Analysis of Fatigue Data with Runouts Based on a Model with Nonconstant Standard Deviation and a Fatigue Limit Parameter

Francis G. Pascual

Department of Statistics
Iowa State University
Ames, IA 50011

William Q. Meeker

Department of Statistics and
Center for Nondestructive Evaluation
Iowa State University
Ames, IA 50011

July 9, 1996

Abstract

The fatigue-limit model studied here contains an unknown fatigue limit parameter. Under this model, specimens tested below this fatigue-limit level of stress will never fail. The model also allows the standard deviation of fatigue life to be a function of stress. Researchers can use this model to describe the standard deviation and stress dependence in fatigue data. To illustrate its application, we use maximum likelihood methods to fit the model to fatigue data on a nickel base superalloy. Modern statistical methods based on likelihood ratio provides confidence intervals for the fatigue limit parameter. We also study the effect that test length has on estimation by analyzing simulated data sets based on this model. Through this simulation study, we gain insight on practical test lengths for future fatigue experiments.

Key Words: fatigue curves, fatigue data, fatigue life distribution, fatigue limit, maximum likelihood methods, nonconstant standard deviation, runouts

Notation

$\beta_0^{[\mu]}, \beta_1^{[\mu]}$	unknown coefficients in the $\mu(x)$ relationship
$\beta_0^{[\sigma]}, \beta_1^{[\sigma]}$	unknown coefficients in the $\sigma(x)$ relationship
$\chi^2_{(k,p)}$	100pth percentile of a chi-square random variable with k degrees of freedom
$F(\ )$	lognormal cumulative distribution function (cdf)
$f(\ )$	lognormal probability density function (pdf)
γ	fatigue limit
$L(\ )$	sample likelihood
$\mathcal{L}(\ )$	sample loglikelihood
$\mu(x)$	mean log cycles to failure at stress level x
n	sample size/number of specimens
$\Phi(\ )$	standard normal cdf
$\phi(\ )$	standard normal pdf
$R(\ )$	profile likelihood
$\sigma(x)$	standard deviation of log cycles to failure at stress level x
$\underline{\theta}$	vector of model parameters
x	stress or pseudostress
x_i	stress on specimen i
x_{minf}	lowest observed stress level yielding a failure
Y	number of cycles to failure
y_i	observed number of cycles to failure of specimen i

1 Introduction

Empirical results from fatigue data, particularly on certain steels and ceramics, suggest that specimens tested below a particular stress level are unlikely to fail. This stress level is called the "fatigue limit" or "threshold stress." Fatigue curve models that include a fatigue limit suggests an alternative, possibly more appropriate, model of fatigue data as a function of stress. Such models help engineers estimate design stress levels below which failure is unlikely to occur during a product's design life.

Nelson in [1] uses maximum likelihood (ML) estimation to analyze censored fatigue data (i.e., data with runouts) on a nickel base superalloy. He models the mean and standard deviation of log life as functions of the stress level. However, those models do not include a fatigue limit parameter. Nelson in pp. 92-93 of [2] presents several statistical models with a fatigue limit parameter. In pp. 268-271 of [2] he performs ML analysis on voltage-endurance test data using a lognormal-power model with a fatigue limit. Hirose in [3] analyzes accelerated life-test data applying ML estimation on Weibull-power models to estimate the threshold stress. One of his models involves different Weibull shape parameters at each level of stress.

We have two main results in this paper. First, in section 2 we present a new fatigue life model that involves a fatigue limit and nonconstant standard deviation. Here we discuss profile and maximum likelihood methods and apply them to the superalloy data. In particular, we use these methods to compute confidence intervals for the fatigue limit. We also use standard residual plots to assess the fit of the model. Second, in section 3 we present simulation studies based on the model to investigate how varying the length (in cycles) of the experiment affects estimation. These simulations provide practical information on removal times that cause runouts (specimens that do not fail within the allotted time for the experiment).

2 Data Analysis and Maximum Likelihood Methods

In this section we introduce the fatigue-limit model and fit it to nickel base superalloy fatigue data using ML estimation. We present inferential methods based on ML methods and illustrate them in the data analysis. We study standard residual plots to assess the fit of the model.

2.1 The Nickel-Based Superalloy Data

Table 1 gives low-cycle fatigue data from a stress-controlled test of a nickel base superalloy. The data are from Nelson [1]. A specimen is said to have failed if a failure occurs in the uniform cross section of the cylindrical specimen. It is called a runout if failures occur in the radius, weld or threads or if no failures occur at all. The data on 26 specimens include pseudo-stress (the specimen's Young's modulus $\times$ strain in ksi, 100 ksi=689.4 MPa), the number of test cycles and failure/runout information. Note that there are 4 runouts. Figure 1 is a log-log S-N plot of the data. A "$\bullet$" and a "$\triangleright$" represent a failure and a runout, respectively.

Nelson [1] uses ML methods to fit several fatigue life models to these data and to obtain estimates and approximate confidence intervals for parameters of interest. He fits fatigue curves with constant or nonconstant standard deviation but with no fatigue limit.

Table 1: Fatigue life of nickel-based superalloy specimens

Stress x_i	Cycles y_i	Observation Type	Stress x_i	Cycles y_i	Observation Type
80.3	211 629	failure	99.8	43 331	failure
80.6	200 027	failure	100.1	12 076	failure
80.8	57 923	runout	100.5	13 181	failure
84.3	155 000	failure	113.0	18 067	failure
85.2	13 949	failure	114.8	21 300	failure
85.6	112 968	runout	116.4	15 616	failure
85.8	152 680	failure	118.0	13 030	failure
86.4	156 725	failure	118.4	8 489	failure
86.7	138 114	runout	118.6	12 434	failure
87.2	56 723	failure	120.4	9 750	failure
87.3	121 075	failure	142.5	11 865	failure
89.7	122 372	runout	144.5	6 705	failure
91.3	112 002	failure	145.9	5 733	failure

2.2 The Fatigue-Limit Model

Let $x_1, \ldots, x_n$ denote pseudo-stress levels of n specimens and let $Y_1, \ldots, Y_n$ be the corresponding numbers of cycles to failure respectively. The values of the random variables $Y_1, \ldots, Y_n$ can either be actual failure times or, in the case of runouts, test termination

