ELSEVIER

Artificial Intelligence 105 (1998) 105–137

Artificial Intelligence

On first-order conditional logics

James P. Delgrande [1]

School of Computing Science, Simon Fraser University, Burnaby, British Columbia, Canada V5A 1S6

Received 5 May 1997; received in revised form 14 July 1998

Abstract

Conditional logics have been developed as a basis from which to investigate logical properties of "weak" conditionals representing, for example, counterfactual and default assertions. This work has largely centred on propositional approaches. However, it is clear that for a full account a first-order logic is required. Existing or obvious approaches to first-order conditional logics are inadequate; in particular, various representational issues in default reasoning are not addressed by extant approaches. Further, these problems are not unique to conditional logic, but arise in other nonmonotonic reasoning formalisms. I argue that an adequate first-order approach to conditional logic must admit domains that vary across possible worlds; as well the most natural expression of the conditional operator binds variables (although this binding may be eliminated by definition). A possible worlds approach based on Kripke structures is developed, and it is shown that this approach resolves various problems that arise in a first-order setting, including specificity arising from nested quantifiers in a formula and an analogue of the lottery paradox that arises in reasoning about default properties.

Keywords: Knowledge representation; Conditional logic; First-order logic; Nonmonotonic reasoning

1. Introduction

Conditional logics have been developed in Artificial Intelligence for representing and reasoning about defeasible conditionals such as "birds fly" [8,11,24]. This work in turn has built on earlier work using conditional logics for dealing with weak subjunctive conditionals, most notably counterfactuals [27,48]. The essential idea is that a sentence representing a default or counterfactual appears as an object in a logical theory; thus it makes sense to talk of a set of defaults entailing others. The semantic theory of such conditionals is typically developed by means of a possible worlds approach based, one

[1] Email: jim@cs.sfu.ca.

0004-3702/98/$ – see front matter
PII: S0004-3702(98)00088-5

way or another, on a preference ordering among possible worlds. This ordering is then used to determine the truth of a conditional. The default statement "birds normally fly" for example may be represented using a weak conditional as $Bird \Rightarrow Fly$. The intended reading is something like "if an individual is a bird, then normally that individual flies". $Bird \Rightarrow Fly$ is true (roughly) if in the *least* (or most normal) worlds in which *Bird* is true, *Fly* is true also. "Penguins are necessarily birds" can be expressed as $\Box(Penguin \supset Bird)$, and along with this we can consistently state that penguins normally do not fly, $Penguin \Rightarrow \neg Fly$. From this one can conclude that birds are normally not penguins (i.e., $Bird \Rightarrow \neg Penguin$ is logically entailed).

Given a logic of default conditionals, one can continue and define forms of defeasible or nonmonotonic inference. The most basic of these is given by:

$$\beta \text{ follows by default from } \alpha \text{ in theory } \mathcal{T} \text{ just when } \mathcal{T} \models \alpha \Rightarrow \beta. \tag{1}$$

Consequently one would conclude that a bird flies (by default) whereas a penguin or bird with a broken wing does not. This form of defeasible inference however is much too weak, a limitation addressed in, for example, [7,18,29,38]. The problem of specifying a (propositional) nonmonotonic reasoning system with "appropriate" properties, including appropriate handling of specificity, inheritance of properties, and accommodating irrelevant properties, remains an area of active research.

However, these considerations aside, there are limitations with a propositional account having to do solely with representational issues. In general, a propositional logic will be inadequate for modelling phenomena of interest, and so ultimately one will require the richer expressiveness of a first-order logic. In the case of flying birds for example, "birds fly" refers (in some sense) to the set or class of birds, and "Tweety is a bird that flies" is an instance. So we might prefer to write something like:

$$\forall x(Bird(x) \Rightarrow Fly(x)), \tag{2}$$

$$Bird(Tweety) \Rightarrow Fly(Tweety). \tag{3}$$

However, if we read (2) as "for every x, if x is a bird then normally x flies" then, without worrying about the precise semantics, this reading is not quite right. For one thing "birds fly" seems to attribute a property to the class of birds; Eq. (2) on the other hand seems to express a normality condition concerning each individual that happens to be a bird.

