TABLE OF CONTENTS

Figures (continued)

A-8. COMBINING AND PORTRAYING RISKS

A-8.1   Key Concepts

After all potential failure modes (PFM) have been identified and described, and their risks have been evaluated, the results need to be combined and portrayed so that the technical reviewers and decision makers can understand and act upon them. This requires attention to detail, and if not undertaken properly, could result in an incorrect portrayal of the risk. This chapter describes some of the details needed to properly do the job.

A risk analysis, whether by a team or by an individual, produces estimates of risk for individual potential failure modes. These estimates might include probability or risk values for different loading conditions, loading ranges, spatial segments, or other situations. The risks from individual potential failure modes are often combined in some way to express their collective effect.

In combining risk estimates, some of the most common problems encountered during risk analyses are related to systems, correlations, and common-cause loading. Although the methods to evaluate these issues can become complex, some simplifications can be applied to situations commonly seen when evaluating risks for dams and levees.

A-8.2   System Considerations

For the purposes of dam and levee safety risk analysis, a dam or levee system is typically defined to include all components of the project that are intended to retain the reservoir (dams) or exclude water from the leveed area. In both cases, the system considers those components that can affect a common consequence center(s) within the associated floodplain.

A-8.3   Mutually Exclusive Events

The probability of the union of two or more mutually exclusive events is equal to the sum of their probabilities. Similarly, risks associated with mutually exclusive events can be directly summed.   In an event tree, all branches originating from the same node are mutually exclusive. This allows the Annualized Failure Probability (AFP) and Annualized Life Loss (ALL) values at the end branches of an event tree to be summed to obtain the total risk for an individual potential failure mode or group of potential failure modes. Similar event tree summations can be performed to obtain the risks associated with a particular loading range (e.g., floods having a return period more frequent than 100 years), a physical

feature (e.g., spillway), or a particular breach mechanism. In summing risk estimates developed using separate event trees, it is common practice to ignore the intersection between the PFMs, even though there usually is one. In most cases this is a reasonable assumption because the "size" of the intersection is small. As a result, the risks associated with floods and earthquakes, for example, can usually be evaluated using separate event trees and the results summed directly to obtain a reasonable estimate of the total risk.

A-8.4   Collectively Exhaustive Events

Events that are collectively exhaustive include all possible outcomes. In an event tree, all branches originating from the same node are collectively exhaustive. The probabilities associated with all branches originating from a particular node must sum to 1.0. This can be used as a check that the event tree has been properly constructed.

A-8.5   Statistically Independent Events

Two or more events are statistically independent (SI) if the occurrence of one event does not affect the probability for occurrence of the other event(s). The definition of statistical independence means that the probability of one event can be estimated without explicitly considering whether the other event(s) has occurred or not. In practice, a potential failure mode is often developed under the assumption that the PFM is statistically independent of other potential failure modes. Correlation and/or common cause effects, if they are relevant, can be accounted for by adjusting the total probability estimates and/or modifying the event tree structure.

A-8.6   Conditional Probability

Two or more events are statistically dependent if the occurrence of one event affects the probability for occurrence of the other event(s). For such events, the occurrence of the conditioning event must be considered when estimating the conditional probability of the dependent event. Conditional probabilities of failure are often dependent on the magnitude of the load because a greater load will typically result in a greater probability of failure. Event trees are constructed such that conditional probabilities are shown to the right of the events on which their probabilities are conditioned. Note that the product of the probability estimates to the right of the load probability is termed the conditional failure probability.

A-8.7   Correlated Events

Correlation is the degree to which the probabilities for two or more events are linearly related. For correlated events, the occurrence of one event is an indication that the other event is also likely to occur (positive correlation or dependence) or likely to not occur (negative correlation or dependence). Probabilities associated with dam/levee components of similar character (e.g., spillway gates, concrete monoliths) might be correlated. If one spillway gate fails to operate, then it may be likely that additional spillway gates will also fail to operate. Correlation can be quantitatively accounted for in the risk analysis using correlation matrices or more qualitatively accounted for by applying expert judgment to the estimated probabilities associated with the responses of groups of similar components.

A-8.8   Total System Probability

The uni-modal bounds theorem (Ang and Tang 1984) stated that for 'n' positively correlated events (E1, E2, E3, …, En) with corresponding probabilities [P(E1), P(E2), P(E3), …, P(En)], the total probability P(E) for the union of the events, which cannot always be calculated directly by using the formula [P(E) = P(E1 $\cup$ E2 $\cup$ E3 …$\cup$ En)], lies between the upper and lower bounds given by the following equation:

$$\max[P(E1), P(E2), P(E3), ..., P(En)] \leq P(E) \leq 1 - \prod_{i=1}^{n}[1 - P(E_i)] \qquad \textbf{Equation A-8.1}$$

The upper bound on the right side of the equation is based on a calculation of the total probability using DeMorgan's rule. The lower bound on the left side of the equation is tied to the event with the largest individual probability.

