Chapter 13
Additional Topics

This chapter begins by introducing Bayesian inference for extreme value processes, such as might be used to model high winds and flooding. It then gives an overview of the Bayesian treatment of expert opinion, and then proceeds to an example pointing out the pitfalls that can be encountered if *ad hoc* methods are employed. We next illustrate how to encode prior distributions into OpenBUGS that are not included as predefined distribution choices. We close this chapter with an example of Bayesian inference for a time-dependent Markov model of pipe rupture.

13.1 Extreme Value Processes

This section gives a brief introduction to statistical models for extreme quantities, such as flooding levels, which are encountered in external events PRA. Much of this material is adapted from [1], which gives an excellent overview of the subject from a mathematical level that should be comfortable to those with a degree in physical science or engineering, although with the exception of a quick overview of Bayesian inference in the last chapter, the treatment of inference in [1] is from a frequentist perspective.

As an example of the type of problem dealt with in this section, consider the following data on annual maximum sea levels, taken from [1]. From these data, one might wish to be able to project, with uncertainty, the maximum level for the next 100 or 500 years. A characteristic feature of this type of problem is the desire to extrapolate beyond the range of observed data.

The extrapolation is based on the so-called *extreme value paradigm*. It works by using asymptotic statistical models to make predictions. For the example above, if X_1, X_2, … are the sequence of daily maximum river levels, then we are interested in the maximum level for an observation period of n days:

D. Kelly and C. Smith, *Bayesian Inference for Probabilistic Risk Assessment*,
Springer Series in Reliability Engineering, DOI: 10.1007/978-1-84996-187-5_13,

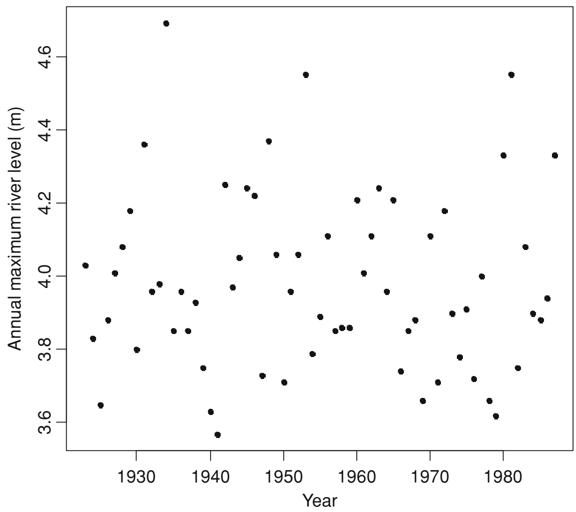

Fig. 13.1 Plot of annual sea level in meters, from [1]

$$M_n = \max\{X_1, \ldots, X_n\} \tag{13.1}$$

If the exact distribution of each daily maximum observation (X_i) were known, then we could find the distribution of M_n. It is given by the following, under the (perhaps dubious) assumption that the X_is are mutually statistically independent:

$$\begin{aligned}\Pr(M_n \le z) &= \Pr(X_1 \le z, X_2 \le z, \ldots, X_n \le z)\\ &= [F(z)]^n\end{aligned}$$

In practice, we do not know the distribution of each X_i. We could estimate this distribution, and find the distribution of M_n using the equation above, but small errors in the estimate of F lead to large errors in F^n. However, we can find the approximate distribution of M_n for large values of n, under certain assumptions. We let $n \to \infty$, and this leads to a family of *extreme value distributions*, whose parameters can be estimated from the observed data.

One can object to this procedure on the basis that the extrapolation to unseen levels is faith-based; such a criticism is easy to make, and there is no real defense against it, except to say that extrapolation is required, and using a method with some rationale (asymptotic theory) is better than any existing alternatives. One must of course be careful about gross violations of the underlying assumptions. For example, what follows is based on the assumption that climate changes do not cause a systematic increase or decrease in the maximum annual river levels. Based on the plot in Fig. 13.1, there does not appear to be any evidence of a change in the pattern of variation over the observation period, but this is no guarantee for the future.

If you look at older references such as [2], you will find a somewhat confusing set of limiting distributions, with names like the Gumbel distribution. Modern references have unified all of these into a single family of limiting distributions, referred to as the *generalized extreme value* (GEV) family of distributions.

For technical reasons, we cannot just look at the limiting behavior of M_n as $n \to \infty$. Instead, we look at $(M_n - a_n)/b_n$, where $\{a_n > 0\}$ and $\{b_n\}$ are sequences of constants that stabilize the location and scale of M_n as n increases, keeping the limiting distribution from becoming degenerate.

The big theorem is that $\Pr[(M_n - a_n)/b_n \leq z] \to G(z)$ as $n \to \infty$, where $G(z)$ has the following form:

$$G(z) = \exp\left\{-\left[1 + \xi\left(\frac{z-\mu}{\sigma}\right)\right]^{-1/\xi}\right\} \tag{13.2}$$

$G(z)$ is defined for $1 + \xi(z - \mu)/\sigma > 0$. The parameters of the GEV distribution satisfy $-\infty < \mu < \infty$, $\sigma > 0$, and $-\infty < \xi < \infty$. The parameter ξ is a shape parameter and determines the tail behavior of $G(z)$.

Quantiles of $G(z)$ are often used, and are given by, for $G(z_p) = 1 - p$

$$z_p = \begin{cases} \mu - \dfrac{\sigma}{\xi}\left\{1 - [-\log(1-p)]^{-\xi}\right\},\ \xi \neq 0 \\ \mu - \sigma \log[-\log(1-p)],\ \xi = 0 \end{cases} \tag{13.3}$$

The quantity z_p is often called the *return level* associated with the *return period* $1/p$; z_p is exceeded, on average, once every $1/p$ years (if we're measuring time in years). Put another way, z_p is exceeded by the annual maximum in any particular year with probability p.

If we define $y_p = -\log(1 - p)$, we can rewrite the quantiles as

$$z_p = \begin{cases} \mu - \dfrac{\sigma}{\xi}\left(1 - y_p^{-\xi}\right),\ \xi \neq 0 \\ \mu - \sigma \log y_p,\ \xi = 0 \end{cases} \tag{13.4}$$

If z_p is plotted against $\log(y_p)$, or z_p is plotted against y_p on a logarithmic scale, the plot will be linear if $\xi = 0$. If $\xi < 0$, the plot approaches the limit $\mu - \sigma/\xi$ as $p \to 0$. If $\xi > 0$, there is no upper bound to z_p. Such a graph is called a *return level plot*.