As shown in the next section, the problems run considerably deeper than this. First, in the obvious, naive, first-order approach there are perfectly reasonable theories that cannot be consistently represented. The difficulty is that a principle of universal instantiation and, in particular, instantiation into the scope of the conditional operator, leads to inconsistency. Second, a representational analogue to the lottery paradox arises in a first-order setting. These problems are not peculiar to conditional logic, but rather appear in other approaches to default reasoning. I argue that the problem arises from assuming an unfettered principle of universal instantiation. An alternative approach is developed where this principle is weakened. In this approach the aforementioned problems are satisfactorily resolved. Crucially, in this approach the domain of individuals may vary across possible worlds and the conditional operator may (effectively) bind variables.

While this paper is motivated by and centred on logics for defeasible conditionals, it is intended to account for first-order issues in conditional logics in general. That is, I address not just first-order conditional logics for defaults, but first-order conditional logics in general, including those intended to represent counterfactuals, notions of obligation, hypotheticals, and other subjunctive conditionals. In addition, the resulting approach is more expressive than found, for example, using nonmonotonic consequence operators [22]. As well, the results given here are relevant to other approaches to nonmonotonic reasoning, such as default logic and circumscription, in that they help explicate the notion of an *exceptional individual.*

Section 2 explores issues in quantification in nonmonotonic reasoning. Section 3 reviews conditional logic and related work. In Section 4 an approach is developed that addresses problems in a naive quantificational approach. Section 5 provides a further exploration of this approach, while Section 6 compares it with related work. Section 7 contains concluding remarks. Proofs of theorems are given in Section 8.

2. Problems with quantification in nonmonotonic reasoning

In this section, I informally consider various problems with the "obvious" approach to quantification in conditional logics, and in nonmonotonic reasoning in general. Consider first a very standard example, where penguins are birds, birds normally fly, whereas penguins do not:

Example 1.

$$\forall x\,(Bird(x) \Rightarrow Fly(x)),$$
$$\forall x\,(Penguin(x) \Rightarrow \neg Fly(x)),$$
$$\forall x\,\Box(Penguin(x) \supset Bird(x)).$$

For the time being, consider $\alpha \Rightarrow \beta$ to represent a conditional in a conditional logic, or a default in some arbitrary approach to nonmonotonic reasoning. So in default logic [43], $\forall x\,(\alpha(x) \Rightarrow \beta(x))$ would be expressed using the rule

$$\frac{\alpha(x) : \beta(x)}{\beta(x)},$$

whereas in a circumscriptive abnormality theory [32] it might be expressed by

$$(\alpha(x) \wedge \neg ab_{\alpha}(x)) \supset \beta(x).$$

In these last two cases "penguins are birds" would be expressed using material implication:

$$\forall x\,(Penguin(x) \supset Bird(x)).$$

Without going into details (but see Section 3), if we have that *Opus* is a penguin, and so a bird, then in a conditional logic (or related approach) we can conclude that *Opus* normally does not fly. Essentially, in a conditional logic, we have the information that being a penguin is a more specific notion than being a bird, and a definition of default inference (as given in (1) as a base case) can make implicit use of this information. In approaches such as default logic and circumscription this is not the case, and other information has to be added to a theory to block the unwanted inference that *Opus* flies by virtue of being a bird. In default logic, for example, one could employ the rule

$$\frac{Bird(x) : \neg Penguin(x) \wedge Fly(x)}{Fly(x)},$$

stating that birds that can be assumed to not be penguins fly.

A second example indicates that perfectly reasonable theories can not be handled by a naive approach to quantification. In this example we have another apparent case of specificity, but where specificity is now given by the nesting of quantifiers [3,28]. Consider where elephants normally like their keepers; however elephants normally do not like keeper Fred, who is mean; but elephant Clyde, who gets along with everyone, normally does like Fred.

