

On the use of risk and decision analysis to support decision-making

T. Aven* and J. Kørte

Stavanger University College, P.O.Box 8002, N-4068 Stavanger, Norway

Terje.aven@tn.his.no, jens.korte@tn.his.no

Abstract

Risk analysis is always part of a decision context. While it is essential that risk analysts understand the decisional context of their work and results, this is often not very clear in risk analysis and engineering texts. Important knowledge for risk analysts are what tools are available for decision analysis and what the methodological and practical strengths and limitations of alternative approaches are.

In the present paper we discuss and give recommendations on the choice of approach to decision problems involving risk and uncertainty. Using an example from investment in risk reducing equipment on offshore helicopters, we illustrate how information from risk analysis is used to reach a decision, combined with cost-effectiveness indexes, cost/benefit calculations and expected utility optimisation. Methodological differences are discussed. It is argued that, for risk related organisational decisions, presentations of consequence and uncertainties from the risk analysis should be highlighted, rather than synthesized measures of utility gains and losses. Results from decision analysis –cost/benefit as well as expected utility— are thus seen as secondary information and strictly as decision aid, rather than methods for prescription of choice. In a model for decision-making, this requires an explicit element of review and judgement by the decision maker. Cost/benefit calculations can provide results that are comprehensible, but will often cause disagreement, about the procedure and reference for the valuation of consequences. We argue that expected utility, although theoretically a more well founded approach, is seen to be of limited practical use.

Keywords: Risk Analysis; Decision analysis; Cost-Benefit; Expected Utility

1. Introduction

The optimisation of the expected utility is the ruling paradigm among economists and decision analysts. The paradigm represents an idealisation of a decision making process and, as thus a reference for other approaches. But in practical application it is just a tool, a normative theory saying how to make decisions strictly within a mathematical framework - it does not replace a management review and judgement process of results produced by a formal analysis. There are factors and issues that go beyond strict mathematics that a management needs to consider. In practice, there will always be constraints and limitations restricting the direct application of the expected utility thinking. Expected utility theory provides guidance on how people ought to make decisions, not how they are made today. Descriptive theory, theory on how people actually behave, can give insights into decision makers behaviour and limitations, but such theory does not replace normative theory. Descriptive theory provides as such little concrete guidance for attacking new, specific decision problems of some complexity. Here, a method, a prescriptive tool is required, which provides a consistent procedure to process the elements of the decision problem and aids in making a decision. Such tools are often identified with normative theory. Expected utility theory can be seen as such a tool and is by many recognized as a good framework for making decisions, as its assumptions and coherency are hard to disagree with on principle.

* Corresponding author. E-mail address: terje.aven@tn.his.no

Cost/benefit analysis is another method for balancing costs and benefits and has been closely tied to decision-making in the ALARP (As Low As Reasonably Practicable) region, cf. e.g. [21, 25]. The idea of the method is to assign monetary values to a list of burdens and benefits, and summarize the "goodness" of an alternative by e.g. the expected Net Present Value (NPV) of costs and equivalent monetary values of benefits and burdens. The method has been subject to strong criticism, due to both philosophical objections, see e.g. [1, 23] and references therein, as well as objections to the formal treatment of values and weights, see e.g. [9, 6, 1].

One major problem is related to the transformation of non-economic consequences such as (expected) loss of lives and damage to the environment, to monetary values. What is the value of a (statistical) life? What is the value of future generations? The answer to these issues is neither trivial nor obvious and all methods of quantifying the value of a risk, such as willingness-to-pay/accept or revealed preferences can be questioned on a number of grounds, see the above references. Cost/benefit analyses are often based on simplifying assumptions and will just focus on certain consequences and ignore others. The probabilistic basis for cost/benefit analysis is seldom clarified, but the classical thinking with a search for correct, objective probability values seems to be dominant. It is common to question the validity of cost/benefit analyses because of their unrealistic assumptions about the availability of data needed to complete the analyses. The underlying philosophy seems to be, that without objective, hard data the analyses break down. Expected utility analysis requires assessments of choices among hypothetical alternatives to define utility functions, and is difficult to carry out in practice. One can question if its assumptions are satisfied in practice and whether the required procedures are practical.

Now, with all these caveats, how should the results of risk analysis be processed and presented to decision makers? For a risk assessor confronted with the fierce pro et con arguments for the available decision analysis approaches, the answer cannot be obvious. Shrader-Frechette [23] rightly states that it is not sufficient to find faults with a technique to reject it. One must identify an alternative that is less defective. We would add that, in addition to satisfying philosophical and methodological requirements, the analysis must be seen as useful by decision makers and stakeholders.

In the present paper, starting from risk analysis, we assess alternative approaches and discuss some of the above issues. Applying decision analysis approaches to an example of a risk related investment from helicopter transport in the offshore oil industry, we argue that the assessment of the attributes associated with the alternatives provides useful insight and decision support in many applications. We will demonstrate our scepticism to a mechanical transformation of consequences to monetary values for both cost/benefit and expected utility analysis; in many cases it is more informative to put attention on each consequence and the associated uncertainty separately. The weighting of consequences and risk is ultimately the responsibility of the decision maker (the 'management'), as part of a less formal review and judgment process. A foundational assumption is that risk expresses uncertainty about the world, or quantities representing the world. Risk analysis produces predictions of these quantities and related uncertainty assessments, expressed by probabilities. The approach is in line with a predictive Bayesian approach, cf. [3, 2]. In Section 2 we present a model for decision-making, following the ideas sketched above, and review the basic elements of cost/benefit analysis and expected utility theory. The presentation of the theory in this section is somewhat elementary, but this is considered necessary, as the readers from the engineering risk community would not, in general, be familiar with the details of the alternative theories. The issue here is not to present technical details of each approach—the literature on these subjects is vast—but to discuss the practical meaning and application of the approaches in 'normal' decision situations.

In Section 3 we discuss the model introduced, the choice of decision analysis tool and the interpretation of risk and decision analysis results. Finally, in Section 4 we draw some conclusions.

