Repair Same as New or Same as Old: Practical Issues in Choosing the Appropriate Stochastic Process to Model Failure

Dana L. Kelly
Dana.L.Kelly@gmail.com

A fundamental assumption of a renewal process is that repair restores a failed component to the state it was in when it was new. This is called *repair same as new*.[1] Times between failures under this assumption are treated as statistically independent observations from a stationary process. The distribution of the times between failures is called the *renewal distribution*. If the failure rate is constant over either operating time or time in standby, depending on which is being modeled, then the renewal distribution is an exponential distribution. A Weibull or gamma distribution allows for monotonically increasing or decreasing failure rates, depending on whether the shape parameter is > 1 or < 1, respectively (when the shape parameter equals 1, both distributions reduce to the exponential). Another popular renewal distribution is the lognormal distribution. The lognormal distribution does not have a shape parameter, and its failure rate increases quickly, and then decreases monotonically with operating or standby time. It can be useful when early failures dominate, causing an initially increasing failure rate.

With a renewal process, the failure rate does not change with *calendar* time, as pointed out above, only with operating time or time in standby, according to the variable being considered. Likewise, *cumulative* times to failure are not the inputs to a statistical analysis; it is the times *between* failures that are treated as a random sample from the renewal distribution. Furthermore, because a renewal process is stationary, a plot of cumulative number of failures versus cumulative failure time will be approximately a straight line, so a cumulative failure plot is not useful in deciding how the failure rate is changing with time. The figures below illustrate this plot for three cases of simulated failure times: constant failure rate, increasing failure rate, and decreasing failure rate. The times in the constant case are simulated from an exponential distribution with mean time to failure of 350 arbitrary time units. The increasing and decreasing cases were simulated with Weibull distributions with shape parameters of 2 and 0.5, respectively, and a scale parameter of 350 arbitrary time units. Note the linearity displayed in all three cases, illustrating the behavior typical of the cumulative failure plot when the times between failures are generated by a renewal (i.e., repair same as new) process.

[1] Note that other terms are used in the literature, such as good as new. We use same as new, because the new state may have a higher failure rate than the state the component was in when failure occurred.

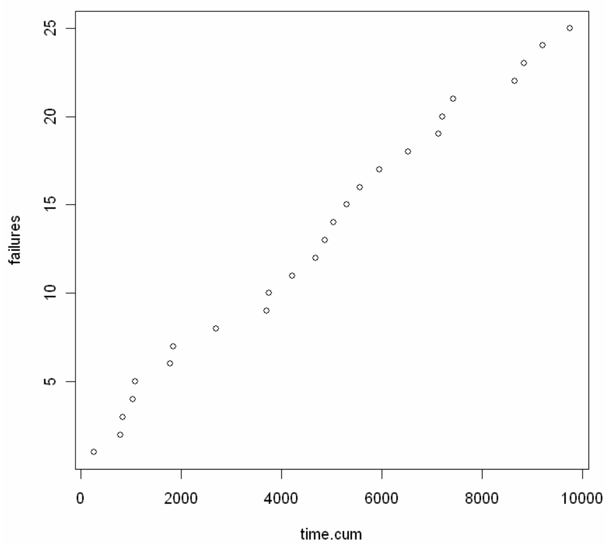

Figure 1 Cumulative failure plot for 25 times between failures from exponential distribution (constant failure rate)

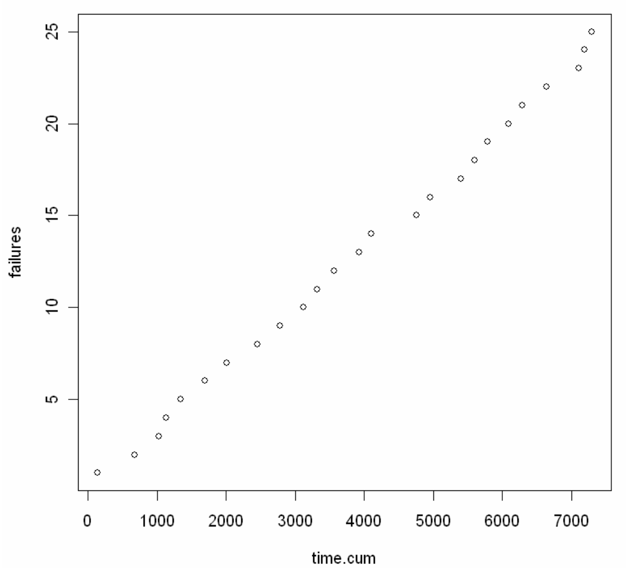

Figure 2 Cumulative failure plot for 25 times between failures for renewal distribution with increasing failure rate

