



HAL
open science

Using possibilistic logic for modeling qualitative decision: Atms-based algorithms

Didier Dubois, Daniel Le Berre, Henri Prade, Régis Sabbadin

► **To cite this version:**

Didier Dubois, Daniel Le Berre, Henri Prade, Régis Sabbadin. Using possibilistic logic for modeling qualitative decision: Atms-based algorithms. *Fundamenta Informaticae*, 1999, 37 (1-2), pp.1-30. 10.3233/FI-1999-371201 . hal-03300799

HAL Id: hal-03300799

<https://univ-artois.hal.science/hal-03300799v1>

Submitted on 1 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Using possibilistic logic for modeling qualitative decision : ATMS-based algorithms*

Didier Dubois, Daniel Le Berre, Henri Prade, Régis Sabbadin

IRIT - Université Paul Sabatier - 31062 Toulouse Cedex (France)

e-mail: {dubois, leberre, prade, sabbadin} @irit.fr

Abstract. This paper describes a logical machinery for computing decisions, where the available knowledge on the state of the world is described by a possibilistic propositional logic base (i.e., a collection of logical statements associated with qualitative certainty levels), and where the preferences of the user are also described by another possibilistic logic base whose formula weights are interpreted in terms of priorities.

Two attitudes are allowed for the decision maker: a pessimistic risk-averse one and an optimistic one. The computed decisions are in agreement with a qualitative counterpart to the classical theory of expected utility, recently developed by three of the authors.

A link is established between this logical view of qualitative decision making and an ATMS-based computation procedure. Efficient algorithms for computing pessimistic and optimistic optimal decisions are finally given in this logical setting (using some previous work of the fourth author).

Keywords: qualitative decision, possibilistic logic, possibility theory, ATMS.

1. Introduction

An increasing interest for qualitative decision has recently appeared in the Artificial Intelligence community : The term “qualitative decision theory” refers to more than one kind of representation. Some approaches consider only all-or-nothing notions of utility and plausibility, for instance (Bonet and Geffner [2]); others use integer-valued functions (Tan and Pearl [34]), (Pearl [30]). Boutilier [3] exploits preference orderings and plausibility orderings by focusing on the most plausible states. In (Dubois and Prade [17]), a qualitative analog to von Neumann and

*This paper is an extended and revised version of a conference paper by the same authors [16].

Morgenstern postulates, intended for rational decision under ordinal uncertainty has been proved to be equivalent to the maximization of a qualitative utility function. Steps to a Savage-like qualitative decision theory are taken by (Dubois, Prade and Sabbadin [21, 22]).

In classical decision theory under uncertainty, the preferences of the decision maker are directly expressed by means of a utility function, while a probability distribution on the possible states of the world represents the available, uncertain information about the situation under consideration. However, it seems reasonable to allow for a more granular and natural expression of both the preferences and the available knowledge about the world, under the form, e.g., of logical statements from which it would be possible to build the utility and the uncertainty functions. The knowledge about the world is supposed to be given in this paper under the form of a set of pieces of knowledge having different levels of certainty, while the preferences are expressed by a set of goals with different levels of priority.

In Section 2 we propose two syntactic approaches based on possibilistic logic, the first one being more cautious than the second, for computing optimal decisions. They are first presented in the case of binary uncertainty and preferences, before considering graded uncertainty and preferences. Here gradual uncertainty and preferences are expressed by means of two distinct possibilistic propositional logic bases (which are stratified bases). Then, the semantics underlying the two syntactic approaches are shown to be in agreement with the two qualitative utility functions advocated in (Dubois and Prade [17]). This section is a revised version of a workshop paper (Dubois et al. [20]).

In Section 3, we recall some background on the ATMS framework, and it is shown how to encode a decision problem as one of label computation. Then a procedure called MPL (French acronym for Literal based Preferred Models), is described for computing optimal decisions in terms of labels. It relies on a modified Davis and Putnam [9] semantic evaluation algorithm, described in (Castell et al. [6]). Two algorithms based on the use of this procedure, are proposed to compute optimistic and pessimistic optimal decisions respectively. An example is given in Section 4, that illustrates the algorithms.

2. Qualitative decision in stratified propositional bases

2.1. Notations

In this article, upper case letters (K, D, P, H, \dots) denote sets of propositional formulas that can possibly be literals. For any set A of formulas, A^\wedge denotes the logical conjunction of the formulas in A , A^\vee denotes the logical disjunction of the formulas in A .

2.2. Binary case

A decision problem under uncertainty can be cast in a logical setting in the following way. A vocabulary of propositional variables contains two kinds of variables: decision variables and state variables. Let D be the set of decision variables. Decision variables are controllable, that

is, their value can be fixed by the decision-maker. Making a decision then amounts to fixing the truth value of every decision variable (or possibly just a part of them). On the contrary state variables are fixed by nature, and their value is a matter of knowledge by the decision maker. He has no control on them (although he may express preference about their values).

Let K be a knowledge base (here in propositional logic) describing what is known about the world including constraints relating the decision variables. Let P be another propositional base describing goals delimiting the preferred states of the world. K , and P are assumed to be finite, as is the logical propositional language L under consideration. Assume K and P are classical logic bases, and preferences are all-or-nothing. The aim of the decision problem, described in the logical setting, is to try to make all formulas in the goal set P true by acting on the truth-value of decision variables which control the models of K and P . A good decision d^\wedge (from a pessimistic point of view) is a conjunction of decision literals that entails the satisfaction of every goal in P , when formulas in K are assumed to be true. Therefore, d^\wedge should satisfy

$$K^\wedge \wedge d^\wedge \vdash P^\wedge. \quad (1)$$

Moreover, $K^\wedge \wedge d^\wedge$ must be consistent, for if it is not the case, (1) is trivially satisfied¹. Under an optimistic point of view, we may just look for a decision d^\wedge which is consistent with the knowledge base and the goals, i.e.

$$K^\wedge \wedge d^\wedge \wedge P^\wedge \neq \perp. \quad (2)$$

This is optimistic in the sense that it assumes that goals will be attained as soon as their negation cannot be proved.

The similarity is striking, between the two modes of decision under uncertainty and the two modes of diagnosis reasoning, namely abductive and consistency-based diagnosis solutions (e.g., Hamscher et al. [28]). It is then tempting to encode a logical decision problem under uncertainty by means of techniques coming from the theory of assumption-based truth maintenance systems (ATMS) initiated by De Kleer [11]. In order to better fit the framework of ATMS (the tool we will use to compute optimal decisions), we have to change the encoding of decisions from conjunctions of literals to sets of *positive* literals. This notion will be explained in greater detail in Section 3.

¹In the following we implicitly assume that the result of the decision d^\wedge does not modify the contents of the knowledge base K , which may include for instance pieces of generic conditional knowledge. Clearly, this is not always the case. Just consider a factual knowledge base describing that either the door or the window is open, and the decision: have the door shut (if it is not already the case); obviously we are here facing an updating problem where we should not conclude that the result of the action makes sure that the window is open. So, in the more general case, $K^\wedge \wedge d^\wedge$ should be changed into $K \circ d^\wedge$, where \circ denotes an updating operation, and $K \circ d^\wedge$ is the result of the updating. The study of such an issue is left for further research. The consistency of $K^\wedge \wedge d^\wedge$ should be restricted to the consistency of d^\wedge and the factual part of K when K includes (consistent) generic knowledge also.

2.3. Stratified case

In the logical form of decision problems, the knowledge base may be pervaded with uncertainty, and the goals may not have equal priority. In classical decision theory, uncertainty is represented by means of a probability distribution over the possible states of the world, and the goal states are ranked according to a real-valued utility function. The decision problem amounts to finding a decision that maximizes an expected utility function.

Let us enrich our logical view of the decision problem, by assigning levels of certainty to formulas in the knowledge base, and levels of priority to the goals. Thus we obtain two stratified logical bases that model gradual knowledge and preferences. It has been shown (e.g., Dubois et al. [15]; see also Section 2.4) that a possibility distribution ranking the possible worlds encodes the semantics of a possibilistic logic base, i.e., a stratified base whose formulas are gathered into several layers according to their levels of certainty or priority. First we focus on how a decision problem can be stated, expressing knowledge and preferences in terms of stratified bases. Then we will show that the corresponding semantics of the decision process can be represented by the qualitative utility introduced in (Dubois and Prade [17]).

In the whole paper we will assume that certainty degrees and priority degrees are commensurate, and assessed on the same (finite, as is the language under consideration) linearly ordered scale S . This assumption will be discussed later on. The top element of S will be denoted $\mathbb{1}$, and the bottom element, $\mathbf{0}$. Knowledge and preferences are stored in two distinct possibilistic bases. The knowledge base is $K = \{(\phi_i, \alpha_i)\}$ where $\alpha_i \in S$ ($\alpha_i > \mathbf{0}$) denotes a degree of certainty, and the ϕ_i 's are formulas in L where decision literals may appear. The base expressing preferences or goals is $P = \{(\psi_i, \beta_i)\}$, where $\beta_i \in S$ ($\beta_i > \mathbf{0}$) is a degree of priority, and the ψ_i are formulas of L (where decision literals may also appear).

A question may be raised as to the meaning of the different levels of preference or certainty that are assigned to each sentence. It is clear that the *preference* ordering can be directly given by the decision maker. The *uncertainty* ordering may be assessed by a unique agent classifying the sentences into layers of different levels of certainty. In case the knowledge is given by multiple sources, we can suppose that they have levels of reliability (which may be different), and thus rank the sentences according to the levels of reliability of the sources which provide them (all the information given by a source having the same reliability). On the contrary if the sources are equally reliable, but each of them has its own ordering, we have to suppose that there exists a common agreement on the meaning of the layers of each source, so as to be able to merge the layers of the different sources. Besides, system Z (Pearl [29]) may also help to rank order pieces of generic conditional knowledge by allowing to take the specificity of formulas into account (Benferhat et al. [1]).

Let K_α (resp. P_β) denote the set of formulas with certainty at least equal to α (resp. the formulas with priority at least equal to β). Note that we only consider layers of K (or P) such that $\alpha > \mathbf{0}$ and $\beta > \mathbf{0}$ since $K_{\mathbf{0}} = P_{\mathbf{0}} = L$. In the following we also use the notations K_α and P_β (with $\alpha < \mathbb{1}$ and $\beta < \mathbb{1}$), for denoting the set of formulas with certainty or priority *strictly* greater than α or β respectively. In particular $K_{\mathbf{0}} = K^*$ and $P_{\mathbf{0}} = P^*$ where K^* and P^*

denote the sets of formulas in K and P respectively, without their certainty levels. We shall notice that since the scale S is finite, $K_\alpha = K_{\alpha'}$, where α' is the level of S just above α (the same property holds for P).

