

From Preference Logics to Preference Languages, and Back

Meghyn Bienvenu

Fachbereich Mathematik und Informatik
Universität Bremen, Germany
meghyn@informatik.uni-bremen.de

Jérôme Lang

LAMSADE
Université Paris-Dauphine, France
lang@lamsade.dauphine.fr

Nic Wilson

Cork Constraint Computation Centre
University College Cork, Ireland
n.wilson@4c.ucc.ie

Abstract

Preference logics and AI preference representation languages are both concerned with reasoning about preferences on combinatorial domains, yet so far these two streams of research have had little interaction. This paper contributes to the bridging of these areas. We start by constructing a “prototypical” preference logic, which combines features of existing preference logics, and then we show that many well-known preference languages, such as CP-nets and its extensions, are natural fragments of it. After establishing useful characterizations of dominance and consistency in our logic, we study the complexity of satisfiability in the general case as well as for meaningful fragments, and we study the expressive power as well as the relative succinctness of some of these fragments.

1. Introduction

Reasoning about preferences on combinatorial domains is an important problem which has been addressed by at least two different streams of work: *preference logics*, mainly studied by philosophers, and *preference representation languages*, mainly studied by researchers in Artificial Intelligence. These two areas have had little interaction so far, partly because they have been studied by two distinct research communities, and also because they focus on different issues: axiomatization and philosophical issues for preference logics; elicitation, compact representation and computational issues for AI¹. Still, both communities have common interests: they both aim at designing complex languages for expressing, and reasoning about, preferences on complex domains; moreover, some key notions appear in both communities, such as the *ceteris paribus* principle for interpreting preference statements.

Bridging these two communities can bring benefits to both. Preference logics can “import” from preference languages some computational results and techniques (such as the complexity of reasoning tasks, or algorithms for optimization), whereas preference languages can gain formal logical foundations and added expressivity. Moreover, as we

show in the paper, some well-known preference languages are fragments of some well-known preference logics.

This paper partly contributes to establishing such a bridge. We start by giving some brief background on preference logics and on AI preference languages. Then we define a prototypical preference logic, obtained by combining ideas from the classical preference logic of von Wright (1963) and from the modal preference logic recently studied in (van Benthem, Girard, and Roy 2009). The basic constructs of the logic are preference statements between propositional formulas, together with a set of formulas that must take identical values when the statement is interpreted. The language is composed of boolean combinations of such preference statements. The semantics is defined in terms of preference relations on worlds (or outcomes, see further). We show that the conjunctive fragment of the language enjoys specific properties, such as the existence of a smallest preference relation satisfying a conjunction of preference statements. We also show how several well-known AI preference languages, such as CP-nets, TCP-nets, CP-theories, CI-nets or prioritized goals, can be recovered as particular fragments of the logic. We next provide a number of complexity results for our logic and some of its fragments, as well as for the logics of von Wright and van Benthem et al. (van Benthem, Girard, and Roy 2009). Finally, we give some expressiveness and succinctness results for different fragments of our logic. Please note that proofs have been omitted or abbreviated for lack of space. Missing proofs can be found in a long version available at <http://www.informatik.uni-bremen.de/~meghyn/pl-long.pdf>.

2. Background

Throughout the paper, we consider a propositional language built from a *finite* set of propositional symbols PS . Under this assumption, we can identify propositional valuations (or worlds) with maximal consistent conjunctions of literals. For instance, if $PS = \{a, b\}$, the valuation where a is assigned true and b is assigned false is identified with the formula $a \wedge \neg b$. Such maximal consistent conjunctions of literals will be called *outcomes* (a.k.a. alternatives). We freely use this correspondence between worlds and outcomes in notations, e.g., $Mod(\varphi)$ will denote the set of outcomes that imply φ (corresponding to the set of worlds that satisfy φ).

Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹We do not mean that compact preference representation languages studied in AI are not based on logic; actually, quite a number of them are (see e.g. (Brewka 2004; Coste-Marquis et al. 2004)); but still, there, the focus is laid on computational issues.

Preference Logics

A preference logic consists of a semantics and/or a formal system meant to interpret dyadic preferences between propositional formulas, or monadic, “absolute” preferences. The starting point of preference logics is that individuals often express relative or absolute preferences that refer not to isolated outcomes, but to logical formulas representing *sets of outcomes*, which are generally not singletons, nor even disjoint subsets. Indeed, interpreting the statement “ φ is preferred to ψ ” is unproblematic when φ and ψ are complete formulas (corresponding each to a unique outcome): such a statement corresponds directly to its semantical counterpart $\omega \succ \omega'$, where $Mod(\varphi) = \{\omega\}$ and $Mod(\psi) = \{\omega'\}$. In a similar way, a statement of indifference between φ and ψ corresponds to $\omega \sim \omega'$. Von Wright’s logic of preference (1963) interprets preferences between two logical formulas φ and ψ as “everything else being equal, I prefer an outcome satisfying $\varphi \wedge \neg\psi$ to an outcome satisfying $\psi \wedge \neg\varphi$ ”².

The main issue is how “everything else being equal” (*ceteris paribus*) should be treated when interpreting φ is preferred to ψ . More precisely, such a preference statement is generally interpreted as follows: φ is preferred to ψ is true whenever for every outcomes ω and ω' such that (i) $\omega \models \varphi \wedge \neg\psi$, (ii) $\omega' \models \psi \wedge \neg\varphi$ and (iii) ω and ω' are *ceteris paribus* with respect to $\varphi \wedge \neg\psi$ and $\psi \wedge \neg\varphi$, then ω is preferred to ω' . The notion of two outcomes being *ceteris paribus* with respect to two formulas has yet to be defined. Let $F(\alpha, \beta)$ be the set of pairs of outcomes that are *ceteris paribus* with respect to α and β (such a function F is called *representation function* in (Hansson 2001)). Several definitions of increasing generality have been proposed:

- choosing $F(\alpha, \beta) = Mod(\alpha) \times Mod(\beta)$ leads to interpreting φ is preferred to ψ as “every model of $\varphi \wedge \neg\psi$ is preferred to every model of $\neg\varphi \wedge \psi$ ” — see (von Wright 1963; Doyle, Shoham, and Wellman 1991).
- in (von Wright 1963), $(\omega, \omega') \in F(\alpha, \beta)$ if ω and ω' give the same truth values to all propositional variables that are mentioned neither in α nor in β .
- in (Doyle, Shoham, and Wellman 1991), and subsequently in (Tan and Pearl 1994), $(\omega, \omega') \in F(\alpha, \beta)$ if ω and ω' give the same truth values to all propositional variables that are irrelevant to α and to β , where propositional variable p is irrelevant to a formula φ if the truth value of φ does not depend on the truth value of p .
- (Doyle and Wellman 1994) take a much more general approach and define $F(\alpha, \beta)$ via a *contextual equivalence*

²Interpreting “ φ is preferred to ψ ” as $\varphi \wedge \neg\psi$ is preferred to $\neg\varphi \wedge \psi$ actually comes back to (Halldén 1957) and (Catañeda 1958). Note also that this principle fails whenever one of these propositions is a logical consequence of the other (since it would then consist in comparing a formula to a logical contradiction). Some natural statements fall in this limit case: consider $\varphi = I \text{ work hard and earn a lot of money}$ and $\psi = I \text{ work hard}$ (Hansson 2001). So as to take this limit case into account, Hansson (1989) proposes a generalization of the latter principle, where a preference for φ over ψ is interpreted as “everything else being equal, I prefer an outcome satisfying φ/ψ to an outcome satisfying ψ/φ ”, where α/β is equal to $\alpha \wedge \neg\beta$ when consistent, and to α otherwise.

relation which maps every set Σ of propositional formulas to an equivalence relation \sim_Σ on the set of outcomes (as their approach is very general, they do not give a specific way of defining \sim_Σ); and $F(\alpha, \beta)$ is simply taken to be the equivalence relation \sim_C assigned to the set $C = \{\alpha, \beta\}$.

- (Hansson 2001) also proposes a very general approach, based on arbitrary representation functions F , and focuses on representation functions based on a *similarity relation* between pairs of outcomes: $(\omega, \omega') \in F(\alpha, \beta)$ if (ω, ω') is a pair of maximally similar outcomes in $Mod(\alpha) \times Mod(\beta)$.

Then a major step is taken in (van Benthem, Girard, and Roy 2009), who define, and axiomatize, a preference logic \mathcal{L}_{CP} where preference is a genuine modality, and where preference statements explicitly mention the set of formulas to be kept constant: if Γ is a set of formulas, then the preference statement $\varphi > \psi \parallel \Gamma$ (the notation is ours) is true whenever for every pair of outcomes ω and ω' such that (i) $\omega \models \varphi \wedge \neg\psi$, (ii) $\omega' \models \psi \wedge \neg\varphi$ and (iii) ω and ω' coincide on all formulas of Γ , then ω is preferred to ω' .³ Having the set of fixed preferences specified in preference statements allows one to recover many of the previous approaches, including von Wright’s, as specific cases.

Other families of preference logics which are not directly relevant to this paper include nonmonotonic logics of preference, where preference statements are interpreted *all other things being normal* instead of *all other things being equal*, e.g. (Boutilier 1994; Lang, van der Torre, and Weydert 2003; Kaci and van der Torre 2008; Girard 2008); logics for preference change, e.g. (van Benthem and Liu 2007; Liu 2008); and logics for monadic preferences, e.g. (Hansson 2001).