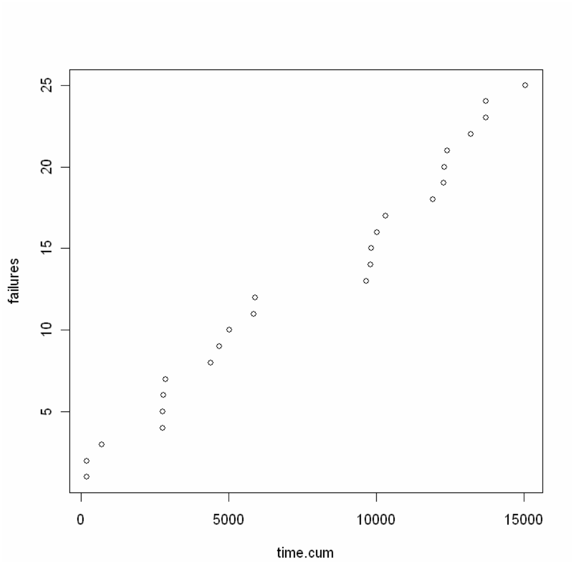

Figure 3 Cumulative failure plot for 25 times between failures for renewal distribution with decreasing failure rate

The plots below show cumulative failures versus cumulative time for 1,000 simulated failure times from two different renewal processes, one in which failure rate is decreasing with increasing operating time, the other where failure rate is increasing with operating time. Note in both cases that the cumulative failure plot produces a straight line, reinforcing the conclusion that this plot cannot detect a time-dependent failure rate under the same-as-new repair assumption.

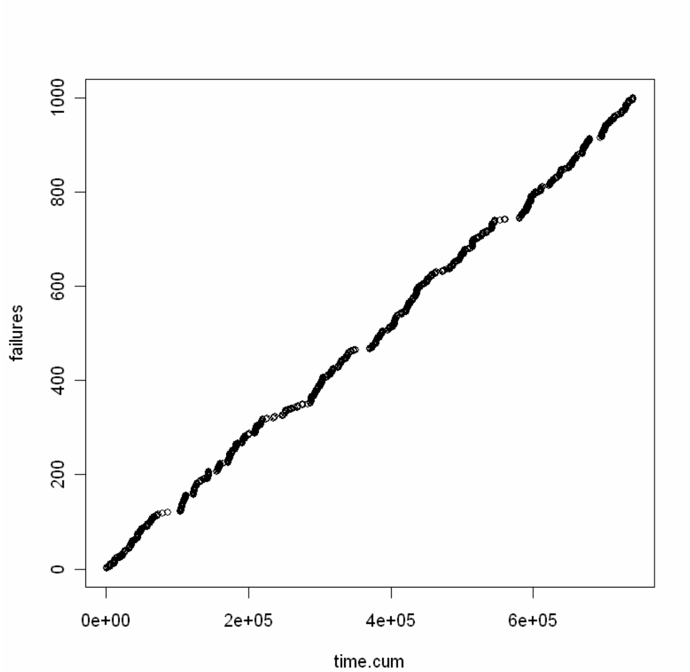

Figure 4 Cumulative failure plot for 1,000 simulated failure times from renewal process with decreasing failure rate

