

Extending CP-Nets with Stronger Conditional Preference Statements

Nic Wilson

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

Abstract

A logic of conditional preferences is defined, with a language which allows the compact representation of certain kinds of conditional preference statements, a semantics and a proof theory. CP-nets can be expressed in this language, and the semantics and proof theory generalise those of CP-nets. Despite being substantially more expressive, the formalism maintains important properties of CP-nets; there are simple sufficient conditions for consistency, and, under these conditions, optimal outcomes can be efficiently generated. It is also then easy to find a total order on outcomes which extends the conditional preference order, and an approach to constrained optimisation can be used which generalises a natural approach for CP-nets. Some results regarding the expressive power of CP-nets are also given.

Introduction

CP-nets (Boutilier *et al.* 1999; a) is a formalism for compactly expressing conditional preferences in multivariate problems. They involve statements of the form: $u : x > x'$, where x, x' are values of a variable X and u is an assignment to a set of variables U (called *parents of X*). The interpretation is that given u , x is (strictly) preferred to x' , *all else being equal* (*ceteris paribus*); that is, for all assignments s to the other variables S , sux is preferred to sux' , where e.g., sux is the outcome (complete assignment) α such that $\alpha(X) = x$, $\alpha(U) = u$ and $\alpha(S) = s$. A set of such statements generates a preference relation on complete tuples, and consistency corresponds to the induced preference relation being acyclic, and hence being a (strict) partial order. An acyclic CP-net involves a set of such statements structured in a particular way, so that the parent-child relation on variables is acyclic. This condition ensures consistency. They also possess a number of other attractive properties.

However, CP-nets are quite restrictive in the kinds of preference statements that they can represent. The *ceteris paribus* statements they are based upon express a very natural kind of preference, but a rather weak one. An agent will sometimes want to express much stronger statements such as *x is preferred to x' irrespective of the values of other variables*, where the variable X is the most important variable, and, for example, x' represents a value which should be

avoided if at all possible. CP-nets cannot generally express such statements (in a compact way).

This paper develops a formalism along similar lines to CP-nets, but where a richer language of preference statements can be expressed: stronger conditional preference statements as well as the usual CP-nets *ceteris paribus* statements, and also allowing locally partially ordered preferences. The language consists of statements of the form $u : x > x' [W]$, (where W is a subset of S) which represents that for all assignments w, w' to W and assignments t to $S - W$, $tuwx$ is preferred to $tuwx'$. So, given u and any t , x is preferred to x' irrespective of the values of W . CP-nets *ceteris paribus* statements are represented by such statements with $W = \emptyset$, and the strong conditional preference statement in the previous paragraph corresponds to $\top : x > x' [V - \{X\}]$, where V is the set of all variables. As in CP-nets, this is a compact representation: each statement typically corresponds to many preferences between outcomes.

Key properties of CP-nets generalise naturally for this formalism. In particular:

- the semantics and complete proof theory;
- there is still a simple and efficient algorithm for generating an optimal outcome;
- there are simple sufficient conditions for consistency based on acyclicity of variable order;
- there exists an efficient algorithm for picking a subset of the optimal solutions, when the set of outcomes is constrained.

The next section introduces the formalism, which can be viewed as a simple logic of conditional preferences. A semantics is given and also a complete proof theory, based on 'swapping sequences' which is the natural generalisation of flipping sequences in CP-nets. The following section examines the relative expressivity of the language as compared with CP-nets. It shows how CP-nets can be represented within the language; however, this stronger kind of preference statement, which can be used, for example, to construct a lexicographic order on outcomes, is not expressible within the language of CP-nets (or TCP-nets).

Some of the main technical results in the paper use what we call a partial conditional lexicographic (pcl) order on out-

comes, which is similar to a standard lexicographic order except that it allows the importance ordering on variables to be only partial, and allows the value orderings to be partial, and conditional on the values of more important variables. If a set Γ of preference statements satisfies a local consistency property and the associated relation on variables is acyclic then the induced preference order $>_{\Gamma}$ is dominated by a pcl order, which implies consistency, as pcl orders are strict partial orders. Generating outcomes in an order consistent with a pcl order is easy, just as it is for a standard lexicographic order; this can then be used, for example, in a constrained optimisation algorithm to generate some of the optimal (i.e., maximal) outcomes with respect to $>_{\Gamma}$, as discussed in the penultimate section along with other applications of the results.

A Logic of Conditional Preferences

In this section a logic of conditional preferences is defined, with a language, semantics and a kind of proof theory. It is strongly related to a system defined in (Lang 2002), where general complexity results are derived. As we will see later, CP-nets can be expressed within this language. The logic has a somewhat restrictive language, but the restrictions entail some nice properties, generalising properties of CP-nets.

The Language. Let V be a set of variables. For each $X \in V$ let \underline{X} be the set of possible values of X . For subset of variables $U \subseteq V$ let $\underline{U} = \prod_{X \in U} \underline{X}$ be the set of possible assignments to U . A *complete tuple* or *outcome* is an element of \underline{V} , i.e., an assignment to all the variables. For complete tuple α and partial tuple $u \in \underline{U}$, we may write $\alpha \models u$ to mean that α projected to U gives u , which can also be written as $\alpha(U) = u$.

The language \mathcal{L}_V (abbreviated to \mathcal{L}) consists of all statements of the form $u : x > x' [W]$, where u is an assignment to a set of variables $U \subseteq V$ (i.e., $u \in \underline{U}$), x, x' are different values of a variable $X \notin U$, and W is some subset of $V - U - \{X\}$. The assignment to the empty set of variables is written \top . If φ is the statement $u : x > x' [W]$, we may write $u_{\varphi} = u, U_{\varphi} = U, x_{\varphi} = x, x'_{\varphi} = x', W_{\varphi} = W$ and $T_{\varphi} = V - (\{X\} \cup U \cup W)$.

