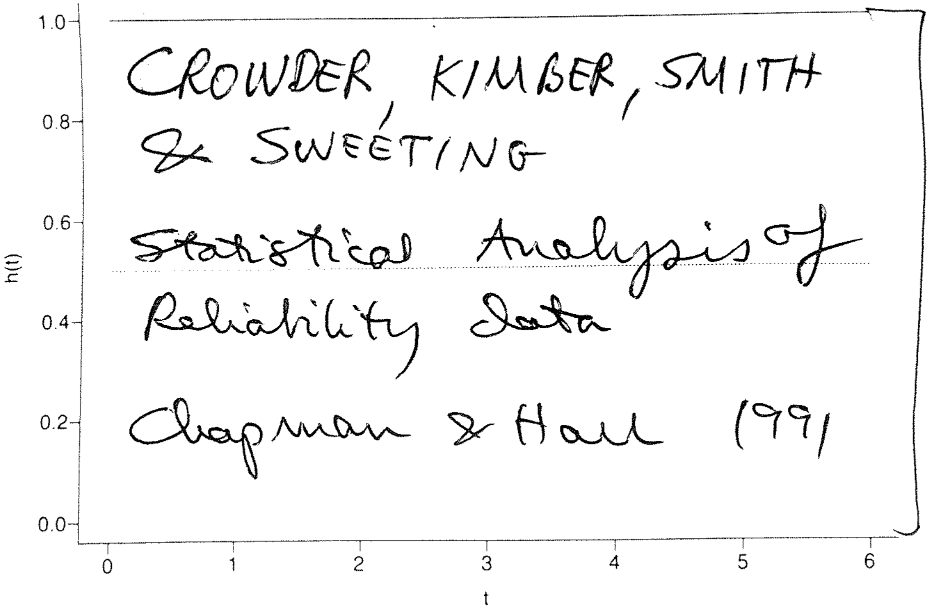

Figure 2.2 Exponential hazard functions with mean 1 (solid) and mean 2 (dotted).

2.4 THE WEIBULL AND GUMBEL DISTRIBUTIONS

A Weibull random variable (after W. Weibull (1939, 1951)) is one with survivor function

$$S(t)=\exp\{-(t/\alpha)^{\eta}\}, \tag{2.3}$$

for $t>0$ and where α and η are positive parameters, α being a scale parameter and η being a shape parameter. Note that when $\eta=1$, we obtain an exponential distribution with $\lambda=1/\alpha$.

The Weibull hazard function is

$$h(t)=\eta\alpha^{-\eta}t^{\eta-1}.$$

This is DFR for $\eta<1$, constant for $\eta=1$ (exponential) and IFR for $\eta>1$. In particular, for $1<\eta<2$, the hazard function increases slower than linearly; for $\eta=2$ the hazard function is linear; and for $\eta>2$ the hazard increases faster than linearly. A selection of Weibull hazard functions is shown in Figure 2.3.

The Weibull density is

$$f(t)=\eta\alpha^{-\eta}t^{\eta-1}\exp\{-(t/\alpha)^{\eta}\}, \tag{2.4}$$

for $t>0$. The mean and variance are given by $\alpha\Gamma(\eta^{-1}+1)$ and $\alpha^2\{\Gamma(2\eta^{-1}+1)-[\Gamma(\eta^{-1}+1)]^2\}$, where Γ is the gamma function

$$\Gamma(x)=\int_0^{\infty} u^{x-1}e^{-u}\,du, \tag{2.5}$$

see, for example, Abramowitz and Stegun (1972, Chapter 6). A Fortran program for computing equation (2.5) is given in Griffiths and Hill (1985, pp. 243–6), which is based on an earlier program of Pike and Hill (1966). When η is large (greater than 5, say), the mean and variance are approximately α and $1.64\alpha^2/\eta^2$ respectively. The shape of the density depends on η. Some Weibull densities are shown in Figure 2.4.

The Weibull distribution is probably the most widely used distribution in reliability analysis. It has been found to provide a reasonable model for lifetimes of many types of unit, such as vacuum tubes, ball bearings and composite materials. A possible explanation for its appropriateness rests on its being an extreme value distribution; see Galambos (1978). Moreover, the closed form of the Weibull survivor function and the wide variety of shapes exhibited by Weibull density functions make it a particularly convenient generalization of the exponential distribution.

The Gumbel (or extreme-value, or Gompertz) distribution has survivor function

$$S(x)=\exp\{-\exp[(x-\mu)/\sigma]\} \tag{2.6}$$

for $-\infty<x<\infty$, where μ is a location parameter and $\sigma>0$ is a scale parameter. This distribution also arises as one of the possible limiting distributions of minima, see Galambos (1978), and has exponentially increasing failure rate. It is sometimes used as a lifetime distribution even though it allows negative values with positive probability. More commonly, however, the Gumbel distribution arises as the distribution of log T. This is equivalent to assuming that T has a Weibull distribution. The relationship between the Gumbel and Weibull parameters is $\mu=\log\alpha$ and $\sigma=1/\eta$.

The Gumbel density function is

$$f(x)=\sigma^{-1}\exp\{(x-\mu)/\sigma\}S(x) \tag{2.7}$$

for $-\infty<x<\infty$, and has the same shape for all parameters. Note that the mean and variance of a Gumbel random variable are $\mu-\gamma\sigma$ and $(\pi^2/6)\sigma^2$ respectively, where $\gamma=0.5772\ldots$ is Euler's constant, and the distribution is negatively skewed. The density and hazard functions for a Gumbel distribution with $\mu=0$ and $\sigma=1$ are shown in Figures 2.5 and 2.6 respectively.

In practice $\hat{H}(t)$ and $\tilde{H}(t)$ are usually close.

At a specified value t^* the standard error of $\hat{H}(t^*)$ and $\tilde{H}(t^*)$ may be approximated, using (2.21) and the delta method, by

$$\text{se}\{\hat{H}(t^*)\} = \text{se}\{\tilde{H}(t^*)\} = \left\{\sum^{(t)} \frac{d_j}{n_j(n_j-d_j)}\right\}^{1/2}.$$

As in sections 2.9 and 2.10 a plot of $\hat{S}(t)$ versus t can be informative. This may be constructed for all t values, giving a step function, or by plotting only the points $(a_j, 1-p_j)$ for $j=1, 2, \ldots, k$ where

$$p_j = 1 - \tfrac{1}{2}\{\hat{S}(a_j) + \hat{S}(a_j+0)\}.$$

Note that $\hat{S}(a_j+0) = \hat{S}(a_{j+1})$ for $j=1, 2, \ldots, k-1$. Similar remarks apply to graphical representation of $\hat{H}(t)$ and $\tilde{H}(t)$.

In the special case of data that are assumed to be Weibull distributed, a plot of the points $(\log a_j, \log\{-\log(1-p_j)\})$ for $j=1, 2, \ldots, k$ should be approximately linear if the Weibull model is appropriate. Similarly a plot of the points $(\log a_j, \Phi^{-1}(p_j))$ should be approximately linear if a lognormal model is appropriate. For both plots rough parameter estimates may be obtained as in section 2.9.

Example 2.3

In an experiment to gain information on the strength of a certain type of braided cord after weathering, the strengths of 48 pieces of cord that had been weathered for a specified length of time were investigated. The intention was to obtain the strengths of all 48 pieces of cord. However, seven pieces were damaged during the course of the experiment, thus yielding right-censored strength-values. The strengths of the remaining 41 pieces were satisfactorily observed. Table 2.1 shows the data in coded units from this experiment.

