

A Postulate-Based Analysis of Comparative Preference Statements

Souhila Kaci

LIRMM

161 rue ADA

F34392 Montpellier Cedex 5 France

Abstract

Most of preference representation languages developed in the literature are based on comparative preference statements. The latter offer a simple and intuitive way for expressing preferences. They can be interpreted following different semantics. This paper presents a postulate-based analysis of the different semantics describing their behavior w.r.t. three criteria: coherence, syntax independence and inference.

Introduction

Preferences are the backbone of various fields as they naturally arise and play an important role in many real-life decisions. Preferences are fundamental in scientific research frameworks as well as applications.

One of the main problems an individual faces when expressing her preferences lies in the number of variables (or attributes or criteria) that she takes into account to evaluate the different outcomes. Indeed, the number of outcomes increases exponentially with the number of variables. Moreover, due to their cognitive limitation, individuals are generally not willing to compare all possible pairs of outcomes or evaluate them individually. These facts have an unfortunate consequence that any preference representation language that is based on the direct assessment of individual preferences over the complete set of outcomes is simply infeasible.

Fortunately, individuals can abstract their preferences. More specifically, instead of providing preferences over outcomes (by pairwise comparison or individual evaluation), they generally express preferences over partial descriptions of outcomes. Often such statements take the form of qualitative comparative preference statements e.g., “I like London more than Paris” and “prefer tea to coffee”. Compact preference representation languages aim at representing such partial descriptions of individual preferences which we refer to as comparative preference statements. They use different completion principles in order to compute a preference relation induced by a set of preference statements.

Comparative preference statements offer an intuitive and natural way to express preferences. Most of the preferences

we express seem to be of this type. Individuals may also wish to consider some factors to express their comparative preference statements, e.g., “If fish is served, then I prefer white wine to red wine”, allowing then to express general preferences (e.g., “I prefer fish to meat”) and specific preferences in particular contexts (e.g., “If red wine is served, I prefer meat to fish”).

An important point we need to fix when handling comparative preference statements occurs when we have to deal with statements which refer to sets of outcomes. For example suppose that one has to choose a menu composed of a main dish (*fish* or *meat*), a wine (*white* or *red*) and a dessert (*cake* or *ice_cream*). If an individual expresses that she prefers *fish* to *meat* then she has to compare between four fish-based menus (*fish - white - cake*, *fish - white - ice_cream*, *fish - red - cake*, *fish - red - ice_cream*) and four meat-based menus (*meat - white - cake*, *meat - white - ice_cream*, *meat - red - cake*, *meat - red - ice_cream*). Different ways are possible to perform such a comparison. They lead to different preference semantics. Mainly, these semantics have their foundation in philosophy and non-monotonic reasoning. So far the main objective in artificial intelligence has been to rank-order the set of outcomes given a set of comparative preference statements and one or several semantics. In this paper we come to this problem from a different angle. We consider a set of postulates studied in preference logics and non-monotonic reasoning. These postulates formalize intuition one may have regarding the behavior of preference statements. We analyze the behavior of the different semantics w.r.t. these postulates.

After necessary background, we discuss the different semantics proposed in the literature. Then we provide a postulates-based analysis of these semantics. Lastly, we conclude.

Background

Let $V = \{X_1, \dots, X_h\}$ be a set of h variables, each takes its values in a domain $Dom(X_i)$. A possible outcome, denoted by ω , is the result of assigning a value in $Dom(X_i)$ to each variable X_i in V . Ω is the set of all possible outcomes. We suppose that this set is fixed and finite. We also suppose that there is no integrity constraint that restricts the set of possible outcomes. Therefore we suppose that all possible outcomes are feasible.

Let \mathcal{L} be a language based on V . $Mod(\alpha)$ denotes the set of outcomes that make the formula α (built on \mathcal{L}) true. It is also called α -outcomes.

A preference relation \succeq on $\mathcal{X} = \{x, y, z, \dots\}$ is a reflexive and transitive binary relation such that $x \succeq y$ stands for x is at least as preferred as y . $x \approx y$ means that both $x \succeq y$ and $y \succeq x$ hold, i.e., x and y are equally preferred. The notation $x \succ y$ means that x is strictly preferred to y . We have $x \succ y$ if $x \succeq y$ holds but $y \succeq x$ does not. \succeq is cyclic iff $\exists x, y \in \mathcal{X}$ such that both $x \succ y$ and $y \succ x$ hold. Otherwise it is acyclic. Given a preference relation \succeq and a formula α , the set of the best (resp. worst) α -outcomes is denoted by $\max(\alpha, \succeq)$ (resp. $\min(\alpha, \succeq)$) and defined as $\{\omega \mid \omega \in Mod(\alpha), \nexists \omega' \in Mod(\alpha), \omega' \succ \omega\}$ (resp. $\{\omega \mid \omega \in Mod(\alpha), \nexists \omega' \in Mod(\alpha), \omega \succ \omega'\}$).