Making a decision amounts to choosing a subset d of the decision set $D = \{l_i\}$ where the l_i are distinguished variables of the language L . The corresponding decision d^\wedge is the logical conjunction of literals in the chosen subset. The variables that are not in D are state variables. Our objective is to rank-order decisions by means of a relation \preceq , which will be done by using a utility function $U : \mathcal{P}(D) \rightarrow S$ such that $d \preceq d' \Leftrightarrow U(d) \leq U(d')$. In the following, we will use two different functions: U_* which agrees with a pessimistic view, and U^* which agrees with an optimistic one.

In the first case (pessimistic view), we are interested in finding a decision d (if it exists) such that

$$K_\alpha^\wedge \wedge d^\wedge \vdash P_\beta^\wedge \quad (3)$$

with α high and β low, i.e., such that the decision d together with the most certain part of K entails the satisfaction of the goals, even those with low priority. d is implicitly assumed to be included in the most certain part of $K \cup d$ (certainty level equal to $\mathbb{1}$). Moreover, $K_\alpha^\wedge \wedge d^\wedge$ should be consistent for the α 's satisfying (3). One way of guaranteeing this consistency requirement is to assume $K_{\mathbf{0}}^\wedge \wedge d^\wedge$ is consistent. By convention, utility $\mathbf{0}$ is assigned to every decision d that is not consistent with K^* . Besides, observe that the β satisfying (3) are necessarily such that $\beta > \mathbf{0}$ (since $P_{\mathbf{0}}^\wedge = L$ is inconsistent).

Let n be the order reversing map of scale S . Namely if S is $\mathbf{0} = \alpha_0 < \dots < \alpha_i < \dots < \alpha_n = \mathbb{1}$ then $n(\alpha_i) = \alpha_{n-i}$, for $i = 0, \dots, n$.

We are interested in finding α as high as possible, and β as low as possible such that : $K_\alpha^\wedge \wedge d^\wedge \vdash P_\beta^\wedge$. Ideally, d , along with the most certain part of K only (i.e. $K_{\mathbb{1}}$), should entail every goal in P , even the least preferred ones ($P_{\mathbf{0}}$). Such a decision should have a maximal utility ($\mathbb{1}$). The worst case would be when a decision is unable, even with the whole knowledge($K_{\mathbf{0}}$), to entail at least the most preferred formulas of P (i.e. $P_{\mathbb{1}}$). Such a decision should have a utility of $\mathbf{0}$.

Suppose now that d is such that $K_\alpha^\wedge \wedge d^\wedge \vdash P_\beta^\wedge$, with $\beta < n(\alpha)$. Then we also have $K_\alpha^\wedge \wedge d^\wedge \vdash P_{n(\alpha)}^\wedge$ because $P_{n(\alpha)} \subseteq P_\beta$, and so $P_\beta^\wedge \vdash P_{n(\alpha)}^\wedge$.

Reciprocally, if $\beta > n(\alpha)$ then $n^{-1}(\beta) < \alpha$. Since $K_\alpha \subseteq K_{n^{-1}(\beta)}$, $P_\beta \subseteq P_\beta$ and $K_\alpha^\wedge \wedge d^\wedge \vdash P_\beta^\wedge$, it follows that $K_{n^{-1}(\beta)}^\wedge \wedge d^\wedge \vdash P_\beta^\wedge$. Letting $\beta = n(\alpha)$, we get $K_\alpha^\wedge \wedge d^\wedge \vdash P_{n(\alpha)}^\wedge$. So it is possible to assume that $\beta = n(\alpha)$ in the maximization problem which amounts to maximizing α such that $K_\alpha^\wedge \wedge d^\wedge \vdash P_{n(\alpha)}^\wedge$, with $\alpha > \mathbf{0}$.

Finally, the pessimistic utility of decision d , defined at the syntactic level, takes the form:

Definition 2.1.

$$U_*(d) = \max_{\alpha / K_\alpha^\wedge \wedge d^\wedge \vdash P_{n(\alpha)}^\wedge, K_\alpha^\wedge \wedge d^\wedge \neq \perp} \alpha.$$

and $U_*(d) = \mathbf{0}$ if $\{\alpha > \mathbf{0}, K_\alpha^\wedge \wedge d^\wedge \vdash P_{n(\alpha)}^\wedge \text{ and } K_\alpha^\wedge \wedge d^\wedge \neq \perp\} = \emptyset$.

If now we take the optimistic, we are interested in finding a decision d such that :

$$K_\alpha^\wedge \wedge d^\wedge \wedge P_\beta^\wedge \neq \perp \quad (4)$$

with α and β as low as possible ($\alpha > \mathbf{0}$, $\beta > \mathbf{0}$ in (4)). That is: the preferred states are among the most plausible ones and are also consistent with the decision. The optimistic utility of d is thus given by

Definition 2.2.

$$U^*(d) = \max_{K_\alpha^\wedge \wedge d^\wedge \wedge P_\alpha^\wedge \neq \perp} n(\alpha),$$

and $U^*(d) = \mathbf{0}$ if $\{\alpha < \mathbb{1}, K_\alpha^\wedge \wedge d^\wedge \wedge P_\alpha^\wedge \neq \perp\} = \emptyset$.

Observe that $U^*(d) = \mathbb{1}$ iff $K_{\mathbf{0}}^\wedge \wedge d^\wedge \wedge P_{\mathbf{0}}^\wedge \neq \perp$, that is if the decision is consistent with every goal and piece of knowledge.

2.4. Possibilistic semantics of decision in stratified bases

Let us present the semantics underlying the logical expression of decision problems we have adopted. Interpreting the α_i 's (which are attached to the layers of K) as the degrees of necessity of the formulas in the corresponding layers of $K \cup d$, we compute a possibility distribution π_{K_d} over Ω (the set of all the interpretations of the language L), expressing the semantics of K (see, e.g., (Dubois et al. [15])) :

$$\forall \omega \in \Omega, \pi_{K_d}(\omega) = \min_{(\phi_i, \alpha_i) \in K/\omega \models \neg \phi_i} n(\alpha_i) \text{ if } \omega \models d^\wedge, \text{ and}$$

$$\pi_{K_d}(\omega) = \mathbb{1} \text{ if } \{\phi_i/\omega \models \neg \phi_i\} = \emptyset \text{ and } \pi_{K_d}(\omega) = \mathbf{0} \text{ if } \omega \not\models d^\wedge.$$

The possibility distribution π_{K_d} rank-orders the interpretations according to their level of possibility/plausibility induced by the levels of certainty of the formulas in K . This semantics agrees with the idea that an interpretation ω is all the less possible as it violates formulas ϕ_i with an higher level of certainty α_i . Note that since $K_{\mathbf{0}}^\wedge \wedge d^\wedge$ is supposed to be consistent, π_{K_d} is normalized, i.e., there exists at least an interpretation ω with degree $\pi_{K_d}(\omega) = \mathbb{1}$.

From P , by interpreting the β_i attached to the layers of P as degrees of priority of the formulas in P , we build a utility function μ over Ω in a similar way (ω is all the more satisfactory as it violates no goal with a high priority):

$$\mu(\omega) = \min_{(\psi_j, \beta_j) \in P, \omega \models \neg \psi_j} n(\beta_j).$$

and $\mu(\omega) = \mathbb{1}$ if $\{\psi_j/\omega \models \neg \psi_j\} = \emptyset$.

The two syntactic utility functions defined in Section 2.3 can be expressed in terms of the possibility distribution π_{K_d} and the utility function μ . We have the following results:

Theorem 2.1. *Semantical expression of $(U_*(d))$.*

Let us assume that $K_{\mathbf{0}}^{\wedge} \wedge d^{\wedge}$ is consistent.

$$U_*(d) = \max_{\alpha/K_{\mathbf{0}}^{\wedge} \wedge d^{\wedge} \vdash P_{n(\alpha)}^{\wedge}} \alpha = \min_{\omega \in \Omega} \max(n(\pi_{K_d}(\omega)), \mu(\omega)).$$

Proof: (for a finite scale)

$$\begin{aligned} & - \forall \omega/\omega \models d^{\wedge}, (\omega \models K_{\alpha} \cup d) \Leftrightarrow (\forall (\phi_i, \alpha_i) \in K, \alpha_i \geq \alpha \Rightarrow \omega \models \phi_i) \\ \Leftrightarrow & (\forall (\phi_i, \alpha_i) \in K, \omega \models \neg \phi_i \Rightarrow \alpha_i < \alpha) \\ \Leftrightarrow & (\min_{(\phi_i, \alpha_i) \in K/\omega \models \neg \phi_i} n(\alpha_i) > n(\alpha)) \\ \Leftrightarrow & \pi_{K_d}(\omega) > n(\alpha). \end{aligned}$$

- In the same way we can prove that $(\omega \models P_{n(\alpha)} \Leftrightarrow (\mu(\omega) \geq \alpha))$.

- We use these results in the following :

$$\begin{aligned} & \forall \alpha/\alpha > \mathbf{0}, (K_{\alpha}^{\wedge} \wedge d^{\wedge} \vdash P_{n(\alpha)}^{\wedge}) \\ \Leftrightarrow & (\forall \omega, \omega \models K_{\alpha} \cup d \Rightarrow \omega \models P_{n(\alpha)}) \\ \Leftrightarrow & (\forall \omega, \pi_{K_d}(\omega) > n(\alpha) \Rightarrow \mu(\omega) \geq \alpha) \\ \Leftrightarrow & (\forall \omega, n(\pi_{K_d}(\omega)) < \alpha \Rightarrow \mu(\omega) \geq \alpha) \\ \Leftrightarrow & (\forall \omega, \max(n(\pi_{K_d}(\omega)), \mu(\omega)) \geq \alpha) \\ \Leftrightarrow & (\min_{\omega \in \Omega} \max(n(\pi_{K_d}(\omega)), \mu(\omega)) \geq \alpha). \end{aligned}$$

- Thus we proved :

$$\forall \alpha > \mathbf{0}, (K_{\alpha}^{\wedge} \wedge d^{\wedge} \vdash P_{n(\alpha)}^{\wedge}) \Leftrightarrow (\min_{\omega \in \Omega} \max(n(\pi_{K_d}(\omega)), \mu(\omega)) \geq \alpha).$$

- It is then obvious to show that :

$$\min_{\omega \in \Omega} \max(n(\pi_{K_d}(\omega)), \mu(\omega)) \geq \max_{\alpha/K_{\mathbf{0}}^{\wedge} \wedge d^{\wedge} \vdash P_{n(\alpha)}^{\wedge}} \alpha, \text{ as a limit case.}$$

- The other inequality may be proved by *reductio ab absurdo*, supposing that $\min_{\omega \in \Omega} \max(n(\pi_{K_d}(\omega)), \mu(\omega)) = \beta > \max_{\alpha/K_{\mathbf{0}}^{\wedge} \wedge d^{\wedge} \vdash P_{n(\alpha)}^{\wedge}} \alpha$.

