Pergamon

PII: S0301-4207(98)00008-7

Resources Policy, Vol. 24, No. 1, pp. 49–57, 1998

Printed in Great Britain
0301-4207/98 $19.00 + 0.00

Risky choice, risk sharing and decision analysis: Implications for managers in the resource sector*

Michael R Walls

Division of Economics and Business, Colorado School of Mines, Golden, CO 80401, USA

Dana R Clyman

Darden Graduate School of Business, University of Virginia, Charlottesville, VA 22906, USA

This paper provides a brief overview of decision and preference analysis concepts and demonstrates an application of these techniques to the project-valuation problem faced by resource managers. Our major focus is on the use of the exponential utility funetion, the utility function most frequently used by resource companies. We discuss the important and practical risk-sharing problem faced by managers in the resource sector, that is, how to choose the optimal share of a risky project. We demonstrate that with decision and preference analysis tools it can be quite straightforward for managers to identify their optimal share in risky projects. We then explore these techniques further and demonstrate that they can lead to some seemingly counter-intuitive results. In particular, we explore how the firm's optimal share changes with exogeneous changes in project parameters. What we find is that while many of the changes in share are intuitive, some are not. In fact, when the firm's estimate of the potential upside payoff upon finding reserves increases, it is sometimes better to decrease the firm's share than it is to increase it. This is important, because by recognizing this counter-intuitive result, we can work to improve our intuition by understanding it. We summarize our findings and offers some guidelines resource managers should consider when considering a choice of utility function.

Introduction

Managers in the resource sector often encounter important choice problems that involve large capital investment commitments and that are characterized by a high degree of uncertainty. Those dimensions of uncertainty may include issues such as reserve recoveries, operating costs, product price and others. Making choices among alternative projects is difficult because the risk characteristics of these projects are often so different. In addition, because resource managers are often constrained in terms of available capital, they are confronted with critical decisions concerning what share of a particular project, if any, a firm should purchase. All of these choices are important because the resulting outcomes can have a significant impact on a firm's performance. In this paper, we demonstrate how decision analysis techniques can go a long way towards addressing these decision problems and show that the implementation of these techniques can have important implications to managers who face choices about risky projects.

What's more, we show that decision-analysis techniques can also help us to make more informed decisions by providing insights that help us to improve our intuition. In particular, this paper will show that our immediate knee-jerk reactions to improvements in certain project parameters—for example, the amount of oil one expects to find should a well be successful—may be wrong, and why. Thus, by using such techniques, not only can we make better

*Due to circumstances beyond the Publisher's control, this article appears in print without author corrections.

decisions, but we can improve our intuition about the decisions that we make.

Increasingly, decision-analysis techniques have been used to aid managerial decision-making in situations like these of substantial risk and uncertainty. Indeed, there have been a wide and diverse range of successful decisions-analysis applications including budget analysis, product planning, corporate strategy, medical diagnosis and treatment, site planning, as well as numerous other private and public sector decision contexts[1].

Both the petroleum and mineral industries represent classic settings for applications of decision analysis (DA) and preference theory. Early work in the petroleum industry includes the paper by Grayson (1960) on decisions about uncertain drilling prospects, the discussion by Raiffa (1968) of the use of the exponential utility function and its application to the risk-sharing problem, the exposition by Spetzler (1968) about measuring utility functions for a large oil company, and the text by Newendorp (1975) on using decision-analysis techniques in petroleum exploration. More recent work includes extensions by Cozzolino (1977) on the use of the exponential utility function in risk-sharing decisions, the discussion by Howard (1988) about measuring risk tolerance for a large oil company, the development and use of a modeling system by Keefer *et al.* (1989) to aid a major oil company in allocating bidding capital, the development by Walls *et al.* (1995) of a computer-based DA model which enabled Phillips Petroleum Company to incorporate its risk tolerance in exploration decisions, the discussion by Walls (1995) of methods for measuring and integrating corporate risk tolerance; and the development by MacKay *et al.* (1996) of a spreadsheet model to assist in risk-sharing decisions at Texaco.

In the mineral industry, Hax and Willig (1977) applied decision analysis techniques to mining-investment decisions and Walls and Eggert (1996) developed a methodology for measuring a mining-company CEO's financial-risk tolerance, surveyed a sample of mining CEOs and provided a comparison of risk tolerance among mining firms.

In this paper we first provide a brief overview of decision analysis and preference analysis concepts and demonstrate an application of these techniques to the project-valuation problem. Our major focus is on the use of the exponential utility function, the utility function most frequently used by resource companies. The third section discusses the important and practical risk-sharing problem faced by managers in the resource sector—how to choose the optimal share of a risky project. Here, we demonstrate that, with decision-analysis tools, it is quite straightforward for exponential utility maximizers to identify their optimal share in risky projects. The fourth section pushes the analysis a bit further and demonstrates that it can lead to some rather counter-intuitive results, even when using an exponential utility function. In particular, we explore how one's optimal share changes with exogeneous changes in project parameters. What we find is that while many of the changes in share are intuitive, some are not. In fact, when one's estimate of the potential upside payoff upon finding oil increases, it is sometimes better to decrease one's optimal share than it is to increase it. Furthermore, we provide examples from the real prospect inventory of one major oil company to demonstrate that it is not unusual for oil companies to be considering projects for which they should decrease their share when the upside potential increases. This is important, because by recognizing this counter-intuitive result, we can work to improve our intuition by understanding it; and, by seeking explanations that help us to improve our understanding, we pre-empt incorrect decisions based on faulty intuition. The final section summarizes our findings and offers some guidelines managers should consider when considering a choice of utility function.