Subsets of \mathcal{L} are called *conditional preference theories* (on V). For $\varphi = u : x > x' [W]$, let φ^* be the set of pairs of outcomes $\{(txw, tux'w') : t \in \underline{T}_{\varphi}, w, w' \in \underline{W}\}$. Such pairs $(\alpha, \beta) \in \varphi^*$ are intended to represent a preference for α over β , and φ is intended as a compact representation of the preference information φ^* . Informally, φ represents that, given u and any t , x is preferred to x' , irrespective of the assignments to W . For conditional preference theory $\Gamma \subseteq \mathcal{L}$, define $\Gamma^* = \bigcup_{\varphi \in \Gamma} \varphi^*$. Γ^* represents a set of preferences. We assume here that preferences are transitive, so it is then natural to define order $>_{\Gamma}$, induced on \underline{V} by Γ , to be the transitive closure of Γ^* . In the next section it is shown that CP-nets can be represented in terms of statements $u : x > x' [W]$ with $W = \emptyset$.

Note that conditional preference theories allow locally partially ordered preferences: we do not need to assume that we can elicit a total order on the values of a variable given

each assignment to its parents. This kind of representation of conditional preferences is very flexible as regards elicitation: we can reason with an arbitrary subset Γ of the language \mathcal{L} , so we can accept any conditional preference statements (of the appropriate form) that the agent is happy to give us (however, to maintain consistency, we may insist on the conditional preference theory Γ having particular properties, such as the conditions of Corollary 1 below). More statements can be added later, and, because the logic is monotonic, all of our previous deductions from Γ will still hold, in particular whether one outcome is preferred to another

We will associate with a conditional preference theory Γ a pair of binary relations on the set of variables V . Let $H(\varphi) = \{(Y, X_{\varphi}) : Y \in U_{\varphi}\}$ and let $H(\Gamma) = \bigcup_{\varphi \in \Gamma} H(\varphi)$; this can be thought of as a directed graph, which contains edge (Y, X) if and only if there is some conditional preference statement $\varphi \in \Gamma$ which makes the preferences for X conditional on Y . Let $G(\varphi) = H(\varphi) \cup \{(X_{\varphi}, Z) : Z \in W_{\varphi}\}$, and define $G(\Gamma) = \bigcup_{\varphi \in \Gamma} G(\varphi)$. $G(\Gamma)$ is $H(\Gamma)$ with extra edges (X, Z) when there is some $\varphi \in \Gamma$ representing some preference for values of X irrespective of the value of Z .

Example. I'm planning a holiday. I can either go next week (n) or later in the year (l). I've decided to go either to Oxford (o) or to Manchester (m), and I can either fly (f) or drive and take a car ferry (d). So there are three variables, X_1, X_2 and X_3 , where $X_1 = \{1, n\}$, $X_2 = \{m, o\}$ and $X_3 = \{d, f\}$. Firstly, I'd prefer to go next week irrespective of the choices of the other variables, as I could do with a break soon. This can be represented by statement φ_1 which equals $\top : n > l [\{X_2, X_3\}]$. This represents a set φ_1^* of pairs of outcomes (nw_1, lw_2) , where w_1 and w_2 are both arbitrary assignments to the set of variables $\{X_2, X_3\}$; e.g., $w_1 = mf$ and $w_2 = of$ gives the pair (nmf, lof) indicating the preference of nmf over lof . φ_1 is a compact way of representing the 16 pairs in φ_1^* . Secondly, all else being equal, I'd prefer to go to Oxford rather than Manchester. This is represented by the statement φ_2 which equals $\top : o > m [\emptyset]$. This is an unconditional *ceteris paribus* statement. It represents set φ_2^* of pairs of outcomes (x_1ox_3, x_1mx_3) meaning outcome x_1ox_3 is preferred to x_1mx_3 , where x_1 is any value of X_1 and x_3 is any value of X_3 .

My preferences on variable X_3 are conditional. I'd prefer to fly rather than drive unless I go later in the year to Manchester, when the weather will be warmer and a car would be useful for touring around. This can be represented by conditional preference statements φ_3, φ_4 and φ_5 defined as follows. φ_3 is $n : f > d [\emptyset]$, and φ_4 is $o : f > d [\emptyset]$. φ_5 is $lm : d > f [\emptyset]$, representing φ_5^* which consists of the single preference of lmd over lmf .

Let $\Gamma = \{\varphi_1, \dots, \varphi_5\}$. $G(\Gamma)$ equals the total order on variables, $\{(X_1, X_2), (X_1, X_3), (X_2, X_3)\}$, and $H(\Gamma) = \{(X_1, X_3), (X_2, X_3)\}$. The statement φ_1 cannot be represented in a CP-net on $V = \{X_1, X_2, X_3\}$. The others all can as they involve empty W and they express locally totally ordered preferences.

The induced partial ordering $>_{\Gamma}$ on outcomes can

be shown to be the transitive closure of: $\text{nof } >_{\Gamma}$ $\{\text{nod}, \text{nmf}\} >_{\Gamma}$ $\text{nmd} >_{\Gamma}$ $\text{lof} >_{\Gamma}$ $\text{lod} >_{\Gamma}$ $\text{lmd} >_{\Gamma}$ lmf , so that $>_{\Gamma}$ is almost a total order, with only the pair of outcomes nod and nmf not being ordered.

