LA-UR -82-1668

Los Alamos National Laboratory
Los Alamos, New Mexico 87545

Los Alamos National Laboratory is operated by the University of

MASTER

LA-UR-

0002 01

CONF-820656--2

TITLE: PREDICTING LOW-PROB

California for the United States Department of Energy under contract W-7405-ENG-36.

AUTHOR(S): H. F. Martz and M

82-1668

1371

SUBMITTED TO: Workshop on Low-Probability/High-Consequence Events
June 15-17, 1982; Arlington, Virginia

tyson

DISCLAIMER

By acceptance of this article the publisher recognizes that the U.S. Government has certain rights in the published form of this contribution or to allow others to do

The Los Alamos National Laboratory requests that the publisher

Probability/High-Consequence Risk Analysis
Arlington, Virginia

Los Alamos

mb

PREDICTING LOW-PROBABILITY/HIGH-CONSEQUENCE EVENTS

by

H. F. Martz and M. C. Bryson
Group S-1, MS#F600
Los Alamos National Laboratory
Los Alamos, New Mexico 87545

Submitted to:

Proceedings of the Workshop on
Low-Probability/High-Consequence Risk Analysis
June 15-17, 1982, Arlington, Virginia

PREDICTING LOW-PROBABILITY/HIGH-CONSEQUENCE EVENTS

ABSTRACT

PRAs often require quantification of the probabilities of various low-probability events, such as accident-initiating events and hardware-fault events. A two-stage Bayes/empirical Bayes data pooling procedure is presented for use in combining as many as five different types of relevant data. A Poisson model is assumed for the event in question. Empirical Bayes methods are used to determine the population variability curve, while Bayesian methods are used to specialize this curve to the specific event in question.

The procedure is illustrated by an example in which we estimate the probability of failure of a hypothetical large dam based on (1) a deductive event-tree-type analysis of the probability, (2) historical US dam failure data, (3) the opinions of a committee of several dam experts, and (4) the operating history for the dam in question. A Stage-2 posterior distribution is produced which incorporates these data sources. Similar distributions are produced for various combinations of data types and used to assess the contribution of each data source.

KEY WORDS: risk assessment; pooling data; Bayesian methods; expert opinion; low-probability events; empirical Bayes; rare events.

1. INTRODUCTION

The prediction of low-probability/high-consequence events is an important but difficult task in probabilistic risk assessment (PRA). Objective frequency data are usually scarce or nonexistent, and scientific cause-effect relationships are unknown. Consequently, assumptive models are frequently used that depend on less objective sources: expert opinion, related or extrapolated data, deductive analysis, etc. The problem, then, reduces to that of modeling and ultimately combining such data. The resulting estimation procedures are addressed here. For convenience, we discuss only the case of a homogeneous Poisson process model, in which case the unknown rate of occurrence

is the parameter to be estimated. However, the general method extends to other models as well.

Various Bayesian methods have been proposed and used to combine relevant types of data in PRAs. References 1-8 describe the use of Bayesian methods in various nuclear reactor applications. The theoretical background for the use of the Bayesian approach in PRA⁹ and the practical steps required^{10,11} have also been described.

In most of these approaches, only two basic types of data are considered. Generic data (nonplant- or nonapplication-specific) are used to determine a suitable "prior" distribution. The prior either reflects a subjective degree-of-belief regarding uncertainty in the unknown occurrence rate, or describes a natural variation in the parameter over some large well-defined population. In the latter case, the prior distribution has a relative frequency interpretation and may be referred to as a population variability curve. Both views are taken in this paper, depending on the available data.

The second data type is plant- or application-specific data. Such data are used in the Bayesian approach to "specialize" the prior to the specific plant or application in the form of the "posterior" distribution.

Kaplan¹² proposes a "two-stage" Bayesian procedure in which up to three types of information (two generic sources and one plant-specific source) are combined. The Zion PRA¹³ also used this two-stage Bayesian approach.

Expert opinion is often used as one source of data for determining a prior distribution on the occurrence rate. Both the Reactor Safety Study¹⁴ and the IEEE Reliability Data Manual¹⁵ use expert opinions. The Reactor Safety Study used them to fit lognormal distributions, which were then interpreted as population variability curves. The IEEE manual formed consensus estimates as

simple unweighted geometric averages of individual expert opinions obtained using Delphi methods.

Two biases are known to accompany expert opinions. Peterson and Miller¹⁶ and Dalkey and Brown¹⁷ determined that a probability of interest was often underestimated and that the 65th percentile of the distribution of expert opinions was often a better estimate of the unknown probability than central measures such as the mean and median. This first bias will be subsequently referred to as the location bias and suitable adjustments will be made in the procedure that follows.

The second bias concerns the accuracy of precision of the expert opinions. Alpert and Raiffa,¹⁸ Lichtenstein et al.,¹⁹ and Slovic et al.²⁰ all concluded that the assessment of percentiles of a distribution tend to be too tight; that is, experts tend to "overvalue" their estimates. Slovic et al. found that, instead of 2% of true values falling outside expert-assessed nominal 98% probability bounds, often 20 to 50% of the true values were outside these limits. Thus, expert-supplied uncertainty bounds tend to be too tight. This scale bias will also be suitably accounted for in the procedure that follows by broadening the corresponding distributions as was also done in the Zion PRA.¹³

A Bayes/empirical Bayes data pooling procedure is described in Sec. 2. Although similar in spirit to Kaplan's procedure,¹² it takes a different approach in permitting the combination of up to five different data types, which are illustrated by two examples. Section 3 illustrates the detailed use of the procedure in estimating dam failure rates. Some conclusions are presented in Sec. 4.

2. A TWO-STAGE BAYESIAN/EMPIRICAL BAYES PROCEDURE

Suppose we are interested in estimating the rate of occurrence λ of some specific event of interest, such as a dam failure rate or the rate of occurrence of some hardware fault in a reactor safety system. With respect to this event, consider these five types of data.

- I_1 : Analysis results regarding the rate of occurrence (event-specific analysis data).
- I_2 : Historical (observed) frequency data on similar events in similar applications (raw generic frequency data).
- I_3 : Expert opinion(s) regarding the rate of occurrence, which may include the analyst's own subjective opinions (event-specific expert opinion data).
- I_4 : Historical (observed) frequency data on the specific event (event-specific historical data).
- I_5 : Generic rate of occurrence estimates for similar (or extrapolated) events (reduced generic data in the form of point and interval estimates).

Data type I_1 may be appropriate for events such as accident initiating events, failure events of man-made systems, etc., where a synthesis-type analysis (such as an event tree analysis of an initiating event), a Monte Carlo simulation experiment, a physical model of the event, or an a priori PRA of the event itself has been constructed. Data type I_1 may not be present when considering basic events in nuclear reactor PRAs involving hardware component failure modes.

Data types I_2 and I_5 are assumed to consist of frequency data and/or occurrence rate estimates in which each element in the data set has an associated underlying true (but unknown) rate of occurrence, considered to be a random value from the population variability curve. Thus I_2 and I_5 are interpreted as generic data and used to fit the Stage 1 prior distribution, which has a corresponding relative frequency interpretation based on the use of empirical Bayes (EB) methods.

