An Argumentation Framework for Deriving Qualitative Risk Sensitive Preferences

Wietske Visser, Koen V. Hindriks, and Catholijn M. Jonker

Man-Machine Interaction Group, Delft University of Technology, The Netherlands
Wietske.Visser@tudelft.nl, K.V.Hindriks@tudelft.nl, C.M.Jonker@tudelft.nl

Abstract. Preferences are derived in part from knowledge. Knowledge, however, may be defeasible. We present an argumentation framework for deriving qualitative, multi-attribute preferences and incorporate defeasible reasoning about knowledge. Intuitively, preferences based on defeasible conclusions are not as strong as preferences based on certain conclusions, since defeasible conclusions may turn out not to hold. This introduces risk when such knowledge is used in practical reasoning. Typically, a risk prone attitude will result in different preferences than a risk averse attitude. In this paper we introduce qualitative strategies for deriving risk sensitive preferences.

1 Introduction

In [1], an argumentation framework for reasoning about qualitative preferences is presented. It introduces a safe and decisive strategy to deal with incomplete information. However, this framework is not able to deal with defeasible information. In this paper we focus on defeasible rather than incomplete information. Adding the means to derive defeasible objective information is quite straightforward, by implementing some kind of defeasible modus ponens. However, deriving preferences from a mixture of certain and defeasible information is not. For this we have to take into account the notion of risk, since with defeasible information there is always a chance that it is incorrect. Multi-attribute utility theory [2], a quantitative approach to preferences, defines two risk attitudes: risk aversion and risk proneness. In this paper we present qualitative strategies for deriving risk sensitive preferences. They are incorporated into an argumentation framework that provides the means to reason about preferences as well as underlying factual information. Argumentation, being a kind of defeasible reasoning, is able to reason with incomplete, uncertain and contradictory information. Besides, because reasoning by means of arguments is a human type of reasoning, argumentation is suitable for explanation of a system's reasoning to a human user.

The topic is related to decision making under uncertainty (e.g. [3]). In DMU, the aim is to find the best decision in case of uncertainty about the current state of the world, and hence about the outcomes of decisions. Our approach is more general and can be applied in different contexts; we compare the preference between abstract 'objects', which could be states of the world, but also e.g. products. One of the challenges of reasoning about preferences is their multi-attribute nature. There are several distinct notions: importance of attributes, degree of satisfaction of attributes, and degree of belief of facts.

In some approaches, these measures are assumed to be commensurate, others (including this paper) suppose non-commensurability. We focus on the case where it is not completely certain which attributes the objects have (there are different degrees of belief), combined with relative importance of attributes. We leave the degree of satisfaction of attributes for future work. [4] present several multi-attribute preference ordering rules, but do not take uncertainty into account. [5] present a qualitative model for decision making with plausibility measures of input situations, but they treat plausible and likely beliefs equally. [6] present an argument-based approach to multi-attribute preferences that does take degree of belief into account, but it is a two-step process in which argumentation is used only for epistemic reasoning. In our approach, we combine reasoning about preferences and knowledge in a single argumentation framework.

In Section 2, we briefly introduce qualititive multi-attribute preferences, and discuss defeasible reasoning and its relation to risk. Section 3 presents a generic framework that provides the means to take the risk introduced by defeasible knowledge into account and introduces qualitative risk-sensitive preferences. Section 4 presents concrete, qualitative preference strategies that provide different ways for handling the risk introduced by reasoning based on defeasible knowledge. Section 5 concludes the paper.

