IIASA
International Institute for Applied Systems Analysis • A-2361 Laxenburg • Austria
Tel: +43 2236 807 • Fax: +43 2236 71313 • E-mail: info@iiasa.ac.at • Web: www.iiasa.ac.at

INTERIM REPORT IR-97-045/October

Catastrophic Risk Evaluation

Love Ekenberg (lovek@dsv.su.se)
Magnus Boman (mab@dsv.su.se)
Joanne Linnerooth-Bayer (bayer@iiasa.ac.at)

Approved by
Gordon MacDonald (macdon@iiasa.ac.at)
Director, *IIASA*

Contents

Abstract

A body of empirical evidence has shown that many managers would welcome new ways of highlighting catastrophic consequences, as well as means to evaluating decision situations involving high risks. When events occur frequently and their consequences are not severe, it is relatively simple to calculate the risk exposure of an organisation, as well as a reasonable premium when an insurance transaction is made. The usual methods rely on variations of the principle of maximising the expected utility (PMEU). When, on the other hand, the frequency of damages is low, the situation is considerably more difficult, especially if catastrophic events occur. When the quality of estimates is poor, e.g., when evaluating low-probability/high-consequence risks, the customary use of quantitative rules together with overprecise data could be harmful as well as misleading.

This work extends the risk evaluation process by the integration of procedures for handling vague and numerically imprecise probabilities and utilities. The shortcomings of PMEU, and of utility theory in general, can in part be compensated for by the introduction of risk constraints. We point out some problematic features of the evaluations performed using utility theory. We also criticise the demand for precise data in situations where none is available. As an alternative to traditional models, we suggest a method for the evaluation of risks when the information at hand is numerically imprecise. The method includes procedures that allow for interval statements and comparisons, and thereby it does not require the use of numerically precise statements of probability, cost, or utility in a general sense. In order to attain a reasonable level of security, and because it has been shown that managers tend to focus on large negative losses, it is argued that a risk constraint should be imposed on the analysis. The strategies are evaluated relative to a set of such constraints considering how risky the strategies are.

Acknowledgments

Love Ekenberg’s work was supported by the Swedish Foundation for International Cooperation in Research and Higher Education (STINT). The authors also gratefully acknowledge the support by the Swedish Council for Planning and Coordination of Research (FRN), Stockholm.

The authors would also like to thank Johan Thorbiörnson, Mid Sweden University, for his valuable comments.

About the Authors

Love Ekenberg and Magnus Boman are from the Department of Computer and Systems Sciences, Stockholm University, Electrum 230, S-164 40 Kista, Sweden. Love Ekenberg was a Guest Research Scholar at IIASA from October 1996 until June 1997. Magnus Boman is also affiliated with The Royal Institute of Technology.

Joanne Linnerooth-Bayer is Co-leader of the Risk, Modeling and Policy Project at IIASA.

Catastrophic Risk Evaluation

Love Ekenberg, Magnus Boman, Joanne Linnerooth-Bayer.

1. Motivation

A variety of methods have been developed for estimating losses and risks (cf., e.g.,[Ekenberg, et al., 1995]). When events occur frequently and when they are not very severe, it is relatively simple to calculate the risk exposure of an organisation, as well as a reasonable premium when, for instance, an insurance transaction is made. Commonly used methods rely on variations of the principle of maximising the expected utility (PMEU). When, on the other hand, the frequency of damages is low, the situation is considerably more difficult; in particular, if catastrophic events can occur. When the quality of the estimates is poor, e.g., when evaluating low-probability/high-consequence events, the customary use of quantitative rules together with overprecise data could be harmful as well as misleading. Nevertheless, PMEU is the basis of many evaluation models, even if the "[...] historical records of catastrophes does not always provide an adequate number of events to estimate reliably the value of some key variables [...]" ([ISO, 1996], p. 36).

Even if the evaluation of a strategy results in an acceptable expected utility, its consequences might be so serious that it should nevertheless be avoided, even if the probability of catastrophic events is low. Therefore, it seems necessary to supplement PMEU by other criteria. Furthermore, since the information of low-probability/high-consequence events is necessarily very uncertain, an important objective is to extend the risk evaluation process by the integration of procedures for handling vague and numerically imprecise information.

In this paper, we scrutinise some problematical features of using utility theory. We also point out some problems with demanding precise data in situations where none is available. As an alternative to traditional models, we suggest a method for the evaluation of risks when the information at hand is numerically imprecise. We begin, in Section 2, with an overview of utility theory and point out some problems with the use of PMEU. Section 3 describes how utility theory can be extended by procedures for evaluating imprecise information, and critically reviews some earlier proposals. Section 4 describes the idea of risk constraints as a means to eliminating alternatives that may lead to catastrophic events.

2. Utility Theory

The utility model has been used for various purposes. Economists use it primarily as a predictive tool [Friedman, 1953], and secondarily as a descriptive tool [Bettman, 1979]. From the predictive viewpoint, the logical basis of the theory is not important. What matters is its ability to predict the behaviour of decision makers. Descriptive studies focus on how well a theory is able to model real-life decision processes. Another perspective is the normative or prescriptive one, reflecting how a rational decision maker should behave. For normative decision models, the logical foundations and the validity of the model are important. Proponents often argue by presenting what they regard as reasonable axioms, from which they deduce criteria for ranking the alternatives. Naturally, this does not persuade everyone. To show the problems connected with the diversity of axiomatisations, [Fischhoff, et al., 1983] investigate some central questions concerning whether people accept the axioms upon which the model is based.

To be more precise, the meaning of an agent accepting the utility principle and PMEU can be formulated in the following way: Given a structured decision situation containing a finite set of alternatives $\{A_i\}$, suppose that each alternative A_i in this set can be represented by a set of consequences and a set of numbers $(\{c_{ij}\},\{p_{ij}\})$, in which $\{c_{ij}\}$, is the set of consequences of A_i, and p_{ij} is the probability that c_{ij} occurs given that A_i is performed.[1]

Definitions: If A_i is $(\{c_{ij}\},\{p_{ij}\})$, and V is a real-valued function on $\{c_{ij}\}$, then A_i has a *value* equal to $\Sigma p_{ij}V(c_{ij})$, sometimes denoted $E_V(A_i)$.[2]

An agent *accepts the utility principle* iff it assigns the value $\Sigma p_{ij}V(c_{ij})$ to A_i, given that it has assigned the value $V(c_{ij})$ to c_{ij}.

