

An Industry Standard Risk Analysis Technique

A.Terry Bahill, PE, University of Arizona
Eric D. Smith, University of Texas at El Paso

Abstract: This article presents the industry standard technique for risk analysis, namely a graph or a plot with frequency of occurrence on one axis and severity of the consequences on the other. It points out the following common errors in implementing this technique: failing to differentiate between types of risk, using estimated probability rather than frequency of occurrence, using ordinal instead of cardinal numbers, using different ranges for severity and frequency, using an inappropriate combining equation, using linear instead of logarithmic scales, explaining only intermediate risks, and confusing risk with uncertainty. Then this article points out ways to ameliorate these mistakes. It presents an algorithm for deriving severity scores. It discusses many factors that affect human perception of risk. Finally, it discusses the mitigation benefit versus cost tradeoff that is always made in risk management.

Keywords: Risk Analysis, Risk Management, Project Management

EMJ Focus Areas: Systems Engineering

History of Risk Analysis

Risk is an expression of the potential harm or loss associated with an activity executed in an uncertain environment. Since 1662, it has been written that risk had at least two components. "Fear of some harm ought to be proportional not only to the magnitude of the harm, but also to the probability of the event" (Arnauld and Nicole, 1996). This is the first use of the words *magnitude of harm and probability of the event*. There are some ancient Greek, Chinese, and biblical sources that have the concept of risk, but they do not have these words. It is unlikely that any older source has the words, because probability was not invented until the seventeenth century (Pascal, 1654).

The **risk matrix** was used in 1973 by Witt and in 1978 by Hussey. Use of the risk matrix was standardized by Mil-Std 882B in 1984 (DoD, 1984).

Risk plots (graphs, charts, figures) showing frequency of occurrence versus the severity of consequences were used in risk assessments of nuclear power systems (Joksimovic, Houghton, and Emon, 1977; Rasmussen, 1981). They were defined by Kaplan and Garrick (1981) and applied to civil engineering projects by Whitman (1984). Use of the risk plot was popularized by the Defense Systems Management College (1986); however, in spite

of its popularity, the execution of this technique often has flaws, among them are:

1. Failing to differentiate between levels and categories of risk,
2. Using estimated probability of the event rather than its frequency of occurrence,
3. Using ordinal numbers instead of cardinal numbers for severity,
4. Using different ranges for severity and frequency,
5. Using an inappropriate combining equation,
6. Using linear scales instead of logarithmic scales,
7. Explaining only intermediate risks while seeming to ignore high and low risks,
8. Ignoring risk interactions and severity amplifiers, and,
9. Confusing risk with uncertainty.

The purpose of this article is to elucidate these methodological mistakes and point out ways to ameliorate them. Its purpose is not to replace this technique with statistics or to fix human decision mistakes in doing a risk analysis.

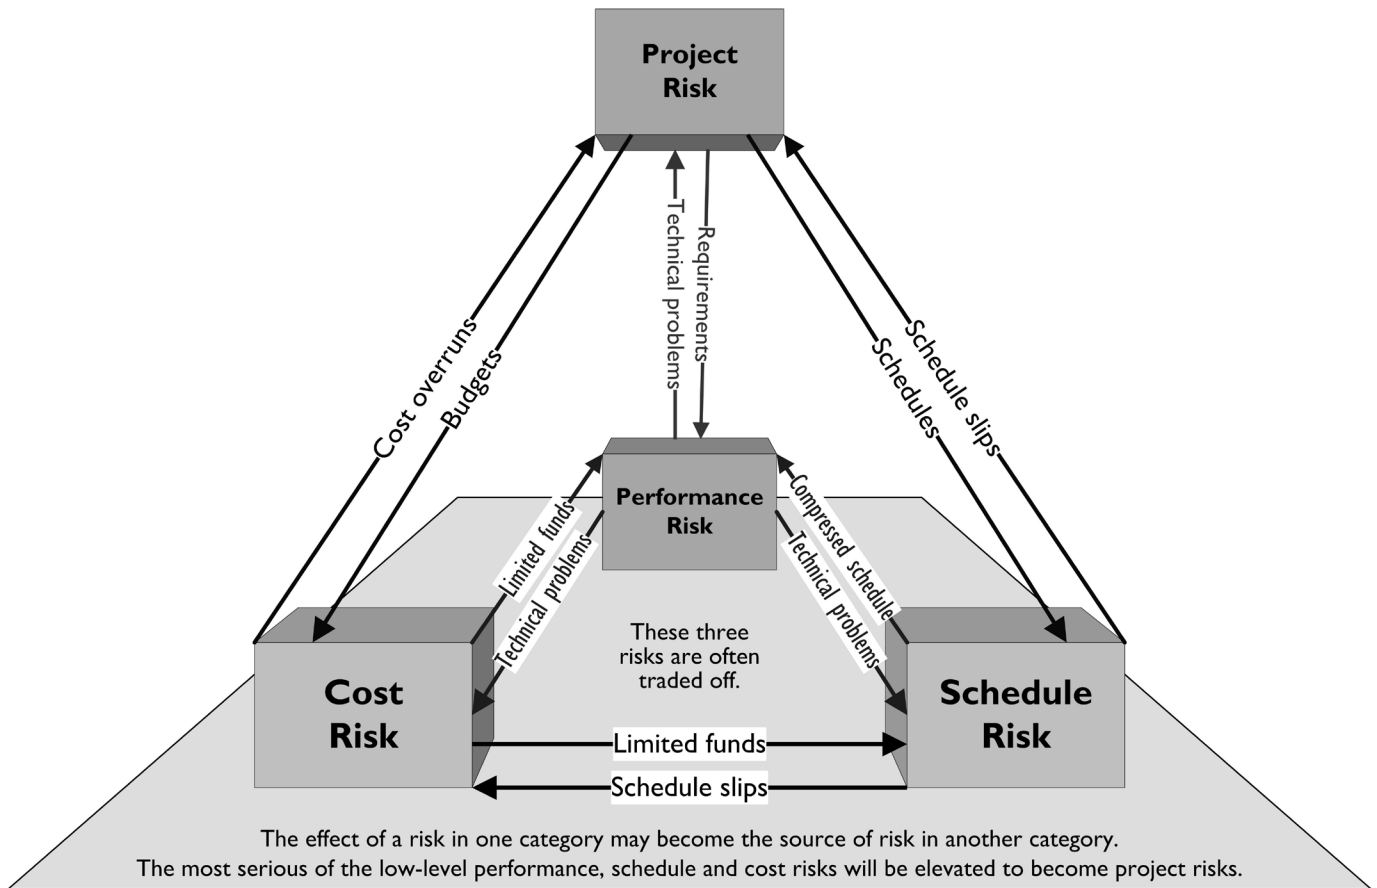
Risk analysis is an important part of engineering. There are general guidelines for how it should be done, but there is no *one correct way* to do a risk analysis (Haimes, 1999; Kirkwood, 1999; Buede, 2000; Carobone and Tipett, 2004; Blanchard and Fabrycky, 2006). This article presents the most common risk analysis technique (Bahill and Karnavas, 2000) and shows problems associated with it. Bernstein (1996) presents the history of mankind's efforts first to understand and then to manage risk.

There are many levels of risk: (1) system risk including performance, cost, and schedule of the product, (2) project risk, (3) business risk including financial and resource risks to the enterprise, and (4) safety, environmental, and risks to the public. The effect of a risk in one category may become the source of risk in another category. Interrelationships between the most common of these risks are shown in Exhibit 1. The low-level system risks of performance, cost, and schedule constitute the higher-level project risks, which might in turn compose enterprise risks. Exhibit 1 and most industry risk analyses only explicitly present the bottom two levels.

Risks can be placed on sequential trees or graphs in order to make their interrelationships clear. Within these structures, the probability of a risk occurring will be dependent on predecessor risks. Another technique for categorizing risks is to use an SE framework to differentiate risks by different framework areas. Risk analyses can then be conducted solely within framework areas, or only across framework areas that are adjacent to one another.

Good risk management will not prevent bad things from happening, but when bad things happen, good risk management will have anticipated them and will reduce their negative effects.