2   Background

Qualitative multi-attribute preferences Qualitative multi-attribute preferences over objects are based on a set of relevant attributes, which are ranked according to their importance. Without loss of generality, we only consider binary (Boolean) attributes (cf. [7]) and assume that the presence of an attribute is preferred over its absence. The importance ranking of attributes is defined by a total preorder (a total, reflexive and transitive relation) $\succeq$. The relation $\succeq$ yields a stratification of the set of attributes into importance levels. Each importance level consists of attributes that are deemed equally important. Together with factual information about which objects have which attributes, the attribute ranking forms the basis on which various object preference orderings can be defined. One of the most well-known preference orderings is the lexicographic ordering, which we will use here. In [8] it was concluded from experiments that among several qualitative approaches to order options based on their positive and negative aspects, cardinality-based approaches such as the lexicographic ordering best predict the actual choices made by humans. The lexicographic preference ordering first considers the highest importance level. If some object has more attributes on that level than another, the first is preferred. If both objects have the same number of attributes on this level, the next importance level is considered, and so on. Two objects are equally preferred if they have the same number of attributes on every importance level. The argumentation framework of [1] models the lexicographic ordering. In this paper we extend that framework to incorporate defeasible knowledge.

Defeasible knowledge Defeasible knowledge is not known for a fact, but derived from defeasible rules. Defeasible rules describe what is 'normally' the case. Using this kind of rules can add some information to an incomplete knowledge base. This can be beneficial in situations where a user does not have certain information, and does not have

the time or resources to verify information. However, defeasibly inferred information is not quite the same as certain information. If we compare two objects based on one desired attribute, and we know for a fact that the first one has this attribute, but we can only defeasibly infer that the second one has it, then we would prefer the first object. This is due to the chance that defeasible information is incorrect. When comparing a single attribute, it is easy to see that a certainly true attribute is preferred over a defeasibly true attribute. It becomes more difficult when we compare multiple attributes. Are three defeasibly true attributes better or worse than two certainly true attributes and one certainly false attribute? This question is related to the notion of risk. The main question of this paper is how we can deal with risk in a qualitative setting. We first look at the way risk is defined in classical quantitative utility theory, and then discuss the risk of reasoning with defeasible knowledge.

Risk In classical utility theory [2], the notions of risk aversion and risk proneness are defined in terms of lotteries. In a lottery, there are several possible outcomes or consequences. Each consequence has a certain probability of occurring. A lottery is called nondegenerate if no single consequence has a probability of 1 of occurring. Risk aversion and risk proneness are defined as follows ([2], p. 149, 150): *'A decision maker is risk averse if he prefers the expected consequence of any nondegenerate lottery to that lottery. A decision maker is risk prone if he prefers any nondegenerate lottery to the expected consequence of that lottery.'* In other words, risk aversion means that when faced with a choice between two alternatives that have the same expected outcome, but in one case this outcome is certain, and in the other case there are multiple possible outcomes, we would prefer the certain alternative. A risk-prone person would prefer the uncertain alternative as there is a chance to obtain a better outcome than the expected outcome.

Defeasible reasoning and risk When there is only certain information, each attribute is either true or false for a given object. No uncertainty is involved. But when defeasible rules are added, uncertain information is introduced. We introduce four abstract 'certainty levels' (degrees of belief): ✓! denotes that it is known for a fact that an object has an attribute, ✓? that it can be defeasibly inferred that an object has an attribute, ✗? that it can be defeasibly inferred that an object does not have an attribute, and ✗! that it is known for a fact that an object does not have an attribute. Of course there are two other options. First, it can be the case that nothing can be derived about the truth of an attribute for a certain object. This situation is treated in [1] and the current framework can be extended in a similar way as presented there. Second, it can happen that it can both be (defeasibly) derived that an object has an attribute and that it does not, so we have contradictory information. In this case, there will be two rebutting arguments and two preferred extensions (see Section 3).

Essentially, a certainty level represents a probability (p) of an attribute being true. In the case of ✓!, this probability is 1; certainly true information is always true. Similarly, the probability is 0 for ✗!. For ✓? and ✗?, no exact probability can be given. However, since defeasible information is supposed to be 'normally' true, we may assume that the probability lies between 0.5 and 1 for ✓? and between 0 and 0.5 for ✗?. We simplify and assume the same probability holds for all attributes, even though in realistic settings the certainty of defeasible knowledge will differ per attribute.