Then, since $\min_{\omega \in \Omega} \max(n(\pi_{K_d}(\omega)), \mu(\omega)) \geq \beta \Rightarrow K_{\beta}^{\wedge} \wedge d^{\wedge} \vdash P_{n(\beta)}^{\wedge}$, we have a contradiction with the assumption we have just made. So we get the result. \square

This result is closely related to an older one by (Prade [31]) expressing the necessity of a fuzzy event in terms of level-cuts of fuzzy sets for the infinite scale $[0,1]$, noticing that Theorem 1 expresses the necessity of a fuzzy event. A similar theorem is easy to prove for the optimistic utility function:

Theorem 2.2. *Semantical expression of $(U^*(d))$.*

$$U^*(d) = \max_{\alpha/K_{\alpha}^{\wedge} \wedge d^{\wedge} \wedge P_{\alpha}^{\wedge} \neq \perp} n(\alpha) = \max_{\omega \in \Omega} \min(\pi_{K_d}(\omega), \mu(\omega)).$$

Proof: (for a finite scale)

$$\begin{aligned} & - \forall \alpha/K_{\alpha}^{\wedge} \wedge d^{\wedge} \wedge P_{\alpha}^{\wedge} \neq \perp, \exists \omega^*/\omega^* \models K_{\alpha} \cup d \cup P_{\alpha} \\ & - \omega^* \models K_{\alpha} \cup d \Leftrightarrow \forall (\phi_i, \alpha_i) \in K, \alpha_i > \alpha \Rightarrow \omega^* \models \phi_i \\ \Leftrightarrow & \forall (\phi_i, \alpha_i) \in K, \omega^* \models \neg \phi_i \Rightarrow \alpha_i \leq \alpha \end{aligned}$$

$$\Leftrightarrow \min_{i, \omega^* \models \neg \phi_i} n(\alpha_i) \geq n(\alpha)$$

$$\Leftrightarrow \pi_{K_d}(\omega^*) \geq n(\alpha)$$

- in the same way we can prove : $\omega^* \models P_\alpha \Leftrightarrow \mu(\omega^*) \geq n(\alpha)$.

Thus, we proved: $\exists \omega^*/\omega^* \models K_\alpha \cup d \cup P_\alpha \Leftrightarrow \min(\pi_{K_d}(\omega^*), \mu(\omega^*)) \geq n(\alpha)$, that is: $\forall \alpha / K_\alpha^\wedge \wedge d^\wedge \wedge P_\alpha^\wedge \neq \perp$, $\max_{\omega \in \Omega} \min(\pi_{K_d}(\omega), \mu(\omega)) \geq n(\alpha)$

- as a limit case: $\max_{\alpha / K_\alpha^\wedge \wedge d^\wedge \wedge P_\alpha^\wedge \neq \perp} n(\alpha) \leq \max_{\omega \in \Omega} \min(\pi_{K_d}(\omega), \mu(\omega))$

- the converse inequality may be proved *ab absurdo* in the same way as in the proof of Theorem 2.1, supposing that

$$\max_{\alpha / K_\alpha^\wedge \wedge d^\wedge \wedge P_\alpha^\wedge \neq \perp} n(\alpha) < \max_{\omega \in \Omega} \min(\pi_{K_d}(\omega), \mu(\omega)) \quad \square$$

The semantical expression of $U_*(d)$ obtained in Theorem 2.1 is exactly the qualitative utility function introduced in (Dubois and Prade [17]). Among the postulates given in [17] so as to justify the pessimistic qualitative utility, some are qualitative counterparts of von Neumann and Morgenstern axioms. Others express the risk aversion of the decision maker. Another one expresses the fact that a one-shot decision is concerned. It emphasizes that the utility of the consequence of the decision, when we know that the state is in A , is of the form $\mu(\omega)$, for some $\omega \in A$ (the worst one for the pessimistic utility). We do not consider (as with expected utility theory), average benefits, gained after repeated actions.

Maximizing $U_*(d)$ means finding a decision d whose highly plausible consequences are among the most preferred ones. The definition of “highly plausible” is decision-dependent and reflects the compromise between high plausibility and low utility expressed by the order-reversing map between the plausibility scale and the utility scale; $U_*(d)$ is small as soon as it exists a possible consequence which is both highly plausible and bad with respect to preferences. This is clearly a risk-averse and thus a pessimistic attitude. When π_{K_d} is the characteristic function of a set A , $U_*(d)$ reduces to:

$$U_*(d) = \min_{\omega \in A} \mu(\omega)$$

which is the Wald criterion, that evaluates the worth of a decision as the worst-case utility. This criterion has been also recently justified in (Brafman and Tennenholtz [4]) in a Savage-like setting. By changing the risk aversion postulate into a risk-prone postulate (Dubois and Prade [18]), the other utility function $U^*(d)$ can be justified. It corresponds to an optimistic attitude since $U^*(d)$ is high as soon as it exists a possible consequence of d which is both highly plausible and highly prized.

3. Computation of decisions

Our purpose is to propose an efficient and unified way of computing both optimistic and pessimistic qualitative decisions. In this section, we give some algorithms based on the use of the MPL procedure (which stands for *Modèles Préférés par leurs Littéraux* in French) described in (Castell et al. [6]) to solve qualitative possibilistic decision problems. The MPL procedure will be briefly described (for a complete description, see (Castell et al. [6])), and it will be shown

how a single pass of MPL allows to compute the decisions consistent with the bases K and P , whereas with two passes of MPL we obtain the “label” of P (which represents the set of “minimal” pessimistic decisions). Notions of ATMS and of a label are described in the following paragraphs.

3.1. ATMS and decision theory

In this section, links between *Assumption-based Truth Maintenance Systems* (ATMS) and pessimistic qualitative decision are formalized. But first of all we restate some basic definitions about ATMS.

3.1.1. Basic definitions of ATMS

The ATMS technique was introduced by (De Kleer [11], [12]). We consider a set of propositional symbols \mathcal{S} divided in two parts, the set of assumptions \mathcal{H} , and the other symbols \mathcal{NH} . A set of assumptions is called an *environment*. An environment E is *inconsistent* for a set of clauses K iff $K \wedge E \vdash \perp$. An environment is *consistent* iff it is not inconsistent.

Definition 3.1. A *nogood* is an inconsistent environment minimal for set-inclusion (i.e., E is a nogood iff $K \wedge E \vdash \perp$ and $\nexists E' \subset E / K \wedge E' \vdash \perp$).

Definition 3.2. The *label* of a formula ψ , denoted $label_K(\psi)$, is the set of all consistent environments E_i minimal for set-inclusion such that $K \wedge E_i \vdash \psi$.

Example 3.1. Let $K = \{A \rightarrow b, B \rightarrow b, c \rightarrow\}$. $\mathcal{S} = \{A, B, b, c\}$. $\mathcal{H} = \{A, B\}$. $\{\{A, B\}\}$ is the set of K 's nogoods. $label_K(b) = \{\{B\}\}$, $label_K(c) = \{\}$, $label_K(A) = \{\{A\}\}$, $label_K(B) = \{\{B\}\}$.

ATMS *assumptions* are useful for computing decisions: assumptions are distinguished positive literals, therefore decisions will be modelled by sets of distinguished *positive* literals in D , consistent with the constraints in K . When d contains several positive literals, d should not be interpreted as a sequence of decisions but as a single decision consisting in assigning simultaneously a positive truth value to every literal in d . For any such decision denoted d , d^\wedge denotes the logical conjunction of the *positive* literals in d . The set of all decisions which obey (1) and such that none of their proper subsets (when they are given in the form of subsets of D) obeys (1), can then be seen as an extension of De Kleer's notion of label of a literal [11], to the notion of label of a conjunction of formulas (of the form $label_K(P^\wedge)$).

Example 3.2. The available decisions are to buy zero, one or two items. The decision set will be $D = \{Zero, One, Two\}$, and K should contain the following constraints expressing the mutual exclusiveness of the available decisions: $\{One \vee Two \vee Zero, \neg One \vee \neg Two, \neg One \vee \neg Zero, \neg Two \vee \neg Zero\} \subseteq K$. For this example, the decisions that are consistent with the constraints are: $d = \{One\}$, $d = \{Two\}$ or $d = \{Zero\}$.

²These definitions are slightly different from De Kleer's definitions. For instance, the label notion was originally defined for a literal.

3.1.2. Using an ATMS in qualitative possibilistic decision theory

We propose to translate our decision problem into a problem tractable by an ATMS. Let us define the set of assumptions symbols $\mathcal{H} = D$. Then, assume that K is the knowledge base of the decision problem in conjunctive normal form and consider the goal base P as a formula P^\wedge . Using the symbols in \mathcal{H} , a decision d is a subset of \mathcal{H} . For any decision d such that $K^\wedge \wedge d^\wedge \vdash P^\wedge$ and $K^\wedge \wedge d^\wedge \neq \perp$ there is at least one element E of $label_K(P)$ according to the assumption set \mathcal{H} such that $E \subseteq d$.

Proposition 3.1.

$$K^\wedge \wedge d^\wedge \vdash P^\wedge \text{ and } K^\wedge \wedge d^\wedge \neq \perp \text{ if and only if } \exists E \in label_K(P) \text{ s.t. } E \subseteq d.$$

Proof: immediate from the definition of a label.

A good (pessimistic) decision is then a superset of an element from $label_K(P)$. In the following we will only look for decisions which are minimal for set inclusion.

Let $K^\wedge = \phi_1 \wedge \phi_2 \dots \wedge \phi_n$ and $P^\wedge = \psi_1 \wedge \psi_2 \wedge \dots \wedge \psi_m$. Finding all decisions d maximizing α such that: $K_\alpha^\wedge \wedge d^\wedge \vdash P_{n(\alpha)}^\wedge$ and $K_\alpha^\wedge \wedge d^\wedge \neq \perp$ is equivalent to finding $label_{K_\alpha}(P_{n(\alpha)}) \neq \{\}$ maximizing α .

