

Schedule Risk Analysis (Presentation to be given at the MPUG web vendor showcase on June 26th 2013.)

Tony Welsh, CEO, Barbecana Inc.

We are going to learn how to actually use schedule risk analysis (SRA) software, but first I want to spend some time discussing why it is so important.

So, why is SRA so important? Have you ever noticed that projects don't always turn out as we expect? (In fact they never turn out as we expect.) Why is this?

"The best laid plans o' mice an' men gang aft a-gley."

Robert Burns

"No plan survives the first contact with the enemy."

Helmuth von Moltke the Elder

Project plans are about the future, and the future is uncertain. So all the information we use to create our plans are at best estimates, or forecasts. The actual outcome of our project plan is dependent on these uncertain data, and if we want to produce good predictions of these outcomes we need to take account of the uncertainty. In particular we need to start asking the right questions:

- NOT, "When will the project be finished?" or "What will it cost?" because these questions imply a precision which we cannot justify. (The answers to these questions, like the deterministic plans we use to deduce them, will certainly be wrong.)
- RATHER, "What is the chance of it being finished by?" or "What is the maximum likely cost?"

Any kind of management is about the future. (After all, we cannot manage the past.) And so the above applies to any kind of management, indeed to any kind of human endeavor. So what is so special about project schedules?

The answer, in two words, is MERGE BIAS. Merge bias is a phenomenon of project networks which means that most deterministic project networks are not just unrealistic in their failure to reflect our uncertainty, but also biased towards undue optimism. So while we might think we have about a 50% chance of meeting these schedule dates – they are based after all on roughly 50th percentile estimates – the actual probability may be much less.

So, what is merge bias? At various points in a project network, two or more paths will merge, which means that some successor task cannot start until all of these paths are complete. Which in turn means that if one happens to take less time than expected while another takes longer, we have to wait for the longer one and get no benefit from the earliness of the early one. We do not gain on the swings what we lose on the roundabouts.

The effect of this is that the average start time of the successor is later than the average finish time of any of the predecessors. This effect tends to build up throughout the project, and there are no countervailing tendencies in the other direction.

Let's consider a simple analogy using dice. Suppose we have a task which could take anywhere from 1 to 6 weeks, represented by the 6 possible outcomes of the throw of a die. What is the expected value of a single throw of a die?

The answer is 3.5. (Statisticians use strange language sometimes. The "expected value" is what most of us call the mean or average; nobody "expects" the die to come up 3.5!)

Now suppose we have two dice, representing the durations of two merging parallel tasks. We cannot start the next task until they are both done, represented by the higher of the two dice scores, so what is the expected value of the higher score?

We can answer that question by considering all possible 36 outcomes, as on this spreadsheet (in the presentation this will be shown live):

	1	2	3	4	5	6
1	1	2	3	4	5	6
2	2	2	3	4	5	6
3	3	3	3	4	5	6
4	4	4	4	4	5	6
5	5	5	5	5	5	6
6	6	6	6	6	6	6

The highlighted row at the top indicates the score from the first die while the column on the left shows the score from the second. The cells in the middle represent the 36 combinations, and what the bigger of the two values is in each case. It is fairly clear that there is only one chance in 36 (less than 3%) that the maximum will be 1 while there are 11 chances (over 30%) that it will be 6.

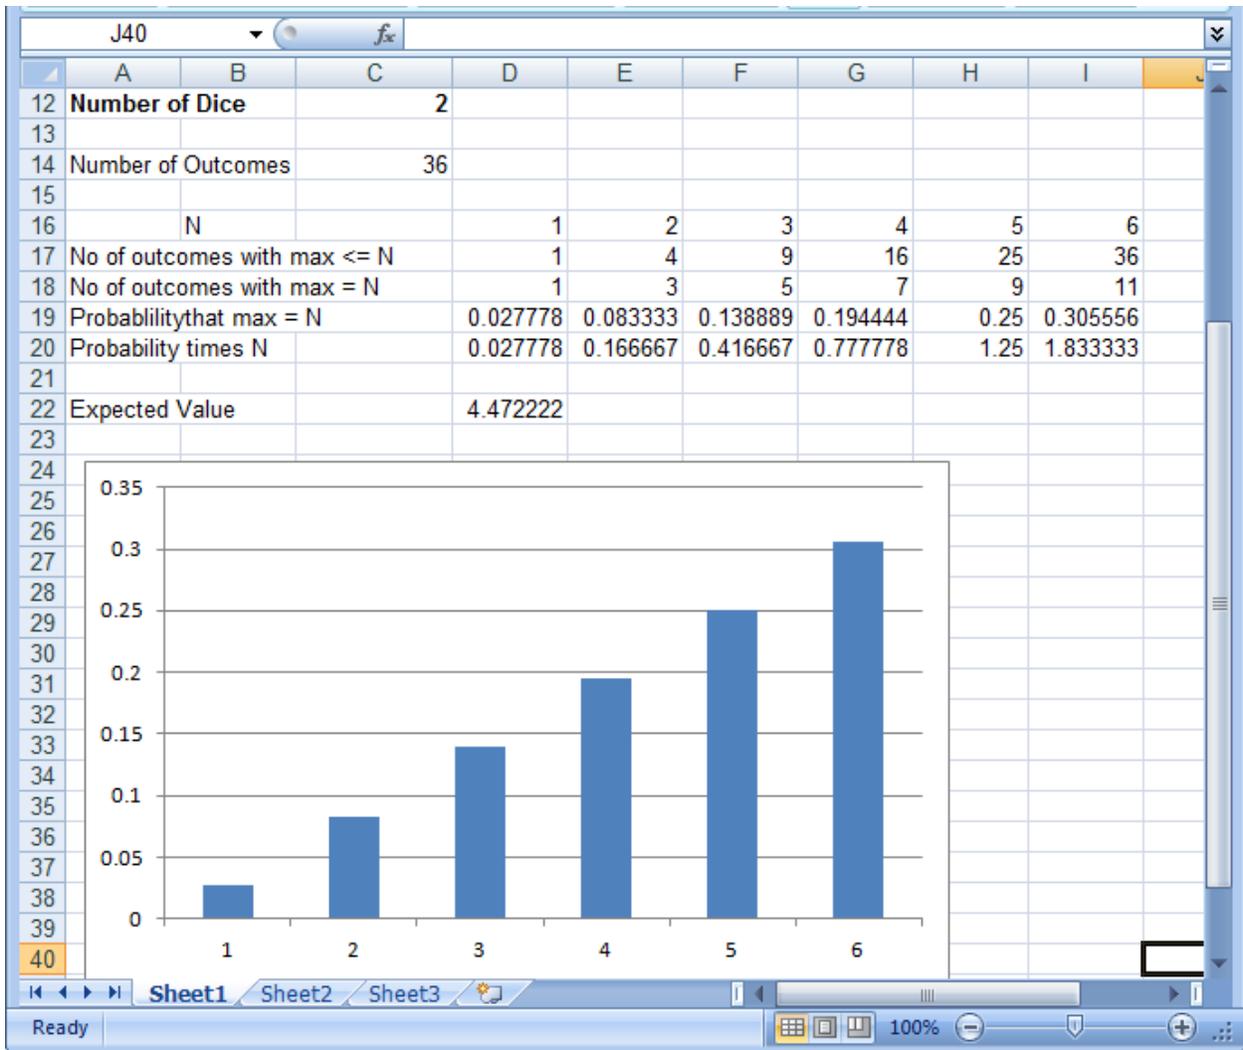
It turns out that the expected value is about 4.5.

If we had more than two paths, it would be worse. This is hard to show graphically but we can do it mathematically in the spreadsheet. (This spreadsheet can be downloaded from our web site, by the way.)

