January 2013

Project management decisions with uncertain targets

Jeffrey Keisler
University of Massachusetts Boston, jeff.keisler@umb.edu

Robert Bordley
Booz-Allen Hamilton, bordley_robert@bah.com

Project Management Decisions with Uncertain Targets

Jeffrey M. Keisler, Professor, University of Massachusetts Boston, Boston, MA 02125,

jeff.keisler@umb.edu

Robert Bordley, Certified Project Management Professional, Fellow, Booz-Allen Hamilton, Troy,

Michigan 48084, bordley_robert@bah.com

Abstract

Formal project management methods are largely organized around ensuring that the uncertain performance of activities suffices to meet fixed targets and requirements. In practice, requirements often change. Qualitative aspects of project management aim to anticipate and respond to the complications these changes create. This paper explicitly treat targets as uncertain. This allows the qualitative challenges to be recast as quantitative decisions under uncertainty. Decision analytic techniques can then be applied. In particular, we can interpret the probability that a project's performance will exceed some uncertain target can be interpreted as the utility for that level of performance. We obtain solutions for a fundamental set of project management problems in terms of uncertain targets. The language of target-oriented utility used here provides a bridge for representing project management problems in decision analytic models, and for translating the solutions of those models into the language of project management.

1 Introduction

1.1 Background

Project management (PM) is a widely practiced set of approaches for applying resources and activities so as to bring about a satisfactory outcome to the project. This paper demonstrates a way to use decision analysis for a set of PM problems that have been resistant to analytic approaches. Before developing machinery, we start with a high level review of PM and its associated decision problems to be modeled. Here, we draw on the standard of the field, the Project Management Body of Knowledge (i.e., PMBOK, Project Management Institute, 2013). PMBOK organizes project management processes into five process groups: initiating, planning, executing, monitoring and closing. An actual project may draw on each of these projects at multiple phases. While there are myriad interrelated tools and structures for coordinating work across these five groups, the following vastly simplified version of the project management process underlying PMBOK captures their key roles.

A project starts with a directive from a project sponsor or client who calls for something to be done, with some idea of who will be responsible (a project manager) and some understanding of the resources that will be made available. The initiation process then identifies stakeholders and other requisite information. Initation bridges to planning as stakeholders decide on the requirements and resources to be used along with deciding on the general strategy of the project.

In planning, especially as relates to managing the timing of the project, the project manager defines the activities necessary in order to complete the project, and the relationship between these activities (e.g., creating a work breakdown structure). The project manager then estimates their durations and resource requirements and decides on their sequence, resources, and schedules, with the aim of meeting project requirements and managing associated risks. The planning group also works with the other groups to revise plans in response to changes.

In execution, project staff perform the planned activities. The activities may go better or worse than expected, and staff must communicate about day to day progress and needs. The project manager may decide to make minor local adjustment to the activity plans as needed, or to request changes if large adjustments seem necessary.

Monitoring of the project occurs with milestones, metrics and systems to check progress and performance. The timing and nature of these checks may be decided in advance. Connected to monitoring is control, whereby needed or approved changes are communicated to the execution group.

Project closure ensures that the product is evaluated and if it is satisfactory, documents client acceptance along with other results of the project. The organization receives the rewards associated with its success, and accounts are closed.

For many of the decisions falling within the process groups above, there are quantitative approaches that yield decision rules. In the project context, it is natural to think in terms of fixed requirements. For example, it may be easier to delegate tactical responsibility down the organization if managers can simply report whether or not activities are done satisfactorily. Decision rules are typically formulated accordingly. For example the critical path method (CPM) supports scheduling of related activities in order to complete a project on time.

Throughout the project life cycle, however, there are often changes to specifications and requirements (e.g., Gjerde et al 2002, Bhattacharya et al, 1998), throwing off plans in non-intuitive ways that cause difficulty and inefficiency. The changes may be due a variety of internal or external conditions. The project manager (and possibly the project sponsor) may decide to change the projects deliverables, schedule or resources. Such changes have consequences that cascade down to project plans in complex ways, e.g., a large exogenous budget cut could present a crisis requiring a complete reorganization of a project. As a result, a lot of PM practice focuses on handling changes of varying degrees as smoothly as possible, and anticipating the possibility of such changes in order to mitigate their consequences. For example, budgets or schedules often build in a specific amount of slack. This aspect of PM has traditionally had a qualitative and organizational focus. Changes in requirements are driven by a wide range of factors. This is somewhat in contrast to what might be thought of as variation in more or less repeatable internal processes that drive performance.

Thus, PM decision approaches may accommodate some uncertainty in project performance, e.g., scheduling algorithms and project risk management in general may incorporate uncertainty about activity completion times. But these approaches do not formally incorporate uncertainty about requirements, and this fact contributes to how changes in requirements complicate PM. In essence, PM requires the project manager to act as if requirements (or targets) are fixed (i.e., as if there is no uncertainty about them) and then to implement an orderly change control process when targets do change.

The decision analytic approach is to explicitly and quantitatively incorporate tradeoffs and uncertainties, even those whose quantification will largely rely on subjective judgment.

Decision analysis (DA) practitioners have applied standard DA tools on traditional PM decisions (Schuyler, 2001, Virine & Trumper, 2007). In fact, PMBOK includes some standard decision analytic techniques, e.g., decision trees for multi-stage decisions under uncertainty, expected monetary values, risk tolerance, and three-point probability assessment techniques. Alternatively, multicriteria models are used in project planning and scheduling (T'Kindt & Billaut, 2001). But PM decision models are already complex and forcing a fuller set of decision analytic constructs onto them can lead to prohibitive modeling challenges. For instance, Goodwin et al (1998) noted that "performing these tradeoffs [for multiattribute utility in PM] requires subtle distinctions that are beyond the fidelity of any utility model that could be constructed with reasonable effort."

Our approach in this paper differs from prior DA efforts by formally introducing uncertainty about requirements (or targets) into PM decisions. The hope is that (1) this will improve PM decisions and their implementation, and (2) that this can be done with minimal disruption to existing PM practice, e.g., by modifying or building on existing calculation and management methods. To this end, there is an emerging literature on decision analysis with uncertain targets (e.g., Castagnoli and LiCalzi, 1996, Bordley & LiCalzi, 2000). In some cases, it is possible to define a utility function (a target-oriented utility function, or TOUF) on the degree of achievement on one or more objectives by relating this achievement to the likelihood of meeting targets. Models using this formulation thus divide the role of the uncertain performance, which is influenced by the decision makers actions, and the uncertain target, which is not. Of course, any problem that could be modeled with target oriented utilities could also be modeled using a more standard decision analysis which simply considers uncertainties on targets the same as uncertainties on any other variables, with endpoint utilities calculated based on the target and performance levels. Thus, we are agnostic about whether target oriented utility (TOU) should actually be interpreted in this context in the same way as classical utility in decision making with its associated behavioral and organizational implications, or simply as a way of encoding information. Using the TOU framework to characterize PM decisions allows for concise expression of various notions that will be important to this effort. To be clear, the proposed integration of uncertain targets to project management decisions is new and the results would be equally correct whether developed with or without TOU.

1.2 Preview

The plan for the paper is to develop decision analytic rules incorporating uncertain targets for specific PM decision problems across the PM process. We consider these problems in an order that allows TOU concepts and simple PM rules to be introduced and woven together in order to build toward larger scale applications. To establish this approach, we touch a lot of the important elementary decision problems spanning the process groups. The sections contain numbered examples for which we derive new heuristics and formal results for PM. These examples demonstrate the approach and its viability in tractable versions of PM problems, e.g., assuming Gaussian distributions to obtain mean-variance results, or using two elements when real-world problems might have any number of elements. Furthermore, the field of PM has a vast literature containing many decision rules and process for many specific problems. Thus, this work is not intended as a final word on PM decision rules, but rather opens up future research avenues to further translate this approach to more practical situations. Specific relevant literature and methods will be discussed in the context of the particular problems. The last part of this introduction (subsection 1.3) lays out our mathematical notation, and the rest of the paper is structured as follows:

Section 2 presents decision analytic concepts and language to be used throughout the paper, introducing them in the context of a generic project with a single dimension (e.g., time) of uncertain performance and an uncertain target. Subsection 2.1 considers the tactical planning problem of selecting a project plan, and uses a TOUF to do so. An immediate application of this to project initiation is developed in subsection 2.2, which shows how explicit treatment of uncertainty on targets can facilitate more effective decisions about project requirements that comprehend the tactical decisions. Subsection 2.3 develops the first example which derives a decision analytic certainty equivalent and risk premium formulations for project performance, which can in turn serve as the basis for various decision rules.

Section 3 extends the approach to planning projects involving multiple interrelated activities. Subsection 3.1 sets up our second example in which a conventional approach to scheduling activities and allocating resources can be used to reduce performance risk along critical paths. With this setup, a modest transformation allows the incorporation of uncertain targets in subsection 3.2, where a modification of the standard decision rule is developed and interpreted, particularly in the context of project execution decisions that adjust plans for one or more activities without changing the overall project resources or deliverables. Subsection 3.3 generalizes activity planning with additional examples intended to serve as building blocks for larger activity planning efforts. Here, a set of activities in series must be completed before some uncertain target date, or either one or both of two activities in parallel must be completed before the uncertain target date.

Section 4 develops the approach for what PMBOK (p.6) describes as balancing the competing project constraints which include, but are not limited to: scope; quality, schedule, budget, resources and risks. Project balancing decisions may be made during project initiation as requirements and resources are defined, as a result of changed conditions that force the reconsideration of requirements and resources. Subsection 4.1 develops an example of project level balancing, where the project manager cannot change the resources or the need to meet the evolving targets, but can shift focus toward improving performance on one dimension or another. To give this example more realism, it is developed into several parts that suggest tradeoffs across the multiple attributes of time, quality and cost, as well as incorporating more flexible assumptions about probability distributions. Subsection 4.2 develops a brief example of business case level balancing, where resources can be shifted toward different dimensions of performance toward the meeting of an overall target.