3.2. The MPL procedure

Let us present here the MPL procedure introduced in (Castell et al. [6]). This procedure does the following: given a logical formula ϕ in conjunctive normal form, involving two types of literals, it computes its projection by restricting to one type of literals; and this projection is the most informative such consequence of ϕ expressed in disjunctive normal form. It is shown that nogoods in an ATMS are easily obtained by means of this procedure. The definitions are slightly different from the ones in (Castell et al. [6]) but the principle is still the same. This new formalization improves the clarity of the proofs. Further information, proofs and links between the two formalisms can be found in (Le Berre [25] and Castell et al. [7]).

3.2.1. Definitions

Let E be a set of literals. $\sim E$ denotes $\{\neg x \mid x \in E\}$ (where $\neg\neg x$ simplifies into x ; e.g.: $E = \{\neg c, a, b\}$, $\sim E = \{c, \neg a, \neg b\}$)³.

Definition 3.3. (implicants, implicates) Let ϕ be a propositional formula. Let E, F be consistent sets of literals.

- The conjunction of literals (or *phrase*) E^\wedge is an *implicant* of ϕ iff $E^\wedge \vdash \phi$. E^\wedge is a *prime implicant* of ϕ iff $E^\wedge \vdash \phi$ and $\nexists F^\wedge$ an implicant of ϕ such that $F \subset E$.

³*Interpretations* and *models* are used in propositional calculus sense. They are represented either by the set of their literals assigned to *True* or as a conjunction of literals.

- The disjunction of literals (or *clause*) E^\vee is an *implicate* of ϕ iff $\phi \vdash E^\vee$. E^\vee is a *prime implicate* of ϕ iff $\phi \vdash E^\vee$ and $\nexists F^\vee$ an implicate of ϕ such that $F \subset E$.

In the following, H will denote a consistent subset of distinguished literals of the language L , that will be used to restrict the language to a subset of propositional variables.

Definition 3.4. (restriction) Let M be a subset of literals corresponding to an interpretation. We call *restriction to H* of M the set $H \cap M$, denoted $R_H(M)$.

Clearly, if M is viewed as a conjunction of literals, $R_H(M)$ is a phrase where only literals in H appear. It can be viewed as a partial interpretation. The following definition replaces the preferred models definition of (Castell et al. [6]). A preferred \mathcal{H} -model was previously defined as a complete assignment of truth value to a subset \mathcal{H} of the symbols, containing a maximal number of “preferred symbols”.

Definition 3.5. (restricted models) Let ϕ be a formula. Let $I \subseteq H$. I is a H -restricted model of ϕ iff $\exists M$ a model of ϕ such that $I = R_H(M)$.

The following results are borrowed from Castell et al. [6][7].

Property 1. I is a H -restricted model of ϕ iff the phrase $(I \cup \sim (H \perp I))^\wedge$ is consistent with ϕ .

Indeed since H is consistent, if $(I \cup \sim (H \perp I))^\wedge$ were inconsistent, it means that there would be a literal both in I and in $H \perp I$. Now, $(I \cup \sim (H \perp I))^\wedge$ consistent with ϕ is equivalent to the existence of a common model ⁴.

Theorem 3.1. I is a H -restricted model of $\phi \Leftrightarrow \exists E^\wedge$ an implicant of ϕ such that $H \cap E = I$.

To prove the \Rightarrow side, just choose as an implicant the phrase built from a common model of I^\wedge and ϕ . Conversely any implicant can be extended to a model. A H -restricted model of ϕ can thus be obtained by considering any of its implicants E^\wedge and masking the literals not in H . Note that the obtained phrase is not an implicant of ϕ .

Definition 3.6. Let ϕ be a formula. Let $I \subseteq H$. I is a H -implicant of ϕ iff $\forall M$, an interpretation such that $I \subseteq M$, M is a model of ϕ .

It means that a H -implicant of ϕ is a H -restricted model which implies ϕ . One can prove indeed that the minimal (for set-inclusion) H -implicants of ϕ are exactly the prime implicants of ϕ included in H . In fact the minimal H -restricted models can be retrieved from a DNF:

Theorem 3.2. *The set of minimal H -restricted models of a DNF $\phi = E_1^\wedge \vee E_2^\wedge \vee \dots \vee E_n^\wedge$ is the set of inclusion-minimal elements of $R_H(E_1), R_H(E_2), \dots, R_H(E_n)$*

⁴Let \mathcal{D} be the set of symbols involved in H . $I \cup \sim (H \perp I)$ is exactly a preferred \mathcal{D} -model of ϕ in the sense of [6].

Proof: It is easy to show that the $R_H(E_i)$ are H -restricted models of ϕ . Suppose now that it exists E , a minimal H -restricted model of ϕ that is not in the set of inclusion-minimal elements of $R_H(E_1), R_H(E_2), \dots, R_H(E_n)$. So, it exists M , model of ϕ , such that $R_H(M) = E$, and $\forall i, R_H(E) \neq R_H(E_i)$. Now, as it is impossible that $R_H(E_i) \subset E$ (minimality of E), there exists a literal $l \in R_H(E_i)$ and $l \notin E$, and this for all i . We have $l \notin M$, because $l \in H$, and $E = R_H(M) = M \cap H$. So, M falsifies the $R_H(E_i)$, and so falsifies ϕ , which is contradictory. \square

So computing the H -restricted models is obvious if the formula ϕ is in disjunctive normal form (DNF).

Example: Let $K = \{\neg A \vee c, \neg A \vee B, \neg B \vee \neg c \vee A\}$ and $H = \{\neg A, \neg B\}$. The disjunctive normal form of K is $(\neg B \wedge \neg A) \vee (\neg A \wedge \neg c) \vee (A \wedge B \wedge c)$. The H -restricted models are $\{\neg A, \neg B\}$, $\{\neg A\}$, and $\{\}$. The only H -implicant of K is $\{\neg A, \neg B\}$.

We can define a function MP_H (resp. PI_H) that from any formula ϕ of the language computes another formula made of literals in H , in DNF form, such that the prime implicants (elements of the disjunction) of $MP_H(\phi)$ (resp. $PI_H(\phi)$) are exactly the minimal (for the inclusion of sets of literals) H -restricted models (resp. H -implicants) of ϕ . In the example, we find respectively $\{\}$ (which denotes the tautology \top) and $\{\neg A, \neg B\}$. In the example, we have $MP_H(\phi) = \top$ and $PI_H(\phi) = \neg A \wedge \neg B$. Clearly, the following properties hold.

Property 2. $\forall \phi, PI_H(\phi) \vdash \phi \vdash MP_H(\phi)$.

Proof: If E^\wedge is an element of the disjunction $PI_H(\phi)$, it is a prime implicant of ϕ , which proves the first \vdash . Now, any model M of ϕ defines a H -restricted model $R_H(M)$. So, there exists E^\wedge an element of the disjunction $MP_H(\phi)$, such that $E \subseteq R_H(M)$, so the second \vdash is proved. \square

Theorem 3.3. Let $E \subseteq H$. $\forall \phi, \phi \vdash E^\vee$ iff $MP_H(\phi) \vdash E^\vee$.

Proof: (\Leftarrow) Let M be a model of ϕ . $R_H(M)$ is a H -restricted model of ϕ , so it contains an element of $MP_H(\phi)$. A fortiori, M contains an element of $MP_H(\phi)$, so M models $MP_H(\phi)$. As $MP_H(\phi) \vdash E^\vee$, we get that M models E^\vee .

(\Rightarrow) We have $\phi \vdash E^\vee$. Suppose it false that $MP_H(\phi) \vdash E^\vee$. So, there exists E_i^\wedge in $MP_H(\phi)$, such that $E_i \cap E = \emptyset$. By definition, $\exists M$, model of ϕ , such that $E_i = R_H(M)$, so $E_i \subseteq M$. As $E \subseteq H$, $E \cap M \subseteq E \cap P \cap M = E \cap E_i = \emptyset$, so M falsifies E^\vee , which is a contradiction. \square

Theorem 3.4. Let $E \subseteq H$. $\forall \phi, E^\wedge \vdash \phi$ iff $E^\wedge \vdash PI_H(\phi)$.

Proof: (\Leftarrow) is obvious.

(\Rightarrow) As $E^\wedge \vdash \phi$, it contains a prime implicant F^\wedge of ϕ . As $E \subseteq H$, F^\wedge is in $PI_H(\phi)$, and so $E^\wedge \vdash PI_H(\phi)$. \square

Corollary 3.1. If all the prime implicants of ϕ are in H then $MP_H(\phi) \equiv PI_H(\phi)$

This corollary is very important because that is the reason why we can easily compute labels and nogoods in an ATMS using the MP_H function when H is a set of hypothese literals. In a classical implicant/implicate calculus, this condition is not satisfied. However, a similar, but extended, approach is introduced in (Castell and Cayrol [5]), to compute prime implicates/implicants.

3.2.2. Principle of the MPL algorithm

The famous problem in complexity theory, called SAT, for SATisfiability of a set of clauses, is NP-complete. Some of the best complete algorithms which solve SAT, (i.e. C-SAT [13], SATZ [27],...), are based on an enumerative approach, the Davis and Putnam algorithm [9] (also called Davis-Putnam-Loveland procedure [10]).

A Davis and Putnam algorithm enumerates the interpretations of a knowledge base K until it finds a model (consistent case) if any (the inconsistent case is when it finds no model). So doing, it is obvious that searching for models is closely related to finding a disjunctive normal form for a knowledge base, since it is easy to exhibit models of a DNF. One of the differences with our algorithm MPL is that MPL does not stop after the first model of K , and rather looks for all the minimal H -restricted models of K .

Davis and Putnam's algorithm builds a binary search tree such that at each node, it branches on the truth value of a literal. So, using a linear ordering relation over the set of literals, we can rank the interpretations I w.r.t. the subset ordering of their restriction to H , $R_H(I)$: at each node, we will first branch to falsify a literal from H . Then, if M is the first interpretation which satisfies the given set of clauses, $R_H(M)$ denotes a minimal H -restricted model.

Now, the problem is “to eliminate the H -restricted models which are not minimal within our algorithm”. Because of the use of subset ordering, the idea is to add a clause that non minimal H -restricted models will falsify. Then, all the H -restricted models returned by our algorithm will be minimal ones. For this purpose, after having found a model M , we add a clause $C = \bigvee_{l \in R_H(M)} \neg l$. Note that each interpretation I such that $R_H(M) \subset R_H(I)$ falsifies C .

