

## I-9 Building the Case

Though many efforts are made during a risk analysis to achieve high quality results, the risk estimates themselves are little more than index values. If arrived at in a consistent manner, they are useful in program management to allow comparisons and rankings between different facilities. However, the Public Protection Guidelines (Reclamation) and the Tolerable Risk Guidelines (USACE) were never intended to be used as rigid decision-making criteria to declare a facility 'safe' solely based on a risk estimate. Since the numbers are neither accurate nor precise measures of risk, and the tolerable risk guidelines themselves are somewhat flexible, reasoning is essential to justify recommended actions. The Dam Safety Case is intended to present rationale in a formal and methodical manner to persuade decision makers to take responsible action.

### Key Concepts

The Dam Safety Case is a logical set of arguments used to advocate a position that either additional safety-related action is justified, or that no additional safety-related action is justified. The arguments string together key evidence regarding the three basic risk components, (i.e. load probability, response probability, and consequences) so as to convince decision-makers that the dam's existing condition and ability to withstand future loading, the risk estimates, and the recommended actions are all coherent. Since uncertainty is inherent in each claim, the arguments should also address whether confidence is high enough to stand on the basis of existing evidence.

The dam safety case and the identification of risk management options are recognized as essential elements in Reclamation's project-ranking efforts to ensure public protection. They represent understanding of existing condition and predicted future behavior stated as objectively as possible. The dam safety case should not be used as a means of back-fitting an argument for design decisions or business decisions that have already been made.

The risk estimates and the dam safety case do not in themselves ensure the safety of a facility. The dam safety case becomes the basis for risk management in the effect it has on the activities and behaviors of the people who interact with the facility. The understanding given to all, from facility operators to caretaker engineers to dam safety program managers to Reclamation directors, by a well-constructed dam safety case is intended to focus attention on behavioral and technical aspects essential to the facility's integrity so that the facility can be operated and maintained in a safe manner.

The process of analyzing safety requires creativity and judgment. It requires an extensive understanding of the facility, its behavior in a variety of conditions, experience of failures in other facilities, and the measures adopted to prevent their recurrence. Creativity is required to notice and identify design, construction, and behavioral weaknesses peculiar to the facility's site, and then to synthesize conclusions and craft the argument in the most



coherent way. Judgment takes the form of event likelihood estimates and is based upon basic knowledge of soil, rock, and structural mechanics. Experience in dam construction and with forensic investigations of dam failures and safety incidents bolsters judgment by providing patterns that can be recognized as good or bad engineering practice.

The dam safety case should be carefully crafted so that all descriptions and terms are easy to understand by the prime audience, all arguments are cogent and coherently developed, all references are easily accessible, and all conclusions are fully supported and follow logically from the arguments.

## **Building from Simple Arguments**

The dam safety case is built up from a number of arguments successively demonstrated to be valid. A simple argument consists of a single claim, evidence to support that claim, and reasoning to suggest how and why the evidence justifies the claim. An example of a simple argument would be to claim that: "Zone 1 core material is very likely filtered by Zone 2 shell material." Evidence would include gradation tests from both materials and the number of tests (among other things). One could infer that the evidence supports the claim if the gradations meet certain filter criteria and if the number of tests would be sufficient to account for variability. All geotechnical engineers should understand filter criteria concepts and most accept that if the gradations are such that one or another of the standard filter criteria is met, it is not very likely material will be capable of moving from one zone to the other. However, if the variability is large enough or if one suspects that poor construction would cause a localized anomaly, one might have to provide reasoning to convince others that the number of tests is statistically significant to infer the likelihood portion of the claim.

The process described in Section 8 on Subjective Probability Elicitation is used to categorize all known evidence about an event tree event into 'likely' and 'unlikely' categories. The act of assigning a probability estimate to the event implies that a side is being taken on the event likelihood (unless 0.5 is chosen). Best practice suggests that the facilitator should have the risk estimating team highlight the key pieces of evidence that made the team choose the number they chose. The risk analysis report's author is then better able to make a simple argument to defend the probability estimate.