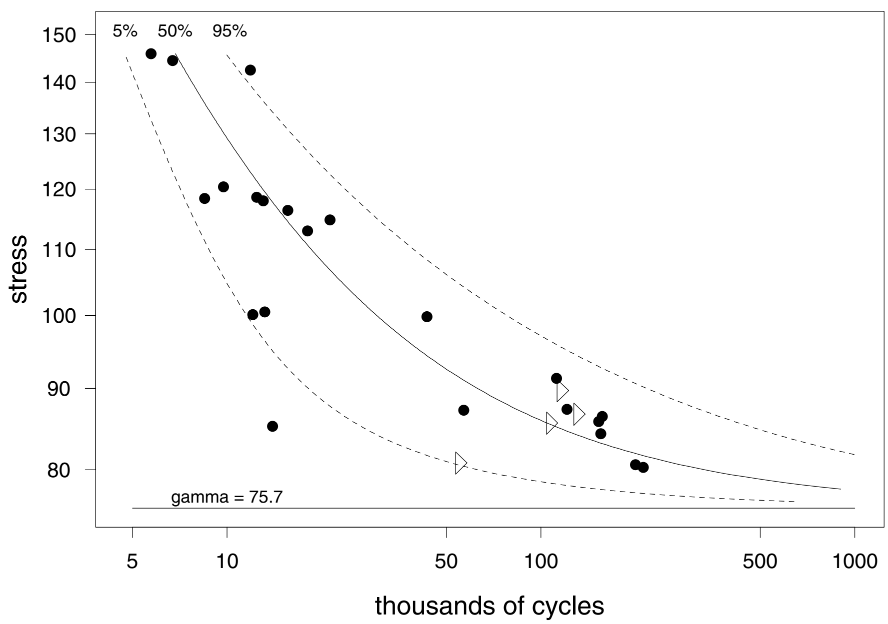

Figure 1: Log-log S-N plot for the superalloy data with ML estimates of the 5, 50 and 95 percentiles from the fatigue limit model (• failure, ▹ runout)

times which may vary from specimen to specimen. Let γ be the fatigue limit. At each pseudo-stress level with $x_i > \gamma$, fatigue life Y_i is modeled with a lognormal distribution, i.e., the cumulative proportion failing function and its derivative are given by

$$\begin{aligned}\Pr(Y_i \le y) = F(y;\mu(x_i),\sigma(x_i)) &= \Phi\left[\frac{\log(y)-\mu(x_i)}{\sigma(x_i)}\right],\\ f(y;\mu(x_i),\sigma(x_i)) &= \frac{1}{\sigma(x_i)y}\phi\left[\frac{\log(y)-\mu(x_i)}{\sigma(x_i)}\right],\ y>0,\end{aligned}$$

where Φ and ϕ are, respectively, the cumulative distribution function and probability density function of the standard normal distribution. This implies that $\log(Y_i)$ is modeled with a normal distribution with mean $\mu(x_i)$ and standard deviation $\sigma(x_i)$. In our fatigue-limit model, these parameters are related to stress according to

$$\mu(x_i) = \mathrm{E}[\log(Y_i)] = \beta_0^{[\mu]} + \beta_1^{[\mu]}\log(x_i-\gamma),\ x_i > \gamma, \tag{1}$$

$$\sigma(x_i) = \sqrt{\mathrm{Var}(\log(Y_i))} = \exp[\beta_0^{[\sigma]} + \beta_1^{[\sigma]}\log(x_i)],\ x_i > \gamma, \tag{2}$$

where $\beta_0^{[\mu]}$, $\beta_1^{[\mu]}$, $\beta_0^{[\sigma]}$, $\beta_1^{[\sigma]}$ and γ are unknown parameters to be estimated from data. There are no restrictions on the values of $\beta_0^{[\mu]}$, $\beta_1^{[\mu]}$, $\beta_0^{[\sigma]}$ and $\beta_1^{[\sigma]}$. If x_{minf} is the smallest observed stress level that yields a failure, then γ must be in the interval $(0, x_{minf})$.

Note that when $\beta_1^{[\sigma]} = 0$, the model has a constant standard deviation. In most fatigue data, the standard deviation decreases as stress increases, which corresponds to $\beta_1^{[\sigma]} < 0$. The plot of the superalloy data in Figure 1 indicates more horizontal scatter at the lower stress levels and less at the higher stress levels.

The size of γ determines the amount of curvature present in the plotted log-log S-N curve. When γ is close to zero, the S-N curve is close to linear. Larger values of γ result in more curvature in the plot. When $\gamma = 0$, the model is equivalent to models (3) and (5) in Nelson [1]. Curvature in Figure 1 suggests the inclusion of a fatigue limit γ in the model. Although a fixed fatigue limit may be unrealistic for describing a population of specimens, the fatigue limit provides a physically appealing alternative to the quadratic term in the $\mu(x_i)$ relationship used by Nelson in [1] for describing S-N curvature.

The maximum likelihood methods described in the next section use the following assumptions: a) specimens are tested independently and b) for $x_i > \gamma$ the times at which observations become runouts are independent of actual failure times if the experiment were to be run until failure.

2.3 Maximum Likelihood Estimation

Based on the parametric statistical model presented above, we estimate parameters by ML estimation. Statistical theory suggests that ML estimators, in general, have favorable asymptotic (large sample) properties. For "large" sample sizes and under certain conditions on the fatigue life distribution, the distribution of a ML estimator is approximately normal with mean equal to the true value being estimated and standard deviation no larger than that of any other competing estimator. See Chapter 5 of [2] for an in-depth discussion of ML estimation.

Aside from estimates, we also compute approximate likelihood-ratio-based confidence intervals for model coefficients. Ostrouchov and Meeker [4] conduct Monte Carlo simulations to compare the accuracy of confidence intervals based on likelihood ratio and those based on asymptotic normal theory for interval censored Weibull and lognormal data. They conclude that likelihood confidence intervals have coverage probabilities closer to nominal confidence levels than those of normal approximation intervals even in small to moderate size samples.