Section 5 builds on the results for single stage decisions in the earlier sections to general multi-stage decisions. Monitoring plans as part of PM can be thought of as decisions about information acquisition (at a cost) preceding later decisions about (also costly) operational plans in response to necessary change or rework. Subsection 5.1 develops an example of a multi-stage decision in which value of information calculations can guide the decisions about monitoring. In our formulation, it is convenient to distinguish between information acquired about performance and information acquired about the evolving requirements. Thus, subsection 5.2 treats the example of planning for verification, which obtains information about performance which, per PMBOK (p.566) is used to whether a product, service, or system complies with a regulation, requirement, specification, or imposed condition. Subsection 5.3 treats the example of planning for validation, which, per PMBOK (p.566) is the assurance that a product, service, or system meets the needs of the customer and other identified stakeholders, which may involve obtaining updated information about those actual needs.

Our contribution thus consists of (1) templates for analyzing a suite of PM decision problems using TOU and (2) tractable illustrative models and decision rules derived for them (which can sometimes be calculated as modifications of familiar PM heuristics. These developments establish connections between TOU, DA and PM as a basis for future practical implementations.

1.3 Notation

A **project** is a set of **activities** to be completed. Project managers make **decisions** about how to complete the project, and the project results in some level of **performance**. Project performance is considered to be against a **target**. The **outcome** is that the project either succeeds or fails in meeting its target.

Performance levels will have n attributes and will be represented by points of $\Re^n$, where $n \geq 1$. We will work with various (n-dimensional) random variables

$X = (X_1, ..., X_n) : \Omega- > \Re^n$,

where X_i is a random variable associated with performance on attribute i. We sometimes write x_i for $X_i(\omega)$ and will suppress the subscript i when there is only one attribute ($n = 1$).

There will be a decision space D. We sometimes introduce an additional argument for a decision d in D, and write

$X(\cdot) : \Omega \times D \rightarrow \Re^n$,

so for each $d \in D$, $X(d) : \Omega \rightarrow \Re^n$ is the n-dimensional random variable associated with performance under decision d.

There will also be a set of $m \geq 1$ activities denoted by $A = \{a_1, ..., a_m\}$. For each activity a_j, the random variable

$X^{(j)} = (X_1^{(j)}, ..., X_n^{(j)}) : \Omega \rightarrow \Re^n$

is associated with project performance on activity a_j. We suppress the superscript when there is only one activity. For each set J of (indexes of) activities, we write

$X^{(J)} = \sum_{j \ \epsilon J} X^{(j)}$

e.g., $X^{(\{1,3,5\})} = X^{(1)} + X^{(3)} + X^{(5)}$. We denote targets with

$T = \{T_1, \ldots, T_n\} : \Omega \rightarrow \Re^n$

where variable T_i is associated with the project target on attribute i; we will write t_i for $T_i(\omega)$. To describe the outcome of a project, we use

$v(\cdot,\cdot) : \Re^n \times \Re^n \rightarrow \{0, 1\}$ assigns value $v(x, t)$ to a project when its performance is x and the target is t. For $n = 1$, if X is a positively oriented metric, i.e., one where a higher score is better, then by convention, $v(x, t) = 1$ if $x \geq t$ and $v(x, t) = 0$ if $x < t$. If X is negatively oriented , e.g., time to completion, then by convention, $v(x, t) = 1$ if $-x \geq -t$ and $v(x, t) = 0$ if $-x < -t$. For $n > 1$, possible rules for v will be discussed. We can think of the binary function v as returning the sponsor's utility for a successful or unsuccessful project.

We also define

$u_T(x) \equiv E[v(x, T)]$,

and we shall typically suppress the subscript and simply write $u(x)$. This notation has two benefits. It allows a more compact expression of some of the quantities and relationships of interest, and, following target-oriented utility theory, it suggests that the decision maker can perform analysis in terms of a utility function over performance levels.

$X_i^{(j)}$, $X_i^{(J)}$, $X_i(d)$, $X_i^{(j)}(d)$, $X_i^{(J)}(d)$, $T^{(j)}$ etc. would have the natural interpretations.

Throughout, e will denote a Gaussian random variable with mean 0 and unit variance. $s(\cdot)$ will denote standard deviation, $V(\cdot)$ will denote variance, and E will denote expectation.

Additional parenthetical descriptors will be used in subscripts for special notations not meant to be part of indexing, e.g., $T_{(av)}$ and $T_{(diff)}$.

2 Deciding on a Plan for a Simple Project

2.1 The TOU frame for project management decisions

Consider a situation where the project manager must select a plan to execute the project. Project performance will be evaluated with a single performance metric, X. The hope is that the project meets or beats its target level of performance T, in this case with known value t. The project manager distinguishes two possible outcomes to the project: success (with $v(x,t) = 1$) when $x \geq t$, and failure ($v(x,t) = 0$) when $x < t$. Performance levels are uncertain, so associated with each plan is a probability density $f(x)$.

The decision problem is simply

1. Select a plan $d \in D$
2. Observe the value of $X = x$
3. Receive value of 1 if $x \geq t$

When the value of T is known to be t, The project manager's utility function for a given plan as defined above is then $u(X) = E[v(X,t)] = \int f(x) \Pr.\{x \geq t\}dx = \Pr.\{x \geq t\}$.

If the project manager is deciding among plans d_1 and d_2, the project manager calculates $E[u(X(d_1)]$ and $E[u(X(d_2))]$ and selects the plan with the higher expected utility, that is, the plan with the greatest probability of meeting the target. Note that if two plans have different probability distributions neither of which stochastically dominates the other, then neither plan will be superior to the other for all values of t.

This becomes more interesting if, rather than treating the target as a constant, we allow T to be a random variable taking the value t. Then the decision problem becomes:

1. Select a plan $d \in D$

2. Observe the value of $X = x$

3. Observe the value of $T = t$

4. Receive value of 1 if $x \geq t$

In this case, the project manager should still select the plan that maximizes $\Pr.\{X \geq T\}$. Her expected utility $E[u(x)] = u(x) = \Pr.\{x \geq T\}$ is monotonic and increasing in x, if T has a right-continuous probability distribution.

As shown in research on TOU, expected utility here can be used like any arbitrary utility function with those properties as long as X and T are independent.

The remaining examples in this paper involve uncertain performance X and uncertain targets T, and in all cases we shall assume X and T are stochastically independent of one another.

2.2 Setting Project Requirements

Note that for a given set of alternatives with known performance probability distributions, maximizing the expected utility of the target- oriented utility function (TOUF) is equivalent to maximizing the expectation of the original binary value function. But there is a very real sense in which these two views differ because of the way they structure information. For example, if a competitor is expected to come out with a product in one year, a project manager may be content to aim for a project deadline of 364 days. The project sponsor may push for a project deadline much earlier than the anticipated competitor launch date to ensure being first to market, while the project manager argues that such an aggressive project timeline is too risky. One proposed use of TOU has been to align organizational incentives (Abbas et al, 2008). This can apply to setting project requirements, which can be thought of as a game of the following sort in which different players have different information sets:

Let D denote the project manager's decision space, i.e., the set of alternative plans d available, and $X(D)$ denotes the project manager's probability distribution on performance levels the plans in D. We borrow from game theory the notion of a project manager's type θ defined by the $X(D)$ she faces. We assume the project sponsor somehow assigns probability distributions over θ and the target T, and sets a reward schedule for the project manager consisting of a requirement level $x_{(req)}$ and an associated bonus b (insignificant relative to the value of the project but sufficient to motivate the manager to even accept the assignment) which is paid to the manager if performance $x \geq x_{(req)}$. The project manager (whom we shall assume risk neutral) knows her type $\theta_{(mgr)}$, but does not know the type of the sponsor, i.e., the sponsor's distribution on T. The project manager is assumed to select d so as to maximizes her expected reward, that is, $b \Pr . X(d) \geq T$, and the sponsor chooses $x_{(req)}$ so as to maximize the expected payoff of the project.

Finally the value of the target is revealed, and the project manager and sponsor both receive the payoff associated with the resulting combination of target, performance and requirement. When the target is known with certainty, setting $x_{(req)} = t$ equal to the target leads to the project manager acting in the best interest of the sponsor. When the target is uncertain, the optimal value of $x_{(req)}$ depends on θ, e.g., for some proposed $x_{(req)}$ and plans d_1 and d_2, for θ_1, $\Pr .\{X(d_1) \geq x_{(req)}\} > \Pr .\{X(d_2) \geq x_{(req)}\}$ and $\Pr .\{X(d_1) \geq T\} > \Pr .\{X(d_2) \geq T\}$, while for θ_2, $\Pr .\{X(d_1) \geq x_{(req)}\} > \Pr .\{X(d_2) \geq x_{(req)}\}$ and $\Pr .\{X(d_1) \geq T\} < \Pr .\{X(d_2) \geq T\}$.

Furthermore, the project manager may not wish to reveal θ if this information could lead the sponsor to set $x_{(req)}$ higher and b lower. Alternatively, the sponsor could try to align incentives by simply rewarding the managing if ultimately $x \geq t$ and announcing to the manager the distribution over T. The problem here is that the sponsor may have reason to misrepresent T, e.g., in order to get the manager to accept a lower b. The project manager may therefore form her own belief about T, and her choice of plan will depend on this belief.

The problem here is designing a mechanism that results in the right incentives for the project manager. A solution is for the sponsor to announce a preliminary reward schedule where for all x, the reward is $bu(x)$ where $u(x) = \Pr.\{x \geq T\}$. The shape and bounds of u depend only on the sponsor's knowledge of T and not on θ. The manager and sponsor must still negotiate an acceptable b, and this may indeed depend in some unspecified way on both sides' knowledge of both T and θ, but this step will not affect actual decisions about the project. The manager's choice depends only on the schedule for $u(x)$ and θ and not on T. The manager who uses this reward schedule as a utility function is positioned to make decisions that also maximize the sponsor's expected utility for the project.

2.3 Certainty Equivalent and Risk Premium

A useful concept in DA is the certainty equivalent, which allows comparison of alternative plans with uncertain results using units of the original performance measure (dollars in many cases) rather than an artificially defined utility scale. Another useful concept is the risk premium, which quantifies the undesirability of the risk in a gamble as the difference between its certainty equivalent and expected value.

