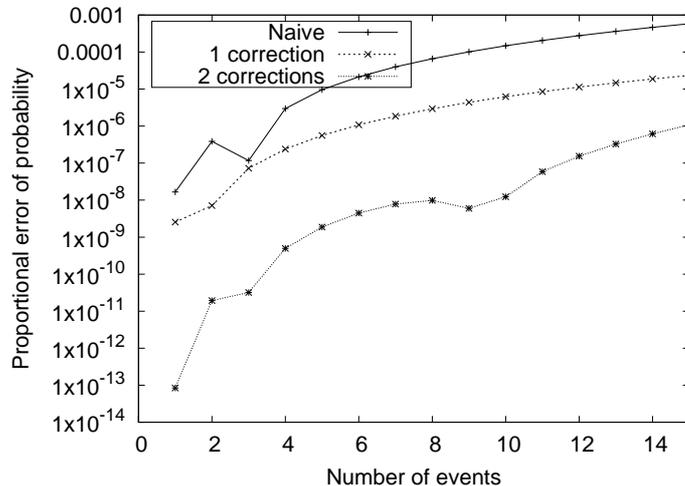


Figure 1: Proportional errors in probabilities for the naïve computation and the two Richardson corrections. Here $\alpha = 1, t = 1, \beta = 1.1$.

Here again the extrapolations progressively reduce error. Figure 4 shows the estimated powers of h for the errors, where now the curves get higher after each extrapolation. Here the initial power is about 1.6, because $1 + \beta < 2$. It then increases to 2, and after applying the second extrapolation, to around 2.6. Finally, Figure 5 shows that the extrapolation works even for a low $\beta = 0.3$.

6 Generalisations

The methodology applies with no change (except the function that provides the survival function) to some generalisations of the Weibull distribution. Thus making the scale α^β a gamma random variate leads to the Burr type XII distribution with survival function

$$S(t) = \frac{1}{(1 + (\alpha t)^\beta)^\nu}, \quad (6)$$

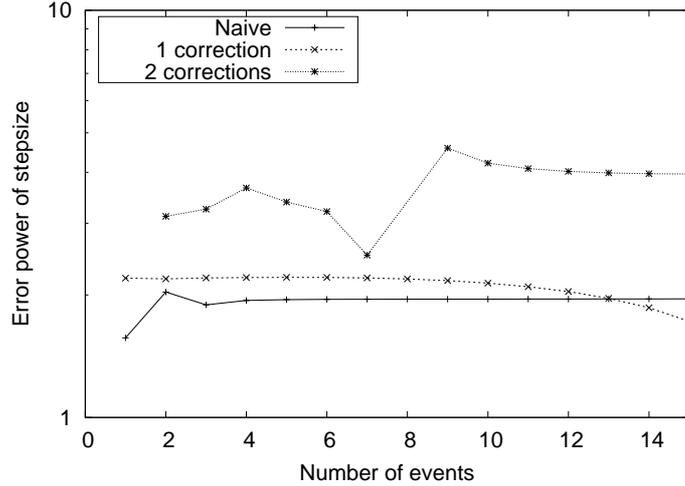


Figure 2: Powers of stepsize h for error in probabilities for the naïve computation and the two Richardson corrections. Here $\alpha = 1, t = 1, \beta = 1.2$.

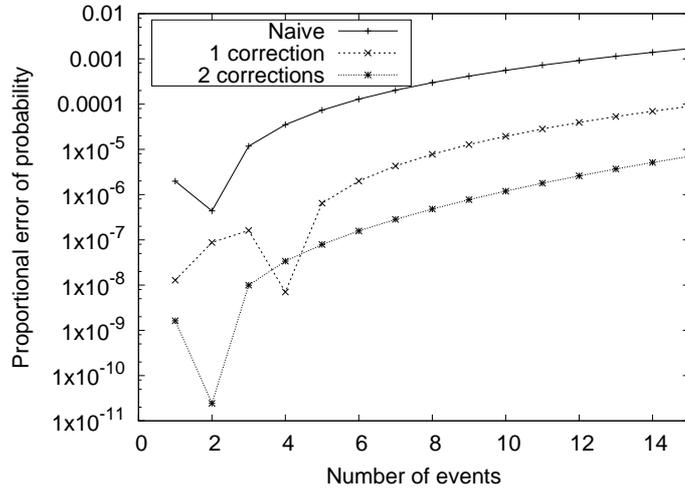


Figure 3: Proportional errors in probabilities for the naïve computation and the two Richardson corrections. Here $\alpha = 1, t = 1, \beta = 0.6$.

where $\nu > 0$. Here α is the scale parameter and β and ν are the shape parameters. When $\beta = 1$ reduces to the Lomax distribution (a shifted Pareto distribution). When $\nu = 1$ this is the log-logistic distribution, and as $\nu \rightarrow \infty$ we regain the Weibull distribution. This distribution addresses the problem of heterogeneity of the hazard function, and is called the heterogeneous Weibull distribution by McShane et al. (2008, Section 3.1, page 374).