Example 2.

$$\forall xy(El(x) \wedge Ke(y) \Rightarrow Likes(x, y)), \tag{4}$$

$$\forall x(El(x) \wedge Ke(Fred) \Rightarrow \neg Likes(x, Fred)), \tag{5}$$

$$El(Clyde) \wedge Ke(Fred) \Rightarrow Likes(Clyde, Fred). \tag{6}$$

Allowing unrestricted universal instantiation, we have the instance of (5):

$$El(Clyde) \wedge Ke(Fred) \Rightarrow \neg Likes(Clyde, Fred). \tag{7}$$

However in a conditional logic, $El(Clyde)$ and $Ke(Fred)$ together with (6) and (7) are inconsistent. This is clearly a significant problem, since Example 2 is straightforward and reasonable. In default logic and circumscription the problem is less severe. In default logic for example, one obtains one *extension*, or set of beliefs, in which $Likes(Clyde, Fred)$ is true and another in which $\neg Likes(Clyde, Fred)$ is true. Again, it is up to the user to hand-tool the default theory to eliminate the unwanted extension.

The difficulty seems to be that in (4) what we mean to express is "elephants like their keepers", while in (5) we have something like "elephants do not like keeper Fred". In a sense we are talking about elephants and keepers as a whole, or as classes. But our treatment of quantification in the above is such that what we really have is a shorthand for expressing something about every individual, and this is what leads us into trouble. The problem is that, one way or another, we want our assertions to say something about elephants and keepers *in general*. In conditional logic, the truth of a conditional is determined by looking at the least worlds in which the antecedent is true. For (4), we should somehow exclude Fred, since he is an exceptional keeper. Similarly for (5) we would like to exclude Clyde in considering those worlds used to determine the truth of the conditional. This is the intuition that will guide the approach of the Section 4; however, first consider another difficulty deriving from universal instantiation involving the conditional operator.

Example 3. An analogue of the lottery paradox, as applied to default properties, crops up in approaches to default reasoning. This is a consequence of the observation that, for example, every species of bird is in some way exceptional with respect to the general class of birds [40]. [2] Thus penguins do not fly, ravens are black, hummingbirds hover, albatrosses are pelagic, etc. If we assert that all species of known birds constitutes the full set of species:

$$\forall x\,(Bird(x) \equiv (Penguin(x) \vee Robin(x) \vee \cdots \vee Raven(x))) \tag{8}$$

then as argued in [40], we essentially lose our "birds fly" default. On the other hand, if we decided that *necessarily* all species of known birds constitutes the full set of species:

$$\forall x\,\Box(Bird(x) \equiv (Penguin(x) \vee Robin(x) \vee \cdots \vee Raven(x))) \tag{9}$$

then if every species is exceptional in some manner, our knowledge base is unsatisfiable (unless there are no birds).

This is a qualitative, or finitary, version of the lottery paradox [23]: given a lottery with some number of players, we assume that a given player typically will not will the lottery, but that someone does win the lottery. There are various ways that this could be stated, the major alternatives being the following:

$$\forall x\,(\top \Rightarrow \neg Winner(x)), \quad \exists x\, Winner(x), \tag{10}$$

$$\forall x\,(\top \Rightarrow \neg Winner(x)), \quad \Box\exists x\, Winner(x), \tag{11}$$

$$\forall x\,(\top \Rightarrow \neg Winner(x)), \quad \top \Rightarrow \exists x\, Winner(x). \tag{12}$$

The first equation seems to mis-state the problem, since we want to say more than that there simply happens to be a winner. In (11) we assert that there must be a winner, while in (12) we assert that normally there is a winner. While (11) and (12) capture two types of lotteries, both are unsatisfiable in a naive approach: we would have that $\forall x(\top \Rightarrow \neg Winner(x))$ logically entails $\top \Rightarrow \forall x \neg Winner(x)$, which is inconsistent with the assertion that (however expressed) there is a winner.

3. Background

The next subsection introduces conditional logic and sketches the obvious extension to a first-order system. The second subsection reviews related work concerning conditional approaches and quantification in nonmonotonic reasoning systems.