For the TOUF described above, the certainty equivalent $x_{(ce)}$ associated with an uncertain payoff X can be determined by solving

$$\Pr.\{x_{(ce)} \geq T\} = \Pr.\{X \geq T\}$$

For T uniform, the certainty equivalent is $E[X]$ with the risk premium (the difference between expected value and certainty equivalent) being trivially equal to zero. We now develop a decision rule based on these concepts for a more interesting case.

Example 1: If T is exponential (corresponding to the exponential utility $u(x) = 1 - \exp(-Rx)$, where 1/R is the risk tolerance) and if X is Gaussian with variance $s^2(X)$, then the certainty equivalent is $E(X) - R\frac{s^2(X)}{2}$ (and the risk premium is $R\frac{s^2(X)}{2}$.)

A TOUF based on a Gaussian distribution for T would have the form

$$u(x) = \int_{\infty}^{x} \exp(-\frac{x - E[T]}{s^2(T)})dt = \int_{\infty}^{x} \exp(-\frac{x}{s^2(T)} + 2x\frac{E[T]}{s^2(T)} - \frac{E[T]}{s^2(T)})$$

which would approach the first TOUF (T uniform) if $s^2(T)$ approaches infinity while $E(T)$ does not, and approaches the second TOUF (exponential) if $E[T]$, and variance, $s^2(T)$, both go to infinity with $\frac{2E[T]}{s^2(T)} = R$ being constant.

For the cumulative normal, the optimal choice maximizes the certainty equivalent

$$E[T] + [\frac{s^2(T)}{s^2(T) + s^2(X)}]^{1/2}(E[X] - E[T])$$

Proof: T has mean $E[T]$ and standard deviation $s(T)$ while X has mean $E[X]$ and standard deviation $s(X)$. Then the certainty equivalent, $x_{(ce)}$, satisfies

$$\Pr.\{x_{(ce)} > T\} = \Pr.\{(c - E[T])/s(T) > e)\} = \Pr.\{X \geq T\} = \Pr.\{\frac{E[X] - E[T]}{(s^2(X) + s^2(T))^{1/2}} > e\}$$

so that $x_{(ce)} = E[T] + \frac{E[X]-E[T]}{(1+s^2(X)/s^2(T))^{1/2}}$ which demonstrates the result.

If we define the *risk-adjusted weighting factor* as

$$r = [\frac{s^2(T)}{s^2(T) + s^2(X)}]^{1/2}$$

then the certainty equivalent becomes $E[T] + r(E[X] - E[T])$ where $r < 1$ implies risk aversion. In the limit, as $s^2(T)$ goes to infinity, r approaches 1, a condition which implies risk-neutrality. It is then meaningful to talk about a *performance risk premium*, $E[X] - x_{(ce)}$ which is just

$$(1 - r)(E[X] - E[T])$$

The amount of expected performance the manager would give up in order to obtain certainty depends on the both the uncertainty associated with project performance and with the performance risk attitude implied by the uncertainty about the target. The z-score commonly used in conventional project management (in which T is assumed known) is given by $z = \frac{E[X]-E[T]}{[s^2(X)+s^2(T)]^{1/2}}$ and corresponds to the standardized gap between certainty equivalent and target, $\frac{x_{(ce)}-E[T]}{s(T)}$. Hence the z-score is an affine transformation of the risk premium.

To gain an understanding of the implications of results such as this, we can formulate the Taylor series expansions and observe that, under some unexceptional assumptions, they approximate simple decision rules.

When $s^2(X) < s^2(T)$, a Taylor Series expansion gives

$$x_{(ce)} = E[T] + (E[X] - E[T])(1 - \frac{s^2(X)}{2s^2(T)} - ...)$$

Since the risk-aversion index for this utility is

$$\frac{-u''}{u'} = -\frac{d\ln(u')}{dx} = (x - E[T])/s^2(T)$$

defining an average risk-aversion index by $R_{(av)} = \frac{E[X]-E[T]}{s^2(T)}$ and substituting into the Taylor Series expansion gives the approximation

$$x_{(ce)} \approx (E[X] - E[T]) - \frac{R_{(av)}s^2(X)}{2}$$

which, when $s^2(T)$ is large (and $R_{(av)}$ is small) leads to the risk-neutral solution. When $s^2(X) > s^2(T)$, we can also write the Taylor Series expansion

$$x_{(ce)} = E[T]+\frac{E[X]-E[T]}{s(X)}s(T)[1-\frac{1}{2}(s^2(T)/s^2(X))+...] = E[T]+\frac{E[X]-E[T]}{s(X)}\frac{s^2(T)}{1+(s^2(T)/s^2(X)))^{1/2}}$$

which, when $s^2(T)/s^2(X)$ is small, is like the conventional z-score heuristic for project management.

The cumulative normal (or Gaussian) utility is concave for values of x less than $E[T]$ and convex otherwise. Thus, when the chances of performance exceeding $E[T]$ are good, the manager will be risk-averse in performance and will only consider a gamble if its expected impact on performance is very positive. But when the chances of exceeding $E[T]$ are poor, this same manager will be more open to taking risks. Note, this is the same pattern of risk-taking behavior as described Kahneman & Tversky's (1979) well-known prospect theory, but with the reference point here ($E[T]$) being unambiguously derivable from the distribution on T.

Thus, the common probability distributions assumed in PM predict that a decision maker using TOU would – and should – display commonly observed behaviors in the pursuit of performance levels. Abbas & Matheson (2005) made a similar observation in the context of aspiration equivalents.

This formulation also has implications for decisions under incremental changes (denoted with δ) induced in $x_{(ce)}$ by changes in $E[X]$ and $s(X)$:

$$\delta x_{(ce)} = s(T)\delta z = r\delta E[X] - (E[X] - E[T]]\delta r = -r[\delta E[X] - (E[X] - E[T])(1 - r^2)\delta s(X)/s(X)]$$

If a manager is considering a decision which could change both the mean and variance of the performance, e.g., cutting corners, that decision will only improve the certainty equivalent if

$$\frac{\delta E[X]}{\delta \ln(s(X))} > (E[T] - E[X])(1 - r^2)$$

Hence the amount of improvement in $E[X]$ required to make a percentage improvement in risk, $s(X)$, viable increases with the amount the project is expected to be below target, $E[T] - E[X]$, and with the amount of risk already in the project.

While many of the examples in this paper assume T Gaussian, this approach can, in principle, be used with any distribution on T. Non-Gaussian distributions could increase the accuracy with which preferences are modeled at the expense of additional modeling complexity.

3 Application to Project Scheduling

Projects are composed of sets of activities related in specified ways. In the "crashing" problem, the alternative plans allocate resources across one or more activities to improve their performance. With a fixed target, Abbasi and Mukattash (2001) incorporated utility functions on this problem. In our setting, as in the case of a project consisting of a single activity, plans should be selected so as to maximize the probability that the combined result of these activities (a function of the performances of the individual activities) will beat some uncertain target.

For instance, in a project with only two activities, a_1 and a_2, there are two uncertainties $X^{(1)}$ and $X^{(2)}$ associated with the execution of the project. The manager might choose a plan that to puts more resources into a_1 and less into a_2, or vice versa, or one that takes a riskier approach with a_1 than a_2, or whatever other possibilities exist.

This is only slightly more complicated than the problem from the previous section.

If we again create a TOUF from the distribution on T, and let g be the function that relates activity performance levels to overall project performance, the project manager maximizes $E[u(g(x^{(1)}(d)), x^{(2)}(d))]$ in the following decision problem:

1. Select a plan $d \in D$
2. Observe the value of $X^{(1)} = x^{(1)}$
3. Observe the value of $X^{(2)} = x^{(2)}$
4. Observe the value of $T = t$
5. Receive value of 1 if $g(x^{(1)}, x^{(2)}) \geq t$

In the remainder of this section, we first develop an example where a TOU-based decision rule can be used with an arbitrary set of connected activities. We then refine this example to obtain more rules for cases where (1) the project consists of parallel activities and (2) the project consists of sequential activities. Finally, we discuss how the latter results relate to the problem of creating suitable reward schedules for decentralized activity managers.

3.1 A Standard Heuristic for Scheduling Activities

We first review the standard heuristic commonly used to solve the fundamental project management problem — how to allocate efforts to activities — and then introduce the simple extension allowed by treating targets as uncertain. The conventional approach involves identifying the final output of the project, listing the various tasks required to reach that output and identifying which activities must be completed before other activities can start. A network is used to describe the precedence relationships between different activities and project completion. Management then specifies the possibly uncertain length of time (and cost) required to complete each activity and how much the expected activity length could be shortened at what cost.

Example 2: Suppose (following Black, 1990) that designing a motorcycle involves five activities labelled $a_1, ..., a_5$. Activity a_1 involves specifying requirements for the axle and wheels and validating the requirements. Activity a_2 involves specifying requirements for the seat and validating the requirements. Activity a_3 involves having a single supplier design the seat, handle, axles and wheels. Activity a_4 involves specifying requirements for the engine and having a supplier design the engine. Activity a_5 involves integrating the engine with the seat/handle/axle/wheel assembly.

The nature of these activities dictates that once the project begins a_1, a_2 and a_4 can start at any time while a_3 can only start when a_1 and a_2 are finished, and a_5 can only start when a_3 and a_4 are finished. The project ends when a_5 is completed. These precedence relationships can be used to define a simple directed graph, often called an activity network in PM.

Let the random variables $X^{(1)}, ..., X^{(5)}$ describe the uncertain times required to complete activities $a_1, ..., a_5$. Then the analyst identifies the three different sequences of consecutive activities (or paths).

We define a path B as a sequence of activities a_j that need to be completed in order before the project is done. In this example, the paths are $\{a_1, a_3, a_5\}$, $\{a_2, a_3, a_5\}$ and $\{a_4, a_5\}$ which have uncertain completion times of $X^{\{1,3,5\}} = X^{(1)} + X^{(3)} + X^{(5)}$, $X^{\{2,3,5\}} = X^{(2)} + X^{(3)} + X^{(5)}$ and $X^{\{4,5\}} = X^{(4)} + X^{(5)}$ respectively. More generally, the uncertain completion time of path B is $X^{(B)} = \sum_{a_j \in B} X^{(j)}$. A standard project management approach is to approximate the $X^{(j)}$ with beta distributions and, applying the central limit theorem, to crudely approximate the sum $X^{(B)}$ as normally distributed.