Comparative preference statements

Individuals express their preferences in different forms. However, often these preferences implicitly or explicitly refer to qualitative comparative preference statements of the form “prefer α to β ”. Handling such a preference statement is easy when both α and β refer to an outcome, e.g. “prefer *fish – white – cake* to *meat – red – ice_cream*” for choosing a menu composed of a main dish (*fish* or *meat*), wine (*white wine* or *red wine*) and dessert (*cake* or *ice cream*). However this task becomes more complex when α and β refer to sets of outcomes, in particular when they share some outcomes. For example the preference statement “prefer *fish* to *red wine*” in the previous example means that we compare the set of menus composed of *fish* and the set of menus composed of *red wine*. The first set is $\Sigma_1 = \{\textit{fish-red-cake}, \textit{fish-red-ice_cream}, \textit{fish-white-cake}, \textit{fish-white-ice_cream}\}$ and the second set is $\Sigma_2 = \{\textit{fish-red-cake}, \textit{meat-red-cake}, \textit{fish-red-ice_cream}, \textit{meat-red-ice_cream}\}$. Therefore the preference statement “prefer *fish* to *red wine*” means that the menus in Σ_1 are preferred to the menus in Σ_2 . However *fish – red – cake* and *fish – red – ice_cream* belong to both Σ_1 and Σ_2 ! In order to prevent this situation Halldén (1957) and von Wright (1963) interpret the statement “prefer α to β ” as a choice problem between $\alpha \wedge \neg\beta$ -outcomes and $\neg\alpha \wedge \beta$ -outcomes. Therefore the statement “prefer *fish* to *red wine*” leads to choose between menus composed of *fish* and *white wine* and menus composed of *meat* and *red wine*. This turns to compare the sets $\Sigma'_1 = \{\textit{fish-white-cake}, \textit{fish-white-ice_cream}\}$ and $\Sigma'_2 = \{\textit{meat-red-cake}, \textit{meat-red-ice_cream}\}$. Particular situations are those when $\alpha \wedge \neg\beta$ (resp. $\neg\alpha \wedge \beta$) is a contradiction or is not feasible in which case it is replaced with α (resp. β). We refer the reader to (von Wright 1963; Hansson 2001) for further details. For simplicity we suppose that both $\alpha \wedge \neg\beta$ and $\neg\alpha \wedge \beta$ are consistent and feasible.

Remark 1 *One may wonder whether “prefer fish to red wine” is a preference statement since it compares the value of two different variables, namely main dish (i.e., fish) and wine (i.e., red wine). This is in fact an importance statement. That is, it is more important for an individual to have a menu composed of fish and not red wine rather than a menu*

composed of red wine and not fish. Therefore menus composed of fish and white wine are preferred to menus composed of meat and red wine. The statement “prefer α to β ” is a preference statement when both α and β refer to the values of the same variable e.g. “prefer fish to meat”. Whatever the statement “prefer α to β ” refers to a preference or an importance, this turns to prefer $\alpha \wedge \neg\beta$ -outcomes over $\neg\alpha \wedge \beta$ -outcomes. For this reason, we do not make a distinction between a preference statement and an importance statement.

Let us now mention that the translation of “prefer α to β ” into a choice between $\alpha \wedge \neg\beta$ -outcomes and $\neg\alpha \wedge \beta$ -outcomes solves the problem of common outcomes; however it does not give an indication on how outcomes are compared. This problem calls for preference semantics.

Preference semantics

We denote by $\alpha \triangleright \beta$ a comparative preference statement “prefer α to β ”. A preference semantics refers to the way $\alpha \wedge \neg\beta$ -outcomes and $\neg\alpha \wedge \beta$ -outcomes are rank-ordered. Different ways have been studied for the comparison of two sets of objects leading to different preference semantics (Boutilier 1994; Wilson 2004; von Wright 1963; Hansson 2001; Boutilier et al. 2004; Pearl 1990; Benferhat et al. 2002; van der Torre and Weydert 2001; Barbera, Bossert, and Pattanaik 2004; Kaci 2011). In this paper, we focus on five semantics. Roughly, these semantics compare two sets of outcomes (namely $\alpha \wedge \neg\beta$ -outcomes and $\neg\alpha \wedge \beta$ -outcomes) on the basis of their best and worst outcomes w.r.t. a given preference relation.