3.1. Conditional logic

Our development follows that of [9] for propositional conditional logic, since arguably it provides the most general framework suitable for representing weak conditionals; see [34] for more on this approach and for a comparison of different conditional logics. In the propositional case, the language of the logic is that of (standard, classical) propositional logic augmented with a binary operator $\Rightarrow$ for the weak conditional, reserving $\supset$ for material implication. So the language for propositional logics of conditionals $\mathcal{L}_C$ is the

[2] Poole [40] also discusses problems that arise in nonmonotonic inference. Our concerns here are restricted to representational issues.

set of formulas constructed from a set of atomic sentences $P = \{p_1, p_2, \ldots\}$, the binary operator $\Rightarrow$, and the standard truth-functional connectives $\neg$ and $\supset$, in the usual manner. The connectives $\wedge$, $\vee$, and $\equiv$ are introduced by definition. In addition $\top$ abbreviates some (classical) tautology and $\bot$ abbreviates $\neg\top$. The operator $\Rightarrow$ is the *weak* or *variable* conditional, expressing a subjunctive relation between its antecedent and consequent. Nestings of the weak conditional are permitted, although this will not be important in our analysis. Sentences are interpreted in terms of a *model* $M = \langle W, R, P\rangle$ where:

(1) W is a set (of worlds),
(2) R is a ternary accessibility relation on worlds, and
(3) P maps atomic sentences onto subsets of W.

$R_w w_1 w_2$ has the informal reading "according to world w, w_2 is accessible from w_1". In logics of defaults, the intent is that w_2 is no more exceptional than w_1; that is, as "further" accessible worlds are considered, these worlds are more "normal". [3]

For all intents and purposes I will treat R as a binary relation on its last two arguments; thus the first argument is subscripted as a reminder that we are (essentially) dealing with a binary relation, but indexed by particular worlds. Informally this means that each world has its own idea as to which worlds are accessible from each other. Thus under a "normality" reading for accessibility, a world where gravity holds will have a different idea of normality than one where it does not. For conditional logics, regardless of application, the accessibility relation (on its last two arguments) is generally a preorder. Hence we assume that R is reflexive and transitive on its last two arguments. [4] I also adopt the Limit Assumption [27], that the order on worlds is well-founded. A necessity operator can be defined by $\Box\alpha = \neg\alpha \Rightarrow \alpha$. That is, $\Box\alpha$ will be true just when at the least $\neg\alpha$ worlds, α is true; since there are no $\neg\alpha$ worlds in which α is true, this means that every accessible world has α true.

In the following definitions, W_w is the set of worlds "visible" from w, $\|\alpha\|^M$ denotes the set of worlds where α is true, while $min(w, \|\alpha\|^M)$ denotes the least exceptional of such worlds visible from w: [5]

Definition 3.1.

$$W_w = \{w_1 \mid \text{ for some } w_2 \in W,\ R_w w_1 w_2\}.$$

Definition 3.2.

$$\|\alpha\|^M = \{w \in W \mid M, w \models \alpha\}.$$

[3] In contrast, in a *centred* logic of counterfactuals [27], Rww_1w_2 means that (seen from w) w_2 is closer to w than w_1.

[4] The resulting logic has been suggested as providing a *conservative core* for defaults: logical relations deriving from this base (so it is argued) ought to be common to *any* logic of defaults. This logic is also appropriate as a base for defining counterfactual conditionals and other subjunctives, with the possible exception of deontic logics. A deontic logic arguably ought not be based on a reflexive accessibility relation. However we retain a reflexive accessibility relation here since it eases the technical development.

[5] The following definitions make reference to $\models$, defined next. While these definitions are interdependent, they are not (viciously) circular.

Definition 3.3.

$$min(w, \|\alpha\|^M) = \{w_1 \in \|\alpha\|^M \mid \text{for every } w_2 \in \|\alpha\|^M \text{ where } R_w w_1 w_2, \text{ we have } R_w w_2 w_1\}.$$

Truth at world w in model M ($\models$) is as for propositional logic, except for the $\Rightarrow$ operator:

(1) $M, w \models p_i$ for $p_i \in \mathbf{P}$ iff $w \in P(p_i)$.
(2) $M, w \models \neg\alpha$ iff not $M, w \models \alpha$.
(3) $M, w \models \alpha \supset \beta$ iff: if $M, w \models \alpha$ then $M, w \models \beta$.
(4) $M, w \models \alpha \Rightarrow \beta$ iff $min(w, \|\alpha\|^M) \subseteq \|\beta\|^M$.