The algorithm described can also cope with many of the Weibull-based distributions described in Lai (2014). It also copes with the gamma distribution, where a function for the gamma survival function is needed. Here of course, an analytic solution is available. Another interesting distribution that could be used with the convolution method is the generalised gamma first introduced by Stacy (1962). This distribution includes the Weibull, gamma and log-normal as special cases. Prentice (1974) proposed an alternative parametrisation which is preferred for computation. In the Prentice (1974) parametrisation, the distribution has three parameters (μ, σ, q) , and its survival function is given by:

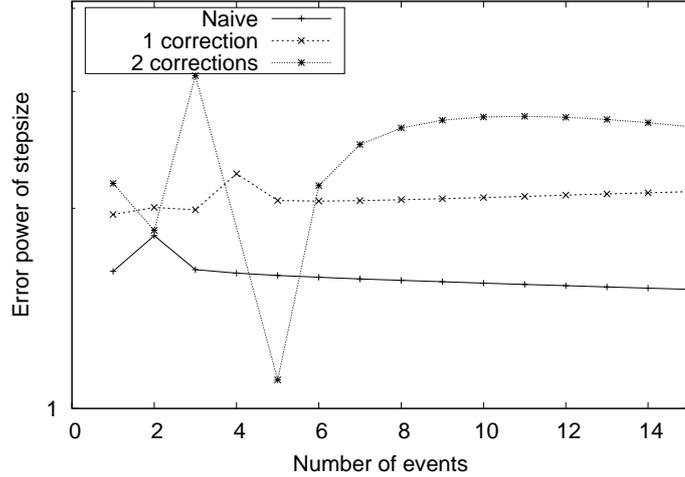


Figure 4: Powers of stepsize h for error in probabilities for the naïve computation and the two Richardson corrections. Here $\alpha = 1, t = 1, \beta = 0.6$.

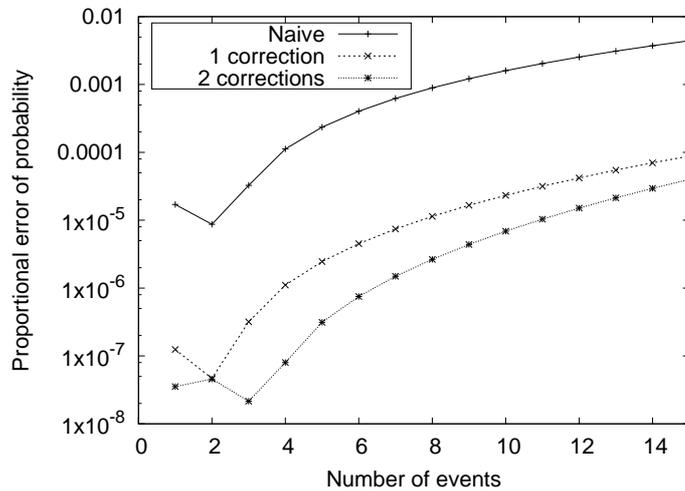


Figure 5: Proportional errors in probabilities for the naïve computation and the two Richardson corrections. Here $\alpha = 1, t = 1, \beta = 0.3$.

$$S(t) = \begin{cases} 1 - I(\gamma, u) & \text{if } q > 0 \\ 1 - \Phi(z) & \text{if } q = 0 \end{cases}$$

where $I(\gamma, u) = \int_0^u x^{\gamma-1} \exp(-x) / \Gamma(\gamma)$ is the regularised incomplete gamma function (the gamma distribution function with shape γ and scale 1), Φ is the standard normal distribution function, $u = \gamma \exp(|q|z)$, $z = (\log(t) - \mu) / \sigma$, and $\gamma = 1/q^2$.

More generally, the convolution step can be applied to any survival distribution. The Richardson improvement of Section 5 requires one to study the first step error to derive a relevant extrapolation. Nevertheless, this extrapolation can be skipped if one is willing to opt for a 'finer' convolution (and hence inevitably longer computation times).

