Weak Conditional Logics of Normality

James P. Delgrande
School of Computing Science,
Simon Fraser University,
Burnaby, B.C., Canada V5A 1S6
jim@cs.sfu.ca

Abstract

A default conditional $\alpha \rightarrow \beta$ has most often been informally interpreted as a defeasible version of a classical conditional, usually the material conditional. That is, the intuition is that a default should behave (implicitly or explicitly) as its (say) material counterpart "by default" or unless explicitly overridden. In this paper, we develop an alternative interpretation, in which a default is regarded more like a rule, leading from premises to conclusion. To this end, a general semantic framework under a "rule-based" interpretation is developed, and a family of weak conditional logics is specified, along with associated proof theories. Nonmonotonic inference is defined very easily in these logics. One obtains a rich set of nonmonotonic inferences concerning the incorporation of irrelevant properties and of property inheritance. Moreover, this interpretation resolves problems that have been associated with previous approaches.

1 Introduction

A major approach in nonmonotonic reasoning has been to represent a default as an object that one can reason about, either as a conditional as part of some object language, or as a nonmonotonic consequence operator. Thus for example "an adult is (typically or normally) employed" might be represented $a \rightarrow e$, where $\rightarrow$ represents a default conditional, distinct from the material conditional $\supset$. Given a suitable proof theory and semantics, one can derive other defaults from a given set of defaults. There has been widespread agreement concerning just what principles should constitute a minimal logic, a suggested "conservative core," of defaults. However, the resulting conditional is quite weak, at least compared with the material conditional, in that it does not (in fact, *should not)* fully support principles such as strengthening of the antecedent, transitivity, and modus ponens.

Since one would want to obtain these latter properties "by default," such logics are extended nonmonotonically by a "closure" operation or step. This closure operation has, for example, been defined by selecting a (preferred) subset of the models of a theory; in the resulting subset of the models one obtains strengthening of the antecedent, transitivity, or (effectively) modus ponens, wherever feasible. Essentially then, there are two components to default reasoning within such a system. First, there is a standard, monotonic logic of conditionals that expresses relations among defaults that are deemed to always hold. Second, there is a nonmonotonic mechanism for obtaining defaults (and default consequences) where justified. Representative (but not even remotely exhaustive) work in this area includes [Geffncr and Pearl, 1992; Pearl, 1990; Kraus *et al,* 1990]. In essence, these approaches treat the default conditional like its classical counterpart, the material conditional, where "feasible" or "by default".

While this work captures an important notion of default entailment - perhaps *the* most important notion - it is not without difficulties. As described in the next section, some principles of the "core" logic are not uncontentious; as well, there are examples of default reasoning in which one obtains undesired results. Lastly, there are more recent approaches, notably addressing causality, in which one requires a weaker notion of default inference, rejecting, for example, contrapositive default inferences. In response to these points, we suggest that there is a second, distinct, interpretation of default conditionals, in which a default is regarded more like a rule, with properties more in line with a rule of inference, than a weakened classical conditional.

In the following sections we describe our proposed approach informally and formally. We begin by proposing an exceptionally weak logic of conditionals; from this basis a family of conditional logics is defined. Given a default conditional $\alpha \rightarrow \beta$, the underlying intuition that is formalised is that α supplies a context in which, all other things being equal, β normally holds or, more precisely, in the context of α, $\alpha \wedge \beta$ is more "normal" than $\alpha \wedge \neg\beta$. Notably, all of the logics that we consider are weaker than the aforementioned "conservative core". It proves to be the case however that a nonmonotonic operation is very easily defined; this nonmonotonic step essentially specifics that a property is irrelevant with respect to a default unless it is known to be relevant. This nonmonotonic step easily admits inferences that in other approaches have taken significant formal machinery to obtain. As well, we show that the aforementioned difficulties that arise in interpreting a default as a weak classical conditional do not arise here.

This distinction between treating a default as a conditional or as a rule has been noted previously; see for example [Geffner and Pearl, 1992]. As well, work on inheritance networks [Horty, 1994] can be viewed as investigating proof theories for the latter interpretation; and work on causality such as [McCain and Turner, 1997] falls in the rule-based framework. However a logic (that is, with both semantics and proof theory) capturing this interpretation has not (to our knowledge) been investigated previously, nor has a fully general nonmonotonic closure operator been developed under this interpretation. Last, we suggest in the conclusion that this alternative interpretation may be widely applicable, extending to areas such as counterfactual reasoning, generally treated via the stronger interpretation.