The following algorithm implements our modified Davis and Putnam procedure. We divide the set of clauses in two parts: K is the original set of clauses and KA is the set of added clauses. This algorithm, invoked with $K, KA = \{\}$ and H as parameters, computes exactly $MP_H(K)$.

Algorithm 1: $MPL(K,KA,H)$

% this function returns the minimal H -restricted models of K ;

Data: K a set of clauses ;

KA a set of added clauses ;

H a consistent set of literals ;

Result: M a set of minimal H -restricted models

begin

$M \leftarrow \{\}$ % M is a local variable ;

 MODEL_PREF LIT($K,KA,\{\},M,H$);

return M ;

end

This algorithm is a modified version of the Davis and Putnam algorithm. It produces a binary search tree, which is developed (the variables in H are instantiated successively) until a model of K is found. The usual algorithm of Davis and Putnam stops at this point, but as we want to find all the minimal H -restricted models of K , we have to go on and search for the

other models of K . But as we search for *minimal* models, when a model M is found, we add a new clause $\{\vee_{l \in M \cap H} \neg l\}$ in the set of added clauses, KA . This insures that when building the other models of K , we will not find one that contains M . In this way, when the algorithm is ran till its end, it produces all the minimal H -restricted models of K , and only them.

In the function `SIMPLIFY()` we use classical techniques of Davis and Putnam based algorithms such as unit propagation, restricted resolution. We don't take into account pure literals⁵, because if we do, the resulting set of clauses is not logically equivalent to the initial one (the consistency is preserved, but it is not sufficient for our purpose).

The function `CHOOSE LITERAL()` returns a literal from H which is not yet assigned if possible, else the word "nothing". We can use the work made in the SAT community about heuristics to choose this literal (for instance, C-SAT [13], Unit Propagation heuristics [27], etc).

Algorithm 2: MODEL PREF LIT(K , var KA , IP , var M, H)


```

Data:  $K$  a set of clauses ;
       $KA$  a set of added clauses ;
       $IP$  a consistent set of literals ;
       $H$  a consistent set of literals ;
Result:  $M$  a set of minimal  $H$ -restricted models
begin
  SIMPLIFY( $K, IP$ ) ;
  %  $\rightarrow$  is the empty clause ;
  if  $(K)_{IP} = \{\}$  and  $\rightarrow \notin (KA)_{IP}$  then
    % we have found an implicant of  $K$  ;
     $M \leftarrow M \cup \{IP \cap H\}$  ;
     $KA \leftarrow KA \cup \{\vee_{l \in IP \cap H} \neg l\}$  ;
  else
    if  $\rightarrow \notin (K \cup KA)_{IP}$  then
      % we look for a literal of  $H$  in  $K$  ;
       $l \leftarrow$  CHOOSE LITERAL( $K, IP, H$ ) ;
      if  $l \neq$  "nothing" then
        % we begin with the negation of the literal ;
        MODEL PREF LIT( $K, KA, IP \cup \{\neg l\}, M, H$ );
        MODEL PREF LIT( $K, KA, IP \cup \{l\}, M, H$ );
      else
        % we just make a consistence test ;
        if  $(K \cup KA)_{IP}$  consistent then
           $M \leftarrow M \cup \{IP \cap H\}$  ;
           $KA \leftarrow KA \cup \{\vee_{l \in IP \cap H} \neg l\}$  ;
    end
  end

```

⁵a literal is *pure* in a formula iff its complementary literal does not appear in this formula.

Let us take an example: let $K = \{B \vee c, \neg c \vee A \vee B\}$. Let $H = \{A, B\}$. The next figure shows the binary tree developed by our algorithm. It can be divided into two steps. First, we use the modified Davis and Putnam algorithm described before, to assign a truth value to all literals from H involved in the formula. Then, we only need to know if the resulting formula is consistent or not, thanks to Theorem 3.1. The consistency test is performed by a usual Davis and Putnam algorithm. If the formula is consistent, we have a minimal H -restricted model. The restriction to H of the first model found is $\{B\}$. So, we add in the knowledge base the clause $\neg B$. In a same way, we add the clause $\neg A$. This clause prunes the last branch. So the two minimal H -restricted models are $\{A\}$ and $\{B\}$.



3.2.3. Application to ATMS

Now, we show how to use the algorithm $MPL()$ in the framework of ATMS [11][12]. Indeed, the basic elements of an ATMS, labels, nogoods, will be efficiently computed by the algorithm, without any minimization step contrary to De Kleer's original one.

Proposition 3.2. *The set of nogoods of a knowledge base K with respect to a set of hypotheses \mathcal{H} is exactly the set of prime implicants of $MP_{\mathcal{H}}(\neg MP_{\sim \mathcal{H}}(K^{\wedge}))$.*

proof: E is a nogood iff $E \subseteq \mathcal{H}$ is minimal such that $K^{\wedge} \wedge E^{\wedge} \vdash \perp$, or $K^{\wedge} \vdash (\sim E)^{\vee}$.

$K^{\wedge} \vdash (\sim E)^{\vee}$ iff $MP_{\sim \mathcal{H}}(K^{\wedge}) \vdash (\sim E)^{\vee}$ (Theorem 3.3),

$MP_{\sim \mathcal{H}}(K^{\wedge}) \vdash (\sim E)^{\vee} \equiv E^{\wedge} \vdash \neg MP_{\sim \mathcal{H}}(K^{\wedge})$

$E^{\wedge} \vdash \neg MP_{\sim \mathcal{H}}(K^{\wedge})$ iff $E^{\wedge} \vdash PI_{\mathcal{H}}(\neg MP_{\sim \mathcal{H}}(K^{\wedge}))$ (Theorem 3.4).

Or equivalently, $E^{\wedge} \vdash MP_{\mathcal{H}}(\neg MP_{\sim \mathcal{H}}(K^{\wedge}))$ (corollary 3.1 since the literals of $\neg MP_{\sim \mathcal{H}}(K^{\wedge})$ are in \mathcal{H}). So, a nogood is exactly a prime implicant of $MP_{\mathcal{H}}(\neg MP_{\sim \mathcal{H}}(K^{\wedge}))$. \square

Proposition 3.3. *Let K be a set of clauses and a set of hypotheses \mathcal{H} . Let ϕ be a formula. The label of ϕ exactly contains the set of prime implicants of $MP_{\mathcal{H}}(\neg MP_{\sim\mathcal{H}}(K \wedge \neg\phi))$ that are not among the nogoods of K .*

The proof is similar to the preceding one since we are then interested in the nogoods of $K \cup \neg\phi$. We have to remove the nogoods from the label because every formula is logical consequence of an inconsistent one. This can be done by initializing KA with the nogoods, in the MPL algorithm ⁶.

Proposition 3.4. *Let K be a set of clauses. Let ϕ and ψ two formulas. The label of $\phi \wedge \psi$ is exactly the prime implicants of $MP_{\mathcal{H}}(\neg MP_{\sim\mathcal{H}}(K \wedge \neg\phi) \wedge \neg MP_{\sim\mathcal{H}}(K \wedge \neg\psi))$ (except Nogoods of K).*

Example 3.1, continued $K = \{\neg A \vee \neg b \vee c, \neg B \vee b, \neg c\}$

$$MP_{\sim\mathcal{H}}(K) = \neg A \vee \neg B \quad (\{\{\neg A\}, \{\neg B\}\}),$$

$$\neg MP_{\sim\mathcal{H}}(K) = A \wedge B \quad (\{\{A\}, \{B\}\}),$$

$$MP_{\mathcal{H}}(\neg MP_{\sim\mathcal{H}}(K)) = A \wedge B \quad (\{\{A, B\}\}).$$

So the set of nogood is $\{\{A, B\}\}$.

Let us point out the fact that *only the algorithm* $MPL()$ is used to compute labels and nogoods. Difficult operations like subsumption are not explicitly performed for these computations.

3.3. Computation of optimistic optimal decisions via MPL

The use of MPL to solve an optimistic decision problem is easy. Assuming that K and P are CNF representations of knowledge and goal bases ⁷ of the decision problem, for a decision d such that $K \wedge d \wedge P \neq \perp$ there is at least one element E of $MPL(K \cup P, \{\}, D)$ such that $E \subseteq d$.

Proposition 3.5.

$$K \wedge d \wedge P \neq \perp \equiv \exists E \in MPL(K \cup P, \{\}, D) \text{ s.t. } E \subseteq d.$$

Proof: immediate from the definition of minimal D -restricted models. \square

A good (optimistic) decision is then a consistent superset of an element from $MPL(K \cup P, \{\}, D)$. In the following we will only look for decisions which are minimal for set inclusion.

Let $K \wedge = \phi_1 \wedge \phi_2 \dots \wedge \phi_n$ and $P \wedge = \psi_1 \wedge \psi_2 \wedge \dots \wedge \psi_m$. Finding d maximizing $U^*(d) = n(\alpha)$ such that: $K \wedge \wedge d \wedge P \wedge \neq \perp$ (indeed, definition 2.2 is equivalent to finding $MPL(K \wedge \cup P \wedge, \{\}, D) \neq \{\}$

⁶We have to compute nogoods before any label computation. The set of added clauses KA contains exactly $\neg MP_{\mathcal{H}}(K \wedge)$ after a call to $MPL(K, KA = \{\}, H)$.

⁷A stratified possibilistic knowledge base can always be put in an equivalent base of weighted clauses [15], since necessity measures are min-decomposable for conjunction.

minimizing α).

This method has only one restriction: $K \wedge P^\wedge$ must be a CNF formula, so P^\wedge must be a CNF formula, that is P must contain only clauses.

The following algorithm computes the best decisions given K and P , according to the optimistic utility:

Algorithm 3: OPTIMISTIC

```

% very simple algorithm to compute optimal optimistic decisions ;
Data: K a stratified knowledge base, P a gradual preference base, D the set of decision
      symbols.
Result: -  $\alpha^*$  such that  $n(\alpha^*)$  is the utility of the best optimistic decision,
        -  $\mathcal{D}$  the set of the best optimistic decisions.
begin
   $\alpha \leftarrow \mathbf{0}$  % we consider all the knowledge ;
   $\mathcal{D} \leftarrow MPL(K_\alpha \cup P_\alpha, \{\}, D)$  ;
  while  $\mathcal{D} = \{\}$  and  $\alpha < \mathbf{1}$  do
     $Inc(\alpha)$  % Eliminate the least certain strate remaining in  $K$  and  $P$  ;
     $\mathcal{D} \leftarrow MPL(K_\alpha \cup P_\alpha, \{\}, D)$  ;
  return  $\langle n(\alpha), \mathcal{D} \rangle$  ;
end

```

3.4. Computation of pessimistic optimal decisions via MPL

It was shown in the previous sections that nogoods and labels of an ATMS set can be obtained after one or two invocations of the MPL algorithm.

