Mixture Densities for Project Management Activity Times: A Robust Approach to PERT

February 4, 2007

Abstract

PERT is a widely utilized framework for project management. However, as a result of underlying assumptions about the activity times, the PERT formulas prescribe a light-tailed distribution with a constant variance conditional on the range. Given the pervasiveness of heavy-tailed phenomena in business contexts as well as inherently differing levels of uncertainty about different activities, there is a need for a more flexible distribution which allows for varying amounts of dispersion and greater likelihoods of more extreme tail-area events. In particular, we argue that the tail-area decay of an activity time distribution is a key factor which has been insufficiently considered previously. We provide a distribution which permits varying amounts of dispersion and greater likelihoods of more extreme tail-area events that is straightforward to implement with expert judgments. Moreover, the distribution can be integrated into the PERT framework such that the classic PERT results represent an important special case of the method presented here.

Keywords: Applied probability; Finite mixture; Beta rectangular distribution; Robust project management; Activity times

1 Introduction

The management of large-scale projects poses numerous challenges. These challenges have led to widespread use of project management techniques such as PERT (Program Evaluation and Review Technique). These project management techniques provide managers with a systematic quantitative framework for scheduling, planning and coordinating the many activities associated with the successful on-time completion of large complex projects. Hence, the goal of PERT and related techniques is to facilitate the management, coordination and control of the various activities involved in a project so that the project itself may be completed successfully. This is accomplished by structuring the activities into a network and examining the time requirements and precedence relationships associated with the activities.

One defining feature of PERT is that the activity times are taken to be stochastic. Thus, in order to facilitate a project's management, it is necessary to elicit activity time distributions for the activities comprising the project. As such, PERT (Malcolm et al., 1959) provides straightforward methods for eliciting activity times and for computing key quantities from the elicited judgments. In particular, four assumptions are made regarding activity time distributions in the context of PERT (Littlefield and Randolph, 1987) that lead to a particular form of beta distribution being utilized for the distribution of activity times. This particular beta distribution can be termed the PERT beta distribution. While visually the PERT beta distribution can take on several shapes, the dispersion of these distributions as measured by the variance is nonetheless taken to be a constant function of the range as a result of the PERT assumptions. This is a limiting characteristic because, in assessing real-world activity times, we need distributions that are capable of handling the potentially differing and substantial amount of uncertainty that may be present. In particular, there is a need to represent varying amounts of expert uncertainty through the dispersion measure.

It is also desirable to have a mechanism for allowing for heavy tails (i.e., greater likelihoods of outlying or extreme events) in activity time distributions given the prevalence of heavy-tailed phenomena in real-world applications. In finance it has long been appreciated that many seemingly anomalous results and incorrect inferences can be traced back to the use of light-tailed distributions such as the normal to model heavy-tailed financial returns (Mandelbrot, 1963; Fama, 1965). These leptokurtic or heavy-tailed returns are known as 'excess returns' and they may instead be t-distributed or follow other non-normal forms. Similar phenomena also occur in economics, thus leading to robust families of models for these situations (e.g., Engle, 1982). We also find these phenomena occurring in project management. For example, Grant et al. (2006) documented the appearance of outliers in project management and suggests that the probabilities of work proceeding more slowly than planned may be underestimated. Mitchell and Zmud (1999) documented an outlier project which experienced a four-year delay due to IT-related issues and a 5-year slippage in scheduled completion (see also Banker and Kauffman, 1991; Raffo and Kellner, 2000). Morgenshtern et al. (2006) empirically found that greater project uncertainty is associated with greater project durations and increasingly erroneous project duration estimations, implying again that project managers tend to give too little weight to the possibility of more extreme delays. This notion is corroborated by Atkinson et al. (2006) who indicate that uncertainty in project management is often played down and insufficiently considered in industry. It is interesting to note that while robust methods are used in many disciplines such as those mentioned above, their application to project management appears to have been almost completely overlooked. To date, no systematic treatment of the need for robust methods in project management has appeared save that of Steele and Huber (2004). Steele and Huber (2004) noted that many project management tools assume certain common distributional forms such as the normal. When the data do not happen to follow these forms, statistical techniques can perform poorly. Steele and Huber (2004) then focus on incorporating methodological robustness into project management using Tukey's methods of exploratory data analysis (EDA). We note here however that Steele and Huber (2004) did not make a direct connection to PERT.

In the statistical literature (e.g., McCullagh and Nelder, 1989), it is recognized that commonly utilized distributions may be unable to adequately represent real-world excess variability and over-occurrence of tail-area events. For example, in conducting Poisson regression it is common to encounter count data that are overdispersed relative to the Poisson. In these cases, continuing to use the typical Poisson formulation leads to problems: test statistics can become greatly inflated, causing large increases in Type I error rates such that incorrect inferences are likely. Here, more flexible distributions such as the negative binomial distribution are often used to correct these problems and represent overdispersed and outlying observations more accurately.

