Standard Deviation

John Owen, *President*
Barbecana, Inc.

Introduction

Here is a typical Probability Distribution Histogram produced by a Monte Carlo Schedule Risk Analysis:

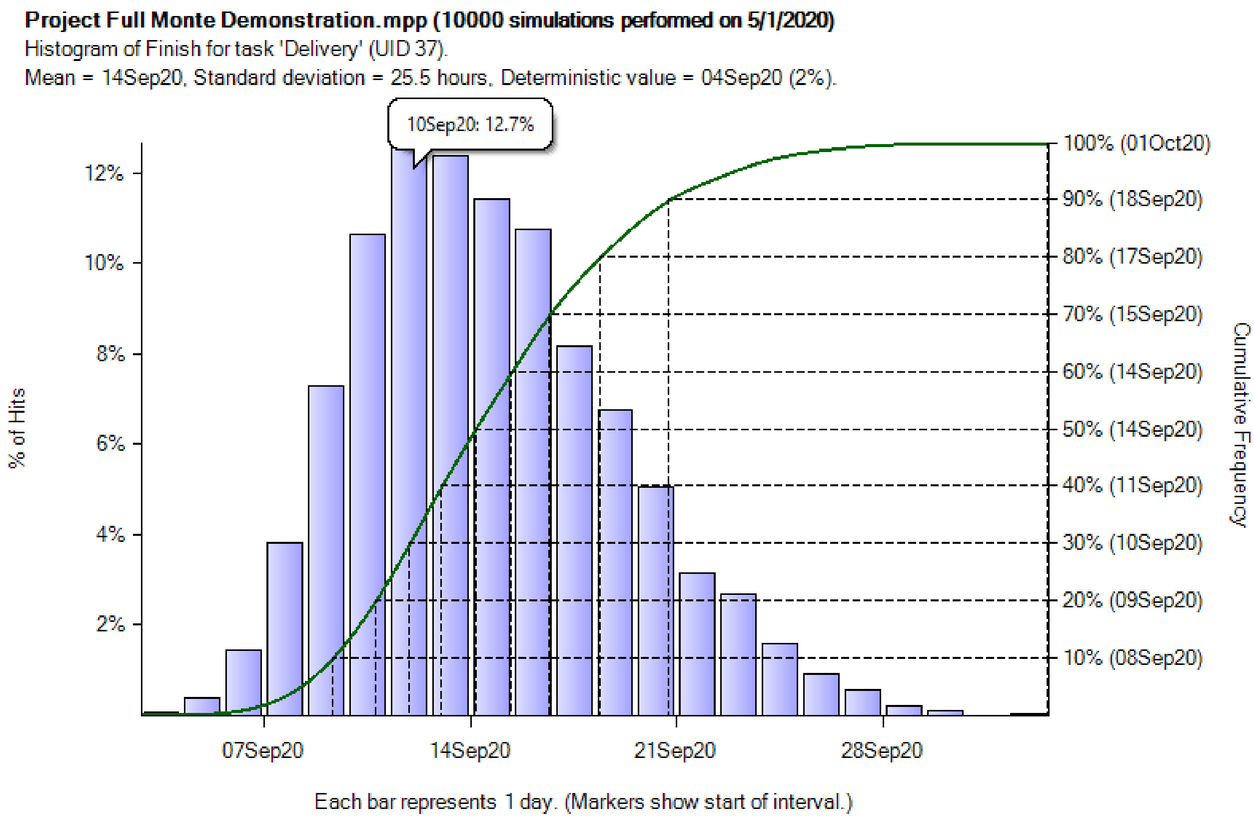

The bars represent the percentage of the simulations that finished in each reported period.

The **Standard Deviation** is shown in the chart header as 25.5 Hours.

The **Mean** finish is 14Sep20. This is an arithmetic average of the sampled finish dates and is usually close to, but not necessarily the same as, the 50% cumulative probability on the right Y-axis. The mean is the central point around which standard deviation is measured.

The **Mode**, the finish date that occurred most often during the simulations, is represented by the tallest bar and was 10Sep20. 12.7% of the simulations finished on that day.

As usual, our primary interest on this chart are the higher confidence dates at around 80% to 90% on the right Y-Axis. In this example, 80% of the simulations projected a finish on or before 17Sep20.

What does this mean? How is it useful?

Standard Deviation

*Note: If you have no interest in the calculation of Standard Deviation, you can skip to the section '**What use is the Standard Deviation**' below.*

The Oxford Dictionary defines Standard Deviation as: "a quantity calculated to indicate the extent of deviation for a group as a whole."

A simpler definition would be: "a measure of how spread out values are".

Technically the Standard Deviation is the Square Root of the Variance of the values where Variance is the Average of the Squared differences from the Mean.

The formula for Standard Deviation is:

$$\sigma = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i - \mu)^2}$$

As usual with Statistics, there are some assumptions that we make, the most important of which is that the probability distribution for which we are reporting Standard Deviation, is approximately normally distributed. This is a reasonable assumption because the **Central Limit Theorem** states that "given a sufficiently large sample size, the sampling distribution of the mean for a variable will approximate a normal distribution regardless of that variable's distribution in the population".

