CP-nets, π-pref nets, and Pareto dominance

Nic Wilson[1], Didier Dubois[2], and Henri Prade[2]

[1] Insight Centre for Data Analytics
School of Computer Science and IT, University College Cork, Ireland
[2] IRIT-CNRS Université Paul Sabatier, Toulouse, France

Abstract. Two approaches have been proposed for the graphical handling of qualitative conditional preferences between solutions described in terms of a finite set of features: Conditional Preference networks (CP-nets for short) and more recently, Possibilistic Preference networks (π-pref nets for short). The latter agree with Pareto dominance, in the sense that if a solution violates a subset of preferences violated by another one, the former solution is preferred to the latter one. Although such an agreement might be considered as a basic requirement, it was only conjectured to hold as well for CP-nets. This non-trivial result is established in the paper. Moreover it has important consequences for showing that π-pref nets can at least approximately mimic CP-nets by adding explicit constraints between symbolic weights encoding the *ceteris paribus* preferences, in case of Boolean features. We further show that dominance with respect to the extended π-pref nets is polynomial.

1 Introduction

Ceteris Paribus Conditional Preference Networks (CP-nets, for short) [5, 6] were introduced in order to provide a convenient tool for the elicitation of multidimensional preferences and accordingly compare the relative merits of solutions to a problem. They are based on three assumptions: only ordinal information is required; the preference statements deal with the values of single decision variables in the context of fixed values for other variables that influence them; preferences are provided all else being equal (*ceteris paribus*). CP-nets were inspired by Bayesian networks (they use a dependency graph, most of the time a directed acyclic one, whose vertices are variables) but differ from them by being qualitative, by their use of the ceteris paribus assumption, and by the fact that the variables in a CP-net are decision variables rather than random variables. In the most common form of CP-nets, each preference statement in the preference graph translates into a strict preference between two solutions (i.e., value assignment to all decision variables) differing on a single variable (referred to as a worsening flip) and the dominance relation between solutions is the transitive closure of this worsening flip relation.

Another kind of conditional preference network, called π-pref nets, has been more recently introduced [1], and is directly inspired by the counterpart of Bayesian networks in possibility theory, called possibilistic networks [3]. A Π-pref net shares with CP-nets its directed acyclic graphical structure between decision variables, and conditional preference statements attached to each variable in the contexts defined by assignments of its parent variables in the graph. The preference for one value against another is captured by assigning degrees of possibility (here interpreted as utilities) to these values. When the only existing preferences are those expressed by the conditional statements (there are no preference statements across contexts or variables), it has been shown that the dominance relation between solutions is obtained by comparing vectors of symbolic utility values (one per variables) using Pareto-dominance.

Some results comparing the preference relations between solutions obtained from CP-nets and π-pref nets with Boolean decision variables are given in [1]. This is made easy by the fact that CP-nets and π-pref nets share the graph structure and the conditional preference tables. It was shown that the two obtained dominance relations between solutions cannot conflict with each other (there is no preference reversal between them), and that ceteris paribus information can be added to π-pref nets in the form of preference statements between specific products of symbolic weights. One pending question was to show that the dominance relation between solutions obtained from a CP-net refines the preference relation obtained from the corresponding π-pref net. In the case of Boolean variables, the π-pref net ordering can be viewed as a form of Pareto ordering: each assignment of a decision variables is either good (= in agreement with the preference statement) or bad. The pending question comes down to prove a monotonicity condition for the preference relation on solutions, stating that as soon as a solution contains more (in the sense of inclusion) good variable assignments than another solution, it should be strictly preferred by the CP-net. Strangely enough this natural question has hardly been addressed in the literature so far (see [2] for some discussion). The aim of this paper is to solve this problem, and more generally to compare the orderings of solutions using the two preference modeling structures.

We further show that dominance with respect to extended π-pref nets can be computed in polynomial time, using linear programming; it thus forms a polynomial upper approximation for the CP-net dominance relation.

The paper is structured as follows: In Section 2 we define a condition, that we call *local dominance*, that is shown to be a sufficient condition for dominance in a CP-net. The follow two sections, Sections 3 and 4, make use of this sufficient condition in producing results that show that a form of Pareto ordering is a lower bound for a lower bound for CP-net dominance. Section 5 then uses the results of Section 4 to show that π-pref nets dominance is a lower bound for CP-net dominance. We also show there that the extended π-pref nets dominance, which is an upper bound for CP-net dominance, can be computed in polynomial time. Section 6 concludes.

2 A Sufficient Condition for Dominance in a CP-Net

We start by recalling the definition of CP-nets and a characterization of the corresponding dominance relation between solutions.

2.1 Defining CP-nets

We consider a finite set of variables $\mathcal{V}$. Each variable $X \in \mathcal{V}$ has an associated finite domain $\mathrm{Dom}(X)$. An outcome (also called a solution) is a complete assignment to the variables in $\mathcal{V}$, i.e., a function w that, for each variable $X \in \mathcal{V}$, $w(X) \in \mathrm{Dom}(X)$.

A CP-net Σ over set of variables $\mathcal{V}$ is a pair $\langle G, P\rangle$. The first component G is a directed graph with vertices in $\mathcal{V}$, and we say that CP-net Σ is *acyclic* if G is acyclic. For variable $X \in \mathcal{V}$, let $\mathcal{U}_X$ be the set of parents of X in G, i.e., the set of variables Y such that (Y, X) is an edge in G. The second component P of Σ consists of a collection of partial orders $\{>^X_u : X \in \mathcal{V}, u \in \mathrm{Dom}(\mathcal{U}_X)\}$, called conditional preference tables; for each variable $X \in \mathcal{V}$ and each assignment u to the parents $\mathcal{U}_X$ of X, relation $>^X_u$ is a strict partial order (i.e., a transitive and irreflexive relation) on $\mathrm{Dom}(X)$. We make the assumption that for each variable X there exists at least one assignment u to $\mathcal{U}_X$ such that $>^X_u$ is non-empty (i.e., for each $X \in \mathcal{V}$ there exists some $x, x' \in \mathrm{Dom}(X)$ and some u such that $x >^X_u x'$).