The number of dice can be changed and the probabilities recalculated and plotted. In the presentation we will start with just one and vary it. The result for two dice is shown on the next page.

(I probably will not have time to explain the formulae, but they are quite simple. The total number of outcomes is 6^M where M is the number of dice. The number of these where the highest value is less than or equal to N is N^M . Hence we can get the number of outcomes where the highest value is exactly N and hence the probabilities. These are plotted on the histogram.)

The result gets more extreme the more parallel paths we have. It should be obvious that as the number of dice increases the probability that the highest value is 6 will approach 1, and so the expected value will approach 6. (Meaning 6 weeks instead of 3.5 weeks in our project example.)



Now, if projects were as simple as this, all we would need is a spreadsheet. But they are not. Tasks combine in complex ways, sometimes in parallel and sometimes in series. Their duration distributions are also more complicated than the throw of a die which has only 6 discrete possible results, and may also be correlated.

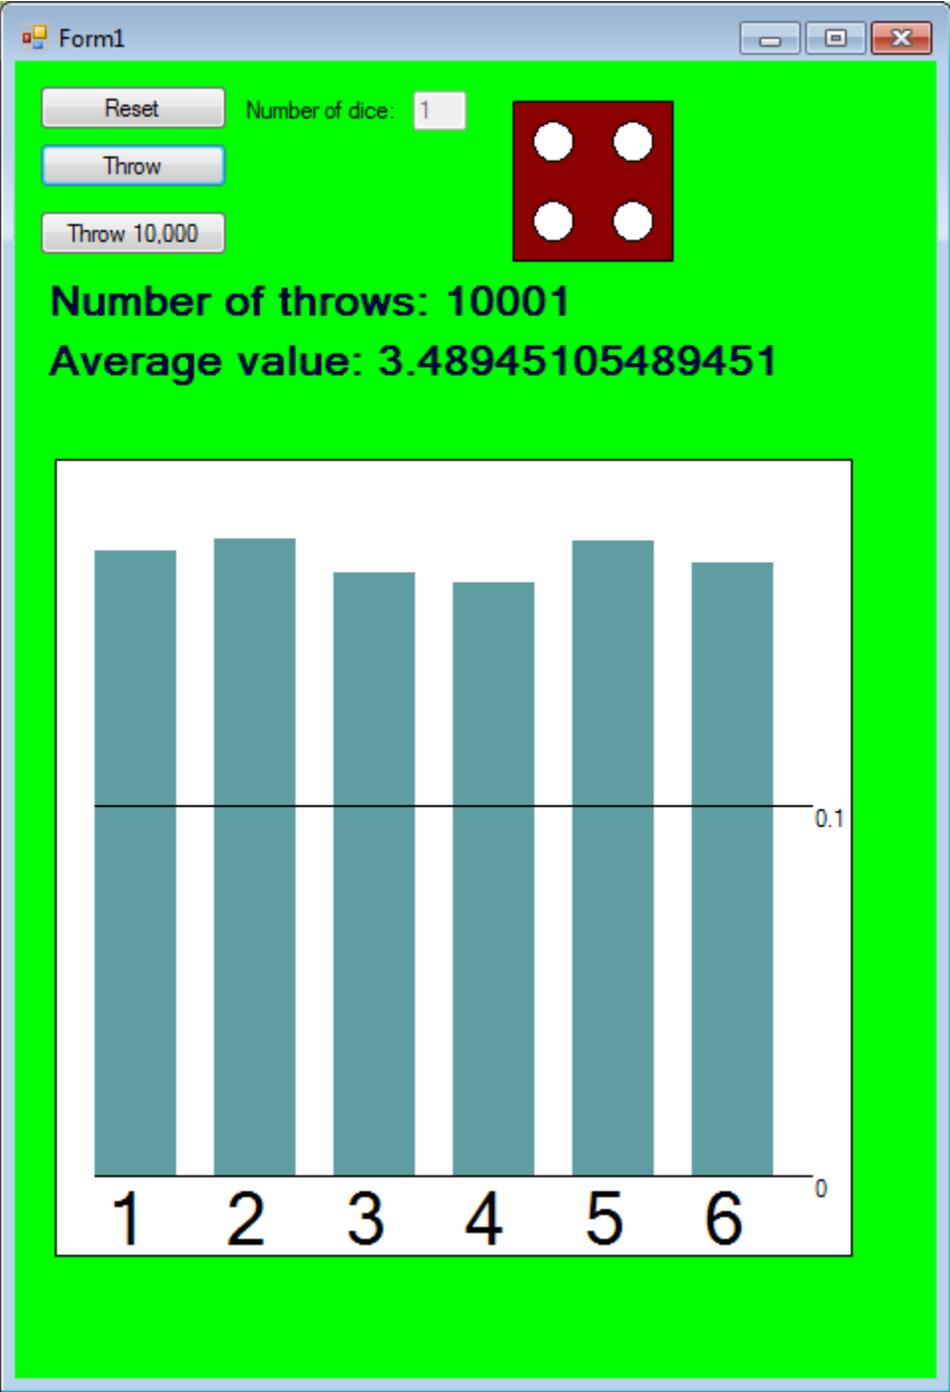
Attempts have been made to solve the problem analytically, starting with PERT and more recently using the Method of Moments. All have to make unjustifiable assumptions in order to stop the math becoming intractable. (PERT in particular takes no account of merge bias and is not much more than a deterministic CPM with slightly different duration estimates.)

The method universally accepted nowadays is Monte Carlo simulation. Essentially we get the computer to pretend to execute the project thousands of times, each time using a different set of durations “sampled” from probability distributions specified for the task durations. From this we can build up a picture of all the possible project outcomes and their relative probabilities.

