PDHonline Course G382 (3 PDH)

Risk-Based Engineering - The New Paradigm

Instructor: Frederic G. Snider, RPG and Michelle B. Snider, PhD

2020

PDH Online | PDH Center

5272 Meadow Estates Drive
Fairfax, VA 22030-6658
Phone: 703-988-0088
www.PDHonline.com

An Approved Continuing Education Provider

Risk-Based Engineering - The New Paradigm

Frederic G. Snider, RPG and Michelle B. Snider, PhD

Introduction

Although the concept of risk assessment for large engineering projects has been around for decades, it has only recently been formalized and found its way into the literature, the topic of sessions at technical conferences and appeared within the regulatory environment. As financial resources become more scarce, owners of engineering projects, including roads and bridges, planes and trains, dams and levees, power plants, structures in seismically active areas, industrial facilities and even office buildings susceptible to terrorist attacks have to decide where best to spend repair, rehabilitation, retrofit and upgrading dollars.

The current formal risk assessment procedure requires familiarity with probability and statistics and some non-traditional thinking (for example, what is the 'acceptable' number of potential deaths?). This course provide an overview of the process and is designed for all types of engineers, architects, project managers, and anyone else who needs to know the basic concepts, the process, the benefits, and the pitfalls.

Beyond large engineering projects, risk-based decision making is becoming more and more common in health care, elder care, homeland security, military campaigns, infrastructure, pharmaceuticals, traffic management, and many other fields that touch all of our lives. As such, familiarity with the type of thinking and methodologies that form the basis of the risk-based decision process should be a requirement for every responsible adult.

Definitions of Key Terms

Many terms are thrown around during any discussions on risk. In order to form a common understanding and facilitate communication, some formal definitions are required, as follows.

Risk - the possibility of suffering loss. The concept of risk includes both the likelihood of an adverse event and the consequences of that event. It can be expressed as a number, such as 0.9 or as a qualitative description, such as 'high'. It is a way to address uncertainty.

Risk Analysis - A quantitative calculation or qualitative evaluation of risk to support decision making. Its purpose is to avoid excessive risk, avoid excessive conservatism, and make the best use of available resources in a systematic, consistent and 'defensible' manner.

Risk Control Measures - Either procedural or remedial actions that can reduce risk. For example, publishing new emergency evacuation procedures or increasing the height of an old levee. In this example, increasing the height of the levee decreases the likelihood of failure, while publishing new emergency evacuation procedures should reduce the consequences of the failure (loss of life).

Risk Assessment - The process of deciding whether an existing risk is acceptable, the risk control measures are adequate, or that additional or alternative risk control measures are necessary to make the risk acceptable.

Acceptable Risk – A nearly impossible term to define, except among a very small-number of like-minded individuals. This is the crux of the entire risk assessment process, and the reason why it is rarely easy.

Components of the Risk-Based Approach

The US Bureau of Reclamation is responsible for the safety of a large number of dams throughout the country. They issued a best practices document in 2010 (U.S. Bureau of Reclamation, 2010). The Bureau's documentation addresses the major components for a formal risk assessment applicable to any engineering project, facility or component.

They identify the major components of risk analysis as follows:

- **Potential Failure Mode Analysis**
- **Event Trees**
- **Load-Frequency Analysis**
- **Probabilistic Analysis and Modeling**
- **Subjective Probability and Expert Elicitation**
- **Consequence Evaluation**

Each of these topics is addressed in the remainder of this course.

Potential Failure Mode Analysis

The Potential Failure Mode Analysis, or PFMA, is the critical first step in development of the risk analysis. This step involves identifying the ways a potential structure or engineered system could fail. Note the PFMA addresses the failure modes, NOT the consequences of the failure. The PFMA must be thorough and complete; otherwise the resulting analyses may yield incorrect and/or incomplete conclusions.

For example, consider an earthen dam holding back a reservoir. In this case, we consider 'failure' as breach of the dam and the draining of the reservoir. Let's look at some of the ways an earthen dam could fail:

- Potential Failure Mode 1: During a major storm, the dam is overtopped and erosion of the dam leads to dam breach, washout and loss of the entire reservoir.

- Potential Failure Mode 2: Over time, seepage through the dam results in weakening of the soils, causing parts or all of the downstream face of the dam to slowly slide downhill, eventually leading to a breach and washout of the dam and loss of the entire reservoir.

- Potential Failure Mode 3: Rotted tree roots and animal burrows in the dam provide open pipes which, if they extend from the downstream to upstream faces of the dam could allow seepage and internal erosion. Such seepage and erosion could then make the passageways bigger and bigger, leading to breach and washout of the dam and loss of the entire reservoir.

Note that we also don't consider time frames at this point. For example, the overtopping scenario could happen immediately during a big storm, but the seepage through the dam might get progressively worse for years or even decades before failure of the dam occurs.

As a different example of timeframes, consider a commercial jet, with failure defined as the plane falling out of the sky. Two of the many possible failure modes could be icing of the wings and metal fatigue in the tail section. Whereas icing could occur immediately under proper weather conditions, metal fatigue and failure of the tail section might not become critical until the plane has flown millions of flight miles.

The PFMA often requires analyzing components or systems on an individual basis and also their interactions with each other. For example, there might be two critical components in a system. If either malfunctions alone there is no failure. If both malfunction at the same time, however, failure ensues. Therefore the PFMA must be both holistic and compartmental. As an added complication, the PFMA must also consider the human factor, whereby all engineered systems are working fine, but human error leads to failure.

Identifying the Potential Failure Modes, or PFMs, is best done in a workshop setting, with a small (usually less than 10) but diverse group of individuals. Care must be taken to remain objective and non-judgmental, especially if the original designers and/or manufacturers are on the team. Operations personnel should also participate in the workshop, as they have the day-to-day knowledge of how the systems actually work, warts and all, and where the vulnerabilities might lie. The PFMA workshop also benefits from a trained facilitator to maintain the record of the workshop, keep people focused, and keep the process on track. Often breaking the workshop over two days is of great help, as the participants get to 'sleep' on the first days

discussions, often leading to improvements in the first days results and increased productivity during the second day. The facilitator is usually tasked with preparing a report identifying the PFMs for review and approval by all the team members shortly following the workshop.

As part of this process, each potential failure mode has to be fully described. Each PFM includes three major parts: the Initiator, the Progression of Events, and the Ultimate Failure. Each is described below.

The Initiator

The initiator is the load or physical condition that starts the failure process. For our dam example, one initiator could be the large flood event, another might be the development of animal burrows.

Depending on the system being evaluated, other initiators might include earthquake ground shaking, terrorist bomb exploding, tsunami, failure of a critical part, wires shorting due to water intrusion, loss of electrical power, excessive vibration due to bearing wear, improper venting of fuel tanks, etc.

Initiators may be a single event or load, or several events or loads happening at the same time or sequentially. Importantly, by definition, the initiator is a single or sequential load or event that, if you can prevent its occurrence, precludes the development of the failure.

Failure Progression

The failure progression is the step-by-step series of additional events, small failures, disruptions, or other processes that eventually lead to the ultimate failure. In addition to listing the individual events and processes, their locations should also be identified, as well as the location of the ultimate failure. For our earth dam example, if one section of the crest is lower than the rest of the dam, this section is the most likely location for the start of erosion and failure during overtopping, although ultimately the entire dam could be washed away.

So a failure progression for the earth dam might be listed as:

1. (initiator) Large flood occurs
2. Flood waters overtop the dam at location “X”
3. Water pouring over location X erodes away the gravel road along crest of dam.
4. Continued erosion removes the soils beneath the road forming a trough, thereby allowing even more water to flow.
5. The increased flow accelerates erosion downward through the dam.

6. As the erosion extends downward, the steep sides of the erosion trough fall in and are washed away, widening the breach.
7. As the breach deepens and widens, flow increases, thereby accelerating erosion.
8. The breach continues to widen and deepen until the entire reservoir is drained.

During this process, it is also worthwhile to discuss the timeline of the failure progression. Does the failure progression happen quickly immediately following the initiator, or does one or more event or process within the failure progression take a longer time to occur? As with the initiator, it is likely that if one or more of the events or processes in the failure progression can be prevented, perhaps by improved inspection, monitoring or upgrade, the related ultimate failure may never occur.

The Ultimate Failure

This is the culmination of the effect of the initiator and the subsequent failure progression. In our commercial jet example, it is the uncontrolled crash of the plane to earth. In our dam example, it is the loss of the dam and draining of the reservoir.

The following is a true story.

Photos - The Failure of Teton Dam, June 5, 1976

The 305 foot-high earthen Teton Dam near Rexburg, Idaho was built by the US Bureau of Reclamation in the early 1970's. The dam was completed in 1976. On June 5, 1976, during the initial filling of the reservoir, the dam suddenly failed. A breach rapidly developed where the dam meets the abutment. The photo on the left shows the beginning of the breach. The photo on the right shows almost half of the dam eroded away and millions of gallons of water pouring through the breach.

When the dam failed, the reservoir was already 270 feet deep and 17 miles long. The entire reservoir drained completely in less than six hours. The rapid draining of the reservoir flooded the canyon downstream and triggered more than 200 landslides. The failure resulted in the loss of 11 (amazingly few) lives and millions of (1976) dollars in property damage. The failure of Teton Dam was the wake-up call to the Bureau and the entire dam engineering profession that initiated the process of thinking about risk in totally different and more formal ways. (Source: US Bureau of Reclamation web site: www.usbr.gov).

