Toward a Logic for Qualitative Decision Theory

Craig Boutilier
Department of Computer Science
University of British Columbia
Vancouver, British Columbia
CANADA, V6T 1Z4
email: cebly@cs.ubc.ca

Abstract

We present a logic for representing and reasoning with qualitative statements of preference and normality and describe how these may interact in decision making under uncertainty. Our aim is to develop a logical calculus that employs the basic elements of classical decision theory, namely probabilities, utilities and actions, but exploits qualitative information about these elements directly for the derivation of goals. Preferences and judgements of normality are captured in a modal/conditional logic, and a simple model of action is incorporated. Without quantitative information, decision criteria other than maximum expected utility are pursued. We describe how techniques for conditional default reasoning can be used to complete information about both preferences and normality judgements, and we show how maximin and maximax strategies can be expressed in our logic.

1 Introduction

We typically expect a rational agent to behave in a manner that best furthers its own interests. However, an artificial agent might be expected to act in the best interests of a user (or designer) who has somehow communicated its wishes to the agent. In the usual approaches to planning in AI, a planning agent is provided with a description of some state of affairs, a *goal state*, and charged with the task of discovering (or performing) some sequence of actions to achieve that goal. This notion of goal can be found in the earliest work on planning and persists in more recent work on intention and commitment [10]. In most realistic settings, however, an agent will frequently encounter goals that it cannot achieve. As pointed out by Doyle and Wellman [12] an agent possessing only simple goal descriptions has no guidance for choosing an alternative goal state toward which it should strive.

Straightforward goal-driven behavior tends to be inflexible: an agent told to ensure that part A and part B are at location L by $5PM$ will be unable to do anything if it cannot locate B or if something prevents it from reaching L by 5PM. One might suppose that the agent should at least deliver A to L as close to 5PM as possible. While such partial fulfillment of deadline goals [16] undoubtedly arises frequently is practice, more general mechanisms will often be required. If A and B can't be delivered, perhaps alternate parts C and D should be; or if the 5PM deadline can't be met, the agent should wait until next week. To this end, a recent trend in planning has been the incorporation of decision-theoretic methods for constructing optimal plans [11]. Decision theory provides most of the basic concepts we need for rational decision making, in particular, the ability to specify arbitrary preferences over circumstances or outcomes. This allows desired outcomes or goals (and hence appropriate behaviors) to vary with context.

Most decision-theoretic analysis is set within the framework of *maximum expected utility* (MEU). One impediment to the general use of such decision-theoretic tools is the requirement to have both numerical probabilities and utilities associated with the possible outcomes of actions. It is quite conceivable that such information is not readily available to the agent. We can often expect users to present information in a *qualitative* manner, including qualitative *preferences* over outcomes (one outcome or proposition is preferred to another) and qualitative *probabilities* (describing the relative likelihood of propositions or outcomes). The ability to reason *directly* with such qualitative constraints is therefore crucial. An appropriate knowledge representation scheme will allow the expression of constraints of this form and allow one to logically derive goals and reasonable courses of action, to the extent the given information allows.[1]

[1]While the foundations of decision theory are, in fact, based on such qualitative preferences [26, 29], the move to numerical utilities (and probabilities) requires that a preferences and likelihoods be calibrated by means of questions concerning acceptable exchanges between outcomes and lotteries. For an agent behaving according to the preferences of some user, this requires that either a) the user's preferences be so completely specified that such calculations can be made; or b) the user (or the source of preference information) be available to be queried about preference information as the need arises. Furthermore, a complete calibration of just the preference ranking, in the most fortunate circumstances, requires a number of queries at least as large as the number of possible worlds (exponential in the number of propositional atoms). Such a mechanism is also often criticized

In this paper, we describe a logic and natural possible worlds semantics for representing and reasoning with qualitative probabilities and preferences, and suggest several reasoning strategies for qualitative decision making using this logic. We can represent *conditional preferences*, allowing (derived) goals to depend on context. Furthermore, these conditional preferences are *defeasible*: I might have a general preference for the proposition A (e.g., that parts be delivered to customers on time) but have a more specific ‘‘defeating’’ preference for $\neg A$ if a customer's account is past due. Semantically, preferences will be captured by an ordering over possible worlds, corresponding to an ordinal value function. The logic that captures such *default* preferences will exactly match existing conditional logics for default reasoning and belief revision [4, 7, 8]. Furthermore, the component of the logic for capturing qualitative probabilities will be isomorphic, with a (separate) *normality* ordering on worlds representing their relative likelihood.

In order to strengthen possible conclusions, we will also present reasoning strategies for completing information about preferences and likelihoods, in essence, making assumptions about unstated constraints. In addition, we describe several ways of making decisions with such completed information. These decision making strategies are motivated by the fact that the scales of normality and preference on which worlds are ranked are incomparable. This reflects the fact that user specified constraints provide qualitative information about the structure of the two rankings, not their relative magnitudes. We will discuss conditions under which decisions are sound in this framework.

In Section 2, we present the basic logic of preferences and its semantics, and show how existing techniques for conditional default reasoning can be used to make various assumptions about incomplete preference orderings. In Section 3, we add normality orderings to our semantics and describe a logic for dealing with both orderings. We describe the derivation of *ideal goal states*, roughly, the best situations an agent can hope for given certain fixed circumstances. This generalizes the usual notion of a goal in AI, for such goals are context-dependent and defeasible, and can be derived from more basic information rather than simply being asserted directly by a user. Such goals do not take into account the ability of an agent to change the fixed circumstances from which they are derived, nor the potential inability of an agent to achieve a goal. In Section 4, we explore a more realistic notion of goal that accounts for a simple form of ability. In planning, as in the decision theory, the ultimate aim is to derive appropriate actions to be performed that will achieve derived goal states. The ability of an agent to affect the world will have a tremendous impact on the *actual goal states* it attempts to achieve. One feature that becomes clear in our model is that, given incomplete knowledge, various behavioral *strategies* can emerge. We show how these can be expressed in our logic. Finally, in Section 5, we point out some related work, and on-going investigations into how the trade-offs between utility and probability can be captured in a qualitative manner. We also point out some interesting connections to deontic logic.

because the queries require answers to which a user does not have ready access or might be uncertain [13].

2 Conditional Preferences

A *goal* is typically taken to be some proposition that we desire an agent to make true. Semantically, a goal can be viewed as a set of possible worlds, those states of affairs that satisfy the goal proposition [10]. Intuitively, if we ignore considerations of ability, the set of goal worlds should be those considered most desirable by an agent (or its designer). To achieve all goals is to ensure that the actual world lies within this desirable set.

Unfortunately, goals are not always achievable. My robot's goal to bring me coffee may be thwarted by a broken coffee maker. Robust behavior requires that the robot be aware of desirable alternatives (‘‘If you can't bring me coffee, bring me tea’’). Furthermore, goals may be defeated for reasons other than inability. It is often natural to specify general goals, but list exceptional circumstances that make the goal less desirable than the alternatives. For instance, a general preference for delivering parts within 24 hours may be overridden when the account is past due (which may in turn be overridden if the customer is important enough). To capture these ideas, we propose a generalization of standard goal semantics. Rather than a categorical distinction between desirable and undesirable situations, we will rank worlds according to their *degree of preference*. The most preferred worlds correspond to goal states in the classical sense. However, when such states are unreachable, a ranking on alternatives becomes necessary. Such a ranking can be viewed as an ordinal value function.

The basic concept of interest will be the notion of *conditional preference*. We write $I(B|A)$, read ‘‘ideally B given A,’’ to indicate that the truth of B is preferred, given A. This holds exactly when B is true at each of the most preferred of those worlds satisfying A. From a practical point of view, $I(B|A)$ means that if the agent (only) knows A, and the truth of A is fixed (beyond its control), then the agent ought to ensure B. Otherwise, should $\neg B$ come to pass, the agent will end up in a less than desirable A-world. The statement can be *roughly* interpreted as ‘‘If A, do B.’’ We propose a bimodal logic CO for conditional preferences using only unary modal operators. The presentation is brief. Further details can be found in [3, 7].

