Vol. 8, No. 3, May 2008, pp. 139–148
ISSN 1532-0545 | 08 | 0803 | 0139

INFORMS Transactions on Education

INFORMS

DOI 10.1287/ited.1080.0013

Teaching Note

Teaching Project Simulation in Excel Using PERT-*Beta* Distributions

Ron Davis
College of Business, San Jose State University, San Jose, California 95192,
davis_r@cob.sjsu.edu, ron@mathproservices.com

This paper presents the methodology for computing the correct general formulas for the PERT-*beta* distribution, and how they are used to carry out stochastic project duration simulations using the built-in tools available in Excel. A comparison with results obtained using the Excel add-ins Crystal Ball, @Risk, RiskSolver, and PopTools is included. Slightly different parameterizations are used by each application, so various forms of the formulas are shown for each case.

Key words: PERT; simulation; beta distribution; beta shape parameters; PERT-beta distributions; project analysis; project duration
History: Received: December 2006; accepted: May 2008. This paper was with the author 11 months for 3 revisions.

1. PERT-*Beta* Distributions for Monte Carlo Project Simulation

Although the beta distribution has received much attention in relation to the PERT methodology since the appearance of the Malcolm et al. (1959) paper that proposed it, there has been a great deal of confusion and misunderstanding in the literature about how to carry out project simulations using the beta distribution based on the PERT paradigm. One factor contributing to this situation is that the creators of the method, and those presenting subsequent clarifications of the PERT rationale such as Clarke (1962) and Littlefield and Randolph (1987), failed to explain exactly how a beta distribution is to be selected for each activity of a project network. Another factor is the early criticism levied against the method (see "Picking on Pert" by Grubbs 1962) that falsely concluded that there were only three beta distributions meeting the PERT assumptions and hence the methodology was simply not well founded. Consequently, the correct PERT-*beta* shape parameter formulas on a general $[a, b]$ interval have not previously been clearly presented in the OR/MS literature.

In order to rectify this unfortunate situation, we proceed to the derivation of the correct beta shape parameter formulas to use when doing a PERT simulation. The traditional PERT formulas give mean and variance statistics for all activity durations, whereas the beta functions are generally expressed in terms of two shape parameters (which are referred to as alpha (α) and beta (β) in the Excel environment). As shown below there are closed form transformations between these two sets of parameters.

For a beta distribution defined on the interval $[a, b]$ with the parameters $[\alpha, \beta, a, b]$ one has

$$\text{MEAN:}\quad \mu = a + (b-a) * \left(\frac{\alpha}{\alpha+\beta}\right) \tag{1}$$

$$\text{VARIANCE:}\quad \sigma^2 = \left(\frac{\alpha}{\alpha+\beta}\right)\left(\frac{\beta}{\alpha+\beta}\right)\left(\frac{(b-a)^2}{\alpha+\beta+1}\right) \tag{2}$$

Solving for alpha and beta in terms of μ and σ^2 is facilitated by noting that the mean value formula implies that $(\alpha/(\alpha+\beta)) = ((\mu-a)/(b-a))$ and $(\beta/(\alpha+\beta)) = ((b-\mu)/(b-a))$. Putting these into the variance formula enables one to compute a quantity representing alpha plus beta:

$$(\alpha+\beta) = \left(\frac{(\mu-a)(b-\mu)}{\sigma^2}\right) - 1 \tag{3}$$

This is then split into the two parts alpha and beta according to

$$\alpha = \left(\frac{\mu-a}{b-a}\right)(\alpha+\beta) \quad \text{or}$$
$$\alpha = \left(\frac{\mu-a}{b-a}\right)\left[\left(\frac{(\mu-a)(b-\mu)}{\sigma^2}\right) - 1\right] \tag{4}$$

$$\beta = \left(\frac{b-\mu}{b-a}\right)(\alpha+\beta) = \alpha\left(\frac{b-\mu}{\mu-a}\right) = (\alpha+\beta)-\alpha \quad \text{or}$$
$$\beta = \left(\frac{b-\mu}{b-a}\right)\left[\left(\frac{(\mu-a)(b-\mu)}{\sigma^2}\right)-1\right]. \tag{5}$$

Equations (3), (4), and (5) are generalizations of those given for parameter estimation by the method of matched moments in the online Engineering Statistics Handbook (see Davis 2008).

Letting $[a, m, b]$ be the three PERT times assumed as given data, the usual PERT formulas give PERT − Mean $= \mu = (a + 4m + b)/6$ and PERT − Variance $= \sigma^2 = (b-a)^2/36$. When these mean and variance formulas are substituted into formulas (3), (4), and (5) above, one gets the unique beta distribution parameters for each project activity, which indeed has the PERT mean and PERT variance as its own mean and variance. Rewriting in terms of $[a, m, b]$, here is what you get after doing the algebra (for the record, in the classroom the unprimed equations are simpler):

$$(\alpha+\beta) = 4 + 16\left(\frac{(m-a)(b-m)}{(b-a)^2}\right) \tag{3′}$$

$$\alpha = \left(\frac{b+4m-5a}{6(b-a)}\right)(\alpha+\beta) \quad \text{or}$$
$$\alpha = \left(\frac{2(b+4m-5a)}{3(b-a)}\right)\left[1+4\left(\frac{(m-a)(b-m)}{(b-a)^2}\right)\right] \tag{4′}$$

$$\beta = \left(\frac{5b-4m-a}{6(b-a)}\right)(\alpha+\beta)$$
$$= \left(\frac{5b-4m-a}{b+4m-5a}\right)\alpha = (\alpha+\beta)-\alpha \quad \text{or} \tag{5′}$$
$$\beta = \left(\frac{2(5b-4m-a)}{3(b-a)}\right)\left[1+4\left(\frac{(m-a)(b-m)}{(b-a)^2}\right)\right]$$