In the above cases, we see that we can often improve upon commonly used light-tailed distributions such as the normal and the Poisson by using more flexible distributions which nest the original distribution as a special limiting case[1]. Interestingly, the more flexible distributions mentioned above both can be conceptualized as mixture distributions. Both the t-distribution and the negative binomial result from specifying that an underlying parameter is itself uncertain and varies probabilistically[2]. Mixture distributions permit increased robustness for inferences as well as more flexible and accurate data fitting (Gelman et al., 2003). In this paper we derive a heavy-tailed alternative to the PERT beta using a mixture distribution. The use of this mixture assigns more weight to extremal tail-area events, and in using the mixture the project manager may assign more probability to the occurrence of outliers and extremal events, thus obtaining inferences that should be better protected against outlying events. In deriving this distribution, the elicitation of the distribution should be straightforward so that experts are able to provide the best possible judgments based on their experience. This is because more straightforward probabilistic judgment tasks are more likely to be performed accurately (Hogarth, 1975). The mixture distribution used here allows for both straightforward elicitation and the incorporation of robustness in PERT.

The contributions of this paper are as follows. We have first examined the need for robust methods in the context of PERT. We then provide a new distribution engineered specifically

[1]As the t-distribution's degrees-of-freedom parameter and the reciprocal of the negative binomial's overdispersion parameter go to infinity, the normal and the Poisson are recovered.

[2]Here coincidentally by using the gamma distribution for the underlying parameter.

for the context of project management and PERT. It permits more robust inferences with increased likelihood of extremal events and a flexible representation of uncertainty while accommodating PERT's classic beta distribution as a special case. Thus, it is useful in mitigating against overly optimistic conclusions regarding project outcomes which may be unwarranted if extreme outcomes are more probable. Finally, we also extend the literature by adopting the perspective that the engineering of distributions to have particular properties (here robustness properties) is relevant for the real-world project management context. This can be contrasted with the use of standard off-the-shelf distributions which may or may not have been specifically designed for a given context and accordingly may or may not possess all the properties desired. The outline of the paper is as follows. Section 2 reviews the use of the PERT beta distribution for project activity time estimates. In Section 3, we derive a heavy-tailed distribution for activity times. In Section 4, we examine the distribution in the context of the classic PERT formulas. In Section 5, we review elicitation for the distribution while in Section 6 we present an empirical application. We draw conclusions in Section 7.

2 Literature Review

The general characterization of the beta distribution having parameters $\alpha > 0$ and $\beta > 0$ is

$$p(y|\alpha,\beta,a,b) = \begin{cases} \dfrac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \dfrac{(y-a)^{\alpha-1}(b-y)^{\beta-1}}{(b-a)^{\alpha+\beta-1}} & \text{if } a \leq y \leq b, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

The distribution in (1) exists on the interval a to b and so is suitable for activities with arbitrary ranges $b-a$. Define $k = \alpha+\beta$. Then, its expectation, variance and mode are α/k, $\alpha\beta/(k^3+k^2)$ and $(\alpha-1)/(k-2)$ respectively.

The use of the beta distribution in the context of PERT can be traced back to the method's origin in Malcolm et al. (1959). Here the modal time m together with a and b can be combined to yield the standard PERT formulas for the expectation and variance of the activity time distribution, namely

$$E(y) = \frac{a+4m+b}{6}, \tag{2}$$

$$\mathrm{Var}(y) = \frac{(b-a)^2}{36}. \tag{3}$$

Littlefield and Randolph (1987) noted that the expressions in (2) and (3) depend on the following four assumptions: the activity time distribution is beta; experts can estimate a, b and

m well; the standard deviation is one-sixth that of the range; and a linear approximation-based result for the mean is acceptable. As for the third assumption, Clark (1962, p. 406), indicated that the derivation of this relationship was as follows: the standard "normal distribution truncated at ± 2.66 has its standard deviation equal to 1/6 the range"[3]. The question of what specific type of beta distributions is implied by (2) and (3) was examined by Sasieni (1986). In the case where $\alpha \neq \beta$, (2) implies a particular subfamily of beta distributions such that $k = \alpha + \beta = 6$. However, when $k = 6$ the relationship in (3) is no longer necessarily exact but becomes generally approximate. Littlefield and Randolph (1987) indicated that if (3) is instead taken to be exact, then the implied values of α and β are found by utilizing the one valid root associated with the solution of a cubic equation. However, then (2) necessarily becomes approximate. So in using the classic PERT formulas, researchers have indicated one can take the expression for the mean to be exact and the expression for the variance to be approximate, or vice versa (Gallagher, 1987; Littlefield and Randolph, 1987). It is then natural to wonder if (2) and (3) may hold simultaneously. Grubbs (1962) originally identified the conditions under which this occurs: these are when $\alpha = \beta = 4$; $\alpha = 3 + \sqrt{2}$, $\beta = 3 - \sqrt{2}$; and $\alpha = 3 - \sqrt{2}$, $\beta = 3 + \sqrt{2}$. The latter two special cases satisfy the $k = 6$ constraint, while the former requires lifting the $\alpha \neq \beta$ restriction. The PERT formulas have been defended by authors such as Clark (1962), Littlefield and Randolph (1987), Kamburowski (1997), and Pleguezuelo et al. (2003) typically by appealing to the normal distribution.

Thus, the beta distributions with $k = 6$ may be called the Type I PERT beta distributions because for these distributions the expectation is exactly described by (2). By extension, the beta distributions with $\alpha = \beta = 4$; $\alpha = 3 + \sqrt{2}$, $\beta = 3 - \sqrt{2}$; and $\alpha = 3 - \sqrt{2}$, $\beta = 3 + \sqrt{2}$ can be termed the Type II PERT beta distributions as the expectation and variance are exactly described by (2) and (3). We will return to these distributions below.