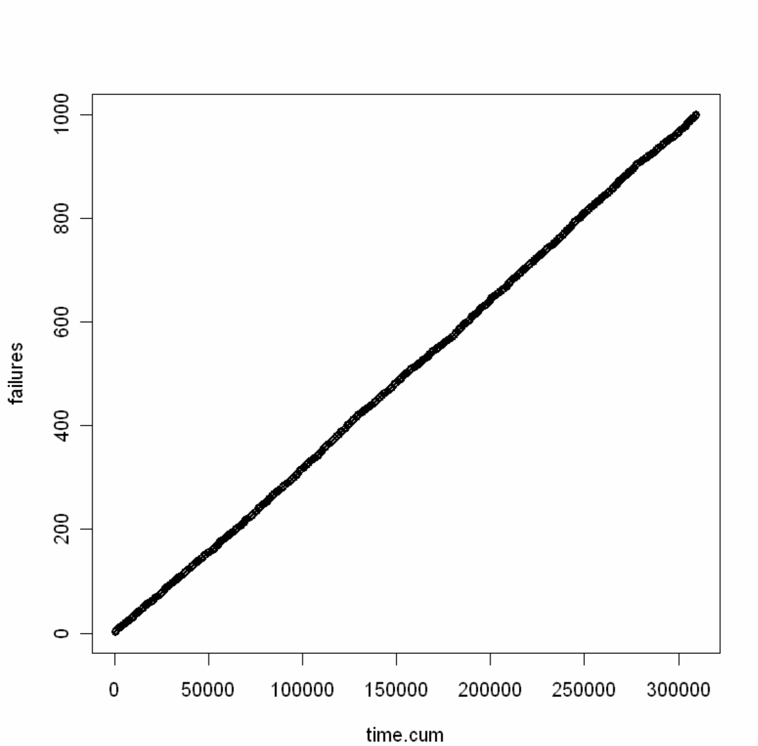

Figure 5 Cumulative failure plot for 1,000 simulated failure times from renewal process with increasing failure rate

A plot that is useful for the renewal process is a *cumulative hazard plot* or one of its close relatives (e.g., Nelson-Aalen plot). If the failure rate is constant in a renewal process, then the times between failures are exponentially distributed. If one plots the *ranked* times between failures on the x-axis, and $1/n_t$ on the y-axis, where n_t is the number of components still operating at time t, the result should be approximately a straight line if the failure rate is constant (i.e., the renewal distribution is exponential). If the slope is increasing (decreasing) with time, this suggests a renewal process whose failure rate is likewise increasing (decreasing) with time. The figures below illustrate this plot for the three cases of simulated failure times described above.

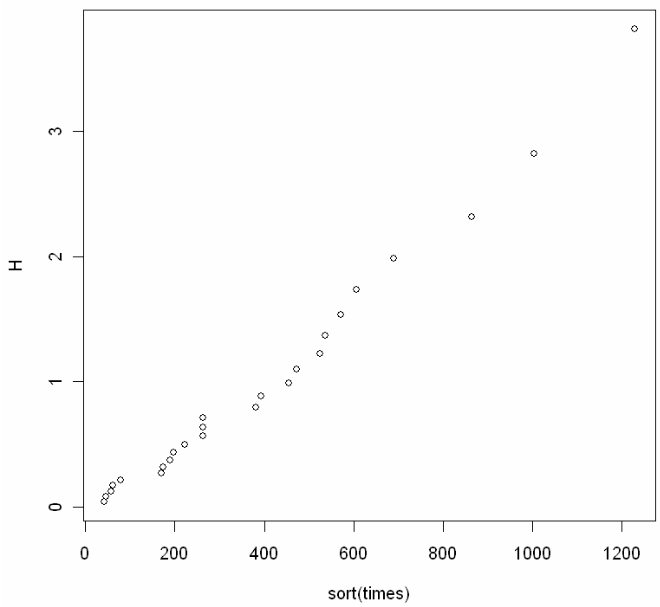

Figure 6 Cumulative hazard plot for 25 times between failures from exponential distribution (constant failure rate)