Now suppose, for the sake of argument, that we know that $p(\checkmark?) = 0.75$, i.e. if it can be defeasibly derived that an attribute is true, there is a 75% chance that it is actually true. Suppose we have an object with four such defeasibly true attributes. Then we can expect to have three true attributes and one false one. This is the same if we would have another object with three certainly true attributes and one certainly false one. The only difference is the certainty: we are only certain of this outcome in the latter case, in the former case we might also end up with four or only two true attributes (or even only one or none). Which of the two objects a person prefers depends on his risk attitude. A risk averse person would prefer the certain case, a risk prone person would prefer the object with four defeasibly true attributes. In general, a risk attitude determines how many defeasibly true attributes are equally preferred as a number of certainly true attributes. The higher this number, the more risk averse a person is. The same holds for the relation between the number of defeasibly false attributes and the number of defeasibly true attributes. In Section 4 we will introduce different strategies with different relative valuations of certainty levels. But first we introduce a generic argumentation framework for risk sensitive preferences.

3   A generic framework for risk sensitive preferences

In the lexicographic ordering, objects are compared w.r.t. attributes on a certain importance level. The highest importance level where a preference can be derived determines the overall preference. In the Boolean case, preference within an importance level is determined solely by the number of true attributes of both objects (the number of false attributes can be ignored because it can be computed when the number of true attributes is known). This comparison is relatively easy, since we only have to compare two numbers. In the case of defeasible knowledge, instead of two possible values for an attribute, we have four. Therefore, preference within an importance level is determined by the number of certainly true, defeasibly true and defeasibly false attributes of both objects (like above, the number of certainly false attributes can be ignored). So in this case, we compare two triples of numbers. Such a comparison can be done in different ways, resulting in different strategies. Abstractly, we can say that the triples are compared by a 'beats' relation B. If on some importance level object a has n certainly true, m defeasibly true and r defeasibly false attributes, object b has n' certainly true, m' defeasibly true and r' defeasibly false attributes, and $\langle n, m, r \rangle >_B \langle n', m', r' \rangle$, then object a is preferred over object b on that importance level.

We first introduce a generic argumentation framework that can derive preferences if a relation B is given. The argumentation framework we present here builds on the framework introduced in [1]. An abstract argumentation framework (AF) [9] is a pair $\langle \mathcal{A}, \rightarrow \rangle$ where $\mathcal{A}$ is a set of arguments, and $\rightarrow$ a binary defeat relation (informally, a counterargument relation) on $\mathcal{A}$. To define which arguments are justified, we use Dung's [9] preferred semantics.

Definition 1. (Preferred semantics) *A preferred extension of an AF $\langle \mathcal{A}, \rightarrow \rangle$ is a maximal (w.r.t. $\subseteq$) set $S \subseteq \mathcal{A}$ such that: $\forall A, B \in S : A \not\rightarrow B$ and $\forall A \in S$: if $B \rightarrow A$ then $\exists C \in S : C \rightarrow B$. An argument is credulously (sceptically) justified w.r.t. preferred semantics if it is in some (all) preferred extension(s).*

Informally, a preferred extension is a coherent point of view that can be defended against all its attackers. In case of contradictory information there will be multiple preferred extensions, each advocating one point of view. The contradictory conclusions will be credulously, but not sceptically justified. The advantage of preferred semantics over e.g. grounded semantics is that a preference based on contradictory information can still be sceptically justified.

We instantiate an abstract AF by specifying the structure of arguments and the defeat relation. Arguments are built from formulas of a logical language, that are chained together using inference steps. Every inference step consists of premises and a conclusion. Inferences can be chained by using the conclusion of one inference step as a premise in the following step. Thus a tree of chained inferences is created, which we use as the formal definition of an argument (similar to e.g. [10]).