13.1.1 Bayesian Inference for the GEV Parameters

We treat the annual maximum river level data in Fig. 13.1 as a random sample from $G(z)$ ($n = 365$ is close to ∞). We will use OpenBUGS to perform the Bayesian inference for the GEV parameters. We will use independent diffuse priors: $\mu \sim$ dflat(), $\sigma \sim$ dgamma(10^{-4}, 10^{-4}), and $\xi \sim$ dflat(). The OpenBUGS script is shown in Table 13.1. The initial values for the three chains were centered on the maximum likelihood estimates of the parameters, which were obtained using the R package [3]. The posterior means and 95% credible intervals are shown in Table 13.2.

Table 13.1 OpenBUGS script for river level example

```
model {
for(i in 1:N) {
level[i] ~ dgev(mu, sigma, eta)
}
sigma ~ dgamma(0.0001, 0.0001)
mu ~ dflat()
xi ~ dflat()
}
data
list(level = c(4.03, 3.83, 3.65, 3.88, 4.01,
    4.08, 4.18, 3.8, 4.36, 3.96, 3.98, 4.69, 3.85, 3.96, 3.85, 3.93,
    3.75, 3.63, 3.57, 4.25, 3.97, 4.05, 4.24, 4.22, 3.73, 4.37, 4.06, 3.71, 3.96, 4.06, 4.55, 3.79, 3.89,
    4.11, 3.85, 3.86, 3.86, 4.21, 4.01, 4.11, 4.24, 3.96, 4.21, 3.74, 3.85, 3.88, 3.66, 4.11, 3.71, 4.18,
    3.9, 3.78, 3.91, 3.72, 4, 3.66, 3.62, 4.33, 4.55, 3.75, 4.08, 3.9, 3.88, 3.94, 4.33), N = 65)
inits
list(mu=3.8, sigma=1, xi=0)
list(mu=3.9, sigma=2, xi=-0.2)
list(mu=3.8, sigma=1.5, xi=0.1)
```

Table 13.2 Posterior summaries for GEV parameters in river level example

Table	Post mean	95% Interval
μ	3.87	(3.82, 3.93)
σ	0.20	(0.17, 0.25)
ξ	−0.03	(−0.21, 0.19)

A 100 year return level can be estimated as follows. This corresponds to $p = 0.01$ in Eq. 13.4. So we add the following line (based on Eq. 13.4) to the OpenBUGS script in Table 13.1:

```
z.01 <- mu - sigma/xi*(1-pow (-log(1-0.01),-xi))
```

The posterior mean is 4.8 m, with a 95% credible interval of (4.5, 5.4). A 500 year return level would be given by a similar line of script for `z.002`. The posterior mean for the 500 year return level is 5.1 m, with a 95% credible interval of (4.6, 6.2).

13.1.2 Thresholds and the Generalized Pareto Distribution

One objection to the use of annual maxima is that it is wasteful of data; we throw away all of the daily readings in each year except the largest one. Also, several values in one year may be larger than the maximum value in another year, but this information is lost as well. A way around these problems is to keep all of the data, and focus on the likelihood of exceeding some predetermined large threshold value.

If we treat X_1, X_2, … as a sequence of independent and identically distributed variables, each with distribution F, then we can use the earlier GEV results as follows, and focus on the tail behavior of the GEV distribution. In other words, we want to find the conditional distribution of z, given that z is in the tail of the distribution, say $z > \mu$. First, we know that

$$\Pr(M_n \leq z) = [F(z)]^n \approx G(z) = \exp\left\{-\left[1+\xi\left(\frac{z-\mu}{\sigma}\right)\right]^{-1/\xi}\right\}$$

Taking logs of both sides gives

$$\log G(z) \approx -\left[1+\xi\left(\frac{z-\mu}{\sigma}\right)\right]^{-1/\xi} \tag{13.5}$$

The Taylor series of $G(z)$ about μ is given by

$$G(z) \approx G(\mu) - g(\mu)(z-\mu) \approx 1 - g(\mu)(z-\mu)$$

where we have used $G(\mu) \approx 1$ because μ is assumed to be large.

The Taylor series for $\log[G(z)]$ is given by

$$\log[G(z)] \approx \log[G(\mu)] - \frac{g(\mu)}{G(\mu)}(z-\mu) \approx -g(\mu)(z-\mu)$$

And so we can write

$$\log[G(z)] \approx -[1-G(z)]$$

Substituting into Eq. 13.5, we find

$$G(z) = 1 - \left[1+\frac{\xi}{\sigma}(z-\mu)\right]^{-1/\xi} \tag{13.6}$$

The distribution in Eq. 13.6 is called a *generalized Pareto distribution* (*GPD*). The parameters (μ, σ, ξ) are the parameters of the GEV distribution discussed earlier. Our argument for deriving the GPD has been approximate, but a more rigorous argument leads to the same result. The tail of the GPD is bounded for $\xi < 0$, and unbounded for $\xi \geq 0$.

The data now consist of a sequence of independent and identically distributed measurements, z_1, z_2,…, z_n. We have to choose a threshold value (μ), and we then keep all z_is that are above this threshold value. The values above this threshold are treated as a random sample from a GPD with parameters μ, σ, and ξ.

As an example, the plot in Fig. 13.2 shows daily rainfall levels (in mm) for a period of 17,531 days, taken from [1]. Assume we have decided that a rainfall of 30 mm or more in a single day is of concern, so we set this as our threshold. The plot in Fig. 13.3 shows the rainfall values greater than this threshold.

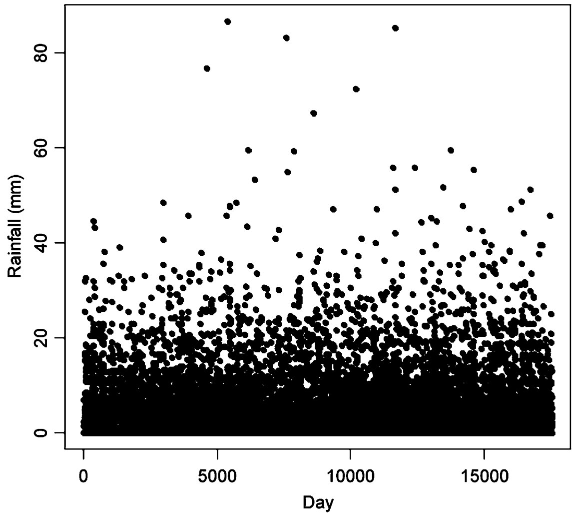

Fig. 13.2 Plot of daily rainfall in millimeter, from [1]