Let w be an outcome and, for variable $X \in \mathcal{V}$, let $u = w(\mathcal{U}_X)$ be the projection of w to the parents set of X. If $x >^X_u x'$ then we shall write, for simplicity, (with the understanding that x and x' are elements of $\mathrm{Dom}(X)$):

$$x > x' \text{ given } w \text{ [with respect to } \Sigma\text{]}.$$

Note that if v is any outcome whose projection to the parents set of X is also u then [$x > x'$ given v] if and only if [$x > x'$ given w]; the values of $w(Y)$ and $u(Y)$ may differ for variables $Y \notin \mathcal{U}_X \cup \{X\}$, but the preference between x and x' in the context u does not depend on Y.

We say that Σ is *locally totally ordered* if each associated strict partial order $>^X_u$ is a strict total order, so that for each pair of different elements x and x' of $\mathrm{Dom}(X)$, we have either $x >^X_u x'$ or $x' >^X_u x$. We say that Σ is *Boolean* if for each $X \in \mathcal{V}$, each domain has exactly two elements: $|\mathrm{Dom}(X)| = 2$.[3]

The dominance relation associated with a CP-net Given a CP-net Σ over variables $\mathcal{V}$, we say that w' is a *worsening flip from w w.r.t. Σ*, if w' and w are outcomes that differ on exactly one variable $X \in \mathcal{V}$ (so that $w'(X) \neq w(X)$ and for all $Y \in \mathcal{V} \setminus \{X\}$, $w'(Y) = w(Y)$), and $w(X) >^X_u w'(X)$, where u is the projection of w (or w') to the parent set $\mathcal{U}_X$ of X.

The set of direct consequences of CP-net Σ are the set of pairs (w, w'), where w' is a worsening flip from w w.r.t. Σ, forming an irreflexive relation:

Definition 1. *The worsening flip relation $>^\Sigma_{wf}$ is defined by $w >^\Sigma_{wf} w'$ if and only if $w(X) >^X_u w'(X)$ where $u = w(\mathcal{U}_X) = w'(\mathcal{U}_X)$. Let the binary relation $>^\Sigma_{\mathrm{cp}}$ on outcomes denote the transitive closure of $>^\Sigma_{wf}$. If $w >^\Sigma_{\mathrm{cp}} w'$ we say that w* [cp-]*dominates w'* [*with respect to Σ*].

[3] If a variable has only one element in its domain, it is a constant, and we could remove it if we wished.

The relation $>^{\Sigma}_{wf}$ is well-defined due to the ceteris paribus assumption. A sequence of outcomes $w_1, \ldots, w_k$ is said to be a worsening flipping sequence [with respect to CP-net Σ] from w_1 to w_k if, for each $i = 1, \ldots, k-1$, w_{i+1} is a worsening flip from w_i. Thus, w cp-dominates w' if and only if there is a worsening flipping sequence from w to w'.

2.2 Some Simple Conditions for cp-Dominance

For outcomes w and v we define $\Delta(w, v)$ to be the set of variables on which they differ, i.e., $\{X \in \mathcal{V} : w(X) \neq v(X)\}$. The following lemma gives two simple sufficient conditions for w to dominate v with respect to a CP-net. In Case (i), for each variable X in $\Delta(w, v)$, there is a worsening flip from w, changing $w(X)$ to $v(X)$. In Case (ii) there is an improving flip from v changing $v(X)$ to $w(X)$.

Lemma 1. *Consider an acyclic CP-net Σ and two different outcomes w and v. Then w cp-dominates v w.r.t. CP-net Σ if either*

(i) for all $X \in \Delta(w, v)$, $w(X) > v(X)$ given w; or
(ii) for all $X \in \Delta(w, v)$, $w(X) > v(X)$ given v.

Proof. Let $k = |\Delta(w, v)|$, which is greater than zero because $w \neq v$. Let us label the elements of $\Delta(w, v)$ as $X_1, \ldots, X_k$ in such a way that if $i < j$ then X_i is not an ancestor of X_j with respect to the CP-net directed graph; this is possible because of the acyclicity assumption on Σ. To prove (i), beginning with outcome w, we flip variables of w to v in the order $X_1, \ldots, X_k$, so that we first change $w(X_1)$ to $v(X_1)$, and then change $w(X_2)$ to $v(X_2)$, and so on. The choice of variable ordering means that when we flip variable X_i the assignment to the parents $\mathcal{U}_{X_i}$ of X_i is just $w(\mathcal{U}_{X_i})$. It can be seen that this is a sequence of worsening flips from w to v, and thus, w cp-dominates v w.r.t. Σ.

Part (ii) is very similar, except that we start with v, and iteratively change X_i from $v(X_i)$ to $w(X_i)$ in the order $i = 1, \ldots, k$. The assumption behind part (ii) implies that we obtain an improving flipping sequence from v to w. □

Lemma 1 can be used to prove a more general form of itself.

Proposition 1. *Consider an acyclic CP-net Σ and two different outcomes w and v. Assume that for each $X \in \Delta(w, v)$ either $w(X) > v(X)$ given w w.r.t. Σ, or $w(X) > v(X)$ given v w.r.t. Σ. Then w cp-dominates v w.r.t. Σ.*

Proof. Define outcome u by $u(X) = v(X)$ if X is such that $w(X) > v(X)$ given w (so $X \in \Delta(w, v)$), and $u(X) = w(X)$ otherwise. Then $\Delta(w, v)$ is the disjoint union of $\Delta(w, u)$ and $\Delta(u, v)$.

For all $X \in \Delta(w, u)$, $w(X) > u(X)$ given w, because $u(X) = v(X)$ and $w(X) > v(X)$ given w. Lemma 1 implies that w cp-dominates u w.r.t. Σ.

For all $X \in \Delta(u, v)$, $u(X) > v(X)$ given v, since $u(X) = w(X)$ and $u(X) > v(X)$ given w. Lemma 1 implies that u cp-dominates v w.r.t. Σ. Thus, w cp-dominates v w.r.t. Σ. □

2.3 The Local Dominance Relation

The conditions of Proposition 1 involve what might be called a local dominance condition.