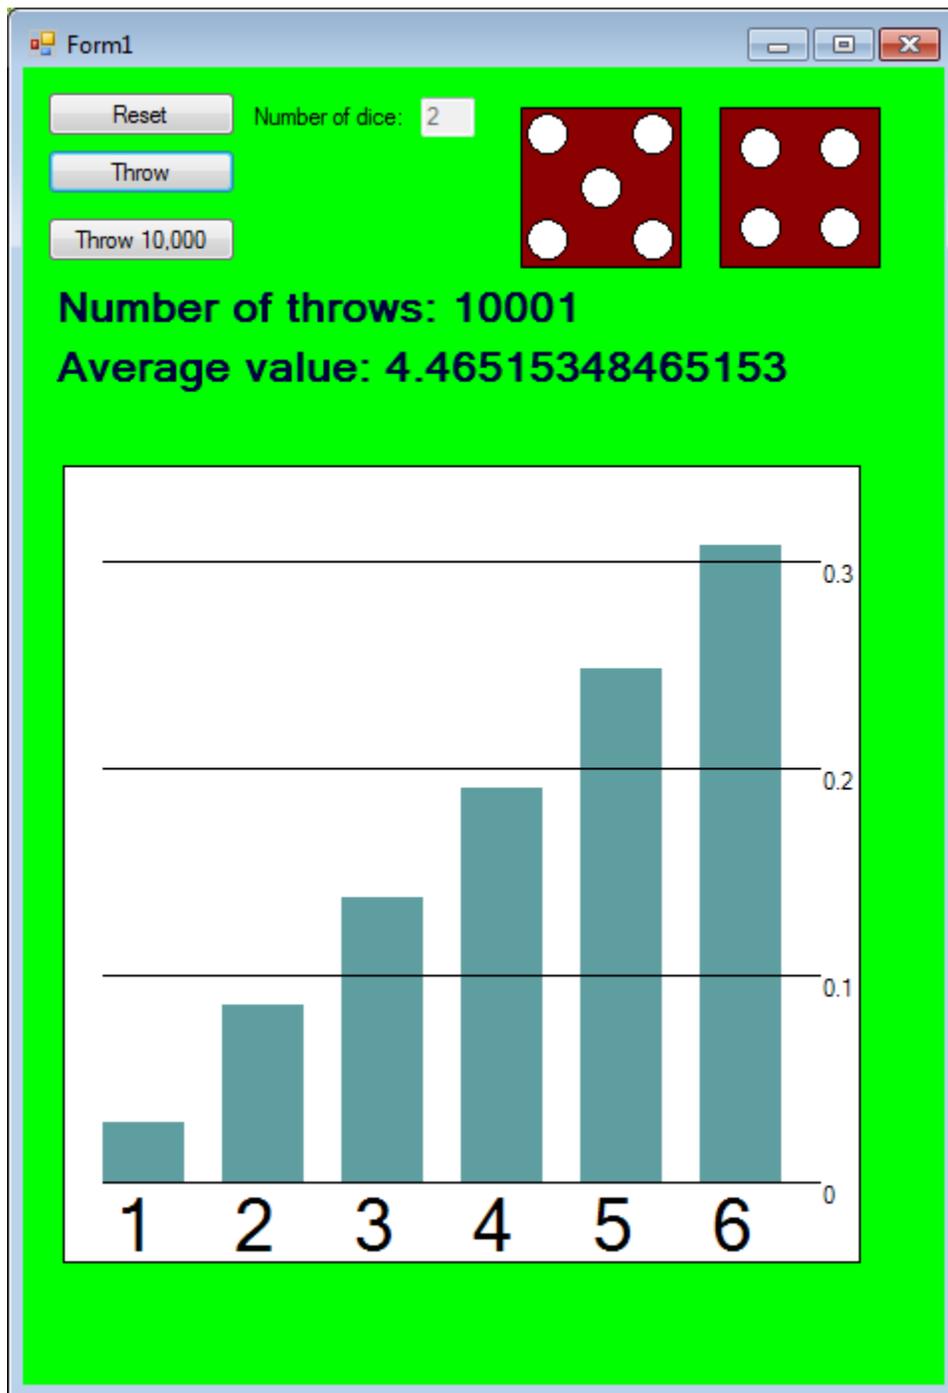
Before moving on to a real project, let’s try Monte Carlo on the dice example for which we know the correct answer.

(The following are shots from a live simulation program which can be downloaded from our web site.)

The first shot shows the result of throwing a single die 10,001 times, the last result happening to be a 4. The histogram shows a roughly uniform distribution as we would expect, and the expected value is close to the theoretical value of 3.5.



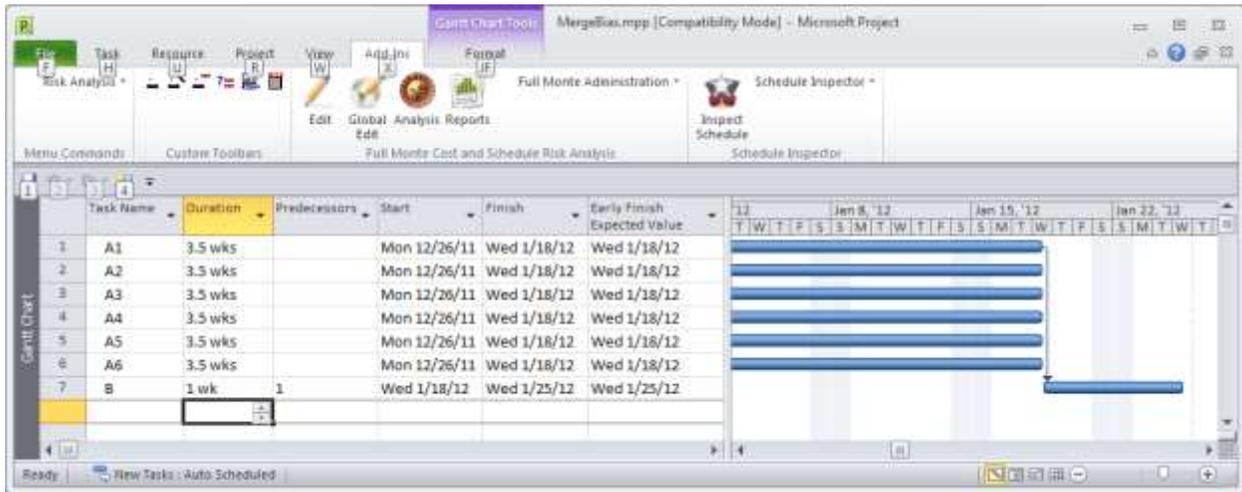
Now try two dice:



Again the result is close to the theoretically correct answer with an expected value of 4.47.

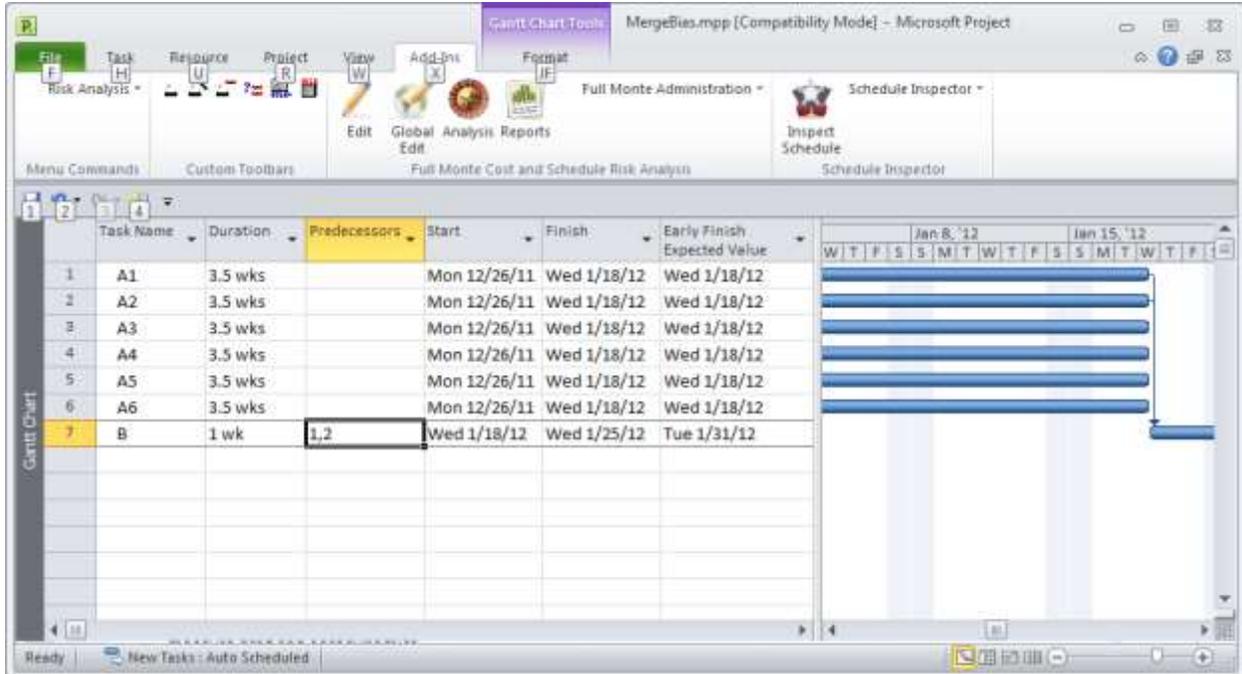
But note that it is not exact. This is because it is based on a random process and relies on the law of large numbers to make the results tend towards the true ones. For this reason, it is important to do a sufficient number of trials, which in turn means that the software needs to be fast.