One of the more notable limitations of PERT beta involves the variance, the most typically encountered measure of uncertainty. By the third assumption of Littlefield and Randolph (1987) above, the variance in (3) is constant conditional on the range. This may be in direct conflict with reality. For example, consider two activities with the same range. While the expert is free to specify m which will change the expected value of the two activity times, both activities are constrained to have identical variances. This is in spite of the fact that considerable differences in uncertainty may exist. For example, one activity may be assigned to a new contractor or workgroup whose abilities may be unknown and untested. Alternatively, the completion of one activity may be advanced or delayed by a complex function of

[3]As pointed out by an anonymous reviewer, it is the standard normal distribution truncated at ± 2.96 that has a standard deviation equal to one sixth of the range.

other contingent factors such that uncertainty is again increased relative to the other better-understood activity. This variance inflexibility of PERT is important because the overall critical path time distribution is a function of the distribution of the path activity times. However, if the variances of the critical path activities are consistently underestimated, then it follows that the variance of the critical path time will be as well. This will lead to falsely precise conclusions regarding the critical path time that are too narrowly centered around the expected value. As a results, errors of the kind described by Atkinson et al. (2006), Grant et al. (2006) and Morgenshtern et al. (2006) would be expected to occur. Additionally, Kamburowski (1997) has indicated that when $k = 6$ and the variance is as in (3) the kurtosis measure is 3 (mesokurtotic). So, as in the case of light-tailed modeling of financial data, utilizing the PERT beta may unknowingly expose the project manager to increased real-world upside and downside risks and extreme events relative to what the method suggests. Project managers may therefore prefer a method that permits both an increased likelihood of tail-area events as well as greater flexibility in the variance specification. In the following section, we provide such a distribution.

3 A Heavy-Tailed Distribution for Activity Times

We employ a finite mixture distribution for the distribution of activity times (Titterington et al., 1985; McLachlan and Basford, 1988). The finite mixture distribution having I components can be written as

$$g(y|\boldsymbol{\Psi}) = \sum_{i=1}^{I} \pi_i f_i(y|\boldsymbol{\psi}_i) \tag{4}$$

where f and g are densities, $\boldsymbol{\psi}_i$ is the parameter vector for the ith component of the mixture, $\boldsymbol{\Psi} = \bigcup_i \boldsymbol{\psi}_i$ is the parameter vector resulting from the union of the I $\boldsymbol{\psi}_i$ vectors, and π_i is a set of non-negative weights which sum to one. Given that more activity-time uncertainty than is allowed for by the PERT beta may exist, the mixture distribution should be able to represent expert beliefs ranging from those specified under the PERT beta conditions to the conditions of maximum uncertainty. The rectangular distribution is the distribution which expresses maximum uncertainty subject to the basic constraint that we have a normalized density. If we consider (1) and the rectangular distribution, then by adding a mixing parameter θ such that $0 \leq \theta \leq 1$ we may write the beta rectangular mixture distribution as

$$p(y|\alpha, \beta, \theta, a, b) = \frac{\theta\, \Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \frac{(y-a)^{\alpha-1}(b-y)^{\beta-1}}{(b-a)^{\alpha+\beta-1}} + \frac{1-\theta}{b-a} \tag{5}$$

Figure 1: Examples of Beta Rectangular Distributions

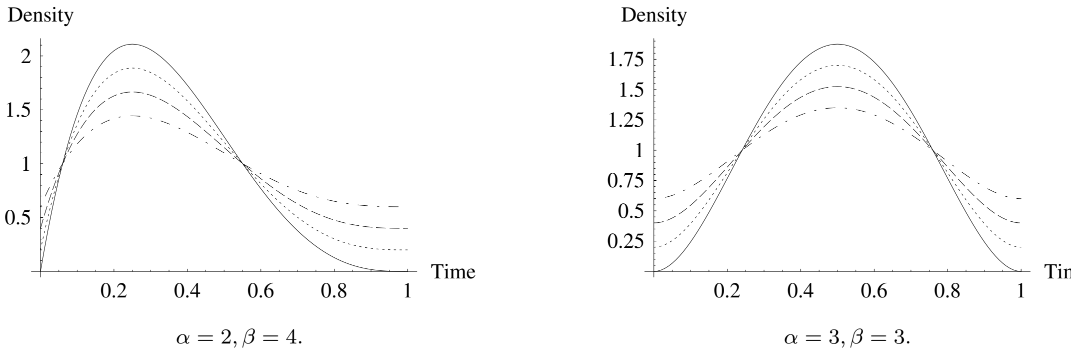

Solid: $\theta = 1$, Dotted: $\theta = 0.8$, Dashed: $\theta = 0.6$, Dash-dotted: $\theta = 0.4$.

Clearly the beta distribution is recovered when $\theta = 1$ and the rectangular distribution is recovered when $\theta = 0$. Thus, as the uncertainty weight $1 - \theta$ increases, overall uncertainty increases. Figure 1 displays the densities of two members of the beta rectangular family under the Type I PERT condition. The special $\theta = 1$ case corresponding to the original non-mixture PERT beta distribution is also shown in the figure via densities with solid lines.