2.1 The Logic CO

We assume a propositional bimodal language $\mathbf{L}_B$ over a set of atomic propositional variables $\mathbf{P}$, with the usual classical connectives and two modal operators $\Box$ and $\overleftarrow{\Box}$. Our possible worlds semantics for preference is based on the class of *CO-models*, of the form $M = \langle W, \leq, \varphi \rangle$, where W is a set of possible worlds, φ is a valuation function, and $\leq$ is a

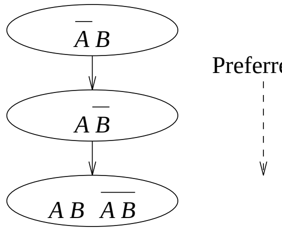

Figure 1: A CO-model

transitive connected binary relation on W.[2] Thus, $\leq$ is a total preorder over W. In other words, W consists of a set of $\leq$-equivalence classes or *clusters* of equally preferred worlds, with these clusters being totally ordered by $\leq$. We take $\leq$ to represent an ordering of preference: $v \leq w$ just in case v is at least as preferred as w. This ordering is taken to reflect the desirability of situations, however this is to be interpreted (e.g., personal utility, moral acceptability, etc.).[3] Figure 1 illustrates a typical CO-model. The truth conditions for the modal connectives are

1. $M \models_w \Box\alpha$ iff for each v such that $v \leq w$, $M \models_v \alpha$.
2. $M \models_w \overleftarrow{\Box}\alpha$ iff for each v such that $w < v$, $M \models_v \alpha$.

$\Box\alpha$ is true at a world w just in case α is true at all worlds at least as preferred as w, while $\overleftarrow{\Box}\alpha$ holds just when α holds at all less preferred worlds. The dual "possibility" connectives are defined as usual: $\Diamond\alpha \equiv_{\text{df}} \neg\Box\neg\alpha$ means α is true at some equally or more preferred world; and $\overleftarrow{\Diamond}\alpha \equiv_{\text{df}} \neg\overleftarrow{\Box}\neg\alpha$ means α is true at some less preferred world. $\overleftrightarrow{\Box}\alpha \equiv_{\text{df}} \Box\alpha \wedge \overleftarrow{\Box}\alpha$ and $\overleftrightarrow{\Diamond}\alpha \equiv_{\text{df}} \Diamond\alpha \vee \overleftarrow{\Diamond}\alpha$ mean α is true at all worlds and at some world, respectively. The logic CO is axiomatized in [3, 7] (see also Section 4).

2.2 Expressing Conditional Preferences

We now define a conditional connective $I(-|-)$ to express conditional preferences. $I(B|A)$ can be read as "In the most preferred situations where A holds, B holds as well," or "If A then ideally B." Intuitively, $I(B|A)$ should hold just when B holds at the most ideal A-worlds.[4] These truth conditions can be expressed in $\mathbf{L}_B$ (see [3, 7]):

$$I(B|A) \;\equiv_{\text{df}}\; \overleftrightarrow{\Box}\neg A \vee \overleftrightarrow{\Diamond}(A \wedge \Box(A \supset B)). \qquad (1)$$

This can be thought of, as a first approximation, as expressing "If A then an agent ought to ensure that B," for making B true ensures an agent ends up in the best possible A-situation. We note that an absolute preference A, capturing the standard unconditional goal semantics, can be expressed as $I(A|\top)$, or equivalently, $\overleftrightarrow{\Diamond}\Box A$. We abbreviate this as $I(A)$ and read this as "ideally A". This can be read as expressing an unconditional desire for A to be true. The model in Figure 1 satisfies $I(B|A)$ and $I(A \equiv B)$.

The dual of preference gives a notion of *toleration* or "don't care conditions." If $\neg I(\neg B|A)$ holds, then in the most preferred A-situations it is not required that $\neg B$. This means there are ideal A-worlds where B holds, or that B is "tolerable" given A. We abbreviate this sentence $T(B|A)$. Loosely, we can think of this as asserting that an agent is *permitted* to do B if A. Unconditional toleration is denoted $T(A)$ and stands for $\neg I(\neg A)$, or equivalently, $\overleftrightarrow{\Box}\Diamond A$.[5] We note that the relative preference of two propositions can be expressed directly in CO. We write $A \leq_P B$ to mean A is at least as preferred as B (intuitively, the best A-worlds are at least as good as the best B-worlds), and define it as:

$$A \leq_P B \equiv_{\text{df}} \overleftrightarrow{\Box}(B \supset \Diamond A)$$

Another useful notion is that of *strict preference*. If some proposition is more desirable than its negation no matter what other circumstances hold (e.g., deliveries to customer C must be on time), we can assert

$$\overleftrightarrow{\Box}(C \supset \Box C)$$

which ensures that every C-world is preferred to any $\neg C$-world. Of course, we cannot *a priori* abolish such situations, for they may occur due to events beyond an agent's control, and the relative preference of these strictly dispreferred worlds is important. But in achieving stated goals condition $\neg C$ will be avoided if at all possible. These strict preferences can also be combined and prioritized [8].

The properties of the connective I are identical to those of the conditional connective $\Rightarrow$ defined in [2, 7] for default reasoning (see also Section 4). They are distinguished simply by their reading and the interpretation of the underlying ordering $\leq$. As one should expect, absolute preferences, as well as preferences in any fixed context, must be consistent, for the following is a theorem of CO (for any possible A):

$$I(B|A) \supset \neg I(\neg B|A)$$

However, an agent's preferences needn't be complete, for $T(B|A) \wedge T(\neg B|A)$ is generally consistent. The property of *preferential detachment* holds in CO:

$$I(B|A) \wedge I(A) \;\supset\; I(B)$$

However, the principle of *factual detachment*

$$I(B|A) \wedge A \;\supset\; I(B)$$

[2] Relation $\leq$ is connected iff $w \leq v$ or $v \leq w$ for all v, w.

[3] While $w < v$ usually means v is a preferred outcome, the usual convention in AI is to "prefer" minimal models, hence we take $w < v$ to mean w is preferred.

[4] Of course, nothing in our models forces the existence of such minimal A-worlds, but our definition is adequate in this case as well [7]. The conditional holds vacuously when A is false at all worlds.

[5] Ideality and toleration are dual in exactly the sense that necessity and possibility are. In deontic contexts, the connectives I and T can be profitably interpreted as expressing some form of obligation and permission, respectively (see Section 5).

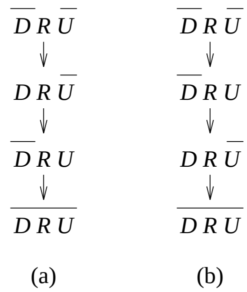

Figure 2: Possible Interpretations of Preferences

is not valid. This has implications for the manner in which an agent should derive its actual preferences in a given situation, as we describe in the next section.

The most important feature is that preferences are conditional and can vary with context. I can consistently assert $I(U|R)$ and $I(\overline{U}|\overline{R})$, that my agent should take an umbrella if it's raining, and leave it home if not. The potential goals or subgoals U or $\neg U$ depend on context and need not be asserted categorically. Furthermore, these conditionals are *defeasible*: I can consistently assert that $I(\overline{U})$ without fear of contradicting $I(U|R)$. Notice that these two statements allow the conclusion $I(\overline{R})$ to be drawn --- the agent can derive its (or *my*) preference for sunny weather.

This defeasibility also allows one to assert, together with the previous conditionals, $I(\overline{U}|R \wedge D)$, that an umbrella is not desired if I drive a car to work (D) instead of walking ($\neg D$). Such a theory induces a partial structure like that illustrated in Figure 2(a). As above, this entails $I(\neg D|R)$. However, is this conclusion truly intended? On the surface, it seems reasonable to accept all three preference statements, but allow the assertion that I prefer to drive when it's raining. Yet $I(D|R)$ contradicts these other premises.

The (intuitive) source of the inconsistency is the statement $I(U|R)$. If I prefer to drive when it's raining, and prefer not to have an umbrella when I drive, then I should not assert that at the most ideal R-worlds, U holds. At the most ideal R-worlds, D (hence $\neg U$) holds. Intuitively, the preference for U given R only holds when I do not drive; thus, $I(U|R \wedge \neg D)$ holds but $I(U|R)$ does not (see Figure 2(b)). Figure 2(a), which validates $I(U|R)$, seems appropriate when I prefer walking to driving, even when it's raining.