OK, now we are ready to look at real project simulation software. We are going to use Full Monte, which works directly on your Microsoft Project schedules. This project is designed to mirror the dice example:

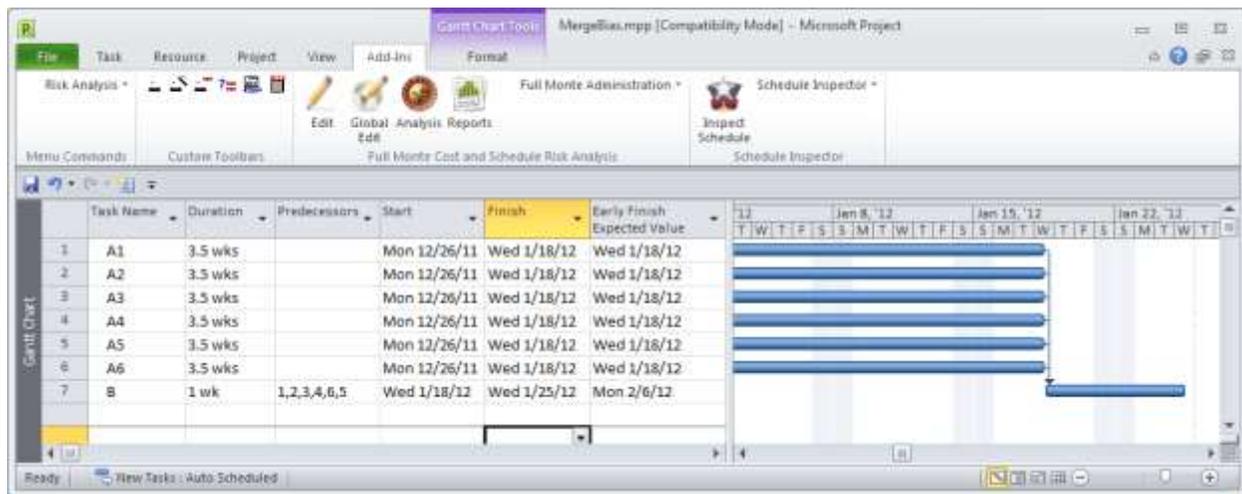


We have 6 tasks each with an expected duration of 3.5 weeks but with a possible range of 1 through 6 weeks. You can see that all the finish dates from Project are identical with the expected values created by Full Monte. That is because right now only one of the parallel tasks is actually linked to a successor.

If we link two of them the situation changes: the expected finish dates of each predecessor is unchanged but the finish date of the successor (and the whole project) is increased by about a week. No swings and roundabouts!



With all six tied in to the project finish the merge bias is more like two weeks:



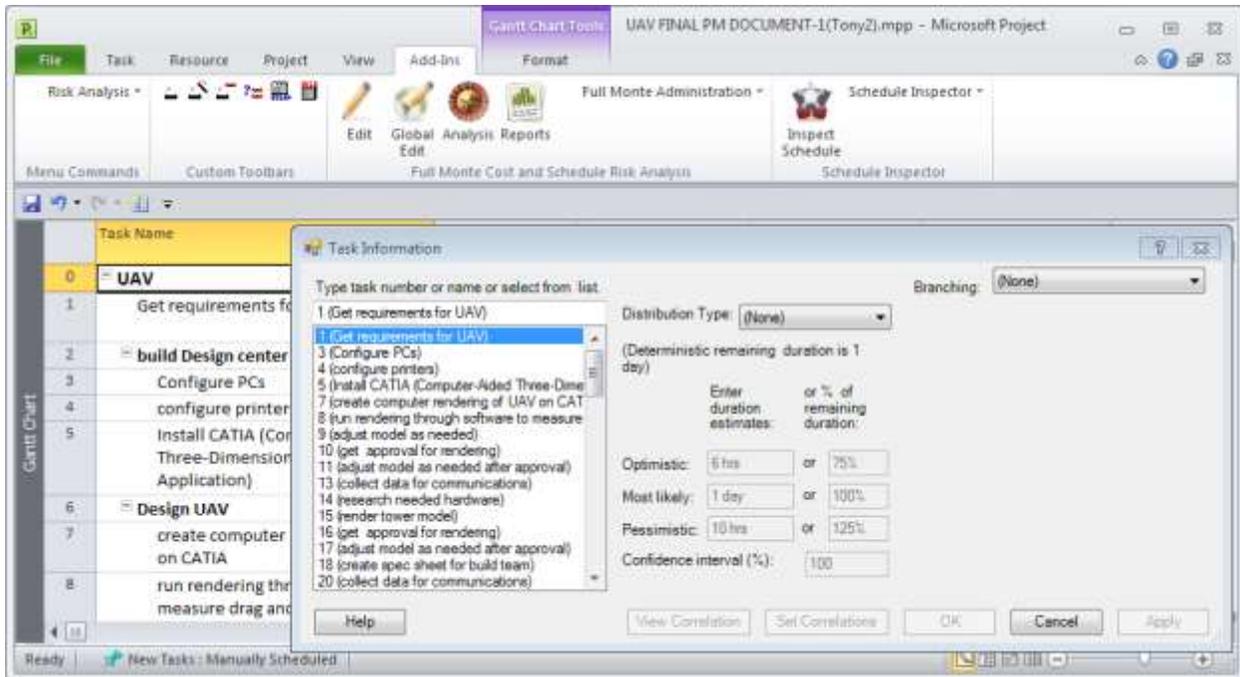
Now this is a more extreme merge than one is likely to get in a real project, for two reasons:

- There are six tasks with identical duration distributions. (Merge bias occurs only when more than one path has a chance of sometimes being the longest, which is less likely if the distributions are different.)
- Each task duration can vary by a factor of six, which is not a likely scenario.

On the other hand this is just one merge point. In a real project there are likely to be many such merge points, and the delays build up throughout the project, making it increasingly difficult to meet project milestones.

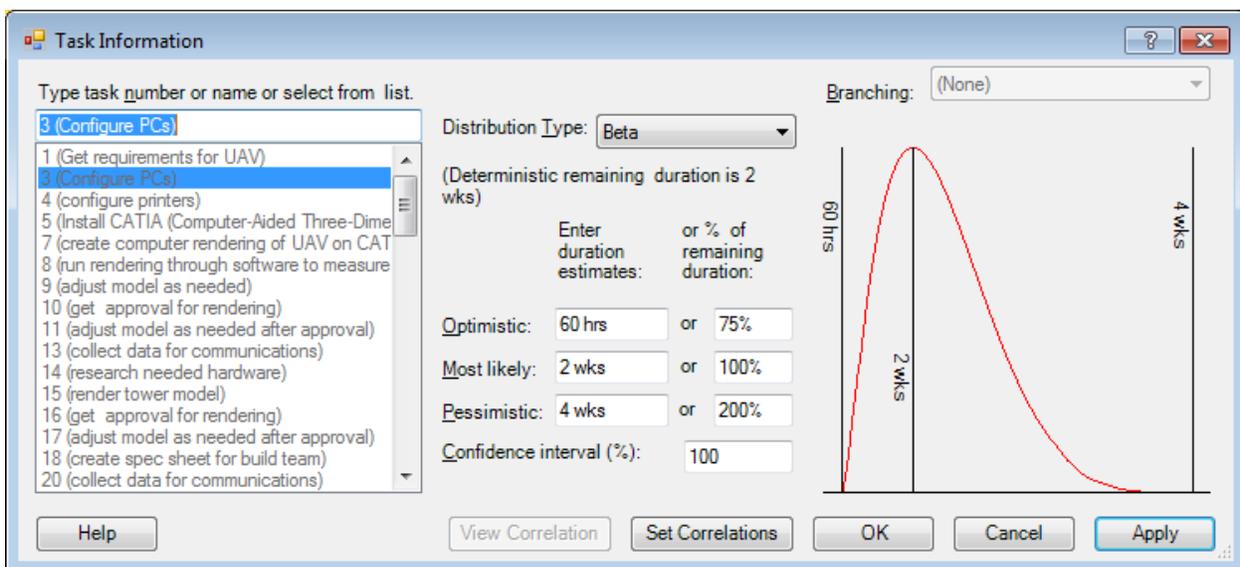
I am now going to move to a more realistic project.