In the general case, the expectation and the variance of (5) are

$$E(y) = a + (b-a)\left(\frac{\theta\,\alpha}{k} + \frac{1-\theta}{2}\right), \tag{6}$$

$$\mathrm{Var}(y) = (b-a)^2\left(\frac{\theta\,\alpha(\alpha+1)}{k(k+1)} + \frac{1-\theta}{3} - \frac{\big(k+\theta(\alpha-\beta)\big)^2}{4k^2}\right). \tag{7}$$

When $\theta > 0$ the mode is

$$m = a + (b-a)\frac{\alpha - 1}{k-2} \tag{8}$$

which clearly is the same mode as that of the underlying beta distribution. There is no unique mode when $\theta = 0$. However, this case poses no special difficulties below as the mode is no longer a quantity of interest when $\theta = 0$.

4 The Beta Rectangular Distribution: Application to PERT

We now consider the beta rectangular distribution in the PERT context. Since the expectation of the mixture is a weighted sum of the expectations of the rectangular and beta components, we may abbreviate the PERT expectation in (2) as $E(y_p)$ and use it to find the PERT-based expectation of the mixture. We denote the mixture expectation by a subscripted pm to indicate it extends a classic PERT formula, here (2), to the mixture case. If y_r has the rectangular distribution on the interval a to b, the PERT mixture expectation

$$\begin{aligned} E(y_{pm}) &= \theta\, E(y_p) + (1-\theta)E\,(y_r) \\ &= \theta \frac{a+4m+b}{6} + (1-\theta)\frac{a+b}{2} \\ &= \frac{\theta(a+4m+b) + 3(1-\theta)(a+b)}{6}. \end{aligned} \tag{9}$$

The result indicates that the mean in (9) is shrunk back toward the center of the interval away from the classic PERT mean in (2) as uncertainty increases with declining θ. We may also insert (8) into (9) and then equate the result to (6) and solve to determine what parameter values are implied. Doing so verifies the $k = 6$ condition is again needed for (6) to equal (9) given $\alpha \neq \beta$.

For the variance, we again utilize the finite mixture formulation in (4). To simplify notation, we write the PERT beta and rectangular components as y_p and y_r. Then we have the following theorem.

Theorem 1. *The PERT mixture variance is*

$$Var(y_{pm}) = \frac{1}{36}\left[\theta\big((a+4m+b)^2+(b-a)^2\big)+12(1-\theta)(a^2+ab+b^2)-\big(\theta(a+4m+b)+3(1-\theta)(a+b)\big)^2\right]. \tag{10}$$

Proof. Note that

$$\begin{aligned} \mathrm{Var}(y_{pm}) &= E(y_{pm}^2) - [E(y_{pm})]^2 && (11) \\ &= \theta\, E[(y_p)^2] + (1-\theta)E[(y_r)^2] - \big[E(y_{pm})\big]^2 && (12) \end{aligned}$$

where the first term in (11) becomes the first two terms in (12) by linearity and independence. The last term in (12) is given by squaring the result in (9). In the middle term of (12) we can show $E[(y_r)^2] = (a^2 + ab + b^2)/3$ by using the rectangular distribution. The first term in (12) is the beta component. We now find an expression for this quantity. Again using the variance formula and inserting the PERT beta expressions we find that

$$\begin{aligned}
\mathrm{Var}(y_p) &= E[(y_p)^2] - [E(y_p)]^2 \\
\frac{(b-a)^2}{36} &= E[(y_p)^2] - \left(\frac{a+4m+b}{6}\right)^2 \\
E[(y_p)^2] &= \frac{(b-a)^2}{36} + \left(\frac{a+4m+b}{6}\right)^2 .
\end{aligned} \tag{13}$$

By putting (13) into (12) we find (10) is obtained. □

If we re-examine the Type II PERT beta conditions, some simplification of (10) is possible. When $\alpha = \beta = 4$, $\mathrm{Var}(y_{pm}) = (b-a)^2(3-2\theta)/36$. Alternatively when $\alpha = 3 \pm \sqrt{2}$, $\beta = 3 \mp \sqrt{2}$,

$$\mathrm{Var}(y_{pm}) = \frac{(b-a)^2(3-2\theta^2)}{36}. \tag{14}$$

The variance in (14) is greater than that obtained in (3) except for $\theta = 1$ where they are equal. Hence, it represents a more conservative value that is amenable to computation as is (3) and that is also consistent with both (2) and (3) and the $k = 6$ condition. For increased precision, however, the expression in (10) would be recommended over (14) in the PERT context. We see from (10) that under the mixture the variance is no longer constant conditional on the range as it was in the original PERT case in (3). We may plot (10) as a function of θ and m to display the flexibility of the variance afforded by the mixture. At right in Figure 2 we see the variance has a minimum of 1/36 when $\theta = 1$ and thereafter increases as a function of θ and m. Figure 2 also shows that the mean becomes more moderated as uncertainty grows and θ declines. As θ tends to 0, the mean tends toward 1/2 given the lack of certainty. Plots for the exact expressions in Equations (6) and (7) for $k = 6$ are extremely similar and are hence omitted.

We now expand our focus from the PERT beta special cases and their associated moment approximations to the general beta distribution in (1) and exact treatments of activity time distributions. In doing so, we compare (1) to the beta rectangular in (5) with a view toward their usage as distributions for activity times. Use of the beta distribution corresponds to a set of assumptions about a representation of expert belief. While the implications of these assumptions regarding the mean and variance of the distributions have been studied in the literature (as described above), the implications of the use of the beta distribution to accurately represent tail-area beliefs do not appear to have been examined, even though tail-area behavior may substantially impact the mean and variance (Berger, 1985).