As mentioned in the introduction, the method can also be applied to a modified or delayed renewal process, where the time to the first event follows a different distribution, with pdf $f_0(x)$. This is useful for modelling distributions where the percentage of zero events is abnormal, and

one uses zero-inflated and hurdle models. When both distributions are exponential, we obtain the ‘burnt fingers’ distribution of Greenwood and Yule (Johnson et al., 2005). For the general case, it is straightforward to tweak the code for finding single probabilities. The algorithm is:

1. if m is 0, return the survival function derived from f_0 ;
2. if m is 1, convolve f_0 with P_0 using (8);
3. for higher m , find $f_{m-1}(u)$ using the previous code, then convolve this with f_0 and finally apply (8).

Note that the convolution method can be readily extended to allow modified renewal processes, whereas series-expansion methods cannot.

The R code available at the time of writing (Baker et al., 2016) allows Weibull, gamma, generalised gamma and Burr distributions to be used, plus user-defined distributions. It also includes an ‘experimental’ version that allows modified (delayed) renewal processes, and of course one can regress on covariates.

7 Estimation and testing

7.1 Data

To illustrate the different algorithms described earlier as well as methods previously suggested in the literature, we use a data set for completed fertility. Completed fertility refers to the total number of children born to a woman who has completed childbearing. The data set considered is the same as the one analysed by Winkelmann (1995) and McShane et al. (2008) and consists of a sample of $n = 1243$ women over 44 in 1985. A more detailed description can be found in Winkelmann (1995). We selected this data set for two main reasons. First, the previous references inspired this research and will be used as a benchmark for our new approach. It was essential to be able to produce results in agreement with previous conclusions and hence re-analysing the same data made sense. Second, this data set is slightly underdispersed (sample variance 2.3 versus the sample mean 2.4) and hence allows us to highlight the flexibility of the new approach compared to the simple Poisson-negative binomial methods. A more precise description of the data is presented in Figure 6 and Table 1. The range of the data is quite

Children	0	1	2	3	4	5	6	7	8	9	10	11
Frequency	76	239	483	228	118	44	30	10	8	3	3	1
Percent	6.1	19.2	38.9	18.3	9.5	3.5	2.4	0.8	0.6	0.2	0.2	0.1
Poisson fitted	9.2	21.9	26.2	20.8	12.4	5.9	2.3	0.8	0.2	0.1	0.0	0.0

Table 1: Number of children in the German fertility dataset.

narrow, with more than 95% of the observations in the range 0-5 and the highest count being 11 in both cases. The data set shows a pronounced mode at 2 children, a number seen as ideal by many families.

7.2 Comparing performance of different methods

In this section, we compare the performance of the various methods using the German fertility data and a univariate Weibull count model, intercept-only. We computed the model

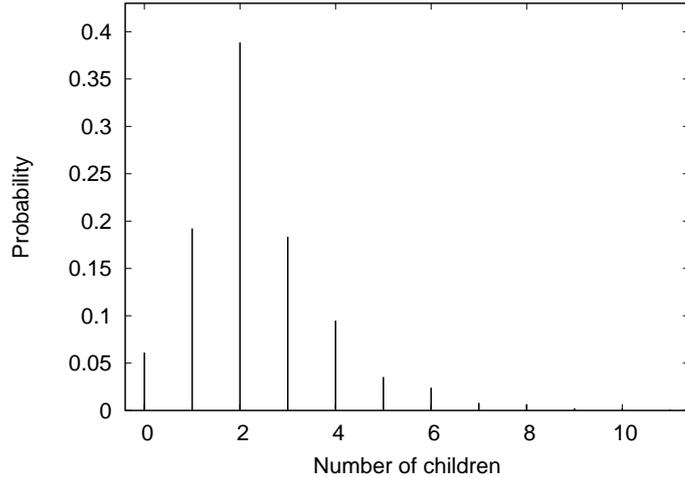


Figure 6: Frequency distribution of the number of children born to a woman who has completed child-bearing in Germany ($n = 1243$)

log-likelihood by a very long convolution (20,000 steps as before), and proportional errors computed taking this as correct after Richardson extrapolation. For each method, we achieved the minimum number of computations to reach an precision (error) of at least 10^{-8} . The computation was repeated 1000 times and execution times measured using routines from the R package `rbenchmark`. The experience was conducted on a 2.6 GHz intel Core i7 computer and results are collected in Table 2. Table 2 suggests that the series expansion methods are almost twice as fast

method	relative	elapsed (in seconds)
Series expansion	1.09	21.74
Accelerated series expansion	1.00	19.86
Direct convolution	8.76	173.98
Direct conv. with extrapolation	1.82	36.09
Naïve convolution	7.57	150.30
Naïve convolution with extrapolation	1.93	38.29
De Pril convolution	5.73	113.72
De Pril conv. with extrapolation	1.93	38.40

