Valuing and Selecting Small Portfolios of Projects

Rick Mitchell, Francis Hunt, Charles Romito and David Probert
Centre for Technology Management, Institute for Manufacturing, Department of Engineering,
University of Cambridge, Mill Lane, Cambridge CB2 1RX, Cambridge, U.K..

Abstract **Valuation methods that involve calculating a statistical expectation value are inadequate and misleading for single projects and small portfolios. For these the full range of outcomes must be considered, balancing the possibility of loss or gain. Two new tools, the Value Range Diagram and the Risk Aversion Diagram are proposed to help managers assess small portfolios.**

I. INTRODUCTION

Technology-based projects often face large uncertainties about the cost and timescale of the work and the timing and magnitude of the benefits that will flow from it. At their outset, and often for much of their lifetime, such projects face a range of possible outcomes. It is conventional to incorporate these uncertainties into the valuation process by multiplying possible outcomes by their probabilities to give a statistical expectation value. Although this is appropriate for valuing a portfolio containing very many projects it is logically incorrect and practically misleading for individual projects and for small portfolios. In these cases a single point valuation is inadequate and the range of possible outcomes must be taken into account. The value of a project will then depend on how managers balance the possibility of a relatively poor outcome against the chance of above-average performance. In this paper we discuss possible approaches to comparing projects, taking account of the spread of possibilities as well as the mean, and propose some new tools for choosing a small portfolio of projects. We argue that the majority of project portfolios are small in this sense so valid valuation is a practical problem for many, perhaps most, companies.

II. SMALL AND LARGE PORTFOLIOS

It is usual [1] to value risky projects by computing the *mean*, or *expected*, financial return. For example if a project has three possible outcomes, with values a, b and c, whose probabilities are estimated to be p_a, p_b and p_c, then the expectation value, E, for the project is:

$$E = a.p_a + b.p_b + c.p_c \qquad (1)$$

This equation, of course, does not predict the value of the project itself, which will be either a, or b, or c. The expectation value, predicts the outcome of a large number of projects, in the sense that the sum of the outcomes of n projects becomes relatively closer and closer to n.E as n increases. Hence, for sufficiently large portfolios the range of possible outcomes of each project may be ignored for the purposes of valuing the total; the individual expectations are sufficient. Small portfolios differ from large ones in that the range of outcomes for each project cannot be ignored. This is the essence of the problem of valuing small portfolios.

The value of a sufficiently large portfolio of projects closely approaches the sum of their expectations because, as is well known, when uncorrelated random variables are added the ratio of the standard deviation to the mean reduces as the square root of the number of variables. Thus the relative spread in the value of a portfolio of 10 projects will be about $1/3^{rd}$ that of a single project; that of 25 projects will be $1/5^{th}$. However, this reduction applies only if the outcomes of all the projects are uncorrelated, a rather unlikely circumstance if they are all conducted in the same organisation. Any correlation between projects makes the uncertainty in the portfolio value greater. Small portfolios such as are considered here may be the rule rather than the exception in real businesses. In any case, large portfolios are made up from small ones so that valuation methods that are valid for small portfolios can also be applied to large (though not vice versa).

III. VALUING A SINGLE PROJECT

A project will usually have a number of phases of work, each with a possible range of costs, and managers will have the option to proceed, or not, at the end of each phase (Fig 1).

There may also be the possibility of changing the direction of the project depending on the outcome of earlier work, so the linear decision tree may branch into a network. All the outcomes that are judged to be plausible are included, with probabilities assigned to them. Costs and incomes– or more generally a range of costs and incomes – can be assigned to each phase.

A decision trees or network diagram can be readily analysed to yield the expectation value (or Expected Commercial Value, [2]) of the project. More generally, Monte-Carlo techniques [3] can be used to explore the range of possible values giving a probability distribution of outcomes, ranging

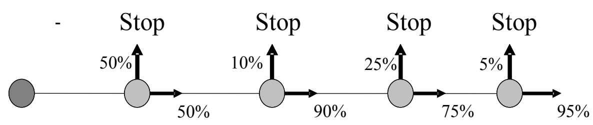

Fig1. Decision tree for a simple multi-phase project.

