Explaining Qualitative Decision Under Uncertainty by Argumentation

Leila Amgoud and **Henri Prade**

IRIT - CNRS

118, route de Narbonne

31062 Toulouse Cedex 9, France

amgoud, prade@irit.fr

Abstract

Decision making under uncertainty is usually based on the comparative evaluation of different alternatives by means of a decision criterion. In a qualitative setting, pessimistic and optimistic criteria have been proposed. In that setting, the whole decision process is compacted into a criterion formula on the basis of which alternatives are compared. It is thus impossible for an end user to understand why an alternative is good, or better than another.

Besides, argumentation is a powerful tool for explaining inferences, decisions, etc. This paper articulates optimistic and pessimistic decision criteria in terms of an argumentation process that consists of constructing arguments in favor/against decisions, evaluating the strengths of those arguments, and comparing pairs of alternatives on the basis of their supporting/attacking arguments.

Introduction

Decision making under uncertainty relies on the comparative evaluation of different alternatives on the basis of a decision criterion, which can be usually justified by means of a set of postulates. This is, for example, the Savage view of decision under uncertainty based on expected utility. Thus, standard approaches for decision under uncertainty consist in defining decision criteria in terms of analytical expressions that summarize the whole decision process. It is then hard for a person who is not familiar with the abstract decision methodology, to understand why a proposed alternative is good, or better than another. Apart from the quantitative expected utility, other examples of such approaches are the qualitative *pessimistic* or *optimistic* criteria, that have been recently proposed and axiomatically justified (Dubois *et al.* 1999b; Dubois, Prade, & Sabbadin 2001; Giang & Shenoy 2004). The qualitative nature of these criteria make them more reliable to be unpacked in order to better understand the underpinnings of the evaluation. Argumentation is the most appropriate way to advocate a choice thanks to its explanatory power.

Argumentation has been introduced in decision making analysis by several researchers only in the last several years (see the section on related works). Indeed, in everyday life, decision is often based on arguments and counterarguments. Argumentation can be also useful for explaining a choice already made. Until recently, argumentation has been mainly developed in AI for inference systems (e.g. (Amgoud & Cayrol 2002; Prakken & Sartor 1997; Simari & Loui 1992)), using a unique type of argument in favor of beliefs. The expression of these arguments directly mirrors the proof-like justification of those beliefs. However, things seem different in a decision making context, as we shall see. Indeed, the definition of an argument is no longer straightforward in this context, and it may even takes different forms. Depending on the decision criterion one wants to use, a particular form of arguments is chosen.

The aim of this paper is to explain optimistic and pessimistic decision criteria in terms of arguments in favor/against each alternative. More generally, the whole decision process will be made explicit in terms of different steps: i) constructing arguments in favor/against each alternative, ii) evaluating the strength of each argument, and iii) comparing pairs of choices on the basis of the quality of their arguments. Moreover, it is shown that there are two types of arguments in favor and two types of arguments against a choice. Depending on the considered criteria, pessimistic or optimistic, the type of argument at work is not the same.

The paper is organized in the following way. First, pessimistic and optimistic criteria in possibility theory-based decision, and their possibilistic logic counterparts are recalled. The next section presents the general argumentation-based decision framework, and discusses different types of arguments. Several instantiations of the general framework are then discussed in the next section, where different preference relations between decisions are defined on the basis of the typology of arguments previously introduced. The links between these preference relations and the pessimistic and optimistic decision criteria are then established. The related work section points out the novelties w.r.t previous works.

Background on Qualitative Decision

A qualitative approach for decision under risk or uncertainty has been justified both in a Von Neumann-Morgenstern-like (Dubois *et al.* 1999b) setting and in a Savage like setting (Dubois, Prade, & Sabbadin 2001). This approach is based on two qualitative criteria, which have respectively a *pessimistic* and an *optimistic* flavor. They are defined in the

following way: given a possibility distribution π_d restricting the plausible states that can be reached when a decision d takes place, and a qualitative utility function μ, a pessimistic qualitative utility can be defined as:

$$E_*(d) = \min_\omega \max\left(\mu(\omega), n(\pi_d(\omega))\right) \qquad (1)$$

where π_d is a mapping from a set of interpretations Ω to a scale S, and μ is a mapping from Ω to a scale U, and n is a decreasing map from S to U such that $n(0) = \bar{1}$ and $n(1) = \bar{0}$, where 0 and 1 (resp. $\bar{0}$, $\bar{1}$) are the bottom and the top elements of S (resp. U). In the following we assume the full commensurateness of the scales (i.e. $S = U$). $\pi_d(\omega)$ (resp. $\mu(\omega)$) is all the greater as ω is more plausible (resp. satisfactory).