Definition 2. *Given an acyclic CP-net Σ we say that outcome w* locally dominates *outcome v [w.r.t. CP-net Σ], written $w >^{\Sigma}_{LD} v$, if for each $X \in \Delta(w,v)$ either $w(X) > v(X)$ given w w.r.t. Σ; or $w(X) > v(X)$ given v w.r.t. Σ.*

Proposition 1 above implies that if w locally dominates v then w cp-dominates v, so that $w >^{\Sigma}_{LD} v$ implies $w >^{\Sigma}_{\mathrm{cp}} v$. In fact, we even have the following result.

Proposition 2. *Given an acyclic CP-net Σ, binary relation $>^{\Sigma}_{\mathrm{cp}}$ is the transitive closure of $>^{\Sigma}_{LD}$.*

Proof. Let $\succ$ be the transitive closure of $>^{\Sigma}_{LD}$. Since, by Proposition 1, $>^{\Sigma}_{LD}$ is a subset of $>^{\Sigma}_{\mathrm{cp}}$, and the latter is transitive, we have that $\succ$ is a subset of $>^{\Sigma}_{\mathrm{cp}}$.

Suppose that w' is a worsening flip from w w.r.t. Σ. Then, $w(X) > w'(X)$ given w and $\Delta(w,w') = \{X\}$, which implies that w locally dominates w'. This shows that $>^{\Sigma}_{LD}$, and thus, $\succ$, contains the worsening flip relation $>^{\Sigma}_{wf}$ induced by Σ. Being transitive, $\succ$ contains the transitive closure $>^{\Sigma}_{\mathrm{cp}}$ of $>^{\Sigma}_{wf}$. We have therefore shown that $>^{\Sigma}_{\mathrm{cp}}$ equals $\succ$, the transitive closure of $>^{\Sigma}_{LD}$. □

3 Pareto Ordering for CP-nets in the general case

A Pareto Ordering between outcomes comes down to saying that w dominates w' if $\forall X \in \mathcal{V}, w(X)$ is at least as good an assignment as $w'(X)$ (and better for some X). However, it is not so easy to define Pareto dominance between outcomes in a CP-net when variables are not Boolean. It is often impossible to compare $w(X)$ and $w'(X)$ directly as there is generally no relation $>^X_u$ that compares them. To perform this kind of comparison in the general case of a dependency graph, we must in some way map the various preference relations $>^X_u$ on $\mathrm{Dom}(X)$ to some common scale, either totally (using a scoring function) or partially on some landmark values (mapping the best choices or the worst choices). We define a somewhat extreme Pareto-like relation, using the latter idea, below. As mentioned in Section 1, and discussed in detail in Section 4, the more natural form of Pareto dominance applies only for the case of Boolean CP-nets.

3.1 A variant of Pareto dominance for CP-nets

We define relation $>^{\Sigma}_{sp}$ on the set of outcomes, which can be viewed as being based on a strong variant of the Pareto condition (with sp standing for *strong Pareto*).

Definition 3 (Fully dominating and fully dominated). *For outcome w, we say that x is* fully dominating in X given w [w.r.t. Σ] *if $x \in \mathrm{Dom}(X)$, and for all $x' \in \mathrm{Dom}(X) \setminus \{x\}$ we have $x > x'$ given w w.r.t. Σ.*

Similarly, we say that x is fully dominated in X given w [w.r.t. Σ] *if $x \in \mathrm{Dom}(X)$, and $|\mathrm{Dom}(X)| > 1$ and for all $x' \in \mathrm{Dom}(X) \setminus \{x\}$ we have $x' > x$ given w w.r.t. Σ.*

Thus, if x is fully dominating in X given w then x is not fully dominated in X given w. Also, there can at most one element $x \in \mathrm{Dom}(X)$ that is fully dominating in X given w, and at most one that is fully dominated in X given w.

We define irreflexive relation $>^{\Sigma}_{sp}$ by, for different outcomes w and v, $w >^{\Sigma}_{sp} v$ if and only if for all $X \in \mathcal{V}$ either $v(X)$ is fully dominated in X given v w.r.t. Σ; or $w(X)$ is fully dominating in X given w w.r.t. Σ.

In the case in which the local relations $>^{X}_{u}$ are total orders, then the definitions can be simplified. Consider any outcome w, and value x in $\mathrm{Dom}(X)$, and let u be the projection of w to the parent set of X. Let x^*_u and x_{u*} be the best and the worst element (respectively) in $\mathrm{Dom}(X)$ for relation $>^{X}_{u}$. Then x is fully dominating in X given w if and only if $x = x^*_u$, and x is fully dominated in X given w if and only if $|\mathrm{Dom}(X)| > 1$ and $x = x_{u*}$. Another way of defining the $>^{\Sigma}_{sp}$ relation then consists, for each relation $>^{X}_{u}$, of mapping $\mathrm{Dom}(X)$ to a three-valued totally ordered scale $L = \{1, I, 0\}$ with $1 > I > 0$ using a kind of qualitative scoring function $f^X_u : \mathrm{Dom}(X) \to L$ defined by $f^X_u(x^*_u) = 1$, $f^X_u(x_{u*}) = 0$, and $f^X_u(x) = I$ otherwise. Note that relation $w >^{\Sigma}_{sp} w'$ expresses a very strong form of Pareto-dominance, since it requires that not only $w \neq w'$ and $f^X_u(w(X)) \geq f^X_u(w'(X))$, but also that either $f^X_u(w(X)) = 1$ or $f^X_u(w(X)) = 0, \forall X \in \mathcal{V}$.