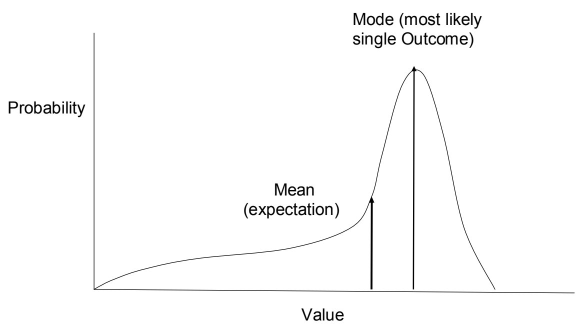

Fig 2. Schematic probability diagram for the outcome of a project

from the least to the most favourable. A schematic version of a probability distribution for a project is shown in Fig 2.

However, there are fundamental problems in assigning probabilities to events in a project plan, as is required in this process. They cannot be measured (as one would for a physical process) by taking samples from a number of identical previous events because there are none; each project is unique. The best one can do is to collect data from similar past projects and hope to deduce from them what may happen in a new project. Inevitably, this involves a great deal of interpretation and personal judgement. In practice this process is often highly flawed, as we discuss further below. Furthermore, there is no way to validate such estimates before or after the event. It cannot generally be proved even in principle that the probability of success of a project actually was (or was not) a particular figure.

Evidently, people do make general judgements about which future events are likely and which are not; life would be impossible otherwise. However, any numerical values placed on the probabilities are at best highly approximate, and possibly meaningless. Hence one should not place much reliance on the details of any probability distributions calculated from them.

IV. ISSUES IN ESTIMATING VALUES AND PROBABILITIES USED IN DECISION TREES.

Research into human estimating skills has shown that our ability to assess probabilities is not impressive. For example, Kahnemann, Slovic and Tversky [4] point out that people generally make such estimates not on the basis of correct statistical reasoning but according to certain heuristics. One such is the representativeness heuristic: we tend to judge whether something belongs to a class simply by the extent to which it resembles members of that class. This leads to a number of problems, of which the most relevant here is insensitivity to sample size [5]. People (even those who are well-versed in the theory of statistics) appear to have a persistent tendency to expect a small sample to be closely representative of the population it comes from. This causes a propensity to over-estimate the significance of a few instances - and hence of our own personal experience.

Another estimating heuristic is termed availability [6]: the tendency to assess the frequency of a class by the ease with which examples of it come to mind. This is reasonable, but may be affected by any bias in the recall process. The emotional force of instances may also make them easier to recall so we tend to over-estimate the frequency of high-impact triumphs and disasters.

A third heuristic is anchoring: the tendency to be unduly influenced by the most recently acquired information.

The biases that may come from these heuristics are apparently innate and not easily corrected. They certainly cast doubt on the accuracy that can be expected of predictions about events in innovation projects. “For anyone who would wish to view man as a reasonable intuitive statistician, such results are discouraging" [7]. In addition predictions may also be adversely affected by social factors such as the influence of powerful or charismatic individuals, or by groupthink [8]: the propensity of tight-knit groups to over-value consensus and so mistake agreement for truth. These, however, may be mitigated by careful design of the estimating process [9].

V. A PORTFOLIO OF 2 PROJECTS

The problem of valuing and comparing a small number of projects is shown most clearly with two projects, such as those whose distributions are shown schematically in Fig 3.

Intuitively, it seems that the different ranges of outcome of these two projects should lead to different valuations, even though their probability distributions have the same mean and the same mode. The issue is to find a compact and logically consistent way to express this. We call this the “Range Valuation” problem.