Overview of decision analysis and preference theory

The foundations of decision analysis are provided by a set of axioms stated alternatively in von Neumann and Morgenstern (1953); Savage (1954); Pratt *et al.* (1964). These axioms, which provide principles for analysing decision problems, imply that the attractiveness of alternatives should depend on: (1) the likelihood of the possible outcomes of each alternative; and (2) the preferences of the decision makers for those outcomes. The technical implications of the axioms are that probabilities and utilities (measures of preference) can be used to calculate the expected utility of each alternative and that alternatives with higher expected utilities should be preferred. The practical implication of the axioms is the provision of a sound basis and general approach for including judgments and values in an analysis of decision alternatives.

The preference-analysis (or expected utility) approach is an extension of the expected-value concept. The result is a more realistic measure of value among competing projects characterized by risk and uncertainty. The expected-utility approach is appealing in that it enables the decision maker to utilize a relatively consistent measure of valuation across a broad range of risky investments. The theory provides a practical basis for the firm to formulate and implement a consistent risk policy that incorporates the firm's attitude about participating in financially risky projects.

The way that preference theory accomplishes this is by encoding the firm's attitude about taking on risks in a utility function. The ability to measure corporate

[1]See Corner and Kirkwood (1991) for a more complete survey of decision-analysis applications.

risk preferences is an important part of both the conceptual and practical views of decision making under risk and uncertainty. To do this, the exponential utility function is often used to measure risk preferences of an individual or organization. One of the reasons that the exponential utility function is so frequently used is that a firm's risk tolerance may be encoded in a single number, which can be estimated (Spetzler, 1968; Walls, 1995; Walls and Eggert, 1996). The exponential utility function has the following mathematical form: $u(x) = a - be^{-x/R}$, where u is utility as a function of the variable of interest, x (e.g. outcomes valued in dollars), e is the exponential constant, a and b are scaling constants, and R is the risk tolerance. The risk tolerance, R, represents the sum of money about which decision makers are indifferent between a 50–50 chance of gaining the whole sum and losing half the sum. Thus, use of the exponential form leads to the characterization of risk tolerance by the single number, R, which, loosely speaking, measures the curvature of the utility function. This functional form is also dominant in both theoretical and applied work in the areas of decision theory and finance for a number of other reasons. It is general enough to treat satisfactorily a wide range of individual and corporate risk preferences (Keeney and Raiffa, 1976). Because it is based on the exponential function, it is highly tractable mathematically, and because it is based on the exponential function it allows us to consider independent projects individually. It has a significant practical advantage over other utility functions in that it does not require that the evaluation consider the firm's entire portfolio of projects.

Whenever we know a firm's utility function, exponential or otherwise, we can compute a risk-adjusted valuation measure for any risky or uncertain investment. This valuation measure is known as the *certainty equivalent* and is defined as that certain value that a decision maker is just willing to accept in lieu of the gamble represented by the uncertain project. It is, in essence, the "cash value" attributed to a project that involves uncertain outcomes. With certainty equivalents, comparisons between projects are easy; higher certainty equivalents are preferred to lower ones. With the exponential utility function, the certainty equivalent is a function of the firm's risk tolerance level, R, and the risk characteristics (probability distribution over the outcomes) of the investment opportunity. For projects modeled as discrete probability distributions, the expression for the certainty equivalent, C_x, is given by the following formula derived from the exponential utility function, which, while previously known, was popularized in the petroleum industry by Cozzolino (1977),

$$C_x = -R\ln\left(\sum_{i=1}^{n} p_i e^{-x_i/R}\right) \qquad (1)$$

where p_i is the probability of outcome i, x_i is the value of outcome i, and ln is the natural-log function.

Consider, for example, the investment projects shown in Figure 1. The payoff of the project upon success, which is denoted in Figure 1 by $x_1 > 0$, is equal to the net present value of all cash flows from the initial investment through the life of the project, given that the project is successful. The total payoff upon failure, which we denote by $x_2 < 0$, is the net present value of all of the cash flows assuming the project fails. The probabilities of success and failure are p_1 and p_2, respectively, with $p_1 + p_2$ equal to one. Once values for the two possible outcomes are estimated, a probability of success is assessed, and a risk tolerance, R, specified, it is a simple and straightforward matter to calculate both the expected value and the certainty equivalent for these projects. In this example we demonstrate that on an expected value basis the firm should prefer project B (EV = US$6.7 MM) over project A (EV = US$4.0 MM). Now let's assume that, in fact, the firm is risk averse and has a financial risk tolerance, R, equal to US$60 MM. By applying equation (1), we can conduct a certainty-equivalent analysis, which considers the firm's attitude about financial risk. When we do, we find the opposite result is true; that is, the less-risky project A (with a C_x of US$1.2 MM) is preferred to the more-risky project B (which has a C_x of − US$2.7 MM). Unlike the expected-value analysis, the certainty-equivalent incorporates the firm's concerns for the "risks" associated with each of these projects. The C_x valuation measures the tradeoffs between potential and uncertain upside gains vs. potential and uncertain downside losses, and it does so with respect to the firm's attitude about taking on risk. Indeed, the C_x analysis explicitly considers the relative magnitudes of capital being exposed to possible loss, the chances of those losses, and the firm's relative risk-attitude about the uncertain financial consequences.