Proposition 3. *Relation $>^{\Sigma}_{sp}$ is transitive, and is contained in $>^{\Sigma}_{LD}$, i.e., $w >^{\Sigma}_{sp} v$ implies $w >^{\Sigma}_{LD} v$, and thus $>^{\Sigma}_{sp} \subseteq >^{\Sigma}_{LD} \subseteq >^{\Sigma}_{\mathrm{cp}}$. Furthermore, we have $>^{\Sigma}_{sp}$ and $>^{\Sigma}_{\mathrm{cp}}$ are equal (i.e., are the same relation) if and only if $>^{\Sigma}_{sp}$ and $>^{\Sigma}_{LD}$ are equal.*

Proof. We will prove transitivity of $>^{\Sigma}_{sp}$ by showing that if $w_1 >^{\Sigma}_{sp} w_2$ and $w_2 >^{\Sigma}_{sp} w_3$ then $w_1 >^{\Sigma}_{sp} w_3$. Consider any $X \in \mathcal{V}$ such that $w_3(X)$ is not fully dominated in X given w_3. Since $w_2 >^{\Sigma}_{sp} w_3$, we have that $w_2(X)$ is fully dominating in X given w_2, and so $w_2(X)$ is not fully dominated in X given w_2. Since $w_1 >^{\Sigma}_{sp} w_2$, we have that $w_1(X)$ is fully dominating in X given w_1. Thus, for all $X \in \mathcal{V}$, if $w_3(X)$ is not fully dominated in X given w_3 then $w_1(X)$ is fully dominating in X given w_1, and hence, $w_1 >^{\Sigma}_{sp} w_3$, proving transitivity.

Now, suppose that $w_1 >^{\Sigma}_{sp} w_2$, and consider any $X \in \mathcal{V}$. Either (i) $w_2(X)$ is fully dominated in X given w_2, and thus, $w_1(X) > w_2(X)$ given w_2; or (ii) $w_1(X)$ is fully dominating in X given w_1, and thus, $w_1(X) > w_2(X)$ given w_1; therefore we have $w_1 >^{\Sigma}_{LD} w_2$.

Clearly if $>^{\Sigma}_{sp}$ and $>^{\Sigma}_{\mathrm{cp}}$ are equal then the inclusions $>^{\Sigma}_{sp} \subseteq >^{\Sigma}_{LD} \subseteq >^{\Sigma}_{\mathrm{cp}}$ imply that $>^{\Sigma}_{sp}$ and $>^{\Sigma}_{LD}$ are equal. Conversely, assume that $>^{\Sigma}_{sp}$ and $>^{\Sigma}_{LD}$ are equal.

We then have that $>^{\Sigma}_{LD}$ is transitive (since $>^{\Sigma}_{sp}$ is transitive), and thus it is equal to its transitive closure, which equals $>^{\Sigma}_{\text{cp}}$ by Proposition 2. □

3.2 Necessary and Sufficient Conditions for Equality of $>^{\Sigma}_{sp}$ and $>^{\Sigma}_{\text{cp}}$

We will show that $>^{\Sigma}_{sp}$ and $>^{\Sigma}_{\text{cp}}$ are only equal under extremely special conditions, including that the CP-net is unconditional and that each domain has at most two elements. We use a series of lemmas to prove the result.

The first lemma follows easily using the transitivity of $>^{\Sigma}_{sp}$.

Lemma 2. *Given CP-net Σ, then we have $>^{\Sigma}_{sp}$ equals $>^{\Sigma}_{\text{cp}}$ if and only if for all pairs (w, w') such that w' is a worsening flip from w we have $w >^{\Sigma}_{sp} w'$.*

Proof. We need to prove that $>^{\Sigma}_{sp}$ equals $>^{\Sigma}_{\text{cp}}$ if and only if $>^{\Sigma}_{sp}$ contains the worsening flip relation $>^{\Sigma}_{wf}$ induced by Σ. Since $>^{\Sigma}_{\text{cp}}$ is the transitive closure of $>^{\Sigma}_{wf}$, if $>^{\Sigma}_{sp}$ equals $>^{\Sigma}_{\text{cp}}$ then $>^{\Sigma}_{sp}$ contains $>^{\Sigma}_{wf}$.

Regarding the converse, assume that $>^{\Sigma}_{sp}$ contains $>^{\Sigma}_{wf}$. Since, by Proposition 3, $>^{\Sigma}_{sp}$ is transitive, then $>^{\Sigma}_{sp}$ contains the transitive closure $>^{\Sigma}_{\text{cp}}$ of $>^{\Sigma}_{wf}$. Proposition 3 implies that $>^{\Sigma}_{sp}$ is a subset of $>^{\Sigma}_{\text{cp}}$, so $>^{\Sigma}_{sp}$ equals $>^{\Sigma}_{\text{cp}}$. □

The definition of $>^{\Sigma}_{sp}$ leads to the following characterisation. Suppose that w' is a worsening flip from w w.r.t. CP-net Σ, with X being the variable on which they differ. Then $w >^{\Sigma}_{sp} w'$ if and only if (a) either $w(X)$ is fully dominating in X given w w.r.t. Σ, or $w'(X)$ is fully dominated in X given w w.r.t. Σ; and (b) for all $Y \in \mathcal{V} \setminus \{X\}$,

(i) if Y is not a child of X then $w(Y)$ is either fully dominated or fully dominating in Y given w w.r.t. Σ; and
(ii) if Y is a child of X then $w(Y)$ is either fully dominating in Y given w w.r.t. Σ or fully dominated in Y given w' w.r.t. Σ.

The above considerations lead to the following result.

Lemma 3. *Consider any $X \in \mathcal{V}$, and any assignment u to the parents of X, and any values $x, x' \in \text{Dom}(X)$ such that $x >^{X}_{u} x'$. Assume that $w >^{\Sigma}_{sp} w'$ whenever (w, w') is an associated worsening flip, i.e., if $w(X) = x$ and $w'(X) = x'$, and w and w' agree on all other variables, and w extends u. Let (v, v') be one such associated worsening flip.*

If variable Z is not a child of X and z is any element of $\text{Dom}(Z)$ then z is either fully dominated or fully dominating in X given v w.r.t. Σ. We have $|\text{Dom}(Z)| \leq 2$.

If variable Y is a child of X and y is any element of $\text{Dom}(Y)$ then y is either fully dominating given v or fully dominated given v'. We have $|\text{Dom}(Y)| \leq 2$.

Note that the condition $|\text{Dom}(Z)| \leq 2$ follows since there can be at most one fully dominated and at most one fully dominating element in X given v.