2.3.1 Parametric Likelihood

For the fatigue-limit model defined by equations (1) and (2) with sample data y_1, ..., y_n, the likelihood is

$$L(\beta_0^{[\mu]}, \beta_1^{[\mu]}, \beta_0^{[\sigma]}, \beta_1^{[\sigma]}, \gamma) = \prod_{i=1}^{n} \left[\frac{1}{\sigma(x_i) y_i} \phi(z_i) \right]^{\delta_i} [1 - \Phi(z_i)]^{1-\delta_i}$$

where $z_i = [\log(y_i) - \mu(x_i)]/\sigma(x_i)$, $\mu(x_i)$ and $\sigma(x_i)$ are given by equations (1) and (2), respectively, and

$$\delta_i = \begin{cases} 1 & \text{if } y_i \text{ is a failure} \\ 0 & \text{if } y_i \text{ is a runout.} \end{cases}$$

Let $\underline{\theta} = (\beta_0^{[\mu]},\ \beta_1^{[\mu]},\ \beta_0^{[\sigma]},\ \beta_1^{[\sigma]},\ \gamma)$ be the vector of model parameters. The function $L(\underline{\theta})$ can be interpreted as being approximately proportional to the probability of observing y_1, ..., y_n, for a given set of parameters $\underline{\theta}$. Generally, it is easier to work with the log-likelihood function

$$\mathcal{L}(\underline{\theta}) = \log[L(\underline{\theta})] = \sum_{i=1}^{n} \mathcal{L}_i(\underline{\theta})$$

where

$$\mathcal{L}_i(\underline{\theta}) = \delta_i\{\log[\phi(z_i)] - \log[\sigma(x_i)y_i)]\} + (1-\delta_i)\log[1-\Phi(z_i)].$$

The ML estimate $\hat{\underline{\theta}}$ of $\underline{\theta}$ is the set of parameter values that maximizes $L(\underline{\theta})$ or $\mathcal{L}(\underline{\theta})$. Table 2 gives the ML estimates of all model parameters resulting from fitting the fatigue-limit model to the data. Figure 1 shows curves of the ML estimates of the 5, 50 and 95 percentiles of fatigue life.

Parameter	Estimate	Normal-Theory Confidence Interval	Likelihood-Ratio Confidence Interval
$\beta_0^{[\mu]}$	14.75	(12.06, 17.44)	(12.90, 21.45)
$\beta_1^{[\mu]}$	−1.39	(−2.02, −0.76)	(−2.81, −0.92)
$\beta_0^{[\sigma]}$	10.97	(3.82, 18.12)	(3.22, 17.90)
$\beta_1^{[\sigma]}$	−2.50	(−4.04, −0.96)	(−3.98, −0.81)
γ	75.71	(67.35, 84.06)	(49.98, 79.79)

Table 2: Maximum likelihood results for the superalloy data

In pp. 72-73 of [1] Nelson comments that some of the fatigue life models that he presents produce percentiles larger at an intermediate stress than at a lower stress which is physically impossible. Although this is also theoretically possible for the fatigue-limit model, it does not occur within the range of interest for these data.

2.3.2 Profile Likelihoods and Likelihood-Ratio-Based Confidence Regions

Here we introduce the profile likelihood and use it to compute approximate confidence intervals for the fatigue limit parameter γ. The profile likelihood is an important tool for making inferences about model parameters or functions of them. Let $\underline{\theta} = (\underline{\theta}_0, \gamma)$ and $\hat{\gamma}$ denote the ML estimate of γ. Aside from $\hat{\gamma}$, we may be interested in other probable values of γ. The profile likelihood can be used to assess the plausibility of other values of γ.

The profile likelihood for γ is defined by

$$R(\gamma) = \max_{\underline{\theta}_0} \left[\frac{L(\underline{\theta}_0, \gamma)}{L(\hat{\underline{\theta}})} \right].$$

A large value (close to 1) of $R(\gamma)$ indicates that the obtained data are highly probable for that value of γ. On the other hand, a small value (close to 0) of $R(\gamma)$ indicates that the observed data are relatively unlikely for the given value of γ. Plotting $R(\gamma)$ against different values of γ yields a profile likelihood plot for γ.

If γ_0 is the true value of γ, then $-2\log[R(\gamma_0)]$ follows approximately, a chi-square distribution with 1 degree of freedom. As a result, an approximate $100(1-\alpha)\%$ likelihood confidence region for γ is given by the set of all γ such that

$$-2\log[R(\gamma)] \leq \chi^2_{(1;1-\alpha)}$$

or, equivalently,

$$R(\gamma) \geq \exp\left[-\frac{\chi^2_{(1;1-\alpha)}}{2}\right]$$

where $\chi^2_{(1;1-\alpha)}$ is the $100 \times (1-\alpha)$ percentile of a chi-square distribution with 1 degree of freedom. Confidence intervals based on the approximate normal distribution of ML estimators can also be computed. See pp. 292-297 of [2]. However, as mentioned earlier likelihood confidence intervals perform better in the sense that coverage probabilities are closer to nominal confidence levels than those of normal approximation intervals.

The confidence intervals in Table 2 indicate that the parameters $\beta_1^{[\mu]}$, $\beta_1^{[\sigma]}$ and γ are statistically significantly different from zero. These intervals indicate that there is a relationship between mean fatigue life and the stress level. The confidence intervals for $\beta_1^{[\sigma]}$ indicate that the standard deviation of fatigue life depends on the stress level and, moreover, that the standard deviation decreases as stress increases, a commonly observed phenomenon in metal fatigue data. The confidence intervals for γ support the inclusion of a fatigue limit as suggested by the curvature in Figure 1.

Figure 2 gives a profile likelihood plot for γ. The scale on the right-hand side of the plot corresponds to confidence levels for confidence intervals based on likelihood ratios. For example, the intersections of the horizontal line through 0.95 and the likelihood curve correspond to lower and upper limits for an approximate 95% confidence interval for γ. For the superalloy data, we are 95% confident that the interval $(49.98, 79.79)$ contains the population fatigue limit.

