ARTICLE IN PRESS

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Information Fusion xxx (2005) xxx–xxx

INFORMATION FUSION

www.elsevier.com/locate/inffus

Bipolar possibility theory in preference modeling: Representation, fusion and optimal solutions ☆

Salem Benferhat [a,*], Didier Dubois [b], Souhila Kaci [c], Henri Prade [b]

[a] Centre de Recherche en Informatique de Lens (C.R.I.L.)—C.N.R.S., Rue Jean Souvraz SP 18, 62307 Lens Cedex, France
[b] Institut de Recherche en Informatique de Toulouse (I.R.I.T.)—C.N.R.S., 118 route de Narbonne, 31062 Toulouse Cedex 4, France
[c] Centre de Recherche en Informatique de Lens (C.R.I.L.)—C.N.R.S., Rue de l'Université SP 16, 62307 Lens Cedex, France

Received 22 October 2003; received in revised form 25 February 2005; accepted 1 April 2005

Abstract

The bipolar view in preference modeling distinguishes between negative and positive preferences. Negative preferences correspond to what is rejected, considered unacceptable, while positive preferences correspond to what is desired. But what is tolerated (i.e., not rejected) is not necessarily desired. Both negative and positive preferences can be a matter of degree. Bipolar preferences can be represented in possibilistic logic by two separate sets of formulas: prioritized constraints, which describe what is more or less tolerated, and weighted positive preferences, expressing what is particularly desirable.

The problem of merging multiple-agent preferences in this bipolar framework is then discussed. Negative and positive preferences are handled separately and are combined in distinct ways. Since negative and positive preferences are stated separately, they may be inconsistent, especially in this context of preference fusion. Consistency can be enforced by restricting what is desirable to what is tolerated.

After merging, and once the bipolar consistency is restored, the set of preferred solutions can be logically characterized. Preferred solutions should have the highest possible degree of feasibility, and only constraints with low priority may have to be discarded in case of inconsistency inside negative preferences. Moreover, preferred solutions should satisfy important positive preferences when feasible (positive preferences may be also inconsistent). Two types of preferred solutions can be characterized, either in terms of a disjunctive combination of the weighted positive preferences, or in terms of a cardinality-based evaluation.

Keywords: Bipolarity; Preferences

1. Introduction

The problem of representing preferences has been considered by various researchers in Artificial Intelligence in the recent past [44,21,14,1,32,6,16]. Indeed this issue is important when we have to represent requirements or desires expressed by users in information systems (e.g. in recommender systems), or to reason about preferences and to solve conflicts in case of inconsistency, as e.g. in multiagent systems.

Preferences over a set of possible choices or solutions are often expressed in two forms: positive and negative aspirations. Indeed, on the one hand, an agent may express what he considers (more or less) unacceptable for him, and on the other hand he may express what he considers as being really satisfactory. The first form of preferences will be called negative preferences and corresponds to constraints that should be respected, while the second form will be called positive preferences,

☆ This paper is an extended version of [5,30].

* Corresponding author. Tel.: +33 3 21 79 17 79; fax: +33 3 21 79 17 70.

E-mail addresses: benferhat@cril.univ-artois.fr (S. Benferhat), dubois@irit.fr (D. Dubois), kaci@cril.univ-artois.fr (S. Kaci), prade@irit.fr (H. Prade).

1566-2535/$ - see front matter
doi:10.1016/j.inffus.2005.04.001

and correspond to appealing states of the world for the agent. Positive preferences can be viewed as "weak" preferences, in the sense that they do not exclude solutions stricto sensu, but only suggest the best ones. Thus solutions falsifying positive preferences may remain tolerated. Of course, among tolerated solutions pointed out by negative preferences, those satisfying positive preferences will be preferred. This distinction does not seem to be accounted for in the above-cited works.

For example, consider a three days summer school, for which each invited speaker is asked to express a preferred time slot for scheduling his talk. We assume that talks can be given either on Monday, or on Tuesday or on Wednesday, and for each day the talk can be either scheduled in the morning, or in the afternoon. The invited speaker may provide two kinds of preferences. First, he specifies negative preferences, which describe unacceptable slots with levels of tolerance. For instance, he may strongly object to working on Monday (e.g., because it is the birthday of his daughter), and weakly refuse to speak on Wednesday. These negative preferences reject options and induce a first ranking on all feasible solutions. For instance, solutions where the talk is scheduled on Tuesday (either morning or afternoon) are the preferred ones, while solutions where the talk is scheduled on Monday are considered as unacceptable. Next, the invited speaker specifies positive preferences. For instance, having a talk in the morning is preferred to having it in the afternoon, and scheduling it in early morning will be even better. These positive preferences will induce a second ranking on all possible solutions. The principal aim of this paper is to propose a tool to represent these two types of preference relations conjointly and compactly, in the possibility theory framework, how to use them in order to select best solutions and lastly how to restore consistency in case of conflicts between them. In our example, if there is no additional expressed preference, then the best solution is to schedule the talk on Tuesday early morning. Indeed, this solution does not violate any invited speaker's constraints and satisfies all of his preferred positive preferences. Now, if for some security reasons, buildings cannot be opened early morning, then scheduling the talk on Tuesday late morning is still a tolerated solution, despite the fact that it does not satisfy one of invited speaker's positive preferences. The situation is different, if for some organisational reasons the talk only can be scheduled on Monday, in which case the proposed solution is unfeasible. Negative preferences act as constraints discarding unacceptable solutions, while positive preferences can only lead to support appealing solutions, and hence can be used to discriminate between tolerated ones.

This representation of bipolar preferences in separate sets is motivated by recent studies in cognitive psychology showing that the distinction between positive and negative preferences makes sense. They are processed separately in the brain, and are felt as different dimensions by people [19,18]. Moreover, referring to [35], the human mind seems to be expert in the processing of bipolar information, and it is very likely that a large part of this expertise is not acquired, but would be the result of the cognitive evolution of the species. It would result from the architecture(s) of the mind and affect system(s) [18,34,36,13,35]. Note that in general there is no symmetry between positive and negative preferences in the sense that positive aspirations do not just mirror what is not rejected. In the summer school example, there are some time slots which are neither rejected nor desired. We suppose that the agent is indifferent w.r.t. these slots.

Describing the set of preferred solutions to a problem may be of interest in decision support systems for communication and explanation purposes. This paper also makes a step towards a machinery capable of computing logical characterizations (i.e., expressed in syntactic terms) of sets of preferred solutions, taking into account bipolar preferences of agents and constraints existing in the environment where the decision takes place. Indeed, a logical handling of preferences is natural for expressing them compactly and for reasoning about them. Here, the worth of solutions according to positive preferences is supposed to be only qualitatively assessed, and priority levels which can be associated with constraints are also assumed to belong to a qualitative scale. This is in contrast to approaches based on utility theory, for instance. This is why the representation framework of possibilistic logic [24] is used, since it allows for a stratified handling of classical logical formulas according to their priority levels. Logical interpretations are then rank-ordered.

A bipolar possibilistic logic framework for modeling preferences was first advocated in [5], where a preliminary version of Sections 2 and 3 of this paper was presented. On the one hand, prioritized logical formulas (weighted in terms of necessity degrees) are used for expressing constraints whose priority is more or less high. It thus delimits the fuzzy set of feasible solutions compatible with the constraints. This corresponds to negative preferences, which are compulsory at least to some degree (what is rejected defines, by complementation, what is tolerated).

On the other hand, other formulas, weighted in terms of a "guaranteed possibility" function, express the level of desirability associated with the satisfaction of those formulas. The second type of formulas corresponds to a "positive" assessment of feasible solutions. The consistency of the two types of preference requires that the fuzzy set of solutions that have some guaranteed satisfaction level be included in the fuzzy set of solutions compatible with the constraints. When positive preferences are not consistent with negative ones, consistency has to be restored by restricting the former to what is allowed by the latter.

This paper is organised as follows. Section 2 presents the framework of bipolar representations. It discusses the possibilistic encoding of positive and negative preferences, and formulates the principle of consistency restoration when they conflict. Section 3 explains how the sets of positive (resp. negative) preferences can be merged at the syntactic level, and how consistency restoration can be syntactically done. Section 4 presents the problems of defining and computing the best solutions to a problem posed in terms of negative and positive preferences. Lastly, Section 5 discusses related works. Appendix A provides a short background on possibilistic logic (see [24] for more details).

2. Representing positive and negative preferences

Consider a propositional language $\mathscr{L}$ over a finite alphabet $\mathscr{P}$ of atoms. S denotes the set of all classical interpretations. $[\![p]\!]$ denotes the set of models of proposition p. Let s be an element in S. In the following, it will be referred to as a solution (to the problem described in terms of positive and negative preferences). The notation $s \models p$ means that s is a model of (a solution to) p, i.e. $s \in [\![p]\!]$. $\vdash$ denotes the classical logic deduction. $\vee$ and $\wedge$ represent the classical disjunction and conjunction respectively, $\neg$ the classical negation, $p \Rightarrow q$ is short for $\neg p \vee q$.

2.1. Syntactic specification of bipolar preferences

In this paper, we propose a bipolar representation of preferences both at the syntactic and at the semantic levels. We introduce the syntactic specification of these preferences in this section. Preferences of an agent will be represented by two different sets of inequality constraints.

The first set expresses positive preferences. It is of the form $\mathbb{W} = \{\Delta(w_j) \geqslant b_j : j = 1, \ldots, m\}$, where w_j is a propositional formula, Δ returns a level of satisfaction, and b_j lies in a finite totally ordered scale L^+, by convention contained in the interval $(0, 1]$. L^+ is a *positive* scale because its top represents maximal possible satisfaction while its bottom is neutral and represents mere tolerance. Thus w_j is supposed to encode a desire or a wish, b_j expresses the minimal level of satisfaction which is guaranteed for a solution where w_j is true. The larger b_j is, the more satisfied is the agent if w_j is true. $\Delta(w_j) = 1$ means that the agent is fully satisfied as soon as w_j is true.

This kind of preferences cannot be directly handled by the standard possibilistic logic machinery [24], based on necessity measures. In fact, positive preferences can be represented using a function of "guaranteed possibility", denoted by Δ, in possibility theory [27]. The syntactic representation of graded information using Δ is dual in some sense to the one used in standard possibilistic logic [23,9]. This is recalled in Section 2.3.

The second set of inequality constraints expresses negative preferences and corresponds to what is not tolerated by the agent. By complementation, it induces what is more or less tolerated, feasible for the agent. This set is of the form $\mathbb{R} = \{\mathscr{R}(r_i) \geqslant a_i : i = 1, \ldots, n\}$, where r_i is a propositional formula which must be violated; we call r_i a rejection statement. $\mathscr{R}$ stands for rejection, and a_i represents the priority level of rejecting r_i for the agent. a_i lies in a finite totally ordered scale L^-, by convention contained in the interval $(0, 1]$. L^- is a *negative* scale because its top represents full strength of rejection while its bottom is neutral and represents feasibility.

$\mathscr{R}(r_i) \geqslant a_i$ expresses that all the models of r_i are solutions unfeasible at least at level a_i. The higher a_i is, the less tolerated are the solutions satisfying r_i. $\mathscr{R}(r_i) = 1$ means that the agent gives the highest priority to the rejection of r_i, and no solution where r_i is true is tolerated by the agent. It is patent that negative preferences can be easily handled using standard possibilistic logic, based on any of the two classical functions of possibility theory: possibility and necessity measures. This will be developed in Section 2.2.