Exhibit 1. Interrelationship of Risks



Frequencies are Better Than Probabilities

The measure of risk is the severity of the consequences times the frequency (or probability) of occurrence. Because humans evaluate probabilities poorly (Gigerenzer, 2002; Clausen and Frey, 2005), we will use frequency of occurrence instead. The following example based on Gigerenzer (1991) shows this superiority of frequencies over probabilities.

A woman has a positive mammogram. What is the probability that she has cancer? The following information is available about women of her age: (a) the probability that a woman tested in this lab has breast cancer is 0.8%; (b) if a woman has breast cancer, the probability that she will have a positive mammogram is 88%; (c) if a woman does not have breast cancer, the probability that she will have a positive mammogram is 7%. This problem seems ideal for Bayes' Rule. Let D be the disease, in this case breast cancer. Let TR be the test result, in this case a positive mammogram. Then

$$P(D | TR) = \frac{P(TR | D) \times P(D)}{P(TR | D) \times P(D) + P(TR | \bar{D}) \times P(\bar{D})}$$

$$P(D | TR) = \frac{0.88 \times 0.008}{0.88 \times 0.008 + 0.07 \times 0.992} = 0.09$$

But this solution is difficult for most people, because people do not think this way. Most people focus on the causally relevant link—(b) if a woman has breast cancer, then the probability that she will have a positive mammogram is 88%.” They ignore (a) the probability that a woman tested in this lab has breast cancer is 0.8%, because it is ‘only statistics,’ not causal. And they ignore (c)

if a woman does not have breast cancer, the probability that she will have a positive mammogram is 7%, because it is not about the category in which we are interested, namely women with cancer. Exhibit 2 shows how most people reason.

- Now let us look at the frequency of occurrence approach:
- (a) eight of 1,000 women tested in this lab have breast cancer
 - (b) of the 8 women with breast cancer, 7 will have a positive mammogram
 - (c) of the remaining 992 women, 70 will have a positive mammogram

Imagine a sample of women who have a positive mammogram. What fraction of them has cancer? The answer obviously is 7/77 = 9%. These three pieces of data (a), (b) and (c) can be represented with a tree, as shown in Exhibit 3.

From this tree, it is easy to see that 7 of the 77 (9%) positive mammograms have cancer. This tree is easier for most people to understand than Bayes' Rule. The frequency of occurrence approach uses a mental operation that people perform quite well: partitioning a set of cases into exclusive subsets. The successive partitioning of sets is well represented by a tree structure. The Bayesian rule approach uses probabilities and is accordingly difficult for those without an education in probabilities.

The Case Study

Since the 1950s over 80 million Cub Scouts have built five-ounce, wooden cars and raced them in pinewood derbies. Bill Karnavas and Terry Bahill ran a Pinewood Derby for their Cub Scout pack for five years. Each year they did a risk analysis, identified the most severe risks and ameliorated them (Bahill and Karnavas,

Exhibit 2. Most People Ignore the Evidence Statements (a) and (c) and Erroneously Conclude that Most Women With Positive Mammograms Have Cancer

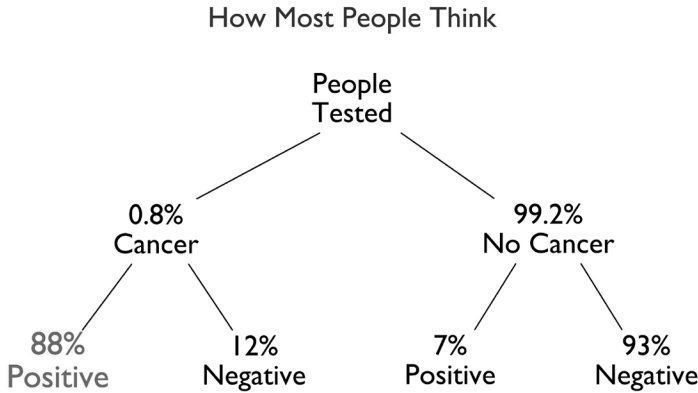
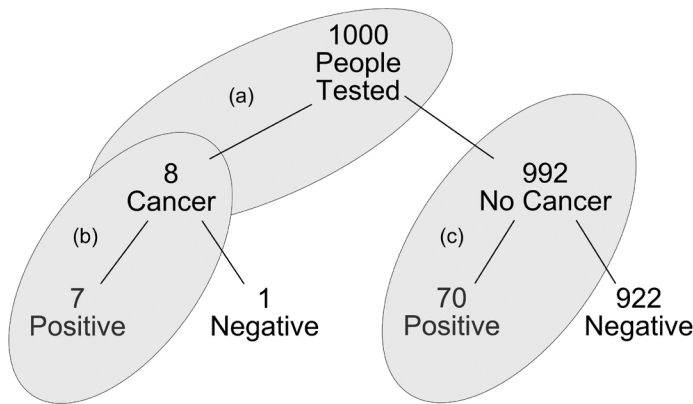


Exhibit 3. Tree Diagram for the Evidence Statements (a), (b) and (c) of the Breast Cancer Example



7 of the 77 positive mammograms have cancer.

2000). For example, in the first two years the most severe risks were lane biases and getting the cars into the wrong lanes. To ameliorate these risks they switched to a round robin race format. This allowed each scout to race often, to race throughout the whole event, and to race more of their friends. They used six

Exhibit 4. Pinewood Derby Performance Risk Matrix

Failure mode	Potential effects	Frequency (failures per 1000 heats)	Severity	Risk
Temporary sensor failure	Heat must be re-run	10	5	50
Lane biases	Some cars have unfair advantage	1000	0	0
Collisions between cars	Heat must be re-run	20	5	100
Mistakes in judging or recording	Wrong car is declared the winner	0.1	1000	100
Human mistakes in				
weighing cars	Some cars have unfair advantage	30	10	300
allowing modifications	Some cars gain unfair advantage	15	10	150
placing cars in wrong lanes	Wrong car is declared the winner	20	1000	20000
lowering the starting gate	Cars can be broken, race could be unfair	1	100	100
resetting finish line switches	Heat must be re-run	10	5	50
wasting time	People get annoyed	10	100	1000

rounds because that gave each car two races in each lane and still kept the whole event reasonably short. The effects of lane biases were eliminated, because each car ran in each lane the same number of times. Getting the cars into the correct lanes is discussed later. Bahill and Karnavas (2000) divided risk into four broad categories: performance, cost, schedule, and project. In Exhibit 4, we look at performance, which is often called technical performance. In Exhibit 4, risk is frequency times severity.

Rationale for the Performance Failure Frequencies

Typically, in each Pinewood Derby, about 250 heats (races) were run in a little over four hours, for an average of one heat per minute. Data were collected in four years, so failures are given in terms of failures per 1000 heats or failures per 1000 minutes.

Temporary failure of sensors at top or bottom of track. This happened a little more than twice a year, yielding 10 failures per 1000 heats.

Lane biases. All Pinewood Derby tracks probably have one lane faster than another, so the frequency is 1000 failures per 1000 heats.

Collisions between cars. Originally, about one out of ten heats had collisions; this was unacceptably high. Collisions were probably caused by imperfections where two sections of track were joined. When a wheel hit such an obstruction, the car bounced out of its lane. By carefully aligning and waxing the joints, the failure rate was reduced to 20 failures per 1000 heats.

Mistakes in judging or recording the results. This was never observed after we switched from human to electronic judging. With computer control and electronic judging, this failure mode is very unlikely. Failures were estimated at 0.1 failures per 1000 heats.

Human mistakes in weighing cars. Every year an over-weight car probably snuck through. This car would then have run in, on average, about 7 heats, meaning 30 failures per 1000 heats.

Human mistakes in allowing modifications (such as adding graphite or weight) after inspection. Each year one scout probably added graphite in the middle of his heats, so the frequency of an unfair heat would be 15 out of 1000.

Human mistakes in placing cars in the wrong lanes. This was observed five times in the first derby, therefore, 20 failures per 1000 heats.

Human mistakes in lowering the starting gate. In the first derby this happened once.

Human mistakes in resetting finish line switches. This happened about 10 times per 1000 heats.

Human mistakes wasting time. This happened about 10 times per 1000 heats.