The uni-modal bounds theorem can be used to obtain the upper and lower bounds of the total probability of failure for a dam or levee system from a set of failure probabilities associated with individual potential failure modes. The lower bound is obtained if the potential failure modes are perfectly correlated. The upper bound is obtained if the potential failure modes are statistically independent. In practice, the upper bound is often used in dam and levee safety risk analysis unless specific knowledge of the degree of positive correlation is available. If one potential failure mode is dominant (i.e., has a probability significantly greater than that of all other failure modes), then the upper bound and lower bound obtained from the above equation will be approximately equal to each other.

If the conditional failure probabilities for individual potential failure modes are small compared to 1, the upper bound of their union can be approximated by summing the individual conditional failure probabilities. For a set of positively correlated PFMs, the maximum possible error in this approximation will be less than about 5% if the sum of the individual conditional failure probabilities is less than about 0.1. In practice, the error is typically much less than this maximum value. For two PFMs with conditional failure probabilities of 0.02 and 0.08, the summation would be 0.1 and the upper bound would be 0.098 resulting in an error of less than 2%. As the sum increases, the error in the approximation also increases. When the sum of the conditional probabilities is equal to 1, the maximum possible error is about 60%. When the sum of the conditional probability estimates becomes greater than 1.0 (a situation that would imply certain failure under the given loading condition), the summing approximation violates the axioms of probability and is probably not appropriate to use.

Consider a dam with the following three seismic potential failure modes: A) sliding within the foundation of a concrete gravity monolith, B) buckling of a spillway gate arm, and C) liquefaction of the embankment foundation leading to crest deformation and overtopping. The conditional probability of failure for each of these potential failure modes given a 0.001 Annual Exceedance Probability (AEP) seismic loading has been estimated independently and is shown below. Note that in practice the seismic loading would typically be divided into multiple loading partitions with conditional failure probabilities for each PFM estimated separately for each partition.

$$P(A) = 0.3$$
$$P(B) = 0.1$$
$$P(C) = 0.2$$

Equation A-8.2

Applying the uni-modal bounds theorem, the total conditional probability of failure for the dam is estimated to be between 0.3 and 0.496.

$$\max\{0.3, 0.1, 0.2\} \leq P(\text{Fail})$$
$$\leq 1$$
$$- (1 - 0.3)(1$$
$$- 0.1)(1 - 0.2)$$

Equation A-8.3

$$0.3 \leq P(\text{Fail}) \leq 0.496$$

These values place bounds on the total conditional probability of failure estimated for the dam. In contrast, a summation of the individual probabilities of failure would result in an estimate of 0.6. Using the summation approximation would in this case overestimate the total conditional probability of failure by about 20%. The total AFP for the dam would be about 5E-4 using the upper bound estimate, 3E-4 using the lower bound estimate, and about 6E-4 using the summation approximation. The difference between the AFP estimates in this example would

typically have a minimal impact on risk informed decisions. In this case, the summation approximation would likely be a reasonable simplification for portraying the total risk in this example.

The upper bound estimate is represented by the shaded area on the Venn diagram in figure A-8-1. The estimate includes the total area enclosed by all of the circles, each representing an individual potential failure mode (from a set of SI PFMs). The uni-modal bounds equation calculates a range of total areas using the limiting cases of circle overlap. The summation approximation, which does not consider the intersection area, is subject to double counting error (i.e., the overlapping areas on the Venn diagram are counted more than once). The double counting of the overlapping area is the source of the 20% error in the example.
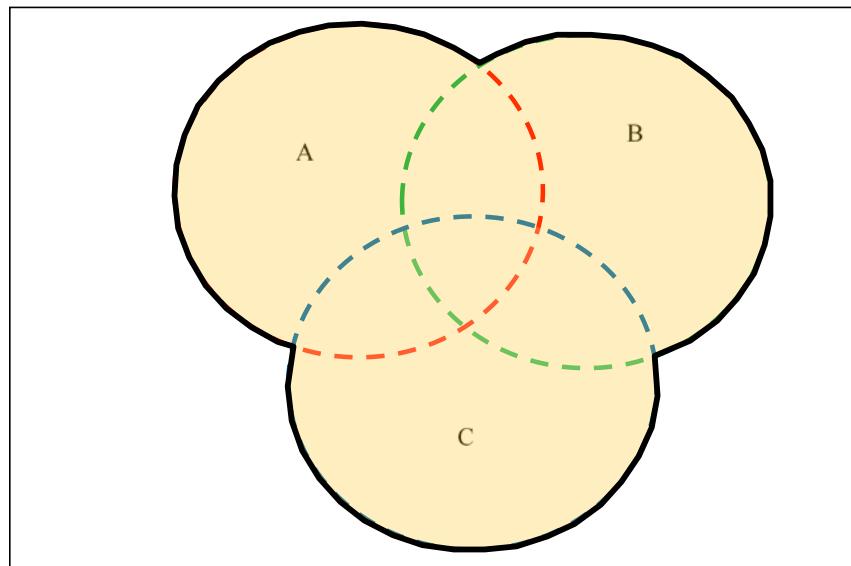