Given the alternatives A_i and A_j, an ordering « of the alternatives is *compatible to PMEU* iff A_i « A_j implies $E_V(A_i) \leq E_V(A_j)$.

An agent *accepts PMEU* iff its ordering(s) of the values of the alternatives is compatible to that principle.

Until the late 1930's, utility theory was not a well-founded subject. This changed with the works of Ramsey, von Neumann and Morgenstern [Neumann and Morgenstern, 1947], and [Ramsey, 1978]. They structured decision analysis by proposing reasonable principles governing decisions, and by constructing a theory out of them. In other words, they (and later many others) formulated sets of axioms meant

---

1 Note that in some formulations the concept of state is used, and that probabilities are assigned to different states instead. However, the two models can be straightforwardly transformed into each other when the number of states and consequences is finite cf., e.g., [Jeffrey, 1983].

2 A survey of different interpretations of the utility principle and PMEU and a more general characterisation of the class of expected utility models is given in [Schoemaker, 1982], p.530 ff. An expected utility model is one which predicts or prescribes that people maximise the expression $\Sigma\Phi(p)U(x)$, where x is an outcome vector. The models differ in (i) how utility is measured, (ii) what kind of concept of probability $\Phi(p)$ is allowed, and (iii) how the outcomes are measured. Schoemaker then examines some frequently used variants of models, according to this structure.

to justify their particular attitude towards the utility principle.[3] Such axiomatic systems usually consist of primitives (such as «, states, sets of states, etc.) and axioms constructed from the primitives. The axioms (ordering axioms, independence axioms, continuity axioms, etc.) imply a numerical representation of preferences and probabilities. Typically implied by the axioms are existence theorems stating that a probability distribution and a utility function exist, and a uniqueness theorem stating that two utility functions, relative to a given preference ranking, are always linear transformations of each other.

It is generally agreed that a large number of axiomatic theories provide justification of PMEU.[4] A common counter argument to this is that the independence axioms of utility theory are fallacious. For instance, [Allais, 1953] has shown that people do not act in accordance with certain independence axioms in the system of [Savage, 1972]. An even more serious problem with the formal justifications of the principle is that even if the axioms in the various axiomatic systems are accepted, the principle does not follow, i.e. the proposed systems are too weak to imply the utility principle and PMEU. Thus, it is highly doubtful whether utility theory can be justified on purely formal grounds. For instance, [Malmnäs, 1994a] demonstrates the weakness in this respect for the systems in [Hernstein and Milnor, 1953], [Oddie and Milne, 1990], and [Savage, 1972]. A fuller account of a wide variety of systems is provided in [Malmnäs, 1990a], which argues that it is implausible that these systems could be extended in any reasonable way to imply the principle. Thus, the logical foundations of utility theory seems to be fundamentally insufficient.

Furthermore, a troubling dilemma with PMEU is that low-probability/high-consequence events, for example, the risks of air travel, will be valued no differently from the same risks generated by independent trials, for example, automobile travel. Indeed, valuing the former more than the latter may even be considered immoral to the extent that one values lives lost collectively more than those same lives lost individually. However, [Linnerooth-Bayer, 1996] argues that this conclusion is misconstrued to the extent that it depends on a strict interpretation of PMEU which does not allow for individual preference on how the odds are generated, independently or collectively.

Another line of criticism is that utility theory is insufficient for modelling risk attitudes adequately. Proponents of utility theory usually present the following argument to show that the concept of utility captures different risk attitudes. Suppose, for example, that all consequence values are monetary. Then associate two expectations with each alternative: its expected monetary value and its expected utility. To each expected utility, there corresponds a certainty monetary equivalent x_{ce}. The decision maker is indifferent to having this monetary value for certain, as compared to the alternative action, i.e., $u(x_{ce}) = \Sigma p_{ij}u(x_{ij})$, where $u(x_{ij})$ is the utility of receiving the monetary value x_{ij}. The risk premium *p* of an act can be defined as the demand that a

---

3 See, e.g., [Savage, 1972], [Hernstein and Milnor, 1953], [Suppes, 1956], [Luce and Krantz, 1971], and [Jeffrey, 1983]. Surveys over a wide variety of axiomatic systems are given in [Fishburn, 1981] and [Malmnäs, 1990a].

4 Recently, several researchers in the area of autonomous agents have equated rationality with the use of PMEU as a prescriptive decision rule (see, e.g., [Gmytrasiewicz and Durfee, 1993]). However, there are many reasons to not identify these (cf. [Boman, 1995, Boman and Ekenberg, 1995]).

decision maker has for carrying out the action, instead of having the monetary equivalent x_{ce} for certain, i.e., $p = \Sigma p_i x_i - x_{ce}$. Decision makers can be divided into three classes with respect to the risk premium p: a decision maker is risk averse if p is positive; risk prone if p is negative; and finally risk neutral if $p = 0$. From this description, it is easy to see how a utility function can be regarded as modelling a decision maker's attitude towards risk. Assume that a decision maker needs a certain amount of money, and a lesser amount would not be very useful to her. The decision maker would normally risk only a certain part of her wealth, but it is not difficult to imagine extreme situations where this does not hold true, for instance when she is desperately in need of money for a medical treatment of a disease that, if not cured, will result in death. Faced with a choice of buying a ticket or not with her last funds that will yield a chance of winning an amount sufficient for paying for the treatment, most rational people would gamble. The risk premium p in such a situation is probably a negative number.

However, the use of a utility function to formalise the decision process with all possible risk attitudes is not possible. The critics point out that most mathematical models of decision analysis are oversimplified and disregard important factors (cf., e.g., [Schoemaker, 1982]). Moreover, individuals tend to avoid the use of precise probability estimates. It also seems that the value of the outcome defines risk for managers, rather than the weighted average of probabilities and utilities [Kunreuther, et al., 1978], [Shapira, 1995]. This is partly a function of the reluctance to take explicit responsibility for losses [Kahneman and Lovallo, 1993].

