Application of Supersoft Decision Theory in Fire Risk Assessment

HENRIK JOHANSSON[1,*] AND PER-ERIK MALMNÄS[2]

[1]*Department of Fire Safety Engineering, Lund University, S-221 00 Lund, Sweden*

[2]*Department of Philosophy, Stockholm University, S-106 91 Stockholm, Sweden*

ABSTRACT: The application of Supersoft Decision theory (SSD) to fire safety problems, and of decision analysis in general to decisions involving a high degree of epistemic uncertainty, are discussed. SSD and two traditional decision analytic methods employed earlier within the context of fire engineering are compared, particularly regarding how uncertainties are dealt with and the robustness of decisions – robustness concerning the likelihood that the alternative adjudged to be best will change when a reasonable degree of change in assessments of either the probabilities or the utilities involved occurs. Substantial differences between the three methods in decision robustness were noted. It was found that, since traditional decision analysis involving precise probability and utility values gives no indication of robustness, it can lead to incorrect conclusions, making it unsuitable in the present context. It is argued that methods not providing the decision maker with information on decision robustness are unsuitable in situations involving a high degree of epistemic uncertainty. A procedure is suggested involving use of Supersoft Decision theory and extended decision analysis to facilitate the choice between different fire protection alternatives for the case of a specific building.

KEY WORDS: decision analysis, fire risk analysis, epistemic uncertainty, supersoft decision theory.

INTRODUCTION

PERFORMING A QUANTITATIVE fire risk analysis for a particular building involves dealing with various uncertainties. First, one needs to address

*Author to whom corresponding should be addressed. E-mail: henrik.johansson@brand.lth.se

1042-3915/04/01 0055–30 $10.00/0 DOI: 10.1177/1042391504035523

the uncertainty regarding the outcome of a fire, which is usually done by constructing a model of various fire scenarios, for example by use of event trees. Uncertainty of this type, termed here aleatory uncertainty, has also been conceptualized in terms of irreducible uncertainty, inherent uncertainty, variability or stochastic uncertainty [1]. Secondly, one has to deal with uncertainty regarding the values of the variables used in the model of the different fire scenarios, the probability values, for example. Uncertainty of this type, based on lack of knowledge or information, is termed here epistemic uncertainty. It has been conceptualised as well in terms of reducible uncertainty, subjective uncertainty or cognitive uncertainty [1]. The present paper is concerned primarily with epistemic uncertainty in decision analysis concerned with investments in fire safety.

Note that the focus here is not on decision analysis concerning a category of buildings, such as, for example, analysis of strategies for reducing residential fire loss generally (see [2], for example). Instead, decision analysis of potential investments in fire safety for a *specific* building is of concern. The difference between analyzing decisions for a category of buildings, and doing so for a specific building, is usually substantial in terms of the amount of information available regarding various parameters of interest. One usually has some information about fires that have already occurred in buildings of a particular category, whereas one may very likely have little or no information about any previous fires in a specific building. Because of this, a decision analysis pertaining to a specific building is likely to involve a high degree of epistemic uncertainty, especially as regards extreme or catastrophic events which even in a large group of buildings occur very seldom. The question is how this uncertainty will affect a decision analysis in a specific building and what method or methods can be used to deal with the large epistemic uncertainties involved. In discussing the analysis of such uncertain decision situations here, an application of a decision analysis method called Supersoft Decision theory (SSD) is presented. This is a method specifically designed to deal with decision situations involving large epistemic uncertainties. The major aims of the paper are to present the conceptual framework of SSD, show how the SSD method can be applied to problems of fire safety, compare SSD with two alternative methods for decision analysis concerned with fire safety employed earlier and provide some general suggestions on how to evaluate decision situations in which an extraordinary degree of epistemic uncertainty exists.

The paper begins with a brief discussion of different criteria that could be useful in evaluating various alternatives in decision analyses concerned with investments in fire safety in a particular building. The SSD method, its theoretical framework and how it can be applied within the context of fire safety engineering are then taken up. The paper continues with a presentation of two examples of how SSD can be used to analyze decision problems concerned with fire safety, each involving a choice between different fire protection alternatives for a particular building. The first example concerns in a basic way the use of SSD, whereas, the second example aims at clarifying differences between the use of SSD and of more traditional decision analysis methods in this context. The paper concludes with a general discussion of decision analysis involving a high degree of epistemic uncertainty.

DECISION ANALYSIS CONCERNED WITH INVESTMENTS IN FIRE SAFETY

In decision analysis one distinguishes between decision making under *risk* and decision making under *uncertainty*. Decision making under risk is characterized by the decision maker's knowing the probabilities of the outcomes of the various decision alternatives exactly, whereas, decision making under uncertainty involves the decision maker's having no information at all about the probabilities of the different possible outcomes. Thus, in terms of epistemic uncertainty, decisions under risk involve no epistemic uncertainty, whereas, decisions under uncertainty involve the maximum epistemic uncertainty possible.

The most common decision criterion in making decisions under risk is the principle of maximizing expected utility (MEU), which has been applied to fire safety problems earlier (see [3,4], for example). In making decisions under uncertainty, there are a number of decision criteria one could choose between. Donegan [5] discusses four such criteria: the Laplace paradigm [6], the Wald paradigm [7], the Savage paradigm [8] and the Hurwicz paradigm [9]. Since it is assumed that the decision maker is completely ignorant with respect to the probabilities of the various outcomes that are possible for the different decision alternatives, each of these criteria involve some form of valuation of the outcomes themselves. In the present context, the decision rule suggested by Laplace implies that the decision maker should choose the decision alternative that minimizes the expected loss, its being assumed that each outcome considered is equally likely to occur. The Wald paradigm involves choosing the alternative for which the loss in case the worst outcome occurs will be lowest (this rule is also called the Maximin rule). The decision rule Savage suggested involves choosing the alternative that would result in the lowest loss possible if the best outcome should occur (this rule is also called the Maximax rule). The Hurwicz decision rule is a combination of the Maximax and Maximin rules. In choosing between different fire protection alternatives for a specific building, none of these decision rules can be considered suitable, however, since they ignore any differences between the alternatives in terms of the probabilities of the different consequences. Since investing in fire protection aims in part at reducing the probability of a serious fire, the benefits of such an investment are not taken into account by any of the decision rules just referred to. An investment in a sprinkler system, for example, reduces the probability of a serious fire but does not reduce the negative effects of the worst possible *consequence*, namely, the complete destruction of the building. Thus, use of the Maximin rule, which simply focuses on the worst possible consequence, would never lead to the recommendation that one makes a fire-safety investment.

The problem of performing a decision analysis concerned with alternative designs for fire protection in a particular building is likely to lie somewhere between decision making under risk and decision making under uncertainty.

Supersoft Decision theory (SSD) was chosen for use in the present, fire-engineering context because of its readily being used in conjunction with a quantitative risk analysis (event trees are used in the paper), and also because its enabling one to compare in a clear way the results obtained with the results of a more conventional decision analysis. Both of these more conventional decision analysis methods are based on Bayesian decision theory, which involves use of the principle of maximizing expected utility as the decision rule. One of these two methods will be termed traditional decision analysis. It involves probabilities and utilities being assigned as precise values. Use of this method allows different decision alternatives to be compared on the basis of expected utilities, the alternative having the highest expected utility being the alternative deemed best. Thus, only *one* value, the expected utility, is used to compare the different decision alternatives. That value (V_T) can be calculated using Equation (1), in which n is the number of possible outcomes of choosing a specific decision alternative that have been identified, P_i is the probability of outcome i occurring, and U_i is the utility associated with the occurrence of the outcome in question.

$$V_T = \sum_{i=1}^{n} (P_i \cdot U_i) \tag{1}$$

The other method of more conventional character is termed extended decision analysis. It involves probabilities and utilities being expressed as probability distributions. In comparing different decision alternatives, it is the expected utilities that are compared, the alternative with the highest expected utility being deemed best. This is almost the same decision rule as that employed in traditional decision analysis, the difference being that in extended decision analysis calculating the expected utility requires taking account of the epistemic uncertainty regarding the probability and utility values. The value (V_E) employed in comparing one alternative with another by use of extended decision analysis can be calculated using Equation (2). There, $f_i(P_i)$ is the probability density function representing the epistemic uncertainty regarding the probability value P_i, and $g_i(U_i)$ is the probability density function representing the epistemic uncertainty regarding the utility value U_i.

$$V_E = \int \cdots \int_D \left(\sum_{i=1}^{n} (P_i \cdot U_i) \cdot f_1(P_1) \cdot \ldots \cdot f_n(P_n) \cdot g_1(U_1) \cdot \ldots \cdot g_n(U_n) \right) \times dP_1 \ldots dP_n dU_1 \ldots dU_n$$

$$D = \left\{ (P_1, \ldots, P_n, U_1, \ldots, U_n); \sum_{i=1}^{n} P_i = 1 \right\} \quad (2)$$

In addition, however, extended decision analysis also involves evaluation of the effect which the epistemic uncertainties have on the expected utility. The idea here is that, since probabilities and utilities are expressed as probability distributions that represent the degree of confidence one has in different values for these, one can also express one's degree of confidence in different expected utility values. Thus, one can relate the epistemic uncertainty pertaining to probabilities and utilities to the expected utility of a decision alternative, which can be expressed as a probability distribution. This probability distribution can be used then to compare different decision alternatives in terms of decision robustness. Robustness has to do with how likely it is that the decision alternative considered best will change if the estimates of the probabilities and the utility values should change. Decision robustness is one of the key topics in this paper and will be taken up shortly.

