Integrated Cost / Schedule Risk Analysis

David T.Hulett, Ph.D.
Hulett & Associates, LLC
(310) 476-7699
david.Hulett@verizon.net
www.projectrisk.com

Bill Campbell, III
Risk Mitigation Associates
(714) 694-0174
bill@riskmitigationassoc.com
www.riskmitigationassoc.com

Abstract. Project costs often exceed their estimates because those estimates do not take into consideration the actual duration of project activities. Cost risk will also be underestimated if it does not take into consideration schedule risk. This paper presents a method of incorporating the uncertainty in activities' durations into the assessment of cost risk. In this method, a Monte Carlo simulation of the schedule provides uncertainty in time. Incorporating the schedule risk results into the cost risk model provides the linkage between schedule and cost risk. Then equivalence must be established between the schedule and network concepts. Uncertainty in costs is then represented by uncertainty in "independent costs" (costs that do not depend on time) and "variable costs" (costs that depend on uncertain time and cost per unit time or "burn rate" and rate of labor compensation.) Simulation of the cost model combines the results from the schedule risk analysis with the uncertainty in the cost assumptions. The results include the probability distribution of total project costs and sensitivity of that distribution to the different inputs. Issues are discussed and simplified examples are provided.

INTRODUCTION

Cost and schedule are related. The estimate at completion (EAC) for a project makes many assumptions. Key among those are the resources that will be working on the project activities, their rate of compensation and their durations for the activities that react to duration. Each of these assumptions contains potential risk, where risk is defined as an uncertain event or condition that, if it occurs has a positive or negative impact on the cost objective. (PMI 2000) Of course, even for those activities where cost does not respond to duration, there is uncertainty.

Integrating risk in cost and schedule. Cost estimates often become disconnected from the facts of the project that are sometimes different from the assumptions. The risk that project cost will exceed the estimate is assessed by taking into account several types of risk, most notably the risk that project activities' durations may differ from those originally assumed (schedule risk). Other cost risks include the resources to be applied per unit time ("burn rate") and compensation per hour for those resources, and uncertainty in the cost of activities whose costs do not depend on elapsed time.

SCHEDULE RISK ANALYSIS

Name	Durati	Start	Finish
Start	0 d	1/2	1/2
Software Subsystem	**265 d**	**1/2**	**9/23**
Design Software	65 d	1/2	3/7
Code Software	150 d	3/8	8/4
Test Software	50 d	8/5	9/23
Hardware Platform	**260 d**	**1/2**	**9/18**
Design Hardware	60 d	1/2	3/2
Build Hardware	165 d	3/3	8/14
QC Hardware	35 d	8/15	9/18
Integration	**120 d**	**9/24**	**1/21**
Integrate Hardware	80 d	9/24	12/12
Test Integrated Syst	40 d	12/13	1/21
Deliver System	0 d	1/21	1/21

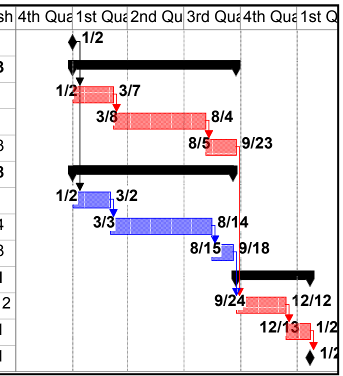

Figure 1. The schedule before risk analysis

Schedule Risk. The strategy for integration of schedule and cost risk begins with an analysis of the risk of the schedule. The dates computed in Figure 1 assume all activity durations are known with certainty. The scheduling software indicates a finish date for system delivery of January 21. Unfortunately we cannot estimate durations with accuracy. Often the uncertainty in the durations is represented by a 3-point estimate with optimistic, most likely and pessimistic scenarios defining the three points. These scenarios are typically discovered by in-depth interviews of those who understand the project and its risks. Figure 2 shows a possible set of 3-point estimates that are assumed to represent triangular distributions.