Definition 2. (Argument) *An argument is a tree, where the nodes are inferences, and an inference can be connected to a parent node if its conclusion is a premise of that node. Leaf nodes only have a conclusion (a formula from the knowledge base), and no premises. A subtree of an argument is also called a* subargument. `inf` *returns the last inference of an argument (the root node), and* `conc` *returns the conclusion of an argument (the conclusion of its last inference).*

Definition 3. (Language) *Let $\mathcal{P}$ be a set of attribute names with typical elements P, Q; $\mathcal{O}$ a set of object names with typical elements a, b. The* input language $\mathcal{L}$ *is defined as* $\varphi \in \mathcal{L} ::= L \mid P \succ Q \mid P \approx Q \mid L_1, \ldots, L_k, \sim L_l, \ldots, \sim L_m \Rightarrow L_n$ *where* $L_i = P(a)$ *or* $\neg P(a)$.

This language allows us to express facts about the attributes that objects (do not) have, statements about the importance ordering of attributes, and defeasible rules. Other formulas, e.g. expressing a preference between two objects, can be derived using inference steps that build up an argument.

Table 1a shows the inference schemes that are used. The first inference scheme is called defeasible modus ponens. It allows to infer conclusions from defeasible rules. To distinguish these conclusions from hard facts, they are labelled with an asterisk. The next two inference rules define the meaning of the weak negation $\sim$. A formula $\sim \varphi$ can always be inferred (inference rule 2), but such an argument will be defeated by an undercutter built with inference rule 3 if φ is the case. Inference scheme 5a is used to count the number of attributes of equal importance as some attribute P that object a certainly has. Similarly, the inference schemes 5b and 5c count the number of defeasibly true and defeasibly false attributes of an object a at some importance level. Inference scheme 4 can be used when an object has no certainly true (defeasibly true/defeasibly false) attributes. It is possible to construct an argument that does not count all attributes that are certainly true (defeasibly true/defeasibly false), a so-called non-maximal count. This can lead to the conclusion of incorrect preferences. To ensure that only maximal counts are used, we provide an inference scheme to construct arguments that undercut non-maximal counts (inference scheme 6). An argument of this type says that any count which is not maximal is not applicable. Inference scheme 7 aggregates the counts of the different certainty levels for an object at some importance level. Inference scheme 8 uses these aggregated counts and the relation B to infer a preference between two objects. This inference can be undercut by an argument using inference scheme 9, stating

$$1 \quad \frac{\varphi_1, \ldots, \varphi_m \Rightarrow \varphi \quad \varphi_1^{(*)} \quad \ldots \quad \varphi_m^{(*)}}{\varphi^*} \; DMP$$

$$11 \quad \frac{has(a,[P],l,n) \quad has(b,[P'],l,m) \quad P \approx P' \quad n > m}{pref(a,b)} \; p(a,b,[P],l)$$

$$2 \quad \frac{}{\sim \varphi} \; asm(\sim \varphi)$$

$$12 \quad \frac{has(a,[P],l',n) \quad has(b,[P'],l',m) \quad P \approx P' \quad l' > l \quad n \neq m}{p(a,b,[P]) \text{ is inapplicable}} \; p(a,b,[P],l)uc_1$$

$$3 \quad \frac{\varphi}{asm(\sim \varphi) \text{ is inapplicable}} \; asm(\sim \varphi)uc$$

$$13 \quad \frac{has(a,[Q],l',n) \quad has(b,[Q'],l',m) \quad Q \approx Q' \succ P \quad n \neq m}{p(a,b,[P]) \text{ is inapplicable}} \; p(a,b,[P],l)uc_2$$

$$4 \quad \frac{}{has(a,[P],l,0)} \; count(a,[P],l,\varnothing)$$

c. *Lexicographic strategy*

$$5a \quad \frac{P_1(a) \quad \ldots \quad P_n(a) \quad P_1 \approx \ldots \approx P_n}{has(a,[P_1],\checkmark!,n)} \; count(a,[P_1],\checkmark!,\{P_1,\ldots P_n\})$$

$$5b \quad \frac{P_1(a)^* \quad \ldots \quad P_n(a)^* \quad P_1 \approx \ldots \approx P_n}{has(a,[P_1],\checkmark?,n)} \; count(a,[P_1],\checkmark?,\{P_1,\ldots P_n\})$$

