Second-Order Risk Constraints

Love Ekenberg[1,2], Aron Larsson[2] and **Mats Danielson[1]**
[1] Dept. of Computer and Systems Sciences, Stockholm University and Royal Institute of Technology
Forum 100, SE-164 40 Kista, Sweden
[2] Dept. of Information Technology and Media, Mid Sweden University
SE-851 70 Sundsvall, Sweden

Abstract

This paper discusses how numerically imprecise information can be modelled and how a risk evaluation process can be elaborated by integrating procedures for numerically imprecise probabilities and utilities. More recently, representations and methods for stating and analysing probabilities and values (utilities) with belief distributions over them (second order representations) have been suggested. In this paper, we are discussing some shortcomings in the use of the principle of maximising the expected utility and of utility theory in general, and offer remedies by the introduction of supplementary decision rules based on a concept of risk constraints taking advantage of second-order distributions.

Introduction

The equating of substantial rationality with the principle of maximising the expected utility (PMEU) is inspired by early efforts in decision theory made by Ramsey, von Neumann, Savage and others. They structured a comprehensive theory of rational choice by proposing reasonable principles in the form of axiom systems justifying the utility principle. Such axiomatic systems usually consist of primitives (such as an ordering relation, states, sets of states, etc.) and axioms constructed from the primitives. The axioms (ordering axioms, independence axioms, continuity axioms, etc.) imply numerical representations of preferences and probabilities. Typically implied by the axioms are existence theorems stating that a utility function exists, and a uniqueness theorem stating that two utility functions, relative to a given preference ranking, are always affine transformations of each other. It is often argued that these results provide justification of PMEU.

However, this viewpoint has been criticised and a common counter-argument is that the axioms of utility theory are fallacious. There is a problem with the formal justifications of the principle in that even if the axioms in the various axiomatic systems are accepted, the principle itself does not follow, i.e. the proposed systems are too weak to imply the utility principle (Malmnäs 1994). Thus, it is doubtful whether this principle can be justified on purely formal grounds and the logical foundations of utility theory seem to be weak. For instance, within the AI agent area, the details of utility-based agent behaviour are usually not formalised, a common explanation being that there are several adequate axiomatisations from which the choice is a matter of taste. In this paper, the generic terms agent and decision-maker are used interchangeably, and include artificial (software) as well as human entities unless otherwise noted.

Critics point out that most mathematical models of rational choice are oversimplified and disregard important factors. For instance, the use of a utility function for capturing all possible risk attitudes is not considered possible (Schoemaker 1982). It has also been shown that people do not act in accordance with certain independence axioms in the system of Savage (Allais 1979). Although descriptive research of this kind cannot overthrow the normative aspects of the system, it shows that there is a need to include other types of functions that can model different types of behaviour in risky situations.

Some researchers have tried to modify the application 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, e.g., (Loomes and Sudgen 1982). Others have tried to resolve the problems mentioned above by having functions modifying both the probabilities and the utilities. But their performances are at best equal to that of the expected value, and at worst inferior, e.g., inconsistent with first-order stochastic dominance (Malmnäs 1996).

Furthermore, the elicitation of risk attitudes from human decision-makers is error prone and the result is highly dependent on the format and method used, see, e.g., (Riabacke, Påhlman, and Larsson 2006). This problem is even more evident when the decision situation involve catastrophic outcomes (Mason et al. 2005). If not being able to elicit a properly reflecting risk attitude, we may have the situation that even if the evaluation of an alternative results in an acceptable expected utility, some consequences might be of a catastrophic kind so the alternative should be avoided in any case. Due to catastrophe aversion, this may be the case even if the probabilities of these consequences are very low. In such cases, the PMEU needs to be extended with other rules, and it has therefore been argued that a useful decision theory should permit a wider spectrum of risk attitudes than by means of a utility function only. A more pragmatic approach should give an agent the means for expressing risk attitudes in a variety of ways, as well as provide procedures for handling both qualitative and quantitative aspects.

