Conditional Preferences Based on Triadic Modal Logic*

Xiaoping Chen
Computer Science Department
University of Science &
Technology of China
Hefei 230027, China

Xiaowu Li
Institute of Logic and Cognition
Zhongshan University
Guangdong 510275
China

Jianmin Ji
Computer Science Department
University of Science &
Technology of China
Hefei 230027, China

Abstract

We concern ourselves in this paper with formalization of *conditional preferences* based on modal logic. We indicate four basic requirements for the task and show that meeting them jointly presents a challenge to existing dyadic modal logics of preferences. As an alternative, we introduce a minimal system, MCP, of triadic modal logic, which is consistent and also frame sound and frame complete with respect to its neighborhood semantics. A simple yet powerful mechanism of preference optimization is also developed on the basis of MCP, which embodies the criteria of maximal partial satisfaction, specificity, and positive utility. Our investigation in the paper shows that this formalization meets all of the requirements to a certain extent.

1 Introduction

Investigations into desires and preferences draw increasingly attention in recent years [Boutilier, 1994; Tan and Pearl, 1994; Brafman and Tennenholtz, 1997; Lang *et al.*, 2003; Eiter *et al.*, 2003; Doyle, 2004; Brewka *et al.*, 2004]. There are two basic issues in the investigation with which we concern ourselves in this paper:

(1) The combination of preferences. Suppose a set of "primitive" preferences are given by the user. What are the "compound" preferences induced by the primitive ones and how can they be derived?

(2) The optimization of preferences. Given a set of preferences. How to derive the most preferred options or desires according to the preference premise?

The first issue is by and large independent of the realization of representing and reasoning about preferences, as pointed in [Delgrande *et al.*, 2004]. The second one can also be investigated in a way independent of realization. In this paper we will introduce a novel modal logic to formalize preference combination of some kind and, on top of it, establish a mechanism for the optimization of preferences.

*This work is supported by NNSF of China (60275024) and 973 program (2003CB317000).

In the literature there are lots of works on dyadic modal logics of desires and preferences [Boutilier, 1994; Tan and Pearl, 1994; Brafman and Tennenholtz, 1997; Doyle and Thomason, 1999; Weydert and van der Torre, 1998; Dastani *et al.*, 2001; Lang *et al.*, 2003]. It turns out, however, that almost all proper axioms of these logics result in some counterintuitive consequences. The thesis of the paper is that dyadic modal logic is unable to support adequate formalization of conditional preferences (eg., "if A then I prefer B to C"), although they work well for that of conditional desires (eg., "if A then I desire B") and that of relative preferences (eg., "I prefer A to B"). Some researchers noticed the issue of conditional preferences (eg., [Tan and Pearl, 1994]). Recently, the necessity of representing conditional preferences explicitly in ternary relations is accepted in the literature [Boutilier *et al.*, 2004; Brewka *et al.*, 2004]. However, there are no existing modal logic systems embodying the realization as far as the authors know.

In this paper we propose introducing triadic modal logic as a new tool for the task. In section 2 we propose four requirements for an adequate formalization of conditional preferences and analyze the reason why dyadic modal logics cannot meet them. In section 3 we introduce a triadic modal logic, TML, with neighborhood semantics and a minimal TML system, MCP, for conditional preferences. On the basis of MCP we establish in section 4 a mechanism for preference optimization, by which one can derive the most preferred options. We also compare our approach with previous ones whenever convenient through the development of the paper.

2 Analysis

Existing modal logics of desires/preferences are generally established in accordance with two semantic principles, (PRP) and (PCD), respectively. Let $\rightarrow$ denote the dyadic modal operator, $|A|$ the set of worlds satisfying the proposition A, and $<_P$ a binary relation over propositions. The two principles can be described as follows.

PCD $A \rightarrow B$ is true if $|A \wedge \neg B| <_P |A \wedge B|$;

PRP $A \rightarrow B$ is true if $|A| <_P |B|$.

According to (PCD), the intuitive meaning of a conditional $A \rightarrow B$ is that within all A-worlds, any B-world is preferred to any $\neg B$-world. Therefore, asserting a conditional $A \rightarrow B$

means establishing a *conditional desire* connecting a condition (or a context) A and a potential goal B. On the other hand, the intuitive meaning of $A \to B$ under (PRP) is that B is preferred to A without conditions; thus to put down a conditional $A \to B$ is to establish an *unconditional relative preference* between A and B. Most of the existing theories are based on (PCD) [Wellman and Doyle, 1991; Boutilier, 1994; Tan and Pearl, 1994; Weydert and van der Torre, 1998; Dastani *et al.*, 2001; Lang *et al.*, 2003], calling them logics of conditional desires and those based on (PRP) logics of relative preferences.