Optimal share analysis

While the previous discussion demonstrated how decision-analysis tools can help a firm to decide whether to accept a project in its entirety, or which of several projects it prefers, one of the more fundamental decisions that managers confront is: what share of the project, if any, should the firm purchase. The certainty-equivalent valuation method also provides guidance to the firm about this risk-sharing decision as well. Unlike expected value analysis, which is an "all or nothing" decision rule, the C_x valuation aids the firm in selecting the appropriate level of participation consistent with the firm's risk propensity. The C_x measure provides a formal means to quantify the advantages of selling down or "spreading the risk".

Consider again our example in Figure 1 where now the payoff upon success (failure) is the ownership share (which we denote by z) times the payoff of the total project upon success, x_1 (or upon failure, x_2). We can now calculate, either analytically or numerically,

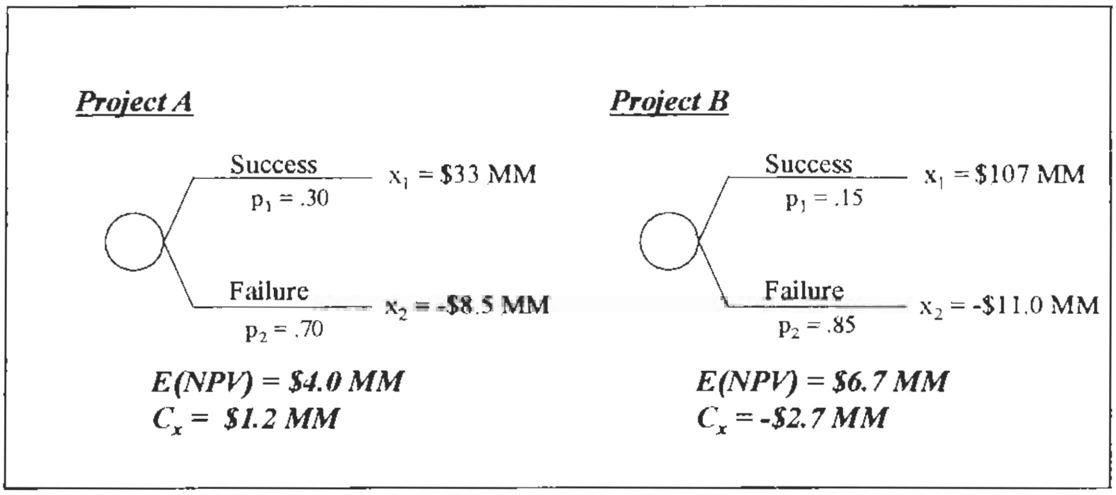

Figure 1 Risky projects A and B are compared on a traditional expected value basis, E(NPV), and a certainty equivalent basis, C_x. The C_x analysis assumes the firm has a financial risk tolerance value, R, of US\$60 million

the firm's optimal share of the project, z^*, where $0 \leq z^* \leq 1$.

Figure 2 displays graphically the results of the numerical calculation of the certainty equivalent for projects A and B at different participation levels (shares, z) for a firm with a risk tolerance of US\$60 million. The optimal share, z^*, of the project occurs where the certainty-equivalent function is maximized, and may be picked off the graph. Alternatively, the optimal share for projects like these may be calculated analytically, and is given by the formula,

$$z^* = \frac{R \, ln\left(-\frac{p_1 x_1}{p_2 x_2}\right)}{x_1 - x_2}. \tag{2}$$

(For a more detailed discussion of this and the other calculations reported in this and the next section, see Clyman *et al.* (1997).)

Once one restricts one's purchase to the optimal amount, z^*, it is a simple matter to calculate the cer-