Algorithm for Computing Severity Values

The following algorithm is used for computing values for the severity of the consequences

1. Assign a frequency of occurrence (F_i) to each failure mode.
2. Find the failure mode that has the most severe consequences, call its value S_{worst} .
3. For each other failure mode, ask: "How many of these failures would be equally painful to the Worst?" Call this N_i . This can be rephrased as, "Cumulatively, how many of these failures would have an equal impact to the Worst?"
4. Compute the severity for each failure mode as $S_i = S_{\text{worst}} / N_i$
5. Normalize the severity values so that their range equals the range of the frequency values.
6. Compute the estimated risk using a combining equation.
7. Prioritize the risks to show which are the most important.

The results of step 1 are shown in the frequency column of Exhibit 4. In our Pinewood example, step 2 identifies "Placing cars in the wrong lanes, thereby declaring the wrong car to be the

winner," to be the most severe failure mode. In step 3 each failure mode is compared to this worst failure mode. For example, our domain expert told us that rerunning 200 heats would be equally painful to declaring the wrong car a winner. Step 4 now produces candidate values for severity.

Severity values should be cardinal measures, not ordinal measures. Cardinal measures indicate size or quantity; ordinal measures merely indicate rank ordering. (A mnemonic for this is that ordinal is ordering, as in rank ordering.) Cardinal numbers do not merely indicate that one failure mode is more severe than another – they also indicate how much more severe. If one failure mode has a weight of 6 and another a weight of 3, then the first is twice as severe as the second. Steps 2 to 4 in the above algorithm help generate these cardinal numbers.

Step 5 requires normalizing the severity values so that their range equals the range of the frequency values. For Pinewood, the frequency values range over five orders of magnitude, therefore the severity values were adjusted so that they also ran over five orders of magnitude, as is shown in the severity column of Exhibit 4. The reason for this is explained with Exhibit 5.

The examples in the left and right halves of Exhibit 5 have the same frequency of failure, but the severity column in the right half has been turned upside down. The risk columns are different, but the rank order columns are identical. Severity had no effect! In general, if two items are being multiplied and they have different ranges, the one with the bigger range has more weight.

This problem of different ranges could be insidious if the severity values were based on physical terms. For example, if the frequency of occurrence for the events ran from one to ten events

Exhibit 5. The Problem With Different Ranges

Example 1				Example 2			
Frequency	Severity	Risk	Rank Order	Frequency	Severity	Risk	Rank Order
10^{-1}	1	1×10^{-1}	1	10^{-1}	6	6×10^{-1}	1
10^{-2}	2	2×10^{-2}	2	10^{-2}	5	5×10^{-2}	2
10^{-3}	3	3×10^{-3}	3	10^{-3}	4	4×10^{-3}	3
10^{-4}	4	4×10^{-4}	4	10^{-4}	3	3×10^{-4}	4
10^{-5}	5	5×10^{-5}	5	10^{-5}	2	2×10^{-5}	5
10^{-6}	6	6×10^{-6}	6	10^{-6}	1	1×10^{-6}	6

Exhibit 6. People Might Not Weigh Severity and Probability Equally

	Severity of Consequences (S)	Probability (P)	Product (SxP)	Perceived risk	$R = S^2 \times P$
Low severity, high probability	0.1	0.9	0.09	Low	0.009
Medium severity, medium probability	0.3	0.3	0.09	Medium	0.027
High severity, low probability	0.9	0.1	0.09	High	0.081

per year and the economic losses ranged from a thousand dollars to a million dollars, then the rank order of the products would be determined strictly by the dollar amounts: the frequency of occurrence would have no effect. This example could be fixed by normalizing the economic loss as follows:

$$\text{normalizedEconomicLoss} = 9 \times 10^{-6} (\text{economicLoss} - 1000) + 1$$

Step 6 is to compute the estimated risk. For Exhibit 4 we computed the risk as the severity of the consequences times the frequency of occurrence. Other formulae have been used. Ben-Asher (2006) says that people do not equally weigh severity of the consequences and probability, as shown in Exhibit 6. For example, people buy collision insurance for their cars, but they seldom insure their tires against premature wear. Therefore, he computes risk as $R = S^2P$.

Some people have computed risk with the Severity plus Probability of Failure minus the product of Severity and Probability of Failure ($R=S+P-SP$) (Kerzner, 2002), but this formula does not perform satisfactorily. For example, if you set the severity to 1 (assuming a range of 0 to 1), then the probability of failure could be reduced from, say, 10^{-1} to 10^{-6} without changing the risk. Furthermore, if either the probability or the severity is zero, then the risk should be zero, but this equation does not produce this result; therefore, we do not use such an equation.

Carl Burgher (personal communication), a UMR project manager, uses Exposure instead of Probability—that is, Risk = Severity x Exposure. The value for exposure should estimate the subject's likely endangerment to the risk. Exposure does not have Probability's time-discounting errors, such as hyperbolic discounting. It would be hard to calculate the probability that a particular person would get flu this season, but it would be easy to estimate that a first-grade teacher would have a greater exposure to viruses than a college professor would.

The Failure Modes and Effects Analysis (FEMA) technique also includes the difficulty of detection in the product (Carbone and Tippett, 2004), so that ($R=S \times P \times DD$). In some cases this is useful, but we will not use it in this article.

All of the following equations have been used in published literature to compute risk.

- Risk = Severity of Consequences × Frequency of Occurrence
- Risk = Severity of Consequences × Likelihood of Occurrence
- Risk = Severity of Consequences × Estimated Probability
- Risk = (Impact + Likelihood)/2
- Risk = Severity + Probability - (Severity × Probability)
- Risk = Severity × Probability × Difficulty of Detection
- Risk = Severity² × Probability
- Risk = Severity × Exposure

Each of these equations is good in certain situations, but for the reasons given above, we will only use the top two equations in this article.

Step 7 of the algorithm prioritizes the risks in order to show which are the most important. Often this process is merely that of selecting the risks with the largest values and pointing them out to the decision maker. Later in the Project Risk Management section of this article, we will show a graphical technique for highlighting the most important risks. Also, there is a general-purpose technique for prioritizing things other than risks—

things like goals, customer needs, capabilities, risks, directives, initiatives, issues, activities, requirements, technical performance measures, features, functions, and value engineering activities (Botta and Bahill, 2007).

Discussion of the Performance Risks

The most important failure mode in this Pinewood Derby risk analysis is *Human mistakes in placing cars in the wrong lanes*. In a previous risk analysis of our Pinewood Derby (Chapman, Bahill, and Wymore, 1992), we found that in order to satisfy our customer, we had to pay the most attention to *Mistakes in judging and recording* and *Lane biases*; therefore, we designed our system to pay special attention to these failure modes. We designed a computerized judging and recording system that had a very low failure frequency; indeed this failure mode is the one least likely shown in Exhibit 1. Furthermore, we designed a round robin tournament where each car races twice in each lane; therefore, lane biases had no effect. When the effect of a failure mode was eliminated, we set the severity to 0. We did not remove the failure mode, because we wanted to show that it was considered; therefore, the failure modes that were very important in our early risk analysis were not important in the risk analysis of this article. In our redesigned derby, humans were the most important element.

Other Columns

A failure modes and effects analysis should include a column indicating who should do what in response to each failure mode. This was not included in this example because all risks were managed by the same person. A root-cause analysis would have included "So what?" columns. For example, the failure mode *Collisions between cars* has the potential effect of "Heat must be re-run." After this we could ask, "So what?" which might produce the response, "This takes additional time." We could ask again, "So what?" producing, "If there were many re-runs the whole Pinewood Derby schedule might slip," etc. Other columns could also be used to document what and when corrective actions were taken.

Risk communication is an important part of a risk analysis. What should be communicated, and to whom? In general, the risks are communicated to the program manager along with proposals for mitigating these risks. In this case, the Pinewood Derby Marshall served as the program manager, so he was told that *Placing cars in the wrong lanes* was the most significant risk. Then he was shown that the race schedules have lane 1 on the left and lane 3 on the right. This matches the perspective of the finish line judges (lane 1 is on their left and lane 3 is on their right), but it is the opposite for the starter at the top of the track; therefore, we suggested printing regular schedules for finish line judges and printing mirror image schedules in another color for use by the starter.