Compact Preference Representation Languages

When the set of outcomes has a combinatorial structure, it is unrealistic to expect the agent to rank them explicitly. Compact preference languages are designed to exploit the structural properties (such as preferential independence or utility independence) that preference relations or utility functions often enjoy so as to elicit, store and process these preferences as compactly as possible. We don’t have enough space for a survey of compact preference languages, so we just name a few which we make reference to later in the paper.

The formalism *CP-nets* (Boutilier et al. 2004) is a graphical language for preference representation based on the notion of preferential independence. A *CP-net* over a set of binary variables PS is a pair $\langle \mathcal{G}, C \rangle$ where \mathcal{G} is an oriented graph whose vertices are PS , and $C = \{C(x) \mid x \in PS\}$ is a set of *conditional preference tables*: for every x , $C(x)$ specifies the preferences on the values of x given the values of its parents in \mathcal{G} . The edges of \mathcal{G} express preferential independencies: every x is preferentially independent of its non-parents in \mathcal{G} given its parents. TCP-nets (Brafman, Domshlak, and Shimony 2006) enrich CP-nets by allowing the expression of relative importance statements between variables. CP-theories (Wilson 2004b) are still more general; they allow

³The logic defined in (van Benthem, Girard, and Roy 2009) is actually considerably more general, as we discuss in Section 3.

conditional preference statements on the values of a variable, together with a set of variables that are allowed to vary when interpreting the preference statement. The language considered in (Wilson 2009) is even more general: there the preference statements do not compare single values of variables but tuples of values of different variables. CI-nets (Bouveret, Endriss, and Lang 2009) express monotonic preferences between sets of goods, *ceteris paribus*.

Other compact representation languages are based on propositional logic, such as prioritized goals (e.g., (Brewka 2004; Coste-Marquis et al. 2004; de Jongh and Liu 2009)).

All languages above are meant to express compactly *ordinal preferences*, that is, rankings on the set of outcomes. Other languages are meant to express *numerical preferences*, that is, utility functions (due to lack of space, and as our focus is on ordinal preferences, we omit references).

3. A Prototypical Preference Logic

We start by defining our “prototypical” preference logic, named *PL*, which can be seen as a cross between von Wright’s classic preference logic and the more recent and more expressive logic \mathcal{L}_{CP} of van Benthem et al. (2009).

Let L be the propositional language built from a finite set of propositional symbols PS , the usual connectives, and the Boolean constants \top , \perp . Let χ , the *integrity constraint*, be a consistent propositional formula in L . This can be used, for example, to simulate multi-valued variables with propositional ones. χ will be fixed throughout the paper. We will use the term *possible outcome* (abbreviated to *outcome*) to refer to conjunctions of propositional literals which contain each variable in PS exactly once, and which are consistent with χ . Possible outcomes correspond to valuations of PS which satisfy χ . Let Ω be the set of (possible) outcomes.

Definition 1 (language of *PL*).

- If α and β are formulas of L , and F is a set of formulas of L then $\alpha \triangleright \beta \parallel F$ and $\alpha \triangleright \beta \parallel F$ are both formulas of *PL*.
- if Φ and Ψ are formulas of *PL* then $\neg\Phi$, $\Phi \wedge \Psi$, $\Phi \vee \Psi$ are formulas of *PL*.

Formulas of the form $\alpha \triangleright \beta \parallel F$ and $\alpha \triangleright \beta \parallel F$ are called strict and non-strict *preference statements* respectively. An arbitrary formula in *PL* is called a *preference formula*, which is a boolean combination of some collection C of preference statements; the elements of C are called the *component preference statements* of the preference formula. It is convenient to define a notation for the constituent parts of a preference statement. If Φ is a preference statement $\varphi \triangleright \psi \parallel S$ then we define $\alpha_\Phi = \varphi$, $\beta_\Phi = \psi$ and $F_\Phi = S$, and similarly if $\Phi = \varphi \triangleright \psi \parallel S$. For instance, if $\Phi = a \triangleright \neg a \parallel \{b, c\}$, then $\alpha_\Phi = a$, $\beta_\Phi = \neg a$, and $F_\Phi = \{b, c\}$. We refer to F_Φ as the set of *fixed formulas* (of Φ). If α_Φ and β_Φ are both outcomes, then Φ is said to be a *basic preference statement*. If F_Φ is empty, we will sometimes abbreviate $\alpha \triangleright \beta \parallel \emptyset$ (resp. $\alpha \triangleright \beta \parallel \emptyset$) to $\alpha \triangleright \beta$ (resp. $\alpha \triangleright \beta$). We say that a preference formula is *free* if $F_\Phi = \emptyset$ for every component preference statement Φ , it is *conjunctive* if it is a conjunction of preference statements, and it is *positive* if it is built from preference statements using only \wedge and \vee . The conjunctive

fragment of *PL*, consisting of only conjunctive preference formulas, is denoted by PL_C .

It is convenient to add the following notation, capturing the usual *ceteris paribus* notion: $\alpha R \beta \parallel CP$ (for $R \in \{\triangleright, \triangleright\}$) is shorthand for $\alpha R \beta \parallel F_{\alpha, \beta}$ where $F_{\alpha, \beta}$ is the set of propositional symbols that do not appear in α nor in β .

Models and satisfaction in *PL* are defined as follows.

Definition 2 (models of *PL*). *A model of PL is a complete pre-order \succsim on the set Ω of outcomes.*

\succ and \sim are the usual strict order and indifference relations induced from \succsim . Given two outcomes $\omega, \omega' \in \Omega$ and a set F of formulas in L , we say that ω and ω' *agree on F*, denoted by $\omega \approx_F \omega'$, if for every $\gamma \in F$ we have $\omega \models \gamma$ if and only if $\omega' \models \gamma$.

Definition 3 (satisfaction in *PL*).

- $\gamma \models \alpha \triangleright \beta \parallel F$ if $\omega \succ \omega'$ holds for all $\omega, \omega' \in \Omega$ such that $\omega \models \alpha$, $\omega' \models \beta$, and $\omega \approx_F \omega'$
- $\gamma \models \alpha \triangleright \beta \parallel F$ if $\omega \succsim \omega'$ holds for all $\omega, \omega' \in \Omega$ such that $\omega \models \alpha$, $\omega' \models \beta$, and $\omega \approx_F \omega'$

Boolean connectives are interpreted in the usual way.

A preference formula Γ is *consistent* if it has at least one model, and inconsistent otherwise. A preference formula Γ *entails* another preference formula Υ , written $\Gamma \models \Upsilon$, if every model of Γ is also a model of Υ , which holds if and only if the preference formula $\Gamma \wedge \neg\Upsilon$ is inconsistent. Equivalence (\equiv) is defined in the usual way.

Discussion and Relation to Existing Preference Logics

Our preference logic, like von Wright’s (1963), consists of Boolean combinations of preference statements. However, in place of von Wright’s preference statements $\alpha P \beta$, which are interpreted as “ $\alpha \wedge \neg\beta$ -outcomes are preferred to $\neg\alpha \wedge \beta$ -outcomes when the variables outside α and β are held constant”, we use more general preference statements of the form $\varphi \triangleright \psi \parallel F$ which state that “ φ -outcomes are preferred to ψ -outcomes provided they agree on the interpretation of formulas in F ”. Thus, von Wright’s preference logic can be straightforwardly encoded in our own simply by considering Boolean combinations of preference statements of the form $\alpha \wedge \neg\beta \triangleright \beta \wedge \neg\alpha \parallel CP$. In a similar manner, the preference statements found in (Doyle, Shoham, and Wellman 1991) and (Tan and Pearl 1994) can be captured in our logic by preference statements of the form $\alpha \wedge \neg\beta \triangleright \beta \wedge \neg\alpha \parallel I$, where I is the set of variables irrelevant to α and β .

Our decision to interpret $\alpha \triangleright \beta \parallel F$ as a comparison between α - and β -outcomes rather than $\alpha \wedge \neg\beta$ - and $\beta \wedge \neg\alpha$ -outcomes was based on the fact that the latter semantics is easily captured in our own framework, whereas a translation in the other direction is cumbersome at best. In practice, this means that preference statements have been preprocessed – for instance, “I prefer an ice-cream (i) to a piece of cake (c), everything else being equal” is directly expressed as $i \wedge \neg c \triangleright \neg i \wedge c \parallel CP$.

Note that preconditions for preferences are easily expressed within our framework: “if γ then α is preferred

to β with F fixed” is expressed using the preference statement $\gamma \wedge \alpha \wedge \neg \beta \triangleright \gamma \wedge \neg \alpha \wedge \beta \parallel F$. Also note that unlike von Wright (and in accordance with van Benthem et al. 2009), we choose to include both strict and non-strict preference statements, since non-strict preference statements are required for the expression of indifference between sets of outcomes. We also point out that our semantics is based on total pre-orders, as are the semantics for most AI preference languages as well as von Wright’s preference logic. The logic of van Benthem et al. employs partial pre-orders instead.

Finally, we remark that while our preference logic is inspired by \mathcal{L}_{CP} (van Benthem, Girard, and Roy 2009), and incorporates some of its features (e.g. use of an explicit fixed set of formulas), it is considerably less general. Indeed, the modal logic of preferences \mathcal{L}_{CP} augments propositional logic with a set of unary (rather than binary) modal operators of preference which are relativised by a fixed set of formulas. Using these unary operators, different types of binary preference statements with various semantics (including our own) can be constructed. However, as we shall see later in the paper, adopting the full expressivity of \mathcal{L}_{CP} also means accepting an increase in the computational complexity of reasoning.

In this paper we make a trade-off between complexity and expressivity, and choose to work with our simple, yet still fairly expressive, preference logic PL , which covers and generalizes many AI preference languages as well as preference logics (including von Wright’s), but also (we believe) will be sufficiently expressive in most practical situations. Investigating the computational properties of logics allowing for preference statements with different semantics is left for further research.