The main advantage of the MPL-technique is its ability to compute the label of a unique literal without computing the labels of the other literals as with De Kleer's technique. Moreover, an MPL-based ATMS can be applied on any set of clauses (CNF formula) and can compute in the same way the label of a literal, a disjunction or a conjunction of literals.

The label computation presupposes a computation of the nogoods, in order to remove from the label the inconsistent environments. Nogoods and labels are computed in the same way, from the knowledge base (K) for no-goods, and from the knowledge base augmented with the negation of the formula, $(K \wedge \neg\phi)$, for the label of this formula.

We are now able to describe an algorithm for the pessimistic case. We need to compute the nogoods, and then the required label. The restriction here is to have $K \wedge (\sim P)^\wedge$ as a CNF formula, so P^\vee being a DNF formula. Since P is a CNF, this procedure will accept only P as a single clause or a single phrase (both are CNF and DNF form). Thus, we have to use a particularity of MPL to compute the label of a conjunction of formulas: the label of a conjunction $\psi \wedge \phi$ can be performed from the two first steps needed to compute the label of both ψ and ϕ (Proposition 3.4). This approach allows to stop label computation as soon as the intermediate

label is empty.

Algorithm 4: COMPUTE PESSIMISTIC DECISION

```

% the final algorithm, based on the MPL algorithm ;
% Let  $P_{(\beta)} = P_{\beta} \perp P_{\beta}$ .  $P_{(\beta)}$  is the set of preferences with priority level  $\beta$ ;
%  $\alpha$  denotes the level (either of certainty or of priority);
% of the next non-empty layer below  $\alpha$  ;
Data:  $K$  the knowledge base,
       $P$  the preference base,
       $D$  the set of decision symbols.
Result:  $\alpha^*$  the utility of the best pessimistic decisions,
         $D$  the set of the best pessimistic decisions.
begin
   $\alpha \leftarrow \mathbb{1}$  % we consider the most certain layer ;
   $S \leftarrow \{\}$  ;
  while  $S = \{\}$  and  $\alpha > \mathbf{0}$  do
    % we must first compute the nogoods of  $K_{\alpha}$  ;
     $NG' \leftarrow MPL(K_{\alpha}, \{\}, \sim D)$  % first call to MPL ;
     $NG \leftarrow \sim MPL(\neg NG', \{\}, D)$  % NG contains the nogoods of  $K_{\alpha}$  ;
     $S' \leftarrow MPL(K_{\alpha} \cup \sim P_{\mathbb{1}}, \{\}, \sim D)$  % ;
     $S \leftarrow \sim MPL(\neg S', \sim NG, D)$  % S contains the label of  $P_{\mathbb{1}}$  ;
     $\beta \leftarrow \mathbb{1}$  ;
    while  $\beta > n(\alpha)$  and  $S \neq \{\}$  do
       $S' \leftarrow S' \wedge MPL(K_{\alpha} \cup \sim P_{(\beta)}, \{\}, \sim D)$  ;
      % S contains the label of  $P_{\beta}$  ;
       $S \leftarrow MPL(\neg S', \sim NG, D)$  ;
       $Dec(\beta)$  ;
    if  $S = \{\}$  then  $\alpha \leftarrow \max(\alpha, n(\beta))$  ;
  return  $\langle \alpha, S \rangle$  ;
end

```

Let us briefly illustrate the behavior of this algorithm on an example:

Example: K and P contain 5 layers (both scales are commensurate). First of all, we consider only the most certain layer of K , and we compute successively the labels of P_4 , P_3 , P_2 , and finally the label of P_1 which is the first one to be found empty (see Figure 1). Now α takes the value $\max(\alpha, n(\beta)) = \max(3, 3) = 3$. Then, we consider K_3 and compute the labels. Once again, the labels of P_4 , P_3 , P_2 , are found non-empty (see Figure 2), then the value of β is 1, which is not strictly greater than $n(\alpha) = 3$. Therefore, α remains unchanged, and as $S \neq \{\}$, the algorithm returns α and S .




Figure 1. Example




Figure 2. Example (continued)

3.5. Conclusions about the implementation

It must be pointed out that the algorithm does not make the assumption that $K_{\mathbf{0}}^{\wedge} \wedge d^{\wedge}$ shall be consistent: For each value of α , the nogoods of K_{α} are computed, and the decisions returned by the algorithm with a utility α are guaranteed to be consistent with K_{α} . Assuming (as is suggested earlier in this paper), that $K_{\mathbf{0}}^{\wedge} \wedge d^{\wedge}$ be consistent would simplify the algorithm, by making the nogoods computation useless. Furthermore, when a new layer of knowledge is considered in the algorithm, it is useless to re-compute the labels of the P_{β} that were previously found non-empty: the augmented knowledge base, as it is consistent, together with the previous *candidate decisions*, keeps on allowing to deduce P_{β} . Then we shall not reinitialize β to $\mathbf{1}$ in the loop. Anyway, the algorithm proposed here, by forgetting the *consistency assumption*, allows to cope with inconsistent knowledge bases.

One of the major advantages of our approach is that we only need to implement the MPL algorithm. An efficient implementation of MPL entails an efficient implementation of the decision algorithm. Thanks to the relation between the MPL algorithm and the Davis and Putnam

algorithm, some improvements on the latter can be used in the former (heuristics for instance). Another advantage of this technique is the ability of an MPL-based ATMS to compute the label of a phrase (a conjunction of literals) or a clause. So we can compute the label of each preference clause in a simple way.

The anytime aspect of the MPL algorithm can be pointed out here. If you stop the algorithm before its normal end, you may obtain a subset of the set of optimal decisions. This can be used for instance if we only need:

- a single optimal decision, or
- the utility of the optimal decision(s).

4. Example

4.1. The problem

Consider (Savage [32], pp. 13-15)'s omelette example. The problem is about deciding whether or not adding an egg to a 5-egg omelette.

The possible states of the world are: *The egg is good* (denoted g), and *The egg is rotten* denoted r . The available acts are: *Break the egg in the omelette* (BIO), *Break it apart in a cup* (BAC), and *Throw it away* (TA). The other literals used for expressing knowledge are: $6e$ (meaning that we obtain a 6-egg omelette), $5e$ (we obtain a 5-egg omelette), wo (the omelette is wasted), we (an egg is wasted), and c (we have a cup to wash). The possible consequences are: $6e$ if g holds and we choose BIO ; $6e \wedge c$ if g holds and we choose BAC , ($6c$); $5e$ if r holds and we choose TA , (5); $5e \wedge c$ if r holds and we choose BAC , ($5c$); $5e \wedge we$ if g holds and we choose TA , ($5w$); and wo if r holds and we choose BIO , (w).

The uncertain part of the knowledge base consists only in our opinion on the state of freshness of the egg.

Concerning the preferences: first of all, we do not want to waste the omelette, then if possible, we prefer not to waste an egg. Then, if possible, we prefer to avoid having a cup to wash if the egg is rotten (that is, if it would have been better to throw it away directly). Finally, if all these preferences are satisfied, then we prefer to have a 6-egg omelette, and the best situation would be to have, in addition, no cup to wash.

4.2. Knowledge and preferences bases

From the expression of the problem given as above, we can construct two stratified bases of formulas: the knowledge base K , and the preferences base P .

Let us use the scale $\{0, 1, 2, 3, 4, 5\}$ for assessing the certainty levels and preferences, where $\mathbb{1} = 5$ and $\mathbb{0} = 0$. Just notice that we could have used linguistic values instead of numbers: only comparison and order-reversing are meaningful operations here. In terms of priority-valued formulas, we get the following base $P = \{(\neg wo, 5), (\neg we, 4), (\neg c \vee \neg 5e, 3), (\neg 5e, 2), (\neg c, 1)\}$.

The preferences could alternately be expressed by the means of a *semantical* utility function μ . Namely, μ can be computed from the priority-valued formulas form of the knowledge base (see Section 2.3, and the Section 4.4 that follows): $\mu(\omega) = \min\{n(\beta_j) \text{ s.t. } (\psi_j, \beta_j) \in P \text{ and } \omega \models \neg\psi_j\}$.

The utilities assigned to the consequences would be, using this property: $\mu(6e) = 5, \mu(6c) = 4, \mu(5e) = 3, \mu(5c) = 2, \mu(5w) = 1, \mu(w) = 0$.

The two following stratified sets of clauses represent knowledge and preferences and can be used as input files for our program.

```
// Knowledge base K
// ##### i represents the ith layer

##### 5

// decisions BIO, BAC and TA are mutually exclusive

-> BIO BAC TA ;
BIO BAC -> ;
BIO TA -> ;
BAC TA -> ;

// we get a 6-egg omelette if and only if the egg is good and we
//break it in the omelette;

g BIO -> 6e ;
g BAC -> 6e ;
6e -> g ;
6e TA -> ;

// if we break the egg apart and it is rotten or if we throw it
//away we get a 5-egg omelette

TA -> 5e ;
r BAC -> 5e ;
BAC -> cw ;
cw -> BAC ;
5e -> TA r ;
5e BIO -> ;

// An egg is wasted if and only if we throw away a good egg

g TA -> we ;
we -> g ;
we -> TA ;

// the omelette is wasted if and only if we break a rotten egg in it
```

```

r BIO -> w ;
w -> BIO ;
w -> r ;

// an egg is either good or rotten

g r -> ;
-> g r ;

##### 2

// in this example, we are slightly convinced that the egg is good

-> g ;

// Preference base P

##### 5

w -> ;

##### 4

we -> ;

##### 3

5e cw -> ;

##### 2

5e -> ;

##### 1

cw -> ;

```

Remarks:

- Notations: these two files are used in the above form as input files for our program. ##### i represents the beginning of the expression of the i th layer, either of the knowledge base, or of the preference base. The scale used for assessing preference and certainty levels is determined from the highest layer number. If it is n , then the scale is $\{0, \dots, n\}$. The text

after // is a commentary. Pieces of knowledge and preference are expressed by the means of clauses, separated by ;, \rightarrow is the material implication, the left part is in conjunctive form, the right part in disjunctive form (e.g.: $BIO \ BAC \rightarrow \equiv \neg(BIO \wedge BAC) \equiv \neg BIO \vee \neg BAC$, $6e \rightarrow BIO \ BAC \equiv \neg 6e \vee BIO \vee BAC$).