We will now take a closer look at how some of these deficiencies can be remedied. The next section introduces a decision tree formalism and corresponding risk constraints. They are followed by a brief description of a theory for representing imprecision using second-order distributions. The last section before the conclusion presents the main contribution of this paper – how risk constraints can be realised in a second-order framework for evaluating decisions under risk. This include a generalisation of risk constraints in a second-order setting and obtaining a reasonable measure of the support for violation of stipulated constraints for each decision alternative.

Modelling the Decision Problem

In this paper, we let an *information frame* represent a decision problem. The idea with such a frame is to collect all information necessary for the model into one structure. Further, the representational issues are of two kinds; a *decision structure*, modelled by means of a decision tree, and *input statements*, modelled by means of linear constraints. A decision tree is a graph structure $\langle V, E \rangle$ where V is a set of nodes and E is a set of node pairs (edges).

Definition 1. *A tree is a connected graph without cycles. A decision tree is a tree containing a finite set of nodes which has a dedicated node at level* 0*. The adjacent nodes, except for the nodes at level* $i-1$*, to a node at level* i *is at level* $i+1$*. A node at level* $i+1$ *that is adjacent to a node at level* i *is a child of the latter. A node at level* 1 *is an* ***alternative****. A node at level* i *is a leaf or* ***consequence*** *if it has no adjacent nodes at level* $i+1$*. A node that is at level* 2 *or more and has children is an* ***event*** *(an intermediary node). The depth of a rooted tree is* $\max(n|$*there exists a node at level n*$)$.

Thus, a decision tree is a way of modelling a decision situation where the alternatives are nodes at level 1 and the set of final consequences are the set of nodes without children. Intermediary nodes are called events. For convenience we can, for instance, use the notation that the n children of a node x_i are denoted $x_{i1}, x_{i2}, \ldots, x_{in}$ and the m children of the node x_{ij} are denoted $x_{ij1}, x_{ij2}, \ldots, x_{ijm}$ and so forth. For presentational purposes, we will denote a consequence node of an alternative A_i simply with c_{ij}.

Over each set of event node children and consequence nodes, functions can be defined, such as probability distributions and utility functions.

Interval Statements

For numerically imprecise decision situations, one option is to define probability distributions and utility functions in the classical way. Another, more elaborate option is to define sets of candidates of possible probability distributions and utility functions and then express these as vectors in polytopes that are solution sets to, so called, *probability* and *utility bases*.

For instance, the probability (or utility) of c_{ij} being between the numbers a_k and b_k is expressed as $p_{ij} \in [a_k, b_k]$ (or $u_{ij} \in [a_k, b_k]$). Such an approach also includes relations – a measure (or function) of c_{ij} is greater than a measure (or function) of c_{kl} is expressed as $p_{ij} \geq p_{kl}$ and analogously $u_{ij} \geq u_{kl}$. Each statement can thus be represented by one or more constraints.

Definition 2. *Given a decision tree* T*, a* ***utility base*** *is a set of linear constraints of the types* $u_{ij} \in [a_k, b_k]$*,* $u_{ij} \geq u_{kl}$ *and, for all consequences* $\{c_{ij}\}$ *in* T*,* $u_{ij} \in [0, 1]$*. A* ***probability base*** *has the same structure, but, for all intermediate nodes* N *(except the root node) in* T*, also includes* $\sum_{j=1}^{m_N} p_{ij} = 1$ *for the children* $\{x_{ij}\}_{j=1,\ldots,m_N}$ *of* N*.*

The solution sets to probability and utility bases are polytopes in hypercubes. Since a vector in the polytope can be considered to represent a distribution, a probability base $\mathcal{P}$ can be interpreted as constraints defining the set of all possible probability measures over the consequences. Similarly, a utility base $\mathcal{U}$ consists of constraints defining the set of all possible utility functions over the consequences. The bases $\mathcal{P}$ and $\mathcal{U}$ together with the decision tree constitute the *information frame* $\langle T, \mathcal{P}, \mathcal{U} \rangle$.

As discussed above, the most common evaluation rules of a decision tree model are based on the PMEU.