4. Preference Languages as Fragments of PL

We show that several well-known preference languages are fragments of our prototypical preference logic PL .

CP-Nets

Let $\mathcal{N} = \langle \mathcal{G}, C \rangle$ be a propositional CP-net. We show that \mathcal{N} can be expressed as a conjunctive preference formula $\Gamma_{\mathcal{N}}$, built as follows. For every variable x , let $U \subseteq PS \setminus \{x\}$ be the parents of x in \mathcal{G} . Then for every entry $u : x \succ \bar{x}$ (resp. $u : \bar{x} \succ x$) in the conditional preference table $C(x)$, we have a conjunct $\gamma_u \wedge x \triangleright \gamma_u \wedge \neg x \parallel CP$ (resp. $\gamma_u \wedge \neg x \triangleright \gamma_u \wedge x \parallel CP$), where γ_u is the conjunction of literals corresponding to assignment u to variables U . For instance, let $PS = \{a, b\}$, let \mathcal{G} contain only a single edge from a to b , and let $C(a) = \{a \succ \bar{a}\}$ and $C(b) = \{a : b \succ \bar{b}, \bar{a} : \bar{b} \succ b\}$; then $\Gamma_{\mathcal{N}} = (a \triangleright \neg a \parallel \{b\}) \wedge (a \wedge b \triangleright a \wedge \neg b) \wedge (\neg a \wedge \neg b \triangleright \neg a \wedge b)$.

It can be shown that for any two outcomes ω, ω' , \mathcal{N} implies $\omega \succ \omega'$ if and only if $\Gamma_{\mathcal{N}} \models \omega \triangleright \omega'$. Therefore CP-nets are expressible within PL and correspond to a syntactic restriction of its conjunctive fragment (made up of conjunctions of preference statements bearing on conjunctions of literals satisfying some specific properties that we won’t state here). Moreover, the CP-net fragment of PL can be encoded in von Wright’s preference logic, since all component preference statements Γ are taken *ceteris paribus* and are such

that α_{Γ} and β_{Γ} are mutually inconsistent.

CP-nets and related formalisms are often based on a set V of multi-valued variables, i.e., not just boolean variables. Each variable x has an associated set of possible values $D(x)$. We can embed multi-valued variables into propositional logic in a standard way⁴: for each $v \in D(x)$ we create a propositional variable x^v . Exhaustivity of $D(x)$ is expressed with the propositional formula $\bigvee_{v \in D(x)} x^v$, and mutual exclusivity with the formula $\bigwedge_{u, v \in D(x), u \neq v} (\neg x^u \vee \neg x^v)$. We can define the integrity constraint χ to be $\bigwedge_x \chi(x)$, where $\chi(x)$ is the conjunction of the exhaustivity formula and the mutual exclusivity formula for variable x . Complete assignments to V are then in one-to-one correspondence with outcomes.

For instance, let x and y two variables, with $D(x) = \{x_1, x_2, x_3\}$ and $D(y) = \{y, \bar{y}\}$. Consider the CP-net whose dependency graph contains only an edge from X to Y , and whose conditional preference tables are $x_1 \succ x_2 \succ x_3$ and $x_1 : y \succ \bar{y}$, $x_2 : \bar{y} \succ y$, $x_3 : y \succ \bar{y}$. This CP-net is expressed within PL by the propositional variables x^1, x^2, x^3 and y , the integrity constraint $\chi = (x^1 \vee x^2 \vee x^3) \wedge (\neg x^1 \vee \neg x^2) \wedge (\neg x^1 \vee \neg x^3) \wedge (\neg x^2 \vee \neg x^3)$, and the preference formula $\Gamma = (x^1 \triangleright x^2 \parallel \{y\}) \wedge (x^2 \triangleright x^3 \parallel \{y\}) \wedge (x^1 \wedge y \triangleright x^1 \wedge \neg y) \wedge (x^2 \wedge \neg y \triangleright x^2 \wedge y) \wedge (x^3 \wedge y \triangleright x^3 \wedge \neg y)$.

Extensions of CP-Nets

A language described in (Wilson 2009) involves preference statements over a set V of multi-valued variables of the form $p \succ q \parallel T$ where p is an assignment to a set of variables P , q is an assignment to set of variables Q , and T is a set of variables. Such a statement expresses a preference for a complete assignment θ over a complete assignment ρ whenever (i) θ extends p , (ii) ρ extends q , and (iii) θ and ρ agree on variables T . The preference language from (Wilson 2004b) involves statements of this form but where $P = Q$, and p and q differ on exactly one variable. CP-nets can be expressed with statements of this form with the additional condition that $T = V \setminus P$. As shown in (Wilson 2004a), TCP-nets can also be expressed with such statements where $|V \setminus (P \cup T)| = 0$ or 1.

Statement $p \succ q \parallel T$ can be easily mapped into an equivalent preference statement⁵ $p' \triangleright q' \parallel T'$ in PL , where assignments p and q are mapped to conjunctions of literals in the obvious way, i.e., $p' = \bigwedge_{x \in P} x^{p(x)}$, $q' = \bigwedge_{x \in Q} x^{q(x)}$, and $T' = \{x^v : x \in T, v \in D(x)\}$.

Conditional Importance Networks

Conditional importance networks (CI-networks) are compact representations of preference orderings over subsets of set V of objects (goods). A CI-network is a set of conditional importance statements (CI-statements). Each CI-statement is of the form $S^+, S^- : S_1 \triangleright S_2$, and it expresses that for each subset S' of T , $S^+ \cup S_1 \cup S'$ is preferred to $S^+ \cup S_2 \cup S'$,

⁴Obviously, shorter translations are possible, expressing each multi-valued variable by a logarithmic number of binary variables.

⁵Or alternatively to $p' \triangleright q' \parallel T'$ if instead one is using a non-strict semantics.

where $\mathcal{T} = V \setminus (\mathcal{S}^+ \cup \mathcal{S}^- \cup \mathcal{S}_1 \cup \mathcal{S}_2)$ is the set of other objects. For instance, a CI-statement $\{b\}, \emptyset : \{a\} \triangleright \{c, d\}$ expresses that the bundle of goods $\{a, b\} \cup X$ is preferred to the bundle of goods $\{b, c, d\} \cup X$, for any X such that $X \cap \{a, b, c, d\} = \emptyset$.

We now show how a CI-network corresponds to a conjunctive preference formula. We map each object p to a propositional variable which we also call p . A set \mathcal{S} of objects then corresponds to an outcome whose positive literals are the elements of \mathcal{S} , and with negative literals for each other propositional variable. In this way a preference ordering over sets of objects corresponds to a preference ordering over outcomes.

The CI-statement $\mathcal{S}^+, \mathcal{S}^- : \mathcal{S}_1 \triangleright \mathcal{S}_2$ can be mapped to the preference statement $\gamma \wedge \delta \wedge \gamma_1 \triangleright \gamma \wedge \delta \wedge \gamma_2 \parallel \mathcal{T}$ in PL , where γ (respectively, γ_1, γ_2) is the conjunction of positive literals in \mathcal{S}^+ (respectively, $\mathcal{S}_1, \mathcal{S}_2$), and δ is the conjunction of negative literals for propositional variables in \mathcal{S}^- .

The preference ordering on sets is assumed to respect monotonicity, so that a set of goods is preferred to any proper subset of it. This can be represented in our language by the conjunction of preference statements $p \triangleright \neg p \parallel V \setminus \{p\}$ over all $p \in V$.

Prioritized Goals

A prioritized goal base is a sequence $G = \langle G_1, \dots, G_n \rangle$ where each G_i is a finite set of propositional formulas. For $\omega \in \Omega$, let $S(\omega, G_i) = \{\varphi \in G_i \mid \omega \models \varphi\}$. Three well-known semantics for inducing a preference relation from G (cf. (Benferhat et al. 1993)) are:

- *best-out*: $\omega \succ_G^{bo} \omega'$ if $\min\{i \mid S(\omega, G_i) \neq G_i\} > \min\{i \mid S(\omega', G_i) \neq G_i\}$;
- *discrimin*: $\omega \succ_G^{disc} \omega'$ if either (a) there is a $k \leq n$ such that $S(\omega', G_k) \subsetneq S(\omega, G_k)$ and for every $i < k$, $S(\omega, G_i) = S(\omega', G_i)$; or (b) $S(\omega, G_k) = S(\omega', G_k)$ for every $k \leq n$.
- *leximin*: $\omega \succ_G^{lex} \omega'$ if either (a) there is a $k \leq n$ such that $|S(\omega, G_k)| > |S(\omega', G_k)|$ and for every $i < k$, $|S(\omega, G_i)| = |S(\omega', G_i)|$; or (b) for every $k \leq n$, $|S(\omega, G_k)| = |S(\omega', G_k)|$.

In the *best-out* case, let γ_i be the conjunction of all formulas in G_i and $\delta_i = \gamma_1 \wedge \dots \wedge \gamma_{i-1} \wedge \neg \gamma_i$ for $1 \leq i \leq n$, and $\delta_{n+1} = \gamma_1 \wedge \dots \wedge \gamma_n$. Then we associate the following formula of PL with the goal base G :

$$\Phi_{G,bo} = \left(\bigwedge_{i=1}^n \gamma_i \wedge \dots \wedge \gamma_i \triangleright \delta_i \parallel \emptyset \right) \wedge \left(\bigwedge_{i=1}^{n+1} \delta_i \triangleright \delta_i \parallel \emptyset \right)$$

Proposition 4. For any $\omega, \omega' \in \Omega$, we have $\omega \succ_G^{bo} \omega'$ if and only if $\Phi_{G,bo} \models \omega \triangleright \omega'$.