2. Features of a decision-making approaches

Basically, we will contrast two different approaches of thinking to reach a good decision;

- i) Decision making as an exercise of modelling alternatives, outcomes, uncertainty and values, and choice of the alternative which maximises/minimises some specified criteria.
- ii) See decision-making as a process with formal risk and decision analyses to provide decision support, followed by an informal managerial judgement and review process resulting in a decision.

The first approach is typical for the pioneers of the economic decision making school [19], and later the Bayesian decision making theorists, see e.g. [16, 18]. In this paper the second approach is adopted as an overall structure, meaning that we see decision analysis strictly as an aid for decision. This is comparable to the "moderate" view as mentioned by [9] and proposed by a number of decision theorists, see e.g. [26, 12]. The decision maker needs to take the results of the decision analysis and make his decision, following a review and judgement process. This does not mean that we can not see examples where approach i) may be appropriate, but considering varying degrees of the informal managerial judgement and review process preceding any decision, we may think of approach i) as a special case of approach ii).

Regardless of the approach, we will not be able to avoid the fact that some decisions will be followed by negative outcomes. But by following a decision-making process in line with the principles in approach ii), we would expect that a collective of decisions will produce positive outcomes in relationship to the objectives. The main features of approach ii) is shown in Figure 1. Variations of this approach are discussed in [14]. The starting point is a decision problem and often this is formulated as a task of choosing among a set of decision alternatives.

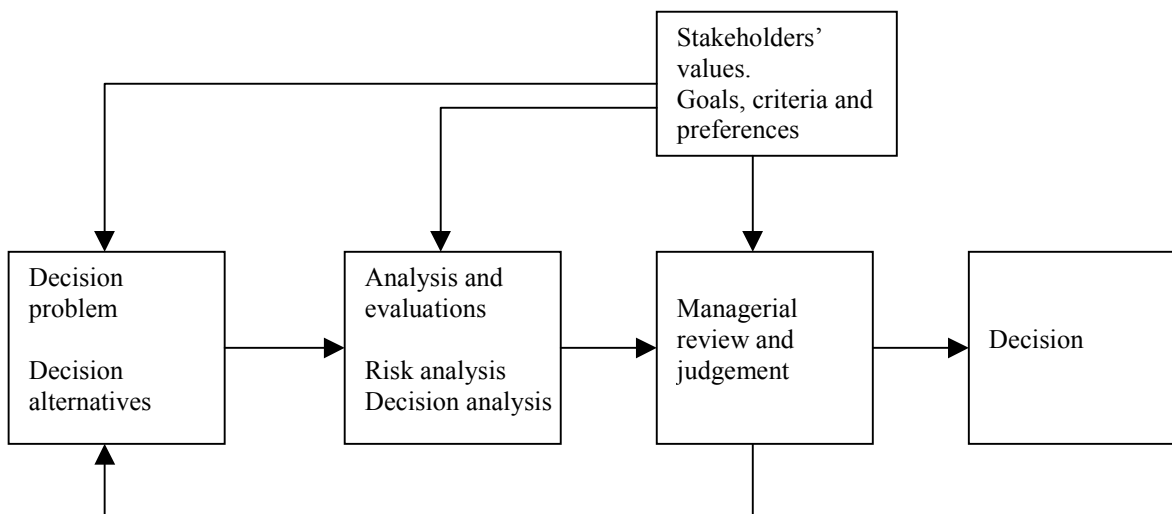


Figure 1 Basic structure of the decision making process

2.1 Defining the decision problem and the alternatives

A decision problem can arise in different ways. Choices among investment alternatives or alternative available technical solutions to a development project are examples. Another example is when the risk of negative consequences of an ongoing activity is such that an action needs to be taken. The latter will be illustrated by an example.

At an early stage a number of possible alternatives may be available. The problem is to identify one or two for further detailing and optimisation. Clearly, much has already been decided when a set of alternatives to be further evaluated has been defined. If, for example for an oil exploration, one were to decide at an early stage to adopt a well-proven technology, one would exclude alternatives involving development of new technology. In a practical setting the number of alternatives to be evaluated need to be manageable, the consequence being that many alternatives could be excluded at an early stage when uncertainties are large. Further studies could have shown that these alternatives are favourable compared to those being evaluated, but this will remain a purely hypothetical problem. The set of alternatives is typically defined through an integrated process involving experts and managers. The experts would often specify an initial list as a basis for discussions. The development of alternatives would to a large extent be driven by the boundary conditions of the decision problem, as judged by the experts and the management. The boundary conditions include stakeholders' values, for example formulated as organisational goals, criteria, standards and preferences, as well as views expressed by politicians, environmentalists and others. Both experts and managers have a background, i.e. knowledge, experience, values, preferences, etc. that could significantly influence the selection process of alternatives. We have to appreciate the subjective element in this creative part of decision-making, establishing an appropriate set of alternatives. People may have personal agendas, but by ensuring that the process involves a sufficient broad group of personnel, the generated alternatives should provide the necessary basis for identifying a good alternative.

Based on the frequency of accidents of helicopters supporting North Sea oil exploration between the 1960s and the 1980s, the risk associated with helicopter transportation was considered high. Independent risk analyses were performed, and investigations into contributing factors. Health monitoring systems (HUMS) were identified as a measure which would reduce risk significantly, see [13, 8, 15]. This conclusion was largely derived through expert judgement by groups of experts from several fields of expertise (technical, operational, legislative and users). The conclusions from the reports focused the attention of the on-following decision process on the most promising alternatives, i.e. HUMS.

The risk analyses were more or less formal, and so was the assessment of risk reducing measures. But, to our knowledge, no formal decision analyses with costs and burdens of the mitigating alternatives were performed. After a process of discussion and negotiation, involving commercial negotiations between system suppliers, helicopter operators and oil companies and their associations, the decision was taken to implement HUMS with the objective of reducing the accident risk.

Based on the reports of the risk analysis and with hindsight knowledge of the decision, it is of interest to study how alternative decision analysis approaches could have resulted in the decision that was actually taken.