Lemma 3 implies that for any variable X, every other variable has at most two values, which immediately implies that every domain has at most two elements:

Lemma 4. *Suppose that* $>^{\Sigma}_{sp}$ *and* $>^{\Sigma}_{\mathrm{cp}}$ *are equal. Then each domain has at most two values.*

Definition 4 (True parents and being unconditional). *Let* Y *be a variable and let* X *be an element of its parent set* $\mathcal{U}_Y$*. We say that* X *is not a* true parent *of* Y *if for all assignments* u *and* u' *to* $\mathcal{U}_Y$ *that differ only on the value of* X*, if* $y >^{Y}_{u} y'$ *then* $y >^{Y}_{u'} y'$*. We say that* Y *is* unconditional *in* Σ *if it has no true parents.*

If X is not a true parent of Y then $>^{Y}_{u}$ does not depend on X. For any CP-net Σ we can generate an equivalent CP-net (i.e., that generates the same ordering on outcomes) such that every parent of every variable is a true parent.

Lemma 5. *Suppose that* $>^{\Sigma}_{sp}$ *and* $>^{\Sigma}_{\mathrm{cp}}$ *are equal. Assume that every parent of variable* Y *is unconditional, and let* X *be one such parent. Suppose that* u *is some assignment to the parents of* Y*, and that* u' *is another assignment that differs from* u *only on the value of* X*. If* $y >^{Y}_{u} y'$ *then* $y >^{Y}_{u'} y'$*.*

Proof. Suppose that $y >^{Y}_{u} y'$. Let v be any outcome extending u and let v' be any outcome extending u'. Lemma 4 implies that X has at most two values. If X had only one value then it is trivially not a true parent of Y, so we can assume that $\mathrm{Dom}(X) = 2$. X is unconditional so it has no parents. Our definition of a CP-net implies that the relation $>^{X}$ is non-empty, so we have $x_1 >^{X} x_2$, for some labelling x_1 and x_2 of the values of X. We first consider the case in which $u(X) = x_1$. Now, y' is not fully dominating given u and so, by Lemma 3, y' is fully dominated given u', which implies $y >^{Y}_{u'} y'$.

We now consider the other case in which $u(X) = x_2$. Then, y is not fully dominated given u, and so, by Lemma 3, y is fully dominating given u', and thus, also $y >^{Y}_{u'} y'$. □

Lemma 5 implies that X is not a true parent of Y. Since X was an arbitrary parent of Y, it then implies that Y has no true parent, so is unconditional. Applying this result inductively then implies that every variable in $\mathcal{V}$ is unconditional with respect to Σ. Along with Lemmas 3 and 4, this leads to the following result.

Proposition 4. *Given CP-net* Σ*, then we have* $>^{\Sigma}_{sp}$ *equals* $>^{\Sigma}_{\mathrm{cp}}$ *if and only if* Σ *be a Boolean locally totally ordered CP-net such that each variable* X *is unconditional in* Σ*.*

4 Pareto Ordering for the Boolean Case

As discussed earlier, there is a natural way of defining a Pareto ordering for the case of Boolean locally totally ordered CP-nets. Basically, if variables are Boolean, each of its values is either fully dominating or fully dominated in each context. So, the relation $>^{\Sigma}_{sp}$ becomes a full-fledged Pareto ordering. In this section we analyse the relationship between this Pareto ordering and the CP-net ordering.

Let Σ be a Boolean locally totally ordered CP-net. Consider any outcome w. We say that variable X is *bad for w* if there is an improving flip of variable X from w to another outcome w'. Define F_w to be the set of variables which are bad for w.

The definition of F_w and of the local dominance relation (see Section 2.3) immediately leads to the following expression of $>^{\Sigma}_{LD}$ in the Boolean locally totally ordered case.

Lemma 6. *Let Σ be a Boolean locally totally ordered CP-net. Then, for different outcomes w and v, we have $w >^{\Sigma}_{LD} v$ if and only if $F_w \cap \Delta(w, v) \subseteq F_v$.*

Proof. $w >^{\Sigma}_{LD} v$ if and only if for each $X \in \Delta(w, v)$ either $w(X) > v(X)$ given w, or $w(X) > v(X)$ given v. For $X \in \Delta(w, v)$, we have $w(X) > v(X)$ given w if and only if $X \notin F_w$; and we have $w(X) > v(X)$ given v if and only if $X \in F_v$. Thus, $w >^{\Sigma}_{LD} v$ if and only if for each $X \in \Delta(w, v)$ [$X \in F_w \Rightarrow X \in F_v$], which is if and only if $F_w \cap \Delta(w, v) \subseteq F_v$. □

We define the irreflexive binary relation $>^{\Sigma}_{par}$ on outcomes as follows.

Definition 5. *For different outcomes w and v, $w >^{\Sigma}_{par} v$ if and only if $F_w \subseteq F_v$, i.e., every variable that is bad for w is also bad for v.*

This can be viewed as a kind of Pareto ordering, and equals the strong Pareto relation $>^{\Sigma}_{sp}$ (see Section 3.1) for the Boolean locally totally ordered case.

Lemma 7. *Let Σ be a Boolean locally totally ordered CP-net. Let w and v be outcomes. Then $w >^{\Sigma}_{sp} v$ if and only if $w >^{\Sigma}_{par} v$.*

Proof. For different w and v, $w >^{\Sigma}_{sp} v$ if and only if for all $X \in \mathcal{V}$ either $v(X)$ is fully dominated in X given v, or $w(X)$ is fully dominating in X given w.

Suppose that $w >^{\Sigma}_{sp} v$ and consider any $X \in \mathcal{V}$. If $X \in F_w$ then $w(X)$ is not fully dominating in X given w, and so $v(X)$ is fully dominated in X given v, which implies that $X \in F_v$. We have shown that $F_w \subseteq F_v$.

Conversely, assume that $F_w \subseteq F_v$, and consider any $X \in \mathcal{V}$. such that $v(X)$ is not fully dominated in X given v. Because Σ is a Boolean locally totally ordered CP-net this implies that X is not bad for v. Since $F_w \subseteq F_v$, this implies that X is not bad for w, and so, $w(X)$ is fully dominating in X given w. This proves that $w >^{\Sigma}_{sp} v$. □

The CP-net relation contains the Pareto relation, with the local dominance relation being between the two.