The degree of belief in the argument's strength is reflected in the probability assigned to the event's likelihood and vice versa. The argument's strength depends on the weight of evidence that is presented to either support or counter the claim. In the example above, having a large number of gradation tests that firmly demonstrate the filter criteria are met would weigh heavily. Evidence this strong would lead to a very low probability estimate for the event: "Zone 2 does not filter Zone 1". Establishing strong belief in this claim also creates a powerful argument in the dam safety case finding that "there is diminished justification to take action to reduce risk" for failure modes involving erosion along seepage paths entirely through the embankment.

Several simple arguments are strung together to form the basic structure of the dam safety case. Arguments are ordered in various ways to help make the argument more sensible or convincing. For example, the event tree is structured to present a number of claims in

series, with each claim having its likelihood or probability estimate. The series structure is most convincing when each probability estimate in the string is well substantiated. Conversely, if only one claim in the string is highly uncertain, the case can become much weaker. A claim that is particularly weak is a good candidate for a sensitivity study. The risk analysis can be evaluated assuming reasonable upper and lower probability estimated bounds to see how the ultimate risk estimate is affected. If the affects are significant, an argument can be made to obtain additional information.

Sometimes, two or more claims can be made, any one of which, if highly convincing, would substantiate the safety case argument. Sometimes, several independent claims combine to justify the recommended action when none by itself can sway the argument. In this structure, the greater the number of claims that can be established, the stronger the case becomes to carry the argument.

Three basic forms that evidence-building take when likelihoods or probabilities are judged are: 1) conclusions based upon numerical models from first principles of physics (theory and analysis), 2) statistical analysis results generated from empirical data, and 3) the informed judgment of experts.

Each evidence-building form has its questionable qualities. Regarding numerical models, Vick (2002, pp 56-64) argues: "Neither theory and analysis nor their predictive results can be taken as uniquely or objectively true because they are unverifiable, not unique, incomplete, indeterminate, and transient." Inferences from statistical regression formulas break down when population or sample sizes are too small and when site-specific characteristics poorly match those of the population upon which forecasting regressions are based. Probability estimates from expert elicitation can be highly questionable if those providing the judgment are not particularly experienced or intuitive, or are not particularly well informed. All three forms fail to predict the unforeseen such as Peck's oddball (Peck, 1998) or Taleb's black swan (Taleb, 2008).

Explicit treatment of uncertainty can help mitigate the overconfidence inherent in single-value estimates. Uncertainty's role in case-building is discussed below.

## **Basic Findings, the Object of the Dam Safety Case**

At Reclamation, a Technical Report of Findings is written for each Issue Evaluation and each Comprehensive Review. A finding is a statement that advocates the position being taken on the basic dam safety question regarding what action(s) should be taken. The findings are couched in risk-based decision terminology and are usually explicit regarding uncertainty. Some typical findings might include:

1. The estimated risk is tolerable, and confidence is high so that no further actions or studies are necessary.
2. The estimated risk is tolerable, but the confidence is low and it is reasonable to expect additional information could increase the perceived risk such that risk reduction actions may be justified.

3. The estimated risk justifies risk reduction measures, but the confidence is low and it is reasonable to expect additional information could decrease the perceived risk such that the perceived risk may be tolerable.
4. The estimated risk justifies expedited action, but the confidence is low and it is reasonable to expect additional information could make the risks either tolerable or such that expedited action is not required.
5. The estimated risk is tolerable, confidence is high, but reasonable and prudent actions are recommended nonetheless.
6. The estimated risk justifies risk reduction measures and confidence is high so that no further studies are necessary before moving to a Corrective Action Study.
7. The estimated risk justifies expedited action and confidence is high so that no further studies are necessary before moving to a Corrective Action Study.

Each finding requires that the Dam Safety Case should establish claims regarding two main issues: First, an author must persuade decision-makers that risk falls within one of the three action-justification categories bounded by the lines on the f/-N diagram (see section on Risk Guidelines). Second, the Dam Safety Case must establish the confidence in the risk category, and whether additional exploration, investigation, or analysis has a reasonable likelihood of changing the perceived risk such that it falls in a different category. It is the rationale and structure of the Dam Safety Case argument that determines whether the risk numbers generated and the actions recommended make sense or 'feel right'.

