SEER by GALORATH

The Risk Driver Impact Approach to Estimation of Project Cost Risks: Clear Thinking about Project Cost Risk Analysis

by

Evin Stump P.E.
August 2000

Perhaps the most commonly used approach to quantitative estimation of project cost risks is to estimate a "nominal" cost for each work element of the project's work breakdown structure, subjectively assign cost risk distributions to those nominal values, then run a Monte Carlo simulation to obtain a risk distribution of the cost sum. When this is done, the Monte Carlo tool used must have the capability to allow the user to assign correlations between the various work element costs, else serious errors can result. In projects, work elements generally are not designed for statistical independence, nor is it likely they could be in any practical sense, even if we wanted to do so. If there are significant correlations and they are not well specified, the result is a serious underestimation of the risks.

Variations on this basic approach replace ordinary Monte Carlo with its latin squares version, or with a tool that uses the method of moments, or less often, numerical convolution integration.

Any of these approaches is valid assuming that 1) the subjective assignment of risk distributions around the nominal cost estimates can be done appropriately, and 2) the correlation between work elements can be correctly assigned. These assumptions are in general quite difficult to meet, and in my opinion seldom are met, resulting in poorly done quantification of cost risk.

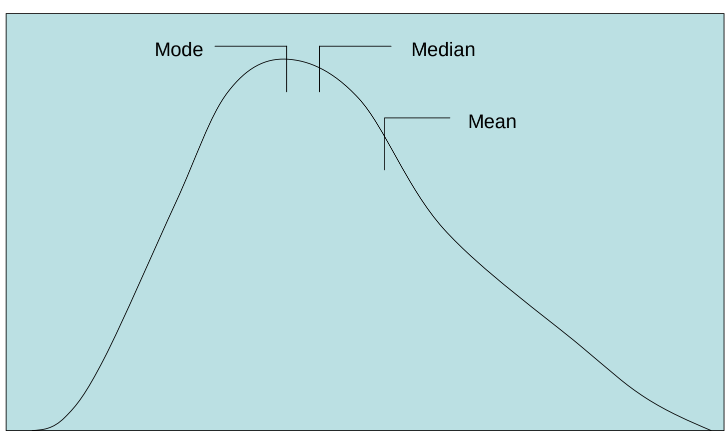

Typical displacements of the mode, median, and mean in a right skewed distribution.

Take first the subjective assignment of risk distributions. I have no quarrel with subjective assignments—we do them all of the time and they are a necessary part of the human experience. I might note that some people of high academic standing claim that there is no such thing as an objective assignment of risk—they are all subjective. I agree. It is true, however, that not all assignments hew closely to the available evidence, and this is a laxity we must guard against.

When we make risk assignments, it is best if we clearly understand what we are assigning values to, and why. For example, what does the nominal cost estimate represent? Sometimes it is said to represent the median cost, that is, the cost that has a 50% chance of being exceeded. Other times it is said to represent the most likely cost, that is, the cost that has the highest probability of occurrence. Still other times it is said to represent the mean or average cost, that is, the average of the costs that would occur if the project were repeated many times (an unlikely event, to be sure!).

Given that subjective assignments of risk distributions are almost invariably right skewed, these quantities--median, mode, and mean—can be significantly different numbers (see figure above). They are generally the same only in symmetrical distributions such as the normal, which is not often used to represent cost risk of a project work element. If the analyst assigning the subjective distribution does not recognize this, and frequently it is not explicitly recognized, the distribution can significantly fail to represent the true situation.

An example occurs when the triangular distribution is used. It typically is assigned based on three points, a, b, and c (see figure at right). The point a is the minimum value deemed likely to occur; the point c is the maximum value deemed likely to occur, and the point b by the geometry of the distribution must be the most likely (modal) value. If the analyst mistakenly deems the b value to be the median or mean, or more likely if the analyst is unclear about what it represents, a serious error can occur.

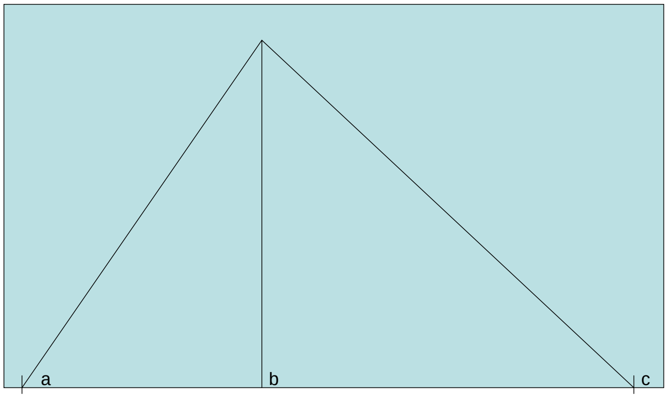

The triangular distribution is frequently used in cost risk analysis. It is defined by three points, a, b, and c.

The error is aggravated the more the distribution is skewed.[1] Often analysts are vague on just what their assignment really represents—mean, median, or mode, or maybe something else entirely.

Another source of error in assigning distributions is the shape of the distribution. The triangular distribution is probably the most widely used distribution for assigning cost risks because it is easy to specify, and it permits easy introduction of skewness, which is a virtual necessity in cost risk analysis. But in principle, the correct distribution to assign is the maximum entropy distribution.[2] This has been defined as the minimally prejudiced distribution that maximizes the entropy subject to constraints supplied by available information. Entropy in turn can be defined as a measure of how weak a hypothesis is compared to any possible competitor.[3]

The maximum entropy distribution in any situation depends on our state of knowledge. Tribus (*op cit*) shows, for example, that if our state of knowledge is limited to a mean value, μ, then the maximum entropy distribution is the exponential distribution with that same mean. He also develops several other maximum entropy distributions for various states of knowledge.