Definition 3. *Given an information frame* $\langle T, \mathcal{P}, \mathcal{U} \rangle$ *and an alternative* $A_i \in A$ *the expression*

$$\mathrm{E}(A_i) = \sum_{i_1=1}^{n_{i_0}} p_{ii_1} \sum_{i_2=1}^{n_{i_1}} p_{ii_1i_2} \cdots \sum_{i_{m-1}=1}^{n_{i_{m-2}}} p_{ii_1i_2\ldots i_{m-2}i_{m-1}} \sum_{i_m=1}^{n_{i_{m-1}}} p_{ii_1i_2\ldots i_{m-2}i_{m-1}i_m} u_{ii_1i_2\ldots i_{m-2}i_{m-1}i_m}$$

where m *is the depth of the tree corresponding to* A_i*,* n_{i_k} *is the number of possible outcomes following the event with probability* p_{i_k}*,* $p_{\ldots i_j \ldots}$*,* $j \in [1, \ldots, m]$*, denote probability variables and* $u_{\ldots i_j \ldots}$ *denote utility variables as above, is the expected utility of alternative* A_i *in* $\langle T, \mathcal{P}, \mathcal{U} \rangle$*.*

The alternatives in the tree are evaluated according to PMEU, and the resulting expected utility defines a (partial) ordering of the alternatives. However, as discussed in the introduction, the use of utility functions to formalise the decision process seem to be an oversimplified idea, disregarding important factors that appear in real-life applications of decision analysis. Therefore, there is a need to permit the use of additional ways to discriminate between alternatives. The next section discusses risk constraints as such a complementary decision rule.

Risk Constraints

The intuition behind risk constraints is that they express when an alternative is undesirable due to too risky consequences. A general approach is to introduce the constraints to provide thresholds beyond which an alternative is deemed undesirable by the decision making agent. Thus, expressing risk constraints is analogous to expressing minimum requirements that should be fulfilled in the sense that a risk constraint can be viewed as a function stating a set of thresholds that may not be violated in order for an alternative to be acceptable with respect to risk (Danielson 2005).

A decision agent might regard an alternative as undesirable if it has consequences with too low a utility and with some probability of occurring, regardless of its contribution to the expected utility being low. Additionally, if several consequences of an alternative A_i are too bad (with respect to a certain utility threshold), the probability of their union must be considered even if their individual probabilities are not high enough by themselves to render the alternative unacceptable. This procedure is fairly straightforward. For an alternative A_i in an information frame $\langle T, \mathcal{P}, \mathcal{U} \rangle$, given a utility threshold r' and a probability threshold s', then

$$\sum_{u_{ij} \leq r'} p_{ij} \leq s'$$

must hold in order for A_i to be deemed an acceptable alternative. In this sense, a risk constraint can be considered a utility-probability pair (r', s'). Then a consequence c_{ij} is violating r' if $u_{ij} > r'$ does not hold. Principles of this kind seem to be good prima facie candidates for evaluative principles in the literature, i.e., they conform well to established practices and enable a decision-maker to use qualitative assessments in a reasonable way. For a comprehensive treatment and discussion, see (Ekenberg, Danielson, and Boman 1997).

However, when the information is numerically imprecise (probabilites and utilities are expressed as bounds or intervals), it is not obvious how to interpret such thresholds. We have earlier suggested that the interval boundaries together with stability analyses could be considered in such cases (Ekenberg, Boman, and Linneroth-Bayer 2001).

Example 1. *An alternative A_i is considered undesirable if the consequence c_{ij} belonging to A_i has a possibility that the utility of c_{ij} is less than 0.45, and if the probability of c_{ij} is greater than 0.65. Assume that the alternative A_i has a consequence for which its utility lies in the interval $[0.40, 0.60]$. Further assume that the probability of this consequence lies in the interval $[0.20, 0.70]$. Since 0.45 is greater than the least possible utility of the consequence, and 0.65 is less than the greatest possible probability, A_i violates the thresholds and is thus undesirable.*

The stability of such a result should also be investigated. For instance, it can be seen that the alternative in Example 1 ceases to be undesirable when the left end-point of the utility interval is increased by 0.05. An agent might nevertheless be inclined to accept the alternative since the constraints are violated in a small enough proportion of the possible values. Thus, the analysis must be refined.