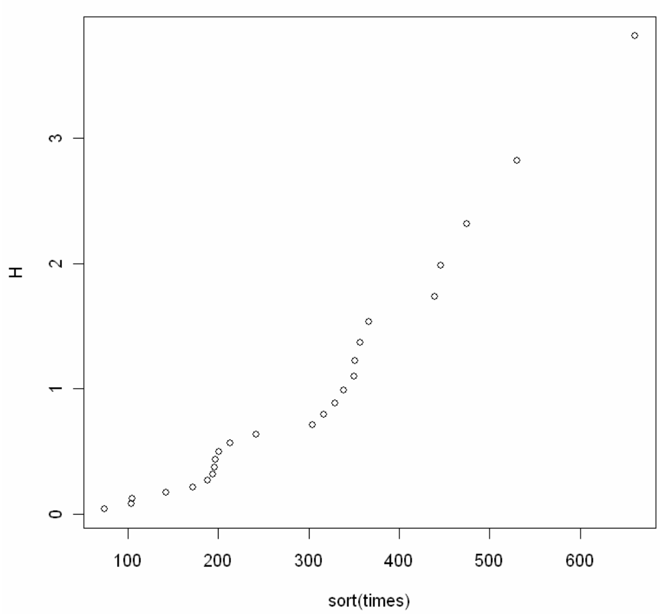

Figure 7 Cumulative hazard plot for 25 times between failures for renewal distribution with increasing failure rate

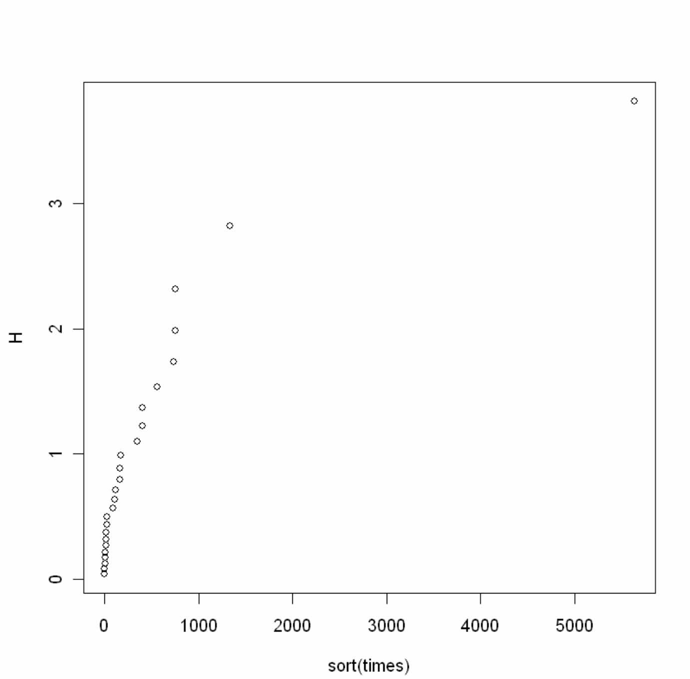

Figure 8 Cumulative hazard plot for 25 times between failures for renewal distribution with decreasing failure rate

As stated above, if repair is same as old and there is an increasing or decreasing trend in the rate of occurrence of failure (ROCOF) over *calendar* time, then the times between failures will *not* be independently and identically distributed, and one cannot simply fit a Weibull, gamma, etc. distribution to the cumulative failure times or the times between failures. There is a non-parametric statistical test that can be helpful in distinguishing between a renewal process (same-as-new repair) and a nonhomogeneous Poisson process (NHPP) representing same-as-old repair. It is very similar to the so-called Laplace test, and it is simple to implement in a spreadsheet or other software. It uses the *cumulative* times to failure. The formula depends on whether the observation period is for a fixed period of time, or terminates at the time of the last failure. If it is for a fixed period of time, τ, and there are n observed failure times, one first computes the Laplace statistic, U, according to the following formula:

$$U = \frac{\sum_{i=1}^{n} t_i / \tau - n/2}{\sqrt{n/12}} \qquad (1)$$

If the observation period is only up until the last observed cumulative failure time, t_n, U is computed by

$$U = \frac{\sum_{i=1}^{n-1} t_i \Big/ t_n - (n-1) \Big/ 2}{\sqrt{\frac{n-1}{12}}} \tag{2}$$

This statistic quickly approaches a standard normal distribution under the null hypothesis of a homogeneous Poisson process. If the process is actually NHPP with increasing ROCOF, too many of the failures will occur after the midpoint of the observation period, and U will be too large. Conversely, if the process is NHPP with decreasing ROCOF, too many failures will occur before the midpoint of the observation period, and U will be too small. At a 5% significance level, we reject the null hypothesis if U is larger than 2 or smaller than -2.