$$5c \quad \frac{\neg P_1(a)^* \quad \ldots \quad \neg P_n(a)^* \quad P_1 \approx \ldots \approx P_n}{has(a,[P_1],\boldsymbol{X}?,n)} \; count(a,[P_1],\boldsymbol{X}?,\{P_1,\ldots P_n\})$$

$$6 \quad \frac{has(a,[P_1],l,m) \quad m > n}{count(a,[P_1],l,\{P_1,\ldots,P_n\}) \text{ is inapplicable}} \; count(a,[P_1],l,\{P_1,\ldots,P_n\})uc$$

$$7 \quad \frac{has(a,[P_1],\checkmark!,n) \quad has(a,[P_2],\checkmark?,m) \quad has(a,[P_3],\boldsymbol{X}?,r) \quad P_1 \approx P_2 \approx P_3}{has(a,[P_1],\langle n,m,r \rangle)} \; aggr(a,[P_1])$$

$$8 \quad \frac{has(a,[P],\langle n,m,r \rangle) \quad has(b,[P'],\langle n',m',r' \rangle) \quad P \approx P' \quad \langle n,m,r \rangle >_B \langle n',m',r' \rangle}{pref(a,b)} \; p(a,b,[P])$$

$$9 \quad \frac{has(a,[Q],\langle n,m,r \rangle) \quad has(b,[Q'],\langle n',m',r' \rangle) \quad Q \approx Q' \succ P \quad \langle n,m,r \rangle \neq_B \langle n',m',r' \rangle}{p(a,b,[P]) \text{ is inapplicable}} \; p(a,b,[P])uc$$

a. *Generic framework*

$$14 \quad \frac{n + v^-(\checkmark?) \cdot m + v^-(\boldsymbol{X}?) \cdot r > n' + v^+(\checkmark?) \cdot m' + v^+(\boldsymbol{X}?) \cdot r'}{\langle n,m,r \rangle >_B \langle n',m',r' \rangle}$$

$$10 \quad \frac{U(n,m,r) > U(n',m',r')}{\langle n,m,r \rangle >_B \langle n',m',r' \rangle}$$

$$15 \quad \frac{\sim \langle n,m,r \rangle =_B \langle n',m',r' \rangle}{\langle n,m,r \rangle \neq_B \langle n',m',r' \rangle}$$

b. *Quantitative approach*

d. *Adaptive strategy*

Table 1. Inference schemes

that there is a difference in preference on a higher importance level. Similar inference schemes for equal preference have been omitted due to space limitations.

To complete our argumentation framework, we need to specify a defeat relation. The most common type of defeat is rebuttal. An argument rebuts another argument if its conclusion is the negation of the conclusion of the other argument. Defeat by rebuttal is mutual if both arguments express defeasible information (we assume that certain information is consistent). An argument claiming a certain fact cannot be defeated by an argument that defeasibly derives the opposite. Another type of defeat is undercut. An undercutter is an argument for the inapplicability of an inference used in another argument. Undercut works only one way. Defeat is defined recursively, which means that rebuttal can attack an argument on all its premises and (intermediate) conclusions, and undercut can attack it on all its inferences.

Definition 4. (Defeat) *An argument A defeats an argument B if (i)* $\text{conc}(A) = \varphi$ *and* $\text{conc}(B) = \neg\varphi$ *(rebuttal) and it is not the case that* $\text{conc}(A)$ *is defeasible and* $\text{conc}(B)$ *is certain, or (ii)* $\text{conc}(A) =$ '$\text{inf}(B)$ *is inapplicable' (undercut), or (iii) A defeats a subargument of B.*