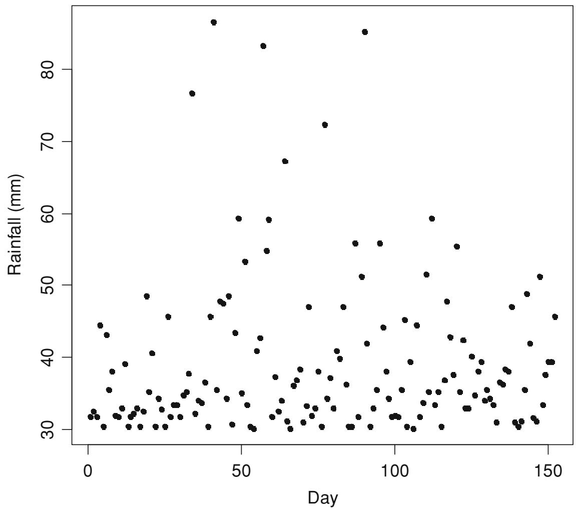

Fig. 13.3 Plot of daily rainfall values from Fig. 13.2 in excess of 30 mm

We model these excess values as a random sample from a GPD with parameters $\mu = 30$, σ, and ξ. We use independent, diffuse priors as before: $\sigma \sim$ dgamma (10^{-4}, 10^{-4}), and $\xi \sim$ dflat(). We could parse the data into values in excess of 30 mm in OpenBUGS, via a line of script such as `rain.large[i] <- rain[i]*step(rain[i] - 30)`. However, because the data file is so large, BUGS runs *very* slowly if we do this. So instead, we created another data file with just the rainfall values in excess of 30 mm. The OpenBUGS script is shown in Table 13.3. The initial values for the two chains were centered on the maximum likelihood estimates of the parameters, which were obtained using the R package [3]. The posterior means and 95% credible intervals are shown in Table 13.4.

Table 13.3 OpenBUGS script for modeling excess rainfall data as GPD

```
model {
for(i in 1:N) {
rain.large[i] ~ dgpar(mu, sigma, xi)
}
mu <- 30
sigma ~ dgamma(0.0001, 0.0001)
xi ~ dflat()
}

list(sigma=7, xi=0.1)
list(sigma=7.5, xi=-0.1)
```

Table 13.4 Posterior summaries for GPD parameters in excess rainfall example

Parameter	Post. mean	95% Interval
σ	7.4	(5.7, 9.4)
ξ	0.21	(0.03, 0.44)

We can also calculate return levels with the GPD model. We are interested in finding z_m, the value that is exceeded on average every m days. Thus, for a 100 year return rainfall level, we would have (ignoring leap years) $m = (365)(100)$. We can write

$$\Pr(Z > z_m) = \Pr(Z > \mu)[1 - G(z_m)]$$

The first term in this equation we can denote as p_μ. The number of daily rainfall values $> \mu$ has a binomial distribution with parameters p_μ and n = number of daily rainfall values, which is 17,531 in our example. We will use the Jeffreys prior for p_μ, which is a beta (0.5, 0.5) distribution.

We then substitute in $G(z)$, which is the GPD, given by Eq. 13.6:

$$\Pr(Z > z_m) = p_\mu \left[1 + \xi \left(\frac{z_m - \mu}{\sigma}\right)\right]^{-1/\xi} = \frac{1}{m}$$

Solving for z_m gives

$$z_m = \mu + \frac{\sigma}{\xi}\left[\left(m p_\mu\right)^{\xi} - 1\right]$$

The OpenBUGS script to calculate the 100 year return level is shown in Table 13.5. The posterior mean and 95% credible interval for the 100 year return level are 119 mm and (82.8, 201.2).

13.2 Treatment of Expert Opinion

The focus of this section is on methods for using information obtained from experts. Whether one has information from one or several experts, one would usually need to develop a representative estimate for use in the analysis. When

Table 13.5 OpenBUGS script to estimate 100 year return level for rainfall, based on GPD model

```
model {
for(i in 1:N) {
rain.large[i] ~ dgpar(mu, sigma, xi)
}
#100 year return level
x.large ~ dbin(p.mu, 17531)
p.mu ~ dbeta(0.5, 0.5)
m <- 100*365
z.m <- mu + (sigma/xi)*(pow(m*p.mu, xi) - 1)
mu <- 30
sigma ~ dgamma(0.0001, 0.0001)
xi ~ dflat()
}

data
list(x.large=152)

list(sigma=7, xi=0.1)
list(sigma=7.5, xi=-0.1)
```

the opinions of several experts are elicited, methods are needed to form the "aggregated" or "consensus" opinion. The formulation is quite simple conceptually. Expert opinion is simply treated as information about the unknown parameter of interest. The information is then used to update the analyst's own (prior) estimate through Bayes theorem. We will describe some of the basic techniques for a number of important classes of problems, but the coverage will not be exhaustive as the techniques for certain classes of problems tend to become very complicated without any assurance of significant improvement in the resulting estimates.

13.2.1 Information from a Single Expert

In this case the expert provides an estimate for an unknown parameter of interest, such as λ in the Poisson distribution. To use this information to update the analyst's prior distribution for λ via Bayes' Theorem, a likelihood function is needed for the information obtained from the expert. When the epistemic uncertainty in parameter values spans several orders of magnitude, as is common in PRA, a lognormal distribution is a convenient likelihood function. The parameter τ (logarithmic precision) in the lognormal distribution will represent the analyst's assessment of the expert's expertise: small values of τ correspond to low confidence (high uncertainty) and vice versa. A bias factor can also be introduced if desired, with bias less than one meaning the analyst believes the expert tends to

Table 13.6 OpenBUGS script for lognormal (multiplicative error) model for information from single expert

```
model {
lambda.star ~ dlnorm(mu, tau) # Lognormal likelihood for information from expert
lambda.star <- 1/MTTF
mu <-log(lambda*bias)
tau <- pow(log(EF.expert)/1.645, -2)
lambda ~ dlnorm(mu.analyst, tau.analyst) # Analyst's lognormal prior for lambda
mu.analyst <- log(prior.median)
tau.analyst <- pow(log(prior.EF)/1.645, -2)
}
data
list(MTTF=500000, bias=0.75, EF.expert=10, prior.median=0.000001, prior.EF=10)
```

underestimate the true value of τ, and a factor greater than one means the expert tends to overestimate the true value.

As an example, assume the analyst's estimate for the failure of a level sensor is 10^{-6}/h, but the analyst is not very confident of this estimate. The analyst adopts a lognormal distribution with this estimate as the median, and an error factor of 10 to describe the uncertainty. The level sensor vendor (the "expert") provides an estimate of the mean time to failure (MTTF) for the level sensor. The vendor's estimate for the MTTF is 500,000 h. We can develop a posterior distribution for the level sensor failure rate that uses these two sources of information.