We refer to beta distributions defined in this fashion as PERT-*beta distributions* (see Regnier 2005) because they are beta distributions that exhibit means and variances as specified by the PERT mean and variance formulas. This is an infinite but proper subset of the beta family, because it turns out that from (3′) one obtains $4 \le \alpha+\beta \le 8$ and the maximum value is achieved only for the symmetric case in which $m = (a+b)/2$ and therefore $\alpha = \beta = 4$. A pictorial representation of the PERT-beta family is shown in Figure 1, which was obtained by letting m vary from 0 to 10 with $a = 0$ and $b = 10$. The points where m assumes integer values from 1 to 9 are labeled for comparison with Figure 2. Points above the 45 degree line are skewed right ($\beta > \alpha$) and those below are skewed left ($\alpha > \beta$). The symmetric case in which $\alpha = \beta = 4$ is on the 45 degree line.

Figure 2 shows various density function shapes that occur as m varies from 1 to 9 when a is 0 and b is 10.

Figure 1 **Locus of PERT-Beta Shape Parameters Mode m Varies from 1 to 9 in** $[a = 0, m, b = 10]$

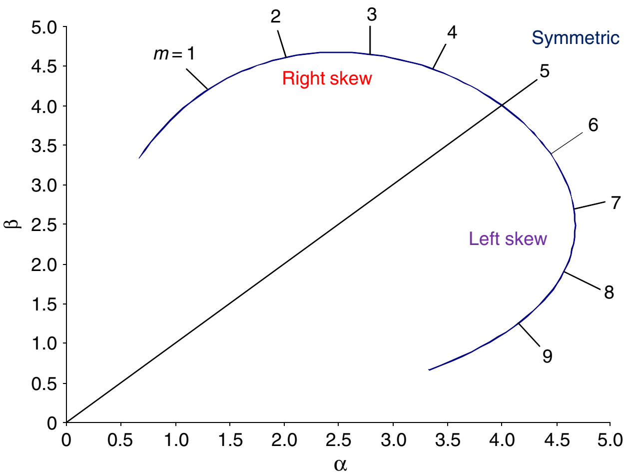

One finds that $\alpha+\beta$ is maximized with value 8 when the distribution is symmetric ($m = 5$) and the sum decreases toward four as m moves farther away from the midpoint of the interval. A sum of 4 results when $m = 0$ or $m = 10$. A lateral shift, or a rescaling of the width of the interval, or a combination of both, have no effect on α, β, or $\alpha+\beta$. Any such affine transformation of a beta distribution leaves its shape, and its shape parameters, unchanged. The family is seen to consist only of unimodal distributions that range from extremely skewed to the left to extremely skewed to the right. Between the extremes there are a continuum of shapes with varying amounts of skew, and a symmetric case in the middle when $m = 5$ and $\alpha = \beta = 4$. The means vary from 10(1/6) to 10(5/6) in this example. The family is therefore seen to be quite flexible and appropriate for use in the PERT context (see Figure 2).

With these shape parameter formulas (4 and 5, or 4′ and 5′), one can generate random activity durations

Figure 2 **PERT-Beta Densities for** $m = 1, 2, 3, 4, 5, 6, 7, 8, 9$ **in** $[a = 0, m, b = 10]$

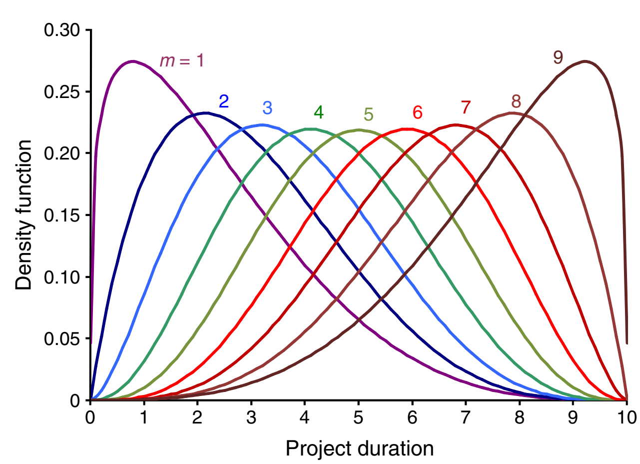

Table 1 Example Project Data for Simulation Analysis

Activities	Pred	ai	mi	bi	Mean	Variance	Alpha	Beta	Min	Max
A		3	6	9	6	1	4	4	3	9
B		2	5	6	4.6667	0.4444	4.6667	2.3333	2	6
C	B	2	3	7	3.5	0.6944	1.968	4.592	2	7
D	A, C	1	3	3	2.6667	0.1111	3.3333	0.6667	1	3
E	D	0	7	8	6	1.7778	4.3125	1.4375	0	8
F	C	1	2	10	3.1667	2.25	1.3434	4.2369	1	10
G	F	3	4	12	5.1667	2.25	1.3434	4.2369	3	12
H	F	1	2	15	4	5.4444	1.0845	3.9767	1	15
I	B	5	10	30	12.5	17.3611	1.968	4.592	5	30
J	E, G, H	1	3	4	2.8333	0.25	4.6173	2.9383	1	4
EOP	I, J	0	0	0						

using the standard uniform random number generator RAND( ) and the percentile point function call = BETAINV(RAND( ), α, β, a, b) in EXCEL (variations for the other add-ins will be shown in §4). The simulation results for each activity will have statistics that converge to the theoretical (PERT) parameter values as the sample size becomes larger and larger.