Definition 1 (Preference semantics) *Let \succeq be a preference relation. Consider $\alpha \triangleright \beta$.*

- **Strong semantics** (Boutilier 1994; Wilson 2004) *\succeq satisfies $\alpha \triangleright \beta$, denoted by $\succeq \models_{\forall\forall} \alpha \triangleright \beta$, iff $\forall \omega \in \min(\alpha \wedge \neg\beta, \succeq), \forall \omega' \in \max(\neg\alpha \wedge \beta, \succeq), \omega \succ \omega'$.*
- **Ceteris paribus semantics** (von Wright 1963; Hansson 2001; Boutilier et al. 2004) *\succeq satisfies $\alpha \triangleright \beta$, denoted by $\succeq \models_{\forall\forall}^{cp} \alpha \triangleright \beta$, iff $\forall \omega \in \min(\alpha \wedge \neg\beta, \succeq), \forall \omega' \in \max(\neg\alpha \wedge \beta, \succeq), \omega \succ \omega'$ if the two outcomes have the same valuation over variables not appearing in $\alpha \wedge \neg\beta$ and $\neg\alpha \wedge \beta$.*
- **Optimistic semantics** (Pearl 1990; Boutilier 1994) *\succeq satisfies $\alpha \triangleright \beta$, denoted by $\succeq \models_{\exists\forall} \alpha \triangleright \beta$, iff $\forall \omega \in \max(\alpha \wedge \neg\beta, \succeq), \forall \omega' \in \max(\neg\alpha \wedge \beta, \succeq), \omega \succ \omega'$.*
- **Pessimistic semantics** (Benferhat et al. 2002) *\succeq satisfies $\alpha \triangleright \beta$, denoted by $\succeq \models_{\forall\exists} \alpha \triangleright \beta$, iff $\forall \omega \in \min(\neg\alpha \wedge \beta, \succeq), \forall \omega' \in \min(\alpha \wedge \neg\beta, \succeq), \omega \succ \omega'$.*
- **Opportunistic semantics** (van der Torre and Weydert 2001) *\succeq satisfies $\alpha \triangleright \beta$, denoted by $\succeq \models_{\exists\exists} \alpha \triangleright \beta$, iff $\forall \omega \in \max(\alpha \wedge \neg\beta, \succeq), \forall \omega' \in \min(\neg\alpha \wedge \beta, \succeq), \omega \succ \omega'$.*

The following proposition gives an equivalent reading of Definition 1. It allows a better understanding of the principles underpinning the semantics (Kaci and van der Torre 2008):

Proposition 1 *Let \succeq be a preference relation and $\alpha \triangleright \beta$ be a comparative preference statement.*

- $\succeq \models_{\forall\forall} \alpha \triangleright \beta$, iff $\forall \omega \in \text{Mod}(\alpha \wedge \neg\beta)$, $\forall \omega' \in \text{Mod}(\neg\alpha \wedge \beta)$, $\omega \succ \omega'$.
- $\succeq \models_{\forall\forall}^{\text{cp}} \alpha \triangleright \beta$, iff $\forall \omega \in \text{Mod}(\alpha \wedge \neg\beta)$, $\forall \omega' \in \text{Mod}(\neg\alpha \wedge \beta)$, $\omega \succ \omega'$ if the two outcomes have the same valuation over variables not appearing in $\alpha \wedge \neg\beta$ and $\neg\alpha \wedge \beta$.
- $\succeq \models_{\exists\forall} \alpha \triangleright \beta$, iff $\exists \omega \in \text{Mod}(\alpha \wedge \neg\beta)$, $\forall \omega' \in \text{Mod}(\neg\alpha \wedge \beta)$, $\omega \succ \omega'$.
- $\succeq \models_{\forall\exists} \alpha \triangleright \beta$, iff $\exists \omega' \in \text{Mod}(\neg\alpha \wedge \beta)$, $\forall \omega \in \text{Mod}(\alpha \wedge \neg\beta)$, $\omega \succ \omega'$.
- $\succeq \models_{\exists\exists} \alpha \triangleright \beta$, iff $\exists \omega \in \text{Mod}(\alpha \wedge \neg\beta)$, $\exists \omega' \in \text{Mod}(\neg\alpha \wedge \beta)$, $\omega \succ \omega'$.