We seldom attempt to define and assign the maximum entropy distribution in cost risk analysis because of practical difficulties in sorting out our true, often convoluted state of knowledge, and performing the often intricate math needed to find the maximum entropy distribution that represents it. Nevertheless, our failure to do so does introduce unknown error.

Another major problem with assigning distributions to work element cost risks, likely the most serious problem, has to do with human visualization of risk limits. It is common knowledge that people underestimate costs more often than they overestimate them. One reason why cost risk distributions are generally skewed right is that the analyst assigning the distribution is concerned that he or she will underestimate a possible overrun, for which there can be adverse consequences, while the consequences of an under-run are generally benign.

[1] The distribution is most likely to be highly skewed in high risk situations. Therefore, in high risk situations it is most important to be clear about what the nominal estimate represents. Unfortunately, in my experience it is precisely in high risk situations that an analyst it most likely to be unclear about this point.

[2] Tribus, Myron, "Rational Descriptions, Decisions and Designs," Pergammon Press, 1969

[3] Mathematically, entropy is defined as $S = -\sum p_i \ln p_i$ where the p_i are assigned probabilities summing to unity.

The difficulty of estimating the upper and lower cost risk limits of a work element in the project environment is highly dependent on how well the work element is defined, and also how similar the work element is to past work elements that have been successfully completed. For example, suppose that a work element designated "Design Actuator" has been estimated as most likely (mode) requiring 250 engineering hours. The estimate was based on two previous actuator designs that required, respectively, 210 hours and 326 hours. The estimate of 250 hours for the instant actuator was based on an engineering assessment of the relative design "complexity" compared to the remembered complexity of the two previous designs.

The estimate has some credibility, but how can risks be assigned to it? The first suspect in an investigation of this issue is the estimate of relative complexity. Could it be off by 10%? 20%? 30%. It often comes down to a gut feel issue, and sometimes engineering gut feel, among experienced engineers, can be quite good. But are there issues other than complexity that bear on this risk? What if, for example, the design effort winds up being hampered by any of the following:

- Unavailability in the appropriate time interval of engineers experienced in actuator design, requiring the design to be done by less experienced engineers who make 20% less in salary and benefits but take 50% more hours to do the job
- A delay in proper specification of the actuator characteristics because of volatility in customer requirements
- An unexpected failure in a test of the actuator prototype requiring some redesign
- A mistake in estimation of the load the actuator must carry, causing some redesign.

The point being made is that the person who made the estimate of 250 hours must, in assigning risk limits to that number, try to visualize a possibly large number of adverse effects, and their likely legitimate consequences. In doing this, he or she cannot simply assume that all of the possible adverse outcomes occur in the worst possible way because this outcome has virtually zero probability of ever happening. What he or she attempts to do instead is visualize some reasonable combination of adverse effects that "might" happen, and quantify their limits.

This kind of mental exercise is very difficult for the human mind, even for a relatively simple and probably well understood task like designing an actuator. It is extremely easy to omit consideration of a possible source of cost impact, and it is even easier to misestimate the practical limit of the combined effect of them. The mental exercise becomes virtually impossible if a work element is defined in such a way that it includes many interacting components that have complex interfaces with each other, because each of the interfaces carries its own risks.

Another major problem with estimating cost risk for a work element is the possible interaction of cost with task duration. If a task is delayed or slowed for any reason, it may not be possible to divert the task resources to other purposes. If so, the resources, particularly labor, may continue to accrue costs. I submit that judging the net effect of this possibility, sometimes referred to as the standing army effect, plus numerous other possibilities, is too convoluted for the unaided human mind to deal with properly.

Yet another potential problem is the extreme nonlinearity of some parameters that affect cost. As an example, consider the cost in labor hours of cleaning up a pool of nuclear waste (this is a real life example). If the pool were not radioactive, we might estimate that the workers could do the job in 1,500 person-hours. We know it is radioactive, but we are not sure how intensely so it is, especially near the bottom of the pool. We also know that when a worker's radiation recording badge shows he has reached a certain exposure limit in a given work week, he must be sent to do "cold work," if it is available. If it is not, he is sent to the recreation hall to play cards or read until the end of the shift, if necessary, at full pay. This is a very nonlinear situation with considerable risk. How do you assign a reasonable risk limit to it? Obviously, the key is to know just how radioactive the pool is, and just how much cold work will be available. But if both are uncertain assigning a risk limit is very tricky.

In explicit recognition of these inherent difficulties, some analysts, after obtaining risk limit estimates from the best available experts, routinely modify the estimates to increase the risk according to some favorite formula. A reasonable question to ask about this practice is, how is the analyst to know that too much or too little risk has been added? How is the analyst to know that he or she has any better perception of reality than the supposed experts?

Previously I mentioned that there are two conditions that a work element based estimate must meet before it can be called valid: 1) the subjective assignment of risk distributions around the nominal cost estimates must be done appropriately, and 2) the correlation between work elements must be correctly assigned. I believe I have shown that there are serious obstacles to the first condition, leading to my conviction that it is seldom met. Now I explore the issue of estimation of correlations between work elements.

A key requirement for the validity of correlations is that all work elements must be estimated consistently. For example, if one nominal estimate is at the modal value, then all estimates must be at the modal value. Further, all risk distributions should at least approximate the maximum entropy distribution. Assuming that these are reasonably well done, the analyst must then develop a correlation matrix that describes the degree with which each work element is correlated with each other work element.

Many modern projects have over 1,000 work elements. A 1,000 by 1,000 correlation matrix requires estimating 499,000 correlation values, at least in principle. Of course, in practice this is never done. What usually happens is that the analyst will attempt to simplify the problem in various ways by making assumptions or by consolidating small work elements into larger ones. Unfortunately, such assumptions or consolidations generally will make it much more difficult to assign valid risk distributions, which can destroy the validity of the correlations.