Data types I_1 , I_3 , and I_4 are data on the specific event of interest and thus are referred to as event-specific data. Data type I_3 is a subjective data source, I_4 is objective, and I_1 may be either subjective, objective, or a combination. The important point is that each data type has the effect of specializing whatever generic data-based prior is available at Stage 1 to the specific event of interest.

In order to more fully understand the five data types, consider two examples. First, suppose that we are interested in estimating the failure rate of a certain large dam as part of a PRA. Second, as part of a PRA for a certain proposed commercial operating LMFBR, suppose that we wish to estimate the hourly failure rate of certain classes of sodium valves. Table I describes several possible data sources categorized according to I_1 - I_5 for these examples. It should be noted that such things as the modeling assumptions, whether or not a data source is generic or event-specific, and so forth, should be used to classify each source of data into type I_1 - I_5 . The first example will be further considered in Sec. 3, while Martz and Bryson²¹ provide a detailed discussion of the second example. Before considering such details, let us first identify the notation and format for each data type.

Table I

2.1 Data Notation and Format

The notation and format for each of the five data types is as follows.

- I_1 : Either $\langle M_1, V_1 \rangle$ or $\langle f_1, T_1 \rangle$, where $M_1 = E(\lambda)$ is the mean rate of occurrence, $V_1 = \text{Var}(\lambda)$ is the variance of the rate of occurrence, and f_1 is the expected number of occurrences of the specific event in total exposure time T_1 .
- I_2 : $\{ \langle f_{2,m}, T_{2,m} \rangle : m = 1, 2, \dots, M \}$, where $f_{2,m}$ is the observed number of occurrences in total exposure time $T_{2,m}$ of the m^{th} similar application and where $M \geq 2$ is the number of similar applications considered.
- I_3 : $\{ \langle \lambda_{3,n}^\alpha, \tilde{\lambda}_{3,n}, \lambda_{3,n}^{1-\alpha} \rangle : n = 1, 2, \dots, N \}$, where $\tilde{\lambda}_{3,n}$ is the subjective estimate of λ provided by the n^{th} expert, $\lambda_{3,n}^\alpha$ and $\lambda_{3,n}^{1-\alpha}$ are the $100\alpha^{\text{th}}$ and $100(1-\alpha)^{\text{th}}$ percentiles, respectively, of the n^{th} expert's underlying subjective distribution of $\tilde{\lambda}_{3,n}$, and $N \geq 1$ is the number of expert opinions available.
- I_4 : $\langle f_4, T_4 \rangle$, where f_4 is the observed number of occurrences of the specific event in total exposure time T .
- I_5 : $\{ \langle \lambda_{5,k}^{\alpha_k}, \tilde{\lambda}_{5,k}, \lambda_{5,k}^{1-\alpha_k} \rangle : k = 1, 2, \dots, K \}$, where $\tilde{\lambda}_{5,k}$ is a generic point estimate from the k^{th} generic data source, $\lambda_{5,k}^{\alpha_k}$ and $\lambda_{5,k}^{1-\alpha_k}$ are the $100\alpha_k^{\text{th}}$ and $100(1-\alpha_k)^{\text{th}}$ percentiles, respectively, from the k^{th} data source, and $K \geq 1$ is the number of generic data sources.

In Sec. 2.3.1, two somewhat different forms for I_5 are also considered. In the first form an error factor, in lieu of percentiles, is reported by some or all of the K generic data sources. In the second, no uncertainties are provided by any of the K sources. The mixed case can also be considered by

either ignoring the uncertainties or assuming uncertainties for those that are missing.

2.2 The Two-Stage Procedure

Fig. 1

Figure 1 illustrates the proposed two-stage procedure. Note that either a noninformative or a lognormal family of prior distributions is used at Stage 1 in which the prior is interpreted as a population variability curve. A noninformative prior distribution is appropriate when generic information is non-existent, unavailable, or deliberately ignored, and the analyst is almost totally ignorant about the value of λ .^{10,22} In this case, I_2 and I_5 are not present and the final Stage 2 posterior distribution will depend almost entirely on the event-specific data (I_1, I_3 , and I_4).

If I_2 and/or I_5 is available, a lognormal family of prior distributions is used at Stage 1. There are several reasons for this choice. A lognormal prior is the conjugate prior for a lognormal likelihood,²² which will be considered in Sec. 2.3.2. The lognormal prior was used in the Reactor Safety Study¹⁴ and thus has historical precedent. Other authors, such as Apostolakis et al.⁶ and Kaplan,¹² have also used a lognormal prior in a Poisson process model.

Sections 2.3 and 2.4 describe the details regarding Stages 1 and 2 of the two-stage procedure shown in Fig. 1.

2.3 Stage 1

Sections 2.3.1-2.3.3 consider the prior, likelihood, and posterior segments of Stage 1 of the two-stage procedure.

2.3.1 Stage 1 Prior Distribution. If I_2 and/or I_5 are unavailable for use in the analysis, the noninformative prior distribution for λ ^{10,22} given by

$$f(\lambda) \propto \lambda^{-1}, \quad 0 < \lambda < \infty, \quad (1)$$

is used as the Stage 1 prior to begin the process. Although this is an improper distribution (it cannot be scaled to integrate to unity), the corresponding Stage 1 posterior distribution will be shown to be proper. Notice also that there are no prior parameters in Eq. (1) that must be estimated, which would be impossible in the absence of generic information (I_2 and I_5).

Next, let us suppose that I_2 and/or I_5 are available. The lognormal family of prior distributions for λ is given by

$$f(\lambda; \kappa n \lambda_0, \psi_0^2) = \frac{1}{\lambda \psi_0 \sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{\kappa n \lambda - \kappa n \lambda_0}{\psi_0} \right)^2 \right], \quad \lambda > 0, \quad (2)$$

which we denote simply as $LN(\kappa n \lambda_0, \psi_0^2)$, and where $E(\kappa n \lambda) = \kappa n \lambda_0$ and $\text{Var}(\kappa n \lambda) = \psi_0^2$. Since the lognormal prior at Stage 1 depends on data types I_2 and/or I_5 , we denote the prior at Stage 1 as $f(\lambda; I_2, I_5)$, realizing that I_2 and I_5 will be subsequently used to estimate the prior parameters $\kappa n \lambda_0$ and ψ_0^2 .

Data Type I_2

First consider data type I_2 . We assume that information I_2 is an independent random sample of occurrence rates $\langle \lambda_{2,m} \rangle$ from the Stage 1 prior distribution (the population variability curve), where $\lambda_{2,m}$ corresponds to the observed Poisson data pair $\langle f_{2,m}, T_{2,m} \rangle$. In addition it is assumed that, conditional on the underlying values of $\lambda_{2,m}$, the Poisson data sets are statistically independent. Note that each underlying $\lambda_{2,m}$ value corresponding to each of the M Poisson data sets will likely be different as these values are assumed to be random selections from the underlying Stage 1 prior distribution. This assumed structural model for data type I_2 is precisely the situation referred to as an

empirical Bayes situation in statistics.²² Thus it is appropriate to estimate λ_0 and ψ_0^2 by use of empirical Bayes estimation methods.