The choice of the index of \models (i.e., $\forall\forall$, $\exists\forall$, $\forall\exists$, $\exists\exists$) refers to the selection of one or all $\alpha \wedge \neg\beta$ -outcomes and $\neg\alpha \wedge \beta$ -outcomes. When there is no ambiguity, we shall abuse notation and write \succeq satisfies $\alpha \triangleright_{xy} \beta$ (with $xy \in \{\exists, \forall\}$) to mean that $\succeq \models_{xy} \alpha \triangleright \beta$. We also use the symbol $\alpha \triangleright_{xy} \beta$ to say that $\alpha \triangleright \beta$ is interpreted following the corresponding semantics.

Proposition 1 reveals that the five semantics express more or less requirements on the way $\alpha \wedge \neg\beta$ -outcomes and $\neg\alpha \wedge \beta$ -outcomes are rank-ordered. As indicated by its name, strong semantics expresses the most requirements. It states that any $\alpha \wedge \neg\beta$ -outcome is preferred to any $\neg\alpha \wedge \beta$ -outcome. We can check that if $\succeq \models_{\forall\forall} \alpha \triangleright \beta$ then $\succeq \models_{\forall\forall}^{\text{cp}} \alpha \triangleright \beta$, $\succeq \models_{\exists\forall} \alpha \triangleright \beta$, $\succeq \models_{\forall\exists} \alpha \triangleright \beta$ and $\succeq \models_{\exists\exists} \alpha \triangleright \beta$.

Being too requiring, strong semantics has been criticized in the literature since it may lead to cyclic preferences when several preference statements are considered. For example there is no acyclic preference relation satisfying both $fish \triangleright_{\forall\forall} meat$ and $white \triangleright_{\forall\forall} red$ since $fish - red$ should be preferred to $meat - white$ given $fish \triangleright_{\forall\forall} meat$ and $meat - white$ should be preferred to $fish - red$ given $white \triangleright_{\forall\forall} red$. Ceteris paribus semantics has been considered as a good alternative. It weakens strong semantics by comparing less outcomes. For example the preference relation $fish - white \succ meat - white$ and $fish - red \succ meat - red$ satisfies $fish \triangleright_{\forall\forall}^{\text{cp}} meat$ but not $fish \triangleright_{\forall\forall} meat$. Also $fish - white \succ meat - white$ and $fish - red \succ meat - red$ and $fish - white \succ meat - red$ and $fish - red \succ meat - red$ satisfies both $fish \triangleright_{\forall\forall}^{\text{cp}} meat$ and $white \triangleright_{\forall\forall}^{\text{cp}} red$.

Optimistic semantics is a left-hand weakening of strong semantics. Instead of requiring that any $\alpha \wedge \neg\beta$ -outcome is preferred to any $\neg\alpha \wedge \beta$ -outcome, it states that at least one $\alpha \wedge \neg\beta$ -outcome should be preferred to any $\neg\alpha \wedge \beta$ -outcome. This reflects a flexibility regarding the outcome(s) which fulfill this requirement. The larger the set $\alpha \wedge \neg\beta$ -outcomes is, the more flexible is the preference statement $\alpha \triangleright \beta$. Flexibility should be understood as the number of possible preference relations satisfying $\alpha \triangleright \beta$. Pessimistic semantics is a right-hand weakening of strong semantics. It requires that at least one $\neg\alpha \wedge \beta$ -outcome should be less preferred to any $\alpha \wedge \neg\beta$ -outcome. Therefore optimistic and pessimistic semantics exhibit a dual behavior. The larger the set $\neg\alpha \wedge \beta$ -outcomes is, the more flexible is $\alpha \triangleright \beta$. Lastly, opportunistic semantics is both left- and right-hand weakening of strong semantics since it requires that at least one $\alpha \wedge \neg\beta$ -outcome should be preferred to at least one $\neg\alpha \wedge \beta$ -outcome.

Beyond semantics

Beyond the technical device of the five semantics regarding the selection of at least one or all $\alpha \wedge \neg\beta$ -outcomes and $\neg\alpha \wedge \beta$ -outcomes, some semantics can be highlighted for their expressive power. More specifically, although strong and ceteris paribus semantics are the most natural among the five semantics, they do not leave room for exceptions. Therefore they are not suitable to reason about defeasible preferences. Suppose that an individual would prefer *fish* to *meat* but if *red wine* is served then her preference is reversed. This means that we have $fish \triangleright meat$ and $red \wedge meat \triangleright red \wedge fish$. Both strong semantics and ceteris paribus semantics return contradictory preferences (i.e., cyclic). More precisely, $fish - red$ is preferred to $meat - red$ w.r.t. $fish \triangleright meat$ and $meat - red$ is preferred to $fish - red$ w.r.t. $red \wedge meat \triangleright red \wedge fish$. This is an undesirable situation because $fish \triangleright meat$ and $red \wedge meat \triangleright red \wedge fish$ are not contradictory. They simply state that an individual has a default preference for *fish* over *meat* but if *red wine* is served then she would prefer *meat*. So in any situation when a wine other than *red wine* is served, *fish* should be preferred to *meat*. Therefore the particular case, i.e., the exception which should be “enforced” occurs when *red wine* is served. In the terminology of defeasible reasoning we say that $red \wedge meat \triangleright red \wedge fish$ is more specific than $fish \triangleright meat$ because the first statement is true in the context *red wine* while the second is expressed in a more general context. Being more specific, the first statement takes precedence over the second one.