As pointed out above, the null hypothesis for the Laplace test is a homogeneous Poisson process. A slight modification to this test allows for a null hypothesis of a renewal process with any renewal distribution, not just the exponential distribution that is the null hypothesis for the Laplace test. We divide U by the estimated coefficient of variation of the *times between failures* (ratio of sample standard deviation to sample mean). Asymptotically, if the renewal distribution is exponential, this new statistic will equal U. Again, this statistic quickly approaches a standard normal distribution under the null hypothesis of a renewal process.

Applying this test to the simulated data shown in Figures 1-3, we obtain the following values of the modified Laplace statistic. The two-sided p-values are shown in parentheses. As expected, we cannot reject the null hypothesis of a renewal process at any reasonable significance level.

Constant failure rate: -0.78 (0.43)
Increasing failure rate: 0.71 (0.48)
Decreasing failure rate: -0.10 (0.92)

There is also a qualitative check for an increasing or decreasing trend in the ROCOF under the same-as-old repair assumption, meaning that our stochastic process for failure is an NHPP. We plot the cumulative number of failures versus the cumulative time to failure. The slope of this plot is an estimate of the ROCOF, so an increasing slope corresponds to aging and vice versa. The figures below show these plots for simulated data from a power-law NHPP with increasing and decreasing ROCOF.

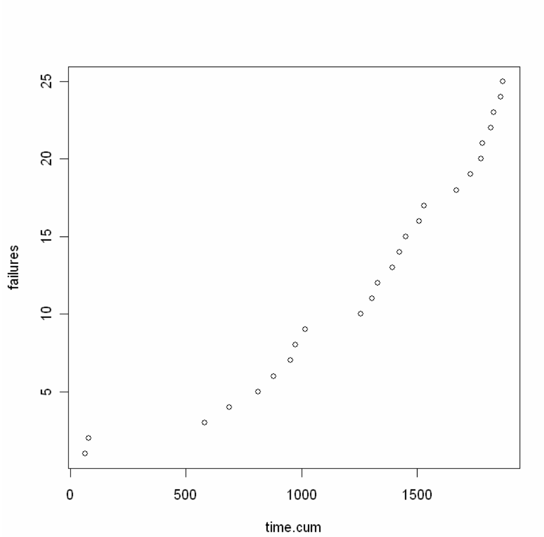

Figure 9 Cumulative failure plot for simulated data from power-law NHPP with increasing ROCOF

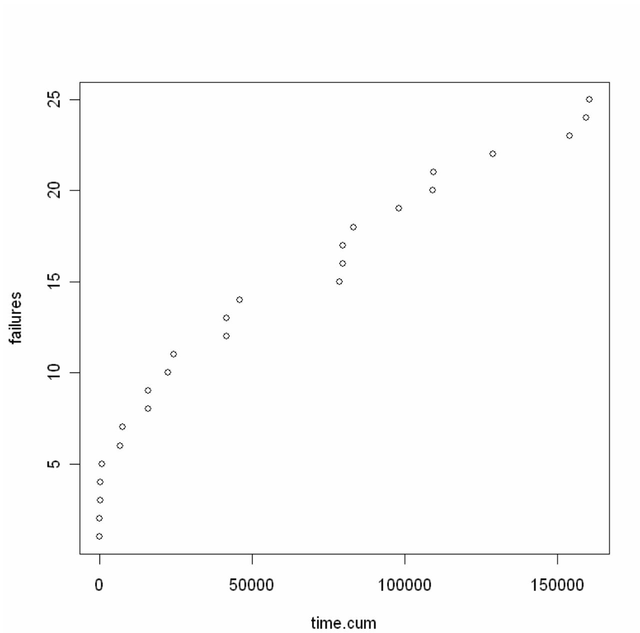

Figure 10 Cumulative failure plot for simulated data from power-law NHPP with decreasing ROCOF

The values of the statistic from our nonparametric test (Laplace U divided by estimated coefficient of variation) are shown below, with the two-sided p-values in parentheses. The null hypothesis of a renewal process would just be rejected at a 0.05 significance level in both of these cases.