On the other hand, we can list four requirements for adequate modal logic of conditional preferences.

Context-dependence The preferences over options should be dependent on the context; ie., generally any preference order $<_X$ is a function of context X. In particular, if both $A <_X B$ and $B <_{X \wedge Y} A$ hold, then the former should be defeated by the latter in context $X \wedge Y$.

Transitivity if $A <_X B$ and $B <_X C$ hold for any X, it should be derived that $A <_X C$.

Composition Under some conditions preferences under different contexts X and Y can be combined to form a new one under context $X \wedge Y$. This is a real challenge to the formalization of conditional preferences.

Conflict tolerance It is well known that desires and preferences can conflict with one another in some sense. Hence adequate formalization should be tolerant of these conflicts, ie., be able to represent them consistently.

We believe that all of these requirements are necessary, although there are some conflicts among them. A basic observation is that dyadic modal logics, whether based on (PCD) or (PRP), do not meet these requirements jointly. Consider the logics of conditional desires. Since the law of transitivity does not hold in these logics, one has to employ that of restricted transitivity instead:

RT $(A \to B) \wedge (A \wedge B \to C) \supset (A \to C)$

However, this would damage context-dependence seriously.

Example 2.1 (Physicist) *Suppose* Γ *consists of*

(1) $p \to w$

(2) $p \to \neg s$

(3) $p \wedge w \to s$

where p*,* w *and* s *represent to be a physicist, to win a Nobel Prize and to study AI, respectively. Then a consequence,* $p \to s$*, can be deduced from (1) and (3) by* **RT** *in these logics, with only exception of [Lang* et al.*, 2003]. This is intuitively irrational. Clearly, an agent with the desires in* Γ *wants to study AI only after he/she has won a Nobel Prize, but the application of* **RT** *destroys the dependence of* s *on its true context* $p \wedge w$ *and fakes unnecessarily a contradiction between* s *and* $\neg s$ *under the context* p*. Of course this defect can be avoided by expressing the preferences in a temporal language. But that is unnecessarily complicated for many applications.*

In the logics of conditional desires, the only way to represent transitivity is to employ the axiom

DT $(B \vee C \to B) \wedge (C \vee D \to C) \supset (B \vee D \to B)$

in its variant form

$$(A \wedge (B \vee C) \to A \wedge B) \wedge (A \wedge (C \vee D) \to A \wedge C) \supset (A \wedge (B \vee D) \to A \wedge B)$$

with A representing the context. However, this simple strategy suffers from some serious shortcomings. A fatal and obvious drawback is that it will inevitably cause confusion between $X \wedge A < X \wedge B$ and $A <_X B$. For example, suppose a user prefers *richness and reputation* to *richness and health*, while preferring *health* to *reputation* in the context of *richness*. Then the two preferences should be expressed as *richness*$\wedge$*health*$<$*richness*$\wedge$*reputation* and *richness*$\wedge$*reputation*$<$*richness*$\wedge$*health*, respectively, according to the strategy. Therefore, similar to the case of Example 2.1, there will be a faked conflict between the two preferences.

Another serious shortcoming of previous modal logics of desires and preferences is their insufficient expressive ability. It is well accepted that desires and preferences can be conflict with one another in some sense, such as that one in Example 2.2. Therefore, these preferences should be expressed as consistent sets in preference logics so that they can be handled in the logics [Bacchus and Grove, 1996]. However, even some common and intuitively tractable conflicts like that in Example 2.2 cannot be dealt with satisfactorily in a self-contained formalism so far developed.

Example 2.2 (Airplane ticket, [Lang *et al.*, 2003]) *Suppose someone is planning for his/her travel during Christmas vacation. He/She wants to have a ticket to Roman or a ticket to Amsterdam, but not both.*

As for composition, the situation is even worse. In dyadic logics of conditional desires, the major rule of composition is

CC $(A \to B) \wedge (A \to C) \supset (A \to B \wedge C)$

Clearly, it is too weak for our purpose—it cannot deal with the composition of contexts in any case. On the other hand, sometimes it will derive irrational consequences. We will discuss this in detail in section 3.

Similar argument applies to the previous logics of relative preferences. Therefore, we think that dyadic modal logic is unable to support adequate formalization of conditional preferences. In fact, if only a dyadic modal operator is expected to represent conditional preference, it will be inevitably overloaded with a twofold function: it has to stand both for a connection between a context and a potential goal and for preference over outcomes. The analysis above reveals that dyadic conditionals cannot afford to take on both of the roles jointly while satisfying all of the four requirements.