Event Trees

Event trees are a graphical representation of a sequence of events and possible outcomes. They are used in many fields to represent many types of sequential processes, including risk assessment. Event trees can be as simple or complex as needed. They are called trees because they typically have a number of branches.

Here is a simple event tree mapping the possible outcomes of flipping a coin three times.

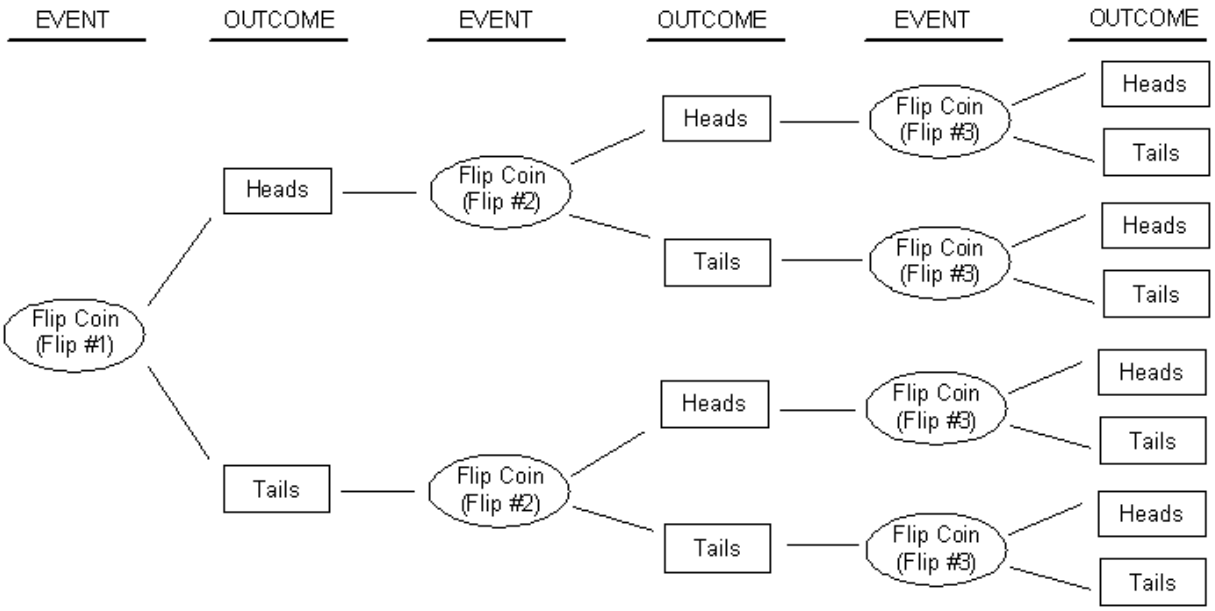

Figure 1 – Event Tree for Flipping a Coin Three Times in a Row

You can trace eight possible paths from the initial event (Flip #1) to the eight final possible outcomes. This means there are eight possible sequences of events. Each of the ovals above is called a 'node' of the event tree, as it is at those points that branches start from.

The coin-flip event tree illustrates two key aspects of all (proper) event trees.

1. The tree always starts with a single event, all the way to the left.

2. Additional events occur along connected branches, from left to right

For each event, there is only one of two possible outcomes. For this reason, event trees are termed **"binary"** or "**Boolean"** trees. If you have an event that has more than two possible outcomes, it must be broken down to smaller sub-events, until each event and sub-event is binary (only two possible outcomes).

Now think of the final outcomes of the coin-flip tree. If you decide that at the final outcome 'heads' means nothing bad happens, and 'tails' means your system failed completely, then the tree shows you that there are four sequences of events (paths) that result in nothing bad happening, and four sequences (paths) that result in failure. At this point, you can examine the intermediate events and see which paths you might preclude by additional engineering, monitoring, or testing.

Probability Analysis

In the coin flip example, note that probability at every event node is fixed, and is independent of whatever happened before. That is, at every node, there is a 50-50 chance of getting a heads or a tails. Since we know the probability at each event, we can calculate the overall probability of any outcome.

By definition, we assign a probability of 1 to any outcome that is absolutely certain to occur and a probability of 0 to an outcome which can't possibly occur. Numbers between 0 and 1 represent the level of uncertainty assigned to a given outcome.

Because we have reduced all our events to binary or Boolean, there can be only one of two outcomes. **Therefore, by definition, the probability of one outcome plus the probability of the other outcome must total to 1.** That is, there is a 100% probability that one of two those outcomes will occur (although we don't know which one).

Let's look at the coin-flip event tree again. For each node, the probability of getting heads is 0.5 (50/50) and the probability of getting tails is 0.5 (50/50). The event meets the requirement that the sum of the outcome probabilities at each node is 1 (0.5+0.5=1).

Here is the coin flip event tree again, with the probabilities shown at each event.

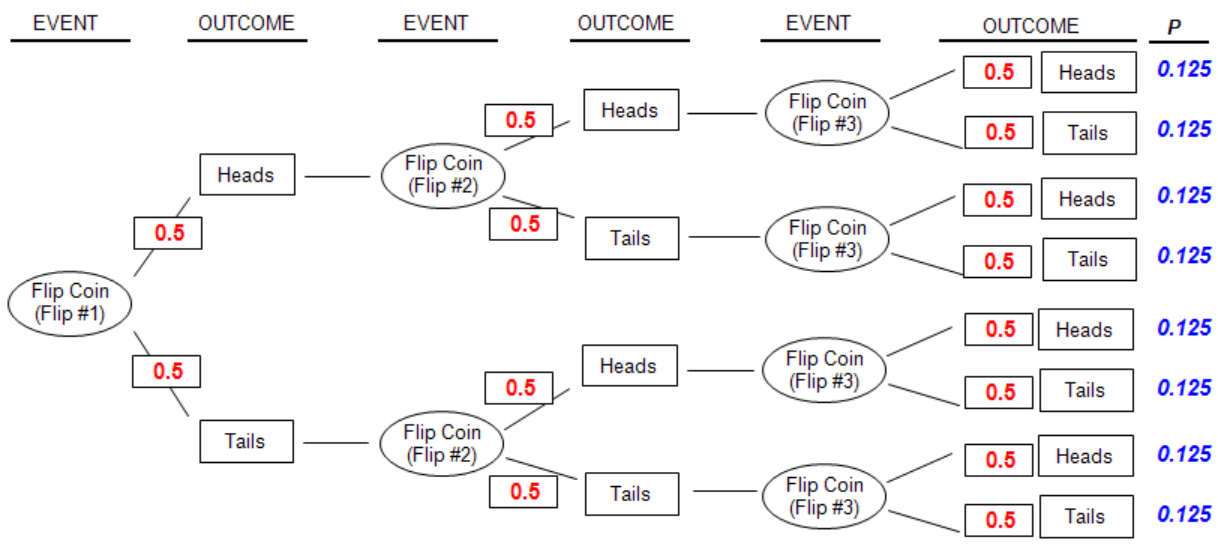

Figure 2 – Event Tree for Flipping a Coin Three Times in a Row with Probabilities Assigned

Now for any particular path through the tree, the probability of the final outcome is calculated by **multiplying** the probabilities at each intermediate outcome along the path. For example, the probability of getting three heads in a row (the topmost path, above) would be flip#1 x flip#2 x flip#3 which is 0.5 x 0.5 x 0.5 = 0.125. Note that 0.125 is the equivalent of 1/8, or in other words, the chances of three heads in a row is 1 out of 8. The final column is the probability of each final outcome, labeled "***P***", the standard symbol for probability expressed as a number between 0 and 1.

By inspection of the event tree with probabilities, we see that the probability of any combination of heads and tails for three flips is 0.125 (12.5%). So there are eight possible final outcomes, each with a probability of 0.125 (12.5%). If we add 0.125 together 8 times, we get 1.0, or 100%, as we must. This says that the probability of one of those outcomes occurring is 1.0 - meaning it is 100% absolutely certain that if we flip a coin three times, one of those outcomes will occur (although we don't know which one). There are no other possible outcomes than those shown on the chart. So at least the event tree is correct mathematically. (No, we have not considered the remote possibility that a coin lands on edge.)

What is the probability of the final outcome being heads by any possible path? There are four paths with the final outcome being heads. Each outcome has a 0.125 probability of occurring. To determine the total probability, we **add** the probability of each path resulting in a final outcome of heads. Adding up the final head outcome paths equals 0.125 + 0.125 + 0.125 + 0.125 = 0.5. Therefore, there is a 50/50 chance that the final outcome will be heads, independent of the path to get there. This is true; even if I flip a coin 100 times and get 100

heads in a row, the chances of the next flip being heads is still 50/50. (If you find this type of thinking interesting, you might consider taking our Probabilities course, also on PDHOnline).

If we don't care what order the heads and tails are in, then there are four possible outcomes:

1. **Three heads in a row**
2. **Two heads and one tail in any order**
3. **Two tails and one head in any order**
4. **Three tails in a row**

By inspecting the event tree, we see that there is only one path for three heads and only one path for three tails. For two heads and one tail in any order there are three paths. Similarly, for two tails and one head in any order there are also three paths. So we add the path probabilities for each combination and get:

1. **Three heads: 0.125 (1 out of 8)**
2. **Two heads and one tail: (0.125 + 0.125 + 0.125) = 0.375 (3 out of 8)**
3. **Two tails and one head: (0.125 + 0.125 + 0.125) = 0.375 (3 out of 8)**
4. **Three tails: 0.125 (1 out of 8)**

This also fits with our intuition. We are three times more likely to get a mix of head and tails than we are to get three heads in a row or three tails in a row.