Graphical Presentation

Risk is often presented graphically with frequency of occurrence on the ordinate axis and severity of consequences on the abscissa axis. Each of the axes is linear with a range of 0 to 1, as in Exhibit 7a. Risks in the upper right corner are the most serious and should be handled first. The problem with this technique is that all of the risks in Exhibit 4 would be squashed onto the x-axis, because they all have frequencies at or below 10^{-1} (except for our anomalous lane bias). If some action resulted in a reduction of frequency of occurrence from 10^{-2} to 10^{-4} , we want the estimated risk to change, and in these graphs, it would not change. Using logarithmic scales as in Exhibit 7b would help, and indeed this may be the best solution if quantitative data are available.

Exhibit 7a. Linear Plot of $xy=0.001$, $xy=0.01$, $xy=0.1$ and $xy=0.6$, With Normalized Data of Exhibit 1

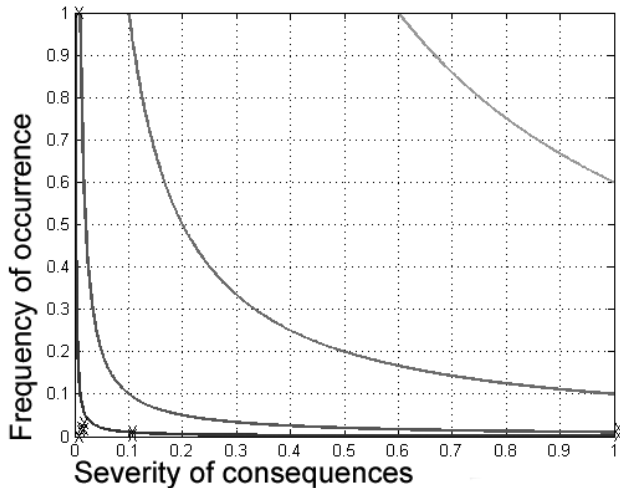


Exhibit 7b. Loglog Plot of $xy=0.001$, $xy=0.01$, $xy=0.1$ and $xy=0.6$, With Normalized Data of Exhibit 1

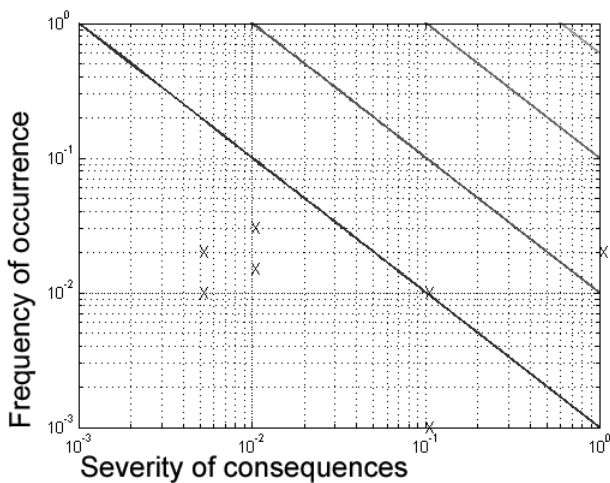


Exhibit 8. Schedule Risk Matrix for a Pinewood Derby

Failure Mode	Potential Effects	Frequency, failures per 1000 minutes	Severity	Risk
Loss of electric power	Delay races until electricity is restored	0.02	100	2
PC hardware failure	Cannot compute, cancel the Derby	1	10000	10000
Server failure	No schedules	0.1	100	10
Failure of commercial software	Cannot compute, cancel the Derby	0.01	10000	100
Failure of custom software	Cannot compute winners	0.05	20	1
Software interface failure	Data is lost and races must be re-run	0.02	10	0.2
Bad weather	Batteries lose power, cannot determine winners	0.008	1000	8
Forgetting equipment	Lose an hour of setup time	0.01	100	1

Schedule Risk

Next, in Exhibit 8, the schedule risk is evaluated. The likelihood that all parts of the system functioned correctly so that any given heat could be run was determined. Typically, in each Pinewood Derby, about 250 heats (races) were run in a little over four hours, for an average of one heat per minute. Data was collected over four years, so we calculated how many minutes of failure should be expected for each 1000 minutes of racing.

Rationale for the Schedule Risk Frequencies

Schedule risk frequencies were computed on a failures per minute basis. Because a heat was run each minute, this is the same as the frequency of delaying an individual heat. Many other criteria could have been used, such as: *Did the Derby start on time? Did it end on time? Did each Divisional Contest start on time? Did each end on time?* The following explains how the schedule risk frequencies were computed.

Loss of electric power. During the 14 years for which we have data, this area of Tucson has lost electric power three times (once when lightning hit a transformer, once when a car hit a pole, and once when the Western Power Grid went down), for a total outage of about three hours. So for any given 1000 minute interval, the power would be expected to be out on average 0.02 minutes. This means that there was a 2% chance of delaying one race during these four years.

Personal computer hardware failure. During the four years of Derby data collection, there were three personal computer hardware failures. It took a total of 45 hours to fix it or get a replacement: therefore, there was one minute of failure per 1000 minutes of operation. This does not count periods of upgrading hardware or software.

Server failure. Schedules were printed weeks in advance, so server failure was unlikely to delay a race, 0.1 minutes of failure per 1000 minutes.

Failure of commercial software. 0.01 minutes of failure per 1000 minutes. This failure frequency is low, because the programs were Unix and DOS-based.

Failure of custom software. 0.05 minutes of failure per 1000 minutes.

Software interface failure. In four years, there was one interface failure: It took an hour to fix, hence 0.02 minutes of failure.

Bad weather. In 14 years, Tucson has had one afternoon where the temperature dropped 30 degrees in two hours. This caused

the NiCad batteries to lose power, and the race was shifted to a manual, paper-based system. The switch took one hour, yielding 0.008 minutes of failure per 1000 minutes.

Forgetting equipment. Over these 14 years, on one road trip, equipment was forgotten and a trip back to the lab was needed. This caused a delay of a little over one-hour, yielding 0.01 minutes of failure per 1000 minutes.

When assigning the severity values we made evaluations such as, “*Forgetting our equipment* 100 times would be equally painful as *Hardware failures* that would cause us to cancel a Derby.”

The probability that this system would be operating mistake free for any given heat is:
 $(1 - 0.00002)(1 - 0.001)(1 - 0.0001)$ etc., which equals 0.99878.

This is good reliability. It resulted from doing studies like these, finding the weak link, and redesigning the system.

As can be seen in Exhibit 8, the present weak link is *Personal computer hardware failure*; therefore, when we put on a Pinewood Derby we keep a spare computer in the car. Actually, we are even more paranoid than this. The last time we put on a Pinewood Derby, we designed it with a round robin (best time) with electronic judging, but we also designed a backup system that required no electricity—it used a round robin (point assignment) with human judging. All of the equipment and forms for the backup system were in the car.

Failure modes and effects analyses usually do not include acts of nature, like our *Bad weather* category. We have included

it in our analysis, because we think that if the process is designed with these items in mind, then although we cannot prevent them the system can better accommodate such failures.

Bahill and Karnavas (2000) also did a cost and safety risk analysis for the Pinewood Derby. That analysis is not included in the present article.

Project Risk Management

Most of this case study has been about analysis of performance and schedule risks; however, project managers are more often worried about project risk which has been defined as “the potential of an adverse condition occurring on a project which will cause the project to not meet expectations” (Georgetown University, 2008).


For the Pinewood Derby, some project risks are (1) *Failure to compute schedules*, (2) *School denies use of facilities*, (3) *Boy Scouts of America (BSA) stops supporting Pinewood Derbies*, (4) *Parents or children lose interest*, and (5) *Lawsuits*. These are presented in Exhibit 9. *Failure to compute schedules* has a large likelihood value, because in our first Pinewood Derby we were not able to compute schedules that met the original requirements, so we had to change the requirements or else run a single elimination tournament, which does not require schedules. *Failure to compute schedules* is an example of a low-level schedule risk that was important enough to be elevated to a project risk. If we have data for failure frequency, we use the term frequency; otherwise, if we are guessing the future, we use the term likelihood.