4 Qualitative strategies

We have seen above that a risk attitude is determined by the relation between the probability of truth of a certainty level and the relative valuations of certainty levels. For example, if $p(\checkmark?) = 0.75$, then a risk neutral person would equally prefer an object with four defeasibly true attributes and an object with three certainly true attributes and one certainly false attribute (at the same importance level), while a risk prone person would prefer the former and a risk averse person the latter. If we would take a quantitative approach and fix the valuation of certainly true attributes at 1 ($v(\checkmark!) = 1$) and the valuation of certainly false attributes at 0 ($v(\boldsymbol{X}!) = 0$), then in this case the risk neutral person would have $v(\checkmark?) = p(\checkmark?) = 0.75$, the risk averse person $v(\checkmark?) < 0.75$ and the risk prone person $v(\checkmark?) > 0.75$. Of course, the same holds for $v(\boldsymbol{X}?)$. Note that if the values of $p(\checkmark?)$ and $p(\boldsymbol{X}?)$ are unknown, we are not able to distinguish this clearly between risk prone and risk averse attitudes, since the risk neutral boundary would be unknown too. However, as the possible values of $p(\checkmark?)$ and $p(\boldsymbol{X}?)$ are bounded ($0.5 < p(\checkmark?) < 1$ and $0 < p(\boldsymbol{X}?) < 0.5$), we can identify some extreme cases. Consider for example the valuations $v(\boldsymbol{X}?) = v(\checkmark?) = v(\checkmark!) = 1$. Essentially, this means that whenever it is not completely certain whether an attribute is true or false, it is assumed to be true. This is an extremely risk prone approach and overly optimistic, since it would mean an equal preference between a certainly true attribute and a defeasibly false one. The opposite case $v(\checkmark?) = v(\boldsymbol{X}?) = v(\boldsymbol{X}!) = 0$ is extremely risk averse and overly pessimistic. Note that both approaches are a waste of defeasible reasoning, since they do not distinguish between defeasibly true and defeasibly false information. A third extreme option is $v(\boldsymbol{X}?) = v(\boldsymbol{X}!) = 0$ and $v(\checkmark?) = v(\checkmark!) = 1$, which shows great confidence in the correctness of defeasible information. But as this approach does not distinguish between defeasible and certain information, it does not take risk into account. For any non-extreme risk sensitive approach, we would have $v(\checkmark!) > v(\checkmark?) > v(\boldsymbol{X}?) > v(\boldsymbol{X}!)$.

The strategies we present in this section all apply the lexicographic ordering in the sense that a preference between two objects is determined at the highest importance level of attributes where a preference can be derived. They differ in the way a preference is determined within one importance level.

In a quantitative approach, one object is preferred over another on a certain importance level if the utility of the first object on that level is higher than the utility of the second. A utility function U takes the number of certainly true (n), defeasibly true (m), and defeasibly false (r) attributes that an object has at a certain importance level and returns that object's utility at that level. A linearly additive utility function is defined as $U(n,m,r) = n \cdot v(\checkmark!) + m \cdot v(\checkmark?) + r \cdot v(\boldsymbol{X}?)$. This can be incorporated into our generic argumentation framework by means of inference rule 10 in Table 1b.

4.1 Lexicographic strategy

In this section we model a strategy in which no number of defeasibly true attributes can be valuated higher than a single certainly true attribute. We do this by applying a lexicographic ordering within one importance level. That is, within one level, we first count all certainly true attributes. If one object has more of those than another, the first object is preferred over the second (within this level). If both objects have the same number, we go on to count the defeasibly true attributes, and so on. When both objects have the same number of certainly true, defeasibly true and defeasibly false attributes, we go on to consider the next importance level.

Instead of the inference schemes 8 and 9 from Table 1a, we use the inference schemes in Table 1c. Inference scheme 11 says that an object a is preferred over an object b if the number of attributes of object a at a certain importance level and with a certain certainty level is higher than the number of attributes of object b on the same importance and certainty levels. It is also required that a and b have the same number of attributes on any higher certainty level within the same importance level, and on any certainty level at importance levels higher than that of P. We model this by defining inference schemes 12 and 13 that undercut scheme 11 if either is not the case. Similar inference schemes can be used for equal preference.