Therefore, we propose introducing triadic modal logic as a new tool for the formalization of conditional preferences. The outline of our proposal is described as follows. First, we establish a triadic modal logic with neighborhood semantics, TML, and a minimal TML system for conditional preferences, MCP. A triadic conditional $A \Rightarrow BC$ is thought of as representing "if A then C is preferred to B". MCP contains only two proper axioms, general conjunction and transitivity of consequents. They prescribe that the operations in MCP are carried out along two dimensions, Boolean and of preference order. Then the most preferred propositions under any context can be derived in the system with some additional machinery or control strategy. This mechanism of optimization embodies three criteria: that of maximal partial satisfaction, specificity, and positive utility. We investigate the performance of and clarify the intuition behind the mechanism by a series of examples.

3 The logic

We will establish TML on the basis of classical proposition logic, denoted by CL, with connectives $\neg$, $\wedge$, $\vee$, $\supset$, and $\equiv$. We introduce a triadic modal operator $\Rightarrow$. The formula of the form $A \Rightarrow BC$ is called a triadic conditional, or conditional for short, which can be informally interpreted as "C is preferred to B if A", where A, B, and C are called antecedent, first consequent, and second consequent of the conditional. The formulae of TML are defined as usual. Sometimes we use the expression of the form $A \Rightarrow B \cdot C$ instead of $A \Rightarrow BC$, especially when B or C is not a literal. The entailment in a logic S, say CL or MCP etc., is denoted by $\vdash_S$. The symbol $\top$ and $\bot$ are taken to be some tautology and contradiction, respectively. In this paper we concentrate on a minimal TML system of conditional preferences, MCP.

Definition 3.1 (MCP) *The axioms and rules of MCP are as follows:*

PT *all propositional tautologies;*

GC $(A_1 \Rightarrow B_1C_1) \wedge (A_2 \Rightarrow B_2C_2) \supset (A_1 \wedge A_2 \Rightarrow B_1 \wedge B_2 \cdot C_1 \wedge C_2)$;

CT $(A \Rightarrow BC) \wedge (A \Rightarrow CD) \supset (A \Rightarrow BD)$;

MP A, $A \supset B$ / B;

RAE $A_1 \equiv A_2$ / $(A_1 \Rightarrow BC) \equiv (A_2 \Rightarrow BC)$;

RFCE $B_1 \equiv B_2$ / $(A \Rightarrow B_1C) \equiv (A \Rightarrow B_2C)$;

RSCE $C_1 \equiv C_2$ / $(A \Rightarrow BC_1) \equiv (A \Rightarrow BC_2)$.

GC generates and strengthens **CC** of traditional conditional logics, which prescribes the law of general conjunction. **CT** describes the law of transitivity of consequents.

The entailment of MCP is defined as usual.

Theorem 3.1 (Consistency) *The system MCP is consistent; ie., there is no TML-formula A such that $\vdash_{MCP} A$ and $\vdash_{MCP} \neg A$.*

Definition 3.2 (TML frames and models) $F = \langle W, N\rangle$ *is a TML-frame, if W is a non-empty set of possible worlds and N is a mapping from W to $P(P(W) \times P(W) \times P(W))$, where $P(W)$ is the power set of W. $M = \langle W, N, V\rangle$ is a TML-model, if $\langle W, N\rangle$ is a TML-frame and V is a mapping from the set of TML-formulae to $P(W)$ such that for all $w \in W$:*

1. $w \in V(\neg A)$ *iff* $w \notin V(A)$;
2. $w \in V(A \wedge B)$ *iff* $w \in V(A)$ *and* $w \in V(B)$;
3. $w \in V(A \Rightarrow BC)$ *iff* $\langle V(A), V(B), V(C)\rangle \in N(w)$.

For convenience we will not distinguish $V(A)$ from A hereafter. The mapping N in the semantics associates with each possible world w a ternary relation $B <_A C$, asserting informally that according to w proposition C is preferred to B under context A. A feature of our logic is that the ternary relation is not defined over worlds, but sets of worlds. We argue that an agent with bounded rationality can only have preferences at best in this coarse level.

Definition 3.3 (MCP frames and models) $F = \langle W, N\rangle$ *is a MCP-frame, if for all $w \in W$ and subsets of W, X, Y, Z, U, with or without subscripts:*

(gc) *if* $\langle X_1, Y_1, Z_1\rangle \in N(w)$ *and* $\langle X_2, Y_2, Z_2\rangle \in N(w)$, *then* $\langle X_1 \cap X_2, Y_1 \cap Y_2, Z_1 \cap Z_2\rangle \in N(w)$;

(ct) *if* $\langle U, X, Y\rangle \in N(w)$ *and* $\langle U, Y, Z\rangle \in N(w)$, *then* $\langle U, X, Z\rangle \in N(w)$.