In the standard PM case, the targeted project completion time is $E[T]$ where T is a constant, and the manager computes the probability of the project being completed by time $E[T]$. If the mean and variance of the completion time for activity a_i are $E[X^{(j)}]$ and $s^2(X^{(j)})$ respectively and if the completion time of each activity is independent, then the mean and variance of each path is easily computed from $E[X^{(B)}] = \sum_{a_j \in B} E[X^{(j)}]$ and $s^2(X^{(B)}) = \sum_{a_j \in B} s^2(X^{(j)})$. The probability of completing the activities on path B by time $E[T]$ is

$$\Pr.\{X^{(B)} \le E[T]\} = \Pr.\{\frac{X^{(B)} - E[X^{(B)}]}{s(X^{(B)})} \le \frac{E[T] - E[X^{(B)}]}{s(X^{(B)})}\}$$

If $X^{(B)}$ is Gaussian, then the random variable $e_B = \frac{X^{(B)} - E[X^{(B)}]}{s(X^{(B)})}$ has the same distribution for all paths. Paths can then be ranked, in order of their probability of not being completed on time by the z-score $\frac{E[T] - E[X^{(B)}]}{s(X^{(B)})}$. In those cases, where one path's probability of completion is significantly lower than all the others, a heuristic can justifiably ignore the other paths and focus efforts on project improvement on this one path — the critical path.

If it is possible to crash certain activities, (typically to shorten their completion time by adding resources), the manager then computes the change in $E[X^{(j)}]$ and $s^2(X^{(j)})$ as well as the associated resource costs of crashing activity i. Since the analyst can compute the resulting impact on the overall probability of project completion, the analyst can then rank initiatives which crash different activities based on the change in the project certainty equivalent per dollar. This is equivalent to ranking activities based on their impact on the z-score per dollar. Thus if a particular activity adjusts the mean completion by $\delta E[X^{(B)}]$ and the standard deviation by $\delta s(X^{(B)})$, then a first-order approximation of the change in z-score is

$$\delta z^{(B)} = -\frac{\delta E[X^{(B)}]}{s(X^{(B)})} + \frac{E[T] - E[X^{(B)}]}{s^2(X^{(B)})}\delta s(X^{(B)}) = -(1/s(X^{(B)}))[\delta E[X^{(B)}] - z^{(B)} ds(X^{(B)})]$$

The right-hand side of the equation above shows whether the change in mean is offset by the change in standard deviation in such a way that the z-score improves and therefore the probability of success improves. Hence the impact of a change on the z-score also defines the amount of added risk $\delta s(X^{(B)})$ which the manager is willing to accept in order to receive an incremental improvement, $\delta E[X^{(B)}]$, in expected performance.

3.2 Modifying the Heuristic to Allow for a Utility Function

As previously noted, utility functions can be introduced into project management by treating T as a random variable with mean $E[T]$ and variance $s^2(T)$. This replacement of T by a random variable can also be motivated by noting that even sponsors who do wish to set fixed targets are sometimes uncertain about what target value is appropriate.

For example, project sponsors in consumer goods industries typically do not know when the bulk of their customers will choose to shop for a new product, although there may be some patterns associated with the seasons. Replacing the fixed target with a random variable changes the probability of completing path B to

$$\Pr.\{X^{(B)} \leq T\} = \Pr.\{\frac{X^{(B)} - T - E[X^{(B)}] + E[T]}{(s^2(T) + s^2(X^{(B)}))^{1/2}} \leq \frac{E[T] - E[X^{(B)}]}{(s^2(X^{(B)}) + s^2(T))^{1/2}}\}$$

so that the z-score can be written as

$$z = \frac{E[T] - E[X^{(B)}]}{(s^2(T) + s^2(X^{(B)}))^{1/2}} \tag{1}$$

which is equivalent to the conventional z-score with the variance of the completion time, $s^2(X^{(B)})$, replaced by the variance of the gap between completion time and target, $s^2(X^{(B)} - T) = s^2(T) + s^2(X^{(B)})$. The first-order change in z-score induced by a change in $E[X^{(B)}]$ and $s(X^{(B)})$ is now

$$\begin{aligned}\delta z &= -\delta E[X^{(B)}]/s(X^{(B)} - T) - \frac{1}{2}\frac{E[T] - E[X^{(B)}]}{[s^2(X^{(B)} - T)]^{3/2}} 2s(X^{(B)})\delta s(X^{(B)}) \\ &= -1/s(X^{(B)} - T)[\delta E[X^{(B)}] - z\frac{s(X^{(B)})}{s(X^{(B)} - T)}\delta s(X^{(B)})]\end{aligned}$$

If $s(X^{(B)} - T) = s(X^{(B)})$, the formula reduces the the previous case of no uncertainty about the target.

However, this extension allows us to handle problems where the uncertainty attached to crashing various activity completion times is different. For instance, suppose the completion times associated with all activities are deterministic and that the project manager expects to fall short of the target. Suppose that crashing activity a_1 will reduce expected completion time by m_1 with a standard deviation of s_1 while crashing activity a_2 will reduce expected completion time by m_2 with a standard deviation of s_2. Suppose s_1 is much smaller than s_2 and m_1 is much smaller than m_2. Then the project manager must trade off expected time reduction against risk. If shortening the completion time by m_1 will guarantee completion of the expected target, the manager — using the conventional project management objective function — should crash activity a_1, even though its expected impact on completion time is less than what he would get from crashing activity a_2. In contrast, introducing enough uncertainty in the target will cause the manager to prefer to crash the second activity, so the uncertain target formulation can lead to different choices and a better chance of meeting the target than would a conventional approach that ignores target uncertainty. This example can be extended to allow for the costs associated with crashing many different activities.

3.3 Generating Utility Functions for Performance on Parallel and Serial Activities

This paper has assumed that the project manager has the same information as activity managers and can unilaterally instruct them. But the same type of problems that can occur in communicating information (and incentives) between project sponsor to project manager may also occur between project manager and activity manager. Then the project manager may want to delegate decision making downward. Again, rather than setting fixed targets for activity managers, the project manager can define new TOUFs with which activity managers can make choices that maximize expected utility from their own, and therefore the project manager's and thus the sponsor's perspective. We consider several elementary cases.

Example 3: Suppose that the activity manager is only rewarded if the project manager is successful. Also suppose that activities $a_1, ...a_n$ are in series. When activity a_k starts, the completion time of preceding activities $a_1, ..., a_{k-1}$ will be known. Let $X^{(B(<k))}$ be the sum of their known completion times. However the sum of the completion times of activities following activity k, will still be uncertain. Let $X^{(B(>k))}$ be the sum of the uncertain completion times of those activities. Then the manager of activity a_k will only be rewarded if activity a_k's completion time, $x^{(k)}$, is less than $T - x^{(B(<k))} - X^{(B(>k))}$. In this case, the activity manager's uncertain target is $T^{(k)} = T - x^{(B(<k))} - X^{(B(>k))}$ with the activity manager's utility function, $u(X^{(k)}) = \Pr.\{X^{(k)} \leq T^{(k)}\}$.

Since the uncertainty in $X^{(B(>k))}$ is greater when k is small, earlier activities will have higher variance targets and more risk-neutral utility functions than later activities. In turn, managers of later activities will have incentives for more pronounced performance risk taking to try to save a lagging project or performance risk aversion to ensure that an ahead-of-schedule project doesn't slip behind. That is, a later stage activity manager might accept gambles that an early stage manager might reject and vice versa. Thus, even if the target itself is known, the nature of the activity network may still induce a utility function on activity managers.

This example focuses on activities in series. But suppose we focus on a single activity of interest. Then we do not need to assume that these other activities are in series. All that we need to assume is that some activities must be finished before the activity of interest starts while other activities cannot start until the activity of interest finishes. We can then consider the first set of activities together as a single higher level activity a_1 preceding the activity of interest, a_2, followed by the subsequent activities also considered together as a single higher level activity a_3. Then the probability of the project finishing is the probability that $x(1) + X^{(3)} + X^{(3)}$ is less than the uncertain deadline which yields a new uncertain target activity utility for the activity of interest, $T^{(2)} = T - x(1) - X^{(3)}$, and thus induces a utility function for this activity $u(X^{(2)}) = \Pr.\{X^{(2)} \leq T^{(2)}\}$.

Example 4: Now consider a different case with two activities in parallel, with unknown and stochastically independent completion times $X^{(1)}$ and $X^{(2)}$, both of which must be completed before the target deadline. Here, the project manager is rewarded only when $t \geq \max(x^{(1)}, x^{(2)})$. The first activity manager's utility for completing a project in time $x^{(1)}$ will be proportional to the probability of T exceeding both $x^{(1)}$ and $X^{(2)}$, i.e.,

$$E[u(x^{(1)})] \propto \Pr.\{x^{(1)} \leq T, X^{(2)} \leq T\} \propto \Pr.\{x^{(1)} \leq T | X^{(2)} \leq T\} \Pr.\{X^{(2)} \leq T\}$$

Hence the first activity manager maximizes $\Pr.\{x^{(1)} \leq T | X^{(2)} \leq T\}$ which exceeds $\Pr.\{x^{(1)} \leq T\}$, and so this manager will be be less aggressive with regard to $X^{(1)}$ than the project manager is with regard to X.

Example 4b: Alternatively suppose the project is successful when either of the activities is completed before the deadline, i.e., the first activity manager gets rewarded if either her completion time $x^{(1)}$ is less than T or the completion time of the second activity is less than T. Thus the first activity manager only fails to get rewarded if both her completion time and the other activity's completion time exceed T. Thus the probability of getting rewarded is for achieving $x^{(1)}$

$$1 - \Pr.\{x^{(1)} \geq T, X^{(2)} \geq T\} = 1 - \Pr.\{x \geq T | X^{(2)} \geq T\} \Pr.\{X^{(2)} \geq T\}$$

Hence the first activity manager maximizes the probability of getting rewarded by minimizing $\Pr.\{x \geq T | X^{(2)} \geq T\}$, i.e., maximizing $\Pr.\{x \leq T | X^{(2)} \geq T\}$ which is less than or equal to $\Pr.\{x \leq T\}$. Hence the manager needs a greater value of x to achieve the target. As a result, the manager's target for success must be more aggressive.