Figure 2: Mean and Variance of y_{pm}

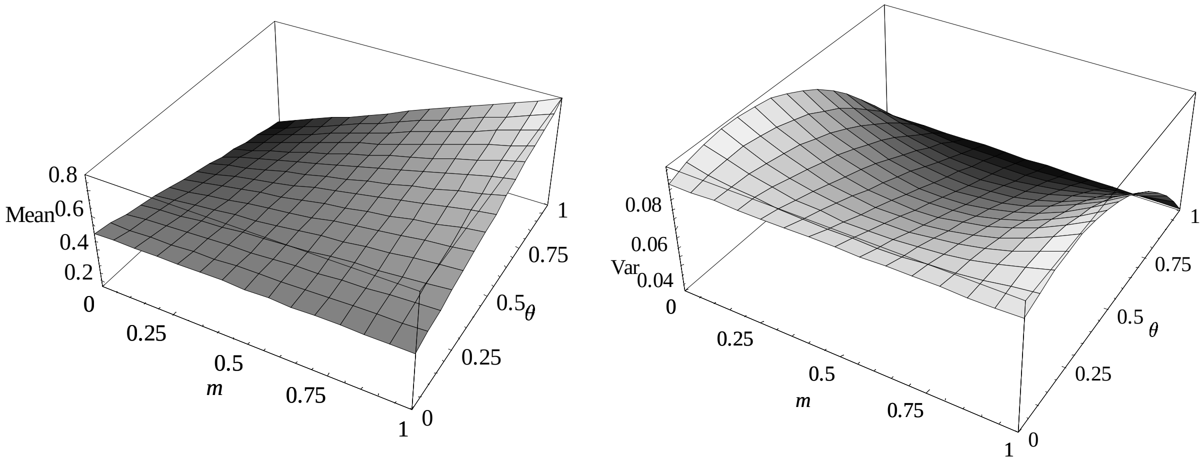

Beta distributions that are unimodal at an interior point ($\alpha > 1, \beta > 1$) vanish in the tails (i.e., their densities converge to zero at the extrema). By construction, the unimodal beta rectangular mixture with $\theta < 1$ apportions more weight to the extremal tail-area regions than does the beta distribution. The density of the unimodal beta rectangular never vanishes at the extrema as the density there can be shown to be $(1-\theta)/(b-a)$. As for the tail-area regions bordering the extrema, beta distributions generally decay rapidly in the tails, and do so increasingly with increasing α and β. The skewed unimodal beta distributions decay especially rapidly in the longer tail even with α and β relatively small. For example, the plot on the left of Figure 1 shows that there is little appreciable density in the top decile region $0.9 \leq y \leq 1$ for the beta distribution ($\theta = 1$) with $\alpha = 2, \beta = 4$. The symmetric beta distribution ($\alpha = 3, \beta = 3, \theta = 1$) on the right of Figure 1 fares better but still an observance of an activity time in the top decile would be rather unlikely. Given that α and β may sometimes be taken to be rather large (e.g., up to 60 in Keefer and Verdini, 1993), the total probability in these areas may become quite small. If the expert's belief regarding the decay of the activity time distribution toward the extrema truly follows that of the beta distribution, then the use of the beta distribution is well-supported. However, this may be challenging to verify. The beta rectangular provides a more gradual decay toward the extrema, and thus may be preferable as a more conservative alternative to the beta[4]. Thus, the beta rectangular may be of more general interest in the context of activity time modeling whether or not the PERT conditions are adopted.

[4]For completeness, we mention that the more conservative nature of the beta rectangular extends directly to the J-shaped distributions and will also moderate the U-shaped distributions should they be put into practice (however, the U-shaped beta distribution seems little used in the literature, perhaps in part because of challenges associated with incorporating the elicitation of distributions with bimodality at the extrema).

5 Elicitation

Direct elicitation of the beta distribution's parameters α and β is difficult although they can be solved for by a method-of-moments approach. However, direct estimation of the variance of a distribution can also be challenging (Clark, 1962; Perry and Greig, 1975), leading the creators of PERT to ask the expert to instead provide a, m and b (Clark, 1962). With regard to distributional elicitation, we note that alternative distributions besides the beta have been used for activity times, such as the triangular (MacCrimmon and Ryavec, 1964), the gamma, the chi-square and the lognormal (Perry and Greig, 1975), the truncated normal (Kotiah and Wallace, 1973), the Weibull (Abdelkader, 2004), the exponential (Kulkarni and Adlakha, 1986), the Erlang (Bendell et al., 1995; Azaron et al., 2006), and piecewise polynomial functions (Schmidt and Grossmann, 2000)[5]. Of these alternatives, it appears that the triangular distribution has been used with the greatest regularity (Megill, 1984; Williams, 1992; Johnson, 1997), the triangular attracting some adherents because of distributional tractability and ease of elicitation. In particular, the triangular distribution shares with the beta the straightforward task of eliciting only a, m and b. The attractiveness of the ease of elicitation property is consistent with the literature on judgment elicitation which points out the importance of straightforward elicitation tasks (Hogarth, 1975).