A concept in line with such stability analyses is the concept of *interval contraction*, investigating to what extent the widths of the input intervals need be reduced in order for an alternative not to violate the risk constraints. The contractions of intervals are done toward a contraction point for each interval. Contraction points can either be given explicitly by the decision making agent or be suggested from, e.g., minimum distance calculations or centre of mass calculations. The level of contraction is indicated as a percentage,

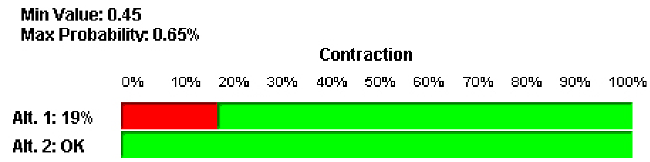

Figure 1: Contraction analysis of risk constraints given in Example 1. Beyond a contraction level of 19%, the constraints are no longer violated for alternative A_1. The constraints for alternative A_2 are never violated.

where at 100% contraction all intervals have been replaced with their contraction points, see Figure 1 for a contraction analysis of the rudimentary problem in Example 1. One refinement is to provide a possibility for an agent to stipulate thresholds for *proportions* of the probability and utility bases, i.e. an alternative is considered unacceptable if it violates the risk constraints at a given contraction level (Danielson 2005).

Including Second-Order Information

The evaluation procedures of interval decision trees yield first-order (interval) estimates of the evaluations, i.e. upper and lower bounds for the expected utilities of the alternatives. An advantage of approaches using upper and lower probabilities is that they do not require taking particular probability distributions into consideration. On the other hand, the expected utility range resulting from an evaluation is also an interval. To our experience, in real-world decision situations it is then often hard to discriminate between the alternatives since the intervals are not always narrow enough. For instance, an interval based decision procedure keeps all alternatives with overlapping expected utility intervals, even if the overlap is small. Therefore, it is worthwhile to extend the representation of the decision situation using more information, such as second-order distributions over classes of probability and utility measures.

Distributions can be used for expressing various beliefs over multi-dimensional spaces where each dimension corresponds to, for instance, possible probabilities or utilities of consequences. The distributions can consequently be used to express strengths of beliefs in different vectors in the polytopes. Beliefs of such kinds are expressed using higher-order distributions, sometimes called hierarchical models. Approaches for extending the interval representation using distributions over classes of probability and value measures have been developed into various such models, for instance second-order probability theory. In the following, we will pursue the idea of adding more information and discuss its implications on risk constraints.

Distributions over Information Frames

Interval estimates and relations can be considered as special cases of representations based on distributions over polytopes. For instance, a distribution can be defined to have a positive support only for $x_i \leq x_j$. More formally, the solution set to a probability or utility base is a subset of a unit cube since both variable sets have $[0,1]$ as their ranges. This subset can be represented by the support of a distribution over the cube.

Definition 4. *Let a unit cube $[0,1]^n$ be represented by $B = (b_1,\ldots,b_n)$. The b_i can be explicitly written out to make the labelling of the dimensions clearer.*

More rigorously, the unit cube is represented by all the tuples $(x_1,\ldots,x_n)$ in $[0,1]^n$.

Definition 5. *By a second-order distribution over a cube B, we denote a positive distribution F defined on the unit cube B such that*

$$\int_B F(x)\,dV_B(x) = 1\,,$$

where V_B is the n-dimensional Lebesque measure on B. The set of all second-order distributions over B is denoted by $BD(B)$.

For our purposes here, second-order *probabilities* are an important sub-class of these distributions and will be used below as a measure of belief, i.e. a second-order joint probability distribution. Marginal distributions are obtained from the joint ones in the usual way.