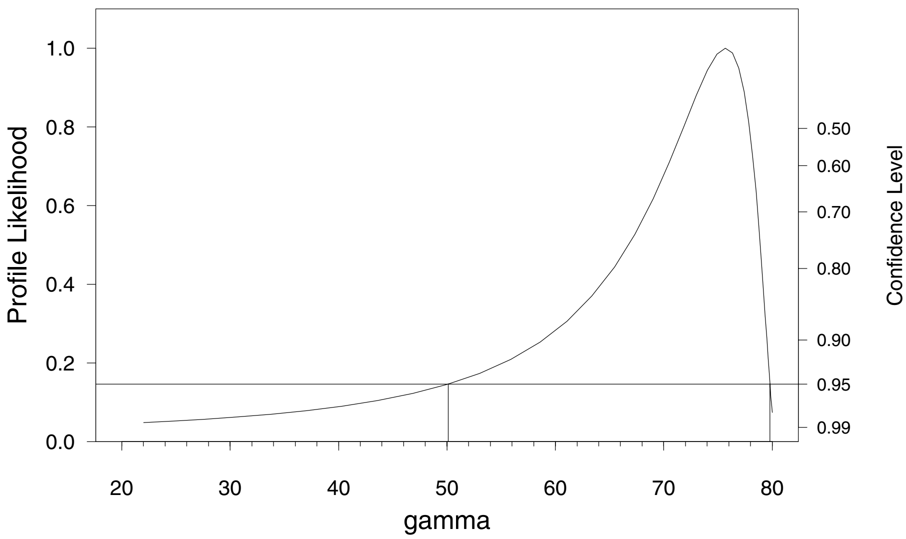

Figure 2: Likelihood plot for the Fatigue Limit γ for the superalloy data

2.4 Residual Plots

To assess the appropriateness of the model, plots of residuals versus standard normal percentiles, the stress levels and the mean estimates (predicted values) are helpful. Our approach parallels that used in [1].

For the stress level x_i, the $\mu(x_i)$ estimate and raw residual are, respectively, given by

$$\hat{\mu}(x_i) = \hat{\beta}_0^{[\mu]} + \hat{\beta}_1^{[\mu]} \log(x_i - \hat{\gamma})$$

and

$$e_i = \log(y_i) - \hat{\mu}(x_i).$$

We compute the ML estimate $\hat{\sigma}(x_i)$ of the standard deviation of log fatigue life at the stress level x_i by evaluating (2) at the ML estimates. Because the standard deviation varies for different stress levels, we use standardized residuals

$$e_i^* = \frac{e_i}{\hat{\sigma}(x_i)}$$

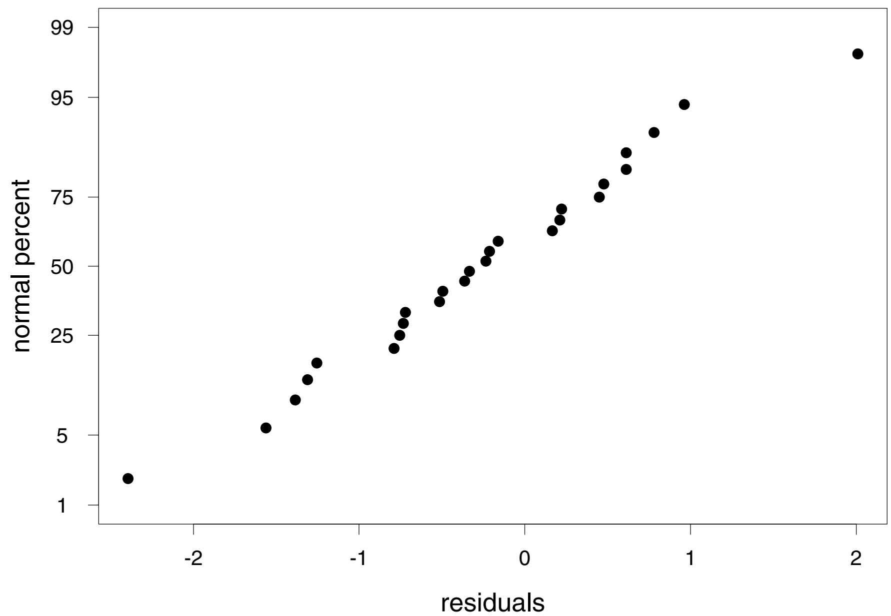

Figure 3: Probability plot of standardized residuals for the superalloy data

in the plots. We will refer to e_i^* as residuals, henceforth.

Since we model fatigue life with a lognormal distribution at every stress level, a normal probability plot of the residuals should appear linear to indicate a good fit. Figure 3 gives a normal probability plot of the residuals. On the other hand, the plots of residuals versus the stress levels and the mean estimates $\hat{\mu}(x_i)$, respectively, should appear patternless. Figures 4 and 5 give the plots of the residuals versus the stress levels and mean estimates, respectively.

The points in Figure 3 lie close to a straight line. This suggests that the lognormal distribution describes the data reasonably. Figures 4 and 5 do not show any clear patterns in the residuals. These plots indicate a plausible fit of the fatigue-limit model to the data.

The fifth observation in Table 1 deviates from the general pattern of the S-N plot in Figure 1. It has the largest negative residual. This point influences the conclusion we make regarding the statistical significance of the coefficient $\beta_1^{[\sigma]}$. In the analysis based on deleting this point, the confidence intervals for $\beta_1^{[\sigma]}$ contain 0. That is, if this point is not

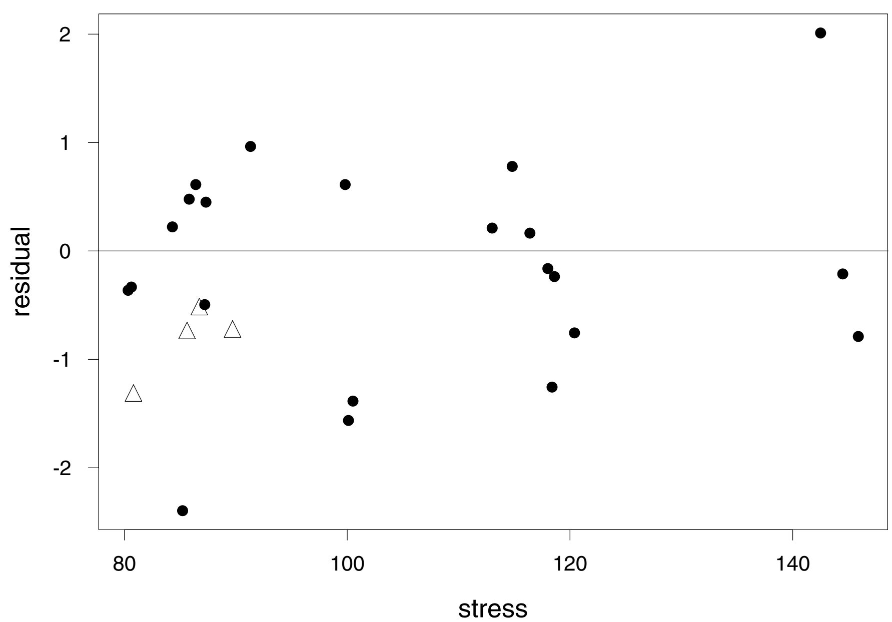

Figure 4: Plot of standardized residuals versus stress for the superalloy data (• failure, △ runout)