Thus, it seems reasonable that a normative decision theory should be very sensitive to different risk attitudes, should give the decision maker the means to express her risk attitudes in a variety of ways, and should provide procedures for handling qualitative aspects as well as quantitative ones. Some researchers have tried to modify the behaviour of PMEU by bringing regret or disappointment into the evaluation to cover cases where numerically equal results are appreciated differently depending on what was once in someone's possession [Loomes and Sudgen, 1982]. Several researchers, among them [Quiggin, 1982], try to resolve the problems mentioned above by requiring functions to modify both the probabilities and the values. However, [Malmnäs, 1996] shows that for these and other proposals, their performance, at best, almost equals that of the expected value and, at worst, is much worse, for example, not being consistent with first order stochastic dominance.

3. Indeterminacy in Decision Situations

The question also arises whether people are capable of providing the inputs which utility theory requires (cf., e.g., [Fischhoff, et al., 1983]). For instance, most people cannot distinguish between probabilities ranging roughly from 0.3 to 0.7 [Shapira, 1995]. Furthermore, even if a decision maker is able to discriminate between different probabilities, very often adequate and precise information is missing. Consequently, there seems to be significant reasons to discriminate between measurable and immeasurable uncertainty. Measurable uncertainty is often referred to as risk and can be represented by precise probabilities. By contrast, immeasurable uncertainty occurs frequently in high consequence situations since there is often a lack of statistical data, and thereby the axioms of Savage and others are not satisfied. [Ellsberg, 1961] proposes a class of choice situations involving immeasurable uncertainty, in which many people do not satisfy the variety of axiomatic systems suggested. Ellsberg is not averse to the use of PMEU, but he suggests that the underlying axiomatic systems should not be applied in situations where the information is numerically imprecise. An alternative is to handle impreciseness using only qualitative data.[5] However, in many cases this seems to be too restrictive, in which case numerical estimates should still play a role. There have been many attempts to express probabilities in terms of intervals.[6] Although a complete discussion goes beyond the scope of this discussion, it might be enlightening to see a few examples of how the strong requirements of utility theory can be relaxed. Before giving the basics of our own approach, we briefly present the approaches of [Gärdenfors and Sahlin, 1982] and [Levi, 1974].

[Gärdenfors and Sahlin, 1982], p.362 state that "[…] the amount and quality of information which the decision maker has concerning the possible states and outcomes of the decision situation in many cases is an important factor when making the decision. […] the information available concerning the possible states and outcomes of a decision situation has different degrees of *epistemic reliability*." The authors assume that the beliefs about the states of nature can be represented by a set P of probability measures, which is supposed to contain all epistemically possible probability measures over the states, i.e., they associate with each state a set of probability values $P_i(s)$, where each $P_i \in P$. The set of all epistemically possible probability measures is then restricted by an epistemic reliability measure ρ, intended to represent the different degrees of reliability of the different probability distributions. The decision problem is then divided into two steps. First, the decision maker has to recognise a set of probability measures, which she has to restrict to a set with a satisfactory degree of epistemic reliability. When this is done, Gärdenfors and Sahlin propose that the agent should obey the maximin criterion for expected utilities. This means that the minimal expected utility of an alternative is computed by taking each probability measure in the set of probability measures with a satisfying degree of epistemic reliability, and then calculating the expected utility in the usual way. The option with the least expected utility is chosen. An agent then obeys the maximin criterion for expected utilities by choosing the alternative with the largest minimal expected utility. They also show that both the classical maximin rule and the PMEU are special cases of their criterion. A problem, however, with this approach is that the minimal expected utility of an alternative is assumed to be on the boundary of the respective probability intervals, where they seem to be least reliable. One could argue that the intention behind the choice of the acceptable set is that its measures are equivalent with respect to ρ. But the properties of ρ are not clear enough to give any hint on how this could be established.

Isaac Levi uses a similar approach, giving three conditions that a set of probability measures B must satisfy.[7] These imply (among other things) that the probability distributions in B for a given state of nature form an interval. He notices that some authors have presupposed such an interval in their theories, but concludes that his own

[5] For instance, this topic has recently been addressed at a recent AAAI Spring Symposium [Doyle et al., 1997].

[6] For instance, [Choquet, 1953/54], [Kyburg, 1961] [Smith, 1961], [Good, 1962], [Dempster, 1967], [Huber, 1973], [Huber and Strassen, 1973], [Shafer, 1976], [Nilsson, 1986].

[7] [Levi, 1974], pp.393–402.

theory "[…] recognizes credal states as different even though they generate the identical valued function–provided they are different convex sets of Q-functions."[8] The significance is emphasised as Levi compares the different alternatives in decision situations. He gives an example when two similar decision situations with different sets of probability measures yield results different from his theory, even if the generated intervals are the same.[9]

Levi also relaxes the Bayesian requirement on representing the utilities of the consequences. He introduces a set G of permissible utility functions, which do not obey the classic Bayesian requirement that all elements in G are linear transformations of each other. He then stipulates the following definitions:

Definition: An alternative A is *E-admissible* iff there is a probability distribution p in B and a utility function u in G, such that E(A), defined relative to p and u, is optimal among all alternatives.

Definition: An alternative A is *S-admissible* iff it is E-admissible and there is a function u in G such that the minimum u-value assigned to some possible consequence is at least as great as the u-values assigned to the consequences of any other of the remaining alternatives.[10]

These definitions seem reasonable, but they have some counter-intuitive implications. They clearly violate the reasonable condition of independence of irrelevant alternatives, i.e. the ordering between the alternatives is not affected by the addition of a new alternative.[11] The theory is also problematic in some respects when confronted with some empirical results, e.g. a situation described in [Ellsberg, 1961], often referred to as Ellsberg's paradox.[12]

[Danielson and Ekenberg, 1997] take another approach. Imprecise probabilities as well as imprecise utilities are handled by modelling a decision situation with numerically imprecise sentences such as "the probability of consequence c_{11} is greater than 5%" and comparative sentences such as "consequence c_{11} is preferred to consequence c_{12}". More generally, imprecise information is represented by vague sentences, interval sentences, and comparative sentences. Examples of *vague sentences* are: "The consequence c_{ij} is probable" or "The event c_{ij} or c_{ik} is possible". Such sentences may be represented by suitable intervals according to the decision maker. Suppose, for example, that a decision maker stipulates that for c_{ij} to be called probable, the probability for it to occur must be greater than 5% but less than 60%. In this case, the translation will be $p_{ij} \in [0.05,0.60]$. This is represented by the two linear inequalities $p_{ij} \geq 0.05$ and $0.60 \geq p_{ij}$. Similar translations apply when representing other vague sentences. *Interval sentences* are of the form: "The probability of c_{ij} lies between the numbers a_k and b_k" and are translated to $p_{ij} \in [a_k,b_k]$. Finally, *comparative sentences* are of the form: "The probability of c_{ij} is greater than the probability of c_{kl}". Such a sentence is translated into an inequality $p_{ij} > p_{kl}$. Each statement is thus represented by

---

8 ibid. p.408. The Q-functions are Levi's probability measures.