$E_*(d)$ is all the greater as *all* the plausible states ω according to π_d are among the most preferred states according to μ. The pessimistic utility $E_*(d)$ is small as soon as there exists a possible consequence of d which is both highly plausible and bad with respect to preferences. This is clearly a risk-averse and thus a pessimistic attitude.

Similarly, the optimistic qualitative criterion is given by

$$E^*(d) = \max_\omega \min\left(\mu(\omega), \pi_d(\omega)\right). \qquad (2)$$

The criterion $E^*(d)$ corresponds to an optimistic attitude since it is high as soon as there exists a possible consequence of d that is both highly plausible and highly prized. $E^*(d)$ is equal to 1 as soon as one fully acceptable choice ω (i.e., such that $\mu(\omega) = 1$) is also completely plausible.

The two formulas 1 and 2 cannot be directly used for providing an end user with explanation, in general, because its knowledge about the current state of the world and its preferences do not take the form of a distribution π and a utility μ in its mind. Rather, the agent knows statements about the world that are more or less certain, and it is pursuing goals that are more or less imperative for it. Fortunately, such epistemic states of the agent are semantically equivalent to mappings such as π and μ as recalled now.

Two possibilistic logic bases are distinguished for encoding knowledge and preferences respectively. Indeed, $\mathcal{K} = \{(k_j, \rho_j); j = 1, l\}$ represents the available knowledge about the world. k_j is a propositional formula of a language $\mathcal{L}$. The pair (k_j, ρ_j) is understood as $N(k_j) \geq \rho_j$, where N is a necessity measure (Dubois, Lang, & Prade 1994). Namely (k_j, ρ_j) encodes that the piece of knowledge "k_j is true" is certain at least at level ρ_j, where ρ_j belongs to a linearly ordered valuation scale R whose top and bottom elements are resp. denoted by 1 and 0. In fact, $R = n(S)$.

The second possibilistic logic base $\mathcal{G} = \{(g_i, \lambda_i); i = 1, m\}$ represents the preferences of the decision-maker under the form of a prioritized set of *goals*, where g_i is a proposition of the language $\mathcal{L}$ and λ_i is the level of priority for getting the goal g_i satisfied. Priority levels take their values on another linearly ordered scale T where $T = n(U)$. $\mathcal{K}^*$ and $\mathcal{G}^*$ denote resp. the corresponding sets of classical propositions when weights are ignored.

Let $\mathcal{D}$ denote the set of all possible decisions. Elements of $\mathcal{D}$ are literals of the language $\mathcal{L}$. Each potential decision $d \in \mathcal{D}$ is represented by a formula $(d, 1)$ to be added to $\mathcal{K}$ when the decision d is chosen. Let $\mathcal{K}_d = \mathcal{K} \cup \{(d, 1)\}$ be the description of what is known about the world when d is applied. Associated with $\mathcal{K}_d$ is the possibility distribution $\pi_{\mathcal{K}_d}$ that defines its semantics counterpart

$$\pi_{\mathcal{K}_d}(\omega) = \min_{j=1,l} \max(v_\omega(k_j), n(\rho_j)),$$

that rank-orders the more or less plausible states of the world when d is chosen, where $v_\omega(k_j) = 1$ if ω is a model of k_j and $v_\omega(k_j) = 0$ if ω falsifies k_j, and $n(R) = S$. Similarly, another possibility distribution μ from Ω to $U = n(T)$ is associated with the layered set of goals $\mathcal{G}$.

It has been shown in (Dubois *et al.* 1999a) that it is possible to compute $E_*(d)$ by only using a classical logic machinery on α-level cuts of $\mathcal{K}_d$ and $\mathcal{G}$.

Proposition 1. $E_*(d)$ *is the maximal value of* α *s.t.*

$$(\mathcal{K}_d)_\alpha \vdash (\mathcal{G})_{\underline{n(\alpha)}} \qquad (3)$$

where $(B)_\alpha$, *resp.* $(B)_{\underline{\alpha}}$ *is the set of classical propositions in a possibilistic logic base B with a level greater or equal to α, resp. strictly greater than α.*

As seen in (3), $E_*(d)$ is equal to 1 ($\alpha = 1$) if the completely certain part of $\mathcal{K}_d$ entails the satisfaction of all the goals, even the ones with low priorities, since $\mathcal{G}_{\underline{0}}$ is just the set of all the propositions in $\mathcal{G}$ with a non-zero priority level. Similarly, the optimistic case can be expressed in logical terms.

Proposition 2. $E^*(d)$ *is equal to the greatest* α *such that* $(\mathcal{K}_d)_{\underline{n(\alpha)}}$ *and* $(\mathcal{G})_{\underline{n(\alpha)}}$ *are logically consistent together.*

Note that this approach does not require a numerical scale. Let's consider the following medical example borrowed from (Fox & Parsons 1997).