The first step is to convert the MTTF estimate provided by the vendor into an estimate of the failure rate. This can be done by taking the reciprocal of the MTTF estimate. The analyst must assess an uncertainty factor on the vendor's estimate, representing their confidence in the estimate provided by the vendor. Assume that the analyst is not very confident in the vendor's estimate, and assesses an error factor of 10. He also believes that the expert tends to overestimate the MTTF, that is, he *underestimates* the failure rate, so a bias factor of 0.75 is assessed by the analyst. The OpenBUGS script in Table 13.6 is used to analyze this problem. Running the script in the usual way gives a posterior mean for λ of 2.7×10^{-6}/h with a 90% credible interval of (3.2×10^{-7}, 8.4×10^{-6}). If the analyst thought the vendor tended to underestimate the MTTF, that is, overestimate λ, he would use a bias factor greater than one. Changing the bias factor to 5, for example, changes the posterior mean to 1.0×10^{-6}/h with a 90% credible interval of (1.3×10^{-7}, 3.2×10^{-6}).

13.2.2 Using Information from Multiple Experts

Cases encountered in practice often involve more than one expert. When multiple experts are involved the main question concerns the method of aggregation or pooling to form a single representative or aggregate estimate from the multiple expert estimates. A number of *ad hoc* approaches have been used for combining information from multiple experts, such as taking the geometric average (arithmetic average

Table 13.7 Expert estimates of pressure transmitter failure rate, from [4]

Expert	Estimate (per hour)	Confidence measure (error factor)
1	3.0×10^{-6}	3
2	2.5×10^{-5}	3
3	1.0×10^{-5}	5
4	6.8×10^{-6}	5
5	2.0×10^{-6}	5
6	8.8×10^{-7}	10

Table 13.8 OpenBUGS script for combining information from multiple experts using multiplicative error model (lognormal likelihood)

```
model {
for(i in 1:N){
        lambda.star[i] ~ dlnorm(mu, tau[i])
        tau[i] <- pow(log(EF[i])/1.645, -2)
        }
mu ~ dflat() # Diffuse prior on mu
lambda <- exp(mu) # Monitor this node for aggregate distribution
}

data
list(lambda.star=c(3.E-6, 2.5E-5, 1.E-5, 6.8E-6, 2.E-6, 8.8E-7), EF=c(3,3,5,5,5,10), N=6)
inits
list(mu=-10)
list(mu=-5)
```

of the logarithms) and taking the low and high estimates as the 5 and 95th percentiles of a lognormal distribution. A justification commonly given for these *ad hoc* approaches is that analytical techniques need not be more sophisticated than the pool of estimates (experts' opinions) to which they are applied. Therefore, a simple averaging technique (equal weights) has often been judged satisfactory as well as efficient, especially when the quantity of information collected is large.

The Bayesian approach of the previous section can be expanded to include multiple experts. Basically, the hierarchical Bayes methods of Chap. 7 can be used to develop a prior distribution representing the variability among the experts. While mathematically cumbersome, this is straightforward to encode in OpenBUGS, as the following example from [4] illustrates.

Six estimates are available for the failure rate of pressure transmitters. These estimates along with the assigned measure of confidence are listed in Table 13.7. The analyst wishes to aggregate these estimates into a single distribution that captures the variability among the experts.

As in the previous section, the likelihood function for each expert will be assumed to be lognormal. A diffuse prior is placed on λ (actually on the logarithm of λ). The OpenBUGS script in Table 13.8 is used to develop a distribution for λ, accounting for the variability among the six experts. Running the script in the

Table 13.9 Hypothetical failure data for fan check valves, from [5]

Record number	Failure mode	Failures	Demands
469	FTO	0	11,112
470	FTO	0	3,493
471	FTO	0	10,273
472	FTO	1	971
473	FTO	0	4,230
474	FTO	0	704
475	FTO	0	7,855
476	FTO	0	504
477	FTO	0	891
478	FTO	0	846
480	FTO	0	572
481	FTO	0	631
482	FTO	0	2,245
488	FTO	0	7,665
532	FTO	0	1,425
534	FTO	0	700
538	FTO	0	716
549	FTO	8	1,236
550	FTO	0	926
551	FTO	1	588
552	FTO	0	856
554	FTO	1	708
569	FTO	0	724
570	FTO	12	8,716
592	FTO	2	632
593	FTO	0	564

usual way gives a posterior mean for λ of 6.5×10^{-6}/h, with a 90% credible interval of $(3.4 \times 10^{-6}, 1.1 \times 10^{-5})$.

13.3 Pitfalls of *ad hoc* Methods

For a group of failure records, one *ad hoc* technique that has been encountered by the authors is a type of data pooling that attempts to approximate the hierarchical Bayes approaches of Chap. 7. For example, each source might be used to generate a mean and variance as follows. If the number of failures (x) is greater than zero, then each source can be used to generate a beta distribution for p with parameters $\alpha = x$ and $\beta = n - x$, where n is the number of demands (the standard conjugate updating approach). From the properties of a beta distribution, the mean is then x/n and the variance is approximately $x(n - x)/n^3$. If no failures were recorded for a particular data source, then α might be taken to be 0.5 (assuming the Jeffreys prior for p is used). For the hypothetical data in Table 13.9, taken from [5], the overall mean and variance can be found to be

Table 13.10 Summary of overall fan check valve prior, average-moment approach

Fitted distribution	Mean	5th Percentile	95th Percentile
Beta	9.6E-4	9.4E-8	4.0E-3
Lognormal	9.6E-4	6.8E-5	3.3E-3

9.6E-4 and 2.8E-6, respectively. From these moments, the parameters of the resulting beta distribution can be found using:

$$\alpha_{tot} \approx \frac{\mu_{tot}^2}{\sigma_{tot}^2}$$

$$\beta_{tot} = \frac{\alpha_{tot}(1 - \mu_{tot})}{\mu_{tot}}$$

A lognormal distribution could also be fit using these overall moments. For the data in Table 13.9, the overall beta distribution has parameters 0.3 and 343.6. The fitted lognormal distribution has mean 9.6E-4 and error factor (EF) of 7.

The overall posterior developed by this *ad hoc* method using the average-moment approach depends on the fitted distribution and is summarized below for two different distributions Table 13.10.

Because of the large source-to-source variability exhibited by the data in Table 13.9, it may be inappropriate to pool the data as was done above. The standard Bayesian approach to such a problem is to specify a hierarchical prior for the demand failure probability, p, as described in Chap. 7. We will compare the results from this approach with the *ad hoc* average-moment approach above. We will analyze two different first-stage priors, beta and logistic-normal, with independent diffuse hyperpriors in both cases. The OpenBUGS script in Table 13.11 was used to carry out the analysis.

13.3.1 Using a Beta First-Stage Prior

The overall average distribution representing source-to-source variability in p has a mean of 1.2×10^{-3}, variance of 1.4×10^{-4}, and a 90% credible interval of $(3.5 \times 10^{-20}, 4.1 \times 10^{-3})$. The very small 5th percentile is an artifact of choosing a beta distribution as a first-stage prior. The posterior mean of α is 0.12, and the average variability distribution has a sharp vertical asymptote at $p = 0$.