Since the second activity manager likewise maximizes $\Pr.\{x \leq T | X^{(1)} \geq T\}$, it is straightforward to compute the probability of the deadline being met given the optimal decisions of the two uncoordinated activity managers. But suppose the project manager were able to coordinate the decisions of both activity managers to optimize the overall probability of meeting the deadline. The improvement obtainable through coordination would then be simply the difference between this probability and the probability when the activity managers' decisions are not coordinated.

Example 4c: Alternatively, if, as in the previous section's example activities a_1 and a_2 are in parallel while the rest of the activities combined have a normal completion time, then TOU could still be used but would be more complicated. Here the distribution of activity completion times is the distribution of $\min(X^{(1)}, X^{(2)})$. Given the normality assumption, closed form expressions exist for the mean and variance of the maximum of two Gaussian random variables (Clark,1961). Specifically if $z = \frac{E[X^{(1)}] - E[X^{(2)}]}{s(X^{(1)} - X^{(2)})}$, then the expected value of the mean (and the square of the mean) of the maximum of two Gaussian variables is

$$E[\max(X^{(1)}, X^{(2)})] \quad = \quad E[X^1] \Pr.\{z \leq 0\} + E[X^{(2)}](1 - \Pr.\{z > 0\})$$

$$+ \quad \frac{s(X^{(1)} - X^{(2)})}{2(\pi)^{1/2}} \exp(-\frac{z^2}{2})$$

and

$$\begin{aligned} E[\max(X^{(1)})^2, X^{(2)}))^2] &= E[(X^{(1)})^2] \Pr.\{z \le 0\} + E[(X^{(2)})^2](1 - \Pr.\{z > 0\}) \\ &+ (E[X^{(1)} + X^{(2)}]) \frac{s(X^{(1)} - X^{(2)})}{2(\pi)^{1/2}} \exp(-\frac{z^2}{2}) \end{aligned}$$

Replacing these two activities with this combined activity with the resulting mean and variance then allows us to use the previous analysis (if it is acceptably accurate to approximate the maximum of two Gaussian random variables with another Gaussian random variable.)

Example 4d: We can generalize these results by considering again sub-networks as higher level activities, and by combining serial and parallel processes. As an example, consider the case where the project can be organized into a sub-network of activities (designated a_1) that must finish before a particular manager's activity (a_2) can begin, and another sub-network of activities (a_3) that can be executed in parallel with a_2. Since the sub-network of activities that finishes before a_2 has a known completion time, $x^{(1)}$, the time from the start of the project to the completion of a_2 is $x^{(1)} + X^{(2)}$. The project finishes on time if both $x^{(1)} + X^{(2)} \le T$ and $X^{(3)} \le T$. Hence, the activity manager should maximize the expectation of $u(X^{(k)}) = \Pr.\{X^{(2)} \le T - x^{(1)} | X^{(3)} \le T - x^{(1)}\}$.

Alternatively, activity a_1 might precede sub-network a_2, while both of these proceed in parallel with sub-network a_3, and the activity manager would maximize the expectation of $u(X^{(1)} = \Pr.\{X^{(1)} + X(2) \le T | X^{(3)} \le T\}$.

By variations on such iterative construction, utility functions could thus be derived for the individual activity managers in an activity network consisting of a mix of parallel and series activities (e.g., in the motorcycle example, activities a_1 and a_2 are in parallel and they are followed by activities a_3 and a_5 in series). Thus, TOU could be used to communicate the sponsor's preferences down through the organization, as responsibilities are delegated through a work breakdown structure.

It allows even activity managers to choose among plans to complete their activities in a way that maximizes the sponsor's expected utility, maximizing their own expected reward, but deciding as if they are rewarded when the larger project succeeds.

4 Project Balancing

To this point, we have discussed situations where the target is a single scalar value. Project management typically focuses on more than one aspect of performance, most often cost, quality and timing. Project managers may be faced with decisions that trade off performance on one of these attributes against the other. The project manager may choose an approach that is likely to be slower that would also be more if it were also likely to result in higher quality. We adapt Verzuh's (2011) distinction between project level balancing (where tradeoffs are only used in deciding how to try to meet the different requirements), and business case level balancing (where it is possible to make tradeoffs around what it means to 'meet' requirements).

As before, we consider a project manager concerned with choosing a plan whose performance is likely to exceed the customer requirements. The manager must choose from a set of alternatives where no alternative dominates the others on all performance dimensions.

4.1 Project Level Balancing: Tradeoffs between Targets are Not Allowed

Let the single uncertainty about project performance X be decomposed into uncertainty about time, quality and cost (random variables X_1, X_2 and X_3), as the uncertainty about the target T is likewise decomposed (into random variables T_1, T_2 and T_3).

The project manager faces the following decision problem:

1. Select a plan $d \in D$

2. Observe the value of $X = (x_1, x_2, x_3)$

3. Observe the value of $T = (t_1, t_2, t_3)$

4. Receive value of 1 if $x_1 \geq t_1, x_2 \geq t_2, x_3 \geq t_3$

Again, we would like to formulate a utility function over multiple performance metrics whose maximization is equivalent to maximizing the probability of a successful project. Here, we explore for several cases how such a utility function over the X terms can be constructed to to reflect the joint distribution of the T terms.

Assume criteria can be transformed so that the total performance measure on each dimension is simply the sum of performance measures in that dimension over each of several activities. For instance, project cost can easily be written as the sum of the cost of each required activity. On the other hand, defining project completion time as the sum of adjusted completion times for each activity requires that we define the adjusted completion time of an activity off the critical path as zero (with adjusted completion time equalling actual completion time for activities on the critical path.) To define an additive measure for project quality, assume that the probability of a project being completed flawlessly is the product of the probability of each activity being completed flawlessly. Then defining quality as the logarithm of the probability of flawlessness implies that project quality is the sum of activity quality.

We first consider a situation where the manager cannot make tradeoffs between performance on these various goals. This can arise when a manager is responsible to different departments for her project's performance on budget, quality and completion time. These goals could apply to a single common dimension for multiple parallel projects, as in Hill and Khosla's (1992) extension with simpler time-cost tradeoffs, or, more generally, different project dimensions such as product quality.

Example 5: Suppose there are only two separate goals (e.g., time, X_1 and cost, X_2) and a known relationship between these random variables with $g(X_1) + X_2 = k$ for some constant k and some strictly monotonic increasing function g. Then we can set $X_2 = k - g(X_1)$ and the utility function is

$$u(X_1, X_2) = \Pr.\{X_1 \geq T_1, k - T_2 \geq g(X_1)\}$$

Defining upper and lower bounds $T_{(U)} = g^{-1}(k - T_2)$ and $T_{(L)} = T_1$ gives

$$u(X_1, X_2) = \Pr.\{T_{(U)} \geq X_1 \geq T_{(L)}\}$$

and the manager maximizes the probability that project completion time will lie in this random interval. Define the midpoint $T_{(av)} = \frac{T_{(U)}+T_{(L)}}{2}$ and the spread $T_{(diff)} = \frac{T_{(U)}-T_{(L)}}{2}$. Then for any strictly monotonic increasing function h,

$$u(X_1, X_2) = \Pr.\{T_{(diff)} \geq X_1 - T_{(av)} \geq -T_{(diff)}\} = \Pr.\{h(|T_{(diff)}|) \geq h(|X_1 - T_{(av)}|)\}$$

Defining $X = h(|X_1 - T_{(av)}|)$ as our performance measure and $T = h(|T_{(diff)}|)$ as our uncertain target gives our utility as $\Pr.\{X \leq T\}$ so that our target T is defined over deviations between X_1 and $T_{(av)}$. If $X - T$ is normal, then we can evaluate scenarios using $\frac{E[X]-E[T]}{s(T-X)}$.[1]

Example 5b: For three performance metrics (defined to be negatively oriented in this example), the utility function corresponding to the probability of meeting all three criteria is the multivariate cumulative probability:

$$u(X_1, X_2, X_3) = \Pr.\{X_i \leq T_i; i = 1...3\}$$

[1]Although in this paper, we assume the T and X terms are independent, T and X are sometimes correlated, e.g., factors that might make performance difficult for the project team could also make performance difficult for competitors, thereby lowering the target. when focusing on utility, assume ind, when focusing on target meeting, then not required. Although we assumed independence for simplicity of results, most of them could be redone without much difficulty to accommodate correlation. While such formulations would typically preclude the TOU interpretation, this would not affect the correctness of the calculated probability that targets will be met.

For analytic simplicity, we return to the simpler assumptions that each T_i follows an exponential distributions with scale parameter β_i and the T_i are are stochastically independent of one another, so that the utility for a known set of performances x_1, x_2, x_3 is:

$$u(x_1, x_2, x_3) = \Pr.\{T_i \geq x_i; i = 1...3\} = \prod_{i=1}^{3} \exp(-\frac{x_i}{\beta_i)}) = \exp(-\sum_i \frac{x_i}{\beta_i)})$$

If $X_i, i = 1, ...3$ are Gaussian with means $E[X_i]$ and covariances V_{ik}, then

$$u(X_1, ...X_n) = 1 - \exp(-\sum_i \frac{E[X_i]}{\beta_i} + \frac{1}{2}\sum_{ij} \frac{V_{ik}}{s(X_i)s(X_k)})$$

Since the T_i are exponential, $u(0, \ldots, 0) = 1$ and $u(\infty, 0, \ldots, 0) = u(\infty, \infty, \ldots, \infty) = 0$, and u is continuous in x. We can define a certainty equivalent [2] by selecting one performance dimension, say the first, and using it as a reference, i.e., $x_{(ce)} = (x_{1(ce)}, 0, \ldots, 0)$ such that $u(x_{(ce)} = E[u(X)]$. Then

$$x_{1(ce)} = \beta_1(\sum_i \frac{E[X_i]}{\beta_i} - \frac{1}{2}\sum_{ij} \frac{V_{ik}}{s(X_i)s(X_k)})$$

Frequently a manager is presented with an option which will improve project's completion time or performance on some other attribute while introducing risk. The manager needs to know the threshold at which the expected improvement is not worth the added risk. This is the point at which the change in mean performance and standard deviation of performance lead to zero improvement in expected utility (and thus in the certainty equivalent.) Thus consider an intervention which leads to a one unit increase in $E[X_j]$ and an increase in V_{ik} of q. The one unit increase in $E[X_j]$ changes the certainty equivalent by $\frac{\beta_1}{\beta_j}$ while the change in V_{ik} changes the certainty equivalent by $-\beta_1 \frac{q}{s(X_i)s(X_k)}$. The project manager should be indifferent about the intervention if $q = \frac{s(X_i)s(X_k)}{\beta_j}$. So in the Gaussian case, determining whether an intervention leads to an overall improvement in the certainty equivalent is straightforward.