At Reclamation, the Report of Findings is presented to the Dam Safety Office at a Dam Safety Advisory Team (DSAT) meeting. After the case has been made, Program Managers will ask questions of the DSAT members. :

What follows in this document is intended to help the Technical Report of Findings author build a persuasive case.

## **Establishing the Dam Safety Case in CR and Issue Evaluation Documentation**

The Technical Report of Findings Portion of the Decision Document Section in either the CR or the Issue Evaluation is where the Dam Safety Case is presented. The other sections validate the evidence and substantiate the claims used to make the case. Thus, copying and pasting from the other sections to the Report of Findings is generally not an effective means for building the case.

In dam safety review reports (e.g. CRs, PA, etc.) various sections of the report are designated to include information on design, construction, analysis, site examination, and structural behavior. These sections provide the background information and evidence to build a case. These sections establish claims which can be verified independently by

others. The claims are used later as evidence in other sections related to potential failure modes, consequences and probability estimation to make likelihood judgments. The validity of these claims is established by directing the reader to the references from which the facts are culled, and by commenting, if necessary, on the quality of the information. These sections also may include conclusions drawn from construction records and evaluations, field investigations, and computational analysis conducted by others. When these conclusions prove crucial to probability estimates, extra effort may be required to substantiate their validity. This is particularly true if the state-of-the-practice has changed since the previous analyses were conducted. While lesser arguments substantiate these claims, they should not all be brought forward verbatim to clutter the Dam Safety Case argument in the decision document. The Dam Safety Case may, however, include a few of the valid claims and analysis interpretations as the key pieces of evidence.

For the most part, the section describing potential failure modes addresses claims in the form of definitions. Definitions require interpretation, place concepts in categories, and provide perspective. The author is asked to provide interpretations regarding the likeliest locations for the failure modes' initiation. Evidence is placed in two categories: those factors which make it more or less likely for the failure mode to take place. The potential failure modes section also lists failure modes that are so much less likely than the others that time need not be spent developing probability estimates for them. Justification to remove a failure mode from further consideration requires an argument. Again, these lesser arguments should not be cut and pasted into the decision document. If well established in the potential failure modes sections, these claims can be made elsewhere without extensive substantiation in a decision document.

Sections describing future performance monitoring related to potential failure modes, consequences, and risk estimation include value judgments. Key evidence from sections that include design, construction, analysis, and structural behavior is used to add weight to the likely or unlikely side of each event in an event tree. Each event's value is established as a probability estimate and range based on the weight of evidence and inferences provided. The section dealing with future performance monitoring establishes the relevance of measurements with regard to their ability to detect failure mode initiation or progression.

## **Making the Case for Risk Belonging in the Proposed Category**

When dam safety risk assessment methodology is strictly followed, and claims regarding risk are well-substantiated, what are the key factors that would make the decision makers believe risk estimate numbers and agree to the recommended actions? A well-argued dam safety case will provide these key factors and will be convincing. Typically, it should not take more than three to five major contentions backed with only the key pieces of valid evidence or strong inferences.

A comprehensive risk analysis typically examines a dozen or so failure modes. Each failure mode may have another dozen or more likely and unlikely factors. If all the failure modes, with all the supporting evidence is cut and pasted into the decision

document, the decision-maker's attention will most certainly be diverted. Therefore, it is important that the findings portion of the decision section, where the Dam Safety Case argument's resolutions are stated, be a carefully constructed synthesis of the most important information. The risk analysis methodology is designed to take apart and examine all the features of a dam and diagnose its particular weaknesses. It is the role of the person making the Dam Safety Case to then interpret and synthesize the information rather than simply report it.

Potential Claims that Help Establish Belief in the Loading Conditions: Return periods or exceedance probabilities associated with extreme storms' peak inflow and volume, along with the reservoir's storage and the spillway's discharge capacities provide understanding regarding behavioral thresholds. So does an interpretation of peak horizontal acceleration levels or other loading parameters where dam responses to extreme earthquake loadings are likely to change, and the return periods associated with these levels. The nature of the storm (general, rain-on-snow, thunderstorm) or earthquake source (random, subduction zone, or active seismogenic crustal fault) helps to understand the nature of the uncertainty portrayed in the hazard curves. The type of empirical information used to generate the hazard curve (e.g. the historical record extends out approximately 150 years, paleo-geologic information contributed non-exceedance levels of \_\_\_ with an approximate return period of \_\_\_) can provide a measure of confidence.