In the world of finance, risky investments (those whose price is expected to fluctuate considerably) attract higher returns than those that are more predictable [10]. Equities on average typically attract a higher yield than government bonds, which are assumed to have effectively zero risk. The “risk premium” for the equity market reflects the risk aversion of the average

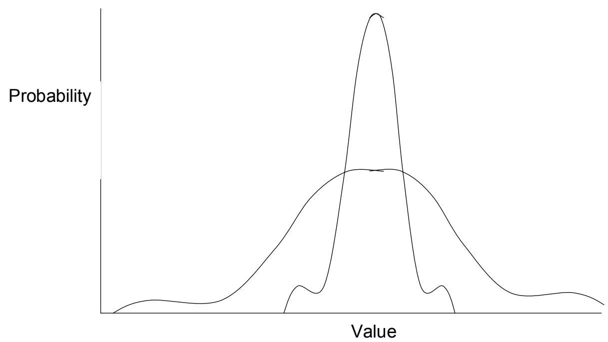

Fig 3. Two projects with the same mean outcome, but different distributions

investor, which is basically a matter of sentiment rather than calculation.

Dembo and Freeman [11] have suggested that competing projects should be valued as the difference between the possible upsides and downsides in relation to a benchmark (such as the bank rate, or some standard for investments set in the company). The value, V, of the project would then be:

$$V = \text{Upside} - \text{Downside} \qquad (2)$$

Here the Upside is the integral of the distribution above the benchmark and the Downside is the integral of the distribution below. Equation (2) would give a different valuation for the two value distributions illustrated in Fig 3, unless the benchmark happened to coincide with the peak. The narrower distribution would be the more valuable if the benchmark is below the peak of the distributions while the broader would be the more valuable if the benchmark is above the peak.

A refinement is to apply different weightings to the Downside and the Upside. Dembo and Freeman suggest that the Upside may be thought of as the chance to participate in a gamble to make more money than the benchmark. A 20% chance of making £1000 is worth something, but for most people the value, or Utility, to them is not as much as the expectation value of £200. Similarly, the utility value to be placed on the downside could be the amount one would be willing to pay to insure against such a loss. Again, this would generally be less than the expectation. The utility to be assigned to the upsides and downsides seems to depend both on the sums involved and the attitude of the investor to the risk. Few people would consider insuring against a potential loss of £1, but few would ignore one of £1M. It seems that very large losses or gains in relation to the benchmark should be weighted differently. Thus the value, V, is:

$$V = x.\text{Upside} - y.\text{Downside} \qquad (3)$$

The utility weightings, x and y, depend on the investor's tolerance of risk, and depend on the upside and downside values themselves.

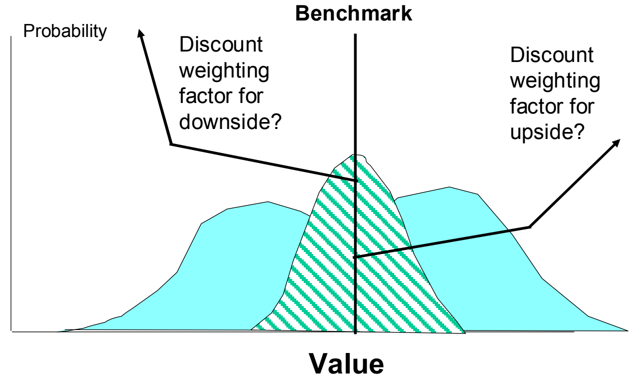

Fig 4. Possible forms for the upside and downside weighting factors

One would expect the factors to increase with distance from the benchmark as shown in Fig 4 because a very large loss may be catastrophic, while a very large gain might have a transforming effect on the organization.
Ignoring their value-dependence, the discounts may be combined into a single figure representing the risk tolerance of the organization.

$$V = \text{Upside} - \lambda.\text{Downside} \qquad (4)$$

Here λ is greater than one if the organization is risk-averse and less than one if not.

This approach offers a promising way of including the spread of possible outcomes into project evaluation. However, it requires a rational way of determining what the discount curves should be linking the for particular projects in particular organizations. More research is needed here. If these weightings could be determined it would be possible, using equation 3, to calculate project values taking proper account of the range of possible outcomes. This would be a solution to the Range Valuation problem mentioned above.