Note that extended decision analysis can be viewed as being in many respects equivalent to what is termed Bayesian analysis (see [10], for example). In Bayesian analysis, epistemic uncertainties are represented by probability distributions and decision alternatives are evaluated on the basis of the expected utility (see Equation (2)). Bayesian analysis differs from extended decision analysis in only the value of the expected utility being used in the evaluation of decision alternatives, whereas, in extended decision analysis the effect which epistemic uncertainties have on the expected utilities, and thus the robustness of the decision, also being taken into account.

All three methods considered above (traditional decision analysis, extended decision analysis and Supersoft Decision analysis) utilize evaluation of expected utilities in one way or another. The use in the present context of expected utility for the evaluation of decision alternatives seems reasonable in view of results that Malmnäs [11] has presented. Malmnäs concludes that any rule that is simpler than that of expected utility performs worse as an evaluator of uncertain decision alternatives than expected utility does. More advanced methods[1] cannot be expected to be substantially better than use of expected utilities, especially when decision situations are involved in which probabilities and utilities cannot be expressed precisely.

Having decided to use expected utilities in evaluating uncertain decision situations leads to the question of how the expected utilities of decision alternatives should be calculated. To do so, one needs to be able to estimate the probability of each possible outcome of an alternative, as well as, to assign a utility value to each of the outcomes in such a way that the utility value arrived at represents the decision maker's preference for the (uncertain) outcome in question. The first problem one encounters here is that of estimating the probabilities of such events as whether the employees in the building will succeed in extinguishing a fire, whether the fire department will succeed in extinguishing it, and the like. Since observations of events of this type are usually rare in any particular building, information regarding them is likely to be scarce, making estimates of the probabilities involved difficult to make. There are different approaches one can take in dealing with the problem of probabilities involved being based on such limited information. The three methods referred to above will be compared in this respect. The first method to consider is that of traditional decision analysis. It involves use of precise values for probabilities and utilities, and consequently of exact values for the expected utilities. In comparing two decision alternatives using this method, one compares two expected utilities, one for each alternative, the alternative with the highest value being regarded as best (see Example 1 in Figure 1). When dealing with problems of the present type, however, it is questionable whether expressing estimates of probabilities as precise values is suitable. This method is taken up in the paper nevertheless in order to compare it with the other methods. Note that it provides the decision maker no information at all on how any uncertainties regarding the probability or utility values affect the decision. The second method to be considered, that of extended decision analysis, involves expressing the uncertainty one has concerning the probability and utility values by use of probability distributions. In comparing two alternatives by use of this method, one compares two distributions of expected utilities, the distribution with the highest expected value representing the decision alternative that is best

[1]The decision rules suggested by Hagen [19], Fishburn [20], Loomes and Sugden [21], Green and Jullien [22], Quiggin [23], and Yaari [24] are evaluated in [11].

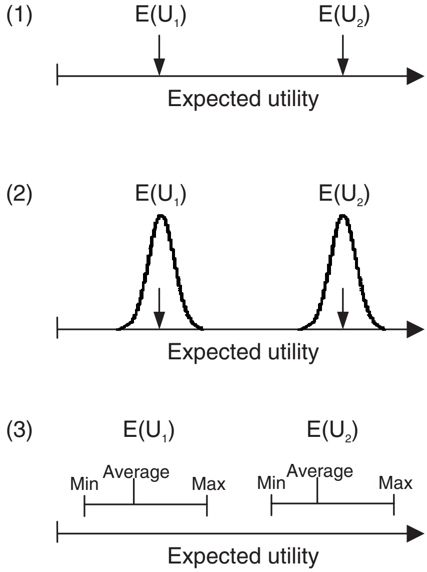

Figure 1. *Illustration of the results of each of three different decision analysis methods provides in comparing two decision alternatives. $E(U_1)$ is the expected utility of alternative 1 and $E(U_2)$ the expected utility of alternative 2: (1) Traditional decision analysis. Each expected utility is expressed as a single value; (2) Extended decision analysis. A given expected utility is expressed as a distribution; (3) Supersoft Decision Theory the maximum, minimum, and average value of an expected utility is used in the evaluation.*

(see Example 2 in Figure 1). Although the decision rule employed for determining the best alternative involves use of only the expected value of each of the two distributions, the form and position of the distributions provides the decision maker information regarding the robustness of the decision, which the first method does not. This method is described in greater detail in [4]. Note that in analysing different investments for a specific building, it may not be possible to use specific probability distributions to represent probability and utility values due to the lack of information. Therefore, one may need to employ a method that can deal with a high degree of epistemic uncertainty regarding the probabilities and utilities. The third method, which is called Supersoft Decision theory (SSD) [12] and is described in greater detail in the next section, is a method able to do this. In comparing alternatives by use of SSD, one compares the maximum and the minimum values of the expected utilities. Besides these two values, one should also take account of a value termed the Average, as will be described in the next section. The maximum and the minimum value form an interval in which the value of the expected utility lies (see Example 3 in Figure 1).

Later in the paper, the three methods described above will be applied to the same decision problem, the results of the analyses of that decision problem providing an additional illustration of the difference between them.

SUPERSOFT DECISION THEORY

Supersoft Decision theory (SSD) [12] allows the decision maker to utilise vague assessments of the values of the probabilities and consequences of interest. "The probability must be somewhere between 0.2 and 0.8" and "The consequence c_1 is at least twice as good as the consequence c_2" are examples of such vague assessments. Vague expressions of this sort are interpreted as inequalities. Thus, the representation of the probability just referred to could be $0.2 < p < 0.8$. Even when utilising such imprecise statements, one can still make use of the same basic model for how a fire in a building can be expected to develop as one does in performing a quantitative risk assessment. The event tree technique, which is useful for modelling possible fire scenarios in a building, will be used here for exemplifying how SSD can be employed for evaluating different fire protection alternatives.

In evaluating a decision situation in terms of SSD, one needs to create a representation of it in terms of a *decision frame*. This consists of the following: the different alternatives that can be chosen ($a_1, \ldots, a_n$), a list of the possible consequences C_i for each alternative, a list of utility statements U_i that pertain to these consequences, and a list of conditional probability statements P_i. The items on the list of consequences could be of the type "Areas 1 and 2 are completely destroyed" and those on the list of utilities could be of the type "The utility of Consequence 1 is at least 20 times as high as the utility of Consequence 2". The items on the list of probabilities can be statements of the type "The probability of event 2 is highly likely given event 1". The event trees (T_1, T_2, etc.), which indicate how the uncertain events are connected with the consequences, represent the last component of the decision frame. Thus, the decision frame can be summarized as consisting of (a_i, C_i, P_i, U_i, T_i). In practice, one should start by clarifying which alternatives are possible to choose between and then to identify for each of the alternatives the various events that can influence the outcome of the decision. The relationship between the occurrence of these events and the consequences should then be described (this can be done by using event trees) and the probability statements for the different events be formulated.