Definition 6. *Let a unit cube $B = (b_1,\ldots,b_n)$ and $F \in BD(B)$ be given. Furthermore, let $B_i^- = (b_1,\ldots,b_{i-1},b_{i+1},\ldots,b_n)$. Then*

$$f_i(x_i) = \int_{B_i^-} F(x)\,dV_{B_i^-}(x)$$

is a marginal distribution over the axis b_i.

A marginal distribution is a special case of an S-projection,

Definition 7. *Let $B = (b_1,\ldots,b_k)$ and $A = (b_{i_1},\ldots,b_{i_s}), i_j \in \{1,\ldots,k\}$ be unit cubes. Let $F \in BD(B)$, and let*

$$F_A(x) = \int_{B\setminus A} F(x)\,dV_{B\setminus A}(x)$$

Then F_A is the S-projection of F on A.

An S-projection of the above kind is also a second-order distribution (Ekenberg and Thorbiörnson 2001). As an information frame has two separated constraint sets, $\mathcal{P}$ holding constraints on probability variables and $\mathcal{U}$ holding constraints on utility variables, it is suitable to distinguish between cubes in the same fashion. A unit cube holding probability variables is denoted by B_P and a unit cube holding utility variables is denoted by B_U.

Example 2. *Given an information frame $\langle T,\mathcal{P},\mathcal{U}\rangle$, constraints in the bases can be defined through a belief distribution. Given a unit cube $U = (u_1,u_2)$ and a distribution G over U defined by $G(u_1,u_2) = 6\cdot\max(u_1-u_2,0)$. Then G is a second-order (belief) distribution in our sense, and the support of G is $\{(u_1,u_2)|0 \le u_i \le 1 \wedge u_1 > u_2\}$. See Figure 2.*

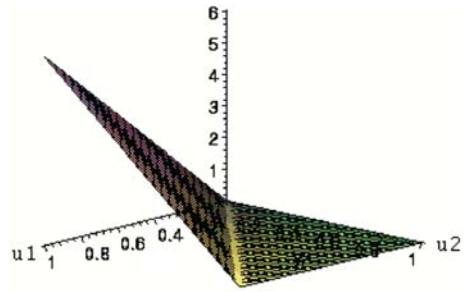

Figure 2: The support of $G(u_1,u_2)$ is the solution set of the set $\{1 \ge u_1 > u_2 \ge 0\}$ of constraints.

As an analysis using risk constraints is done investigating one alternative at a time, we let a utility cube with respect to an alternative A_i be denoted by B_{U_i} and a probability unit cube with respect to A_i be denoted by B_{P_i}. Hence, B_{U_i} is represented by all the tuples $(u_{i1},\ldots,u_{in})$ in $[0,1]^n$ and B_{P_i} is represented by all the tuples $(p_{i1},\ldots,p_{in})$ in $[0,1]^n$ when A_i has n consequences. The normalisation constraint for probabilities imply that for a belief distribution over B_{P_i} there can be positive support only for tuples where $\sum p_{ij} = 1$.

Definition 8. *A probability unit cube for alternative A_i is a unit cube $B_{P_i} = (p_{i1},\ldots,p_{in})$ where $F_i(p_{i1},\ldots,p_{in}) > 0 \Rightarrow \sum_{j=1}^n p_{ij} = 1$. A utility unit cube for A_i, B_{U_i}, lacks this latter normalisation.*

One candidate for serving as a belief distribution over B_{P_i} is the Dirichlet distribution.

Example 3. *The marginal distribution $f_{i1}(p_{i1})$ of the uniform Dirichlet distribution in a 4-dimensional cube is*

$$f_{i1}(p_{i1}) = \int_0^{1-p_{i1}} \int_0^{1-p_{i2}-p_{i1}} 6\,dp_{i3}\,dp_{i2} = 3(1-2p_{i1}+p_{i1}^2)$$

$$= 3(1-p_{i1})^2\,.$$

Evaluation of decision trees with respect to PMEU using second-order distributions is discussed in (Ekenberg et al. 2007). The result is a method that can offer more discriminative power in selecting alternatives where overlap prevails, as the method may compare expected utility sub-intervals where the second-order belief mass is kept under control. With respect to the input statements of this model, there are similarities with the additional input required for conducting probabilistic sensitivity analyses, which aims at an analysis of post hoc robustness, see, e.g., (Felli and Hazen 1998). However, the primary concern herein is to take such input into account already in the evaluation rules. The next section discusses how this may be done for risk constraints.