Based on the use of the method of moments,²² empirical Bayes estimates for λ_0 and ψ_0^2 become (see Appendix A)

$$\hat{\lambda}_{0,2} = 2 \sum a_1 - 0.5 \sum a_2 \quad (3)$$

and

$$\hat{\psi}_{0,2}^2 = \sum a_2 - 2 \sum a_1, \quad (4)$$

where a_1 and a_2 are calculated from the I_2 data by means of

$$\begin{aligned} a_1 &= \sum f_{2,m} / \sum T_{2,m}, \\ a_2 &= 2(f_{2,m} / T_{2,m})(f_{2,m} - 1) / T_{2,m}, \end{aligned} \quad (5)$$

where the indicated summations are from $m = 1$ to $m = M$. The "hats" in Eqs. (3) and (4) denote "estimates for" and the subscripts "2" denote estimates based on the use of I_2 . From Eq. (5) it follows that we must require that at least one $f_{2,m} > 1$; otherwise, the prior distribution either degenerates to 0 or has non-estimable parameters in the empirical Bayes approach.

Note also that the variance estimator in Eq. (4) can yield negative estimates, which sometimes occurs when the prior distribution is highly concentrated (nondiffuse) and the values of $f_{2,m}$ are small. A conservative estimator then for ψ_0^2 is

$$\hat{\psi}_{0,2}^L = \frac{1}{2} T_{2,m} (\ln \hat{\lambda}_{2,m} - \overline{\ln \hat{\lambda}_{2,m}})^2 / \frac{1}{2} T_{2,m} , \quad (6)$$

where

$$\overline{\ln \hat{\lambda}_{2,m}} = \frac{1}{2} T_{2,m} \ln \hat{\lambda}_{2,m} / \frac{1}{2} T_{2,m} ,$$

and where $\hat{\lambda}_{2,m} = f_{2,m}/T_{2,m}$. If $\hat{\lambda}_{2,m} = 0$, then $\hat{\lambda}_{2,m}$ is replaced by an arbitrary small constant. However, Eqs. (3) and (4) usually provide satisfactory estimates of the lognormal prior parameters and in several example applications of the procedure, Eq. (4) has never produced negative estimates.

Data Type I₅

For data type I₅, the same underlying assumptions are made as for I₂; namely, the true (but unknown) occurrence rate $\lambda_{5,k}$ underlying the data supplied by generic data source k is a random value from the assumed lognormal population variability distribution (the prior at Stage 1). In addition, the k data sources that constitute I₅ are assumed to be conditionally statistically independent. It is realized that this assumption may be invalid to some degree in practice, since many reported generic estimates are derived from the same underlying raw data sources. Nevertheless, this simplifying assumption will be made in our initial attempt to address the general pooling problem described in Sec. 1.

As with I₂, we use an empirical Bayes method of moments approach for estimating $\ln \lambda_0$ and ψ_0^2 . However, we assume now that the generic point estimate $\hat{\lambda}_{5,k}$ has a conditional $LN(\ln \lambda_{5,k}, \sigma_k^2)$ distribution, where $\lambda_{5,k}$ has a $LN(\ln \lambda_0, \psi_0^2)$ prior distribution. The parameter σ_k measures the uncertainty associated with $\hat{\lambda}_{5,k}$ and is given by

$$\sigma_k = (\sum_{i=1}^n \lambda_{5,k}^{1-\alpha_k} - \sum_{i=1}^n \lambda_{5,k}^{\alpha_k}) / 2z_{1-\alpha_k} \quad , \quad (7)$$

where $z_{1-\alpha_k}$ is the $100(1-\alpha_k)^{th}$ percentile of a standard normal distribution. If, instead of two percentiles, an error factor EF_k is reported by the k^{th} generic data source, then

$$\sigma_k = \sum_{i=1}^n EF_k / z_{1-\alpha_k} \quad , \quad (8)$$

for either a stated or assumed value of α_k .

The empirical Bayes estimators of $\sum_{i=1}^n \lambda_0$ and ψ_0^2 become (see Appendix B)

$$\hat{\sum_{i=1}^n \lambda_{0,5}} = \sum_{k=1}^K \sigma_k^{-1} \sum_{i=1}^n \tilde{\lambda}_{5,k} / \sum_{k=1}^K \sigma_k^{-1} \quad (9)$$

and

$$\hat{\psi}_{0,5}^2 = s_{\lambda}^2 - \sum_{k=1}^K \sigma_k^2 / K \quad , \quad (10)$$

where

$$s_{\lambda}^2 = \sum_{k=1}^K (\sum_{i=1}^n \tilde{\lambda}_{5,k} - \overline{\sum_{i=1}^n \lambda})^2 / K, \quad \overline{\sum_{i=1}^n \lambda} = \sum_{k=1}^K \sum_{i=1}^n \tilde{\lambda}_{5,k} / K \quad , \quad (11)$$

and all indicated summations are from $k = 1$ to $k = K$.

As with Eq. (4), Eq. (10) may sometimes yield a negative estimate. A conservative solution is then to use the estimator of ψ_0^2 given by

$$\hat{\psi}_{0,5}^2 = s_{\lambda}^2 \quad , \quad (12)$$

which tends to overestimate the variance of the prior distribution, thus making it more diffuse than it may actually be. In essence, this estimator simply ignores the uncertainty in each generic point estimate $\tilde{\lambda}_{5,k}$.

If $K = 1$, then a different approach must be used to estimate $\mu\lambda_0$ and ψ_0^2 . The recommended approach is to assume that $\tilde{\lambda}_{5,k}$ is the lognormal prior median from which it follows that $\hat{\mu\lambda}_{0,5} = \mu\tilde{\lambda}_{5,k}$. This estimate is then used in conjunction with either the stated percentiles or error factor to directly solve for the estimate $\hat{\psi}_{0,5}^2$ of ψ_0^2 .

Now consider the case in which only a set of K generic point estimates are available for I_5 . In this case, each point estimate is treated as though it were exactly equal to the corresponding value $\lambda_{5,k}$ (since no measure of uncertainty exists) and ψ_0^2 and $\mu\lambda_0$ are estimated by means of the estimators

$$\hat{\mu\lambda}_{0,5} = \overline{\mu\lambda} \quad , \quad \hat{\psi}_{0,5}^2 = s_{\lambda}^2 \quad . \quad (13)$$

Both Data Types I_2 and I_5

If both data types I_2 and I_5 are used in the analysis, they can be combined to estimate a single lognormal prior distribution at Stage 1 by combining the individual lognormal prior distributions based on I_2 and I_5 . By considering an estimator formed as the pooled estimate of the common variance ψ_0^2 based on two independent variance estimates ($\hat{\psi}_{0,2}^2$ and $\hat{\psi}_{0,5}^2$), each based on data from a population with different means ($\mu\lambda_{0,2}$ and $\mu\lambda_{0,5}$), we obtain the combined weighted average estimator of ψ_0^2 given by

$$\hat{\psi}_0^2 = w_2 \hat{\psi}_{0,2}^2 + w_5 \hat{\psi}_{0,5}^2 + w_2 w_5 (\hat{x}n\lambda_{0,2} - \hat{x}n\lambda_{0,5})^2, \quad (14)$$

where w_2 and w_5 are weights assigned by the analyst to I_2 and I_5 , respectively, with $w_2 + w_5 = 1$. In the absence of any preferences, the selection $w_2 = M/(M + K)$ and $w_5 = K/(M + K)$ is an obvious choice, as is $w_2 = w_5 = 0.5$.