We notice, however, that the assertion $I(U|R)$, I prefer an umbrella when it's raining, seems (potentially) appropriate even when Figure 2(b) is the intended model. This might be the case if I am usually unable to drive to work. Even if I prefer to drive, I probably won't be able to, so my stated preference for U given R might reflect this fact. In this case, the *typical* R-world is one in which $\neg D$ holds, and hence one in which U should hold: my robot *should* bring an umbrella along. Very often stated preferences do not express *ideal* preferences. Rather, they may incorporate into the stated context (here, R) certain assumptions or default conclusions (such as $\neg D$), and thus express a preference conditioned on this extended context ($R \wedge \neg D$). The intended assertion $I(U|R \wedge \neg D)$ is perfectly consistent with Figure 2(b), but it may be abbreviated as $I(U|R)$ if the default conclusion $\neg D$ is understood. It is therefore crucial to realize that linguistically stated preferences can be come in different varieties. A statement $I(U|R)$ expresses an ideal preference: in the *best possible* R-worlds U is true. Other varieties, such as those where the user has considered the default consequences of a proposition *before* expressing conditional preference, require the additional machinery we introduce in Sections 3 and 4.[6]

2.3 Defeasible Reasoning with Preferences

The conditional logic of preferences we have proposed above is similar to the (purely semantic) proposal put forth by Hansson [17] for deontic reasoning (reasoning about obligation and permission). In our logic, one may simply think of $I(B|A)$ as expressing a conditional obligation to see to it that B holds if A does. Loewer and Belzer [22] have criticized this semantics ‘‘since it does not contain the resources to express actual obligations and no way of inferring actual obligations from conditional ones.’’ In particular, they argue that any deontic logic should validate something like factual detachment, not just deontic detachment (the deontic analog of preferential detachment). The criticism applies equally well to our preference logic --- one cannot logically derive actual preferences --- because the principle of factual detachment does not hold. Factual detachment expresses the idea that if there is a conditional preference for B given A, and A is *actually* the case, then there is an actual preference for B. While the inference is a reasonable one, we do not expect, nor do we want it to hold logically because it threatens the natural defeasibility of our conditionals. For instance, if R and $I(U|R)$ entailed U, so too would R, D, $I(U|R)$ and $I(\overline{U}|R \wedge D)$. Defeasible conditional preferences could not be expressed.

Various logics have been proposed to capture factual detachment in the deontic setting, and recently several complex default reasoning schemes have been applied to this problem [18, 20]. We propose a simple solution based on the following observation: to determine preferences based on certain actual facts, we consider only the *most ideal* worlds satisfying those facts, rather than *all* worlds satisfying those facts. Let *KB* be a knowledge base containing statements of conditional preference and actual facts. Given that such facts actually obtain, the ideal situations are those most preferred worlds satisfying *KB*. This suggests a straightforward mechanism for determining actual preferences. We simply ask for those propositions α such

[6] Similarly, one can impose this alternate interpretation on direct statements of preference $A <_P B$, as Jeffrey [19] does. On our definition, $A <_P B$ means the *best* A-worlds are preferred, whereas Jeffrey defines such a statement to mean the expected utility of *all* A-worlds is greater than that for B.

that

$$\vdash_{CO} I(\alpha|\mathit{KB})$$

This is precisely the preliminary scheme for conditional default reasoning suggested in [3]. This mechanism unfortunately has a serious drawback: seemingly *irrelevant* factual information, or information about the consequences of actions, can paralyze the derivation of actual preferences.

Example Let P denote that a certain part is painted, B that it's blemished, and S that it's destined for shipment to a specific warehouse. Let D, E and F denote possible locations for a certain piece of equipment. If

$$\mathit{KB} = \{I(P|\overline{B}), \overline{B}\}$$

then the actual preference P is derivable using the scheme suggested above. However, it is not derivable from $\mathit{KB}' = \mathit{KB} \cup \{S\}$. Because conditionals are defeasible, it is consistent (with KB') to assert $I(\overline{P}|\overline{B} \wedge S)$, although intuitively S is irrelevant to this preference.

Again consider *KB* with actual preference P. Suppose a painting action that achieves P requires the equipment in question to be moved, making either D, E or F true. Even though not stated, one can consistently assert $I(P \wedge \overline{D}|\overline{B})$, $I(P \wedge \overline{E}|\overline{B})$ or $I(P \wedge \overline{F}|\overline{B})$. Thus the agent cannot show that any of the moves D, E or F is tolerated --- it cannot decide what to do.

In this example, the fact that $I(P|\overline{B})$ is the only stated preference suggests that other factors are irrelevant to the relative preference of situations. Intuitively, these factors should be discounted. Unless stated otherwise, the part should be painted regardless of its destination and the manner in which P is achieved (D, E or F) is not of concern.

One possible way to deal with this difficulty is to make certain assumptions about the preference ordering. In particular, it is possible to adopt the default reasoning scheme System Z [23] in this context. Given a set of conditional constraints, System Z enforces the assumption that worlds are assumed to be as preferred as possible consistent with these constraints. In other words, worlds are pushed down as far as possible in the preference ordering, ''gravitating'' toward absolute preference. In our example, the model induced by this assumption is shown in Figure 3. (For convenience we assume that $I(\overline{P}|B)$ and that D, E and F are mutually exclusive.) *Any* $\neg B$-world that satisfies P is deemed acceptable, regardless of the truth of the irrelevant factors. The technical details of System Z may be found in [23], and in [3] we describe how the Z-model for any conditional theory can be axiomatized in CO. The important features of this model are: a) the assumption induces a unique, ''most compact'' preference ordering; and b) the consequences associated with these assumptions can sometimes be efficiently computed.

Is the assumption that worlds are preferred unless stated otherwise reasonable? For instance, Tan and Pearl [28]

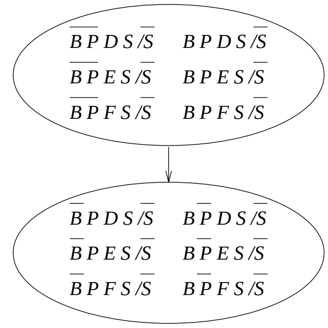

Figure 3: The Compact Preference Ordering

argue that worlds should gravitate toward ''indifference'' rather than preference. We cannot, of course, make sense of such a suggestion in our framework, since we do not have a bipolar scale (where outcomes can be good, bad or neutral).[7] However, even if an ''assumption of indifference'' were technically feasible, we claim that the ''assumption of preference'' is the the right one in our setting.

Recall that we wish to use preferences to determine the set of goal states for a given context C. These are simply the most preferred C-worlds according to our ranking; call this set *Pref*(C). If the agent brings about *any* of these situations, it will have behaved correctly. A conditional preference $I(A|C)$ constrains the set *Pref*(C) to contain only A-worlds. Thus an agent will attempt to bring about some $A \wedge C$-world when C holds. But which $A \wedge C$-world is the right one? With no further information, System Z will set *Pref*$(C) = ||A \wedge C||$; all $A \wedge C$-worlds will be assumed to be equally acceptable. This seems to be appropriate: with no further information, any course of action that makes A true should be judged to be as good as any other. Any other assumption, such as gravitation of worlds toward indifference, must make the set *Pref*(C) smaller than $||A \wedge C||$. For example, if we rule out worlds satisfying α from *Pref*(C), then *Pref*$(C) = ||A \wedge C \wedge \neg\alpha||$. This requires that an agent striving for *Pref*(C) make $\neg\alpha$ true as well as A. This imposes unnecessary and unjustified restrictions on the agent's goals, or on the manner in which it decides to achieve them.

Notice that when *worlds* gravitate toward preference, our agent becomes indifferent toward most *propositions*. By maximizing the size of *Pref*(C) (subject to the constraint that A be true), we minimize the number of propositions an agent will care about or attempt to make true in context C. In our example, if $A \wedge C \not\vdash \alpha$ and $A \wedge C \not\vdash \neg\alpha$, then $T(\alpha|C)$ and $T(\neg\alpha|C)$ will both be true in the Z-model. Such indifference toward propositions in a given context seems to be the most appropriate assumption.