Second-Order Risk Constraints

The generalisation of risk constraints in second-order decision analysis is rather straightforward. The basic idea is to consider the actual proportions of the resulting distributions that the thresholds cut off.

In the following, let $\langle T, \mathcal{P}, \mathcal{U} \rangle$ be an information frame. A prima facie solution is then to let $f_{ij}(p_{ij})$ and $g_{ij}(u_{ij})$ be marginal second-order distributions over the probabilities and utilities of a consequence c_{ij} in the frame. Then, given thresholds r' and s' and second-order thresholds r'' and s'', where $s', r', s'', r'' \in [0, 1]$, if

$$\int_0^{r'} g_{ij}(u_{ij})\, du_{ij} \geq r''$$

and

$$\int_{s'}^{1} f_{ij}(p_{ij})\, dp_{ij} \geq s''$$

holds the alternative is deemed undesirable. Note that r' and s' are limits on actual utilities and probabilities respectively but r'' and s'' are limits on their distributions.

However, as for ordinary risk constraints, it is also necessary to take into account the way in which subsets of consequences, i.e. events, together can make an alternative undesirable. If we would have independent distributions in the probability base, this would be accomplished by using standard convolution, utilizing the product rule for standard probabilities. Due to normalization and possible inequality constraints, this approach must be modified.

Let $\{g_{ij}(u_{ij})\}_{j=1}^{n}$ be marginal second-order distributions with respect to consequences $\{c_{ij}\}$ of an alternative A_i in an information frame $\langle T, \mathcal{P}, \mathcal{U} \rangle$. Let Φ_i be the consequence set such that

$$c_{ij} \in \Phi_i \iff \int_0^{r'} g_{ij}(u_{ij})\, du_{ij} \ \geq r''$$

Further, let P_i be the set of possible (joint) probability distributions $(p_{i1} \ldots, p_{in})$ over the consequences of an alternative A_i, let F_i be a belief distribution over P_i, and let

$$t'' = \int_{\Gamma_{s'}} F_i(p_{i1}, \ldots, p_{in}) dV_{B_{P_i}}$$

where

$$\Gamma_{s'} = \left\{ P_i : \sum_{c_{ij_k} \in \Phi_i} p_{ij_k} \geq s' \right\}$$

Then the inequality

$$t'' \leq s'' \tag{1}$$

must hold for the alternative to be acceptable. This is a straightforward generalisation of the risk constraint concept utilising second-order information. In addition to the utility-probability threshold pair (r', s'), we also use a pair (r'', s'') acting as thresholds on the belief mass violating r' and s' respectively.

Belief in Risk Constraint Violation

Given the proportions that the risk constraints specify, we can derive a measure $\tau_i \in [0, 1]$ of to what extent the input statements support a violation of a risk constraint (r', s') for a given alternative A_i. The rationale behind such a measure is that it delivers further information to a decision-maker when more than one alternative violate stipulated risk constraints. This is especially important for cases when only some consistent probability-utility assignments (i.e. subsets of the polytopes) violate the risk constraints.

If an alternative do not, for any consistent probabilities or utilities in the information frame, violate the risk constraint, this yields a violation belief measure of zero. On the other hand, if all consistent probabilities and utilities violate the risk constraint, a violation belief of one is obtained.

For such a measure to be meaningful, it should as a *minimum* requirement fulfil the following desiderata. In the following, $\tau_{(i,r',s')}$ denote the violation belief of a risk constraint (r', s') for an alternative A_i.