Ok, enough with the coin flips. Let us create an event tree for a more important situation:

We have a factory, and we would like to address the scenario of a fire breaking out on the factory floor. We recently installed a good sprinkler system and we test it every month. If the sprinkler system triggers, it also automatically calls the fire department.

The purpose of any sprinkler system is to buy a little time until the fire department arrives. As our backup plan in case the automatic calling fails, we are counting on at least one of our employees calling 911 to get the fire department.

We consider that we have three possible final outcomes. If the sprinkler system sprays water and the fire department arrives, we will likely have minimal damage. If the sprinkler system doesn't spray water, and neither the system nor an employee calls the fire department, we figure it's a total loss. If the sprinkler sprays water but neither it nor our employees calls the fire

department, or if the system doesn't spray water but it or an employee does call the fire department, we expect partial damage.

Did you get all that? This is one of those cases where a picture is worth a thousand words. The picture will be an event tree that looks like this:

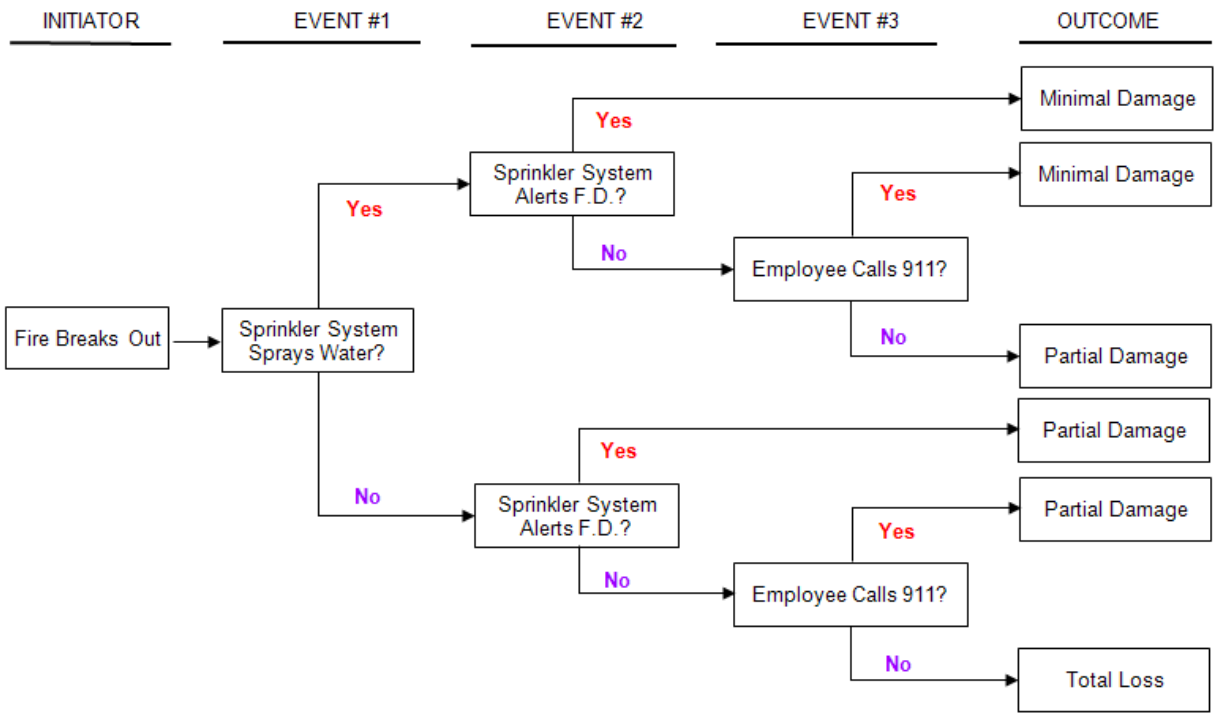

Figure 3 – Event Tree for the Fire Problem

So if fire breaks out, there are six possible paths to the final outcomes, and six final outcomes. Even though the final outcomes have duplicates, we leave them as separate paths, and combine them later after we calculate the probabilities. There are several things to note here:

1. **The initiator event is shown on the left and labeled to clarify what this event tree addresses**

2. **There are no branches to the event "Employee Calls 911?" in the two cases where the sprinkler system alerts the fire department. If the sprinkler system called the fire department, it doesn't matter whether an employee also calls. Therefore that event is irrelevant and not shown on the tree.**

3. **You might be tempted to combine the outcomes, since, for example, there are three "Partial Damage" outcomes. Although having only one "Partial Damage"**

box and having three lines point to it might simplify the drawing, it will be confusing when we start assigning probabilities. Leave each path separate.

Assigning Probabilites

As with the coin flip analysis, the next step is to assign a probability to each intermediate event and outcome in order to determine the probabilities of the final outcomes. However, in this case assigning the probability that, say, an employee will call 911 is not a straightforward exercise. For real-life probability analyses, we need tools to help us assign probabilities to these events. There are several possible approaches, described below.

Deterministic Approach

The deterministic approach is to just look up the probability of the outcomes at each node in a book or on the internet. This works for situations where experience or lab testing or some other method has established a commonly accepted probability for a particular event. For example, we know that the probability of head or tails on a coin flip is 50/50. As another example, actuarial tables of life expectancy of people in the U.S. are used by life insurance companies. The federal government publishes the actuarial tables. So I can look up the odds of dying at age 85.

These examples don't help us in our Fire Problem. But maybe extensive research by an independent laboratory has established that the probability of the sprinkler system not spraying water is 83.5%. Then you could use that number. However, in most cases, we don't have many deterministic values.

Subjective Probability

The next level approach to the Fire Problem is to assign a single probability to each event based on our gut feeling. As the owner of the factory, I have assigned probabilities based on the following reasoning.

- **I think chances of the sprinkler system activating and spraying water is 90% (0.9). Note that since the sum must equal 1.0 (100%), this is the same as saying I think the chances of the sprinkler system not spraying water is 10% or 0.1.**
- **If the sprinkler system successfully activates the water spray, I think there is an 85% chance that the automated call to the fire department will go through.**
- **If the sprinkler system fails to activate the water spray, I think there is only a 20% chance the system will successfully call the fire department. After all, it didn't work correctly on the spraying part.**

- **If the sprinkler system sprays water, I think there is a 25% probability someone will call 911, since the employees have been told that the sprinkler system automatically calls the fire department if there is fire, so why do they have to do it?**

- **If the sprinkler system doesn't activate, I think there is a 75% chance that someone will call 911, since obviously the automatic water spraying isn't working. I only assigned 75% because there are parts of the factory where people rarely go, and without the water spraying, no one would think to look until things were out of control with smoke everywhere.**

So now I have one probability for each event outcome, and I just jot them down on the event tree and do the math as in the coin flip example. Since there are not many calculations to do, I could use any number of calculation tools, from hand calculations to desk calculators to spreadsheets.

On the event tree below, the probabilities have been assigned at each node, based on the rationale explained above, and hand calculated the resulting probabilities of the different outcomes.

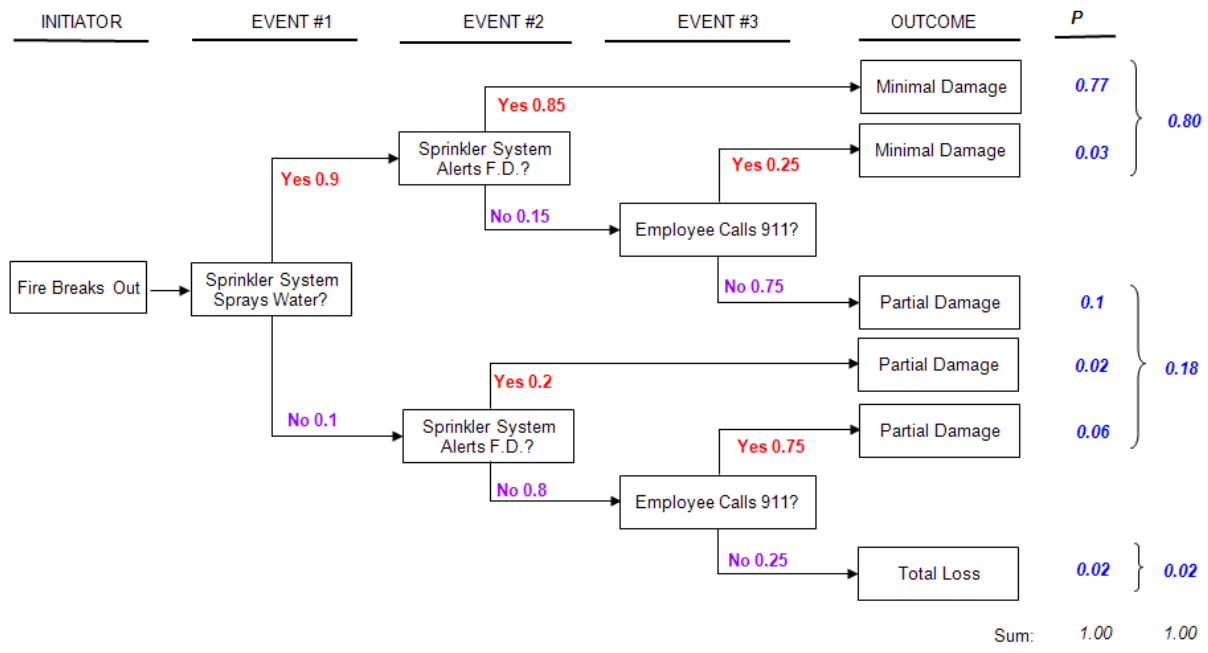

Figure 4 – Event Tree for the Fire Problem with Single Probabilities

Note that for each node, the outcome probabilities add to 1.00, as they must if we did our setup and calculations correctly. Then, as with the coin flip, to determine the probability of each final outcome, we **multiply** the probabilities along each branch. So for the uppermost path – the sprinkler sprays water (0.9) and the sprinkler system calls the fire department (0.85) the probability of this path occurring is 0.77, or 77% (0.9 times 0.85). Each outcome has the probability shown under the "P" column.