And then, there is the problem of estimating the correlation values themselves. Consider what the imagination is called upon to do. With respect to any pair of work elements, the analyst must consider what will happen to the cost of the other if the cost of the first goes up or down. If there is a direct cause and effect relationship, the problem is simple—the correlation is either 1 or -1 depending on whether both costs move in the same direction or in opposite directions. If there is no interaction at all, then the correlation is zero. Where it gets tricky is assigning intermediate values. One then has to imagine that if one cost goes up, the other may be "sort of" likely to go up (or down), and then guess just how "sort of" is "sort of," anyway? This is conceptually very difficult, and I submit that the chance of error is large.

A set of number pairs exhibits a correlation between the two sides of the pair that can be calculated using accepted mathematical techniques. For example, this set of number pairs (by inspection, without doing the math) has a correlation of 1.

1	2	3	4	5	6
5	10	15	20	25	30

But what do you think the mathematical correlation of the following number pairs is, just be looking at it? Make your best guess, then look for the correct answer on the last page of this document.

1	2	3	5	7	11	13
1	3	7	18	33	75	101

How close was your guess in terms of percent error? For most people, it's fairly large. Correlation is a tricky concept we are not used to estimating. For people with little or no background in statistics, the concept can be virtually meaningless

What if the above problem had been posed as shown in the table below, with the numbers in random order? The correlation is the same, but do you think you would have been able to estimate it as well? What if, instead of six pairs of numbers, there were 20?

3	5	13	1	7	2	11
7	18	101	1	33	3	75

So, if I am correct as to the difficulties of making a valid estimate of project cost risk by assigning risk distributions and correlations to project work elements, am I advocating that there is no point in doing quantitative project cost risk analyses? No, I'm merely setting the stage for suggesting what I believe to be a better approach than the traditional one described at the beginning of this paper.

The basic problem with assigning risks directly to work elements is that work elements are not risk drivers. That is, they are not the causative agents of uncertainty. The causative agents of uncertainty are things like vague knowledge; variability in quality, manufacturing methods, dimensions, people, customers, laws, customs, mistakes, the environment and so on. A given project work element may be subject to perhaps a dozen or more risk drivers, though in practice the number of material ones is usually well under ten. As I have said, the imagination can find it difficult to assess their combined effect when attempting to set risk limits.

Let me pause for a moment and give a careful definition of a risk driver. This is important for future discussion.

A risk driver is any root cause that MAY force a project to have outcomes different than the plan.

Several things are worth noting about this definition. First is the emphasis on the word "may." "May" indicates uncertainty. If something is currently causing or already has caused a deviation from the plan, it is not a risk driver. It is a problem that must be dealt with. If there is no chance of a certain situation arising, it obviously is not a risk driver. To prevent cluttering the risk management agenda with a lot of things not worth considering, project teams generally should limit consideration of risk drivers to those they believe have a probability between 10% and 90%. If the probability is less than 10%, a situation can usually be ignored. If it is above 90%, it should be treated as a real problem and dealt with. Moreover, teams should not waste time with drivers that could only result in small changes to the plan. It's up to the team to decide what "small" means.

Note the phrase "root cause." Why is it in the definition? For an excellent reason! One of the biggest obstacles to effective quantification of project risks assigned directly to work elements is the problem of correlation, as previously discussed. It is all but impossible to design project work elements that are uncorrelated. However, that is not nearly so true of risk drivers. Consider this candidate risk driver that has been proposed by a member of a project team:

"The delivery of the 50 hp compressor from Acme Corp. may be late."

On its face, it sounds like a reasonable risk driver. Maybe it is, but is it a root cause? To answer that question, we must ask WHY the compressor might be late. If we have no answer for the "why" question, then it can be declared a root cause. It simply could be late because that kind of thing happens sometimes and the reason(s) why are not discernable, at least to us at this time. As a further test of whether it is a root cause, we might ask what we can do to mitigate the risk. For example, a team member might suggest sending someone to the Acme plant to cut any red tape that might arise, make sure Acme gets paid in a timely manner according to the agreement with them, and make sure that proper arrangements are made for shipment of the compressor. Ah! Is it possible that red tape at Acme a deeper root cause? Have we had problems in our accounting department with timely payment? Is there any reason to believe shipment could be a problem?

We should stop asking "why" about a potential root cause when we reach the metaphysical level. When the answer to "why" becomes because 1) God, 2) chance, 3) the fates 4) human nature (take your pick) have ordained it, we have gone one step too far. We must stop at the last "why" that results in a cause that human beings might be able to do something about. Analysis of risk at the micro level is pointless if we can't do anything to mitigate it.

The "why" process should also stop short of arriving at a state of blame or finger pointing, in which the policies of a person or institution are attacked. Blame can be counterproductive and can spawn enemies for the project. Identification of the policy that causes the problem is generally sufficient. It may or may not be worthwhile to work to change it.

The beautiful thing about true root causes is that they are statistically independent, by definition. If you are dealing with root causes, correlation is not an issue. Another beautiful thing is that you are not necessarily limited to treating symptoms. You may be able to cure the disease.[4] Yet another is that you often find that a number of proximate causes have the same root cause, reducing the number of potential risk drivers you have to deal with. By dealing with one root cause, you may fix several potential problems all at once.

I have heard this argument: "But isn't it a lot of work to dig for root causes?" Yes, it is a challenging but fruitful mental exercise. But compare it to assigning correlations, which is so challenging it often degenerates into a mere guessing game. Digging for root causes is real work, but it is fruitful work and can develop great understanding in the project team of what they are doing, and how they should go about it. Moreover there is a great and easy to use tool for this purpose called the cause and effect diagram. (See Appendix A.)

The risk driver definition implicitly assumes that there is a serious project plan and the intent is to follow it. A project plan can have many elements, but the key ones from a risk standpoint are usually goals, budget, and schedule, often in that order of importance, but not always. The goals are what we want the project to accomplish, sometimes called positive goals. The budget and the schedule are constraints on the project, sometimes called negative goals. They represent what we are willing to give up in real resources and in time to achieve the goals. Risk drivers can potentially cause us to fail by not achieving the goals, or by violating either the budget or the schedule constraints, or both. In a project, as in a machine, there can be many modes of failure, but only one or at most a very few modes of success.