2.2 Decision support tools

Now, suppose we have given a set of decision alternatives, A and B. Before a decision can be made, some support is needed, as a basis for the decision. The consequences of choosing one alternative instead of another need to be known. Risk analysis provides such decision support. It gives predictions of the performance of the various alternatives with related uncertainty assessments. This information is then linked to costs and benefits associated with the various alternatives. There are many ways of thinking on how to integrate risk, costs and benefits in problems involving multiple attributes. All (rational) approaches involve prediction of future consequences. They differ in the degree to which they include uncertainty, valuation and trade-off of attributes and representation of risk attitude, as well as methods for prescribing choice. The

historical and theoretical foundation of the methods varies. In the following we will discuss two typical approaches :

1. Presentation of consequences and risk of selected attributes for the various alternatives, with flexibility in attribute weighting and trade-off.
2. Formal Bayesian decision analysis.

Consider again the HUMS case introduced above. According to approach 1, the idea is to analyse and evaluate factors such as the impact on safety, the investment and operational costs, market impact and operational regularity and maintenance impact. For costs, safety, maintenance and operational issues, quantitative analyses could be conducted producing predictions and uncertainty assessments and cost/benefit trade-offs. For other important aspects such as political or public opinion pressure, only qualitative analyses and evaluations would normally be performed. The total of these analyses and evaluations is used as a decision basis. Before a decision is made, the management reviews and evaluates the decision support information, and relates it to values formulated as goals, criteria and preferences. There is no strict procedure on how to do this managerial process. It is an individual process based on the constraints of the structure given by Figure 1. This structure can be interpreted as a prescription, but it is not very detailed and specific.

Returning to our example, the decision makers, in this case the Norwegian oil operators, were faced with a risk which could be presented as the probability distribution of helicopter accident fatalities for the next years, as shown in Figure 2.

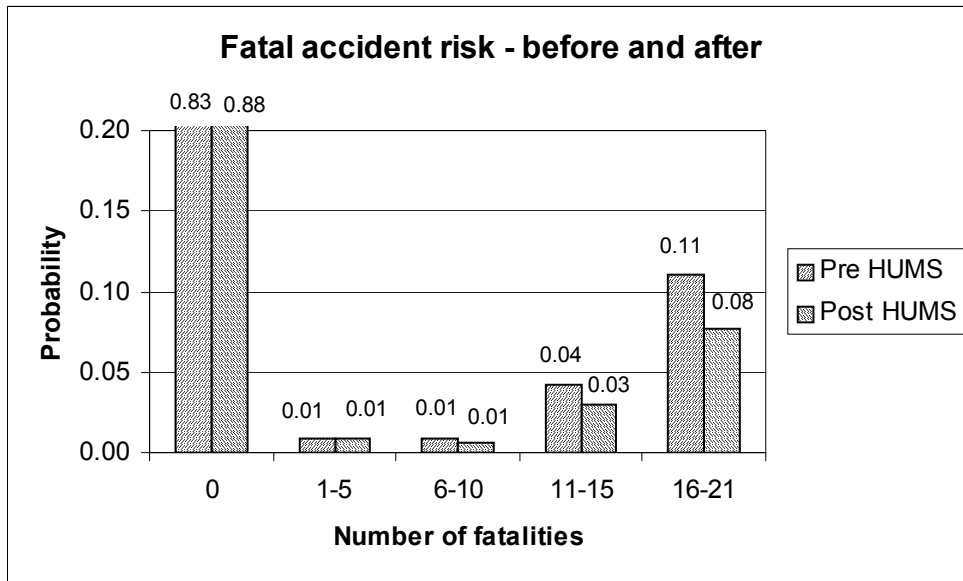


Figure 2 Fatal accident risk, expressed as the probability distribution for the number of fatalities in one year; before and after HUMS, for operations on the Norwegian shelf

The shown distribution is not part of the referred report, but has been constructed on the basis of approximate figures. The figures are based on the expectancy, a FAR of 400 (FAR is defined as the expected number of fatalities per 10^8 exposure hours), a prediction of 40000 flight hours per year, and an average of 17 passengers and crew per flight. As an alternative or as a supplement to the distribution, the expectancy of fatalities for one year could be stated, which for this case is 2.7. With the installation of HUMS on the whole fleet, the reduction in fatality risk, expressed as the

expected number of fatalities, was judged to be from 15% to roughly 50%. In the following we will assume that the risk analysis used as a decision basis showed a decrease in fatality risk as shown in Figure 2. The resulting reduction in the number of expected fatalities could be stated as 1.9, i.e. a reduction. The distribution provides additional decision support, as it informs the decision maker that the expected reduction is to be highest for accidents with high numbers of fatalities (this is in accordance with the expert panels' assessments in the SINTEF report for technical failures in-flight, see Ingstad et al. [1990]).

The presentation of consequence distributions in such a manner can be regarded as the most fundamental basis for decision makers. It is the basis for any associated formal decision analysis and should be part of the decision support information. In some cases it may be the only information required. But, comparing alternatives in this way can be challenging, as a number of factors of different dimensions need to be considered. How should we balance for example cost and safety. We will treat the presented example initially using a cost/benefit approach. Cost/benefit, in a narrow sense, is the expected net present value of costs and the cost trade-offs of other attributes, cf. e.g. [25, 20]. The ALARP applies a principle of "reasonableness", the operational interpretation of which is closely tied to the narrow interpretation of cost/benefit. This is in line with standard textbooks on cost/benefit analysis, see e.g. [24]. As used in the paper here, cost/benefit analysis also encompasses a class of analyses labelled "cost-effectiveness" analysis, see e.g. [7], [9], i.e. we treat cost/benefit in a somewhat wider sense.

Cost/benefit analyses

Let us look at a simplified case of the HUMS investment, with two alternatives compared; alternative A: reject the investment in HUMS and, alternative B: invest in HUMS and receive the effect on risk. The decision makers in this case were the Norwegian oil companies, who had to make a decision, based on the risk analysis and a cost proposal from the helicopter operators. Both profit and cost are assumed to be fixed, i.e. no uncertainty is associated with these. The cost of HUMS used here, is a yearly cost of \$ 3 million, reflecting both annualised capital cost and operational cost.