The *discrimin* and *leximin* cases are more interesting, since the translation into PL makes use of fixed formulas. For the sake of simplicity we first focus on prioritized goal bases where each G_i is a singleton $\{\gamma_i\}$. In this case, \succ_G^{disc} and \succ_G^{lex} coincide and $\omega \succ_G^{disc} \omega'$ if either (a) there is a $k \leq n$ such that $\omega \models \gamma_k$, $\omega' \models \neg \gamma_k$ and for every $i < k$, $\omega \models \gamma_i$ if and only if $\omega' \models \gamma_i$ or (b) for every $i \leq n$, $\omega \models \gamma_i$ if and only if

$\omega' \models \gamma_i$. In this case, the formula of PL associated with G is:

$$\bigwedge_{i=1}^n (\gamma_i \triangleright \neg \gamma_i \parallel \{\gamma_1, \dots, \gamma_{i-1}\}) \wedge (\top \triangleright \top \parallel \{\gamma_1, \dots, \gamma_n\})$$

For instance, the goal base $G = \langle \{a\}, \{b \vee c\}, \{b \wedge c\} \rangle$ is mapped to $(a \triangleright \neg a \parallel \emptyset) \wedge (b \vee c \triangleright \neg(b \vee c) \parallel \{a\}) \wedge (b \wedge c \triangleright \neg(b \wedge c) \parallel \{a, b \vee c\}) \wedge (\top \triangleright \top \parallel \{a, b \vee c, b \wedge c\})$ in PL .

In the general case where the G_i 's are not necessarily singletons, a translation to PL is possible for *discrimin* and *leximin*, but leads to an exponential blow-up.

In the *discrimin* case, we can use the following formula to encode the goal base G :

$$\Phi_{G,disc} = \bigwedge_{i=1}^n \bigwedge_{\substack{S \subseteq G_i \\ S \neq \emptyset}} \left(\bigwedge_{\varphi \in S} \varphi \wedge \bigwedge_{\psi \in G_i \setminus S} \neg \psi \triangleright \bigvee_{\varphi \in S} \neg \varphi \wedge \bigwedge_{\psi \in G_i \setminus S} \neg \psi \parallel \bigcup_{k=1}^{i-1} G_k \right) \wedge (\top \triangleright \top \parallel \bigcup_{i=1}^n G_i)$$

Proposition 5. For any $\omega, \omega' \in \Omega$, we have $\omega \succeq_G^{disc} \omega'$ if and only if $\Phi_{G,disc} \models \omega \triangleright \omega'$.

For the *leximin* case, we will make use of the following abbreviation

$$\geq k : S \stackrel{\text{def}}{=} \bigvee_{\substack{S' \subseteq S \\ |S'| \geq k}} \left(\bigwedge_{\varphi \in S'} \varphi \right)$$

which expresses that at least k formulas in S hold true. The formula $\Phi_{G,lex}$ of PL associated with G can then be defined as follows:

$$\Phi_{G,lex} = \bigwedge_{i=1}^n \bigwedge_{j=1}^{|G_i|} (\geq j : G_i \triangleright \neg(\geq j : G_i) \parallel F_{i-1}) \wedge (\top \triangleright \top \parallel F_n)$$

where $F_i = \{\geq \ell : G_k \mid 1 \leq k \leq i, 1 \leq \ell \leq |G_k|\}$.

Proposition 6. For any $\omega, \omega' \in \Omega$, we have $\omega \succeq_G^{lex} \omega'$ if and only if $\Phi_{G,lex} \models \omega \triangleright \omega'$.

For instance, let $G = \langle G_1, G_2 \rangle$ with $G_1 = \{p \vee q, r\}$ and $G_2 = \{\neg p, \neg q\}$. We have

$$\begin{aligned} \Phi_{G,disc} = & ((p \vee q) \wedge r \triangleright (\neg p \wedge \neg q) \vee \neg r \parallel \emptyset) \\ & \wedge ((p \vee q) \wedge \neg r \triangleright \neg p \wedge \neg q \wedge \neg r \parallel \emptyset) \\ & \wedge (r \wedge \neg p \wedge \neg q \triangleright \neg r \wedge \neg p \wedge \neg q \parallel \emptyset) \\ & \wedge (\neg p \wedge \neg q \triangleright p \vee q \parallel \{p \vee q, r\}) \\ & \wedge (\neg p \wedge q \triangleright p \wedge q \parallel \{p \vee q, r\}) \\ & \wedge (p \wedge \neg q \triangleright p \wedge q \parallel \{p \vee q, r\}) \\ & \wedge (\top \triangleright \top \parallel \{p \vee q, r, \neg p, \neg q\}) \end{aligned}$$

and

$$\begin{aligned} \Phi_{G,lex} = & ((p \vee q) \wedge r \triangleright (\neg p \wedge \neg q) \vee \neg r \parallel \emptyset) \\ & \wedge (p \vee q \vee r \triangleright \neg p \wedge \neg q \wedge \neg r \parallel \emptyset) \\ & \wedge (\neg p \wedge \neg q \triangleright p \vee q \parallel \{(p \vee q) \wedge r, p \vee q \vee r\}) \\ & \wedge (\neg p \vee \neg q \triangleright p \wedge q \parallel \{(p \vee q) \wedge r, p \vee q \vee r\}) \\ & \wedge (\top \triangleright \top \parallel \{(p \vee q) \wedge r, p \vee q \vee r, \neg p \wedge \neg q, \neg p \vee \neg q\}) \end{aligned}$$

Two remarks are in order. First, for both *discrimin* and *leximin* semantics, the exponential blow-up is with respect to the cardinality of the largest priority class G_i , rather than the entire goal base. Thus, if each priority class contains only a few formulae, the translation to *PL* will not be too onerous. Second, we remark that in the case of *leximin*, a polynomial translation into *PL* is possible provided we introduce new propositional variables to define the cardinality formulas $\geq k : S$, cf. (Benhamou, Sais, and Siegel 1994; Liu and Truszczynski 2003).

Discussion

Not only have we shown that some well-known preference languages happen to be fragments of *PL*, but we can define other useful fragments that do not correspond to any existing languages so far. In particular, it seems natural to *combine* different preference languages, which could be captured in our framework but not in the individual preference languages. We can think, for example, of expressing preferences on a complex domain where preference statements in some parts of the domain are of the CP-net kind, whereas some other parts make use of prioritized goals. Apart from the approach in (Brewka, Niemelä, and Truszczynski 2005), we are not aware of such combinations of languages.

Example 7. *Consider an agent who is considering buying some items on the internet. The items she can possibly purchase are the outgoing flight (o), the return flight (r), a hotel night (h), and a book (b). Her preferences about the flight tickets are best modelled in a prioritized goal fashion: better both tickets than none, and better none than just one; tickets being much more expensive than the rest, these preferences override everything else. Now, if she buys the flight tickets, then she wants a hotel night, otherwise she doesn't. Lastly, she wants to buy the book in any case, whether she goes on a trip or not. Her preferences can be expressed by the following preference formula*

$$\begin{aligned} \Phi = & \quad o \wedge r \triangleright \neg o \wedge \neg r \parallel \emptyset \\ \wedge & \quad \neg o \wedge \neg r \triangleright o \leftrightarrow \neg r \parallel \emptyset \\ \wedge & \quad o \wedge r \wedge h \triangleright o \wedge r \wedge \neg h \parallel \emptyset \\ \wedge & \quad \neg(o \wedge r) \wedge \neg h \triangleright \neg(o \wedge r) \wedge h \parallel \{o, r\} \\ \wedge & \quad b \triangleright \neg b \parallel CP \end{aligned}$$

Disjunctions of preference statements are needed for expressing structural properties of the agent's preference relation, such as preferential independence or relative importance. For instance, let $PS = \{a, b, c\}$; we can express that b is preferentially independent of c given a by the preference formula $((a \wedge b \geq a \wedge \neg b \parallel \{c\}) \vee (a \wedge \neg b \geq a \wedge b \parallel \{c\})) \wedge ((\neg a \wedge b \geq \neg a \wedge \neg b \parallel \{c\}) \vee (\neg a \wedge \neg b \geq \neg a \wedge b \parallel \{c\}))$. (Given either a or $\neg a$, whether or not b is preferred to $\neg b$ does not depend on the value of c .) We can express that the propositional variable a is more important than the other propositional variables b and c (in the lexicographic sense) by the preference formula $(a \geq \neg a \parallel \emptyset) \vee (\neg a \geq a \parallel \emptyset)$, implying that if a is preferred to $\neg a$ then any outcome satisfying a is preferred to any outcome satisfying $\neg a$; similarly, if $\neg a$ is preferred to a .

Negation and disjunction also allow implications between preference statements ($\Phi \rightarrow \Psi$ being a shorthand for $\neg\Phi \vee \Psi$). Implication is useful for expressing *domain restrictions*, such as single-peakedness (a key notion in social choice). For instance, when reasoning about the preferences of an agent about curries, $(mild \triangleright hot \parallel CP) \rightarrow (hot \triangleright very-hot \parallel CP)$ (where χ implies that *mild*, *hot* and *very-hot* are mutually exclusive) expresses that if the agent prefers a mild curry to a hot curry *ceteris paribus*, then she prefers a hot curry to a very hot curry, *ceteris paribus*.

Free preference statements allow clear and strong statements to be made about an agent's preferences. To continue with our curry example, an agent who has a low tolerance for spice might express with the free preference statement $mild \triangleright very-hot$ that any meal containing a mild curry is preferred to any meal with very hot curry. Free preferences are obviously related to the $\forall\forall$ interpretation of preferences between formulas (see *e.g.*, (Kaci and van der Torre 2008; van Benthem, Girard, and Roy 2009)).