[2] Abbas & Matheson (2009) give a more general formulation of the notion of a certainty equivalent in the context of multi-attribute TOU.

The next variation is motivated by the fact that Gaussian distributions may be too restrictive for the range of possible performance measures. We weaken assumptions about the distribution by using the Weibull (whose density function is unimodal with a mode that can be greater than zero) and still obtain tractable results comparable to those of the (somewhat less flexible) exponential distribution whose density function is unimodal with its mode at zero.

Example 5c: If T_i is approximated with a shifted Weibull distribution with lower bound γ_i, scaling factor β_i and shape parameter α_i, then

$$u(x_1, x_2, x_3) = 1 - \prod_i \exp(-(\frac{x_i - \gamma_i}{\beta_i})^{\alpha_i}) = 1 - \exp(-\sum_i (\frac{x_i - \gamma_i}{\beta_i})^{\alpha_i})$$

Example 5d: Suppose in addition that each T_i has been linearly transformed so that $\gamma_i = 0$ and $\beta_i^{\alpha_i} = 2$, i.e., that T_i follows a Rayleigh distribution. The disutility is $1 - u(x_1, x_2, x_3)$ is just $\exp(-\sum_i x_i^{\alpha_i}/2)$. Assuming X_i Gaussian (as before) and defining I as the identity matrix, $\mathbf{V}$ as the covariance matrix of X, W as $[I + \mathbf{V}^{-1}]^{-1}$ and H as $W\mathbf{V}^{-1}E[X]$, the expected disutility is then

$$\begin{aligned}
1 - E[u(X_1, X_2, X_3)] &= |V^{-1}| \int \exp(-\frac{1}{2}[x^T x + (x - E[X])\mathbf{V}^{-1}(x - E[X])])dx \\
&= |V^{-1}| \int \exp(-\frac{1}{2}[x^T[I + \mathbf{V}^{-1}]x - 2x^T V^{-1}E[X] + E[X^T]\mathbf{V}^{-1}E[X]) \\
&= |\mathbf{V}^{-1}| \int \exp(-\frac{1}{2}[(x - H)^T W^{-1}(x - H) - HW^{-1}H + E[X^T]\mathbf{V}^{-1}E[X]]dx \\
&= |\mathbf{V}^{-1}W| \exp(\frac{1}{2}E[X^T][\mathbf{V}^{-1}WW^{-1}W\mathbf{V}^{-1} - \mathbf{V}^{-1}]E[X]) \\
&= [I + \mathbf{V}]^{-1} \exp(\frac{1}{2}E[X^T]\mathbf{V}^{-1}[W\mathbf{V}^{-1} - I]E[X]) \\
&= [I + \mathbf{V}]^{-1} \exp(-\frac{1}{2}[EX^T\mathbf{V}^{-1}WE[X]) = [I + \mathbf{V}]^{-1} \exp(-\frac{1}{2}E[X^T](I + \mathbf{V})^{-1}E[X])
\end{aligned}$$

This formula for utility of multi-dimensional performance generalizes the Gaussian case and, per Tsetlin & Winkler (2007), accounts for the potentially critical correlation between attributes.

4.2 Business-Case-Level Balancing: Tradeoffs between Targets are Allowed

In the previous section, a customer's tradeoffs between different levels of performance were based on the relative importance of meeting different requirements on time, cost, performance, etc. But as Tsetlin and Winkler (2006) noted, there are some multi-attribute preferences which cannot be described in this way. Instead a value function must be defined over known levels of cost, completion time and performance which reflects tradeoffs. (There is an extensive literature on estimating weights in the value function.) A utility function is then defined over the value function reflecting the customer's attitudes to risk. Specifically suppose the value, $X_{(total)}$, is a weighted average of performance on the three dimensions of cost, completion time and performance.

This leads to a different type of TOUF for cost, quality and time along the lines of which will allow utility maximization to solve the following decision problem:

1. Select a plan $d \in D$
2. Observe the value of $X = (x_1, x_2, x_3)$, and calculate $x_{(total)} = w_1 x_1 + w_2 x_2 + w_3 x_3$
3. Observe the value of $T = (t_1, t_2, t_3)$ and calculate $t_{(total)} = w_1 t_1 + w_2 t_2 + w_3 t_3$
4. Receive value of 1 if $x_{(total)} =\geq t_{(total)}$

If $X^{(A)}$ is the sum of the scores over various activities j, $a_j \in A$, then $X_i^{(A)} = \sum_j X_i^{(j)}$ and $X_{(total)}^{(A)} = \sum_{ij} X_i^{(j)} w_i$. The mean and variance of X is

$\sum_{ij} E[X_i^{(j)}] w_i$ and $\sum_{ijk} Cov(X_i^{(j)} w_i, X_k^{(j)} w_k)$.

If there is no correlation between different activities, the variance simplifies to $\sum_{ijk} w_i Var(X_i^{(j)}, X_k^{(j)}) w_k$.

If there is no correlation between attributes, this further simplifies to $\sum_{ij} w_i^2 Var(X_i^{(j)})$.

As before, we introduce a random variable T which represents the uncertain level of $X_{(total)}^{(A)}$ which the manager is trying to achieve (e.g., the overall performance of a key competitor.) We can now apply as a decision rule the solution of section 3 (for the manager trying to meet an uncertain deadline) using $X_{(total)}^{(A)}$ to represent the overall performance measure instead of the project completion time X.

With the results from this section, project plans can be modified to increase chances of success. This can be done at the start of the project (more likely as business case level balancing: picking a different combination of performance levels in order to match the overall performance target), or in the middle of a troubled project (more likely as project-level balancing) by changing the relative likelihood of meeting the various targets.

5 Multi-stage Decision-Making

To this point, we have focused on using utility functions to compare and select alternatives in single-stage decision problems by properly accounting for risk and uncertainty. In multi-stage decisions, DA techniques, especially the concept of "value of information," can also help decision makers anticipate and properly account for information that might be obtained before some later stage choices. This applies in project management where it is often possible to build in additional decision points allowing for course corrections during the project. For instance, in product development projects, initial plans may contain valuable real options (Huchzermeier & Loch, 2001, Keisler & Mang, 2011) if they build in flexibility to modify the product design or even abandon the project if it becomes known that it is too technically challenging or that the business environment changes before completion. We convert the project from section 2.1 into a two-stage (or multi-stage) decision problem where it is possible to acquire information in between the stage 1 choice and the later stage choice(s). In general, the decision problem will be as follows:

1. Select a first stage strategy

2. If the strategy includes information acquisition, obtain partial or perfect information Ψ and update distributions on X and or T

3. Select a second stage decision

4. Observe the value of $X = x$

5. Observe the value of $T = t$

6. Receive value of 1 if $x \geq t$

In 5.1, we describe how expected TOU maximization can work in a generic case where the project manager anticipates downstream project decisions based on information that may be acquired. In 5.2, we develop a realistic example where it is possible to develop a testing plan to, at a cost, obtain information about system performance during the project. In 5.3, we assume system requirements are being met, but consider how the project might plan to validate the system against customer requirements, e.g., through focus groups.

5.1 Downstream Decisions and Value of Information in Project Management

Decision analysis can support the project manager in doing more than merely deciding to approve an initial project plan. For instance, if certain aspects of the project are running late, management may reallocate resources in order to accelerate various later stage activities (Kavadias & Loch, 2003). Management can also order rework, i.e., allocate resources to reworking activities previously done (Ahmadi & Wang, 1999, Smith & Eppinger, 1997). Such decisions will only be made with the addition of the information set (Ψ) available at the kth stage.

Example 7: We consider a project's stages as sequential activities $J = \{a_1, ..., a_n\}$, and a_j denotes the jth activity. Thus at the start of the project, the firm's uncertainty about the project's future performance upon completion is described by the random variable $X^{(J)} = \sum_{j=1}^{n} X^{(j)}$. At the kth stage, based on monitoring of earlier stages of the project , the manager may now believe that $\sum_{j=1}^{k} X^{(j)}$, which previously had mean $\sum_{j=1}^{k} E[X^{(j)}]$ and variance $\sum_{j=1}^{k} s^2(X^{(j)})$, now has mean $\sum_{j=1}^{k} E[X^{(j)}|\Psi]$ and variance $\sum_{j=1}^{k} s^2[X^{(j)}|\Psi]$. The manager can now recompute the probability of meeting project objectives as the probability that $z(\Psi) > e$ where

$$z(\Psi) = \frac{\sum_{j=1}^{k} E[X^{(j)}|\Psi] + \sum_{j>k} E[X_j] - E[T]}{[s^2(T) + \sum_{j=1}^{k} s^2(X^{(j)}|\Psi) + \sum_{j>k} s^2(X^{(j)})]^{1/2}}$$

If we let $s^2(X)$ and $E[X]$ be the original mean and variance for the project at time 0 and define $E[X|\Psi] = \sum_{j=1}^{k} [E[X^{(j)}|\Psi - E[X^{(j)}]]$ and $s^2(X|^{\Psi}) = \sum_{j=1}^{k} [s^2(X^{(j)}) - s^2(X^{(j)}|\Psi)]$, then

$$z = \frac{E[X] - E[T]}{[s^2(T) + s^2(X)]^{1/2}}$$

Recalling the risk factor from section 3,

$$r = \frac{s(T)}{\sqrt{s^2(X) + s^2(T)}}$$

we can also write

$$s(T)z = r(E(X) - E(T))$$

Similarly if

$$z(\Psi) = \frac{E[X] - E[T] + E[X|\Psi]}{[s^2(T) + s^2(X) - s^2(X|\Psi)]^{1/2}}$$

then defining $\lambda = [1 - \frac{s^2(X|\Psi)}{s^2(X)+s^2(T)}]^{-1/2}$ so that

$$r\lambda = \frac{s(T)}{[s^2(T) + s^2(X) - s^2(X|\Psi)]^{1/2}}$$

implies

$$s(T)z(\Psi) = \frac{E[X] - E[T] + E[X|\Psi]}{[s^2(T) + s^2(X) - s^2(X|\Psi)]^{1/2}} = r\lambda[E[X] - E[T] + E[X|\Psi]$$

As a result,the change in z-score can be written as

$$s(T)z(\Psi) = (E[X] - E[T])r\lambda + E[X|\Psi]r\lambda = s(T)[z + r[\lambda E[X|\Psi] + (1-\lambda)(E[T] - E[X])]$$

Likewise the change in certainty equivalent is $s(T)(z(\Psi) - z)$ or

$$r[\lambda E[X|\Psi] + (1-\lambda)(E[T] - E[X])]$$

Since $s^2(X|\Psi) > 0$, λ will exceed one.