Example 1. *The example is about having or not a surgery, knowing that the patient has colonic polyps. The knowledge base is* $\mathcal{K} = \{(sg \rightarrow se, 1), (\neg sg \rightarrow \neg se, 1), (sg \rightarrow ll, 1), (ca \wedge \neg sg \rightarrow ll, 1), (cp, 1), (\neg ca \rightarrow \neg ll, 1), (cp \rightarrow ca, \lambda)\}$ *($0 < \lambda < 1$) with se: having side-effect, ca: cancer, ll: loss of life, sg: having a surgery, cp: having colonic polyps. λ refers to a lack of complete certainty. The goals base is* $\mathcal{G} = \{(\neg ll, 1), (\neg se, \sigma)\}$ *with ($0 < \sigma < 1$). We do not like to have side effects after a surgery, but it is more important to not lose life. The set of decisions is $\mathcal{D} = \{sg, \neg sg\}$.*

The best pessimistic decision is to have surgery with $E_*(sg) = n(\sigma)$. Moreover, $E_*(\neg sg) = 0$, $E^*(sg) = n(\sigma)$ and $E^*(\neg sg) = n(\lambda)$. Thus the best decision in the optimistic case depends on the values λ and σ.

Argumentation-based Decision Framework

A decision problem amounts to defining a pre-ordering, usually a complete one, on a set $\mathcal{D}$ of possible alternatives, on the basis of the different consequences of each alternative. Argumentation can be used for defining such a pre-ordering. The idea is to construct arguments in favor of and against each alternative, to evaluate such arguments, and finally to apply some principle for comparing pairs of alternatives on the basis of the quality or strength of their arguments.

General Setting

The main ingredients that are involved in the definition of an argumentation-based decision framework are gathered in the following definition:

Definition 1 (Argumentation-based decision framework). *An* argumentation-based decision framework *is a tuple* $<\mathcal{D}, \mathcal{A}, \succeq, \triangleright_{Princ}>$ *where:*

- $\mathcal{D}$ *is the set of possible decisions.*
- $\mathcal{A}$ *is a set of arguments supporting/attacking elements of* $\mathcal{D}$.
- $\succeq$ *is a (partial or complete) pre-ordering on* $\mathcal{A}$.
- $\triangleright_{Princ}$ *is a (partial or complete) pre-ordering on* $\mathcal{D}$.

$d_1 \triangleright_{Princ} d_2$ means that the alternative d_1 is *at least as preferred as* the alternative d_2 w.r.t. the *principle Princ*. Note that different definitions of $\succeq$ and $\triangleright_{Princ}$ may lead to different decision frameworks that may not return the same results. The relation $\succeq$ reflects the strengths of the arguments that depend on the certainty degrees of the pieces of knowledge in $\mathcal{K}$ used in the argument, and the priority degrees of the goals in $\mathcal{G}$ associated with the arguments. $\succ$ is the strict part of $\succeq$.

Typology of Arguments

Given two *consistent* bases $\mathcal{K}^*$ and $\mathcal{G}^*$, we will define in a systematic way the different types of arguments that can be built for a decision in $\mathcal{D}$. A decision may have arguments in its favor (called PROS), and arguments against it (called CONS). In the following, an argument is associated to a decision, and always refers to one goal.

Arguments PROS point out the existence of good consequences, or the absence of bad ones for a given decision. More precisely, we can distinguish between two types of good consequences, namely the guaranteed satisfaction of a goal when $\mathcal{K}^* \cup \{d\} \vdash g$, and the possible satisfaction of a goal when $\mathcal{K}^* \cup \{d\} \not\vdash \neg g$, with $d \in \mathcal{D}$ and $g \in \mathcal{G}^*$. Note that this latter situation corresponds to the existence of an interpretation that satisfies $\mathcal{K}^*$, d, and g. Formally:

Definition 2 (Types of arguments PROS). *Let* $\mathcal{T} = \langle \mathcal{D}, \mathcal{K}, \mathcal{G} \rangle$ *be a theory. An* argument in favor of *a decision d is a tuple* $A = \langle S, d, g \rangle$ *such that:*

1. $S \subseteq \mathcal{K}^*, d \in \mathcal{D}, g \in \mathcal{G}^*, \mathcal{K}^* \cup \{d\}$ *is consistent*
2. • $S \cup \{d\} \vdash g$, *and S is minimal for set* $\subseteq$ *among subsets of* $\mathcal{K}^*$ *satisfying the above criteria (Type P_1), or*
   • $S \cup \{d\} \not\vdash \neg g$, *and S is maximal for set* $\subseteq$ *among subsets of* $\mathcal{K}^*$ *satisfying the above criteria (Type P_2)*

d is called the conclusion *of the argument.* AP_1 *(resp.* AP_2*) denotes the set of all arguments of type P_1 (resp. of type P_2) that can be built from the theory* $\mathcal{T}$.