ID	Name	Duration	@RISK: Functions
1	Start	0 d	
2	**Software Subsystem**	**265 d**	**Finish=RiskOUTPUT()**
3	Design Software	65 d	Duration=RiskTRIANG(50,65,80)
4	Code Software	150 d	Duration=RiskTRIANG(140,150,16
5	Test Software	50 d	Duration=RiskTRIANG(45,55,70)
6	**Hardware Platform**	**260 d**	**Finish=RiskOUTPUT()**
7	Design Hardware	60 d	Duration=RiskTRIANG(50,60,80)
8	Build Hardware	165 d	Duration=RiskTRIANG(135,165,19
9	QC Hardware	35 d	Duration=RiskTRIANG(30,40,60)
10	**Integration**	**120 d**	**Finish=RiskOUTPUT()**
11	Integrate Hardware / Soft	80 d	Duration=RiskTRIANG(65,80,100)
12	Test Integrated System	40 d	Duration=RiskTRIANG(30,40,65)
13	Deliver System	0 d	

Figure 2. Add risk to each activity

Schedule risk results. The risk results provide possible dates and their probabilities for system delivery. This distribution of completion dates is found by simulating the schedule many times. In each iteration a duration is picked at random from the distributions shown in Figure 2 for each activity. Because we do not know which combination of durations will actually occur, we tell the computer to conduct, say, 3,000 iterations, each with durations selected at random according to the probability distributions for each task's duration. A probability distribution of possible delivery dates results and is shown in Figure 3.

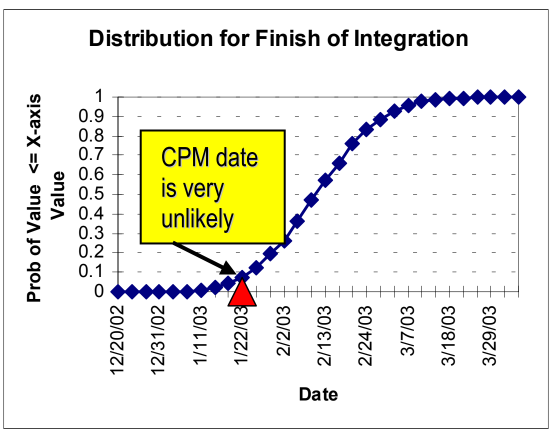

Figure 3. Distribution of delivery dates

These results indicate that the date of 2/21 is very unlikely. There are other results from the schedule risk analysis. The software shows which activities and paths are most likely to delay the project to assist risk response planning. Figure 4 shows that Integration is always on the critical path and that software is the next most likely to delay the project. This information is quite helpful to the project manager.

Activity	Risk Critical Index
Design Software	58%
Code Software	58%
Test Software	58%
Design Hardware	42%
Build Hardware	42%
QC Hardware	42%
Integrate Hardware / Software	100%
Test Integrated System	100%
Deliver System	100%

Figure 4. Risk criticality of activities/paths

INTEGRATION OF SCHEDULE RISK INTO THE COST ESTIMATE

Matching cost and schedule concepts. The cost estimates for each activity should be driven by a realistic estimate of the duration of the tasks with which they are connected. Similarly, the *risk* in the cost estimate is driven by the *risk* in the estimate of duration for the same activities. To drive the cost risk for each activity by the correct schedule risk, we need to match the cost estimates to their corresponding schedule items.

Finding schedule activities that are the same as the entries in the cost estimate is not as easy as it should be. Often the cost estimate is related to the work breakdown structure (WBS) but the schedule is not. We found that it was easier to take the dollars in the cost estimate and apportion them into the schedule summary tasks than *vice versa*. The uncertainty in each summary activity can then drive the appropriate number of dollars. This mapping is shown in Figure 5.

Map to Schedule Summary Activity:	Cost Element	($ 000)	Schedule Element		($ 000)
A	S/W	4,594	A	S/W	5,359
B	H/W	3,804	B	H/W	4,569
C	Integ.	2,077	C	Integ.	2,842
A/B/C	Mat.	1,976			
A/B/C	P.M & Support	320			
Total Cost per EAC		12,770	Total Cost per Sch.		12,770

Figure 5. Map Costs to Schedule Activities

Schedule element input to the cost model. Total cost is the sum of variable and independent costs. Each of these has uncertainty. Variable costs are the product of three items and each is uncertain.