With the precise specification of the beta distribution to use for each activity given above, the Monte Carlo procedure for analyzing the distribution of project duration given PERT-*beta* distributed activity times can be specified in terms of a combination of the critical path method (CPM) implementation augmented with the appropriate function calls for generating random beta distributed activity durations having the desired mean and variance properties. The Data Table command can then be used for generating and tabulating multiple simulation trials that can be analyzed statistically using the Data Analysis ToolPak after the fact.

Heretofore the appropriate beta shape parameter formulas that make this possible ((3), (4), and (5)) have not been presented in OR/MS textbooks. One hopes that future textbook treatments will show these formulas and encourage students to use them. They should also be offered in the commercial software products for OR/MS applications, such as from Palisade, Frontline Systems, and Decisioneering.

2. Example Simulation Results Using PERT-*Beta* Distributions

To see how this spreadsheet implementation works in practice we have simulated a modification of the example presented in Ragsdale (2003) using 64,000 trials in order to get a quite accurate result. The modifications involved providing min and max times for each activity as required by PERT, and increasing the time estimate data for a few project activities in order to make several different paths have very similar mean path durations. This implies that during the simulation there is meaningful "competition" between them for longest path. The modified project data, associated precedence structure, and secondary parameters are shown in Table 1.

The Mean and Variance columns are obtained using the standard PERT formulas, and the alpha and beta shape parameters are obtained using the formulas (4) and (5), applied once in each row, for each activity. The EOP row is for the End of Project milestone node that denotes the end of the project when there would otherwise be multiple loose ends. The two results saved during the simulation are computed in this row (see Davis 2006).

In this case, using the CPM with mean durations, the PERT project critical path is found to be BCDEJ with mean duration 19.67. It has a critical path variance of 3.28 and a critical path standard deviation of 1.81 days. Under the PERT approximation, these values are taken as applicable to the total project duration, whereas in fact they pertain only to the BCDEJ critical path.

Only activity A is symmetric (with alpha and beta equal to 4) whereas all the other activities are skewed right or left with various degrees of skewness. The last four columns of the table are the beta distribution parameters for the activity durations given in the sequence that they are expected in the EXCEL built-in functions, BETADIST and BETAINV. Using the RAND( ) function for a uniformly distributed cumulative probability value (between zero and one) as the first argument in the BETAINV function call, one can then generate beta distributed activity durations for each of the activities and parameter sets listed above with the function call

$$= \text{BETAINV}(\text{RAND}(\,),\ \text{alpha},\ \text{beta},\ a,\ b).$$

Application of the standard forward and backward pass recursion formulas (given in Davis 2006) give rise to a set of additional columns for which simulation results can be tabulated using the Data Table command. For sample random activity duration times, a CPM table is shown as Table 2. The critical path in this case is seen to be ADEJ, and the duration

Table 2 CPM Computations Based on Randomized Activity Durations

Activity	RANDDur	ES	EF	LF	LS	Slack	Critical?
A	7.4681	0	7.4681	0	7.4681	0	A
B	3.7166	0	3.7166	0.4965	4.2131	0.4965	
C	2.0266	3.7166	5.7432	4.2131	6.2397	0.4965	
D	2.5291	7.4681	9.9971	7.4681	9.9971	0	D
E	6.0266	9.9971	16.0237	9.9971	16.0237	0	E
F	2.9149	5.7432	8.6581	6.2397	9.1546	0.4965	
G	6.8691	8.6581	15.5272	9.1546	16.0237	0.4965	
H	4.3728	8.6581	13.0310	11.6509	16.0237	2.9928	
I	12.5475	3.7166	16.2641	6.9681	19.5156	3.2515	
J	3.4919	16.0237	19.5156	16.0237	19.5156	0	J
EOP	0	19.5156	19.5156	19.5156	19.5156	0	ADEJ

Table 3 Descriptive Statistics for PERT-*Beta* Simulation

Project duration stats		Continued	
Mean	21.5412		
Standard error	0.0093	Range	20.3480
Median	21.2630	Minimum	13.8212
Mode	20.3930	Maximum	34.1692
Standard deviation	2.3445	Sum	1,378,637.065
Sample variance	5.4965	Count	64,000
Kurtosis	0.9458	Largest(16,000)	22.8261
Skewness	0.7070	Smallest(16,000)	19.9628

is at the base of the EF column. These two highlighted results in the EOP row will change from one iteration to the next as the random activity times change, and will be saved in two columns of a Data Table during the simulation process so as to enable statistical analysis of the results.

For this study, a Monte Carlo simulation in Excel was run having 64,000 simulation trials so that all mean, variance and skew estimates would be quite accurate. The Data Analysis Toolpak was then invoked to get the Descriptive Statistics Report and the Histogram Chart shown in Table 3 and Figure 3.