2 Background

In recent years, much attention has been paid to conditional systems of default reasoning. Such systems deal with defeasible conditionals based on notions of preference among worlds or interpretations. Thus, the default that a bird normally flics can be represented propositionally as $b \rightarrow f$.[1] These approaches are typically expressed using a modal logic in which the connective $\rightarrow$ is a binary modal operator. The intended meaning of $\alpha \rightarrow \beta$ is approximately "in the least worlds (or most preferred worlds) in which α is true, β is also true". Possible worlds (or, again, interpretations) are arranged in at least a partial preorder, reflecting a metric of "normality" or "preferredness" on the worlds. Given a set of defaults Γ, default entailment with respect to Γ, $\mid\sim_\Gamma$, can be defined via:

$$\text{If} \quad \Gamma \vdash \alpha \rightarrow \beta \quad \text{then} \quad \alpha \mid\sim_\Gamma \beta.$$

There has been a remarkable convergence or agreement on what inferences ought to be common to all nonmonotonic systems; space considerations preclude a full listing of approaches and references. The resulting set of principles has been called the *conservative core* in [Pearl, 1989]. It was originally considered in [Adams, 1975], and has been studied extensively, as the system P, in [Kraus *et at.,* 1990]. One expression of the logic of conditionals is as follows. The logic includes classical prepositional logic and the following rules and axioms:[2]

RCEA/LLE: From $\vdash \alpha \equiv \beta$ infer $\vdash (\alpha \rightarrow \gamma) \equiv (\beta \rightarrow \gamma)$.

RCM/RW: From $\vdash \beta \supset \gamma$ infer $\vdash (\alpha \rightarrow \beta) \supset (\alpha \rightarrow \gamma)$

ID/Ref: $\alpha \rightarrow \alpha$

RT/Cut: $((\alpha \rightarrow \beta) \wedge (\alpha \wedge \beta \rightarrow \gamma)) \supset (\alpha \rightarrow \gamma)$

ASC/CM: $((\alpha \rightarrow \beta) \wedge (\alpha \rightarrow \gamma)) \supset (\alpha \wedge \beta \rightarrow \gamma)$

CA/Or: $((\alpha \rightarrow \gamma) \wedge (\beta \rightarrow \gamma)) \supset (\alpha \vee \beta \rightarrow \gamma)$

These principles are not uncontentious; for example, [Poole, 1991] can be viewed as arguing against a derived principle CC/And (viz. $((\alpha \rightarrow \beta) \wedge (\alpha \rightarrow \gamma)) \supset (\alpha \rightarrow \beta \wedge \gamma)$ obtained via RT and ASC). Likewise, [Neufeld, 1989] suggests against CA in some cases.

Nonetheless, the resulting logic (and proposed strengthenings of the logic) are weak. A central difficulty is that seemingly *irrelevant* properties will block a desired inference. For example, given that a bird is asserted to fly by default, one cannot thereby conclude that a green bird flies by default. The problem, essentially, is that there is nothing requiring preferred worlds in which birds fly to include among them green-bird worlds. Consequently, various means of strengthening the logic to incorporate irrelevant properties in a principled fashion have been proposed. Thus for "birds fly," since being green is presumably irrelevant with respect to flight, one would want to have, among the preferred worlds in which birds fly, a (preferred) subset in which there are green birds.

Rational closure [Lehmann and Magidor, 1992], for example, assumes that a world is ranked as unexceptional as possible. In *conditional entailment* [Geffner and Pearl, 1992], a partial order on possible worlds is determined by ranking worlds based on the highest-ranked default that distinguishes between the worlds. In both cases, defaults are evaluated with respect to the resulting ranking. Thus, since there is no reason to suppose that greenness has any bearing on flight, one assumes that greenness has no effect on flight.