[7] Note that in classical decision theory, such distinctions do not exist. An outcome cannot be good or bad, nor can an agent be indifferent toward an outcome, in isolation; it can only be judged *relative* to other outcomes. An agent can adopt an attitude of indifference toward a *proposition*, as we explain below.

In [3, 4] we characterize System Z, in a default reasoning context, as embodying the principle of *conditional only knowing*. When certain beliefs are stated, either actual or conditional, System Z ensures that only propositions that can be shown to be believed (in a given context) are actually believed. We show this to be a generalization of the notion of *only knowing* often adopted in belief logics [21] that accounts for defeasible beliefs. In the preference setting, System Z captures the analogous assumption of ''only preferring.'' Those preferences that can be derived in a given context C are assumed to be the *only* propositions the agent prefers or cares about in that context.

Certain problems with System Z have been shown to arise in default reasoning. These problems occur when reasoning about preferences as well. For example, if we have two independent (absolute) preferences $I(A)$ and $I(B)$, System Z will sanction both $T(A|\neg B)$ and $T(\neg A|\neg B)$; once the preference for B has been violated, one cannot ensure that A is still preferred. Various modifications to System Z have been proposed to deal with such problems, for instance, the ''rule counting'' systems of [15, 5]. Such solutions can be applied in this setting as well, but the assumption of ''only preferring'' lies at the heart of these solutions as well.

We should point out that, while our presentation will assume a unique preference ordering, the definitions to follow do not require this assumption. We are typically given a set of conditional premises of the form $I(B|A)$, plus other modal sentences constraining the ordering. Unless these premises form a ''complete'' theory, there will be a space of permissible orderings. A defeasible reasoning scheme such as System Z can be used to complete this ordering, but we do not *require* the use of a single ordering --- the definitions presented below can be re-interpreted to capture truth in all permissible orderings (i.e., consequence in QDT).

3 Default Knowledge

We should not require that goals be based only on ''certain'' beliefs in *KB*, but on reasonable default conclusions as well. Consider the following preference ordering with atoms R (it will rain), U (have umbrella) and C (it's cloudy). Assuming $\overline{C} \wedge R$ is impossible, we have:

$$\{\overline{CRU}, C\overline{RU}\} < CRU < \{\overline{CR}U, C\overline{R}U\} < CR\overline{U}$$

Suppose, furthermore, that it usually rains when its cloudy. If $KB = \{C\}$, according to our notion of actual preference in the last section, the agent prefers $\overline{R}$ and $\overline{U}$ --- in the best *KB*-world it doesn't rain despite the clouds. However, we cannot use factual preferences (given *KB*) directly to determine goals. Ideally, the agent would like to ensure that it doesn't rain and that it doesn't bring its umbrella. However, clearly the agent can do nothing to make sure $\overline{R}$ holds (we return to this in the next section). Given this, the ''goal'' $\overline{U}$ seems to be wrong. Once C is known, the agent should *expect* R and act accordingly.

As in decision theory, actions should be based not just on preferences (utilities), but also on the likelihood (probability) of outcomes. In order to capture this intuition in a qualitative setting, we propose a logic that has two orderings, one for preferences and one representing the degree of *normality* or *expectation* associated with a world.

The logic QDT, a step toward a qualitative decision theory, is characterized by the class of *QDT-models*, of the form

$$M = \langle W, \leq_P, \leq_N, \varphi \rangle$$

where W is a set of worlds (with valuation function φ), $\leq_P$ is a transitive, connected *preference ordering* on W, and $\leq_N$ is a transitive, connected *normality ordering* on W. We interpret $w \leq_P v$ as above, and take $w \leq_N v$ to mean w is at least as normal a situation as v (or is at least as expected). The submodels formed by restricting attention to either relation are clearly CO-models. The language of QDT contains four modal operators: $\Box_P$, $\overleftarrow{\Box}_P$ are given the usual truth conditions over $\leq_P$; and $\Box_N$, $\overleftarrow{\Box}_N$ are interpreted using $\leq_N$. The conditional $I(B|A)$ is defined as previously, using $\Box_P$, $\overleftarrow{\Box}_P$. A new *normative conditional* connective $\Rightarrow$ is defined in exactly the same fashion using $\Box_N$, $\overleftarrow{\Box}_N$:

$$A \Rightarrow B \ \equiv_{\text{df}} \ \overleftrightarrow{\Box}_N \neg A \vee \overleftrightarrow{\Diamond}_N (A \wedge \Box_N (A \supset B)) \qquad (2)$$

The sentence $A \Rightarrow B$ means B is true at the most normal A-worlds, and can be viewed as a default rule. This conditional is exactly that defined in [3, 7], and the associated logic is equivalent to a number of other systems (e.g., the qualitative probabilistic logic of [14]). QDT can be axiomatized using the following axioms and inference rules for both the preference operators $\Box_P$, $\overleftarrow{\Box}_P$ and the normality operators $\Box_N$, $\overleftarrow{\Box}_N$:

K $\Box(A \supset B) \supset (\Box A \supset \Box B)$

K′ $\overleftarrow{\Box}(A \supset B) \supset (\overleftarrow{\Box} A \supset \overleftarrow{\Box} B)$

T $\Box A \supset A$

4 $\Box A \supset \Box\Box A$

S $A \supset \overleftarrow{\Box} \Diamond A$

H $\overleftrightarrow{\Diamond}(\Box A \wedge \overleftarrow{\Box} B) \supset \overleftrightarrow{\Box}(A \vee B)$

Nec From A infer $\overleftrightarrow{\Box} A$.

MP From $A \supset B$ and A infer B

We require the following axiom to capture their interaction:

PN $\overleftrightarrow{\Box}_N A \equiv \overleftrightarrow{\Box}_P A$

Theorem 1 *The logic QDT is sound and complete with respect to the class of QDT-models.*

Given a QDT-model and a (finite) set of facts *KB*, we define the *default closure* of *KB* to be (where $\mathbf{L}_{CPL}$ is our propositional sublanguage)

$$Cl(KB) = \{\alpha \in \mathbf{L}_{CPL} : KB \Rightarrow \alpha\}$$

That is, those propositions α that are normally true given *KB* form the agent's set of default conclusions. As with preferences, we base our presentation on a unique model determining a unique set of default conclusions. For instance, System Z is one mechanism for defining a unique normality ordering. However, as with preferences, this assumption is not necessary. We assume (for simplicity of presentation) that $Cl(KB)$ is finitely specifiable and take it to be a single propositional sentence.[8]

An agent ought to act not as if only *KB* were true, but also as if its default beliefs $Cl(KB)$ were true. Given a model M, as a first approximation of a definition of goal, we define an *ideal goal* (w.r.t. *KB*) to be any $\alpha \in \mathbf{L}_{CPL}$ such that

$$M \models I(\alpha | Cl(KB))$$

The *ideal goal set* is the set of all such α. Intuitively, the ideal goals are those sentences that must be true if the agent is to find itself in a best possible situation satisfying $Cl(KB)$. In our previous example, where $KB = \{C\}$, we have that $Cl(KB) \equiv C \wedge R$ and the agent's goals are those sentences entailed by $C \wedge R \wedge U$. It should be clear that ideal goals are *conditional* and *defeasible*; for instance, given $C \wedge \overline{R}$, the agent has the ideal goal $\overline{U}$.

This formulation does not provide any indication as to what an agent should *do* in order to achieve these ideal goals. This will require the introduction of actions and ability (see the next section). For instance, notice that the ideal goal set is deductively closed, and we should not expect an agent to have to consider each member of this set individually. The notion of a *sufficient condition* for achieving all ideal goals can be defined in QDT and will prove useful later.

Definition Let X be some proposition. C is a *sufficient condition* given X iff $C \wedge X$ is satisfiable and $M \models \overleftrightarrow{\Box}_P(X \supset \overleftarrow{\Box}_P(X \supset \neg C))$.

Intuitively, a sufficient condition C guarantees that an agent is in some best possible X-world. Thus, if X is some fixed, unchangeable context, ensuring proposition C means the agent has done the best it could.