Table 2.1 Strengths in coded units of 48 pieces of weathered braided cord

Uncensored observations									
36.3	41.7	43.9	49.9	50.1	50.8	51.9	52.1	52.3	52.3
52.4	52.6	52.7	53.1	53.6	53.6	53.9	53.9	54.1	54.6
54.8	54.8	55.1	55.4	55.9	56.0	56.1	56.5	56.9	57.1
57.1	57.3	57.7	57.8	58.1	58.9	59.0	59.1	59.6	60.4
60.7									
Right-censored observations									
	26.8	29.6	33.4	35.0	40.0	41.9	42.5		

Two aspects were of particular interest. First, the experimenter felt that it was important that even after weathering the cord-strength should be above 53 in the coded units. Thus, an estimate of $S(53)$ is required. Secondly, previous experience with related strength-testing experiments indicated that a Weibull model might be appropriate here. Thus, a check on the adequacy of the Weibull model is needed, together with estimation of the Weibull parameters if appropriate.

These questions will be addressed further in Chapter 3, but for now we use the PL estimate for the data to obtain $\hat{S}(53)$ and to investigate the Weibull model graphically.

Table 2.2 shows the first few lines of the calculation of $\hat{S}$ and related quantities for illustration. We see that, from equation (2.20) and Table 2.2,

$$\hat{S}(53) = 0.6849.$$

Table 2.2 Sample calculations of $\hat{S}$ and related quantities for Example 2.3

j	a_j	n_j	d_j	$(n_j-d_j)/n_j$	$\hat{S}(a_j+0)$	$d_j/\{n_j(n_j-d_j)\}$
0	$-\infty$	48	0	1.0000	1.0000	0.0000
1	36.3	44	1	0.9773	0.9773	0.0005
2	41.7	42	1	0.9762	0.9540	0.0006
3	43.9	39	1	0.9744	0.9295	0.0007
4	49.9	38	1	0.9737	0.9051	0.0007
5	50.1	37	1	0.9730	0.8806	0.0008
6	50.8	36	1	0.9722	0.8562	0.0008
7	51.9	35	1	0.9714	0.8317	0.0008
8	52.1	34	1	0.9706	0.8072	0.0009
9	52.3	33	2	0.9394	0.7583	0.0020
10	52.4	31	1	0.9677	0.7338	0.0011
11	52.6	30	1	0.9667	0.7094	0.0011
12	52.7	29	1	0.9655	0.6849	0.0012
13	53.1	28	1	0.9643	0.6605	0.0013

The approximate standard error of $\hat{S}(53)$ is, from equation (2.21) and Table 2.2,

$$\text{se}\{\hat{S}(53)\} = 0.6849\{0.0112\}^{1/2} = 0.0725.$$

Thus an approximate 95% confidence interval for $S(53)$ is

$$0.6849 \pm 1.96 \times 0.0725.$$

That is,

$$(0.54, 0.83).$$

Note that the four smallest censored values have been in effect ignored in this analysis as they were censored before the first uncensored observation.

Figure 2.17 shows a plot of $(\log a_j, \log\{-\log(1-p_j)\})$. The bulk of this plot seems to be linear. However, the points corresponding to the three lowest strengths lie considerably above the apparent line. Whilst these points have great visual impact because they are somewhat isolated from the rest of the plot, they are also the least reliable since they correspond to the extreme lower tail where the data are somewhat sparse. In this particular example the extreme points are made even less reliable in view of the relatively large amount of censoring at low strength-levels. The overall lack of curvature in the plot apart from the three isolated points suggests that a Weibull model should not be ruled out. If we assume a Weibull model for the moment, fitting a line by eye to the plot (ignoring the isolated points) gives slope 18.5 and intercept -75. Thus, rough estimates of η and α are 18.5 and 57.6 respectively.

Figure 2.18 shows a plot of $(\log a_j, \Phi^{-1}(p_j))$. Here there is some slight indication of curvature in the plot, even after ignoring the three isolated points.

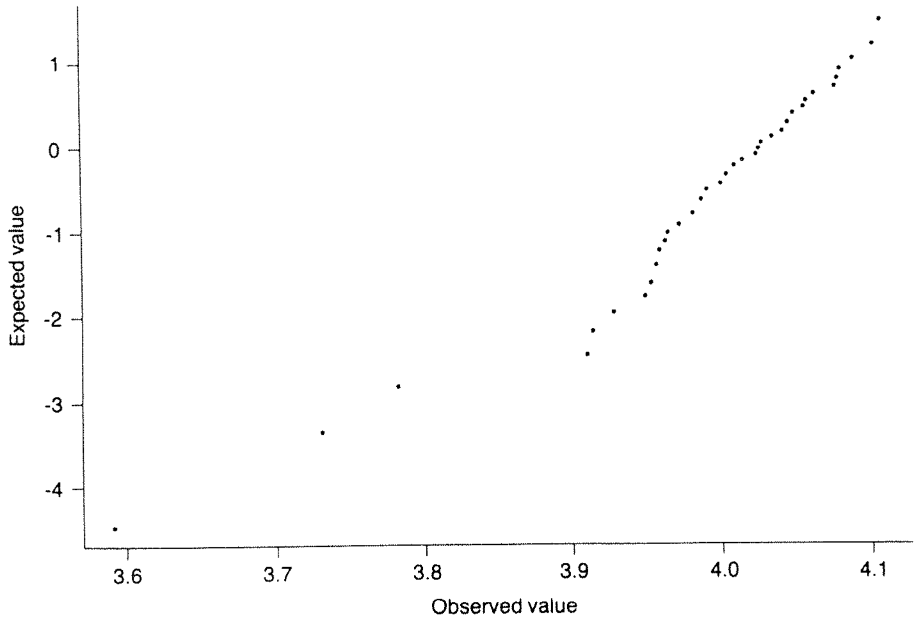

Figure 2.17 Weibull plot for the cord strength data.

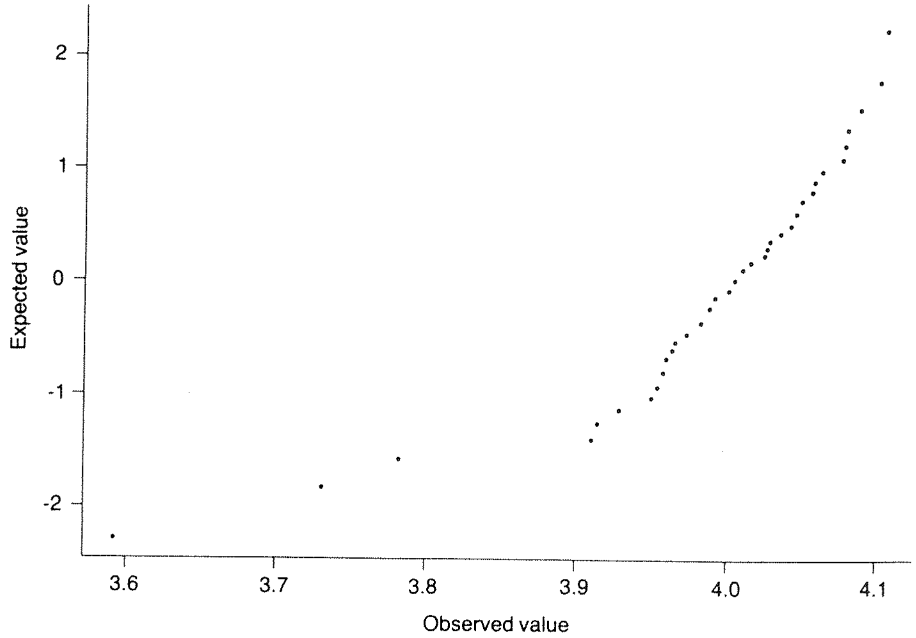

Figure 2.18 Lognormal plot for the cord strength data.

Hence a lognormal model appears to be somewhat less satisfactory than the Weibull model.

In this section we have discussed non-parametric estimation, that is estimation which does not require specification of a particular parametric model, of the survivor function and the cumulative hazard function in the presence of right-censored observations. A related approach for estimating quantiles, such as the median and the quartiles, in the presence of right-censoring is given in Kimber (1990).