Balancing the cost and accident risk is a management task, which is based on goals, criteria and preferences, but in most cases there is no obvious line to a specific decision. In our example, alternative A means a reduced cost compared to B, but a higher probability of a fatality. How should one value the difference in probability distributions or the expected number of fatalities? The statistical expected fatalities for the two alternatives are 2.7 and 1.9, respectively. For the sake of simplicity we assume that there are no other factors to consider. What alternative should be chosen? Of course, no general answer exists to this, but we can compute a cost index expressing cost per expected life saved, which gives a reference and a link between the two dimensions costs and safety. We see that the cost index, which makes the two alternatives equal, in this case is $3/(2.7-1.9)$, which is approximately equal to 3.7 ($\$ \times 10^6$) per statistically lost life. The reasoning is as follows: To go from alternative A to B, it would cost 3, and the expected number of saved lives would be 2.7-1.9. In the terminology of some authors, see above, this is termed a cost-effectiveness calculation. Studying a number of activities, we can make comparisons and conclude whether an cost-effectiveness index of 3.7 is high or low.

In many situations, this type of information, combined with the results of the risk analysis is what is needed for decision makers in comparing the effects of alternative actions and choosing among alternative mitigating actions, when resources are limited. In addition to the cost-effectiveness index, risk analysis provides the basis for a decision, as the decision makers might have a special interest to reduce the probability for accidents with high numbers of fatalities.

Does this approach represent a complete decision analysis? It enables the comparison of a set of alternatives for a given objective. It does not enable these alternatives to be measured against an

absolute reference for acceptance or rejection of risk. A cost/benefit analysis assumes this measure to be expected monetary value and trades other attributes or objectives against gain in expected monetary value. This implies that we must explicitly state the exchange in expected money for expected value of all other objectives we might have.

Following our example, we will limit the analysis to the objectives safety and profit, i.e. monetary gain. With alternative A, we associate a gain of 1000 (i.e., a cost -1000), the profit of the same industry using helicopters and with alternative B, a reduced gain of $1000 - 3 = 997$ (i.e., a cost of -997). We may think of one cost unit as one million dollars.

For the cost/benefit analysis to be complete, we are required to state a cost value for a statistical life, or here, the cost that is set per expected saved life. This value should not be confused with the cost-effectiveness index calculated above. Let the value we assign for the cost per expected life be \$7 million. Then the total statistical expected gain associated with alternative A would be

$$1000 - 7 \cdot 2.7 = 981,$$

for alternative B, the corresponding value would be

$$997 - 7 \cdot 1.9 = 984.$$

The conclusion would thus be that alternative B is preferable as the gain is 984 compared to a gain of 981 for alternative A. In practice one might need to take into account time and the discounting of cash flow, but the above calculations show the main principles of this way of balancing cost and benefit. Note that by using the total profit of the companies as the other consequence, we are scaling the cost benefit calculation. In this case the relative difference in total expected cost/benefit is small. The structure of the decision-making process when cost/benefit analyses is used is shown in Figure 3.

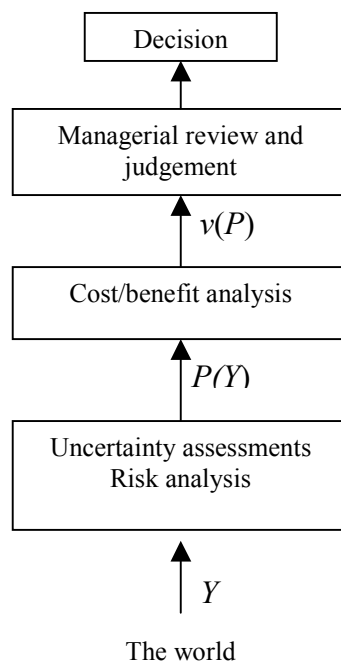


Figure 3 Basic structure of the decision-making process when cost/benefit analysis is used

The starting point is the world and observable quantities, here represented by Y . Risk and uncertainty analyses are conducted producing probabilities, denoted $P(Y)$. Based on these analyses, cost/benefit analyses are carried out resulting in performance measures $v(P)$, for example expected cost per expected saved life. These measures are reviewed and a decision is made.

Note that in the analysis we have not specified the value of a life, but probabilities and expected numbers of fatalities. This might appear to be solely semantic, but we would like to see it as an important ethical discrimination. The number used for a statistical life or, to be precise, one expected fatality, does not imply that anyone in reality would trade a specific life for a higher value. In our example, if a company knew that a specific person would die with certainty, but his/her life could be saved for a \$7 million that could be made available, could there be doubt? The value of a life is infinite. Using cost/benefit numbers in such a way would be a cynicism not warranted by the idea behind the cost/benefit approach. It merely reflects recognition (resignation?) that in any situation, resources are limited and potential actions to reduce uncertainty and risk are indefinite. Cost/benefit values of expected fatalities, in this situation, provide an instrument to relate to uncertainty and risk in decision making. The chosen values, in our example \$ 7 million, would always refer to specific methods or standards of comparing risks and benefits. It is a value that can only be related to a large population of people exposed to risk, not to decide on life or death of individuals. For a discussion of the principles of placing value on risk, see [9]. Often, these will not be obvious, but they are important underlying principles. Recently, attempts have been made to establish context-independent measures for life and life-saving measures, such as the life quality index, see e.g. [17]. Although such measures provide interesting reference material, we cannot see that such measures overcome the fundamental problems discussed. On the contrary, we are concerned that such technical measures may lead people into mechanical trading of safety measures and costs.

In this paper we focus on tools for decision-making, and cost/benefit analysis as described above is just an example of such a tool. It provides input to the decision-maker, but it does not prescribe the decision. By presenting the results of the analysis as a function of the consequences, here the value of a statistical life, we can demonstrate the sensitivity of the conclusions of the analysis, see Figure 4.