Potential Claims that Help Establish Belief in the Dam's Response to Loading: Key evidence of exceptionally bad behavior from piezometers, structural measurement points or seepage weirs would be highly convincing. Factors that demonstrate poor design or faulty construction techniques are also important. The Case should include mention of design details or features that indicate most-likely locations of vulnerabilities. Strong evidence from laboratory or in situ tests that show low strength is also helpful. Conversely, evidence of good structural performance based on monitoring, attention to sound design and construction details, and/or testing that indicates high strengths would also be important.

Potential Claims that Help Establish Belief in the Life Loss Estimate: It is not sufficient to briefly describe the features inundated downstream and report a life loss number. A life loss estimate claim would be established by including pertinent facts regarding PAR clusters categorized by floodwave travel time downstream and anticipated flood severity, listing key factors affecting the failure initiation and breach development time, and briefly commenting on ease of evacuation. Also important are the factors that would increase or decrease the life loss from the best estimate, what the range might be, and why the best estimate is more likely to fall at some point in the range.

Usually it is informative to demonstrate where the highest contributions to risk are coming from, and to highlight the strongest evidence to support why this contributor is doing so. The potential failure modes that contribute most to the risk should be identified, and the reasons why this is believed to be so recounted. Sometimes the highest contributions come from a particular loading branch of the event tree. It can be informative to report a conditional probability such as the probability that the dam will fail given that an earthquake of a particular load range has happened. A contention of a clear and present danger (failure mode in progress) would be essential to justify taking immediate risk reduction action. Strong discomfort with the uncertainty regarding a piece of information that has a high degree of influence on the outcome is worth arguing.

Every effort should be made to avoid using conservative values, but reasoning should be provided any time they are used.

## **Assessing Whether Actions to Reduce Uncertainty Are Justified**

Each event and each state of nature comprising an event tree model has uncertainty in its likelihood. Issue evaluation risk analyses typically use Monte Carlo simulation to treat the uncertainty in these variables explicitly. A range of values and a probability distribution shape are elicited for each probability estimate. Monte Carlo simulation results are depicted in scatterplots, where each point represents a single combination of possible values chosen from each variable's probably distribution (see section on Combining and Portraying Risks).

While the scatterplot is an effective way to demonstrate uncertainty visually, the whole plot is not particularly useful for decision-making when deciding about risk tolerability. The outlying points represent situations where extreme values are chosen from each variable's probability distribution. Clearly, it is possible that each variable could obtain at its extreme value, but if the events are all independent, the likelihood that all are at extreme values at the same time is extremely unlikely. Should these outlier points influence decisions?

Reclamation has chosen to characterize 'Risk' as a mean value or as an expected value. Presumably, if each risk analysis team uses similar methodology to estimate event probability ranges, and they manipulate the uncertainty in a similar fashion, a mean value would have a degree of consistency when calculated for each dam and for each failure mode. Decision-makers then have an index number that can be used to rank risk amongst the dams in their portfolio. It is reasonable to use this index number to answer the question: "should we take action, today, at this dam or should we take action at some other dam first?"

Explicit treatment of uncertainty demonstrates that it may not be wise to ask the following question: "If no further action is taken for the present moment at this dam, precisely what is the Annualized Probability of Life Loss that Reclamation is incurring for this dam?" A single value such as the mean of the scatterplot is not sufficient to answer this question. It is presumptuous to believe that we are capable of obtaining a single, precise number that truly represents risk, as we have neither sufficient information on the true states-of-nature nor fully-reliable predictive physical models.

The deterministic methods in engineering design profession recognize these predictive shortfalls with the use of safety factors. Parameter combinations that obtain a safety factor of 1.0 are good to know, but it is rarely the goal to design to these combinations. Uncertainty in process and information is treated by prescribing various safety factor values greater than 1.0. The greater the uncertainty inherent to the problem or model being analyzed, the larger the prescribed safety factor required. Extreme (conservative) values for design parameters are rarely used because the design becomes cost-prohibitive (unless the organization involved has very deep pockets).