The *School denies use of facilities* failure mode includes the school demanding an exorbitant fee, having another big school

Exhibit 9. Project Risk Matrix for a Pinewood Derby

Id	Failure Mode	Potential Effects	Likelihood Value	Severity Value	Risk	Assigned to
1	Failure to compute schedules	Run a single elimination tournament	0.8	0.4	0.32	Bill
2	School denies use of facilities	Run the Derby in someone’s backyard	0.7	0.9	0.63	Terry
3	BSA stops supporting Derbies	Buy commercial car kits	0.1	0.1	0.01	Harry
4	Parents or scouts lose interest	Cancel the Derby	0.1	1.0	0.10	Harry
5	Lawsuits	Buy insurance	0.9	0.9	0.81	Terry

Exhibit 10. Project Risk Chart for a Pinewood Derby



event scheduled for the same day, and the school not giving access to electric power outlets. We also considered unavailability of equipment, but this is not a big risk because the Boy Scout pack owns everything we need except for a PC. We considered loss of knowledge, but our system is well documented. We considered loss of key people, but we have many qualified parents.

For project risk, we do not have accurate numbers, so 0 to 1 scales are appropriate. The likelihood values are *not* probabilities; rather they are relative measures indicating likelihood: the school denying use of their facilities is considered more likely than the *Boy Scouts of America stopping its support of Pinewood Derbies*. We are not, however, saying that seven out of ten times we expect the school to turn down our request. *Id* is the identification number that will be used in Exhibits 10 and 12. The failure modes are assigned to Bill Karnavas, Terry Bahill, and Harry Williams, the Cub Master. Exhibit 9 is analogous to a risk register.

In Exhibit 10, these data are put into a risk management tool that is commonly used in industry, namely a plot of relative likelihood versus severity of failure. In such plots, the ordinate is often labeled probability or potential probability, but they are *not* probabilities. Uncertainty and unknown unknowns prevent calculation or even estimation of real probabilities. If this scale were probability, then any item with a probability greater than, say, 0.6 would not be a risk—it would be a mandate for a design change. We think the best label for this axis is *relative likelihood*. The word *relative* emphasizes that it is the relationships between (or ranking of) risks that is being illustrated. The word *likelihood* is not used in a mathematical sense, but rather in its dictionary sense to indicate the risks that are most likely. These *relative likelihood* estimates are typically very conservative. They are derived with the following attitude: “If we built the system today, with lofty capabilities, without improving the design, doing nothing about

this risk item, and if the worst possible circumstances occurred, then what is the possibility that this risk item will harm us?” Initially, some risk items are expected to be in the red zone (the upper right corner); otherwise, program managers will think that the designers are not trying hard enough to get more performance or less cost. The relative likelihood scale is not linear—a risk with a relative likelihood value of 0.8 is not considered twice as probable as one with a value of 0.4.

The description just given is the way this tool is actually used in industry and government (British Cabinet Office, 2008). Statisticians are appalled, because of its use of ordinal numbers instead of cardinal numbers. Often, qualitative scales are given to help assign values. The numbers given in Exhibit 11 are typical. Clearly, these are not cardinal numbers.

In Exhibit 10, risk number 5, *Lawsuits*, is the most important risk; therefore, we buy insurance to transfer this risk to an insurance company. Risk number 2, *School denies use of facilities*, is the next most important risk. We should spend a lot of time making sure it does not happen and preparing backup plans. The third most important risk is number 1—*Failure to compute schedules*—so we should keep an eye on this. The curves should be rectangular hyperbolas, not circles, because they are iso-risk contour lines for severity times likelihood.

The most important project risks are put in the project risk register and are analyzed monthly and at all design reviews. For this Pinewood Derby example, *School denies use of facilities* is an important risk, so we tried to mitigate it. We talked to the school principal and also found an alternate site. (In one year, our pack ran the derby at the Northwest Emergency Center.) This ameliorated the risk, and at the next review we presented the risk plot of Exhibit 12. The forte of this risk management tool is showing changes in risk items.

Exhibit 11a. Qualitative Scale for Frequency of Failure

How often will this failure occur?	Your processes ...	Value
Almost never	will effectively avoid or mitigate this risk based on standard practices.	0.2
Occasionally	have usually mitigated this type of risk with minimal oversight in similar cases.	0.4
Sometimes	may mitigate this risk, but workarounds will probably be required	0.6
Frequently	cannot mitigate this risk, but a different approach might	0.8
Almost always	cannot mitigate this risk, no known processes or workarounds are available.	1.0

Exhibit 11b. Qualitative Scale for Severity of Schedule Slippage

If the risk event happens, what would be the severity of the impact on the schedule?	Value
Minimal impact, slight changes could be compensated by available program slack	0.2
Additional activities required, you would be able to meet key dates	0.4
Minor schedule slip, you would miss a minor milestone	0.6
Program critical path affected	0.8
You would fail to achieve a key program milestone	1.0

Exhibit 12. Pinewood Derby Project Risk Plot After Risk Mitigation





Exhibit 13. Hydraulic Model for Risk of a Commercial Airplane Flight



In most risk analyses plots that we have seen in industry, most of the risks were in the center of the plot. Most of the risks that would have been in the upper-right corner had already been ameliorated, whereas most of the risks in the lower left corner were considered unworthy of the tracking effort. This distribution could also be caused by a common human decision mistake. People overestimate events with low probabilities (e.g., being killed by a terrorist, dying in an airplane crash, or winning the state lottery) and they underestimate high probability events (e.g., adults dying of cardiovascular disease) (Tversky and Kahneman, 1992).

Discussion

The highest priority risks are often tracked with technical performance measures (TPMs) (Oakes, Botta, and Bahill, 2006). TPMs are critical technical parameters that a project monitors to ensure that the technical objectives of a product will be realized. Typically, TPMs have planned values at defined time increments, against which the actual values are plotted. Monitoring TPMs allows trend detection and correction, and helps identify possible performance problems prior to incurring significant cost or schedule overruns. The purpose of TPMs is to (1) provide visibility of actual versus planned performance, (2) provide early detection

or prediction of problems requiring management attention, (3) support assessment of the impact of proposed changes, and (4) reduce risk.

Preemptive system design can reduce risk. Preemptive design includes these practices: risk management, rapid prototyping, failure modes and effects analysis, sensitivity analysis, identification of cost drivers, reliability analysis, anticipating problems, quality engineering, variation reduction, complexity analysis, 6 Sigma, poka-yoke (translate as fool proofing), design for manufacturability, design for test, design for assembly, and design for cost (Bahill and Botta, 2008). The Pinewood Derby design used the first seven of these practices.

The project risk list should not be synonymous with the requirements. (1) Requirements have attributes such as priority, stability and difficulty; however, the difficulty of satisfying a requirement should not be confounded with project risk. For example, one of the original requirements for the Pinewood Derby was that in a 12-car, 6-heat, round-robin race every car was supposed to race every other car exactly once (Bahill and Karnavas, 2000). This requirement was impossible to satisfy and it had to be relaxed. A performance risk notes that there is a chance that the requirement cannot be met. When it is confirmed that the requirement cannot be met, it now becomes a problem not a risk; therefore, this problem was handled with requirements management, not with risk management. (2) Risks identify unknown effects. If something is spelled out as a requirement, then it is not an unknown. (3) Risks also arise from uncertainty in the environment and external entities; however, a requirement cannot apply to things outside the boundary of its system. (4) Requirements are orders of magnitude more numerous than risks. (5) Risks should be related to the mission, high-level goals or capabilities, and not to numerous low-level requirements.

Exhibit 13 shows a summarizing example of risk management: the hydraulic model for risk of a commercial airplane flight. The total risk is the sum of latent hazards (e.g., poor quality maintenance, poor employee moral) and direct hazards (e.g., lightning, rain, microbursts). The risk is largest during takeoff and landing. The dispatcher's job is that of balancing risk with resources.

Here is a scenario for risk management in a commercial aircraft situation. You are the dispatcher in Salt Lake City. It's a clear morning. Temperature is in the low 30s °F. A dozen skiers are scheduled to fly to Jackson Hole, but there is a 1/2" of ice on their runway. How would you manage this risk? Possible risk management actions are:

Transfer:	divert to Idaho Falls and bus the skiers, thus transferring the risk to Greyhound Bus Lines
Eliminate:	use a helicopter
Accept:	send the plane as scheduled
Mitigate	Request removal of ice from runway Change runways (orthogonal direction) Change equipment (different type of aircraft) Change crew (switch the lander from co-pilot to pilot)

These actions produce the acronym TEAM. There is, however, an additional complication: the icy runway hazard object interacts with temperature and cross winds. The hazard is amplified by temperatures just above freezing and it is amplified by cross winds. If there are cross winds, the hazard might be

mitigated by using a runway co-linear with the wind. How about ignoring the risk? This is not acceptable: there is no I in TEAM.