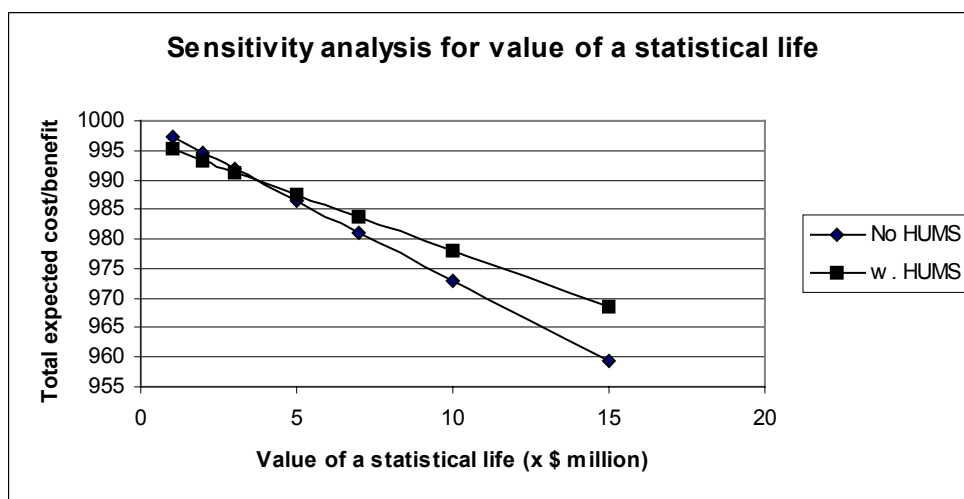


Figure 4 Sensitivity of expected cost/benefit depending on the value of a statistical life

The results of the decision support analyses need to be evaluated in light of the premises, assumptions and limitations of these analyses. The analyses are based on background information that must be reviewed, together with the results of the analyses. Considerations need to be given to factors such as

- The decision alternatives being analysed.
- The performance measures analysed.
- The fact that the results of the analyses represent judgements, to a large extent conducted by some experts.
- The difficulty of assessing values for costs and benefits, and uncertainties.
- The fact that the analysis results apply to a model, i.e., a *simplification* of the world, and not the world itself.

Such considerations are an essential element of the managerial review and judgment process. For example, in the model we may have defined just five possible consequences of an event—for the HUMS example only two—even though in the real world there were many aspects and facets that need consideration. In real life cases there will always be factors of relevance that go beyond the results of the formal analyses.

We have used the term cost/benefit analysis as a category of analysis techniques which is based on calculating expected costs and benefits. In the cost/benefit analysis, expectations of all consequences are integrated into one unit, typically the expected net present value (NPV). This requires a transformation function for all consequences into monetary units. For NPV calculations, time discounting factors are required as well. In such a setting, the analyses provide decision support, not hard recommendations. Thus we may for example consider different values of a statistical life, to provide some insight for the decision. Searching for a correct, objective number, is meaningless, as no such number exist; the statistical life used in the analysis is a value that represents an attitude to risk and uncertainty, and that attitude may vary and depend on the context. When using a one-dimensional scale, uncertainty of observable quantities are mixed with value statements about how to weigh the different assessed uncertainties. In a multi-stakeholder decision process, we would expect to encounter difficulties in obtaining consensus about the recommendations provided by the cost/benefit analysis, as there are always different opinions about how to look at risk in a society, and the result of the analysis hides many assumptions and calculations. Do we expect a group of stakeholders with different agendas, e.g. management, union and public authority representatives to agree on a decision based on cost/benefit values of 984 and 981 alone, for the two alternatives considered above? –Hardly. For example, in a traditional cost/benefit analysis, an alternative with a high expected cost/benefit is preferred to an alternative with a lower cost/benefit value, even if the latter alternative would mean a negligible probability of a serious hazard, whereas this could not be disregarded in the former case. In traditional cost/benefit analysis it has also been common to discount the values of statistical lives on the same scale as monetary discounting, the results being often that negligible weight is put on consequences affecting future generations. Clearly, it is of paramount importance that the cost/benefit analyses are being reviewed and evaluated, as we cannot replace difficult ethical and political deliberations with a mathematical one-dimensional formula, integrating complex value judgements.

Formal Bayesian decision analysis. Maximization of expected utility

The expected utility approach does not require monetary value as an assumption. Utility is derived from preferences among outcomes. It uses a principle of indifference between certain outcomes and uncertain outcomes to define a subjective value scale –utilities—to different outcomes, as will be illustrated below. The cost/benefit analysis approach requires balancing various assessed

uncertainties and expected values, i.e., in the example costs and accident risk, whereas, when using a Bayesian utility approach explicit trading of costs and consequences, here the number of fatalities is required. Using the same example as above, the possible consequences for the two alternatives are $(1000, X)$ and $(997, X)$, where the first component of (\cdot, \cdot) represents the benefit and X represents the number of fatalities which for our example is taken as either the number of fatalities, given an accident, 15, or 0.

Now, what is the utility value of each of these consequences? Well, the best alternative would obviously be $(1000,0)$, so let us give this consequence the utility value 1. The worst consequence would be $(997,15)$, so let us give this consequence the utility value 0. It remains to assign utility values to the consequences $(1000,15)$ and $(997,0)$.

Consider the standard of balls in an urn with u being the portion of balls that are white. Let a ball be drawn at random; if the ball is white the consequence $(1000,0)$ results, otherwise, the consequence is $(997,15)$. We refer to this lottery as " $(1000,0)$ with a chance of u ". Now, how does " $(1000,0)$ with a chance of u " compare to achieving the consequence $(997,0)$ with certainty? If $u = 1$, the gamble is clearly better than $(997,0)$, if $u = 0$ it is worse. If u increases, the gamble gets better. Hence there must be a value of u such that you are indifferent between " $(1000,0)$ with a chance of u " and a certain outcome $(997,0)$, call this number u_0 . Were $u > u_0$, the urn gamble would improve and be better than $(997,0)$; with $u < u_0$ it would be worse. This value u_0 is the utility value of the consequence $(997,0)$. Similarly, we assign a value to $(1000,1)$, say u_1 . As a numerical example, we may think of $u_0 = 99/100$, and $u_1 = 1/10$, reflecting that we consider a life to have a high value relative to the gain difference. Now, according to the utility-based approach, that decision should be chosen which maximizes the expected utility.