5. The Conjunctive Fragment of *PL*

We have seen in Section 4 that many AI preference languages can be encoded in the conjunctive fragment PL_C of *PL*. The aim of the current section is to provide a characterization of consistency and dominance for conjunctive preference formulas, which can be seen as a generalization of “flipping sequences” for CP-nets (Boutilier et al. 2004).

We start by introducing some notation. If Φ is a preference statement $\alpha R \beta \parallel F$, where R is either \triangleright or \geq , then Φ^* is defined to consist of all pairs of outcomes (ω, ω') such that $\omega \models \alpha$, $\omega' \models \beta$, and ω and ω' agree on F . It follows from the semantics that the model \succcurlyeq satisfies a strict preference statement Φ if and only if \succcurlyeq contains Φ^* , and that \succcurlyeq satisfies a non-strict preference statement Φ if and only if \succcurlyeq contains Φ^* . For a conjunctive preference formula Γ , we define Γ^* to be the union of Φ^* over all conjuncts Φ of Γ .

The transitive relation $>_\Gamma$ on the set of outcomes is defined as follows: for outcomes ω and ω' , we have $\omega >_\Gamma \omega'$ if and only if there exists a sequence $\omega_1, \dots, \omega_k$ of outcomes with $\omega_1 = \omega$ and $\omega_k = \omega'$ and a sequence $\Phi_1, \dots, \Phi_{k-1}$ of conjuncts of Γ such that for some i , Φ_i is a strict preference statement, and for all $i = 1, \dots, k-1$, $(\omega_i, \omega_{i+1}) \in (\Phi_i)^*$.

The relation \geq_Γ is defined in just the same way, except without the condition that some Φ is strict, so that \geq_Γ is the transitive closure of Γ^* . Clearly, \geq_Γ extends $>_\Gamma$. We also have that if $\omega >_\Gamma \omega' \geq_\Gamma \omega''$ then $\omega >_\Gamma \omega''$, and if $\omega \geq_\Gamma \omega' >_\Gamma \omega''$ then $\omega >_\Gamma \omega''$.

The following lemma, which is used in the proof of Proposition 9 below, shows that satisfaction for conjunctive preference formulas can be expressed in terms of the relations \geq_Γ and $>_\Gamma$.

Lemma 8. *A model \succcurlyeq satisfies a conjunctive preference formula Γ if and only if \succcurlyeq contains \geq_Γ and \succcurlyeq contains $>_\Gamma$.*

The following result characterizes consistency and dominance for a conjunctive preference formula Γ . It shows that the relations \geq_Γ and $>_\Gamma$ represent what is entailed by a conjunctive preference formula Γ .

Proposition 9. *Let Γ be a conjunctive preference formula.*

1. Γ is inconsistent if and only if \succ_{Γ} contains (ω, ω) for some outcome ω .
2. If Γ is consistent and ω, ω' are outcomes, then
 - (a) Γ entails $\omega \triangleright \omega'$ if and only if \succ_{Γ} contains (ω, ω') ;
 - (b) Γ entails $\omega \trianglerighteq \omega'$ if and only if \geq_{Γ} contains (ω, ω') .
3. Let Γ, Φ be consistent preference formulas. Then, $\Gamma \equiv \Phi$ if and only if $\geq_{\Gamma} = \geq_{\Phi}$ and $\succ_{\Gamma} = \succ_{\Phi}$.

Proposition 9 shows for instance that proving $\omega \trianglerighteq \omega'$ comes down to exhibiting a sequence of outcomes from ω to ω' such that each outcome is at least as preferred as its successor according to one of the preference statements.

Example 10. Let $PS = \{a, b\}$ and $\Gamma = (a \trianglerighteq \neg a \parallel \{b\}) \wedge (a \wedge b \triangleright a \wedge \neg b) \wedge (\neg a \wedge \neg b \trianglerighteq \neg a \wedge b)$. \geq_{Γ} is the complete pre-order $ab \succ ab \succ \bar{a}\bar{b} \succ \bar{a}b$, whereas \succ_{Γ} is the strict partial order $\{ab \succ \bar{a}\bar{b}, ab \succ \bar{a}b, ab \succ \bar{a}b\}$. From this we can draw the following inferences: $\Gamma \models a \wedge b \triangleright \neg a \wedge b$, $\Gamma \models \neg a \wedge \neg b \trianglerighteq \neg a \wedge b$, $\Gamma \models a \wedge \neg b \trianglerighteq \neg a \wedge \neg b$; but $\Gamma \not\models \neg a \wedge \neg b \triangleright \neg a \wedge b$, and $\Gamma \not\models a \wedge \neg b \triangleright \neg a \wedge \neg b$.

6. Computational Complexity

In this section, we study the complexity of two fundamental reasoning tasks for preferences, consistency and dominance. The latter is the problem of deciding whether a preference formula entails $\omega \triangleright \omega'$ for a pair of outcomes ω, ω' .

Reduction to Conjunctive Formulas

We begin by showing how reasoning with general or positive preference formulas in PL can be transformed into reasoning about conjunctive preference formulas.

We first remark that we can assume without loss of generality that preference formulas are built from preference statements and negations of preference statements using only conjunction and disjunction. This is because arbitrary preference statements can be put in this form simply by applying the standard (linear-time) negation normal form transformation from propositional logic, treating preference statements as propositional atoms. We assume henceforth that all preference formulas are in this form.

Next, we note that by distributing disjunction over conjunction, we can transform any preference formula into an equivalent disjunction of conjunctions of preference statements and negations of preference statements. We will use $\text{DNF}(\Gamma)$ to refer to the result of applying this procedure to a preference formula Γ . Recall from propositional logic that this transformation may yield an exponential number of disjuncts, but that each disjunct has only polynomial size as it is a conjunction of preference statements (or their negations) appearing in the original formula.

Finally we remark that negations of preference statements can be replaced by equivalent disjunctions of (positive) preference statements. Indeed, it follows from the semantics that $\neg(\omega \triangleright \omega') \equiv \omega' \trianglerighteq \omega$ and $\neg(\omega \trianglerighteq \omega') \equiv \omega' \triangleright \omega$. This means that if Φ is a strict preference statement, then

$$\neg\Phi \equiv \bigvee_{(\omega, \omega') \in \Phi^*} (\omega' \trianglerighteq \omega)$$

and if Φ is non-strict, then

$$\neg\Phi \equiv \bigvee_{(\omega, \omega') \in \Phi^*} (\omega' \triangleright \omega)$$

Putting this all together, we get:

Proposition 11. Let Γ be a preference formula, and let D_{Γ} be the set of disjuncts of $\text{DNF}(\Gamma)$. Create a new set S_{Γ} composed of all formulas obtained from a formula in D_{Γ} by replacing each component strict negative preference statement $\neg\Phi$ by a formula $\omega' \trianglerighteq \omega$ such that $(\omega, \omega') \in \Phi^*$, and each non-strict negative preference statement $\neg\Phi$ by a formula $\omega' \triangleright \omega$ such that $(\omega, \omega') \in \Phi^*$. Note that each formula in S_{Γ} is a conjunctive preference formula.

1. Let Υ be some preference formula. Then, $\Gamma \models \Upsilon$ if and only if for all $\Psi \in S_{\Gamma}$, $\Psi \models \Upsilon$.
2. Γ is consistent if and only if some $\Psi \in S_{\Gamma}$ is consistent.

Complexity of PL

Using the embedding of CP-nets into PL shown earlier, we can obtain a PSPACE lower bound for reasoning in PL , using results in (Goldsmith et al. 2008). For the conjunctive fragment of PL , thanks to our characterization of consistency and dominance (Proposition 9), we can show membership in PSPACE by adopting a similar strategy to that used for general CP-nets. We can then extend the PSPACE membership result to arbitrary preference formulas in PL by leveraging the fact that arbitrary preference formulas can be rewritten as disjunctions of conjunctive preference formulas (Proposition 11). This leads to the following result:

Proposition 12. Consistency and dominance in PL are PSPACE-complete. Hardness holds even for PL_C .

Since we can encode (binary) CP-nets in the logic of von Wright, and we can encode von Wright's logic in our own, we obtain the following corollary. Note that von Wright's own decision procedures require exponential space.

Corollary 13. Consistency and dominance in von Wright's logic of preference are PSPACE-complete.

In a similar manner, we obtain PSPACE-completeness for the logics of (Doyle, Shoham, and Wellman 1991) and (Tan and Pearl 1994).

One might wonder whether it is possible to extend Proposition 12 even further to the full modal preference logic \mathcal{L}_{CP} of van Benthem et al. We show that this is not the case, as consistency becomes NEXPTIME-hard (and it is generally believed that $\text{PSPACE} \neq \text{NEXPTIME}$). The culprit is their conditionalized universal modality $[\Gamma]$, which allows one to specify properties that hold at all worlds which give the same values to the formulas in Γ as a given world. This enables an elegant reduction from the NEXPTIME-complete $2^n \times 2^n$ tiling problem to consistency in \mathcal{L}_{CP} .

Proposition 14. Consistency in \mathcal{L}_{CP} is NEXPTIME-hard.