To evaluate the different alternatives, so as to identify which one is best, the qualitative statements of the decision frame need to be transformed into quantitative ones. A qualitative statement of the type "Event 1 is highly likely, given Event 2" can be translated, for example, into a quantitative statement of the type "$0.85 \leq P(E_1 \mid E_2) \leq 0.95$". Note that SSD does not prescribe any rules for how qualitative statements regarding probabilities are to be translated into intervals. Considering the empirical evidence (see [13], for example) indicating a great between-subject variability in the probability values assigned to verbal statements, use of such fixed transformation rules is probably not a very good idea. Using the verbal statements here, only as points of departure in determining the intervals of the probabilities, can be suggested instead. In doing this, the analyst and the decision maker needs to work together to find suitable intervals to represent the verbal statements made. Note that in determining the intervals of the probabilities, the decision maker is asked to *exclude* probability and utility values he/she considers too unlikely to be worth considered. This makes SSD different from other types of probability estimation techniques, in which either a single probability value (examples of such methods are provided in [14]) or a single interval containing the most likely values is to be estimated. The analyst might ask a decision maker who states "Event 1 is highly likely, given Event 2" whether it would be possible to exclude probability values of less than 0.05 for this conditional event. If this seems reasonable to the decision maker, one can continue and ask him/her whether it would be reasonable to exclude values of less than 0.1 and so on. In the end, an interval is established, that can be used to represent as adequately as possible the probability in question. For a more comprehensive discussion of this matter, see [12,16].

Examples of the types of statements that can be employed and of the respective inequalities are shown in Table 1.

Table 1. Applicable statements and their corresponding mathematical representations.

Statement	**Representation**
The probability of event E is equal to the probability of event F.	$P(E) = P(F)$
The probability of event E is less than x.	$P(E) < x$
The probability of event E is greater than x.	$P(E) > x$
The probability of event E lies between x and y.	$y < P(E) < x$
The probability of event E is at least i times as probable as event F.	$i * P(E) > P(F)$
The utility of consequence c_1 is higher than that of consequence c_2.	$U(c_1) > U(c_2)$
The utility of consequence c_1 is at least i times as high as that of consequence c_2.	$U(c_1) > i * U(c_2)$
The utility of consequence c_1 is equal to that of consequence c_2.	$U(c_1) = U(c_2)$

EVALUATION OF ALTERNATIVES IN SUPERSOFT DECISION THEORY

In employing SSD to evaluate the different decision alternatives, use is made of their expected utilities. However, since the decision frame contains statements in which the probabilities and the utilities are not assigned precise values or single probability distributions, the decision criterion of maximizing expected utility cannot be employed directly. Instead, the evaluation of alternatives in SSD is based on three criteria presented in Equations (3)–(5). $E(U,P)$ is the expected utility of the alternative in question (which is a function of the probabilities, P, and the utilities of the consequences, U).

$$\mathrm{Min}\,(E(U,P)) = \mathrm{Min}_{P,\,U}\left(\sum_{i=1}^{n}(P_i \cdot U_i)\right) \tag{3}$$

$$\mathrm{Max}\,(E(U,P)) = \mathrm{Max}_{P,\,U}\left(\sum_{i=1}^{n}(P_i \cdot U_i)\right) \tag{4}$$

$$\mathrm{Average}\,(E(U,P)) = \frac{\int \ldots \int \left(\sum_{i=1}^{n}(P_i \cdot U_i)\right) dP_1 \ldots dP_n\, dU_1 \ldots dU_n}{\int \ldots \int 1\, dP_1 \ldots dP_n\, dU_1 \ldots dU_n} \tag{5}$$

The $\mathrm{Min}\,(E(U,P))$ and $\mathrm{Max}\,(E(U,P))$ criteria are the lowest and highest expected utility values that satisfy the decision frame. In this case, satisfying the decision frame means that a solution to the inequalities is found within the decision frame. For example, assume that in the decision frame it is stated that the probability of a particular sprinkler system extinguishing a fire is somewhere between 0.8 and 0.9 ($0.8 \leq P\,(Sprinkler) \leq 0.9$), and also that this probability is at least 2 times as great as the probability that the employees will extinguish the fire ($2 * P\,(Employee) \leq P\,(Sprinkler)$). If the probability that the employees will extinguish the fire were 0.6, there would be no solution within the decision frame, since this implies that the probability of the sprinkler system's extinguishing the fire would be higher than 1, which is impossible. However, if the probability that the employees will extinguish the fire were 0.42, then the decision frame would be satisfiable, since this implies that the probability that the sprinkler system will extinguish the fire is greater than 0.84, which satisfies the inequality

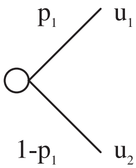

O: $p_1 \in [0.85, 0.95]$ $u_1 \in [0.9, 1]$ $u_2 \in [0, 0.5]$

Figure 2. *Illustration of an uncertain situation.*

$0.8 \leq P\,(Sprinkler) \leq 0.9$. Throughout the paper, Min will be used to denote Min $(E(U,P))$ and Max to denote Min $(E(U,P))$.

For simple problems, the calculation of Min and Max is not very complicated. One such simple problem is illustrated in Figure 2. The problem involves there being two possible outcomes if a particular decision alternative is chosen, one of them with a utility value of between 0.9 and 1 (u_1) and the other with a utility value between 0 and 0.5 (u_2). The best consequence (u_1) occurs with the probability p_1, which is judged to be somewhere between 0.85 and 0.95, and the other consequence with the probability $1-p_1$.

In order to evaluate the decision alternative shown in Figure 2, one needs to analyse the expression for the expected utility ($E(U,P)$), as given in Equation (6). The *Min*-value of Equation (6), given the constraints (O, in Figure 2), is found by setting p_1, u_1, and u_2 to their lowest values. The resulting value is 0.765. The Max-value is calculated using the same procedure but setting the parameters (p_1, u_1, and u_2) at their highest values. This yields a value of 0.975.

$$E(U,P) = p_1 \cdot u_1 + (1-p_1) \cdot u_2 \tag{6}$$

Although calculating the Max and Min-values may be a simple task in the case of such limited problems as that illustrated in Figure 2, as the scope of the problem increases the expression for the expected utility becomes more complicated, making the calculations much more complex. The complexity is due to the fact that the problem of calculating Min and Max is a nonlinear multivariable optimisation problem with a set of inequalities as constraints. $E(U,P)$ is the objective function that one seeks to minimize or maximize, the inequalities found in the decision frame represents the constraints. Such a problem, except for one of the simplest type, is difficult to solve by hand, but there are computer programs that can solve them[2].

The Max and Min-criteria alone are not ideal for comparing decision alternatives, since they are very sensitive to changes in values near the edges of the decision frame, these being the values the decision maker is most

[2]In order to solve the first example that is found later in the present paper, the "fmincon" function in the Optimization Toolbox for MATLAB [17] was used.

likely to be uncertain about. It is useful, therefore, to also employ the Average $(E(U,P))$-criterion shown in Equation (5). The Average $(E(U,P))$-criterion can be seen as the expected value of $E(U,P)$ when the probability and utility values are treated as being uniform distributions that extend between their maximum and minimum values. An example of an evaluation using the Average $(E(U,P))$-criterion is given in Equation (7), where calculation of the Average $(E(U,P))$-value of the decision alternative shown in Figure 2 is presented. Throughout the paper, Average will be used to denote Average $(E(U,P))$.

$$\text{Average}\,(E(U,P)) = \frac{\iiint\limits_{O} p_1 \cdot u_1 + (1-p_1)\cdot u_2 dp_1 du_1 du_2}{\iiint\limits_{O} 1 dp_1 du_1 du_2} = \frac{0.0044}{0.005} = 0.88 \quad (7)$$

In calculating the Average-value for decision alternatives involving only parameters (probabilities and utilities) that are independent of each other, which was the case in the example presented above, one can utilize the fact that the expected value of the sum of two stochastic variables is equal to the sum of the expected value of each, as well as that the expected value of the product of two *independent* stochastic variables is equal to the product of the expected value of each of the variables. Using this method, the Average-value that was calculated in Equation (7) can be calculated instead using Equation (8), without any integrals needing to be solved. This simpler form of calculation becomes increasingly useful as the scope of the problem increases and the integrals in Equation (5) becomes more cumbersome to solve. Note that when analyzing uncertain situations represented by event trees having chance nodes with three or more branches, the probabilities of the different branches are not independent and therefore the problem becomes more difficult to solve. It is possible, however, to evaluate even this type of problem using SSD, a computer algorithm has been developed so as to help the decision maker in doing this [15].

$$\begin{aligned}\text{Average}\,(E(U,P)) &= E(p_1)\cdot E(u_1) + (1-E(p_1))\cdot E(u_2)\\ &= 0.9\cdot 0.95 + 0.1\cdot 0.25 = 0.88 \qquad (8)\end{aligned}$$

Although the Average-criterion is also sensitive to changes near the edges of the decision frame, this criterion, if used in isolation, is nevertheless a natural candidate for use as the general decision criterion. It can be reasonable, however, to employ a set of criteria by combining all three criteria. Thus, if one wanted to evaluate the decision alternative which the situation shown in Figure 2 represents, one would use the three values Min $(E(U,P)) = 0.765$, Max $(E(U,P)) = 0.975$ and Average $(E(U,P)) = 0.88$

and compare these with the comparable values obtained in analyzing whatever other decision alternatives are involved.