For this simple example, we see that the expected utility for alternative A is equal to

$$1.0 \cdot P(X = 0) + u_1 \cdot P(X = 15) = 1.0 \cdot 0.83 + 0.1 \cdot 0.17 = 0.839$$

whereas for alternative B,

$$u_0 \cdot P(X = 0) + 0 \cdot P(X = 1) = 0.9 \cdot 0.88 + 0 \cdot 0.12 = 0.836.$$

The conclusion from this analysis is that alternative A is to be preferred. Observe that the computed expected values are in fact equal to the probability of obtaining the best consequence, namely a gain of 1000 and no fatalities. To see this, note that for alternative A, the consequence $(1000,0)$ can be obtained in two ways, either if $X=0$, or if $X=15$ and we draw a white ball in the lottery. Thus by the law of total probability, the desired results follow for alternative A. Analogously we establish the result for alternative B.

We conclude that maximizing the expected gain would produce the highest probability of the consequence $(1000,0)$ and, as the alternative is the worst outcome, $(997,15)$, we have established that maximizing the expected utility value gives the best decision. This is an important result. We refer to [18] for further details on the expected utility approach. Based on requirements of consistent (coherent) comparisons for events and for consequences, we are led to the inevitability of using the expected utility as a criterion for choosing decisions among a set of alternatives.

The structure of the decision-making process when utilities are used is shown in Figure 5. As for the cost/benefit case, the starting point is the world, represented by a set of consequence variables, or attributes Y . Uncertainty assessments are conducted, i.e., risk analysis resulting in probabilities $P(Y)$, and utilities $u(Y)$ are elicited. Combining $P(Y)$ and $u(Y)$, we derive at the expected value, $Eu(Y)$, and maximization of this measure gives the optimal decision alternative within the given framework. Based on managerial review and judgment, a decision can now be made.

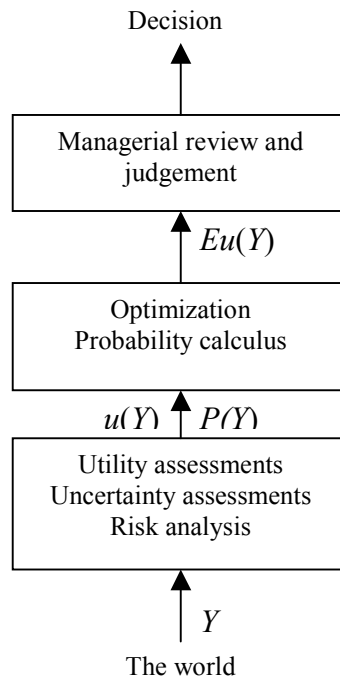


Figure 5 Basic structure of the decision-making process when utilities are used

What have we gained compared to the cost/benefit approach? We have not assumed monetary values for all consequences (except those that are expressed as money, i.e. the profit). We have assessed our preferences by comparing complete outcomes (all consequences) and have expressed the degree of preference by probabilities. The result is a utility calculation on a normalised scale, pretending to represent a total measure of good. Along the same argument as for cost/benefit: do we expect a group to reach agreement on a decision based on utility values of 0.839 and 0.836 alone? The assessment process of objectives and value trade-offs in the expected utility approach combined with the use of sensitivity analyses could provide knowledge and insight into aspects of the decision, but the utility approach is time and resource consuming, for both experts and management. Assessing utilities is not a trivial task. We always need to question whether the extra costs of formal benefit/utility assessments can be justified against the benefit of producing better decisions. The analysis needs to be evaluated, in light of its premises, assumptions and limitations, as discussed above for the cost/benefit analyses. In practice, decisions are often taken by a group of people with their own array of probabilities and utilities, but the expected utility approach is only valid for a single decision-maker. No coherent approach exists for the multiple decision maker problem. Of course, if the group can reach consensus on the judgments, probabilities and utilities, we are back to the single decision-maker situation. Unfortunately, life is not that simple in most cases; people have different views and preferences. Reaching a decision then is more about discourse and negotiations than mathematical optimisation.

3. Discussion

We have considered different approaches for aiding decision-making involving costs, benefits and risk, cost/benefit analyses and maximization of expected utility. Now, should any single approach be taken as prescriptive for all decisions?

The expected utility approach is attractive as it provides recommendations based on a logical, mathematically coherent basis. If a person is coherent both in his preferences amongst consequences and in his beliefs about uncertain quantities, it can be shown that the rational way for him to proceed, is by maximizing expected utility.

For a person to be coherent in his assessment of uncertainties of events, the requirement is that he follows the axioms of probability. For consequences, coherence means that if c_1 is preferred to c_2 , which is in turn preferred to c_3 , then c_1 is preferred to c_3 . What we are doing is making an inference according to a principle of logic, namely that preferences should be transitive. Given the framework in which such maximization is conducted, this approach has a strong appeal to our sense of rationality. The question remains whether it provides information to decision makers which is perceived as relevant and a legitimate basis for the decision. Some of the problems with this approach have been discussed already. An important point when comparing it with the cost/benefit analyses as decision aid, is that preferences have to be specified for *all* consequences, which is a difficult task in practice, and more important, not necessarily something that a management would like to do. Refer to the example above where utilities were established for cost and loss of lives. Specifying a value of a life is required. Using a cost/benefit analysis the value of a statistical life is of interest, not the value of life. The utility approach requires values (utilities) to be assessed for quantities of all consequences, Y . For the cost/benefit approach, the value judgements to be made by the decision-makers relate to $P(Y)$, the probability assignments, and not Y . Thus in the cost/benefit case we assess values in a "world" constructed by the analysts, not the real world as in the utility based approach. For many consequences, e.g. consequences sensitive to moral judgement, it is easier to relate to this constructed world, as we can make use of appropriate summarizing performance measures. The simple example discussed in the previous section, demonstrates this. The use of reference gambling to produce the utilities is hard to carry out in practice, in particular when there are many relevant factors, or attributes, measuring the quality of an alternative. Various approaches of multi-attribute utility theory, e.g. [16, 10, 5] offer simplification in the assumptions and advice on practical methodology. In theory, the meaning of the total expected utility is clear from the assumptions, definitions and the derivations as originally lined out in [19, 22] and operationalised by e.g. the above mentioned authors. In practice, we doubt that practitioners of significant decisions will be able to relate to the resulting utility numbers as a basis for decisions. The utility must always be seen in relation to the defined worst and best outcomes. The implications are not easy to grasp for the relatively trivial example used here. For a truly multi-objective problem it will be even more difficult to comprehend the meaning. Decision makers are not omnipotent actors, dealing in isolation. They are accountable to others, often with even less capacity than themselves to study underlying assumptions. Communication of the basis for a decision is important. Theoretical soundness of the analysis approach may not be the ultimate requirement for applied decision analysis. Are there decision situations where this type of approach does have a role? We can see the expected utility approach to be of relevance in repeated, standardised decision situations, where the consequences/attributes of the decision are the same, the stakeholders are few and the values and trade-offs do not vary from situation to situation. An example could be e.g. a bank's decisions to grant mortgages based on risk assessment of the loaner and the project. A number of factors could be relevant to consider. Standardisation of the appreciation of the individual factors is desirable and a certain non-linearity of the utility of some factors might be considered. The stakeholders will normally be the bank and the loaner. In such settings, the expected utility approach should be a strong candidate.