Proof sketch. We recall that an instance of the $2^n \times 2^n$ bounding tiling problem is given by a triple (\mathcal{T}, H, V) with \mathcal{T} a non-empty, finite set of tile types including an initial tile

t_1 to be placed on the lower left corner, $H \subseteq \mathcal{T} \times \mathcal{T}$ a horizontal matching relation, and $V \subseteq \mathcal{T} \times \mathcal{T}$ a vertical matching relation. A tiling for (\mathcal{T}, H, V) is a map $f: \{0, \dots, 2^n - 1\} \times \{0, \dots, 2^n - 1\} \rightarrow \mathcal{T}$ such that $f(0, 0) = t_1$, $(f(i, j), f(i + 1, j)) \in H$ for all $i < 2^n - 1$, and $(f(i, j), f(i, j + 1)) \in V$ for all $j < 2^n - 1$. It is NEXPTIME-complete to decide whether a tiling problem admits a tiling.

Given a tiling problem (\mathcal{T}, H, V) , we construct a formula $\Phi \in \mathcal{L}_{CP}$ such that Φ is consistent if and only if there is there exists a tiling for (\mathcal{T}, H, V) . The formula Φ will utilize the variables $x_1, \dots, x_n, y_1, \dots, y_n, t_1, \dots, t_k$ (where $\mathcal{T} = \{t_1, \dots, t_k\}$). The x_i and y_i will encode the binary representation of horizontal and vertical coordinates, and the t_i represent the tile types. Our encoding will utilise \mathcal{L}_{CP} -formulas of the form $[\Gamma]\varphi$. Such a formula is verified at a world w if φ is satisfied by every world v with $w \sim_{\Gamma} v$. The modality $\langle \Gamma \rangle$ is the dual modality to $[\Gamma]$, so $\langle \Gamma \rangle \varphi$ holds at w if there is some v with $w \sim_{\Gamma} v$ that verifies φ . The modality \square is an abbreviation for $[\emptyset]$ and corresponds to the standard universal modality. We define Φ as the conjunction of the following formulas, which we annotate with their purpose:

- Exactly one tile per world.

$$\square(t_1 \vee \dots \vee t_k) \wedge \bigwedge_{i=1}^k \square(t_i \rightarrow \bigwedge_{j \neq i} \neg t_j)$$
- Worlds with the same coordinates have the same tile.

$$\bigwedge_{i=1}^k \square(t_i \rightarrow [x_1, \dots, x_n, y_1, \dots, y_n] t_i)$$
- Specification of the lower left corner.

$$\neg x_1 \wedge \dots \wedge \neg x_n \wedge \neg y_1 \wedge \dots \wedge \neg y_n \wedge t_1$$
- Horizontal grid successors exist.

$$\bigwedge_{m=1}^n \square((x_1 \wedge \dots \wedge x_{m-1} \wedge \neg x_m) \rightarrow \langle x_{m+1}, \dots, x_n, y_1, \dots, y_n \rangle (\neg x_1 \wedge \dots \wedge \neg x_{m-1} \wedge x_m))$$
- Vertical grid successors exist.

$$\bigwedge_{m=1}^n \square((y_1 \wedge \dots \wedge y_{m-1} \wedge \neg y_m) \rightarrow \langle y_{m+1}, \dots, y_n, x_1, \dots, x_n \rangle (\neg y_1 \wedge \dots \wedge \neg y_{m-1} \wedge y_m))$$
- Horizontal matching relation is satisfied.

$$\bigwedge_{i=1}^k \bigwedge_{m=1}^n \square((x_1 \wedge \dots \wedge x_{m-1} \wedge \neg x_m \wedge t_i) \rightarrow [x_{m+1}, \dots, x_n, y_1, \dots, y_n] ((\neg x_1 \wedge \dots \wedge \neg x_{m-1} \wedge x_m) \rightarrow \bigvee_{t_j: (t_i, t_j) \in H} t_j))$$
- Vertical matching relation is satisfied.

$$\bigwedge_{i=1}^k \bigwedge_{m=1}^n \square((y_1 \wedge \dots \wedge y_{m-1} \wedge \neg y_m \wedge t_i) \rightarrow [y_{m+1}, \dots, y_n, x_1, \dots, x_n] ((\neg y_1 \wedge \dots \wedge \neg y_{m-1} \wedge y_m) \rightarrow \bigvee_{t_j: (t_i, t_j) \in V} t_j))$$

For a proof that the formula Φ has the desired properties, refer to the long version of the paper. \square

Complexity of Free Preferences

We know from CP-nets that reasoning is PSPACE-hard for the conjunctive fragment even under the restriction that component preference statements Φ are such that F_{Φ} contains only variables and α_{Φ} and β_{Φ} are cubes, i.e., conjunctions of literals. Thus, in order to isolate fragments of PL with lower complexity, we consider the case when F is empty.

Our first result provides a characterization of consistency and dominance for conjunctions of free preference statements. It makes reference to the graph G_{Γ} associated with a conjunctive preference formula Γ which is defined as follows: the vertices of G_{Γ} are the component preference statements of Γ , and there is an edge from Φ to Ψ whenever $\beta_{\Phi} \wedge \alpha_{\Psi}$ is compatible with the integrity constraint χ .

Lemma 15. *Let Γ be a conjunction of a set S of free preference statements, and let θ and ρ be outcomes. Then*

1. \succ_{Γ} contains (θ, ρ) if and only if there exists $\Phi, \Psi \in S$ with $\theta \models \alpha_{\Phi}$, and $\rho \models \beta_{\Psi}$, and there exists a directed path from Φ to Ψ in G_{Γ} which includes at least one strict preference statement.
2. \succeq_{Γ} contains (θ, ρ) if and only if there exists $\Phi, \Psi \in S$ with $\theta \models \alpha_{\Phi}$, and $\rho \models \beta_{\Psi}$, and there exists a directed path from Φ to Ψ in G_{Γ} .

Using Lemma 15 and Proposition 9, we obtain the following characterization of consistency and dominance:

Proposition 16. *Suppose that Γ is a conjunction of a set S of free preference statements. Then the following hold:*

1. Γ is inconsistent if and only if G_{Γ} contains a cycle involving at least one strict preference statement.
2. Let ω and ω' be outcomes. If Γ is consistent, then
 - (a) Γ entails $\omega \triangleright \omega'$ if and only if there exists $\Phi, \Psi \in S$ with $\omega \models \alpha_{\Phi}$, and $\omega' \models \beta_{\Psi}$, and there exists a directed path from Φ to Ψ in G_{Γ} involving at least one strict preference statement.
 - (b) Γ entails $\omega \triangleright \omega'$ if and only if there exists $\Phi, \Psi \in S$ with $\omega \models \alpha_{\Phi}$, and $\omega' \models \beta_{\Psi}$, and there exists a directed path from Φ to Ψ in G_{Γ} .

Example 17. *Let $\chi = \top$ and $\Gamma = \Phi_1 \wedge \Phi_2 \wedge \Phi_3$ where $\Phi_1 = (a \wedge \neg c \triangleright b \parallel \emptyset)$, $\Phi_2 = (\neg a \wedge b \triangleright \neg a \wedge \neg b \parallel \emptyset)$, $\Phi_3 = (\neg b \triangleright a \wedge b \wedge c \parallel \emptyset)$. The graph G_{Γ} contains precisely the edges (Φ_1, Φ_1) , (Φ_1, Φ_2) and (Φ_2, Φ_3) . Hence, by Proposition 16, Γ is consistent. We also obtain that Γ entails $abc \triangleright abc$, because $abc \models \alpha_{\Phi_1}$, $abc \models \beta_{\Phi_3}$, there is a path from Φ_1 to Φ_3 in G_{Γ} , and at least one of the preference statements involved in the path is strict (i.e., Φ_2).*

Proposition 16 implies that determining consistency of a conjunction Γ of free preference statements – as well as checking dominance for pairs of outcomes w.r.t. Γ – is polynomial whenever checking the compatibility of $\beta_{\Phi} \wedge \alpha_{\Psi}$ with the integrity constraint χ is polynomial. This holds, in particular, if $\chi = \top$ and all the α 's and β 's are cubes.

Proposition 18. *When $\chi = \top$, dominance and consistency are both tractable for conjunctions of preference statements of the form $\alpha \triangleright \beta$ where α, β are both cubes.*

In the more general case in which the preference statements may contain arbitrary formulas the complexity jumps to co-NP for consistency (and NP for dominance).

Proposition 19. *Consistency (resp. dominance) is co-NP-complete (resp. NP-complete) for conjunctions of free preference statements.*

Proof. We consider only consistency, but dominance is handled similarly. Co-NP-hardness is straightforward: a propositional CNF φ is unsatisfiable if and only if the preference statement $\varphi \triangleright \varphi \parallel \emptyset$ is consistent. For membership in co-NP, we have, by Proposition 16, that the formula Γ is inconsistent if and only if G_Γ is cyclic. To show cyclicity of G_Γ , we guess a sequence of (distinct) preference statements Ψ_1, \dots, Ψ_m from Γ , together with a sequence of valuations v_1, \dots, v_m . We then verify in polynomial time that $v_m \models \beta_{\Psi_m} \wedge \alpha_{\Psi_1}$ and for each $1 \leq i \leq m-1$, we have $v_i \models \beta_{\Psi_i} \wedge \alpha_{\Psi_{i+1}}$. If this is the case, there must be edges between Ψ_i and Ψ_{i+1} ($1 \leq i \leq m-1$) and between Ψ_m and Ψ_1 , so G_Γ contains a cycle. Moreover, if G_Γ is cyclic, we can clearly guess a sequence of preference statements and valuations satisfying the desired conditions. This yields membership in NP for cyclicity of G_Γ , and hence a co-NP upper bound for consistency of Γ . \square

For boolean combinations of free preference statements where the α_Φ, β_Φ are cubes, the complexities of the two problems are swapped: it is consistency which is NP-complete, and dominance which is co-NP-complete.