Due to "competition" between the paths, simulated mean duration is almost two days longer than the PERT project mean result, and the simulation variance is rather larger than the PERT project variance result. The skewness is significantly to the right.

By recording the critical path name computed in the lower right corner of the CPM table along with its duration from the bottom of the EF column, one

Figure 3 Histogram of Project Duration. Mean = 21.5412; Standard Deviation = 2.3445

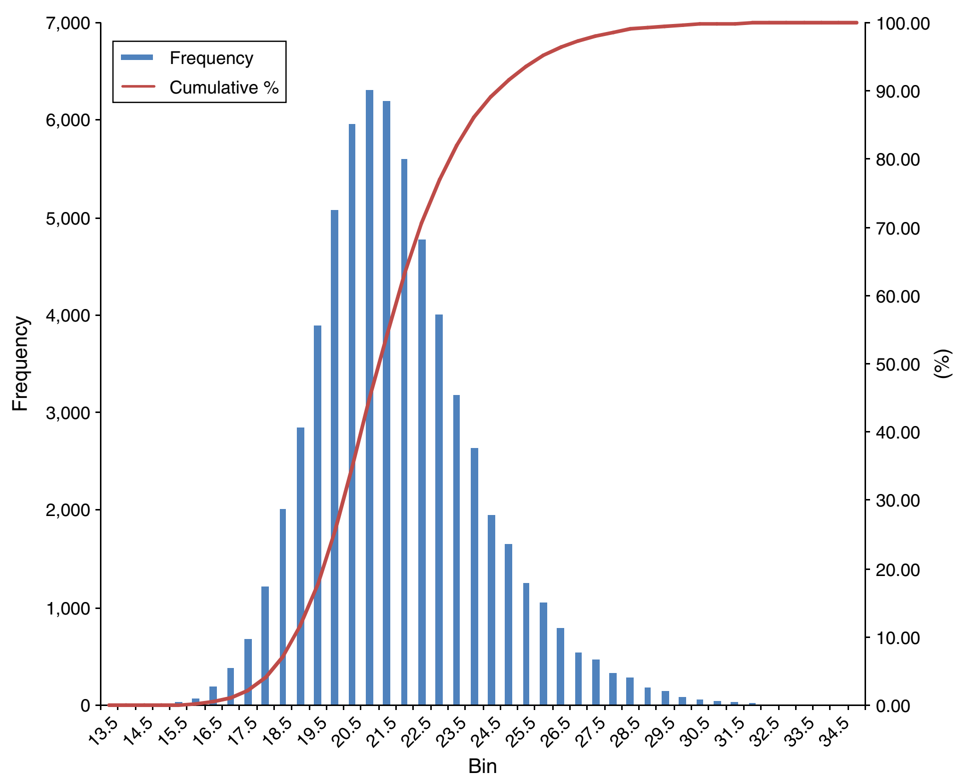

Table 4 **Path Probabilities**

Path	RelFreq (%)
A-D-E-J	2.8766
B-C-D-E-J	34.5078
B-C-F-G-J	26.1500
B-C-F-H-J	15.8734
B-I	20.5922

Table 5 **Percentile Comparisons**

Percentile	90%	95%	99%
PERT normal	21.9869	22.6446	23.8784
PERT-*beta* simulation	24.6498	25.8890	28.3841

can construct a relative frequency tabulation for the various paths that are critical on the various simulation trials. Table 4 shows the results obtained for this experiment. As can be seen, all paths are critical for some random samples, but they are not equally likely to be critical. Interestingly, the PERT critical path BCDEJ occurs only about 1/3 of the time. Hence something other than BCDEJ is critical about 2/3 of the time.

Finally, the first quartile of the simulation (19.96) is larger than the PERT mean duration estimate, showing the significant extent to which the project duration distribution has been "shifted to the right" by the durations of other paths.

The distribution for total project duration, based on the interaction of all five project paths, assumes a shape that looks very much like a beta distribution with a substantial right skew. The right skew derives from the skew of the individual activity durations and the effect of the maximum value computations for the early start (ES) times throughout the project. The path probability table shows that each of the paths is critical a significant portion of the time, but some paths are much more likely than others. The right skew of the simulation result implies that percentiles for percents close to 100 will lie further to the right (i.e. be larger) than for the symmetric PERT approximation. To emphasize this point, consider the comparison in Table 5. For this example, the 90th percentile of the simulation result is larger than the 99th percentile of the PERT normal approximation.

Another way of viewing the comparison is by means of the plots of cumulative percent from the simulation result and cumulative probability under the normal PERT approximation. This CDF comparison is shown in Figure 4.

Over the range where there is a visible difference between the curves, the normal PERT approximation lies above and to the left of the simulation CDF. This means that the normal PERT approximation consistently underestimates the time to a given percentile, and over estimates the probability for any given time. In other words, it gives an optimistic estimate, especially above the means where questions are most often asked.