Since $s^2(X|\Psi) > 0$, λ will exceed one, if there were no change in mean performance, the information would increase the certainty equivalent if the project were expected to exceed its target $(E[X] - E[T])$ and would reduce the certainty equivalent if the project were expected to fall short of its target. If the manager expects to meet (miss) the target, small decreases (increases) in mean performance could be offset by greater certainty.

Of course, when there is no decision involved, the expected value does not change merely because we explicitly enumerate the possible results of uncertainty resolution. For example, if the manager expects to miss the target, then as uncertainty is resolved, the change in expected value is asymmetric in the change in the expected performance — an increase in expected performance leads to a greater change in expected value than the loss associated with the same magnitude of decrease in expected performance, and the additional size of the potential gain also offsets the decrease in expected value that occurs in the case when uncertainty is reduced and the mean remains unchanged.

Reworking earlier activities will typically increase the mean completion time and cost of the project while possibly improving quality. Computing the time, quality and cost implications of rework is not straightforward since reworking an activity may involve reworking all later activities which depended on inputs from the activity. Furthermore, reworking those later activities may then induce new errors. The Design Structure Matrix (DSM, see Steward, 1981, Browning, 2001) provides a natural way of estimating the cost, time and quality implications of rework. A typical DSM might specify the probability that reworking a specific activity (e.g., engine design) requires rework in each of the other activities in the project (e.g., chassis design).

Once these potential decisions are translated into changes in the mean and variance of X, formulas discussed in the previous section describe how we can potentially evaluate alternative decisions. If $z(\Psi)$ has substantially worsened compared to z, then it is possible that no decision will be able to improve $z(\Psi)$ so as to yield an acceptable probability of success to justify the cost of continuing. In this case, the project manager might choose to cancel the project. Or, alternatively, $z(\Psi)$ might have improved so much relative to z as to make rework unnecessary (although management might choose to adjust the resources allocated to activities in later stages.) Thus in some cases, we can define an upper bound and a lower bound which define three ranges corresponding to: advancing the project to the next stage, requesting project revision, or canceling the project.

5.2 Engineering Testing

In system development projects, testing allows managers to learn how close the system is to meeting its specifications. The earlier that problems are identified, the more easily they can be corrected. Because testing takes time and money, it should only be done if it has enough to improve the project's trajectory. This is essentially a value of information problem, as in 5.1. Partial information will be obtained about the performance of the project X prior to selecting a plan to finish the project that maximizes $\Pr.\{X \geq T\}$.

We previously considered the case where management at the jth stage received information for free, and how this information would be anticipated in earlier stage decisions. For instance, a manager typically learns a lot about activity completion times by observing which activities were actually completed by the time of the jth stage review.

But in the case of quality, management typically does not have have clear information about the quality of the deliverables resulting from activity k. Sometimes such information can be gathered by commissioning a peer review. Sometimes it can be gathered by actually creating a prototype of the partly completed deliverable and subjecting it to tests.

Peer review will certainly have direct cost for the staff involved. If the validation activity can be conducted off the project's critical path, then it will not impact project completion time. But in many cases, a good prototype can only be constructed after completion of activities on the critical path. Thus the validation activity typically occurs on the project's critical path and then necessarily consumes time. As a result, validation can increase both cost and time in order to provide information at stage j which may lead to decisions improving product quality that are superior to those that would have been made without this additional information.

Note, for instance, that the performance of a product is a combination of the performance of the product if manufactured exactly to the engineered specifications and the deviations from that performance arising from inherent manufacturing variability. Similarly the performance of a prototype is a combination of its performance if manufactured exactly to the required specifications and the deviations arising from manufacturing variability. Hence if the prototype fails an engineering test, this failure could reflect imperfections in the prototype and not in the engineering design. But if the prototype is created to parallel how the product will be manufactured, failure in the prototype may suggest that the design is not easily manufactured,e.g., it requires unreasonably tight tolerances or involves assembly procedures that are overly demanding for human assemblers.

We now develop a specific example for which we can derive decision rules for planning tests to validate the technical performance of a product. Such tests will not give management perfect information on the product although they should the variance associated with management's prediction of product performance.

Example 8: Suppose the manufacturing variability for a single test item is e_m with mean zero and variance σ^2. Then the variability in the performance of the prototype design is $X + e_m$ where X reflects uncertainty in the performance of the design. If the variance on X was $s^2(X)$ prior to the test, and if tests allow direct measurement of the performance of N independent prototypes, then the new uncertainty about X is $s^2(X_{new}) = (\frac{1}{s^2(X)} + \frac{N}{\sigma^2})^{-1}$. Defining $w = \frac{s^2(X)}{\frac{\sigma^2}{N}+s^2(X)} = \frac{N}{N+\frac{\sigma^2}{s^2(X)}}$ implies $s^2_{new}(X) = (1-w)s^2(X)$.

Given an average test result of $\hat{x}$, the new mean estimate of X is

$$\frac{\frac{E[X]}{s^2(X)} + \frac{\hat{x}}{\sigma^2/N}}{\frac{1}{s^2(X)} + \frac{1}{\sigma^2/N}} = E[X] + (\hat{x} - E[X])\frac{s^2(X)}{s^2(X) + (\sigma^2/N)} = E[X] + w(\hat{x} - E[X])$$

Define $\Delta = \hat{x} - E[X]$ and

$$z(Delta) = \frac{E[T] - E[X] - w(\hat{x} - E[X])}{[(1-w)s^2(X) + s^2(T)]^{1/2}}$$

The expected utility given the test result is then $\Pr.\{z(\Delta) \geq e\}$. Integrating this expected utility, over all possible test results, yields the original expected utility prior to doing the test.

Example 8b: Now suppose that doing rework changes X to $X_{(Q)}$ which is Gaussian with mean $E[X_{(Q)}]$ and variance $s^2(X_{(Q)})$. Let $z(Q)$ be the z-score associated with rework, i.e.,

$$z(Q) = \frac{E[T] - E[X_{(Q)}]}{(s^2(X_{(Q)}) + s^2(T))^{1/2}}$$

So the utility of doing rework is $\Pr.\{z(Q) > e\}$. Since we only do rework if $z(Q) > z(\Delta)$, our expected utility if we observe the test result Δ and then act optimally is

$$\Pr.\{\max(z(Q), z(\Delta) > e\}$$

This will, of course, not be less than the utility of always doing rework, $\Pr.\{z(Q) > e\}$ or the utility of observing the test result and proceeding with our original plans, $\Pr.\{z > e\}$.

The certainty equivalent $x_{(ce*)}$ will satisfy

$$\begin{aligned}\Pr.\{\frac{x_{(ce*)} - E[T]}{s(T)} > e\} = \Pr.\{\max(z(\Delta), z(Q)) > e_m\} &= 1 - \Pr.\{\max(z(\Delta), z(Q)) < e\} \\ &= 1 - \Pr.\{z(\Delta) < e\}\Pr.\{z(Q) < e\}\end{aligned}$$

Let $G(y) = \Pr.\{e > y\}$. Then

$$1 - G(\frac{x_{(ce*)} - E[T]}{s(T)}) = 1 - G(z(\Delta))G(z(Q))$$

The inverse of G is a readily accessible function so that we can compute

$$x_{(ce*)} = E[T] + s(T)G^{-1}[G(z(\Delta))G(z(Q))]$$

Since the certainty equivalent before validation was just $x_{(ce)} = E[T] + \frac{E[X]-E[T]}{(s^2(X)+s^2(T))^{1/2}}$, the difference between the two gives the value of information.

Example 8c: Suppose in the above case the cost of running N validation tests is $K(N)$. To determine how much validation to conduct, we need to adjust $z(Q)$ and $z(\Delta)$ as follows

$$z(Q;N) = \frac{E[T] - E[X_{(Q)}] - K(N)}{(s^2(X_{(Q)}) + s^2(T))^{1/2}}, \quad z(\Delta;N) = \frac{E[T] - E[X] - w\Delta - K(N)}{[(1-w)s^2(X) + s^2(T)]^{1/2}}$$

The optimal number of tests is that N optimizing $X_{(ce*)}$ and thus $[G(z(\Delta;N)G(z(Q;N)]$. If the test yields information on multiple dimensions, we can still use the vector of G and G^{-1} values corresponding to these dimensions, with a correlation term if necessary.

In addition to specifying how many tests to run, this calculation can determine when the tests should be run (following Ha & Porteus, 1995). If a test is run early, then the cost of rework is small. However the accuracy of the test is lower since the design being tested is more unfinished. In this case, w is smaller and the test is less credible. Alternatively, running the test later ensures a more credible test (larger w) but the costs of rework are higher. Thus, by including the effect of product design on the project's expected utility, it becomes possible to perform value of information calculations which can influence both the timing and the extent of validation.

5.3 Market Testing: Validating that the Solution Satisfies Customer Requirements)

As the project advances toward completion, new information may become available about whether the product specifications are still likely to meet customer requirements. Conceptually, this is similar to the system testing example, except that here the project manager faces a decision about whether to obtain information about the uncertain target T rather than about X before selecting a plan for the remainder of the effort.