Theorem 1. *Let Σ be a Boolean locally totally ordered CP-net. Relation $>^{\Sigma}_{par}$ is transitive, and is contained in $>^{\Sigma}_{LD}$, i.e., $w >^{\Sigma}_{par} v$ implies $w >^{\Sigma}_{LD} v$, and thus $>^{\Sigma}_{par} \subseteq >^{\Sigma}_{LD} \subseteq >^{\Sigma}_{\mathrm{cp}}$. Furthermore, we have $>^{\Sigma}_{par}$ and $>^{\Sigma}_{\mathrm{cp}}$ are equal (i.e., are the same relation) if and only if $>^{\Sigma}_{par}$ and $>^{\Sigma}_{LD}$ are equal, which happens only if every variable of the CP-net is unconditional.*

Proof. Theorem 1 follows immediately from Propositions 3 and 4 and Lemma 7. □

As a consequence, we get that CP-nets are in agreement with Pareto ordering in the case of Boolean locally totally ordered variables: for any variable X and any configuration u of its parents, consider the mapping $f_u^X : \mathrm{Dom}(X) \to \{0,1\}$ such that $f_u^X(x^*) = 1$ and $f_u^X(x_{u*}) = 0$. For any two distinct outcomes w and w', we have that $\forall X \in \mathcal{V}, f^X_{w(U_X)}(w(X)) \geq f^X_{w'(U_X)}(w'(X))$ if and only $F_w \subseteq F_{w'}$, which is Pareto-ordering $>^{\Sigma}_{par}$.

We emphasise the following part of the theorem:

Corollary 1. *Let Σ be a Boolean locally totally ordered CP-net, $w >^{\Sigma}_{par} w'$ implies $w >^{\Sigma}_{\mathrm{cp}} w'$.*

As shown in the previous section, it does not seem straightforward to extend this Pareto ordering in a natural way to non-Boolean variables without using scaling functions that map all partial orders $(\mathrm{Dom}(X), >^X_u), u \in \mathcal{U}_X$ to a common value scale, unless the variables are all preferentially independent from one another. In this case, $\mathcal{U}_X = \emptyset, \forall X$, and $>^X_u = >^X, \forall X \in \mathrm{Dom}(X)$. We could then define the Pareto dominance relation $>^{\Sigma}_{par}$ on outcomes as $w >^{\Sigma}_{par} w'$ if and only if $w \neq w'$ and $w(X) >^X w'(X)$ or $w(X) = w'(X)$ for all $X \in \mathcal{V}$.

5 π-pref nets

Possibility theory [8] is a theory of uncertainty devoted to the representation of incomplete information. It is maxitive (addition is replaced by maximum) in contrast with probability theory. It ranges from purely ordinal to purely numerical representations. Possibility theory can be used for representing preferences [9]. It relies on the idea of a possibility distribution π, i.e., a mapping from a universe of discourse Ω to the unit interval $[0,1]$. Possibility degrees $\pi(w)$ estimate to what extent the solution w is not unsatisfactory. π-pref nets are based on possibilistic networks [3], using conditional possibilities of the form $\pi(x|u) = \frac{\Pi(x \wedge u)}{\Pi(u)}$, for $u \in \mathrm{Dom}(\mathcal{U}_X)$, where $\Pi(\varphi) = \max_{w \models \varphi} \pi(w)$. The use of product-based conditioning rather than min-based conditioning leads to possibilistic nets that are more similar to Bayesian nets.

The ceteris paribus assumption of CP-nets is replaced in possibilistic networks by a chain rule like in Bayesian networks. It enables one to compute, using an aggregation function, the degree of possibility of solutions. However it is supposed that these numerical values are unknown and represented by symbolic weights. Only ordering between symbolic values or products thereof can be expressed. The dominance relation between solutions is obtained by comparing products of symbolic utility values computed for them from the conditional preference tables.

Definition 6. *([1]) A Boolean* possibilistic preference network (π-pref net) *is a preference network, where $|Dom(X)| = 2, \forall X \in \mathcal{V}$, and each preference statement $x >^X_u x'$ is associated to a conditional possibility distribution such that*

$\pi(x|u) = 1 > \pi(x'|u)) = \alpha^u_X$*, and* α^u_X *is a non-instantiated variable on* $[0, 1)$ *we call a* symbolic weight*. One may also have indifference statements* $x \sim^X_u x'$*, expressed by* $\pi(x|u) = \pi(x'|u) = 1$*.*

π-pref nets induce a partial ordering between solutions based on the comparison of their degrees of possibility in the sense of a joint possibility distribution computed using the product-based chain rule: $\pi(x_i, \ldots, x_n) = \prod_{i=1,\ldots,n} \pi(x_i|u_i)$. The preferences between solutions are of the form $w \succ_\pi w'$ if and only if $\pi(w) > \pi(w')$ *for all instantiations* of the symbolic weights.

5.1 π-pref nets vs CP-nets

Let us compare preference relations between solutions induced by both CP-nets and π-pref nets. It has been shown [2] that the ordering between solutions induced by a π-pref net corresponds to the Pareto ordering between the vectors $w = (\theta_1(w), \ldots, \theta_n(w))$ where $\theta_i(w) = \pi(w(X_i)|w(\mathcal{U}_{X_i})), i = 1, \ldots, n$.

As symbolic weights are not comparable across variables, it is easy to see that the only way to have $\pi(w) \geq \pi(w')$ is to have $\theta_k(w) \geq \theta_k(w')$ in each component k of w and w'. Otherwise the products will be incomparable due to the presence of distinct symbolic variables on each side. So, if $w \neq w'$,

$$w \succ_\pi w' \text{ if and only if } \theta_k(w) \geq \theta_k(w'), k = 1, \ldots, n \text{ and } \exists i : \theta_i(w) > \theta_i(w').$$

It is then known that the π-pref net ordering between solutions induced by the preference tables is refined by comparing the sets F_w of bad variables for w:

$$w \succ_\pi w' \Rightarrow F_w \subset F_{w'}$$

since if two solutions contain variables having bad assignments in the sense of the preference tables, the corresponding symbolic values may differ if the contexts for assigning a value to this variable differ. It has been shown that if the weights α^u_X reflecting the satisfaction level due to assigning the bad value to X_i in the context u_i do not depend on this context, then we have an equivalence in the above implication:

$$\text{If } \forall X \in \mathcal{V}, \alpha^u_X = \alpha_X, \forall u_i \in \mathrm{Dom}(\mathcal{U}_X), \text{ then } w \succ_\pi w' \iff w >^{\Sigma}_{par} w'.$$

As a consequence, using Corollary 1, it is clear that $w \succ_\pi w'$ implies $w >^{\Sigma}_{\mathrm{cp}} w$ so that the CP-net preference ordering refines the one induced by the corresponding Boolean π-pref net. It suggests that we can try to add ceteris paribus constraints to a π-pref net and so as to capture the preferences expressed by a CP-net.