The cost/benefit approach is, in principle not very different to the expected utility approach. But many arguments exist against cost/benefit analysis, also from a methodological point of view. For the theme of this paper we find the arguments of [11] interesting, as they are comprehensive and view cost/benefit in a dichotomy to expected utility theory. We agree with French et al. in the critique raised over the traditional cost/benefit analysis interpretations, but we would claim that the arguments brought forward in favour of expected utility theory could more or less be applied also to cost benefit theory.

Traditional cost/benefit analysis is based on a notion of objective market prices. For most attributes this requires application of some method of contingent valuation (willingness-to-pay/accept), which invariably generates fundamental disagreement about validity. Multi attribute utility theory, on the other hand, is claimed to be explicitly subjective with respect to the values included. The question of objective pricing is clearly an issue of interpretation. We favour an interpretation of non-monetary benefit-costs or burden-costs to be interpreted as subjective magnitudes, which need to be seen on the background of the assessors and the method of assessment. We would like to emphasise here that cost/benefit assessments are related to future outcomes and are valuations of the expected values. These expected values are a weighted means over subjective probability distributions.

Discounting of benefits and burdens over time in cost/benefit analysis results in unreasonably low and even negligible regard of burdens for distant future generations. Expected utility theory in its original version addresses single time periods. In an extension, outcome streams of utilities over time can be considered in an almost identical manner to cost/benefit analysis. References [11] and [4] offer alternative approaches to better represent the effects and interests of future generations. As mentioned in the cited references, utility theory and cost benefit theory both apply discount rates to future values. If this is an unreasonable strategy then, alternative strategies and methods to account for future effects and consequences need to be developed in either case. We do not see this as an argument in favour of either cost/benefit or expected utility, but rather a fundamental discussion of how to value future non-monetary effects and rights of future generations.

In cost/benefit analysis "disproportionality factors" can be used, to account for non-linear risk preference, i.e. risk aversion. Multi attribute utility theory, on the contrary, uses a method of weighting attributes and subsequent averaging (calculation of expected value). Both methods can be seen as an attempt to model decision makers' valuing of uncertain outcomes different than the expected value. In cost/benefit approaches, valuing uncertainty is not an inherent element as it is for expected utility. As a result, there does not exist a unique approach of how to take account of risk aversion. But, in principle, there is no reason why value functions could not be transformed to account for risk aversion in cost/benefit analyses. This would come very close to the formulation of utility functions.

The arguments against cost/benefit analysis discussed here, appear to be formed by the perspective that expected utility theory, derived from measurement of preferences over uncertain consequences, represents a sounder model of human preferences than does cost/benefit. We argue, simply, that none of the analyses presented gives "objective" results. Both theories are seen as methodological frameworks to structure degrees of beliefs over outcomes and subjective values of these. The results produced have to be interpreted considering the method used as background knowledge. The results must be subjected to judgement before decisions are taken.

The review and judgement process involves taking into account factors outside the boundaries of the model and formal analysis. The decision scenario discussed above involves accident risk. A company may be under scrutiny from the public, authorities or unions for their safety policy and management. In this case, the perception of risk by the public and associated negative attention could influence the decision maker – regardless of the results of the formal analysis. If we take the case of an oil company, conservatism on safety issues might be seen as an advantage for future

exploration licenses. Such political considerations could be difficult to capture in the formal analysis, but would still be an important element of the decision making process.

The process of managerial review and judgement relies on meaningful information from the risk and decision analyses. It is therefore natural to ask how such complementary considerations communicate with the results of alternative assessment approaches. A thorough discussion is outside the scope of this paper, but it appears that the higher the degree of abstraction of the results, the more difficult it is to judge the influence of such factors. It would therefore appear that it would be more difficult to judge how the mentioned considerations would affect the results of a utility analysis, compared with a cost-benefit analysis.

As we discussed in the section on cost/benefit analyses above, the primary basis for decisions is the risk and uncertainty analysis. The various attributes: costs, safety, political aspects, etc. are analysed separately. One alternative, and not an uncommon one, is to leave it a management task to make a decision balancing the costs and benefits. The decision basis could be improved by supplying cost-effectiveness indexes for the risk reducing measures, as exemplified. No values of statistical lives or other unwanted consequences would be provided, as these would have to always refer to some more or less arbitrarily chosen standard. Would that mean lack of coherence in decision-making. Yes, it could in some cases. The ideal is not always attainable. If mathematical coherence means that we lose important aspects for the decision, then the result of a "coherent" decision analysis would become meaningless to the decision maker. The rationality of the result could certainly be questioned, and a more flexible approach should be chosen.

Rationality has a lot to do with consistency. If we adopt and declare the rules to which our statements or actions should conform, we act in a way that is consistent with them. In this sense, rationality appears to be close to the ideas of [26]. Rules are defined by the people involved in the decision (including stake-holders). This means that whether a behaviour is rational will depend on the rules adopted. We find Bayesian decision theory a sensible and (more than anything else) consistent framework for decision analysis. But it follows from our definition of rationality that people who do not abide by the precepts of Bayesian decision theory may be perfectly rational if they abide by their rules for decision making, which could involve informal review and judgement of the results of the risk and decision analysis. Consequently, if one adopts the structure for decision-making presented here, one would behave rationally, according to the rules set by this structure.