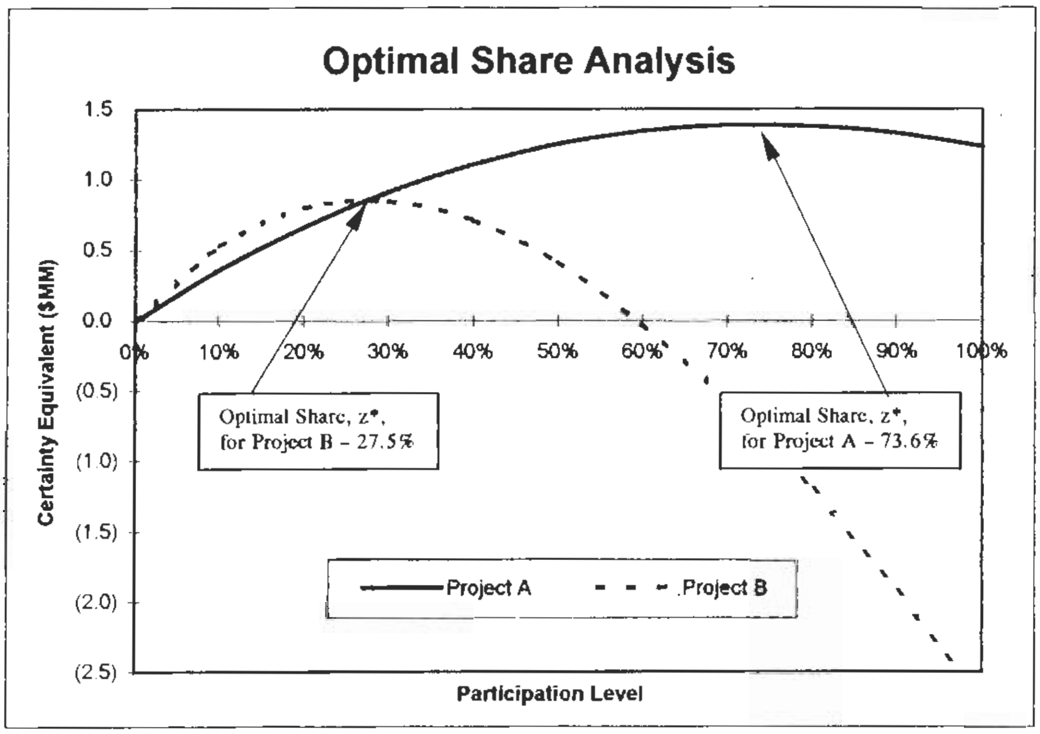

Figure 2 This optimal share analysis graph identifies the optimal share for projects A (73.6%) and b (27.5%) for a firm with a US\$60 million risk tolerance, R. The optimal share is that share with the highest certainty equivalent value, C_x

tainty equivalent of the optimal holding by replacing the payoffs (x_i in equation (1)) from owning the entire project with the payoffs z^*x_1 and z^*x_2 from owning the optimal share. Utilizing equation (2), we compute the optimal share, z^*, given the firm's risk tolerance, R, of US$60 million. The optimal shares are 73.6% for project A and 27.5% for project B. Note also in Figure 2 that at levels of participation up to just a little more than 27.5% in either project, the firm should always prefer project B over project A since it has a greater certainty equivalent. However, for participation levels greater than that, Project A is the preferred alternative as it possesses a higher certainty equivalent. Most important, each project should be valued at its own optimal participation level, and when one does so, the certainty-equivalent of project A at its best participation level is higher than the certainty equivalent of project B valued at its best participation level.

The important implication in this analysis is that the firm has a formal means of measuring the value of diversification. Note, for example, that the certainty equivalent for either project at its optimum participation level (US$1.4 million for project A at 73.6% and US$0.9 million for project B at 27.5%) is greater than the sum of the certainty equivalents for both projects at 100% participation (– US$1.5 million). Also note that participation greater than 60% in project B has a negative certainty equivalent, which implies that this project is too risky for the firm at these participation levels.

For managers of petroleum and mineral firms who are concerned with integrating a comprehensive discounted-cash-flow analysis with a corporate-risk-management decision model, the preference theory approach has important implications. The certainty equivalent valuation ensures a consistent risk attitude in project evaluation. In addition, it provides a reasonable and rational analysis of a wide range of risky investment alternatives and the capability to provide management a value maximizing set of projects consistent with the firm's willingness to take on financial risk. This approach can increase management's awareness of risk and risk tolerance, provide insight into the relative financial risks associated with its set of investment opportunities, and provide the company a formal decision model for allocating scarce capital.

A counter-intuitive result and more informed intuition

The authors of this paper are strong proponents of decision analysis and its application. Besides helping managers solve the more routine questions of whether to invest and how much, decision analysis can be used to improve our intuition about the decisions we make, and it can be used to prevent us from making misinformed decisions. Nonetheless, the users must exercise care because embedded in the choice of utility function is a hard-to-decipher statement about the user's risk preferences. And while that may sound innocuous, those embedded risk preferences can lead to some highly surprising, and often counter-intuitive, choices.

Clyman (1995), for example, examined the portfolio managers' rebalancing decision in a particularly simple structure where he could ask the question: do you want to hold more, less, or the same amount of a security whose prospects have improved (but whose price has not yet changed)? Surprisingly, he demonstrated that all three answers are consistent with expected-utility theory. In particular, he demonstrated that if the portfolio-manager's utility function is sufficiently risk–averse then a decision-analytic model could say sell now, before the market price increases. Thus, decision analysis does not necessarily require that portfolio managers prefer to increase their portfolio's holdings of a security whose prospects have improved.

