Merging Optimistic and Pessimistic Preferences

Souhila Kaci
CRIL
Rue de l'Université SP 16
62307 Lens Cedex, France
kaci@cril.univ-artois.fr

Leendert van der Torre
ILIAS
University of Luxembourg
Luxembourg
leon.vandertorre@uni.lu

Abstract - ***In this paper we consider the extension of non-monotonic preference logic with the distinction between controllable (or endogenous) and uncontrollable (or exogenous) variables, which can be used for example in agent decision making and deliberation. We assume that the agent is optimistic about its own controllables and pessimistic about its uncontrollables, and we study ways to merge these two distinct dimensions. We also consider complex preferences, such as optimistic preferences conditional on an uncontrollable, or optimistic preferences conditional on a pessimistic preference.***

Keywords: Preference logic, preference merging, non-monotonic reasoning.

Introduction

In many areas such as cooperative information systems, multi-databases, multi-agents systems, information comes from multiple sources. The multiplicity of sources providing information makes that information is often contradictory which requires conflict resolution. This problem has been widely studied in literature where implicit priorities, based on Dalal's distance, (Lin 1996; Lin & Mendelzon 1998; Konieczny & Pérez 1998; Revesz 1993; 1997) or explicit priorities (Benferhat *et al*. 1999; 2002) are used in order to solve conflicts.

Our concern in this paper is the merging of preferences of a single agent when they are expressed in a logic of preferences. Logics of preferences attract much attention in knowledge representation and reasoning, where they are used for a variety of applications such as qualitative decision making (Doyle & Thomason 1999). In this paper we oppose to the common wisdom that the very efficient specificity algorithms used in some non-monotonic preference logics are too simple to be used for knowledge representation and reasoning applications. In that logics we distinguish minimal and maximal specificity principles which correspond to a gravitation towards the ideal and the worst respectively. We counter the argument that a user is forced to chose among minimal and maximal specificity by introducing the fundamental distinction between controllable and uncontrollable variables from decision and control theory, and merging preferences on the two kinds of variables as visualized in Figure 1. Our work is based on the hypothesis that each set of preferences on controllable and uncontrollable variables is consistent. The merging process aims to cohabit controllable and uncontrollable variables in an intuitive way. Preferences on controllable variables are called optimistic preferences since minimal specificity principle is used for such variables. This principle is a gravitation towards the ideal and thus corresponds to an optimistic reasoning. Preferences on uncontrollable variables are called pessimistic preferences since maximal specificity principle is used for such variables. This principle is a gravitation towards the worst and thus corresponds to an pessimistic reasoning.

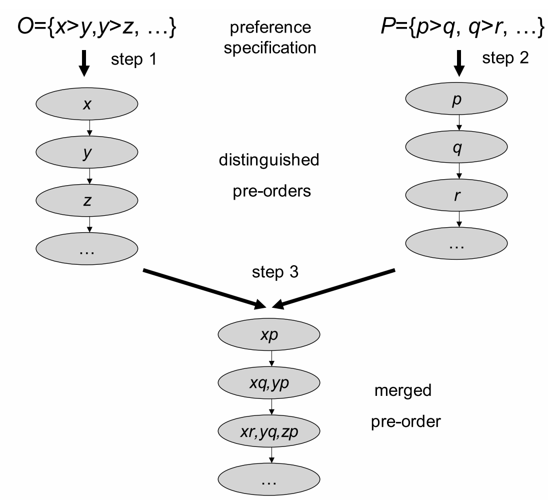

Figure 1: Merging optimistic and pessimistic preferences.

A preference specification contains optimistic preferences (O) defined on controllables $x, y, z, \ldots$, and pessimistic preferences (P) defined on uncontrollables $q, r, t, \ldots$, which are interpreted as constraints on total pre-orders on worlds. The efficient specificity algorithms (step 1 and 2 in Figure 1) calculate unique distinguished total pre-orders, which are thereafter merged (step 3) by symmetric or a-symmetric mergers. If the optimistic and pessimistic preferences in Figure 1 are defined on separate languages, then for step 1 and 2 we can use existing methods in preference logic, such as (Kaci & van der Torre 2005a). In this paper we also consider more general languages, in which preferences on controllables are conditional on uncontrollables, or on preferences on uncontrollables (or vice versa).

The remainder of this paper is organized as follows. After a necessary background, we present a logic of optimistic preferences defined on controllable variables and a logic of pessimistic preferences defined on uncontrollable variables. Then we propose some merging approaches of optimistic and pessimistic preferences. We also introduce a logic of preferences where pessimistic and optimistic preferences are merged in the logic itself. Lastly we conclude with future research.

Background

Let W be the set of propositional interpretations of $\mathcal{L}$, and let $\succeq$ be a total pre-order on W (called also a preference order), i.e., a reflexive, transitive and connected ($\forall \omega, \omega' \in W$ we have either $\omega \succeq \omega'$ or $\omega' \succeq \omega$) relation. We write $w \succ w'$ for $w \succeq w'$ without $w' \succeq w$. Moreover, we write $\max(x, \succeq)$ for $\{w \in W \mid w \models x, \forall w' \in W : w' \models x \Rightarrow w \succeq w'\}$, and analogously we write $\min(x, \succeq)$ for $\{w \in W \mid w \models x, \forall w' \in W : w' \models x \Rightarrow w' \succeq w\}$.

The following definition illustrates how a preference order can also be represented by a well ordered partition of W. This is an equivalent representation, in the sense that each preference order corresponds to one ordered partition and vice versa. This equivalent representation as an ordered partition makes the definition of the non-monotonic semantics, defined later in the paper, easier to read.

Definition 1 (Ordered partition) *A sequence of sets of worlds of the form $(E_1, \ldots, E_n)$ is an ordered partition of W iff*

- $\forall i$, E_i *is nonempty,*
- $E_1 \cup \cdots \cup E_n = W$ *and*
- $\forall i, j$, $E_i \cap E_j = \emptyset$ *for* $i \neq j$.

An ordered partition of W is associated with pre-order $\succeq$ on W iff $\forall \omega, \omega' \in W$ with $\omega \in E_i, \omega' \in E_j$ we have $i \leq j$ iff $\omega \succeq \omega'$.

Preferences for controllables

Reasoning about controllables is optimistic in the sense that an agent or decision maker can decide the truth value of a controllable proposition, and thus may expect that the best state will be realized.