When data contain left-censored or interval-censored observations, results analogous to the PL estimator of the survivor function are available. However, they are considerably more complex than the PL estimator and they have not been used much in the reliability context. One reason for this is the relative rarity of the occurrence of left-censored observations in reliability. In addition, unless the intervals are very wide, interval-censoring is generally relatively unimportant in practical terms. In fact all observations of continuous variables, such as time and strength, are interval-censored since data are only recorded to a finite number of decimal places. However, this aspect is usually ignored in analyses. For further information on non-parametric estimation of the survivor function in the presence of left-censored or interval-censored observations the reader is referred to Turnbull (1974, 1976).

3

Statistical methods for single samples

3.1 INTRODUCTION

Towards the end of Chapter 2 some simple statistical methods were introduced. These methods are such that they may be used before embarking on a more formal statistical analysis. In this chapter we first discuss methods for obtaining statistical inferences in the reliability context. In section 3.2 the method of maximum likelihood estimation is discussed in general terms. Some particular illustrations are given in section 3.3. Likelihood-based methods for hypothesis testing and confidence regions are then introduced in section 3.4. We then make some general remarks on likelihood-based methods in section 3.5. Finally we discuss in section 3.6 some methods that may be applied after fitting a parametric model, such as a Weibull distribution, in order to assess the adequacy of the fitted model.

3.2 MAXIMUM LIKELIHOOD ESTIMATION: GENERALITIES

In this section we give details of a general method of estimation of parameters, called maximum likelihood estimation (ML). To fix ideas suppose we have a sample of observations $t_1, t_2, \ldots, t_n$ from the population of interest. For the moment we assume that none of the observations is censored. In the reliability context it is reasonable to assume that the t_i are lifetimes. Suppose also that they can be regarded as observations with common density function $f(t; \theta_1, \theta_2, \ldots, \theta_m)$ where the form of f is known but where the parameters $\theta_1, \theta_2, \ldots, \theta_m$ are unknown. So, for example, we may perhaps assume that the observations are Weibull-distributed with unknown η and α. For brevity we shall denote dependence on $\theta_1, \theta_2, \ldots, \theta_m$ by θ, so that the common density may be written $f(t; \theta)$. Then the likelihood of the observations is defined by

$$L(\theta) = \prod_{i=1}^{n} f(t_i; \theta).$$

More generally, suppose that some of the observations are right-censored. Then we can split the observation numbers $1, 2, \ldots, n$ into two disjoint sets, one, U say, corresponding to observations that are uncensored, the other, C say, corresponding to right-censored observations. Then the likelihood in this case is defined by

$$L(\theta) = \left\{ \prod_{i \in U} f(t_i; \theta) \right\} \left\{ \prod_{i \in C} S(t_i; \theta) \right\}. \tag{3.1}$$

Thus, for a right-censored observation the density has been replaced by the survivor function. In a similar manner for a left-censored observation the density should be replaced by the distribution function. For an interval-censored observation the density should be replaced by the distribution function evaluated at the upper end-point of the interval minus the distribution function evaluated at the lower end-point of the interval, thus yielding the probability of occurrence of a lifetime within the interval.

It is almost always more convenient to work with the log-likelihood, $l(\theta)$ defined by

$$l(\theta) = \log L(\theta).$$

The maximum likelihood estimates (MLEs) $\hat\theta_1, \hat\theta_2, \ldots, \hat\theta_m$ of $\theta_1, \theta_2, \ldots, \theta_m$ are those values that maximize the likelihood, or equivalently, the log-likelihood. Alternatively, and more usually, the MLEs may be found by solving the likelihood equations

$$\frac{\partial l}{\partial \theta_j} = 0 \quad (j = 1, 2, \ldots, m).$$

Both approaches will usually involve numerical methods such as Newton or quasi-Newton algorithms. For most of the problems covered in this book the necessary numerical methods will be available in a package such as GLIM (see also Aitkin and Clayton, 1980), in subroutine libraries such as NAG and IMSL, and in the literature; see Press *et al.* (1986). In most simple situations (e.g. fitting a two parameter Weibull distribution) direct maximization of L or l will yield identical results to solving the likelihood equations. However, there exist situations where one or other of the two methods is unsatisfactory; see section 3.5.

Suppose that $\hat\theta = (\hat\theta_1, \hat\theta_2, \ldots, \hat\theta_m)$ has been calculated. We may be interested in some function of the unknown parameters such as

$$\phi = g(\theta)$$

where g is a specified one-to-one function. Then the MLE of ϕ is $\hat{\phi}$ defined by

$$\hat{\phi}=g(\hat{\theta}).$$

For example, one is often interested in estimating a quantile of the lifetime distribution; that is, estimating

$$q(p)\equiv q(p;\theta),$$

satisfying

$$\Pr\{T\geq q(p)\}=S\{q(p)\}=p,$$

where $0<p<1$ is specified. Hence the MLE of the quantile is just $q(p;\hat{\theta})$.

Furthermore, from asymptotic theory the precision of the MLEs may in many cases be estimated in a routine way. Consider the $m\times m$ observed information matrix $\mathbf{J}$ with entries

$$\frac{-\partial^2 l}{\partial\theta_j\partial\theta_k}\quad (j=1,2,\ldots,m;\ k=1,2,\ldots,m) \tag{3.2}$$

evaluated at $\hat{\theta}$. Then the inverse of $\mathbf{J}$ is the estimated variance-covariance matrix of $\hat{\theta}_1,\hat{\theta}_2,\ldots,\hat{\theta}_m$. That is, if $\mathbf{V}=\mathbf{J}^{-1}$ has entries v_{jk}, then v_{jk} is the estimated covariance between $\hat{\theta}_j$ and $\hat{\theta}_k$. In particular an estimate for the standard error of $\hat{\theta}_j$ $(j=1,2,\ldots,m)$ is just $v_{jj}{}^{1/2}$.

In addition, if $\phi=g(\theta)$, then the standard error of $\hat{\phi}$ may be estimated by

$$\mathrm{se}(\hat{\phi})=\left\{\sum_{j=1}^{m}\sum_{k=1}^{m}(\partial g/\partial\theta_j)(\partial g/\partial\theta_k)v_{jk}\right\}^{1/2}, \tag{3.3}$$

where the partial derivatives are evaluated at $\hat{\theta}$. This procedure is often referred to as the **delta method**. In the special case where $m=1$, so that θ is a scalar parameter, equation (3.3) reduces to

$$\mathrm{se}(\hat{\phi})\simeq\left|\frac{\mathrm{d}g}{\mathrm{d}\theta}\right|\sqrt{v_{11}}, \tag{3.4}$$

where $\mathrm{d}g/\mathrm{d}\theta$ is evaluated at $\hat{\theta}$. It is this equation (3.4) that was used in section 2.9 to obtain equation (2.16) from equation (2.14) with $g(\theta)=-\log\theta$. A more detailed discussion of the delta method is given in the Appendix at the end of this book.

Whilst construction of standard errors on the basis of equations (3.2), (3.3) and (3.4) is usually straightforward, the method does have certain drawbacks. These are discussed more fully in section 3.4.

Of course, ML is not the only estimation method available. We have already seen some ad hoc methods in sections 2.9 to 2.11. However, from the point of view of the user, ML has several major advantages. First its generality ensures that most statistical problems of estimation likely to arise in the reliability context may be dealt with using ML. Many other methods, such as those based on linear functions of order statistics (see David, 1981), are very simple to use in some univariate problems but are difficult or impossible to generalize to more complex situations. In addition, the generality of ML is an advantage from a computational point of view since, if desired, essentially the same program may be used to obtain MLEs whatever the context. Secondly, the functional invariance property of MLEs ensures that, having calculated $\hat{\theta}$, one may obtain the MLE of $g(\theta)$ immediately without having to restart the estimation process. Thirdly, approximate standard errors of MLEs may be found routinely by inversion of the observed information matrix.

From the theoretical point of view ML also has some properties to recommend it. Under mild regularity conditions MLEs are consistent, asymptotically Normal and asymptotically efficient. Roughly speaking these results mean that if the whole population is observed ML will give exactly the right answer, and that in large samples a MLE is approximately Normally distributed, approximately unbiased and with the smallest attainable variance. For technical details the reader should consult Cox and Hinkley (1974).