Table 2: Performance measure of the different computation methods available for the Weibull count (German fertility data). The series expansion is as described in McShane et al. (2008) programmed in vectorised form, and the series expansion method was also accelerated by the Euler and van-Wijngaarden transformations. The direct convolution algorithm is as described in Section 3 and the naïve and De Pril methods are described in Section 4. Convolution methods are tested with and without Richardson extrapolation.

as the convolution methods and more than 5 times faster than convolutions without Richardson correction. Surprisingly, the De Pril method (with correction) performed slightly worse than the direct approach and similarly to the naïve approach. The reason is that this method needed slightly more steps to reach the desired accuracy.² However, the De Pril method has been found to be slightly more accurate than all other methods including series expansion for large

²When extrapolation was applied, the De Pril approach needed 36 steps when the other methods required only 24. If no extrapolation was applied, all methods used 132 steps. In this case, the De Pril method was found to be faster (32 % faster than the naïve approach and 53% faster compared to the direct approach).

counts (larger than 10). Given that the testing data set we use has a narrow range of (low) counts, the added value of the method was not seen.

In order to highlight the improvement introduced by the De Pril approach, we slightly modified the German fertility data set by 'artificially' adding some large counts. The results were more in accordance with what we expect. The De Pril approach was three times faster than the naïve approach and more than four times faster than the direct approach.

Nevertheless, it was still slower than the series approach, the accelerated approach still being slightly faster than the vectorial approach. It is not surprising that a 'tailored' method such as the series expansion outperforms a generic method such the convolution method described in this paper. Nevertheless, computation times are comparable and the convolution approach has the advantage of being more much flexible as it allows any survival distribution, and can be adapted for modified renewal processes. One pays the price for this flexibility in slightly increased computation time.

7.3 Univariate models

The first family of models considered is an intercept-only (no individual covariates) version of several renewal processes with different distributions for the inter-arrival times. Table 3 presents values of model-choice criteria for the various models. First, we note from Table 1 that

Variable	Poisson		Weibull		Gamma		gen. Gamma	
	Coef	SE	Coef	SE	Coef	SE	Coef	SE
scale	2.38	0.02	2.64	0.03	0.35	0.06	0.64	0.09
shape			1.12	0.03	1.16	0.06	1.93	0.07
shape2							2.29	0.38
log-likelihood	-2186.78		-2180.36		-2182.53		-2167.18	
AIC	4375.55		4364.71		4369.06		4340.37	
BIC	4380.68		4374.97		4379.31		4355.74	
χ^2	126.16		111.79		115.53		87.29	
df	6		5		5		4	
p-value	8.2×10^{-25}		1.7×10^{-22}		2.7×10^{-23}		4.9×10^{-18}	

Table 3: German fertility data: Model choice criteria for the various models.

the Poisson model over-fits the zero count and under-fits the peak at 2.

The log-likelihood values reported in Table 3 show best fit by the generalised gamma, which is clearly preferred according to AIC and BIC. Significant improvements are confirmed by likelihood ratio tests over Poisson (-2LR = 39.2) and gamma (-2LR = 30.7) at any conventional level of significance. The result is similar for the Weibull process model (-2LR = 26.3) compared with Poisson. It is also worth mentioning here that the chi-squared goodness of fit test rejects the null hypothesis (that the empirical data comes from the claimed distribution) at any conventional level of significance for the four models suggesting that these simple models (with no covariates) fail to capture the data generating process. A closer investigation of the table of observed and expected frequencies tells us that all models under-estimate the peak at 2 children. Nevertheless, as mentioned earlier, it made sense to analyse this dataset in order to be able to validate and compare the results to what have been suggested in the literature.

One can also note that the log-likelihood value presented in Table 3 computed with the convolution method is identical to the one in Winkelmann (1995, Table 1) and McShane et al.

(2008, Table 1), thus validating the accuracy of our computation. The standard errors are obtained from numerical computation of the Hessian matrix at the fitted value of the parameters.