The consistency of $\mathcal{K}^* \cup \{d\}$ means that d is applicable in the context $\mathcal{K}^*$, in other words that we cannot prove from $\mathcal{K}^*$ that d is impossible. This means that impossible alternatives w.r.t. $\mathcal{K}^*$ have already been taken out from $\mathcal{D}$.

Arguments CONS highlight the existence of bad consequences for a given decision, or the absence of good ones. As in the case of arguments PROS, there are a strong form

and a weak form of both situations. Namely, arguments CONS are defined either by exhibiting a goal that is violated for sure, or a goal that might not be satisfied. Formally:

Definition 3 (Types of arguments CONS). *Let* $\mathcal{T} = \langle \mathcal{D}, \mathcal{K}, \mathcal{G} \rangle$ *be a theory. An* argument against *a decision d is a tuple* $A = \langle S, d, g \rangle$ *such that:*

1. $S \subseteq \mathcal{K}^*, d \in \mathcal{D}, g \in \mathcal{G}^*, \mathcal{K}^* \cup \{d\}$ *is consistent*
2. • $S \cup \{d\} \vdash \neg g$, *and S is minimal for set* $\subseteq$ *among subsets of* $\mathcal{K}^*$ *satisfying the above criteria (Type C_1), or*
   • $S \cup \{d\} \not\vdash g$ *and S is maximal for set* $\subseteq$ *among subsets of* $\mathcal{K}^*$ *satisfying the above criteria (Type C_2)*

d is called the conclusion *of the argument.* AC_1 *(resp.* AC_2*) denotes the set of all arguments of type C_1 (resp. of type C_2) that can be built from the theory* $\mathcal{T}$.

In what follows, Conc is a function that returns the conclusion of an argument, i.e. the decision supported by that argument.

Note that arguments of type 1 are particular cases of arguments of type 2, since if $S \cup \{d\} \vdash g$, then $S \cup \{d\} \cup \{g\}$ is obviously consistent.

Property 1.

- *If* $\exists \langle S, d, g \rangle \in AP_1$, *then* $\exists \langle S', d, g \rangle \in AP_2$.
- *If* $\exists \langle S, d, g \rangle \in AC_1$, *then* $\exists \langle S', d, g \rangle \in AC_2$.

Let us now define the notion of argument strength.

Definition 4 (Strength of an Argument). *Let* $A = \langle S, d, g \rangle \in \mathcal{A}$. *The* strength *of A is a pair (*$\text{Lev}(A), \text{Wei}(A)$*) s.t.*

- *The* certainty level *of the argument is* $\text{Lev}(A) = min\{\rho_i \mid k_i \in S$ *and* $(k_i, \rho_i) \in \mathcal{K}\}$. *If* $S = \emptyset$ *then* $\text{Lev}(A) = 1$
- *The* weight *of the argument is* $\text{Wei}(A) = \alpha$ *s.t* $(g, \alpha) \in \mathcal{G}$

The strengths of arguments make it possible to strictly compare pairs of arguments. Note that strong arguments are not compared in the same way as the weak ones. In the case of strong arguments an argument is all the better as it uses more certain beliefs and refers to an important goal. In contrast, a weak argument is all the better as it is based on a larger subset of $\mathcal{K}$ involving less certain beliefs that is consistent with a goal (resp. its negation). Formally, using the usual Pareto partial order between vectors:

Definition 5 (Comparing arguments of type 1). *Let A and B be two arguments of type AP_1 or AC_1. A is* stronger *than B, denoted* $A \succ B$ *iff (*$\text{Lev}(A), \text{Wei}(A)) >_{Pareto} (\text{Lev}(B), \text{Wei}(B))$.

Definition 6 (Comparing arguments of type 2). *Let A and B be two arguments of type AP_2 or AC_2. A is* stronger *than B, denoted* $A \succ B$ *iff*
$(\text{Lev}(A), n(\text{Wei}(A))) <_{Pareto} (\text{Lev}(B), n(\text{Wei}(B)))$.

Particular Frameworks

This section outlines different argumentation frameworks that handle different situations of various levels of generality, namely the case where both $\mathcal{K}$ and $\mathcal{G}$ are flat (without priorities), the case where only one of $\mathcal{K}$ and $\mathcal{G}$ is prioritized, and finally the general case where both bases are prioritized. In each case, we emphasize the existence of a *cautious* and

a *bold* way for ranking decisions on the basis of previously identified types of arguments.

Let $\text{Goals}_X(d)$ be a function that returns for a given decision d, all the goals for which there exists an argument of type X with conclusion d.

Flat Bases In this case, the two bases $\mathcal{K}$ and $\mathcal{G}$ are supposed to be flat. Consequently, all the arguments are equally preferred, and thus the relation $\succ$ is empty.