2.2. Modeling negative preferences in possibilistic logic

Rejection is graded on the scale of dissatisfaction L^-. Let $du_{\mathbb{R}}(s)$ (*du* for disutility function) be the level of dissatisfaction with solution s. The larger $du_{\mathbb{R}}(s)$, the greater the dissatisfaction. Clearly, an agent when expressing his negative preferences does not provide $du_{\mathbb{R}}$, but rather more compact statements of the form $\mathscr{R}(r_i) \geqslant a_i$, to be understood as $\forall s \models r_i, du_{\mathbb{R}}(s) \geqslant a_i$. This parallels the understanding of positive preferences.

By order-reversing the scale L^-, we obtain a feasibility scale L defining a (negative) possibility distribution $\forall s, \pi_{\mathbb{R}}(s) = 1 - du_{\mathbb{R}}(s)$. Then, $\mathscr{R}(r_i) \geqslant a_i$ translates into $\forall s \models r_i, \pi_{\mathbb{R}}(s) \leqslant 1 - a_i$, i.e. $\max_{s \models r_i} \pi_{\mathbb{R}}(s) \leqslant 1 - a_i$, where we recognize the expression of a constraint on the possibility measure (see Appendix A) associated with $\pi_{\mathbb{R}}$, namely $\Pi_{\mathbb{R}}(r_i) \leqslant 1 - a_i$. The latter constraint is itself equivalent to $N_{\mathbb{R}}(\neg r_i) \geqslant a_i$, where $N_{\mathbb{R}}$ is the dual measure of necessity associated with $\Pi_{\mathbb{R}}$ (see Appendix A). We thus recognize the necessity-based constraints that are underlying possibilistic logic [24]. Thus, rejecting r_i to some extent amounts to enforcing the validity of $\neg r_i$ to the same extent.

This shows that possibilistic logic provides a natural framework for modeling negative preferences. Tolerability or feasibility can be represented, at the semantic level, by the total pre-order on the set of all possible solutions encoded by $\pi_{\mathbb{R}}$, from what is tolerated for an agent to what is considered as unacceptable for him.

$\pi_{\mathbb{R}}(s)$ represents the degree of tolerability of a solution s given agent's preferences. $\pi_{\mathbb{R}}(s) = 1$ means that s is fully tolerated, $\pi_{\mathbb{R}}(s) = 0$ means that s is completely unacceptable (rejected), and more generally, $\pi_{\mathbb{R}}(s) > \pi_{\mathbb{R}}(s')$ means that s is more tolerated than s'.

In practice an agent cannot provide the whole possibility distribution $\pi_{\mathbb{R}}$, but only the set $\mathbb{R}$ of negative preferences, and their level of priority. Let $\mathbb{R} = \{\mathscr{R}(r_i) \geqslant a_i : i = 1, \ldots, n\}$. The question is then how to obtain $\pi_{\mathbb{R}}$. Let us illustrate the construction of $\pi_{\mathbb{R}}$ when we only have one strong constraint $\mathscr{R}(r) = 1$. Let s be a solution. Intuitively, if s falsifies r then s is fully tolerated by the agent, i.e., $\pi_{\mathbb{R}}(s) = 1$, and if s satisfies r then s is fully unacceptable, i.e. $\pi_{\mathbb{R}}(s) = 0$. Now, assume that r is weakly prioritized, i.e. $\mathscr{R}(r) \geqslant a$ with $a < 1$. Again, if s falsifies r then it is fully tolerated. If s satisfies r then the higher is a, the less tolerated is s. One way to achieve this constraint is to assign the value $1 - a$ to $\pi_{\mathbb{R}}(s)$. More formally, the possibility distribution associated with $\mathbb{R} = \{\mathscr{R}(r) \geqslant a\}$ is

$$\pi_{\mathbb{R}}(s) = \begin{cases} 1 & \text{if } s \nvDash r, \\ 1 - a & \text{if } s \models r. \end{cases}$$

Now assume that the agent expresses two prioritized rejection statements r_1 and r_2 with $\mathscr{R}(r_1) \geqslant a$ and $\mathscr{R}(r_2) \geqslant b$ respectively. Then, we have three cases:

- $s \nvDash r_1$ and $s \nvDash r_2$, then s is fully feasible. Hence $\pi_{\mathbb{R}}(s) = 1$.
- $s \models \neg r_1 \wedge r_2$ (resp. $s \models r_1 \wedge \neg r_2$) then the higher is b (resp. a), the less tolerated is s. Hence $\pi_{\mathbb{R}}(s) = 1 - b$ (resp. $1 - a$).
- $s \models r_1 \wedge r_2$ then the higher is a or b, the less tolerated is s. Hence $\pi_{\mathbb{R}}(s) = 1 - \max(a, b)$.

More generally, we have

Definition 1. The possibility distribution $\pi_{\mathbb{R}}$ associated with a set of negative preferences $\mathbb{R} = \{\mathscr{R}(r_i) \geqslant a_i : i = 1, \ldots, n\}$ is:

$$\pi_{\mathbb{R}}(s) = 1 - \max\{a_i : s \models r_i, \mathscr{R}(r_i) \geqslant a_i \in \mathbb{R}\},$$

with $\max\{\emptyset\} = 0$.

Clearly, this definition coincides with the semantics of a possibilistic logic base in standard possibilistic logic, for inducing a possibility distribution π_Σ from a possibilistic knowledge base Σ, expressing priority levels of constraints in terms of a necessity measure. Thus, a set of negative preferences $\mathbb{R} = \{\mathscr{R}(r_i) \geqslant a_i : i = 1, \ldots, n\}$ can be simply encoded as a set of weighted formulas in possibilistic logic i.e., $\{(\neg r_i, a_i) : i = 1, \ldots, n\}$, where r_i is a rejection statement and a_i is now the priority level associated with the induced constraint of having to satisfy $\neg r_i$.

This result is important since it means that the classical possibilistic logic machinery can be used for handling negative preferences (or negative preferences) and drawing inferences from them. Moreover, the complexity of possibilistic logic is only slightly higher than the one of classical logic [24].

In the rest of this paper, we will encode a set of rejection statements by a possibilistic logic base $\mathbb{R} = \{(\neg r_i, a_i): i = 1, \ldots, n\}$, where $(\neg r_i, a_i)$ represents the constraint $\mathscr{R}(r_i) \geqslant a_i$ and stands for "*if a solution s satisfies r_i then it is tolerated at most to a degree $1 - a_i$*". Mind that we continue to name $\mathbb{R}$ the set of standard possibilistic logic formulas providing an equivalent encoding, in terms of tolerability, of the initial set of negative preferences.

Moreover, tolerated solutions induced by negative preferences are not necessarily feasible w.r.t. existing integrity constraints. Suppose that our aim is to buy a small apartment. Then the solution "large and small apartment" may be tolerated, however it is not feasible since there is an integrity constraint which says that the surface of the apartment is either small or large but not both. Feasibility is induced by integrity constraints which delimit the set of potential solutions of the problem. We assume that this set is represented in *standard* possibilistic logic by a set of completely certain weighted formulas of the form $\mathbb{F} = \{(f_k, 1): k = 1, \ldots, s\}$, where f_k are propositional formulas. The degree 1 is associated with the formulas in order to express that the f_k's are hard (or strict) constraints. A solution s is said to be feasible (or compatible) if it satisfies all formulas f_k's of $\mathbb{F}$. Thus, tolerated solutions are computed among feasible ones.

Example 1. Let p_1, p_2 and p_3 be three propositional symbols representing three categories of price, which stand respectively for "cheap", "reasonable" and "expensive". Let s, m, l be three kinds of surfaces which stand respectively for small, medium and large. By "garden" (g for short), we denote the fact that the house or the apartment has a garden.

Assume that we have the following set of integrity constraints, including domain exclusion constraints, regarding a house or an apartment:

$$\begin{aligned}\mathbb{F} = \{&(l \Rightarrow \neg p_1, 1), (s \vee m \vee l, 1), (\neg s \vee \neg m, 1), \\ &(\neg s \vee \neg l, 1), (\neg m \vee \neg l, 1), (p_1 \vee p_2 \vee p_3, 1), \\ &(\neg p_1 \vee \neg p_2, 1), (\neg p_2 \vee \neg p_3, 1), (\neg p_1 \vee \neg p_3, 1), \\ &(g \Rightarrow p_3, 1)\}.\end{aligned}$$

Suppose now that the agent expresses one rejection statement $\mathbb{R} = \{(g, .5)\}$, namely having a garden with a priority .5.

The set $\mathbb{F}$ expresses mutual exclusivity constraints about categories of price, of surface and an integrity constraint about the existence of gardens.

Table 1
The possibility distribution associated with $\mathbb{R}$

s	$\pi_{\mathbb{R}}(s)$
p_3sg	1
p_3mg	1
p_3lg	1
$p_1s\neg g$	.5
$p_1m\neg g$	.5
$p_2s\neg g$	.5
$p_2m\neg g$	.5
$p_2l\neg g$	.5
$p_3s\neg g$	.5
$p_3m\neg g$	.5
$p_3l\neg g$	.5
Other feasible solutions	0

Feasible solutions are models of the set $\mathbb{F}$ namely models of $(\neg l \vee \neg p_1) \wedge (s \vee m \vee l) \wedge (\neg s \vee \neg m) \wedge (\neg s \vee \neg l) \wedge (\neg m \vee \neg l) \wedge (p_1 \vee p_2 \vee p_3) \wedge (\neg p_1 \vee \neg p_2) \wedge (\neg p_2 \vee \neg p_3) \wedge (\neg p_1 \vee \neg p_3) \wedge (\neg g \vee p_3)$.

Table 1 gives the possibility distribution[1] associated with $\mathbb{R}$ applying Definition 1.

2.3. Representing positive preferences in the logic of guaranteed possibility

Positive preferences of an agent can be described, at the semantical level, in terms of the linearly ordered positive scale L^+, encoding satisfaction levels.

They give birth to a distribution of positive weights, that we shall call positive possibility distribution, over the set of the logical interpretations that represent possible choices or solutions. Let $\delta_{\mathbb{W}}$ be this distribution. $\delta_{\mathbb{W}}(s) > \delta_{\mathbb{W}}(s')$ means that s is more satisfactory than s'. Usually an agent is not able to provide such a distribution explicitly.

Rather, the agent provides a set of positive preferences of the form $\mathbb{W} = \{\Delta(w_j) \geqslant b_j : j = 1, \ldots, m\}$ with the intended meaning that $\forall s \models w_j$, $\delta_{\mathbb{W}}(s) \geqslant b_j$, which is indeed a more compact way of expressing sets of solutions that are satisfactory to some extent. When processed syntactically as logical expressions, $\mathbb{W}$ will be denoted $\{[w_j, b_j] : j = 1, \ldots, m\}$. Thus, $[w_j, b_j]$ encodes the information that the solution will be satisfactory for the agent, at least to a degree b_j if w_j is satisfied. Clearly, a solution s, such that $s \models w_1$ and $s \models w_2$ with $[w_1, b_1] \in \mathbb{W}$ and $[w_2, b_2] \in \mathbb{W}$, will be such that $\delta_{\mathbb{W}}(s) \geqslant \max(b_1, b_2)$.