Then, the probabilities of the identical outcomes are **added** to determine the overall probability for each outcome, shown after the brackets, above. So the overall probability of minimal damage is 80%, odds of partial damage are 18% and odds of a total loss are 2%. Can you live with that? If not, should you get a backup system to call the fire department in case the sprinkler system fails instead of counting on an employee? Good question, and partially answerable using the event tree as shown, as I can see where the weak points in my system are.

But what happens if I'm just wrong about the probabilities I've assigned? As computer folks don't remind us enough: "garbage-in yields garbage-out." Read on….

Expert Elicitation

Clearly, the subjective probability approach described above is limited by how good I am at guessing the individual probabilities for each event and outcome. The need for "good guessing" is common in most risk assessment projects.

Recall the famous quote: "It's hard to make predictions, especially about the future". (According to Wikipedia, this quote is attributed to many people, including Yogi Berra, Niels Bohr, Samuel Goldwyn, Robert Storm Petersen, and Mark Twain). There is always uncertainty in life. In this case, there is even uncertainty about who said the famous quote about uncertainty.

One approach to addressing the uncertainty in our assignment of probabilities is using a panel of experts. This is called using Expert Elicitation, and is called a "degree-of-belief" probability method. The process is usually best accomplished in a workshop setting to encourage brainstorming and enhancing synergy among a group of experts.

Experience has shown that for best results, a verbal mapping scheme should be used, rather than ask experts to assign a numeric probability. Then, based on the words used, we assign probabilities or ranges of probabilities for each node in the event tree. Here is an example of a verbal mapping scheme adapted from Vick (2002) and Reagan (1989).

Verbal Descriptor	Suggested Probability	Probability Range
"**Virtually Impossible**" due to known physical conditions or processes that can be described and specified with almost complete confidence	0.01 (1%)	0.00 to 0.05 (0% to 5%)
"**Very Unlikely**" although the possibility cannot be ruled out.	0.10 (10%)	0.02 to 0.15 (2% to 15%)
"**Equally Likely**" with no reason to believe that one outcome is more or less likely than the other.	0.5 (50%)	0.45 to 0.55 (45% to 55%)
"**Very Likely**" but not completely certain	0.9 (90%)	0.75 to 0.9 (75% to 90%)
"**Virtually Certain**", due to known physical conditions or processes that can be described and specified with almost complete confidence	0.99 (99%)	0.9 to 0.995 (90% to 99.5%)

The table above is based partially on experiments by Reagan (1989). The experiments showed that, at least within reasonable limits, people are fairly well calibrated and consistent relative to known probabilities when describing events in these verbal terms.

However, a key finding of the experiments was that people's ability to judge likelihood does not extend very far out on either end of the probability scale; that is below 1% or above 99%, even when words such as "virtually certain" are proposed. This is likely due to the fact that human experience rarely allows us to conceptualize likelihoods at extreme probabilities. Therefore, we don't have adequate words to describe them. "Once in a blue moon" or "when pigs fly" doesn't quite cut it mathematically.

It is also difficult for most people to conceptualize how often extreme events happen and in what time frames. For example, all earth scientists would probably agree that it is "**Virtually Certain**" that the earth will be hit by a large meteorite at some time in the future – it has happened many times in our geologic past. But most would probably agree that it is "**Very Unlikely**" to happen in our lifetime. So we also need to make sure we are all on the same page with respect to time frames.

Adverse and Favorable Factors

In order to help facilitate the verbal mapping process, part of the workshop should be devoted to listing adverse and favorable factors for each potential failure mode identified.

For the Fire Problem, adverse factors for the “Sprinkler System Sprays Water?” node might be “unreliable municipal water supply pressure” and/or ”poor coverage of sprinkler heads through the building”. Favorable factors could be the young age of the sprinkler system and monthly testing practices. As these factors are listed and discussed, the participants will gravitate towards the verbal mapping description that seems most appropriate (like “very unlikely”).

If facilitated properly, the process of listing the adverse and favorable factors for each failure mode will typically generate lively discussion among the entire team and hopefully form the basis for a consensus view of the assigned likelihoods. Once the facilitator senses the discussion of one potential failure mode is over, the group is queried for the most appropriate verbal descriptor. After the descriptors are assigned, the key adverse and favorable factors and justifications are captured and documented.

The facilitator helps guide the process and attempts to deal with biases and adverse group interactions along the way. In some cases it may be beneficial to poll the team anonymously for decision making, especially if discussions are dominated by a few strong, and loud, individuals.

One interesting aspect is whether the facilitator is attempting to assign probabilities by majority rules (voting), tries to achieve consensus by additional discussion, or assigns a weighted score for each probability based on the number of people who voted for it.

Or should all opinions be weighted equally? As my father was so fond of saying, “disagreeing with everyone doesn’t make me wrong, just outnumbered”. If ten years ago you argued that a tsunami in Japan would bring about the virtual end of nuclear power as an ‘acceptable’ energy source, you certainly would have been a lone voice in the wilderness. Who knew?

Clearly, there is no right or wrong way to run the workshop, just some ways that take more time than others. For engineered systems, the workshop process can also be useful to identify ways that adverse factors could be mitigated or additional favorable factors introduced by testing, monitoring, repairs or upgrading.

The Workshop

As the owner of our factory, I decided to convene a workshop with a group of experts to more fully evaluate my fire scenario. I’ve invited a manufacturer’s rep from the sprinkler company, the local fire chief, a couple of other factory owners who have the same system, and my plant manager. I hired a facilitator to make it go smoothly and make sure we attain our goals.

The facilitator suggested we use the verbal description approach described above.

We convened on a Wednesday morning. We discussed adverse and favorable factors for each node of the event tree and after lively discussions reached the following consensus. Note we only have to decide on the likelihood of the “Yes” branches.

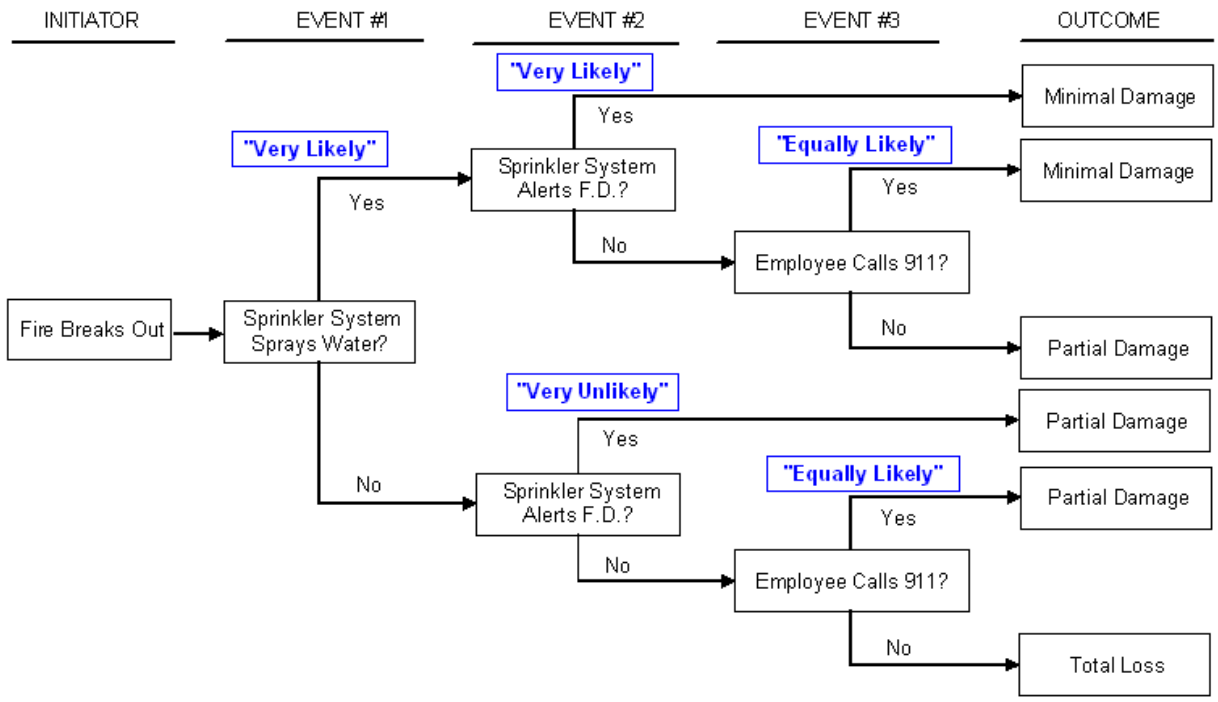

Figure 5 – Verbal probabilities assigned to the Fire Problem nodes at the workshop

Interesting, but probably not surprising, was the amount of discussion reaching consensus on the two “Employee calls 911” nodes. We finally just decided that we would assign the equally likely (50/50 chance) to the employees, although I personally feel they would be much more responsive. Others felt they would be in such a panic to get out that no one would even think to call until it was too late.