13.3.2 Using a Logistic-Normal First-Stage Prior

The logistic-normal distribution is constrained to lie between 0 and 1, and because the density function goes to 0 at both 0 and 1, it avoids the vertical asymptote at $p = 0$ from which the above beta distribution suffers. For small values of p, the logistic-normal and lognormal distributions are very close; we chose to use the

Table 13.11 OpenBUGS hierarchical Bayes analysis of data in Table 13.9

```
model {
for(i in 1:N) {
        x[i] ~ dbin(p[i], n[i]) # Binomial model for number of events in each source
        p[i] ~ dbeta(alpha, beta) # First-stage beta prior
        #p[i] <- exp(p.norm[i])/(1 + exp(p.norm[i])) # Logistic-normal first-stage prior
        #p.norm[i] ~ dnorm(mu, tau)
        x.rep[i] ~ dbin(p[i], n[i]) # Replicate value from posterior predictive distribution
        #Generate inputs for Bayesian p-value calculation
        diff.obs[i] <- pow(x[i] - n[i]*p[i], 2)/(n[i]*p[i]*(1-p[i]))
        diff.rep[i] <- pow(x.rep[i] - n[i]*p[i], 2)/(n[i]*p[i]*(1-p[i]))
        }
p.avg ~ dbeta(alpha, beta)   #Average beta population variability curve
#p.avg ~ dlnorm(mu, tau)
#p.norm.avg ~ dnorm(mu, tau)
#p.avg <- exp(p.norm.avg)/(1 + exp(p.norm.avg))
#Compare observed failure total with replicated total
x.tot.obs <- sum(x[])
x.tot.rep <- sum(x.rep[])
percentile <- step(x.tot.obs - x.tot.rep)  #Looking for values near 0.5
# Calculate Bayesian p-value
chisq.obs <- sum(diff.obs[])
chisq.rep <- sum(diff.rep[])
p.value <- step(chisq.rep - chisq.obs)  #Mean of this node should be near 0.5
# Hyperpriors for beta first-stage prior
alpha ~ dgamma(0.0001, 0.0001)
beta ~ dgamma(0.0001, 0.0001)
#mu ~ dflat()
#tau <- pow(sigma, -2)
#sigma ~ dunif(0, 20)
}

inits
list(alpha=1, beta=100) #Chain 1
list(alpha=0.1, beta=200) #Chain 2
list(mu=-11, sigma=1)
list(mu=-12, sigma=5)
```

logistic-normal distribution because, with such large variability, the Monte Carlo sampling in OpenBUGS can generate values of $p > 1$, and these have the potential to skew the results, particularly the mean.

With a logistic-normal first-stage prior, we found the overall average distribution representing source-to-source variability to have a mean of 0.01, variance of 8.7E-3, and 90% credible interval of (7.6E-11, 1.2E-2).

Table 13.12 Posterior results for the *ad hoc* versus Bayesian method comparison

Method	Mean	5th Percentile	95th Percentile
Ad hoc (beta)	1.3E-4	1.4E-8	5.9E-4
Ad hoc (lognormal)	3.0E-4	4.3E-5	8.3E-4
Hierarchical Bayes (beta)	4.9E-5	5.9E-23	2.9E-4
Hierarchical Bayes (logistic-normal)	4.7E-5	2.6E-11	2.5E-4
Jeffreys prior	2.5E-4	9.8E-7	9.6E-4

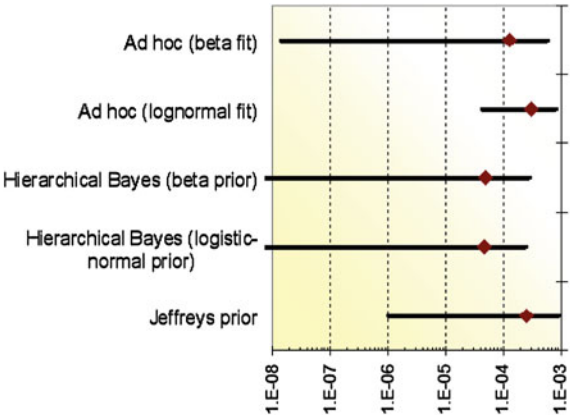

Fig. 13.4 Plot of the posterior results for the Bayesian versus *ad hoc* method comparison

13.3.3 Update with New Data

Assume we would like to update the overall check valve prior with new data, which we take to be 0 failures in 2,000 demands. The results of the four different update possibilities are shown in Table 13.12. As a reference point, we include an update of the Jeffreys prior.

As shown in Fig. 13.4, a perhaps surprising outcome is that updating the lognormal distribution fit with the *ad hoc* average-moment approach gives about the same mean and 95th percentile for p as simply updating the Jeffreys prior (however, the 5th percentile differs between the two posteriors). The beta prior would give about the same result if there were not a vertical asymptote at $p = 0$, causing excess shrinkage of the mean toward 0.

Both hierarchical Bayes analyses give similar means and 95th percentiles; the 5th percentiles differ because of the vertical asymptote in the beta first-stage prior. Hierarchical Bayes allows the large number of sources with zero failures to more strongly influence the result than the average-moment *ad hoc* approach. With no failures in 2,000 demands, the posterior mean is pulled more towards a value of zero in the hierarchical Bayes analysis, giving a less conservative result.

Table 13.13 Model validation results for the *ad hoc* versus Bayesian method comparison

Method	Total replicated failures (mean)	Bayesian p-value
Ad hoc (beta)	61.0	0.001
Ad hoc (lognormal)	67.3	0.001
Hierarchical Bayes (beta)	25.1	0.44
Hierarchical Bayes (logistic-normal)	25.0	0.38

13.3.4 Model Checking

We can generate replicate failure counts for the data sources in Table 13.9, and then use the Bayesian p-value calculated from the chi-square summary statistic described in Chap. 4 to compare models. Table 13.13 shows the results of this model-checking calculation. The *ad hoc* distributions derived from the average-moment approach are poor at replicating the observed data: they over-predict the total number of failures (the observed total was 25) and they under-predict the variability in the failure count, leading to a very low Bayesian p-value. In contrast, the hierarchical Bayes models have much better predictive validity.

13.4 Specifying a New Prior Distribution in OpenBUGS

There are times when an analyst may wish to use a distribution that is not available directly as a choice in OpenBUGS. We have already encountered two instances of this. In Chap. 3, we say how to specify a logistic-normal prior for p in the binomial distribution by specifying the underlying normal distribution and then transforming. In Chap. 9 we saw how to specify an aleatory model for failure with repair in the case of a power-law nonhomogeneous Poisson process. There, we used the `dloglik()` distribution, which requires us to specify the logarithm of the likelihood function. In this section, we show how to use the `dloglik()` distribution to enter a prior distribution, where there is no underlying distribution to exploit via a transformation, as we were able to do in the case of the logistic-normal distribution.