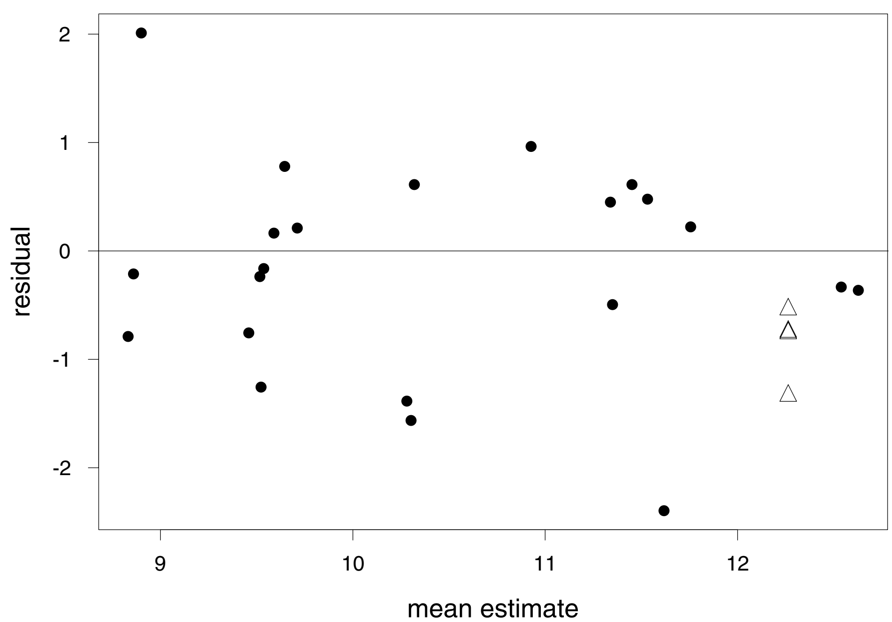

Figure 5: Plot of standardized residuals versus $\hat{\mu}(x_i)$ for the superalloy data ($\bullet$ failure, $\triangle$ runout)

representative of the experiment, then there is less evidence for a nonconstant standard deviation. It would be important to determine whether or not this observation is valid. The 24th observation produces the largest positive residual but it is not influential in the same manner as the fifth observation.

3 Results of Simulations

We now present results of analyses of simulated data to study the effect that varying the test length has on the performance of ML estimates of model parameters in practical testing situations. To achieve a desirable level of precision in estimation, we need to observe at least a certain number of failures in the fatigue experiment. This means that for stress levels close to the fatigue limit, we need to run tests for long periods of time and this may not be practical. On the other hand, we may not observe enough failures if we shorten test runs. The simulation results provide us insight to the trade-offs between test lengths and estimation precision. This insight serves as a guideline for designing future fatigue tests.

We consider both the constant and nonconstant standard deviation cases. In each case, we fit the fatigue-limit model to the data and compare ML estimates and likelihood confidence intervals for the fatigue limit γ for different test lengths.

3.1 Example 1: Simulated Data with Constant Standard Deviation

The uncensored simulated data of size 35 in Table 3 were generated using the fatigue-limit model with $\beta_0^{[\mu]} = 15$, $\beta_1^{[\mu]} = -1.5$, $\beta_0^{[\sigma]} = -0.9$, $\beta_1^{[\sigma]} = 0$ and $\gamma = 76$. Since $\beta_1^{[\sigma]} = 0$, the standard deviation of fatigue life is constant for all stress levels. The stress levels are equally-spaced in the log scale. We fit the fatigue-limit model to the data for different test lengths and present ML results below.

We vary the test length to study its effect on the estimation of the fatigue limit γ. We consider test lengths of 100, 70, 30 and 20 thousand cycles. Fatigue lives exceeding the test length are runouts. Shortening the test length simulates stopping the test earlier. Figures 6 and 7 give S-N plots for test lengths 100 and 20 thousand cycles with the ML estimates of the 5, 50 and 95 percentiles of fatigue life distribution. Again, "$\bullet$" and "$\triangleright$"

Table 3: Example 1: Simulated data with constant standard deviation

Stress x_i	Cycles y_i	Stress x_i	Cycles y_i
80.00	646 484	110.80	15 810
81.46	298 375	112.82	21 490
82.95	181 605	114.88	18 401
84.46	139 041	116.98	8 182
86.00	76 762	119.11	9 446
87.57	54 177	121.29	9 848
89.17	118 582	123.50	8 574
90.80	52 888	125.76	11 671
92.46	39 621	128.06	5 311
94.15	57 573	130.39	7 592
95.87	20 174	132.77	6 427
97.62	35 983	135.20	3 407
99.40	15 619	137.67	4 468
101.21	12 975	140.18	7 692
103.06	31 611	142.74	5 539
104.94	22 472	145.35	4 554
106.86	15 591	148.00	7 789
108.81	11 522		

represent a failure and a runout, respectively.

For each test length, Table 4 gives the ML estimate and a 95% confidence interval for the fatigue limit. Figure 8 gives the relative likelihood profile plots for γ.