What this means in practice is that, even if you apply a Triangular distribution to all your uncertainty estimates, the histogram for completion will approximate a Normal distribution.

Note: The example histogram in the introduction above was obtained from a small project, where all tasks had a Triangular probability distribution applied. It demonstrates that the resulting finish distribution appears more 'Normal' than 'Triangular', albeit somewhat skewed.

For an approximately Normal distribution, the following statements can be made:

68.2% of the distribution values fall within one standard deviation of the mean
95.4% of the distribution values fall within two standard deviations of the mean
99.7% of the distribution values fall within three standard deviations of the mean.

This means that 99.7% of the simulated results are expected to fall within a range of **six** times (+/- 3 times) the standard deviation.

What use is the Standard Deviation?

For our purposes, modelling the likely range of dates that a project (or deliverable) will complete, the Standard Deviation is simply used to judge the 'reasonableness' of the result.

Consider the histogram above. The standard deviation was shown as 25.5 Hours.

Six times 25.5 hours is 153 hours. Divide this by the number of working hours per day (8) and we get 19.125 Days.

So, we would expect 99.7% of the simulated results to fall within just over 19 working days. This means we would expect, rounding up, 20 bars to be displayed on the graph (the chart is scaled such that each bar represents one day).

If we count the actual bars, we can see there are 22 days included on the chart. The extra bars are accounted for by the remaining 0.3% (100% - 99.7%) of the simulations and, in this case, by the fact that the distribution is somewhat skewed and therefore not as close an approximation to a Normal distribution as would be ideal. As with all statistics, the larger the sample size, the better the result.

So, the question is: "Is a range of just over 19 days in the expected result reasonable?"

That all depends on how far in the future we are predicting a range of 19 days and the riskiness of the project. The sample project has an overall remaining duration of 70 days. This is roughly three working months in the future while the range is just under 1 working month (based on 22 working days / month).

If the project involved some new product development (greater uncertainty) than this might be a reasonable range.

However, if the delivery was 5 years in the future, then a range of roughly one working month in 5 years' time is probably unrealistic (too small) and uncertainty assumptions and critical path driving the delivery should be examined/reevaluated.

The Effect of Task Calendars

Standard Deviation is based on the working time periods of the task calendar. For example, for a 20-day task with +/- 25% uncertainty the standard deviation will always be approximately 16 hours regardless of the calendar used. Here are three histograms of a 20-day task using different working period calendars.

Example based on a 5-day, 8-hour calendar.