Let us see now how to build the positive possibility distribution $\delta_{\mathbb{W}}$, associated with $\mathbb{W}$. We first consider the case where we only have one formula $[w, b]$. Then if a given solution s satisfies w, the associated level of satisfaction will be equal to b. Otherwise, this level will be equal to 0, since no information is given by $[w, b]$ for the models of $\neg w$. In the general case, given a set of positive preferences, the level of satisfaction associated with s is equal to the highest level of a formula appearing in $\mathbb{W}$ satisfied by s:

Definition 2. The positive possibility distribution $\delta_{\mathbb{W}}$ associated with a set of positive preferences $\mathbb{W} = \{[w_j, b_j] : j = 1, \ldots, m\}$ is

$$\delta_{\mathbb{W}}(s) = \max\{b_j : s \models w_j \text{ and } [w_j, b_j] \in \mathbb{W}\},$$

with $\max\{\emptyset\} = 0$.

Note that the addition of positive preferences in $\mathbb{W}$ can only lead to an increase of the satisfaction level of solutions.

The meaning of $\pi_{\mathbb{R}}(s)$ is different from the one of $\delta_{\mathbb{W}}(s)$ in the previous section. Indeed $\delta_{\mathbb{W}}(s)$ evaluates to what degree s is *satisfactory* for the agent, while $\pi_{\mathbb{R}}(s)$ evaluates to what degree s is *tolerated* by the agent. In particular $\delta_{\mathbb{W}}(s) = 1$ means that the agent is fully satisfied, while $\delta_{\mathbb{W}}(s) = 0$ simply means that the agent is indifferent, while $\pi_{\mathbb{R}}(s) = 0$ means that s is impossible.

Example 2. Let $\mathbb{W}$ be the following set of positive preferences, concerning the choice of a house or an apartment

$$\mathbb{W} = \{[m, 1], [\neg apart \wedge l, 1], [\neg g \wedge l, .6]\}.$$

The first expression means that the agent is fully satisfied when the surface is medium (m). He is also fully satisfied with a large (l) house (expressed by the second formula). The last expression means that the agent is fairly satisfied if the surface is large and there is no garden (g), whether it is a house or an apartment.

Table 2 gives the positive possibility distribution associated with $\mathbb{W}$ following Definition 2.

The set of positive preferences cannot be directly handled by standard possibilistic logic. This contrasts with the handling of negative preferences, in the next section. This is because a weighted formula like $[w, b]$ cannot be

Table 2
The positive possibility distribution associated with $\mathbb{W}$

s	$\delta_{\mathbb{W}}(s)$
$m\neg apart\neg g$	1
$m\neg apart\, g$	1
$m\, apart\neg g$	1
$m\, apart\, g$	1
$l\neg apart\neg g$	1
$l\neg apart\, g$	1
$l\, apart\neg g$	.6
Other feasible solutions	0

[1] Strictly speaking, solutions given in Table 1 are only partial. In fact for the sake of simplicity, when we write p_3sg, we mean the interpretation $p_3\neg p_1\neg p_2 s\neg m\neg lg$. In the rest of the paper, we use this simplified notation for sake of clarity.

directly expressed by means of possibility or necessity measures. Namely, constraints like $[w, b]$ can be represented using a third set-function called *guaranteed possibility*, denoted by Δ [27], defined by

$$\Delta(w) = \min_{s \models w} \delta_{\mathbb{W}}(s).$$

The expression $\Delta(w) \geqslant b$ means that any solution where w is true, has a possibility (here: satisfaction) level at least equal to b. This is exactly what is intended by the kind of information encoded in $\mathbb{W}$.

See Appendix A for details on the Δ function and its logic. See also [23,9] for detailed presentations.

2.4. Coherence between positive and negative preferences

Even if independently specified negative and positive preferences must nevertheless be in agreement with each other: a solution cannot be at the same time unacceptable and desired by the same agent. Let $(\mathbb{W}, \mathbb{R})$ be the preferences for an agent. Intuitively, if $\mathbb{R} = \{(\neg r_i, 1): i = 1, \ldots, n\}$ and $\mathbb{W} = \{[w_j, 1] : j = 1, \ldots, m\}$ are classical logic bases (with maximal priority or satisfaction weights) then $\mathbb{W}$ and $\mathbb{R}$ are coherent if

$$\bigvee_{j=1,\ldots,m} w_j \vdash \bigwedge_{i=1,\ldots,n} \neg r_i,$$

namely any solution verifying at least one positive preference of $\mathbb{W}$ should also satisfy all formulas in $\mathbb{R}$, which represent constraints induced by negative preferences. More generally, any solution satisfactory to degree a (w.r.t. $\mathbb{W}$) should be at least tolerated to degree a (w.r.t. $\mathbb{R}$).

Definition 3. Let $\delta_{\mathbb{W}}$ and $\pi_{\mathbb{R}}$ be the two possibility distributions representing the positive and negative preferences of an agent respectively. Then, $\delta_{\mathbb{W}}$ and $\pi_{\mathbb{R}}$ are said to be coherent iff

$$\forall s, \quad \delta_{\mathbb{W}}(s) \leqslant \pi_{\mathbb{R}}(s).$$

Coherence checking can also be done syntactically using the bases $\mathbb{W}$ and $\mathbb{R}$.

Proposition 1. *Let $\mathbb{W}$ and $\mathbb{R}$ be respectively the sets of positive preferences and rejections of an agent. Then, $\mathbb{W}$ and $\mathbb{R}$ are said to be coherent (in the sense of Definition 3) iff*

$$\forall a \geqslant 0, \quad \bigvee_{[w_j, a_j] \in \mathbb{W}, a_j \geqslant a} w_j \vdash \bigwedge_{(\neg r_i, a_i) \in \mathbb{R}, a_i > 1-a} \neg r_i.$$

The proof can be obtained by noticing that Definition 3 is equivalent to $\forall a \geqslant 0$, $\{s: \delta_{\mathbb{W}}(s) \geqslant a\} \subseteq \{s: \pi_{\mathbb{R}}(s) \geqslant a\}$, and by noticing that $[\![\bigvee_{[w_j, a_j] \in \mathbb{W}, a_j \geqslant a} w_j]\!] = \bigcup_{a_j \geqslant a} [\![w_j]\!] = \{s: \delta_{\mathbb{W}}(s) \geqslant a\}$ and $[\![\bigwedge_{(\neg r_i, a_i) \in \mathbb{R}, a_i > 1-a} \neg r_i]\!] = \bigcap_{a_i > 1-a} [\![\neg r_i]\!] = \{s: \pi_{\mathbb{R}}(s) \geqslant a\}$.

3. Merging multiple agents preferences in a bipolar representation

This section recalls methods for the possibilistic merging of agents' preferences from the semantic and the syntactic points of view. The result of the merging process will also be a pair $(\mathbb{R}_{\oplus_{\mathbb{R}}}, \mathbb{W}_{\oplus_{\mathbb{W}}})$, where $\mathbb{R}_{\oplus_{\mathbb{R}}}$ is the result of merging negative preferences expressed by several agents, and $\mathbb{W}_{\oplus_{\mathbb{W}}}$ is the result of merging the agents' positive preferences. In the spirit of Cumulative Prospect Theory [42], merging positive preferences and negative ones are processed in separate steps, using generally different merging operators. Then the paper discusses how to revise the set of positive preferences when it is not coherent with the negative ones.

3.1. Fusion of negative preferences

Let $\{\mathbb{R}_1, \ldots, \mathbb{R}_n\}$ be a set of negative preference bases provided by n agents to be merged with some merging operator $\oplus_{\mathbb{R}}$. $\oplus_{\mathbb{R}}$ is a function from $[0, 1]^n$ to $[0, 1]$. Since negative preferences have an immediate encoding in terms of possibilistic knowledge bases, we can apply the merging procedures of possibilistic knowledge bases [7] for merging negative preference bases. In particular, this allows to have the syntactic counterpart of any semantic merging operator, satisfying minimal properties (see [7] for details).

In this section, we define a class of operators which seem to be appropriate for merging negative preferences. The idea is that since negative preferences determine the level of tolerability of solutions then if some solution s is somewhat rejected by some agent (i.e., $\pi_{\mathbb{R}_i}(s) < 1$) then this solution is somewhat rejected in the result of merging. Such kind of behaviour is captured by *conjunctive* operators.

Let s be a solution. To see if s is explicitly stated as tolerated (or not) and with what degree of tolerability in $\mathbb{R}_{\oplus_{\mathbb{R}}}$ (the result of merging), we compute from each $\mathbb{R}_i$ its level of tolerability a_i. Then, s will be tolerated in $\mathbb{R}_{\oplus_{\mathbb{R}}}$ with a level $\oplus_{\mathbb{R}}(a_1, \ldots, a_n)$.

Natural properties to require for $\oplus_{\mathbb{R}}$ are:

(i) $\oplus_{\mathbb{R}}(1, \ldots, 1) = 1$.
If a solution is fully tolerated by all agents then it should be fully tolerated in the merging result $\mathbb{R}_{\oplus_{\mathbb{R}}}$.

(ii) If $\forall i = 1, \ldots, n$, $a_i \geqslant b_i$ then $\oplus_{\mathbb{R}}(a_1, \ldots, a_n) \geqslant \oplus_{\mathbb{R}}(b_1, \ldots, b_n)$ (monotonicity property).

(iii) $\oplus_{\mathbb{R}}(1, \ldots, 1, a, 1, \ldots, 1) = a$.
Namely, if a solution is partially rejected by at least one agent then its level of tolerability should not increase in the result of the merging process. In fact, it will be tolerated to the same level.

Note that (ii) and (iii) imply:

$$\text{if } a_i = 0 \text{ for some } i \text{ then } \oplus_{\mathbb{R}}(a_1, \ldots, a_n) = 0,$$

which corresponds to say that if a solution is unacceptable for at least one agent then it is also unacceptable in the result obtained after merging.

Let $\mathbb{R}_1 = \{(\neg r_i, a_i): i = 1, \ldots, n\}$ and $\mathbb{R}_2 = \{(\neg r'_j, b_j): j = 1, \ldots, m\}$ be two sets of rejection statements, and $\pi_{\mathbb{R}_1}$ and $\pi_{\mathbb{R}_2}$ be their associated possibility distributions respectively. Let $\oplus_{\mathbb{R}}$ be a merging operator satisfying the properties (i) and (ii). Then, the syntactic counterpart of $\pi_{\oplus_{\mathbb{R}}} = \oplus_{\mathbb{R}}(\pi_{\mathbb{R}_1}, \pi_{\mathbb{R}_2})$ is

$$\begin{aligned}\mathbb{R}_{\oplus} = &\{(\neg r_i, 1 - \oplus_{\mathbb{R}}(1 - a_i, 1)):(\neg r_i, a_i) \in \mathbb{R}_1\} \\ &\cup \{(\neg r'_j, 1 - \oplus_{\mathbb{R}}(1, 1 - b_j)):(\neg r'_j, b_j) \in \mathbb{R}_2\} \\ &\cup \{(\neg r_i \vee \neg r'_j, 1 - \oplus_{\mathbb{R}}(1 - a_i, 1 - b_j)): \\ &(\neg r_i, a_i) \in \mathbb{R}_1 \text{ and } (\neg r'_j, b_j) \in \mathbb{R}_2\}.\end{aligned}$$

If $\oplus_{\mathbb{R}}$ also satisfies (iii) then $\mathbb{R}_{\oplus_{\mathbb{R}}}$ is of the form:

$$\begin{aligned}\mathbb{R}_{\oplus_{\mathbb{R}}} = \mathbb{R}_1 \cup \mathbb{R}_2 \cup &\{(\neg r_i \vee \neg r'_j, 1 - \oplus_{\mathbb{R}}(1 - a_i, 1 - b_j)): \\ &(\neg r_i, a_i) \in \mathbb{R}_1 \text{ and } (\neg r'_j, b_j) \in \mathbb{R}_2\}.\end{aligned}$$