However there are difficulties with both approaches. Rational closure employs a very strong minimization criterion; see [Geffner and Pearl, 1992] for a number of problematic examples. As well, consider an elaboration of an example given by John Horty:

$$\mathsf{T} \rightarrow \neg f,\ a \rightarrow f,\ \mathsf{T} \rightarrow n,\ o \rightarrow \neg n. \tag{1}$$

(Normally one does not eat with the fingers (f), but one does when eating asparagus (a); normally one uses a napkin (n), but not when one is out of napkins (o).) The rational closure of these conditionals gives that, if one is not out of napkins $(\neg o)$, one is not eating asparagus $(\neg a)$. Clearly this interaction between unrelated defaults is undesirable. As well, in neither conditional entailment nor rational closure does one obtain inheritance of properties (however see [Benferhat *et ai,* 1993]). In addition, consider the following example [Geffner and Pearl, 1992]:

$$a \rightarrow e,\ u \rightarrow a,\ u \rightarrow \neg e,\ f \rightarrow a. \tag{2}$$

(That is, adults are normally employed, university students are normally adults but are not employed, and Frank Sinatra fans are normally adults.) In both conditional entailment and rational closure we obtain the default inference that Frank Sinatra fans are not university students. But this is a curious inference, since there is nothing in the example that would seem to relate Frank Sinatra fans to university students.

We suggest that at least some of these examples do not necessarily reflect a problem with the approaches per se. Rather, our thesis is that there are (at least) two distinct interpretations that can be given to a default. First, there is the intuition

[1] An alternative is to treat the conditional as a nonmonotonic inference operator, $b \mid\sim f$. In a certain sense these approaches can be considered equivalent; here, for simplicity, we remain within the conditional logic framework.

[2] Two systems of nomenclature have arisen, one associated with conditional logic and one with nonmonotonic consequence operators. We list both (when both exist) when first presenting an axiom or rule; for example the conditional logic rule for substitution of logical equivalents in the antecedent is called RCEA; its nonmonotonic consequence operator, Left Logical Equivalence is abbreviated LLE. Hence we first list the rule as RCEA/LLE.

that a default is essentially a weak version of the material conditional (or, in more recent approaches, necessary entailment), and should behave as such a conditional, except that it is defeasible. This intuition is seen most clearly in the expression of defaults by circumscriptive abnormality theories [McCarthy, 1986]. In this case a default $\alpha \rightarrow \beta$ is represented as the formula $\alpha \wedge \neg Ab_i \supset \beta$. The circumscription of Ab_i asserts that Ab_i is false (roughly) if consistently possible. Obviously, if Ab_i is asserted to be false, the result is exactly the material conditional. As well, conditional entailment adopts this intuition: the default $\alpha \rightarrow \beta$ is basically the same $\top \rightarrow (\alpha \supset \beta)$ (and so $\top \rightarrow (\neg\beta \supset \neg\alpha)$), together with specificity information implicit in α [Geffner and Pearl, 1992, p. 232]. This interpretation also underlies approaches that assume that defaults are founded (formally or informally) on notions of probability.[3]

The second interpretation regards a default more as an (object-level) *rule*, whose properties are closer to those of a rule of inference. Thus, given a conditional $\alpha \rightarrow \beta$, if the antecedent *a* happens to be true, we conclude β by default. Given $\neg\beta$ we specifically do not want to conclude $\neg\alpha$. This latter interpretation covers an important class of defaults that have not been addressed as a logic of conditionals per se (although this distinction has been noted in [Geffner and Pearl, 1992]; and inheritance networks [Horty, 1994] can be viewed as proof theoretic accounts of this interpretation). A further motivation for exploring this interpretation is that there has been recent interest in conditional accounts of causality (for example [McCain and Turner, 1997]), in which reasoning via a default contrapositive is explicitly rejected. Thus, from "a causes b" we don't want to conclude "$\neg b$ causes $\neg a$."

3 Defaults as Rules

The general approach is the same as those described in the previous section: we begin by specifying a logic of defaults and subsequently provide a principled, nonmonotonic, means to extend the logic to account for irrelevant properties. Our point of departure is that we informally treat defaults more like rules of inference; in particular, defaults are intended to be applied in a "forward" direction only. Our interpretation, roughly, is that the antecedent of a default establishes a context in which the consequent (normally) holds, or holds all other things being equal. Thus, for default $\alpha \rightarrow \beta$, our interpretation is roughly that, in the context established by a, it is the case that β is more normal (typical, etc.) than $\neg\beta$. We express this semantically by

$$\|\alpha \wedge \neg\beta\|^M < \|\alpha \wedge \beta\|^M \quad (3)$$

that is, the proposition (see the next $\|\alpha\wedge\beta\|^M$ is more normal (typical, etc.) than $\|\alpha \wedge \neg\beta\|^M$. It seems reasonable that our binary relation of relative normality $<$ be asymmetric and transitive, and so we generally assume that these conditions hold. We note that the form of (3) has appeared regularly in the literature, going back at least to [Lewis, 1973].[4]

[3] We note in passing that there are examples of defaults that have nothing to do with probability, for example (1).