Input Item	Uncertain	Source of Data
Duration	Yes	Schedule Risk Anaylsis Histogram
"Burn Rate"	Yes	Interviews
Compensation rate	Yes	Interviews

Figure 6. Components of Variable Cost

The first problem is to find a common denominator for these three items. Hours can be used, but integration with the schedule is easier if they are translated to days.

The second problem is to determine the uncertain duration of the schedule summary elements. Each of these inputs is the probability distribution of the appropriate summary activity, such as those listed in Figure 5, above calculated from the simulation of the schedule. This distribution or "histogram" shows the different duration that the summary activity could take and their relative likelihood.

Calculation of the number of days the Software Subsystem and Hardware Platform take can be found directly from the simulation of the schedule because these two summary activities start on a fixed day, the first day of the project. (We have used 7-day weeks in the schedule because the days will be translated to Excel that does not distinguish weekends from weekdays in computing the date mathematics. Thus the schedule date will be wrong, unless it is a plant turnaround with 7-day schedules but the duration in days should be accurate.)

The histogram from the schedule simulation for the Integration summary task will not represent faithfully the number of days of integration work because its starting date is not fixed but depends on the completion of the longer of its two uncertain predecessor paths, software or hardware. We simulated the Integration tasks separately, starting from an artificial fixed date, to develop the histogram of uncertain integration days that abstracts from the uncertain start date for those tasks.

Integration in the Original Schedule				
Integration Summary	120 d	9/24	1/21	Finish=RiskOUTPUT()
Integrate H/W & S/W	80 d	9/24	12/12	Duration=RiskTRIANG(65,80,100)
Test System	40 d	12/13	1/21	Duration=RiskTRIANG(30,40,65)
Deliver	0 d	1/21	1/21	
Integration Simulated Separately to Derive a Histogram in Elapsed Days				
Integration Summary	120 d	1/2	5/1	Finish=RiskOUTPUT()
Integrate H/W & S/W	80 d	1/2	3/22	Duration=RiskTRIANG(65,80,110)
Test System	40 d	3/23	5/1	Duration=RiskTRIANG(30,40,65)
Deliver	0 d	5/1	5/1	

Figure 7. Simulate a downstream path

The histogram is transferred to the cost model from the schedule simulation easily because we used the same maker's software for each of the simulations. Being specific, the schedule was simulated using @RISK for (Microsoft) Project that produced a histogram as an output. We used a 30-point distribution in this analysis for accuracy. An example of that histogram might be the following: Duration=RiskHISTOGRM(242,383,{0.001,0.004,0.009,0.017,0.0255,0.032,0.053,0.0715,0.0655,0.0895,0.088,0.09,0.0955,0.0885,0.056,0.0545,0.0485,0.032,0.027,0.019,0.013,0.0115,0.0015,0.0045,0.0015,0,0.001,0,0,0}).

The cost risk model was programmed in Microsoft (MS) Excel. The simulation program for the cost risk model was @RISK for Excel, which accepts the histograms from its cousin in Project. If the software that simluates the Excel model does not accept the outputs of the schedule risk model, for instance if Risk+ is used to simulate MS Project and Crystal Ball simulates the Excel model, an approximation of the schedule risk outputs must be made and input into the cost risk model, introducing some error that may be large or small. A larger error might occur if there are probabilistic branches in the schedule and the histogram from the schedule risk analysis has two maxima. It would be difficult to estimate a single distribution or to make a custom distribution in Excel for such a shape.

Collecting variable cost risk data. The data collection process starts by defining the components of cost that are needed to simulate the cost risk analysis. These items are all uncertain unless they are in the past. In fact, if the analysis is done in the middle of the project, there should be a history of "actuals" that would be certain and added onto any uncertain estimate of costs yet to go.

The first element to be collected from the project team leaders is the "burn rate" or number of resources (hours, heads) applied to the task per unit time (day). The burn rate (e.g., heads per day) will be multiplied by the duration (e.g., days) to compute the resources (e.g., total hours) for the schedule summary task.