9 ibid. pp. 416, 417.

10 We deviate from Levi in that he also uses the concept P-admissible. However, this concept is not defined and we omit it from this presentation.

11 For further details, see [Gärdenfors and Sahlin, 1982], p.379.

12 In [Malmnäs, 1990b] and [Gärdenfors and Sahlin, 1982].

one or more constraints. The conjunction of constraints of the four types above, together with $\Sigma\ p_{ij} = 1$ for each strategy A_i involved, is called the *probability base* (P). The *value base* (V) consists of similar translations of vague and numerically imprecise value estimates. The collection of probability and value statements constitutes the *information frame*.

Definition: An *information frame* is a structure $\langle\{C_1,\ldots,C_m\},P,V\rangle$, where each C_i is a finite set of consequences $\{c_{i1},\ldots,c_{ih_i}\}$. P is a finite list of linear constraints in the probability variables and V is a finite list of linear constraints in the value variables.

For instance, the expected value of an action with respect to an information frame can be expressed by the following definition.

Definition: Given an information frame $\langle\{C_1,\ldots,C_m\},P,V\rangle$, the *expected value* of an action A_i can be expressed as $E(A_i) = p_{i1}\cdot v_{i1} +\ldots+ p_{in}\cdot v_{in}$, where v_{ij} denotes the value of the consequence c_{ij}, and p_{ij} denotes the probability of c_{ij} occurring given that action A_i is taken.

If the expected value in the definition above seems to be very similar to the expected value as defined in Section 2, it is important to bear in mind that this is evaluated with respect to the solution sets of the information frames rather than to precise numbers. The expected value is now a range of possible values. A quite uncontroversial strategy of evaluation is “never to eliminate or disqualify an action that might be the best one”. However, the only option becomes then “never to eliminate any strategy”. Needless to say, this is a very weak decision strategy.

One less careful attempt to discriminate between different actions is to introduce the concept of admissibility [Lehmann, 1959]. Intuitively, an action can be discarded if it is always worse than all other actions, i.e., an admissible alternative is in some sense a non-dominated alternative.[13] This is (informally) described by the following:

Definitions: Given an information frame $\langle\{C_1,\ldots,C_m\},P,V\rangle$, A_i is *at least as good as* A_j iff $E(A_i) - E(A_j) \geq 0$, for all instances of the probability and utility variables that are solutions to P and V.

A_i is *better than* A_j iff A_i is at least as t-good as A_j, and $E(A_i) - E(A_j) > 0$, for some instances of the probability and utility variables that are solutions to P and V.

A_i is an *admissible alternative* iff no other alternative is better.

However, the concept of admissibility is based on PMEU, and thus the approach of considering only admissible actions cannot be entirely uncontroversial. Since the idea to dismiss a clearly inferior action seems to be reasonable, we have to be careful about how we measure this inferiority.

Furthermore, the imprecision represented in the information frames often, viz. in most non-trivial situations, results in the ranges of the expected value of some actions being overlapping. Thus, the set of candidates will often be far too large. Consequently, even if we adhere to PMEU, we still have weak guidance on how to choose. Thus, we

---

13 For instance, Levi argues ([Levi, 1992], §4) that to choose rationally is to restrict the choice to an admissible option.

need further discrimination principles. One way to proceed is to determine the stability of the relation between the actions under consideration. Values near the boundaries of the intervals are probably less reliable than more central values due to interval statements being deliberately imprecise. This can be taken into account by measuring the dominated regions indirectly with the use of the concept of contraction, which is motivated by the difficulties of performing sensitivity analyses in several dimensions simultaneously [Ekenberg, et al., 1996], [Danielson, 1997], [Danielson and Ekenberg, 1997]. It can be difficult to gain a real understanding of the solutions to large decision problems using only one-dimensional analyses, since different combinations of dimensions can be critical to the evaluation results. Investigating all possible combinations would lead to a computationally demanding procedure. Using contractions, this difficulty is circumvented. The idea behind the principle is to investigate how much the different intervals can be decreased before an expression such as $E(A_i) - E(A_j) > 0$ ceases to be consistent. At the same time, we must avoid the complexity inherent in combinatorial analyses, but still be able to study the stability of a result by gaining a better understanding of how important the interval boundary points are.

4. Risk Constraints

A key observation in this paper is that there is no perfect evaluation rule, although PMEU is at least as good as many of its contenders. To improve that rule (or any other numeric rule), we might extend it with supplementary rules rather than engaging in further modifications of replacement rules in pursuit of the perfect rule. Thus, there is no reason to completely reject the use of PMEU, but a reasonable method should provide possibilities to evaluate decision situations in several respects, and the purpose of this section is to propose how this can be done by a class of decision rules called risk constraints.