[4] We use $<$ in the opposite sense of Lewis and many other authors. For our interpretation, it seems to make more sense to express The difference is that usually the interpretation of (3) is along the lines of "the *least worlds* $\alpha\wedge\neg\beta$ is true are less normal than the least $\alpha \wedge \beta$ worlds"; our interpretation refers to the *proposition* expressed by these formulas. Consequently, our relation $<$ is not an accessibility relation in the normal sense, since it is a relation on *sets* of worlds.

Filling in the (formal) details yields a weak logic of conditionals, significantly weaker than the so-called "conservative core". We also consider various strengthenings of the logic, but these strengthenings are still weaker than this "core" set of defaults. We subsequently define a notion of nonmonotonic inference with respect to these logics. It proves to be the case that this is very easy to do in our approach. Basically, the (semantic) relation $X < Y$ asserts that in the "context" (set of possible worlds) $X \cup Y$, partitioned by X,Y, we have that Y is more normal than X. Our nonmonotonic assumption is that this obtains in all "feasible" subcontexts. That is, for proposition Z, unless there is reason to conclude otherwise, we assert that $Z \cap X < Z \cap Y$. The next section develops the formal details.

4 The Approach

4.1 The Base Logic

Let $\mathcal{L}_{PC}$ be the language of propositional logic defined, for simplicity, over a finite alphabet $P = \{a, b, c, \ldots\}$ *of propositional letters* or *atomic propositions,* and employing the logical symbols $\neg$, $\vee$, $\wedge$, $\supset$, and $\equiv$. The symbol $\top$ is taken to be some propositional tautology, and $\bot$ is defined as $\neg\top$. The language $\mathcal{L}$ is $\mathcal{L}PC$ extended with the binary operator $\rightarrow$ as a *weak conditional.* For convenience, arguments of $\rightarrow$ are members of $\mathcal{L}_{PC}$; hence, we do not allow nested occurrences of $\rightarrow$. Formulas are denoted by the Greek letters α, β, ... and sets of formulas by upper case Greek letters $\Gamma, \Delta,$ Sentences are interpreted with respect to a *comparative conditional model* $M = \langle W, <, P\rangle$ where:

1. W is a set (of states or possible worlds);
2. $< \subseteq W \times 2^W \times 2^W$ with properties described below;
3. P: $\mathbf{P} \mapsto 2^W$.

We use the upper case letters *X, X1, Y,* ...to denote sets of possible worlds. *P* maps atomic sentences onto sets of worlds, being those worlds at which the sentence is true. The relation $<$ associates with each world $w \in W$ a binary notion of *relative normality* between propositions; we write $X <_w Y$ to assert informally that, according to world *w,* proposition Y is more normal than X. That is, given a partition $\{X, Y\}$ of a context $X \cup Y$, the relation $X < Y$ asserts that Y is more normal (unexceptional, etc.) than X. We assume that $<$ is a strict partial ordering on its last two arguments, that is for $w \in W$, $<_w$ is asymmetric and transitive. As well, we assume that the incoherent proposition is maximally abnormal:

$$X \neq \emptyset \quad \text{iff} \quad \emptyset <_w X. \quad (4)$$

Truth of a formula at a world in a model is as for propositional logic, with an addition for $\rightarrow$:

"Y is more normal than X" by A' $<$ *Y* rather than by $Y < X$.

Definition 4.1

1. $\models^M_w p$ *for* $p \in \mathbf{P}$ *iff* $w \in P(p)$.
2. $\models^M_w \alpha \wedge \beta$ *iff* $\models^M_w \alpha$ *and* $\models^M_w \beta$.
3. $\models^M_w \neg\alpha$ *iff* $\not\models^M_w \alpha$.
4. $\models^M_w \alpha \rightarrow \beta$ *iff* $\|\alpha \wedge \neg\beta\|^M <_w \|\alpha \wedge \beta\|^M$ *or* $\|\alpha\|^M = \emptyset$.

A formula α is *valid,* written $\models \alpha$, just if it is true at every world in every model. We identify the proposition expressed by a sentence α with the set of worlds in which α is true, denoted $\|\alpha\|^M$, that is,

$$\|\alpha\|^M = \{w \mid \models^M_w \alpha\}.$$

Thus $\alpha \rightarrow \beta$ is true just if the proposition expressed by $\alpha \wedge \beta$ is more normal than that expressed by $\alpha \wedge \neg\beta$. α is necessarily true, $\Box\alpha$, is defined as $\neg\alpha \rightarrow \alpha$.