Similarly, the weighted estimator of $xn\lambda_0$ given by

$$\hat{x}n\lambda_0 = w_2 \hat{x}n\lambda_{0,2} + w_5 \hat{x}n\lambda_{0,5} \quad (15)$$

is the corresponding estimator based on the combined use of I_2 and I_5 .

2.3.2 Stage 1 Likelihood. The event-specific expert opinion data I_3 is incorporated into the overall procedure in the Stage 1 likelihood function, denoted by $L(I_3|\lambda)$.

Some experimental evidence suggests that the lognormal distribution is appropriate for modeling expert opinions. Dalkey²³ and Martino²⁴ both concluded that the lognormal distribution represents a good model for combined (mixed) expert opinions. This conclusion suggests that a lognormal likelihood in conjunction with a lognormal prior is a suitable combination, which is what is considered here. Mosleh and Apostolakis²⁵ also use a lognormal model to represent expert opinions. The likelihood function in this section is similar to the model proposed by Mosleh and Apostolakis, having been significantly motivated by their work.

Let us assume a $LN[\psi(\lambda), \sigma_n^2]$ distribution for the n^{th} expert's subjective estimate $\bar{\lambda}_{3,n}$ of λ . In terms of the two percentiles $\lambda_{3,n}^\alpha$ and $\lambda_{3,n}^{1-\alpha}$ supplied by the n^{th} expert, we have that

$$\sigma_n^2 = (xn\lambda_{3,n}^{1-\alpha} - xn\lambda_{3,n}^\alpha) / 2z_{1-\alpha}. \quad (16)$$

Recall from Sec. 1 that location and scale biases often accompany expert opinions. To account for the location bias, let us assume that the unknown value λ is the $100\delta^{\text{th}}$ percentile of the $LN[\xi(\lambda), \sigma_n^2]$ distribution, where δ accounts for the bias. The larger (or smaller) the value of δ , the larger the location bias in the model. For example, if $\delta = 0.65$, the actual (unknown) value of λ is located at the 65^{th} percentile of the expert's subjective distribution for λ . It follows that $\xi(\lambda) = \mu n \lambda - \sigma_n^2 z_\delta$.

We account for the scale bias by assuming that the tail-area probabilities to the left and right of $\lambda_{3,n}^u$ and $\lambda_{3,n}^{1-u}$ are actually $\rho = k\alpha$, rather than α as claimed. As in the case of δ , ρ must also be specified by the analyst. For example, if $\alpha = 0.05$ and $\rho = 0.15$, then instead of providing 5^{th} and 95^{th} percentiles as believed, the expert is unknowingly providing 15^{th} and 85^{th} percentiles instead.

Although the biases may actually differ between experts, the identification of such individual differences is beyond the state-of-the-art in cognitive processes such as these. Consequently, we consider identical values of δ and ρ for each expert.

For simplicity, let us initially assume that all expert point estimates are statistically independent. Dependencies may be accounted for by the use of a multivariate lognormal model, as considered by Mosleh and Apostolakis.²⁵ The joint subjective distribution of all the expert-supplied estimates of λ thus becomes

$$f(\tilde{\lambda}_3; \lambda) = \frac{1}{(n\tilde{\lambda}_{3,n})(n\sigma_n)(2\pi)^{n/2}} \exp \left[-\frac{1}{2} \left(\frac{\mu n \tilde{\lambda}_{3,n} - \mu n \lambda + \sigma_n^2 z_\delta}{\sigma_n} \right)^2 \right], \quad (17)$$

where $\tilde{\lambda}_3 = (\tilde{\lambda}_{3,1}, \dots, \tilde{\lambda}_{3,N})$ is the vector of expert-supplied estimates of λ and where

$$\sigma_n = (\ln \lambda_{3,n}^{1-\alpha} - \ln \lambda_{3,n}^\alpha) / 2z_{1-\beta} \quad (18)$$

In this section all products and summations are from $n = 1$ to $n = N$.

Martz and Bryson²⁶ show that the maximum likelihood estimate of λ based on I_3 data is given by

$$\hat{\lambda}_3 = \exp \left(\sum_{n=1}^N w_n \hat{\lambda}_{3,n} \right) \quad (19)$$

where $\hat{\lambda}_{3,n} = \exp(\ln \tilde{\lambda}_{3,n} + \sigma_n z_\delta)$ is the "adjusted" subjective estimate of λ for the n^{th} expert, and where w_n is the exponential "weight" applied to the n^{th} expert's adjusted estimate defined by $w_n = \sigma_n^{-2} / \sum \sigma_n^{-2}$. Thus, $\hat{\lambda}_3$ is a weighted geometric mean of the adjusted expert-supplied estimates of λ .

If all experts give the same range of uncertainty on their estimate and further if there is no location bias ($\delta = 0$), then $\hat{\lambda}_3$ reduces to the simple geometric mean of the expert estimates as used in the compilation of the IEEE Nuclear Reliability Data Manual.¹⁵ Martz and Bryson²⁶ compare both the weighted and unweighted geometric mean estimators and conclude that the estimator in Eq. (19) is generally significantly superior in mean-squared-error performance.

It follows that the Stage 1 likelihood function indicated in Fig. 1 may be expressed in terms of $\hat{\lambda}_3$ in the form of a lognormal distribution as

$$L(I_3 | \lambda) = \frac{1}{\hat{\lambda}_3 \left(\sum w_n \sigma_n^2 \right)^{1/2} \sqrt{2\pi}} \exp \left\{ -\frac{1}{2} \left[\frac{\ln \hat{\lambda}_3 - \ln \lambda}{\left(\sum w_n \sigma_n^2 \right)^{1/2}} \right]^2 \right\} \quad (20)$$

2.3.3 Stage 1 Posterior Distribution. Bayes' theorem is used to obtain the Stage 1 posterior distribution as shown in Fig. 1. In the case of a noninformative prior at Stage 1, the Stage 1 posterior distribution of λ , $f(\lambda|I_3)$, is an $LN(\kappa\hat{\lambda}_3, \frac{1}{\kappa}w_{nn}^2)$ distribution.²²

The mean and median of $f(\lambda|I_3)$ are $\hat{\lambda}_3 \exp(\frac{1}{2}w_{nn}^2)$ and $\hat{\lambda}_3$, respectively, which may be used as point estimates of λ based on I_3 . A 100% two-sided Bayes uncertainty interval for λ is given by $(\hat{\lambda}_3/F, \hat{\lambda}_3 \cdot F)$, where the factor F is given by

$$F = \exp[z_{1-\gamma/2}(\frac{1}{\kappa}w_{nn}^2)^{1/2}] \quad . \quad (21)$$

Thus, $1 - \gamma$ is the subjective probability that λ is contained in the interval, and the probability excluded by each endpoint is $\gamma/2$.