Some frequently asked questions come to mind:

- Is everything that changes the project plan a risk driver come to fruition? Not necessarily. Sometimes in the course of a project, desirable changes to the plan become apparent and are executed for a variety of reasons. Whenever we voluntarily change the plan, we must re-examine the risk drivers. They could have more or less effect than before the change. Risk mitigation is actually a voluntary change of plan designed to reduce the effect of risk drivers. It makes the plan more robust.
- Is the action of a risk driver always adverse to the plan? Again, not necessarily. Unfortunately, the word "risk" has become associated with bad outcomes, so the first thing many of us think about when the word risk is mentioned is that something bad happening. This view is clearly skewed. Nothing great and good is achieved in this

[4] Of course, you are still free to treat symptoms if a cure is not possible. At least you know where you stand.

world without taking risks, so risk obviously has two faces. One and the same risk driver MAY have different outcomes, some good and some bad.
- Can a risk driver affect more than one work element? Yes, and this is a key idea in the methodology I will disclose in this paper. Instead of trying to mentally visualize risk limits on the cost of a work element, we will determine them rather directly by first identifying all key risk drivers that we can think of, then analyzing their effects on each work element individually. The summation of these effects will help us define reasonable risk limits. Note that this approach offers a new possibility: We can directly estimate the overall project impact of each individual risk driver, thus greatly aiding decisions about potential risk mitigation, and the appropriate level of resources to devote to that. That is much harder to do when using the method of assigning risks directly to work elements.

My proposed approach I call the risk driver impact approach, hence the title of this paper. It attempts two things: 1) Essentially getting rid of the sticky correlation problem by working with root cause risk drivers, and 2) making the assignment of risk distributions much simpler and more intuitive. The distributions are not assigned directly to work elements, but to the risk drivers, and regulate their impact on the work elements. However, the ultimate result is a cost risk distribution for each affected work element, based on the cumulative effects of all risk drivers that impact it. We also can obtain a cost risk distribution for the entire project or for any project phase or cost type, such as labor and material.[5] As mentioned earlier, we can also directly measure the total project impact of any one risk driver.

The methodology I propose has a common architecture for treatment of each risk driver. These are its main features, which will be addressed in the order shown:

- Statement of risk
- Statement of outcomes and probabilities
- Statements of work element cost impacts
- Monte Carlo treatment

<u>Statement of risk.</u> The statement of risk is a well thought out, simple declarative sentence about a root cause that could force the project away from its plan. The statement should also say, in broad terms, not in detail, what the adverse effects might be (there may be

---

[5] So far in this paper I have treated cost risk as a subject mostly if not entirely separable from schedule risk and what some call technical or goal risk. Of course, it is not. The main purpose of the paper is to deal with cost risk, however Appendix B advances some ideas for addressing the often complex interactions between cost, schedule, and technical risk.

possible beneficial effects, but for reasons of economy of description they are generally not included in the statement).

It can be helpful to supplement the statement of risk with broader explanations of possible impacts in a document one might call a “risk impact dictionary” or a “risk register.”

The statement of risk should not anticipate any proposed but not yet agreed to mitigation measures, but should anticipate the estimated future effects of mitigation measures already in work or in place. The statement should contain the word “may” or some other conditional word such as “could” or “might” to make clear that there is significant uncertainty. (It can sometimes be easy to confuse risk drivers with problems.)

An example statement of root cause risk is:

“Too small an effort for near term personnel search and hiring could result in a shortage of engineers in the design phase.”

The statement of risk is intended to clarify for all concerned the true nature of the risk and its possible importance to the project. It also serves as a starting point for thinking about risk mitigation.

Statement of outcomes and probabilities. This is essentially a short list of the possible outcomes due to the risk driver, and the assigned (agreed) probability of each. Although possible outcomes may actually be infinite in number, we typically limit this list to five or fewer outcomes for practical reasons. This often requires “grouping” of certain closely related outcomes that have roughly the same potential impact on the project.

For the sample risk driver presented above, this might be the agreed statement of outcomes and probabilities:

Outcome	Probability %
Plan met: 50-60 engineers hired	50
Engineers hired short by ~10	30
Engineers hired short by ~20	10
Engineers hired short by >30	10

The following comments about this design are important:

- The first outcome listed always must be that the plan is met. Logically, that could be assigned a zero probability, if that is really the case. But if it is the case, clearly the plan needs to be rethought.
- The outcomes listed must be exhaustive by implication if not by specific statements. Exhaustive means that all possibilities are accounted for. The above outcomes are not exhaustive based on the specific statements given, so they must be interpreted as exhaustive by implication. By implication, considering all of the statements taken as a whole, the entry "Engineers hired short by ~20" means that the number hired will be about 35, plus or minus a few. The basic idea is that the effects on the project will not be much different if 33 are hired or 38 are hired.
- The statements must not overlap. Here is an example of two overlapping statements: 1) delivery may be between 1 and 3 weeks late, and 2) delivery may be between 2 and 4 weeks late.
- The probabilities assigned must sum to 100%, or else they must be normalized to sum to 100%. Sometimes people find it easier to assign probabilities in terms of relative likelihoods of outcomes, and these do not necessarily sum to 100. Example: Outcome A is 3 times as likely as outcome B, and 5 times as likely as outcome C. The corresponding probability assignments for A, B, and C would be 65.2%, 21.7%, and 13%. These both sum to 100% and meet the relative likelihood assignments.