For the beta rectangular distribution discussed here, we also need the expert to provide θ in addition to a, m and b and we now describe how he/she can provide this. Since m is the likeliest, most frequently-occurring activity time, it is the quantity that the expert should be able to assess with the greatest certainty. The location of the mode also has a major impact on the shape of the overall distribution. As such, the less certainty there is about the mode, the less certainty there is likely to be about the shape of the distribution. Thus, lack of certainty about the mode is a clear indication that expert opinion is at least somewhat vague and diffuse, especially relative to the fairly specific assumptions underlying the PERT beta. As uncertainty about the mode gets particularly pronounced, the density of the distribution should become relatively flat to reflect the fact that other nearby parts of the distribution have almost the same relative likelihood as the mode. The beta rectangular distribution possesses this property.

Elicitation of θ is therefore relatively direct and can be accomplished by one of the following two methods or their conjunction. The expert can be asked to recall m. Then, using a scale from zero to 10 (or zero to 100), he/she is asked to indicate how certain he/she is that the mode is truly m, where the highest value on the scale indicates he/she

[5] An anonymous reviewer has also pointed out the uniform has been considered by Heller (1981).

is completely certain and the lowest that he/she is completely uncertain. We obtain θ by taking the number given and dividing it by the highest value on the scale (10 or 100), yielding $0 \leq \theta \leq 1$. Alternatively, it is easy to elicit θ using a computer-mediated interactive graphical method where the distributions associated with different values of θ are displayed and key quantities calculated such as in Kadane et al. (2006). A third approach that may be useful to the practitioner is as follows. The practitioner provides a, m, and b. They are then told that these estimates in the context of the PERT beta can be thought of as the opinion of Expert 1. Then, they are told to imagine that there is another expert, Expert 2, who agrees with a and b but is otherwise uncertain about the time and so gives equal weight to all times between a and b. The practitioner now has two 'opinions' about the distribution of the activity time (represented by the standard PERT beta and the rectangular respectively). For this particular activity, the practitioner is asked to consider the two experts' opinions and think of which is the more likely to best correspond to the true state of affairs regarding the distribution of times. He or she is then asked to give his/her personal odds as to which of these two opinions will likely be the more correct. θ and its complement can then be immediately obtained from the odds, such that if the odds ratio provided is $O_{\text{beta}}/O_{\text{rectangular}}$, then $\theta = O_{\text{beta}}/(O_{\text{beta}} + O_{\text{rectangular}})$. For example, if the odds are 2/1 that Expert 1 is the more correct, then we have $\theta = 2/3$. If the odds formulation is not desired, the question can be recast to ask the practitioner to indicate how much personal percentage weight he/she would give to these two opinions based on how likely each is to be the more correct. The percentage weight for Experts 1 and 2 respectively correspond to θ and $1 - \theta$.

6 Empirical Application

We now provide an example of the practical implications of the use of the method in an real-world example for the purpose of illustrating the outcomes and the inferential implications of the method in a real-world context. We examine the performance of the beta rectangular mixture vis-à-vis the beta formulation in the context of a real-world electronic module development project appearing in Moder et al. (1983, p. 294). The project consisted of 29 activities in a project development network having multiple paths. One key area of interest associated with managing such a project involves the obtaining the distribution of the total time required for the entire project to be completed, which is also known as the critical path time. In obtaining this distribution a manager is a position to forecast the likely time-to-completion of the project and may obtain other quantities of interest such as the probability the project will be completed by a particular date. The distribution of total project time, T, in stochastic project management techniques is often obtained by Monte

Carlo simulation because simulation allows the exact distribution to be arbitrarily closely approximated by using a large number of simulated realizations from the distribution of interest (Bowman, 1995).

We turn now to the distribution of T for the full electronic module development project. Plots of the distributions of T under the Type I condition appear in Figure 3. We see that as θ declines, the variance increases and weight shifts toward the tails, particularly the upper one. The mean project completion times are 47.8, 49.6, 51.3 and 53.0 days for the respective cases of the standard beta and the mixture given $\theta = 0.75, \theta = 0.5$, and $\theta = 0.25$. The respective standard deviations are 3.4, 4.6, 5.2 and 5.6, while the respective 95% probability intervals are 42.2 to 55.3 days, 42.5 to 60.1 days, 43.0 to 62.3 days, and 43.6 to 63.9 days. Under the standard beta distribution, the probability of T exceeding 60 is rather small at less than 0.05%. When $\theta = 0.75$ however, this probability rises to 2.6%. For $\theta = 0.5$ and $\theta = 0.25$, the respective probabilities are 7% and 13%. Again, more conservative project completion time estimates are obtained even with modest increases in the amount of expert uncertainty. Assuming that the standard beta is insufficient to describe the true uncertainty about activity times, these more conservative estimates of the total project length will help to reduce the probability of a failure to complete the project on time.

7 Conclusions

The beta rectangular mixture distribution allows for the representation of judgments ranging from those corresponding exactly to an arbitrary beta distribution to those in which the conditions are of maximum uncertainty. The distribution permits flexibility in the variance specification as well as the existence of heavier tails. The results of Section 4 allow the PERT beta conditions to be incorporated by using the PERT parameters a, m and b as well as θ.