Example 1 Consider the situation to the right. Argument A claims that b is preferred over a because it has more defeasibly true attributes at the second importance level. This argument is undercut, by argument B stating that since a and b have a different number of certainly true attributes at this importance level, preference cannot be derived on the basis of the number of defeasibly true attributes. The only justified preference argument is C, which states that a is preferred over b because it has more certainly true attributes at the second importance level.

	P	$\approx$	Q	$\succ$	R	$\approx$	S	$\approx$	T
a	✗?		✓!		✓!		✓!		✗?
b	✓!		✗?		✓?		✓?		✓?

A:
$$\frac{\dfrac{R(b)^*\quad S(b)^*\quad T(b)^*\quad R\approx S\approx T}{has(b,[R],✓?,3)}\qquad \dfrac{}{has(a,[R],✓?,0)}\qquad 3>0}{pref(b,a)}\,\beta$$

B:
$$\frac{\dfrac{R(a)\quad S(a)\quad R\approx S}{has(a,[R],✓!,2)}\qquad \dfrac{}{has(b,[R],✓!,0)}\qquad ✓!>✓?\quad 2\neq 0}{\beta \text{ is inapplicable}}$$

C:
$$\frac{\dfrac{R(a)\quad S(a)\quad R\approx S}{has(a,[R],✓!,2)}\qquad \dfrac{}{has(b,[R],✓!,0)}\qquad 2>0}{pref(a,b)}$$

With the quantitative approach described above, we can model the same preferences if we choose the right valuations for the certainty levels. In this case, the valuations for ✓? and ✗? depend on the maximal number of attributes at an importance level, n. Since no number of defeasibly true attributes is enough to beat one certainly true attribute, we have $v(✓?) < \frac{v(✓!)}{n}$. Similarly, we have $v(✗?) < \frac{v(✓?)}{n}$. So, for example, $v(✓?) = \frac{1}{n+1}$ and $v(✗?) = \frac{1}{(n+1)^2}$ would satisfy these constraints and produce the same preferences as the lexicographic strategy. Since the maximum number of attributes at one importance level is typically more than one, both $v(✓?)$ and $v(✗?)$ would be low, certainly below $p(✓?)$ and most probably below $p(✗?)$, which means that this strategy is risk averse.

4.2 Adaptive strategy

Up till now we have been specific in assigning valuations to certainty levels. But the probability of truth associated with a certainty level is not so specific. So maybe a user is hesitant to specify exact values for certainty levels, especially if he is risk neutral or only moderately risk averse or risk prone. In that case the valuation of a certainty level would have to be the same as or close to the probability of truth, which is unknown. We do know some lower and upper bounds: $0.5 < p(\checkmark?) < 1$ and $0 < p(\textbf{X}?) < 0.5$. We can use these to construct worst and best scenarios. In the worst case, $v(\checkmark?) = 0.5$ and $v(\textbf{X}?) = 0$, and in the best case, $v(\checkmark?) = 1$ and $v(\textbf{X}?) = 0.5$. Now we can use the following intuition. If the worst possible case for object a is still preferred over the best possible case for object b, then a has to be preferred over b (on some importance level). This is formalised by inference rule 14. Two objects are equally preferred if they are equally preferred in every possible scenario, which can only be the case if the probabilities of truth are known exactly, or neither object has any defeasibly true or false attributes. There are also cases in which no preference can be derived. To facilitate the use of the undercutting scheme 9 we use inference scheme 15 which states that unless two triples of numbers can be derived to be equal according to B, they are inequal.

It may happen that the user is able to give more information regarding the probabilities of truth of the defeasible certainty levels. If for example a user knows that $p(\checkmark?)$ lies between 0.8 and 0.85, the upper and lower values for $\checkmark?$ could be adapted accordingly. This makes this strategy adaptive to the background knowledge of the user. The strategy is less decisive than the others we presented, since it is not always able to derive a preferene between two objects. But it can also be considered safer because it takes into account the uncertainty of the probability of truth of defeasible information.