It is also easy to check that $\oplus_{\mathbb{R}}(a_1, \ldots, a_n) \leqslant \min(a_1, \ldots, a_n)$. Indeed, from (ii) and (iii), we have $\oplus_{\mathbb{R}}(a_1, \ldots, a_n) \leqslant \oplus_{\mathbb{R}}(1, \ldots, 1, a_i, 1, \ldots, 1) = a_i$. Hence, $\oplus_{\mathbb{R}} = \min$ represents the most cautious merging operator, in the sense that a solution is not tolerated more than what is stated by the most demanding agent (who rejects s the most strongly). In this case, we simply have (for two bases):

$$\mathbb{R}_{\min} = \mathbb{R}_1 \cup \mathbb{R}_2.$$

Another conjunctive operator is the *product* (then the toleration scale is numerical). Using this operator, a solution which is weakly tolerated by different agents can be strongly unacceptable in the merging result. At the syntactic level, $\mathbb{R}_{\oplus_{\mathbb{R}}}$ is of the form:

$$\begin{aligned}\mathbb{R}_{\oplus_{\mathbb{R}}} = \mathbb{R}_1 \cup \mathbb{R}_2 \cup &\{(\neg r_i \vee \neg r'_j, a_i + b_j - a_i b_j): \\ &(\neg r_i, a_i) \in \mathbb{R}_1 \text{ and } (\neg r'_j, b_j) \in \mathbb{R}_2\}.\end{aligned}$$

Note that the product operator allows the value $a_i + b_j - a_i b_j$ to go beyond $\max(a_i, b_j)$ without reaching 1 if both a_i and b_j differ from 1. That is if a formula r is inferred from $\mathbb{R}_1$ and $\mathbb{R}_2$ with the weights a_1 and a_2 respectively then it is inferred from $\mathbb{R}_{\oplus_{\mathbb{R}}}$ with a weight greater than $\max(a_1, a_2)$, without reaching the highest priority namely 1, which expresses a reinforcement on the priority of r.

Example 3. Let $\mathbb{R}_1$ be the base given in Example 1, namely $\mathbb{R}_1 = \{(g, .5)\}$. Let $\mathbb{R}_2 = \{(p_1, .5)\}$. Let $\oplus_{\mathbb{R}} = \min$. Then, $\mathbb{R}_{\min} = \{(p_1, .5), (g, .5)\}$.

3.2. Fusion of positive preferences

We now discuss the merging of positive preferences. We first recall a recent general result on the syntactic fusion of positive preferences [9], similar to the one in [7] for standard possibilistic knowledge bases.

Let $\mathbb{W}_1, \ldots, \mathbb{W}_m$ be m bases of positive preferences and $\delta_{\mathbb{W}_1}, \ldots, \delta_{\mathbb{W}_m}$ be their associated positive possibility distributions given by Definition 2. Let $\oplus_{\mathbb{W}}$ be a merging operator satisfying the following requirements:

- $\oplus_{\mathbb{W}}(0, \ldots, 0) = 0$.
- If $\forall j = 1, \ldots, m,\ a_j \geqslant b_j$ then $\oplus_{\mathbb{W}}(a_1, \ldots, a_m) \geqslant \oplus_{\mathbb{W}}(b_1, \ldots, b_m)$.

The first requirement expresses that if a solution is not satisfactory for any agent then it should not be satisfactory in the result of the merging. The second property is simply the monotonicity property.

Let us restrict ourselves, for the sake of simplicity, to the case of two bases. Then, the following proposition gives the positive preference base associated with $\oplus_{\mathbb{W}}(\delta_{\mathbb{W}_1}, \delta_{\mathbb{W}_2})$:

Proposition 2. *Let* $\mathbb{W}_1 = \{[w_i, a_i] : i = 1, \ldots, n\}$ *and* $\mathbb{W}_2 = \{[w'_j, b_j] : j = 1, \ldots, m\}$ *be two bases of positive preferences. Let* $\delta_{\mathbb{W}_1}$ *and* $\delta_{\mathbb{W}_2}$ *be their associated positive possibility distributions respectively. Let* $\oplus_{\mathbb{W}}$ *be a merging operator. Then, the base of positive preferences associated with* $\oplus_{\mathbb{W}}(\delta_{\mathbb{W}_1}, \delta_{\mathbb{W}_2})$ *is*:

$$\begin{aligned}\mathbb{W}_{\oplus_{\mathbb{W}}} = &\{[w_i, \oplus(a_i, 0)] : [w_i, a_i] \in \mathbb{W}_1\} \cup \{[w'_j, \\ &\oplus(0, b_j)] : [w'_j, b_j] \in \mathbb{W}_2\} \cup \{[w_i \wedge w'_j, \\ &\oplus(a_i, b_j)] : [w_i, a_i] \in \mathbb{W}_1 \text{ and } [w'_j, b_j] \in \mathbb{W}_2\}.\end{aligned}$$

The solution of the merging operator for combining $\delta_{\mathbb{W}_1}$ and $\delta_{\mathbb{W}_2}$ is less constrained since several merging operators can be considered such as conjunctive, disjunctive and also "intermediary" operators which reinforce what is common and discount the positive preferences which only concern one agent (see [26] for a representation of such operators). If the agents are highly cooperative then we can say that an agent adds to its positive preferences those of the other agent, provided that they do not conflict with what is tolerated for him. In this case $\oplus_{\mathbb{W}} = \max$ is recommended, and

$$\mathbb{W}_{\max} = \mathbb{W}_1 \cup \mathbb{W}_2.$$

Example 4. Let us consider the two following bases of positive preferences: $\mathbb{W}_1 = \{[m, 1], [\neg apart \wedge l, 1], [\neg g \wedge l, .6]\}$ and $\mathbb{W}_2 = \{[\neg apart, 1], [p_1 \wedge (apart \vee s), 1], [\neg g \wedge apart \wedge l, .4]\}$.

Let $\mathbb{W}_{\max}$ be the result of combining $\mathbb{W}_1$ and $\mathbb{W}_2$ with $\oplus_{\mathbb{W}} = \max$. Then, $\mathbb{W}_{\max} = \{[m, 1], [\neg apart \wedge l, 1], [\neg g \wedge l, .6]\} \cup \{[\neg apart, 1], [p_1 \wedge (apart \vee s), 1], [\neg g \wedge apart \wedge l, .4]\}$ which is semantically equivalent to

$$\mathbb{W}_{\max} = \{[m, 1], [\neg apart, 1], [p_1 \wedge (apart \vee s), 1], [\neg g \wedge l, .6]\}.$$

Indeed, the formula $[\neg g \wedge apart \wedge l, .4]$ is subsumed by $[\neg g \wedge l, .6]$ since $\neg g \wedge apart \wedge l \vdash \neg g \wedge l$.

Note that this mode of merging corresponds to considering that what is satisfactory for one agent is also satisfactory for the others.

3.3. Restoring consistency as characterizing best solutions

In the merging process, the consistency of each pair $(\mathbb{W}_i, \mathbb{R}_i)$ does not guarantee the coherence of the pair $(\mathbb{W}, \mathbb{R})$, where $\mathbb{W}$ (resp. $\mathbb{R}$) is the result of merging $\mathbb{W}_i$'s (resp. $\mathbb{R}_i$'s), for most of the operators $(\oplus_{\mathbb{W}}, \oplus_{\mathbb{R}})$.

Example 5. Let $\mathbb{R}_{\min}$ and $\mathbb{W}_{\max}$ be respectively the sets of merged negative and merged positive preferences computed in Examples 3 and 4 respectively. We have $\mathbb{R}_{\min} = \{(p_1, .5), (g, .5)\}$, and $\mathbb{W}_{\max} = \{[m, 1], [\neg apart, 1], [p_1 \wedge (apart \vee s), 1], [\neg g \wedge l, .6]\}$.

At the semantic level, consider the solution $s_0 = p_1 m \neg apart \neg g$. Then, we can check that $\pi_{\mathbb{R}}(s_0) = .5$ while $\delta_{\mathbb{W}}(s_0) = 1$.

Observe that $\delta_{\mathbb{W}}(s_0) \not\leqslant \pi_{\mathbb{R}}(s_0)$, then $\delta_{\mathbb{W}}$ and $\pi_{\mathbb{R}}$ are not coherent.

When the coherence condition is not satisfied by the results of the merging process, it means that the set of positive preferences resulting from merging the positive preferences of the agents is not compatible with what is tolerated by the agents. One way to restore the coherence in the sense of Definition 3 is to revise either the positive possibility distribution $\delta_{\mathbb{W}}$ or the possibility distribution $\pi_{\mathbb{R}}$. We choose to revise $\delta_{\mathbb{W}}$ since in practice it is more difficult to question $\pi_{\mathbb{R}}$, as it expresses what remains tolerated, according to agents' negative preferences. The revision of $\delta_{\mathbb{W}}$ in this case consists in decreasing the possibility degree of each solution in $\delta_{\mathbb{W}}$ to the possibility degree of this solution in $\pi_{\mathbb{R}}$. In other terms, restoring coherence leads to revise $\delta_{\mathbb{W}}$ into $\delta_{\mathbb{W}_{\text{rev}}}$ as follows:

$$\forall s, \delta_{\mathbb{W}_{\text{rev}}}(s) = \min(\delta_{\mathbb{W}}(s), \pi_{\mathbb{R}}(s)).$$

At the syntactic level, this leads to discounting positive preferences by decreasing the levels of satisfaction associated with some positive preference statements in $\mathbb{W}$.

In order to syntactically compute the revised set of positive preferences $\mathbb{W}_{\text{rev}}$ associated with $\delta_{\mathbb{W}_{\text{rev}}}$, we need the following proposition which gives an equivalent translation of a set of negative preferences into a set of positive preferences [9]:

Proposition 3. *Let* $\mathbb{R} = \{(\neg r_i, a_i): i = 1, \ldots, n\}$ *be a set of negative preferences where each level contains one formula,*[2] *and such that* $1 \geqslant a_1 > \cdots > a_n > 0$ *and we let* $a_{n+1} = 0$. *Let* $\mathbb{W}_{\mathbb{R}}$ *be a set of positive preferences defined as follows*:

$$\mathbb{W}_{\mathbb{R}} = \{[\neg r_1 \wedge \cdots \wedge \neg r_i, 1 - a_{i+1}] : i = 1, \ldots, n\} \cup \{[\top, 1 - a_1]\}.$$

Then, $\mathbb{R}$ *and* $\mathbb{W}_{\mathbb{R}}$ *are semantically equivalent i.e.,* $\forall s, \pi_{\mathbb{R}}(s) = \delta_{\mathbb{W}_{\mathbb{R}}}(s)$.

The proof of this proposition can be found in [9]. Given this proposition the syntactic computation of $\mathbb{W}_{\text{rev}}$ can be achieved in two steps: (i) first compute $\mathbb{W}_{\mathbb{R}}$ associated with $\mathbb{R}$, (ii) then apply Proposition 2 to $\mathbb{W}_{\mathbb{R}}$ and $\mathbb{W}$ for $\oplus_{\mathbb{W}} = \min$. Note that when $\oplus_{\mathbb{W}} = \min$, Proposition 2 becomes:

$$\mathbb{W}_{\oplus_{\min}} = \{[w_i \wedge w'_j, \min(a_i, b_j)] : [w_i, a_i] \in \mathbb{W}_1 \text{ and } [w'_j, b_j] \in \mathbb{W}_2\}.$$

Example 6. Let us again consider the possibility distributions $\delta_{\mathbb{W}}$ and $\pi_{\mathbb{R}}$ associated with $\mathbb{W}$ and $\mathbb{R}$ computed in Examples 3 and 4 respectively. We have $\mathbb{R} = \{(p_1, .5), (g, .5)\}$ and $\mathbb{W} = \{[m, 1], [\neg apart, 1], [p_1 \wedge (apart \vee s), 1], [\neg g \wedge l, .6]\}$.