This section of the paper examines a similar problem of particular interest to the resource industry where project risk-sharing and joint ventures are common. How does the mining or petroleum firm's optimal share of an investment project change as the firm's evaluation of the project's parameters change through either better assessments of the project's chances of success or better estimates of the project's potential payoffs? In answering this question, we find that our intuition is often correct—we always want more of a project when the chance of success increases, and we always want less when the costs of failure increase—but we also find that it is perfectly consistent with expected-utility theory for expected-utility maximizers sometimes to want a smaller share of the project when the potential upside payoff increases. Furthermore, it is easy to demonstrate that this seemingly counter-intuitive desire for less of an improved project is possible when using an exponential utility function—the utility function most frequently recommended and used by analysts in the resource sector.

For instance, assume that new information indicates that, if you discover petroleum or mineral reserves, the amount you will find will be far larger than previously anticipated, and hence, the potential return on your investment far greater. The knee-jerk reaction is to buy a bigger share of this obviously better prospect. However, certainty-equivalent analysis may suggest otherwise—that the certainty equivalent increases and becomes best, not when one buys more of the prospect, but when one buys less. In other words, *ceteris paribus, decision analysis can lead you to choose a smaller optimal share upon an increase in the success-state payoff.* What's more, wanting less of an improved project should not be an uncommon event. By examining the prospect portfolio of one major US oil and gas company for the year 1996, we find that a surprisingly high percentage of its prospects are ones for which the firm should choose a smaller share should its estimate of the project's upside potential increase.

Let us consider these issues in the context of the two investment projects that we examined earlier. The data for projects A and B are summarized in Table 1. Consider project A. Using equation (2), the firm's optimal share is 73.6%. That is, this is the share of project A that results in the highest certainty equivalent (or cash value) for the firm. Adjusting the payoff structure for this project using z^*, we solve, using equation (1), for the certainty equivalent of US$1.4 million at the optimal share. Now note that if we increase the upside payoff (by 20% to US$39.6 million), we discover the rather intuitive result that the optimal share increases to 86.3%, and the certainty equivalent also increases to US$2.4 million. And again when we increase the upside payoff by 100% (from US$33 million to US$66 million), the optimal share again increases to 96.8%. However, if we increase the original upside payoff by 400% from US$33 million to US$132 million, the optimal share drops to 73.3%. Thus, for project A, as the upside potential gets better, the optimal share increases for a while and then begins to decline. It is important to note, however, that even when the optimal share is declining, the certainty equivalent (the cash value of the drilling project to the firm) continues to increase. All of these numbers are presented in Table 1, along with similar numbers for project B.

We note from examining Table 1 that for both projects the optimal share increases for a while before beginning to decline. It is possible, if given a different set of project characteristics that the optimal share would begin to decline immediately upon an increase in x_1. For example, consider a project C where p_1 is 20%, p_2 is 80%, x_1 is US$133 MM and x_2 is equal to − US$6 MM. For this project, a 1% increase in the success payoff, x_1, leads to a decrease in the optimal share, z^*. Notwithstanding this case, it is important to note that in all cases it is possible to find an upside-payoff level upon which further increases cause the then existing optimal share to begin to decline.

Table 1 **The optimal share for projects A and B first increases and then decreases even as the potential upside payoffs, x_1, for these projects continue to improve**

	Project A	Project B
p_1—Success	30%	15%
p_2—Failure	70%	85%
x_1—Success NPV (US$ MM)	33.0	107.0
x_2—Cost NPV (US$ MM)	−8.5	−11.0
Optimal share, z^*	73.6%	27.5%
C_e at optimal share	$1.4 MM	$0.9 MM
z^* w/20% increase in x_1	86.3%	31.1%
C_e at z^* w/20% increase in x_1	$2.4 MM	$1.4 MM
z^* w/100% increase in x_1	96.8%	32.9%
C_e at z^* w/100% increase in x_1	$5.8 MM	$3.1 MM
z^* w/400% increase in x_1	73.3%	23.6%
C_e at z^* w/400% increase in x_1	$12.2 MM	$5.9 MM

In Figure 3, we present this effect graphically and in greater detail for project A, by showing the effect of increases in the upside payoff (x_1) on both the optimal share and resulting certainty equivalent. In Figure 3 the solid line shows how the optimal share (z^*) changes as the success payoff (x_1) changes; the bold dotted line shows the effect of success-payoff changes on the certainty equivalent at the optimal share, z^*. The fine dotted line indicates the certainty equivalent of the 100% share as the success payoff changes. What is striking about this example is that the optimal share increases for a while and then begins to decrease. Once it starts decreasing, it continues to decrease, and continues to decrease upon ever greater improvements in the upside payoff, x_1. The certainty equivalents, on the other hand, continue to grow, but at ever decreasing rates.