In the original account of SSD [12], the quantitative evaluations which the paper takes up (Equations (3)–(5)) are conceived as being employed in combination with qualitative methods. Also, the method described in [12] does not prescribe the preferred alternative needing to be best in terms of *all* of the criteria presented above, although in the present paper we do treat an alternative as being best only if it is the best in terms of all three criteria.

ANALYZING A SMALL DECISION PROBLEM USING SSD

In order to exemplify the use of SSD within the context of fire risk management, a simple hypothetical example will be used. The aim of this example is to show how a decision problem can be analyzed using SSD.

The hypothetical example concerns the decision of whether a particular investment in fire safety should be made, one rather modest in cost. Assume that the risk manager of a company has found there to be a room containing electrical equipment in which a fire might readily start, a fire that could be very severe in its effects. The risk manager wants to determine whether the decision to invest in a CO_2 system to be installed there would be a good one. Since both the occurrence of fire in that room and the spread of fire from it if a fire should occur are very uncertain, the risk manager decides to use SSD to evaluate the different alternatives.

Denote the expected utility of the alternative of investing in a CO_2 system as $E_1(U,P)$ and the expected utility of the alternative of not investing as $E_2(U,P)$. Although the analysis of the decision here will only involve use of the Min and Max criteria, $\text{Min}(E_1(U,P)\text{-}E_2(U,P))$ and $\text{Max}(E_1(U,P)\text{-}E_2(U,P))$ will be examined, rather than $\text{Min}(E_1(U,P))$ being compared with $\text{Min}(E_2(U,P))$ and $\text{Max}(E_1(U,P))$ with $\text{Max}(E_2(U,P))$. This allows account to be taken of the fact that the probability of a severe fire occurring in the room and the probability that such a fire, if it does occur, will be contained there (within the room of origin) are estimated to be the same for both alternatives. If both $\text{Min}(E_1(U,P)\text{-}E_2(U,P))$ and $\text{Max}(E_1(U,P)\text{-}E_2(U,P))$ give positive values, it can be concluded that alternative 1 is best, whereas if both evaluations yield negative values, it can be concluded that alternative 2 is best. Note that the approach of evaluating the *difference* in expected utility between two alternatives resembles that of the Delta-method [16].

The first thing to do is to set up the decision frame. Two alternatives have been identified, one involving the company's investing in the CO_2 system (a_1) and the other the company's not investing in it (a_2). Assume, so as to simplify the problem, that a fire in the room can only have three possible consequences: (1) its being too small to have any appreciable impact on the company (consequences $c_{1,1}$, $c_{1,2}$, and $c_{2,1}$ in the event trees shown in Figures 3 and 4), (2) the fire destroys everything in the room of origin but is contained in that room (consequences $c_{1,3}$ and $c_{2,2}$ in the event trees shown in the two figures), and (3) the fire is spreading from the room of origin and destroying the entire factory (consequences $c_{1,4}$ and $c_{2,3}$ in the event trees shown in the same two figures).

In order to continue with the analysis, the decision maker needs to make a statement regarding the probability that a fire with the potential to become severe will occur in the room of interest within the period planned for, which is assumed to be ten years. The risk manager estimates this probability as being lower than 0.2 but not lower than 0.05. For simplicity, assume that during the period planned for only one severe fire can occur. This probability and the remaining statements regarding probabilities are summarized in Table 2.

To arrive at statements regarding the utility values, it can be helpful to begin by visualising the relative positions of the various consequences on a utility scale. Figure 5 shows the utilities of the different consequences ($u_{1,1}, \ldots, u_{2,3}$). Note that the distances between the utilities of the different consequences, as shown in the figure, are not correct, only their relative

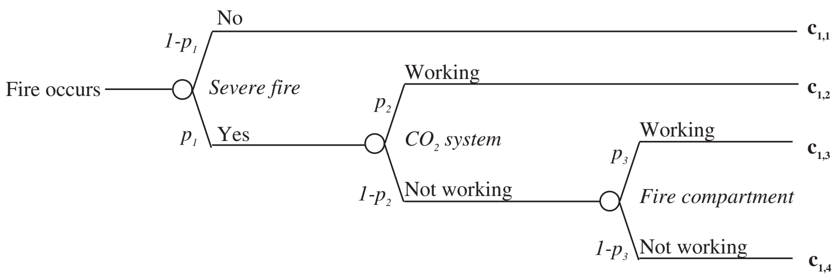

Figure 3. *Event tree representing the possible consequences of a fire in the electrical equipment room, given that alternative 1 is chosen.*

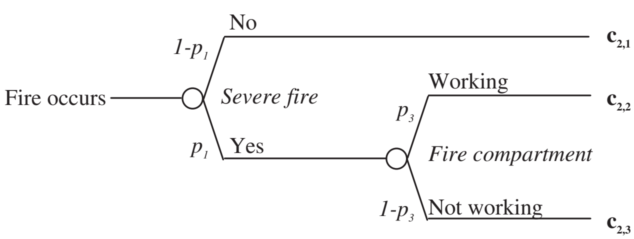

Figure 4. *Event tree representing the possible consequences of a fire in the electrical equipment room, given that alternative 2 is chosen.*

Table 2. The probability statements and their representation in the form of inequalities.

Statement	Representation
The probability of a potentially severe fire within the next 10 years is between 0.05 and 0.2.	$0.05 \leq p_1 \leq 0.2$
The probability that the CO_2 system will be working and will extinguish the fire is between 0.7 and 0.95.	$0.7 \leq p_2 \leq 0.95$
The probability that the fire will be contained in the room of origin given that the fire has not been extinguished, is between 0.5 and 0.95.	$0.5 \leq p_3 \leq 0.95$

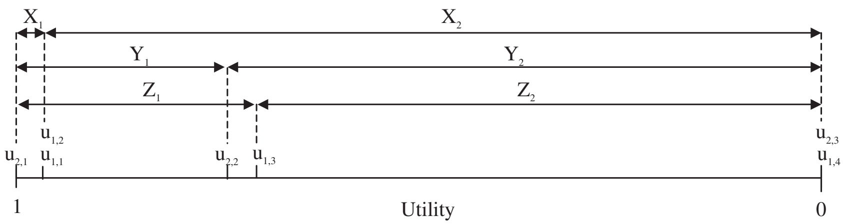

Figure 5. *Diagram of the utilities of the different consequences. Note that the value distances as shown are not correct, only the relative positions of the consequences on the utility scale being correct.*

positions being correct. The utility statements and the inequality representations of the utility values are shown in Table 3. Note that the reason for consequence $c_{1,1}$ and $c_{1,2}$ not being as good as $c_{2,1}$, despite none of them involving any serious fire occurring, is that $c_{1,1}$ and $c_{1,2}$ involve the company's having invested in a CO_2 system, whereas $c_{2,1}$ does not. The same reasoning applies to the difference found between $c_{1,3}$ and $c_{2,2}$. However, $c_{1,4}$ and $c_{2,3}$ can be judged to be equally bad due to the costs of the CO_2 system being small compared with the costs of a total loss of the factory.