The analogous extreme situation in a probabilistic analysis is that it is imprudent to believe the outlier risk values in the scatterplot should be given strong weight, and it is equally imprudent to declare a dam 'safe' forever and for all time based on the mean value. Reclamation struggled with structuring tolerable risk guidelines around both a mean value and some measure of variance to the risk estimate. However, no rational basis could be developed, and Reclamation guidelines currently do not include a consistent requirement to portray and evaluate variance. The scatterplot has been used to characterize the full uncertainty. The mean has been used for program management. The recurrent nature of CRs ensures that a decision regarding risk be made every six years. Essentially, the question "How safe is safe enough?" will never be answered (unless the dam safety program is ended by an external policy decision).

Confidence that the risk number falls in the right guideline category can be addressed using the Monte Carlo analysis results. For instance, what would happen if more information was gathered, and the information proved favorable or unfavorable? A sensitivity study can be useful to investigate this question. Plausible upper and lower bound values for a variable in question can be chosen and held constant in a Monte Carlo simulation while all other variables are allowed their full ranges. When this test causes the scatterplot mean to move significantly, there may be justification to obtain additional information. A move is significant if it changes the risk guideline category. A person making this claim must provide additional reasoning to show why they believe the upper or lower bound values are plausible. They also provide reasoning as to why the additional information being requested is likely to reduce the uncertainty.

It should be understood that the risk guideline categories are not clean and crisp. They are not defined by where a "point" plots relative to a "line". Risks that plot just below a risk guideline are essentially the same as risks that plot just above a guideline, as the estimates are too imprecise to conclude anything else. Thus, it may be necessary to argue as to when it is believed that an assumption results in a change in risk guideline category.

Since there are several variables in any event tree, each variable that is explicitly treated for uncertainty (each variable depicted as a probability distribution using an @Risk formula) is a candidate for sensitivity analysis. There is a function in @Risk which, during the Monte Carlo simulation, keeps track of which variables most affect the ultimate risk estimate value in terms of "Sensitivity Rank Coefficients". The higher the rank coefficient, the stronger affect that variable had on the results. A negative rank coefficient simply means the results are negatively correlated with that parameter. This function does not provide information that can be used directly to assess the significance of total risk with respect to risk guidelines, but it can be used to discover those variables that create the largest variation in risk. These variables would be candidates for the sensitivity study described above involving plausible upper and lower bounds.

Reclamation Methodology requires that risk be reported as a best estimate and a range for each potential failure mode. There is a temptation to use the bounds of the scatterplot to depict the variation in risk. The range of risk estimates in a summary table are supposed to show the range of the mean, not the full range of the Monte Carlo simulation depicted in a scatterplot. Sensitivity studies provide a means to argue rationally about the range of means.

## The Dam Safety Case

Gathering evidence, stringing lines of reasoning together, and justifying a course of action for the Dam Safety Case are skills that require a significant amount of creative thinking; so where do these thoughts originate? Steve Vick, in his book “Degrees of Belief”, suggests judgment plays the major role, and then goes on to describe what judgment is and how judgment affects all manner of engineering thinking. Mr Vick presents the concept of an “hourglass model of situation awareness” which breaks the creative thinking process into three phases. The kinds of thinking he describes as taking place during each phase are essential to the case-building process. The reader is encouraged to go to the source to obtain a more thorough understanding of Mr Vick’s behavioral model.

First there is a Diagnosis Phase, where the engineer asks a series of key questions and formulates hypothetical answers to these questions:

- What conditions might lead to adverse performance?
- What aspects of the geology, design, or construction would lead one to think there might be a flaw or weakness?
- What field investigations, laboratory testing, or computer analyses have been done or should have been done?
- What can these investigations, along with monitoring program data, suggest about influential physical features which lie hidden within the dam or foundation?

By asking questions like these, there is an indistinct, before-thinking awareness that draws patterns which are recognized from the risk analyst’s past experience, and then leads to intuition about what the problems are or how the problems should be framed. Information and observations are compared to get a feel for a proper ordering of importance and to sense how the situation at hand is similar to or different from previously experienced situations. A mental image of conditions and failure processes begins to form. Inductive reasoning is used primarily during this phase. This phase is very subjective and is greatly enhanced by including several people with vast experience who may have different perspectives and who, together in lively debate, can create an exhaustive hypothesis set and thus a comprehensive diagnosis.