Proposition 2 *Let C be a sufficient condition given X and let $w \models C \wedge X$. Then $v <_P w$ only if $v \not\models X$.*

With respect to $Cl(KB)$, ideal goals are necessary conditions for ensuring the best situation. A sufficient condition C for $Cl(KB)$ guarantees the entire ideal goal set is satisfied.[9]

Proposition 3 *If C is a sufficient condition for $Cl(KB)$, then $M \models C \wedge Cl(KB) \supset \alpha$ for all ideal goals α.*

We will explore a detailed example in the next section. We also examine the ''priority'' given to defaults over preferences implicit in this scheme, where $Cl(KB)$ is constructed before the preference ranking is consulted.

[8] A sufficient condition for this property is that each ''cluster'' of equally normal worlds in $\leq_N$ corresponds to a finitely specifiable theory. This is the case in, e.g., System Z [3].

[9] Hector Levesque (personal communication) has suggested that sufficiency is the crucial ''operator.''

4 Ability and Incomplete Knowledge

The definition of an ideal goal given *KB* embodies the idea that an agent should attempt to achieve the best possible situation consistent with what it knows (as well as what it conjectures by default). However, as we have emphasized, this is suitable only when *KB* is fixed. If the agent can change the truth of certain elements in *KB*, ideal goals may be too restrictive. Thus, some notion of action and ability must come into play in goal derivation. Actions must also play a role if we are to derive what an agent should *do*, rather than simply what it should *achieve*. Indeed, the term ''goal'' is often interpreted in this way. This is especially important when we notice that the set of propositions an agent should achieve will always be deductively closed. Finally, actions must play a role in factoring out unachievable desires. For instance, an agent might prefer that it not rain; but this is something over which it has no control. Though it is an ideal outcome, to call this a goal is unreasonable.

4.1 Controllable Propositions

To capture distinctions of this sort, we introduce a simple model of action and ability and demonstrate its influence on conditional goals. We ignore the complexities required to deal with effects, preconditions and such, in order to focus attention on the structure and interaction of ability and goal determination.

We partition our atomic propositions into two classes: $\mathbf{P} = \mathcal{C} \cup \overline{\mathcal{C}}$. Those atoms $A \in \mathcal{C}$ are *controllable*, atoms over which the agent has direct influence. The only actions available to the agent are $do(A)$ and $do(\overline{A})$, which make A true or false, for every $A \in \mathcal{C}$. To keep the treatment simple, we assume actions have no effects other than to change the truth value of A. The atom U (have umbrella) is an example of a controllable atom. Atoms in $\overline{\mathcal{C}}$ are *uncontrollable*, for example, R (it will rain).

Definition For any set of atomic variables $\mathcal{P}$, let $V(\mathcal{P})$ be the set of truth assignments to this set. If $v \in V(\mathcal{P})$ and $w \in V(\mathcal{Q})$ for disjoint sets $\mathcal{P}$, $\mathcal{Q}$, then $v; w \in V(\mathcal{P} \cup \mathcal{Q})$ denotes the obvious extended assignment.

We can now distinguish three types of propositions:

Definition A proposition α is *controllable* iff, for every $u \in V(\overline{\mathcal{C}})$, there is some $v \in V(\mathcal{C})$ and $w \in V(\mathcal{C})$ such that $v; u \models \alpha$ and $w; u \models \neg\alpha$.

A proposition α is *influenceable* iff, for some $u \in V(\overline{\mathcal{C}})$, there is some $v \in V(\mathcal{C})$ and $w \in V(\mathcal{C})$ such that $v; u \models \alpha$ and $w; u \models \neg\alpha$.

α is *uninfluenceable* iff it is not influenceable.

Intuitively, since atoms in $\mathcal{C}$ are within complete control of the agent, it can ensure the truth or the falsity of any controllable proposition α, according to its desirability, simply by bringing about an appropriate truth assignment. If $A, B \in \mathcal{C}$ then $A \vee B$ and $A \wedge B$ are controllable. If α

$$RC\overline{T}\overline{H} \quad R\overline{C}\overline{T}\overline{H} \quad \overline{R}\overline{C}TH \quad \overline{R}\overline{C}\overline{T}H$$

$$\downarrow$$

$$RC\overline{T}H \quad R\overline{C}\overline{T}H \quad \overline{R}\overline{C}T\overline{H} \quad \overline{R}\overline{C}\overline{T}\overline{H}$$

$$\downarrow$$

$$RCT\overline{H} \quad R\overline{C}T\overline{H} \quad \overline{R}CTH \quad \overline{R}C\overline{T}H$$

$$\downarrow$$

$$RCTH \quad R\overline{C}TH \quad \overline{R}CT\overline{H} \quad \overline{R}C\overline{T}\overline{H}$$

Figure 4: User Preferences

is influenceable, we call the assignment u to $\overline{\mathcal{C}}$ a *context* for α; intuitively, should such a context hold, α can be controlled by the agent. If $A \in \mathcal{C}$, $X \in \overline{\mathcal{C}}$ then $A \vee X$ is influenceable but not controllable: in context X the agent cannot do anything about the truth of $A \vee X$, but in context $\overline{X}$ the agent can make $A \vee X$ true or false through $do(A)$ or $do(\overline{A})$. Note that all controllables are influenceable. In this example, X is uninfluenceable. The category of controllability into which a proposition falls is easily determined by writing it in minimal DNF. Let $PI(\alpha)$ denote the set of prime implicants of α.

Proposition 4 *a) α is controllable iff each clause in $PI(\alpha)$ contains some literal from $\mathcal{C}$ and some clause contains only literals from $\mathcal{C}$. b) α is influenceable iff some literal from $\mathcal{C}$ appears in $PI(\alpha)$. c) α is uninfluenceable iff no literal from $\mathcal{C}$ appears in $PI(\alpha)$.*

4.2 Complete Knowledge

Given the distinction between controllable and uncontrollable propositions, we want to define goals so that an agent is required to do only those things within its control. A first attempt might simply be to restrict the ideal goal set as defined above to controllable propositions. The following example shows this to be inadequate.

Example Consider five atoms: O, it is overcast; R, it will rain; C, I have coffee; T, I have tea; and H, my office thermostat is set high. My robot has the default information $O \Rightarrow R$. The robot knows *KB* $= \{O, \overline{H}, \overline{C}, \overline{T}\}$: it is overcast and the thermostat is turned down. Its closure is *Cl*(*KB*) $= \{O, R, \overline{H}, \overline{C}, \overline{T}\}$. It can control the three atoms C, T and H. Its preference ordering is designed to respect my preferences: when it's raining I prefer tea when I arrive and the thermostat set high, otherwise I prefer coffee and the thermostat set low. Thus, we have the preference ordering illustrated in Figure 4. (We assume O, R do not contribute directly to preference, and that priority has been given to C and T over H. We also allow the possibility that both C and T together satisfy a preference for either.) The robot has to decide what to do before I arrive at the office.

It should be clear that the robot should not determine its goals by considering the ideal situations satisfying *Cl*(*KB*). In such situations, since $\overline{H}$ is known, $\overline{H}$ is true and indeed, it is a simple theorem of QDT that $I(\alpha|\alpha)$. Thus, the robot concludes that $\overline{H}$ *should be true*. This is clearly mistaken, for considering only the best situations in which one's knowledge of controllables is true prevents one from determining whether changing those controllables could lead to a better situation. Since any controllable proposition can be changed if required, we do not require an agent to restrict attention to those situations where *KB* or *Cl*(*KB*) is true. The fact that $\overline{H}$ is known should not unduly influence what are considered to be the best alternatives --- H can be made true if that is what's best.

Of course, the goals of an agent must still be constrained by known *uninfluenceable* propositions. The agent should not reject all of its knowledge. For example, if the preference ordering above were modified to reflect my preference for $\overline{R}$, the agent should not base its goals on this preference if it knows R. Making R false is beyond its control, and it goals should determined by restricting attention to $\overline{R}$-worlds. Thus we insist that the best situations satisfying known *uninfluenceable* propositions be considered.