Recall from the verbal description table that each description has an assigned range of probabilities, as summarized here:

Verbal Description	Range of Probabilities
Very Likely	0.75 to 0.90 (75% to 90%)
Equally Likely	0.45 to 0.55 (45% to 55%)
Very Unlikely	0.02 to 0.15 (2% to 15%)

We now have a range of probabilities for each node, rather than just one like in my original “gut feeling” approach. We now could just calculate the outcome probabilities using the averages, or the worst case or the best case combinations, but we would not get the whole picture. Note that this is a simple example – as the number of nodes increases, the mathematics quickly gets unwieldy.

So how do we approach the calculations if all our probabilities have ranges? For that we turn to the experts in Las Vegas.

Monte Carlo Simulations

A mathematical approach to the range-of-probabilities issue commonly applied to the modern risk-assessment process is called the **Monte Carlo Simulation**, based on the mathematics of winning (or losing) at blackjack, roulette, and other casino games.

The concept is to simulate many initiator events (in our case: “Fire breaks out”), and for each simulation, select a random probability within the allowable range for each node. For each run, then calculate the probabilities for each outcome with the selected probabilities for each node and store the final answers.

For example, let’s look at just the top path: sprinkler system sprays water (Node 1) and sprinkler system alerts fire department (Node 2), leading to minimal damage.

From our verbal description table, we assigned the probability of Node 1 being yes as the range from 0.75 to 0.90. Similarly, the probability of Node 2 being yes also ranges from 0.75 to 0.90, since both were considered “very likely”.

Fire up the computer (pun intended), and let’s see what happens.

We start the simulation. For the first run, the computer picks random numbers between 0.75 and 0.90 for both nodes. Let’s say it picked 0.82 for Node 1 and 0.76 for Node 2. The computer then multiplies them together to get the probability of minimal damage, resulting in an answer of **0.62**. That number is stored.

Now the computer runs the simulation again. The second time the computer picks two different random numbers between 0.75 and 0.90, say 0.78 for Node 1 and 0.89 for Node 2. Then it multiplies these two together to get a probability of **0.69** for minimal damage along this path. The number is also stored.

This seems simple, and it is simple, but as the number of nodes and branches increases, the number of calculations increases exponentially. That’s why we use a computer. There are

software packages available to run these simulations – it can't really be done easily on a spreadsheet or by hand.

Note that doing this twice doesn't really tell us much. So we tell the computer to run the simulation, for example, 10,000 times. That is, assuming my factory caught on fire 10,000 times, what would be the outcomes I could expect? Of course I don't expect my factory to catch on fire 10,000 times, but the simulation will show the **range** of outcomes I could expect even if it catches on fire only once.

Note there was no magic to selecting 10,000 runs. You can run the simulations a million times if you want. Or just 50 times. The important thing is to run it enough times that you develop meaningful results. You can tell when you've run it enough times when if you run it more times you get essentially the same answer.

After running the simulation 10,000 times and storing all the probabilities for each outcome, we have a large set of numbers. The best way to visualize the results is with a graph, as follows:

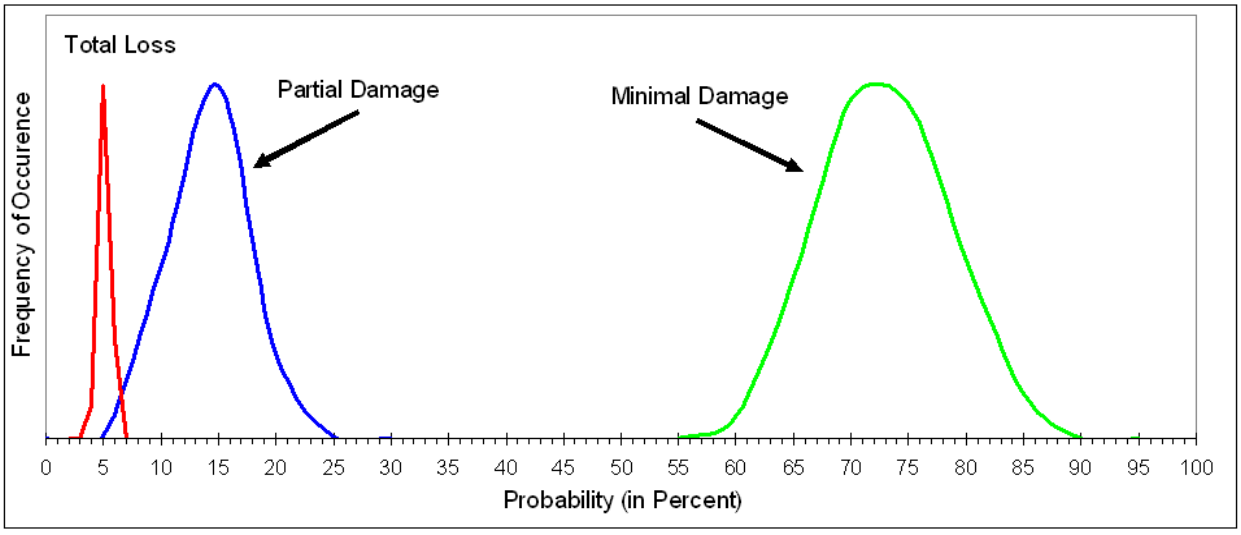

Figure 6 – Probability Distributions of the Fire Problem

Reading the Graph

So how do we read the graph? Probability is on the X axis, and the Frequency of Occurrence is on the Y axis. The scale of the Y axis doesn't matter. You can think of it in this case as "Number of Fires", since that was the initiator of the simulation which we ran 10,000 times. We have one curve for Total Loss, one for Partial Damage, and one for Minimal Damage. For our example, the shape of each is called a "bell curve" since is looks somewhat like a bell.

Let's look first at the Minimal Damage curve. We interpret it this way: the odds of only minimal damage in a fire ranges from 55% to 90%, but most of the time it ranges from about 65% to 80%. My worst case scenario is that the odds of only minimal damage is about 55% to 65%. My best case scenario is that my odds of only minimal damage is about 80% to 90%. Overall, my 'average' odds of only minimal damage is about 73% (the peak of the green curve).

In contrast, look at the Total Loss curve. The odds of a total loss ranges from 3% to 7%, with an 'average' probability of 5%. Note the range of probabilities is much narrower than the Minimal Damage curve. I feel pretty good that I've got the total loss issue well covered, since this is a small probability.

The Partial Damage curve is intermediate to the other two, with an 'average' odds of 15%, with most of the data ranging from about 10% to 20%.

(For the mathematically inclined, all three of these curves are ideally bell curves (normal distributions), since the range of probabilities for each node are assumed to be uniformly distributed. We can also apply statistics, such as mean and standard deviation, to characterize the results. If this interests you, the authors of this course also have a Probability course and a Statistics course on PDHOnline which explore these concepts in more depth.)

Consequence Evaluation

So now what? The next step is a consequence evaluation. **A consequence is the measurable loss.** Most of the time the consequence is in dollars, but it could also be in, for example, human lives.

For our example, let's talk dollars.

So when I say "Minimal Damage" what do I mean? How about-$20,000 in damage and lost inventory. What is "Partial Damage"? How about $150,000 in building damage, damaged equipment and lost inventory.

So the ultimate questions come down to things like: If there is a 15% percent chance of a fire doing $150,000 in damage, should I spend $5,000 in a backup system to call the Fire Department to improve my odds? Or, what face value and deductible makes sense for my fire insurance policy?

These are the business questions. We have outlined the process to establish risk – the business decisions are addressed by others once the risk assessment is completed.

We did not include in our example the risk that people could die in our fire, but that could certainly be addressed. Factors for that analysis could include such things as distance to the nearest exit, distribution of exit signage, degree of smoke generated, number of people in the building, degree of training and number of fire drills per year, etc. The business decisions get more complex (and certainly more emotional), but they are still, ultimately business decisions.

Summary of the Risk Assessment Steps

To summarize, a formal risk assessment includes the following major components:

A ***Potential Failure Mode Analysis*** workshop or study, where the critical failure modes of the structure, project or component are identified. For each potential failure mode, the initiator is identified, and then the progression of steps leading to failure are developed.

Event Trees are then constructed for each potential failure mode, where the steps leading to failure are graphically portrayed as nodes and branches of a tree. Complex events are further broken down until each node has only two possible outcomes.

Probabilities at each node are assigned. Several options are available from deterministic (if available), to subjective probabilities, all the way to expert elicitation, or any combination.

If ranges of probabilities are selected, a ***Monte Carlo Simulation*** is run to simulate the outcomes of the event tree a large number of times.

A ***Consequence Evaluation*** is then done to determine what, if anything, to do next in terms of repair, upgrades, rehabilitation, changing designs, training, signage - whatever will reduce risk and the probabilities of failure.