You enter your project data in the usual way, and then add data specific to Full Monte. The main Full Monte functions are accessed through four icons on the add-in page. (This is Project 2010; it looks a little different in 2007.) We will start with Edit.



On the left of the dialog we can navigate to the task we want by typing, clicking, using up/down arrows, etc.

The center and right hand side of the dialog are about the selected task. (Once we start modifying data here we cannot navigate to another task until we commit the changes or cancel them.) Here we have set up a beta distribution for task 3:



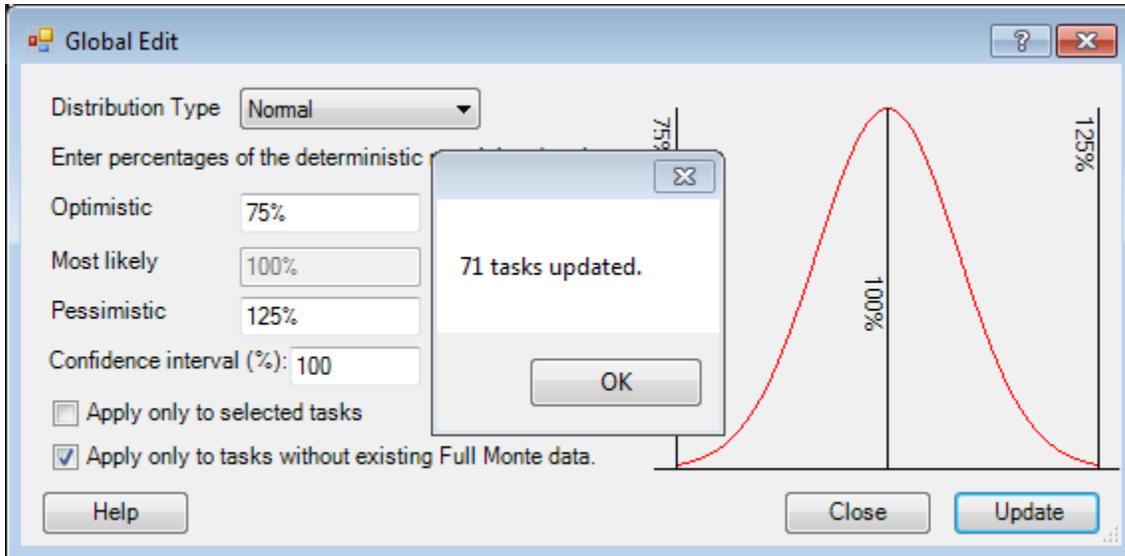
(In the live presentation I will show the various distribution options and if time permits also deal with correlations and branching.)

Note that we have to enter an optimistic and pessimistic value of the duration and maybe a most likely value too. Some people object to risk analysis because they find this onerous, but in reality it is easier to provide a range of values than to specify one, which is bound to be wrong.

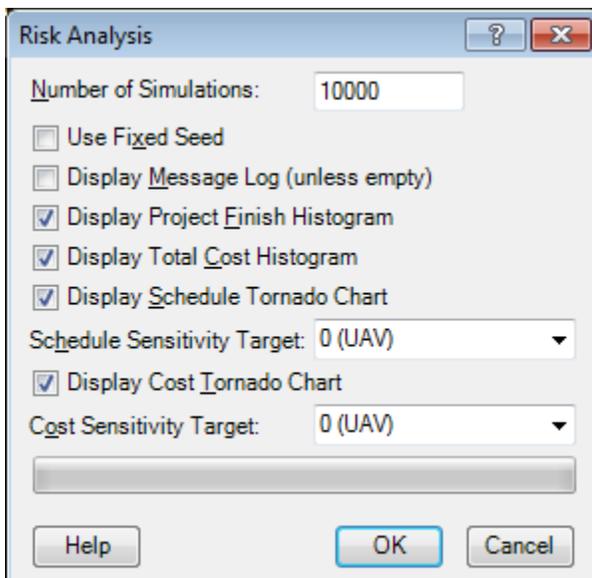
“He saw no fundamental irony in his position: Because he believed he did not have enough data to estimate a range, he had to estimate a point.” Douglas Hubbard, in the Failure of Risk Management.

To illustrate this, how tall am I? You could guess say 5’10” but would it not be easier and more reliable to say with some certainty (even without meeting me) that I am between 5’ and 7’?

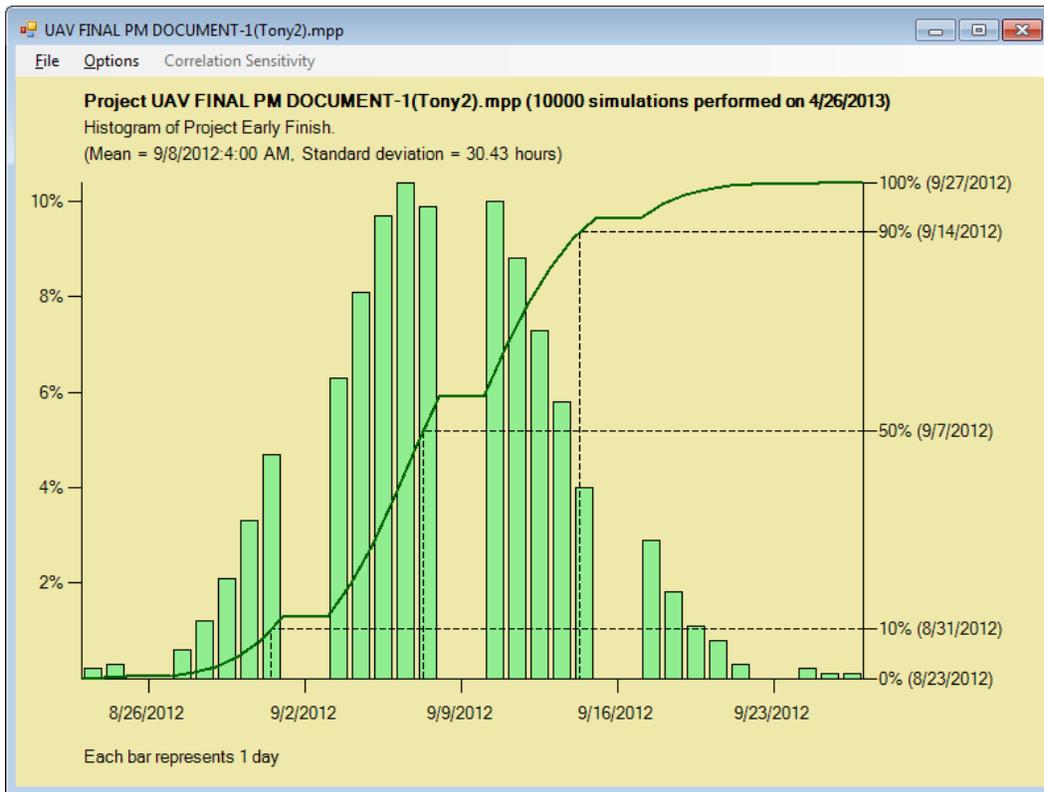
To save time I will enter most of the data using the global edit function. Note that there is an option not to overwrite data put in specifically for the most important tasks in the dialog above, so in this case it updated only 71 out of the 72 detail tasks.