<u>Statements of work element cost impacts.</u> In principle, any number of work elements can be impacted by a particular risk driver. In general, the impacts will have different effects on different work elements. To keep things manageable, the method for describing each impact is based on a state of knowledge where the upper and lower bounds of the impact are assumed known. Tribus (*op cit*) has shown that the maximum entropy distribution for this state of knowledge is the uniform distribution. We therefore use that distribution consistently in this part of the analysis.

The following table shows statements of work element impacts for three hypothetical work elements (WE3, WE9, and WE14) appended to the table of statements of outcomes and probabilities shown above.

		WE3		**WE9**		**WE14**	
Outcome	**Probability %**	**Min K$**	**Max K$**	**Min K$**	**Max K$**	**Min K$**	**Max K4**
Plan met: 50-60 engineers hired	50	0	0	0	0	0	0
Engineers hired short by ~10	30	20	30	20	45	55	65
Engineers hired short by ~20	10	30	45	45	65	65	90

Engineers hired short by >30	10	45	60	65	80	90	120

Note that in the first row meeting the plan is assumed to have no cost impact on any WBS element. This should be generally true, but there can arise situations where the plan is deemed under- or over-funded and it is reasonable to show an impact. In this table, the cost impacts are all shown as positive, meaning a potential increase. It is in principle possible to have savings (negative costs) in these tables. Some risk drivers can produce either negative or positive impacts depending on the outcome.

This table has been constructed to show costs in K$. Costs can also be shown in labor hours or in other ways, but eventually all costs and risks should be converted to equivalent currency units. Note that one can, if desired, add more columns to the table to indicate costs and risks to components of a work element, such as material cost and labor hours, if such distinctions are deemed significant.

This method makes no assumption about whether the nominal cost estimate is a mode, median, mode or anything else. It merely states that if a particular outcome happens, the impact to the cost will be within a certain stated range. Just where in that range is uncertain.

Monte Carlo treatment. The Monte Carlo process can be used to establish a distribution of incremental costs for each work element due to each risk driver. This can be presented as a set of moments of the distribution about its mean, or as a pdf or a cdf, or all of these. The process can also concurrently accumulate statistics and distributions for the entire project, or for phases of the project.

Techniques of Monte Carlo programming are fairly well known and will not be described in detail here except to say the following.

- The first step in each Monte Carlo iteration is to select the outcome for each risk driver. This is done by drawing a single (separate and different) 0-1 uniform random number for each risk driver.
- Once the outcome is randomly selected, another (separate and different) 0-1 uniform random number is drawn for each risk driver to determine the point in the impact range. This is done using the equation Impact = Min + Rand(Max – Min).
- The results are collected for statistical presentation as may be desired.

Appendix A – Finding Root Causes

Some seeming “risk drivers” are in reality just symptoms of the true root cause. How can you distinguish a symptom from the true root cause? It isn’t always easy, but most of the time it is worth the effort, because it gives great insight into the nature of the risk drivers and the consequences that could flow from them. This, in turn promotes the possibility of more comprehensive risk mitigation.

In principle, the method of finding a root cause is simple. The basic idea is that you first describe the most obvious manifestation of the potential problem, such as a possible late delivery of some newly coded software. This manifestation is generally the one that first occurs to the project team, because their thoughts are typically focused on events and their proximate causes, as opposed to root causes.

Once you have stated a symptom or proximate cause of a potential problem, you ask the question, “How could that happen?” You may come up with just one answer, or there may be several answers. For each of these answers, you can ask again, “How could that happen?” And again, you may come up with even deeper answers. At some point you decide that you don’t really know a deeper cause, or that the only explanation lies in some mystical expression such as “It’s a mystery—God must have willed it”. Such expressions are not useful as root causes because they put mitigation out of reach of mere mortals. (Sometimes it’s out of reach even when the expression is not at all mystical!)

It is not hard to see that this process of reaching for deeper causes could lead to the discovery of entirely new risk drivers. These, in turn, can be analyzed for potential impacts. This is yet another good reason for a careful evaluation of root causes. Cause and effect diagramming is a method for illustrating the relationships between a perceived risk and the many factors (risk drivers) that can cause it. The method can be used by an individual or by a group. When used by three or more persons, one should act as facilitator and recorder. The method uses brainstorming to help identify factors populating a cause and effect diagram, as demonstrated in the diagram below.
Here are the typical steps in the process.

- The facilitator presents a risk driver symptom for which root causes are to be found (in the diagram, the perceived symptom is a possible schedule slip in the delivery of the “ABC” system).
- The facilitator explains the cause and effect process

- The group brainstorms and constructs a “fishbone” diagram on board or on paper
- The group identifies cause factors and adds them to fishbone structure
- Then group identifies the most significant causes and circles them
- The group collects additional information if necessary to refine the diagram
- From the process should emerge a clearer vision of the nature of the risk drivers and how to describe them

Note in the diagram that four important risk drivers have been identified as contributing to the symptom shown (assuming that “cost containment” is actually two separate risk drivers). An interesting possibility is that each of these drivers may have other outcomes that should be considered. Another interesting possibility is that other “minor” drivers not considered important at this time, such as poor estimating, or supplier labor difficulties, have been identified and can be put on a watch list.

The cause and effect diagram is a powerful tool for analyzing risk symptoms to find root causes. Interestingly, it has at least one other benefit in project risk management: Once the chain of causality is established, it is easier to visualize the most cost effective approach to risk mitigation.