Advanced Topics

There are a number of additional aspects to the risk assessment process that we have yet to discuss. A couple of these aspects are presented below.

Load Frequency Analysis

Our example risk assessment was based using a fire as the initiator of the event tree. We made the assumption that a fire broke out, and then evaluated the possible outcomes. A more complete analysis would include an assessment of the probability of fire breaking out in the first place. Such an assessment is termed a **Load Frequency Analysis**. That is, how frequently does the “load” of a fire affect our factory?

For our discussion of Load Frequency Analysis, let's shift gears somewhat and talk about hurricanes in New Orleans. Over the long term, we consider that hurricanes making landfall in New Orleans is a random process. That is, the weather has no memory of past events, so each year starts with a clean slate. Such a process is called "Poisson". No, not as in a French fish; as in Siméon Denis Poisson (1781–1840) who first published the concept in 1837.

For New Orleans, let us define a hurricane as a storm with a central pressure of less than 955 millibars. We have instrumental data for about the last 50 years. Over the past 50 years there have been three hurricanes that met the pressure criterion and made landfall. Three in 50 years means we 'statistically' expect one hurricane every 17 years. The 17 years is called the **return period** for major hurricanes making landfall in New Orleans. The return period is used to compute the following probability curve.

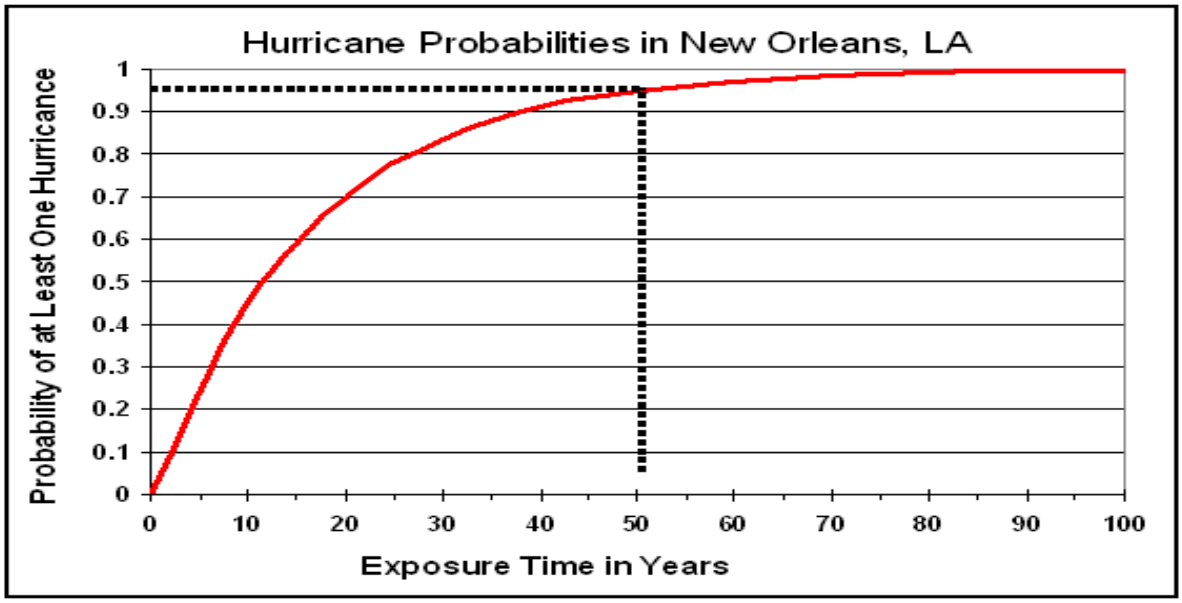

Figure 7 – The Probability of a Hurricane Making Landfall in New Orleans

This curve tells us the probability that our levees will be "loaded" by a hurricane over any specific time period. For example, if I am responsible for upgrading the levees to perform for the next 50 years, I see that there is a 95% probability of at least one major hurricane making landfall during that period, as shown by the dotted line on the graph. This curve would be considered one Load Frequency Curve to be used in the risk assessment.

Photos - Levee Failures in New Orleans and Missouri

We could also approach this issue another way. Levees often fail when they are overtopped. Therefore, in addition to the frequency of hurricanes, I might also be very interested in how high the water got during the historical hurricanes, since this would take into account tides and storm surge as well as rain.

The following curve was developed using the New Orleans historical data, and shows the probability in any given year of a specified increase in water level, where 0 is normal sea level.

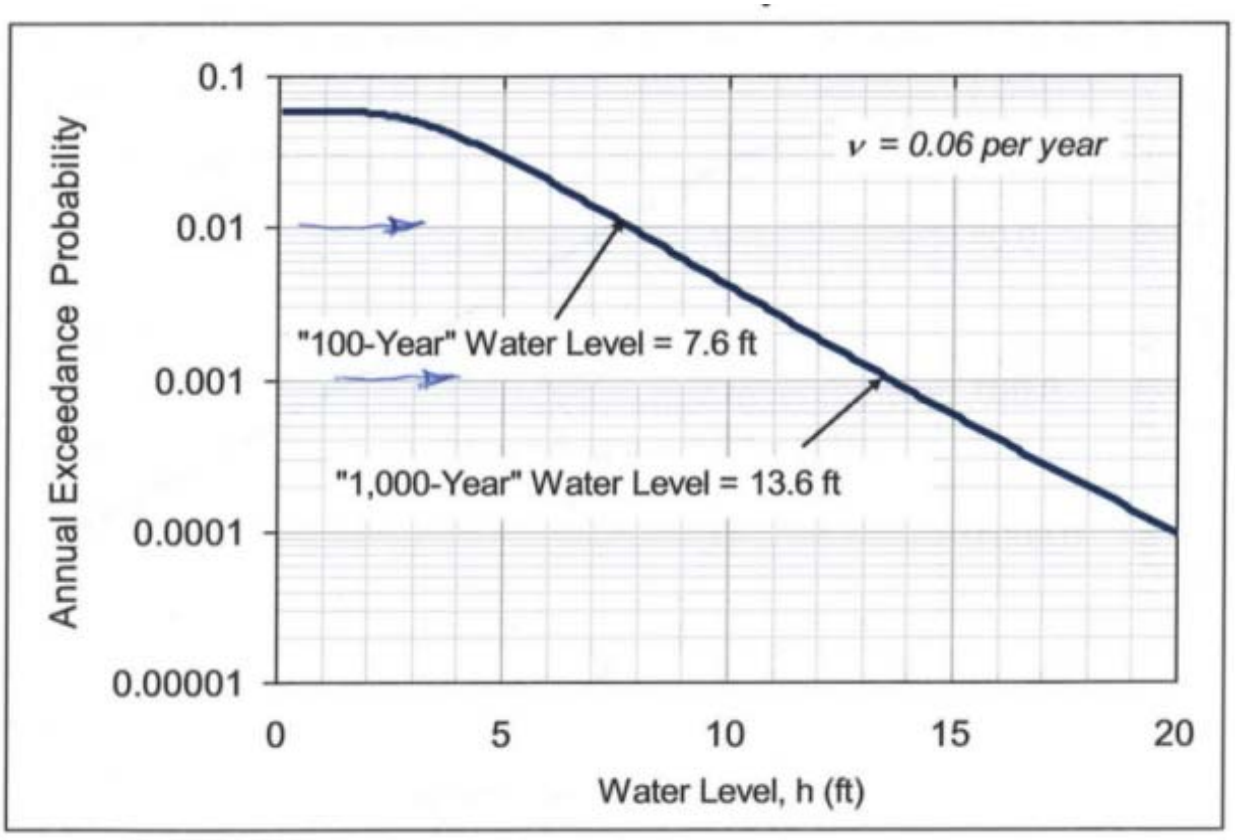

Figure 8 – Annual Exceedance Probabilities for Floodwater Levels - New Orleans

The graph shows us, for example, that for any given year, the chances of the 100 year flood level of 7.6 feet is .01, or 1 in 100 (that is, after all, the definition of the 100 year flood). We can also see that the annual probability of a 20 foot deep flood is 1 in 10,000 – so that would be the "10,000 year flood". During Hurricane Katrina in 2005, flood levels reached over 10 feet, which would be about the "500 year flood", based on the graph.

So why did levees overtop during Katrina? Because they were designed to withstand the 100-year flood level, not the 500-year flood level. A trade-off was made originally by the US Corps of Engineers between the probability of overtopping and the cost to construct, since higher levees are more expensive to build.

This curve could be used as a Load Frequency Curve for the New Orleans levees, since it tells us how often a "load" of a certain flood level might occur. I could also use both the Figure 7 and Figure 8 curves together to get the probability that a major hurricane will occur AND overtop the levees.

The information in one or both curves can be included mathematically into the risk assessment process using the Monte Carlo simulation technique we discussed earlier. What we are doing here is to assign a probability to an initiator, instead of just assuming it will occur, as we did for the fire problem.

Fragility Curves

Fragility curves are sometimes used to represent the probability of a failure or adverse outcome as a function of a certain load or other parameter. This can best be explained by examples.

Back to our levees for a minute. Here is a fragility curve for levees in general.