Semantics. We define models of \mathcal{L} to be strict total orders on \underline{V} , i.e., irreflexive¹ transitive binary relations $>$ on \underline{V} such that for all α and β in \underline{V} , with $\alpha \neq \beta$, either $\alpha > \beta$ or $\beta > \alpha$.² For such a total order $>$, we say $> \models \varphi$ if and only if $> \supseteq \varphi^*$, so that $>$ is a model of $\varphi = u : x > x' [W]$ if and only if for all $t \in \underline{T}$ and $w, w' \in \underline{W}$, $twxw > twx'w'$. For $\Gamma \subseteq \mathcal{L}$ we say $> \models \Gamma$ if and only if for all $\varphi \in \Gamma$, $> \models \varphi$, which is if and only if $> \supseteq \Gamma^*$. For $\Gamma \subseteq \mathcal{L}$ and $\varphi \in \mathcal{L}$, we define the semantic entailment relation by $\Gamma \models \varphi$ if and only if $> \models \varphi$ for all $>$ such that $> \models \Gamma$. For $\alpha, \beta \in \underline{V}$ we also say that $\Gamma \models (\alpha, \beta)$ if $\alpha > \beta$ holds for all models $>$ of Γ . We say that Γ is *consistent* if it has a model, i.e., if there exists strict total order $>$ with $> \models \Gamma$. The construction of semantic entailment relation \models ensures that it is monotonic.

In the example, Γ is consistent; in fact, there are two total orders $>$ on outcomes which satisfy Γ (i.e., contain $>_{\Gamma}$); they only differ according to whether they have $\text{nod} > \text{nmf}$ or $\text{nmf} > \text{nod}$.

Proof theory. Let $\alpha, \beta \in \underline{V}$ be two outcomes. We say that β is a *worsening swap from* α with respect to conditional preference theory Γ if and only if $(\alpha, \beta) \in \Gamma^*$, i.e., iff there exists $\varphi = (u : x > x' [W]) \in \Gamma$ such that $\alpha \models u$, $\beta \models u$, $\alpha(X) = x$, $\beta(X) = x'$, and $\alpha(T_{\varphi}) = \beta(T_{\varphi})$. We say that β can be reached from α with a *worsening swapping sequence* (with respect to Γ) if there exists a sequence $\alpha = \alpha_1, \dots, \alpha_l = \beta$ with for each $k = 1, \dots, l-1$, α_{k+1} is a worsening swap from α_k , i.e., $(\alpha_k, \alpha_{k+1}) \in \Gamma^*$. Clearly, then (α, β) is in the transitive closure $>_{\Gamma}$ of Γ^* . Conversely, if (α', β') is in the transitive closure of Γ^* then there exists a sequence $\alpha' = \alpha_1, \dots, \alpha_l = \beta'$ with for each $k = 1, \dots, l-1$, $(\alpha_k, \alpha_{k+1}) \in \Gamma^*$. In fact we have the following result which is a soundness and completeness result for worsening swapping sequences.

Theorem 1 *Let Γ be a conditional preference theory on V and let $\alpha, \beta \in \underline{V}$ be outcomes. Then $\alpha >_{\Gamma} \beta$ if and only if there exists a worsening swapping sequence with respect to Γ from α to β . Also Γ is consistent if and only if $>_{\Gamma}$ is irreflexive. If Γ is consistent then $\Gamma \models (\alpha, \beta)$ if and only if $\alpha >_{\Gamma} \beta$.*

CP-nets and Expressibility

In this section we show how CP-nets can be expressed as conditional preference theories, using statements $u : x > x' [W]$ with $W = \emptyset$. It is also shown that the language is a good deal more expressive than CP-nets.

¹Relation $>$ on set A is irreflexive if and only if for all $a \in A$, it is not the case that $a > a$. It is acyclic if and only if its transitive closure is irreflexive, so that there are no cycles $a > a' > a'' > \dots > a$.

²Binary relation \succ on set \underline{V} is defined to be a subset of $\underline{V} \times \underline{V}$; the notations “ $(\alpha, \beta) \in \succ$ ” and “ $\alpha \succ \beta$ ” are used interchangeably.

Expressing CP-nets in the language

A *CP-net over V* is defined (see (Boutilier *et al.* 1999) and especially definitions 1, 2 and 3 of (Boutilier *et al.* a)) to be a pair $N = (H, \text{CT})$ where H is a (binary) relation on V (which is conventionally thought of as a directed graph) and CT is a function which assigns a *conditional preference table* to each variable $X \in V$. The conditional preference table $\text{CT}(X)$ is defined to be a function which assigns to each³ $u \in \text{Pa}_H(X)$ a strict total order \succ_u^X on \underline{X} .

Let $>$ be a (strict) total order on \underline{V} . Let X be a variable and let $u \in \text{Pa}_H(X)$ be an assignment to the parents of X . Let $T = \underline{V} - \text{Pa}_H(X) - \{X\}$. $>$ is said to satisfy \succ_u^X if $twx > twx'$ holds for all $t \in \underline{T}$ and for all $x, x' \in \underline{X}$ such that $x \succ_u^X x'$.

$>$ is said to satisfy CP-net $N = (H, \text{CT})$ if for all $X \in V$, and all $u \in \text{Pa}_H(X)$, $>$ satisfies \succ_u^X (where $\succ_u^X = \text{CT}(X)(u)$). CP-net N is said to be *satisfiable* if there exists some $>$ which satisfies N . There is a simple sufficient condition for satisfiability of a CP-net N : that its associated relation H is acyclic.

For CP-net N define relation \succ_N on \underline{V} as follows. For $\alpha, \beta \in \underline{V}$, $\alpha \succ_N \beta$ if and only if $\alpha > \beta$ for all total orders $>$ satisfying N . Therefore \succ_N is the intersection of all $>$ satisfying N .