3.3 MAXIMUM LIKELIHOOD ESTIMATION: ILLUSTRATIONS

In this section we illustrate the calculation of MLEs in certain special cases. Throughout we shall assume that we have a single sample of observations, possibly right-censored, and that these observations are identically distributed. The case in which the parameters of interest depend on some explanatory or regressor variables will be covered in Chapters 4 and 5.

We assume that in the sample of possibly right-censored lifetimes $t_1,t_2,\ldots,t_n$, there are r uncensored observations and $n-r$ right-censored observations. We also define $x_i=\log t_i$ $(i=1,2,\ldots,n)$.

Exponential distribution

The log-likelihood is, from equations (2.2) and (3.1),

$$l(\lambda)=r\log\lambda-\lambda\sum_{i=1}^{n}t_i.$$

Hence

$$\frac{dl}{d\lambda}=\frac{r}{\lambda}-\sum_{i=1}^{n} t_i.$$

This may be set to zero and solved immediately to give

$$\hat{\lambda}=\frac{r}{\sum_{i=1}^{n} t_i}. \tag{3.5}$$

Notice that the denominator is the total time on test(TTT). Also,

$$\frac{-d^2 l}{d\lambda^2}=\frac{r}{\lambda^2}. \tag{3.6}$$

Hence the estimated standard error of $\hat{\lambda}$ is $\hat{\lambda}/\sqrt{r}$. Notice that it is necessary that $r>0$, that is, at least one lifetime must be uncensored.

As special cases, if all the observations are uncensored then $\hat{\lambda}$ is just the reciprocal of the sample mean, whereas if only the r smallest lifetimes $t_{(1)}<t_{(2)}<\cdots<t_{(r)}$ have been observed (simple Type II censoring) then

$$\hat{\lambda}=r\bigg/\left\{\sum_{i=1}^{r} t_{(i)}+(n-r)t_{(r)}\right\}. \tag{3.7}$$

The reciprocal of the right-hand side in (3.7) is sometimes known as the one-sided Winsorized mean.

Weibull distribution

The log-likelihood for a Weibull sample is, from equations (2.3), (2.4) and (3.1)

$$l(\eta,\alpha)=r\log\eta-r\eta\log\alpha+(\eta-1)\sum_u \log t_i-\alpha^{-\eta}\sum_{i=1}^{n} t_i^{\eta}.$$

Alternatively, letting $x_i=\log t_i$ and using a Gumbel formulation of the problem, we obtain from equations (2.6), (2.7) and (3.1)

$$l(\mu,\sigma)=-r\log\sigma+\sum_u (x_i/\sigma)-(r\mu/\sigma)-\sum_{i=1}^{n}\exp\{(x_i-\mu)/\sigma\}.$$

Thus

$$\sigma\frac{\partial l}{\partial\mu}=-r+\sum_{i=1}^{n}\exp\{(x_i-\mu)/\sigma\}$$

$$\sigma^2\frac{\partial l}{\partial\sigma}=-r\sigma-\sum_u x_i+r\mu+\sum_{i=1}^{n}\exp\{(x_i-\mu)/\sigma\}(x_i-\mu).$$

Hence

$$\hat{\mu}=\hat{\sigma}\log\left\{\frac{1}{r}\sum_{i=1}^{n}\exp(x_i/\hat{\sigma})\right\} \tag{3.8}$$

and

$$\frac{1}{r}\sum_u x_i+\hat{\sigma}-\sum_{i=1}^{n} x_i\exp(x_i/\hat{\sigma})\bigg/\sum_{i=1}^{n}\exp(x_i/\hat{\sigma})=0 \tag{3.9}$$

Note that equation (3.9) does not involve $\hat{\mu}$. So the problem of obtaining $\hat{\mu}$ and $\hat{\sigma}$ reduces simply to finding $\hat{\sigma}$, after which $\hat{\mu}$ may be found directly from equation (3.8). The solution to equation (3.9) must be found numerically. Routines for solving such non-linear equations are readily available in subroutine libraries such as NAG and IMSL. See also Press *et al.* (1986). A further possibility is to find two values of σ by trial and error which give opposite signs to the left side of equation (3.9). These may be used as starting values in a repeated bisection scheme. This can easily be programmed on even a very small computer.

The second derivatives of l are

$$\frac{-\partial^2 l}{\partial\mu^2}=\frac{1}{\sigma^2}\exp(-\mu/\sigma)\sum_{i=1}^{n}\exp(x_i/\sigma)$$

$$\frac{-\partial^2 l}{\partial\mu\partial\sigma}=\frac{-r}{\sigma^2}+\frac{1}{\sigma^3}\exp(-\mu/\sigma)(\sigma-\mu)\sum_{i=1}^{n}\exp(x_i/\sigma)+\frac{1}{\sigma^3}\exp(-\mu/\sigma)\sum_{i=1}^{n} x_i\exp(x_i/\sigma)$$

$$\frac{-\partial^2 l}{\partial\sigma^2}=\frac{-r}{\sigma^2}-2\sum_u\frac{x_i}{\sigma^3}+\frac{2r\mu}{\sigma^3}+\frac{1}{\sigma^4}\sum_{i=1}^{n}\exp(x_i/\sigma)(x_i-\mu)\{2\sigma+x_i-\mu\}.$$

These expressions simplify considerably when they are evaluated at $(\mu,\sigma)=$

$(\hat{\mu}, \hat{\sigma})$ to give

$$\frac{-\partial^2 l}{\partial \mu^2} = \frac{r}{\hat{\sigma}^2}$$

$$\frac{-\partial^2 l}{\partial \mu \partial \sigma} = \sum_{i=1}^{n} \left(\frac{x_i - \hat{\mu}}{\hat{\sigma}}\right) \exp\left\{\left(\frac{x_i - \hat{\mu}}{\hat{\sigma}}\right)\right\}$$

$$\frac{-\partial^2 l}{\partial \sigma^2} = r + \sum_{i=1}^{n} \left(\frac{x_i - \hat{\mu}}{\hat{\sigma}}\right)^2 \exp\left\{\left(\frac{x_i - \hat{\mu}}{\hat{\sigma}}\right)\right\}.$$

Example 3.1

Consider again the data first encountered in Example 2.2. For illustration purposes we now fit an exponential model using ML. Here $n = 13$, $r = 10$ and

$$\sum_{i=1}^{n} t_i = 23.05$$

Hence, using equations (3.5) and (3.6) for the exponential model the MLE for λ is $\hat{\lambda} = 10/23.05 = 0.434$, with standard error $\hat{\lambda}/\sqrt{r} = 0.137$. Figure 3.1 shows a plot of the log-likelihood as a function of λ. Note that there is a single maximum and that the function is rather skewed.

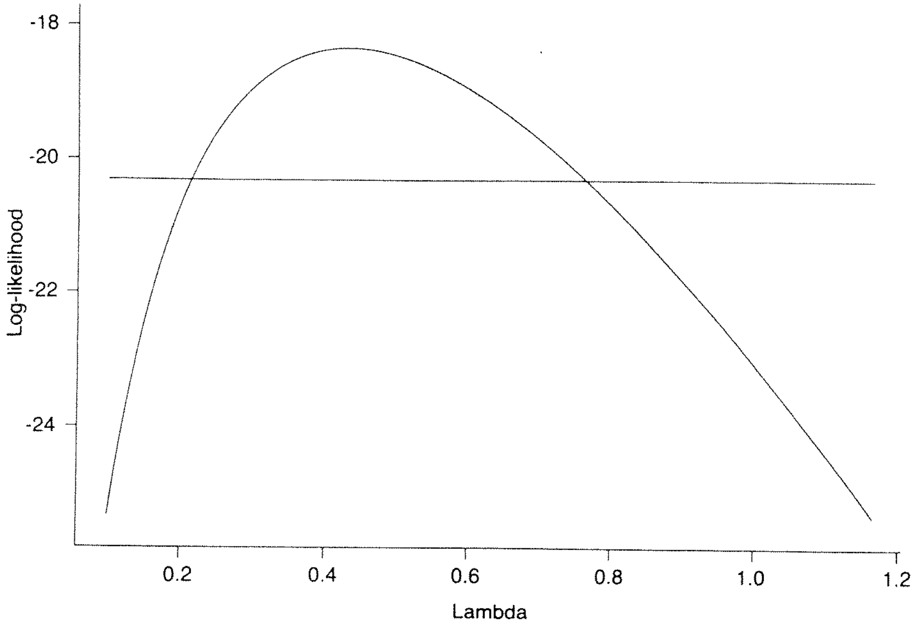

Figure 3.1 Log-likelihood for the aircraft components data with an exponential model. The line $l(\hat{\lambda})$-1.92 has been marked to show the 95 per cent confidence interval for λ based on W.