The burn rate is a partly-new concept for the teams because they tend to think in time-phased terms. We could not incorporate any time phasing of staffing in the simulation model, so the concept had to be the equivalent based on an assumption of level staffing. It turned out that this concept was easy enough to estimate once it was explained and the team was given encouragement.

The team leader was expected to meet with the team on all of these data collection tasks and develop 3-point estimates that could be turned into probability distributions for simulation. For burn rate we provided data on the estimate of burn rate computed from the cost estimate. This is found by computing the total hours for the tasks in the base estimate and dividing by the number of days assumed in that estimate. The result is either in hours or headcount. The team fills out the table with their first estimate of optimistic (low), most likely and pessimistic (high) number of hours or heads per day.

Note that the number of resources that emerges from this exercise need not be the number in the base estimate. Sometimes the base estimate is out of date. Sometimes it was "squeezed" to fit some target handed down by management or the customer and was never a true estimate of cost. Given these factors, it is not surprising that the "most likely" estimate of resources per day is often not that implied by the original estimate. Infrequently, but sometimes, the optimistic (low) estimate of resources is above that implied by the base estimates.

The team is then interviewed on their first guess at the three-point estimates. The most common problem with these estimates is that they are too narrow. It is well-established that people who have no or little experience in providing data on risk will initially underestimate the true uncertainty in the numbers. For this reason, an in-depth interview is needed to check both their understanding of the concepts and also their representation of the optimistic, pessimistic and most likely estimate. The following table shows the result of one of those interviews. The hours and equivalent headcount are both provided and teams may differ in the way they want to provide the data. The translation can be made in the simulation.

Uncertainty in Variable Hours				
	Nominal	Low	Most Likely	High
Duration Dependent Labor				
EAC Burn Rate (hours/day)	118.5	115.0	120.0	165.0
(Equivalent headcount)	14.8	14.4	15.0	20.6

Figure 8. Interview data on "burn rate"

To review, the total variable cost is found by multiplying (1) the number of days' duration by (2) the hours per day and then by (3) the compensation per hour. Each input comes from a specific source as shown in Figure 6, above.

The compensation per hour really represents the skill mix of the people presumed to be on the job. If skill levels 2, 3 and 4 are used in the original estimate, an average compensation can be computed from the weighted average of those to be assigned in the baseline plan. During the teams' deliberations they often noticed that there were a predominance of levels 4 and 5 on the job. This finding, or projection, implied that the rate per hour in the estimate was not accurate. In several instances, the skill mix assumed in the initial baseline estimate was quite different from that elicited from the interviews. The table below is a typical entry for the uncertain compensation rate for variable costs.

Uncertainty in Labor Rate				
Labor Rate				
Dollars/Hour (Using current data)	$122.12	$110.00	$130.00	$160.00

Figure 9. Interview data on labor rate

Collecting risk data on independent costs. The term "independent cost" refers to the cost of the project elements that do not rely on the duration of the activity. Cost elements in this category might include software releases that are provided by a vendor, a component or piece of equipment provided by a subcontractor, or a piece of purchased equipment. The basic characteristic of these could be that their cost, while uncertain, does not depend on the date of arrival. Thus, while we might not know how much a piece of equipment may cost at an early point in the project before we get a firm quotation, its cost will not necessarily be higher if it arrives later rather than earlier in the project.

There were really two types of independent costs elements in the analysis. One, the simpler, would be represented by a piece of equipment or purchased item that had no time dimension at all. The 3-point estimate for this element of cost has the same concept as is usually found in direct estimation of cost risk. The interviewee prepares three scenarios representing the low, most likely and high estimates of cost. These scenarios are discussed and sometimes challenged during the risk interview and three estimates of cost are developed that correspond to the scenarios. These three estimates are used to develop a probability distribution, often a triangular distribution, of the risk that the element will cost more or less than estimated. The cost model simulation uses samples directly from this distribution.