$M = \langle W, N, V\rangle$ *is a MCP-model, if* $\langle W, N\rangle$ *is a MCP-frame.*

Intuitively, the ternary relation $B <_A C$ can be taken to be the qualitative and order-preserving representation of some utility function μ_A defined over possible outcomes and parameterized in A. The only constraints on the relation in any MCP frame are (gc) and (ct). The latter specifies that the order $B <_A C$ is transitive for any parameter A and the former that the order is closed under intersections over its parameter A and two variables, B and C. While (ct) reflects directly the corresponding characteristic of classical utility functions, the synthesis of (ct) and (gc) constitutes the truly unique assumption of our approach that the preferential order is closed under transitivity and the intersection described above, denoted as (ATI) hereafter. This assumption is rather strong indeed and not necessarily true of all situations. However, it is not stronger or less rational than the corresponding rule, **CC/And**, in traditional conditional logics of desires and preferences. More importantly, we will discover later that (ATI) provides a solid foundation for an appropriate formalism of preference optimization.

Any TML-formula A is valid in a MCP-model $M = \langle W, N, V\rangle$, denoted by $M \models A$, if $V(A) = W$. A is valid in a MCP-frame $F = \langle W, N\rangle$, denoted by $F \models A$, if $M \models A$ for any $M = \langle W, N, V\rangle$. A is valid in fr(MCP), denoted by fr(MCP) $\models A$, if $F \models A$ for any $F \in$ fr(MCP), where fr(MCP) denotes the class of all MCP-frames.

Theorem 3.2 (Frame soundness) *The system MCP is sound in fr(MCP); ie., for any TML-formula A, if $\vdash_{MCP} A$ then fr(MCP) $\models A$.*

Theorem 3.3 (Frame completeness) *The system MCP is complete with respect to fr(MCP); ie., for any TML-formula A, if fr(MCP) $\models A$ then $\vdash_{MCP} A$.*

Obviously, the logic is rather simple. But one can see that it is very powerful with respect to our purpose. Consider the "puzzle of airplane ticket" mentioned above.

Example 3.1 (Airplane ticket, continued**)** *In our logic these conditional preferences can be expressed as $\Gamma = \{\top \Rightarrow \neg rr, \top \Rightarrow \neg aa, \top \Rightarrow r \wedge a \cdot \neg(r \wedge a)\}$, where r and a stand for having an airplane ticket to Roman* *and having an airplane ticket to Amsterdam, respectively. Then one can derive in our logic following consequences:* $\top \Rightarrow \neg r \wedge \neg a \cdot r \wedge a$, $\top \Rightarrow r \wedge a \cdot \neg r \wedge a$ *and* $\top \Rightarrow r \wedge a \cdot r \wedge \neg a$.

At first glance, it seems that the first consequence does not completely adhere to our intuitive understanding of Γ. This may be controversial. We give some justification as follows. (1) In our logic $r \wedge a$ is not the most preferred option in the example. In fact, one can derive in our logic that most preferred options are $\neg r \wedge a$ and $r \wedge \neg a$. Clearly, this conforms to our intuition completely. (2) Even the first consequence is certainly more rational in intuition than that derived by **CC** in existing conditional logics, where the desires of the planner are expressed as $\Gamma^* = \{\top \rightarrow r, \top \rightarrow a, \neg(\top \rightarrow r \wedge a)\}$. By **CC** from the first two conditionals one can derive $\top \rightarrow r \wedge a$. This causes two serious problems. Firstly, Γ^* is inconsistent and thus cannot be represented in these logics. And the situation would be the same if the third conditional was expressed as $\top \rightarrow \neg(r \wedge a)$. Secondly, according to (PCD), $\top \rightarrow r \wedge a$ means that $r \wedge a$ is preferred to $\neg(r \wedge a)$; in other words, any $r \wedge a$-world is preferred to any $\neg r$-world or $\neg a$-world. But this is absolutely irrational, because the most preferred in this example should be $\neg r \wedge a$-worlds or $r \wedge \neg a$-worlds.

(ATI) needs to be further justified when one of consequents of a conditional is a contradiction. Suppose both $A \Rightarrow BC$ and $A \Rightarrow B\neg C$ are given as an agent's conditional preferences. Then one can derive in our logic a consequence $A \Rightarrow B \cdot \neg C \wedge C$. Since $\neg C \wedge C$ is interpreted in our semantics as an empty set of possible worlds, it corresponds to zero utility. Then the consequence can be interpreted as specifying that B corresponds to some minus utility. This conforms completely to our intuition to the case: if an agent prefers both C and $\neg\, C$ to B under the same context, then he/she will reject B under the context. Correspondingly, our machinery for the optimization of preferences will reject all of the options with minus utilities. This is *the criterion of positive utility*. Therefore, non-positive utility provides a simple and flexible conceptual tool for coping with contradictory options. More significantly, the discussion above also reveals that the essence of (ATI) is the assumption that the derivation of MCP will preserve the order of utilities of comparative options.