Figure A-8-1.—Venn diagram for upper bound estimate.

A-8.9   Common Cause Adjustment

In typical dam and levee safety risk analyses, intersection events representing the occurrence of two or more potential failure modes are not explicitly evaluated in the event trees. This is usually, but not always, a reasonable simplification. If the probabilities of the intersection events are small relative to the probabilities of each potential failure mode, then the intersection event probabilities can be ignored. This allows the potential failure mode probabilities to be summed to obtain a reasonable approximation of the total probability of failure. When the intersection probabilities are not small, adjustments to account for the over counting of the intersection probability may need to be made so that the correct total probability of failure can be obtained. The term "common cause adjustment" refers to any of the methods that can be used to make this correction.

In dam safety risk analysis, the AFP associated with a simple adverse event chain (e.g., A happens, B happens, C happens, D happens ...) is calculated as the probability of the intersection event P(ABCD...). For a dam susceptible to multiple adverse event chains (i.e., potential failure modes), the total probability of failure would be properly calculated as the probability of the union event. For example, given potential failure modes (PFMs) 1 and 2 with calculated annualized failure probabilities AFP1 and AFP2, the total AFP would be calculated as:

$$P(PFM\ 1 \cup PFM\ 2) = P(PFM\ 1) +$$
$$P(PFM\ 2) - P(PFM\ 1 \cap PFM\ 2)$$
$$= AFP1 + AFP\ 2 - \epsilon$$

Equation A-8.4

where ϵ is a number between zero and the smaller of (AFP1, AFP2). If PFMs 1 and 2 are SI, ϵ is usually small enough to be ignored without inflating the total risk estimate (<u>not</u> because SI events have trivial intersections, but because P(PFM 1 $\cap$ PFM 2) typically reduces to the product of two small numbers for SI PFMs). However, if PFMs 1 and 2 are not SI at the level of the overall sample space (for example, if both PFMs involve the occurrence of a flood or earthquake), it may not be possible to simply assume that ϵ is negligible.

Consider a set of seismic related potential failure modes, PFMs 1 and 2. The risks for both PFMs are controlled by the occurrence of a 50,000-year earthquake. Given the occurrence of the earthquake (Event Q), the probability of breach due to embankment liquefaction (response event A) is estimated to be 0.5 (= P[A|Q]) and the probability of breach due to the gravity section sliding (response event B) as 0.9 (= P[A|Q]). Because the response events involve completely different mechanisms of failure, they could be considered independent given the occurrence of triggering event, or "conditionally independent" (Galic 2017). Since the rules of probability theory still apply within a reconditioned sample space, the fact that the conditional probabilities of A and B sum to greater than 1 implies that within the reconditioned sample space, there is intersection between the events. Note that this does not imply there is anything "wrong" with the team's conditional probability estimates.

Figure A-8-2 shows the Venn diagram for the above example, both before (left) and after (right) the sample space transformation associated with the occurrence of the 50,000-year earthquake. Prior to the occurrence of the quake, the "area" of S occupied by Q is relatively small (since the earthquake has only about a 1/50,000 chance of occurring). However, once it is known that the earthquake has occurred, the sample space changes from all of S to only the region bounded by Q. Although the "size" of AB given Q is not obvious from the sum of the response probabilities, the conditional probability of A $\cap$ B can be estimated using the information already known, provided that A and B are statistically independent within the reconditioned sample space (i.e., conditionally independent). Once quantified, the conditional probability of the intersection can be used to correct the total AFP estimate. The process is as follows:

Step 1.   Calculate P(AB|Q): P(AB|Q) = P(A|Q)*P(B|Q) = 0.5*0.9 = 0.45

Step 2.   Calculate the uncorrected total AFP: AFP = P(Q)*P(A|Q)+P(Q)*P(B|Q)
          = 1/50,000*(0.5+0.9) = 1.4/50,000