In order to deal with defeasible preferences interpreted following ceteris paribus semantics, Tan and Pearl (1994) rank-order comparative preference statements w.r.t. their specificity. Thus ceteris paribus semantics is first applied to most specific preferences. Less specific preferences are then considered as soon as they do not lead to contradiction. Therefore we have $meat - red \succ fish - red$ (given $red \wedge meat \triangleright red \wedge fish$ since it is more specific than $fish \triangleright meat$). Then we have $fish - white \succ meat - white$ (given $fish \triangleright meat$) but $fish - red \succ meat - red$ is not accepted. van Benthem et al. (2009) speak about normal situations. That is, $fish \triangleright meat$ is applied in normal situation, namely when $\neg red$ is true. Therefore we have $\neg red \wedge fish \triangleright \neg red \wedge meat$ and $red \wedge meat \triangleright red \wedge fish$. Note however that in both works we need additional information about the specificity between preference statements and normal situations.

Besides, let us mention that optimistic and pessimistic semantics have been proposed in non-monotonic reasoning to deal with defeasible knowledge (Pearl 1990; Benferhat et al. 2002). Given that optimistic semantics requires that at least one $\alpha \wedge \neg\beta$ -outcome should be preferred to any $\neg\alpha \wedge \beta$ -outcome, it leaves room for exceptions. Therefore $fish \triangleright meat$ and $red \wedge meat \triangleright red \wedge fish$ can be consistently handled together. For example the preference relation $fish - white \succ meat - white \approx meat - red \succ fish - red$ satisfies both statements w.r.t. optimistic semantics. Pessimistic semantics also deals with defeasible preferences. It works in a dual way w.r.t. optimistic semantics. The preference relation $meat - red \succ fish - white \approx fish - red \succ$

meat – white satisfies the two preference statements w.r.t. pessimistic semantics. Opportunistic semantics is the weakest semantics. Nevertheless it is not less useful. We refer the reader to (van der Torre and Weydert 2001) where an example shows that a preference relation can be derived given opportunistic semantics but not the other semantics.

What do we know about semantics?

Among the five semantics, *ceteris paribus* has attracted much attention of artificial intelligence researchers, philosophers and psychologists. In contrast to this semantics, strong, optimistic, pessimistic and opportunistic semantics (in particular the latter three) have attracted less attention in the preference representation community. Nevertheless they have been studied from algorithmic perspective. Specifically, algorithms have been developed to compute a distinguished preference relation associated with a set of preference statements and a given semantics (Pearl 1990; Benferhat et al. 2002). However much less is relatively known about their properties. For example it is not known which of pessimistic, optimistic or opportunistic semantics to use when dealing with defeasible preferences. In the next section, we study the behavior of the five semantics w.r.t. a set of postulates.

Postulate-based analysis of preference semantics

Our aim in this section is to make bridge between intuition and theoretical results. More precisely, we consider some postulates proposed in the literature for preference logics and check whether they are satisfied by the semantics.

Postulates

One may imagine a multitude of postulates for comparative preference statements. Nevertheless we will focus on a set of postulates which are in accordance with intuition behind the semantics. They also refer to comparative preference statements that are inferred given one or more comparative preference statements. Let us be more precise. Recall that a comparative preference statement $\alpha \triangleright \beta$ leads to the comparison of two sets, namely $\alpha \wedge \neg\beta$ -outcomes and $\neg\alpha \wedge \beta$ -outcomes. Then each semantics selects at least one or all $\alpha \wedge \neg\beta$ - and $\neg\alpha \wedge \beta$ -outcomes. For example $\succeq \models_{\exists\forall} \alpha \triangleright \beta$ means that at least one $\alpha \wedge \neg\beta$ -outcome is preferred w.r.t. \succeq to any $\neg\alpha \wedge \beta$ -outcome. Therefore if we are provided with another preference statement $\alpha' \triangleright \beta'$ such that $Mod(\alpha \wedge \neg\beta) \subset Mod(\alpha' \wedge \neg\beta')$ and $Mod(\neg\alpha \wedge \beta) \subseteq Mod(\neg\alpha' \wedge \beta')$ we can ensure that $\succeq \models_{\exists\forall} \alpha' \wedge \neg\beta'$. This means that optimistic semantics is tolerant for expanding the set of $\alpha \wedge \neg\beta$ -outcomes and reducing the set of $\neg\alpha \wedge \beta$ -outcomes. We now give a formal definition of tolerance for expansion/reduction.