Increasing ROCOF: 1.99 (0.046)
Decreasing ROCOF: -1.99 (0.046)

The assumption made regarding repair (same as old versus same as new) is crucial to the analysis. Consider times to failure being produced by a process with increasing ROCOF, corresponding to aging with repair same as old. As time progresses, times between failure will tend to decrease, and there will be a preponderance of short times between failures in a sample. If the process is assumed to be a renewal process, with times between failures described by, for example, a Weibull distribution, the preponderance of short times between failures will cause the estimate of the Weibull shape parameter to be less than one, corresponding to an apparent *decreasing* failure rate, the opposite of what is actually happening. Unfortunately, Bayesian posterior predictive checks may not help much in deciding which model is better, because both models can replicate the observed data quite well.

This can be illustrated by simulation. We generated 1,000 cumulative failure times for a system whose repair is same-as-old, described by a power-law process with shape parameter of 2 and scale parameter of 350. The cumulative failure plot below shows the increasing trend in ROCOF with time.

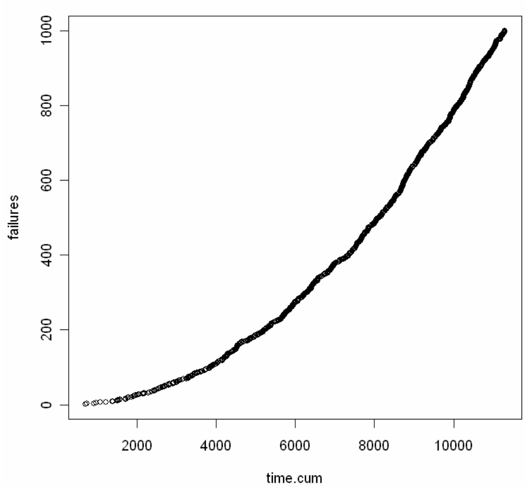

Figure 11 Cumulative failure plot for 1,000 times simulated from power-law process with shape parameter of 2, illustrating increasing ROCOF

The histogram below of the times between failures shows the preponderance of short times between failures caused by the increasing ROCOF.

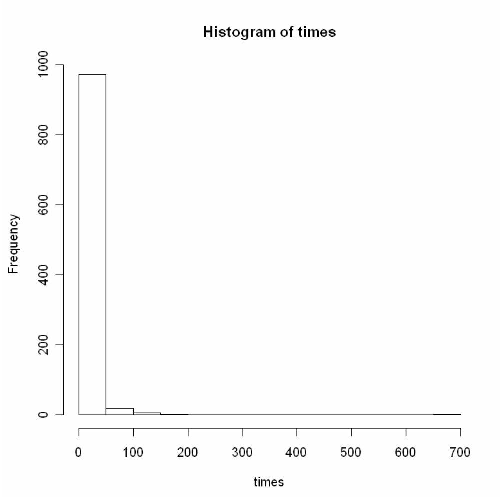

Figure 12 Histogram of times between failures for simulated failure times from power-law process with increasing ROCOF

Assuming the repair is same-as-new instead of same-as-old and fitting a Weibull distribution to these times between failures, one estimates a Weibull shape parameter of about 0.8, which would incorrectly suggest a failure rate that is decreasing with time.

If the repair is same-as-new, and the failure rate increases with operating time or time in standby (whichever is being modeled), an assumption of same-as-old repair will, as suggested by Figure 4, lead to an estimate near one for the shape parameter of the power-law process. In this case, the Bayesian posterior predictive check can be helpful, as a power-law process with shape parameter near one cannot replicate data from a renewal process with increasing failure rate very well.