3. Fitting a Beta Distribution to Project Duration Statistics

Over the 15 years or so that I have taught PERT simulation, I have noticed that the project duration distribution almost always has a noticeable amount of

Figure 4 **CDF Comparison: PERT Normal vs. PERT-Beta Simulation**

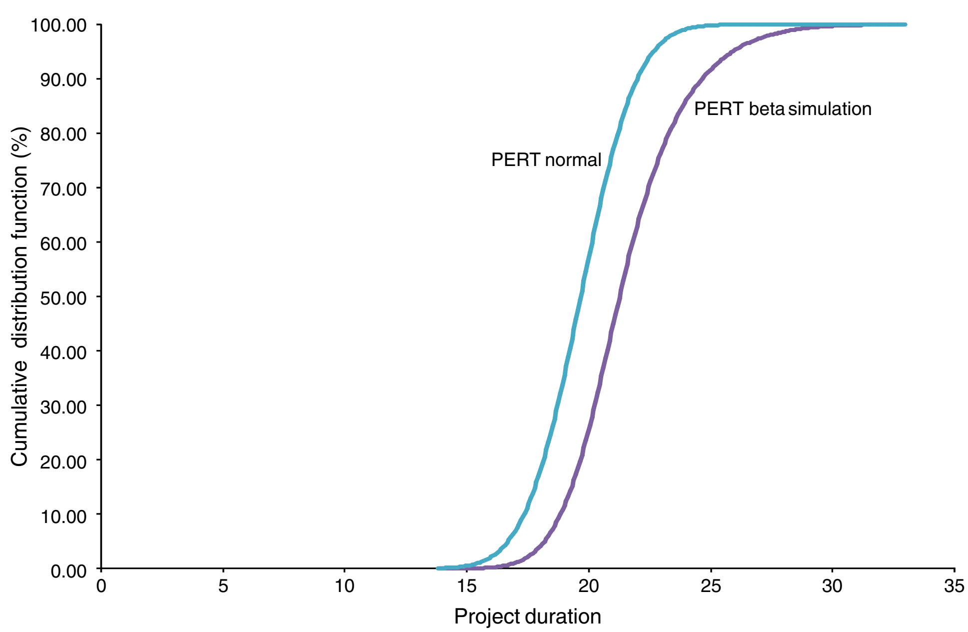

Table 6 **Beta-fit for Project Duration CDF**

Beta-fit	Parameters		
alpha_fit	6.3498	min_fit	13.8212
beta_fit	10.3867	max_fit	34.1692

skew, one way or the other. So as an extra credit assignment, I ask students to fit a beta distribution to their simulation results. Using the method of moments, they can use the mean and variance results and the min and max times observed during the simulation as a basis for computing a preliminary beta fit for the simulation result. Quite simply, it is the beta on the given range that matches the first two moment statistics given in the Descriptive Statistics Report shown in Table 3. The parameter estimates for the fitted beta obtained in this way are shown in Table 6.

The CDF of this fitted beta and the cumulative percentage of the simulation result are then plotted on the same axes to get a measure of how well the fitted beta matches the results from the simulation across the entire range of project durations observed. As is seen in Figure 5, the fit is very close, with the fitted beta points lying very close to the cumulative percentage points across the entire range of the distribution. This leads to the conjecture that perhaps *all* of the early finish time distributions are closely approximated by beta distributions. For the present example, that turns out to be the case, although we do not take the space to show all the CDF plots that indicate this is so.

The "gap" between these two curves can be reduced further by letting both min and max parameters go free and optimizing all four beta distribution parameters to minimize the root mean squared (RMS) deviation of the fitted fractiles in comparison with the fractile locations of the simulation result on some grid of fractiles (1% to 99% was used in this study with increments of 1%). This is done using the nonlinear optimizer provided by the SOLVER add-in subject to the constraints that the mean and standard deviation of the beta-fit match the statistics for the simulation results. The fitting procedure could therefore be called moment-constrained fractile fitting. The result of this refinement is shown in Figure 6.

Although not perfect, the beta distribution fit nevertheless matches the simulation result to an error tolerance that would be satisfactory for most practical applications.

The significance of this result is two fold. First it means that probability and percentile questions posed about project duration can be readily answered using the beta-fit for the PERT-*beta* simulation results, rather than using the histogram of the simulation results directly. In particular, the probability of not completing the project by any given target time "t" is estimated by

$$1 - \text{BETADIST}(t, \text{alpha\_fit}, \text{beta\_fit}, \text{min\_fit}, \text{max\_fit}),$$

and the time required to achieve a probability "p" of timely completion is estimated by

$$\text{BETAINV}(p, \text{alpha\_fit}, \text{beta\_fit}, \text{min\_fit}, \text{max\_fit}).$$

Secondly, and more importantly, it suggests the idea of developing beta distributions for *all* of the early finish (EF) times during the forward pass in an analytical way, so that the Monte Carlo simulation process can be omitted altogether. This approach to project duration analysis is given in a separate paper to be presented elsewhere.

4. Parametric Variations for Add-Ins

The simulation example reported here was repeated four more times with the Crystal Ball, @Risk, Risk-Solver, and PopTools add-ins. The internal consistency between the results is a welcome indication that