Consider the logic closed under classical propositional logic along with the following rules of inference and axioms:

RCEA/LLE: From $\vdash \alpha \equiv \alpha'$ infer $\vdash (\alpha \rightarrow \beta) \equiv (\alpha' \rightarrow \beta)$

RCECA:

From $\vdash \alpha \supset (\beta \equiv \beta')$ infer $\vdash (\alpha \rightarrow \beta) \equiv (\alpha \rightarrow \beta')$

RI/SupraCI: From $\vdash \alpha \supset \beta$ infer $\vdash \alpha \rightarrow \beta$

WeakCEM: $\neg(\alpha \rightarrow \bot) \supset ((\alpha \rightarrow \beta) \supset \neg(\alpha \rightarrow \neg\beta))$

Trans: $((\alpha \vee \beta \rightarrow \beta) \wedge (\beta \vee \gamma \rightarrow \gamma)) \supset (\alpha \vee \gamma \rightarrow \gamma)$

We call the smallest logic based on the above axiomatisation C. Theoremhood of a formula *a* is denoted, as usual, by $\vdash \alpha$. This system is quite weak; however in the full paper we consider the even weaker logic C- consisting of propositional logic and RCEA, RCECA, and RI. RCEA asserts that conditionals with the same consequent and equivalent (in propositional logic) antecedents are equivalent. RCECA asserts the same thing with respect to consequents, but is somewhat more general, in that the consequents need be equivalent just in the "context" given by the antecedent. RI asserts that if β is logically implied by α, then it is also normally implied. WeakCEM gives a weak version of the excluded middle for a weak conditional; in the semantics this is reflected by asymmetry of $<$. Similarly Trans reflects transitivity of $<$ in the semantics. We obtain the following basic results:

Theorem 4.1

1. $\alpha \rightarrow \alpha$
2. $(\alpha \rightarrow \beta) \supset (\alpha \rightarrow (\alpha \wedge \beta))$
3. $(\alpha \rightarrow \beta) \supset (\alpha \rightarrow (\alpha \supset \beta))$
4. *if* $\vdash \beta$ *then* $\vdash \alpha \rightarrow \beta$

As well, we obtain:

Theorem 4.2 *a is valid in the class of comparative conditional mciff* $\vdash \alpha$ *in* C.

For those familiar with conditional logic, the formulas RCM, RT, ASC, CC, and CA (and their nonmonotonic consequence operator counterparts: Right Weakening, Cut, Cautious Monotonity, And, and Or) are not valid in C. Nonetheless, despite its (monotonic) inferential weakness, the logic already allows a rich set of nonmonotonic inferences, as covered in the next section. However, before proceeding we first consider various extensions to the logic.

4.2 Extensions to the Logic

In the logic C most properties of the relation $<_w$ stem from its being a strict partial order (viz. asymmetric and transitive). We also look at strengthening $<_w$ by considering properties that seem reasonable for a notion of normality. Consider the following:

CD: If $X <_w Y$ then $X \setminus Z <_w Y$.

CU: If $X <_w Y$ then $X <_w Y \cup Z$ provided $X \cap Z = \emptyset$.

UD: If $X <_w Y$ then $X \setminus Z <_w Y \cup Z$.

WDjU: If $X_1 <_w Y_1$, $X_2 <_w Y_2$, $Y_1 \cap Y_2 = \emptyset$ and $(X_1 \cup X_2) \cap (Y_1 \cup Y_2) = \emptyset$ then $X_1 \cup X_2 <_w Y_1 \cup Y_2$.

For CD, if *X* is less normal than *Y,* then a stronger proposition than *X* (viz. *X\Z)* is also less normal than *Y. CD* is a dual: if X is less normal than Y, then a weaker proposition than *Y* (viz. *Y* $\cup$ *Z)* is also more normal than *X.* UD combines these conditions, and WDjU is a weak version of disjoint unions, described below. Interestingly, the first two of these conditions have appeared in the belief revision literature. Our relation $<$ is what [Alchourron and Makinson, 1985] call a *(transitive) hierarchy;* while $<$ with CD and CU is a *regular* hierarchy. Their interpretation of $<$ echoes ours for $X < y$, that *"X* is less secure or reliable or plausible … than *Y"* [Giirdenfors and Rott, 1995, p. 75].