From a *cautious* point of view, a decision d is "good" iff it leads to the satisfaction of all the goals for sure, i.e. it has arguments of type P_1 for all the goals in $\mathcal{G}$. d is "bad" (not good) as soon as one goal may be missed, i.e. there exists g such that $\exists S \subseteq \mathcal{K}^*$, $S \cup \{d\} \not\vdash g$. In other words, d is bad as soon as there exists an argument of type C_2 against it. The decision strict criterion that captures the above idea is the following: Let $d, d' \in \mathcal{D}$.

$$d \triangleright_{caut_1} d' \text{ iff } \text{Goals}_{C_2}(d) = \emptyset, \text{ and } \text{Goals}_{C_2}(d') \neq \emptyset \text{ (4)}$$

Indeed, the above definition is equivalent to the following one:

$$d \triangleright_{caut_1} d' \text{ iff } \forall g \in \mathcal{G}, g \in \text{Goals}_{P_1}(d), \text{ and } \exists g \in \mathcal{G} \text{ s.t}$$
$$g \in \text{Goals}_{C_2}(d')$$

The argumentation-based decision framework that is used in the above case is $<\mathcal{D}, AC_2, \emptyset, \triangleright_{caut_1}>$, or $<\mathcal{D}, AP_1, \emptyset, \triangleright_{caut_1}>$.

Remark: A first possible refinement of the above strict criterion is the following:

$$d \triangleright_{caut_2} d' \text{ iff } \text{Goals}_{P_1}(d) \supset \text{Goals}_{P_1}(d') \text{ (5)}$$

This partial preorder can be further refined into a complete preorder as follows:

$$d \triangleright_{caut_3} d' \text{ iff } |\text{Goals}_{P_1}(d)| > |\text{Goals}_{P_1}(d')| \text{ (6)}$$

Similarly, from a *bold* point of view, a decision d is bad only if there exists a goal g that is surely missed, i.e. if there exist arguments of type C_1 against d. On the contrary it is "good" iff d is consistent with $\mathcal{K}$ and $\mathcal{G}$, i.e. there exist arguments of type P_2 for all the goals in $\mathcal{G}$. Formally:

$$d \triangleright_{bold_1} d' \text{ iff } \text{Goals}_{C_1}(d) = \emptyset \text{ and } \text{Goals}_{C_1}(d') \neq \emptyset. \text{ (7)}$$

This relation is equivalent to the following one:

$$d \triangleright_{bold_1} d' \text{ iff } \forall g \in \mathcal{G}, g \in \text{Goals}_{P_2}(d), \text{ and } \exists\, g \in \mathcal{G} \text{ s.t}$$
$$g \in \text{Goals}_{C_1}(d').$$

The argumentation-based decision framework that is used in this case is $<\mathcal{D}, AC_1, \emptyset, \triangleright_{bold_1}>$. Inclusion and cardinality-based refinements can be defined as in the cautious case.

Flat Knowledge Base and Prioritized Goals Base Since goals are prioritized, then the arguments may not have equal strength. Indeed, the strength of an argument A is equal to $(1, \text{Wei}(A))$. The decision criterion (4) is generalized into:

$$d \triangleright_{caut_4} d' \text{ iff } \exists A \in AC_2 \text{ with } \text{Conc}(A) = d' \text{ s.t.}$$
$$\forall B \in AC_2 \text{ with } \text{Conc}(B) = d, \text{ then } A \succ B. \text{ (8)}$$

It can be shown that the above relation is equivalent to the following one:

$$d \triangleright_{caut_4} d' \text{ iff } \exists \alpha \text{ s.t. } \forall g \in \mathcal{G}_\alpha, g \in \text{Goals}_{P_1}(d), \text{ and}$$
$$\exists g \in \mathcal{G}_\alpha \text{ s.t. } g \in \text{Goals}_{C_2}(d').$$

This means that d leads for sure to the satisfaction of all goals having a priority at least equal to α, whereas d' may miss a goal in the same category. The extension of criteria (5) and (6) is straightforward by taking into account only arguments whose goals are in $\mathcal{G}_\alpha$.