Proposition 20. *Suppose $\chi = \top$. Consistency (resp. dominance) is NP-complete (resp. co-NP-complete) for preference formulas built from free preference statements Φ such that α_Φ and β_Φ are conjunctions of literals.*

Proof. We consider only consistency; dominance is treated similarly. For membership in NP, we know from Lemma 11 that Γ is consistent if and only if some $\Psi \in S_\Gamma$ is consistent. The idea then is to guess some $\Psi \in S_\Gamma$ and to verify its consistency. It follows from the definition of S_Γ that Ψ is a (polysize) conjunction of free preference statements Φ such that α_Φ and β_Φ are conjunctions of literals. It follows then from Proposition 18 that consistency of Ψ can be decided in polytime.

For hardness, consider an instance of SAT on variables x_1, \dots, x_n , represented as a conjunction of disjunction of literals. For $j = 1 \dots, 2n$, let θ_j be different outcomes. Let P_j be the preference statement $\theta_{2j-1} \triangleright \theta_{2j}$, and let N_j be the preference statement $\theta_{2j} \triangleright \theta_{2j-1}$. Generate a preference formula Γ by replacing each positive literal x_j by P_j , and each negative literal $\neg x_j$ by N_j . It can be seen that Γ is consistent if and only if the SAT instance is satisfiable. \square

For the case where arbitrary formulas may appear in the preference statements, we give only upper bounds.

Proposition 21. *Consistency is in Σ_2^P for free preference formulas, whereas dominance belongs to Π_2^P .*

Proof Sketch. We combine the ideas from Propositions 19 and 20. To decide consistency of a free preference formula Γ , we first guess some $\Psi \in S_\Gamma$, and then we use an NP-oracle to verify the consistency of Ψ . \square

We conjecture that these problems are in fact Σ_2^P - and Π_2^P -complete.

Note that Propositions 18, 19, 20, and 21 can be used to derive corresponding results for the bounded fragment

of PL in which $|F|$ is bounded by some constant k . To do so, we simply translate a formula $\varphi \triangleright \Psi \parallel \{\chi_1 \dots, \chi_k\}$ into the conjunction of all of the free preference statements $\mu \wedge \varphi \triangleright \mu \wedge \Psi$ such that $\mu = \lambda_1 \wedge \dots \wedge \lambda_k$ where $\lambda_i \in \{\chi_i, \neg \chi_i\}$ for every $1 \leq i \leq k$.

Overall, the above results show that the complexity of reasoning with arbitrary free preference formulas is much lower than for conjunctions of even very simple preference statements without restrictions on the number of fixed formulas.

7. Expressiveness and Succinctness

Given two languages L_1 and L_2 , we recall from (Cadoli et al. 2000) that L_1 is *at least as succinct as* L_2 if there exists a polysize, equivalence-preserving translation f from L_2 to L_1 , i.e., (i) for any $\Phi \in L_2$, $f(\Phi) \in L_1$ is equivalent to Φ , and (ii) there exists a polynomial p s.t. for any $\Phi \in L_2$, $|f(\Phi)| \leq p(|\Phi|)$. Moreover, L_1 is *at least as expressive as* L_2 if there exists an equivalence-preserving translation f from L_2 to L_1 .

We have shown earlier that the positive fragment of PL is just as expressive as the whole language. However, excluding negation incurs a potential exponential blowup in formula size, as the following result attests.

Proposition 22. *Given $\chi = \top$, the positive fragment of PL is strictly less succinct than PL .*

Proof Sketch. Let Φ be the preference formula $\neg \Gamma$, where Γ is a preference statement $x_1 \triangleright \neg x_1 \parallel \emptyset$ and x_1 is a propositional variable. Φ is equivalent to a disjunction Φ' of the exponentially many basic preference statements of the form $\theta \triangleright \tau$ for all θ extending $\neg x_1$ and all τ extending x_1 . It can be shown (with a bit of work) that there is no more compact representation of Φ as a positive preference formula. \square

Moving from preference statements to conjunctive preference formulas, and also on to positive preference formulas, increases the expressivity of the language:

Proposition 23. *The conjunctive fragment of PL is strictly less expressive than the positive fragment of PL .*

Proof Sketch. Consider the positive preference formula $\Psi = (\theta_1 \triangleright \theta_2) \vee (\theta_2 \triangleright \theta_3)$, where θ_1, θ_2 and θ_3 are three different outcomes. It can be shown that Ψ does not entail any non-tautologous conjunctive preference formula, hence there can be no conjunctive preference formula equivalent to Ψ . \square

Proposition 24. *The preference statement fragment of PL is strictly less expressive than the conjunctive fragment of PL .*

Proof Sketch. Let Ψ be the conjunctive preference statement $(\theta_1 \triangleright \theta_2) \wedge (\theta_2 \triangleright \theta_3)$, where θ_1, θ_2 and θ_3 are different outcomes. It can be shown that no consistent preference statement implies Ψ , which means there exists no preference statement equivalent to Ψ , as required. \square

We now consider sublanguages of the conjunctive fragment PL_C . For the sake of simplicity, we consider only strict preference statements (results can easily be extended to the general case but are a little bit tedious to state). Let PL_C^S be the resulting language.

A first way of obtaining interesting fragments of PL_C^s consists in imposing syntactic restrictions on α_Φ and β_Φ . Given a subset \mathcal{T} of L , the fragment $PL_{C,\mathcal{T}}^s$ is the set of all conjunctions of preference statements of the form $\alpha \triangleright \beta \parallel F$, where $\alpha \in \mathcal{T}$ and $\beta \in \mathcal{T}$ (we do not restrict F). Let Lit , Cl , and Cb be respectively the sets of literals, clauses, and cubes built from PS .

For every consistent Φ and Ψ in PL_C^s , using Proposition 9, we have that Φ and Ψ are equivalent if and only if $\triangleright_\Phi = \triangleright_\Psi$. For a sublanguage L of PL_C^s , let $Exp(L) = \{\triangleright_\Phi \mid \Phi \text{ is consistent and } \Phi \in L\}$. We have that, given two sublanguages L_1 and L_2 of PL_C^s , L_1 and L_2 are equally expressive if $Exp(L_1) = Exp(L_2)$ and L_1 is strictly less expressive than L_2 if $Exp(L_1) \subsetneq Exp(L_2)$.

Proposition 25 (expressiveness inside PL_C^s).

- $PL_{C,Cb}^s$ is equally as expressive as PL_C^s .
- If $\chi = \top$, then $PL_{C,Cb}^s$ is strictly more expressive than $PL_{C,Lit}^s$.
- $PL_{C,Cl}^s$ is equally as expressive as $PL_{C,Lit}^s$.

Proof of Point 2. We aim to show that $Exp(PL_{C,Lit}^s) \subsetneq Exp(PL_{C,Cb}^s)$. Let $PS = \{a, b\}$, and consider the following free preference statement from $PL_{C,Cb}^s$:

$$\Phi = a \wedge b \triangleright \neg a \wedge \neg b$$

Now suppose there exists a formula Ψ in $PL_{C,Lit}^s$ which is equivalent to Φ . Then by Proposition 9, we have that $\triangleright_\Psi = \triangleright_\Phi$. As $ab \triangleright_\Phi \bar{a}\bar{b}$, we must also have $ab \triangleright_\Psi \bar{a}\bar{b}$. By Proposition 9, there must exist a sequence of outcomes $\omega_1, \dots, \omega_n$ such that $\omega_1 = ab$, $\omega_n = \bar{a}\bar{b}$, and for every $i < n$, (ω_i, ω_{i+1}) is sanctioned by a preference statement in Ψ . Note however that if $n > 2$, then we have $\omega_1 \triangleright_\Psi \omega_{n-1} \triangleright_\Psi \omega_n$, which cannot be since \triangleright_Φ (hence \triangleright_Ψ) contains a single tuple, namely $(ab, \bar{a}\bar{b})$. Thus, it must be the case that $n = 2$, and there is a single preference statement $\ell_1 \triangleright \ell_2 \parallel F$ from $PL_{C,Lit}^s$ which sanctions $ab \triangleright_\Psi \bar{a}\bar{b}$. It follows that the literal ℓ_1 is either a or b , and ℓ_2 is either $\neg a$ or $\neg b$; we consider only the case where $\ell_1 = a$ and $\ell_2 = \neg a$, but the other three cases proceed similarly. Now we also know that $ab \approx_F \bar{a}\bar{b}$. But any formula on PS which is interpreted identically by ab and $\bar{a}\bar{b}$ must be equivalent to \top , \perp , or $a \leftrightarrow b$. It follows that F only contains formulae of these three types. But that means that the preference statement sanctions $\bar{a}\bar{b} \triangleright_\Psi ab$ since $\bar{a}\bar{b} \models \ell_1$, $\bar{a}\bar{b} \models \ell_2$, and $\bar{a}\bar{b}$ and ab agree on the interpretation of \top , \perp , and $a \leftrightarrow b$. This contradicts our assumption that $\triangleright_\Psi = \triangleright_\Phi$. \square

Writing $L <_e L'$ for L strictly less expressive than L' and $L \sim_e L'$ for L equally as expressive as L' , Proposition 25 yields:

$$PL_{C,Lit}^s \sim_e PL_{C,Cl}^s <_e PL_{C,Cb}^s \sim_e PL_C^s$$

The reason why we need the condition $\chi = \top$ in some of the results above is that otherwise it is possible to restrict the set of possible outcomes in a drastic way, so that both

languages are equally expressive or equally succinct (for instance, in the extreme case where χ has a unique model, then all fragments of PL_C are equally expressive).