The action *eliminate* includes reducing the likelihood of the risk event; however, in this case we cannot do much to reduce the likelihood of occurrence. Mother Nature, and therefore, the temperature at Jackson Hole, is not under our control.

This example shows that risk changes with time, that risks are traded off with available resources, and the common ways of managing risks.

In this article, we treated risk as severity of consequences times frequency of occurrence. There are many other factors that humans use when assessing risk. In particular, the following are *severity amplifiers*: lack of control, lack of choice, lack of trust, lack of warning, lack of understanding, manmade, newness, dreadfulness, personalization, recallability and immediacy. *Lack of control*: a man may be less afraid driving his car up a steep mountain road at 55 mph than having his wife drive him to school at 35 mph. *Lack of choice*: we are more afraid of risks that are imposed on us than those we take by choice. *Lack of trust*: we are less afraid listening to the head of the Center for Disease Control explain anthrax than listening to a politician explain it. *Lack of warning*: people dread earthquakes more than hurricanes, because hurricanes give hours or days of warning. *Lack of understanding*: we are more afraid of ionizing radiation from a nuclear reactor than of infrared radiation from the sun. *Manmade*: we are more afraid of nuclear power accidents than solar radiation. *Newness*: we are more afraid when a new disease (e.g., swine flu, SARs or bird flu) first shows up in our area than after it has been around a few years. *Dreadfulness*: we are more afraid of dying in an airplane crash than of dying from heart disease. *Personalization*: a risk threatening you is worse than that same risk threatening someone else. *Recallability*: we are more afraid of cancer if a friend has recently died of cancer. We are more afraid of traffic accidents if we have just observed one. *Immediacy*: A famous astrophysicist was explaining a model for the life cycle of the universe. He said, "In a billion years our sun will run out of fuel and the earth will become a frozen rock." A man who was slightly dozing awoke suddenly, jumped up, and excitedly exclaimed, "WHAT did you just say?" The astrophysicist repeated, "In a billion years our sun will run out of fuel and the earth will become a frozen rock." With a sigh of relief, the disturbed man said, "Oh thank God. I though you said in a *million* years."

Errors in Numerical Values

The data used in a risk analysis have different degrees of uncertainty: some of the values are known with precision, others are wild guesses; however, in addition to uncertainty, all data have opportunity for errors. For example, Ord, Hillerbrand, and Sandberg (2008) calculated that one in a thousand papers published in peer reviewed journals have been or should have been retracted, because they contained serious errors.

There are three types of errors that can be made in assessing risk: theory, model, and calculation (Ord, Hillerbrand, and Sandberg 2008). (1) Using Newtonian mechanics instead of Einstein's relativity theory to describe merging black holes would be an error of using the wrong theory. (2) A modeling error was the direct cause of the failure on the Mars Climate Orbiter (NASA, 2000). The prime contractor, Lockheed Martin, used English units onboard the spacecraft, whereas JPL used metric units in the ground-based model. (3) Errors in calculating values are common. Several hospital studies have shown that one to two percent of drug administrations were the wrong dosage. A study

of papers from *Nature* and the *British Medical Journal* found that about 11% of the statistical results were wrong (Ord, Hillerbrand, and Sandberg 2008). To validate numerical values in a risk analysis, you should analyze the possibility of errors in choosing the theory, errors in making the model, and errors in estimating and computing numerical values.

In cases where risks can be reduced, the level of effort is driven by a cost/benefit analysis. If we assign expected costs to the risks, then we can see our greatest vulnerabilities, but sheer magnitude is not the whole story. We might also look at the sensitivity of total risk dollars to risk handling dollars. By calculating the cost of implementing a particular strategy versus the reduced expected cost of risk, we can make good decisions. In the Pinewood Derby example of this article, we can see that spending hundreds of dollars on equipment boxes to protect computers that are not very likely to get broken gives a very poor cost reduction to spending ratio. If instead we spend a couple dollars on duct tape for the extension cords to avoid human injury with large potential cost risk, the net savings is great. This also works comparing the cost of insurance against the costs of eliminating a risk. For instance, we may decide that the cost risk reduction of buying insurance does not compare to spending the same money for nets or pads to protect derby cars around elevated areas of track.

History teaches us that managers have frequent and common shortcomings when addressing risks. (1) They underestimate the probability of bad things happening. (Prior to the *Challenger* accident, NASA officials put the risk of a shuttle accident at 1 in 100,000: afterwards Richard Feynman computed the odds at 1 in 100.) (2) They underestimate the consequence and chaos caused by the bad things when they do happen. (3) They ignore or fail to recognize that the occurrence of certain risks can trigger the occurrence of other risks. (4) They ignore or fail to recognize that many risks are not independent and may be coupled with other risks. (5) They inadequately investigate unknown unknowns. Experience teaches that no matter how competent and thorough our technical project team, and no matter how competent and thorough our technical peer review team, Mother Nature still provides a few surprises.

Future Work

We are investigating the effects of interdependencies and interactions in risk analyses and tradeoff studies. With interdependencies, a change in some external factor can cause changes in the frequencies or severities of two or more failure modes. For example, in the Pinewood Derby, switching the electronics from DC battery power to AC electric power decreased the severity of *Bad weather* and increased the severity of *Loss of electric power*. As a second example, moving the race venue from Tucson to a remote desert community increased the probability of *Loss of electric power* and increased the severity of *Forgetting equipment*. With interactions, a change in the frequency or severity of one failure mode might cause a change in the frequencies or severities of other failure modes. For example, decreasing the frequency of occurrence of *Lane bias* decreased the severity of *Placing cars in the wrong lanes*.

Summary

The risk assigned to each possible cause is determined by the risk analysis; however, this analysis is not static. Each time the system is changed or a relevant test is performed, the risk analysis needs to be reviewed. As risks are eliminated or reduced, other risks will

increase in relative importance. As a system is deployed, the actual risks will become better known. The Pinewood Derby analysis is presented in this article after a dozen years of data acquisition and is, therefore, far better than our initial analysis. In our first attempts, our biggest risks were believed to be equipment failures. Inclement weather was not actually considered (we were in Tucson after all). And as the system evolved with new instrumentation, the risks changed. The first derbies were not computer based, so extension cord trip hazards or computer failures were not relevant. As the system evolved, the risks evolved with it and so did the risk analysis.

Risk managers should manage the identified risks by regularly reviewing and re-evaluating them, because risk identification and data gathering is not an end in itself. Also, a word of caution: More than one project has made the costly mistake of being too comfortable with their risk ranking, which led to ignoring those areas they initially called low risk.

Risk analysis is not an exact science—there are many ways to do a risk analysis, the risks change with time, there is uncertainty in the data. In this article, we showed the simplest and most common industry risk analysis technique, namely plotting the frequency of occurrence versus the severity of consequences. We showed mistakes that are often made when using this technique and ways to ameliorate them. The purpose of risk analysis is to give the risk manager insight into the project and to communicate his or her understanding with others. The definition of risk was stated in the late 17th century by Arnauld and Nicole (1996): “In order to decide what we ought to do to obtain some good or avoid some harm, it is necessary to consider not only the good or harm in itself, but also the probability that it will or will not occur.” Since then, the engineering literature has embraced this definition stating that risk is the product of frequency of occurrence and severity of consequences.

The object of risk analysis is not just to assess the likelihood of completing a properly functioning project on schedule and cost, but to give managers insight into where to invest resources to reduce risk to an acceptable level, how much time they can take out of a schedule without increasing risk, how much they can add before the gain is marginal, etc. It is, or should be, a decision tool. Risk analysis is an ongoing process that must be redone regularly, or whenever significant changes occur. It is hard to overemphasize that point. It is where many programs falter. Most programs do a risk analysis, fewer build a good risk reduction plan, and fewer still keep it current. Many managers believe they are so up to date on their programs that the time spent is not worth the return. Too bad, because it is one of the most effective ways to show the customer you have your arms around his problem – and life really gets ugly if you cannot show how you tried to anticipate and solve problems, because they tend to happen!