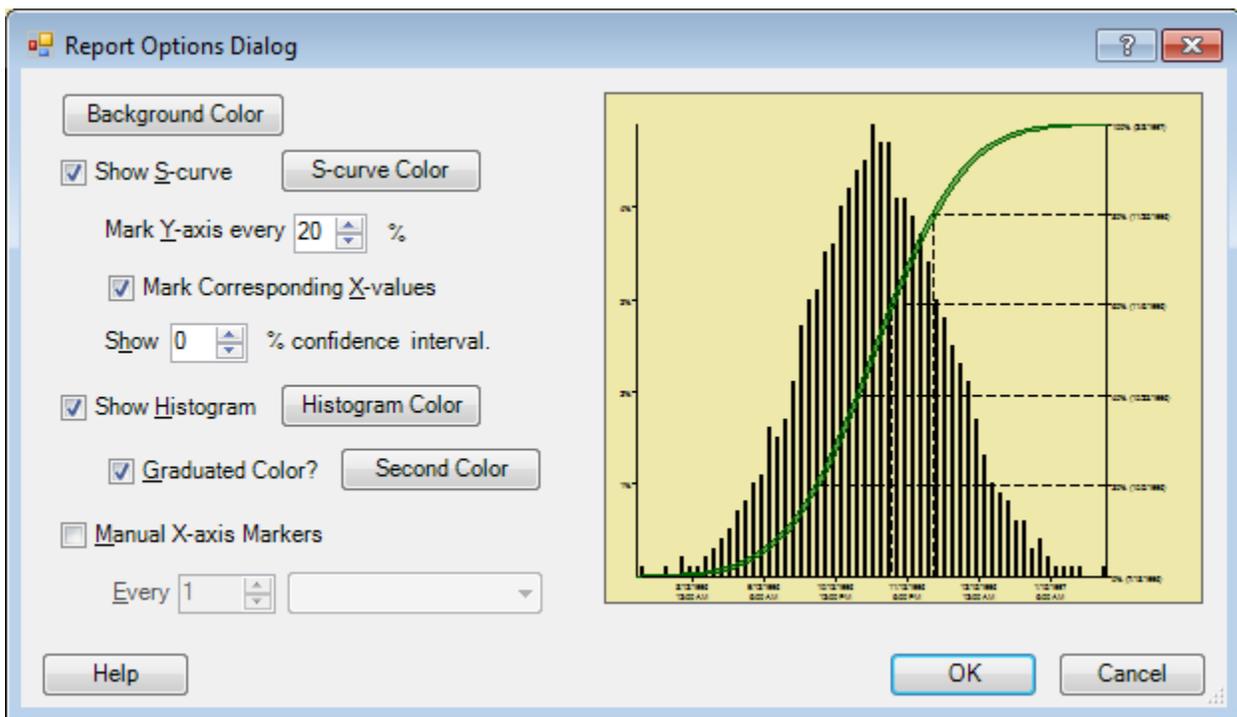
It is now time to do the analysis. The dialog remembers settings last used on this particular project. This is set up to do 10,000 trials and to display a few reports at the end. (As we shall see, the reports are available later whether or not we view them here.)



Histogram of project early finish:

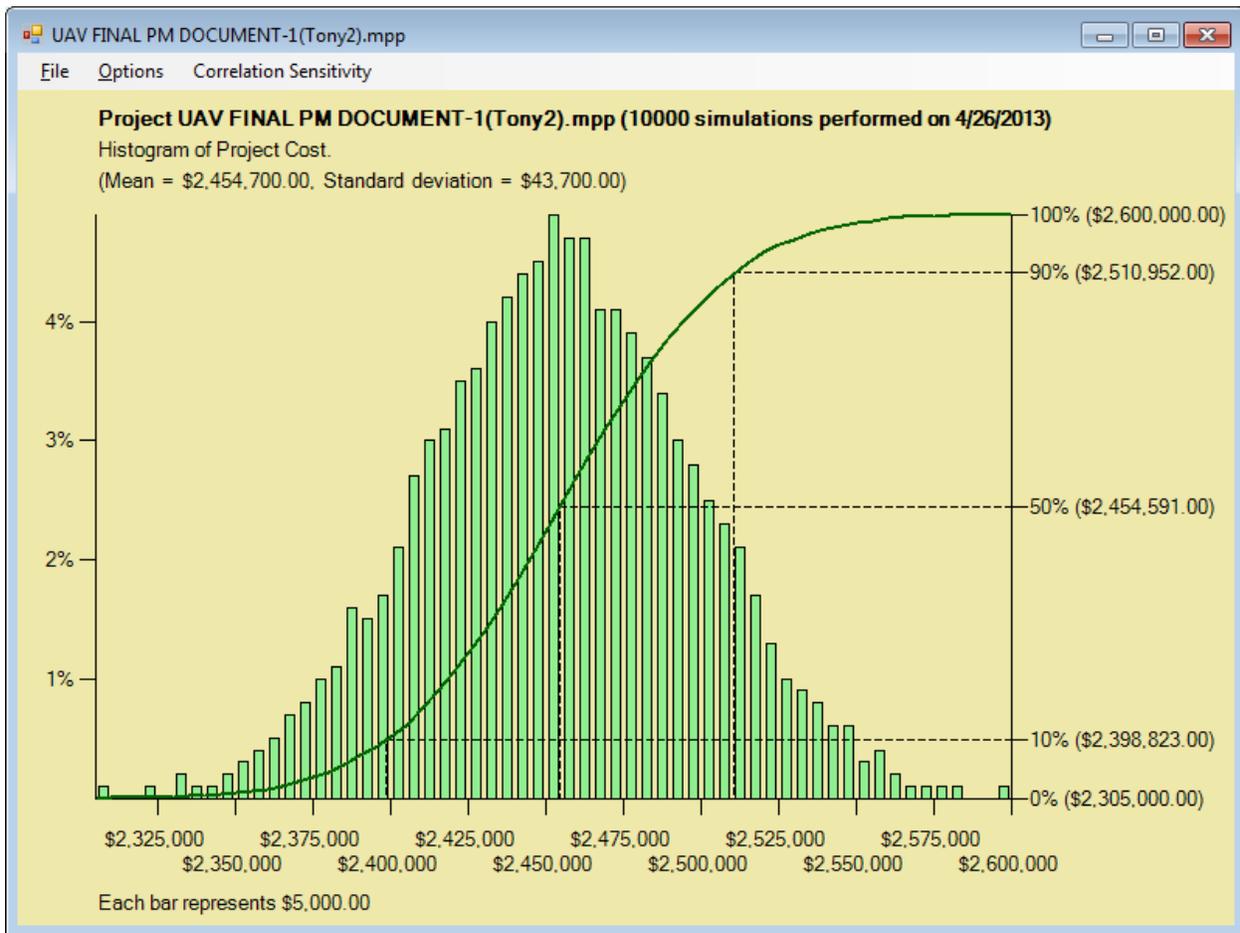


(Probably will not have time to deal with the options on this graphic, as shown in the dialog below.)



(Or to show how graphic can be printed and/or copied to clipboard.)