Three-point estimates for other so-called independent costs were really derived from assumptions about how long activities would take. Some of these independent cost items were actually estimated with assumptions of hours, but not necessarily calendar duration. This category included items that did not have any clear linkage to an available schedule concept. (Risk analysis is often the art of the possible, and sometimes even the best concepts cannot be implemented.)

Correlation of uncertain cost concepts. Correlation often occurs in projects when the elements' costs might be expected to move together in real projects. Elements' costs could move together if a common outside influence is driving them.

One example of correlation between elements' costs can be expected when a particular task or group of tasks is more difficult than expected in the original estimate. For instance, suppose software-coding activities are taking more than 265 days in the current schedule. This might occur if the software turns out to be more difficult than originally expected, or if the estimates of software coding durations were unwarrantedly short as in a "success-oriented" schedule. Perhaps the original assumptions, both for cost and schedule, included reuse of existing software, but that turns out not to be feasible.

If software (or hardware or integration) tasks are taking longer several things might happen: (1) the schedule is in the high end of its estimated distribution, (2) project management puts more people on the activity increasing the per/time unit burn rate, and (3) the people added are more skilled than the average assumed in the estimate, which drives up the compensation or labor rate.

In this scenario, it would be possible and even expected that the three items in question would all be above their level in the original cost estimate. That is, the time would be longer, the hours per day more and the labor rate higher, and these items would all be expected together in the same project. The effect could also be expected to operate in reverse. That is, if the software (hardware, integration) tasks turn out to be easier than originally estimated, then duration would be shorter, the head count lower and the skill level (hence the labor rate) lower together.

It would therefore be said that these three items are positively correlated. That is, if many projects were done and some had harder software (hardware, integration) tasks than originally estimated, they would find that all three types of cost elements were moving against them. These simultaneous adverse events would reinforce each other when multiplied for variable costs to produce a sort of magnified high cost of the project. The same effect would work in reverse to produce significantly low cost results. This is the working of correlation.

In this demonstration cost risk model we installed correlations between the three pairs of elements as follows.

Sample Correlation	
Correlation Pairs	Correlation
Duration: Burn Rate	0.5
Duration: Labor Rate	0.5
Burn Rate: Labor Rate	0.5

Figure 10. Example Correlations

In the real project the independent costs could be correlated with these as well. Also, in the real project the teams were queried about the correlations. This concept was unfamiliar to them and had to be explained at some length. Some interviewees did not expect to see any correlation and others did.

RESULTS

Main findings. The cost estimation model was simulated using uncertainty in the schedule duration in days represented by the histogram from the schedule risk analysis, variable cost uncertainties represented by burn rate and labor rate uncertainties, and independent cost uncertainties. Correlations were inserted. The results show that it is very important to represent uncertain durations along with the traditional cost risk analysis variables.

The model was set up so that different combinations of assumption uncertainties could be turned on and off. This allowed us to identify sensitivities and compare the impact of cost risk variables against the schedule-related uncertainty. For instance, in the first sensitivity, one of the scenarios, the "Schedule Risk Only" scenario, looked like this:

Risk-Type Variables	Turned On
Schedule Variation	Yes
Independent Hour Variation	No
Burn Rate Variation	No
Labor Rate Variation	No
Material Variation	No

Figure 11. Specifying partial scenarios: Schedule-risk-only- scenario

The first question was whether the traditional cost risk items such as independent cost, burn rate, labor rate or material variations were more important in determining the overall cost risk than schedule risk (histograms from the schedule risk analysis) alone. For this sensitivity, we specified the scenario in Figure 11 above, than reversed the items for the cost risk analysis alone. These results are shown in Figure 12.