Equation A-8.5

Step 3.   Calculate the corrected total AFP: AFP = P(Q)*[P(A|Q)+P(B|Q) –
          P(AB|Q)] = 1/50,000 (0.5+0.9-0.45) = 0.95/50,000

Step 4.   Calculate the percent error: Error = 1 – (corrected AFP/uncorrected
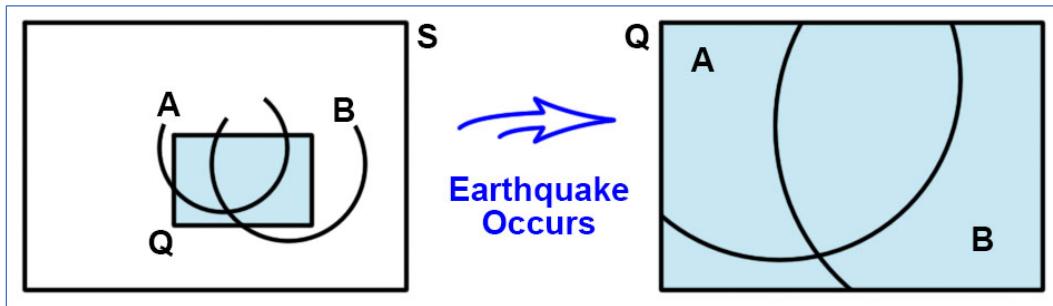   AFP) = 1– 0.95/1.4 = 0.32

Figure A-8-2.—The probabilities of events A (e.g., breach due to liquefaction mechanism) and B (e.g., breach due to gravity section sliding mechanism) within the overall sample space (left panel) and the reconditioned sample space associated with the occurrence of event Q, a 50,000-year earthquake (right panel).

The final step would be to decide whether a formal AFP correction needed to be reported. Although the error calculated in the example may appear significant, the potential impact on the dam safety decision (or lack thereof) must be taken into account (as well as the fact that risks are typically plotted on a log scale). In this example, correction for the 32% error would result in a relatively small difference in the plotting position if an fN chart ("little f-n") were being used to plot the risks. If it was decided to report the correction the final step in the process would be to change the automatically summed AFP total to the number obtained in Step 3. The individual PFM risk estimates are already "correct", and would not need to be adjusted on the fN chart. The process outlined in the example can be generalized to any number of PFMs, or any number of load ranges (for multiple controlling load ranges, a separate adjustment would be performed for each load range). It should be stressed that the process outlined in the example is only applicable to conditionally independent PFMs (or those that are essentially so).

The considerations that sometimes lead to a total AFP adjustment are also applicable to the calculation of total An ALL. However, because the ALL is a different kind of mathematical object than the AFP (an expectation versus a probability) the process used to adjust the ALL is slightly different, and depends

on whether the PFMs involve different or similar life loss estimates, and on whether the populations at risk overlap or are unique to each PFM (Galic 2017). For the purposes of this introductory course on risk analysis, the ALL Common Cause Adjustment can be summarized as follows:

- When the life loss estimates for the PFMs over which the adjustment is being performed are identical or reasonably similar, and when the inundation area associated with each of the PFMs is identical or reasonably similar, the corrected total ALL can be obtained simply by multiplying the corrected total AFP by the estimated life loss. The percent error in the ALL correction will be the same as for the total AFP in this case.

- When the life loss estimates associated with the PFMs over which the adjustment is being performed are additive (for example, when considering a pair of dikes constructed over different drainages), no ALL adjustment is required, and the total ALL can be calculated from the individual PFM risk estimates (which are already "correct", and therefore not subject to adjustment).

- When the life loss estimates for the PFMs over which the adjustment is being performed differ significantly, a specific value of life loss must be assigned to the intersection event (Galic 2017). However, in most cases, the percent error for this situation would be relatively small, and there would typically not be a compelling reason to perform a common cause adjustment for the total ALL.

As demonstrated above, the intersection between PFMs may support a decision to adjust the total AFP and ALL estimates. In typical risk assessments, only the total AFP and ALL should be adjusted. Individual PFM risk estimates are typically not adjusted. However, if the intersection event has the potential to be a significant contributor to the total AFP or ALL with potential impacts to a decision, the risk analyst might consider portraying the intersection event separately. In this situation, the individual PFMs A and B would need to be adjusted so as to not include the intersection event, which could be done with the help of the complementary events ¬A (read "not A") and ¬B. In the example, if the life loss is estimated to be about 10 for event A¬BQ, 20 for event ¬ABQ, and 50 for event ABQ, then the intersection event accounts for about 47% of the total AFP and about 80% of the total ALL. The risk analyst could choose to portray event A¬BQ with AFP of 9E-6 and ALL as 9E-5; event ¬ABQ with AFP of 1E-6 and ALL of 2E-5, and event ABQ with AFP of 9E-6 and ALL of 5E-4. This would convey that the intersection event (e.g., multiple failures occur as a result of the earthquake) is the most likely outcome and has the highest risk.