- The two bases, and especially the knowledge base in this example, may express more than what is really necessary for computing optimal decisions. Anyway, usually the decision maker is not able to distinguish the knowledge that will be useful for the decision problem, from the one that will not be of any use. The decision maker is only concerned with giving *enough* knowledge for the program to compute optimal (pessimistic or optimistic) decisions. Furthermore, the knowledge or the preference base may be redundant, which is of no importance for the decision problem.

In the following paragraph, we will see how the algorithm works on this example.

4.3. Computation with MPL

Optimistic case

First step : $\alpha = \mathbb{0} = 0$.

The Computation $MPL(K_{\mathbb{0}} \cup P_{\mathbb{0}}, \{\}, D)$ gives us one solution $\{BIO\}$.

The optimistic utility of this solution is $n(\alpha) = \mathbb{1}$.

Pessimistic case

First step : $\alpha = \mathbb{1} = 5$.

Computation of the nogoods : $NG = \{\{BIO, BAC\}, \{BIO, TA\}, \{BAC, TA\}\}$.

Computation of the label of $\neg w \wedge \neg we \wedge (\neg 5e \vee \neg cw) \wedge \neg 5e \wedge \neg cw$ ($n(\alpha) = 1$):

Label $_{K_{\alpha}}$ ($\neg w$) = $\{\{BAC\}, \{TA\}\}$. $\beta = 5$.

Label $_{K_{\alpha}}$ ($\neg w \wedge \neg we$) = $\{\{BAC\}\}$. $\beta = 4$.

Label $_{K_{\alpha}}$ ($\neg w \wedge \neg we \wedge (\neg 5e \vee \neg cw)$) = $\{\}$. $\beta = 3$.

Stop. Try with $\alpha = n(\beta) = 2$ (the new value of α is $\max(\alpha, n(\beta)) = 2$, since $K_{(4)}$ and $K_{(3)}$ are empty).

Second Step: $\alpha = 2$.

Computation of the nogoods : $NG = \{\{BIO, BAC\}, \{BIO, TA\}, \{BAC, TA\}\}$.

Computation of the label of $\neg w \wedge \neg we$ ($n(\alpha) = 4$):

Label $_K$ ($\neg w$) = $\{\{\}\}$. $\beta = 5$.

Label $_K$ ($\neg w \wedge \neg we$) = $\{\{BAC\}, \{BIO\}\}$. $\beta = 4$.

Stop. $n(\alpha) > 3 = \beta$.

So, the pessimistic optimal decisions for $N(G)=2$ are BIO and BAC for a pessimistic utility $\alpha = 2$.

4.4. Semantics

In this paragraph, we show how the preceding example can be dealt with in a semantical way. Of course, we will see that both approaches lead to the same result.

First of all, we shall notice the correspondence between the representations of the preferences in terms of prioritized formulas and in terms of utility functions over the consequences of actions. Indeed, a utility function μ , such as the above one, can be always put under the form:

$$\mu(\omega) = \max_j \min(\nu_\omega(q_j), \alpha_j)$$

with $\forall \omega \models q_j, \mu(\omega) = \alpha_j$, and where $\nu_\omega(q_j) = 1$ if $\omega \models q_j$, and $\mathbf{0}$ if not. This max-min form can be turned into the equivalent min-max form $\mu_P(\omega) = \min_i \max\{\nu_\omega(p_i), n(\alpha_i)\}$, where we recognize the standard semantics of a stratified possibilistic knowledge base $P = \{(p_i, \alpha_i)\}$, used in Section 2.

A part of the knowledge base K is certain ($\mathbb{1} =$ level 5), including constraints over the decision set: $\{BIO \vee BAC \vee TA, \neg BIO \vee \neg BAC, \dots\}$, factual knowledge $\{g \vee r, \neg g \vee \neg r\}$, and knowledge about the system, e.g. $\{g \wedge BIO \rightarrow 6e, g \wedge TA \rightarrow we, \dots\}$. The only part of the knowledge base that may be uncertain is about the state of freshness of the egg (represented by a necessity valued literal : $(g, N(g))$ or $(r, N(r))$).

In this example, the possibility distribution π_{K_d} restricting the more or less plausible consequences of a decision d , depends only on the possibility distribution on the two possible states g and r , namely, on $\Pi(g)$ and $\Pi(r)$. Let $N(g) = n(\Pi(r))$ and $N(r) = n(\Pi(g))$ (the certainty or necessity of an event is the impossibility of the opposite event). Note that $\min(N(g), N(r)) = \mathbf{0}$, where $\mathbf{0}$ is here the bottom element of our scale (since the possibility distribution over $\{g, r\}$ should be normalized whatever decision d).

The pessimistic utilities of the possible decisions, given by U_* are the following, according to the levels of certainty of g and r :

$$- U_*(BIO) = \min(\max(n(\Pi(r)), \mu(w)), \max(n(\Pi(g)), \mu(6))),$$

which simplifies into $U_*(BIO) = N(g)$.

$$- U_*(BAC) = \min(\max(n(\Pi(r)), \mu(5c)), \max(n(\Pi(g)), \mu(6c))).$$

Thus, $U_*(BAC) = \min(\max(N(g), 2), 4)$.

$$- U_*(TA) = \min(\max(n(\Pi(r)), \mu(5)), \max(n(\Pi(g)), \mu(5w))).$$

Thus, $U_*(TA) = 1$ if $N(g) > 0$ and $\min(3, \max(N(r), 1))$ if not.

The best decisions are therefore:

- Break the egg in the omelette if $N(g) = 5$ (we are sure that the egg is good).
- Break it in the omelette or apart if $N(g) \in \{2, 3, 4\}$ (we are rather sure that the egg is good).
- Break it apart in a cup if $N(g) < 2$ and $N(r) < 2$ (we are rather ignorant on the quality of the egg).
- Throw it away or break it apart if $N(r) = 2$ (we have a little doubt on its quality).
- Throw it away if $N(r) > 2$ (we do not think that the egg is good).

Notice the importance of the commensurability assumption in the computation of U_* where both degrees of certainty and preferences are involved. Note also the qualitative nature of the

approach, since the results depend only on the ordering between the levels in the scale.

4.5. Calculation with symbolic levels

Another solution for computing the pessimistic utility of a decision that combines the syntactic and the semantic approaches can be adopted. We can translate K into another knowledge base K' using additional symbols:

- A_i which will express the fact that we need pieces of knowledge belonging to the i^{th} layer of K to reach the goal,
- P_j which will express the fact that some goal in layer j cannot be reached.
- p_j^i which are non-assumption symbols, representing the individual preferences in layer j .

Let K' be the non-stratified knowledge base obtained from K such that each clause C from the first layer is replaced by $A_1 \rightarrow C$, while the second layer is replaced by the two clauses $A_2 \rightarrow g, A_3 \rightarrow r$.

```
// Knowledge base K'
```

```
// we get a 6-egg omelette if and only if the egg is good and  
// we break it in the omelette;
```

```
A1 g BIO -> 6e ;  
A1 g BAC -> 6e ;  
A1 6e -> g ;  
A1 6e TA -> ;  
...  
A2 -> g ;  
A3 -> r ;
```

A preference layer j is considered as “satisfied” if and only if every preference in this layer is satisfied. Therefore, we construct P' the non-stratified preference base obtained from P in the following way: a clause is added for expressing the condition under which a layer j is satisfied: $C_j^1 C_j^2 \dots \rightarrow P_j$, plus the clause $P_1 \dots P_5 \rightarrow goal$. In this way, if an environment of the *label* contains the symbol P_j , it means that there is at least one preference in P_j that is not satisfied. On the contrary, if it does not, we are sure that every preference in P_j is satisfied.

In our example, each layer contains only one clause, so P' becomes:

```
// Preference base P'
```

```
-> P5 w ;  
-> P4 we ;  
-> P3 5e ;  
-> P3 cw ;
```

```

-> P2 5e;
-> P1 cw;
P1 P2 P3 P4 P5 -> goal ;

```

Label $_{K' \wedge P'}$ (goal) gives us the following result:

{ Environments Label(goal) (from ATMS)	Associated pessimistic utility (non-automatic process)	Decision concerned
{P1 P2 P3 P4 P5}	0	BAC,BIO,TA
{A1 A2 P1 P2 P4}	$\min(N(g), 1)$	BAC,BIO,TA
{A1 A3 P1 P2 P3 P5}	0	BAC,BIO,TA
{A1 P1 P2 P3 BAC}	2	BAC
{A1 A2 P1 BAC}	$\min(N(g), 4)$	BAC
{A1 P2 P4 TA}	1	TA
{A1 A3 P2 TA}	$\min(N(r), 3)$	TA
{A1 P5 BIO}	0	BIO
{A1 A2 BIO}	$N(g)$	BIO

The pessimistic utility of a decision can be obtained from the “best way” it allows to reach the goal, where a “way” is an environment from the label of *goal*. If for example we choose decision *BIO*, then the possible “ways” for reaching the goal are : {P1 P2 P3 P4 P5}, {A1 A2 P1 P2 P4}, {A1 A3 P1 P2 P3 P5}, {A1 P5 BIO}, {A1 A2 BIO}. The “utility of a way” depends on the least certain assumption it involves, and on the goal with the highest priority that has to be assumed true (thus, not provable). For instance, {A1 A2 P1 BAC} depends on assumption *A2* (which level of certainty is $N(g)$), and assumes that *P1* (of priority 1) is true. Therefore, the utility of this “way” (assuming that *BAC* is performed) is $\min(N(g), n(1)) = \min(N(g), 4)$.

Notice that the environment (or “way”) {P1 P2 P3 P4 P5} is *meaningless* in so far as it is an artificial “way” to reach the goal, assuming that every preference is satisfied, even if no decision at all is taken.

Each of the available decisions can thus be evaluated:

$$u_*(BIO) = \max(0, \min(N(g), 1), 0, 0, N(g)),$$

$$u_*(TA) = \max(0, \min(N(g), 1), 0, 1, \min(N(r), 3)),$$

$$u_*(BAC) = \max(0, \min(N(g), 1), 0, 2, \min(N(g), 4)).$$

These expressions are the max-min equivalent forms of the min-max expressions given in the preceding paragraph.