Let $\delta_{\mathbb{W}_{\text{rev}}} = \min(\delta_{\mathbb{W}}, \pi_{\mathbb{R}})$.

First, we use Proposition 3 to equivalently translate $\mathbb{R}$ into a new set of positive preferences, we get:

$$\mathbb{W}_{\mathbb{R}} = \{[p_1 \wedge g, 1], [\top, .5]\}.$$

Now the application of Proposition 2 between $\mathbb{W}_{\mathbb{R}}$ and $\mathbb{W}$ using $\oplus_{\mathbb{W}} = \min$ gives:

$$\mathbb{W}_{\text{rev}} = \{[p_1 \wedge g \wedge m, 1], [p_1 \wedge g \wedge \neg apart, 1], [p_1 \wedge g \wedge (apart \vee s), 1], [p_1 \wedge g \wedge \neg g \wedge l, .6], [m, .5], [\neg apart, .5], [p_1 \wedge (apart \vee s), .5], [\neg g \wedge l, .5]\}.$$

In Example 5 we have seen that $\delta_{\mathbb{W}}$ and $\pi_{\mathbb{R}}$ are not coherent, since: $\delta_{\mathbb{W}}(s_0) = 1 > \pi_{\mathbb{R}}(s_0) = .5$ for $s_0 = p_1 m \neg apart \neg g$.

Now it can be checked that $\delta_{\mathbb{W}_{\text{rev}}}$ and $\pi_{\mathbb{R}}$ are coherent. In particular, $\delta_{\mathbb{W}_{\text{rev}}}(s_0) = .5$.

4. Finding best solutions according to negative and positive preferences

Already a long time ago, Bellman and Zadeh [2] proposed a graded view both for criteria and constraints. This was the starting point for introducing the idea of flexible (or soft) constraint, which was further developed

[2] The fact that we assume that each layer is composed of a unique formula is not a limitation. Indeed, a set of formulas having the same weight can equivalently be replaced by a unique formula, with the same weight, and corresponding to the conjunction of these formulas.

in a constraint satisfaction problem perspective later (e.g. [37,22,11]). However, Bellman and Zadeh's approach tends to somewhat abolish the distinction between criteria and constraints, by modeling both of them by means of fuzzy sets and by combining all these fuzzy sets together. This is debatable, since a constraint is something which should be satisfied (at least to some extent for a flexible constraint, or as far as it does not lead to an empty set of feasible solutions for a constraint whose priority is not maximal). On the contrary, there is no idea of (minimal) requirement associated with a criterion.

Flexible constraints modeled by fuzzy sets can be approximated by a collection of nested prioritized constraints or negative preferences [22]. Each prioritized constraint corresponds then to a level cut of the fuzzy set. Such "fuzzy" constraints correspond to "negative" preferences in the sense that their complements define fuzzy sets of values that are rejected as being non-acceptable. These constraints should be combined in a conjunctive manner, thus acknowledging the fact that they are constraints. However, criteria correspond to another type of preferences, which can be qualified as "positive" ones. These preferences are not constraints, but can be expressed in terms of positive preferences, which are more or less strong (in the sense that it is more satisfactory to satisfy some positive preferences). If at least some of these positive preferences are satisfied, it should give some bonus to the corresponding solutions (provided that they also satisfy the constraints). Positive preferences are not compulsory and may be combined disjunctively. This view has been recently applied in flexible querying systems where the retrieved items are first rank-ordered w.r.t. the satisfaction degree of the flexible requirements contained in the request. In case of ties, the items are further discriminated on the basis of the additional positive preferences (also included in the request) that they satisfy [28].

Thus, the problem of computing the best solutions after merging the negative and the positive preferences separately can be also viewed as an optimization problem involving the sets $\mathbb{F}$, $\mathbb{R}$ and $\mathbb{W}$. So our aim in this section is, given the triple $(\mathbb{F}, \mathbb{R}, \mathbb{W})$, to compute the logical description of the set of "preferred" solutions which do not violate integrity constraints, allow to respect all agent's negative preferences $\mathbb{R}$, and satisfy as many as possible agent's positive preferences $\mathbb{W}$. For this, we will proceed by steps, first considering the case of *flat* bases $\mathbb{R}$ and $\mathbb{W}$, and then dealing with the general case of *prioritized* bases.

4.1. Binary negative and positive preferences

In this section all the negative preferences are supposed to be compulsory (i.e. $\forall i$, $a_i = 1$) and the positive preferences have the same importance (i.e. $\forall j$, $b_j = 1$); so we can omit the weights. Moreover, in the following as well as in the next section, the aggregation of $(\mathbb{F}, \mathbb{R}, \mathbb{W})$ first considers the feasible solutions obtained from $\mathbb{F}$, then should give priority to $\mathbb{R}$, before refining obtained solutions with respect to $\mathbb{W}$. We assume that $\mathbb{R}$ is consistent, while the consistency condition is not required for $\mathbb{W}$. Therefore, the existence of preferred solutions depends on whether $\mathbb{R}$ is consistent or not. In case of inconsistency, agent's preferences are declared to be unfeasible, and no solution is considered as tolerated.

Now, let us assume that $\mathbb{R}$ is consistent, and we denote by $\mathbb{S}_{\mathbb{R}}$ the set of all solutions satisfying (i.e., models of) $\mathbb{R}$. In this case, they are the solutions satisfying agent's constraints induced by negative preferences, and the "preferred" solutions w.r.t. positive preferences should be looked for inside $\mathbb{S}_{\mathbb{R}}$. One may think of the three following approaches for selecting *preferred* solutions. In all these approaches, the aim is to try to verify as many agent's positive preferences as possible.

(i) *Conjunctive selection: Satisfying all positive preferences*: A preferred solution s is a preferred solution if it satisfies all agent's positive preferences. This approach is too requiring, would abolish the distinction between negative and positive preferences, and is not desirable since it can lead to the absence of solutions (remind that $\mathbb{W}$ is not even required to be consistent), while as soon as $\mathbb{F} \wedge \mathbb{R}$ is consistent, tolerated solutions exist w.r.t. agent's preferences, and thus should be proposed.

(ii) *Disjunctive selection: Verifying at least one positive preference*: A tolerated solution s is preferred as soon as it satisfies at least one positive preference. More precisely, let $\mathbb{R}^* = \bigwedge_{(\neg r_i, a_i)\in\mathbb{R}} \neg r_i$ and $\mathbb{W}_* = \bigvee\{w_j \in \mathbb{W}\}$. Then, solutions to agent's preferences $(\mathbb{F}, \mathbb{R}, \mathbb{W})$ are

$$\begin{cases} [\![\mathbb{F} \wedge \mathbb{R}^* \wedge \mathbb{W}_*]\!] & \text{if } \mathbb{F} \wedge \mathbb{R}^* \wedge \mathbb{W}_* \text{ is consistent,} \\ [\![\mathbb{F} \wedge \mathbb{R}^*]\!] & \text{otherwise.} \end{cases}$$

This solution guarantees the existence of solutions (up to inconsistency) but does not make any distinction between solutions verifying one positive preference, with the ones verifying several agent's positive preferences (if any). Let us illustrate the approach by again an example about an agent looking for a house or an apartment.

Example 7

- Assume that integrity constraints are:
  - The price of house or apartment of surface l is either p_2 or p_3, but cannot be p_1.
  - The price of a house or apartment exclusively belongs to one of the three price categories, and the surface is exclusively in one of the three surface categories (s, m, l).

Formally, $\mathbb{F}$ is written as:

$$\mathbb{F} = \{l \Rightarrow \neg p_1, s \vee m \vee l, \neg s \vee \neg m, \neg s \vee \neg l, \neg m \vee \neg l, p_1 \vee p_2 \vee p_3, \neg p_1 \vee \neg p_2, \neg p_2 \vee \neg p_3, \neg p_1 \vee \neg p_3\}.$$

Indeed, feasible solutions satisfy $(\neg l \vee \neg p_1) \wedge (s \vee m \vee l) \wedge (\neg s \vee \neg m) \wedge (\neg s \vee \neg l) \wedge (\neg m \vee \neg l) \wedge (p_1 \vee p_2 \vee p_3) \wedge (p_1 \vee p_2) \wedge (p_1 \vee p_3) \wedge (p_2 \vee p_3)$.

- Assume now that agent's negative preferences are:
  - (i) he can only accept houses or apartments with a price p_1,
  - (ii) he rejects small apartments.

  Note that $\neg apart$ encodes "having a house".

  Then the set of negative preferences $\mathbb{R}$ is written as: $\mathbb{R} = \{p_1, \neg apart \vee \neg s\}$.
- Assume that agent's positive preferences are:
  - (i) he prefers apartments or houses with surface m,
  - (ii) he prefers an apartment with a price p_1,
  - (iii) he prefers a house without garden and to pay the price p_1.

  This is formally written as:

$$\mathbb{W} = \{m, p_1 \wedge apart, p_1 \wedge \neg apart \wedge \neg g\}.$$

We assume that the agent is indifferent w.r.t. having a garden or not.

Let us compute the set of preferred solutions. Feasible solutions should first satisfy $\mathbb{R}^*$ which is equivalent to satisfying $p_1 \wedge (\neg apart \vee m) \wedge (s \vee m) \wedge (\neg s \vee \neg m)$.

Moreover, preferred solutions should also satisfy at least one of agent's positive preferences. Namely they satisfy $p_1 \wedge (\neg apart \vee m) \wedge (s \vee m) \wedge (\neg s \vee \neg m) \wedge (m \vee (p_1 \wedge apart) \vee (p_1 \wedge \neg apart \wedge \neg g))$ which is equivalent to $p_1 \wedge (\neg apart \vee m) \wedge (m \vee \neg g) \wedge (s \vee m) \wedge (\neg s \vee \neg m)$.

Indeed preferred solutions are: s_1: $p_1 \neg apart\, m \neg g$, s_2: $p_1 \neg apart\, m\, g$, s_3: $p_1 \neg apart\, s \neg g$, s_4: $p_1\, apart\, m \neg g$ and s_5: $p_1\, apart\, m\, g$ i.e., $(p_1 \wedge m) \vee (p_1 \wedge \neg apart \wedge s \wedge \neg g)$.

(iii) *Cardinality-based selection: Satisfying as many positive preferences as possible:* A third approach is to consider that the more positive preferences a solution satisfies, the better it is. More precisely, s is preferred to s' if $|s|_{\mathbb{W}} > |s'|_{\mathbb{W}}$, where $|s|_{\mathbb{W}}$ is the number of positive preferences satisfied by s. A solution s is said to be a "cardinality-based preferred" solution, if there is no s' such that $|s'|_{\mathbb{W}} > |s|_{\mathbb{W}}$. The cardinality-based approach can be seen as a weakening of the conjunctive approach when there is no tolerated solution which satisfies $\mathbb{R}^* \wedge \mathbb{W}^*$, where $\mathbb{W}^* = \{w_j : [w_j, b_j] \in \mathbb{W}\}$.

Example 7 (*continued*). Let us consider again the above example. We observe that s_1, s_4 and s_5 satisfy two positive preferences: w_1 and w_3 for s_1 and w_1, w_2 for s_4 and s_5 while s_2 and s_3 satisfy only one positive preference: w_1 and w_3 respectively. Hence, using the cardinality-based selection mode, only s_1, s_4 and s_5 are preferred.