We now show that using clauses rather than literals does not result in a more succinct language. We conjecture that $PL_{C,Cb}^s$ is strictly less succinct than PL_C^s , but so far we don't have a proof.

Proposition 26. $PL_{C,Cl}^s$ and $PL_{C,Lit}^s$ are equally succinct.

A second way of obtaining fragments of PL_C consists in restricting the allowed sets of fixed formulas F_Φ . We consider three options: restrictions to propositional symbols ($PL_{C,FPS}^s$), to free preference statements ($PL_{C,FREE}^s$), and to *ceteris paribus* preference statements ($PL_{C,CP}^s$). Note that $PL_{C,FREE}^s, PL_{C,CP}^s \subsetneq PL_{C,FPS}^s \subsetneq PL_C^s$.

We first note that the encoding of a preference relation in $PL_{C,Cb}^s$ (by explicitly comparing maximal cubes) belongs both to $PL_{C,FREE}^s$ and to $PL_{C,CP}^s$, which shows that these fragments are fully expressive.

Regarding succinctness, we have the following:

Proposition 27.

- $PL_{C,FREE}^s$ is strictly less succinct than $PL_{C,FPS}^s$;
- $PL_{C,FPS}^s$ is strictly less succinct than PL_C^s .

Proof of Point 1. Let $\Phi = (x_1 \triangleright \neg x_1 \parallel \{x_2, \dots, x_n\})$. We have that $\omega \triangleright_\Phi \omega'$ if and only if $\omega = x_1 \wedge \gamma$ and $\omega' = \neg x_1 \wedge \gamma$ for some maximal consistent cube γ built on $\{x_2, \dots, x_n\}$. Let $\Psi \in PL_{C,FREE}^s$ such that $\triangleright_\Psi = \triangleright_\Phi$. From Proposition 16, $\omega \triangleright_\Psi \omega'$ if and only if there exists a sequence of preference statements $\phi_1 \triangleright \psi_1, \phi_2 \triangleright \psi_2, \dots, \phi_q \triangleright \psi_q$ in Ψ such that for every $i < q$, $\psi_i \wedge \phi_{i+1}$ is satisfiable, and $\omega \models \phi_1$, $\omega' \models \psi_q$. Suppose $q \geq 2$; then $\psi_1 \wedge \phi_2$ is satisfiable, and for every $\omega'' \models \psi_1 \wedge \phi_2$ we have $\omega \triangleright_\Psi \omega'' \triangleright_\Psi \omega'$. This contradicts $\triangleright_\Psi = \triangleright_\Phi$, because there is no triple $(\omega, \omega', \omega'')$ in \triangleright_Φ such that $\omega \triangleright_\Phi \omega'' \triangleright_\Phi \omega'$. Therefore, $q = 1$, and this is true for every maximal consistent γ built on $\{x_2, \dots, x_n\}$. Thus for every such γ there must exist a preference statement $(\phi_\gamma \triangleright \psi_\gamma \parallel \emptyset)$ in Ψ such that $x_1 \wedge \gamma \models \phi_\gamma$ and $\neg x_1 \wedge \gamma \models \psi_\gamma$. Now, assume that $(\phi_\gamma, \psi_\gamma) = (\phi_{\gamma'}, \psi_{\gamma'})$ holds for some γ and $\gamma' \neq \gamma$. Then we have $x_1 \wedge \gamma \models \phi_\gamma$ and $\neg x_1 \wedge \gamma' \models \psi_\gamma$, therefore $x_1 \wedge \gamma \triangleright_\Psi \neg x_1 \wedge \gamma'$, which contradicts $\triangleright_\Psi = \triangleright_\Phi$. Therefore, we have a distinct preference statement $\phi_\gamma \triangleright \psi_\gamma \parallel \emptyset$ in Ψ for every γ , which implies that Ψ contains at least 2^{n-1} preference statements. \square

8. Conclusion

The contributions of this paper are twofold. First, from an AI point of view, we provide a more expressive language for preference representation, whose worst-case complexity is not worse than the worst-case complexity of CP-nets. The results we obtained concerning the different fragments of our logic go some way to generating a map for choosing the right trade-off between expressiveness, complexity, and succinctness. Second, the computational aspects of preference logics had almost never been dealt with before, and our

work include significant results for some important preference logics. This paper contributes to bringing together two research areas which have so far had little interaction.

Further work includes a similar study for other fragments of the logic of van Benthem et al., in particular when the interpretation of preference statements uses a different alternation of quantifiers than $\forall\forall$. Also, we intend to explore the connections between our results and complexity, expressiveness and succinctness results for logics for cardinal preference representation (e.g., (Uckelman et al. 2009)). Finally, an interesting topic for further research would consist in using our preference logic — and preference logics in general — on dynamic environments including actions, where preferences over states could be lifted to preferences over actions or courses of actions.

Acknowledgements

The work of the third author is supported by the Science Foundation Ireland under Grant No. 08/PI/I1912.

References

- Benferhat, S.; Cayrol, C.; Dubois, D.; Lang, J.; and Prade, H. 1993. Inconsistency management and prioritized syntax-based entailment. In *Proc. of IJCAI*, 640–645.
- Benhamou, B.; Sais, L.; and Siegel, P. 1994. Two proof procedures for a cardinality based language in propositional calculus. In *Proc. of STACS*, 71–82.
- Boutilier, C.; Brafman, R.; Domshlak, C.; Hoos, H.; and Poole, D. 2004. CP-nets: a tool for representing and reasoning with conditional ceteris paribus statements. *Journal of Artificial Intelligence Research* 21:135–191.
- Boutilier, C. 1994. Toward a logic for qualitative decision theory. In *Proc. of KR*, 75–86.
- Bouveret, S.; Endriss, U.; and Lang, J. 2009. Conditional importance networks: A graphical language for representing ordinal, monotonic preferences over sets of goods. In *Proc. of IJCAI*, 67–72.
- Brafman, R.; Domshlak, C.; and Shimony, S. E. 2006. On graphical modeling of preference and importance. *Journal of Artificial Intelligence Research* 25:389–424.
- Brewka, G.; Niemelä, I.; and Truszczyński, M. 2005. Prioritized component systems. In *Proc. of AAI*, 596–601.
- Brewka, G. 2004. A rank-based description language for qualitative preferences. In *Proc. of ECAI*, 303–307.
- Cadoli, M.; Donini, F. M.; Liberatore, P.; and Schaerf, M. 2000. Space efficiency of propositional knowledge representation formalisms. *Journal of Artificial Intelligence Research* 13:1–31.
- Catañeda, H. 1958. Review of Sören Halldén, “On the logic of ‘better’”. *Philosophy and Phenomenological Research* 19:266.
- Coste-Marquis, S.; Lang, J.; Liberatore, P.; and Marquis, P. 2004. Expressive power and succinctness of propositional languages for preference representation. In *Proc. of KR*, 203–212.
- de Jongh, D., and Liu, F. 2009. Preference, priorities and belief. In *Preference Change: Approaches from Philosophy, Economics and Psychology*. Theory and Decision Library. 85–108.
- Doyle, J., and Wellman, M. P. 1994. Representing preferences as ceteris paribus comparatives. In *Working Notes of the AAI Symposium on Decision-Theoretic Planning*.
- Doyle, J.; Shoham, Y.; and Wellman, M. P. 1991. A logic of relative desire. In *Proc. of ISMIS*, 16–31.
- Girard, P. 2008. *Modal logic for belief and preference change*. Ph.D. Dissertation, Stanford University.
- Goldsmith, J.; Lang, J.; Truszczyński, M.; and Wilson, N. 2008. The computational complexity of dominance and consistency in CP-nets. *Journal of Artificial Intelligence Research* 33:403–432.
- Halldén, S. 1957. *On the logic of ‘better’*. Library of Theoria, Lund.
- Hansson, S. O. 1989. A new semantical approach to the logic of preferences. *Erkenntnis* 31:1–42.
- Hansson, S. O. 2001. *The structure of values and norms*. Cambridge University Press.
- Kaci, S., and van der Torre, L. 2008. Reasoning with various kinds of preferences: logic, non-monotonicity, and algorithms. *Annals of OR* 163(1):89–114.
- Lang, J.; van der Torre, L.; and Weydert, E. 2003. Hidden uncertainty in the logical representation of desires. In *Proc. of IJCAI*, 685–690.
- Liu, L., and Truszczyński, M. 2003. Local-search techniques for propositional logic extended with cardinality constraints. In *Proc. of CP*, 495–509.
- Liu, F. 2008. *Changing for the Better. Preference Dynamics and Agent Diversity*. Ph.D. Dissertation, ILLC, University of Amsterdam.
- Tan, S., and Pearl, J. 1994. Specification and evaluation of preferences for planning under uncertainty. In *Proc. of KR*, 530–539.
- Uckelman, J.; Chevaleyre, Y.; Endriss, U.; and Lang, J. 2009. Representing utility functions via weighted goals. *Mathematical Logic Quarterly* 55(4):341–361.
- van Benthem, J., and Liu, F. 2007. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logic* 17(2).
- van Benthem, J.; Girard, P.; and Roy, O. 2009. Everything else being equal: A modal logic for ceteris paribus preferences. *Journal of Philosophical Logic* 38(1):83–125.
- von Wright, G. H. 1963. *The logic of preference*. Edinburgh University Press.
- Wilson, N. 2004a. Consistency and constrained optimisation for conditional preferences. In *Proc. of ECAI*, 888–892.
- Wilson, N. 2004b. Extending CP-nets with stronger conditional preference statements. In *Proc. of AAI*, 735–741.
- Wilson, N. 2009. Efficient inference for expressive comparative preference languages. In *Proc. of IJCAI*, 961–966.