Notice that we should not ignore the truth of controllables when making default predictions. The prior truth value of a controllable might provide some indication of the truth of an uncontrollable; and we *must* take into account these uncontrollables when deciding which alternatives are *possible*, before deciding which are best. In this example, we might imagine that the default $O \Rightarrow R$ doesn't hold, but that $O \wedge H \Rightarrow R$ does: if it is overcast, then the thermostat is set high because I anticipated rain before I left last night. Our agent must use the truth of this controllable atom H to determine the truth of the uncontrollable R, which in turn will influence its decisions.[10] Once accounted for in forming *Cl*(*KB*), H can safely be ignored.

This leads to the following formulation of goals that account for ability. We again assume a QDT-model M and sets $\mathcal{C}$, $\overline{\mathcal{C}}$. The *uninfluenceable belief set* of an agent is

$$UI(KB) = \{\alpha \in Cl(KB) : \alpha \text{ is uninfluenceable}\}$$

For the time being, we assume that *UI*(*KB*) is complete: the truth value of all uncontrollable atoms is known. This set of beliefs determines an agent's goals.

Definition Proposition α is a *complete knowledge (CK) goal* iff $M \models I(\alpha|UI(KB))$ and α is controllable.

[10] If a controllable provides some indication of the truth of an uncontrollable or another controllable, (e.g., $H \Rightarrow R$) we should think of this as an *evidential rule* rather than a *causal rule*. Given our assumption about the independence of atoms in $\mathcal{C}$, we must take all such rules to be evidential (e.g., changing the thermostat will not alter the chance of rain). This can be generalized using a more reasonable conditional representation, and ultimately should incorporate causal structure. Note the implicit temporal aspect here; propositions should be thought of as *fluents*. We can avoid an explicit temporal representation by assuming that preference is solely a function of the truth values of fluents.

As with ideal goals, the set of CK-goals is deductively closed and should be viewed as a set of necessary conditions in any rational course of action. Of course, goals can only be affected by atomic actions, so we will typically be interested in a set of actions that is guaranteed to achieve each CK-goal. An (atomic) *action set* is any set of controllable literals. If $\mathcal{A}$ is such a set we use it also to denote the conjunction of its elements. An *atomic goal set* is any action set $\mathcal{A}$ that guarantees each CK-goal; that is

$$M \models \mathit{UI}(\mathit{KB}) \wedge \mathcal{A} \supset \alpha$$

for each CK-goal α. Clearly, such any atomic goal set determines a reasonable course of action. Of course, such action sets can be determined by appeal to sufficiency.

Theorem 5 *Let $\mathcal{A}$ be some atomic action set. Then $\mathcal{A}$ is a goal set iff $\mathcal{A}$ is a sufficient condition for $\mathit{UI}(\mathit{KB})$.*

In our example above, where the robot knows O, possible atomic goal sets are $\{T, H\}$ and $\{C, T, H\}$. Typically, we will be interested in minimal goals sets, since these require the fewest actions to achieve ideality. We may impose other metrics and preferences on goals sets as well (e.g., associating costs with various actions). Notice that the preference for tea does not prevent the robot from bringing coffee. However, such constraints can easily be imposed on the preference ordering. Furthermore, disjunctive goals and ''integrity constraints'' pose no difficulty. If I have no preference for coffee or tea, but prefer exactly one of the two, the generated atomic goal sets will be $\{C, \overline{T}\}$ and $\{\overline{C}, T\}$. The set $\{\overline{C}, \overline{T}\}$ is not a goal set in this case.

With default information and controllability in place, we can briefly return to the alternative interpretation of preference statements suggested in Section 2. The assertion ''I prefer an umbrella when it's raining'' can now be interpreted as $I(U|\mathit{UI}(\{R\}))$. Together with the ''pure'' preferences $I(D|R)$ and $I(\overline{U}|D)$ (and other background information as before), one can conclude $R \Rightarrow \neg D$.

In our goal derivation scheme, a certain priority is given to defaults over preferences. Goals are determined by first constructing the default consequences of *KB*, and then deciding what to do based on this knowledge as if it were certain. In a truly decision-theoretic setting acting on the basis of uncertain information is a function not only of its likelihood, but also the consequences of being incorrect. For instance, in our framework we might have the default rule $R \Rightarrow S$, if I run across the freeway I will cross safely. If this allows me to arrive at my destination five minutes sooner than had I crossed at a crosswalk, the default assumption S will ensure that I run across the freeway: I won't (by default) get hit by a car and I will arrive sooner. In general, the (drastic) consequences of being wrong in this regard must be traded off against the probability of being right. If the five minutes saved is not worth the risk, then I decide to go to the crosswalk.

To express this tradeoff we must assume that the qualitative scales of preference and normality are calibrated somehow, and nothing in the constraints expressed by the user in our setting allows such an assumption. In the concluding section we discuss ''qualitative'' ways around this problem. However, the scheme presented here has a certain naive appeal, which may be partly due to the observation that defaults are usually expressed with such considerations in mind [27, 25]. Furthermore, the scheme is conceptually simple in that it embodies a principle analogous to the *separability* of state estimation and control [11]. An agent can calculate what is (probably) true of the world and subsequently and independently base its decisions upon these beliefs. Finally, our scheme is applicable when likelihood and preference information is truly qualitative and explicit calibration of the orderings is not feasible. We can describe some conditions under which the assumption of separability is appropriate.

The logic of conditional normality statements can be given a probabilistic interpretation as described in [7]. In particular, the purely conditional fragment is equivalent to Adam's system of ε-*semantics*, which has also been applied to the representation of defaults [14]. This means that there is a probability assignment that ensures that every default rule $A \Rightarrow B$ corresponds to an assertion of high conditional probability $P(B|A) > 1 - \varepsilon$, for any $\varepsilon > 0$. Thus, we may assume that a user chooses default rules with such a parameter in mind, and that $P(\mathit{Cl}(\mathit{KB})|\mathit{KB}) > 1 - \varepsilon$. We can also assume that the preference ordering is ''constructed'' by clustering together worlds that have actual utility within some reasonably small range, and treating distinct clusters as separated by a reasonably large gap in utility. Thus, the user can treat certain outcomes as having (more or less) indistinguishable utility. Outcomes in different clusters have sufficiently different utilities. To analyze the appropriateness of our goal derivation scheme, we make this assumption precise by assigning a point utility δ_i to each cluster in the preference ordering. Let δ denote the smallest gap $\delta_i - \delta_{i+1}$ between any two adjacent point utilities (the ''smallest perceptible change'' in utility) and let $= \delta_0 - \delta_n$ denote the magnitude of the possible range in utility.

Goals (or decisions) are determined with respect to a given *KB*, which induces a decision problem in the obvious fashion: given $\mathit{UI}(\mathit{KB})$ what is an optimal action set? Let U^* denote the expected utility of an optimal action under the assumptions above, and let $EU(\mathcal{A})$ denote the expected utility of arbitrary action set $\mathcal{A}$. For any goal set $\mathcal{A}$, we want to compare $EU(\mathcal{A})$ to U^*. We consider a special case first. A *degenerate KB* is one for which every action set applied to $\mathit{UI}(\mathit{KB})$ leads to an equally desirable outcome --- $\mathit{UI}(\mathit{KB})$ allows no decisions to be distinguished. Since only unlikely circumstances (that contradict default conclusions) can influence the choice of action, our scheme cannot generally be optimal in this case, but the error is bounded by the probability of default violation:

Proposition 6 *If KB induces a degenerate decision problem, then $U^* - EU(\mathcal{A}) \leq \varepsilon$ for any goal set $\mathcal{A}$.*

Degenerate problems will be rare: we imagine some differentiation among decisions is possible most of the time. If this is the case, then we have $U^* - EU(\mathcal{A}) \leq \varepsilon - (1-\varepsilon)\delta$.

Proposition 7 *If KB is nondegenerate, any goal set $\mathcal{A}$ is an optimal decision if $\delta(1-\varepsilon) \geq \varepsilon$.*

This gives some idea of the circumstances under which the assumption of separability is sound. Of course, it is unreasonable to only reason with qualitative constraints that meet these stringent requirements. But they do suggest useful abstractions for ordinary goal derivation, and the degree to which these conditions are approximated gives reasonable assurance of good decisions. Thus, the separability assumption provides a computationally manageable procedure for finding ''satisficing'' solutions.