Suppose that the analyst wishes to use a truncated exponential distribution for p in a binomial aleatory model, this being a type of maximum entropy prior for p, if the analyst knows a mean value, μ, and lower and upper bounds a and b, respectively. As discussed in [6], the density function is given by

$$f(p) = \frac{\beta e^{\beta p}}{e^{\beta b} - e^{\beta a}}$$

where β is determined by the specified mean constraint, μ. As an example, assume that p is known to lie between a = 0.1 and b = 0.8. Assume that the mean is specified as μ = 0.7. The parameter β is found to be 4.5 by numerically solving the following equation:

Table 13.14 OpenBUGS script to specify maximum entropy prior for *p*

```
model {
p ~ dunif(0.1, 0.9)
zero <- 0
zero ~ dloglik(phi)
phi <- log(beta) + beta*p - log(exp(beta*b) - exp(beta*a))
beta <- solution(F(s), 2, 8, 0.1)
F(s) <- (b*exp(s*b) - a*exp(s*a))/(exp(s*b) - exp(s*a)) - 1/s - mu
x ~ dbin(p, n)
}
data
list(a=0.1, b=0.9, mu=0.7)
list(x=5, n=9)
```

$$\frac{be^{\beta b} - ae^{\beta a}}{e^{\beta b} - e^{\beta a}} - \frac{1}{\beta} - \mu = 0$$

The OpenBUGS script in Table 13.14, solves for β using the `solution()` function, specifies the maximum entropy prior for p, and updates it with 5 events in 9 trials, giving a posterior mean for p of 0.63, which we note lies between the prior mean of 0.7 and the MLE of 0.56.

13.5 Bayesian Inference for Parameters of a Markov Model

Markov models are occasionally encountered in PRA applications, especially when time-dependence is an explicit concern. In this section, we illustrate the capability to simultaneously perform Bayesian inference for the Markov model parameters and solve the system of Markov ordinary differential equations (ODEs) within OpenBUGS.

We take as our example the Markov model used in [7] to estimate piping rupture frequency. This model is shown in Fig. 13.5.

13.5.1 Aleatory Models for Failure

The occurrences of failures as a result of stress corrosion cracking (SC) and design and construction defects (DC) were assumed to be described by independent Poisson processes, with rates λ_{SC} and λ_{DC}, respectively. The occurrence of failures overall is then described by a Poisson process with rate $\lambda = \lambda_{SC} + \lambda_{DC}$. Data consist of the number of SC and DC failures, n_{SC} and n_{DC}, observed over specified exposure times.

Reference [7] accounted for uncertainty in the exposure times via a discrete distribution, with nine components for SC, and three for DC. Lognormal priors for

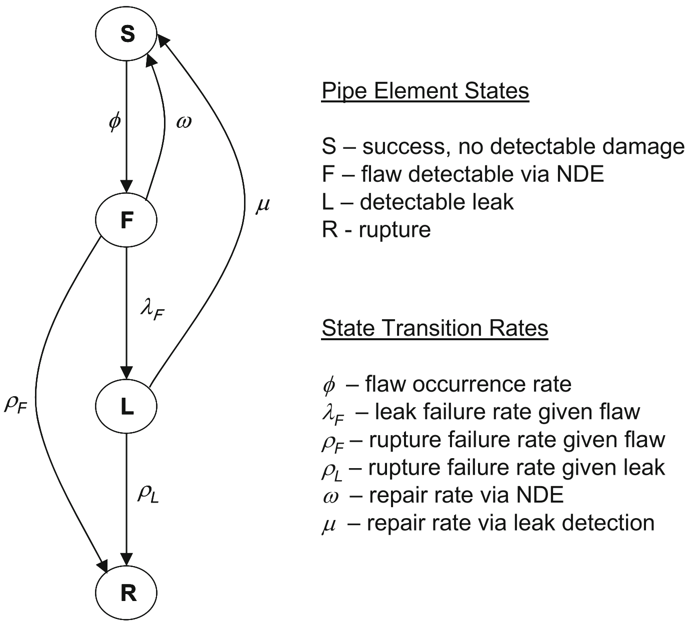

Fig. 13.5 Markov model for piping rupture, taken from [7]

λ_{SC} and λ_{DC} were updated with $n_{SC} = 8$ and $n_{DC} = 4$, with the exposure times and weights given in Fig. 13.6.

The posterior distributions for λ_{SC} and λ_{DC} are obtained by averaging over the weights given in Fig. 13.6. The overall failure rate, λ, is then found by summing λ_{SC} and λ_{DC}. The posterior distribution for λ is multimodal.

Occurrences of ruptures conditional upon failures were described by a binomial distribution with parameters $P(R_l)$ and 12 (sum of n_{SC} and n_{DC}). Each $P(R_l)$ has a lognormal prior distribution, as given in Table 13.15.

The frequency of pipe rupture of a given size is found by multiplying λ by $P(R_l)$.

13.5.2 Other Markov Model Parameters

Uncertainties for the other parameters were represented as described in [7], with the exception of P_{FD} and P_{LD}, for which [7] used a triangular distribution. Because the triangular distribution is not implemented in OpenBUGS, a beta distribution was used over the range given in [7], with a mean value approximately equal to the mode of the triangular distribution.

Base Exposure = (Reactor-Years)x(Welds per Reactor)=3088.6x366= 1.13E6 weld years

Welds	366
Rx-yrs	3089
Base Exposure	1,130,428

Weld Count Uncertainty	Fraction of Welds Susceptible to Stress Corrosion Cracking (SC)	Exposure Case Probability	Exposure Multiplier	Exposure
p=.25 High (1.5 X Base)	p=.25 High (.25 X Base)	0.0625	0.375	423,910 weld-yrs
	p-.50 Medium (.05 X Base)	0.125	0.075	84,782 weld-yrs
	p=.25 Low (.01 X Base)	0.0625	0.015	16,956 weld-yrs
p=.50 Medium (1.0 X Base)	p=.25 High (.25 X Base)	0.125	0.25	282,607 weld-yrs
	p-.50 Medium (.05 X Base)	0.25	0.05	56,521 weld-yrs
	p=.25 Low (.01 X Base)	0.125	0.01	11,304 weld-yrs
p=.25 High (0.5 X Base)	p=.25 High (.25 X Base)	0.0625	0.125	141,303 weld-yrs
	p-.50 Medium (.05 X Base)	0.125	0.025	28,261 weld-yrs
	p=.25 Low (.01 X Base)	0.0625	0.005	5,652 weld-yrs

Fig. 13.6 Discrete distribution model for service data exposure uncertainty

Table 13.15 Lognormal prior distributions for pipe rupture probabilities conditional upon pipe failure

Symbol	Break size (in.)	Lognormal prior distribution Mean	Range factor
$P(R_1)$	0.032	0.13	2.8
$P(R_2)$	0.10	0.06	3.4
$P(R_3)$	0.32	0.021	4.5
$P(R_4)$	1.00	5.00E-03	6.6
$P(R_5)$	3.16	8.00E-04	10.5
$P(R_6)$	10.0	1.80E-04	15.1
$P(R_7)$	31.62	3.80E-05	21.8
$P(R_8)$	42.4	2.11E-05	25.1

13.5.3 Markov System Equations

The Markov model is described by a set of four coupled linear first-order ordinary differential equations (ODEs). The equations give the rate of change of the four components of the state probability vector in terms of the Markov parameters and the state probability vector components, which are time-dependent. The initial condition necessary for the solution of the vector equation is $P(0) = (1, 0, 0, 0)^T$.