Indeed, it is possible to demonstrate mathematically that once exponential-utility-function maximizers begin to want less of a project as the upside payoff improves, they will always continue to want less as the upside continues to increase. (Again, see Clyman *et al.* (1997) for a mathematical treatment of these ideas.) Indeed, because the optimal share, upon increases in x_1, may increase for a while, eventually begins to decrease, and never again begins increasing, it is possible to derive a breakeven equation. That is, as long as x_1 is less than some amount, the optimal share will increase when x_1 increases, but once x_1 exceeds that amount, the optimal share will always decrease. Exponential-utility-function maximizers will want more, the same amount, or less of a project as the upside payoff increases depending on whether

$$1 - \frac{x_2}{x_1} \gtreqless ln\left(-\frac{p_1 x_1}{p_2 x_2}\right). \tag{3}$$

Using equation (3), it is possible to calculate (numerically) the breakeven value of x_1 as a function of the other parameters. Table 2 shows the breakeven success payoffs, x_1, for each of the three projects that we have discussed. For example, we find for project A that exponential-utility-function maximizers will continue to want more of project A until x_1 equals US$63 million, and will then want less as x_1 increases further. In the case of project B, the breakeven value is US$181 million, and in the case of project C the breakeven value is US$71 million. The breakeven point is above the current value of x_1 in two of the three projects (A and B), and so increases in the upside cause the firm to want more of these two projects for a while. On the other hand, the upside potential already exceeds the breakeven point in project C, and so any additional increase in its value causes the firm to want less.

equation (3) can also be used to calculate the breakeven x_1 as a function of the cost of failure (holding the probability of success constant) or as a function of

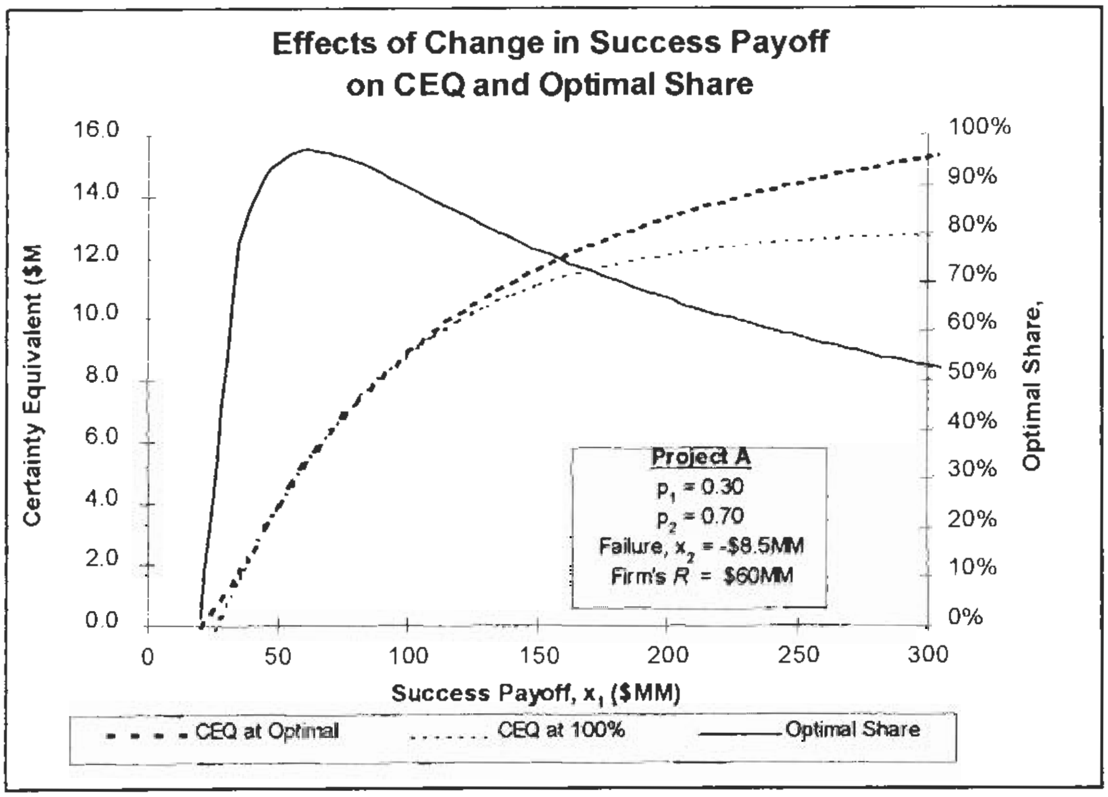

Figure 3 Certainty equivalent (CEQ) comparisons are shown between the optimal share and a 100% participation basis for project A. Note the turnover in the optimal share line as the success payoff, x_1, continues to increase. This suggests the firm should reduce its share at some point, even as the prospect's upside potential continues to improve

Table 2 Breakeven values for each project indicate the point at which any additional increase in the value of the upside potential, x_1, will cause the firm to want less of the project

	Project A	Project B	Project C
Current upside (x_1)	33	107	133
Breakeven x_1	63	181	71

the probability of success (holding the cost of failure constant). For example, Figure 4 presents for project A the breakeven x_1 as a function of the cost of failure, holding the probability of success constant. That is, it defines for this particular project two regions: region 1, where an increase (decrease) in the success payoff leads to a decrease (increase) in the optimal share, z^*; and region 2, where an increase (decreases) in the success payoff leads to an increase (decrease) in the optimal share.

Examination of the 121 drilling projects that were under serious consideration from the 1996 prospect inventory of one major US-based oil company[2] showed that a surprisingly high number of them were "good enough" to be in region 1; that is, further increases in the upside payoff would lead to decreases in the optimal share preferred.