In order to compare **MCP** with previous conditional logics of desires and preferences, consider following theorems.

Proposition 3.1

RTC $(A \Rightarrow BC) \wedge (A \Rightarrow B \wedge C \cdot D) \wedge (A \Rightarrow BB) \supset (A \Rightarrow BD)$;

ACC $(A \Rightarrow B \cdot C \wedge D) \wedge (A \Rightarrow CD) \wedge (A \Rightarrow DD) \supset (A \Rightarrow B \wedge C \cdot D)$;

LCC $(A \Rightarrow B \wedge C \cdot D) \wedge (A \Rightarrow BC) \wedge (A \Rightarrow BB) \supset (A \Rightarrow B \cdot C \wedge D)$.

RTC, **ACC**, and **LCC** are the corresponding form of **RT**, **AC**, and the conjunction rule of Lang et al. [Lang *et al.*, 2003], respectively. Note that all antecedents of the conditionals appeared in each theorem are fixed. This feature makes them free from those deficiencies of their traditional counterparts described in section 2.

Example 3.2 (Physicist, continued**)** *Now the desires are re-expressed in MCP as* $\Gamma' = \{p \Rightarrow \neg ww,\ p \Rightarrow s\neg s,\ p \wedge w \Rightarrow \neg ss\}$. *Then following consequences can be derived, among which there is no counterintuitive case:* $p \Rightarrow \neg w \wedge s \cdot w \wedge \neg s$, $p \wedge w \Rightarrow \neg w \wedge \neg s \cdot w \wedge s$, *and* $p \wedge w \Rightarrow \bot\bot$. *Consequently,* $w \wedge \neg s$ *is most preferred under context* p *and* $w \wedge s$ *under* $p \wedge w$.

A second feature of **RTC**, **ACC** and **LCC** is that they all need additional conditions in the form of $A \Rightarrow BB$. In fact, since our logic is neutral to reflexivity/irreflexivity, any formula in this form means more than identity in our logic. Basically, it means that B is not unacceptable as a most preferential desire or "goal". For example, according to our definitions of goals in the next section, B is a goal under $\langle\{A \Rightarrow BB\}, A\rangle$, while there is none under $\langle\emptyset, A\rangle$.

We do not believe that MCP can be weaken while still meet the four requirements to the least extent. On the other hand, one can add more axioms such as reflexivity or irreflexivity into our logic. In this sense MCP is a minimal TML system of conditional preferences.

4 Optimization of preferences

Suppose Γ is a set of TML-formulae consisting of triadic conditionals and static domain knowledge expressed in CL-formulae, and θ is a CL-formula representing the belief about the current situation. In order to find the most preferential options, called goals hereafter, with respect to Γ and θ one needs some machinery for optimization of preferences. Generally, this machinery can be regarded as a meta-level control strategy of some preference logic, by which a goal can be deduced as a consequence from the premises under some restrictions. In this section, we specify these restrictions for the optimization of preferences in MCP. We clarify the intuition behind the definitions through a series of examples.

Example 4.1 (Raining and taking umbrella) *Let* $\Gamma_1 = \{r \Rightarrow ud, r \Rightarrow du\}$*, where* r*,* d *and* u *stand for raining, driving a car and taking umbrella, respectively. Are the preferences in* Γ_1 *conflict with each other?*

One can give two answers to the question. First, Γ_1 can be understood as specifying that the preference between u and d is the same. This means that there is no conflict between them. Hence both u and d are the maxima of Γ_1 under $<_r$. Generally, let $B <_{A,\Gamma} C$ denote $\Gamma \vdash_{MCP} A \Rightarrow BC$ and $B \leq_{A,\Gamma} C$ denote $B <_{A,\Gamma} C$ or $\vdash_{CL} B \equiv C$. We often omit the subscript Γ for short.

Definition 4.1 (Weak $<_A$-maxima) *The set of weak* $<_A$*-maxima of* Γ *is defined as*

$$WMax(\Gamma, <_A) =_{df} \{\, x \not\leq_A \bot \mid \text{there is } y \text{ such that } y <_A x \text{ and for all } z,\ z <_A x \text{ if } x <_A z \,\}.$$

Therefore, we have that $\text{WMax}(\Gamma_1, <_r) = \{u, d, u \wedge d\}$. The restriction $x \not\leq_A \bot$ implies that any $<_A$-maximum cannot have minus utility. We will take $<_A$-maximum as a necessary condition of weak goals, which excludes any proposition with minus utility from being a weak goal. This is the so-called criterion of positive utility.