Typical Cause and Effect Diagram

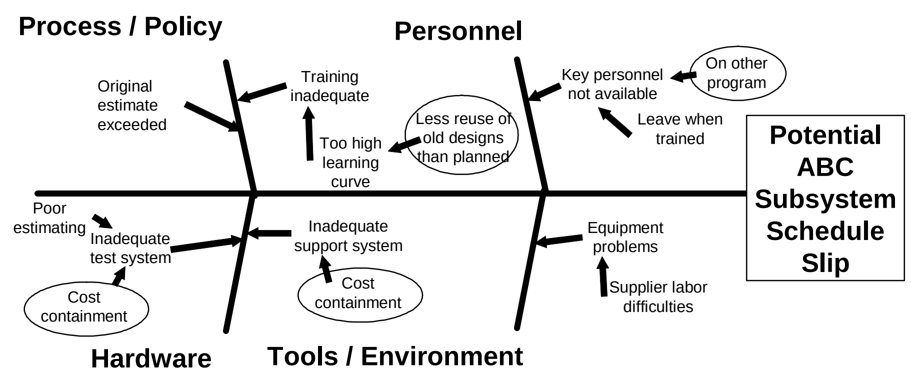

Appendix B – Risk interactions

People in projects often speak of cost risk, schedule risk, and goal (or technical) risk as if they were completely separate things, and even try to analyze them as such. Mostly this is done because a comprehensive analysis of the interactions of these risk "categories" can be extremely complex. The model can become so cumbersome that one would despair of using it for practical analysis.

To get some insight into the problem, and perhaps get clues as to how to build relatively simple yet reasonable models, let's look separately at the goal and schedule risk categories and see what makes them tick.

Goal risk. Many call this technical risk, but I prefer the term goal risk because project goals are not always about technology. However, in discussing this risk area it is nevertheless useful to focus on a technology example because that is generally the most difficult to deal with. Many other goal risks can be treated in an analogous manner.

Technical risk is generally about the possibility that product performance will be less than desired. So I use the term performance risk interchangeably with technical risk. Our approach will be to assume that, in the absence of identified risk drivers, the project's product, whatever that may be, will comply with all of the goals set for it by the project's sponsor. If there is any significant chance of that not happening, it should be attributable to risk drivers that are mostly identifiable.

When you boil it all down, performance impacts come in only three categories. One is the situation where the product fails to meet its goals due to the action of one or more risk drivers, and that is the end of the matter. Nothing can be done to repair the situation, or else we are unwilling to try, for one reason or another. The project soon terminates. There may be further termination activities and costs, but they are not directed toward product performance. I will call this Type 1 performance risk.

The second is the situation where the product also fails to meets its goals due to one or more risk drivers, and we have decided we want to (or must) spend more time and money in the hope of making the product goal compliant. In the Type 1 category, we are defeated and accept defeat. In this second category, we are damaged, but refuse to accept defeat. But in refusing to accept defeat, we may spend more money and more time than our plan calls for. We expect to finally prevail and have a successful product. I will call this Type 2 performance risk.

The third is the situation where, again, the product fails, and again we spend time and money trying for success, but after a time we fail and concede defeat. I will call this Type 3 performance risk. Let's look at examples of each of these situations.

Consider one type of project commonly worked by the National Aeronautics and Space Administration (NASA). The product is a body of scientific data, typically gathered by instruments carried aboard a spacecraft. NASA characteristically writes success criteria for these projects. These criteria are the project's positive goals. NASA also creates a schedule and a budget, which are the negative goals (constraints). A failure of the launch vehicle (the rocket that boosts the instrument-laden spacecraft into earth orbit) will result in loss of the spacecraft and all of its instruments. This is a total mission failure. Once it happens, nothing can be done to repair the situation. It is the end of the matter. This is a Type 1 performance failure.

Suppose that the mission is to go to Mars, land, take pictures, and measure soil composition and atmospheric properties. The launch vehicle succeeds, the interplanetary flight succeeds, the Mars landing succeeds, but one of the cameras fails, reducing the amount of data gathered somewhat below the level contemplated by the project goals. The loss is less than a total mission failure, but again, nothing can be done about it. We can't (yet) send a repair crew to Mars. This is a Type 1 performance failure, but is not a total failure.

To generate an example of the second type, suppose that during the development phase for the vehicle that will go to Mars, one of the key scientific instruments failed a test designed to demonstrate its ability to withstand the severe vibration environment encountered during launch. Taking that instrument out of the launch manifest would have meant not being able to meet all of the project goals. Leaving it in would require some redesign and retest of the instrument. Assuming that we must hold to the positive goals, and the money to do the added work is not in the current plan, it would mean a cost overrun and possibly a schedule overrun as well.[6] The project sponsor would be obliged to augment the funds to accommodate the overrun, unless the project already had contingency funds that would cover the loss.

[6] Schedule overruns can be a serious matter for space vehicle projects, because they might cause missing a launch window.

With regard to the third type, suppose that we have set out to build a fuel cell powered car that will sell for under $20k. We overrun the initial project budget and schedule, but decide to put additional funds into the project and try again. At length, we fail and quit trying.

In the second and third types of performance risk, the technical failure is not the ultimate impact. The ultimate impact is a cost and possibly a schedule impact. We can boil the previous few paragraphs down into the following three definitions:

- **Type 1** performance impact is a failure of the product to meet project goals that cannot or will not be reversed. There is no secondary cost or schedule impact aimed at fixing the failure (there could be secondary impacts such as project cancellation charges).
- **Type 2** performance impact is a failure of the product to meet project goals, it being both possible and likely that such failure will be reversed, resulting in a secondary cost or schedule impact, or both. The impacts are due to work needed to repair the failure.
- **Type 3** performance impact is a failure of the project to meet project goals, and we assume the failure is reversible and continue trying. But ultimately we concede defeat and end the project.

These definitions clarify the nature of what some call “technical risk,” an unfortunately ambiguous and often-misused term. As I have shown, there are three kinds of “technical” risk, and two of them are just special cases of cost or schedule risk.

I have shown that Type 1 performance risks can result in either partial or total compromise of project goals. It would seem natural to think of each discrete positive performance goal as having a “percentage importance,” so that all of them taken together add up to 100% of what is wanted. Assigning quantitative relative importance measures to goals is a very useful practice in all projects, and especially in projects that have Type 1 performance risks. It focuses attention and effort on what is important, and reduces the likelihood that the project team will spend inappropriately large amounts of time and money on things that are not very important. It also promotes the likelihood of customer satisfaction.