For $X \in V$ and $u \in \text{Pa}_H(X)$, let $\Gamma_N^{X,u} \subseteq \mathcal{L}$ be the set of statements $\{(u : x > x' [\emptyset]) : x, x' \in \underline{X}, x \succ_u^X x'\}$. Let conditional preference theory Γ_N be the union of sets $\Gamma_N^{X,u}$ over all $X \in V$ and $u \in \text{Pa}_H(X)$. Note that the construction of Γ_N is linear in the size of the conditional preference table.⁴ Now, $> \models \Gamma_N^{X,u}$ if and only if $>$ satisfies \succ_u^X . So $> \models \Gamma_N$ if and only if $>$ satisfies N . Using Theorem 1 this leads to:

Proposition 1 *Let N be a CP-net, and $\Gamma_N \subseteq \mathcal{L}$ (as defined above) be its associated conditional preference theory. Then N is satisfiable if and only if Γ_N is consistent. If N is satisfiable, then $>_{\Gamma_N} = \succ_N$.*

This shows that a CP-net can be represented within the language \mathcal{L} , with the same associated order on outcomes.

Representing lexicographic orders

For set of variables V , a lexicographic order on \underline{V} involves an ordering X_1, \dots, X_n of the variables V , and for each X_i a total order $>_i$ on the set of values \underline{X}_i of X_i . Define relation $>_{lex}$ as follows. For $\alpha, \beta \in \underline{V}$, $\alpha >_{lex} \beta$ if and only if $\alpha \neq \beta$ and $\alpha(X_i) >_i \beta(X_i)$, where X_i is the first variable (i.e., with minimal i) such that $\alpha(X_i) \neq \beta(X_i)$. The lexicographic order $>_{lex}$ is a strict total order on \underline{V} .

The following proposition shows that lexicographic orders can be represented by conditional preference theories, i.e., for any lexicographic order $>_{lex}$, there exists Γ such that its associated order $>_{\Gamma}$ equals $>_{lex}$.

³ $\text{Pa}_H(X)$, the parents of X with respect to H , is the set of all Y such that $(Y, X) \in H$.

⁴If the domain of variable X is large, one might represent total order \succ_u^X by a sub-relation whose transitive closure is \succ_u^X ; the sub-relation could then also be used in the definition of $\Gamma_N^{X,u}$.

Proposition 2 For each variable X_i , let Γ_i be the set of all statements $\top : x > x' [\{X_{i+1}, \dots, X_n\}]$, where $x, x' \in X_i$ are such that $x >_i x'$. Let $\Gamma = \Gamma_1 \cup \dots \cup \Gamma_n$. Then the associated order $>_\Gamma$ equals $>_{lex}$.

The following lemma is useful for revealing the limited expressiveness of CP-nets (and TCP-nets (Brafman, Domshlak, & Shimony 2002)). We say that α covers β with respect to a (transitive) relation \succ on \underline{V} if $\alpha \succ \beta$ and there does not exist $\gamma \in \underline{V}$ with $\alpha \succ \gamma \succ \beta$.

Lemma 1

- (i) Let Γ be a conditional preference theory. Suppose α covers β with respect to $>_\Gamma$. Then β is a worsening swap from α .
- (ii) Let N be a CP-net. Suppose α covers β with respect to \succ_N . Then α and β differ on precisely one variable. In other words, there exists $X \in V$ with $\alpha(X) \neq \beta(X)$ and for all $X' \in V - \{X\}$, $\alpha(X') = \beta(X')$.
- (iii) Let M be a TCP-net, with associated relation \succ_M (so that $\gamma \succ_M \delta$ if and only if $\gamma \succ \delta$ is a consequence of M). Suppose α covers β with respect to \succ_M . Then α and β differ either on one variable or on two variables.

All three parts follow easily from the appropriate completeness theorems for swapping/flipping sequences: Theorem 1 above for (i); Theorem 8 (the CP-nets completeness result for flipping sequences) of (Boutilier *et al.* a) for (ii); and for (iii): the TCP-nets completeness result: see lemma 5 of (Brafman, Domshlak, & Shimony 2002). For example, to prove (iii): α covers β with respect to \succ_M implies, by lemma 5 of (Brafman, Domshlak, & Shimony 2002), that there exists a worsening flipping sequence from α to β ; but since α covers β , there can be no element in the sequence between α and β , so β is a worsening TCP-net flip from α . Therefore α and β differ on either one or two variables, according to whether it's a CP-flip or an I-flip (see definition 4 of (Brafman, Domshlak, & Shimony 2002)).

In the example, there are a pair of outcomes, nmd and $1of$, which are consecutive in the preference order $>_\Gamma$ that differ on all three variables. The lemma then implies that the preferences in the example cannot be represented by a CP-net or TCP-net, i.e., there's no CP-net or TCP-net N on V with $>_N = >_\Gamma$.

A consequence of the above lemma is that, except in some trivial cases, if N is a CP-net or a TCP-net, then \succ_N is never a lexicographic order. This is because lexicographic orders on n variables include consecutive elements that differ on all n variables (assuming the domain of each variable has more than one element). To illustrate this, consider the case of boolean variables and the order on complete tuples being just the usual order of binary numbers. Then $(1, 0, 0, \dots, 0)$ and $(0, 1, 1, \dots, 1)$ are consecutive in the order, but they differ on all the variables. Therefore, by the lemma the order cannot be generated by a CP-net if $n > 1$ and the order cannot be generated by a TCP-net if $n > 2$.

Proposition 3 Let $>_{lex}$ be a lexicographic order (as defined above) on \underline{V} , where the domain of each variables contains more than one element, i.e., for all $X \in V$, $|\underline{X}| > 1$. Then (a) if $|V| > 1$, there exists no CP-net N on V with $\succ_N =$

$>_{lex}$; (b) if $|V| > 2$, there exists no TCP-net M on V with $\succ_M = >_{lex}$.