In the case where I_2 or I_5 (or both) is used, the Stage 1 posterior distribution of λ , denoted by $f(\lambda|I_2, I_3, I_5)$, is likewise an $LN(\kappa\lambda_1, \psi_1^2)$ distribution in which

$$\kappa\lambda_1 = v_1 \kappa\hat{\lambda}_0 + v_2 \kappa\hat{\lambda}_3 \quad (22)$$

and

$$\psi_1^2 = v_1 v_2 (\frac{1}{\kappa}w_{nn}^2 + \hat{\psi}_0^2) \quad , \quad (23)$$

where

$$v_1 = \frac{1}{\kappa}w_{nn}^2 / (\frac{1}{\kappa}w_{nn}^2 + \hat{\psi}_0^2) \quad (24)$$

and $v_2 = 1 - v_1$. Thus, λ_1 is the weighted geometric mean of $\hat{\lambda}_0$ and $\hat{\lambda}_3$ given by

$\lambda_1 = \hat{\lambda}_0^v \hat{\lambda}_3^v$. This Stage 1 posterior distribution may or may not be event-specific, depending upon the use of the event-specific data I_3 . If I_3 is not used in the analysis, then the Stage 1 posterior distribution, represented by $f(\lambda; I_2, I_5)$, is not event-specific since it is simply the Stage 1 prior distribution estimated from only the non-event-specific data I_2 and/or I_5 .

The mean and median of $f(\lambda; I_2, I_3, I_5)$ are $\lambda_1 \exp(\psi_1^2/2)$ and λ_1 , respectively, which may be used as point estimators of λ based on I_2 , I_3 , and I_5 . A 100% two-sided Bayes uncertainty interval for λ is given by $(\lambda_1/F_1, \lambda_1 \cdot F_1)$, where the factor F_1 is given by

$$F_1 = \exp(z_{1-\gamma/2} \psi_1) \quad . \quad (25)$$

2.4 Stage 2

The Stage 2 prior distribution is the Stage 1 posterior distribution discussed above. Sections 2.4.1 and 2.4.2 consider the likelihood and posterior segments for Stage 2 of the two-stage procedure.

2.4.1 Stage 2 Likelihood. In the case where the I_1 data is in the form $\langle f_1, T_1 \rangle$, we can regard both data types I_1 and I_4 as Poisson data for the common rate of occurrence λ . Thus, the Stage 2 likelihood in this case becomes

$$L(I_1, I_4 | \lambda) = (\lambda T)^f \exp(-\lambda T) / f! \quad , \quad (26)$$

where $f = f_1 + f_4$ and $T = T_1 + T_4$.

In the case where I_1 is of the form $\langle M_1, V_1 \rangle$ a different approach must be considered. Let us "convert" $\langle M_1, V_1 \rangle$ into nearly equivalent Poisson data $\langle f_1, T_1 \rangle$ by temporarily assigning a noninformative prior to λ and using a Poisson likelihood to produce a temporary posterior distribution via Bayes'

theorem. We can then equate the mean and variance of the temporary posterior distribution to $\langle M_1, V_1 \rangle$ and solve for the corresponding nearly equivalent Poisson data $\langle f_1, T_1 \rangle$. These "new" data $\langle f_1, T_1 \rangle$ are nearly equivalent to $\langle M_1, V_1 \rangle$ in the Bayesian sense that little additional information is being interjected into the results through the noninformative prior. Thus, the Poisson data provides virtually all of the information in the Bayesian posterior.

The temporary posterior distribution is a gamma distribution with shape parameter $f_1 + 0.5$ and scale parameter T_1 ,²² which we denote by $G(f_1 + 0.5, T_1)$. Setting

$$M_1 = (f_1 + 0.5)/T, \quad V_1 = (f_1 + 0.5)/T_1^2$$

yields the nearly equivalent Poisson data

$$f_1 = M_1^2/V_1 - 0.5, \quad T_1 = M_1/V_1. \quad (27)$$

Thus, in this case, let us define

$$f = \max(0, M_1^2/V_1 + f_d - 0.5), \quad (28)$$

where f should also be rounded to the nearest integer, and

$$T = M_1/V_1 + T_d. \quad (29)$$

These values are then used in Eq. (26).

2.4.2 Stage 2 Posterior Distribution. Since a Poisson likelihood and log-normal prior are not conjugate to each other, numerical integration must be

used to evaluate the Stage 2 posterior distribution which, according to Bayes' theorem, may be expressed as

$$f(\lambda | I_1, I_2, I_3, I_4, I_5) = \frac{f(\lambda | I_2, I_3, I_5) L(I_1, I_4 | \lambda)}{\int_0^{\infty} f(\lambda | I_2, I_3, I_5) L(I_1, I_4 | \lambda) d\lambda} \quad (30)$$

The well-known Gauss-Laguerre quadrature formulas are convenient for numerically performing the integration in Eq. (30). Likewise, numerical integration can be used to determine the mean and median, as well as for determining a Bayesian two-sided 100% uncertainty interval, from Eq. (30).

However, numerical evaluation of Eq. (30) is not always necessary. For all possible data-type combinations, Table II gives the corresponding Stage 1 distributions, as well as the final Stage 2 posterior distribution. Cases in which numerical integration must be used are so indicated, and it is observed that this is not always the case.

Apostolakis et al.⁶ conclude that often a lognormal distribution may be used to closely approximate the posterior distribution derived from a lognormal prior and a Poisson likelihood. Thus, it is likely that some or all of the numerically expressed posterior distributions indicated in Table II should be well-approximated by a lognormal distribution. In the example in Sec. 3, a very good lognormal approximation was obtained by matching the means and variance of the numerical and lognormal distributions, as will be shown.

3. ESTIMATING DAM FAILURE RATES

The following hypothetical example illustrates the use of the two-stage Bayes/empirical Bayes data pooling procedure. Suppose that we have a large

concrete and earthen dam, built in 1942, over which some concern has been expressed. As part of an overall estimation of the public risk, the dam failure probability is to be estimated. Assuming failure to be a Poisson process (exponentially distributed time to failure), this is equivalent to estimating the failure rate parameter λ .

Historical data from other dams represents one source of information. Table III presents actual failure data from 1900-1965 on US dams of more than 50-foot height, by ten-year intervals.^{27,28} The classification by time groups is somewhat arbitrary. It seems reasonable, though, that difference in construction materials, design philosophies, regulations, and so forth, would create some sources of variation that may be thus identified. This, then, constitutes the I_2 data in our example, and year of completion will be used to express variation in the population variability curve. It may be noted, incidentally, that since we know the date of construction of the dam in question, that we cannot regard it as a random selection from the I_2 population. In fact, there seems to be a slight trend (as would be expected from technological progress) toward lower failure rates in later time periods, so that the dam being evaluated should be expected to have a slightly lower failure rate than a randomly selected dam from the Table III population. Accordingly, our analysis will be subsequently shown to err in the conservative direction insofar as the I_2 data component is concerned. The alternative is to either ignore I_2 or to obtain a different categorization for I_2 . Since an alternative classification was unavailable, the effect of ignoring I_2 will be subsequently investigated.

To obtain some dam-specific information, suppose that a panel of five experts provide their best subjective estimates of the failure rate for this particular dam. Their hypothesized results are given in Table IV, and

Table IV constitute the I_3 data set. Based on the discussion in Sec. 1, $\phi = 0.65$ and $\rho = 0.15$ are chosen (see Sec. 2.3.2). The effect of different choices of ϕ and ρ on the overall posterior was also found to be rather slight.

Since the dam was built in 1942, it has withstood 40 years of operation without failure. This constitutes the I_4 Poisson data: zero failures in $T_4 = 40$ years.