7.4 Regression models using renewal processes

We turn now to the analysis of the model with individual covariates. The explanatory variables available are the woman's general education (given by the number of years of school), nationality (a dummy, either German or not), university access (yes or no), rural or urban dwelling, religion (a categorical variable with levels Catholic, Protestant, and Muslim, with others being the reference group), year of birth and the year of marriage). Results are collected in Table 4.

Variable	Poisson		Weibull		Gamma		gen. Gamma	
	Coef	SE	Coef	SE	Coef	SE	Coef	SE
scale	3.150	0.302	4.044	0.315	0.211	0.252	-1.087	0.252
German	-0.200	0.072	-0.223	0.072	-0.190	0.059	-0.190	0.059
Years of schooling	0.034	0.032	0.039	0.033	0.032	0.027	0.032	0.026
Vocational training	-0.153	0.044	-0.173	0.044	-0.144	0.036	-0.144	0.036
University	-0.155	0.159	-0.181	0.160	-0.146	0.130	-0.146	0.129
Catholic	0.218	0.071	0.242	0.070	0.206	0.058	0.206	0.058
Protestant	0.113	0.076	0.123	0.076	0.107	0.062	0.107	0.062
Muslim	0.548	0.085	0.639	0.087	0.523	0.070	0.523	0.069
Rural	0.059	0.038	0.068	0.038	0.055	0.031	0.055	0.031
Year of birth	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002
Age at marriage	-0.030	0.007	-0.034	0.006	-0.029	0.005	-0.029	0.005
shape			1.236	0.034	1.439	0.071	2.211	0.031
shape2							1.121	0.169
log-likelihood	-2101.8		-2077.0		-2078.2		-2076.7	
AIC	4225.6		4178.0		4180.5		4179.6	
BIC	4281.980		4240		4242		4246.2	

Table 4: Regression model results for German fertility data

One can also note here that the values of the log-likelihood are in accordance with the previously mentioned literature. The value of the coefficients are not exactly identical but are within the same confidence region. The generalised gamma distribution still provides the best likelihood, but with a higher AIC, so the Weibull model would be (slightly) preferred. One may conclude that the introduction of individual covariates improves the data description rather more than a more flexible hazard model (as introduced by the generalised gamma).

We would also like to mention here that we tried to reproduce the heterogeneous-Weibull (6) described in McShane et al. (2008, Table 2). We found similar results using the series expansion methods when we used 50 terms to expand the series but different results were obtained (with smaller log-likelihood values) when more terms were used. We think that the series expansion may need more than 50 terms to converge in the heterogeneous-Weibull case and hence the conclusion of McShane et al. (2008, Table 2) should be interpreted with care. Although the series expansion method works smoothly in the simple Weibull case (around 20 terms are usually enough to ensure convergence), for more complicated distributions more terms may be needed. On the other hand, due to the use of the gamma function, there is a limitation on the maximum number of terms that could be numerically computed. The convolution method described in

this paper does not suffer from this limitation and hence can be seen as more robust as well as being more flexible.

8 Conclusions

A fast and flexible method is presented for computing the probabilities of discrete distributions derived from renewal and modified renewal processes. This should pave the way for more widespread use of this type of model in management science, and wherever count data needs to be modelled. Where the data arise from a stochastic process, such as football goals or hospital visits, the renewal model can have a strong basis in fact. It can however be applied to any count data, such as number of bacteria seen under a microscope, using the renewal framework purely as a mathematical device.

This class of models is we think tractable enough for use by practitioners. Computation of probabilities of numbers of events is essential for likelihood-based inference, and we have focused on this. Tests are often also needed, e.g. for under or overdispersion. If fitting a Weibull model, as the shape parameter β determines under or overdispersion, we simply need to test that $\beta = 1$. Computing the log-likelihood with β ‘floating’ and fixed to unity, twice the increase in log-likelihood on floating β is asymptotically distributed as $X^2[1]$, a chi-squared with one degree of freedom. For small samples, one can find the distribution of this statistic under H_0 more accurately by using the parametric bootstrap. We would thus claim that these distributions are tractable where it matters: computation of moments for example is difficult, but is not needed for inference. We would suggest that a Monte-Carlo simulation would be easy to program and fast enough for the modest accuracy required.