Definition 2 (Expansion/reduction tolerance) *Let \succeq be a preference relation and $\alpha \triangleright \beta$ be a comparative preference statement. Let $x, y \in \{\exists, \forall\}$.*

- *A semantics is left- (resp. right-) expansion tolerant iff $\forall \succeq$, if $\succeq \models_{xy} \alpha \triangleright \beta$ then $\succeq \models_{xy} \alpha' \triangleright \beta'$ with $Mod(\alpha \wedge \neg\beta) \subset Mod(\alpha' \wedge \neg\beta')$ (resp. $Mod(\neg\alpha \wedge \beta) \subset Mod(\neg\alpha' \wedge \beta')$).*

- *A semantics is left- (resp. right-) reduction tolerant iff $\forall \succeq$, if $\succeq \models_{xy} \alpha \triangleright \beta$ then $\succeq \models_{xy} \alpha' \triangleright \beta'$ with $Mod(\alpha' \wedge \neg\beta') \subset Mod(\alpha \wedge \neg\beta)$ (resp. $Mod(\neg\alpha' \wedge \beta') \subset Mod(\neg\alpha \wedge \beta)$).*

It is worth noticing that the construction of $\alpha' \triangleright \beta'$ is not an end-point for itself. We aim to construct such a statement in a way that coincides with intuition and serves for real applications. For example given two preference statements $\alpha \triangleright \gamma$ and $\alpha \triangleright \beta$, one would intuitively expect that $\alpha \triangleright \beta \vee \gamma$ and/or $\alpha \triangleright \beta \wedge \gamma$ holds. We aim to check whether the semantics validate this intuition or not. A typical application of such inferences is recommender systems when, based on previous preferences of a user, we try to refine them by inferring new preferences.

In addition to postulates related to reduction and expansion principles, we also consider postulates related to the coherence and syntax independence. We first list the postulates.

Coherence - P1: if $\alpha \triangleright \beta$ then $not(\beta \triangleright \alpha)$

Syntax independence - P2: if $\alpha \equiv \alpha'$ and $\alpha \triangleright \beta$ then $\alpha' \triangleright \beta$
- if $\beta \equiv \beta'$ and $\alpha \triangleright \beta$ then $\alpha \triangleright \beta'$

Left composition - P3: if $\alpha \triangleright \gamma$ and $\beta \triangleright \gamma$ then $\alpha \vee \beta \triangleright \gamma$

Left decomposition - P4: if $\alpha \vee \beta \triangleright \gamma$ then $(\alpha \triangleright \gamma$ and $\beta \triangleright \gamma)$

Right composition - P5: if $\alpha \triangleright \beta$ and $\alpha \triangleright \gamma$ then $\alpha \triangleright \beta \vee \gamma$

Right decomposition - P6: if $\alpha \triangleright \beta \vee \gamma$ then $(\alpha \triangleright \beta$ and $\alpha \triangleright \gamma)$

Preference independence - P7: if $\alpha \triangleright \beta$ then $\alpha \vee \gamma \triangleright \beta \vee \gamma$

Left weakening - P8: if $Mod(\alpha') \subset Mod(\alpha)$ and $\alpha \triangleright \beta$ then $\alpha' \triangleright \beta$

Right weakening - P9: if $Mod(\beta') \subset Mod(\beta)$ and $\alpha \triangleright \beta$ then $\alpha \triangleright \beta'$

These postulates have been borrowed or adapted from (van Bentheim, Girard, and Roy 2009; Kraus, Lehmann, and Magidor 1990; Barbera, Bossert, and Pattanaik 2004).