Optimistic reasoning semantics

A preference statement is a comparative statement "x is preferred to y", with x and y propositional sentences of a propositional language on a set of controllable propositional atoms. A reasoning about a preference can be optimistic or pessimistic with respect to both its left hand side and right hand side, indicated by o and p respectively. Formally we write $x \, {}^{a}{>}^{b} y$, where $a, b \in \{o, p\}$. An optimistic reasoning focuses on the best worlds while a pessimistic reasoning focuses on the worst worlds. For example, the preference $x \, {}^{p}{>}^{o} y$ indicates that we are drawing a pessimistic reasoning with respect to x, and an optimistic reasoning with respect to y. This means that we deal with the worst x-worlds i.e. $\min(x, \succeq)$ and the best y-worlds i.e. $\max(y, \succeq)$.
An optimistic reasoning on a preference statement over controllable variables consists of an optimistic reasoning w.r.t. its right and left hand side. This also includes the case where the reasoning is pessimistic w.r.t. its left hand side and optimistic w.r.t. its right hand side. This will be explained later in this subsection. For the sake of simplicity, such a preference is called *optimistic*. Indeed we define an *optimistic* preference specification as a set of strict and non-strict optimistic preferences:

Definition 2 (Optimistic preference specification) *Let $\mathcal{L}_{\mathcal{C}}$ be a propositional language on a set of controllable propositional atoms $\mathcal{C}$. Let $\mathcal{O}_{\rhd}$ be a set of optimistic preferences of the form $\{x_i \rhd y_i \mid i = 1, \cdots, n, x_i, y_i \in \mathcal{L}_{\mathcal{C}}\}$. A preference specification is a tuple $\langle \mathcal{O}_{\rhd} \mid \rhd \in \{ {}^{p}{>}^{o}, {}^{p}{\geq}^{o} \; {}^{o}{>}^{o}, {}^{o}{\geq}^{o} \} \rangle$.*

We define preferences of x over y as preferences of $x \wedge \neg y$ over $y \wedge \neg x$. This is standard and known as von Wright's expansion principle (Wright 1963). Additional clauses may be added for the cases in which sets of worlds are nonempty, to prevent the satisfiability of preferences like $x > \top$ and $x > \bot$. To keep the formal exposition to a minimum, we do not consider this borderline condition in this paper.

Definition 3 (Monotonic semantics) *Let $\succeq$ be a total preorder on W.*

$\succeq \models x \, {}^{o}{>}^{o} y$ *iff* $\forall w \in \max(x \wedge \neg y, \succeq)$ *and* $\forall w' \in \max(\neg x \wedge y, \succeq)$ *we have* $w \succ w'$

$\succeq \models x \, {}^{o}{\geq}^{o} y$ *iff* $\forall w \in \max(x \wedge \neg y, \succeq)$ *and* $\forall w' \in \max(\neg x \wedge y, \succeq)$ *we have* $w \succeq w'$

$\succeq \models x \, {}^{p}{>}^{o} y$ *iff* $\forall w \in \min(x \wedge \neg y, \succeq)$ *and* $\forall w' \in \max(\neg x \wedge y, \succeq)$ *we have* $w \succ w'$

$\succeq \models x \, {}^{p}{\geq}^{o} y$ *iff* $\forall w \in \min(x \wedge \neg y, \succeq)$ *and* $\forall w' \in \max(\neg x \wedge y, \succeq)$ *we have* $w \succeq w'$.

A total pre-order $\succeq$ is a model of an optimistic preference specification $\mathcal{O}_{\rhd}$ if it is a model of each $p_i \rhd q_i \in \mathcal{O}_{\rhd}$.

Note that $x \, {}^{p}{>}^{o} y$ means that each x-world is preferred to all y-worlds w.r.t. $\succeq$. This preference can be equivalently written as a set of optimistic preferences of the form $\{x' \, {}^{o}{>}^{o} y : x' \text{ is a } x - \text{world}\}$. This is also true for $x \, {}^{p}{\geq}^{o} y$ preferences.

Example 1 *Consider an agent organizing his evening by deciding whether he goes to the cinema (c), with his friend (f) and whether he alos goes to the restaurant (r). We have $\mathcal{O} = \langle \mathcal{O}_{{}^{o}{>}^{o}}, \mathcal{O}_{{}^{p}{>}^{o}}, \mathcal{O}_{{}^{p}{\geq}^{o}} \rangle$, where $\mathcal{O}_{{}^{o}{>}^{o}} = \{c \wedge f \, {}^{o}{>}^{o} \neg(c \wedge f)\}$, $\mathcal{O}_{{}^{p}{>}^{o}} = \{c \wedge r \, {}^{p}{>}^{o} c \wedge \neg r\}$, $\mathcal{O}_{{}^{p}{\geq}^{o}} = \{c \wedge r \, {}^{p}{\geq}^{o} \neg c \wedge r\}$. The strict preference $c \wedge f \, {}^{o}{>}^{o} \neg(c \wedge f)$ means that there is at least a situation in which the agent goes to the cinema with his friend which is strictly preferred to all situations where the agent does not go to the cinema with his friend. The strict preference $c \wedge r \, {}^{p}{>}^{o} c \wedge \neg r$ means that each situation in which the agent goes to the cinema and the restaurant is strictly preferred to all situations in which the agent goes to the cinema but not to the restaurant. Finally the non-strict preference $c \wedge r \, {}^{p}{\geq}^{o} \neg c \wedge r$ means that* *each situation in which the agent goes to the cinema and the restaurant is at least as preferred as all situations in which the agent goes to the restaurant but not to the cinema.*

We compare total pre-orders based on the so-called specificity principle. Optimistic reasoning is based on the minimal specificity principle, which assumes that worlds are as good as possible.

Definition 4 (Minimal specificity principle) *Let $\succeq$ and $\succeq'$ be two total pre-orders on a set of worlds W represented by ordered partitions $(E_1, \cdots, E_n)$ and $(E'_1, \cdots, E'_m)$ respectively. We say that $\succeq$ is at least as specific as $\succeq'$, written as $\succeq \sqsubseteq \succeq'$, iff $\forall \omega \in W$, if $\omega \in E_i$ and $\omega \in E'_j$ then $i \leq j$. $\succeq$ belongs to the set of the least specific pre-orders among a set of pre-orders $\mathcal{O}$ if there is no $\succeq'$ in $\mathcal{O}$ s.t. $\succeq' \sqsubset \succeq$, i.e., $\succeq' \sqsubseteq \succeq$ holds but $\succeq \sqsubseteq \succeq'$ does not.*

Algorithm 1 gives the (unique) least specific pre-order satisfying an optimistic preference specification. All the proofs can be found in (Kaci & van der Torre 2006).