We have chosen to implement what seemed the most direct method of computing probabilities, after ruling out Monte-Carlo integration on the grounds that regular quadrature methods are better for one-dimensional integrals. The method given can be applied as it stands to a variety of generalisations of the Weibull distribution, and can be applied in outline to other survival distributions, such as the lognormal. The `Countr` R package that allows the Weibull, gamma and few other distributions is available from the CRAN archive. Details are given in the electronic companion to this paper.

This is an area where much further work could be done. There is a bewildering variety of possible approaches to computing the probabilities, and the successful use of Laplace or Fourier transforms is surely a possibility. However, the disadvantage of direct methods, that computation time goes as N^2 for N steps, is much ameliorated by using Richardson extrapolation, so that N can be small. The Weibull distribution has a virtue for the direct convolution approach adopted here, in that the distribution function is easy to compute. However, it has the disadvantage for transform methods that the transform $M(s)$ cannot be found analytically, but must be evaluated numerically for each value of s , where the transform is $M(s) = \int_0^\infty \exp(-st) dF(t)$. The present method, which already gives adequate performance, would be a useful benchmark for developers of more advanced methods to compare with. We conjecture that great improvements in speed are not possible, but hope to be proved wrong here.

Perhaps of greater interest than further speeding up computation is gaining experience with the expanded range of renewal-type models that can now be feasibly used. This includes modified renewal processes, where the time to the first event follows a different distribution to later events. This for example yields a natural class of hurdle models, where the first event is slow to happen, but later events follow more quickly. Conversely, this class includes distributions where there are very few occurrences of zero events. It will be interesting to see how useful practitioners find these new models.

Appendices

A Addition chain method for computing probabilities

The aim is to find the m th convolution of the pdf in as few convolutions as possible. The method works by convolving the pdf f_i of i events occurring, using

$$f_{i+j}(t) = \int_0^t f_i(u)f_j(t-u) du \quad (7)$$

and finally

$$P_m(t) = \int_0^t f_m(u)P_0(t-u) du \quad (8)$$

We need two work arrays: one (pdfn) for the n -th convolution of the pdf, initially set to pdfn[j] = $(F((j-1)h) - F(jh))/h$, an approximation to f_1 , and repeatedly overwritten, the other, q, to hold what will become the final pdf as it is being updated. Two routines are needed to do the convolving: one for convolving the m th order pdf with itself, the other for convolving two pdfs of different order. The symmetry of the integrand means that only half the multiplications are required when doubling the order of the pdf.

To organise the calculation, we first find the binary decomposition of m . For example, with $m = 21$, we would have $21 = 1 + 2^2 + 2^4$. This can be translated into code as:

- set q to f_1 ,
- apply (7) to obtain f_2 ,
- then apply (7) to f_2 to obtain f_4 ,
- convolve q with f_4 to obtain q as f_5 ,
- apply (7) again to f_4 to obtain f_8 and f_{16} ,
- then convolve q with f_{16} to obtain f_{21} .
- Finally, apply (8) to obtain $P_{21}(t)$.

This has required 6 convolutions and one evaluation, instead of 20 convolutions and one evaluation.

The best case occurs when $m = 2^k$, when k convolutions are needed, all order doublings. The worst case occurs when $m = 2^k - 1$, when $m = \sum_{j=0}^{k-1} 2^j$. Here all the pdfs $f_1, f_2 \cdots f_{k-1}$ must be convolved, giving a total of $2(k-1)$ convolutions. This is still $O(\ln_2(m))$.

B Richardson extrapolation

This technique can substantially reduce the required number of steps N . To derive a useful extrapolation we start by considering the error of the extended midpoint approximation. The error E_j is given by

$$E_j = \int_{(j-1)h}^{jh} g(u) dF(u) - g\{(j-1/2)h\}(F\{jh\} - F\{(j-1)h\}). \quad (9)$$

Expanding the integrand in a Taylor series $g(u) \simeq g(u_0) + g' \cdot (u - u_0) + (1/2)g'' \cdot (u - u_0)^2$, where $u_0 = (j - 1/2)h$ and the derivatives are taken at u_0 . Writing similarly the pdf $f(u) = f(u_0) + f' \cdot (u - u_0) + (1/2)f'' \cdot (u - u_0)^2$, we have for the step error to the lowest order in h ,

$$E_j = h^3 \{f'g'/12 + fg''/24\}. \quad (10)$$

The proof follows:

We have that

$$g(u) \simeq g(u_0) + g' \cdot (u - u_0) + (1/2)g'' \cdot (u - u_0)^2,$$

so that

$$\begin{aligned} E_j &= \int_{(j-1)h}^{jh} g(u) dF(u) - g\{(j-1/2)h\}(F\{jh\} - F\{(j-1)h\}) \\ &= \int_{(j-1)h}^{jh} (g' \cdot (u - u_0) + (1/2)g'' \cdot (u - u_0)^2) f(u) du. \end{aligned} \quad (11a)$$

Expanding

$$f(u) \simeq f(u_0) + f' \cdot (u - u_0) + (1/2)f'' \cdot (u - u_0)^2$$

and substituting in (11a) we obtain

$$E_j \simeq \int_{(j-1)h}^{jh} \{g' \cdot (u - u_0) + (1/2)g'' \cdot (u - u_0)^2\} \{f(u_0) + f' \cdot (u - u_0) + (1/2)f'' \cdot (u - u_0)^2\} du.$$

The integrand $I(u)$ is:

$$I(u) \simeq g(u_0)(u - u_0)\{g' \cdot f(u_0)\} \quad (12a)$$

$$+ (u - u_0)^2 \{g' \cdot f' + (1/2)g'' \cdot f(u_0)\} \quad (12b)$$

Then we need to integrate each term in the previous equation between $(j - 1)h$ and jh :

- Integration of Equation (12a) gives 0 by symmetry.
- Integration of Equation (12b) gives $h^3/12 \times \{g' \cdot f' + (1/2)g'' \cdot f(u_0)\}$

Therefore, using the definition of E_j in (9), we get the result in Equation (10).

Since there are $N = t/h$ terms, this gives an error of $O(h^2)$. However, the first step cannot be treated in this way, because u^β has a singularity at $u = 0$, which is therefore at the radius of convergence of the Taylor expansion. We instead consider the error of the first term when $F(u)$ is approximated as $(\alpha u)^\beta$, i.e. at small times u . Then the error E_1 can be found from (9) without expanding out f as

$$E_1 \simeq g'k_1(\beta)(\alpha h)^{\beta+1}/\alpha + g''k_2(\beta)(\alpha h)^{\beta+2}/\alpha^2$$

where $k_1(\beta)$ and $k_2(\beta)$ are some functions of β that could be found exactly. This is $O(h^{\beta+1})$. For $\beta > 1$, the $O(h^2)$ error dominates, but for $\beta < 1$ the error is $O(h^{\beta+1})$. Higher order errors are of type $O(h^{\beta+n})$ and $O(h^{n\beta+1})$ for $n > 1$.

This affects what can be achieved by Richardson extrapolation. Two steps are advocated using (5), so that 3 sets of convolutions are done with series lengths $N, 2N, 4N$. Let a particular probability be A_1, A_2 and A_3 from the convolutions (in order of increasing length). Then the

extrapolation used is: Define $\gamma_1 = \beta + 1$, $\gamma_2 = 2$ (the order does not matter). Compute $B_1 = (2^{\gamma_1} A_2 - A_1)/(2^{\gamma_1} - 1)$, $B_2 = (2^{\gamma_1} A_3 - A_2)/(2^{\gamma_1} - 1)$. Finally, the extrapolated probability is $C_1 = (2^{\gamma_2} B_2 - B_1)/(2^{\gamma_2} - 1)$. We have removed the two errors, leaving higher order errors: $O(h^{\beta+2})$ and $O(h^4)$. When $\beta > 1/2$, two extrapolations leave an error of order $\min(1 + 2\beta, 2 + \beta, 4)$, which is at least $O(h^3)$. When β is small, say 0.1, there are many errors of similar orders, and Richardson extrapolation, although it can improve accuracy, can not remove the low-order error. However, we believe that the procedure recommended will generally be satisfactory, and if not, for low β one would have to increase N .

The code that carries out the extrapolation also computes the minimum number of exponentiations, because some of those for $4N$ can be re-used for $2N$ and N .

For studying the order of error, a very long convolution was used, with 20000 steps, and errors computed taking this as correct (after Richardson extrapolation). The order of error can be studied by carrying out three convolutions with $N, 2N, 4N$, and solving the 3 equations for γ . We then find

$$\gamma = \ln \frac{S_2 - S_1}{S_3 - S_2} / \ln(2), \quad (13)$$

where $S_1 = S + ah^\gamma$ etc. The extrapolated value S_1^e is

$$S_1^e = \frac{S_1 S_3 - S_2^2}{S_1 + S_3 - 2S_2}.$$

This is in fact the ‘Aitken acceleration’ of S_1 , sometimes used to speed up convergence of series, where S_1, S_2, S_3 would be successive partial sums. Press et al. (2007) recommend writing it in the form

$$S_1^e = S_1 - (S_1 - S_2)^2 / (S_1 + S_3 - 2S_2), \quad (14)$$

which reduces rounding error.