More generally, the beta distribution in (1) is often characterized as being a flexible distribution. The distribution can take on several shapes; however, its tail-area behavior is actually rather limited. By construction, beta distributions tend to vanish rapidly in the tails. In particular, it is impossible to have appreciable density at both endpoints simultaneously unless the particular beta distribution that corresponds to the rectangular ($\alpha = 1, \beta = 1$) or one of its close neighbors is used (i.e., a beta($\alpha = 1 \pm \delta, \beta = 1 \pm \epsilon$) distribution, with δ, ϵ small). Thus, the beta distribution in general is not particularly flexible on this key consideration of tail areas. However, the beta rectangular mixture is more flexible in this regard. Further, the beta distribution, by vanishing in the tails, heavily downweights the probability of extremal events according to the restrictions of its functional form. The beta

Figure 3: Distributions of T: Electronic Module Development Project

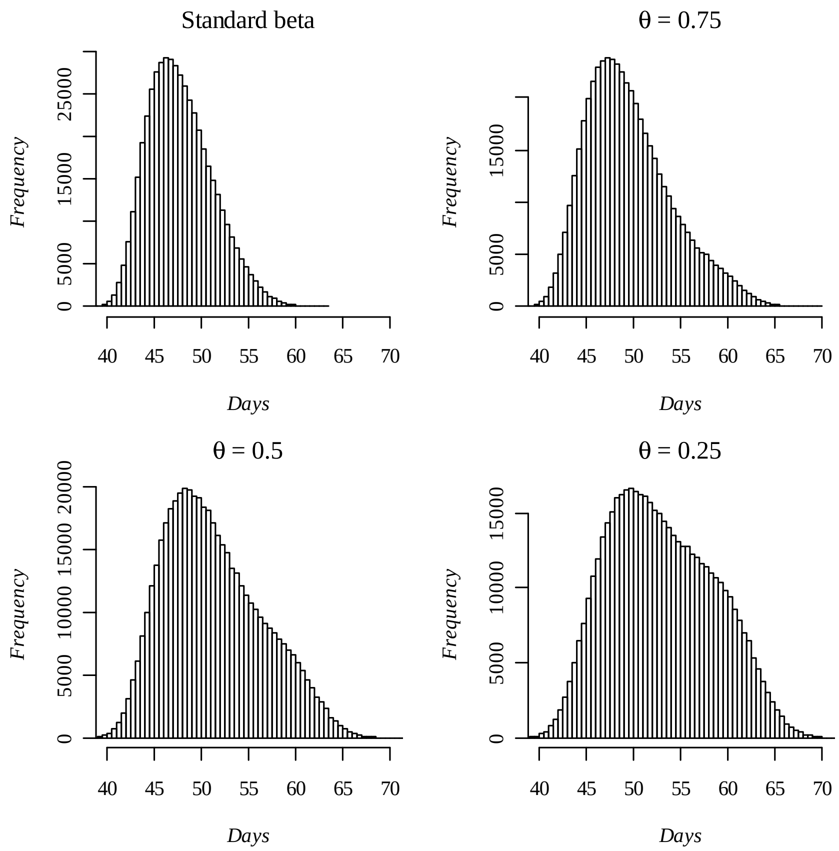

rectangular mixture, by contrast, addresses this problem by allowing the expert to provide additional information that influences the likelihood of extremal events.

Since the pioneering development of PERT, the literature to date has often proceeded along one of two lines. In the first, various well-known distributions have been used in PERT and the focus has been on deepening understanding of how these distributions' characteristics are related to properties of the estimates (Perry and Greig, 1975; Sasieni, 1986; Littlefield and Randolph, 1987; Gallagher, 1987; Kamburowski, 1997). In the second, the focus has been on refining the capability to make more accurate estimates of distributional estimands (Keefer and Bodily, 1983; Farnum and Stanton, 1987; Golenko-Ginzburg, 1988; Keefer and Verdini, 1993; Premachandra, 2001). The current work departs from these streams and instead argues that a key issue that has been overlooked is that experts may have considerably more uncertainty than is allowed for by the use of well-known but thin-tailed distributions. We then examine the characteristics of the situation to formulate a distribution that is more tailored to the situation and utilize a mixture distribution to do so.

This paper has implications for the following audiences. First, the general framework of mixture distributions and the associated newly-developed project time distributions presented here provide a new and richer framework for specification of underlying models for stochastic project management phenomena. Second, researchers may use the current work in the context of existing research streams such as in formulating more accurate estimates of distributional estimands to further extend these literatures. Finally, project management practitioners may readily apply these methods for improved description of project duration times. Moreover, the methods can be easily combined with Monte Carlo simulation techniques to provide more exact results regarding other key quantities such as the distribution of critical path times.

8 Appendix

As the CDF of the newly-described beta-rectangular distribution does not appear in the literature and as the CDF is useful for Monte Carlo simulation common in project management, it is included here. It is

$$F(y|\alpha,\beta,\theta,a,b) = \begin{cases} 1 & \text{if } y > b, \\ \dfrac{\theta\,\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \displaystyle\int_0^{(y-a)/(b-a)} t^{\alpha-1}(1-t)^{\beta-1}\,dt + \dfrac{(1-\theta)(y-a)}{b-a} & \text{if } a \le y \le b, \\ 0 & \text{otherwise.} \end{cases}$$

References

Abdelkader YH. Evaluating project completion times when activity times are Weibull distributed. European Journal of Operational Research 2004;157(3); 704–715.

Atkinson R, Crawford L, Ward S. Fundamental uncertainties in projects and the scope of project management. International Journal of Project Management 2006;24(8); 687–698.