Note that the set of preferred solutions w.r.t. the cardinality-based selection can be computed syntactically using results in [10]. The main idea behind the algorithm is to replace lexicographic entailment from the original base by classical entailment from a new "compiled" base, which contains either formulas from the original base or formulas subsumed by the original ones, obtained from the disjunction of some original formulas. More precisely, when there is a conflict in a knowledge base, then pieces of information are weakened instead of being removed. So instead of removing all formulas involved in conflicts, we take their disjunctions pairwise. If the result is consistent, then the algorithm stops. If the result is still inconsistent, then we replace the formulas in conflict by all possible disjunctions involving three formulas in the conflict sets and so on: if the result is still inconsistent we successively consider disjunctions of size 4, 5, etc. (see [10] for more details).

More precisely, when $\mathbb{F} \wedge \mathbb{R} \wedge \mathbb{W}$ is inconsistent, we compute the set $\mathbb{F} \wedge \mathbb{R} \wedge d_2(\mathbb{W})$, where $d_2(\mathbb{W})$ is the set of all possible disjunctions of size 2 between formulas of $\mathbb{W}$. If this set is consistent then its models are the preferred solutions w.r.t. the cardinality-based selection. Otherwise, we compute $\mathbb{F} \wedge \mathbb{R} \wedge d_i(\mathbb{W})$ $(i \geqslant 3)$ until reaching consistency. If the disjunction of all formulas in $\mathbb{W}$ is still inconsistent with $\mathbb{R}$ then the set of preferred solutions is simply the set of feasible tolerated solutions i.e. models of $\mathbb{F} \wedge \mathbb{R}$.

Example 7 (*continued*). Recall that tolerated solutions are models of $\mathscr{T} = p_1 \wedge (\neg apart \vee m) \wedge (s \vee m) \wedge (\neg s \vee \neg m)$. Let $\mathbb{W} = \{m, p_1 \wedge apart, p_1 \wedge \neg apart \wedge \neg g\}$. $\mathscr{T} \wedge \mathbb{W}$ is inconsistent so we compute $d_2(\mathbb{W})$, the set of all possible disjunctions pairwise between different formulas of $\mathbb{W}$. We have $d_2(\mathbb{W}) = \{m \vee (p_1 \wedge \neg apart), m \vee (p_1 \wedge \neg apart \wedge \neg g), (p_1 \wedge apart) \vee (p_1 \wedge \neg apart \wedge \neg g)\}$. $\mathscr{T} \wedge d_2(\mathbb{W})$ is consistent and equivalent to $p_1 \wedge m \wedge (apart \vee \neg g)$. Then the preferred solutions are indeed s_1, s_4 and s_5.

4.2. Prioritized negative and positive preferences

This section extends results of the previous section to the case where both negative and positive preferences are prioritized. Recall that negative preferences are the agent's constraints, which provide additional restrictions on the set of feasible solutions. Any tolerated solution satisfying all agent's negative preferences (i.e., satisfying no rejection statement) is entirely tolerated. Solutions not respecting hard negative preferences are completely unacceptable. Any solution verifying at least one agent's rejection statement is unacceptable to some degree. The higher the priority of a verified rejection statement, the less tolerated is the solution. On the contrary, positive preferences represent agent's desires. If a solution satisfies them, then it increases the degree of satisfaction of the agent. However if a solution does not satisfy the agent's positive preferences, then it does not mean that such a solution should be excluded.

4.2.1. Finding preferred solutions among best tolerated ones

The selected tolerated solutions are the feasible ones which maximize $\pi_{\mathbb{R}}$. Namely,

$$\mathscr{T}_{\mathbb{R}} = \left\{ s\colon s \models \mathbb{F} \text{ and } \pi_{\mathbb{R}}(s) = \max_{s' \models \mathbb{F}} \pi_{\mathbb{R}}(s') \right\}.$$

At the syntactic level, $\mathscr{T}_{\mathbb{R}}$ is the set of models of $\mathbb{F} \wedge \mathbb{R}_{>a}$, where $a = \max\{a_i : \mathbb{F} \wedge \mathbb{R}_{\geqslant a_i} \text{ is inconsistent}\}$, with $\max\{\emptyset\} = 0$, and where $\mathbb{R}_{>a} = \{\neg r_i : (\neg r_i, a_i) \in \mathbb{R} \text{ and } a_i > a\}$ and $\mathbb{R}_{\geqslant a} = \{\neg r_i : (\neg r_i, a_i) \in \mathbb{R} \text{ and } a_i \geqslant a\}$ (see Appendix A for more details).

Now, we apply a selection mode to positive preferences. A disjunctive selection mode could be used where preferred solutions are those which belong to $\mathscr{T}_{\mathbb{R}}$ and satisfy at least one important positive preference, i.e., a positive preference having the highest degree in $\mathbb{R}$. If such a solution does not exist, the preferred solutions are those which belong to $\mathscr{T}_{\mathbb{R}}$ and satisfy a positive preference with the second highest degree (if any) in $\mathbb{W}$. And if such a solution does not exist, we consider positive preferences having the third highest degree (if any), and so on. More formally, the preferred solutions are defined as follows:

$$\left\{ s\colon s \in \mathscr{T}_{\mathbb{R}} \text{ and } \delta_{\mathbb{W}}(s) = \max_{s' \in \mathscr{T}_{\mathbb{R}}} \delta_{\mathbb{W}}(s') \right\}.$$

At the syntactic level, preferred solutions can be easily defined. Let $\mathbb{W}_{*\alpha} = \bigvee\{w_j : [w_j, b_j] \in \mathbb{W} \text{ and } b_j = \alpha\}$. Then, preferred solutions are:

$$\begin{cases} [\![\mathbb{F} \wedge \mathbb{R}_{>a}]\!] & \text{if } \nexists\alpha, \mathbb{F} \wedge \mathbb{R}_{>a} \wedge \mathbb{W}_{*\alpha} \text{ is consistent,} \\ [\![\mathbb{F} \wedge \mathbb{R}_{>a} \wedge \mathbb{W}_{*\alpha}]\!] & \text{if } \mathbb{F} \wedge \mathbb{R}_{>a} \wedge \mathbb{W}_{*\alpha} \text{ is consistent and} \\ & \quad \nexists\beta > \alpha, \mathbb{F} \wedge \mathbb{R}_{>a} \wedge \mathbb{W}_{*\beta} \text{ is consistent.} \end{cases}$$

Example 8. We consider the set $\mathbb{F}$ given in Example 7 and the following bases $\mathbb{R}$ and $\mathbb{W}$: $\mathbb{R} = \{(p_1, 1), (\neg apart \vee \neg s, 1), (l, .7), (\neg p_1 \vee s, .4)\}$ and $\mathbb{W} = \{[m, 1], [p_1 \wedge apart, 1], [p_1 \wedge \neg apart \wedge \neg g, .8]\}$.

First, the set of tolerated solutions is the set of feasible solutions which satisfy $\mathbb{R}_{>.7}$ which is equivalent to satisfying $p_1 \wedge (\neg apart \vee m) \wedge (s \vee m) \wedge (\neg s \vee \neg m)$.

Now preferred solutions having the highest degree in $\delta_{\mathbb{W}}$ are those which satisfy $m \vee (p_1 \wedge apart)$ which corresponds to the disjunction of the highest positive preferences in $\mathbb{W}$. Indeed preferred solutions are models of $p_1 \wedge (\neg apart \vee m) \wedge (s \vee m) \wedge (\neg s \vee \neg m) \wedge (m \vee \ \ apart)$, which is equivalent to $p_1 \wedge m$, namely: s_1: $p_1 \neg apart\, m \neg g$, s_2: $p_1 \neg apart\, m g$, s_3: $p_1\, apart\ m \neg g$ and s_4: $p_1\, apart\, m g$.

4.2.2. Cardinality-based selection mode

Let us now consider the more refined selection mode based on cardinality selection. The idea is to apply the cardinality-based selection mode by both maximizing the number of falsified rejection statements and maximizing the number of satisfied positive preferences. Namely, the procedure for selecting preferred solutions follows two steps:

(i) First, select solutions falsifying as many prioritary constraints as possible. Namely, let s and s' be two feasible solutions. Then [3]:

Definition 4. $s >_{\text{leximin},\mathbb{R}} s'$ iff:
if $\exists a_1, \lceil c\rceil_{\mathbb{R}}^{a_1} > \lceil c'\rceil_{\mathbb{R}}^{a_1}$ then $\forall a_2 > a_1, \lceil c\rceil_{\mathbb{R}}^{a_2} = \lceil c'\rceil_{\mathbb{R}}^{a_2}$,
where $\lceil c\rceil_{\mathbb{R}}^{a_i}$ is the number of constraints in $\mathbb{R}$ satisfied by s to a degree a_i.

And $s =_{\text{leximin},\mathbb{R}} s'$ $\forall a_i, \lceil s\rceil_{\mathbb{R}}^{a_i} = \lceil s'\rceil_{\mathbb{R}}^{a_i}$.

(ii) Then, among solutions selected at the first step, choose those satisfying a maximal number of positive preferences (in the leximax style).

More precisely, let s and s' be two feasible solutions. Then, s is preferred to s' if:

(i) $s >_{\text{leximin},\mathbb{R}} s'$ (Definition 4), or
(ii) $s =_{\text{leximin},\mathbb{R}} s'$ and $c >_{\text{leximax},\mathbb{W}} s'$, where $s >_{\text{leximax},\mathbb{W}} s'$ iff $\exists b_1, \mid c|_{\mathbb{W}}^{b_1} > \mid s'|_{\mathbb{W}}^{b_1}$ and $\forall b_2 > b_1, \mid c|_{\mathbb{W}}^{b_2} = \mid s'|_{\mathbb{W}}^{b_2}$.

Example 8 (*continued*). Solutions induced by leximin ordering are feasible solutions which satisfy $\mathbb{R}_{>.7} \wedge (\neg p_1 \vee s)$ which is equivalent to satisfying $p_1 \wedge \neg apart \wedge s$. Indeed, there are two tolerated solutions s_1: $p_1 \neg apart\, s \neg g$ and s_2: $p_1 \neg apart\, s g$. Let $W = p_1 \wedge \neg apart \wedge s$ and $\mathbb{W}_\alpha = \{w_j : [w_j, b_j] \in \mathbb{W} \text{ and } b_j = \alpha\}$. W is inconsistent with both $\mathbb{W}_1$ and $d_2(\mathbb{W}_1)$ so we move to $\mathbb{W}_{.8}$. Now, $W \wedge \mathbb{W}_{.8}$ is consistent and equivalent to $p_1 \wedge \neg apart \wedge s \wedge \neg g$ which is the preferred solution.

Example 9. Let us consider the invited speaker's example provided in the introduction section. Let M, T, W, Mo, Af, E be literals which stand respectively for 'Monday', 'Tuesday', 'Wednesday', 'morning', 'afternoon' and 'early'. We have:

$$\mathbb{F} = \{\neg M \vee \neg T, \neg M \vee \neg W, \neg T \vee \neg W, M \vee T \vee W,$$
$$\neg Mo \vee Af, Mo \vee Af, \neg E \vee \neg L, E \vee L\},$$

$$\mathbb{R} = \{(\neg M, .9), (\neg W, .7)\} \text{ and}$$

$$\mathbb{W} = \{[Mo, .5], [E \wedge Mo, .8]\}.$$

Applying the cardinality-based selection mode leads to select $E\ Mo\ T$.