Figure 5 **CDF Comparison: PERT-Beta Simulation vs. Simple Beta-fit**

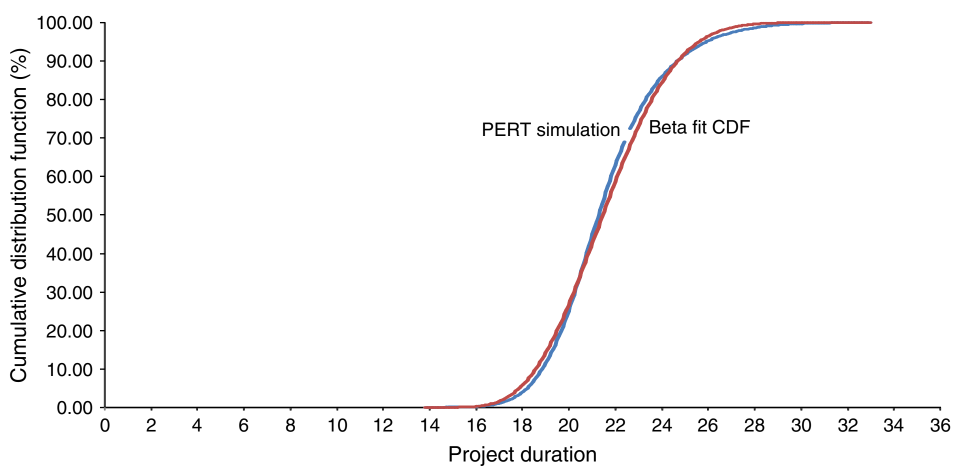

Figure 6 CDF Comparison: PERT-Beta Simulation vs. RMS Beta-fit

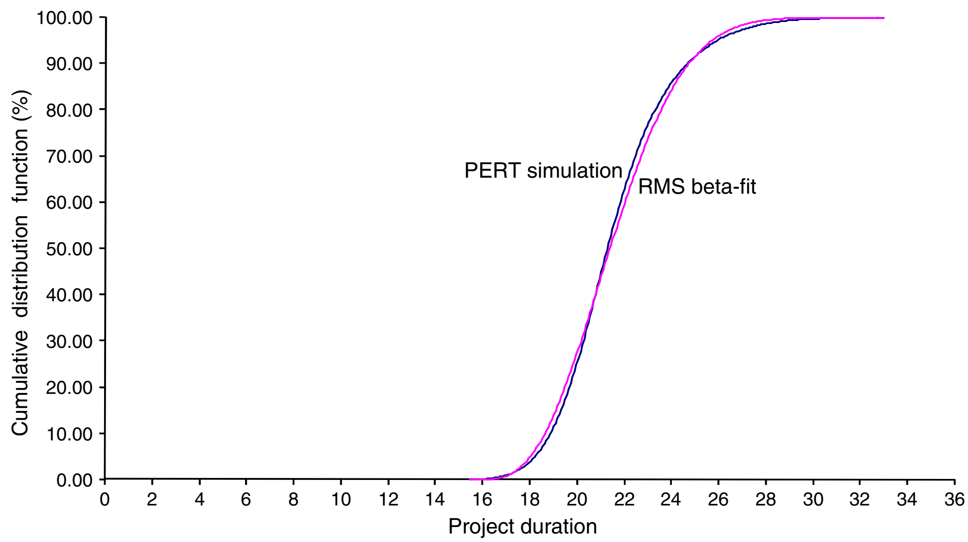

all implementations have been carried out correctly. However, the implementation details vary slightly due to differences in the way that the beta distribution has been scaled and parameterized. For clarity, we spell out these implementation differences here in terms of activity C from Table 1 having $[a, m, b] = [2, 3, 7]$ and (alpha, beta) = (1.968, 4.592).

Crystal Ball (Decisioneering)
First, the Crystal Ball implementation assumes that the min location parameter "a" is zero, and uses a "scale" parameter to represent our $(b - a)$ interval width. Hence the output of their random beta variate must by "shifted" to the right by "a" to get the desired activity time. This is done by simple addition of a to the random beta variate. So for activity C, scale would be $7 - 2 = 5$ and the random activity time would be $= 2 +$ CB.beta(1.968, 4.592, 5).

@RISK and RISKOptimizer (Palisade)
Second, the @Risk implementation for the Beta distribution has several different forms (RISKOptimizer uses the distributions from @RISK so the same comments apply to both). A beta on the unit interval [0, 1] is provided by the RiskBeta(alpha, beta) function. The beta on a general [A min, B max] interval is provided by the RiskBetaGeneral(alpha, beta, A min, B max) function. This latter is the one to use, with alpha and beta derived from the PERT-beta formulas given above. It uses all four of the parameters used in the Excel implementation, in the same order, and for the example at hand would appear as = RiskBetaGeneral(1.968, 4.592, 2, 7).

In addition, there is a RiskBetaSubj distribution based on mean, mode, min, and max. They also have a RiskPert distribution, but it is not for the PERT-*beta* distribution developed here. It is based on matching mean and mode instead of mean and variance (see Vose 1996), and hence is not correctly named since it does not supply the correct variance. Their RiskPert distribution was not used in this study, as a result, and is not recommended for PERT simulation studies. (I will attempt to get the implementation of the RiskPert distribution changed so that it conforms to the PERT-beta formulas given here, but at the present time RiskBetaGeneral is the function to use to get true PERT-beta distributions.)

RiskSolver (Frontline Systems)
The RiskSolver package provides the same four options as the RiskOptimizer does, with slightly different names. For the unit interval, its PsiBeta(alpha1, alpha2) where alpha1 takes the place of alpha, and alpha2 takes the place of beta; the general $[a, b]$ interval is PsiBetaGen(alpha1, alpha2, a, b). They also supply a PsiBetaSubj(min, mode, mean, max) function for the betas that are specified in terms of mode and mean instead of mean and variance, and also a PsiPert(min, mode, max) function which matches PERT mean value, but not PERT variance. Hence, again, to obtain the PERT-beta distributions at the present time, one must skip over the PsiPert function and use PsiBetaGen with the arguments obtained using the PERT-beta formulas given here. Hence for activity C in our example one would use PsiBetaGen(1.968, 4.592, 2, 7).