The second answer to the question in Example 4.1 is that the preferences between u and d in Γ_1 conflict with each other and thus both are invalid. This notion of goal leads to the definition of strong $<_A$-maxima.

Definition 4.2 (Strong $<_A$-maxima) *The set of strong $<_A$-maxima of Γ is defined as*

$$SMax(\Gamma, <_A) =_{df} \{ x \in WMax(\Gamma, <_A) \mid \text{there is no } y \in WMax(\Gamma, <_A) \text{ such that } x <_A y \text{ and } \nvdash_{CL} x \equiv y \}.$$

Hence $\text{SMax}(\Gamma_1, <_A) = \{u \wedge d\}$. Generally, we have that $\text{SMax}(\Gamma, <_A) \subseteq \text{WMax}(\Gamma, <_A)$. Typically Γ contains conditionals with different strength of antecedents. This implies that Γ can specify multiple preference orders over propositions parameterized in different conditions.

Example 4.2 (Raining and taking umbrella, continued) *Suppose $\Gamma_2 = \Gamma_1 \cup \{r \wedge h \Rightarrow u \vee d \cdot \neg(u \vee d)\}$, where h stands for staying home. Then we have that $SMax(\Gamma_2, <_{r\wedge h}) = \{\neg(u \vee d)\} \neq \{u \wedge d\} = SMax(\Gamma_2, <_r)$.*

Intuitively, goals under given Γ and θ should satisfy two conditions. Firstly, they should be $<$-maxima with respect to Γ and θ; secondly, the antecedents of their corresponding conditionals should be activated by θ. The first condition involves *the criterion of specificity*. Assume $\theta = r \wedge h$. Then the strong $<_{r\wedge h}$-maximum in Example 4.2 should override strong $<_r$-maxima and be selected as the strong goal under Γ and θ. The situation is similar to weak goals. For any CL-formulae θ, A and B, A is called θ-stronger than B if that $\theta \vdash_{CL} A \supset B$ and properly θ-stronger than B if A is θ-stronger than B and B is not θ-stronger than A.

Definition 4.3 (Weak and strong goals) *Given any Γ and θ. The sets of weak and strong goals of Γ under θ, respectively, are defined as*

$$WG(\Gamma, \theta) =_{df} \{ x \in WMax(\Gamma, <_A) \mid \theta \vdash_{CL} A \text{ and for all } A' \text{ properly } \theta\text{-stronger than } A,\ WMax(\Gamma, <_{A'}) = \emptyset \text{ or } \theta \nvdash_{CL} A' \};$$

$$SG(\Gamma, \theta) =_{df} \{ x \in SMax(\Gamma, <_A) \mid \theta \vdash_{CL} A \text{ and for all } A' \text{ properly } \theta\text{-stronger than } A,\ SMax(\Gamma, <_{A'}) = \emptyset \text{ or } \theta \nvdash_{CL} A' \}.$$

For instance, $\text{WG}(\Gamma_1, r) = \{u, d, u \wedge d\}$, $\text{SG}(\Gamma_1, r) = \{u \wedge d\}$, $\text{WG}(\Gamma_2, r \wedge h) = \text{SG}(\Gamma_2, r \wedge h) = \{\neg(u \vee d)\}$.

The mechanism of preference optimization defined so far is very simple. However, it is not powerful enough to deal with some complicated cases.

Example 4.3 (Airplane ticket, continued) *Reconsidering $\Gamma_3 = \{\top \Rightarrow \neg rr, \top \Rightarrow \neg aa, \top \Rightarrow r \wedge a \cdot \neg(r \wedge a)\}$. Then we have that $WMax(\Gamma_3, <_\top) = SMax(\Gamma_3, <_\top) = \{r,\ a,\ \neg(a \wedge r),\ r \wedge \neg a,\ a \wedge \neg r\} = WG(\Gamma_3, <_\top) = SG(\Gamma_3, <_\top)$. Intuitively, it is perfect to select $r \wedge \neg a$ and $a \wedge \neg r$ as goals, although r and a are also acceptable.*

For this purpose we introduce the concept of supporting sets of strong goals. Informally, a supporting set of a strong goal C is the minimal subset of Γ which contains C as one of its strong goals.