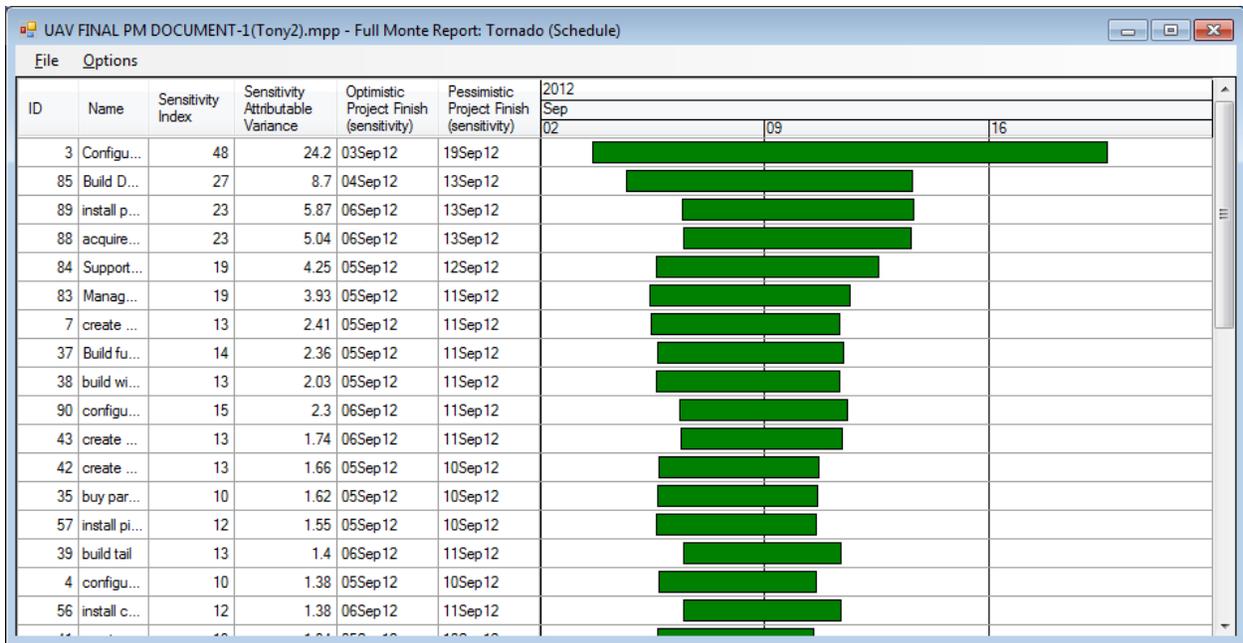
Histogram of total project cost:



The tornado chart on the next page shows estimated sensitivity. This is calculated during the simulation, which means that it takes into account merge bias from all the other tasks. It tells us the tasks where uncertainty has the greatest impact on project finish (or other important milestone) and where it might be worth spending time narrowing down our estimates.

"The word 'chance' then expresses only our ignorance of the causes of the phenomena that we observe to occur and to succeed one another in no apparent order. Probability is relative in part to this ignorance, and in part to our knowledge."
Pierre-Simon Laplace

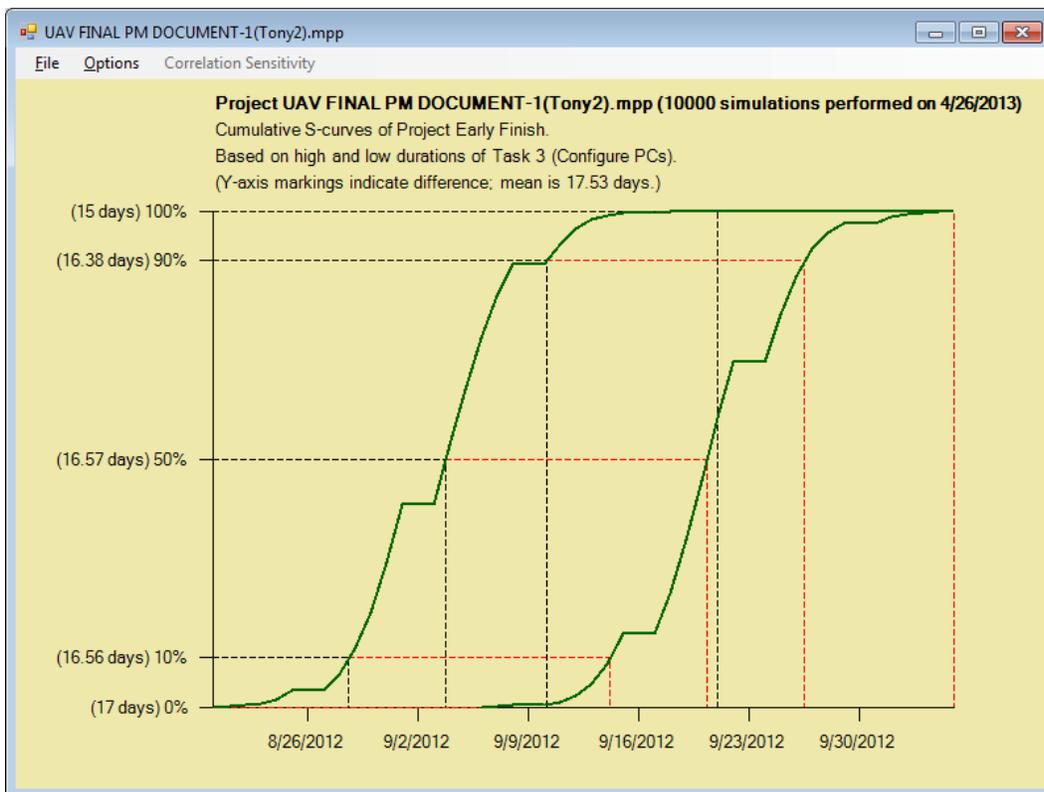
We can sometimes reduce our ignorance by doing more research, but only at a cost, so it is worth knowing where it is most worthwhile.



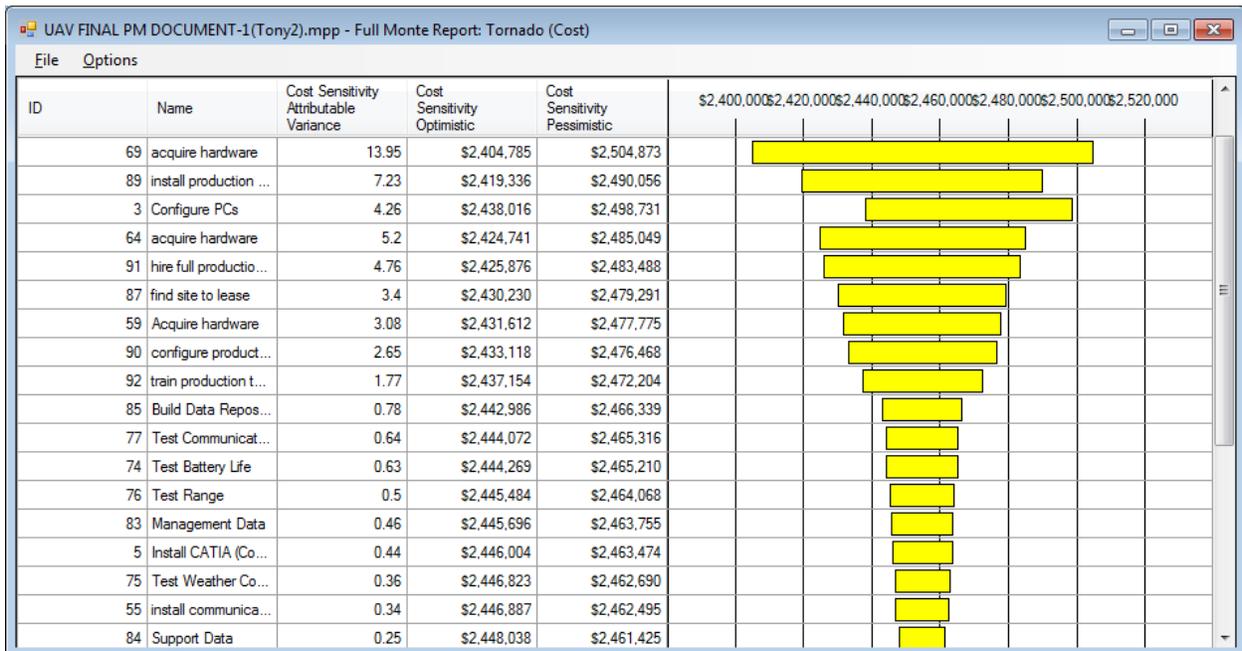
(Note that it is not symmetrical, and in general it should not be.)