Remarks

(1) Note that even if the sets $\mathbb{F}$ and $\mathbb{R}$ are both modelled in possibilistic logic, we must keep them separate in the scope of the last procedure. Suppose that $\mathbb{F}$ is inconsistent then in this case, the set of preferred solutions is empty. Suppose now that we build a single set $\mathbb{F} \cup \mathbb{R}$ and apply the lexicographic-based selection mode to compute tolerated solutions. Then possible solutions will be ranked w.r.t. $\mathbb{R}$ even if $\mathbb{F}$ is inconsistent.

(2) Positive preferences are not equivalent to a low priority constraint from negative preferences. Consider for instance $\mathbb{R} = \{(\neg p, 1), (\neg q, a), (\neg r, b)\}$ and $\mathbb{W} = \{[w, d]\}$ with $1 > a > b > d$, where p, q, r, w are propositions. A solution which satisfies $\neg p$ and violates $\neg q$ has a possibility degree equal to $1 - a$ even if it fulfils $\neg r$. A solution which satisfies $\neg p$ and w but *not* $\neg q$ will be preferred to a solution which satisfies $\neg p$ and $\neg r$, but satisfies q and violates w. This indicates that positive preferences are used to discriminate among solutions having the same level of tolerability (computed on the basis of the rejections that they falsify). In case $[w, d]$ would be a constraint (i.e., (w, d)), a solution satisfying $\neg p \neg q \neg r$ but *not* w would be preferred to a solution satisfying $\neg p \neg q w$ and r, for instance. Thus the use of the positive preference $[w, d]$ enables us to discriminate between solutions satisfying only important constraints, while considering (w, d) as a constraint only provides a way for choosing between solutions satisfying all the other constraints.

5. Related works

The idea of bipolar representations of preferences was considered in different frameworks, such as Cumulative Prospect Theory [42] multicriteria decision making [29], or in Artificial Intelligence works oriented towards qualitative decision [43,31]. However these works usually assume that the positive and negative parts of the preferences, once they have been specified and represented, can be combined and processed together, rather than separately as it is done in this paper. However classical theories of decision, like multiattribute utility theory of Keeney and Raiffa, or Simon's satisficing approach, do not account for separate treatment of positive and negative preferences. It is handled in a very special way in operations research where it is usual to optimise a criterion or a set of criteria under rigid constraints. Constraints express negative preferences and the criterion embodies the positive preferences of the decision-maker. The usual averaging of several ratings according to various criteria is clearly in the spirit of positive preferences. However constraint satisfaction problems envisage decisions from a pure negative preference point of view. Nevertheless our negative preferences exactly corresponds to what is known in the literature as soft (prioritized) constraints [12]; hence our framework is an extension of the soft constraint approach to positive preferences.

In [32,33,43], the authors propose to view conditional desires as constraints on utility functions. Intuitively, $D(p|q)$ stands for: the q-worlds with highest utility satisfy p. Contrary to Boutilier's approach, Lang et al. [33] use numerical utilities, where the expression $D(p|q)$ induces a loss of utility if it is falsified, and/or a gain of utility if it is satisfied. Then an additive utility function is used to aggregate (sum up) losses and gains induced by each conditional desire. There are at least three differences between Lang et al. [33]'s approach and the one proposed in this paper. First, we start with two different sets of preferences, while in Lang et al. [33]'s approach only one set of preference is used, which corresponds to our negative preference. Second, even if Lang et al. [33]'s approach to preferences are expressed by means of two parameters: utility gain and utility loss, there is no separate treatment of positive and negative preferences since the two parameters are aggregated into one value called *polarity* parameter which is used to rank-order the set of solutions contrary to our approach. Lastly, Lang et al. [33]'s approach is based on some additive aggregation mode (basically, the sum operator), while most of the methods proposed in this paper are qualitative, in the sense that only the ordering between negative (respectively positive) preferences is important.

In [14] (see also [41]), a logic for expressing qualitative preferences was proposed. More precisely, some of the expressed preferences are the following: a *conditional preference* expressing that in the most preferred worlds where p holds, q holds also, a *toleration* expressing that the agent is permitted to do q if p is true, a *relative preference* of two propositions expressing that the best models of p are at least as good as the best models of q, and a *strict preference* expressing that every model of t is preferred to any of its countermodels. It turns out that all these preferences can be expressed in the possibility theory setting by means of constraints on Δ and Π measures. More precisely, a conditional preference is expressed in possibility theory by means of a constraint on Π of the form $\Pi(p \wedge q) > \Pi(p \wedge \neg q)$ which means that in the best models when p is true, it is better to satisfy q than to falsify it. The toleration and the relative preference are modelled by the following constraints $\Delta(p \wedge q) > 0$ and $\Pi(p) \geqslant \Pi(q)$ respectively. Lastly the strict preference is modelled by the constraint $\Delta(t) > \Pi(\neg t)$ which means that any model of t is preferred to all its countermodels. Note that these constraints on Δ and Π may also be expressed by means of prioritized bases as given in this paper. Similar conditional preferences (also called conditional desires) based on Δ and Π constraints were also proposed in [43]. More details on the possibilistic handling of this representation format and their bridges with logical representation of preferences can be found in [6,8].

In [16,17,15] a graphic representation, directly reminiscent of Bayes nets, for conditional preferences among feature values under the *ceteris paribus* [21,45] principle is proposed, together with corresponding algorithms. Basically, the authors use a directed acyclic graph to express conditional preferences: for each variable p, and for each instance of variables that are parents of p, an agent should express preferences between different instances of p. From these conditional preferences, and with help of "ceteris paribus" principle, a partial ordering between possible solutions is defined, and algorithms are provided for computing optimal solutions. In our framework a total ordering between solutions is induced instead of a partial ordering as in [16,17,15]. Moreover the use of minimum and maximum operations make the representation of preference fail with respect to *ceteris paribus* constraints. However the use of lexicographic selection modes restore preferential independence between partial possibility distributions induced by each weighted preference statement.

6. Concluding remarks

This paper advocates the handling of bipolar preferences bases with a separate treatment of positive and negative information under the form of two sets of weighted logical formulas having different semantics, both of them being encoded in the framework of possibility theory. The representation framework remains simple (although each of the two sets could be equivalently represented in a graphical form, or as a set of conditionals [4]). The possibilistic logic approach enables us to compute a syntactic description of the set of preferred solutions, in agreement with the semantics. Besides, the proposed model remains qualitative, since only the ordering between the satisfaction levels or the rejection levels is meaningful, but presupposes the existence of two ordinal value scales for positive and negative preferences.

This kind of representation of bipolarity could accommodate more quantitative frameworks, like penalty logic [20], when the weights are thought of as violation fines to be paid (see also [43,25]). In the penalty logic framework, one could also distinguish between two bases: one representing constraints whose valuation has a cost, and one representing desires whose satisfaction produces a reward.

Another quantitative framework is evidence theory [38,40]. Belief functions have already been used for representing negative preferences. Here we could use the so-called "commonality" function Q for representing positive preferences. Indeed, a commonality function generalizes the guaranteed possibility measure Δ. Particularly when focal elements are nested, then the commonality function simply coincides with a guaranteed possibility measure.

Our bipolar representation is also of interest when representing knowledge rather than preferences, as discussed in [23], where the negative parts correspond to what is known as being (more or less impossible), while the positive parts gather solutions which are guaranteed to be feasible because they have been observed or reported or are supported by empirical evidence. This is why integrity constraints which are pieces of knowledge of the first kind, can be considered here, when necessary, on top of negative preferences, leading to the specification of what is tolerated because it respects integrity constraints or reflects the requirements of a user.

Appendix A. A brief refresher on possibilistic logic

The basic concept in possibilistic theory is the notion of possibility distribution, denoted by π, which is simply a function from the set of interpretations to the unit interval $[0,1]$ viewed as an ordinal scale, or to any bounded totally ordered scale L, finite or not. Thus we can use only a finite, totally ordered set of qualitative levels if necessary. In the usual acception, the top 1 of L stands for *fully possible*, *plausible*, *normal*, *preferred* according to the context. The bottom value 0 means *fully impossible*, *abnormal*, *rejected* likewise. Function π may encodes an agent's epistemic state, or a qualitative counterpart of a utility function, according to the context.

Given a possibility distribution π, two standard measures are defined for formulas:

- the possibility (or consistency) measure of a formula φ:

  $\Pi(\varphi) = \max\{\pi(s)\colon s \in S \text{ and } s \models \varphi\};$

  it evaluates the extent to which φ is consistent with the available information expressed by π, and
- the necessity (or certainty) measure of a formula φ:

  $N(\varphi) = 1 - \Pi(\neg\varphi),$

  which evaluates the extent to which φ is entailed by the information expressed by π.

A third set function Δ can be defined:

$$\Delta(\varphi) = \min\{\pi(s)\colon s \in S \text{ and } s \models \varphi\};$$

it evaluates the extent to which all models of φ are possible according to π.

The two first set-functions satisfy the following properties: $\Pi(w_1 \vee w_2) = \max(\Pi(w_1), \Pi(w_2))$, $N(w_1 \wedge w_2) = \min(N(w_1), N(w_2))$, and are monotonic increasing with entailment. On the contrary, $\Delta(w_1 \vee w_2) = \min(\Delta(w_1), \Delta(w_2))$, so Δ is decreasing with respect to entailment. Given a possibility distribution π, we define the core of π as the set of interpretations having the highest possibility degree in π. Formally,

Definition 5. The core of a possibility distribution π, denoted by *core*(π), is defined by:

$$\text{core}(\pi) = \{s\text{: } s \in S, \nexists s' \in S, \pi(s') > \pi(s)\}.$$

We now define the contextual core as follows:

Definition 6. The contextual core of a possibility distribution π given a formula φ (φ represents the context), denoted by $\text{core}_\varphi(\pi)$, is defined by:

$$\text{core}_\varphi(\pi) = \{s\text{: } s \in S, s \models \varphi, \nexists s' \in S, s' \models \varphi \text{ and } \pi(s') > \pi(s)\}.$$

Two kinds of semantic inference relations can be defined from π, in the same spirit as [39]:

Definition 7. Plausible and preferential inferences.

Let π be a possibility distribution. The formula ψ is said to be a plausible consequence of π, denoted by $\pi \models_P \psi$, iff

$$\text{core}(\pi) \subseteq [\![\psi]\!].$$

ψ is said to be a preferential consequence of π given the formula φ, denoted by $\pi \models_\varphi \psi$, iff

$$\text{core}_\varphi(\pi) \subseteq [\![\psi]\!].$$

The syntactic representation of prioritized information in the possibilistic logic framework is given by means of a set of weighted formulas, called *a possibilistic logic base*, of the form $\Sigma = \{(\varphi_i, a_i) : i = 1, \ldots, n\}$, where φ_i is a propositional formula and a_i, called the certainty degree of φ_i and belongs to a totally ordered scale such as $[0, 1]$. (φ_i, a_i) means that the necessity degree of φ_i is at least equal to a_i, i.e., $N(\varphi_i) \geqslant a_i$.

The fuzzy set of models of a possibilistic logic base is a possibility distribution induced by the weighted formulas:

$$\pi_\Sigma(s) = \min\{1 - a_i : s \models \neg\varphi_i, (\varphi_i, a_i) \in \Sigma\}$$

with $\min\{\emptyset\} = 1$.

Definition 8. Let Σ be a possibilistic knowledge base, and $a \in [0, 1]$. We call the a-cut (resp. strict a-cut) of Σ, denoted by $\Sigma_{\geqslant a}$ (resp. $\Sigma_{>a}$), the set of classical formulas in Σ having a certainty degree at least equal to a (resp. strictly greater than a).