Following the Diagnosis Phase is an Analysis Phase, where the risk analyst attempts to assess the probable truth of the hypotheses formulated during the Diagnosis Phase. Failure modes are decomposed into event sequences or probable states of nature. Evidence, information, and underlying knowledge from different sources are gathered regarding each event in the sequence or each state of nature described. Physical evidence is compared, and is given relative weights to represent how likely or unlikely the event behavior or conditional state may be. Empirical evidence is gathered, grouped appropriately, and counted to provide relative frequency information along with statistical measures of central tendency, range, and skewness. Numerical models, formulated from principals of physics, are used to predict behavior. Model results should include not only a ‘best’ or ‘reasonably conservative’ answer, but should also include answers obtained when model parameters are varied within their potential ranges. This additional information can demonstrate how sensitive the results can be to particular parameters, and thus demonstrate how important it might be to gather additional in-situ or laboratory measurements to strengthen confidence in the values chosen for the critical parameters. Throughout this phase, assumptions must be verified and referenced as, all

too often, strong beliefs about the probable truth of evidence are generated when plausible but unsubstantiated hearsay is repeated until most who hear it believes it to be fact. Everything must be questioned!

Vick's third phase is an Interpretation Phase, where a critical review of all analyses is conducted to evaluate whether the results 'make sense'. Again, the results make sense when the recommended actions, the risk estimate numbers, and the understanding of the dam or levee's geologic setting, design, construction, and structural behavior all fit together coherently. The key evidence and the reasoning that warrants various conclusions must be drawn together to establish meaning and content. The most significant factors contributing to failure probability or risk are brought forward. These might include particular load ranges, populations at risk in certain locations, or assumptions regarding critical material properties. Key uncertainties should be identified, their significance demonstrated, and any additional information that could reduce the uncertainties should be presented along with reasoning to support a belief that additional investigations have a good chance of actually reducing the uncertainty. Arguments recommending that additional information be gathered can be given strong justification when the range for plausible variable input parameters can be shown to cause risk estimates to cross Public Protection Guideline action classification categories. Key reasons for believing there is or is not justification to take action can be provided and potential risk-reduction options which might be inexpensive and simple to implement can be identified. Adverse or favorable conditions can be compared to conditions previously interpreted in published case histories. How recommended actions will improve various situations should be projected.

There have been a number of schemes devised to simplify risk estimation during the Analysis Phase. These simplified methods are typically a part of screening-level risk assessments, but are sometimes used during what are supposed to be more comprehensive issue evaluation analyses. Frequently, these schemes take the form of spreadsheets in which a number of prescribed procedures have been coded, which may or may not fit well with the particular site under investigation. While these schemes can help the risk estimating team by suggesting a list of failure modes for consideration, along with important factors affecting each failure mode, there are dangers that can limit thinking during both the Diagnosis Phase and the Interpretation Phase. Site-specific failure modes may include events or states of nature which do not correspond to the limited number of prescribed failure modes. Those entering data into the spreadsheets tend to spend their effort trying to interpret vaguely worded instructions instead of thinking through site-specific features and processes for themselves. These instructions have to be general, and often do not lead one to picture a model that matches very well the particular situations at a given site. A spreadsheet user might interpret the instructions to find a good match for one event tree event at a particular location, but if the instructions for the next event do not continue to make sense, they might look for a good match at a different location.

The Dam Safety Case is a structured argument developed to have the facility's condition, risk estimates, and recommended actions make sense together. The risk estimate itself is not the sole basis for a particular decision. The case builder must show why it is reasonable to believe the risk and annual failure probability estimates. The Dam Safety Case should be backed by sound reasoning and solid evidence; it should not simply be a plausible story. The justification to take action (or not) must be fully developed. Sensitivity of the risk estimate to key parameters, a change in the Dam Safety

Classification when the key parameters are assigned one end or the other of their potential ranges, and the likelihood of a successful investigation should be addressed when recommending additional studies to reduce uncertainty.