For instance, a risk constraint can provide thresholds beyond which a strategy is undesirable. If a decision maker can work with numerically precise information only, such exclusions can be dealt with by specifying constraints using two thresholds—one for the probability and one for the value. Then a strategy would be undesirable if it violates both of these thresholds. For instance, a strategy could be considered undesirable if a consequence with a probability above 0.15 also has a value below 0.10.[14]

When the information is numerically imprecise, the meaning of such thresholds is not obvious. In [Ekenberg, et al., 1997] it is suggested that the interval limits together with stability analyses should be considered in such cases. For example, assume that a company desires to decrease its exposure to risk by installing more (but quite expensive) protective equipment. Furthermore, suppose that the management of the company has stipulated that a strategy A_i is undesirable if the consequence c_{ij} belongs to A_i. There is a possibility that the value of c_{ij} is less than 0.45, and that the probability of c_{ij} is greater than 0.65. Now, assume that the strategy “Install protective equipment” has a consequence “reduction of mean time between failure by less than 12 hours”. The management estimates that the value of this consequence is in the interval [0.40, 0.60],

[14] This kind of risk constraint is very similar to the concept of acceptable level of risk.

and that the probability is in the interval [0.20, 0.70]. Then this strategy is below the threshold, and thus undesirable. To further refine the analysis, it could be investigated where the risk constraints are violated. In this way, the stability of the result can be studied. For instance, it can be seen that the strategy ceases to be undesirable when the left end-point of the value interval is increased by 0.05. Thus, exclusions of this kind can be dealt with by specifying a risk constraint using a number of thresholds, including one for the probability and one for the value. A strategy is considered undesirable if it violates these thresholds.[15]

However, this is still insufficient. In many situations there are trade-offs between the value of a possibly disastrous event and the probability with which it can occur. This is particularly true in situations where all strategies are undesirable, but we nevertheless have to choose. For instance, a strategy with a low-valued consequence that may occur might be perceived as better than a strategy with a slightly better but considerably more probable consequence. Furthermore, as [Linnerooth-Bayer, 1996], pp.135–138, points out, there is sometimes a need to take the distribution of risk into account, i.e.; it may appear unfair that a few people are exposed to the full risk. Consider the following situation. From a group of 1000 people, two people are exposed to an extremely high probability of dying, and the others face no risk at all. Given reasonable conditions of the utility function, this would have the same expected utility as a situation when there is a 1/1000 chance of a disastrous event taking the lives of all 1000 people. Nevertheless it may be perceived that the first situation is unfair, in particular if the benefits are equally distributed.

Below, the concept of a risk constraint is generalised. We also show how this concept can be used for handling a number of different types of constraints that can be imposed on a decision situation.

First, the concept of a cube is defined. Intuitively, a cube is nothing more than a unity cube with the same dimension as the number of consequences in a given decision situation.

Definition: Given an information frame $F = \langle\{\{c_{ij}\}_{j = 1,...,h_i}\}_{i = 1,...m},P,V\rangle$, a *cube* with respect to F is the set of equations $\{0 \leq p_{ij} \leq 1\}_{j = 1,..., h_i, i = 1,...m}$ or $\{0 \leq v_{ij} \leq 1\}_{j = 1,..., h_i, i = 1,...m}$. Such a cube will be denoted $B = (b_1,\ldots,b_k)$ below.[16]

Just for completeness, a definition of the solution set to a set of linear constraints L is included. A solution set for L consists of vectors consistent with L.

Definitions: Given a cube $B = (b_1,\ldots,b_k)$ and a set L of linear constraints of B. A list of numbers $[n_1,\ldots,n_k]$ is a *solution vector* to L if the substitution of n_i for b_i, for all $1 \leq i \leq k$, in L does not yield a contradiction.[17] The set of solution vectors to L constitutes the *solution set* for L and will be denoted s(L) below.

---

[15] [Malmnäs, 1994b] supplements PMEU by some qualitative evaluations. An example is the qualitative sorting function, further developed in [Ekenberg, et al., 1997].

[16] We will omit the reference to the frame below, since the frame will be clear from the contexts.

[17] To be precise, a list of numbers is a solution vector to L with respect to B. However, for the same reason as in the previous footnote, we will use the more compact notation above.

Next, the concept of a risk constraint as well as what is meant by an information frame violating a risk constraint are defined. Intuitively, a risk constraint is a function that states a set of constraints that must be fulfilled for an information frame to be acceptable. The definition is general and the examples below show how it can be used to state a variety of decision rules.

Definition: Given an information frame $F = \langle C,P,V\rangle$, let $\{p_1,...,p_k\}$ and $\{v_1,...,v_k\}$ be the sets of variables in P and V, respectively.
Let $P' = (p_1,...,p_k)$ and $V' = (v_1,...,v_k)$ be two bases. A *risk constraint* with respect to F is a function S from $2^{s(V')}$ into $2^{s(P')}$.

Definition: Given an information frame $F = \langle C,P,V\rangle$, and a risk constraint S with respect to F, F *violates* S iff
there is an element (α,β) of S, such that not $s(V) \subseteq \alpha$ and not $s(P) \subseteq \beta$.

Since the operations on the different bases are independent of each other, determining whether an information frame violates a risk constraint is simple.

In the following examples, we will use the following notation: Given a (probability or utility) base B and a constraint *e*, $s(\{e\})$ denotes the set $s(\{e\} \cup B)$. We will also use notations from first-order logic for notational convenience. These have their usual set-theoretic counterparts. Thus, for example, given a base B and two constraints *e* and *d*, $s(\{e \,\&\, d\})$ denotes the set $\{\alpha : \alpha \in s(\{e\} \cup B) \cap s(\{d\} \cup B)\}$. Analogously, $s(\{\neg e\})$ denotes the set $\{\alpha : \alpha \in s(\{B\}) - s(\{e\} \cup B))\}$.

Furthermore, given an information frame $F = \langle\{\{c_{ij}\}_{j = 1,...,h_i}\}_{i = 1,...m},P,V\rangle$, then $S(s(\{f(v_{ij})\})) = s(\{g(p_{ij})\})$ is used to denote $S(s(\{f(v_{11})\})) = s(\{g(p_{11})\}),\ldots, S(s(\{f(v_{mh_m})\})) = s(\{g(p_{mh_m})\})$.

In the following example, the risk constraint states that the value of consequence c_{11} must not possibly be less than 0.1, and at the same time possibly occurring with a probability greater than 30%.

Example: Given the information frame $F = \langle\{\{c_{11},c_{12},c_{13}\}, \{c_{21},c_{22}\}\}, \{p_{11}\geq 0.1, p_{11}\leq 0.5, p_{12}\geq 0.5\}, \{v_{11}\leq 0.2, v_{12}\geq 0.1\}\rangle$.
Define a risk constraint with respect to F, $S(s(\{v_{11}\geq 0.1\})) = s(\{p_{11}\leq 0.3\})$.[18] Now, F clearly violates S, since, for example, $v_{11} = 0$ is a possible substitution in the value base, and $p_{11} = 0.4$ is a possible substitution in the probability base. Both these are outside the elements in the risk constraint.