Following Definition 2 an optimistic preference specification contains the following sets of preferences:

$\mathcal{O}_{{}^o\!>^o} = \{C_{i_1} : x_{i_1} \; {}^o\!>^o y_{i_1}\}$,

$\mathcal{O}_{{}^o\!\geq^o} = \{C_{i_2} : x_{i_2} \; {}^o\!\geq^o y_{i_2}\}$,

$\mathcal{O}_{{}^p\!>^o} = \{C_{i_3} : x_{i_3} \; {}^p\!>^o y_{i_3}\}$,

$\mathcal{O}_{{}^p\!\geq^o} = \{C_{i_4} : x_{i_4} \; {}^p\!\geq^o y_{i_4}\}$.

Moreover, we refer to the constraints of these preferences by

$$\overline{\mathcal{C}} = \bigcup_{k=1,\cdots,4} \{\overline{\mathcal{C}}_{i_k} = (L(C_{i_k}), R(C_{i_k}))\},$$

where the left and right hand side of these constraints are $L(C_{i_k}) = |x_{i_k} \wedge \neg y_{i_k}|$ and $R(C_{i_k}) = |\neg x_{i_k} \wedge y_{i_k}|$ respectively; $|\phi|$ denotes the set of interpretations satisfying ϕ.

The basic idea of the algorithm is to construct the least specific pre-order by calculating the sets of worlds of the ordered partition, going from the ideal to the worst worlds.

At each step of the algorithm, we look for worlds which can have the current highest ranking in the preference order. This corresponds to the current minimal value l. These worlds are those which do not falsify any constraint in $\overline{\mathcal{C}}$. We first put in E_l worlds which do not falsify any strict preference. These worlds are those which do not appear in the right hand side of the strict preferences $\overline{\mathcal{C}}_{i_1}$ and $\overline{\mathcal{C}}_{i_3}$. Now we remove from E_l worlds which falsify constraints of the non-strict preferences $\overline{\mathcal{C}}_{i_2}$ and $\overline{\mathcal{C}}_{i_4}$. Constraints $\overline{\mathcal{C}}_{i_2}$ are violated if $L(C_{i_2}) \cap E_l = \emptyset$ and $R(C_{i_2}) \cap E_l \neq \emptyset$, while the constraints $\overline{\mathcal{C}}_{i_4}$ are violated if $L(C_{i_4}) \not\subseteq E_l$ and $R(C_{i_4}) \cap E_l \neq \emptyset$. Once E_l is fixed, satisfied constraints are removed. Note that constraints $\overline{\mathcal{C}}_{i_k}$ s.t. $k \in \{1, 2\}$ are satisfied if $L(C_{i_k}) \cap E_l \neq \emptyset$ since in this case, worlds of $R(C_{i_1})$ are necessarily in E_h with $h > l$ and worlds of $R(C_{i_2})$ are in $E_{h'}$ with $h' \geq l$. However constraints $\overline{\mathcal{C}}_{i_k}$ with $k \in \{3, 4\}$ are satisfied only when $L(C_{i_k}) \subseteq E_l$ otherwise they should be replaced by $(L(C_{i_k}) - E_l, R(C_{i_k}))$.

Algorithm 1: Handling optimistic preferences.

Data: An optimistic preference specification.
Result: A total preorder $\succeq$ on W.
begin
  $l \leftarrow 0$;
  while $W \neq \emptyset$ **do**
    – $l \leftarrow l + 1$, $j \leftarrow 1$;
    /** strict constraints **/
    – $E_l = \{\omega : \forall \overline{\mathcal{C}}_{i_1}, \overline{\mathcal{C}}_{i_3} \in \overline{\mathcal{C}}, \omega \notin R(C_{i_1}) \cup R(C_{i_3})\}$;
    while $j = 1$ **do**
      $j \leftarrow 0$;
      for *each* $\overline{\mathcal{C}}_{i_2}$ *and* $\overline{\mathcal{C}}_{i_4}$ *in* $\overline{\mathcal{C}}$ **do**
        /** constraints induced by non-strict preferences **/
        if *$(L(C_{i_2}) \cap E_l = \emptyset$ and $R(C_{i_2}) \cap E_l \neq \emptyset)$ or $(L(C_{i_4}) \not\subseteq E_l$ and $R(C_{i_4}) \cap E_l \neq \emptyset)$* **then**
          $E_l = E_l - R(C_{i_k})$;
          $j \leftarrow 1$
    if $E_l = \emptyset$ **then** Stop (inconsistent constraints);
    – from W remove elements of E_l ;
    /** remove satisfied constraints induced by ${}^o\!>^o$ preferences **/
    – from $\overline{\mathcal{C}}$ remove $\overline{\mathcal{C}}_{i_k}$ $k \in \{1, 2\}$ such that $L(C_{i_k}) \cap E_l \neq \emptyset$;
    /** update constraints induced by ${}^p\!>^o$ constraints **/
    – replace constraints $\overline{\mathcal{C}}_{i_k}$ $(k \in \{3, 4\})$ by $(L(C_{i_k}) - E_l, R(C_{i_k}))$;
    /** remove satisfied constraints induced by ${}^p\!>^o$ preferences **/
    – from $\overline{\mathcal{C}}$ remove $\overline{\mathcal{C}}_{i_k}$ $(k \in \{3, 4\})$ with empty $L(C_{i_k})$.
  return $(E_1, \cdots, E_l)$
end

Example 2 *Let us consider again the optimistic preference specification given in Example 1.*
Let $W = \{\omega_0 : \neg c \neg f \neg r, \omega_1 : \neg c \neg f r, \omega_2 : \neg c f \neg r, \omega_3 : \neg c f r, \omega_4 : c \neg f \neg r, \omega_5 : c \neg f r, \omega_6 : c f \neg r, \omega_7 : c f r\}$.
We have $\overline{\mathcal{C}} = \{(\{\omega_6, \omega_7\}, \{\omega_0, \omega_1, \omega_2, \omega_3, \omega_4, \omega_5\})\} \cup \{(\{\omega_5, \omega_7\}, \{\omega_4, \omega_6\})\} \cup \{(\{\omega_5, \omega_7\}, \{\omega_1, \omega_3\})\}$.
We put in E_1 all worlds which do not appear in the right hand side of strict constraints, we get $E_1 = \{\omega_7\}$. The constraint induced by $c \wedge r \; {}^p\!\geq^o \neg c \wedge r$ is not violated. The constraint induced by $c \wedge f \; {}^o\!>^o \neg(c \wedge f)$ is satisfied while the ones induced by $c \wedge r \; {}^p\!>^o c \wedge \neg r$ and $c \wedge r \; {}^p\!\geq^o \neg c \wedge r$ are not. So $\overline{\mathcal{C}} = \{(\{\omega_5\}, \{\omega_4, \omega_6\})\} \cup \{(\{\omega_5\}, \{\omega_1, \omega_3\})\}$.
We repeat this process and get $E_2 = \{\omega_0, \omega_1, \omega_2, \omega_3, \omega_5\}$ and $E_3 = \{\omega_4, \omega_6\}$.