## Example Dam Safety Case Arguments

Case 1. The risk justifies corrective action and confidence is high:

An extensive field exploration program established that the representative SPT blow count for an area at least twice the dam height is in the range 5 to 8. By the end of the exploration program, there was a combination of SPT and CPT boreholes on approximately 50 foot centers along the embankment toe and perpendicular to the dam axis. The borehole grid gave conclusive evidence that material with blowcounts in the 5 to 8 range was repeatedly found approximately 10 to 20 feet beneath the original ground surface. An impervious clay cap immediately above the low blowcount material is expected to keep any excess pore pressure generated by earthquake shaking from dissipating, and it may allow a particularly low shear strength zone to form if water from lower zones migrates up to the contact region.

An active fault, approximately 40 kilometer long, is located 12 kilometers from the dam. Trenching of this fault has identified displacements of at least a meter. A magnitude 5.75 to 6.5 earthquake is expected during a major event. A fault rupture analysis concluded that the shear wave velocity profile and the geometry of the basin in the dam's vicinity are such that significant amplification of ground motions is expected. Peak horizontal accelerations of approximately 1.0g are predicted at the 10,000 year return period, 0.5 g is at the 500 year return period, and 0.3g is predicted at a 170 year return period. Ground motions above 0.5 will liquefy the material represented by 5 to 8 blowcounts, ground motions between 0.3g and 0.5g might have a 50% liquefaction likelihood.

Material represented by this blowcount range is expected to have a post-liquefaction residual shear strength of from zero to 400 psf. The reservoir is a regulating basin for several major canals, and is operated in a narrow band constantly producing a freeboard of from 9 to 15 feet. Flac was not available but limit equilibrium analysis showed safety factors less than 1.0 would include the embankment crest and slip surfaces would intersect the upstream dam slope below the water line. It was considered between highly likely and neutral that deformation would exceed freeboard (0.8 if the shear strength is 0 to 200, 0.6 if the shear strength is 200 to 400).

A community of approximately 1900 people live less than 2 miles downstream from the dam. About 40 percent of these people leave the area for work during the day, but during the day, the transient population is increased by the presence of restaurants and service stations at an exit for a major interstate located within the flood inundation boundaries. Dambreak flooding would put fast-flowing water, approximately 20 feet deep through a trailer park, so about 30 percent of the PAR would be exposed to moderate severity flooding while the other 70 percent would be exposed to low severity flooding. The PAR has less than a mile to travel to be out of the inundation boundaries. There would be little warning if the dam was immediately overtopped and if a breach were to form quickly. The embankment is constructed of very dense, medium-plasticity sandy clay and is

expected to be highly resistant to erosion. The county sheriff mans a substation in the middle of the trailer park during the day on weekdays, but when not manned, the response time could be as much as an hour. The potential life loss given sudden failure is estimated to range between 50 and 400, depending on how deep the water is initially at the point of overtopping, and upon how fast the full breach could form.

Given an earthquake capable of producing  $>0.5g$ , the likelihood of failure by overtopping is estimated to be 0.62. Given an earthquake between 0.3 and 0.5g, the likelihood of failure by overtopping is 0.38. The annual probability of failure from these two branches alone is 0.0035. If 180 is the life loss, the annual probability of life loss is 0.63. If the breach formation is not so rapid, a life loss of 50 might be more appropriate, and the annual probability of life loss would be about 0.18.

Case 2. The risk does not justify corrective action, but confidence is low and action could be justified if investigations prove adverse conditions.

Dam B is an older Reclamation dam, constructed by dumping upstream and downstream shell material from railroad trestles and then washing fine-grained material into a settling pond between the coarse-grained ridges. The left abutment is very steep, and a spillway structure is perched on its rim. Construction photographs show a complex timber frame structure on the embankment side of the spillway wall, presumably placed to help with construction of the spillway walls. Exploratory drilling in other locations at the dam encountered wood intervals, so it is believed that the timber framing for the trestles was left in place. It is not known if the timber frame structure on the back side of the spillway walls was removed. Depressions have been observed in fill material in the vicinity of the spillway wall. Asphalt has been placed several times to bring the dam crest road level with the spillway wall.