In calculating the expected utilities of the two alternatives, the event trees shown in Figures 3 and 4 are employed. These result in Equations (9) and (10), which represent expressions for the expected utilities of the two alternatives.

$$E_1(U,P) = (1-p_1)\cdot u_{1,1} + p_1 \cdot p_2 \cdot u_{1,2} + p_1 \cdot (1-p_2)\cdot p_3 \cdot u_{1,3} + p_1 \cdot (1-p_2)\cdot(1-p_3)\cdot u_{1,4} \quad (9)$$

$$E_2(U,P) = (1-p_1)\cdot u_{2,1} + p_1 \cdot p_3 \cdot u_{2,2} + p_1 \cdot (1-p_3)\cdot u_{2,3} \quad (10)$$

One can calculate the difference in expected utility between the two alternatives by use of Equation (11). Several of the expressions appearing in

Table 3. The utility statements and their representation in the form of inequalities and equalities.

Statement	Representation
Consequence $c_{2,1}$ is the best consequence.	$u_{2,1} = 1$
Consequence $c_{1,4}$ is the worst consequence.	$u_{1,4} = 0$
Consequence $c_{1,1}$ and consequence $c_{1,2}$ are equally good.	$u_{1,1} = u_{1,2}$
Consequence $c_{1,4}$ and consequence $c_{2,3}$ are equally bad.	$u_{1,4} = u_{2,3}$
The utility difference between $c_{1,1}$ and $c_{1,4}$ (X_2) is at least 1000 times as great, and not more than 10,000 times as great, as the difference between $c_{2,1}$ and $c_{1,1}$ (X_1).	$0.9999 \geq u_{1,1} \geq 0.999$
The utility difference between $c_{2,2}$ and $c_{1,4}$ (Y_2) is at least 100 times as great, and not more than 1000 times as great, as the difference between $c_{2,1}$ and $c_{2,2}$ (Y_1).	$0.999 \geq u_{2,2} \geq 0.99$
The utility difference between consequence $c_{1,3}$ and consequence $c_{2,2}$ is equal to the difference between consequence $c_{1,1}$ and $c_{2,1}$ or less, and the utility of consequence $c_{2,2}$ is equal to or better than that of consequence $c_{1,3}$.	$u_{2,2} \geq u_{1,3} \geq (u_{2,2} - (u_{2,1} - u_{1,1}))$

the column labeled "Representation" in Table 3 have been used in arriving at Equation (11).

$$E_1(U,P) - E_2(U,P) = u_{1,1} - p_1 \cdot p_2 \cdot p_3 \cdot u_{1,3} + p_1 \cdot p_2 \cdot u_{1,1} + p_1 \cdot p_3 \cdot (u_{1,3} - u_{2,2} + u_{1,4}) + p_1 \cdot (1 - u_{1,1} - u_{1,4}) - 1 \quad (11)$$

Equation (11) represents the difference in expected utility between Alternatives 1 and 2. This is the equation one seeks to calculate the Min and Max Equation (11) can be regarded as the objective function in a nonlinear multivariable optimisation problem. The constraints of the problem are given by the decision frame. The constraints are presented in the "Representation" columns in Tables 2 and 3.

Although the present optimization problem is relatively simple to solve, and can readily be solved using only hand calculations, a function called "fmincon" in the Optimization Toolbox for MATLAB [17] was used to solve it. The results obtained were that Min($E_1(U,P)$-$E_2(U,P)$) is 7.84×10^{-4} and Max($E_1(U,P)$-$E_2(U,P)$) is 9.59×10^{-2}. Since both the Min and the Max-evaluations result in positive values, one can conclude that the best alternative is to invest in a CO_2 system (Alternative 1) without performing an evaluation of the Average-value (see Case 1 in Figure 9). Note that, despite one's not knowing the exact value of the expected utility of either Alternative 1 or 2, it can be shown that the expected utility of Alternative 1 is higher than that of Alternative 2. Thus, the decision to invest in a CO_2 system is robust. This implies that the decision alternative is the best, regardless of which values of the uncertain variables are "correct" (assuming the values to be contained within the decision frame). This is an important principle, one that will be discussed in detail later in the paper.

COMPARING SSD WITH BAYESIAN DECISION ANALYSIS: AN EXAMPLE OF A REAL-WORLD EXAMPLE

The aims of the example that follows are to show how SSD can be applied to a real-world decision problem and to compare the results of the SSD analysis with those of two other decision analysis methods. The two methods used for purposes of comparison are those referred to in a previous section as traditional decision analysis and extended decision analysis.

The decision problem here concerned the question of whether an investment in a water sprinkler system for a factory should be made. Since the production in the factory, which belonged to the firm ABB, was very important for the company, a serious fire in the building would have had extremely negative consequences. The decision problem was analyzed earlier by use of extended decision analysis, an analysis described in greater detail [3,4]. The alternatives the decision maker (the company) could choose between were to invest in a water sprinkler system (a_1) and to not invest in it (a_2). Evaluation of the alternatives was performed by use of an event tree technique aimed at modeling different fire scenarios that were possible. Since the building to which the analysis referred was large ($55{,}000\,m^2$), an extensive decision analysis was required. In order to simplify the presentation of the problem and the comparison of the results, it was decided here to only carry out a comparative analysis for one of the fire compartments in the building (that was approximately $5{,}500\,m^2$ in size). This fire compartment was treated as if it were a separate building, one for which the decision maker was to decide whether to invest in a water sprinkler system.

In the original analysis, the consequences were expressed in terms of monetary losses and, since the monetary losses associated with any given fire scenario were uncertain, they were expressed by use of triangular probability distributions. The same approach was employed for the probabilities used in the model. The minimum, the most likely, and the maximum values for the probability distributions are presented in Tables 4 and 5. These probability

Table 4. The minimum, the most likely and the maximum value of the different probabilities used in the model. All the probabilities are conditional upon the event that a fire has occurred in the building.

Probability	Abbreviation	Minimum	Most Likely	Maximum
The probability that the growth potential of the fire will be low.	P(Low)	0.55	0.80	0.85
The probability that the fire detection system will detect the fire.	P(Alarm)	0.90	0.98	0.99
The probability that the employees will succeed in extinguishing the fire given that the fire detection system has detected it.	P(Emp./Alarm)	0.70	0.80	0.95
The probability that the employees will succeed in extinguishing the fire given that the fire detection system has *not* detected it.	P(Emp./No Alarm)	0.20	0.50	0.60
The probability that the fire department will succeed in extinguishing the fire before it destroys the fire compartment, given that the fire detection system has detected the fire.	P(Fire dept./Alarm)	0.50	0.70	0.85
The probability that the fire department will succeed in extinguishing the fire before it destroys the fire compartment, given that the fire detection system has *not* detected the fire.	P(Fire dept./No Alarm)	0.20	0.30	0.50
The probability that the water sprinkler system will extinguish the fire.	P(Sprinkler)	0.90	0.95	0.98

distributions were arrived at in discussions between the analyst and personnel both from ABB and from the fire department. The monetary sums given in this section, originally in Swedish crowns (SEK), were converted to US dollars at the rate of $1 to 10 SEK.

The event tree used to describe each of the fire scenarios considered is presented in Figure 6 (note that in the alternative of there being no sprinkler system in the building the probability that the water sprinkler system will succeed in extinguishing the fire is considered to be 0). Using this event tree in combination with the different estimates of the probabilities and the consequences allows one to calculate the expected utility of investing in a water sprinkler system. Assume exactly one fire to be the number of fires

Table 5. The minimum, the most likely and the maximum value of the different monetary consequences used in the model.

		Monetary Loss Associated with a Particular Fire Scenario ($ Thousand)		
Fire Scenario	**Abbreviation**	**Minimum**	**Most Likely**	**Maximum**
The fire will be extinguished by the employees or by the sprinkler system.	C_1	5	10	15
A fire with low growth potential will be extinguished by the fire department.	C_2	25	50	75
A fire with high growth potential will be extinguished by the fire department.	C_3	533	1,067	1,933
The fire will completely destroy the fire compartment.	C_4	25,920	32,000	38,720

expected to occur during the time for which the decision maker wishes to take account of the benefits the sprinkler system would provide. Certain assumptions are also made concerning the decision maker's preferences regarding the occurrence of more than one fire during a given period of time. In particular, it is assumed that the expected utility of k fires occurring during a given period, each fire having the expected utility of $E(u)$, is $kE(u)$. This implies that the utility of any given fire is not affected by how many other fires occurred during the period in question. The assumptions just mentioned allow the expected utility for each of the two decision alternatives to be calculated by multiplying the expected utility of one fire by the number of fires expected to occur during that time period. In the present case, 1 fire is expected to occur.