4. Conclusions

The issue of the present article has been to place risk analysis in a decision context and discuss alternative approaches to decision making and support. The focus is on decisions involving risk and uncertainty with the character of being 'one-time', i.e. not repeated situations of similar character.

There does not appear to exist consensus on the norms and approaches to valuation of consequences or complete methods of decision making. Neither does there appear to be hope that this will ever be achieved. Some issues in risk analysis and decision making are more contentious than others.

Although some authors express a general concern about the risks present in modern society, we believe that no real dispute exists about the need for risk and uncertainty analysis. For decisions involving risk, the primary input is explicit expressions about risk and uncertainty. Ideally this would be uncertainty distributions as illustrated in the present article; alternatively, expected values of consequences could be presented, although such summarising values often hide important aspects of the involved risk.

The next step in an analysis dealing with risk, could be to present cost-effectiveness index numbers, e.g. cost per saved expected fatality, for alternative risk measures. These do not express

preferences, but indicate the effectiveness of resource allocations to reduce risk and uncertainty. The result of these assessments could form the decision bases. The decision makers, in a process that we term review and judgement, would have to take into account the risk analysis model, the data, the analysis competence and other background variables. Combined with his/her assessment of other relevant factors and the value/preference associated with the risk picture, a decision would be taken. This is a process with no precise formal definition. On some occasions, more formal demonstrations of the valuation of consequences and risk are required, and even requirements for optimum solutions. This requires transformations of consequences and risk onto one scale and, for the latter requirement, synthesising of all measures into a single measure. In the present paper, we have discussed cost/benefit analysis, using monetary unit as the common unit, and expected utility, using an abstract dimension of utility as the measure.

Cost/benefit has been controversial due to philosophical, theoretical and methodological objections. The theoretical critique is based on its less stringent foundation, compared to the expected utility approach. Although this may be valid, we believe these objections to be of little practical importance. The methodological objections are related to the fundamental difficulties in who and how values of consequences and uncertainty should be established. We argue that in the case where decisions are based on cost/benefit assessments, the risk analysis results should be a visible part of the decision basis. Further it needs to be acknowledged that the decisions cannot follow directly from the cost/benefit results. These are input to the decision makers. They are processed during review and judgement. The decision does not need to reflect the cost/benefit optimum.

Although by many considered a normative reference for decision making, expected utility does not appear to be very widespread as an applied approach. We believe there are several reasons for this. The meaning of choices between lotteries is difficult to comprehend and it is difficult to perform in a practical setting. The expected utility as the ultimate, operational criterion for the decision is difficult to interpret and communicate. We do not see results of expected utility analysis being used to substantiate public risk related decisions. Communication of expected utility results would be a difficult challenge. As mentioned, we could see a role of expected utility for repeated decisions in situations with few stakeholders.

Acknowledgements

The authors would like to thank the referees for valuable comments and suggestions.

References

- [1] Adams J. Risk. London: UCL Press, 1995.
- [2] Aven T. How to Approach Risk and Uncertainty to Support Decision Making (to appear). New York: John Wiley & Sons Ltd., 2002.
- [3] Aven T. Risk analysis - a tool for expressing and communicating uncertainty. In: Proceedings of Esrel 2000. Edinburgh, 2000, A.A.Balkema, pp.21-28
- [4] Bedford T, Atherton E and French S. Time at risk: ALARP trade-offs over time. In: Proceedings of ESREL Safety and reliability, 1999, Balkema,
- [5] Bedford T and Cooke R. A new generic model for applying MAUT. European Journal of Operational Research 1999; 118 589-604.
- [6] Bedford T and Cooke R. Probabilistic Risk Analysis. Cambridge: Cambridge University Publishing Ltd, 2001.

- [7] Bell D E, Keeney R and Raiffa H, eds., *Conflicting Objectives in Decisions*, John Wiley & Sons Ltd., New York, 1977.
- [8] Bristow Helicopters Ltd. *Integrated health and usage monitoring systems*. Bristow, Redhill, 1989.
- [9] Fischhoff B, Lichtenstein S, Slovic P, Derby S and Keeney R. *Acceptable Risk*. New York: Cambridge University Press, 1981.
- [10] French S. *Decision Theory - An Introduction to the Mathematics of Rationality*. Chichester: Ellis Horwood, 1986.
- [11] French S, Bedford T and Atherton E. Supporting ALARP decision-making by Cost Benefit Analysis and Multi Attribute Utility Theory. *Journal of Risk Research* 2002, to appear.
- [12] French S and Rios Insua D. *Statistical Decision Theory*. London: Arnold, 2000.
- [13] Helicopter Airworthiness Review Panel. *Review of the helicopter airworthiness*. CAA-UK, London, 1984.
- [14] Hertz D B and Thomas H. *Risk analysis and its applications*. Chichester: John Wiley & Sons Ltd., 1983.
- [15] Ingstad O and al. e. *Helicopter Safety Study*. Sintef, Trondheim, 1990.
- [16] Keeney R and Raiffa H. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge: Cambridge University Press, 1993.
- [17] Lind N. Social and economic criteria of acceptable risk. *Reliability Engineering & System Safety* 2002; 78 21-25.
- [18] Lindley D V. *Making Decisions*. London: John Wiley & Sons Ltd., 1985.
- [19] Neumann J v and Morgenstern O. *Theory of Games and Economic Behaviour*. Princeton: Princeton University Press, 1944.
- [20] NORSOK. *Risk and emergency preparedness analysis - Norsok Standard*. Standard, Norwegian Technology Standards Institution, Oslo, 1998.
- [21] Pape R P. Developments in the tolerability of risk and the application of ALARP. *Nuclear energy* 1997; 36(6): 457-463.
- [22] Savage L. *The foundations of statistics*. New York: Dover, 1972.
- [23] Shrader-Frechette. *Risk and Rationality*. Oxford: University of California Press, 1992.
- [24] Sugden R and Williams A. *The Principles of Practical Cost-Benefit Analysis*. Oxford: Oxford University Press, 1978.
- [25] UKOOA. *A framework for risk related decision support - Industry guidelines*. UK Offshore Operators Association, 1999, pp. 31.
- [26] Watson S R and Buede D M. *Decision Synthesis*. Cambridge: Cambridge University Press, 1987.