Acknowledgements

We thank George Dolan, Bob Sklar, Gary Lingle, Bruce Gissing, Al Chin, and Frank Dean for comments and discussion about this article. We thank John Michel for pointing out the Arnauld reference.

References

Arnauld, Antoine, and Pierre Nicole, *Logic, or, the Art of Thinking: Containing, Besides Common Rules, Several New Observations Appropriate for Forming Judgment*, 5th Ed., 1662, translated from French in 1996 by Jill Vance Buroker, Cambridge University Press.

- Bahill, A. Terry, and William J. Karnavas, "Risk Analysis of a Pinewood Derby: A Case Study," *Systems Engineering*, 3:3 (2000), pp. 143-155.
- Bahill, A. Terry, and Rick Botta, "Fundamental Principles of Good System Design," *Engineering Management Journal*, submitted for the special issue on Systems of Systems (2008).
- Battalio, Raymond C., John H. Kagel, and Don N. MacDonald, "Animals' Choices over Uncertain Outcomes: Some Initial Experimental Results," *The American Economic Review*, 75:4 (1985), pp. 597-613.
- Ben-Asher, Joseph, "Development Program Risk Assessment Based on Utility Theory," *Proceedings of the 16th Annual International Symposium of INCOSE* (2006).
- Bernstein, Peter L., *Against the Gods: The Remarkable Story of Risk*, John Wiley and Sons (1996).
- Blanchard, Benjamin S., and Wolt J. Fabrycky, *Systems Engineering and Analysis*, Prentice-Hall (2006).
- Botta, Rick, and A. Terry Bahill, "A Prioritization Process," *Engineering Management Journal*, 19:3 (2007), pp. 31-38.
- British Cabinet Office, *National Risk Register*, (2008), http://www.cabinetoffice.gov.uk/reports/national_risk_register.aspx, accessed September 2008.
- Buede, Dennis M., *The Engineering Design of Systems: Models and Methods*, John Wiley and Sons, Inc. (2000).
- Carbone, Thomas A., and Donald D. Tippett, "Project Risk Management: Using the Project Risk FMEA," *Engineering Management Journal*, 16(4) (2004), pp. 28-35.
- Chapman, William L., A. Terry Bahill, and A. Wayne Wymore, *Engineering Modeling and Design*, CRC Press Inc. (1992).
- Clausen, Don, and Daniel D. Frey, "Improving System Reliability by Failure-Mode Avoidance Including Four Concept Design Strategies," *Systems Engineering*, 8:3 (2005), pp. 245-261.
- Defense Systems Management College, *Risk Management: Concepts and Guidance*, Fort Belvoir, VA: DSMC, (1986).
- Department of Defense, *Military Standard 882B* (1984).
- Georgetown University, University Information Services, "Data Warehouse: Glossary" <http://uis.georgetown.edu/departments/eets/dw/GLOSSARY0816.html>, accessed September 2008.
- Gigerenzer, Gerd, "How to Make Cognitive Illusions Disappear: Beyond 'Heuristics and Biases,'" *European Review of Social Psychology*, 2:4 (1991), pp. 83-115.
- Gigerenzer, Gerd, *Reckoning the Risk*, Penguin Books (2002).
- Haimes, Yacov Y., "Risk Management," in *Handbook of Systems Engineering and Management*, Andrew P. Sage, and W. B. Rouse (Eds.), John Wiley & Sons (1999), pp. 137-174.
- Hsu, Ming, Meghana Bhatt, Ralph Adolphs, Daniel Tranel, and Colin P. Camerer, "Neural Networks Responding to Degrees of Uncertainty in Human Decision-Making," *Science*, 310 (2005), pp. 1680-1683.
- Hussey, David E, "Portfolio Analysis: Practical Experience with Directional Policy Matrix," *Journal for Long Range Planning*, 11:4 (1978), pp. 2-8.
- Joksimovic, V., W.J. Houghton and D.E. Emon, "HTGR Risk Assessment Study," in Jerry B. Fussell and G.R. Burdick (Eds.), *Nuclear Systems Reliability Engineering and Risk Assessment*, Society for Industrial and Applied Mathematics (1977), pp. 167-190
- Kaplan, Stanley, and B. John Garrick, "On the Quantitative Definition of Risk," *Risk Analysis*, 1:1 (1981) pp. 11-27.
- Kerzner, Harold, *Project Management: A Systems Approach to Planning, Scheduling and Controlling*, 8th Ed., John Wiley & Sons (2002).
- Kirkwood, Craig W., "Decision Analysis," in *Handbook of Systems Engineering and Management*, Andrew P. Sage, and W.B. Rouse (Eds.), John Wiley & Sons (1999), pp. 1119-1145.
- NASA, *Mars Program Independent Assessment Team Summary Report*, March 14, 2000, ftp://ftp.hq.nasa.gov/pub/pao/reports/2000/2000_mpiat_summary.pdf
- Ord, Toby, Rafaela Hillerbrand, and Anders Sandberg, "Probing the Improbable: Methodological Challenges for Risks with Low Probabilities and High Stakes," <http://www.arXiv.org/abs/0810.5515> (2008).
- Martino De, Benedetto, Dharshan Kumaran, Ben Seymour, and Raymond J. Dolan, "Frames, Biases, and Rational Decision-Making in the Human Brain," *Science* 313:5787 (2006), pp. 684-687.
- Pascal, Blaise, "Fermat and Pascal on Probability," (1654) www.york.ac.uk/depts/maths/histstat/pascal.pdf accessed October 27, 2008.
- Rasmussen, Norman C., "Application of Probabilistic Risk Assessment Techniques," *Annual Review of Energy*, 6 (1981) pp. 123-128.
- Real, Leslie A., "Uncertainty and Pollinator-Plant Interactions: The Foraging Behavior of Bees and Wasps on Artificial Flowers," *Ecology* 62 (1981), pp. 20-26.
- Shafir, Sharoni, Tom A. Waite, and Brian H. Smith, "Context-Dependent Violations of Rational Choice in Honeybees (*Apis Mellifera*) and Gray Jays (*Perisoreus Canadensis*)," *Behavioral Ecology and Sociobiology* 51 (2002), pp. 180-187.
- Tversky, Amos, and Daniel Kahneman, "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5 (1992), pp. 297-323.
- Whitman, Robert V., "Evaluating Calculated Risk in Geotechnical Engineering," *Journal of Technical Engineering*, 110:2 (1984) pp. 145-188.
- Witt, Robert C., "Pricing and Underwriting Risk in Automobile Insurance: A Probabilistic View," *Journal of Risk and Insurance*, 40:4 (1973), pp. 509-531.

About the Authors

Terry Bahill is a Professor of Systems Engineering at the University of Arizona in Tucson. He received his PhD in electrical engineering and computer science from the University of California, Berkeley, in 1975. Bahill has worked with BAE Systems in San Diego, Hughes Missile Systems in Tucson, Sandia Laboratories in Albuquerque, Lockheed Martin Tactical Defense Systems in Eagan MN, Boeing Information, Space and Defense Systems in Kent, WA, Idaho National Engineering and Environmental Laboratory in Idaho Falls, and Raytheon Missile Systems in Tucson. For these companies he presented seminars on Systems Engineering, worked on system development teams and helped them describe their Systems Engineering process. He holds a U.S. patent for the Bat Chooser, a system that computes the Ideal Bat Weight for individual baseball and softball batters. He received the Sandia National Laboratories Gold President's Quality Award. He is a Fellow of the Institute of Electrical and Electronics Engineers (IEEE), of Raytheon, and of the International Council on Systems Engineering (INCOSE). He is the Founding Chair Emeritus of the INCOSE Fellows Selection Committee. Bahill is a registered professional engineer in California and Pennsylvania. His picture is in the Baseball Hall of Fame's exhibition "Baseball as America." You can view this picture at <http://www.sie.arizona.edu/sysengr/>.