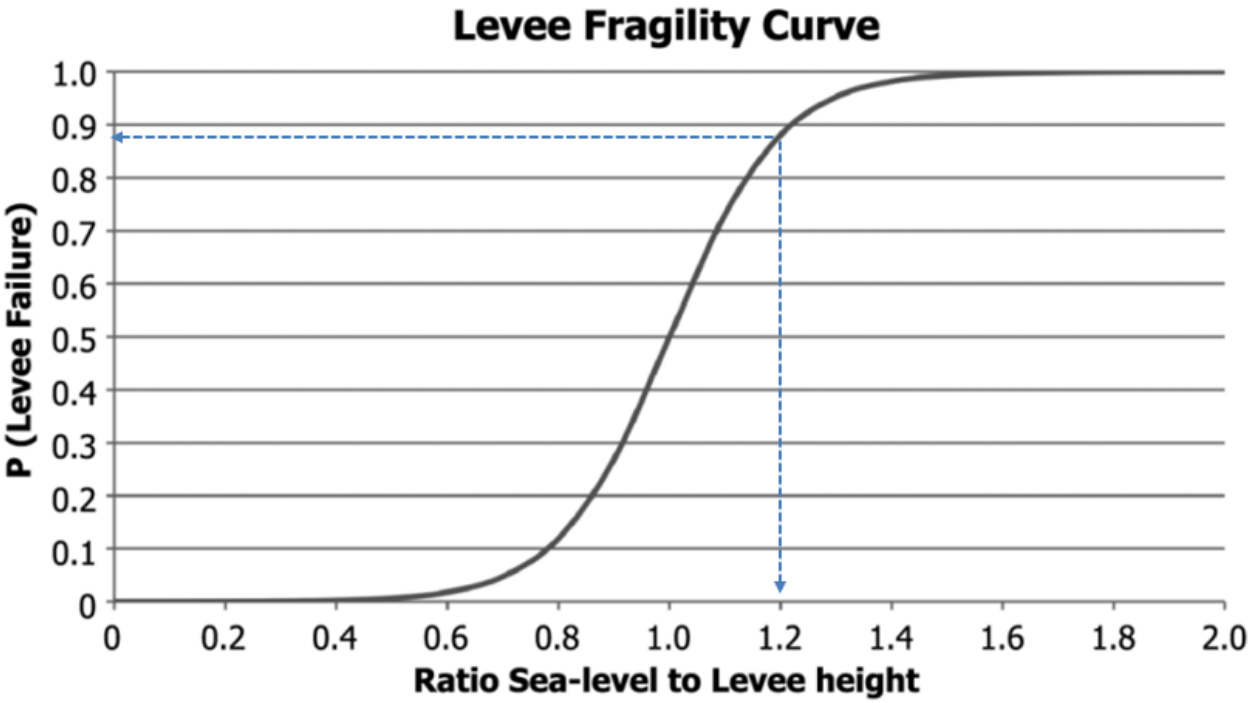

Figure 9 – Fragility Curve for a New Orleans Levee

This curve tells us that if sea level during a hurricane exceeds the levee height by about 20% (a ratio of 1.2), there is a about an 88% chance that the levee will fail. It is also interesting to note that a there is a finite probability of failure even if the water never overtops (ratio of 1.0 and less). These failures are due to other processes, like seepage through the levee.

Fragility Curves can also be created for any manufactured component. Consider a steel beam used in the floor of an office building. Let's say for my risk assessment for this building, I need to consider the possibility that the steel beams may buckle under load. I do know that all steel beams are not created equal, even for the same dimensions and material. So for a given load near the buckling point, some may buckle and some may not. In my research I came across the following Fragility Curve for the type of steel beams in the office building.

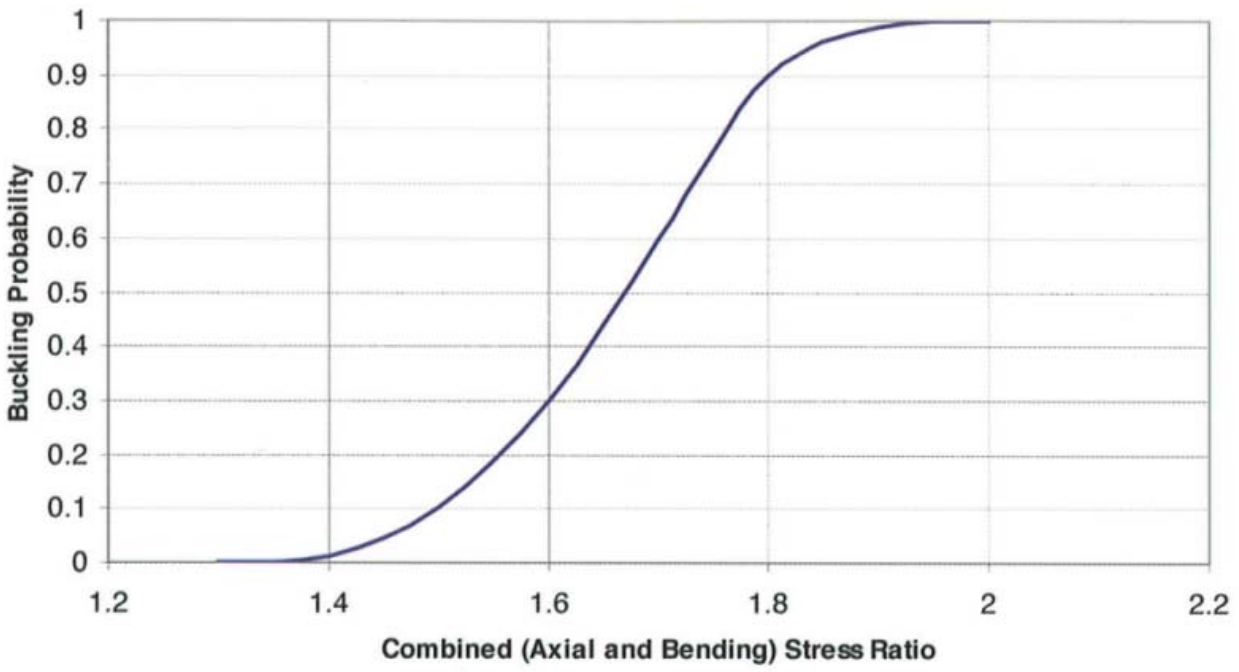

Figure 10 – Fragility Curve for a Particular Steel Beam

The curve illustrates the probability that a beam will buckle relative to the combined stress ratio - the load in my building. For example, I see from the curve that if the combined stress ratio reaches 1.6, then there is a 30% chance that the beam will buckle.

If, however, the combined stress ratio hits 1.8, then there is a 90% chance the beam will buckle.

I can put this curve into my event tree where I am evaluating the loads the floor may encounter. It may turn out that the loads are small enough that the probability of buckling failure is near zero, or it may turn out that I really need bigger beams to lower the risk to an acceptable level. Or, typically, somewhere in the middle.

Fragility curves can be determined directly from lab testing, field experience, or can be created using Monte Carlo Simulations, mathematical modeling, or even the judgment of a panel of experts. They can be applied to almost anything we manufacture or construct, from steel beams to levees in New Orleans to individual parts in the Space Shuttle. The curves can be included before the event tree as part of the initiator description, and/or within the event tree at the appropriate nodes. A number of different curves could be used in different parts of the event tree.

You can see why this discussion is listed under ‘advanced topics’. It can get very complicated very quickly.

There is, however, a risk inherent in this level of complexity. There is always the danger that when the complete analysis is done and the results are in, I won't understand what the source of my risk really is. That is, it may not be clear where I should bring resources to bear to improve the outcomes. A good risk assessment, therefore, balances all the possible level of details against the KISS principle: Keep it Simple, Stupid! If I can't understand the results, then what good was it?

What is an Acceptable Risk?

We end this course with an important yet difficult discussion. Clearly everything we do, manufacture, or construct brings with it some risk. During this course we have limited ourselves to discussing risks that can usually be quantified or translated into dollars.

But what about human life? Driving a car, smoking cigarettes, flying in an airplane, crossing a busy street all carry risk of bodily harm. Risk is inescapable. So as engineers and engineering managers, what metric do we use to balance the risk of bodily harm versus the needs of our society versus unreasonable cost or complexity?

Since risk is inescapable and we don't have an infinite number of dollars, we are forced to define "acceptable risk" as probabilities. The failure of Teton Dam in 1976, as shown on the photos earlier in this course, led the Bureau of Reclamation to formalize the concept of risk. Out of their initial work came the following figure. The figure defines what is considered an "acceptable" risk of fatalities often used in the risk analyses for industrial facilities, nuclear power plants, and dams (Georisk, 2011).