Representing stronger conditional preferences

Lexicographic orders are a very special type of order, but the kind of statements they represent can be natural. Let \succ be a strict partial order (i.e., a transitive irreflexive relation) on \underline{V} . Let $X \in V$ and $W \subseteq V - \{X\}$ and let $T = V - \{X\} - W$, so that $\{X\}$, W and T partition V . Let $>_X$ be a non-empty partial order on \underline{X} . We say X (unconditionally) dominates W with respect to $(\succ, >_X)$ if the following condition holds: for $\alpha, \beta \in \underline{V}$, $\alpha \succ \beta$ holds whenever α and β are such that: $\alpha(X) >_X \beta(X)$ and $\alpha(T) = \beta(T)$. In particular, if X dominates $W = V - \{X\}$ with respect to $(\succ, >_X)$, then a sufficient condition for $\alpha \succ \beta$ is $\alpha(X) >_X \beta(X)$. This is a stronger form of preference statement than *ceteris paribus* statements. It represents a situation where the value of variable X is much more important than the values of any other variable; we prefer any outcome that does better on variable X .

This kind of condition is naturally represented within the language \mathcal{L} . Let $\Theta = \{(\top : x > x' [W]) : x >_X x'\}$. Then, if $\Gamma \supseteq \Theta$, X dominates W with respect to $(>_\Gamma, >_X)$. Such statements can be used to represent a lexicographic order, as shown above.

This type of variable dominance is not at all natural for CP-nets and TCP-nets, as the following two propositions indicate. But it is easy to construct consistent Γ which satisfy the hypotheses of the two propositions (e.g., $\Gamma = \Theta$ for the representation Θ above), or extensions of Θ , in particular, a lexicographic order).

Proposition 4 Consider any satisfiable CP-net N on $V = \{X_1, \dots, X_n\}$ ($n \geq 2$) such that X_2 has no parents and $|\underline{X}_2| > 1$ with associated order \succ_N on \underline{V} . Then for no (non-empty) $>_1$ on \underline{X}_1 is it the case that X_1 dominates $\{X_2, \dots, X_n\}$ with respect to $(\succ_N, >_1)$.

In the example, X_1 dominates $\{X_2, X_3\}$ with respect to $(>_\Gamma, >_1)$, where $n >_1 1$; also X_2 has no parents. The proposition then implies (without looking at the level of outcomes) that there's no CP-net N on V with $\succ_N = >_\Gamma$. It also implies that the same would hold if we were to change the preferences on X_3 in any way.

There is a similar result for TCP-nets:

Proposition 5 Consider any TCP-net M on $V = \{X_1, \dots, X_n\}$ ($n \geq 3$) such that X_2 has no parents and X_3 has no parents, $|\underline{X}_2|, |\underline{X}_3| > 1$ and the associated relation \succ_M on \underline{V} is acyclic. Then for no total order $>_1$ on \underline{X}_1 is it the case that X_1 dominates $\{X_2, \dots, X_n\}$ with respect to $(\succ_M, >_1)$.

Generating precisely a total order on outcomes We finish this section with an expressibility result that can be proved with the help of Theorem 2 and Lemma 1, illustrating how hard it is to generate a CP-net associated to a total order of outcomes. (However, one should not usually expect an agent's preferences to generate a total order, so this should not be considered as a damning criticism of CP-nets.) It shows that once one removes the obvious symme-

tries concerned with variable and value ordering, there is a unique CP-net on a set of boolean variables which generates a total order of outcomes. This contrasts with the situation for conditional preference theories, where there are precisely $2^{2^n - n - 1}$ total orders $>$ with maximum element $(1, 1, \dots, 1)$ on \underline{V} with are equal to some $>_\Gamma$, for Γ such that $G(\Gamma)$ is consistent with the variable ordering X_1, \dots, X_n .

Proposition 6 *There is a unique CP-net N on boolean variables $V = \{X_1, \dots, X_n\}$ satisfying the following properties: (i) the CP-net order $>_N$ is a total order of outcomes with maximum element $(1, \dots, 1)$; (ii) the variable ordering is consistent with the relation H on V associated with N , i.e., $(X_j, X_i) \in H$ implies $j < i$.*

It can be shown furthermore that H is maximally large: $H = \{(X_j, X_i) : j < i\}$ so that the parents set $\text{Pa}(X_i)$ of X_i is $\{X_1, \dots, X_{i-1}\}$. The conditional preference tables (when written out explicitly) are therefore of exponential size. They can be expressed compactly as follows: for each $i = 1, \dots, n$, and assignment u to $\text{Pa}(X_i)$, $1 \succ_u^{X_i} 0$ if and only if u (viewed as a sequence of boolean values) contains an even number of zeros.

Ensuring Consistency

The main purpose of this section is to give simple and natural sufficient conditions for a conditional preference theory to be consistent. There is a clear necessary condition, which we call local consistency, based on considering each variable separately (that always holds for CP-nets). The main result of this section is that given $G(\Gamma)$ is acyclic, Γ is consistent if and only if Γ is locally consistent. (This generalises the acyclicity condition for CP-nets since, for CP-nets, the relations G and H (of the corresponding conditional preference theory) are equal.) To prove this we use a more general form of lexicographic order, which we call a pcl order: a partial conditional lexicographic order. It is like a lexicographic order except that the importance order between variables need only be a partial order, and the order on values of a variable can be partial and conditional on the values of more important variables. A pcl order is a strict partial order. We show that for any locally consistent conditional preference theory Γ such that $G(\Gamma)$ is acyclic, the associated order $>_\Gamma$ on outcomes is a subrelation of a pcl order; this then implies that $>_\Gamma$ is irreflexive, so Γ is consistent by Theorem 1.

Local consistency

In certain cases, it's clear that Γ is not consistent, by just looking at local conditions: if there's a sequence of worsening swaps from some α to itself, which just change the values of a single variable X .

Fix conditional preference theory Γ on V , and consider outcome $\alpha \in \underline{V}$ and variable $X \in V$. Say that pair (x, x') of values of X is validated by α if there exists some statement $(u : x > x' [W]) \in \Gamma$ with $\alpha \models u$ (i.e., u is a projection of α). Define relation $>_\alpha^X$ on \underline{X} to be the transitive closure of the set of all pairs (x, x') validated by α . We say that Γ is locally consistent if for all α and X , $>_\alpha^X$ is irreflexive.