A seepage path along the embankment side of the concrete spillway wall was judged to be the most likely location for initiation of a failure mode involving internal erosion. A continuous flaw would have to form along through core materials adjacent to the spillway. Seepage through the flaw would have to be of sufficient velocity to begin erosion and to transport core materials. The downstream shell zone would be ineffective as a filter, and erosion would continue until a 'pipe' develops in the soils adjacent to the wall. The upstream zones would fail fill in the pipe and lodge against the downstream shell zone to form filter and limit seepage flows. As the pipe enlarges, increased seepage velocities would result in more erosion. If intervention failed to stop the process, the expanding pipe would ultimately form a breach through the dam crest, which would then expand rapidly.

The embankment design and construction was inadequate by modern standards. Some of the embankment materials appear to be silty. In addition, the embankment materials were not compacted and therefore are considered quite erodible. The abutment profile is relatively steep, which raises the possibility of embankment cracking. The spillway walls provide upstream to downstream continuity through the embankment. The observed depressions have been filled previously and thus may be recurring. New piezometers by the spillway and through the crest respond to reservoir at reservoir water surface elevations 4405 and above. Plumes of moist air have been observed on the downstream slope may indicate changing conditions. There is some indication that these plumes were first observed a few years ago. There is no engineered filter zone at the dam. The downstream rockfill zone could mask seepage. There are no apparent flow limiters if

internal erosion were to initiate. Evidence of trestles, wood forms, and construction debris was encountered in the upstream slope test pits. Some small pockets or voids were observed. A significant portion of the reservoir (20,000 acre-feet) is above the downstream 8:1 slope about 95% of the time.

The dam has more than 80 years of performance history without a significant seepage incident. The reservoir has a long history of operations 5-10 feet above reservoir levels in recent years without any apparent problems. There are no known case histories of dams constructed by hydraulic fill methods failing from internal erosion or piping through the embankment. Another Reclamation dam was constructed with similar methods and materials as this dam, and seepage rates there were estimated to be about 20 ft<sup>3</sup>/s without failure. Observed seepage at this dam is very minor. There is no observable seepage on the downstream slope. The construction methods used may have created a 'natural filter' between the puddled core and the shells. The spillway has a very large release capacity and can draw the reservoir down relatively rapidly, which makes successful intervention likely if detection occurs. The reservoir would not have to be drawn-down very far to prevent a developing internal erosion failure mode from progressing to a breach. Internal erosion is not the only possible explanation for the low spots along the spillway (it was thought that settlement into the valley a very likely candidate). A potential transverse crack that intercepts the reservoir probably has a very small aperture at depth, which would suggest limited potential for seepage velocities capable of causing internal erosion.

Though there has never been a failure of a dam constructed using hydraulic fill methods, the base rate frequency for initiation of an 'incident' was judged to be very low.

Approximately 1100 permanent residents are located less than 2 miles downstream from the dam. A State Military Reservation is also located within the inundation boundaries. The number of people stationed at the base changes throughout the year, but when the reservoir is full and most likely to fail from internal erosion, the PAR is approximately 1,500. Several towns would be inundated farther downstream, but the long warning time and favorable evacuation factors would significantly reduce the potential for life loss. At the nearest town downstream, dambreak flooding would put fast-flowing water, up to 60 feet deep through the portions of town closest to the river. Approximately 80 percent of the PAR would be exposed to moderate severity flooding while the other 20 percent would be exposed to low severity flooding. The PAR has less than a mile to travel to be out of the inundation boundaries. There would be little warning if a breach were to form quickly, though this was not considered likely. The facility has a power plant and there is frequent recreational usage immediately downstream from the dam, so it is fairly certain a slowly developing seepage problem would be noticed early. The potential life loss given rapid failure is estimated to range between 50 and 400, depending on how fast the full breach could form.

## **Workshop Problem**

Consider a facility for which you have considerable first-hand knowledge. In fifteen minutes, develop the case for a recommended action touching upon only the strongest

evidence you might have with respect to the loading, response to the loading, and consequences.

## References

Peck, R.B., "The Risk of the Oddball," Reclamation Workshop on Risk of Piping and Internal Erosion, December 2-5, 1997.

Taleb, Nassim B., "The Black Swan: The Impact of the Highly Improbable", Random House, April, 2007.

Vick, S.G. (2002), "Degrees of Belief - Subjective Probability and Engineering Judgment", ASCE Press, VA, 2002.