Naturally, the restriction to c_{11} is not necessary. A more natural restriction is that no consequence with these attributes may occur. This is easily stated as in the following example.

---

[18] Since P is an probability base, the constraints $\Sigma p_{ij}=1$, for each consequence set, and $\forall p_{ij}(p_{ij}\geq 0)$ is implicitly included in the information frame: Similarly, $\forall v_{ij}(v_{ij}\geq 0)$ and $\forall v_{ij}(v_{ij}\leq 1)$ are included.

Example: Given the information frame $F = \langle\{\{c_{11},c_{12},c_{13}\}, \{c_{21},c_{22}\}\}, \{p_{11}\geq 0.1, p_{11}\leq 0.5, p_{12}\geq 0.5\}, \{v_{11}\leq 0.2, v_{12}\geq 0.1\}\rangle$.
Define a risk constraint with respect to F, $S(s(\{v_{ij}\geq 0.1\})) = s(\{p_{ij}\leq 0.3\})$. According to the notation introduced above, the latter is the same as $S(s(\{v_{11}\geq 0.1\})) = s(\{p_{11}\leq 0.3\})$, $S(s(\{v_{12}\geq 0.1\})) = s(\{p_{12}\leq 0.3\})$, $S(s(\{v_{13}\geq 0.1\})) = s(\{p_{13}\leq 0.3\})$, $S(s(\{v_{21}\geq 0.1\})) = s(\{p_{21}\leq 0.3\})$, $S(s(\{v_{22}\geq 0.1\})) = s(\{p_{22}\leq 0.3\})$.

More generally, we also take into account how subsets of consequences, i.e. events, can make a strategy undesirable. This means that if several consequences of a strategy are too dire (with respect to a certain value threshold), their total probability should be considered even if their individual probability is too low to make the strategy undesirable. The intuition is that all subsets of consequences $\{c_{ij_1},\ldots,c_{ij_k}\}$ of a strategy A_i are taken into account. Each subset is checked for consistency with the expressions in the information frame in order to see if the value of each element in the subset is less than a certain value threshold. If this is the case, the probability of this subset is checked, i.e. if the sum of the probabilities for all elements is greater than a certain probability threshold.[19]

Example: Given the information frame $F = \langle\{\{c_{11},c_{12},c_{13}\}, \{c_{23}\}\}, \{p_{11}\geq 0.1, p_{11}\leq 0.7, p_{12}\geq 0.3\}, \{v_{11}\leq 0.2\}\rangle$.
Define a risk constraint with respect to F by $S(s(\{v_{11}\geq 0.1 \;\&\; v_{12}\geq 0.1\})) = s(\{p_{11}+p_{12}\leq 0.6\})$. Now, F violates S since, for example, $v_{11} = 0$, $v_{12} = 0$ is a possible substitution in the value base, and $p_{11} = 0.6$, $p_{12} = 0.4$ is a possible substitution in the probability base.

Similarly, a functional relationship can be defined between the value and the probability bases. This can be used to take account of situations where there are a relation between the probability of a possibly disastrous event and the probability with which it can occur.

Example: Given the information frame $F = \langle\{\{c_{11},c_{12},c_{13}\}, \{c_{23}\}\}, \{p_{11}\geq 0.1, p_{11}\leq 0.7, p_{12}\geq 0.3\} \{v_{11}\leq 0.2\}\rangle$. Let $g: [0,0.7]\rightarrow[0.3,1]$, such that $g(x) = x+0.3$.
Define a risk constraint with respect to F by $S(s(\{v_{11}>x\})) = s(\{p_{11}<g(x)\})$. Now, F violates S, since for example $x = 0.1$ is a possible substitution, $v_{11} = 0$ is a possible substitution in the value base, and $p_{11} = 0.5$ is a possible substitution in the probability base.

It could also be the case that an alternative with a negative consequence is acceptable since there might be consequences that are very positive.

Example: Given the information frame $F = \langle\{\{c_{11},c_{12}\}, \{c_{23}\}\}, \{p_{11}\geq 0.1, p_{11}\leq 0.7, p_{12}\geq 0.3\}, \{v_{11}\leq 0.2, v_{12}\geq 0.9\}\rangle$.
Define a risk constraint with respect to F by $S(s(\{\neg(v_{11}\leq 0.1 \;\&\; v_{12}\leq 0.8)\})) = \phi$. Now, F does not violate S, since v_{12} is always greater than 0.8.

Similarly, a decision maker could require that a consequence must possible be more probable than a worse consequence of the same alternative.

---

[19] This seems to be very demanding computationally, but it is sufficient to check the largest subset. All sets of consequences that have a possible value below the threshold are subsets of this set.

Example: Given the information frame F = $\langle\{\{c_{11},c_{12}\}, \{c_{23}\}\}, \{p_{11}\geq 0.1, p_{11}\leq 0.3, p_{12}\geq 0.7\}, \{v_{11}\geq 0.2, v_{12}\leq 0.1\}\rangle$.
Define a risk constraint with respect to F by $S(s(\{v_{ij}\leq v_{ik}\})) = s(\{p_{ij}\leq p_{ik}\})$. Now, F violates S, since c_{11} is a better consequence occurring with lower probability than c_{12}.

Also the risk distribution can be handled by imposing risk constraints. As mentioned above, in certain decision situations it seems unfair that a few people would endure the full risk.

Example: Reconsider the situations from [Linnerooth-Bayer, 1996] described at the beginning of this section. By modifying the value range to a range from [0,1000][20], the situations can be modelled in a information frame F = $\langle\{\{c_{11}, c_{12}\},\{c_{21},c_{22}\}\}\{p_{11}>0.9, p_{21}=1/1000\}\ \{v_{11}=2, v_{21}=1000\}\rangle$.[21]
If a decision maker, for instance, requires that no less than 100 persons should share the risk, she might define a risk constraint with respect to F by $S(s(\{v_{ij}\geq 100\})) = s(\{p_{ij}\leq 0.9\})$. Now, F violates S.

As argued above, quantitative methods are applicable only at information frames that do not violate the risk constraint. Thus, the first step is to eliminate consequence sets that cause the violation. The meaning of a consequence set violating a risk constraint can be stated as the following definition.