If Γ were not locally consistent then there would exist outcome α , variable X and a sequence x_1, \dots, x_k of values of

X with associated statements in Γ , $(u_i : x_i > x_{i+1} [W_i])$, such that $\alpha \models u_i$, and $\alpha(X) = x_1 = x_k$. This would give a worsening swapping sequence from α to α (only involving changing variable X), thus implying that Γ is not consistent, by Theorem 1. Therefore local consistency is a necessary condition for consistency.

The set of statements Γ in the example is easily seen to be locally consistent. However, if φ_5 were changed to $\varphi'_5 = m : d > f [\emptyset]$ then Γ would no longer be locally consistent as φ'_5 and $\varphi_3 = n : f > d [\emptyset]$ would give conflicting preferences for X_3 under the conditions nm. Let $\alpha = nmd$. Then $>_\alpha^{X_3}$ is not irreflexive since $d >_\alpha^{X_3} f$ using φ'_5 and $f >_\alpha^{X_3} d$ using φ_3 , so $d >_\alpha^{X_3} d$. Γ would no longer be consistent as $>_\Gamma$ is no longer irreflexive: we have $nmd >_\Gamma nmf >_\Gamma nmd$ so $\alpha >_\Gamma \alpha$.

For $X \in V$, let $U_X = \text{Pa}_{H(\Gamma)}(X)$ be the parents⁵ of X with respect to $H(\Gamma)$. Then $>_\alpha^X$ equals $>_\beta^X$ whenever $\alpha(U_X) = \beta(U_X)$, so for any $u \in \underline{U_X}$ we can define relation $>_u^X$ to be $>_\alpha^X$ for any α such that $\alpha(U_X) = u$. Local consistency can then be expressed in terms of these relations: Γ is locally consistent if and only if for all $X \in V$ and $u \in \underline{U_X}$, $>_u^X$ is irreflexive.

Often checking local consistency will be easy; in particular, when the sets $\text{Pa}_{H(\Gamma)}(X)$ are small (as in intended applications of CP-nets) one can efficiently construct each relation $>_u^X$ explicitly. (CP-nets assume these $>_u^X$ relations have already been computed, or directly elicited; they are also assumed to be total orders, so local consistency is guaranteed.) To give another example, when all the variables are binary, local consistency can be determined in time proportional to $|\Gamma|^2|V|$.

Partial conditional lexicographic orders

For $\alpha, \beta \in \underline{V}$, define $\Delta(\alpha, \beta)$ to be the set of variables where α and β differ, i.e., $\{X \in V : \alpha(X) \neq \beta(X)\}$. For relation G on V , let G° be the transitive closure of G , and for $U \subseteq V$, define $\text{min}_{G^\circ}(U)$ to be $\{X \in U : \forall Y \in U, (Y, X) \notin G^\circ\}$, that is, the set of undominated variables in U with respect to G° , i.e., the set of variables in U which have no ancestors in U (with respect to G). If G is acyclic then for any non-empty U , $\text{min}_{G^\circ}(U)$ is non-empty.

A pcl structure on set of variables V is defined to be a tuple $\langle G, H, \{\succ_u^X : X \in V, u \in \text{Pa}_H(X)\} \rangle$, where G and H are acyclic relations on V with $G \supseteq H$, and for each variable $X \in V$, and each $u \in \text{Pa}_H(X)$, \succ_u^X is a (possibly empty) strict partial order (a transitive irreflexive relation) on \underline{X} . Associated with a pcl structure p is a binary relation \succ_p on \underline{V} , called the pcl order, given as follows:

For $\alpha, \beta \in \underline{V}$, $\alpha \succ_p \beta$ if and only if $\alpha \neq \beta$ and for all $X \in \text{min}_{G^\circ}(\Delta(\alpha, \beta))$, $\alpha(X) \succ_u^X \beta(X)$, where $u = \alpha(\text{Pa}_H(X))$ (which equals $\beta(\text{Pa}_H(X))$) since any variable Y in $\text{Pa}_H(X)$ is an element of $\text{Pa}_G(X)$ so is not in $\Delta(\alpha, \beta)$ by minimality of X , so $\alpha(Y) = \beta(Y)$.

⁵The set of parents $\text{Pa}_H(X)$ of X with respect to relation H is defined to be the set of all Y such that (Y, X) is in H .

The idea is that $\min_{G^\circ}(\Delta(\alpha, \beta))$ is the set of most important variables where α and β differ. α is preferred to β if α is better than β on each of these variables. A standard lexicographic order compares two outcomes α and β by considering the most important variable X on which α and β differ, and preferring α to β if $\alpha(X)$ is preferred to $\beta(X)$. So a pcl order generalises a standard lexicographic order by allowing: (i) there to be more than one best variable on which α and β differ, because G° is only a partial order; and (ii) the local preference of $\alpha(X)$ over $\beta(X)$ to be partial and conditional on some of the earlier variables.

Proposition 7 *For any pcl structure p , the associated order \succ_p is transitive.*

Hence \succ_p is a partial order, since it is also irreflexive.

Generating a Partial Conditional Lexicographic Order from Γ . If Γ is locally consistent and such that $G(\Gamma)$ is acyclic then we can generate a pcl structure from Γ . Abbreviate $\text{Pa}_{H(\Gamma)}(X)$ to U_X . For each variable $X \in V$, and each $u \in U_X$, define relation \succ_u^X on X to be just $>_u^X$ (see Local Consistency section). Because Γ is locally consistent, each \succ_u^X is a (strict) partial order, so $\langle G(\Gamma), H(\Gamma), \{\succ_u^X : X \in V, u \in U_X\} \rangle$ is a partial conditional lexicographic structure, and we write the associated order \succ_p as $\succ_{p(\Gamma)}$.