Note that one could discount losses occurring in the future in the same way as is done in capital investment analysis. Since the focus here, however, is on the decision rules used in the different methods and on how to deal with epistemic uncertainty, no attempt is made to discount future losses. For further information on this matter, see [18].

For simplicity, assume that the decision maker's preferences with respect to uncertain monetary outcomes can be described by a risk-neutral utility function, which means that he/she evaluates uncertain monetary outcomes exactly according to their monetary values. Since in the present context only potential losses are being analyzed, it might be regarded as being more appropriate to use the term disutility in discussing the decision maker's preferences with respect to different losses. Nevertheless, the term utility will

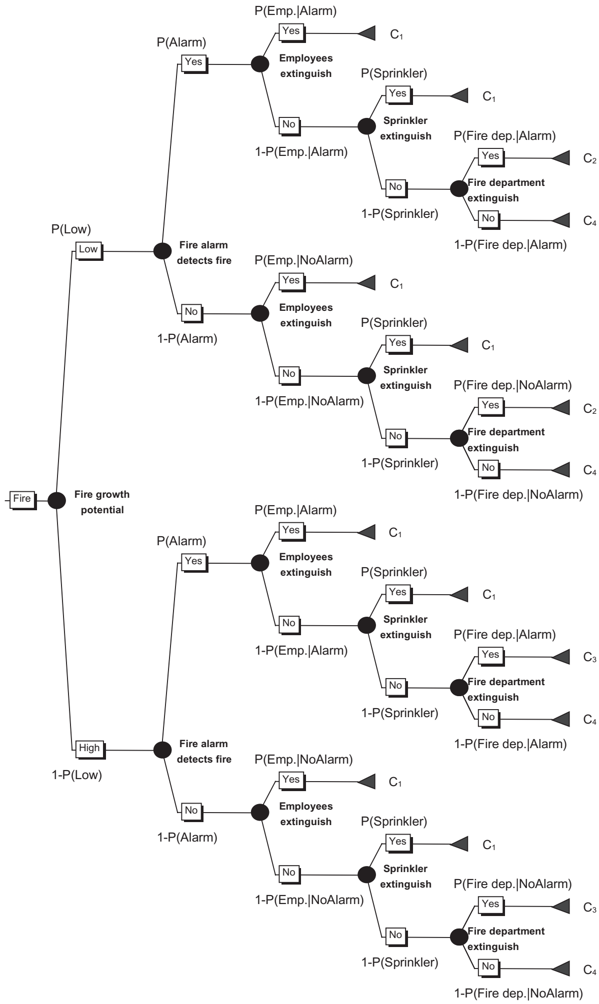

Figure 6. *Event tree showing the different fire scenarios considered in the analysis of whether to invest in a sprinkler system.*

be used here, the worst consequence being assigned a utility value of −1 and the best consequence a utility value of 0. In assigning utility values to the different consequences, we use the monetary outcomes reported in Table 5, assigning a utility value of −1 to the consequence involving a loss of $38,820,000 (loss of the entire factory, including the sprinkler system) and a utility value of 0 to a loss of $0. Since we assume that the decision maker is risk neutral, each of the other monetary outcomes can easily be translated into utility values of between 0 and −1.

In addition, we assume that the sprinkler system costs \$100,000, which is a reasonable assumption in view of the fact that in the original analysis the total cost of the sprinkler system for the building as a whole was approximately \$1,000,000. This \$100,000 cost is taken account of when one calculates the expected utility of the alternative of investing in a sprinkler system.

One aim of this example is to clarify the relationship between decision analyses performed using Bayesian decision theory, as described in [3,4], and one performed using Supersoft Decision theory. In one of the analyses using Bayesian decision theory (the extended decision analysis method), two Monte Carlo-simulations were carried out so as to be able to relate the uncertainty contained in the probability and consequence estimates to the expected utility of the different alternatives. In each of the two Monte Carlo-simulations, 5000 iterations were performed. The distribution representing the expected utility of a particular alternative can then be compared with the results generated by analyzing the alternative by use of SSD. The results of the Monte Carlo-simulations, expressed in terms of expected utility, and of the analysis using SSD are presented in Figures 7 and 8. The vertical lines are the results of the SSD evaluation using the Max and Min values given in Tables 4 and 5. The line associated with the highest utility is the result of the Max-evaluation (see Equation (4)), the line associated with the lowest expected utility being the result of the Min-evaluation (see Equation (3)). The line in the middle is the result of the Average-evaluation. The Average-evaluation was performed using the expected values of the probabilities and the utilities, assuming there to be a uniform distribution between the highest and lowest values, and utilizing the same technique as illustrated by Equation (8).

The results of the SSD-evaluation are presented in Table 6. A decision analysis using exact values for the probabilities and utilities (traditional

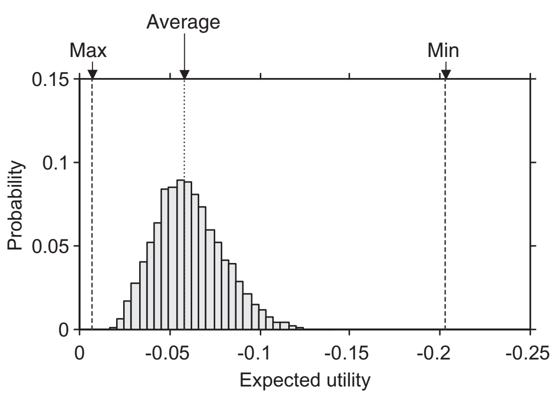

***Figure 7.** Results of the analysis of the alternative of keeping the building in its present condition. Note that the horizontal scale is not the same as the one used in Figure 8.*

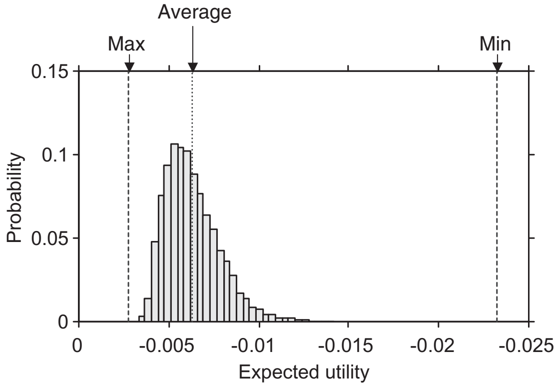

Figure 8. *Results of the analysis of the alternative of investing in a water sprinkler system for the fire compartment in question. Note that the horizontal scale is not the same as the one used in Figure 7.*

Table 6. Results of the SSD-evaluation.

	Expected Utility		
Alternative	**Min**	**Average**	**Max**
Sprinkler	−0.0232	−0.0067	−0.0028
No sprinkler	−0.2023	−0.0641	−0.0065

Table 7. Results of the evaluation using exact probability and utility values.

Alternative	**Expected Utility**
Sprinkler	−0.00559
No sprinkler	−0.05536

decision analysis) was also performed. The values for the probabilities and the consequences as given in the "Most likely" column in Tables 4 and 5 were used to calculate the expected utilities of the two decision alternatives. The results are presented in Table 7.

Comparing the results of evaluating the two decision alternatives allows one to conclude that, in terms of an extended decision analysis, the alternative of investing in a sprinkler system is best, since the mean value of the distribution shown in Figure 8, which represents the alternative of investing in the sprinkler system, is higher than the mean value of the distribution shown in Figure 7, which represents the alternative of not investing in the sprinkler system. The results SSD provides allow the same conclusion to be drawn, the alternative of investing in the sprinkler system being best, since the Min-evaluation gives a higher expected value for the sprinkler alternative, as well as for the Max-evaluation and the Average-evaluation. The results obtained using traditional decision analysis (see Table 7) imply the same conclusion to be reached there as in the other two analyses, namely that the decision alternative of investing in a sprinkler system is best.

DECISION ROBUSTNESS

There are differences between the three approaches to evaluation in terms of decision robustness. Here, decision robustness is used to denote how likely it is that the best alternative would change if a reasonable degree of change were to be made in the estimation of either the probabilities or the utilities. Since the probabilities and utilities are expressed as exact values when traditional decision analysis is employed, that method provides no information concerning robustness.