Definition 4.4 (Supporting sets) *Given any $x \in SG(\Gamma, \theta)$. $\Delta \subseteq \Gamma$ is a supporting set of x, if $x \in SG(\Delta, \theta)$ and $x \notin SG(\Delta', \theta)$ for any proper subset Δ' of Δ.*

In Example 4.3, for instance, $\{\top \Rightarrow \neg rr\}$ and $\{\top \Rightarrow \neg rr, \top \Rightarrow r \wedge a \cdot \neg(r \wedge a)\}$ are supporting sets of r and $r \wedge \neg a$, respectively, and the latter covers the former. Intuitively, a supporting set of a strong goal C reflects the partial desire specified by Γ that C must satisfy at least. And thus the larger supporting sets of a strong goal are, the more desires of Γ the goal satisfies. Based on this observation, r, a and $\neg(a \wedge r)$ should be defeated by $r \wedge \neg a$ or $a \wedge \neg r$ according to a more strict goal definition conforming to *the criterion of maximal partial satisfaction*. However, only considering subset relation between supporting sets is not sufficient for our purpose, although it is for this example. Note that this example is very difficult to be treated in previous work [Lang *et al.*, 2003], but we want to solve it together with more complicated cases like following one by means of the same mechanism.

Example 4.4 (Airplane ticket, complicated) *Suppose $\theta = d \wedge s$ and $\Gamma_4 = \{d \Rightarrow \neg vv,\ s \Rightarrow \neg ff,\ d \wedge s \Rightarrow v \vee f \cdot b,\ \neg(v \wedge b) \wedge \neg(f \wedge b)\}$, where d and s stand for travelling with daughter and travelling with son, respectively, and v, f and b for visiting Venice, Florence and Beijing, respectively. Then both $v \wedge f$ and b are strong goals with supporting sets $\{d \Rightarrow \neg vv, s \Rightarrow \neg ff\}$ and $\{d \wedge s \Rightarrow v \vee f \cdot b\}$, respectively; but these two supporting sets cannot be distinguished from each other by set inclusion. However, it is easy to see that the former should be defeated by the latter since the antecedent of the third conditional, which belongs to the latter's supporting set, is properly closer to θ than that of the first and the second ones, which constitute the supporting set of the former.*

It follows that one supporting set can be overridden by another according to the conditionals they contain. Now we define covering of supporting sets to capture the overriding between supporting sets and then define superior goals as those strong ones with strongest supporting sets.

Definition 4.5 (Covering) *Given any Γ, θ and $\Delta_1, \Delta_2 \subseteq \Gamma$. Δ_1 covers Δ_2, if (1) $\Delta_2 \subseteq \Delta_1$; or (2) there is a conditional in Δ_1 such that its antecedent is properly θ-stronger than that of some conditional in Δ_2 and there is no conditional in Δ_2 such that its antecedent is properly θ-stronger than that of any conditional in Δ_1.*

Definition 4.6 (Superior goals) *Given any Γ and θ. The set of superior goals under Γ and θ is defined as*

$$UG(\Gamma, \theta) =_{df} \{ x \in SG(\Gamma, \theta) \mid \text{any supporting set of } x \text{ is not covered by any supporting set of } y \text{ for any } y \in SG(\Gamma, \theta) \}.$$

It is easy to verify that $\text{UG}(\Gamma_3, \top) = \{r \wedge \neg a, a \wedge \neg r\}$ and $\text{UG}(\Gamma_4, d \wedge s) = \{b\}$, both of which completely conform to our intuition and the principles we prescribed for reasoning with conditional preferences.

Property 4.1 *Weak, strong and superior goals are not monotonic, ie., $\Gamma \subseteq \Gamma'$ does not implies $XG(\Gamma, \theta) \subseteq XG(\Gamma', \theta)$, where XG stands for WG, SG or UG.*

5 Conclusions

The main contributions of the paper include three aspects. First, we show that dyadic modal logic is unable to be an adequate logical tool for the formalization of conditional preferences, because it does not meet four basic requirements demanded by the task and causes a lot of confusions and muddles. Second, we introduce triadic modal logic as an alternative and establish a minimal TML system, MCP, for conditional preferences. The design of the system is based on a novel point of view on the relation between the qualitative representation and its quantitative counterpart: regarding desire/preference conflicts as non-positive utilities. As a result, one of the most difficult and crucial issues in previous work, conflict resolving, becomes much easier to be treated in our theory. Third, a very simple yet powerful mechanism of preference optimization is established on the basis of MCP, which embodies the criteria of maximal partial satisfaction, specificity, and positive utility. In summary, our formalization of conditional preferences, as a self-contained theory, meets all of the requirements we prescribe for the task to a certain extent.