Definition: Given an information frame F = $\langle C,P,V\rangle$, and a risk constraint S with respect to F, F *violates* S iff there is an element (α,β) of S, such that not $s(V) \subseteq \alpha$ and not $s(P) \subseteq \beta$.

Definition: Given an information frame F = $\langle\{\{c_{ij}\}_{j = 1,\ldots,h_i}\}_{i = 1,\ldots m},P,V\rangle$, and a risk constraint S with respect to F, C_i *violates* S iff there is an element (α,β) of S, such that not $P^{V_i}(s(V)) \subseteq P^{V_i}(\alpha)$ and not $P^{P_i}(s(P)) \subseteq P^{P_i}(\beta)$. $P^{V_i}(x)$ is the natural projection of *x* on $\{v_{ij}\}_{j = 1,\ldots,h_i}$, and $P^{P_i}(x)$ is the natural projection of *x* on $\{p_{ij}\}_{j = 1,\ldots,h_i}$.

By using this definition it is possible to check whether there are consequence sets that do not violate the risk constraint.

Example: Given the information frame F = $\langle\{\{c_{11},c_{12},c_{13}\}, \{c_{21},c_{22}\}\}, \{p_{11}\geq 0.1, p_{11}\leq 0.5, p_{12}\geq 0.5\}\ \{v_{11}\leq 0.2, v_{12}\geq 0.1, v_{21}\geq v_{22}, v_{22}\geq 0.1\}\rangle$.
Define a risk constraint with respect to F, $S(s(\{v_{ij}\geq 0.1\})) = s(\{p_{ij}\leq 0.3\})$. Now, F violates S, but the consequence set $\{c_{21},c_{22}\}$ does not, since the value base includes the constraints $v_{21}\geq v_{22}$. Thus, $v_{21}\geq v_{22}\geq 0.1$, which assures that the values of consequences c_{21} and c_{22} both are above the threshold 0.1.

5. Conclusions and Further Research

We have extended the risk evaluation process by the integration of procedures for handling vague and numerically imprecise probabilities and utilities. The shortcomings of PMEU, and of utility theory in general, can in part be compensated for by the

[20] Naturally, we could instead have used a linear mapping from this interval on the [0,1] interval.

[21] For simplicity c_{12} and c_{22} denote all other possible consequences in the situations.

introduction of risk constraints. A body of empirical evidence has shown that many managers would welcome new ways of highlighting catastrophic consequences, as well as means to evaluating decision situations involving high risks. For example, the possibility to state threshold values, which are then used directly in the evaluation, may be appreciated. We have given a theoretical motivation for the methodology required for such evaluations involving vague and numerically imprecise data. The main contribution of this paper lies in its analysis of the concept of risk constraint violation.

To refine the analysis, contractions could also be introduced in the model. These indicate how much the different intervals in the information frame or the risk constraint can be reduced before consequence sets in the information frame cease to violate the risk constraint. In this way, it is possible to investigate critical variables and the stability of the evaluations. The inclusion of contractions in this context is a straightforward generalisation of the results in [Danielson, 1997], [Danielson and Ekenberg, 1997], and [Ekenberg, et al., 1997].

The notions discussed above will be included in DELTA [Danielson, 1997], a tool for evaluation that may use a number of different decision criteria. DELTA has been tested in several real-life applications, and we are currently preparing a case study involving catastrophic risks, and hence risk constraints.

References

[Allais, 1953] M. Allais, *Fondements d'une Théorie Positive des Choix Comportant un Risque et Critique des Postulats et Axioms de L'Ecole Americane*: D. Reidel Publishing Company., 1953.

[Bettman, 1979] J. Bettman, *An Information Processing Theory of Consumer Choice*: Addison-Wesley, 1979.

[Boman, 1995] M. Boman, "Rational Decisions and Multi-Agent Systems," Proceedings of Working Notes for the AAAI Fall Symposium on Rational Agency, MIT Technical Report, 1995.

[Boman and Ekenberg, 1995] M. Boman and L. Ekenberg, "Decision Making Agents with Relatively Unbounded Rationality," Invited paper, Proceedings of DIMAS'95, pp. I/28–I/35, 1995.

[Choquet, 1953/54] G. Choquet, "Theory of Capacities," *Ann. Inst. Fourier*, vol. 5, pp. 131_295, 1953/54.

[Danielson, 1997] M. Danielson, "Computational Decision Analysis," PhD-Thesis, Department of Computer and Systems Sciences: Royal Inst. of Technology, 1997.

[Danielson and Ekenberg, 1997] M. Danielson and L. Ekenberg, "A Framework for Analysing Decisions under Risk," *European Journal of Operational Research*, 1997.

[Dempster, 1967] A. P. Dempster, "Upper and Lower Probabilities Induced by a Multivalued Mapping," *Annals of Mathematical Statistics*, vol. xxxviii, pp. 325–339, 1967.

[Doyle et al., 1997] Doyle et al.:Eds., Qualitative Preferences in Deliberation and Practical Reasoning, Working notes, Stanford University, Stanford, California 1997.

[Ekenberg, et al., 1995] L. Ekenberg, S. Oberoi, and I. Orci, "A Cost Model for Managing Information Security Hazards," *Computers & Security*, vol. 14, pp. 707–717, 1995.

[Ekenberg, et al., 1996] L. Ekenberg, M. Danielson, and M. Boman, "From Local Assessments to Global Rationality," *International Journal of Intelligent and Cooperative Information Systems*, vol. 5, nos. 2 & 3, pp. 315–331, 1996.

[Ekenberg, et al., 1997] L. Ekenberg, M. Danielson, and M. Boman, "Imposing Security Constraints on Agent-Based Decision Support," *Decision Support Systems International Journal*, 1997.

[Ellsberg, 1961] D. Ellsberg, “Risk, Ambiguity, and the Savage Axioms,” *Quaterly Journal of Economics*, vol. 75, pp. 643–669, 1961.

[Fischhoff, et al., 1983] B. Fischhoff, B. Goitein, and Z. Shapira, “Subjective Expected Utility: A Model of Decision-Making,” in *Decision making under Uncertainty*, R. W. Scholz, Ed.: Elsevier Science Publishers B.V. (North-Holland), 1983, pp. 183–207.

[Fishburn, 1981] P. Fishburn, “Subjective Expected Utility: A Review of Normative Theories,” *Theory and Decision*, vol. 13, pp. 139–199, 1981.