Desideratum 1. *Given an information frame with an alternative A_i and risk constraints $(r'_1, s'), (r'_2, s')$. Then $r'_1 > r'_2 \Rightarrow \tau_{(i,r'_1,s')} \geq \tau_{(i,r'_2,s')}$.*

Desideratum 2. *Given an information frame with an alternative A_i and risk constraints $(r', s'_1), (r', s'_2)$. Then $s'_1 < s'_2 \Rightarrow \tau_{(i,r',s'_1)} \geq \tau_{(i,r',s'_2)}$.*

Desideratum 3. *Given an information frame with an alternative A_i and a risk constraint (r', s') and let k be a consequence index c_{ik}. Let $I \neq \emptyset$ be the index set of consequences violating r', yielding $\tau_{(i,r',s')}$ when $k \notin I$. If the information frame is modified only with respect to the utility u_{ik} leading to $k \in I$ yielding $\tau^*_{(i,r',s')}$, then $\tau^*_{(i,r',s')} > \tau_{(i,r',s')}$.*

In essence, Desiderata 1-2 say that given an information frame, more demanding risk constraints should not yield lower belief in their violation, and Desideratum 3 says that we wish to take into account the way in which subsets of consequences together can make an alternative undesirable.

One proposal is to select the resulting value of the integral on the left hand side of inequality (1) as a measure of violation belief. Although this would fulfil the minimum requirements stipulated in Desiderata 1-3, one would need to choose a second-order threshold r'' and the result would be sensitive with respect to this assignment. Another disadvantage with this approach is that it would discriminate between smaller and larger violations of r''. However, since this technique operates on the marginals $g_{ij}(u_{ij})$, it might be preferred due to its intuitive appeal. Another proposal is given below, operating on global belief distributions and not utilising second-order thresholds.

Define $B_R = B_{P_i} \times B_{U_i}$, existing of all tuples (p, u), i.e. $(p_{i1}, u_{i1} \ldots, p_{in}, u_{in})$. Let F_i be a belief distribution on B_{P_i} and G_i be a belief distribution on B_{U_i}, then it follows that

$$\int_{B_R} F_i(p) \cdot G_i(u)\, dV_{B_R}(p, u) = 1 \tag{2}$$

See, e.g., (Danielson, Ekenberg, and Larsson 2007).

Definition 9. *Given an information frame, the violation belief τ_i of A_i violating (r', s') is*

$$\tau_i = \int_{\mathcal{R}} F_i(p) \cdot G_i(u)\, dV_{\mathcal{R}}(p, u)$$

where $\mathcal{R}$ is the set of points $(p_{i1}, u_{i1}, \ldots, p_{in}, u_{in}) \in B_R$, such that $\sum_{j \in K} p_{ij} > s'$ where K is the index set given by $u_{ij} \leq r' \Leftrightarrow j \in K$.

Proposition 1. *$\tau_i \in [0, 1]$ and fulfils Desiderata 1-3.*

Proof. Since $\mathcal{R} \subseteq B_R$ then $\tau_i \in [0, 1]$. When distributions F_i, G_i are given (the information frame is given) the violation belief τ_i depends only on the proportion of $\mathcal{R}$ relative to B_R. If no consistent probability and utility assignments can violate (r', s'), then $\mathcal{R} = \emptyset$ yields $\tau_i = 0$. If all consistent probability and utility assignents violate (r', s'), then $\mathcal{R} = B_R$ yield $\tau_i = 1$ from (2). For Desideratum 1, we have risk constraints (r'_1, s') and (r'_2, s'). As $\mathcal{R}$ is bounded above by $u_{ij} \leq r'$, $r'_1 > r'_2$ cannot result in a lower proportion for r'_1 than for r'_2. Hence, Desideratum 1 is fulfilled. For Desideratum 2, essentially the same reasoning applies. For Desideratum 3, let $\mathcal{R}$ denote the domain when $k \notin I$, and denote it by $\mathcal{R}'$ when $k \in I$. Then since $\mathcal{R} \subset \mathcal{R}'$ and $\mathcal{R}' \setminus \mathcal{R} \neq \emptyset$ is a convex subset of B_R Desideratum 3 is satisfied. □

Summary and Conclusions

The most often used decision rules in formal models of decision making are based on the principle of maximising the expected utility. However, the various axiomatic theories proposed to support this principle are insufficient and have been subject to severe criticism. Therefore, it seems reasonable to supplement frameworks based on the utility principle with other decision rules taking a wider spectrum of risk attitudes into account. One such supplement is the inclusion of thresholds in the form of risk constraints.