Once they have a list of performance goals that may be subject to Type 1, 2, or 3 performance risks, the project team will in due course develop a baseline design for meeting them. Given a specific design, the team can do analyses called failure mode analyses. These are searches for ways the product might fail. A failure mode analysis is essentially a form of search for root causes that can cause product failure or degraded performance. The more sophisticated forms of this type of analysis for advanced hardware and software systems are beyond the scope of this paper. Engineering specialists should always perform them. But we can illustrate the general approach here.

We have previously described cost impacts in terms of money, and schedule impacts in terms of a span of time. These dimensions are readily quantifiable. But how should we describe Type 1 performance impacts? Keep in mind that these impacts are final outcomes for the project, and that they are presumed to be irreversible. Not all goals are easily quantifiable, so the approach I recommend here is to quantify goals in terms of their perceived importance.

For example, suppose that a Mars project has just these three goals:

- Goal 1: Take and transmit to Earth a minimum of 500 high resolution digital pictures of the Mars surface around Landing Site X
- Goal 2: Determine the gaseous content of the Mars atmosphere at Landing Site X in terms of the percentage of each component, and transmit the data to Earth
- Goal 3: Perform a complete chemical analysis of the soil at Landing Site X, and transmit the data to Earth.

The physical reality is that none of these goals can be accomplished unless a spacecraft is successfully directed to and lands on Mars. The exhibit below shows the necessary sequence of events. It also shows the assigned mission values in percentages for the three goals (summing to 100%), and shows nine hypothetical risk drivers (RD) and the performance areas they can impact.

Let’s now build a table, also shown below, of the risk drivers and their probabilities of occurrence. (We assume that reliability engineers have assigned these probabilities based on a pool of historical data.)

Example of Treatment of Type 1 Performance Risks

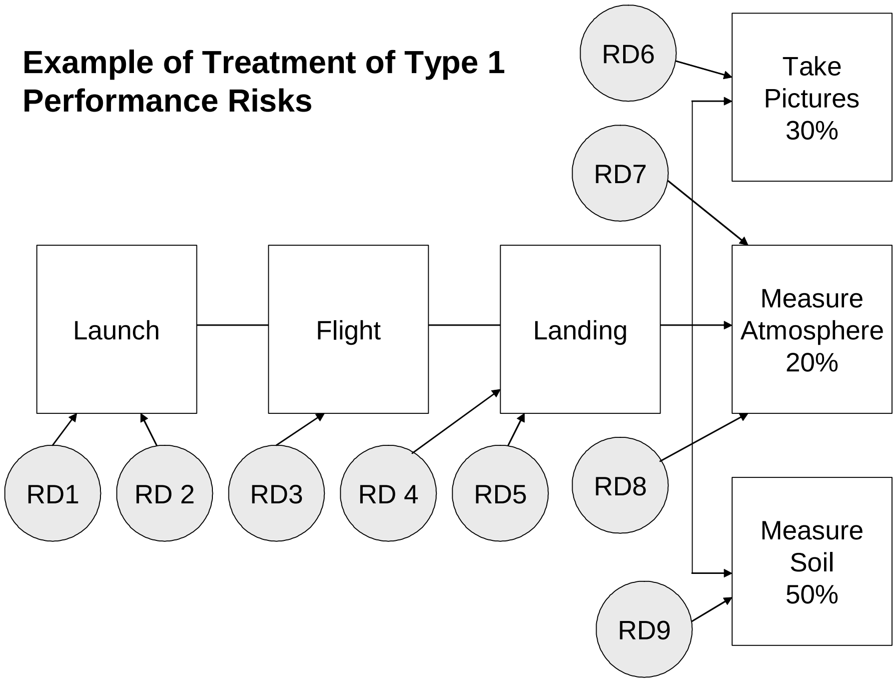

Risk Driver Occurrence Probabilities

1	0.001
2	0.032
3	0.011
4	0.0025
5	0.008
6	0.003
7	0.002
8	0.06
9	0.0021

We first need to find the probability of reaching Mars successfully. That is the same thing as the probability that none of the first five risk drivers happens. Mathematically, it can be calculated as:[7]

(1-0.001)(1-0.032)(1-0.011)(1-0.0025)(1-0.008) = 0.9464

(NASA would probably never accept a mission probability this low, but remember, this is just an example. Also, note that the NASA success rate with Mars missions is actually lower!)

The probability of total success in taking pictures is given by:

(0.9464)(1-0.003) = 0.9436

Note that the second factor is the probability that risk driver #6 will not occur. The probability of total success in measuring the atmosphere is given by:

(0.9464)(1-0.002)(1-0.06) = 0.9162

Note that the second and third factors are, respectively, the probabilities that risk drivers #7 and #8 do not occur. The probability of success in measuring the soil is given by:

(0.9464)(1-0.0021) = 0.9444

Finally, we calculate the expected project value with respect to the mission goals. This is a weighted sum, wherein the goal values are multiplied by their respective probabilities of achievement and then are summed:

Project figure of success = (30)(0.9436) + (20)(0.9161) +(50)(0.9444) = 93.85%

To recapitulate, some projects have risk drivers that can create Type 1 performance risks. These are essentially irreversible failures. I have suggested that a good way to describe these is to allocate values as percentages to all of the positive goals that can be affected by these failures, such that these percentages add up to 100%. Each goal value is multiplied by its probability of

[7] This calculation uses two rules of probability. The first is that if the probability of something happening is P, a number between zero and one, then the probability of it not happening is 1-P. The second is that if the probability of success of the first of a series of independent events is X, the probability of the second is Y, the probability of the third is Z, etc., then the probability of them all succeeding is the product XYZ.

occurrence, and these are summed. The result is a project score or grade, with the highest possible grade being 100% if nothing can fail.