Flat Goals Base and Prioritized Knowledge Base Here again, the arguments may not have equal strength, which is in this case is equal to $(\text{Lev}(A), 1)$ for a given argument A. The decision criterion (4) is generalized into:

$$d \triangleright_{caut_4} d' \text{ iff } \exists A \in AC_2 \text{ with } \text{Conc}(A) = d' \text{ s.t.}$$
$$\forall B \in AC_2 \text{ with } \text{Conc}(B) = d, \text{ then } A \succ B. \text{ (9)}$$

It can also be checked that relation (9) is equivalent to the following one:

$$d \triangleright_{caut_5} d' \text{ iff } \exists \beta \text{ s.t. } \forall g \in \mathcal{G}, \exists A =< S, d, g > \in AP_1 \text{ s.t.}$$
$$\text{Lev}(A) \geq \beta, \text{ and } \exists g \in \mathcal{G} \text{ s.t } \exists B =< S', d', g >\in AC_2, \text{ s.t.}$$
$$\text{Lev}(B) \geq \beta.$$

General Case Let us now analyze the general case where the two bases $\mathcal{K}$ and $\mathcal{G}$ are supposed to be weighted. The decision criterion (4) is generalized into:

$$d \triangleright_{caut_4} d' \text{ iff } \exists A \in AC_2 \text{ with } \text{Conc}(A) = d' \text{ s.t.}$$
$$\forall B \in AC_2 \text{ with } \text{Conc}(B) = d, \text{ then } A \succ B. \text{ (10)}$$

Thus, decisions that have no strong argument of type C_2 against them are favored. The argumentation-based decision framework that is used in the above case is $<\mathcal{D}, AC_2, \succ, \triangleright_{caut_4}>$.

Now let us come back to the bold point of view, which reads in the general case as follows:

$$d \triangleright_{bold_2} d' \text{ iff } \forall A \in AP_2 \text{ s.t. } \text{Conc}(A) = d', \exists B \in AP_2 \text{ s.t.}$$
$$\text{Conc}(B) = d,$$
$$\text{then } B \succ A. \text{ (11)}$$

Here decisions that have stronger arguments of type P_2 in favor of them are preferred. The argumentation-based decision framework that is used in the above case is $<\mathcal{D}, AP_2, \succ, \triangleright_{bold_2}>$.

Linking Argumentation-based Decision with Qualitative Decision

We are now in a position for establishing that the cautious point of view agrees with the pessimistic criterion, while the bold relation corresponds to the optimistic one.

Proofs are omitted for the sake of brevity. They heavily rely on Propositions 1 and 2.

Pessimistic Criterion

In the pessimistic view, as pointed out by Proposition 1, we are interested in finding a decision d (if it exists) such that $\mathcal{K}_\alpha \cup \{d\} \vdash \mathcal{G}_\beta$ with α high and β low, i.e. such that the decision d together with the most certain part of $\mathcal{K}$ entails the satisfaction of the goals, even those with low priority. To capture the result of Proposition 1, we need the following argumentation framework $<\mathcal{D}, AC_2, \succ, \triangleright_{caut_4}>$. Indeed, the following results can be proved:

Theorem 1. *Let $d \in \mathcal{D}$. If $\exists A \in AC_2$ s.t. $\text{Conc}(A) = d$, then $E_*(d) \leq max(\text{Lev}(A), n(\text{Wei}(A)))$.*

Theorem 2. *Let $d \in \mathcal{D}$. If $\forall g \in \mathcal{G}_\alpha$, $\exists A \in AP_1$ s.t. $\mathtt{Conc}(A) = d$, and $\mathtt{Lev}(A) \geq \beta$, then $E_*(d) \geq min(\beta, n(\alpha))$.*

Example 2 (Cont.)**.** *In the above example, there is an argument of type P_1 in favor of sg: $A = <\{sg \rightarrow \neg ll\}, \{\neg ll\}, sg>$, and there is an argument of type P_1 in favor of $\neg sg$: $B = <\{\neg sg \rightarrow \neg se\}, \{\neg se\}, \neg sg>$.*

The strength of A is $<1, 1>$, whereas the strength of B is $<1, \sigma>$. Thus, A is preferred to B. Consequently, the decision sg is preferred to the decision $\neg sg$.

The agreement between the qualitative decision criterion and the argument-based view is due to a decomposability property of arguments of type P_1 w.r.t the conjunction of goals. Namely, $\mathcal{K}_\beta \cup \{d\} \vdash g$ and $\mathcal{K}_\beta \cup \{d\} \vdash g'$ is equivalent to $\mathcal{K}_\beta \cup \{d\} \vdash g \wedge g'$. However, things are not as simple with consistency since one may have $\mathcal{K}_\beta \cup \{d\}$ consistent with both g and g' separately without having it consistent with $g \wedge g'$. This means that arguments of type P_2 are only necessary conditions for consistency w.r.t the whole set of goals. Thus, the optimistic criterion can only be approximated.

Optimistic Criterion

In the optimistic point of view, we are interested in finding a decision d (if it exists) which is consistent with the knowledge base and the goals (i.e. $\mathcal{K}^* \wedge \{d\} \wedge \mathcal{G}^* \neq \bot$). This is optimistic in the sense that it assumes that goals may be attained as soon as their negation cannot be proved. In order to capture the result of Proposition 2, the following argumentation framework will be used $<\mathcal{D}, AP_2, \succ, \rhd_{bold_2}>$.