4.3 Incomplete Knowledge

The goals described above seem reasonable, in accord with the general maxim ''do the best thing possible consistent with your knowledge.'' We dubbed such goals ''CK-goals'' because they seem correct when an agent has complete knowledge of the world (or at least of uncontrollable atoms). But CK-goals do not always determine the best course of action if an agent's knowledge is *incomplete*. Consider the preferences in the umbrella example and an agent with an empty knowledge base. For all the agent knows it could rain or not (it has no indication either way). Using CK-goals, the agent ought to $do(\overline{U})$, for the best situation consistent with $KB = \emptyset$ is $\overline{RU}$. Leaving its umbrella is the best choice should it turn out not to rain; but should it rain, the agent has ensured the *worst* possible outcome. It is not clear that $\overline{U}$ should be a goal. Indeed, one might expect U to be a goal, for no matter how R turns out, the agent has avoided the worst outcome.

In the MEU framework, once can deal with such uncertainty easily; but qualitatively, when trying to do as much as possible with strictly ordinal value information, a different approach is required. The scales of preference and normality are unknown and incomparable. It is clear, in the presence of incomplete knowledge, that there are various *strategies* for determining goals. CK-goals form merely one alternative. Such a strategy is opportunistic, or optimistic. Clearly it *maximizes potential gain*, for it allows the possibility of the agent ending up with the best possible outcome. In certain domains this might be a prudent choice (for example, where a cooperative agent determines the outcome of uncontrollables). Of course, another strategy might be the cautious strategy that *minimizes potential loss*.[11] This too can be captured in our logic.

Let a *complete action set* be any complete truth assignment to the atoms in $\mathcal{C}$. These are the alternative courses of action available. To minimize potential loss, we must consider the worst possible outcome for each alternative, and pick those with the ''best'' worst outcomes. If $\mathcal{A}_1$, $\mathcal{A}_2$ are complete action sets, $\mathcal{A}_1$ is *as good as* $\mathcal{A}_2$ ($\mathcal{A}_1 \leq \mathcal{A}_2$) iff

$$M \models \overleftrightarrow{\Diamond}_P(\mathcal{A}_2 \wedge \mathit{UI}(\mathit{KB}) \wedge \neg\overleftarrow{\Diamond}_P(\mathcal{A}_1 \wedge \mathit{UI}(\mathit{KB})))$$

Intuitively, if $\mathcal{A}_1 \leq \mathcal{A}_2$ then the worst worlds satisfying $\mathcal{A}_1$ are at least as preferred as those satisfying $\mathcal{A}_2$ (in the context $\mathit{UI}(\mathit{KB})$). It is not hard to see that $\leq$ forms a transitive, connected preference relation on action sets. The *best* actions sets are those minimal in this ordering $\leq$. To determine the best action sets, however, we do not need to compare all action sets in a pairwise fashion:

Theorem 8 *$\mathcal{A}_i$ is a best action set iff $M \models \mathcal{A}_i \leq \neg\mathcal{A}_i$.*

This holds because the negation of a complete action set (a complete conjunction of literals) is consistent with any other complete action set. If an agent chooses other than a best action set, it opens the possibility for a worse outcome:

Theorem 9 *Let $\mathcal{A}_i$ be a best action set for KB and $\mathcal{A}_j$ be any complete action set. For any $w \models \mathit{UI}(\mathit{KB}) \wedge \mathcal{A}_i$, there is some $v \models \mathit{UI}(\mathit{KB}) \wedge \mathcal{A}_j$ such that $w \leq_P v$.*

Now, we say α is a *cautious goal* iff

$$\vee\{\mathcal{A}_i : \mathcal{A}_i \text{ is a best action set }\} \models \alpha$$

In this way, if (say) $A \wedge B$ and $A \wedge \neg B$ are best action sets, then A is a goal but B is not. Simply doing A (and letting B run its natural course) is sufficient. This notion of goal has controllability built in (ignoring tautologies). In our example above, U is a cautious goal.

We cannot expect best action sets, in general, to be sufficient in the same sense that CK-goal sets are. The potential for desirable and undesirable outcomes makes it impossible to ensure best outcomes consistent with $\mathit{UI}(\mathit{KB})$. However, we can show that if there is some action set that is sufficient for *KB* then all best action sets will be sufficient.

Proposition 10 *If some action set $\mathcal{A}$ for KB is CK-sufficient for KB, then every best action set is CK-sufficient.*

Hence, CK-sufficiency can be applied even in the case of incomplete knowledge. Its applicability implies that possible outcomes of unknown uncontrollables have no influence on preference: all *relevant* factors are known.

The cautious strategy seems applicable in a situation where one expects the worst possible outcome, for example, in a game against an adversary. Once the agent has performed its action, it expects the worst possible outcome, so there is no advantage to discriminating among the candidate (best) action sets: all have equally good worst outcomes. However, it's not clear that this is the best strategy if the outcome of uncontrollables is essentially ''random.'' If outcomes are simply determined by the natural progression of events, then one should be more selective. We think of nature as neither benevolent (a cooperative agent) or malevolent (an adversary). Therefore, even if we decide to

[11] These alternatives are analogs of the *maximax* and *maximin* decision criteria for decision making without outcome probabilities (under *strict* uncertainty [13]).

be cautious (choosing among *best* action sets), we should account for the fact that a worst outcome might not occur: we should choose the action sets that take advantage of this fact.

Observations

It should be clear that if an agent can *observe* the truth values of certain unknown propositions before it acts, it can improve its decisions. In many cases, it will make the worst outcomes better and change the actions chosen. To continue the ''umbrella'' example, suppose R and C are unknown. The agent's cautious goal is then U. If it were in the agent's power to determine C or $\overline{C}$ before acting, its actions could change. Observing $\overline{C}$ indicates the impossibility of R, and the agent could then decide to $do(\overline{U})$.

Space limitations preclude a deep discussion, but briefly, we can distinguish two types of uncontrollable atoms: *observables* and *unobservables*. Suppose *KB* determines a best action set $\mathcal{A}_B$. Intuitively, the observation of some unknown uncontrollable atom O is worthwhile if it can potentially change the agent's goal set. Cautious and optimistic goals must be treated differently. Assume first a cautious strategy. Note that a goal set accounts for some worst outcome which must include either O or $\overline{O}$. Thus, an observation can never be guaranteed to change the agent's decision: it may ''validate'' its cautious approach. In our example, observing C will not change the agent's decision, but observing $\overline{C}$ will. We say atom O has *value* if $\mathcal{A}_B$ is not a best action set for one of $KB \cup \{O\}$ or $KB \cup \{\overline{O}\}$. In this case, observing O is worthwhile since it *might* (depending on its actual truth value) change the agent's goal set. This is a qualitative analog of *value of information*. Of course, we cannot quantify the potential value of making an observation; but we may compare the relative values of two pieces of information O and P. For simplicity, assume that positive observations O and P are the ''improving'' outcomes. Let $\mathcal{A}_O$ and $\mathcal{A}_P$ be best action sets for O and P. The value of O is as great as that of P just when

$$M \models \overleftrightarrow{\Diamond}_P(\mathcal{A}_P \wedge UI(KB \cup \{P\}) \wedge \neg \overleftarrow{\Diamond}_P(\mathcal{A}_O \wedge UI(KB \cup \{O\})))$$

A similar treatment of optimistic goals can be given, where the valuable observations are *undesired* outcomes that change appropriate action. Observation O has value iff $\neg I(\mathcal{A}_B | UI(KB \cup \{O\})$ or $\neg I(\mathcal{A}_B | UI(KB \cup \{\overline{O}\})$ hold.