Preferences for uncontrollables

Reasoning about uncontrollables is pessimistic in the sense that an agent cannot decide the truth value of a uncontrollable proposition, and thus may assume that the worst state will be realized (known as Wald's criterion).

Pessimistic reasoning semantics

A pessimistic preference specification contains four sets of preferences, which are pessimistic on their left and right hand side. This also includes the case where preferences are pessimistic with respect to their left hand side and optimistic with respect to their right side (as in optimistic reasoning semantics). This will be explained later in this section.

Definition 5 (Pessimistic preference specification) *Let $\mathcal{L}_{\mathcal{U}}$ be a propositional language on a set of uncontrollable propositional atoms $\mathcal{U}$. Let $\mathcal{P}_{\rhd}$ be a set of pessimistic preferences of the form $\{q_i \rhd r_i \mid i = 1, \cdots, n, q_i, r_i \in \mathcal{L}_{\mathcal{U}}\}$. A preference specification is a tuple $\langle \mathcal{P}_{\rhd} \mid \rhd \in \{ {}^{p}{>}^{o}, {}^{p}{\geq}^{o}, {}^{p}{>}^{p}, {}^{p}{\geq}^{p} \} \rangle$.*

Definition 6 (Monotonic semantics) *Let $\succeq$ be a total preorder on W.*

$\succeq \models q\ {}^{p}{>}^{p}\ r$ *iff* $\forall w \in \min(q \wedge \neg r, \succeq)$ *and* $\forall w' \in \min(\neg q \wedge r, \succeq)$ *we have* $w \succ w'$

$\succeq \models q\ {}^{p}{\geq}^{p}\ r$ *iff* $\forall w \in \min(q \wedge \neg r, \succeq)$ *and* $\forall w' \in \min(\neg q \wedge r, \succeq)$ *we have* $w \succeq w'$

$\succeq \models q\ {}^{p}{>}^{o}\ r$ *iff* $\forall w \in \min(q \wedge \neg r, \succeq)$ *and* $\forall w' \in \max(\neg q \wedge r, \succeq)$ *we have* $w \succ w'$

$\succeq \models q\ {}^{p}{\geq}^{o}\ r$ *iff* $\forall w \in \min(q \wedge \neg r, \succeq)$ *and* $\forall w' \in \max(\neg q \wedge r, \succeq)$ *we have* $w \succeq w'$

A total pre-order $\succeq$ is a model of $\mathcal{P}_{\rhd}$ iff $\succeq$ satisfies each preference $q_i \rhd r_i$ in $\mathcal{P}_{\rhd}$.

Note that $q\ {}^{p}{>}^{o}\ r$ can be equivalently written as $\{q\ {}^{p}{>}^{p}\ r' : r'$ is a $r-$world$\}$. This is also true for $q\ {}^{p}{\geq}^{o}\ r$ preferences.

Pessimistic reasoning is based on the maximal specificity principle, which assumes that worlds are as bad as possible.

Definition 7 (Maximal specificity principle) *$\succeq$ belongs to the set of the most specific pre-orders among a set of preorders $\mathcal{O}$ if there is no $\succeq'$ in $\mathcal{O}$ such that $\succeq \sqsubset \succeq'$.*

Algorithm 2 gives the (unique) most specific preorder satisfying a pessimistic preference specification. It is similar to Algorithm 1.

This algorithm is based on the following four sets of preferences:

$$\mathcal{P}_{{}^{p}{>}^{p}} = \{C_{i_1} : q_{i_1}\ {}^{p}{>}^{p}\ r_{i_1}\},$$

$$\mathcal{P}_{{}^{p}{\geq}^{p}} = \{C_{i_2} : q_{i_2}\ {}^{p}{\geq}^{p}\ r_{i_2}\},$$

$$\mathcal{P}_{{}^{p}{>}^{o}} = \{C_{i_3} : q_{i_3}\ {}^{p}{>}^{o}\ r_{i_3}\},$$

$$\mathcal{P}_{{}^{p}{\geq}^{o}} = \{C_{i_4} : q_{i_4}\ {}^{p}{\geq}^{o}\ r_{i_4}\}.$$

Let $\overline{\mathcal{C}} = \bigcup_{k=1,\cdots,4}\{\overline{\mathcal{C}}_{i_k} = (L(C_{i_k}), R(C_{i_k}))\}$, where $L(C_{i_k}) = |q_{i_k} \wedge \neg r_{i_k}|$ and $R(C_{i_k}) = |\neg q_{i_k} \wedge r_{i_k}|$.

Algorithm 2: Handling pessimistic preferences.

Data: A pessimistic preference specification.
Result: A total pre-order $\succeq$ on W.

```
begin
  l ← 0;
  while W ≠ ∅ do
    l ← l + 1, j ← 1;
    E_l = {ω : ∀C̄_{i1}, C̄_{i3} in C̄, ω ∉ L(C_{i1}) ∪ L(C_{i3})};
    while j = 1 do
      j ← 0;
      for each C̄_{i2} and C̄_{i4} in C̄ do
        /** constraints induced by non-strict preferences **/
        if (L(C_{i2}) ∩ E_l ≠ ∅ and R(C_{i2}) ∩ E_l = ∅) or
           (L(C_{i4}) ∩ E_l ≠ ∅ and R(C_{i4}) ⊈ E_l) then
          E_l = E_l − L(C_{ik}), j ← 1
    if E_l = ∅ then Stop (inconsistent constraints);
    – From W remove elements of E_l;
    /** remove satisfied constraints induced by p>p preferences **/
    – From C̄ remove C̄_{ik} (for k ∈ {1,2}) s.t. E_l ∩ R(C_{ik}) ≠ ∅;
    /** update constraints induced by p>o preferences **/
    – Replace C̄_{ik} (for k ∈ {3,4}) in C̄ by (L(C_{ik}), R(C_{ik}) − E_l);
    /** remove satisfied constraints induced by p>o preferences **/
    – From C̄ remove C̄_{ik} (k ∈ {3,4}) with empty R(C_{ik});
  return (E'_1, ···, E'_l) s.t. ∀1 ≤ h ≤ l, E'_h = E_{l−h+1}
end
```

Merging optimistic and pessimistic preferences

In this section we consider the merger of the least specific pre-order satisfying the optimistic preference specification, and the most specific pre-order satisfying the pessimistic preference specification. From now on, let $\mathcal{L}$ be a propositional language on disjoint sets of controllable and uncontrollable propositional atoms $\mathcal{C} \cup \mathcal{U}$. A preference specification $\mathcal{PS}$ consists of an optimistic and a pessimistic preference specification, i.e., optimistic preferences on controllables and pessimistic preferences on uncontrollables. In general, let $\succeq$ be the merger of $\succeq_o$ and $\succeq_p$. We assume that Pareto conditions hold:

Definition 8 *Let $\succeq_o$, $\succeq_p$ and $\succeq$ be three total pre-orders on the same set. $\succeq$ is a merger of $\succeq_o$ and $\succeq_p$ if and only if the following three conditions hold:*
If $w_1 \succ_o w_2$ and $w_1 \succ_p w_2$ then $w_1 \succ w_2$,
If $w_1 \succeq_o w_2$ and $w_1 \succeq_p w_2$ then $w_1 \succeq w_2$.