Consider next the following rules and formulas:

WeakRCM: If $\vdash \beta \supset \gamma$ then $\vdash (\alpha \rightarrow \beta) \supset (\alpha \wedge \gamma \rightarrow \beta)$

CM/Cond: $(\alpha \wedge \beta \rightarrow \gamma) \supset (\alpha \rightarrow (\beta \supset \gamma))$

RCM: If $\vdash \beta \supset \gamma$ then $\vdash (\alpha \rightarrow \beta) \supset (\alpha \rightarrow \gamma)$

D: $((\alpha \wedge \beta \rightarrow \gamma) \wedge (\alpha \wedge \neg\beta \rightarrow \gamma)) \supset (\alpha \rightarrow \gamma)$.

Space considerations preclude a full discussion of these conditions. WeakRCM is a weaker version of both the rule RCM and the formula ASC in conditional logic. CM ("Conditionalisation" in nonmonotonic consequence relations) gives a conditional version of one half of the deduction theorem. Combining WeakRCM and CM yields the rule RCM, allowing weakening of the consequent of a conditional. D supplies a certain "reasoning by cases" for the conditional.

We obtain the following correspondence between semantic conditions and the axiomatisation:

Theorem 4.3 C + CD (CU, UD, WDjU) *is complete with respect to the class of comparative conditional models closed under* WeakRCM (CM, RCM, *D)*

We could go on and add other conditions in the semantics. Again space considerations dictate against a lengthy discussion, but two conditions are worth noting here:

Disjoint Union: If $(X \cup Y) \cap Z = \emptyset$ then:
$X <_w Y$ iff $X \cup Z <_w Y \cup Z$.

Connectivity: **For** $X, Y \subseteq W$, **either** $Y <_w X$ **or** $X <_w Y$.

Disjoint union has appeared frequently in the literature, for example [Savage, 1972; Fine, 1973; Dubois *et al,* 1994]. The addition of disjoint union requires that the notion of a model be altered slightly (from a relation $<$ **to** $\leq$**);** the resultant semantic framework would correspond to the basic definition

of a plausibility structure [Friedman and Halpern, 2001]. The addition of connectivity would make $<_w$ a *qualitative probability* in the terminology of [Savage, 1972].

5 Nonmonotonic Reasoning

We claimed at the outset that the logic C and its strengthenings would allow a simple approach to nonmonotonic inference, having just the "right" properties for a rule-based interpretation of a conditional. For the logics, the central idea was that, given a partition $\{X, Y\}$ of a context $X \cup Y \subseteq W$, the relation $X < Y$ asserts that Y is more normal (unexceptional, etc.) than X. To obtain nonmonotonic inference, we simply assume that this relation holds in any subcontext, that is $X \cap Z < Y \cap Z$ wherever "reasonable". More formally, we have the following:

Definition 5.1 *Let $M = \langle W, <, P\rangle$ be a comparative conditional model in* C.
Define $M^* = \langle W, <^*, P\rangle$, the augmentation *of M, by:*
$X <^*_w Y$ *iff there are $X' \supseteq X$, $Y' \supseteq Y$ such that*
1. $X' <_w Y'$ *and*
2. *for every X'', Y'' where*
$X \subseteq X''$, $Y \subseteq Y''$ *and* $Y'' <_w X''$
we have:
$X' \subseteq X''$, $Y' \subseteq Y''$.

In the augmentation of a model, we will have $X <^*_w Y$ if $X <_w Y$. Otherwise, $X <^*_w Y$ if there are X', Y' where $X' <_w Y'$ for $X \subseteq X'$ and $Y \subseteq Y'$, and there is no X'', Y'' where $X \subseteq X'' \subseteq X'$ and $Y \subseteq Y'' \subseteq Y'$ and $Y'' <_w X''$ in the original model.

Theorem *If $M = \langle W, <, P\rangle$ is a comparative conditional model then so is $M^* = \langle W, <^*, P\rangle$.*

We define $\models^*$ as validity in the class of augmented comparative conditional models; that is $\models^* \alpha$ iff a is true at every world in every augmented comparative conditional model. Nonmonotonic inference is defined as follows:

Definition 5.2 *Let $\Gamma \subseteq \{\alpha \rightarrow \beta \mid \alpha, \beta \in \mathcal{L}_{PC}\}$.*
$\alpha \mid\!\sim_\Gamma \beta$ *iff* $\models^* \Gamma \supset (\alpha \rightarrow \beta)$.