In the following, we highlight local constraints between each node and its children that enable ceteris paribus to be simulated. Ceteris paribus constraints are of the form $w >^{\Sigma}_{\mathrm{cp}} w'$ where w and w' differ by one flip. For each such statement (one per variable), we add the constraint on possibility degrees $\pi(w) > \pi(w')$. Using the chain rule, it corresponds to comparing products of symbolic weights. Let $\mathrm{Dom}(\mathcal{U}_X) = \times_{X_i \in \mathcal{U}_X} Dom(X_i)$ denote the Cartesian product of domains of variables in $\mathcal{U}_X$, $\alpha_X^u = \pi(x^-|u)$, where x^- is bad for X and $\gamma_Y^{u'} = \pi(y^-|u')$. Suppose a CP-net and a π-pref net built from the same preference statements. It has been shown in [2] that the worsening flip constraints are all induced by the conditions: $\forall\ X \in \mathcal{V}$ s.t. X has children $\mathcal{C}h(X) \neq \emptyset$:

$$\max_{u \in \mathrm{Dom}(\mathcal{U}_X)} \alpha_X^u < \prod_{Y \in \mathcal{C}h(X)} \min_{u' \in \mathrm{Dom}(\mathcal{U}_Y)} \gamma_Y^{u'}$$

Let $\succ_\pi^+$ be the resulting preference ordering built from the preference tables and applying constraints of the above format between symbolic weights, then, it is clear that $\omega \succ_{cp} \omega' \Rightarrow \omega \succ_\pi^+ \omega'$: relation $\succ_\pi^+$ is a bracketing from above of the CP-net ordering.

5.2 Relation $\succ_\pi^+$ as a Polynomial Upper Bound for CP-net Dominance

In this section we give a characterisation of the relation $\succ_\pi^+$ in terms of deduction of linear constraints, which implies that determining dominance with respect to $\succ_\pi^+$ is polynomial. It is thus a polynomial upper bound for CP-net dominance.

We list all the different symbolic weights (not including 1) as $\alpha_1, \ldots, \alpha_m$, and let α represent the whole vector of symbolic weights $[\alpha_1, \ldots, \alpha_m]$.

Let a *weights vector* z be a vector of m real numbers $[z_1, \ldots, z_m]$ (with each z_i in $\{-1, 0, 1\}$). For each such weights vector z, we associate the product $\alpha_1^{z_1} \cdots \alpha_m^{z_m}$, which we abbreviate as $R_\alpha[z]$.

A comparison between products of symbolic weights can be encoded as a statement $R_\alpha[z] > 1$. For example, a comparison $\alpha_1 > \alpha_2\alpha_3$ is equivalent to $R_\alpha[z] > 1$ where $z = [1, -1, -1, 0, 0, \ldots]$, since $R_\alpha[z] = \alpha_1^1\alpha_2^{-1}\alpha_3^{-1}$ and so $R_\alpha[z] > 1 \iff \alpha_1^1\alpha_2^{-1}\alpha_3^{-1} > 1 \iff \alpha_1 > \alpha_2\alpha_3$. In this way, every ceteris paribus statement corresponds to a set of statements $R_\alpha[z] > 1$ for different vectors z.

For each $i = 1, \ldots, m$, define the vector $z^{(i)}$ as $z_i^{(i)} = -1$ and for all $j \neq i$, $z_j^{(i)} = 0$. $R_\alpha[z^{(i)}] > 1$ expresses that $\alpha_i^{-1} > 1$, i.e., $\alpha_i < 1$. For a CP-net Σ let $Z(\Sigma)$ be the set of weights vectors associated with symbolic weights comparisons for each ceteris paribus statement, plus for each $i = 1, \ldots, m$, the element $z^{(i)}$.

Similarly, every solution is associated with a product of symbolic weights, so a comparison $w > w'$ between solutions w and w' corresponds to a statement pertaining to a weights vector z'. The definitions lead easily to the following characterisation of this form of dominance.

Proposition 5. *Consider any CP-net Σ with associated set of weights vectors $Z(\Sigma)$, and let w and w' be two solutions, where the comparison $w > w'$ has associated vector z'. We have that $w \succ_\pi^+ w'$ if and only if $\{R_\alpha[z] > 1 : z \in Z(\Sigma)\}$ implies $R_\alpha[z'] > 1$, i.e., if one replaces the values of symbolic weights α_i by any real values such that $R_\alpha[z] > 1$ holds for each $z \in Z(\Sigma)$ then $R_\alpha[z'] > 1$ also holds.*

We can write $\log R_\alpha[z]$ as $z_1\lambda_1 + \cdots + z_m\lambda_m = z \cdot \lambda$, where λ is the vector $(\lambda_1, \ldots, \lambda_m)$ and each $\lambda_i = \log \alpha_i$. Thus, $R_\alpha[z] > 1 \iff \log R_\alpha[z] > 0 \iff z \cdot \lambda > 0$. By Proposition 5 this implies that $w \succ_\pi^+ w'$ if and only if for vectors λ, $\{z \cdot \lambda > 0 : z \in Z(\Sigma)\}$ implies $z' \cdot \lambda > 0$.

Using a standard result from convex sets, this leads to the following result, which gives a somewhat simpler characterisation that shows that dominance is polynomial. It also suggests potential links with Generalized Additive Iindependent (GAI) value function approximations of CP-nets [4, 7].