Hill et al (2003) have proposed a simplified approach for adjusting the system response probabilities for each potential failure mode as a way of adjusting the total AFP and ALL. The method redistributes the overlapping area in figure A-8-1 to each individual failure mode (or response event, in the case of a reconditioned sample space). The magnitude of the redistribution is proportional to the estimated probability of failure for each potential failure mode (or the conditional probability of failure for a common load condition). Events with larger probabilities of failure receive a larger portion of the overlapping area. The approach is implemented using the following equation, where p_j is the unadjusted probability of failure for potential failure mode (or response event) j and p'_j is the adjusted probability of failure.

$$p'_j = p_j \frac{1 - \prod_{i=1}^{n}[1 - p_i]}{\sum_{i=1}^{n} p_i} \qquad \textbf{Equation A-8.6}$$

For the example on page A-8-4, the adjusted probabilities of failure are

$$p'_A = 0.3 \frac{1 - (1 - 0.3)(1 - 0.1)(1 - 0.2)}{0.3 + 0.1 + 0.2} = 0.248$$

$$p'_B = 0.1 \frac{1 - (1 - 0.3)(1 - 0.1)(1 - 0.2)}{0.3 + 0.1 + 0.2} = 0.083$$

$$p'_C = 0.2 \frac{1 - (1 - 0.3)(1 - 0.1)(1 - 0.2)}{0.3 + 0.1 + 0.2} = 0.165$$

The adjusted failure probabilities could now be added to calculate a total AFP, or multiplied by the associated life loss estimates and added to obtain a total ALL to portray on an fN chart. In typical risk assessments, adjustments to the individual PFMs would not be portrayed on fN charts but, the adjusted individual failure probabilities are required when an FN chart is being used to portray the risks.

Additional details regarding this approach are provided in Hill et al (2003).

A-8.10   Length Effects

Dam and levee systems may be comprised of significant lengths of constructed embankments or walls extending thousands of feet to hundreds of miles. This may result in considerable uncertainty about the loadings, performance, and consequences for sections within the system.

Systems fail at locations where loads are high and strengths are low. If these critical locations are known and identified ahead of time, traditional methods can be used to analyze stability and estimate probabilities of failure. In such

situations, the overall length of the system is immaterial, because the weakest spots have been identified, and the performance of the system depends on the probability of failure for the weak spots. The more common situation is that the system is not characterized with enough detail for the risk analyst to know with reasonable certainty where the weakest spots are. In this case, any section of the system has some probability of experiencing higher than average loads and/or lower than average strengths. Since these locations cannot be uniquely identified before a failure occurs, a longer system length results in a greater probability of a failure somewhere.

A detailed discussion of length effects is beyond the scope of this manual. Risk analysts should consult with appropriate experts when estimating risks for long systems or for systems with many components (e.g., a levee with many pipe penetrations).

A-8.11   Cascading Events

Individual dams and levees are often part of larger infrastructure systems. Within these watershed systems, risk is typically attributed to the specific dam or levee that is the source of the risk.

Estimating and attributing risk to an individual dam or levee can sometimes be complicated by system effects. Failure of a dam or levee might impact the performance of other dams or levees in the system. Breach of an upstream levee might reduce the loading on a downstream levee. Failure of an upstream dam might result in overtopping of a downstream dam. In these situations, agency specific policy and methodology will dictate the scenarios that need to be evaluated in order to estimate, attribute, and portray risks.

To support portfolio prioritization decisions or to communicate the flood risk from multiple flooding sources, there may be a benefit in estimating the risk from a river systems perspective. These analyses can support improved prioritization decisions within the larger watershed to obtain more efficient and effective total risk reduction across the portfolio. In these situations, it may be appropriate to evaluate the cascading impacts of failure.

A-8.12   Intervention

The ability to intervene to mitigate or prevent failure of a dam or levee is an important consideration in risk analysis. The potential for successful intervention can be important in setting priorities across a portfolio and in developing specific

risk reduction actions at a dam or levee. Risk analysts should refer to agency specific policy and methodology for information on how intervention should be considered in the portrayal of risk.