The above gives us a shortlist of where to look but various factors beyond the scope of this paper can make the tornado chart misleading. Clicking on a bar gives a definitive result based on two complete simulations.

Result of clicking on the first bar:

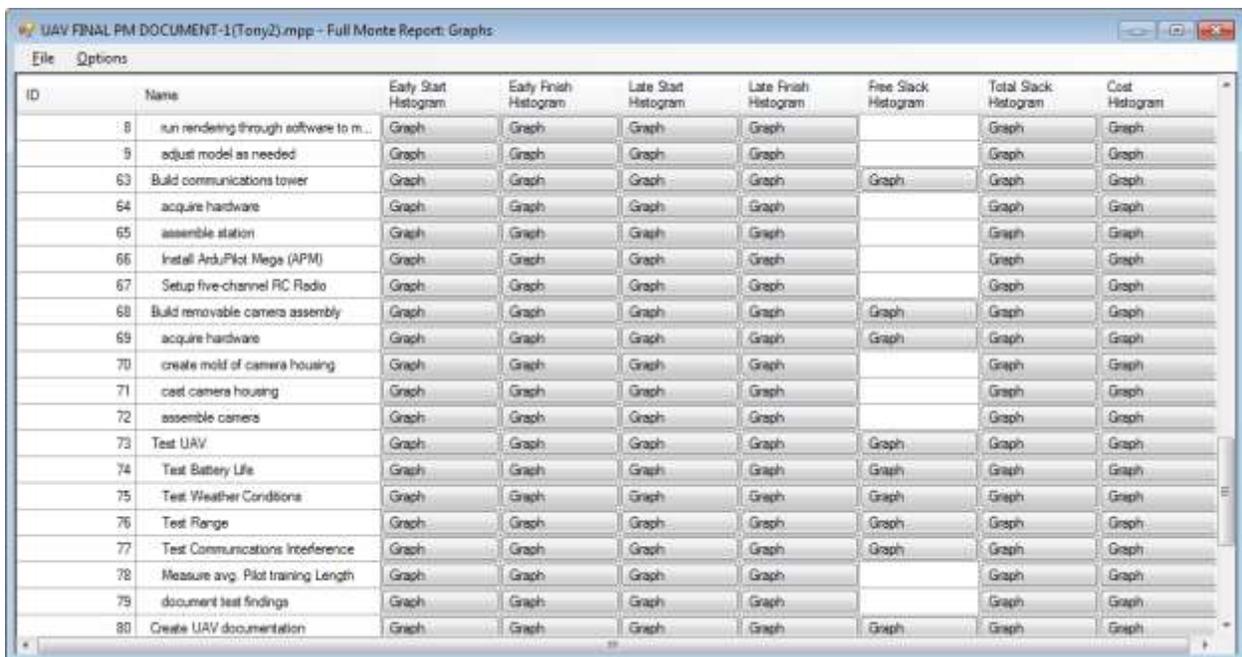


Cost tornado chart. (Again, all the comments above apply here too.)



Finally I want to go to reporting. As noted before, all the data produced by an analysis run are available subsequently for reporting on demand. And unlike some other software, Full Monte generates a full range of reports for every task in the network.

As we have already seen, many of the Full Monte results (e.g. the expected early finish shown in screen shots above) can be displayed in regular Project views, but Full Monte also has its own powerful reporting capability, accessed from the Reports icon. This initially brings up the last report viewed. Below is a standard report we call Graphs, which just comprises a set of buttons to give access to the various histograms.



(Note that histograms are not produced for items which don't vary, for example some of the free slacks in this case. Hovering over a cell without a button will show the specific reason.)

Clicking on a button brings up the histogram. (Similar to those on pages 10 and 11. In the live presentation I will show one or two.) Clicking anywhere else on the grid brings up a synopsis of the results for that task, from which you can also bring up the histograms:

The screenshot shows a 'Results Dialog' window with the following data:

	Expected	Standard Deviation	Histograms
Task Name	acquire hardware		
Active %	100	Critical % 0	Merge Delay NA
Early Start	4/2/2012 7:40:00 AM	NA	Graph
Late Start	9/1/2012 4:00:00 AM	30.53 hrs	Graph
Early Finish	4/9/2012 7:21:00 AM	20.68 hrs	Graph
Late Finish	9/8/2012 4:00:00 AM	30.43 hrs	Graph
Total Slack	21.69 wks	27.28 hrs	Graph
Free Slack	22 days	11 hrs	Graph
Cost	\$199,900	\$16,300	Graph
Sensitivity of Finish of Project 'UAV'			
Optimistic	9/8/2012 4:00:00 AM	Sensitivity Index	0
Pessimistic	9/8/2012 4:00:00 AM	% Variance Attributable	0
Sensitivity of Cost of Project 'UAV'			
Optimistic	\$2,404,785		
Pessimistic	\$2,504,873	% Variance Attributable	13.95

Buttons at the bottom: Help, Show Field Names, Close.

All reports can be printed and/or copied to the clipboard. They are highly customizable, and report designs can be saved for future use. Let's just add a couple of columns. Here I have added columns for the 80th percentile of the early finish date and the probability of the task being critical. You can also add Project columns and define histograms. (In the live demonstration, I will show the range of data available and how the new report specification can be saved under a name of the user's choosing.)

ID	Name	Early Finish Percentile (80%)	Percent Critical	Early Start Histogram	Early Finish Histogram	Late Start Histogram	Late Finish Histogram	Free Slack Histogram
8	run rendering through software to m...	17Jan12	63.89	Graph	Graph	Graph	Graph	
9	adjust model as needed	24Jan12	63.89	Graph	Graph	Graph	Graph	
63	Build communications tower	08Mar12	0	Graph	Graph	Graph	Graph	Graph
64	acquire hardware	08Feb12	0	Graph	Graph	Graph	Graph	
65	assemble station	22Feb12	0	Graph	Graph	Graph	Graph	
66	Install ArduPilot Mega (APM)	01Mar12	0	Graph	Graph	Graph	Graph	
67	Setup five-channel RC Radio	08Mar12	0	Graph	Graph	Graph	Graph	
68	Build removable camera assembly	14May12	36.12	Graph	Graph	Graph	Graph	Graph
69	acquire hardware	12Apr12	0	Graph	Graph	Graph	Graph	Graph
70	create mold of camera housing	19Apr12	36.12	Graph	Graph	Graph	Graph	
71	cast camera housing	03May12	36.12	Graph	Graph	Graph	Graph	
72	assemble camera	14May12	0	Graph	Graph	Graph	Graph	
73	Test UAV	23Aug12	63.89	Graph	Graph	Graph	Graph	Graph
74	Test Battery Life	07Aug12	0	Graph	Graph	Graph	Graph	Graph
75	Test Weather Conditions	07Aug12	0	Graph	Graph	Graph	Graph	Graph
76	Test Range	07Aug12	0	Graph	Graph	Graph	Graph	Graph
77	Test Communications Interference	07Aug12	0	Graph	Graph	Graph	Graph	Graph
78	Measure avg. Pilot training Length	04Jul12	63.89	Graph	Graph	Graph	Graph	
79	document test findings	23Aug12	0	Graph	Graph	Graph	Graph	

I hope that gives you a taste at least of how we can use Monte Carlo simulation to improve our expectations about our projects and to modify our plans to make them more achievable.

"The really nice thing about not planning is that failure comes as a complete surprise and is not preceded by long periods of worry and depression."
 Sir John Harvey-Jones

From our web site you can download:

- The spreadsheet used in this presentation.
- The dice simulation used in this presentation.
- The Full Monte software for 30-day free trial.

The address is www.barbecana.com.