Theorem 2. *Consider any CP-net with associated set of weights vectors $Z(\Sigma)$, and let w and w' be two different solutions, where $w > w'$ has associated vector z'. We have that $w \succ_\pi^+ w'$ if and only if there exist non-negative real numbers r_z for each $z \in Z(\Sigma)$ such that $\sum_{z \in Z(\Sigma)} r_z z = z'$. Hence, whether or not $w \succ_\pi^+ w'$ holds can be checked in polynomial time.*

Proof. As argued above, $w \succ_\pi^+ w'$ holds if and only if for vectors λ, the set of inequalities $\{z \cdot \lambda > 0 : z \in Z(\Sigma)\}$ implies $z' \cdot \lambda > 0$. We need to show that this holds if and only if there exist non-negative real numbers r_z for each $z \in Z(\Sigma)$ such that $\sum_{z \in Z(\Sigma)} r_z z = z'$. Firstly, let us assume that there exist non-negative real numbers r_z for each $z \in Z(\Sigma)$ such that $\sum_{z \in Z(\Sigma)} r_z z = z'$. Consider any vector λ such that $z \cdot \lambda > 0$ for all $z \in Z(\Sigma)$. Then $z' \cdot \lambda = \sum_{z \in Z(\Sigma)} r_z z \cdot \lambda$ which is greater than zero since each r_z is non-negative, and at least some $r_z > 0$ (else z' is the zero vector, which would contradict $w \neq w'$).

Conversely, let us assume that there do not exist non-negative real numbers r_z for each $z \in Z(\Sigma)$ such that $\sum_{z \in Z(\Sigma)} r_z z = z'$. To prove that the set of inequalities $\{z \cdot \lambda > 0 : z \in Z(\Sigma)\}$ does not imply $z' \cdot \lambda > 0$, we will show that there exists a vector λ with $z \cdot \lambda > 0$ for all $z \in Z(\Sigma)$ but $z' \cdot \lambda \leq 0$. Let C be the set of vectors of the form $\sum_{z \in Z(\Sigma)} r_z z$ over all choices of non-negative reals r_z. Now, C is a convex and closed set, which by the hypothesis does not intersect with $\{z'\}$ (i.e., does not contain z'). Since $\{z'\}$ is closed and compact we can use a hyperplane separation theorem to show that there exists a vector λ and real numbers $c_1 < c_2$ such that for all $x \in C$, $x \cdot \lambda > c_2$ and $z' \cdot \lambda < c_1$. Because C is closed under strictly positive scalar multiplication (i.e., $x \in C$ implies $rx \in C$ for all real $r > 0$) we must have $c_2 \leq 0$, and $x \cdot \lambda \geq 0$ for all $x \in C$, and in particular $z \cdot \lambda \geq 0$ for all $z \in Z(\Sigma)$. Also, $z' \cdot \lambda < c_1 < c_2 \leq 0$ so $z' \cdot \lambda \leq 0$, as required.

The last part follows since linear programming is polynomial. □

6 Summary and Discussion

In this paper we have compared CP-nets and π-pref nets, two qualitative counterparts of Bayes nets for the representation of conditional preferences. We have studied them from the point of view of their rationality, namely whether they respect Pareto dominance between multiple Boolean variable solutions to a decision problem expressed by such graphical models. While π-pref nets naturally respect this property, strangely enough, it was previously unknown whether the preference ordering induced by CP-nets respects it or not. For more general (non-Boolean) variables, it seems difficult to extend this notion of Pareto-dominance for a CP-net in an entirely natural way. Besides, it was shown previously that the ordering induced by π-pref nets is weaker than the one induced by CP-nets, but ceteris paribus constraints can be added to a π-pref net in the form of constraints between products of symbolic variables. Here we show the polynomial nature of this encoding. Thus we get a bracketing of the CP-net preference ordering by bounds which are apparently easier to compute than standard CP-net preferences. Further research includes constructing an example that explicitly proves that the upper approximation of the CP-net ordering is not tight; moreover the case of non-Boolean variables deserves further investigation.

Acknowledgements

This material is based upon works supported by the Science Foundation Ireland under Grants No. 12/RC/2289 and No. 12/RC/2289-P2 which are co-funded under the European Regional Development Fund.

References

1. N. Ben Amor, D. Dubois, H. Gouider, H. Prade, Possibilistic preference networks, Information Sciences, 460-461: 401-415 (2018).
2. N. Ben Amor, D. Dubois, H. Gouider, H. Prade, Expressivity of Possibilistic Preference Networks with Constraints. Scalable Uncertainty Manag., LNCS 10564, Springer, 163-177 (2017).
3. S. Benferhat, D. Dubois, L. Garcia, H. Prade, On the transformation between possibilistic logic bases and possibilistic causal networks, Int. J. of Approximate Reasoning, 29(2): 135-173 (2002).
4. C. Boutilier, F. Bacchus, R.I. Brafman, UCP-Networks: A Directed Graphical Representation of Conditional Utilities. Proc. of the 17th Conf. on Uncertainty in AI, Seattle, Washington, USA, 56-64 (2001).
5. C. Boutilier, R. I. Brafman, H. H. Hoos, D. Poole, Reasoning With Conditional Ceteris Paribus Preference Statements, Proc. of the 15th Conf. on Uncertainty in AI, Stockholm, Sweden, 71-80 (1999).
6. C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, D. Poole, CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements. J. Artificial Intelligence Research, 21: 135-191 (2004).
7. R.I. Brafman, C. Domshlak, T. Kogan, Compact Value-Function Representations for Qualitative Preferences. Proc. of the 20th Conf. on Uncertainty in AI, Banff, Canada, 51-59 (2004).
8. D. Dubois and H. Prade, Possibility Theory: An Approach to Computerized Processing of Uncertainty, Plenum Press, (1988).
9. D. Dubois and H. Prade: Possibility theory as a basis for preference propagation in automated reasoning. Proc. of the 1st IEEE Inter. Conf. on Fuzzy Systems, San Diego, Ca., 821-832 (1992).
10. C. Eichhorn, M. Fey, and G. Kern-Isberner, CP- and OCF-networks: a comparison, Fuzzy sets and Systems, 298:109-127 (2016).