Azaron A, Katagiri H, Sakawa M, Kato K, Memariani A. A multi-objective resource allocation problem in PERT networks. European Journal of Operational Research 2006;172(3); 838–854.

Banker RD, Kauffman RJ. Reuse and productivity in integrated computer-aided software engineering: An empirical study. MIS Quarterly 1991;15(3); 375–401.

Bendell A, Solomon D, Carter JM. Evaluating project completion times when activity times are Erlang distributed. Journal of the Operational Research Society 1995;46; 867–882.

Berger JO. Statistical Decision Theory and Bayesian Analysis. Springer, New York, 1985.

Bowman RA. Efficient estimation of arc criticalities in stochastic activity networks. Management Science 1995;41(1); 58–67.

Clark CE. The PERT model for the distribution of an activity time. Operations Research 1962;10; 405–406.

Engle RF. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. Econometrica 1982;50(4); 987–1007.

Fama EF. The behavior of stock-market prices. Journal of Business 1965;38(1); 34–105.

Farnum NR, Stanton LW. Some results concerning the estimation of beta distribution parameters in PERT. Journal of Operational Research Society 1987;38(3); 287–290.

Gallagher C. A note on PERT assumptions. Management Science 1987;33(10); 1360.

Gelman A, Carlin JB, Stern HS, Rubin DB. Bayesian Data Analysis. Chapman & Hall/CRC, Boca Raton, FL, 2nd edn., 2003.

Golenko-Ginzburg D. On the distribution of activity time in PERT. Journal of the Operational Research Society 1988;39(8); 767–771.

Grant KP, Cashman WM, Christenson DS. Delivering projects on time. Research Technology Management 2006;49(6); 52–58.

Grubbs FE. Attempts to validate certain PERT statistics or 'Picking on PERT'. Operations Research 1962;10; 912–915.

Heller U. On the shortest overall duration in stochastic project networks. Methods of Operations Research 1981;42; 85–104.

Hogarth RM. Cognitive processes and the assessment of subjective probability distributions. Journal of the American Statistical Association 1975;70; 271–289.

Johnson D. The triangular distribution as a proxy for the beta distribution in risk analysis. The Statistician 1997;46(3); 387–398.

Kadane JB, Shmueli G, Minka TP, Borle S, Boatwright P. Conjugate analysis of the Conway-Maxwell-Poisson distribution. Bayesian Analysis, 2006;1(2); 363–374.

Kamburowski J. New validations of PERT times. Omega 1997;25(3); 323–328.

Keefer DL, Bodily SE. Three-point approximations for continuous random variables. Management Science 1983;29(5); 595–609.

Keefer DL, Verdini WA. Better estimation of PERT activity time parameters. Management Science 1993;39(9); 1086–1091.

Kotiah T, Wallace N. Another look at the PERT assumptions. Management Science 1973; 20(1); 44–49.

Kulkarni VG, Adlakha VG. Markov and Markov-regenerative PERT Networks. Operations Research 1986;34(5); 769–781.

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

MacCrimmon KR, Ryavec CA. An analytical study of the PERT assumptions. Operations Research 1964;12(1); 16–37.

Malcolm DG, Roseboom JH, Clark CE, Fazar W. Application of a technique for research and development program evaluation. Operations Research 1959;7; 646–669.

Mandelbrot BB. The variation of certain speculative prices. Journal of Business 1963;36(4); 394–419.

McCullagh P, Nelder JA. Generalized Linear Models. Chapman & Hall, London, 2nd edn., 1989.

McLachlan GJ, Basford KE. Mixture Models: Inference and Applications to Clustering. Marcel Dekker, New York, 1988.

Megill RE. An Introduction to Risk Analysis. PennWell Books, Tulsa, OK, 2nd edn., 1984.

Mitchell VL, Zmud RW. The effects of coupling IT and work process strategies in redesign projects. Organization Science 1999;10(4); 424–438.

Moder JJ, Phillips CR, Davis EW. Project Management with CPM, PERT and Precedence Diagramming. Van Nostrand Reinhold, New York, 3rd edn., 1983.

Morgenshtern O, Raz T, Dvir D. Empirical analysis in software process simulation modeling. Information and Software Technology 2006;; in press.

Perry C, Greig ID. Estimating the mean and variance of subjective distributions in PERT and decision analysis. Management Science 1975;21(12); 1477–1480.

Pleguezuelo RH, Pérez JG, Rambaud SC. A note on the reasonableness of PERT hypotheses. Operations Research Letters 2003;31(1); 60–62.

Premachandra IM. An approximation of the activity duration distribution in PERT. Computers and Operations Research 2001;28(5); 443–452.

Raffo DM, Kellner MI. Empirical analysis in software process simulation modeling. Journal of Systems and Software 2000;53(1); 31–41.

Sasieni MW. A note on PERT times. Management Science 1986;32(12); 1652–1653.

Schmidt CW, Grossmann IE. The exact overall time distribution of a project with uncertain task durations. European Journal of Operational Research 2000;126(3); 614–636.

Steele MD, Huber WA. Exploring data to detect project problems. AACE International Transactions 2004;32(12); PM.21.1–PM.21.7.

Titterington DM, Smith AFM, Makov UE. Statistical Analysis of Finite Mixture Distributions. Wiley, Chichester, 1985.

Williams TM. Practical use of distributions in network analysis. Journal of the Operations Research Society 1992;43; 265–270.