One hopes that in time both the RiskPert and PsiPert functions will give PERT-beta distributions rather than the betaPERT distributions introduced by Vose (1996). Since Vose's distributions are derived by matching mean and mode instead of mean and variance, in my view they are an aberration away from the correct formulas. The variance of the triangular distribution INCREASES as M moves away from the middle towards the extremes, whereas the variance of the betaPERT DECREASES towards the extremes. The PERT-beta distributions have CONSTANT variance as M moves from one extreme to the other, and that is the way it should be, in my opinion, if one is going to do project simulation based on the PERT paradigm.

PopTools

The PopTools beta function is unusual in that it does not require the shape parameters customary in most other packages. It also assumes the interval of uncertainty is the unit interval $[0, 1]$. The arguments it does require are the mean and standard deviation of the distribution, given that it is on the $[0, 1]$ interval. Hence for the present application, the beta distribution for each activity must be mapped to the $[0, 1]$ interval first to get the corresponding mean and standard deviation to use as arguments to this function; then the output of the random number generator needs to be mapped back to the $[a, b]$ interval for use in the simulation. This is accomplished as follows. The arguments to the *d*BetaDev function are given by

$$\mu_i' = (\mu_i - a_i)/(b_i - a_i) \quad \text{and}$$
$$\sigma_i' = \sqrt{\frac{\sigma_i^2}{(b_i - a_i)^2}} = \frac{\sigma_i}{(b_i - a_i)} = 1/6. \tag{6}$$

The random activity time is then given by

$$a_i + (b_i - a_i) * d\text{BetaDev}(\mu_i', \sigma_i'). \tag{7}$$

For activity C, μ' would be $(3.5-2)/5$ or 0.3 and $\sigma' = (5/6)/5 = 1/6$. Hence the random time for this activity would be $= 2 + 5 * d\text{BetaDev}(0.3, 1/6)$. In fact, *all* PERT-beta distributions have a standard deviation of 1/6 when normalized to the unit interval, and that is the major part of what makes them PERT-beta (the rest of it being that $1/6 \le \mu' \le 5/6$).

5. Results Summary

A comparison of the simulation results obtained using the five beta functions included in this study is shown in Table 7. This consistency of results tends to reinforce the prior notion that all reported implementations are valid and equivalent.

Table 7 Simulation Results Comparison

				Percentiles		
Method	Mean	StDev	Skew	90%	95%	99%
Excel alone	21.5412	2.3445	0.7070	24.6498	25.8890	28.38411
Crystal Ball	21.5545	2.3605	0.7125	24.7012	25.9143	28.40154
@Risk	21.5531	2.3491	0.7236	24.6509	25.9446	28.49812
RiskSolver	21.5607	2.3450	0.7208	24.6513	25.9424	28.43702
PopTools	21.5612	2.3504	0.7045	24.7054	25.9377	28.45213

Table 8 RMS of Relative Errors after 64,000 Trials

Method	Error in means	Error in standard deviations
Excel (alone)	0.000627	0.000378
Crystal Ball	0.000400	0.000481
@Risk	0.000000147921	0.00000166663
RiskSolver	0.000471	0.000385
PopTools	0.000743	0.000417

Another consistency check was done by keeping statistics also on the individual activity durations which were selected randomly throughout the simulation. Theoretically, the mean of the simulated durations for an activity should approach the mean of the PERT-*beta* distribution being used for the activity (i.e. the PERT mean for the activity), and the sample variance and standard deviation of the simulated durations should approach the PERT variance and standard deviation of the activity. Since the intervals of uncertainty vary across the activities, the tabulation in Table 8 is in terms of "relative error" in which the indicated difference is divided by the associated interval width. An RMS of the relative errors was computed across the ten input quantities in each case.

The small size of the @Risk errors was due to the Latin Hypercube sampling procedure used, in comparison with the other methods based on simple Monte Carlo sampling. Since the standard deviation of all PERT-*beta* distributions on the unit interval is 1/6 and there were 64,000 samples taken, the relative standard error of the sample mean should be (1/6)/SQRT(64,000) according to theory, which is 0.0006588. Since there is only one RMS reported that is just slightly larger than this, all others being less, we again find that the computational results reinforce the prior notion that all implementations are correctly done.

Since the RiskPert distribution (supplied by @Risk) and the PsiPert distribution (supplied by RiskSolver) have standard deviations that deviate from the required $1/6 * (b - a)$, they were not used in this simulation experiment. In addition, I believe that they should not be used for PERT simulations until they have been changed to an implementation of the PERT-*beta* distribution. For the example reported here, we used the RiskBetaGeneral distribution (and PsiBetaGen distribution) with shape parameters determined by the PERT-*beta* formulas given in this paper, and we recommend that the reader do the same until the RiskPert and PsiPert implementations have been corrected.