With the example, we have e.g., the following conditional preferences for X_3 : $\succ_{n\circ}^{X_3} = \succ_{nm}^{X_3} = \succ_{1\circ}^{X_3}$ defined by $f \succ_{n\circ}^{X_3} d$, and also $d \succ_{1m}^{X_3} f$. The relation $\succ_{p(\Gamma)}$ is in this case a total order, which extends $>_\Gamma$ with the additional preference $n\circ d \succ_{p(\Gamma)} nmf$, because variable X_2 is more important than X_3 according to $G(\Gamma)$. As we see below, it is a general result that $\succ_{p(\Gamma)}$ extends $>_\Gamma$. It is also the case that given conditional preference theory $\Gamma \subseteq \mathcal{L}$, we can extend Γ to a set $\bar{\Gamma}$ making the associated ordering on outcomes equal to the pcl ordering.

Fix $\Gamma \subseteq \mathcal{L}$. For $\varphi = (u : x > x' [W]) \in \Gamma$ let $\bar{\varphi}$ be $(u : x > x' [W'])$ where W' is the set of descendants of X in $G(\Gamma)$ (i.e., the set of variables Y such that (X, Y) is in the transitive closure of $G(\Gamma)$). Define $\bar{\Gamma} = \{\bar{\varphi} : \varphi \in \Gamma\}$.

Theorem 2 *If Γ is locally consistent and $G(\Gamma)$ is acyclic then $\succ_{p(\Gamma)} = \succ_{p(\bar{\Gamma})} = >_{\bar{\Gamma}} \supseteq >_\Gamma$.*

The most important part of this theorem is the result that $\succ_{p(\Gamma)} \supseteq >_\Gamma$, i.e., that if $\alpha >_\Gamma \beta$ then $\alpha \succ_{p(\Gamma)} \beta$. This can be proved by showing that $\succ_{p(\Gamma)} \supseteq \Gamma^*$, which implies the result, since $\succ_{p(\Gamma)}$ is transitive (by Proposition 5) and $>_\Gamma$ is the transitive closure of Γ^* .

If locally consistent conditional preference theory Γ is such that $G(\Gamma)$ is acyclic then by Theorem 2 $\succ_{p(\Gamma)} \supseteq >_\Gamma$ which implies that $>_\Gamma$ is irreflexive since $\succ_{p(\Gamma)}$ is irreflexive, and hence by Theorem 1, Γ is consistent. So we have the following result (since local consistency is a necessary condition for consistency), which generalises the consistency result for acyclic CP-nets (i.e., CP-nets whose associated relation H is acyclic).

Corollary 1 (Consistency) *Let conditional preference theory Γ be such that $G(\Gamma)$ is acyclic. Then Γ is consistent if and only if Γ is locally consistent.*

This shows that, as long as Γ is chosen so that the associated order $G(\Gamma)$ on variables is acyclic, and the local consistency property is confirmed, then Γ is guaranteed to be consistent and the associated order on outcomes will be irreflexive (and hence acyclic). It gives an agent considerable flexibility in making their preference statements, without risking inconsistency.

Further Applications

We consider some further consequences of our results.

Choosing a total order compatible with Γ For some applications, one does not need to determine $>_\Gamma$ precisely; it is sufficient to be able to list outcomes in an order compatible with $>_\Gamma$, i.e., compatible with the preferences expressed by Γ . Theorem 2 shows that this is easy (given that Γ is locally consistent and $G(\Gamma)$ is acyclic), since we can pick an ordering compatible with $\succ_{p(\Gamma)}$, and hence compatible with $>_\Gamma$. We list the variables in an order X_1, \dots, X_n compatible with $G(\Gamma)$ and we extend each local partial order relation $>_u^X$ to a total order on X (this can be done implicitly using a default ordering on X). We can then generate outcomes lexicographically: to start with, we pick the best value x_1 of X_1 first and then pick the best value x_2 of X_2 conditional on x_1 , etc. Note that it is also very easy to check, for this total order $>$, if $\alpha > \beta$ for outcomes α and β .

Finding optimal (maximal) outcomes Clearly, one can use the above algorithm to efficiently find an outcome which is $>_\Gamma$ -maximal. It can also be done more directly. Assume that $H(\Gamma)$ is acyclic and Γ is locally consistent (actually, the latter condition can be considerably weakened; c.f. the proof of Theorem 2 in (Domshlak *et al.* 2003)). Write V as $\{X_1, \dots, X_n\}$ where the variable ordering is consistent with $H(\Gamma)$. We say that $x' \in X_i$ is *undominated given y* if there does not exist statement $u : x > x' [W]$ in Γ such that $y \models u$ (i.e., such that u is a projection of y). We can generate (without backtracking) a (in fact, any) $>_\Gamma$ -maximal outcome α as follows: for $i = 1, \dots, n$, we let $\alpha(X_i) = x'$ for any $x' \in X_i$ undominated given $(\alpha(X_1), \dots, \alpha(X_{i-1}))$ (local consistency ensures that there always is such an x'). This can be done in time approximately linear in the size of Γ .

Constrained Optimisation The approach to constrained optimisation described in (Boutilier *et al.* b) can be easily generalised to finding maximal outcomes with respect to $>_\Gamma$ that satisfy a set of constraints. Furthermore, the following result shows that each outcome which is $\succ_{p(\Gamma)}$ -maximal among those satisfying a set of constraints, is also maximal with respect to $>_\Gamma$.