Example 3.2

Consider again the data first encountered in Example 2.1. We shall initially be concerned with fitting a parametric model to these data using for illustration first the Weibull distribution and then the lognormal distribution. For convenience we shall work with the logged data which amounts to fitting a Gumbel distribution and a Normal distribution.

For the Gumbel distribution we used a simple iterative scheme via the NAG library to obtain $\hat{\mu} = 4.405$, $\hat{\sigma} = 0.476$ with estimated variance-covariance matrix

$$\mathbf{V}_{\text{Gumbel}} = \begin{pmatrix} 0.01104 & -0.00257 \\ -0.00257 & 0.00554 \end{pmatrix}.$$

For the Normal distribution $\hat{\mu}$ and $\hat{\sigma}$ are respectively the sample mean and sample standard deviation (with divisor n rather than $n-1$), giving $\hat{\mu} = 4.150$ and $\hat{\sigma} = 0.522$ and

$$\mathbf{V}_{\text{Normal}} = \begin{pmatrix} 0.01184 & 0 \\ 0 & 0.00592 \end{pmatrix}.$$

Note that the parameters μ and σ are used here as generic notations for location and scale parameters for log-lifetimes. There is no reason why, say, μ in the Gumbel formulation should be equal to μ in the Normal formulation.

Suppose that various quantiles are of interest: the median, the lower 10% point and the lower 1% point. For both the distributions fitted to the log-lifetimes a quantile $q(p)$ is of the form $\mu + \sigma a(p)$, where $a(p)$ is readily available in standard statistical tables in the Normal case and where $a(p) = \log(-\log p)$ in the Gumbel case. Table 3.1 gives the numerical values of $a(p)$ for these two distributions for the required values of p, namely 0.5, 0.9 and 0.99. Using the given values of $a(p)$ and the calculated MLEs we can obtain the estimated quantiles for the log-lifetimes.

Table 3.1 Values of $a(p)$ in the Normal and Gumbel cases with $p = 0.5$, 0.9, 0.99

p	Gumbel $a(p)$	Normal $a(p)$
0.5	−0.367	0
0.9	−2.250	−1.282
0.99	−4.600	−2.326

These may be transformed back to the original scale of measurement by exponentiating. For example, the estimated median of the log-lifetimes in the Gumbel case is given by $4.405 - 0.367 \times 0.476 = 4.230$. Thus, in the Gumbel case the estimated median lifetime is exp(4.320) = 68.7 million revolutions. Table 3.2 shows the estimated quantiles for the lifetimes using the two parametric formulations, together with standard errors using equation (3.3) with $g(\mu, \sigma) = \exp\{\mu + \sigma a(p)\}$.

Table 3.2 Quantile estimates (in millions of revolutions) for Weibull and lognormal models, together with their standard errors

Quantile	Weibull estimate	Lognormal estimate
Median	68.7 (8.0)	63.4 (6.9)
Lower 10%	28.1 (6.3)	32.5 (4.8)
Lower 1%	9.2 (3.6)	18.1 (3.9)

Examination of Table 3.2 shows that, especially in the case of the most extreme quantile, the two parametric models appear to give very different estimates. This situation is not unusual.

3.4 TESTS AND CONFIDENCE REGIONS BASED ON LIKELIHOOD

In Example 3.2 above we have seen that fitting different parametric models may give very different estimates of a quantity of interest, such as a quantile. Thus, some methods for choosing between parametric models are clearly required. One approach is based on asymptotic properties of the likelihood function. These properties also enable confidence intervals, or more generally, confidence regions, to be calculated. First, we shall state the main general results. Then we shall give some simple illustrations using the data already discussed above.

We begin by supposing that the parametric model of interest depends on parameters $\theta_1, \theta_2, \ldots, \theta_m$. Suppose we are interested in testing or making confidence statements about a subset $\theta^{(A)}$ of these parameters. Label the remaining parameters $\theta^{(B)}$. Of course, $\theta^{(A)}$ may contain all of the m parameters, so that $\theta^{(B)}$ is empty. Let $(\hat{\theta}^{(A)}, \hat{\theta}^{(B)})$ be the joint MLE of $(\theta^{(A)}, \theta^{(B)})$. Let $\hat{\theta}^{(B)}(A_0)$ be the MLE of $\theta^{(B)}$ when $\theta^{(A)}$ is fixed at some chosen value $\theta_0^{(A)}$, say. Then two likelihood-based methods for testing and constructing confidence regions are as follows:

1. Let

$$W(\theta_0^{(A)}) = W = 2\{l(\hat{\theta}^{(A)}, \hat{\theta}^{(B)}) - l[\theta_0^{(A)}, \hat{\theta}^{(B)}(A_0)]\}.$$

Then under the null hypothesis $\theta^{(A)} = \theta_0^{(A)}$, W has approximately a chi-squared distribution with m_a degrees of freedom, where m_a is the dimension of $\theta^{(A)}$. Large values of W relative to $\chi^2(m_a)$ supply evidence against the null hypothesis. The corresponding $1-\alpha$ confidence region for $\theta^{(A)}$ is

$$\{\theta^{(A)}: W(\theta^{(A)}) \leq \chi_\alpha^2(m_a)\},$$

where $\chi_\alpha^2(m_a)$ is the upper 100α percentage point of $\chi^2(m_a)$.

2. Suppose $\mathbf{V} = \mathbf{V}(\hat{\theta}^{(A)}, \hat{\theta}^{(B)})$ is the variance-covariance matrix for $(\hat{\theta}^{(A)}, \hat{\theta}^{(B)})$ evaluated at the MLE, as in section 3.3. Let $\mathbf{V}_A = \mathbf{V}_A(\hat{\theta}^{(A)}, \hat{\theta}^{(B)})$ be the leading submatrix of $\mathbf{V}$ corresponding to $\hat{\theta}^{(A)}$. That is, $\mathbf{V}_A$ is the submatrix of $\mathbf{V}$ corresponding to the estimated variance and covariance of $\hat{\theta}^{(A)}$. Then

$$W^*(\theta_0^{(A)}) = (\hat{\theta}^{(A)} - \theta_0^{(A)})^T \mathbf{V}_A^{-1} (\hat{\theta}^{(A)} - \theta_0^{(A)})$$

also has an approximate $\chi^2(m_a)$ distribution under the null hypothesis $\theta^{(A)} = \theta_0^{(A)}$. The corresponding approximate $1-\alpha$ confidence region for $\theta^{(A)}$ is given by

$$\{\theta^{(A)}: W^*(\theta^{(A)}) \leq \chi_\alpha^2(m_a)\}.$$

In the special case when $\hat{\theta}^{(A)}$ is a scalar this leads to a symmetric $1-\alpha$ confidence interval

$$\hat{\theta}^{(A)} \pm z_{\alpha/2} V_A^{1/2}$$

where $z_{\alpha/2}$ is the upper $100\alpha/2$ percentage point of the standard Normal distribution.