This paper discusses how numerically imprecise information can be modelled and evaluated, as well as how the risk evaluation process can be elaborated by integrating procedures for handling vague and numerically imprecise probabilities and utilities. The shortcomings of the principle of maximising the expected utility, and of utility theory in general, can in part be compensated for by the introduction of the concept of risk constraint violation. It should be emphasised that this is not the only method of comparing the risk involved in different alternatives in imprecise domains. However, it is based on a well-founded model of imprecision and meeting reasonable requirements on its properties.

Using risk constraint violation, a general model can be constructed for representing various risk attitudes and providing alternative means for expressing such. The definitions are computationally meaningful, and are therefore also well suited to automated decision making. Rules have been suggested for sorting out undesirable decision alternatives, rules which should also serve as a tool for guaranteeing that certain norms are not violated, even when it is desirable (or necessary) that the agents be able to maintain their autonomy.

References

Allais, M. 1979. The Foundations of a Positive Theory of Choice involving Risk and a Criticism of the Postulates and Axioms of the American School. In Allais, M., and Hagen, O. eds. *Expected Utility Hypotheses and the Allais Paradox*. D. Reidel Publishing Company.

Danielson, M. 2005. Generalized Evaluation in Decision Analysis. *European Journal of Operational Research* 162(2): 442–449.

Danielson, M. and Ekenberg, L. 2007. Computing Upper and Lower Bounds in Interval Decision Trees. *European Journal of Operational Research* 181(2): 808–816.

Danielson, M., Ekenberg, L., and Larsson, A. 2007. Distribution of Expected Utility in Decision Trees. *International Journal of Approximate Reasoning* 46(2): 387–407.

Ekenberg, L., Danielson, M., and Boman, M. 1997. Imposing Security Constraints on Agent-based Decision Support. *Decision Support Systems* 20(1): 3–15.

Ekenberg, L., Boman, M., and Linneroth-Bayer, J. 2001. General Risk Constraints. *Journal of Risk Research* 4(1): 31–47.

Ekenberg, L. and Thorbiörnson, J. 2001. Second-Order Decision Analysis. *International Journal of Uncertainty, Fuzziness, and Knowledge-Based Systems* 9(1):13–38.

Ekenberg, L., Andersson, A., Danielson, M., and Larsson, A. 2007. Distributions over Expected Utilities in Decision Analysis. *Proceedings of the 5th International Symposium on Imprecise Probabilities and their Applications*: 175–182.

Felli, J. C. and Hazen, G. B. 1998. Sensitivity Analysis and the Expected Value of Perfect Information. *Medical Decision Making* 18(1): 95–109.

Loomes, G. and Sugden, R. 1982. Regret Theory: An Alternative Theory of Rational Choice under Uncertainty. *The Economic Journal* 92: 805–824.

Malmnäs, P.-E. 1994. Axiomatic Justifications of the Utility Principle: A Formal Investigation. *Synthese* 99: 233–249.

Malmnäs, P.-E. 1996. Evaluations, Preferences, Choice Rules. Research Report, Dept. of Philosophy, Stockholm University.

Mason, C. F., Shogren, J., Settle, C., and List, A. J. 2005. Environmental Catastrophes and Non-Expected Utility Maximization: An Experimental Evaluation. *Journal of Risk and Uncertainty* 31(2): 187–215.

Riabacke, A., Påhlman, M., and Larsson, A. 2006. How Different Choice Strategies Can Affect the Risk Elicitation Process. *IAENG International Journal of Computer Science* 32(4): 460–465.

Schoemaker, P. J. H. 1982. The Expected Utility Model: Its Variants, Purposes, Evidence and Limitations. *Journal of Economic Literature* 20: 529–563.