One easy way to determine whether further increases in upside payoff would lead to increases or decreases in optimal share, other than solving numerically for the breakeven point, is to increase the upside payoff by some percentage and recalculate the optimal share using equation (1). When we did this for the 121 exploration prospects in this firm's prospect inventory, we found that a 20% increase in upside payoff led to decreases in optimal share in 83 (about 70%) of the prospects. In other words, only 30% of the drilling prospects under consideration were located sufficiently far into region 2 of Figure 4 to support a continued increase in investment upon a 20% increase in the upside payoff. When we doubled the success payoff (increased x_1 by 100%), we discovered that the firm would want a decreased share from current levels in 104 (or 86%) of the original 121 prospects. The immediate reaction that we should want more of a project whose payoff potential has increased is not correct, assuming that the exponential utility function accurately represents our risk preferences.

For us to gain the most from insights like these, we must now see if we can also improve our intuition so that results like these begin to make sense. The

[2]Because of the confidential nature of these drilling prospects and the disclosure of the company's financial risk tolerance (R), we are unable to release the company's identity.

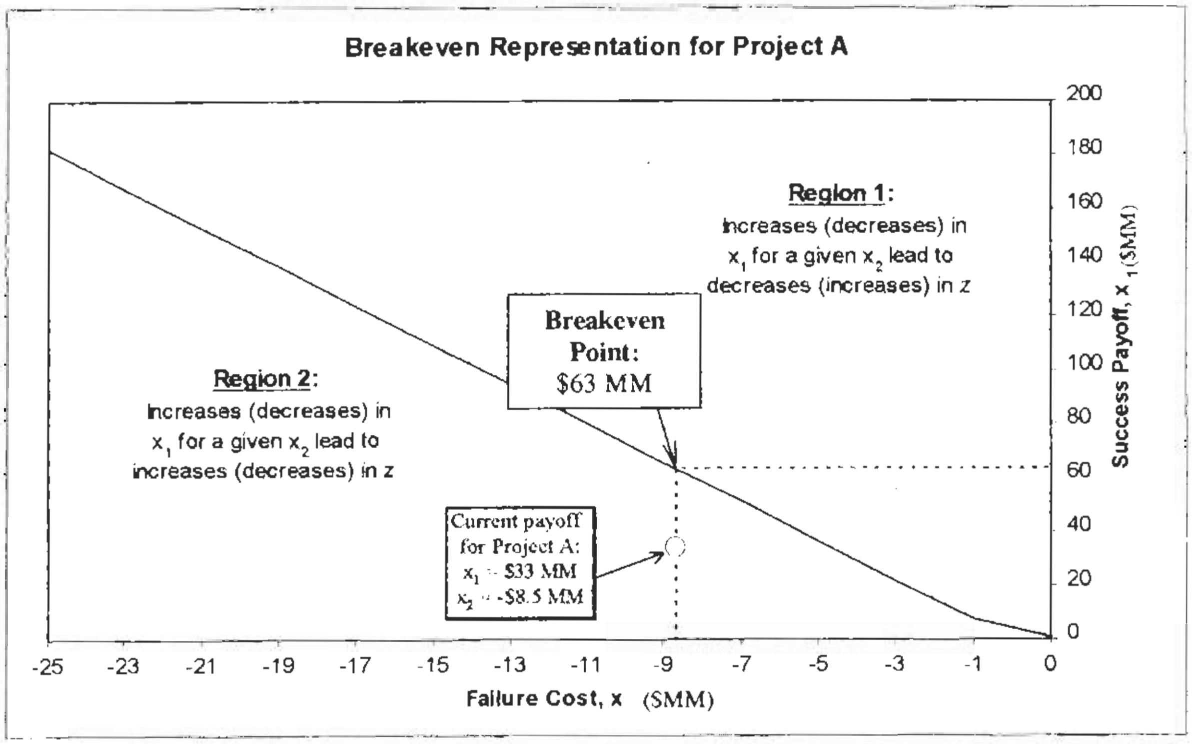

Figure 4 Breakeven representation shows location of investment projects with respect to region 1, where increases in the upside payoff lead to decreases in the optimal share, and region 2, where increases in the upside payoff lead to increases in the optimal share. Note that project A lies in region 2 such that further increases in the upside payoff, up to the breakeven value of US$63 million, will lead to increases in the optimal share

"knee-jerk" intuition that we should always want more of a project whose upside has improved derives, we believe, from the fact that we assume that the decision maker is already prepared to expose the firm to the downside payoff (− US$8.5 million in project A). Therefore, the manager should be willing to withstand at least that high a downside, or possibly even more, if the upside increases. However, it may be just as accurate to assume that the manager should focus on the upside (US$33 million in project A). When the upside potential increases, the firm can do just as well as before with a smaller share, and in the process decrease its exposure should the project fail. Thus, by decreasing one's holdings, one can realize as great an upside gain and lessen the exposure. What's more, as is shown in Clyman *et al.* (1997), the optimal share never declines so much as to end up with a lower upside payment. That is, the product of the new (decreased) optimal share times the new (increased) upside payoff is always greater than the old optimal share times the old upside payoff.