[Friedman, 1953] M. Friedman, *Essays in Positive Economics*: University of Chicago Press, 1953.

[Gmytrasiewicz and Durfee, 1993] P. J. Gmytrasiewicz and E. H. Durfee, “Elements of a Utilitarian Theory of Knowledge and Action,” Proceedings of 13th IJCAI, pp. 396–402, 1993.

[Good, 1962] I. J. Good, “Subjective Probability as the Measure of a Non-measurable Set,” in *Logic, Methodology, and the Philosophy of Science*, Suppes, Nagel, and Tarski, Eds.: Stanford University Press, 1962, pp. 319–329.

[Gärdenfors and Sahlin, 1982] P. Gärdenfors and N.-E. Sahlin, “Unreliable Probabilities, Risk Taking, and Decision Making,” *Synthese*, vol. 53, pp. 361–386, 1982.

[Hernstein and Milnor, 1953] I. N. Hernstein and J. Milnor, “An Axiomatic Approach to Measurable Utility,” *Econometrica*, vol. 21, pp. 291_297, 1953.

[Huber, 1973] P. J. Huber, “The Case of Choquet Capacities in Statistics,” *Bulletin of the International Statistical Institute*, vol. 45, pp. 181–188, 1973.

[Huber and Strassen, 1973] P. J. Huber and V. Strassen, “Minimax Tests and the Neyman-Pearsons Lemma for Capacities,” *Annals of Statistics*, vol. 1, pp. 251–263, 1973.

[ISO, 1996] ISO, “Managing Catastrophe Risk,” Insurance Services Office, Inc., ISO Insurance Issues Series 1996.

[Jeffrey, 1983] R. Jeffrey, *The Logic of decision*: University of Chicago Press, 1983.

[Kahneman and Lovallo, 1993] D. Kahneman and D. Lovallo, “Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking,” *Management Science*, vol. 39, pp. 17–31, 1993.

[Kunreuther, et al., 1978] H. Kunreuther, R. Ginsberg, L. Miller, P. Slovic, B. Borkan, and N. Katz, *Disaster Insurance Protection: Public Policy Lessons*: John Wiley and Sons, 1978.

[Kyburg, 1961] H. E. Kyburg, *Probability and the Logic of Rational Belief*: Wesleyan University Press, 1961.

[Lehmann, 1959] E. L. Lehmann, *Testing Statistical Hypothesis*: John Wiley and Sons, 1959.

[Levi, 1974] I. Levi, “On Indeterminate Probabilities,” *The Journal of Philosophy*, vol. 71, pp. 391–418, 1974.

[Levi, 1992] I. Levi, “Feasibility,” in *Knowledge, Belief, And Strategic Interaction*, Bicchieri and D. Chiara, Eds.: Cambridge University Press, 1992, pp. 1–20.

[Linnerooth-Bayer, 1996] J. Linnerooth-Bayer, “Does Society Mismanage Risk?,” in *Wise Choices: Decisions, Games, and Negotiations*, R. J. Zeckhauser, R. L. Keeney, and J. K. Sebenius, Eds. Boston, Massachusetts: Harward Business School Press, 1996.

[Loomes and Sudgen, 1982] G. Loomes and R. Sudgen, “Regret Theory: an Alternative Theory of Rational Choice under Uncertainty,” *The Economic Journal*, vol. 92, pp. 805–824, 1982.

[Luce and Krantz, 1971] R. D. Luce and D. Krantz, “Conditional Expected Utility,” *Econometrica*, vol. 39, pp. 253–271, 1971.

[Malmnäs, 1990a] P. E. Malmnäs, Axiomatic Justification of the Utility Principle, HSFR 677/87, 1990.

[Malmnäs, 1990b] P. E. Malmnäs, Real-Life Decisions, Expected Utility and Effective Computability, HSFR 677/87, 1990.

[Malmnäs, 1994a] P. E. Malmnäs, “Axiomatic Justification of the Utility Principle,” *Synthese*, vol. 99, pp. 233–249, 1994.

[Malmnäs, 1994b] P. E. Malmnäs, “Towards a Mechanization of Real Life Decisions,” in *Logic and Philosophy of Science in Uppsala*, Prawitz and Westerståhl, Eds.: Kluwer Academic Publishers, 1994.

[Malmnäs, 1996] P. E. Malmnäs, “Evaluations, Preferences, Choice Rules,” Research Report, Dept. of Philosophy, Stockholm University, 1996.

[Neumann and Morgenstern, 1947] J. v. Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, 2 ed: Princeton University Press, 1947.

[Nilsson, 1986] N. Nilsson, “Probabilistic Logic,” *Artificial Intelligence*, vol. 28, pp. 71–87, 1986.

[Oddie and Milne, 1990] G. Oddie and P. Milne, “Act and Value,” *Theoria*, vol. LVII, pp. 42–76, 1990.

[Quiggin, 1982] J. Quiggin, “A Theory of Anticipated Utility,” *Journal of Economic Behavior and Organization*, vol. 3, pp. 323_343, 1982.

[Ramsey, 1978] F. P. Ramsey, “Truth and Probability,” in *Foundations: Essays in Philosophy, Logics, Mathematics and Economics*, Mellor, Ed.: Routledge & Kegan Paul, 1978, pp. 58–100.

[Savage, 1972] L. Savage, *The Foundations of Statistics*, 2 ed: Dover, 1972.

[Schoemaker, 1982] P. Schoemaker, “The Expected Utility Model: Its Variants, Purposes, Evidence and Limitations,” *Journal of Economic Literature*, vol. XX, pp. 529–563, 1982.

[Shafer, 1976] G. Shafer, *A Mathematical Theory of Evidence*: Princeton University Press, 1976.

[Shapira, 1995] Z. Shapira, *Risk Taking: A Managerial Perspective*: Russel Sage Foundation, 1995.

[Smith, 1961] C. A. B. Smith, “Consistency in Statistical Inference and Decision,” *Journal of the Royal Statistic Society, Series B*, vol. xxiii, pp. 1–25, 1961.

[Suppes, 1956] P. Suppes, “The Role of Subjective Probability and Utility Maximization,” Proceedings of Third Berkely Symposium on Mathematical Statistics and Probability 1954-55, pp. 113–134, 1956.