As far as the authors know, the only theory in the previous work that can also direct represent conditional preferences is LPOD [Brewka *et al.*, 2004]. However, there are some significant differences between theirs and ours. First, in LPOD the user's preferences are actually expressed in terms of their "disappointment degrees". For instance, if there is a rule $A \times B \leftarrow C$, then A and B have disappointment degrees 1 and 2, respectively, no matter whether there are any other rules with the same body in the program. Therefore, each "complete preferential chain" $A_1 > \cdots > A_n$ in some context C must be specified with a rule $A_1 \times \cdots \times A_n \leftarrow C$. While in our logic the same "chain" can be represented with several conditionals such as $C \Rightarrow A_1A_2$, $C \Rightarrow A_2A_3$, etc. Thus users can easily revise and maintain their conditional preferences in our logic. Secondly, LPOD employs *not* and corresponding machinery to explicitly represent the abnormal conditions of rules and thus realize the principle of specificity with respect to contexts, while our logic does not. With help of this machinery, the computation of LPOD will be more efficient. On the other hand, there seems no obstacle to prevent one from employing the similar machinery to realize the computation of MCP.

This work is only a first step toward adequate formalization of conditional preferences in triadic modal logic. There is a lot of work to do in the future. For example, besides what mentioned above, it is needed to investigate the performance of our theory in more complicated situations. Also, the minimal system MCP should be expanded to become more powerful. Most importantly, we should clarify the whole principle governing the composition of contexts in reasoning with conditional preferences.

Acknowledgments

The authors thank Fangzhen Lin, Jiehui Jiang, Yi Zhou and anonymous reviewers for their invaluable comments on the draft of the paper.

References

[Bacchus and Grove, 1996] Fahiem Bacchus and Adam J. Grove. Utility independence in a qualitative decision theory. In *Proceedings of KR-96*, pages 542–552, 1996.

[Boutilier *et al.*, 2004] Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos, and David Poole. Cp-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res. (JAIR)*, 21:135–191, 2004.

[Boutilier, 1994] Craig Boutilier. Toward a logic for qualitative decision theory. In *Proceedings of KR-94*, pages 75–86, 1994.

[Brafman and Tennenholtz, 1997] Ronen I. Brafman and Moshe Tennenholtz. Modeling agents as qualitative decision makers. *Artif. Intell.*, 94(1-2):217–268, 1997.

[Brewka *et al.*, 2004] Gerhard Brewka, I. Niemelä, and T. Syrjänen. Logic programming with ordered disjunction. *Computational Intelligence*, 20(2):335–357, 2004.

[Dastani *et al.*, 2001] Mehdi Dastani, Zhisheng Huang, and Leendert van der Torre. Dynamic desires. In *Game Theory and Decision Theory in Agent-Based Computing*, volume 5, pages 65–79, 2001.

[Delgrande *et al.*, 2004] James P. Delgrande, Torsten Schaub, and Hans Tompits. Domain-specific preferences for causal reasoning and planning. In *Proceedings of KR-04*, pages 673–682, 2004.

[Delgrande, 2003] James P. Delgrande. Weak conditional logics of normality. In *Proceedings of IJCAI-03*, pages 873–878, 2003.

[Doyle and Thomason, 1999] Jon Doyle and Richmond H. Thomason. Background to qualitative decision theory. *AI Magazine*, 20(2):55–68, 1999.

[Doyle, 2004] Jon Doyle. Prospects for preferences. *Computational Intelligence*, 20(2):111–136, May 2004.

[Eiter *et al.*, 2003] Thomas Eiter, Wolfgang Faber, Nicola Leone, and Gerald Pfeifer. Computing preferred answer sets by meta-interpretation in answer set programming. *TPLP*, 3(4-5):463–498, 2003.

[Lang *et al.*, 2003] Jérôme Lang, Leendert W. N. van der Torre, and Emil Weydert. Hidden uncertainty in the logical representation of desires. In *Proceedings of IJCAI-03*, pages 685–690, 2003.

[Tan and Pearl, 1994] Sek-Wah Tan and Judea Pearl. Specification and evaluation of preferences under uncertainty. In *Proceedings of KR-94*, pages 530–539, 1994.

[Wellman and Doyle, 1991] Michael P. Wellman and Jon Doyle. Preferential semantics for goals. In *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI-91)*, volume 2, pages 698–703, Anaheim, California, USA, 1991. AAAI Press/MIT Press.

[Weydert and van der Torre, 1998] Emil Weydert and Leendert W. N. van der Torre. Goals, desires, utilities and preferences. In Jerome Lang, editor, *Proceedings of the ECAI-98 Workshop : Decision theory meets artificial intelligence - qualitative and quantitative approaches*, pages 57–64, Brighton, United Kingdom, 1998. ECAI.