Theorem 3. *Let $d \in \mathcal{D}$. If $\exists A \in AC_1$ s.t. $\mathtt{Conc}(A) = d$, then $E^*(d) \leq max(n(\mathtt{Lev}(A)), n(\mathtt{Wei}(A)))$.*

Example 3 (Cont.)**.** *In the above example, there is one argument against the decision 'sg': $<\{sg \rightarrow se\}, \{\neg se\}, sg>$. There is also a unique argument against the decision $\neg sg$: $<\{cp, cp \rightarrow ca, ca \wedge \neg sg \rightarrow ll\}, \{\neg ll\}, \neg sg>$.*

The level of the argument $<\{sg \rightarrow se\}, \{\neg se\}, sg>$ is 1 whereas its weight is σ. Concerning the argument $<\{cp, cp \rightarrow ca, ca \wedge \neg sg \rightarrow ll\}, \{\neg ll\}, \neg sg>$, its level is λ, and its weight is 1.

In this example, the comparison of the two arguments amounts to compare σ with λ. Namely, if σ (the priority of the goal "no side effect") is small then the best decision will be to have a surgery. If the certainty degree λ of having cancer in presence of colonic polyps for the particular patient is small enough then the best optimistic decision will not be to have a surgery.

Related Works

As said in the introduction, some works have been done on arguing for decision. In (Fox & Parsons 1997), no explicit distinction is made between knowledge and goals. However, in their examples, values (belonging to a linearly ordered scale) are assigned to formulas which represent goals. These values provide an empirical basis for comparing arguments using a symbolic combination of strengths of beliefs and goals values. This symbolic combination is performed through dictionaries corresponding to different kinds of scales that may be used. Only one type of arguments in favor of or against is used.

In (Bonet & Geffner 1996), Bonet and Geffner have also proposed an original approach to qualitative decision, inspired from Tan and Pearl (Tan & Pearl 1994), based on "action rules" that link a situation and an action with the satisfaction of a *positive* or a *negative* goal. However in contrast with the previous work and the work presented in this paper, this approach does not refer to any model in argumentative inference. In their framework, there are four parts:

1. a set $\mathcal{D}$ of *actions* or *decisions*.

2. a set $\mathcal{I}$ of input propositions defining the possible input situation. A degree of plausibility is associated with each input. Thus, $\mathcal{I} = \{(k_i, \alpha_i)\}$ with $\alpha_i \in \{likely, plausible, unlikely\}$.

3. a set $\mathcal{G}$ of *prioritized goals* such that $\mathcal{G} = \mathcal{G}^+ \cup \mathcal{G}^-$. $\mathcal{G}^+$ gathers the *positive* goals that one wants to achieve and $\mathcal{G}^-$ gathers the *negative* goals that one wants to avoid. Thus, $\mathcal{G} = \{(g_i, \beta_i)\}$ with $\beta_i \in [0, 1, ...N]$.
   Note that in our framework what they call here negative goals are considered in our goal base as negative literals.

4. a set of *action rules* $\mathcal{AR} = \{(A_i \wedge C_i \Rightarrow x_i, \lambda_i), \lambda_i \geq 0\}$, where A_i is an action, C_i is a conjunction of input literals, and x_i is a goal. Each action rule has two measures: a *priority* degree which is exactly the priority degree of the goal x_i, and a *plausibility* degree. This plausibility is defined as follows: A rule $A \wedge C \Rightarrow x$ is likely if any conjunct of C is likely. A rule $A \wedge C \Rightarrow x$ is unlikely if some conjunct of C is unlikely. A rule $A \wedge C \Rightarrow x$ is plausible if it is neither likely nor unlikely.

In this approach only input propositions are weighted in terms of plausibility. Action rules inherit these weights through the three above rules in a rather empirical manner which depends on the chosen plausibility scale. The action rules themselves are not weighted since they are potentially understood as defeasible rules, although no non-monotonic reasoning system is associated with them. In contrast, our approach makes use of an abstract scale. Moreover, weighted possibilistic clauses have been shown to be able to properly handle non-monotonic inference in the sense of Kraus, Lehmann and Magidor (Kraus, Lehmann, & Magidor 1990)' preferential system augmented with rational monotony (see (Benferhat, Dubois, & Prade 1992)). So a part of our weighted knowledge may be viewed as the encoding of a set of default rules.