Example 2 Consider the same situation again. If nothing specific is known about the probabilities of truth of $\checkmark?$ and $\textbf{X}?$, no preference can be derived, because the best scenario for a is better than the worst scenario for b and the best scenario for b is better than the worst scenario for a. But if we would know that $0.8 < p(\checkmark?) < 0.9$ and $0.1 < p(\textbf{X}?) < 0.2$, we can construct the following argument.

$$\frac{has(a,[R],\langle 2,0,1\rangle) \quad has(b,[R],\langle 0,3,0\rangle) \quad \dfrac{2+0+0.1 < 0+2.4+0}{\langle 2,0,1\rangle <_B \langle 0,3,0\rangle}}{pref(b,a)}$$

Finally, if we knew exactly that $p(\checkmark?) = 0.75$ and $p(\textbf{X}?) = 0.25$, then we could derive an equal preference between a and b.

The lexicographic strategy results in a complete preorder of preference between objects. The preference order induced by the adaptive strategy is not necessarily complete. Preference is determined by the 'beats' relation B. In the lexicographic strategy, $\langle n,m,r\rangle \geq_B \langle n',m',r'\rangle$ if $n > n'$, or $n = n'$ and $m > m'$, or $n = n'$ and $m = m'$ and $r > r'$, or $n = n'$ and $m = m'$ and $r = r'$. Since this relation is complete (for any n,m,r,n',m',r' we have either $\langle n,m,r\rangle \geq_B \langle n',m',r'\rangle$ or $\langle n',m',r'\rangle \geq_B \langle n,m,r\rangle$), the resulting preference order is also complete. In the adaptive strategy, triples are mapped to a range between the worst case and the best case. Since ranges may overlap, this ordering is not complete. One case where no preference can be derived was shown in Example 2.

5 Conclusion

In this paper we have made the following contributions. We discussed the risk involved in basing preferences on defeasible information and identified the relation between this type of risk and risk as defined in utility theory. We then proposed a generic argumentation framework to reason about multi-attribute preferences in the context of defeasible information. Finally, we presented several qualitative strategies to deal with risk sensitive preferences, which were incorporated into the generic framework.

In our future work we would like to evaluate the strategies in more detail. It would be interesting to understand the properties of the adaptive strategy better. We would also like to distinguish more explicitly between mental attitudes such as beliefs, goals, desires and preferences. This will also allow us to reason about these attitudes, for example that a certain preference we have is based on some specific beliefs. We hope to gain insight from modal preference languages with belief operators. Other interesting areas for future work include the representation of dependent preferences, different degrees of satisfaction of attributes, and preferences based on underlying interests.

References

1. Visser, W., Hindriks, K.V., Jonker, C.M.: Argumentation-based preference modelling with incomplete information. Proc. CLIMA-X. (2009) 156–171
2. Keeney, R.L., Raiffa, H.: Decisions with multiple objectives: preferences and value trade-offs. Cambridge University Press (1993)
3. Dubois, D., Fargier, H., Perny, P.: Qualitative decision theory with preference relations and comparative uncertainty: An axiomatic approach. Artif. Intell. **148** (2003) 219–260
4. Dubois, D., Fargier, H., Bonnefon, J.F.: On the qualitative comparison of decisions having positive and negative features. J. Artif. Intell. Res. **32** (2008) 385–417
5. Bonet, B., Geffner, H.: Arguing for decisions: A qualitative model of decision making. Proc. UAI. (1996) 98–105
6. Amgoud, L., Prade, H.: Using arguments for making and explaining decisions. Artif. Intell. **173**(3-4) (2009) 413–436
7. Brewka, G.: A rank based description language for qualitative preferences. Proc. ECAI. (2004) 303–307
8. Bonnefon, J.F., Fargier, H.: Comparing sets of positive and negative arguments: Empirical assessment of seven qualitative rules. Proc. ECAI. (2006) 16–20
9. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artif. Intell. **77** (1995) 321–357
10. Vreeswijk, G.A.W.: Abstract argumentation systems. Artif. Intell. **90**(1-2) (1997) 225–279