We say that β is a nonmonotonic inference from α with respect to Γ, or just β is a nonmonotonic inference from α if the set Γ is clear from the context of discussion.

We illustrate nonmonotonic inference first by a familiar example:

$$b \rightarrow f, \quad (5)$$
$$b \rightarrow w, \quad (6)$$
$$\Box(p \supset b), \quad (7)$$
$$p \rightarrow \neg f. \quad (8)$$

Thus birds fly and have wings, and penguins are (necessarily) birds that do not fly. We obtain the following:

$$b \wedge w \mid\!\sim f, \qquad p \wedge w \mid\!\sim \neg f,$$
$$b \wedge \neg w \mid\!\sim f, \qquad p \wedge b \mid\!\sim \neg f,$$
$$b \wedge \neg p \mid\!\sim f, \qquad p \wedge b \wedge w \mid\!\sim \neg f.$$

We also obtain $b \wedge x \wedge y \wedge z \mid\!\sim w$ for $x \in \{\top, g, \neg g\}, y \in \{\top, p, \neg p\}, z \in \{\top, f, \neg f\}$. Thus green ($g$) birds have wings, as do non-green flying penguins. As well $p \mid\!\sim w$, and so penguins inherit the property of having wings by virtue of necessarily being birds. Note that if we replaced (7) by $p \rightarrow 6$, we would no longer obtain $p \mid\!\sim w$; however we would obtain the weaker $b \wedge p \mid\!\sim w$. We justify this by noting that a normality conditional $\alpha \rightarrow \beta$ does not imply a strict specificity relation between α and β whereas $\Box(\alpha \supset \beta)$ does.

The next example further illustrates reasoning in the presence of exceptions.

$$q \rightarrow p, \quad r \rightarrow \neg p, \quad q \rightarrow g \quad (9)$$

So Quakers are pacifists while Republicans are not, and Quakers are generous. We obtain $q \wedge \neg r \mid\!\sim p$ and $q \wedge r \mid\!\sim g$. Thus in the last case, while Quakers that are Republican are, informally, exceptional Quakers, they are nonetheless still generous by default.

Concerning our original motivating examples, in (1) we do not obtain the undesirable inference $\neg o \mid\!\sim \neg a$, and in (2) we do not obtain $f \mid\!\sim \neg u$. Last, we note that while we obtain full incorporation of irrelevant properties, we do not obtain full default transitivity. Thus

$$a \rightarrow b, \quad b \rightarrow c$$

does not yield $a \mid\!\sim c$ (nor, incidentally, do we obtain $\neg b \mid\!\sim \neg a$). However we do get $a \wedge b \mid\!\sim c$. If we replaced $a \rightarrow b$ with $\Box(a \supset b)$ we would get $a \mid\!\sim c$. If we replaced $b \rightarrow c$ with $\Box(b \supset c)$ we would again get $a \mid\!\sim c$, in **C** (in fact we could derive $a \rightarrow c$ in those logics containing **RCM**, as given in Theorem 4.3).

6 Discussion

We have argued that there are two interpretations of a default conditional: as a weak (typically material) implication, or as something akin to a rule of inference. The former interpretation is explicit in, for example, circumscriptive abnormality theories, and implicit in an approach such as conditional entailment. It is clear that there are many, and varied, applications in which the first interpretation is appropriate. However we have also noted that there are various reasons to suppose that this is not the only such interpretation: First, work such as [Poole, 1991] and [Neufeld, 1989] can be viewed as arguing against principles of the "core" logic underlying this first interpretation (the former arguing against the principle CC/And and the latter against CA/Or). Second, there are examples of inferences in approaches such as rational closure or in conditional entailment that are either too weak or too strong. Last, there are emerging areas (such as causal reasoning) in which a "weak material implication" interpretation is not appropriate. While this distinction has been recognized previously, what is new here is the development of a family of logics, with a novel semantic theory and proof theory, along with a specification of nonmonotonic inference, for the "rule-based" interpretation.

All of the logics presented here are quite weak, at least compared to the "conservative core" or, equivalently, the system P of [Kraus *et al,* 1990]. We argue however that such lack of inferential capability is characteristic of a "rule-based" interpretation of a conditional. Moreover it proves to

be the case that nonmonotonic reasoning is defined very easily in these logics, and allows a rich set of inferences concerning the incorporation of irrelevant properties and of property inheritance.