The width of the confidence interval depends heavily on the length of the fatigue tests. As expected, longer test lengths result in shorter confidence intervals. Test lengths of 70 to 100 thousand cycles yield more meaningful lower confidence bounds for γ. The shorter test lengths give lower confidence bounds equal to zero, the smallest possible value for the fatigue limit. Notice that the estimates are fairly accurate even for the shorter test lengths. It appears that for this data set there is no significant loss of accuracy but there is an

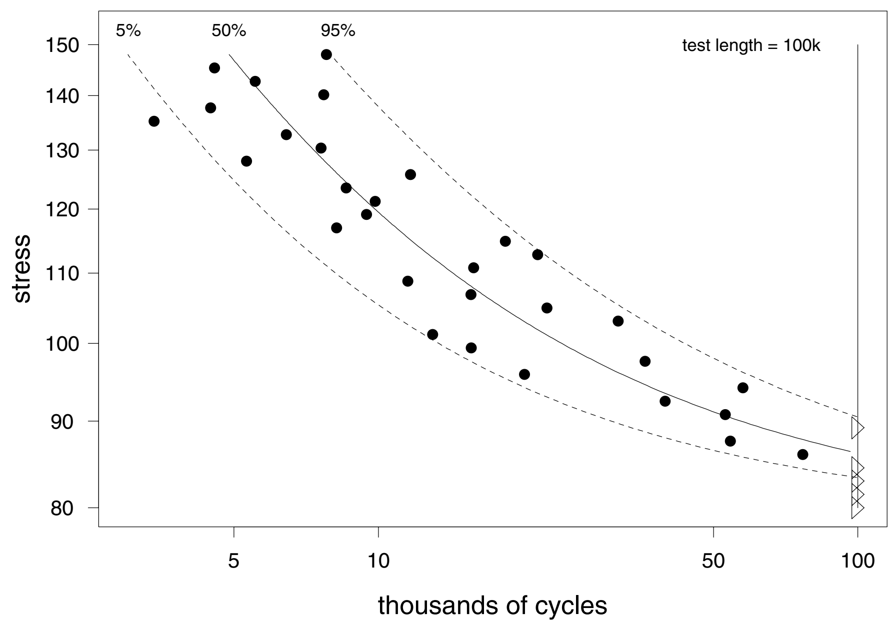

Figure 6: Log-log S-N plot for simulated data (constant standard deviation) with the ML estimates of the 5, 50 and 95 fatigue life percentiles for test length of 100 thousand cycles (• failure, ▹ runout)

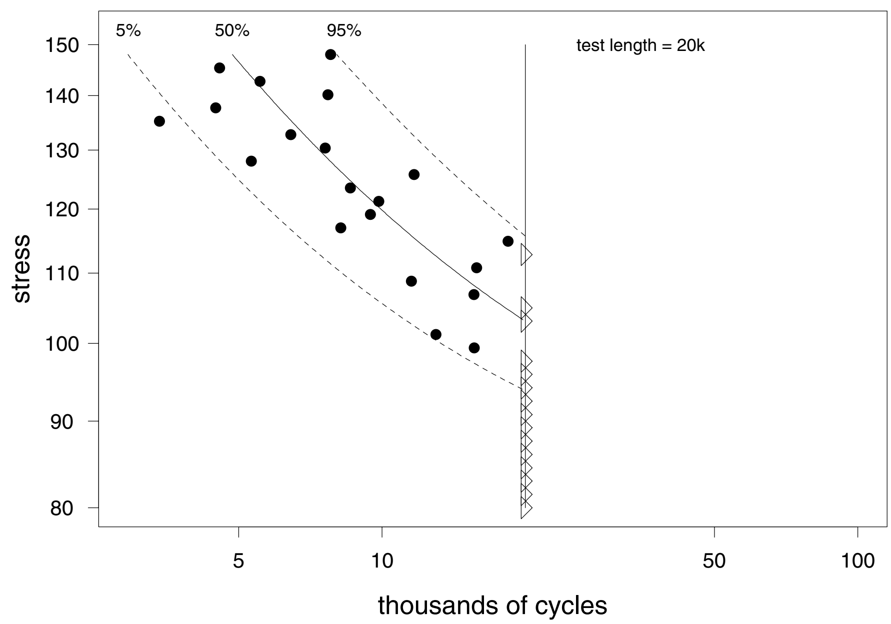

Figure 7: Log-log S-N plot for simulated data (constant standard deviation) with the ML estimates of the 5, 50 and 95 fatigue life percentiles for test length of 20 thousand cycles (• failure, ▹ runout)

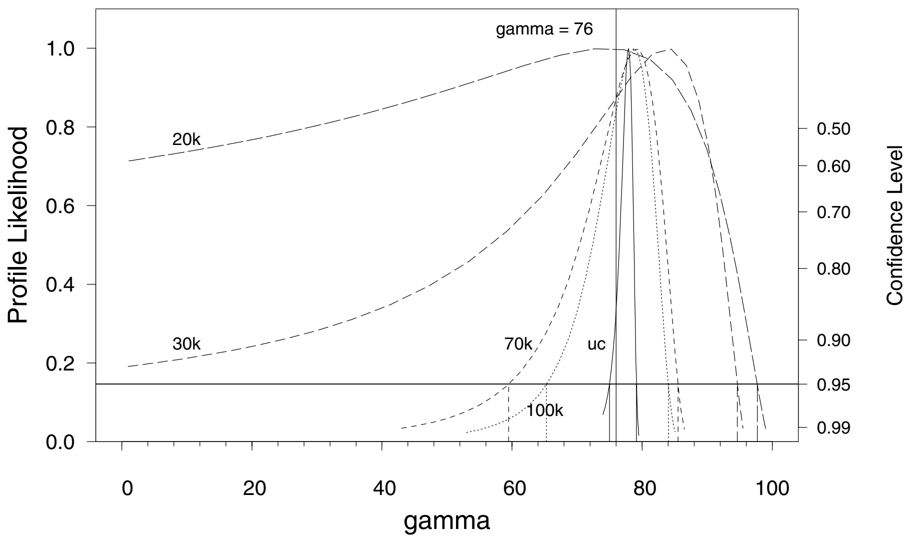

Figure 8: Relative likelihood plots for simulated data with constant standard deviation

increased uncertainty in estimation when we terminate testing earlier.

Test Length (thousands of cycles)	ML Estimate for γ	Likelihood-Ratio Confidence Interval for γ
uncensored	77.90	[74.97, 79.14]
100	78.74	[65.28, 84.04]
70	79.14	[59.45, 85.52]
30	83.75	[0, 94.64]
20	74.67	[0, 97.69]

Table 4: Maximum likelihood results for simulated data with constant standard deviation

3.2 Example 2: Simulated Data with Nonconstant Standard Deviation

The uncensored simulated data in Table 5 were obtained using the fatigue-limit model with $\beta_0^{[\mu]} = 15$, $\beta_1^{[\mu]} = -1.4$, $\beta_0^{[\sigma]} = 10$, $\beta_1^{[\sigma]} = -2.3$ and $\gamma = 76$. Since $\beta_1^{[\sigma]} = -0.7$, the standard deviation of fatigue life is a decreasing function of the stress level. The stress levels are equally-spaced in the log scale.