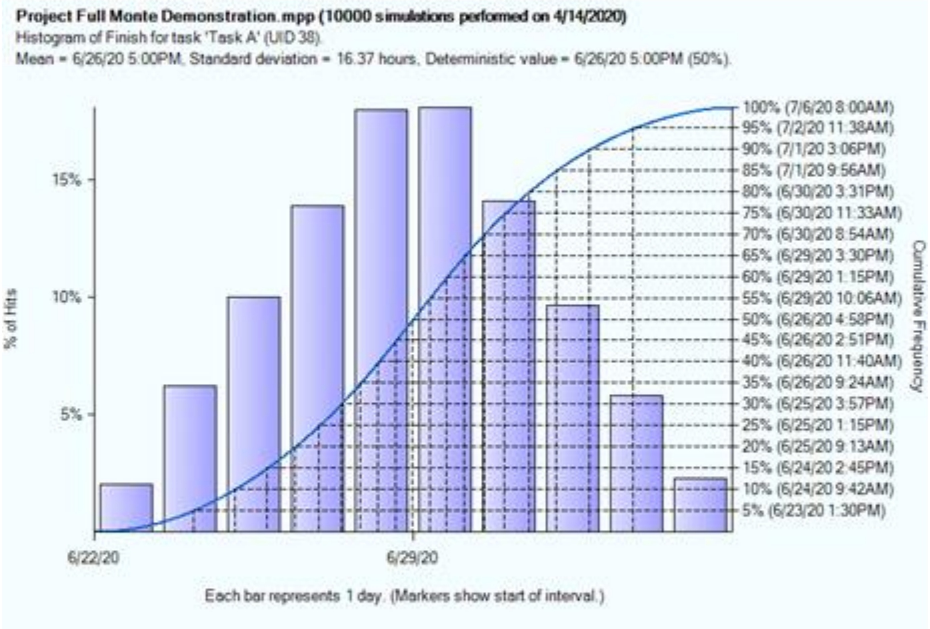

Example based on a 7-day, 8-hour calendar

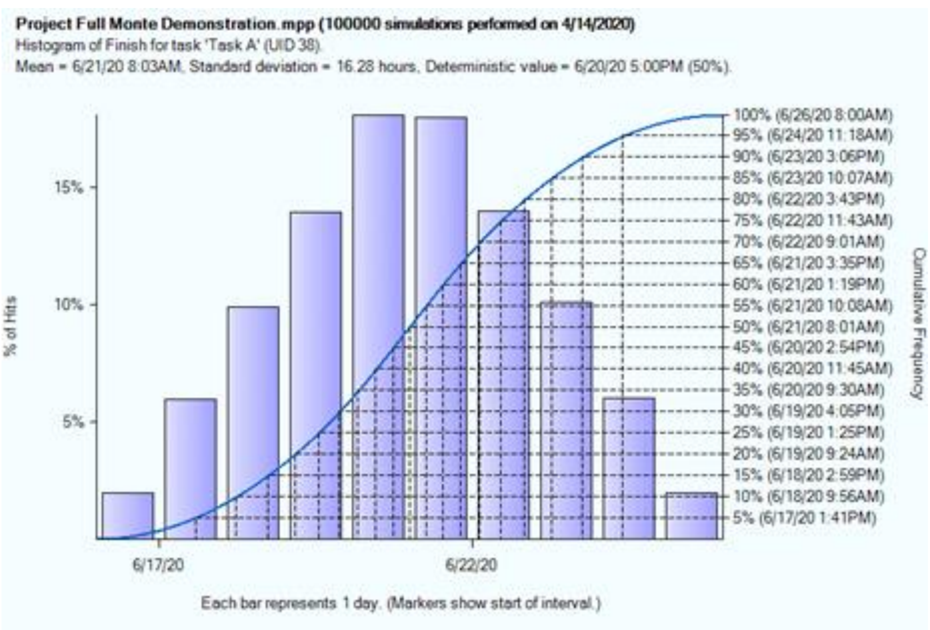

Example based on a 24/7 calendar

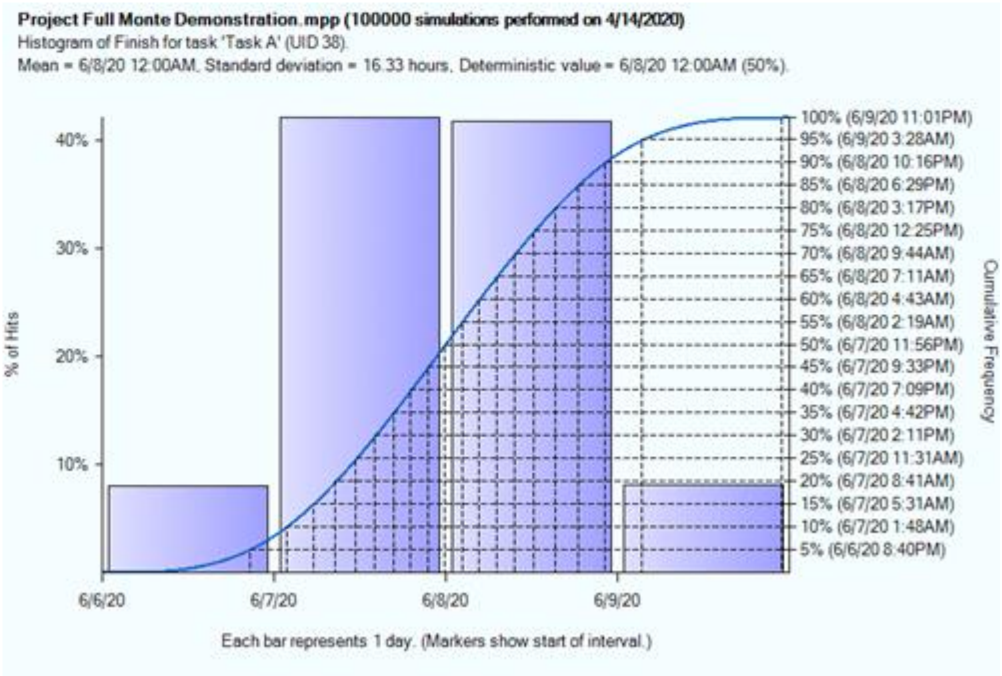

For a standard deviation of 16.33 hours, then 99.7% of the simulations should finish in +/- 3 * 16.33 hours. In this last example this means that 99.7% of the simulations finished in a range of 97.98 hours. We can validate this by seeing the earliest bar started at 6/6/20 12:00am and the latest bar finished at 6/9/20 11:01pm. That is roughly 4 complete days. 4*24 = 96 hours which is close to the estimated range of 97.98 hours.

Summary

Six times the Standard Deviation reported by Full Monte is approximately the range of results obtained during the simulations for the selected task or milestone. This can be used to assess the reasonableness of the range, given how far in the future the predicted finish is.

There is no fixed guidance as to what is reasonable since it must depend on your assessment of the uncertainty or risk in the estimated durations based on the type of project.

For a project based on work that has been performed before, there will be less uncertainty and therefore a smaller standard deviation may be reasonable. For new product development, we would expect greater uncertainty in the task duration estimates and therefore a larger standard deviation.

The standard deviation is also likely to get larger the further into the future we are trying to predict.