Finally, suppose that a quantitative event tree analysis is conducted, yielding a mean failure rate prediction of 8.0×10^{-5} failures/year with a variance of 4.0×10^{-10} (failures/year)². This is the I_1 data. According to (27), it corresponds to Poisson data consisting of roughly 16 failures in 200,000 dam-years of service. Since we have no generic data in this example, there is no I_5 data component.

Table V With four data types, there are 15 possible combinations that we could use in the analysis. Table V summarizes those 15 analyses. Dashes in the "mode" column indicate that the mode is nonexistent or undetermined. Case 16 represents a lognormal distribution that approximates the numerically-derived Case 15 using all four data types. Fig. 2 illustrates the numerical posterior and the lognormal approximation (obtained by matching means and variances). In Case 16, the lower and upper 5% uncertainty bounds may be obtained by using the median value together with an error factor of 1.4. Thus, our final result would be reported as: median failure rate estimate = 8.7×10^{-5} , error factor = 1.4. It may be of interest to note that this corresponds closely with the commonly reported estimate of 10^{-4} for large dams.^{29,30}

Comparisons of the different rows in Table V provide substantial insight into the data-pooling procedure. Notice that the estimates based only on I_4 (Case 4) are so diffuse as to be practically worthless. All other three data types, especially I_1 , are more valuable. The effect of this is seen in later

rows of the table: adding I_4 to any other data combination makes practically no difference in the posterior distribution parameters. The I_2 (historical) data also yield a diffuse distribution and therefore have only small effects when combined with other data types; however, since the parameters for the I_2 distribution are somewhat larger than the others, incorporation of I_2 data tends to shift distributions slightly to the right. In view of the earlier-noted conservatism of the I_2 procedure in this case, one might want to regard such a shift with some skepticism; on the other hand, the I_2 and I_4 data types are the most objective of the four, so their incorporation into the analysis does lend credibility from the standpoint of providing basically consistent results. Most of our information is contained in the (I_1, I_3) data set (Case 5); adding I_2 and I_4 (Case 15 or 16) has no perceptible effect on the variability, but does introduce a note of conservatism by shifting the entire distribution slightly toward higher failure rates. Introducing I_2 by itself, of course (Case 12) has a virtually identical effect since I_4 is contributing so little information.

4. CONCLUSIONS

A two-stage Bayes/empirical Bayes procedure has been presented for use in combining as many as five different types of data regarding low-probability events. The five data types represent both generic and event-specific types of data.

The prior distribution at Stage 1 is considered to be a population variability curve that embodies a relative frequency notion of probability. Empirical Bayes methods are used in conjunction with either generic raw frequency data (I_2) or generic reduced frequency estimates (I_5) or both for estimating the Stage 1 prior distribution.

The remaining data types (I_1 , I_3 , and I_4) are event-specific data, which are used to specify the generic Stage 1 prior to the specific event of interest. The expert opinion data (I_3) are adjusted for two biases that often accompany such data.

An example application of the procedure is provided in which the method is used to estimate a hypothetical dam failure rate. The estimated failure rate was derived based on the results of an event-tree analysis, expert opinions, historical dam failure data in the US, and the operating history of the dam in question. The median estimated failure rate was found to be 8.7×10^{-5} failures per dam-year with an associated error factor of approximately 1.4.

The method conforms to our intuitive notions about how diverse, but nonetheless relevant, data should be combined. The sample application supports our conclusion that the procedure is conceptually pleasing and easily applied in practice.

APPENDIX A

EMPIRICAL BAYES ESTIMATORS FOR λ_0 AND ψ_0^2 BASED ON POISSON (I_2) DATA

Since $f_{2,m}$ is conditionally Poisson with mean $\lambda_{2,m} T_{2,m}$ for a given $\lambda_{2,m}$, it follows that

$$E(f_{2,m}) = T_{2,m} E(\lambda_{2,m}) = T_{2,m} E(\lambda)$$

and

$$\begin{aligned}\text{Var}(f_{2,m}) &= E(\lambda_{2,m} T_{2,m}) + \text{Var}(\lambda_{2,m} T_{2,m}) \\ &= T_{2,m} E(\lambda) + T_{2,m}^2 \text{Var}(\lambda) \quad .\end{aligned}$$

Also,

$$E(f_{2,m}^2) = T_{2,m}^2 E(\lambda^2) + T_{2,m} E(\lambda) \quad .$$

For each $\lambda_{2,m}$, the maximum likelihood estimator is

$$\hat{\lambda}_{2,m} = f_{2,m}/T_{2,m} \quad , \quad m = 1, 2, \dots, M \quad .$$

It is readily shown that the first two moments of λ may be expressed as

$$E(\lambda) = E(\hat{\lambda}_{2,m}) \tag{A-1}$$

and

$$E(\lambda^2) = E(\hat{\lambda}_{2,m}^2 - \hat{\lambda}_{2,m}/T_{2,m}) \quad . \tag{A-2}$$

Using $T_{2,m}$ as the obvious weighting factor, we can estimate the expected values in Eqs. (A-1) and (A-2) by means of the sample averages given by

$$\hat{E}(\lambda) = \sum T_{2,m} \hat{\lambda}_{2,m} / \sum T_{2,m}$$

and

$$\hat{E}(\lambda^2) = \sum T_{2,m} (\hat{\lambda}_{2,m}^2 - \hat{\lambda}_{2,m}/T_{2,m}) / \sum T_{2,m} \quad .$$

Since $E(\lambda) = \lambda_0 \exp(\psi_0^2/2)$ and $E(\lambda^2) = \lambda_0^2 \exp(2\psi_0^2)$, we can equate corresponding moments and solve for the empirical Bayes estimators of λ_0 and ψ_0^2 , thus yielding the estimators given in Eqs. (3) and (4) in Sec. 2.3.1.

APPENDIX B

EMPIRICAL BAYES ESTIMATORS FOR λ_0 AND ψ_0^2 BASED ON LOGNORMAL (I_5) DATA

Unconditionally, we have that

$$E(\lambda_{5,k}) = E(\lambda_{5,k}) = E(\lambda) = \lambda_0$$

and

$$\text{Var}(\lambda_{5,k}) = E[\text{Var}(\lambda_{5,k} | \lambda_{5,k})] + \text{Var}[E(\lambda_{5,k} | \lambda_{5,k})] = \sigma_k^2 + \psi_0^2.$$

Now, the average unconditional variance of $\lambda_{5,k}$ (averaged over all K data sources) becomes

$$\frac{1}{K} \sum \text{Var}(\lambda_{5,k}) = \bar{\sigma}^2 + \psi_0^2, \quad (B-1)$$

where $\bar{\sigma}^2 = \sum \sigma_k^2 / K$. For $K \geq 2$, consider the sample variance s_λ^2 defined by Eq. (11). It can be shown that

$$E(s_{\lambda}^2) = \frac{1}{K} \sum \text{Var}(x\tilde{\lambda}_{5,k}) \quad .$$

Thus, following the method of moments procedure, we equate s_{λ}^2 to Eq. (B-1) to find the empirical Bayes estimator of ψ_0^2 given by $\hat{\psi}_{0,5}^2 = s_{\lambda}^2 - \bar{\sigma}^2$. Similarly, since $E(\hat{x}\lambda_{0,5}) = x\lambda_0$, where $\hat{x}\lambda_{0,5}$ is defined in Eq. (9), $\hat{x}\lambda_{0,5}$ is the empirical Bayes estimator of $x\lambda_0$ based on I_5 data.