An open question concerns how informal, commonsense defaults should be classified - whether as a defeasible classical conditional or as a rule. Certainly past work has favoured the "defeasible classical conditional" interpretation. However, a case can be made that many examples formerly interpreted as belonging to the first category are better interpreted as belonging to the "rule" category. Consider Lewis' approach to counterfactuals [Lewis, 1973] in which the following example, concerning a past party, is given: "If John had gone it would have been a good party" and "If John and Mary had gone it would have not been a good party". From this we deduce that "if John had gone, Mary would not have gone". This, to most readers, is a strange result: John's going and Mary's going are (presumably) independent events. Arguably this result ought not to obtain, and so perhaps counterfactuals, as previously modelled by Lewis' sphere semantics, may be better interpreted via the "rule" interpretation.

Acknowledgements

1 thank the reviewers for their useful suggestions. This work was supported by a Natural Sciences and Engineering Council of Canada Discovery Grant.

References

[Adams, 1975] E. Adams. *The Logic of Conditionals.* D. Reidel Publishing Co., Dordrecht, Holland, 1975.

[Alchourron and Makinson, 1985] C.E. Alchourron and D. Makinson. On the logic of theory change: Safe contraction. *Studia Logica,* 44:405-422,1985.

[Benferhat el al., 1993] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proceedings of the International Joint Conference on Artificial Intelligence,* pages 640-645, Chambery, Fr., 1993.

[Dubois *et ah,* 1994] D. Dubois, J. Lang, and H. Prade. Possibilistic logic. In Dov M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Nonmonotonic Reasoning and Uncertain Reasoning,* volume 3 of *Handbook of Logic in Artificial Intelligence and Logic Programming,* pages 439-513. Oxford, 1994.

[Fine, 1973] T.L. Fine. *Theories of Probability: An Examination of Foundations.* Academic Press, 1973.

[Friedman and Halpern, 2001] N. Friedman and J. Halpern. Plausibility measures and default reasoning. *JACM,* 48(4):649-685,2001.

[GardenforsandRott, 1995] P. Gardenfors and H. Rott. Handbook of logic in artificial intelligence and logic programming, volume 4 of *Epistemic and Temporal Reasoning.* Oxford Science Publications, 1995.

[Geffnerand Pearl, 1992] H. Geffner and J. Pearl. Conditional entailment: Bridging two approaches to default reasoning. *Artificial Intelligence,* 53(2-3):209-244,1992.

[Horty, 1994] J.F. Horty. Some direct theories of nonmonotonic inheritance. In D.M. Gabbay, C.J. Hogger, and J.A. Robinson, editors, *Nonmonotonic Reasoning and Uncertain Reasoning,* volume 3 *of Handbook of Logic in Artificial Intelligence and Logic Programming,* pages 111-187. Oxford: Clarendon Press, 1994.

[Krause et al., 1990] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence,* 44(1-2): 167-207, 1990.

[Lehmann and Magidor, 1992] D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence,* 55(1): 1-60, 1992.

[Lewis, 1973] D.Lewis. *Counterfactuals.* Harvard University Press, 1973.

[McCain and Turner, 1997] N. McCain and H. Turner. Causal theories of action and change. In *Proceedings of the AAAI National Conference on Artificial Intelligence,* pages 460-465, Providence, RI, 1997.

[McCarthy, 1986] J. McCarthy. Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence,* 28:89-116, 1986.

[Neufeld, 1989] E. Neufeld. Defaults and probabilities; extensions and coherence. In *Proceedings of the First International Conference on the Principles of Knowledge Representation and Reasoning,* pages 312-323, Toronto, Canada, 1989.

[Pearl, 1989] J. Pearl. Probabilistic semantics for nonmonotonic reasoning: A survey. In *Proceedings of the First International Conference on the Principles of Knowledge Representation and Reasoning,* pages 505-516, Toronto, May 1989. Morgan Kaufman.

[Pearl, 1990] J. Pearl. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In R. Parikh, editor, *Proc. of the Third Conference on Theoretical Aspects of Reasoning About Knowledge,* pages 121-135, Pacific Grove, Ca., 1990. Morgan Kaufmann Publishers.

[Poole, 1991] D.L.Poole. The effect of knowledge on belief: Conditioning, specificity, and the lottery paradox in default reasoning. *Artificial Intelligence,* 49(l-3):281-307,1991.

[Savage, 1972] L.J. Savage. *The Foundations of Statistics.* Dover Publications Inc., New York, second edition, 1972.