Eric D. Smith earned a B.S. in Physics in 1994, an MS in Systems Engineering in 2003, and his PhD in Systems and Industrial Engineering in 2006, from the University of Arizona, Tucson. His dissertation research lay at the interface of systems engineering, multi-criterion decision-making, judgment and decision-making, and cognitive science. He taught for two years in The Boeing Company's Systems Engineering Graduate Program at Missouri University of

Science and Technology. He is currently an assistant professor in the Department of Industrial Engineering, in the Systems Engineering Program, at the University of Texas at El Paso. His interests include risk management, cognitive biases, and complex systems engineering.

Contact: A. Terry Bahill, PE, Systems and Industrial Engineering, University of Arizona, Tucson, AZ, 85721-0020, phone: 520-621-6561, terry@sie.arizona.edu

Appendix: Ambiguity, Uncertainty and Hazards

In economics and in the psychology of decision-making, risk is defined as the variance of the expected value, the uncertainty. Exhibit A1 explains three bets: A, B and C. The p 's are the probabilities, the x 's are the outcomes, μ is the mean and σ^2 is the variance. This exhibit shows, for example, that half the time bet C would pay \$1 and the other half of the time it would pay \$19. Thus, this bet has an expected value of \$10 and a variance of \$81. This is a comparatively big variance, so the risk is said to be high. Most people prefer the A bet, the certain bet.

In choosing between alternatives that are identical with respect to quantity (expected value) and quality of reinforcement, but that differ with respect to probability of reinforcement, humans, rats (Battalio, Kagel, and MacDonald, 1985), bumblebees (Real, 1981), honeybees (Shafir, Waite, and Smith, 2002), and gray jays (Shafir, Waite, and Smith, 2002) prefer the alternative with the lower variance. To model this risk averseness across different situations the coefficient of variability is often better than variance, Coefficient of variability = (Standard Deviation) / (Expected Value).

Engineers use the term *risk* to evaluate and manage bad things that could happen, hazards. Risk is defined as the severity of the consequences times the frequency of occurrence.

To avoid confusion, in the rest of this section, the engineers' risks will be called *hazards*, and decision theory's risks will be called *uncertainty*. We will now explain three important concepts in decision-making: ambiguity, uncertainty, and hazards.

People do not like to make decisions when there is ambiguity, uncertainty, or hazards, so, for the most part, they prefer alternatives that have low ambiguity, low uncertainty, and few hazards.

As we gain information, we progress from ambiguity to uncertainty. Ambiguity means we know very little about the situation. Uncertainty means we have learned enough about the system that we can estimate the mean and variance, although we

definitely do not have to use those terms: in fact, engineers might say that variance is analogous to entropy.

A little while ago, a wildfire was heading toward Bahill's house. He packed his car with his valuables, but he did not have room to save everything, so he put his wines in the swimming pool. He put the dog in the car and drove off. When he came back, the house was burned to the ground, but the swimming pool survived; however, all of the labels had soaked off the wine bottles. Tonight he is giving a dinner party to celebrate their survival. He is serving mushrooms that he picked in the forest while he was waiting for the fire to pass. There may be some hazard here, because he is not a mushroom expert. Guests will drink some of his wine; therefore, there is some uncertainty here. You can be assured that none of his wines are bad, but some are much better than others. Finally, he tells you that his sauce for the mushrooms contains saffron and oyster sauce. This produces ambiguity, because you probably do not know what these ingredients taste like. How would you respond to each of these choices?

Hazard: Would you prefer his forest picked mushrooms or portabella mushrooms from the grocery store?

Uncertainty: Would you prefer one of his unlabeled wines or a Robert Mondavi Napa Valley merlot?

Ambiguity: Would you prefer his saffron and oyster sauce or marinara sauce?

Decisions involving these three concepts are probably made in different parts of the brain. Hsu, Bhatt, Adolphs, Tranel, and Camerer (2005) used the Ellsberg paradox to explain different brain processing of ambiguity and uncertainty. They gave their subjects a deck of cards and told them it contained 10 red cards and 10 blue cards (the uncertain deck). Another deck had 20 red or blue cards but the percentage of each was unknown (the ambiguous deck). The subjects could take their chances drawing a card from the uncertain deck: if the card were the color they predicted they won \$10; otherwise, they got nothing. They could

Exhibit A1. An Explanation of Three Bets: A, B and C

Bet	p_1	x_1	p_2	x_2	μ	σ^2	Risk, uncertainty
A	1.0	\$10			\$10	\$0	None
B	0.5	\$5	0.5	\$15	\$10	\$25	Medium
C	0.5	\$1	0.5	\$19	\$10	\$81	High

just take \$3 and quit. Most people picked a card. Then the subjects were offered the same bets with the ambiguous deck. Most people took the \$3 avoiding the ambiguous decision. Both decks had the same expected value, \$5, because the subjects picked the color. The paradox is that they accepted the uncertain deck, but rejected the ambiguous deck. Hsu et al., recorded functional magnetic resonance images (fMRI) of the brain while their subjects made these decisions. While contemplating decisions about the uncertain deck, the dorsal striatum showed the most activity and while contemplating decisions about the ambiguous deck, the amygdala and the orbitofrontal cortex showed the most activity.

In another fMRI study of ambiguous decisions, when subjects were given a choice of gaining \$20 or spinning a gambling wheel, most chose to keep the \$20. When they were given a choice of

losing \$30 or spinning the wheel, most chose to spin the wheel (Martino De, Kumaran, Seymour, and Dolan, 2006). The fMRI scans suggested that amygdala activity is an emotional signal that pushes subjects to keep sure money and gamble instead of taking a loss; however, not all subjects succumbed to this emotional signal. Those who overcame it had high levels of activity in the orbital and medial prefrontal cortex. De Martino speculates that the prefrontal cortex integrates emotional signals from the amygdala with other cognitive information: people who are more rational do not perceive emotion less, they just regulate it better.

Ambiguity, uncertainty and hazards are three different things, and people prefer to avoid all three. The purpose of this section was to introduce the reader to other meanings of the word of risk.

Engineering Management Tenure-Track Position

College of Engineering and Computer Science

The University of Tennessee at Chattanooga

The College of Engineering and Computer Science invites applications for a full-time, tenure-track, faculty position at the Assistant Professor level in Engineering Management, beginning ideally January 1, 2010 or August 1, 2010 depending on selected candidate's availability. Review of applications will begin immediately and will continue until the position is filled. Applicants must have a bachelor's degree in engineering; master's degree in engineering management, engineering, business, or construction management; and Ph.D. in engineering management or related field. The successful candidate will have the ability to teach undergraduate Engineering Technology Management courses with concentrations in engineering management and construction management and graduate Engineering Management courses both on campus and online, conduct research, advise students, chair graduate students' projects and/or theses, and interact with industry. Relevant industrial experience is preferred. Professional engineering registration or significant progress towards licensure is desirable.

The Engineering Management and Technology Department offers a B.S. in Engineering Technology Management with concentrations in Construction Management and Engineering Management, and M.S. in Engineering Management, which is offered both on campus and fully online, for people with engineering or science backgrounds who have moved or expect to move into managerial responsibilities.

UTC (www.utc.edu) is a metropolitan university, committed to excellence in teaching and applied research. With more than 10,000 students, UTC is the second largest institution in the University of Tennessee system. The College of Engineering and Computer Science (<http://www.utc.edu/EngineeringAndComputerScience>) offers both undergraduate and graduate degrees in Engineering, Technology Management, and Computer Science. The College is also home to the SimCenter with graduate programs (M.S./Ph.D.) in Computational Engineering.

Chattanooga is a thriving, mid-sized city located on the Tennessee River near the center of the Tennessee Valley Technology Corridor. With nearby mountains, trout streams, hiking trails, excellent schools and a revitalized downtown waterfront, Chattanooga offers ample opportunities for balanced family living.

Applications must include a cover letter that explains applicant's interest in the position, a statement of teaching philosophy and research interests, a complete resume, including names, addresses, e-mails, telephone numbers of at least three references, and copies of undergraduate and graduate (M.S. and Ph.D.) transcripts. These documents should be e-mailed to: facultyvitae@utc.edu. Please reference search number F09-2-004 Engineering.

The University of Tennessee at Chattanooga is an equal employment opportunity/affirmative action/Title VI & IX/Section 504 ADA/ADEA institution. Qualified women and minorities are encouraged to apply.