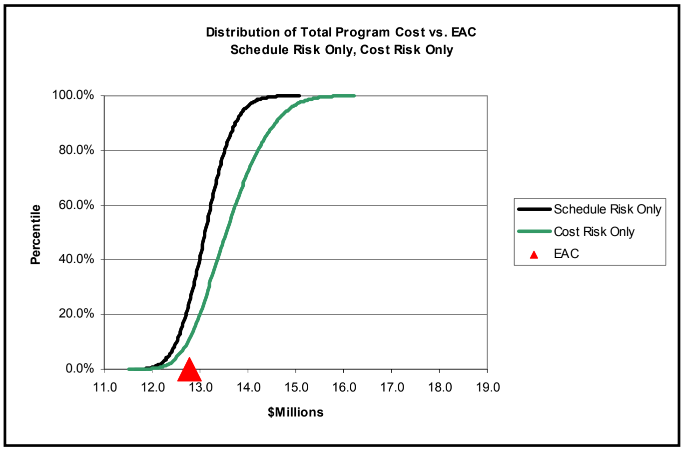

Figure 12. Scenarios of cost risk and schedule risk separately

The curves show that the cost risk scenario has the greater impact on the uncertainty in project costs, with some iterations being somewhat below and others being quite a bit higher than the EAC, which is shown as a triangle. This shows that even if only cost risk items were considered and schedule assumptions were held fixed at their baseline levels, costs could be two million higher at least 5% of the time. The curve for the schedule risk (histograms) only is much steeper than that for cost risk elements only.

The next question is whether the interaction between cost and schedule risk adds much to the results when using cost risk. The answer to this question will demonstrate whether the integration of cost and schedule risk elements into one model and simulation has been worth the efforts described above. This issue is addressed in Figure 13 below.

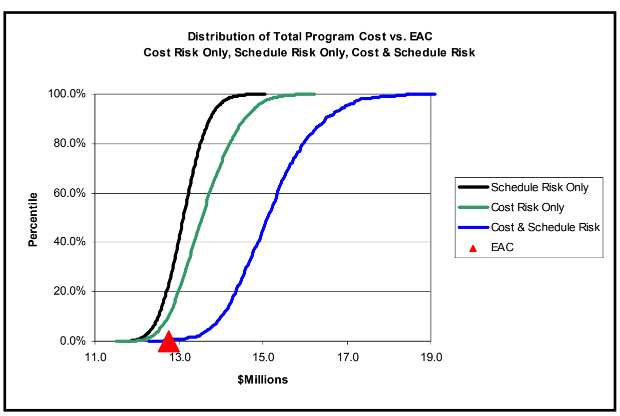

Figure 13. Simulating schedule and cost risk elements together

The results in Figure 13 indicate that the total cost risk of this project would be incompletely specified if only cost risk elements or only schedule risk elements were used to represent project risk. The cumulative distribution of project cost uncertainty is further to the right than either cost or schedule elements alone would indicate. Of course it is very unlikely that a cost risk analyst would incorporate only the uncertain schedule data in a cost risk, but it has been known to happen. Usually the cost risk is represented by the uncertain cost risk elements alone. If schedule risk is not taken into account the cost risk will be underestimated, in some cases by a substantial margin.

Figure 14 below indicates the degree to which the integrating schedule risk makes a difference.

Compare Cost, Schedule and Integrated Risk			
	Cost Risk Only	Schedule Risk Only	Cost & Schedule Risk
	($millions)		
Minimum	11.5	11.6	12.3
Maximum	16.2	15.1	19.1
Mean	13.9	13.4	15.8
10%	12.8	12.5	14.0
20%	13.0	12.7	14.4
30%	13.2	12.9	14.6
40%	13.4	13.0	14.9
50%	13.6	13.1	15.1
60%	13.7	13.2	15.4
70%	14.0	13.4	15.6
80%	14.2	13.5	16.0
90%	14.6	13.7	16.5

Figure 14. Components of project cost risk

One measure is the likelihood of overrunning the EAC, which was $12.8 million in this example. The table shows that this value could possibly be achieved with either cost risk elements or schedule risk elements alone, but it is not very likely with both cost and schedule risk elements taken into account.

Another measure of the contribution of integrating cost and schedule risk elements into the cost risk estimate is the mean of the distributions. The EAC of $12.8 million increases to $13.6 million when all cost risk elements are considered, but it increases to $14.3 million when schedule risk elements are added. Thus, the overrun at the mean is only $0.6 million with cost risk elements alone, but it is $1.5 million, more than twice that, when schedule risk elements are added.