One way of evaluating the robustness of a decision is to compare the resulting distributions of the expected values for all of the alternatives and to note whether these distributions overlap to an appreciable extent. Looking at the results of the extended decision analysis of the ABB-case, it is clear that the decision to invest in the sprinkler system can be considered robust, since the expected utility distributions for the two alternatives under consideration do not overlap. Although an SSD evaluation does not result in a probability distribution, one can investigate whether the intervals defined by Min and Max overlap. In the ABB-case, one finds that the intervals for the two alternatives do overlap. The fact that the two approaches to evaluation differ in robustness (i.e. in the extent to which the intervals or distributions overlap) is due to the SSD approach being conservative in terms of its manner of dealing with robustness. More specifically, since the endpoints of the interval representing the possible values for the expected utility of an alternative obtained by use of SSD usually has little credibility in terms of extended decision analysis, they are not part of the distributional results presented in Figures 7 and 8. Comparing the results of the two approaches highlights the fact that SSD is coarser in its treatment of robustness than the extended decision analysis method.

The robustness of a decision is very important when a high degree of epistemic uncertainty is involved, since taking account of robustness can enable the decision maker to reach a definite conclusion there nevertheless regarding which alternative is best. When the robustness of a decision problem has been analyzed, such a conclusion can be drawn, provided one of the decision alternatives has the highest expected utility and its expected utility is clearly separated from those of the other alternatives (i.e. if there is no overlap between the distributions or intervals). Note that to conclude by use of this method that a particular alternative is the best requires (1) that one accepts expected utility as being the basis for evaluation and (2) that the decision frame (i.e. the basis for analysis) contains all the plausible values of both the probabilities and the utilities. How the decision situation is structured when the expected utilities are separated from each other can be illustrated by Case 1 in Figure 9, in which the results are presented in the manner typical for SSD.

If the two alternatives are not completely separated in terms of the plausible values of the expected utilities assigned to them, one needs an evaluation criterion appropriate for determining the best alternative under such conditions. Two such criteria have been presented in the paper, the maximization of expected utility (MEU) and the SSD criteria (Equations (2)–(5)). An extended decision analysis uses the MEU criterion but also utilizes the distributions of the expected utilities in determining the degree of robustness, i.e., how much the distributions of the expected utilities overlap. The situation in which the two distributions of expected utilities are not entirely separated and an appropriate decision rule for determining the best alternative under such conditions is needed and is illustrated by Case 2 in Figure 9. There, in terms of SSD, Alternative 1 is the best since each of the three criteria (Equations (3)–(5)) gives a higher value ($E(U_1)$) for Alternative 1 than for 2 ($E(U_2)$).

One can end up, however, with a situation in which the evaluation criteria taken up in the present paper cannot provide any recommendation of which alternative to select. This can happen when the possible values of the

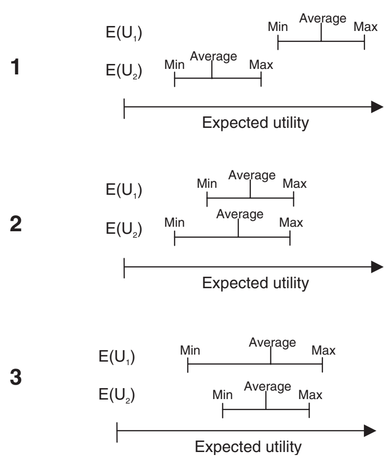

Figure 9. *Illustration of three different results of an evaluation of two decision alternatives using SSD.*

expected utilities for the different alternatives are too close together and none of the intervals of possible expected utility values are clearly higher in terms of expected utility values. This is illustrated by Case 3 in Figure 9, in which the results of an SSD evaluation of two decision alternatives are presented. Since the Average and Max values for Alternative 1 are higher than the corresponding values for Alternative 2, one might think that Alternative 1 is better. However, since the Min value of Alternative 1 is lower than that of Alternative 2, one cannot conclude that Alternative 1 is best. Instead, one needs to obtain more information regarding the problem, so as to hopefully reduce the uncertainties sufficiently to arrive at a clear conclusion. Note that traditional decision analysis utilising exact values of the probabilities and the utilities could not have distinguished between the three cases shown in Figure 9. That method would have indicated Alternative 1 being the best in all three cases.

DISCUSSION ON THE APPLICATION OF DECISION ANALYSIS TO A POSSIBLE FIRE PROTECTION INVESTMENT IN A SPECIFIC BUILDING

In the previous sections, three methods for decision analysis were compared in terms of their applicability to a problem involving a high degree of epistemic uncertainty. One can conclude that these methods differ substantially in how precise one needs to be when estimating the probabilities and evaluating the consequences. It is doubtful whether traditional decision analysis, which requires that the parameters involved be assigned exact values, has any practical usefulness in the present context, since in dealing with possible fire protection investments in a particular building one is not likely to have the amount of information needed to assign exact values. Using such a method in the present context could in fact be very misleading, since the results could give the impression that those decision alternatives which are *not* identified as being the best are *clearly* inferior to the alternative deemed best whereas in reality the obtaining of additional information might well lead to the results one arrives at changing easily. For a context such as the present one, therefore, a method which involves expressing the probabilities and utilities as being uncertain is more appropriate.

From a practical standpoint, methods such as SSD which involve interval statements seem attractive. In a practical decision situation there may be several stakeholders and thus several "decision makers". Under such conditions, it may be impossible for the stakeholders to agree upon a specific distribution for the probability and utility values to be employed. Instead, each stakeholder could assign the parameters of interest a maximum and a minimum value. One could then employ the lowest of the stakeholders' minimum values together with the highest of the stakeholders' maximum values in the analysis to be carried out. This would result in a decision frame that includes all the stakeholders' estimates. If the analysis resulted in a robust decision, all the stakeholders should be satisfied with the decision alternative that was recommended (provided they accepted expected utility as a reasonable means of evaluating the decision alternatives).

Another important aspect of decision analysis in the present context is that in many applications one is only interested in determining whether a particular alternative is better or worse than the other alternatives, not in *exactly* how much better or worse it is. This can be exemplified by the two examples included in the present paper in which the question was which of two alternatives one should choose. In the first example, it was shown that to answer this question one did not need to know exactly how much better or worse the alternatives were in comparison to each other. Instead, it was sufficient to show within the decision frame at hand, that for all plausible values of the evaluation criteria (in this case, expected utility) one alternative was better than the other. This was almost true, but not quite, in the second example, where an extended decision analysis showed there to be no overlap between the distributions of expected utility, but evaluation by SSD showed there to be some possible values of the probabilities and utilities for which the decision alternative to be recommended changed.

This brings up one point concerning differences between SSD and extended decision analysis. As was mentioned earlier, SSD is conservative with respect to robustness. This means that, even if SSD indicates there to be an overlap of the intervals of the expected utilities, the decision situation may very well be considered robust in terms of extended decision analysis since the endpoints indicated by an SSD analysis have so little credibility, the values involved being so unlikely that in practice the decision can be considered robust. It is clear, therefore, that in decision situations in which the information available regarding the probabilities, for example, is sufficient to justify expressing epistemic uncertainty by use of specific probability distributions, extended decision analysis is better to use than SSD. However, the salient argument in favor of SSD is that, since it requires no probability distributions for representing epistemic uncertainty, it can be used for decision analysis in cases in which the information regarding the probabilities, for example, is too vague to justify using specific probability distributions. In such cases, extended decision analysis cannot be employed. Thus, SSD and extended decision analysis should not be viewed as competing, but rather as complementary methods, the one being useful when precision in terms of robustness is sought and when the information at hand justifies epistemic uncertainties being expressed as specific probability distributions, the other being useful when the information at hand is vague and does not justify expressing uncertainties as specific probability distributions. Also, note that SSD and extended decision analysis will provide results that are very similar when use is made of the Average evaluation in connection with SSD (Equation (5)) and the expected utility evaluation in connection with extended decision analysis (Equation (2)), the results of the two being identical, in fact, if uniform distributions are used to represent the epistemic uncertainty involved in the case of extended decision analysis.

It could appear that SSD involves more complicated and time-consuming calculations in evaluating decision alternatives than extended decision analysis. Use of computers, however, can make the calculations SSD would require no more cumbersome than those involved in Monte Carlo-simulations, SSD calculations probably being faster, in fact, than Monte Carlo-simulations because of one's not having to define probability distributions in SSD.