From the above four bases, reasons are constructed for (against) actions in (Bonet & Geffner 1996). Indeed, goals provide reasons for (or against) actions. Positive goals provide reasons *for* actions, whereas negative goals provide reasons *against* actions. The basic idea behind this distinction is that negative goals should be discarded, and consequently any action which may lead to the satisfaction of such goals should be avoided. However, the approach makes no distinction between what we call pessimism and optimism. The definition of a 'reason' in (Bonet & Geffner 1996) is quite different from our definition of an argument. Firstly, a reason considers only one goal and secondly, the

definition is poor since it only involves facts. Finally, in Bonet and Geffner's framework, decisions which satisfy the most important goals are privileged. This is also true in our approach, but the comparison between decisions can be further refined, in case of several decisions yielding to the satisfaction of the most important goals, by taking into account the other goals which are not violated by these decisions.

Amgoud and Prade in (Amgoud & Prade 2004) have already proposed an argumentation-based reading of possibilistic decision criteria. However, their approach has some drawbacks from a pure argumentation point of view. In their approach, there was only one type of arguments PROS and one type of arguments CONS. Moreover, these arguments were taking into account the goal base as a whole, and a consequence for a given decision there was at most a unique argument PROS and a unique argument CONS. This does not really fit with the way human are discussing decisions, for which there are usually several arguments PROS and CONS, rather than a summarized one. On the contrary in this paper, we have discussed all the possible types of arguments PROS and CONS in a systematic way, and each argument pertains to only one goal.

Conclusion

The paper has sketched a method, agreeing with qualitative possibility-based decision, which enables us to compute and justify best decision choices. We have shown that it is possible to design a logical machinery which directly manipulates arguments with their strengths and compute acceptable and best decisions from them. The approach can be extended in various directions. The computation of the strengths of arguments PROS and CONS can be refined by using vectors of values rather than scalar values for refining max and min aggregation (Fargier & Sabbadin 2003), in order to take into account the presence of several arguments PROS and CONS with the same strength, for instance. This will correspond to extensions of the refinement we proposed in the flat case. Another extension of this work consists of allowing for inconsistent knowledge or goal bases. It can be already noticed that several of the preference relations introduced between decisions still make sense when the set of goals is inconsistent. The approach can be transposed to multiple criteria decision making from the one proposed here for decision under uncertainty, taking advantage of the close relation between both areas (Dubois *et al.* 2000). Criteria satisfaction can then be explained in terms of nested goals with higher priority for the less restrictive goals.

Acknowledgments

References

Amgoud, L., and Cayrol, C. 2002. Inferring from inconsistency in preference-based argumentation frameworks. *International Journal of Automated Reasoning* Volume 29, N2:125–169.

Amgoud, L., and Prade, H. 2004. Using arguments for making decisions. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, 10–17.

Benferhat, S.; Dubois, D.; and Prade, H. 1992. Representing default rules in possibilistic logic. In *Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning*, 673–684.

Bonet, B., and Geffner, H. 1996. Arguing for decisions: A qualitative model of decision making. In *Proceedings of the 12th Conference on Uncertainty in Artificial Intelligence*, 98–105.

Dubois, D.; Berre, D. L.; Prade, H.; and Sabbadin, R. 1999a. Using possibilistic logic for modeling qualitative decision: Atms-based algorithms. *Fundamenta Informaticae* 37:1–30.

Dubois, D.; Godo, L.; Prade, H.; and Zapico, A. 1999b. On the possibilistic decision model: from decision under uncertainty to case-based decision. *Int. J. of Uncert., Fuzziness and Knowledge-Based Syst.* 7:6:631–670.

Dubois, D.; Grabisch, M.; Modave, F.; and Prade, H. 2000. Relating decision under uncertainty and multicriteria decision making models. *International Journal of Intelligent Systems* 15:967–979.

Dubois, D.; Lang, J.; and Prade, H. 1994. Automated reasoning using possibilistic logic: semantics, belief revision and variable certainty weights. *IEEE Trans. on Data and Knowledge Engineering* 6:64–71.

Dubois, D.; Prade, H.; and Sabbadin, R. 2001. Decision-theoretic foundations of qualitative possibility theory. *European Journal of Operational Research* 128:459–478.

Fargier, H., and Sabbadin, R. 2003. Qualitative decision under uncertainty: back to expected utility. In *Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI'03*.

Fox, J., and Parsons, S. 1997. On using arguments for reasoning about actions and values. In *Proceedings of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning, Stanford*.

Giang, P. H., and Shenoy, P. 2004. Two axiomatic approaches to decision making using possibility theory. *European Journal of Operational Research* 162:2:450–467.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Non-monotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44:167 – 207.

Prakken, H., and Sartor, G. 1997. Argument-based extended logic programming with defeasible priorties. *Journal of Applied Non-Classical Logics* 7:25–75.

Simari, G. R., and Loui, R. P. 1992. A mathematical treatment of defeasible reasoning and its implementation. *Artificial Intelligence and Law* 53:125–157.

Tan, S. W., and Pearl, J. 1994. Qualitative decision theory. In *Proceedings of the 11th National Conference on Artificial Intelligence*, 928–933.