6. Conclusions

The principal analytical contribution of this paper is to show exactly how to select the beta distribution for each activity in a project model, based on the PERT paradigm for activity time modeling. We have shown in fact that there is a UNIQUE beta that conforms to the Min, Max, PERT-mean, and PERT-variance parameters for each activity, and that the family of all such PERT-beta distributions is comprised of unimodal distributions that includes a symmetric case as well as distributions that are skewed right or skewed left with various degrees of skewness. Contrary to the false impression created in Grubbs (1962), the PERT-beta family is quite flexible and quite suitable for the purpose for which it was created. It is unclear why the relevant formulas have not been presented previously (in the general form given here), but it is also quite clear that the formulas presented here are correct. We hope that Palisade, Frontline Systems, and Decisioneering will soon provide PERT distributions based on these formulas rather than the ones they are now using.

My classroom use of the methodology presented in this paper has led to some other conclusions, as well. The first two of these have been noted by other authors earlier on (see Schonberger 1981) and are reiterated here for the sake of completeness.

1. The simple PERT approximation (without simulation) contains a systematic bias that makes it a "rough first cut" approximation that should not be relied upon in practice. It systematically underestimates the mean duration of the project, and systematically overestimates the probability of completion at any given time. Hence it gives a "rosy" or optimistic view of the situation that should not be taken literally. The degree to which the simulation CDF will be shifted to the right relative to the PERT normal CDF depends on the specific project structure and activity duration distributions, but the example results shown in Figure 4 and Table 5 are enough to see that the results from simulation may be significantly different than what is indicated by PERT without simulation.

2. The Central Limit Theorem, which is often invoked in relation to individual paths of a project network, is not applicable to the total project duration (except in two special cases, i.e. linear project structure or dominant critical path). Because of the intervention of the MAX function for computation of each and every EST, there are no long sums of random variables involved in the definition of the project duration, except as arguments of a MAX function. And the MAX of a set of random variables is usually skewed, even when the arguments are symmetric. Therefore (except in the two special cases) there is no a priori reason to expect that project duration distributions will tend to be either symmetric or normal in shape.

3. Simulations based on the beta distribution may be carried out in Excel without use of simulation add-ins. The results obtained with simulation add-ins, at least for those considered in this paper, given sufficiently large simulation trial counts, will be very much the same as what can be achieved in Excel without the add-ins. While there may be run time and convenience advantages for the commercial add-ins, it is not necessary for students to buy such software in order to do project duration simulation in the spreadsheet. Due to the many unfamiliar Excel details involved in the process, however, I find it extremely helpful to prepare demonstration videos using SnagIt that show students the procedures to follow. These can be converted to flash movies that students can view through the browser (see Davis 2007).

4. In practice, skewness to a degree clearly visible in the simulation histogram is the rule, not the exception. Moreover, project duration distributions are usually well described by fitted beta distributions, and are almost never symmetric. One can even use the fitted betas as a means of answering probability and percentile questions about project duration, rather than the simulation result itself. I sometimes use an expression "beta in | beta out" to describe this situation.

5. Schonberger ends his 1981 paper by saying "my suggestion is that students—future project managers and staff—should be schooled in the network simulation and project lateness phenomena, but that project managers may employ the insights gained without the need for actually conducting expensive simulations of large project networks." Obviously, times have changed. With the enormous growth in computer power and personal computer software sophistication, it is no longer advisable to omit the simulation process. Except for extremely large projects, it is no longer too expensive or too time consuming to actually carry out simulations on realistic project networks. We hope that reading this paper has enhanced your understanding of how to do this and what to expect from your results.

References

Clark, C. E. 1962. The PERT model for the distribution of an activity time. *Oper. Res.* **10**(3) 405–406.

Davis, R. E. 2006. A spreadsheet CPM implementation for PERT simulation. http://www.mathproservices.com/CPM-PERT/cpmspread.html (last accessed on May 31, 2008).

Davis, R. E. 2007. CPM/PERT Spreadsheet Implementation. http://www.mathproservices.com/cpmmenu.html (last accessed on May 31, 2008).

Engineering Statistics Handbook. 2008. *NIST/SEMATECH e-Handbook of Statistical Methods.* http://www.itl.nist.gov/div898/handbook/eda/section3/eda366h.htm (last accessed on May 31, 2008).

Grubbs, F. E. 1962. Attempts to validate certain PERT statistics or "picking on PERT." *Oper. Res.* **10**(6) 912–915.

Jensen, P. O. 2007. Operations management/industrial engineering: Computation: Project management: Beta distribution. http://www.me.utexas.edu/~jensen/ORMM/omie/computation/unit/project/beta.html (last accessed on May 31, 2008).

Littlefield, T. K. Jr., P. H. Randolph. 1987. An answer to Sasieni's question on PERT times. *Management Sci.* **33**(10) 1357–1359.

Malcolm, D. G., J. J. Roseboom, C. E. Clark, W. Fazar. 1959. Application of a technique for research and development program application. *Oper. Res.* **7**(5) 646–669.

Ragsdale, C. T. 2003. A new approach to implementing project networks in spreadsheets. *INFORMS Trans. Ed.* **3**(3) 76–85. http://ite.pubs.informs.org/Vol3No3/Ragsdale.

Regnier, E. 2005. Activity completion times in PERT and scheduling network simulation, Part II. *DRMI Newsletter* (April 8) 1, 4–9.

Schonberger, R. J. 1981. Why projects are "always" late: A rationale based on manual simulation of a PERT/CPM network. *Interfaces* **11**(5) 66–70.

Vose, D. 1996. *Quantitative Risk Analysis, A Guide to Monte Carlo Simulation Modeling*. John Wiley and Sons, New York.