4.6. Remarks

As we have seen on the well known Savage's omelette example, a qualitative, possibilistic decision problem can be described either syntactically by the means of two stratified bases, or semantically by the means of a possibility distribution and a qualitative utility function. The main feature to notice is that both approaches are equivalent. We have proposed and implemented two original algorithms for treating the syntactical case. A question is to see which of the syntactical or the semantical representation is the more appropriate for a given problem. In the Savage's omelette example, it is not clear whether a decision maker would give a logic representation of the problem, or would more willingly give a semantic representation under the form of a utility function. In the general case, the appropriate representation would depend on the particular decision problem under consideration, and on the decision maker's habits. However, in problems involving a large number of states, one may expect that the logical representation of partial belief about the world, and preferences on goals would be more economic than an explicit enumeration of states with their levels of plausibility and of preference. We have also pointed out that it is possible to pass from one representation to the other, and how it can be done.

5. Concluding remarks

The main contribution of this paper has been to describe a logical machinery for decision-making, implementing the qualitative possibilistic utility theory, in the framework of possibilistic logic. A link between this logical machinery and the ATMS framework has been pointed out, which has allowed to adapt some efficient algorithms proposed in this framework to possibilistic qualitative decision making.

One strong assumption has been made in this paper, which is that certainty levels and priority levels be commensurate. An attempt to relax this assumption has been made in (Dubois, Fargier and Prade [14]). These authors point out that working without the commensurability assumption leads them to a decision method close to rational inference machinery in non-monotonic reasoning. Unfortunately, that method also proves to be either very little decisive or to lead to very risky decisions.

Besides, the links between possibilistic qualitative decision making and diagnosis (abductive and consistency-based) may be further explored: (Cayrac et al. [8]) have proposed a way to handle uncertainty in model-based diagnosis which is technically very close to the one exposed here in the decision framework. In (Le Berre and Sabbadin [26]), a logical machinery similar to the one exposed here has been presented, in the diagnosis and repair framework. This machinery is also based on ATMS techniques. However, the actions under consideration are repair-actions, preferences are expressed by the means of real-valued goals (where the value of a goal is its

utility in the classical sense of decision theory) of a specific kind, and uncertainty is modeled by “probability-valued” assumptions. Methods (also based on the MPL procedure) are given for computing the *belief-based expected utility* of a decision (a counterpart of classical expected utility, in the Dempster-Shafer theory).

Finally, we can think of dealing with possibilistic logic formulas involving time instants (e.g., as in Dubois and Prade [19]) in order to extend the syntactical approach presented here to multiple-stage possibilistic decision (Fargier et al. [23]). Such an extended framework will be also useful if the computation of the result of the decision requires an updating of K .

References

- [1] S. Benferhat, D. Dubois, H. Prade. Representing default rules in possibilistic logic. In *Proc. 3rd Inter. Conf. on Principles of Knowledge Representation and Reasoning* (KR'92) (B. Nebel, C. Rich, W. Swartout eds.), pp. 673-684, Cambridge, MA, Oct. 25-29, 1992.
- [2] B. Bonet, H. Geffner. Arguing for decisions : a qualitative model of decision making. In *Proc. of the 12th Conf. on Uncertainty in Artificial Intelligence* (UAI'96) (E. Horwitz, F. Jensen, eds.), pp. 98-105, Portland, Oregon, July 31-Aug. 4 1996.
- [3] C. Boutilier. Toward a logic for qualitative decision theory. In *Proc. 4th Inter. Conf. on Principles of Knowledge Representation and Reasoning* (KR'94) (J. Doyle, E. Sandewall, P. Torasso, eds.), pp. 75-86, Bonn, Germany, May 24-27, 1994.
- [4] R. I. Brafman, M. Tennenholtz. On the foundations of qualitative decision theory. In *Proc. 12th National Conf. on Artificial Intelligence* (AAAI'96), pp. 1291-1296, Portland, Oregon, Aug. 4-8, 1996.
- [5] T. Castell, M. Cayrol. Computation of prime implicates and prime implicants by the Davis and Putnam procedure. In *Proc. ECAI'96 Workshop on Advances in Propositional Deduction*, pp. 61-64, Budapest, 1996.
- [6] T. Castell, C. Cayrol, M. Cayrol, D. Le Berre. Efficient computation of preferred models with Davis and Putnam procedure. In *Proc. European Conference on Artificial Intelligence* (ECAI'96), pp. 354-358, 1996.
- [7] T. Castell, C. Cayrol, M. Cayrol, D. Le Berre. Modèles P-restreints. Applications à l'inférence propositionnelle. In *Proc. Reconnaissance des Formes et Intelligence Artificielle* (RFIA'98), Clermont-Ferrand, France, Jan 20-22, pp. 205-214, 1998.
- [8] D. Cayrac, D. Dubois, D. Prade. Practical model-based diagnosis with qualitative possibilistic uncertainty. In *Proc. of the 11th Conf. on Uncertainty in Artificial Intelligence* (UAI'95), pp. 68-76, Montreal, Canada, Aug. 1995.
- [9] H. Davis and L. Putnam. A computing procedure for quantification theory, *Journal of the ACM* (7), pp 201-215, 1960.
- [10] H. Davis, G. Logemann and D. Loveland. A machine program for theorem proving, *Commun. ACM* (5), pp 394-397, 1962.
- [11] J. De Kleer. An assumption based truth maintenance system. In *Artificial Intelligence*, (28) pp. 127-162, 1986.

- [12] J. De Kleer. Extending the ATMS. In *Artificial Intelligence*, (28) pp. 163-196, 1986.
- [13] O. Dubois, P. André, Y. Boufkhad, J. Carlier. Sat vs. unsat. In *2nd DIMACS Implementation Challenge*, volume 26 of DIMACS series, pp 415-436. American Mathematical Society, 1996.
- [14] D. Dubois, H. Fargier, H. Prade. Decision-making under ordinal preferences and uncertainty. In *Proc. of 13th Nat. Conf on Uncertainty in Artificial Intelligence (UAI'97)*, pp. 157-164, Providence R.I., 1997.
- [15] D. Dubois, J. Lang, H. Prade. Automated reasoning using possibilistic logic: Semantics, belief revision, and variable certainty weights. *IEEE Trans. on Knowledge and Data Engineering*, 6(1):64-69, 1994.
- [16] D. Dubois, D. Le Berre, H. Prade, R. Sabbadin. Logical representation and computation of optimal decisions in a qualitative setting. In *Proc. 14th Nat. Conf on Artificial Intelligence (AAAI'98)*, pp. 588-593. Madison, Wisconsin, July 26-30, 1998.
- [17] D. Dubois, H. Prade. Possibility theory as a basis for qualitative decision theory. In *Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence (IJCAI'95)*, pp. 1925-1930, Montreal, Canada, Aug. 20-25, 1995.
- [18] D. Dubois, H. Prade. Possibilistic mixtures and their applications to qualitative utility theory. Part II: decision under incomplete knowledge. In *Proc. of Foundations and Applications of Possibility Theory (FAPT'95)*, (G. de Cooman, Da Ruan, E.E. Kerre, eds.), pp. 256-266, Ghent, Belgium, Dec. 13-15, 1995.
- [19] D. Dubois, H. Prade. Combining hypothetical reasoning and plausible inference in possibilistic logic. *Multiple Valued Logic*, 1, pp. 219-239, 1996.
- [20] D. Dubois, H. Prade, R. Sabbadin. A possibilistic logic machinery for qualitative decision. In *Proc. of the AAAI 1997 Spring Symposium Series (Qualitative Preferences in Deliberation and Practical Reasoning)*, Stanford University, California, March 24-26, 1997.
- [21] D. Dubois, H. Prade, R. Sabbadin. Towards axiomatic foundations for decision under qualitative uncertainty. In *Proc. of the 7th Conf. International Fuzzy Systems Association (IFSA'97)*, pp. 441-446, Prague, Czech Republic, June 25-29, 1997.
- [22] D. Dubois, H. Prade, R. Sabbadin. Qualitative decision theory with Sugeno integrals. In *Proc. 14th Conf. Uncertainty in artificial Intelligence (UAI'98)*, pp. 121-128. Madison, July 24-26, 1998.
- [23] H. Fargier, J. Lang, R. Sabbadin. Towards qualitative approaches to multistage decision making. In *Proc. of the 6th Int. Conf. on Information Processing and Management of Uncertainty in Knowledge-based Systems (IPMU'96)*, pp. 31-36, Granada (Spain), July 1-5, 1996.
- [24] J. Lang. Possibilistic logic as a logical framework for min-max discrete optimization problems and prioritized constraints. In *Proc. Fundamentals of Artificial Intelligence Research*. (P. Jorrand, J. Kelemen eds.) LNAI 535, pp. 112-126, 1991.
- [25] D. Le Berre. Le calcul de modèles possibles basé sur la procédure de Davis et Putnam, et ses applications. Technical report IRIT / 96-32-R, September 1996.

- [26] D. Le Berre, R. Sabbadin. Decision-theoretic diagnosis and repair: representational and computational issues. In *Proc. of the 8th International Workshop on Principles of Diagnosis (DX'97)*, pp. 141-145, Le Mont-Saint-Michel, France, September 15-18, 1997.
- [27] Chu Min Li, Anbulagan. Heuristics based on unit propagation for satisfiability problems. In *Proc. of the 15th Inter. Joint Conf. on Artificial Intelligence (IJCAI'97)*, pp. 366-371, 1997.
- [28] W. Hamscher, L. Console, J. de Kleer (eds). *Readings in Model-Based Diagnosis*. Morgan Kaufmann Publishers, San Mateo, CA, 1992.
- [29] J. Pearl. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proc of the 3rd Conf. on Theoretical Aspects of Reasoning About Knowledge (TARK'90)*, pp. 121-135, Morgan Kaufmann, 1990.
- [30] J. Pearl. From conditional oughts to qualitative decision theory. In *Proc. of the 9th Conf. on Uncertainty in Artificial Intelligence (UAI'93)*, (D. Heckerman, A. Mamdani eds.), pp. 12-20, Washington, DC, July 9-11, 1993.
- [31] H. Prade. Model semantics and fuzzy set theory. In *Fuzzy Sets and Possibility theory. Recent Developments* . R.R Yager Ed, Pergamon Press, pp. 232-246, 1982.
- [32] L.J. Savage. *The Foundations of Statistics*. Dover, New York, 1972.
- [33] R. Schrag. Compilation for critically constrained knowledge bases. In *Proc. 12th National Conf. on Artificial Intelligence (AAAI'96)*, pp. 510-515, 1996.
- [34] S. W. Tan and J. Pearl. Qualitative decision theory. In *Proc. 11th National Conf. on Artificial Intelligence (AAAI'94)*, pp. 928-933, Seattle, WA, July 31-Aug. 4, 1994.