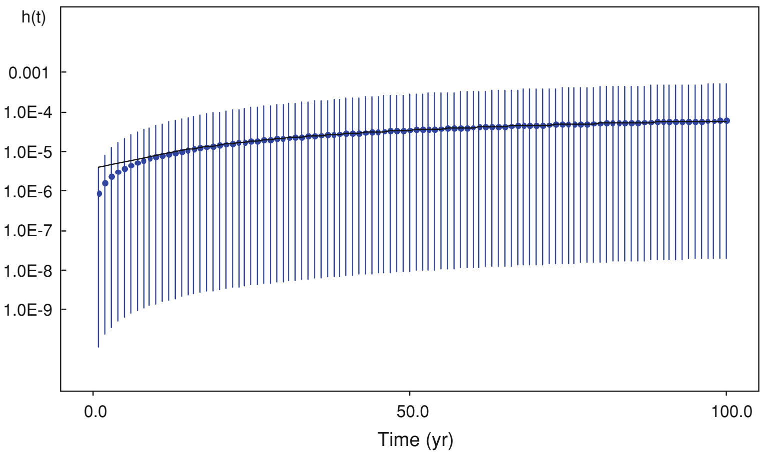

Fig. 13.7 System hazard rate as a function of time for 10 inch rupture. Line is posterior mean and bars illustrate 95% credible intervals

The system of ODEs can be written as

$$A = \begin{pmatrix} -\phi & \omega & \mu & 0 \\ \phi & -(\lambda_F + \rho_F + \omega) & 0 & 0 \\ 0 & \lambda_F & -(\rho_L + \mu) & 0 \\ 0 & \rho_F & \rho_L & 0 \end{pmatrix}$$

With rupture (state 4) defined as failure for the system, the system hazard rate is given by

$$\begin{aligned} h(t) &= \frac{f(t)}{1 - F(t)} = -\frac{1}{r(t)}\frac{dr}{dt} \\ &= \frac{\rho_F P_2(t) + \rho_L P_3(t)}{\sum_{i=1}^{3} P_i(t)} \end{aligned}$$

13.5.4 Implementation in OpenBUGS

In OpenBUGS, a numerical ODE solver is used, via the `ode()` function. The OpenBUGS script in Table 13.16 implements this example. Note the use of a transformation to encode a beta distribution over a range other than [0, 1]. The model converged quickly, but for conservatism the first 1,000 iterations were discarded for burn-in. Parameter estimates are based on 50,000 iterations after burn-in, resulting in a Monte Carlo error of 1% or less of the overall standard deviation for each parameter. Posterior summaries for the input parameters

Table 13.16 OpenBUGS script for Markov pipe rupture model

```
model {
#Aleatory models for number of failures
for(i in 1:9) {
        n.sc[i] ~ dpois(mean.sc[i])
        mean.sc[i] <- lambda.sc[i]*time.sc[i]
        lambda.sc[i] ~ dlnorm(mu.sc, tau.sc)
        }
for(j in 1:3) {
        n.dc[j] ~ dpois(mean.dc[j])
        mean.dc[j] <- lambda.dc[j]*time.dc[j]
        lambda.dc[j] ~ dlnorm(mu.dc, tau.dc)
        }
for(k in 1:8) {
        x.R[k] ~ dbin(p.R[k], n.R)
        p.R[k] ~ dlnorm(mu.R[k], tau.R[k])
        mu.R[k] <- log(prior.mean.R[k]) - pow(sigma.R[k], 2)/2
        sigma.R[k] <- log(RF.R[k])/1.645
        tau.R[k] <- pow(sigma.R[k], -2)
#rho.rupt[k] is used to calculate the transition rate rho.F below
        rho.rupt[k] <- lambda*p.R[k]
        }
####################################
#Weighted-average posterior distributions
lambda.sc.avg <- lambda.sc[r.sc]
r.sc ~ dcat(w.sc[])
lambda.dc.avg <- lambda.dc[r.dc]
r.dc ~ dcat(w.dc[])
lambda <- lambda.sc.avg + lambda.dc.avg
########################################
#Other Markov model parameters
phi <- m.f*lambda
m.f.trunc ~ dbeta(1, 2)
m.f <- (10 - 1)*m.f.trunc + 1
lambda.F <- lambda*f.L/f.f
f.L ~ dbeta(1, 4)
f.f ~ dbeta(1, 2)
rho.L ~ dlnorm(mu.rho.L, tau.rho.L)
mu.rho.L <- log(mean.rho.L) - pow(sigma.rho.L, 2)/2
sigma.rho.L <- log(RF.rho.L)/1.645
tau.rho.L <- pow(sigma.rho.L, -2)
mu <- P.LD/(T.LI + T.R)
#Triangular distribution replaced with beta distribution over same range
P.LD.trunc ~ dbeta(9, 1)
P.LD <- (0.99-0.5)*P.LD.trunc + 0.5
omega <- P.I*P.FD/(T.FI + T.R)
```

(continued)