5 Concluding Remarks

Related Work

Other attempts to define goals using preferences bear some relationship to our system. Doyle and Wellman [12] define goals that exhibit a conditional aspect like ours. Roughly, B is a goal given A just when $A \wedge B$ is preferred to $A \wedge \neg B$ for any *fixed* circumstance. For instance, if such a relationship holds $A \wedge B$ should be preferred given C, given $\neg C$, and so on. Such goals incorporate a *ceteris paribus* assumption: B is preferred to $\neg B$ given A, *all else being equal*. This guarantees that doing B will lead to a better situation whenever A holds. Our conditional goals are much weaker. No such assurances can be provided. Intuitively, if B is a goal given A, then doing B will lead to a better situation, *all else being normal*. However, this permits defeasible goals, affording greater flexibility and naturalness of expression. Only factors directly relevant to utility need be stated, and others are assumed to be irrelevant. In addition, our goals incorporate elements of controllability.

Pearl [24] has proposed a system using much the same underlying logical apparatus as ours. However, conditional statements are taken to impose specific constraints on utility and probability distributions, allowing expected utility calculations (with ''order of magnitude'' values) to be performed. While this allows stronger conclusions to be reached in general, it makes stronger demands on the input information as well. Thus, the system cannot be construed as truly qualitative, so in a sense the aim here is different. Tan and Pearl [28] introduce a somewhat more qualitative system. It handles quantified conditional desires (adopting the machinery of qualitative probability [14]). To account for likelihood, they adopt our model of closing under default consequence before consulting preferences. Incompletely specified preferences induce a ''compact'' model where worlds gravitate toward neutrality, but as noted earlier, this is not an obviously useful strategy. Furthermore, conditional preferences are given a *ceteris paribus* interpretation along the lines of Doyle and Wellman. Aside from the unknown impact on the computation of compact rankings, their particular semantics is of questionable value for representing conditional preferences. For example, a preference for A given $A \vee B$ requires that $\neg A \wedge \neg B$ be dispreferred. In our semantics, a conditional preference given any α imposes no constraints on the degree of preference of $\neg\alpha$-worlds.

Our representation of preferences draws much from work on deontic logic, where preference may be determined by some legal or moral code. Indeed, our logic can be applied to such problems [6]. However, the slogan that characterizes ideal goals, ''do the best given what you know,'' is accepted in much work on the derivation of obligations. Just as in the derivation of goals, such a mechanism is not generally appropriate. Some work in deontic logic has recently begun to incorporate, as we do here, default information [20, 1].

Summary

We have presented a logic QDT for representing qualitative preference and likelihood information. We have shown how defeasible conditional preferences can be expressed, and described several methods for goal derivation based on the assumption that priority be given to defaults. There are a number of ways in which this work can be extended. Clearly, the account of action and ability is naive. An object-level characterization of actions with true causal structure can be added to the conditional framework [24] to make goal derivation more realistic.

The assumption of separability and priority of default information must be relaxed in many circumstances. In order to allow reasonable decisions to be made, a logic that allows tradeoffs of likelihood and preference to be expressed in a qualitative fashion is desirable. For instance, if I instruct my robot that it should run across the street (instead of crossing at the crosswalk) to save three minutes while fetching my coffee, it can safely deduce that running across the street is worth the risk if a courier deadline is involved. I have implicitly calibrated part of its preference and normality rankings with each other. We are currently exploring how such mechanisms to reason directly with such qualitative tradeoff information [9]. This can be viewed as a mechanism to deal with *imperatives*, and propagate the implicit knowledge in such commands to other contexts.

Related to this is a fuller investigation of the different forms preference information might take in such a setting. As mentioned earlier, user preferences might be stated independently of typicality information, or might incorporate expected circumstances and controllability information. A well-developed logic for these and other ‘‘entangled’’ constraints is certainly worth pursuing.

References

[1] Carlos E. Alchourron. Philosophical foundations of deontic logic and the logic of defeasible conditionals. In *Workshop on Deontic Logic in Computer Science*, Amsterdam, 1991.

[2] Craig Boutilier. Conditional logics of normality as modal systems. In *Proc. of AAAI-90*, pages 594--599, Boston, 1990.

[3] Craig Boutilier. Inaccessible worlds and irrelevance: Preliminary report. In *Proc. of IJCAI-91*, pages 413--418, Sydney, 1991.

[4] Craig Boutilier. Normative, subjunctive and autoepistemic defaults: Adopting the Ramsey test. In *Proc. of KR-92*, pages 685--696, Cambridge, 1992.

[5] Craig Boutilier. What is a default priority? In *Proceedings of Canadian Society for Computational Studies of Intelligence Conference*, pages 140--147, Vancouver, 1992.

[6] Craig Boutilier. A modal characterization of defeasible deontic conditionals and conditional goals. In *AAAI Spring Symposium on Reasoning about Mental States: Formal Theories and Applications*, pages 30--39, Stanford, 1993.

[7] Craig Boutilier. Conditional logics of normality: A modal approach. *Artificial Intelligence*, 68:87--154, 1994.

[8] Craig Boutilier. Unifying default reasoning and belief revision in a modal framework. *Artificial Intelligence*, 68:33--85, 1994.

[9] Craig Boutilier and David Poole. Qualitative lotteries: Capturing expected utility calculations in a qualitative framework. Technical report, University of British Columbia, Vancouver, 1994. (Forthcoming).

[10] Philip R. Cohen and Hector J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213--261, 1990.

[11] Thomas Dean and Michael Wellman. *Planning and Control*. Morgan Kaufmann, San Mateo, 1991.

[12] Jon Doyle and Michael P. Wellman. Preferential semantics for goals. In *Proc. of AAAI-91*, pages 698--703, Anaheim, 1991.

[13] Simon French. *Decision Theory*. Halsted Press, New York, 1986.

[14] Moises Goldszmidt. Qualitative probabilities: A normative framework for commonsense reasoning. Technical Report R-190, University of California, Los Angeles, October 1992. Ph.D. thesis.

[15] Moises Goldszmidt, Paul Morris, and Judea Pearl. A maximum entropy approach to nonmonotonic reasoning. In *Proc. of AAAI-90*, pages 646--652, Boston, 1990.

[16] Peter Haddawy and Steve Hanks. Representations for decision-theoretic planning: Utility functions for deadline goals. In *Proc. of KR-92*, pages 71--82, Cambridge, 1992.

[17] Bengt Hansson. An analysis of some deontic logics. *Noûs*, 3:373--398, 1969.

[18] John F. Horty. Moral dilemmas and nonmonotonic logic. *J. of Philosophical Logic*, 1993. To appear.

[19] Richard C. Jeffrey. *The Logic of Decision*. McGraw-Hill, New York, 1965.

[20] Andrew J. I. Jones and Ingmar Porn. On the logic of deontic conditionals. In *Workshop on Deontic Logic in Computer Science*, Amsterdam, 1991.

[21] Hector J. Levesque. All I know: A study in autoepistemic logic. *Artificial Intelligence*, 42:263--309, 1990.

[22] Barry Loewer and Marvin Belzer. Dyadic deontic detachment. *Synthese*, 54:295--318, 1983.

[23] Judea Pearl. System Z: A natural ordering of defaults with tractable applications to default reasoning. In M. Vardi, editor, *Proceedings of Theoretical Aspects of Reasoning about Knowledge*, pages 121--135. Morgan Kaufmann, San Mateo, 1990.

[24] Judea Pearl. A calculus of pragmatic obligation. In *Proc. of UAI-93*, pages 12--20, Washington, D.C., 1993.

[25] David Poole. Decision-theoretic defaults. In *Proceedings of Canadian Society for Computational Studies of Intelligence Conference*, pages 190--197, Vancouver, 1992.

[26] Frank P. Ramsey. Turth and probability. In R. B. Braithwaite, editor, *The Foundations of Mathematics and Other Logical Essays*, pages 156--198. Kegan Paul, London, 1931.

[27] Yoav Shoham. *Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press, Cambridge, 1988.

[28] Sek-Wah Tan and Judea Pearl. Specification and evaluation of preferences for planning under uncertainty. In *Proc. of KR-94*, pages 530--539, Bonn, 1994.

[29] John von Neumann and Oskkar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, 1944.

Acknowledgements

Thanks to Carlos Alchourron, Jon Doyle, Keiji Kanazawa, Judea Pearl, David Poole and Michael Wellman for very valuable discussion of this topic. This research was supported by NSERC Research Grant OGP0121843.