VI. INCREMENTAL VALUATION

The early stages of a technology project are usually devoted to improving knowledge and understanding. This serves to clarify the possibilities and options available and so to narrow the range of likely outcomes of the project. A solution to the Range Valuation problem would make it possible to place a value on this reduction in the range of likely outcomes and so to place a value on each phase of preliminary work. This would be a new and valuable insight. When choices have to be made between projects it may be better to choose not the project that has the higher overall expectation value, but the one whose next phase of work adds the most value.

VII. RISK-REWARD DIAGRAMS FOR SMALL PORTFOLIOS

A. *Previous Methods*

In the absence of a solution to the Range Valuation problem we must examine how familiar tools of portfolio management should be adapted for use with small portfolios. The most widely-used of these is the risk-reward diagram [12]. In this, projects are plotted on a two-dimensional grid where one dimension represents the value of the project and the other its perceived risk. The idea is that managers can readily review the balance of risk and reward in the portfolio and ensure it conforms to their judgement of what is required.

In the probability-distribution approach to project valuation just outlined, risk is not a separate parameter but merely an expression of the range of plausible outcomes. A single project does not have a single outcome and a greater or smaller level of risk; it just has a greater or lesser range of possible outcomes. This perception accords well with our experience in collecting the information for drawing risk-reward diagrams. Participants often find the concept of risk too imprecise. A project may have a low risk of outright technical failure but a high risk of falling somewhat short of expectations. Is this high or low risk? This has led to the proposal to use ellipses instead of circles on a risk-reward diagram [13], where the width of the ellipse represents the range of possible rewards, with the risk dimension representing only the probability of technical failure. This is a step forward, but it assumes that the technical phase will end either in complete success or in complete failure, which is not necessarily the case. Furthermore it does not distinguish projects where the risk of failure can be tested early and at low cost from those that retain a significant possibility of late, and expensive, cancellation.

B. *The Value Range Diagram*

We propose that for small portfolios the risk-reward diagram should be replaced by a Value Range Diagram, in which the two scales are the highest and lowest values expected for the project. These can be obtained by constructing a decision tree for the project including all phases, costs and decision points, and drawing a probability diagram for the outcomes (for example using Monte-Carlo techniques). The Highest Likely Value (HLV) and the Lowest Likely Value (LLV) are then selected and plotted on a Value Range Diagram such as Fig 5. In practice managers will not necessarily choose the absolute extremes of the probability diagram for the HLV and LLV but will use some judgement, bearing in mind that the uncertainty of the data and the actions that they may themselves take to limit the extreme outcomes as the project develops.

The two axes of a Value Range Diagram are the highest likely value, HLV, and the lowest likely value, LLV. The scales may absolute values such as NPV, or performance ratios such as return on investment. The horizontal and vertical scales will generally be different, because HLVs will usually be significantly larger than LLVs. A performance benchmark may be included, as in Fig 5. The shaded area on the upper left is the illogical domain where worst returns exceed best.

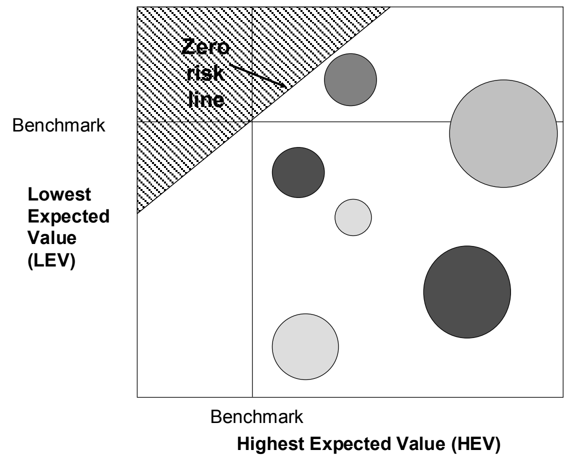

Fig. 5. The Value Range diagram

The bold diagonal line represents the case of zero uncertainty, where HLV=LLV.