Table 5: Example 2: Simulated data with nonconstant standard deviation

Stress x_i	Cycles y_i	Stress x_i	Cycles y_i
80.00	475 733	111.59	11 683
81.49	43 662	113.67	11 348
83.01	113 523	115.79	29 415
84.56	569 601	117.95	14 242
86.14	296 853	120.15	16 732
87.75	238 892	122.40	29 686
89.39	86 889	124.68	11 059
91.05	186 854	127.01	14 750
92.75	30 758	129.38	9 586
94.48	49 030	131.79	9 221
96.25	20 009	134.25	9 226
98.04	63 342	136.76	12 170
99.87	53 142	139.31	14 627
101.74	62 442	141.91	9 934
103.63	42 383	144.55	9 160
105.57	54 222	147.25	7 060
107.54	12 615	150.00	9 159
109.54	18 067		

We consider test lengths of 100, 60, 20 and 15 thousand cycles. Figures 9 and 10 give S-N plots for test lengths of 100 and 20 thousand cycles with the ML estimates of the 5, 50 and 95 percentiles of fatigue life.

Table 6 gives the ML estimate and a 95% confidence interval for γ for each test length.

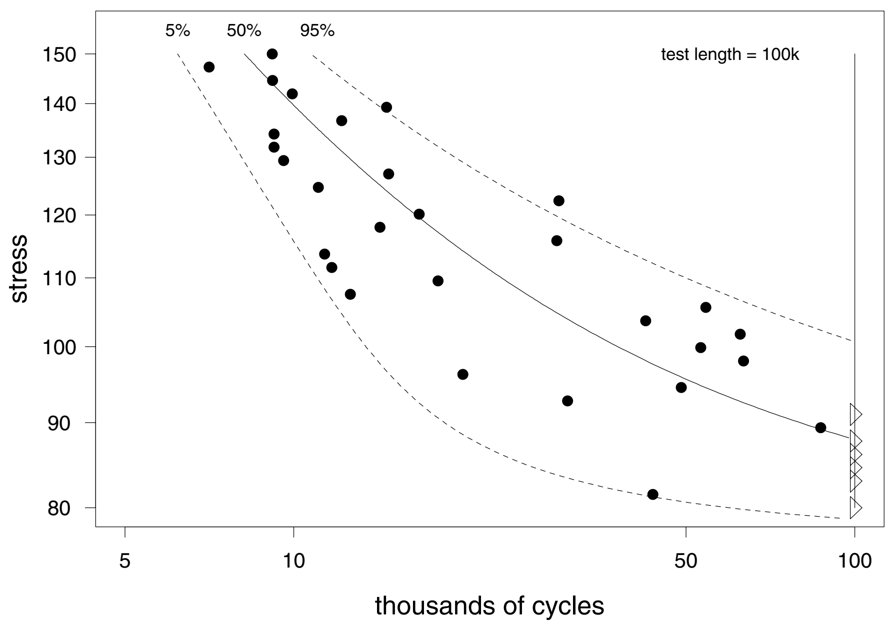

Figure 9: Log-log S-N plot for simulated data (nonconstant standard deviation) with the ML estimates of the 5, 50 and 95 fatigue life percentiles for test length of 100 thousand cycles (• failure, ▹ runout)

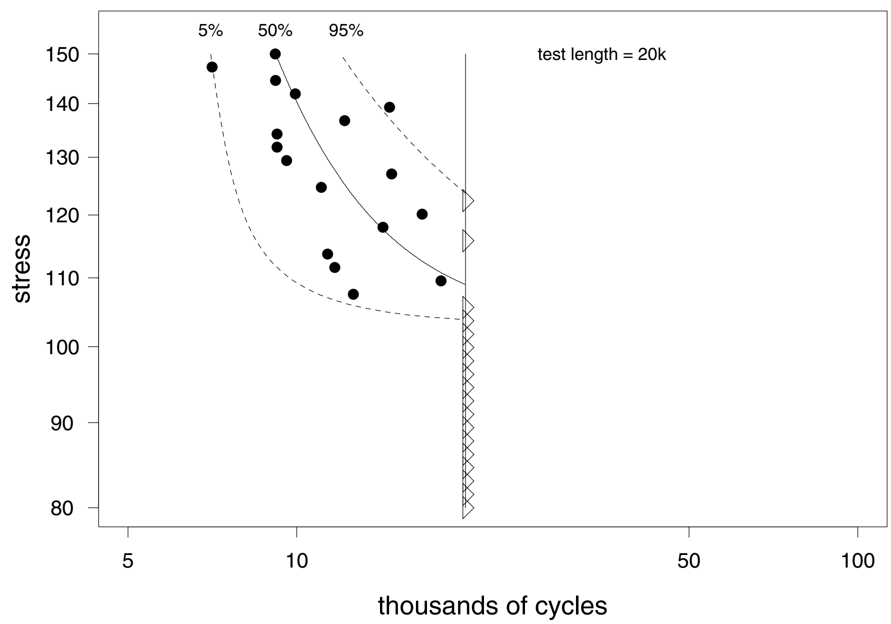

Figure 10: Log-log S-N plot for simulated data (nonconstant standard deviation) with the ML estimates of the 5, 50 and 95 fatigue life percentiles for test length of 20 thousand cycles (● failure, ▷ runout)