ACKNOWLEDGMENT

Development of the methods in this paper has been supported by the Nuclear Regulatory Commission under Los Alamos National Laboratory Project R665, Statistical Techniques for Risk Analysis, and by the US Department of Energy under Los Alamos Project D023, R&D Integration-Integrated Analysis (LMFBR).

REFERENCES

1. G. Apostolakis and A. Mosleh, Expert opinion and statistical evidence: An application to reactor core melt frequency, Nucl. Sci. Eng., 70, 135 (1979).
2. S. Kaplan and B. J. Garrick, On the use of a Bayesian reasoning in safety and reliability decisions--three examples, Nucl. Tech., 44, 231 (1979).
3. G. Apostolakis, S. Kaplan, B. J. Garrick, and W. Dickter, Assessment of the frequency of failure to scram in light-water reactors, Nucl. Safety, 20, 690 (1979).
4. G. Apostolakis, Bayesian methods in risk assessment, Adv. in Nucl. Sci. and Tech., 13, 415 (1981).
5. G. Apostolakis, Probability and risk assessment: The subjectivistic viewpoint and some suggestions, Nucl. Safety, 19, 305 (1978).
6. G. Apostolakis, S. Kaplan, B. J. Garrick, and R. J. Daphily, Data specialization for plant specific risk studies, Nucl. Eng. Design, 56, 321 (1980).
7. M. Kazarians and G. Apostolakis, Some aspects of the fire hazard in nuclear power plants, Nucl. Eng. Design, 47, 157 (1978).
8. G. Apostolakis, Data analysis in risk assessments, submitted to Nucl. Eng. Design, 1981.

9. T. W. Parry and P. W. Winter, Characterization and evaluation of uncertainty in probabilistic risk analysis, Nucl. Safety, 22, 28 (1981).
10. US Nuclear Regulatory Commission, A Guide to the Performance of Probabilistic Risk Assessments for Nuclear Power Plants, NUREG/CR-2300, Vol. 1, Rev. 1 (April 1982).
11. Pickard, Lowe, and Garrick, Inc., Irvine, CA, Methodology for Probabilistic Risk Assessment of Nuclear Power Plants, PLG-0209 (June 1981).
12. S. Kaplan, On a 'two-stage' Bayesian procedure for determining failure rates from experimental data, IEEE Trans. Power Apparatus and Design, to appear.
13. Pickard, Lowe, and Garrick, Inc., Irvine, CA, Zion/Indian Point Probabilistic Safety Analysis (1981).
14. US Nuclear Regulatory Commission, Reactor Safety Study, Appendix III--Failure Data, WASH-1400 (NUREG-75/014) (1975).
15. IEEE Nuclear Reliability Data Manual, IEEE Std-500 (John Wiley, New York, 1977).
16. C. Peterson and A. Miller, Mode, median, and mean as optimal strategies, J. Exp. Psych., 68, 363 (1964).
17. N. Dalkey and B. Brown, Comparison of Group Judgment Techniques with Short-Range Predictions and Almanac Questions, RAND Corp., R-578-ARPA (1971).
18. M. Alpert and H. Raiffa, A progress report on the training of probability assessors, Harvard University, unpublished, 1969.
19. S. Lichtenstein, B. Fischhoff, and L. D. Phillips, "Calibration of Probabilities: The State of the Art," in Decision Making and Human Affairs (Reidel, Holland, 1977), p. 275.
20. P. Slovic, B. Fischhoff, and S. Lichtenstein, "Fact Versus Fears: Understanding Perceived Risk," in Societal Risk Assessment (Plenum, New York, 1980).
21. H. Martz and M. Bryson, On combining data for estimating the frequency of low-probability events with application to sodium valve failure rates, submitted for publication, 1982.
22. H. Martz and R. Waller, Bayesian Reliability Analysis (John Wiley, New York, 1982).
23. N. C. Dalkey, An Experimental Study of Group Opinion, RAND Corp., RM-5848-PW (1969).
24. J. P. Martino, Lognormality of Delphi estimates, Tech. Forecasting, 1, 355 (1970).

25. A. Mosleh and G. Apostolakis, "Models for the Use of Expert Opinions," presented at the Workshop on Low-Probability/High-Consequence Risk Analysis, June 15-17, 1982, Arlington, VA.
26. H. F. Martz and M. C. Bryson, A statistical model for combining biased expert opinions, submitted for publication, 1982.
27. A. Biswas and K. Chatterjee, Dam disasters: An assessment, Eng. J., 3 (March 1971).
28. N. Schnitter, A short history of dam engineering, Water Power, 4, 2 (April 1967).
29. P. Gast, "Divergent Public Attitudes toward Nuclear and Hydro-electric Plant Safety," presented at the ANS Meeting, June 1973, Chicago, IL.
30. W. Baldewicz, "Dam Failures: Insights to Nuclear Power Risk," presented at the Workshop on Low-Probability/High-Consequence Risk Analysis, June 15-17, 1982, Arlington, VA.

TABLE I. Sample Data Sources for Dam Failure Rate and Sodium Valve Failure Rate Estimation

<u>Data Type</u>	<u>Dam Failure Rate Estimation</u>	<u>Sodium Valve Failure Rate Estimation</u>
I ₁	Event and fault-tree analysis results.	- - -
I ₂	Historical dam failure data categorized by year of construction, method of construction, type of dam, etc.	Data from various sodium test loops and/or foreign LMFBFR operating experience.
I ₃	Expert opinions of a panel of experts in dam engineering, construction, safety, operation, etc.	Opinions of sodium valve experts.
I ₄	Number of years of operation without failure.	Population and failure data for operating LMFBFRs, such as EBR-II.
I ₅	- - -	LWR failure rates of similar valves.