A-8.13   Incremental Risk

Risks can be estimated and attributed to a PFM based on the total consequences associated with the failure of a particular dam or levee. For example, the risk for a levee overtopping PFM might include an AFP of 0.0002, a total life loss of 25, and an ALL of 0.005. In this scenario, all of the risk associated with the failure event would be attributed to the levee.

Risks can also be estimated and attributed based on the incremental consequences associated with a failure. The incremental consequences are calculated as the difference between the total consequences that occur with a failure and the consequences that would have occurred for the same event if the structure had not failed. If the same levee overtops and does not fail, the life loss estimate may be 10. The incremental risk associated with the failure would include an AFP of 0.0002, an incremental life loss of 15, and an ALL of 0.003. In this scenario, only the risk associated with the actual failure of the levee is attributed to the levee.

Risk analysts should refer to agency specific policy and methodology for whether or not to consider incremental consequences in the estimation and portrayal of risk.

A-8.14   System Response Curves

In dam and levee risk analysis, probability of failure estimates associated with system response curves are often interpreted as the products of conditional breach probabilities and specific load probabilities. This is a simplification that is made to facilitate event tree analysis and event tree calculations. This approach assumes that breach will occur at the maximum loading during a flood or earthquake event. It is important to recognize that a system response curve is actually a cumulative distribution function for the capacity or strength of the dam or levee. In practice, this means that the system response curve gives the probability that a failure or breach will occur at a load that is less than or equal to the specified load. This means that breach can occur at a loading that is less than the maximum loading experienced during a specific flood or earthquake event.

Characterizing the system response curve as a cumulative distribution function provides an important and powerful tool for simulation based risk models. For each simulation, the capacity of the system can be randomly sampled from the system response curve prior to the start of the simulation. During the simulation,

the modeled failures or breaches occur when the demand reaches the sampled capacity. In this framework, it is not necessary to assume that failure or breach always occurs at the peak demand. Not making this assumption allows for explicit consideration of the temporal aspects of loading events such as floods and explicit analysis of scenarios where a failure or breach can occur at a water level that is less than the peak water level.

The branches in an event tree do not explicitly consider the temporal aspects of events like floods, and it is often assumed that failure or breach will occur at the peak load. The risk analyst should keep in mind that failure or breach can also occur at loads less than the peak load for the event. This simplification is usually not a significant issue when estimating or portraying the total annual probability of failure or breach. It can sometimes be an issue when estimating or portraying the total ALL. In certain situations, failure at a load less than the peak load could result in different consequences and different risks. In these situations, the risk analyst should consider whether or not further refinement of the risk model is needed to obtain a more accurate portrayal of the risk for decision makers.

A-8.15   Monte Carlo Simulation

In dam safety risk analysis, the intersection formula or "multiplication rule" is used to compute the Annualized Failure Probability (AFP) associated with each PFM. The intersection formula uses conditional probabilities as its primary inputs, and since the probability estimates developed throughout a risk analysis are conditional probability estimates, the AFP (e.g., for a given PFM load range) often reduces to a simple product of the probabilities estimated by the team. However, due to the uncertainty inherent in subjective probability estimation, the conditional probability estimates are usually developed as ranges or distributions, rather than as single values. In this case, the probability estimates are themselves random variables, and the AFP obtained from the intersection formula is a product function of the n random variables that comprise the PFM.

In order to obtain an estimate of the mean of a product distribution, it would first be necessary to derive an equation for the PDF of the distribution of the product function. Unfortunately, this is difficult to accomplish unless the input random variables are jointly lognormal. Since this is not typically the case with subjective probability estimates, Monte Carlo simulation (e.g., Ang and Tang 1984) is used instead to approximate the analytical product distribution. The process can be implemented through the use of specialized software, often available in the form of an add-on to Microsoft Excel, that allows input distributions to be entered in terms of their key parameters. For a given trial, the Monte Carlo simulation process consists of randomly sampling the input probability distributions associated with a PFM, and passing each set of sampled values through the

intersection formula to obtain a trial value of AFP. Over thousands of trials, an AFP output distribution, whose mean can be calculated numerically, is built up.