Although this extrapolation improves the results when $\beta < 1$, the procedure recommended is sometimes more accurate. However, one could use either. It can be seen from (14) that unlike the recommended procedure, longer convolutions do not have more weight, and that there is the potential for divide overflow and loss of accuracy in computing S .

It is possible in the same way to go further, and remove the next power of error, $\beta + 2$. Equation (13) was applied to the probabilities C_1, C_2, C_3 . This requires 5 initial computations, of $A_1 \cdots A_5$. The power of h remaining was roughly $\beta + 2$, but decreased below this when $\beta < 0.5$. However, application of the Richardson extrapolation will reduce error, even if the power of h used, γ_2 , is not correct, and the true power is γ_1 . It is easy to show that error is reduced if $\gamma_2 \geq \gamma_1$. Hence this third Richardson step will always reduce the error further.

References

- Baker, R., Boshnakov, G., Kharrat, T., and McHale, I. (2016). *Countr: an R package to generate flexible count models*. *Journal of Statistical Software*, forthcoming.
- Boshnakov, G., Kharrat, T., and McHale, I. G. (2017). A bivariate Weibull count model for forecasting association football scores. *International Journal of Forecasting*, 33:458–466.
- Cameron, A. C. and Trivedi, P. K. (2013). *Regression Analysis of Count Data*, volume 53. Cambridge University Press.
- Chaudhry, M. L., Yang, X., and Ong, B. (2013). Computing the distribution function of the number of renewals. *American Journal of Operations Research*, 3(03):380.

- Cox, D. R. (1962). *Renewal Theory*, volume 4. Methuen London.
- De Pril, N. (1985). Recursions for convolutions of arithmetic distributions. *Astin Bulletin*, 15(02):135–139.
- Feller, W. (1970). *An introduction to probability theory and its applications*, volume 2. John Wiley & Sons.
- Gonzalez-Prida, V. (2015). Modelling the repair warranty of an industrial asset using a non-homogeneous poisson process and a general renewal process. *IMA Journal of Management Mathematics*, 26:171–183.
- Greene, W. H. (2011). *Econometric Analysis (7th ed.)*. Prentice Hall.
- Johnson, N. L., Kemp, A. W., and Kotz, S. (2005). *Univariate Discrete Distributions*. John Wiley & Sons.
- Jose, K. and Abraham, B. (2013). A counting process with Gumbel inter-arrival times for modeling climate data. *Journal of Environmental Statistics*, 4(5).
- Jose, K. K. and Abraham, B. (2011). A count model based on Mittag-Leffler inter-arrival times. *Statistica*, 71(4):501–514.
- Lai, C.-D. (2014). *Generalized Weibull Distributions*. Springer.
- Larsen, C., Seiding, C., Teller, C., and Thorstenson, A. (2008). An inventory control project in a major Danish company using compound renewal demand models. *IMA Journal of Management Mathematics*, 19:145–162.
- Lomnicki, Z. (1966). A note on the Weibull renewal process. *Biometrika*, 53(3-4):375–381.
- Marquez, A. C., Iung, B., Macchi, M., and Kobbacy, K. A. H. (2015). Editorial: Special Issue: Maintenance Modelling and Engineering. *IMA Journal of Management Mathematics*, 26:121.
- McShane, B., Adrian, M., Bradlow, E. T., and Fader, P. S. (2008). Count models based on Weibull inter-arrival times. *Journal of Business & Economic Statistics*, 26(3).
- Mullahy, J. (1986). Specification and testing of some modified count data models. *Journal of Econometrics*, 33:341–365.
- Prentice, R. L. (1974). A log gamma model and its maximum likelihood estimation. *Biometrika*, 61(3):539–544.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (2007). *Numerical Recipes: The art of scientific computing (3rd ed.)*. Cambridge University Press.
- R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Stacy, E. W. (1962). A generalization of the gamma distribution. *The Annals of Mathematical Statistics*, pages 1187–1192.
- Winkelmann, R. (1995). Duration dependence and dispersion in count-data models. *Journal of Business & Economic Statistics*, 13(4):467–474.
- Winkelmann, R. (2013). *Econometric Analysis of Count Data*. Springer Science & Business Media.