Both methods based on W and W^* respectively are asymptotically equivalent, and often give very similar results in practice. However, large discrepancies are possible. In such cases the method based on W is preferable because the results are invariant to reparametrization and the shape of the confidence region is essentially decided by the data. Confidence regions based on W^* are necessarily elliptical in the parametrization used but will yield non-elliptical regions under non-linear parameter transformations.

Example 3.1 (continued)

Having fitted an exponential model to the Mann and Fertig data, we shall now calculate a 95% confidence interval for λ. In the notation we used in the general case $\theta^{(A)} = \lambda$, $\theta^{(B)}$ is redundant, and $m_a = 1$. Here

$$l(\lambda) = 10 \log \lambda - 23.05\lambda$$

Example 2.1

The following data from Lieblein and Zelen (1956) are the numbers of millions of revolutions to failure for each of 23 ball bearings. The original data have been put in numerical order for convenience.

17.88	28.92	33.00	41.52	42.12	45.60	48.40	51.84
51.96	54.12	55.56	67.80	68.64	68.64	68.88	84.12
93.12	98.64	105.12	105.84	127.92	128.04	173.40	

The data are clearly positively skewed. The sample mean and standard deviation $\bar{t}$ and s_t are given by

$$\bar{t} = \sum_{i=1}^{23} t_i/23 = 72.22$$

and

$$s_t = \left\{ \sum_{i=1}^{23} (t_i - \bar{t})^2/22 \right\}^{1/2} = 37.49.$$

$$\frac{\bar{t}}{s_t} \approx 2$$

Tilpass Weibull $\quad R(t) = e^{-(\lambda t)^{\alpha}}$

SME: $\left.\begin{array}{l} \lambda^* = 0.0122 \\ \alpha^* = 2.100 \end{array}\right\} \Rightarrow MTTF^* = \frac{\Gamma(\frac{1}{\alpha^*}+1)}{\lambda^*} = \underline{72.6}$

Test $\underline{H_0: \alpha = 1}$ mot $H_1: \alpha \neq 1$

Kan vise

$\ell(\alpha^*, \lambda^*) = \ell(2.1, 0.0122) = \underline{-18.24}$

Når $\alpha = 1$ blir SME for λ lik $\hat{\lambda} = \frac{1}{\bar{T}} = 0.0138$

$\ell(\alpha_0, \hat{\lambda}) = \ell(1, 0.0138) = \underline{-25.99}$

Gir

$W(\alpha_0 = 1) = 2[\ell(\alpha^*, \lambda^*) - \ell(1, \hat{\lambda})]$

$= 2[-18.24 + 25.99] = \underline{15.50} > z_{0.05,1}$

Forkast H_0, og påstå at vi ikke har eksponensial fordeling and

$$l(\hat{\lambda}) = -18.35.$$

Hence

$$W(\lambda) = 2\{-18.35 - 10 \log \lambda + 23.05\lambda\}.$$

Thus, a 95% confidence interval for λ based on W is $\{\lambda: W(\lambda) \le 3.84\}$ since 3.84 is the upper 5% point of $\chi^2(1)$. In other words the required confidence interval consists of all those values of λ such that $l(\lambda)$ is within 1.92 of $l(\hat{\lambda})$. This interval is marked on Figure 3.1 and corresponds to [0.22, 0.76]. Note that because of the skewness of the log-likelihood, this interval is not centred at the MLE, 0.434.

A second method of calculating a 95% confidence interval for λ is to use W^*. This gives limits $0.434 \pm 1.96 \times 0.137$, where 1.96 is the upper 2.5% point of the standard Normal distribution and 0.137 is the standard error of $\hat{\lambda}$. Thus, the interval based on W^* is [0.17, 0.70]. To show the dependence of W^* on the particular parametrization used, consider first reparametrizing the exponential distribution in terms of $\alpha = 1/\lambda$. Then the MLE for α is $\hat{\alpha} = 1/\hat{\lambda} = 1/0.434 = 2.305$. Further, the standard error of $\hat{\alpha}$ may be found using equation (3.4), giving

$$\mathrm{se}(\hat{\alpha}) = \frac{\hat{\alpha}}{\sqrt{r}} = 0.729.$$

Hence, using W^*, a 95% confidence interval for α is $2.305 \pm 1.96 \times 0.729$; that is, [0.88, 3.73]. Hence the corresponding confidence interval for λ is [1/3.73, 1/0.88] i.e. [0.27, 1.14].

Thus we have three quite different confidence intervals for λ. As remarked above, the interval based on W is generally to be preferred. In the special case of the exponential distribution some exact distribution theory for the MLE is available, which gives an exact 95% confidence interval for λ as [0.21, 0.74], which is in good agreement with the w-based interval. The exact interval is based upon the fact that $2r\lambda/\hat{\lambda}$ has a $\chi^2(2r)$ distribution. Note also that, had a higher degree of confidence been required, the W^*-based intervals would have contained negative values, clearly a nonsense since λ is necessarily positive. The occurrence of negative values in a confidence interval for a necessarily positive parameter is not possible when using the W-based method.

Example 3.2 (continued)

Earlier we fitted a Weibull model to the Lieblein and Zelen data. Could we simplify the model and assume exponentiality? Equivalently, working with the log-lifetimes, could we fix $\sigma = 1$ in the Gumbel model? We can take $\theta^{(A)} = \sigma$, $\theta^{(B)} = \mu$

and $m_a = 1$ in the general framework and test the null hypothesis $\sigma = 1$.

We begin by considering the testing procedure based on W^*, results for which are virtually immediate since we have already estimated the variance-covariance matrix of $(\hat{\mu}, \hat{\sigma})$ as $\mathbf{V}_{\text{Gumbel}}$. Here

$$W^*(1) = (0.476 - 1)^2/0.00554 = 49.56.$$

This is highly significant as $\chi^2(1)$, indicating that an exponential model is inappropriate. The result is so clear cut that it is probably not worth using the more accurate test based on w in this particular numerical example. However, we shall do so in the interest of illustrating the method. Let $\hat{\mu}_0$ be the MLE of μ when $\sigma = 1$ is fixed. Then

$$l(\mu, 1) = \sum_u x_i - r\mu - \sum_{i=1}^{n} \exp(x_i - \mu)$$

$$\frac{dl}{d\mu} = -r + \sum_{i=1}^{n} \exp(x_i - \mu).$$

Hence

$$\hat{\mu}_0 = \log\left\{\frac{1}{r} \sum_{i=1}^{n} \exp(x_i)\right\}.$$

In this example, $\hat{\mu}_0 = 4.280$, so that

$$l(\hat{\mu}_0, 1) = \sum_u x_i - r\hat{\mu}_0 - r = -25.99.$$

Also, $l(\hat{\mu}, \hat{\sigma}) = -18.24$, giving

$$W(1) = 2\{l(\hat{\mu}, \hat{\sigma}) - l(\hat{\mu}_0, 1)\} = 15.50.$$

Again, this is highly significant as $\chi^2(1)$, leading to the same inference as that based on W^*, though the result is not so extreme. These results confirm the earlier informal analysis discussed in Example 2.1.

Example 3.2 (continued) gives an illustration of testing nested hypotheses: the model specified in the null hypothesis is a special case of the model specified in the alternative. In the example the exponential model is a special case of the Weibull model. Other examples include testing:

1. an exponential model against a gamma model;
2. a Weibull model against a generalized gamma model; and
3. an exponential model with specified λ against an exponential model with unspecified λ.

This situation also arises commonly in regression analysis (see Chapters 4, 5, and 6) in which the effect of an explanatory variable may be assessed by testing that its coefficient is zero against the alternative that its coefficient is non-zero.

What of the situation in which the hypotheses to be tested are non-nested? In the context of this chapter, where we are concerned with methods for univariate data, such a situation might arise in attempting to choose between two separate families of distributions. For example, we might wish to choose between a Weibull model and a lognormal model. Unfortunately the distribution theory of section 3.4 no longer holds in general in this situation. Thus, one is faced with the choice of either circumventing the problem or using theory specifically designed for this particular problem. To fix ideas we stay with the choice between Weibull and lognormal models.