Table 13.16 (continued)

```
P.FD.trunc ~ dbeta(9, 1)
P.FD <- (0.99-0.1)*P.FD.trunc + 0.1
#Sum over appropriate rupture frequency range for given rupture size
#0.1-inch break
#rho.F <- sum(rho.rupt[2:8])/f.f
#10-inch break
rho.F <- sum(rho.rupt[6:8])/f.f
#####################################
#  Markov system equations
solution[1:n.grid, 1:dim] <- ode(init[1:dim], times[1:n.grid], D(P[1:dim], t), origin, tol)

D(P[1], t) <- -phi*P[1] + omega*P[2] + mu*P[3]
D(P[2], t) <- phi*P[1] - lambda.F*P[2] - omega*P[2] - rho.F*P[2]
D(P[3], t) <- lambda.F*P[2] - mu*P[3] - rho.L*P[3]
D(P[4], t) <- rho.F*P[2] + rho.L*P[3]
#System hazard rate
#This is the rate at which ruptures occur divided by the probability of not being in the ruptured
state at time j
for(j in 1:n.grid) {
          h.sys[j] <- (rho.F*solution[j,2] + rho.L*solution[j,3])/(solution[j,1] + solution[j,2] +
 solution[j,3])
          }
####################################
#Prior distribution parameters
mu.sc <- log(prior.mean.sc) - pow(sigma.sc, 2)/2
sigma.sc <- log(RF.sc)/1.645
tau.sc <- pow(sigma.sc, -2)
mu.dc <- log(prior.mean.dc) - pow(sigma.dc, 2)/2
sigma.dc <- log(RF.dc)/1.645
tau.dc <- pow(sigma.dc, -2)
}

data
list(
n.grid = 100, dim = 4, origin = 0, tol = 1.0E-8, init=c(1,0,0,0),
n.sc=c(8,8,8,8,8,8,8,8,8),
n.dc=c(4,4,4),
prior.mean.sc=4.27E-5, RF.sc=100,
prior.mean.dc=1.24E-6, RF.dc=100,
time.sc=c(423910, 84782, 16956, 282607, 56521, 11304, 141303, 28261, 5652),
w.sc=c(0.0625, 0.125, 0.0625, 0.125, 0.25, 0.125, 0.0625, 0.125, 0.0625),
time.dc=c(1695641, 1130428, 565213.8),
w.dc=c(0.25, 0.50, 0.25),
P.I=1, T.FI=10, T.R=200,
T.LI=1.5,
prior.mean.R=c(0.13, 0.06, 0.021, 5.0E-3, 8.0E-4, 1.8E-4, 3.8E-5, 2.11E-5),
RF.R=c(2.8, 3.4, 4.5, 6.6, 10.5, 15.1, 21.8, 25.1),
x.R=c(0,0,0,0,0,0,0,0), n.R=12,
mean.rho.L=0.02, RF.rho.L=3
)
```

Table 13.17 Posterior summaries of rate parameters in Markov model

	5th	50th	Mean	95th
λ_{SC}	1.7E-5	1.3E-4	2.7E-4	1.0E-3
λ_{DC}	8.7E-7	2.8E-6	3.5E-6	8.5E-6
$\lambda = \lambda_{SC} + \lambda_{DC}$	2.0E-5	1.3E-4	2.7E-4	1.0E-3
$\lambda_F = \lambda f_L/f_f$	2.9E-6	7.7E-5	1.6E-3	2.0E-3

Table 13.18 Posterior summaries of conditional rupture probabilities based on updating lognormal priors with 0 ruptures in 12 trials

Rupture size (in.)	5th	50th	Mean	95th
0.032	0.03	0.07	0.08	0.16
0.10	0.01	0.03	0.04	0.10
0.32	2.8E-3	0.01	0.02	0.05
1.00	3.8E-4	2.4E-3	4.3E-3	0.01
3.16	2.7E-5	2.8E-4	7.9E-4	2.9E-3
10.0	3.1E-6	4.7E-5	1.8E-4	6.9E-4
31.62	3.1E-7	6.7E-6	3.6E-5	1.4E-4
42.4	1.2E-7	3.1E-6	2.5E-5	8.0E-5

Table 13.19 Posterior summaries of rupture frequencies λ times conditional rupture probabilities from Table 13.18

Rupture size (in.)	5th	50th	Mean	95th
0.032	1.2E-6	9.4E-6	2.1E-5	8.4E-5
0.10	5.1E-7	4.6E-6	1.1E-5	4.5E-5
0.32	1.5E-7	1.6E-6	4.4E-6	1.8E-5
1.00	2.3E-8	3.3E-7	1.2E-6	4.8E-6
3.16	1.9E-9	3.9E-8	2.1E-7	8.1E-7
10.0	2.3E-10	6.3E-9	4.7E-8	1.8E-7
31.62	2.4E-11	9.0E-10	9.3E-9	3.5E-8
42.4	9.8E-12	4.2E-10	7.0E-9	2.0E-8

Table 13.20 Posterior summaries of remaining Markov parameters

	5th	50th	Mean	95th
$\phi = m_f\lambda$	5.4E-5	4.65E-4	1.1E-3	4.3E-3
μ	4.2E-3	4.7E-3	4.7E-3	4.9E-3
ω	3.5E-3	4.4E-3	4.3E-3	4.7E-3
ρ_L	5.3E-3	1.6E-2	2.0E-2	4.8E-2
ρ_F (0.1-in.)	2.6E-6	3.1E-5	8.1E-4	6.5E-4
ρ_F (10-in.)	2.2E-9	5.2E-8	2.0E-6	1.6E-6

Table 13.21 Results for system hazard rate (yr^{-1})

h(t, size)	5th	Median	Mean	95th
h(1, 0.1)	2.1E-10	1.5E-8	8.7E-7	1.6E-6
h(10, 0.1)	2.6E-9	1.9E-7	7.4E-6	1.9E-5
h(40, 0.1)	1.3E-8	1.0E-6	2.8E-5	9.9E-5
h(100, 0.1)	3.8E-8	2.8E-6	6.0E-5	2.5E-4
h(1, 10)	3.6E-12	3.6E-10	3.2E-8	4.9E-8
h(10, 10)	2.3E-10	2.7E-8	1.8E-6	3.9E-6
h(40, 10)	2.7E-9	3.2E-7	1.4E-5	4.3E-5
h(100, 10)	1.1E-8	1.3E-6	3.9E-5	1.5E-4

to the ODEs are shown in Tables 13.17, 13.18, 13.19, and 13.20. The hazard rate results for 10 inch ruptures are plotted in Fig. 13.7 and tabulated in Table 13.21.

References

1. Coles S (2001) An introduction to statistical modeling of extreme values. Springer, Berlin
2. Gumbel EJ (1958) Statistics of extremes. Columbia University Press, New York
3. R Development Core Team (2011) R: A language and environment for statistical computing. Vienna, Austria
4. Mosleh A (1992) Bayesian modeling of expert-to-expert variability and dependence in estimating rare event frequencies. Reliab Eng Syst Saf 38:47–57
5. Dezfuli H, Kelly DL, Smith C, Vedros K, Galyean W (2009) Bayesian inference for NASA probabilistic risk and reliability analysis. NASA, Washington
6. Siu NO, Kelly DL (1998) Bayesian parameter estimation in probabilistic risk assessment. Reliab Eng Syst Saf 62:89–116
7. Fleming KN (2004) Markov models for evaluating risk-informed in-service inspection strategies for nuclear power plant piping systems. Reliab Eng Syst Saf 83:27–45