Given two arbitrary pre-orders, there are many possible mergers. We therefore again consider distinguished preorders in the subsections below. The desideratum of a merger operator is that the merger satisfies, in some sense, most of the preference specification. However, it is clearly unreasonable to ask for an operator that satisfies the whole preference specification. For example, we may have strong preferences $x\ {}^{p}{>}^{o}\ \neg x$ and $p\ {}^{p}{>}^{o}\ \neg p$, which can be satisfied by a minimal and maximal specific pre-order separately, but which are contradictory given together. This motivates the next definition of partial satisfaction, which only considers some of the preference types.

Definition 9 *A pre-order partially satisfies a preference specification $\mathcal{PS}$ when it satisfies $\mathcal{PS}_{\rhd}$ with $\rhd \in \{ {}^{o}{>}^{o}, {}^{o}{\geq}^{o}, {}^{p}{>}^{p}, {}^{p}{\geq}^{p} \}$.*

The merging operators in this section satisfy our desideratum that the merger partially satisfies the preference specification, as a consequence of the following lemma. The two minimal and maximal specific pre-orders of optimistic and pessimistic preference specifications satisfy the property that no two sets are disjoint.

Lemma 1 *Let $(E_1, \cdots, E_n)$ and $(E'_1, \cdots, E'_m)$ be the ordered partitions of $\succeq_o$ and $\succeq_p$ respectively. We have for all $1 \leq i \leq n$ and all $1 \leq j \leq m$ that $E_i \cap E'_j \neq \emptyset$.*
Proof. Due to the fact that $\succeq_o$ and $\succeq_p$ are defined on disjoint sets of variables.

Symmetric mergers

Let $\succeq$ be the merger of $\succeq_o$ and $\succeq_p$. The least and most specific pre-orders $\succeq$ satisfying Pareto conditions, are *unique* and *identical*, and can be obtained as follows. Given Lemma 1, thus far nonempty sets E''_k do not exist, but they may exist in extensions discussed in future sections.

Proposition 1 *Let $(E_1, \cdots, E_n)$ and $(E'_1, \cdots, E'_m)$ be the ordered partitions of $\succeq_o$ and $\succeq_p$ respectively. The least/most specific merger of $\succeq_o$ and $\succeq_p$ is $\succeq = (E''_1, \cdots, E''_{n+m-1})$ such that if $\omega \in E_i$ and $\omega \in E'_j$ then $\omega \in E''_{i+j-1}$, and by eliminating nonempty sets E''_k and renumbering the nonempty ones in sequence.*

The symmetric merger, called also the least/most specific merger, is illustrated by the following example.

Example 3 *Consider the optimistic preference specification $p \; {}^{o}{>}^{o} \neg p$ and the pessimistic preference specification $m \; {}^{p}{>}^{p} \neg m$, where p and m stand respectively for "I will work on a* project *in order to get money" and "my boss accepts to give me* money *to pay the conference fee".*
Applying Algorithm 1 and Algorithm 2 on $p \; {}^{o}{>}^{o} \neg p$ and $m \; {}^{p}{>}^{p} \neg m$ respectively gives $\succeq_o = (\{mp, \neg mp\}, \{m\neg p, \neg m\neg p\})$ and $\succeq_p = (\{mp, m\neg p\}, \{\neg mp, \neg m\neg p\})$. The least/most specific merger is $\succeq = \{\{mp\}, \{\neg mp, m\neg p\}, \{\neg m\neg p\}\}$.

Proposition 2 *The least/most specific merger of two pre-orders satisfying Lemma 1 partially satisfies the preference specification.*

Proposition 3 *The least/most specific merger is not complete, in the sense that there are pre-orders which cannot be constructed in this way.*

Proof. Consider a language with only one controllable x and one uncontrollable p. The minimal and maximal specific pre-orders consist of at most two equivalence classes, and the least/most specific merger consists therefore of at most three equivalence classes. Hence, pre-orders in which all four worlds are distinct cannot be constructed.

We can also consider the product merger, which is a symmetric merger, defined by: if $\omega \in E_i$ and $\omega \in E'_j$ then $\omega \in E''_{i*j}$.

Dictators

We now consider dictator mergers that prefer one ordering over the other one. The *minimax merger* gives priority to the preorder $\succeq_o$ associated to the optimistic preference specification, computed following the minimal specificity principle, over $\succeq_p$ associated to the pessimistic preference specification, computed following the maximal specificity principle. Dictatorship relation of $\succ_o$ over $\succ_p$ means that worlds are first ordered with respect to $\succeq_o$ and only in the case of equality $\succeq_p$ is considered.

Definition 10 $w_1 \succ w_2$ *iff* $w_1 \succ_o w_2$ *or (*$w_1 \sim_o w_2$ *and* $w_1 \succ_p w_2$*).*

The minimax merger can be defined as follows.

Proposition 4 *Let $(E_1, \cdots, E_n)$ and $(E'_1, \cdots, E'_m)$ be the ordered partitions of $\succeq_o$ and $\succeq_p$ respectively. The result of merging $\succeq_o$ and $\succeq_p$ is $\succeq = (E''_1, \cdots, E''_{n*m})$ such that if $\omega \in E_i$ and $\omega \in E'_j$ then $\omega \in E''_{(i-1)*m+j}$.*

Example 4 *(continued) The minimax merger of the preference specification is $\{\{mp\}, \{\neg mp\}, \{m\neg p\}, \{\neg m\neg p\}\}$.*

The principle of the *maximin merger* is similar to minimax merger. The dictator here is the pre-order associated to the pessimistic preference specification and computed following the maximal specificity principle.