P1 is intuitively natural. It says that if an individual expresses a strict preference for a statement against another statement then this means that she does not strictly prefer the latter to the former. P2 expresses a syntax independence w.r.t. both α and β . P3 and P5 express composition of preferred formulas or less preferred ones. At first sight, P4 may not appear natural because it departs from $\alpha \vee \beta \triangleright \gamma$ and ends up with $\alpha \triangleright \gamma$ and $\beta \triangleright \gamma$ (and not $\alpha \triangleright \gamma$ or $\beta \triangleright \gamma$). The idea behind this postulate is the following. Since $\alpha \vee \beta \triangleright \gamma$ turns out to prefer $(\alpha \vee \beta) \wedge \neg\gamma$ -outcomes then we also prefer the two sets $\alpha \wedge \neg\gamma$ -outcomes and $\beta \wedge \neg\gamma$ -outcomes taken separately. Nevertheless inferring $\alpha \triangleright \gamma$ or $\beta \triangleright \gamma$ is also meaningful. It is however captured by P8 since $Mod(\alpha) \subset Mod(\alpha \vee \beta)$. A similar reasoning is drawn in P6. P7 expresses that if α is preferred to β then the preference holds between two statements that extend them with the same formula. P8 says that if α is preferred to β then a subset of α in terms of outcomes

Table 1: Left/Right expansion/reduction principles involved in the postulates.

	Left-expansion	Left-reduction	Right-expansion	Right-reduction
P1	-	-	-	-
P2	-	-	-	-
P3	✓			✓
P4		✓	✓	
P5		✓	✓	
P6	✓			✓
P7		✓		✓
P8		✓	✓	
P9	✓			✓

is still preferred to β . P9 applies the same principle to less preferred formulas.

All the postulates but P1 and P2 refer to expansion and/or reduction principles. Consider for example the postulate P3. Given $\alpha \triangleright \gamma$ and $\beta \triangleright \gamma$ we want to check whether $\alpha \vee \beta \triangleright \gamma$ holds. The statement $\alpha \triangleright \gamma$ (resp. $\beta \triangleright \gamma$) compares $\alpha \wedge \neg\gamma$ -outcomes and $\neg\alpha \wedge \gamma$ -outcomes (resp. $\beta \wedge \neg\gamma$ -outcomes and $\neg\beta \wedge \gamma$ -outcomes). On the other hand, $\alpha \vee \beta \triangleright \gamma$ compares $(\alpha \wedge \neg\gamma) \vee (\beta \wedge \neg\gamma)$ -outcomes and $\neg\alpha \wedge \neg\beta \wedge \gamma$ -outcomes. Therefore it leads to a left-expansion of $\alpha \triangleright \beta$ and $\alpha \triangleright \gamma$ and their right-reduction. Table 1 summarizes the principles involved in each postulate. Given this table, we can already state an impossibility result:

if there is no semantics which is tolerant for left-expansion, left-reduction, right-expansion and right-reduction together then the postulates cannot be satisfied all together.

Nevertheless we have the following proposition which gives sufficient conditions to satisfy subsets of postulates.

Proposition 2 • *If a given semantics is left-expansion and right-reduction tolerant then it satisfies P3, P6 and P9.*

- *If a given semantics is left-reduction and right-expansion tolerant then it satisfies P4, P5 and P8.*
- *If a given semantics is left-reduction and right-reduction then it satisfies P7.*

The above proposition is intended to have a general characterization of any semantics (not necessarily one of the five semantics). This is why it works in one direction (if then) providing sufficient but not necessary conditions. In the next subsection, we instantiate these results, with additional results, on the five semantics.

Focus on the five semantics

Let us now consider again our five semantics. The following proposition gives the tolerance of each semantics w.r.t. reduction/expansion principles.

Proposition 3 *Table 2 summarizes the tolerance of each semantics for left/right expansion/reduction.*

Given Proposition 2 and Table 2, Table 3 reports the satisfaction or not of each postulate by the five semantics. A satisfaction means that any preference relation \succeq which satisfies the antecedent of “If” then it also satisfies its consequence. For example a given semantics satisfies P1 if for all \succeq such that $\succeq \models_{xy} \alpha \triangleright \beta$ then \succeq does not satisfy $\beta \triangleright \alpha$. Table 3 ensures that if a semantics is tolerant to a reduction/expansion and such a reduction/expansion is involved in a postulate then the semantics satisfies the postulate in question. For example optimistic semantics is left expansion and right reduction tolerant. The latter are involved in P3, P6 and P9. Therefore optimistic semantics satisfies postulates P3, P6 and P9. YES that are marked with * do not follow from Proposition 2.

What should we conclude?