Thus, by utilizing the decision-analytic techniques available to us to aid in making decisions about risky projects, we can not only make better decisions, but also inform our intuition so that we do not seize upon poor choices thinking they are obviously right.

Summary and conclusions

We have shown that for managers in resource firms who are concerned with integrating a comprehensive discounted cash flow model with a corporate risk management decision model, the decision-analytic approach has important implications. Utilizing these tools, resource managers are able consistently and systematically to evaluate a wide range of risky investment opportunities. These techniques enable managers to rank projects and select participation levels consistent with the firm's willingness to take on risk.

When utilizing the exponential utility function, at first glance we appear to have uncovered a somewhat counter-intuitive finding with regard to choices about project share. That is, it may appear to be somewhat counter-intuitive that a manager would want less of a project if the only change was an increase in the success payoff, x_1

However, as we have shown, this seemingly counter-intuitive result may actually make a great deal of sense, especially if one is more concerned about achieving a particular upside while minimizing the downside, than one is about maximizing the upside without so much regard for downside consequences.

What this finding tells us is that while we may want less of an improving prospect, we never want too much less. In other words, though we may decrease our optimal share a little bit, thereby decreasing the optimal cost upon failure, we still allow the optimal payoff upon success to increase.

Experimental and survey data in behavioral decision making also show that the preference analysis approach is consistent with how managers make decisions under uncertainty. According to Wehrung (1989) study of 127 executives from 29 oil and gas firms, more than half of the executives gave responses that were fully consistent with this approach, and an additional quarter of executives were consistent within a 10% margin of error. This behavior, as well as other published work mentioned earlier, suggests that managers would benefit from using decision-analytic techniques to help make project-selection choices and risk-sharing decisions. It is very important, therefore, that managers pause and take the time necessary to do appropriate analysis when making project-selection and risk-sharing decisions.

Decision-analysis techniques offer great benefits for systematic, scientifically grounded decision making. These techniques increase management's awareness of risk and risk tolerance, provide insight into the relative financial risks associated with its set of investment opportunities, help to inform and improve our intuition, and provide the company a formal and consistent approach for allocating scarce capital.

References

Clyman, D R (1995) Unreasonable rationality?. *Management Science* **41**, 1538–1548.

Clyman, D R, Walls, M R and Dyer, J S (1997) Too much of a good thing? Under review at Operations Research.

Corner, J L and Kirkwood, C W (1991) Decision analysis applications in the operations research literature 1970–1989. *Operations Research* **39**, 206–219.

Cozzolino, J (1977) A simplified utility framework for the analysis of financial risk. Economics and Evaluation Symposium of the Society of Petroleum Engineers, Dallas, Texas.

Grayson, J C (1960) Decisions under uncertainty: drilling decisions by oil and gas operators. Harvard Business School, Division of Research, Boston, MA.

Hax, A C and Willig, K M (1977) The use of decision analysis in capital investment problems. In *Conflicting Objectives in Decisions*, ed. Bell, D E, Keeney, R L and Raiffa, H. John Wiley and Sons, New York.

Howard, R A (1988) Decision analysis: practice and promise. *Management Science* **34**, 679–695.

Keefer, D L, Smith, F B and Back, H B (1989) Development and use of a modeling system to aid a major oil company in allocating bidding capital. Technical Report DIS 88/89-11, College of Business, University of Arizona.

Keeney, R L and Raiffa, H (1976) *Decisions with Multiple Objectives*. John Wiley, New York.

MacKay, J A, Lerche, I and Rose, P R (1996) Modifications to risk aversion in high gain situations. Manuscript submitted to the Society of Petroleum Engineers.

Newendorp, P D (1975) *Decision Analysis for Petroleum Exploration*. Penn-Well Books, Tulsa, Oklahoma.

Pratt, J W, Raiffa, H and Schlaiffer, R O (1964) The foundations of decision under uncertainty: an elementary exposition. *Journal of American Statistical Association* **59**, 353–375.

Raiffa, H (1968) *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*. Addison-Wesley, Reading, MA.

Savage, L J (1954) *The Foundation of Statistics*. Wiley, New York.

Spetzler, C S (1968) The development of a corporate risk policy for capital investment decisions. IEEE Transactions on Systems Science and Cybernetics Institute of Electrical and Electronics Engineers SSC-4.

von Neumann, J and Morgenstern, O (1953) *Theory of Games and Economic Behavior*, 3rd edn. Princeton University Press, Princeton, NJ.

Walls, M R (1995) Corporate risk tolerance and capital allocation: a practical approach to implementing an exploration risk policy. *Journal of Petroleum Technology* **4**, 307–311.

Walls, M R and Eggert, R (1996) Managerial risk-taking—a study of mining CEOs. *Mining Engineering* **3**, 61–67.

Walls, M R, Morahan, G T and Dyer, J S (1995) Decision analysis of exploration opportunities in the onshore US at Phillips Petroleum Company. *Interfaces* **25**, 39–56.

Wehrung, D A (1989). Risk taking over gains and losses: a study of oil executives. In *Choice Under Uncertainty, Annals of Operations Research*, Vol. 19. eds. Fishburn, P C and LaValle, I H J C. Baltzer Scientific Publishing Company, Basel.