The Value Range diagram is used in a similar way to a Risk-Reward diagram. It displays the current state of the portfolio in a way that helps managers to appreciate their exposure to gains and losses, leaving to their judgement whether the balance is appropriate. As they progress, projects will be expected to migrate upwards towards the area where both HLV and LLV are above the benchmark. As with the risk-return diagram, the size of the circles can represent the investment required to complete the project, and colour or shading can indicate the type of project or its maturity.

C. Risk Tolerance Diagram

An alternative way to display the probability diagram information is to calculate the risk tolerance parameter, λ, in (4), that would bring the value of the project to the benchmark. We call this the Risk Tolerance Value for the project. We would emphasise that the Risk Tolerance Value is not the risk of the project in the conventional sense of probability of failure. It is the risk tolerance the organisation would have to have for it to regard the spread of possible returns above and below the benchmark as just acceptable. The concept is analogous to the Internal Rate of Return used in financial analysis, which is the discount rate that would reduce the calculated NPV to zero. Projects may be displayed for decision making on a Risk Tolerance diagram, with the Risk Tolerance value as one dimension and the mean return, or profitability, on the other.

VIII. CONCLUSIONS

We believe that the use of a single mean value for valuing projects is logically flawed and seriously misleading unless the portfolio is very large, probably larger than is the case for most businesses. For small portfolios the range of possible outcomes cannot be neglected. The value of a project will then depend on how managers balance the possibility of a relatively poor outcome against the chance of above-average performance. No coherent way appears to be available to say what risk tolerance an organisation ought to have for its projects and so it is not (yet) possible to propose a single measure of value that is valid for single projects and small portfolios. Further investigation is required here.

We present two new tools, the Value Range Diagram and the Risk Tolerance Diagram which provide a straightforward and logically coherent approach to assessing the balance of projects in a small portfolio, though still leaving the final choice to management judgement.

REFERENCES

[1] F. P. Boer, "Risk-adjusted valuation of R and D Projects" *Research - Technology Management* vol.46 Sept-Oct 203 pp50-8

[2] R. G. Cooper, S. J. Edgett and E. J. Kleinschmidt, *Portfolio Management for new products*. Cambridge Mass.: Perseus Books, 2nd ed., 2001, p34

[3] R. Razgaitis, *Dealmaking Using Real Options and Monte-Carlo analysis*. Hoboken New Jersey: John Wiley, 2003.

[4] D. Kahneman, P. Slovic, and A. Tversky, *Judgement under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press, 1982 Ch. 1.

[5] A. Tversky and D. Kahneman, "Belief in "The law of Small Numbers"". *Psychological Bulletin* 1971;2:105-110.

[6] Kahneman, Slovic and Tversky, op. cit. ch 1

[7] ibid.. ch 10

[8] I. L. Janis, *Groupthink.* Boston: Houghton-Mifflin, 1972

[9] J. A. Keizer, J., J. I. M. Halman and M. Song, "From experience: applying the risk diagnosing methodology" *Journal of Product Innovation Management (2002)* Vol 19 pp 213-232

[10] J. F. Weston, S. Besley and E. F. Brigham, *Essentials of Managerial Finance*. 11th ed. Fort Worth: The Dryden Press, 1996.

[11] R. Dembo and A. Freeman, *Seeing Tomorrow: Rewriting the Rules of Risk*. New York: John Wiley, 1998.

[12] R.G. Cooper, S. J. Edgett and E. J. Kleinschmidt *Portfolio Management for new products*. Cambridge Mass.: Perseus Books, 2nd ed., 2001.

[13] G. L. Tritle, F. V. Scriven, and A. R. Fusfeld, "Resolving Uncertainty in R&D Portfolios", *Research-Technology Management*, Vol. 43, no. 6, (2000), pp. 47-55.