Corollary 2 (of Theorem 2) *Let Γ be a locally consistent conditional preference theory such that $G(\Gamma)$ is acyclic. For $\Omega \subseteq \underline{V}$, if α is $\succ_{p(\Gamma)}$ -maximal in Ω then α is $>_\Gamma$ -maximal in Ω .*

Finding the $\succ_{p(\Gamma)}$ -maximal outcomes satisfying a set of constraints is relatively easy. In particular, we can modify the complete algorithm given in section 3.1 of (Boutilier *et al.* b) by replacing each (generally hard) dominance test

$\alpha \succ_N \beta$ by the test $\alpha \succ_{p(\Gamma)} \beta$, which is easy because of the lexicographic-style construction of $\succ_{p(\Gamma)}$. To determine if $\alpha \succ_{p(\Gamma)} \beta$ or not, we consider the set $\Delta(\alpha, \beta)$ of all variables on which α and β differ; we find all variables X which are minimal in $\Delta(\alpha, \beta)$ with respect to the transitive closure of $G(\Gamma)$, and we check the local condition $\alpha(X) \succ_u^X \beta(X)$, where $u = \alpha(\text{Pa}_H(X))$.

For moderate to large problems, often even the number of $\succ_{p(\Gamma)}$ -maximal outcomes will be very large, so that we can enumerate more than enough $\succ_{p(\Gamma)}$ -maximal outcomes; in such cases there may be little advantage in using the much less efficient complete algorithm. These remarks of course also apply when we restrict to CP-nets.

Searching for Swapping Sequences If one wants to prove that $\alpha \succ_\Gamma \beta$, one may well need to search for a worsening swapping sequence from α to β . Since these generalise flipping sequences, this can be a hard problem, as shown in (Boutilier *et al.* a; Brafman & Domshlak 2002). The following proposition is useful in restricting the swaps that one need consider, as only certain variables need be changed.

Proposition 8 (prefix and suffix fixing) *Let conditional preference theory Γ be locally consistent and such that $G = G(\Gamma)$ is acyclic. Suppose $\alpha \succ_\Gamma \beta$. Define R to be the set of variables $X \in V$ such that $\alpha(X) = \beta(X)$ and for all ancestors Y in G of X , $\alpha(Y) = \beta(Y)$. Define S to be the set of variables $X \in V$ such that $\alpha(X) = \beta(X)$ and for all descendants Z in G of X , $\alpha(Z) = \beta(Z)$. Then there exists a worsening swapping sequence from α to β in which the values of $R \cup S$ remain constant. Furthermore, in any worsening swapping sequence $\alpha = \alpha_1, \dots, \alpha_l = \beta$ the values of R remain constant, i.e., for all $k = 1, \dots, l$, and for all $X \in R$, $\alpha_k(X) = \alpha(X) = \beta(X)$.*

This means that when searching for a worsening swapping sequence from α to β we need only consider swaps that don't change the values of $R \cup S$. Suffix fixing generalises the CP-nets property described in (Boutilier *et al.* 1999; a). Prefix fixing also generalises a property of CP-nets, and can be seen to be revealing regarding the structure of \succ_Γ .

Conclusion

In this paper, a logic of conditional preferences is defined, with a language which allows the compact representation of certain kinds of statements of conditional preference. It is shown that the language can express CP-nets, and that the semantics and proof theory generalise those of CP-nets. The formalism also generalises other important properties of CP-nets; maximal outcomes can be efficiently generated, and there are simple sufficient conditions for consistency. It is also easy, under such conditions, to find a total order on outcomes compatible with the conditional preference order, and a similar approach to constrained optimisation can be used as for CP-nets. Along the way, a number of results were given illustrating the restrictive expressive power of CP-nets and TCP-nets.

Despite being a substantially more expressive language than CP-nets, it is still quite restrictive (with these restric-

tions allowing some good properties). There are many natural ways of augmenting the language to allow the compact representation of other kinds of preferences. For example, one could allow statements of the form $u : s > s' [W]$ where s and s' are assignments to a set of variables S , rather than just a single variable X . One could allow statements of the form $u : s > * [W]$ meaning that conditional on u , s is the most preferred assignment of variables S . The language might also be extended to allow the representation of indifference between values or partial tuples.

This paper has focused mainly on CP-nets rather than their extension TCP-nets. (Wilson 2004) considers the relationship with TCP-nets, which can be expressed with statements $u : x > x' [W]$ with $|W| = 0$ or 1 , and derives similar results to this paper under weaker acyclicity conditions.

Acknowledgements

This work has received support from Science Foundation Ireland under Grant 00/PL1/C075. I am grateful to Rick Wallace, Paul Snow and Barry O'Sullivan for helpful discussions.

References

- Boutilier, C.; Brafman, R. I.; Domshlak, C.; Hoos, H.; and Poole, D. CP-nets: A tool for reasoning with conditional *ceteris paribus* preference statements. *To appear in JAIR*.
- Boutilier, C.; Brafman, R. I.; Domshlak, C.; Hoos, H.; and Poole, D. Preference-based constrained optimization with CP-nets. *To appear in Computational Intelligence, Special Issue on Preferences*.
- Boutilier, C.; Brafman, R. I.; Hoos, H. H.; and Poole, D. 1999. Reasoning with conditional *ceteris paribus* preference statements. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence (UAI99)*, 71–80.
- Brafman, R. I., and Domshlak, C. 2002. CP-nets—reasoning and consistency testing. In *Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning (KR2002)*.
- Brafman, R. I.; Domshlak, C.; and Shimony, E. 2002. Introducing variable importance tradeoffs into CP-nets. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI02)*.
- Domshlak, C.; Rossi, F.; Venable, K.; and Walsh, T. 2003. Reasoning about soft constraints and conditional preferences: complexity results and approximation techniques. In *Proc. IJCAI03*.
- Lang, J. 2002. From preference representation to combinatorial vote. In *Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning (KR2002)*, 277–288.
- Wilson, N. 2004. Consistency and optimisation for conditional preferences. *Unpublished report*.