An informal approach is to use data analytic methods such as those outlined in sections 2.9 and 2.11. Another approach is to use goodness-of-fit tests for the Weibull distribution and for the lognormal distribution. A major reference in this field is D'Agostino and Stephens (1986). Yet another method is to fit a more comprehensive model which contains the competing models. Possibilities include the generalized gamma model and a mixture model, both discussed in section 2.7. The hypotheses Weibull versus general alternative and lognormal versus general alternative can then be tested. If only one of the competing models can be rejected, then the other one is preferable. However, it is perfectly possible that neither model can be rejected.

On the other hand we can face up squarely to the problem of testing separate families of distributions. First, it must be decided which distribution should be the null hypothesis. For example, if in the past similar data sets have been well fitted by a Weibull model, then it is natural to take a Weibull distribution as the null model. Pioneering work on this topic appears in Cox (1961, 1962b) and Atkinson (1970). In the reliability context, see Dumonceaux and Antle (1973) and Dumonceaux, Antle and Haas (1973). Typically, however, these tests, which are based on maximized likelihoods, have rather low power unless the sample size is large. When there is no natural null hypothesis a pragmatic approach is to test Weibull versus lognormal and lognormal versus Weibull. Once again, it is possible that neither distribution may be rejected in favour of the other.

A different approach is to force a choice between competing models; see Siswadi and Quesenberry (1982). This procedure, however, is not recommended, especially if estimation of extreme quantiles is the aim of the analysis. To force a choice between two or more models, which all fit a set of data about equally well in some overall sense, can lead to spuriously precise inferences about extreme quantiles (see Examples 2.9 and 3.2).

Example 3.2 (continued[2])

We have seen in Example 2.1 that on the basis of simple plots both the Weibull and lognormal models appear to fit the data well. However, in Example 3.2 we have seen that estimates of extreme quantiles for the two models are very different.

The maximized log-likelihoods are, working with the log-lifetimes,

Weibull	−18.24
lognormal	−17.27.

Thus, on the basis of maximized log-likelihood the lognormal model, as with the plots, appears slightly better than the Weibull model. However, using the results of Dumonceaux and Antle (1973), we cannot reject the Weibull model in favour of the lognormal model or vice versa (using a 5% significance level for each test). So we are still unable to choose between the two models on statistical grounds.

If, for example, the median is of major concern, the choice between the models is relatively unimportant. Alternatively, a non-parametric estimate could be used; see sections 2.9 and 2.11 and Kimber (1990). However, if the lower 1% quantile is of prime interest the non-parametric estimate is not an option. If a choice must be made between models, then the more pessimistic one is probably preferable (though this depends on the context of the analysis). However, the real message is that the data alone are insufficient to make satisfactory inferences about low quantiles. If quantifiable information is available in addition to the sample data, then a Bayesian approach (see Chapter 6) may be fruitful. Another way round the problem is to collect more data so that satisfactory inferences about low quantiles may be drawn. The approach of Smith and Weissman (1985) to estimation of the lower tail of a distribution using only the smallest k ordered observations is another possibility, thus side-stepping the problem of a fully parametric analysis.

3.5 REMARKS ON LIKELIHOOD-BASED METHODS

Likelihood-based procedures have been discussed above in relation to parametric analyses of a single sample of data. These results generalize in a natural way for more complex situations, such as when information on explanatory variables is available as well as on lifetimes. These aspects are covered in later chapters.

The asymptotic theory which yields the Normal and χ^2 approximations used above requires certain conditions on the likelihood functions to be satisfied. The situation most relevant to reliability in which the regularity conditions on the likelihood do not hold occurs when a guarantee parameter must be estimated. This can cause some problems, which have been addressed by many authors including Smith and Naylor (1987) for the three-parameter Weibull distribution, Eastham *et al.* (1987) and LaRiccia and Kindermann (1983) for the three-parameter lognormal distribution, and Kappenman (1985) for three-parameter Weibull, lognormal and gamma distributions. Cheng and Amin (1983) also discuss problems with guarantee parameters.

For the relatively simple models discussed so far (e.g. exponential, Weibull, Gumbel, Normal, lognormal) the regularity conditions on the likelihood do hold. For some of these models (e.g. Normal and lognormal with no censoring) closed form MLEs exist. However, in most likelihood-based analyses of reliability data some iterative scheme is needed. Essentially one requires a general program to handle the relevant data, together with a function maximizing procedure, such as quasi-Newton methods in NAG or Press *et al.* (1986). To fit a specific model all that is necessary is to 'bolt on' a subroutine to evaluate the relevant log-likelihood function, and possibly the first and second derivatives of the log-likelihood. For some models some numerical integration or approximation may be needed (e.g. polygamma functions for the gamma distribution, Normal survivor function for the Normal distribution with censored observations). This may be programmed from the relevant mathematical results (Abramowitz and Stegun, 1972) via available algorithms (Press *et al.* 1986; Griffiths and Hill, 1985) or using subroutines available in libraries such as NAG and IMSL. In any event, it is worth trying several different initial values to start an iterative scheme in order to check the stability of the numerical results.

Of course, even when regularity conditions on the likelihood function are satisfied, the asymptotic results may give poor approximations in small samples or samples with heavy censoring. In cases of doubt one can either search for exact distribution theory or adopt a more pragmatic approach and use simulation to examine the distribution of any appropriate estimator or test statistic; see Morgan (1984).

Other references of interest are DiCiccio (1987) and Lawless (1980) for the generalized gamma distribution and Shier and Lawrence (1984) for robust estimation of the Weibull distribution. Cox and Oakes (1984, Chapter 3) give a good general discussion of parametric methods for survival data. Johnson and Kotz (1970) give details of a whole range of estimation methods for the basic distributions discussed here. The paper by Lawless (1983), together with the resulting discussion is also well worth reading.

3.6 GOODNESS-OF-FIT

As part of a statistical analysis which involves fitting a parametric model, it is always advisable to check on the adequacy of the model. One may use either

a formal goodness-of-fit test or appropriate data analytic methods. Graphical procedures are particularly valuable in this context.

Before discussing graphical methods for model checking, we mention formal goodness-of-fit tests. One approach is to embed the proposed model in a more comprehensive model. Tests such as those outlined in section 3.4 may then be applied. For example, one might test the adequacy of a exponential model relative to a Weibull model, as in Example 3.2 (continued). In contrast, the proposed model may be tested against a general alternative. This is the classical goodness-of-fit approach. An example is the well-known Pearson χ^2-test. The literature on goodness-of-fit tests is vast, though the tendency has been to concentrate on distributions such as the Normal, exponential and Weibull. A key reference is D'Agostino and Stephens (1986), but see also Lawless (1982, Chapter 9). Within the reliability context, however, tests based on the methods of section 3.4, combined with appropriate graphical methods will almost always be adequate.

In order to discuss graphical methods for checking model adequacy we shall use the notation introduced in section 2.11. That is, we suppose there are k distinct times $a_1 < a_2 < \cdots < a_k$ at which failures occur. Let d_j be the number of failures at time a_j and let n_j be the number of items at risk at a_j. In addition we shall use the plotting positions

$$p_j = 1 - \tfrac{1}{2}\{\hat{S}(a_j) + \hat{S}(a_j + 0)\},$$

where $\hat{S}$ is the PL estimator as in equation (2.20). Kaplan-Meier

In section 2.11 we introduced a plotting procedure for models that depend only on location and scale parameters, μ and σ say, such as Gumbel and Normal distributions. In reliability this amounts to plotting the points

$$(\log a_j, F_0^{-1}(p_j)),$$

where F_0 is the distribution function of the proposed model with μ and σ set to 0 and 1 respectively. If the model is appropriate, the plot should be roughly linear. This type of plot, called the quantile–quantile (QQ) plot can be applied before a formal statistical analysis is attempted. In addition, rough parameter estimates for the proposed model may be obtained from the slope and intercept of the plot. These estimates may be of interest in their own right or may be used as starting values in an iterative scheme to obtain ML estimates. However, the applicability of this type of plot is limited. For example, it cannot be used for the gamma model. Moreover, the points on the plot which usually have the greatest visual impact, the extreme points, are those with the greatest variability.