Table II. Prior and Posterior Distributions for Each Combination of Data Types

Data Type ^a					Stage 1 Prior	Stage 1 Posterior	Stage 2 Posterior
I_1	I_2	I_3	I_4	I_5			
+	-	-	-	-	Noninformative	Noninformative	$G(f_1+0.5, T_1)$
-	+	-	-	-	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$
-	-	+	-	-	Noninformative	$LN(\hat{\kappa n \lambda}_3, \hat{\omega}_{n \sigma n}^2)$	$LN(\hat{\kappa n \lambda}_3, \hat{\omega}_{n \sigma n}^2)$
-	-	-	+	-	Noninformative	Noninformative	$G(f_4+0.5, T_4)$
-	-	-	-	+	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$
+	+	-	-	-	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	(Numerical)
+	-	+	-	-	Noninformative	$LN(\hat{\kappa n \lambda}_3, \hat{\omega}_{n \sigma n}^2)$	(Numerical)
+	-	-	+	-	Noninformative	Noninformative	$G(f+0.5, T)$
+	-	-	-	+	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	(Numerical)
-	+	+	-	-	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	$LN(\hat{\kappa n \lambda}_1, \hat{\psi}_1^2)$	$LN(\hat{\kappa n \lambda}_1, \hat{\psi}_1^2)$
-	+	-	+	-	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	(Numerical)
-	+	-	-	+	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$
-	-	+	+	-	Noninformative	$LN(\hat{\kappa n \lambda}_3, \hat{\omega}_{n \sigma n}^2)$	(Numerical)
-	-	+	-	+	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	$LN(\hat{\kappa n \lambda}_1, \hat{\psi}_1^2)$	$LN(\hat{\kappa n \lambda}_1, \hat{\psi}_1^2)$
-	-	-	+	+	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	(Numerical)
+	+	+	-	-	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	$LN(\hat{\kappa n \lambda}_1, \hat{\psi}_1^2)$	(Numerical)
+	+	-	+	-	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	(Numerical)
+	+	-	-	+	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	(Numerical)
+	-	+	+	-	Noninformative	$LN(\hat{\kappa n \lambda}_3, \hat{\omega}_{n \sigma n}^2)$	(Numerical)
+	-	+	-	+	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	$LN(\hat{\kappa n \lambda}_1, \hat{\psi}_1^2)$	(Numerical)
+	-	-	+	+	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	(Numerical)
-	+	+	+	-	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	$LN(\hat{\kappa n \lambda}_1, \hat{\psi}_1^2)$	(Numerical)
-	+	+	-	+	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	$LN(\hat{\kappa n \lambda}_1, \hat{\psi}_1^2)$	$LN(\hat{\kappa n \lambda}_1, \hat{\psi}_1^2)$

Table II (continued)

Data Type ^a							
<u>I₁</u>	<u>I₂</u>	<u>I₃</u>	<u>I₄</u>	<u>I₅</u>	Stage 1 Prior	Stage 1 Posterior	Stage 2 Posterior
-	+	-	+	+	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	(Numerical)
-	-	+	+	+	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	$LN(\kappa n \lambda_1, \psi_1^2)$	(Numerical)
+	+	+	+	-	$LN(\hat{\kappa n \lambda}_{0,2}, \hat{\psi}_{0,2}^2)$	$LN(\kappa n \lambda_1, \psi_1^2)$	(Numerical)
+	+	+	-	+	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	$LN(\kappa n \lambda_1, \psi_1^2)$	(Numerical)
+	+	-	+	+	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	(Numerical)
+	-	+	+	+	$LN(\hat{\kappa n \lambda}_{0,5}, \hat{\psi}_{0,5}^2)$	$LN(\kappa n \lambda_1, \psi_1^2)$	(Numerical)
-	-	+	+	+	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	$LN(\kappa n \lambda_1, \psi_1^2)$	(Numerical)
+	+	+	+	+	$LN(\hat{\kappa n \lambda}_0, \hat{\psi}_0^2)$	$LN(\kappa n \lambda_1, \psi_1^2)$	(Numerical)

^a+ = Data type present; - = data type absent.

Table III. Failure Data for US Dams of More than 50-ft Height by Year of Completion for the Period 1900-1965

<u>Year of Completion</u>	<u>Total Number of Dams Built</u>	<u>Number of Failures</u>	<u>Total Dam-Years* of Operation Through 1965</u>
1900-1909	100	9	6017
1910-1919	220	12	11639
1920-1929	280	6	12367
1930-1939	280	1	9765
1940-1949	240	0	6000
1950-1959	530	2	7932

*Uses the lower endpoint of the completion interval as the approximate year of completion.

Table IV. Expert Opinions Regarding the Failure Rate of the Dam in Question

<u>Expert No.</u>	<u>5-th Percentile*</u>	<u>Best Estimate</u>	<u>95-th Percentile</u>
1	$1.0 \times 10^{-5**}$	$1.0 \times 10^{-4**}$	$1.0 \times 10^{-3**}$
2	1.0×10^{-6}	5.0×10^{-6}	1.0×10^{-5}
3	1.0×10^{-5}	3.0×10^{-5}	8.0×10^{-5}
4	7.0×10^{-7}	7.0×10^{-5}	7.0×10^{-3}
5	5.0×10^{-5}	1.0×10^{-4}	5.0×10^{-4}

*Measures the uncertainty in each expert's best estimate.

**Failures per dam per year.

Table V. Comparison of Dam Failure Rate Results for All Data Combinations

Case No.	Data Considered	Stage 1 Distributions		Stage 2 Posterior	Stage 2 Posterior Distribution ^a				
		Prior	Posterior		5th Per-centile	Mode	Median	Mean	95th Per-centile
1	I ₁	Noninformative	Noninformative	G(16.5,200000)	5.22	7.75	8.08	8.25	11.85
2	I ₂	LV(-7.71,0.66)	LV(-7.71,0.66)	LV(-7.71,0.66)	15.08	28.91	44.84	55.85	133.33
3	I ₃	Noninformative	LV(-10.01,0.59)	LV(-10.01,0.59)	1.70	3.16	4.48	5.33	11.83
4	I ₄	Noninformative	Noninformative	G(0.5,40)	4.92	---	568.67	1250.00	4801.68
5	I ₁ ,I ₃	Noninformative	LV(-10.01,0.59)	(Numerical)	4.73	---	7.18	7.33	10.45
6	I ₁ ,I ₄	Noninformative	Noninformative	G(16.5,200040)	5.22	7.75	8.08	8.25	11.85
7	I ₁ ,I ₂	LV(-7.71,0.66)	LV(-7.71,0.66)	(Numerical)	6.61	---	9.61	9.76	13.45
8	I ₂ ,I ₄	LV(-7.71,0.66)	LV(-7.71,0.66)	(Numerical)	15.07	---	44.84	55.85	133.29
9	I ₃ ,I ₄	Noninformative	LV(-10.01,0.59)	(Numerical)	1.70	---	4.48	5.33	11.83
10	I ₂ ,I ₃	LV(-7.71,0.66)	LV(-8.99,0.44)	(-8.99,0.44)	6.01	10.22	12.42	13.69	25.64
11	I ₁ ,I ₂ ,I ₄	LV(-7.71,0.66)	LV(-7.71,0.66)	(Numerical)	6.61	---	9.60	9.76	13.45
12	I ₁ ,I ₂ ,I ₃	LV(-7.71,0.66)	LV(-8.99,0.44)	(Numerical)	6.11	---	8.77	8.91	12.19
13	I ₁ ,I ₃ ,I ₄	Noninformative	LV(-10.01,0.59)	(Numerical)	4.73	---	7.18	7.33	10.45
14	I ₂ ,I ₃ ,I ₄	LV(-7.71,0.66)	LV(-8.99,0.44)	(Numerical)	6.01	---	12.42	13.69	25.64
15	I ₁ ,I ₂ ,I ₃ ,I ₄	LV(-7.71,0.66)	LV(-8.99,0.44)	(Numerical)	6.11	---	8.77	8.91	12.19
16	I ₁ ,I ₂ ,I ₃ ,I ₄	LV(-7.71,0.66)	LV(-8.99,0.44)	LV(-9.35,0.21) Approximation	6.21	8.36	8.72	8.91	12.21

^aAll entries times 10⁻⁵ f/dam-year.

Fig. 1. The two-stage Bayes/empirical Bayes data pooling procedure.

Fig. 2. The numerically-derived posterior distribution and lognormal approximation based on the combined use of I_1 , I_2 , I_3 , and I_4 .