As information is revealed over time, $E[T]$ changes and $s(T)$ typically shrinks. Sometimes this information is free, e.g., a competitor announces its intention to market a product with a certain level of quality and cost. Sometimes companies will conduct real-time market research, e.g., customer clinics or focus groups. Just as updating information on product performance causes management to adjust its decisions, and possibly rework new activities, so this updated information on market needs can cause a revisiting of all these decisions. The formulas previously developed allow us to make tradeoffs between possible courses of action. As always, the critical question is properly quantifying the costs and time-impacts of changes versus the benefits of better aligning performance with updated market needs.

Updating the uncertain threshold, T, is theoretically equivalent to modifying the utility function applied to the project. Typically, information which reduces the uncertainty about market requirements should make the manager more risk-sensitive, i.e., more inclined to avoid gambles if the product seems to be exceeding requirements and more inclined to take gambles if the product seems to fall short of requirements.[3] But while the utility function becomes more risk-sensitive

[3] A highly engaged project sponsor might at this point want to change the utility function to be used in the remaining decisions, in which case the project manager could receive credit for earned value (Kim et al, 2003), and then be rewarded based on the new utility function for the remainder of the project.

with time, this is coupled with decreasing risk in the product which is being managed.

Example 9: Suppose that new information on product performance and on the market requirement arrives in such a way that their expected values follow a random walk. Then the reduction in variance on both will decrease linearly as the project progresses. If σ_X and σ_T are the volatility associated with the performance and the requirements respectively and if $s^2_{\tau 0}(X)$ and $s^2_{\tau 0}(T)$ denote the initial uncertainty about the performance and requirement, then $s^2_{\tau}(X) = s^2_{\tau 0}(X) - \tau\sigma_X$ and $s^2_{\tau}(T) = s^2_{\tau 0}(T) - \tau\sigma_T$. Hence the certainty equivalent at time τ is an affine transformation of

$$\frac{E[X] - E[T]}{1 + \frac{s^2_{\tau}(X)}{s^2_{\tau}(T)}} = \frac{E[X] - E[T]}{1 + \frac{s^2_{\tau 0}(X) - \tau\sigma_X}{s^2_{\tau 0}(T) - \tau\sigma_T}}$$

Since the uncertainty about the requirements $s^2_{\tau}(T)$ changes over time, this implies that the utility function defined by T also changes over time. While there is nothing inconsistent about utility preferences predictably changing over time, unstable preferences differ from what is commonly assumed in practice. But if all uncertainty disappeared at some time period τ_F, then $s^2_{\tau 0}(X) = (\tau_F - \tau_0)\sigma_X$ and $s^2_{\tau 0}(T) = (\tau_F - \tau_0)\sigma_T$. As a result, the certainty equivalent would be an affine transformation of

$$\frac{E[X] - E[T]}{1 + \frac{\sigma_X}{\sigma_T}}$$

which does not change systematically over time. Hence, if a project manager in practice were to create project planning decision trees, it might be reasonable to assume a stable utility function.

Even if this condition doesn't hold exactly, the competing effects will mean that original utility function will usually be fairly close to the post-test utility function. Thus, practical decision tree analysis can be conducted early in the project using the original TOU function, with confidence that the project manager's actions will continue to maximize the project sponsor's expected utility throughout the project.

6 Conclusion

The project management process features a variety of planning decisions, ranging from upfront design to in-process adjustments. These decisions often require tradeoffs involving the probability distribution for performance within and across dimensions. Uncertain requirements can render existing fixed-requirement PM tools inadequate. Our contribution starts with the insight that TOU fits rather naturally with PM, and thus might facilitate the inclusion DA within PM. We propose that PM add assessment of uncertain targets on top of its current planning techniques, e.g., defining work breakdown structures and activity networks, and estimating mean and variance of performance for activities (or more generally, full distributions). A utility function for performance is then based on the distribution of the target, replacing the implicit binary function determined by success in meeting fixed targets. Many resulting heuristics for maximizing expected utility can be understood as modified versions (including suitable adjustments for uncertainty) of heuristics for maximizing the chance of meeting fixed requirements.

We conceive of several common PM problems as sequential decision problems containing uncertainty about performance and targets, and identify statistics and decision rules for simple (variants of Gaussian) versions of these problems:

- selecting among alternative project plans, by calculating expected utility, certainty equivalent and risk premium;
- selecting among alternative plans (or changes to plans) affecting separate activities in a generic network, in parallel, or in series, possibly allowing for distributed PM;
- making tradeoffs in achieving performance on cost, quality and time aiming to meet either a set of separate targets or an aggregate performance target;

- planning to obtain mid-project information about performance and requirements, and optimal testing, by calculating the expected value of information.

Future work would move beyond the illustrative examples in this paper to more flexible formulations capable of capturing more of the considerations that arise in realistic settings. Specific developments might include:

- multilevel uncertain targets, e.g., where there is a value of 1.0 for meeting a stretch goal, 0.5 for meeting a lower threshold, and 0 for failing to meet that lower threshold;
- distributions other than Gaussian;
- other network topologies (which might be easier to model with non-Gaussian distributions such as the Gumbel distribution);
- different performance measures for the primary attributes of cost, quality and timing, or other primary attributes in addition to the above;
- aspiration equivalent-based heuristics;
- characterization of the resolution of uncertainty as a function of time (e.g., to identify the optimal time for testing);

By explicitly treating targets as uncertain variables, we can formalize many of the practical challenges of PM. In particular, this creates interesting theoretical and applied opportunities to enhance PM with powerful decision analytic methods.

References

Abbas, A., J. Matheson. (2005). Normative Target-Based Decision Making *Managerial and Decision Economics* 26:373-385.

Abbas, A., J. Matheson. (2009). Normative Decision Making with Multiattribute Performance Targets. *Journal of Multi-Criteria Decision Analysis* 16:67-78.

Abbas, A., J. Matheson, R. Bordley. (2009). Effective Utility Functions Induced by Organizational Target-based Incentives. *Managerial and Decision Economics* 30(4):235-251.

Abbasi, G. Y., A.M. Mukattash. (2001). Crashing PERT networks using mathematical programming. *International Journal of Project Management.* 19:181-188.

Ahmadi, R.,R. Wang. (1999) Managing Design Risk in Product Design Processes. *Operations Research.* 47(2):235-246.

Bhattacharya, S., V. Krisnan, V. Mahajan (1998). Managing new product definition in highly dynamic environments. *Management Science.* 44(11):850-864.

Black, T. (1990). A systems design methodology application to automotive brake design. *Masters Thesis.* MIT Dept of Mgt & Mech. Engr. Cambridge,Mass.

Bordley, R., C. Kirkwood. (2004). Multiattribute preference analysis with performance targets. *Operations Research.* 52(6):823-835.

Bordley, R., M. LiCalzi. (2000). Decision theory using targets instead of utility functions. *Decisions in Economics and Finance.* 23:53-74.

Browning, T. (2001). Applying the design structure matrix to system decomposition and integration problems: a review and new directions. *IEEE Transactions on Engr.Mgt.* 48(3):292-306.

Castagnoli, E., M. Li Calzi.(1996). Expected utility without utility. *Theory and Decision.* 41:281-301.

Clarke, C. (1961). The Greatest of a Finite Set of Random Variables. *Operations Research.* 9,2 (145-162).

Duncan, W. (1996). A Guide to The Project Management Body of Knowledge - Third Edition. Project Management Institute., Pennsylvania.

Gjerde, K,, S. Slotnick, M. Sobel. (2002). New product innovation with multiple features and technology constraints. *Management Science.* 48(10):1268-1284.

Ha, A., E. Porteus. (1995). Optimal Timing of Review with Concurrent Design for Manufacturability. *Management Science.* 41(9):1431-1447.

Hill, A.V., I. S. Khosla. (1992). Models for optimal lead time reduction. *Production and Operations Management.* 1(2):185-197.

Huchzermeier, A., C. Loch. (2001). Project Management Under Risk: Using the Real Options Approach to Evaluate Flexibility in R&D. *Management Science.* 47(1):85-101.

Kahneman, D., A. Tversky. (1979). Prospect theory: an analysis of decisions under risk. *Econometrica.* 47:263-291.

Kavadias, S., C. H. Loch. (2003). Optimal project sequencing with recourse at a scarce resource. *Production & Operations Management.* 12(4):433-444.

Keisler, J., P. Mang (2011). Applying value of information and real options to R & D and new product development. *Wiley Encyclopedia of Operations Research and Management Science* Vol. 1:138-158.

Kim, E., W. Wells, M. Duffey (2003, July). A Model for Effective Implementation of Earned Value Management Methodology. *International Journal of Project Management.* 21(5):375-382.

Ozgur, C. (1998). Capacity constrained resource scheduling: A decision utility approach. *International Journal of Operations and Quantitative Management.* 4(3):1-21.

A Guide to the Project Management Body of Knowledge (2013). Fifth Edition. Project Management Institute, Newtown Square, Pennsylvania

Schuyler, J. (2001). *Risk and Decision Analysis for Projects.* (2nd ed.) Project Project Management Institute.

Smith, R., S. Eppinger. (1997). Identifying controlling features of engineering design iteration. *Management Science.* 4(3):276-293.

Steward, D.(1981). The design structure matrix: A method for managing the design of complex systems. *IEEE Transactions on Engineering Mgt.* EM-28(3).

Thomke,S., D. Bell. (2001). Sequential testing in product development. *Management Science.* 47(2):308-323.

Tsetlin, I. & R. Winkler (2006). On Equivalent Target-Oriented Formulations for Multiattribute Utility. *Decision Analysis* 55:226-233.

Tsetlin, I. & R. Winkler (2007). Decision Making with Multiattribute Performance Targets: The Impact of Changes in Performance and Target Distributions. *Ops Research* 55:226-233.

Verzuh, E. (2011). *Fast Forward MBA in Project Management* Fourth edition, Wiley, Hoboken,

Virine, L., M. Trumper (2007) *Project Decisions: The Art and Science* Management Concepts, Vienna, VA.