A different graphical display, but which uses the same ingredients as the QQ plot, is the probability (PP) plot. This involves plotting the points

$$(p_j, F(a_j; \hat{\theta})),$$

where $F(a_j; \hat{\theta})$ denotes the distribution function of the proposed model, evaluated at the point a_j and with the parameters of the model set to 'reasonable' estimates (usually the ML estimates). Again, linearity of the plot is indicative of a good agreement between fitted model and data. Since estimates are used, the PP plot can only usually be constructed after fitting the model. However, its use is not limited to models with location and scale parameters only.

In the PP plot the extreme points have the lowest variability. A refinement which approximately stabilizes the variability of the plotted points is the stabilized probability (SP) plot; see Michael (1983). This involves plotting the points

$$\left(\frac{2}{\pi}\sin^{-1}\left(p_j^{1/2}\right), \frac{2}{\pi}\sin^{-1}\left\{F^{1/2}\left(a_j; \hat{\theta}\right)\right\}\right).$$

Alternatively, simulated envelopes may be used to aid interpretation of QQ and PP plots; see Atkinson (1985). The construction of formal goodness-of-fit tests which are based on plots has also been investigated by Michael (1983), Kimber (1985) and Coles (1989).

Examples 3.3 (Example 2.3 continued)

The Weibull QQ plot for the strengths of 48 pieces of cord is shown in Figure 2.17. The overall impression of this plot is that it is basically linear except for the three smallest observed strengths. We now investigate the PP and SP plots for these data for a Weibull model.

Maximum likelihood estimation for a Gumbel distribution applied to the log-strengths gives $\hat{\mu}=4.026$, and $\hat{\sigma}=0.0613$ with variance-covariance matrix for the parameter estimates $(\hat{\mu}, \hat{\sigma})$

$$10^{-4}\begin{pmatrix} 0.9985 & -0.2191 \\ -0.2191 & 0.5889 \end{pmatrix}. \tag{3.10}$$

The PP and SP plots are shown in Figures 3.2 and 3.3 respectively. The visual impact of the three smallest points is much less than in Fig. 2.17. Overall the satisfactory fit of the Weibull model is confirmed.

Returning to the original purpose of analyzing these data, we wish to estimate $S(53)$. In Example 2.3 we found that the PL estimate is 0.685 with a

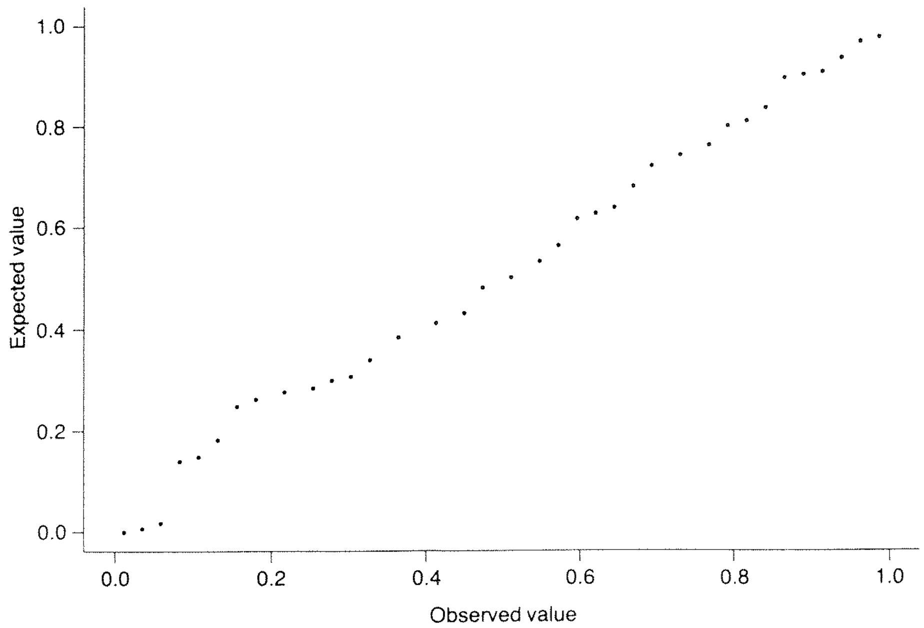

Figure 3.2 Weibull *PP* plot for the cord strength data.

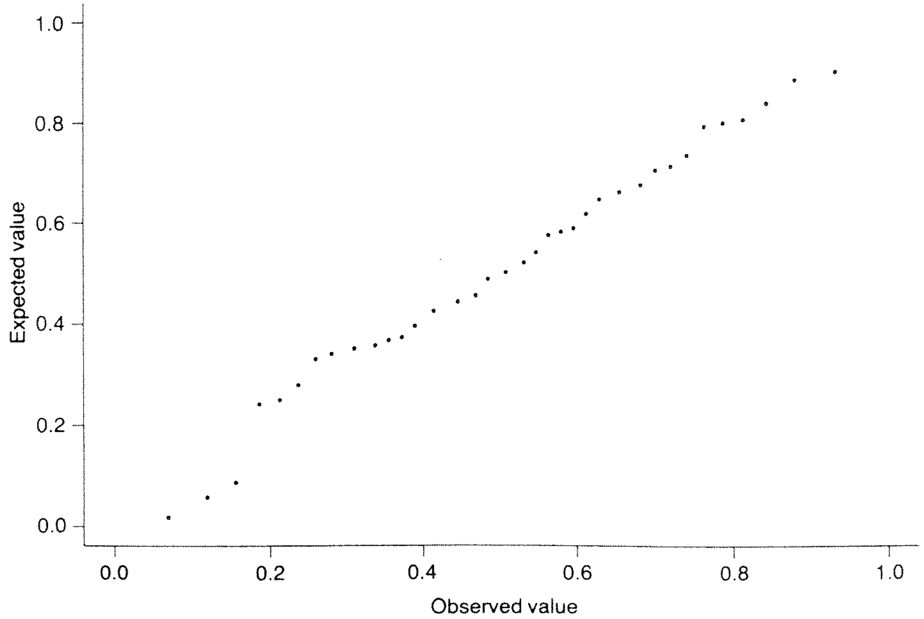

Figure 3.3 Weibull *SP* plot for the cord strength data.

standard error of 0.073. Using the Weibull model, the ML estimate of $S(53)$ is

$$\exp\left\{-\exp\left(\frac{\log 53-4.026}{0.0613}\right)\right\}=0.668.$$

Using equation (3.3) in conjunction with the above variance-covariance matrix in expression (3.10), the standard error of the ML estimate is 0.046. This leads to a somewhat shorter approximate 95% confidence interval, [0.58, 0.76], than that obtained using the PL estimate, which was [0.54, 0.83].

2.10 DATA ANALYTIC METHODS: TYPE II CENSORING

Consider the situation in which n units are put on test and observation continues until r units have failed. In other words we have Type II censoring: the first r lifetimes $t_{(1)}<t_{(2)}<\cdots<t_{(r)}$ are observed, but it is known only that the remaining $n-r$ lifetimes exceed $t_{(r)}$. Because $n-r$ observations are 'incomplete', it is impossible to calculate the sample moments. So standard moment-based methods cannot be used. However, all the results based on the empirical survivor function as discussed in section 2.9 still hold. The only difference is that $\hat{S}(t)$ is not defined for $t>t_{(r)}$.

Example 2.2
Mann and Fertig (1973) give details of a life test done on thirteen aircraft components subject to Type I censoring after the tenth failure. The ordered data in hours to failure time are:

0.22	0.50	0.88	1.00	1.32
1.33	1.54	1.76	2.50	3.00