The Monte Carlo generated AFP output distribution is often approximately log-normal (figure A-8-3), consistent with the predictions of the Central Limit Theorem for product functions of random variables (e.g., Ang and Tang 1975).
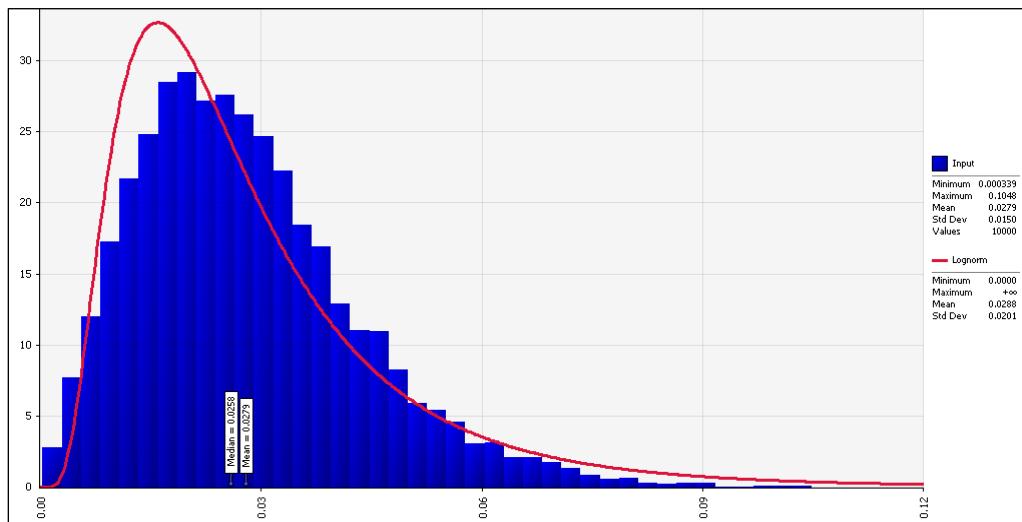

Figure A-8-3.—Comparison of a Monte Carlo AFP output distribution (blue) and the PDF of an analytical lognormal distribution (red line).

However, this may not always be the case, especially if the event tree contains many "short" branches; when summation, rather than multiplication, is the dominant event tree operation, the output distribution tends toward a normal distribution shape. For consequence estimates, the shape of the output distribution can be log-normal (as in the case of an ALL distribution obtained from a simple event tree), normal (e.g., when the ALL is obtained as a sum across many branches), or neither (e.g., when a life loss distribution is back-calculated from the ALL and AFP).

A-8.16   Informing Decisions

Risk analysis results are typically portrayed on an f-N or F-N chart. The usual format for an f-N plot features an annual failure probability on the vertical axis and the estimated life loss by PFM on the horizontal axis. Note that the latter is not the same as the ALL, whose value is read in the up-right diagonal direction. The F-N plot features an annual probability of N or more consequences on the vertical axis and the magnitude of consequences on the horizontal axis. In both plots, the axes are shown with a log scale.

Another type of risk plot is the scatter plot (e.g., figure A-8-4), which uses the same log-scale axes as the fN chart and can also be generated using a standardized template. In order to obtain the data required for a scatter plot, a Monte Carlo simulation must first be performed, as discussed below. Figure A-8-4 is an example of a scatter plot produced from the results of individual Monte Carlo trials for a dam where the slip rate of a local fault was unknown. In this case, a recommendation was made not to investigate the fault any further because its existence would not have a significant effect on the final decision. The point clouds for both slip rate assumptions (red versus blue) are similarly shaped and plot over nearly the same area, suggesting that further refinement of the seismology would not change the decision.
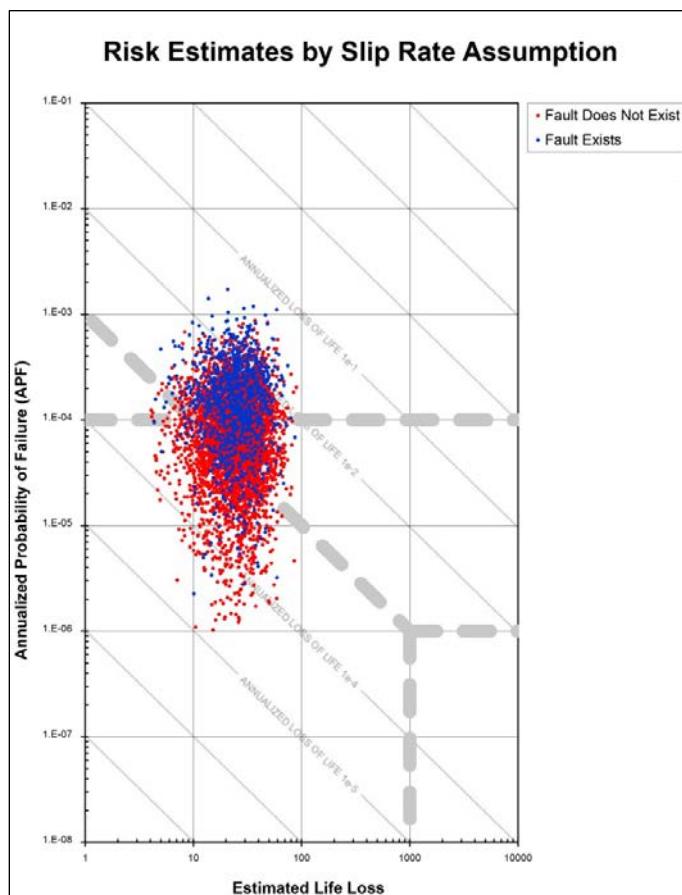

Figure A-8-4.—f–N scatter plot used to make a case against fault investigation.