Definition 11 $w_1 \succ w_2$ *iff* $w_1 \succ_p w_2$ *or (*$w_1 \sim_p w_2$ *and* $w_1 \succ_o w_2$*).*

Example 5 *(continued) The maximin merger of the preference specification is $\{\{mp\}, \{m\neg p\}, \{\neg mp\}, \{\neg m\neg p\}\}$.*

Conditional preferences

The drawback of handling preferences on controllable and uncontrollable variables separately is the impossibility to express *interaction* between the two kinds of variables. For example my decision on whether I will work hard to finish a paper (which is a controllable variable) depends on the uncontrollable variable "money", decided by my boss. If my boss accepts to pay the conference fees then I will work hard to finish the paper. We therefore consider in the remainder of this paper preference formulas with both controllable and uncontrollable variables.

A general approach would be to define optimistic and pessimistic preference specifications on any combination of controllables and uncontrollables, such as an optimistic preference $p \; {}^{o}{>}^{o} x$ or even $q \; {}^{o}{>}^{o} r$. However, this approach blurs the idea that optimistic reasoning is restricted to controllables, and pessimistic reasoning is restricted to uncontrollables. We therefore define conditional preferences. Conditional optimistic and pessimistic preferences are defined as follows.

Definition 12 (Conditional optimistic preference specification) *Let $\mathcal{O}_{\rhd}$ be a set of conditional optimistic preferences of the form $\{q_i \rightarrow (x_i \rhd y_i) \mid i = 1, \cdots, n, q_i \in \mathcal{L}_{\mathcal{U}}, x_i, y_i \in \mathcal{L}_{\mathcal{C}}\}$,* *where* $q \rightarrow (x \rhd y) = (q \wedge x) \rhd (q \wedge y)$. *A conditional optimistic preference specification is a tuple* $\langle \mathcal{O}_\rhd \mid \rhd \in \{ {}^{p}{>}^{o}, {}^{p}{\geq}^{o} {}^{o}{>}^{o}, {}^{o}{\geq}^{o}\}\rangle$.

Definition 13 (Conditional pessimistic preference specification) *Let* $\mathcal{P}_\rhd$ *be a set of conditional pessimistic preferences of the form* $\{x_i \rightarrow (q_i \rhd r_i) \mid i = 1, \cdots, n, x_i \in \mathcal{L}_\mathcal{C}, q_i, r_i \in \mathcal{L}_\mathcal{U}\}$, *where* $x \rightarrow (q \rhd r) = (x \wedge q) \rhd (x \wedge r)$. *A conditional pessimistic preference specification is a tuple* $\langle \mathcal{P}_\rhd \mid \rhd \in \{ {}^{p}{>}^{o}, {}^{p}{\geq}^{o}, {}^{p}{>}^{p}, {}^{p}{\geq}^{p}\}\rangle$.

In the following examples we merge the two pre-orders using the symmetric merger operator since there is no reason to give priority neither to $\succeq_o$ nor to $\succeq_p$. We start with some simple examples to illustrate that the results of the merger behaves intuitively.

Example 6 *The merger of optimistic preference* $m \rightarrow (p\ {}^{o}{>}^{o} \neg p)$ *and pessimistic preference* $\neg m\ {}^{p}{>}^{p} m$ *is the merger of* $\succeq_o = (\{mp, \neg mp, \neg m\neg p\}, \{m\neg p\})$ *and* $\succeq_p = (\{\neg mp, \neg m\neg p\}, \{mp, m\neg p\})$, *i.e.,* $\succeq = (\{\neg m\neg p, \neg mp\}, \{mp\}, \{m\neg p\})$.

The merger of optimistic preference $m \rightarrow (p\ {}^{o}{>}^{o} \neg p)$ *and pessimistic preference* $m\ {}^{p}{>}^{p} \neg m$ *is the merger of* $\succeq_o = (\{mp, \neg mp, \neg m\neg p\}, \{m\neg p\})$ *and* $\succeq_p = (\{mp, m\neg p\}, \{\neg mp, \neg m\neg p\})$, *i.e.,* $\succeq = (\{mp\}, \{\neg mp, m\neg p, \neg m\neg p\})$.

The merger of optimistic preference $m \rightarrow (p\ {}^{o}{>}^{o} \neg p)$ *and pessimistic preference* $p \rightarrow (m\ {}^{p}{>}^{p} \neg m)$ *is the merger of* $\succeq_o = (\{mp, \neg mp, \neg m\neg p\}, \{m\neg p\})$ *and* $\succeq_p = (\{mp\}, \{\neg mp, m\neg p, \neg m\neg p\})$, *i.e.,* $\succeq = (\{mp\}, \{\neg mp, \neg m\neg p\}, \{m\neg p\})$.

Proposition 5 *The most specific merger of two minimal and maximal pre-orders of conditional preference specifications does* not *necessarily partially satisfy the preference specification.*

Proof. *The merger of optimistic preference* $m \rightarrow (p\ {}^{o}{>}^{o} \neg p)$ *and pessimistic preference* $\neg p \rightarrow (m\ {}^{p}{>}^{p} \neg m)$ *is the merger of* $\succeq_o = (\{mp, \neg mp, \neg m\neg p\}, \{m\neg p\})$ *and* $\succeq_p = (\{m\neg p\}, \{mp, \neg mp, \neg m\neg p\})$, *i.e.,* $\succeq = (\{mp, m\neg p, \neg m\neg p, \neg mp\})$. *The merger is the universal relation which does not satisfy any non-trivial preference.*

We now consider an extension of our running example on working and money.

Example 7 *Let's consider another controllable variable* w *which stands for "I will* work *hard on the paper". Let*
$\mathcal{O} = \{money \rightarrow (work\ {}^{o}{>}^{o} \neg work),$
$\neg money \rightarrow (\neg work\ {}^{o}{>}^{o} work),$
$\neg money \rightarrow (project\ {}^{p}{>}^{o} \neg project)\}$.
This is equivalent to
$\{money \wedge work\ {}^{o}{>}^{o} money \wedge \neg work,$
$\neg money \wedge \neg work\ {}^{o}{>}^{o} \neg money \wedge work,$
$\neg money \wedge project\ {}^{p}{>}^{o} \neg money \wedge \neg project\}$.
Applying Algorithm 1 gives
$\succeq_o = (\{\neg m\neg wp, mwp, mw\neg p\}, \{m\neg w\neg p, m\neg wp, \neg mwp\},$
$\{\neg m\neg w\neg p, \neg mw\neg p\})$.