It may be possible to improve the mission score by mitigating risk drivers. Mitigation may cost money, and may also add time to the critical path. So a judgmental tradeoff is possible between cost, schedule and project value. The tools described in the above discussion make these tradeoffs possible on a rational basis.

Please don't think that the process for Type 1 performance risk drivers described here applies only to NASA projects that launch spacecraft. It can apply to any mission-oriented project that can have irreversible failures of the product that essentially kill or damage the project. For example, a project to develop a new financial instrument could have an irreversible failure if investors did not accept it, resulting in a shutdown of the project. The probability of success might be estimated using focus groups.

Type 2 and 3 performance risk drivers have further consequences beyond the technical failure itself that can be measured as cost or schedule impacts of both. Note that the failure itself is not the risk driver, the cause of the failure is. There is a great temptation when describing this type of risk driver to write a simple statement such as "The gizmo might fail its acceptance tests." I have done this myself. The conscientious risk analyst will try to write such a statement for every known failure mode of the gizmo. Problem is this might result in 50 risk driver statements, which is more trouble than it is worth. My advice: do the best you can to maintain fidelity in risk description, but don't try to make risk management the only activity in the project. If you do, you will soon run out of friends.

Type 2 and 3 performance risk drivers can easily be included in cost and schedule risk impact tables. We have talked already about cost risk tables. We turn now to schedule risk.

<u>Schedule risk.</u> Meeting a prescribed schedule is frequently a cherished goal of project sponsors. There can be many reasons for this. Among the possibilities are getting to market ahead of the competition, meeting schedules on other, related projects, meeting scheduled occupancy or availability dates, or getting the project completed within a funding cycle, while money is available.

Cost risk impacts are relatively simple to deal with, as illustrated in a previous section. So-called technical risks can be dealt with as described in the previous section. Schedule impacts are a bit more complex. To illustrate the nature of this complexity for even a simple project, consider the diagram below. It shows the schedule of a hypothetical project, we'll call it Project X, in network format.

Schedule for Project X

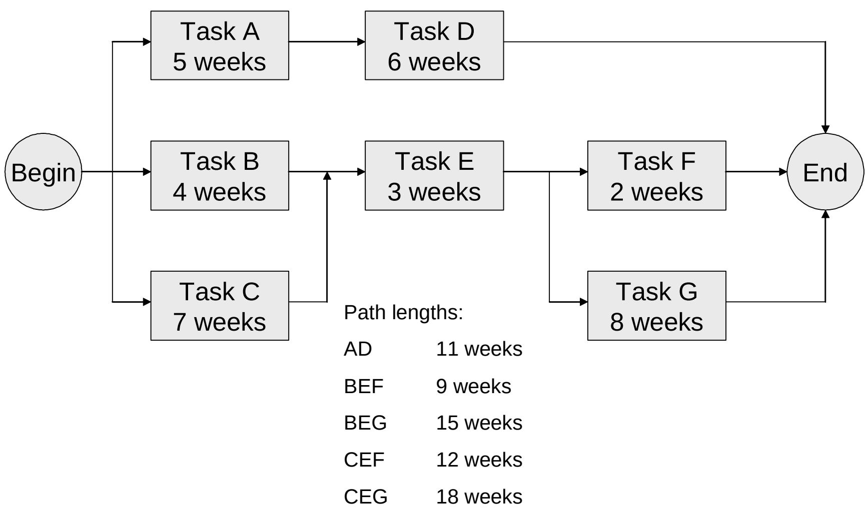

There are five paths through this network, AD, BEF, BEG, CEF, and CEG, with path lengths as noted above. The longest path, CEG, is called the critical path; its length is the duration of the project.

Now consider a risk driver impact of five weeks on Task A, doubling it to ten weeks. The path AD will now be 16 weeks long. The duration increase on task A may also increase the cost of task A, but it has no effect on the duration of the overall project, which is still 18 weeks. Clearly, risk driver impacts off the critical path do not affect the project overall duration *if they are small enough.*

Now consider a risk driver impact of five weeks on task C. This lengthens the critical path by five weeks, causing the overall project duration to increase from 18 to 23 weeks. Clearly, risk driver impacts on the critical path always affect the project duration directly.

Finally, consider a risk driver impact of four weeks on Task B. This will change the critical path from CEG to BEG, and will increase its length by one week to 19 weeks.

This gives you a notion of the complexities inherent in schedule risk analysis. Some sophisticated methods are available for dealing with this complexity. Most involve Monte Carlo simulation.

It should be noted that some large projects build huge "working level" schedule networks that may have thousands of detailed tasks. Schedule risk analysis at this level of detail is tedious, partly because of the necessity of assigning risks to so many tasks.

What can be done to make the work easier is to break schedule risk analysis into two parts. In the first part, one considers schedule interactions that can directly affect work element cost. As an example, in many tasks if the schedule increases, some of the costs may increase also, labor in particular. A way to handle this is illustrated in the following table.

		WE3 Cost		WE3 Sch	
Outcome	**Probability %**	**Min K$**	**Max K$**	**Min Wks**	**Max Wks**
Plan met: 50-60 engineers hired	50	0	0	0	0
Engineers hired short by ~10	30	20	30	2	4
Engineers hired short by ~20	10	30	45	4	8
Engineers hired short by >30	10	45	60	8	14

This table is adapted from a table shown previously in connection with cost risk analysis. Note that for work element 3 (WE3) range of impact columns for schedule (duration, actually) have been added. The basic idea is to pair up cost and schedule impacts. The cost impacts can be considered at least roughly linear functions of the schedule impacts. As such, they are highly correlated. This is easily handled in the Monte Carlo by using the same 0-1 random number at each iteration to estimate both the cost and the duration impact.

The second part is to put these random impacts into the context of a scheduling tool so that schedule slack and other issues are properly considered. Tools such as Microsoft Project can be used for this purpose.

[Correlation = 0.983. Did you guess a lot lower? Most do.]