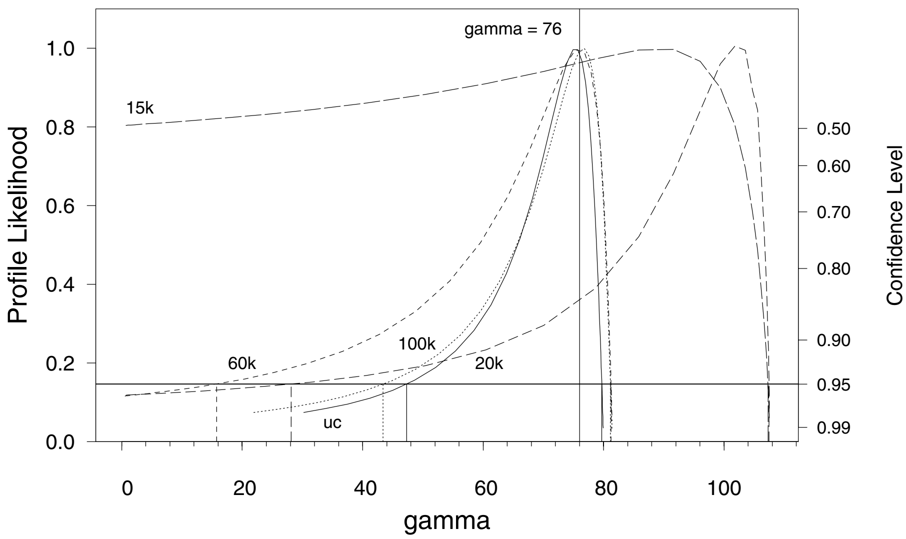

Figure 11: Relative likelihood plots for simulated data with nonconstant standard deviation

Figure 11 gives the relative likelihood profiles for the different test lengths.

Table 6: Maximum likelihood results for simulated data with nonconstant standard deviation

Test Length (thousands of cycles)	ML Estimate for γ	Likelihood-Ratio Confidence Interval for γ
uncensored	75.34	[47.28, 79.70]
100	76.66	[43.39, 81.14]
60	75.94	[15.73, 81.25]
20	103.07	[28.25, 107.49]
15	89.33	[0, 107.30]

As in the previous example, the results indicate that the confidence intervals for γ are narrower for longer test lengths. Note that when we shorten the test length from 100 to 20 thousand cycles, the ML estimate of γ shifts from 75.94 to 103.07. The profile likelihood for the 20 thousand cycles data shows, however, that the other ML estimates are consistent as they lie within the 20 thousand cycles confidence interval for γ.

3.3 Discussion

In both simulated data sets, it is clear that we get wider confidence intervals for γ when we shorten the test, as expected. This reflects the loss of information from the data due to runouts. Shortening test lengths obviously results in more runouts. At a fixed stress level, a runout at 100 thousand cycles is more informative than one at 70 thousand cycles. Shorter tests result in higher values of x_{minf}, the lowest stress level at which a failure occurs. This widens the interval of possible values for γ because γ can be any value between 0 and x_{minf}.

In the simulation examples above, the amount of curvature in the S-N plot of the data is closely related to the uncertainty in estimating γ as reflected in the width of the confidence intervals. As can be seen in the S-N plots, there is a fair amount of curvature in longer tests for which narrower confidence intervals are obtained. There is less observed curvature in the plot in shorter tests. Here, the confidence intervals are wider.

These results are relevant to life-test experiments designed to estimate, usually by extrapolation, proportions failing or fatigue life percentiles at low levels of stress. In the presence of a fatigue limit, extrapolation may yield inaccurate estimates especially if the S-N plot of the data does not show enough curvature. Thus, the tasks of choosing stress levels and lengths of tests involved in designing experiments require careful thought and should take into account the appropriateness of assuming a fatigue limit.

4 Conclusion and Suggestions for Further Research

The fatigue-limit model is motivated by the possible existence of the threshold stress below which testing yields no failure. Although the existence of such a limit may be questioned, this model still provides an alternative to a quadratic fatigue curve model for describing curvature in S-N plots. Quadratic models have been shown to produce larger fatigue life percentiles at higher stress levels than at lower levels which is not physically possible.

Although this may happen with the fatigue-limit model, it does not occur within the range of the superalloy data.

In the simulation studies above, we have studied the effects of different test lengths on the estimation of the fatigue limit. We can quickly obtain results by running short tests on specimens but probably at the expense of estimation precision. Running long tests may give us the desired precision but this may not be feasible for practical purposes. Such simulation studies provide experimenters with information on possible trade-offs between test lengths and estimation precision.

The simulated data sets studied here included equally-spaced stress levels (in log scale) with one observation at each level. Alternative experimental designs might perform better with respect to the objectives of the experiment. The researcher or engineer, for example, may be interested in estimating certain functions of $\underline{\theta}$ (e.g., a fatigue life percentile at a specified stress level x_0; or a stress level X such that, given p and y_0, the $100p$th percentile of fatigue life at this stress level is $Y_p(X) = y_0$). Interest may be in predicting the fatigue life at a given stress level of new specimens. With these and other possible considerations in mind, methods for comparing and choosing among different experimental designs should be studied.

In this paper, we considered only the case when fatigue life follows a lognormal distribution at a given stress level. Other distributions such as the Weibull and logistic distributions need to be studied.

The fatigue-limit model, however, has some shortcomings. From a practical point of view, it does not seem plausible to think of all specimens having a common fatigue limit. In pp. 93-95 of [2] Nelson suggests modeling the fatigue limit with a distribution called the "strength distribution," i.e., each specimen has a different fatigue limit according to this distribution. If this is indeed true, the fatigue-limit model can still be applied when the spread of the strength distribution is small.

We are currently investigating fatigue life models that involve fatigue curves with fatigue limits and strength distributions. The S-N plots of simulated data sets based on these models exhibit curvature and nonconstant standard deviation that we likewise observe in the superalloy data. The inclusion of a strength distribution allows us to model nonconstant standard deviation without directly writing the standard deviation as a function of stress.

Acknowledgments

We would like to thank Chuck Annis, Luis A. Escobar and two anonymous referees for providing helpful comments on an earlier version of this paper.

REFERENCES

[1] Nelson, W. (1984). "Fitting of Fatigue Curves with Nonconstant Standard Deviation to Data with Runouts." *Journal of Testing and Evaluation* 12, 69-77.

[2] Nelson, W. (1990), *Accelerated Testing: Statistical Models, Test Plans, and Data Analyses*, New York: John Wiley & Sons.

[3] Hirose, H. (1993). "Estimation of Threshold Stress in Accelerated Life-Testing." *IEEE Transactions on Reliability* 42, 650-657.

[4] Ostrouchov, G. and Meeker, W. (1988). "Accuracy of Approximate Confidence Bounds Computed from Interval Censored Weibull and Lognormal Data." *Journal of Statistical Computing and Simulations* 29, 43-76.