The postulate-based analysis given in the previous section is intended to be a descriptive analysis which helps understand the behavior of the different semantics. It turns that opportunistic has bad properties given our postulates. This is not surprising as it is the weakest semantics. Nevertheless it is still useful in other frameworks, e.g. interval orders. Therefore this semantics calls for further investigation of its properties. Clearly, the choice of a semantics to use may be made on the basis of postulates we aim to satisfy. From Table 3, we know that strong semantics is coherent, syntax independent and it ensures that (i) $\alpha \vee \beta \triangleright \gamma$ entails $\alpha \triangleright \gamma$ and $\beta \triangleright \gamma$, (ii) $\alpha \triangleright \beta$ and $\alpha \triangleright \gamma$ entail $\alpha \triangleright \beta \vee \gamma$ and (iii) $\alpha \triangleright \beta$ entails $\alpha \vee \gamma \triangleright \beta \vee \gamma$.

Ceteris paribus does not satisfy much postulates. It only ensures coherence and preference independence. Thus this semantics does not allow any decomposition/composition.

Lastly we previously said that optimistic and pessimistic semantics exhibit a dual behavior. This is also reflected in Table 3. While optimistic semantics allows left composition, right decomposition and right weakening, pessimistic semantics allows left decomposition, right composition and left weakening.

Table 2: Left/Right expansion/reduction tolerance of the semantics.

	Left-expansion	Left-reduction	Right-expansion	Right-reduction
Strong	NO	YES	NO	YES
Ceteris Paribus	NO	YES	NO	YES
Optimistic	YES	NO	NO	YES
Pessimistic	NO	YES	YES	NO
Opportunistic	YES	NO	YES	NO

Table 3: Postulates satisfaction.

Postulates	Strong	Ceteris Paribus	Optimistic	Pessimistic	Opportunistic
P1	YES	YES	YES	YES	NO
P2	YES	NO	YES	YES	YES
P3	YES*	NO	YES	NO	NO
P4	NO	NO	NO	YES	NO
P5	YES*	NO	NO	YES	NO
P6	NO	NO	YES	NO	NO
P7	YES	YES	NO	NO	NO
P8	NO	NO	NO	YES	NO
P9	NO	NO	YES	NO	NO

Conclusion

Comparative preference statements represent a common ingredient of different preference representation languages. They have been studied both in artificial intelligence and philosophy. Different semantics (strong, ceteris paribus, optimistic, pessimistic, opportunistic) have been studied in the literature. In this paper, we provided a postulate-based analysis of the semantics. These postulates are a formal description of intuition one might have about the composition and decomposition of comparative preference statements. This analysis should give an indication of which semantics to be used depending on the properties we aim to satisfy about such composition and decomposition. It is also useful in recommender systems when we need to infer user's preferences on the basis of her previous preferences.

As we previously said, the five semantics we studied in this paper have been separately addressed in the literature. The present work offers a complete picture of all these semantics which permits to choose the one to be used for which purpose (i.e., which properties we would like to satisfy).

References

Barbera, S.; Bossert, W.; and Pattanaik, P. 2004. Ranking sets of objects. In *Chapter 17 of the Handbook of Utility Theory, vol 2*.

Benferhat, S.; Dubois, D.; Kaci, S.; and Prade, H. 2002. Bipolar representation and fusion of preferences in the possibilistic logic framework. In *KR'02*, 421–432.

Boutilier, C.; Brafman, R.; Domshlak, C.; Hoos, H.; and Poole, D. 2004. CP-nets: A tool for representing and reason-

ing with conditional ceteris paribus preference statements. *JAIR* 21:135–191.

Boutilier, C. 1994. Toward a logic for qualitative decision theory. In *KR'94*, 75–86.

Halldén, S. 1957. On the logic of “Better”. In *Library of Theoria, Lund*.

Hansson, S. 2001. The structure of values and norms. In *Cambridge University Press*.

Kaci, S., and van der Torre, L. 2008. Reasoning with various kinds of preferences: Logic, non-monotonicity and algorithms. *Annals of Operations Research* 163(1):89–114.

Kaci, S. 2011. *Working with Preferences: Less Is More*. Cognitive Technologies. Springer.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intel.* 44(1-2):167–207.

Pearl, J. 1990. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *TARK'90*, 121–135.

Tan, S., and Pearl, J. 1994. Specification and evaluation of preferences under uncertainty. In *KR'94*, 530–539.

van Bentheim, J.; Girard, P.; and Roy, O. 2009. Everything Else Being Equal: A Modal Logic for Ceteris Paribus Preferences. *Journal of Philosophical Logic* 38:83–125.

van der Torre, L., and Weydert, E. 2001. Parameters for utilitarian desires in a qualitative decision theory. *Applied Intel.* 14(3):285–301.

von Wright, G. H. 1963. *The Logic of Preference*. University of Edinburgh Press.

Wilson, N. 2004. Extending CP-nets with stronger conditional preference statements. In *AAAI'04*, 735–741.