The same story is told at a greater level of safety, taking the 80% level. Here cost risk elements alone have an estimate of $14.2 million but a complete analysis with schedule risk elements shows $15.0 million.

Clearly the inclusion of schedule risk in the cost risk analysis is meaningful and necessary. Without the schedule uncertainty, the cost risk analysis would underestimate the risk of meeting cost objectives and not provide the project manager complete information.

Using the results to assist risk response planning. The cost risk model was configured to allow each type of risk to be simulated separately or to be subtracted separately from the total cost risk. The first example of this use of the model is to see which type of cost risk variable provides the most risk. There are two ways to look at risk in this context: (1) the greatest increase in cost risk from the EAC representing overruns, and (2) the greatest spread of risk from lowest to highest representing uncertainty or imprecision in any particular estimate of cost.

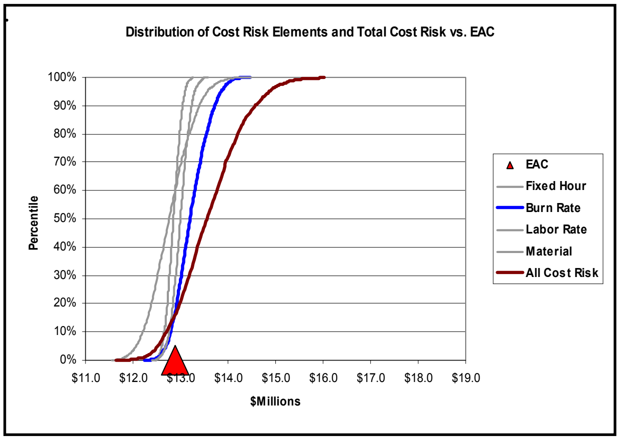

Figure 15. Burn rate contributes the most cost risk difference from the EAC

Figure 15 above shows that the burn rate of all cost risk elements contributes the most to increases of the risk over the EAC. Control of the rate of resources per unit time may help to manage the risk of cost overruns the most. Of course, controlling this factor may be difficult and the ratio benefit / cost of burn rate control might be very low.

The cost risk element that contributes the greatest total uncertainty from positive (under runs) to negative (over runs) is the labor rate in compensation per hour. Uncertainty about the skill mix of people on the job can cause the cost to under run or over run the estimate as shown in Figure 16.

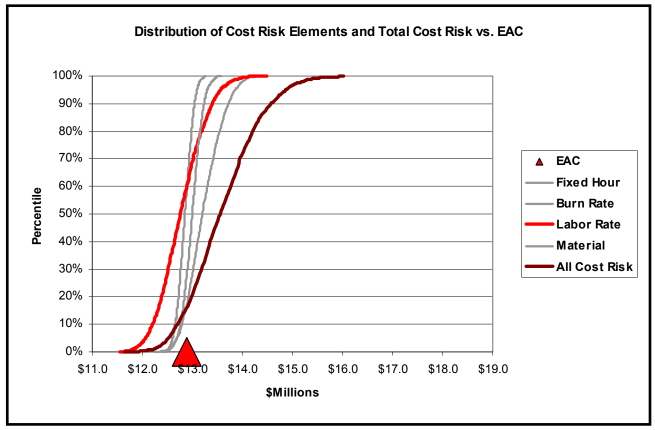

Figure 16. Labor rate contributes the most total uncertainty in the cost risk

Another way to examine the benefit of risk mitigation is to subtract each cost risk type from the total risk representing all cost risk types and total schedule risk. In this way the project manager may see what risk responses might reduce total project cost risk and prioritize risk mitigation actions. The information in Figure 17 shows that the greatest improvement in total cost risk might come from improving the burn rate. It shows what would happen to total cost risk if the project manager might hold the burn rate to baseline levels. Whether this is possible or not is another story. Still, controlling the number of resources applied per unit time would seem to offer significant cost risk reduction.