All preferences are true in $\succeq_o$. *According to these preferences, the best situations for the agent are when there is money and she works hard on the paper, or when there is no money, she works on a project but does not work hard on the paper. This is intuitively meaningful since when there is money the agent is motivated to work hard on the paper however when there is no money, it becomes necessary to work on a project which prevents her to work hard on the paper. The worst situations (as one would expect) are when there is no money and she does not work on a project.*

Example 8 *Let*
$\mathcal{P} = \{\neg project \rightarrow (money\ {}^{p}{>}^{o} \neg money),$
$\neg work \rightarrow (\neg money\ {}^{p}{>}^{p} money)\}$.
This is equivalent to
$\{\neg project \wedge money\ {}^{p}{>}^{o} \neg project \wedge \neg money,$
$\neg work \wedge \neg money\ {}^{p}{>}^{p} \neg work \wedge money\}$.
Applying Algorithm 2 gives
$\succeq_p = (\{mw\neg p, m\neg w\neg p\}, \{\neg m\neg w\neg p, \neg m\neg wp\},$
$\{\neg mw\neg p, \neg mwp, m\neg wp, mwp\})$.

Now given a preference specification $\mathcal{PS} = \mathcal{O} \cup \mathcal{P}$, the associated total pre-order is the result of combining $\succeq_o$ and $\succeq_p$ using the symmetric merger.

Example 9 *The merger of* $\succeq_o$ *and* $\succeq_p$ *given in Examples 7 and 8 respectively is* $\succeq = (\{mw\neg p\}, \{\neg m\neg wp, m\neg w\neg p\}, \{mwp\}, \{m\neg wp, \neg mwp, \neg m\neg w\neg p\}, \{\neg mw\neg p\})$. *The best situation is when there is money, the agent works hard on the paper and does not work on a project and the worst situation is when the agent works hard on the paper but unfortunately neither she works on a project nor there is money.*

The following example illustrates how our approach can be used in qualitative decision making. The distinction between controllable and uncontrollable variables exists in many qualitative decision theories, see e.g. (Boutilier 1994), and most recently preference logic for decision has been promoted in particular by Brewka (Brewka 2004). We use Savage's famous egg breaking example (Savage 1954), as also used by Brewka (Brewka 2004) to illustrate his extended logic programming approach in decision making.

Example 10 *An agent is preparing an omelette. 5 fresh eggs are already in the omelette. There is one more egg. She does not know whether this egg is fresh or rotten. The agent can (i) add it to the omelette which means the whole omelette may be wasted, (ii) throw it away, which means one egg may be wasted, or (iii) put it in a cup, check whether it is ok or not and put it to the omelette in the former case, throw it away in the latter. In any case, a cup has to be washed if this option is chosen.*
There is one controllable variable which consists in putting the egg $in\_omelette$, in_cup *or* $throw\ away$. *There is also an uncontrollable variable which is the state of the egg* $fresh$ *or* $rotten$. *The effects of controllable and uncontrollable variables are the following:*

$5\_omelette \leftarrow throw\_away,$
$6\_omelette \leftarrow fresh, in\_omelette$
$0\_omelette \leftarrow rotten, in\_omelette,$
$6\_omelette \leftarrow fresh, in\_cup,$
$5\_omelette \leftarrow rotten, in\_cup,$

$\neg wash \leftarrow$ not in_cup,
$wash \leftarrow in\_cup$.

Agent's desires are represented as follows:

$\neg wash \times wash$
$6\_omelette \times 5\_omelette \times 0\_omelette$.

We used here notations of logic programming (Brewka 2004). For example $5\_omelette \leftarrow throw\_away$ is interpreted as: if the egg is thrown away then the agent will get an omelette with 5 eggs. The desire $6\_omelette \times 5\_omelette \times 0\_omelette$ is interpreted as: preferably $6\_omelette$, if not then $5\_omelette$ and if neither $6\_omelette$ nor $5\_omelette$ then $0\_omelette$.

Possible solutions are:

$S_1 = \{6\_omelette, \neg wash, fresh, in\_omelette\}$,
$S_2 = \{0\_omelette, \neg wash, rotten, in\_omelette\}$,
$S_3 = \{6\_omelette, wash, fresh, in\_cup\}$,
$S_4 = \{5\_omelette, wash, rotten, in\_cup\}$,
$S_5 = \{5\_omelette, \neg wash, fresh, throw\_away\}$,
$S_6 = \{5\_omelette, \neg wash, rotten, throw\_away\}$.

Each solution is composed of an instantiation of decision variables and the satisfied desires.

Let us run this example following Brewka's approach (Brewka 2004).

Example 10 *(Continued) Brewka generates a preference order on the solutions (called answer sets in his framework) following agent's desires. Indeed S_1 is the single preferred solution. S_5 and S_6 are equally preferred. They are preferred to S_2 and S_4 but incomparable to S_3. S_3 is preferred to S_4 and incomparable to S_5, S_6 and S_2. Lastly S_2 and S_4 are incomparable.*

In our approach, controllable and uncontrollable variables are dealt with separately, respecting their distinct nature in decision theory. Our approach uses also various kinds of preferences, and non-monotonic reasoning (based on specificity algorithms) to deal with under-specification.

Example 10 *(Continued) Let us consider the following preferences on controllable and uncontrollable variables:*

$$\mathcal{O} = \begin{cases} fresh \rightarrow in\_omelette > in\_cup \\ fresh \rightarrow in\_cup > throw\_away \\ rotten \rightarrow throw\_away > in\_cup \\ rotten \rightarrow in\_cup > in\_omelette \end{cases}$$

$$\mathcal{P} = \begin{cases} in\_omelette \rightarrow fresh > rotten \\ in\_cup \rightarrow fresh > rotten \\ throw\_away \rightarrow rotten > fresh \end{cases}$$

The set of possible alternatives is $W = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6\}$ where
$\omega_1 = fresh \wedge in\_omelette$,
$\omega_2 = rotten \wedge in\_omelette$,
$\omega_3 = fresh \wedge in\_cup$,
$\omega_4 = rotten \wedge in\_cup$,
$\omega_5 = fresh \wedge throw\_away$ and
$\omega_6 = rotten \wedge throw\_away$.
We apply Algorithm 1 on the set $\mathcal{O}$ of optimistic preferences, we get $(\{\omega_1, \omega_6\}, \{\omega_3, \omega_4\}, \{\omega_2, \omega_5\})$.
We apply Algorithm 2 on the set $\mathcal{P}$ of pessimistic preferences, we get $(\{\omega_1, \omega_3, \omega_6\}, \{\omega_2, \omega_4, \omega_5\})$.
We merge the two preorders using the symmetric merger, we get $(\{\omega_1, \omega_6\}, \{\omega_3\}, \{\omega_4\}, \{\omega_2, \omega_5\})$.
Now agent's desires may be used to discriminate ω_1 and ω_6. Both satisfy $\neg wash$ however ω_1 satisfies $6\_omelette$ while ω_6 satisfies $5\_omelette$ so ω_1 is preferred to ω_6.
Concerning ω_2 and ω_5, ω_5 is preferred to ω_2. Indeed solutions of the previous example are ordered as follows in our framework: $S_1 \succ S_6 \succ S_3 \succ S_4 \succ S_5 \succ S_2$.