Many types of charts can be generated depending on the information needing to be conveyed. Some of the common charts used as supplementary information are:

- Charts that show the relative contribution to the total risk from each potential failure mode

- Charts that show the contribution to risk from each load range or slip rate assumption (as in figure A-8-5 for example)

- Charts that show effects of applying different flood-related assumptions (showing the value of additional hazard studies)

- Charts that show the effects of different foundation assumptions (showing the value of additional geologic investigation)
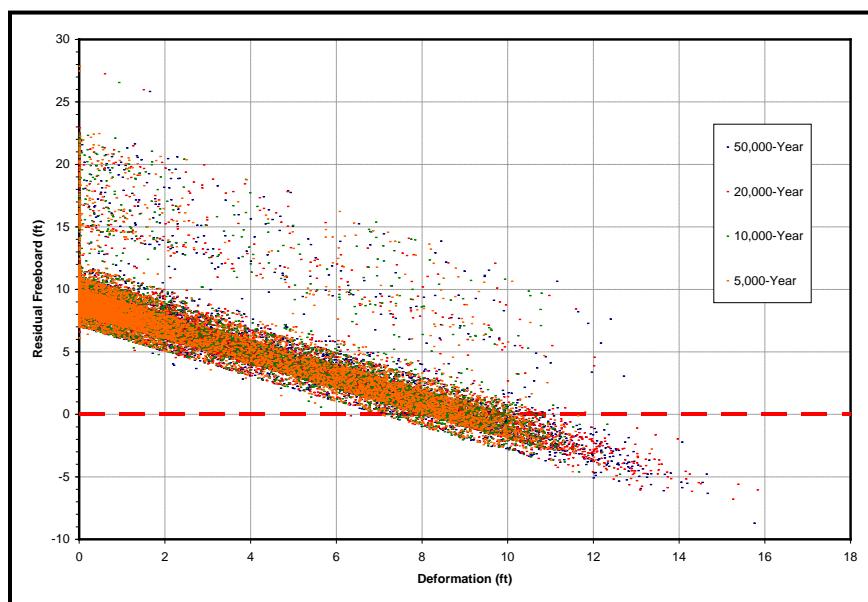

Figure A-8-5.—Contributions to risk by load range.

The spreadsheets used to create the standard risk plots often require the user to enter plausible ranges for the probabilities of failure and consequences. The lower and upper bounds of these ranges can be estimated using several different approaches, and their use helps communicate the uncertainty of the risk estimates. In each case, the meaning of and reason for the uncertainty bounds used should be explained in the report. Some of the more commonly used lower and upper bounds are:

- +/- 1 standard deviation

- The 5th and 95th percentiles of the Monte Carlo output

- The 1st and 99th percentiles of the Monte Carlo output

- The absolute bounds of the entire Monte Carlo range (i.e., 0 and 100 percentiles)

- +/- one order of magnitude from the mean estimate

- The range over which the mean estimate could reasonably change with additional information or with more refined analysis

The report should always explain how and why a particular set of uncertainty bounds was selected. The reporting of uncertainty bounds is always recommended, but they do not by themselves make a compelling case for the risk estimates. Identifying the separate components of risk, discussing the meaning and importance of a particular load range, and describing the sensitivity of the results to a key probability estimate are examples of other things could be done to help build the case for a particular interpretation of risk. The key question would be whether the uncertainty of the probability estimates has the potential to affect the overall dam safety case.

A-8.17   References

Ang, A. and W.H. Tang. 1975. Probability Concepts in Engineering, Vol.1. John Wiley and Sons, New York.

_____. 1984. Probability Concepts in Engineering, Vol.2. John Wiley and Sons, New York.

Galic, D. 2017. "The Common Cause Adjustment in Dam Safety Risk Analysis: Scope, Purpose, and Applicability." *Proceedings of the 2017 Annual USSD Conference*. Anaheim, California.

Hill, P., D. Bowles, P. Jordan, and R. Nathan. 2003. Estimating Overall Risk of Dam Failure: Practical Considerations in Combining Failure Probabilities. Australian National Committee on Large Dams.