In practice, the choice of a method for analyzing different fire protection alternatives for a specific building is not always an easy one, its depending very much on the situation at hand. One can conclude, however, that traditional decision analysis, if used in isolation, is clearly unsuitable for decision problems involving a high degree of epistemic uncertainty, due to its inability to provide or utilise information regarding the epistemic uncertainty of the results. This can be seen as applicable to Bayesian analysis as well. There, although probabilities and utilities are expressed as probability distributions, only one value of the expected utility is used when alternatives are compared. In many situations, extended decision analysis is probably useful, since it provides information regarding the robustness of a decision in a way that is readily grasped (yielding a distribution of expected utilities). The decision maker who finds it difficult to assign probability distributions to uncertain parameters or lacks the time to do so can use the SSD method instead.

In practice, a possible procedure in analysing a decision problem would be to start using a rough model involving use of SSD. If the results indicate that the decision problem to be of Types 1 or 2 as shown in Figure 9, one can readily conclude which alternative is best. One could then take the analysis one step further if one wished, using extended decision analysis (if the information justified its use) to analyze the robustness of the decision as adequately as possible. If, on the other hand, the results of the initial SSD evaluation indicate the decision problem to be of Type 3, one would need to collect more information about the problem at hand so as to either be able to conclude which alternative was best through use of SSD or use extended decision analysis for determining the robustness of the decision problem.

Note that the analysis just presented of the differences between the three different methods discussed also applies to situations in which determining which alternative is best is based not on expected utilities but on other measures of evaluation. If decision alternatives are to be screened, for example, by excluding from further analysis any alternatives for which the probability of an extreme event occurring is too high, one can make use of exact probabilities, probability distributions or SSD, the latter two approaches allowing one to analyse the robustness of the screening process.

CONCLUSIONS

Decision analysis of fire safety decisions applying to the choice of a possible fire protection investment was discussed, an investment applying not to buildings in general or some basic category of buildings but to a specific building, particular attention being directed at situations in which one can expect there to be a considerable degree of epistemic uncertainty. A new decision analysis method termed Supersoft Decision theory (SSD) was introduced and was applied to problems of fire safety. Two concrete applications of SSD in a fire safety context were described. The first case discussed was a hypothetical decision situation of limited scope aimed at illustrating some of the calculations to be made when evaluating alternatives by use of SSD. The second case involved part of a real-world decision problem that had been analyzed earlier by use of Bayesian decision theory. The SSD analysis of the second case was compared with two other types of decision analysis utilising Bayesian decision theory, one of them termed traditional decision analysis and the other termed extended decision analysis. The traditional decision analysis utilised exact values of the probabilities and utilities involved, whereas the extended decision analysis used probability distributions to represent the uncertainty regarding the probabilities and utilities.

It was concluded that a decision analysis involving use of precise values for probabilities and utilities can easily be misleading, since the results obtained would provide no indication of the robustness of the decision, i.e., of how readily the alternative judged to be best could change if a reasonable degree of change in the probability or utility values should occur.

It was also argued that the robustness of a decision is an important consideration when the decision situation involves a high degree of epistemic uncertainty. Methods in which no evaluation of the robustness of the decision is provided are not suitable for analysing decision situations in which there is a high degree of epistemic uncertainty.

Supersoft Decision analysis and extended decision analysis are not viewed as competing methods. Rather, they are seen as complementing each other, the one being able to deal better than the other with situations of certain types. Use of extended decision analysis is appropriate when epistemic uncertainty can be quantified in terms of specific probability distributions, whereas Supersoft Decision analysis is appropriate when no such quantification is possible. Extended decision analysis provides more precise information on the robustness of the results (since it provides a distribution for each expected utility), whereas, Supersoft Decision analysis provides only an interval for the expected utilities. In practice, one could start analyzing a problem by use of a rough model based on SSD. Depending on the results of this initial analysis, one could then, if desired and regarded as possible, continue with a more refined analysis involving use of extended decision analysis.

The examples presented indicate Supersoft Decision theory (SSD) to be a tool that can be used for analyzing practical decision problems. The main advantage of using SSD is that it allows the decision maker to employ imprecise statements concerning the probability and utility values used in the model. This also makes SSD particularly suitable for difficult decision problems in which the decision maker is not a single person, but a group of persons. The use of imprecise statements allows the interval used to denote the probabilities and utilities to be sufficiently broad to encompass the estimates and views of all the group members.

ACKNOWLEDGMENTS

The authors would like to thank The Swedish Fire Research Board (BRANDFORSK) for funding the research on which the present paper is based.

REFERENCES

1. Oberkampf, W.L., DeLand, S.M., Rutherford, B.M., Diegert, K.V. and Alvin, K.F., "Error and Uncertainty in Modeling and Simulation," Reliability Engineering and System Safety, Vol. 75, No. 3, 2002, pp. 333–357.
2. Gomberg. A., Buchbinder, B. and Offensend, F.J., Evaluating Alternative Strategies for Reducing Residential Fire Loss – The Fire Loss Model, U.S. Department of Commerce, Washington, National Bureau of Standards, NBSIR 82-2551, 1982.
3. Johansson, H., Decision Making in Fire Risk Management, Report 1022, Lund, Dept. of Fire Safety Engineering, Lund University, 2001.
4. Johansson, H., "Decision Analysis Concerned with Investments in Fire Safety," 7th International Symposium on Fire Safety Science, Worcester, 2002.

5. Donegan, H.A., "Decision Analysis," SFPE Handbook of Fire Protection Engineering, 3rd ed., Quincy, National Fire Protection Association, 2002.
6. Laplace, P.S., Essai Philosophique sur les Probabilites, 5th ed., Paris, Translation by Dover, New York, 1952.
7. Wald, A., Statistical Decision Functions, Wiley, New York, 1950.
8. Savage, L.J., "The Theory of Statistical Decision," J. Amer. Statist. Assoc., Vol. 46, 1951, pp. 55–67.
9. Hurwicz, L., "Optimality Criteria for Decision Making Under Ignorance," Discussion paper No. 370, Cowles Commission, 1951.
10. Berger, J.O., Statistical Decision Theory – Foundations, Concepts and Methods, Springer-Verlag, New York, 1950.
11. Malmnäs, P.-E., Foundations of Applicable Decision Theory, Department of Philosophy, Stockholm University, Stockholm, 1999.
12. Malmnäs, P.-E., Methods of Evaluation in Supersoft Decision Theory, Department of Philosophy, Stockholm University, Stockholm, 1995.
13. Wallsten, T.S., Budescu, D.V., Rapoport, A., Zwick, R. and Forsyth, B., "Measuring the Vague Meanings of Probability Terms, " Journal of Experimental Psychology: General, Vol. 115, 1986, pp. 348–365.
14. Spetzler, C.S. and Stael Von Holstein, C.-A.S., "Probability Encoding in Decision Analysis," Management Science, Vol. 22, No. 3, 1975, pp. 340–358.
15. Malmnäs, P.-E. (ed.), Hypermjuk Beslutsteori Och Ekonomisk Optimering av det Industriella Brandskyddet, Stockholm, Thales, 2002.
16. Danielsson, M., Computational Decision Analysis, Doctoral Thesis, Department of Computer and Systems Sciences, Stockholm University and Royal Institute of Technology, Stockholm, 1997.
17. Coleman, T., Branch, M.A. and Grace, A., Optimization Toolbox User's Guide, Version 2, The Math Works Inc., 1999.
18. Johansson, H., "Investment Appraisal Using Quantitative Risk Analysis," Journal of Hazardous Materials, Vol. 93, 2002, pp. 77–91.
19. Hagen, O., "Towards a Positive Theory of Preferences Under Risk," In: Allais, M. and Hagen, O. (eds), Expected Utility and the Allais Paradox, Dordrecht, Reidel, 1979, pp. 271–302.
20. Fishburn, P., "Transitive Measurable Utility," Journal of Economic Theory, Vol. 31, 1983, pp. 293–317.
21. Loomes, G. and Sugden, R., "Regret Theory: an Alternative Theory of Rational Choice Under Uncertainty," The Economic Journal, Vol. 92, 1982, pp. 805–824.
22. Green, J. and Jullien, B., "Ordinal Independence in Nonlinear Utility Theory," Journal of Risk and Uncertainty, Vol. 1, 1988, pp. 355–387.
23. Quiggin, J., "A Theory of Anticipated Utility," Journal of Economic Behaviour and Organization, Vol. 3, 1982, pp. 323–343.
24. Yaari, M., "The Dual Theory of Choice Under Risk," Econometrica, Vol. 55, 1987, pp. 95–115.