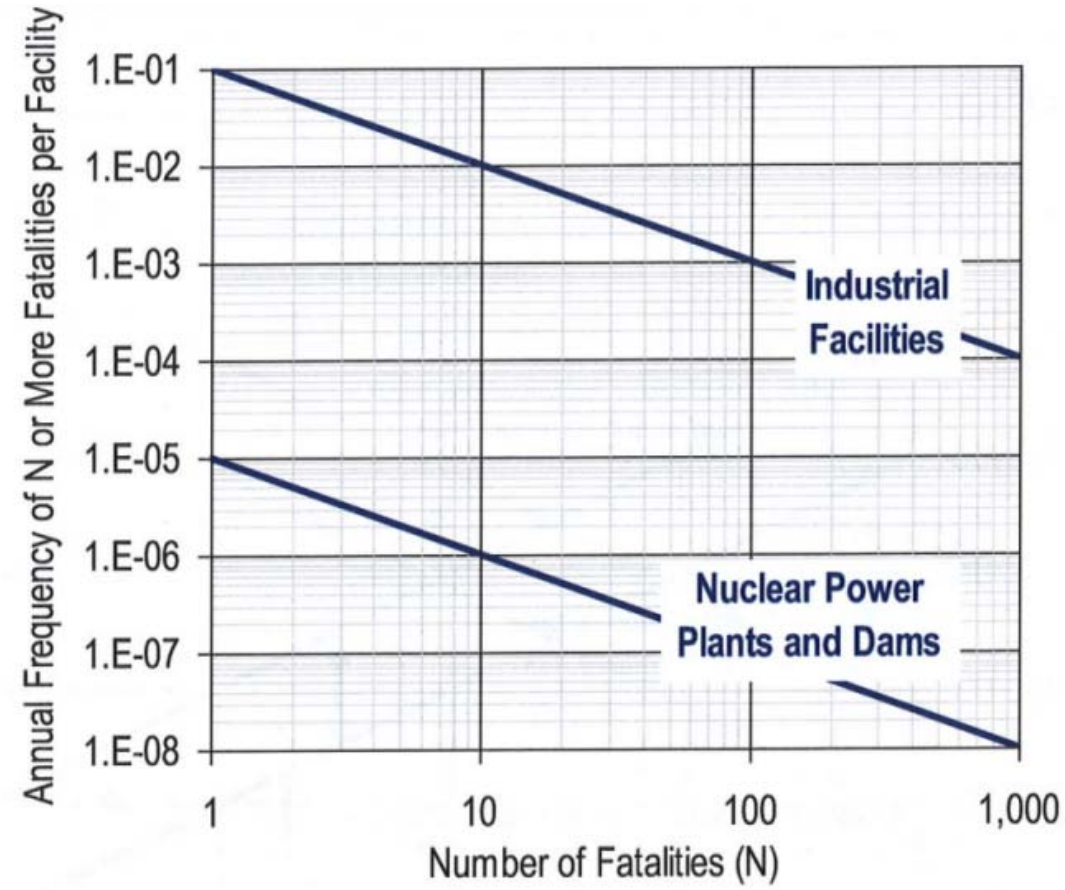

Figure 11 – Acceptable Probabilities of Fatalities at Various Facility Types

For example, the plot tells us that for a nuclear power plant or dam, the "acceptable probability" of 10 fatalities per facility per year is 1.E-06 or 1 in 1,000,000 – one in a million. However, for industrial facilities, the "acceptable probability" of 10 fatalities per facility per year is only 1 in 100. What this says is that dams and nuclear power plants must be 10,000 times safer than industrial facilities to meet the criteria shown on this plot.

Of course we would all prefer that every facility had a zero risk of fatalities. But if we enforced that concept, there would be no power plants, no dams, no industrial facilities. At the same time, you wouldn't be allowed to drive a car, fly in a plane, or ride a train. Or have a gas furnace, a fireplace, a stove, a clothes dryer. Or eat solid food or receive medical treatment. Not to mention how many people get hurt falling down stairs every year.

The Bureau of Reclamation has been completing risk assessment screening for all of its "high hazard" dams against the plots shown above. From that they derived an action plan based on the probability of failure calculated for each dam, as follows:

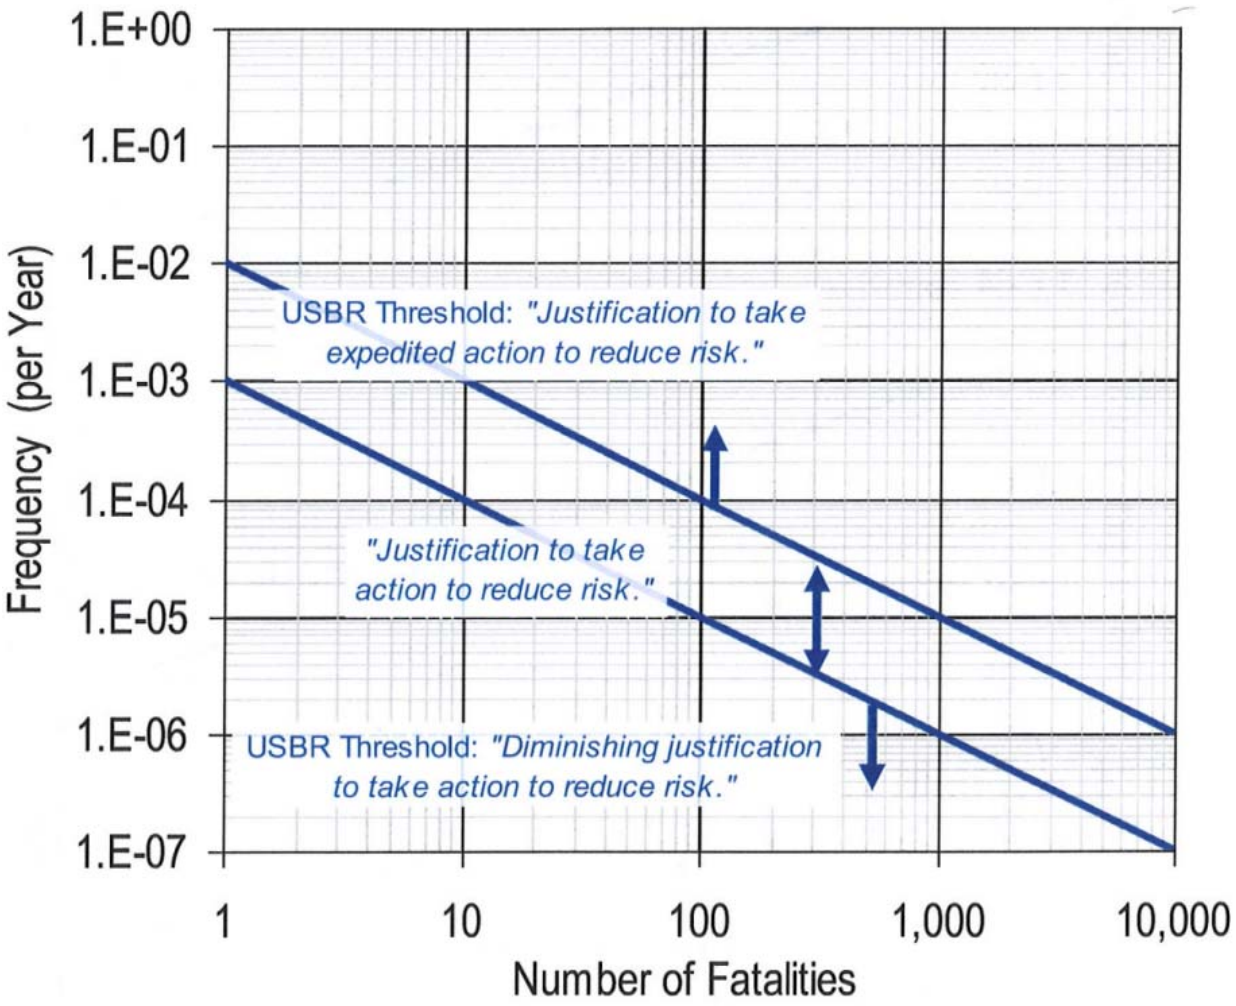

Figure 12 – US Bureau of Reclamation - Guidelines for Dams

As shown on the plot, the USBR uses the risk assessment process to determine where to place scarce resource dollars for repair, rehabilitation or upgrade of its dams. The thresholds shown determine the level of justification for risk-reduction actions. This doesn't mean anything will happen fast, just that there may be justification or expedited justification for spending effort and money to reduce risks.

Conclusion

The material presented in this course provides an overview of the formalized risk assessment process. The formal process is currently used primarily to identify weak points in large engineered structures or groups of structures such as dams and levees. Its purpose is to best define where limited repair, rehabilitation or upgrade dollars should be spent to get the best

"bang for the buck". Bang for the buck could translate into increased reliability, lower operation and maintenance costs, or lowered risk of loss of human life. Every situation is different.

As time passes, the principles and methodologies described here will trickle into more and more engineering fields, and will likely not only be used to evaluate risks associated with existing components and structures, but also be used more up-front, during the design of new structures. Engineers by their very nature pride themselves on conservatism in design. But as materials and labor become more and more expensive, engineers will be put under tremendous pressures to design to be "good enough". On the surface, the risk assessment process described here seems to provide a way to quantify what "good enough" means in terms of probability of failure.

Because the process can be made complex and convoluted and the purview of relatively few experts, however, abuse or misuse is certainly possible. Will the risk-assessment paradigm go the way of "value engineering" as a euphemism for "make it cheaper"? Time will tell. As engineers, we must remain eternally vigilant.

References Cited:

Georisk-2011; Conference on Geotechnical Risk Assessment and Management, June 26-28, 2011; Atlanta, GA. American Society of Civil Engineers

Reagan, R., Mostellar, F., and Youtz, C.; 1989; *Quantitative Meanings of Verbal Probability Expressions*, Journal of Applied Psychology Vol 74 Number 3, pages 433-442.

U.S. Bureau of Reclamation, 2010; *Best Practices in Dam Safety Risk Analysis*, Version 1.3, Denver, CO, Feb 2010

Vick, S.G.; 2002, *Degrees of Belief, Subjective Probability and Engineering Judgment*; Reston, VA. American Society of Civil Engineers