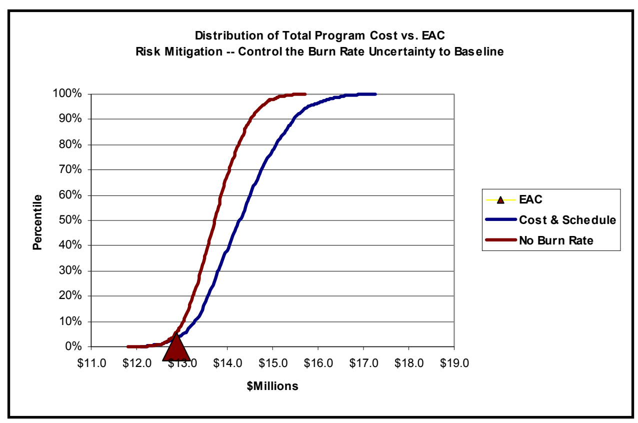

Figure 17. Controlling the burn rate could contribute the most to reducing cost risk

OBSERVATIONS

Methodology. There are several issues with the methodology of integrating cost and schedule risk analysis.

♦ Cost and schedule structures can become disjoint, making linkage of the two difficult. Serious effort was required to link the structures so the costs could be allocated to schedule concepts. The recommendation is to structure both at a common starting point, probably the WBS.

♦ Modelling issues were serious as well. Decomposing the baseline estimate into its components and discovering what must have been the baseline assumptions was tricky. Keeping the estimates clean was not possible, and many approximations were needed. Making the model work when many different cost and schedule elements are present, well beyond the simple example described here, implies a very complicated spreadsheet model. Including the facility to simulate one or another type of variable alone was complex modelling as well.

Data collection issues. Collecting high-quality data that are as accurate as possible is the most important part of the analysis.

♦ There are new concepts for the interviewees, such as the burn rate, average compensation and correlation between elements. We simplified the burn rate to a level-loading concept, which further complicated it for the interviewees. Still, the teams were able to understand and respond.

♦ The explicit use of burn rate and labor rate highlights the basic forces of the cost risk elements of variable costs. It was found that the participants could look at both elements separately, which we believe improved the representation of variable cost risk.

♦ Data collection is most difficult when the participants have never been part of a risk analysis, as is often the case. Expert facilitation is often necessary. The good news is that this is not brain surgery, and participants can learn while doing.

♦ Individual and organizational maturity comes from treating data collection seriously. It includes preparing the teams and giving them a sense of importance for the effort. It includes providing enough time for the exercise and making it part of the regular work, not an extra. It also includes recognizing that the team leaders and members may not provide perfect data the first time they are interviewed. They may answer one way the first time and a different way in a later interview. Usually the second time more risk will be reported. The mature organization will understand that people become more comfortable discussing risk as they learn how to do it.

CONCLUSIONS

This paper has demonstrated that a good cost risk analysis needs to incorporate schedule risk as well as the traditional cost risk elements. Integrating the results from the schedule risk into the cost risk model provides a more complete picture of cost risk than if it were excluded. The picture of cost risk that emerges will depend significantly on the degree of schedule risk that is found in the project.

Often cost risk analysis looks at individual cost elements and tries to figure out what drives their uncertainty. Sometimes participants in a traditional cost risk analysis reveal uncertainty in activity durations, but those concerns are usually *ad hoc* and incomplete. The disciplined approach of this paper makes the consideration of schedule risk explicit and unavoidable.

Often a “troubled project” is having as much difficulty keeping on schedule as it is on budget. These two problems reinforce each other, magnifying the problems. This magnification is represented in the present model by two things: (1) multiplication of time, and therefore time uncertainty by burn rate and therefore burn rate uncertainty, and (2) correlation between burn rate and duration.

The cost estimates often lose track of schedule realities, causing the estimates of cost to be unrealistically low in many cases. Estimators feel that they did not have a full kit of information when the facts of the schedule are made clear to them, making them look bad at estimating and causing them to lose credibility with the project executives. Insisting on a close communication between the scheduling and the cost estimation functions will help to reduce this problem. Insisting on looking at the schedule risk when evaluating the cost risk will improve the accuracy of the estimates of project cost and of cost risk.