$\text{Inc}(\Sigma) = \max\{a_i : \Sigma_{\geqslant a_i}$ *is inconsistent*$\}$ denotes the inconsistency degree of Σ. When Σ is consistent, we have $\text{Inc}(\Sigma) = 0$, which is equivalent to $\pi_\Sigma(s) = 1$ for some s. It is also equivalent to the consistency of the classical base associated with Σ when the weights are ignored. The knowledge base $\{(\varphi_1, b), (\varphi_2, b)\}$ is semantically equivalent to the formula $[\varphi_1 \wedge \varphi_2, b]$. $[\varphi, a] \in \Sigma$ is said to be subsumed if there exists $[\varphi', b] \in \Sigma$ such that $a \geqslant b$ and $\varphi' \vdash \varphi$. Adding or removing a subsumed formula preserves semantic equivalence i.e. does not change the induced possibility distribution. Full details on possibilistic logic are in [24].

The set-function Δ gives birth to a different form of possibilistic logic where the semantic entailment goes opposite to the classical logic one [23]. Namely, if all the models of w_1 are possible, we can conclude from this piece of information that all the models of w_2 are possible only if the entailment $w_2 \models w_1$ holds, and nothing is said about the interpretations violating w_1. This dual possibilistic logic handles weighted formulas $\Sigma = \{[\varphi_i, a_i] : i = 1, \ldots, n\}$, but here $[\varphi_i, a_i]$ is interpreted as $\Delta(\varphi_i) \geqslant a_i$, and the corresponding possibility distribution is

$$\delta_\Sigma(s) = \max\{a_i : s \models \varphi_i, [\varphi_i, a_i] \in \Sigma\}.$$

The knowledge base $\{[\varphi_1, b], [\varphi_2, b]\}$ is semantically equivalent to the formula $[\varphi_1 \vee \varphi_2, b]$. $[\varphi, a] \in \Sigma$ is said to be subsumed if there exists $[\varphi', b] \in \Sigma$ such that $b \geqslant a$ and $\varphi \vdash \varphi'$. Adding or removing a subsumed formula preserves semantic equivalence i.e. does not change the induced possibility distribution. The logical machinery of the Δ-weighted formulas is governed by a cut rule of the form [23]:

$$\Delta(w_1 \wedge w_2) \geqslant a,$$
$$\Delta(\neg w_1 \wedge w_3) \geqslant b \vdash \Delta(w_2 \wedge w_3) \geqslant \min(a, b),$$

which is the counterpart to possibilistic resolution [24] changing a necessity measure into a guaranteed possibility measure and the disjunctions into conjunctions.

A complete presentation of the syntactic inference machinery underlying the logic of Δ-weighted formulas can be found in [9].

References

[1] F. Bacchus, A.J. Grove, Utility independence in a qualitative decision theory, in: Proceedings of the 5th International Conference on Principles of Knowledge Representation, (KR'96), 1996, pp. 542–552.

[2] R.E. Bellman, L.A. Zadeh, Decision-making in a fuzzy environment, Management Sciences 17 (4) (1970) 141–164.

[3] S. Benferhat, D. Dubois, C. Cayrol, J. Lang, H. Prade, Inconsistency management and prioritized syntax-based entailment, in: 13th International Joint Conference on Artificial Intelligence IJCAI'93, 1993, pp. 640–645.

[4] S. Benferhat, D. Dubois, S. Kaci, H. Prade, Bridging logical, comparative and graphical possibilistic representation frameworks, in: 6th European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU'01), Lecture Notes in Computer Science LNCS 2143, 2001, pp. 422–431.

[5] S. Benferhat, D. Dubois, S. Kaci, H. Prade, Bipolar representation and fusion of preferences in the possibilistic logic framework, in: 8th International Conference on Principle of Knowledge Representation and Reasoning (KR'02), 2002, pp. 421–432.

[6] S. Benferhat, D. Dubois, H. Prade, Towards a possibilistic logic handling of preferences, Applied Intelligence 14 (3) (2001) 303–317.

[7] D. Benferhat, H. Dubois, M. Prade, S. Williams, A practical approach to fusing and revising prioritized belief bases, in: Proceedings of EPIA 99, LNAI No 1695, Springer-Verlag, 1999, pp. 222–236.

[8] S. Benferhat, S. Kaci, A possibilistic logic handling of strong preferences, in: International Fuzzy Systems Association (IFSA'01), Vancouver, IEEE Operation Center, 2001, pp. 962–967.

[9] S. Benferhat, S. Kaci, A logical representation and fusion of prioritized information based on guaranteed possibility measures: Application to the encoding of distance-based merging for classical bases, Artificial Intelligence 148 (2003) 291–333.

[10] S. Benferhat, S. Kaci, D. Leberre, M. Williams. Weakening conflicting information for iterated revision and knowledge integration, in: Proceedings of Seventeenth International Joint Conference on Artificial Intelligence (IJCAI'01), 2001, pp. 109–115.

[11] S. Bistarelli, Ph. Codognet, F. Rossi, Abstracting soft constraints: Framework, properties, examples, Artificial Intelligence 139 (2) (2002) 175–211.

[12] S. Bistarelli, H. Fargier, U. Montanari, F. Rossi, T. Schiex, G. Verfaillie, Semiring-based csps and valued csps: Frameworks, properties, and comparison, Constraints: An International Journal 4 (1999) 199–240.

[13] J.C. Borod, The Neuropsychology of Emotion, Oxford University Press, 2000.

[14] C. Boutilier, Toward a logic for qualitative decision theory, in: Proceedings of the 4th International Conference on Principles of Knowledge Representation, (KR'94), 1994, pp. 75–86.

[15] C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, D. Poole, CP-nets: A tool for representing and reasoning with conditional *ceteris paribus* preference statements, Journal of Artificial Intelligence Research 21 (2004) 135–191.

[16] C. Boutilier, R.I. Brafman, H. Hoos, D. Poole, Reasoning with conditional *ceteris paribus* preference statements, in: Proceedings of 15th Conference on Uncertainty in Artificial Intelligence (UAI'99), 1999, pp. 71–80.

[17] R.I. Brafman, C. Domshlak. Introducing variables importance tradeoffs into CP-Nets, in: Proceedings of 18th Conference on Uncertainty in Artificial Intelligence (UAI'02), 2002, pp. 69–76.

[18] J.T. Cacioppo, G.G. Bernston, The affect system: Architecture and operating characteristics, Current Directions in Psychological Science 8 (5) (1999) 133–137.

[19] J.T. Cacioppo, W.L. Gardner, G.G. Bernston, Beyond bipolar conceptualizations and measures: The case of attitudes and evaluative space, Personality and Social Psychology Review 1 (1) (1997) 3–25.

[20] F. Dupin de Saint-Cyr, J. Lang, T. Schiex, Penalty logic and its link with Dempster-Shafer theory, in: Proceedings of 10th International Conference on Uncertainty in Artificial Intelligence (UAI'94), 1994, pp. 204–211.

[21] J. Doyle, Y. Shoham, M. Wellman, A logic of relative desire, in: Proceedings of the Sixth International Symposium on Methodologies for Intelligent Systems, 1991, pp. 16–31.

[22] D. Dubois, H. Fargier, H. Prade, Possibility theory in constraint satisfaction problems: Handling priority, preference and uncertain, Applied Intelligence 6 (1996) 287–309.

[23] D. Dubois, P. Hajek, H. Prade, Knowledge-driven versus data-driven logics, Journal of Logic, Language, and Information 9 (2000) 65–89.

[24] D. Dubois, J. Lang, H. Prade, Possibilistic logic, in: D. Gabbay, et al. (Eds.), Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 3, Oxford University Press, 1994, pp. 439–513.

[25] D. Dubois, H. Prade, Toll sets and toll logic, in: R. Lowen, M. Roubens (Eds.), Fuzzy Logic, State of the Art, Kluwer Academic Publishers, 1993, pp. 169–177.

[26] D. Dubois, H. Prade, Possibility theory and data fusion in poorly informed environments, Control Engineering Practice 2 (5) (1994) 811–823.

[27] D. Dubois, H. Prade, Possibility theory: Qualitative and quantitative aspects, in: D. Gabbay, Ph. Smets (Eds.), Handbook of Defeasible Reasoning and Uncertainty Management Systems, in: Ph. Smets (Ed.), Quantified Representations of Uncertainty and Imprecision, vol. 1, Kluwer, Dordrecht, 1998, pp. 169–226.

[28] D. Dubois, H. Prade, Bipolarity in flexible querying, in: Proceedings of the 5th International Conference Flexible Query Answering Systems, FAQS 2002, LNAI 2522, 2002, pp. 174–182.

[29] M. Grabisch, C. Labreuche, The Sipos integral for the aggregation of interacting bipolar criteria, in: 8th International Conference on Information Processing and Management of Uncertainty in Knowledge-based Systems (IPMU'00), Madrid, 2000, pp. 395–409.

[30] S. Kaci, H. Prade, Bipolar goals. A possibilistic logic characterization of preferred choices, in: Proceedings of Workshop on Local Computation for Logics and Uncertainty, in conjunction with ECAI'04, Valence, August 2004.

[31] C. Lafage, Représentation de préférences en logique: application à la décision de groupe, Thèse de doctorat. Université Paul Sabatier, Toulouse, 2001.

[32] J. Lang, Conditional desires and utilities—and alternative logical approach to qualitative decision theory, in: Proceedings of 12th European Conference on Artificial Intelligence (ECAI'96), 1996, pp. 318–322.

[33] J. Lang, L. Van Der Torre, E. Weydert, Utilitarian desires, Autonomous Agents and Multi-Agent Systems 5 (2002) 329–363.

[34] J.T. Larsen, P. McGraw, J.T. Cacioppo, Can people feel happy and sad at the same time? Journal of Personality and Social Psychology 81 (2001) 684–696.

[35] R. Neves, E. Raufaste, A psychological study of bipolarity in the possibilistic framework, in: 10th International Conference on Information Processing and Management of Uncertainty in Knowledge-based Systems (IPMU'04), Perugia, 2004, pp. 975–980.

[36] E.T. Rolls, Precis of brain and emotion, Behavioral and Brain Sciences 23 (2) (2000) 177–234.

[37] T. Schiex, Possibilistic constraints satisfaction problems or How to handle soft constraints, in: Proceedings of 15th Conference on Uncertainty in Artificial Intelligence (UAI'92), 1992, pp. 268–275.

[38] G. Shafer, A Mathematical Theory of Evidence, Princeton University Press, New Jersey, 1976.

[39] Y. Shoham, Reasoning about Change, MIT Press, MA, 1988.

[40] P. Smets, R. Kennes, The transferable belief model, Artificial Intelligence 66 (1994) 191–234.

[41] S. Tan, J. Pearl, Qualitative decision theory, in: Proceedings of the National Conference on Artificial Intelligence (AAAI'94), 1994, pp. 928–933.

[42] A. Tversky, D. Kahneman, Advances in prospect theory: Cumulative representation of uncertainty, Journal of Risk and Uncertainty 5 (1992) 297–323.
[43] L. van der Torre, E. Weydert, Parameters for utilitarian desires in a qualitative decision theory, Applied Intelligence 14 (2001) 285–301.
[44] M.P. Wellman, J. Doyle, Y. Shoham, Preferential semantics for goals, in: Proceedings of the 9th Conference on Artificial Intelligence (AAAI'91), 1991, pp. 698–703.
[45] G.H. Von Wright, The Logic of Preference, Edinburg University Press, 1963.