Our approach may be viewed as an extension of Brewka's approach where preferences among alternatives are used in addition to preferences among desires.

Concluding remarks

The distinction between controllable and uncontrollable propositions is fundamental in decision and control theory, and in various agent theories. Moreover, various kinds of optimistic and pessimistic reasoning are also present in many decision theories, for example in the maximin and minimax decision rules. However, their role seems to have attracted less attention in the non-monotonic logic of preference (Boella & van der Torre 2005; Dastani *et al.* 2005; Kaci & van der Torre 2005a; Lang 2004), despite the recent interest in this area, and the recent recognition that preference logic plays a key role in many knowledge representation and reasoning tasks, including decision making.

In this paper we study non-monotonic preference logic extended with the distinction between controllable and uncontrollable propositions. We illustrate how the logic can be used in decision making where preferences on controllables and preferences on uncontrollables have to be merged.

Our approach may also be used in more complex merging tasks such as social and group decision making. For example, one such extension are preferences on controllable variables conditional on preferences on uncontrollable variables, i.e. $(q \rhd_p r) \rightarrow (x \rhd_o y)$, or conversely, i.e. $(x \rhd_o y) \rightarrow (q \rhd_p r)$. This extension can be used for social decision making where an agent states its preferences given the preferences of another agent.

The following example illustrates how such social preferences can be used. Roughly, for a conditional optimistic preference $(q \rhd_p r) \rightarrow (x \rhd_o y)$, we first apply the pessimistic ordering on uncontrollables and then use the result to incorporate preferences on controllables, combining the two using the maximin merger.

Example 11 *Carl and his girlfriend Sandra go the restaurant. Menus are composed of meat or fish, wine or jus and dessert or cheese. Sandra is careful about her fitness so each menu without cake is preferred for her to all menus with cake. Even if Carl likes dessert, he does want to attempt Sandra by choosing a menu composed of a cake so, to compensate, he states that there is at least one menu composed of wine and cheese which is preferred to all menus composed of neither cake nor wine. Let $W = \{\omega_0 : \neg d \neg w \neg m, \omega_1 : \neg d \neg wm, \omega_2 : \neg dw \neg m, \omega_3 : \neg dwm, \omega_4 : d \neg w \neg m, \omega_5 : d \neg wm, \omega_6 : d \neg w \neg m, \omega_0 : dwm\}$ be the set*

of possible menus where m, w *and* d *stand for* meat, wine *and* dessert *respectively.* $\neg m$, $\neg w$ *and* $\neg d$ *stand for* fish, jus *and* cheese *respectively.*
Sandra's preferences give the following preorder $\succeq = (\{\omega_0, \omega_1, \omega_2, \omega_3\}, \{\omega_4, \omega_5, \omega_6, \omega_7\})$ *and Carl's preferences give the following preorder* $\succeq' = (\{\omega_2, \omega_3, \omega_4, \omega_5, \omega_6, \omega_7\}, \{\omega_0, \omega_1\})$. *We use the maximin merger and get:* $(\{\omega_2, \omega_3\}, \{\omega_0, \omega_1\}, \{\omega_4, \omega_5, \omega_6, \omega_7\})$.

Given a set of preferences of the form $\{q_j \rhd_p r_j \to x_i \rhd_o y_i\}$, one may be tried to compute the preorders associated to $\{q_j \rhd_p r_j\}$ and $\{x_i \rhd_o y_i\}$ and then to merge them. However this way is misleading since each set of preferences may be inconsistent. The correct way would be to compute the preorder associated to each rule $q_j \rhd_p r_j \to x_i \rhd_o y_i$ as explained above and then to merge the different preorders using the symmetric merger since there is no reason to give priority to any preorder. The investigation of this idea is left to a further research.

Other topics for further research are preference specifications in which strong preferences ${}^p\!>^o$ are defined on both controllables and uncontrollables to define a stronger notion than weak satisfiability of a preference specification, the extension with beliefs, and ceteris paribus preferences (see (Kaci & van der Torre 2005b)).

References

Benferhat, S.; Dubois, D.; Prade, H.; and Williams, M. 1999. A practical approach to fusing and revising prioritized belief bases. In *Proceedings of EPIA 99, LNAI* n^o *1695, Springer Verlag*, 222–236.

Benferhat, S.; Dubois, D.; Kaci, S.; and Prade, H. 2002. Possibilistic merging and distance-based fusion of propositional information. In *Annals of Mathematics and Artificial Intelligence*, volume 34(1-3), 217–252.

Boella, G., and van der Torre, L. 2005. A nonmonotonic logic for specifying and querying preferences. In *Proceedings of IJCAI'05*.

Boutilier, C. 1994. Toward a logic for qualitative decision theory. In *Proceedings KR94*, 75–86.

Brewka, G. 2004. Answer sets and qualitative decision making. In *Synthese*.

Dastani, M.; Governatori, G.; Rotolo, A.; and van der Torre, L. 2005. Preferences of agents in defeasible logic. In *Proceedings AI'05*. Springer.

Doyle, J., and Thomason, R. 1999. Background to qualitative decision theory. *AI Magazine* 20(2):55–68.

Kaci, S., and van der Torre, L. 2005a. Algorithms for a nonmonotonic logic of preferences. In *Eighth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'05)*, 281–292.

Kaci, S., and van der Torre, L. 2005b. Non-monotonic reasoning with various kinds of preferences. In *IJCAI'05 Multidisciplinary Workshop on Advances in Preference Handling*.

Kaci, S., and van der Torre, L. 2006. Merging Optimistic and Pessimistic Preferences. In *C.R.I.L., Technical report*.

Konieczny, S., and Pérez, R. P. 1998. On the logic of merging. In *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98), Trento*, 488–498.

Lang, J. 2004. A preference-based interpretation of other agents' actions. In *Proceedings of KR04*, 644–653.

Lin, J., and Mendelzon, A. 1998. Merging databases under constraints. *International Journal of Cooperative Information Systems* 7(1):55–76.

Lin, J. 1996. Integration of weighted knowledge bases. *Artificial Intelligence* 83:363–378.

Revesz, P. Z. 1993. On the semantics of theory change: arbitration between old and new information. In 12^{th} *ACM SIGACT-SIGMOD-SIGART symposium on Principles of Databases*, 71–92.

Revesz, P. Z. 1997. On the semantics of arbitration. *International Journal of Algebra and Computation* 7(2):133–160.

Savage, L. 1954. The foundations of Statistics. *Dover, New York*.

Wright, G. V. 1963. *The Logic of Preference*. Edinburg. University Press.