Thus $\alpha \Rightarrow \beta$ is true at w when β holds at the least α worlds according to R and w. So for the conditional *Bird* $\Rightarrow$ *Fly*, we look at the "least" (if such exist) worlds in which *Bird* is true, and if *Fly* is true at these worlds then *Bird* $\Rightarrow$ *Fly* is true at w. The intuition is that for determining the truth of a weak conditional, we factor out exceptional circumstances such as having a broken wing, being a penguin, etc.

The resulting logic has some quite reasonable properties. For example, the following sets of sentences are nontrivially and simultaneously satisfiable:

$$\{Bird \Rightarrow Fly,\ Bird,\ \neg Fly\},$$

$$\{Bird \Rightarrow Fly,\ Penguin \Rightarrow \neg Fly,\ \Box(Penguin \supset Bird)\},$$

$$\{Bird \Rightarrow Fly,\ Bird \wedge BW \Rightarrow \neg Fly\}.$$

Hence, first, a bird may normally fly although in actual fact it does not; second, birds fly, penguins do not, but penguins are birds; and third, while birds fly, birds with broken wings do not. The underlying modal logic (treating R as a binary relation) is S4 [19]; Lamarre [24] presents a closely-related logic. Optionally we could have defined $\Rightarrow$ in term of $\Box$ [8]. The resulting conditional logic $\mathcal{S}$ is characterised as follows.

Axioms consist of tautologies of classical propositional logic together with:

ID. $\alpha \Rightarrow \alpha$.

ASC. $\alpha \Rightarrow \beta \supset (\alpha \Rightarrow \gamma \supset \alpha \wedge \beta \Rightarrow \gamma)$.

CA. $(\alpha \Rightarrow \gamma \wedge \beta \Rightarrow \gamma) \supset (\alpha \vee \beta \Rightarrow \gamma)$.

In addition to modus ponens, there are rules of inference:

RCEA. From $\alpha \equiv \alpha'$ infer $(\alpha \Rightarrow \beta) \equiv (\alpha' \Rightarrow \beta)$.

RCK. From $(\beta_1 \wedge \cdots \wedge \beta_n) \supset \beta$ infer $(\alpha \Rightarrow \beta_1 \wedge \cdots \wedge \alpha \Rightarrow \beta_n) \supset \alpha \Rightarrow \beta$ for $n \geqslant 0$.

This axiomatisation is taken from [34] since it is somewhat more perspicuous than that of [9]. The names of axioms and rules of inference are relatively standard in conditional logics; see [10,35]. Notions of theoremhood, consistency, etc., are standard, as is the notion of deductive consequence.

Definition 3.4. A formula α is a *deductive consequence* of a set of formulas Γ, written $\Gamma \vdash \alpha$, iff there are $\{\gamma_1, \ldots, \gamma_n\} \subseteq \Gamma$ such that $\vdash \gamma_1 \wedge \cdots \wedge \gamma_n \supset \alpha$.

In our original example, we obtain:

$$\{Bird \Rightarrow Fly,\ Penguin \Rightarrow \neg Fly,\ \Box(Penguin \supset Bird)\} \vdash Bird \Rightarrow \neg Penguin$$

and so we can deduce that birds are not normally penguins, given our initial conditionals.

The obvious extension to quantification is relatively straightforward. I describe it here in order to provide a point of departure for the subsequent approach. "Birds fly" is represented as $\forall x(Bird(x) \Rightarrow Fly(x))$; this formula can be regarded, essentially, as standing for all its instances. So if this formula is true then we would expect $Bird(t) \Rightarrow Fly(t)$ to also be true, where t is an arbitrary term. However, as described in the previous section, this approach is not satisfactory. Moreover, the problem lies not with conditional logic per se, but rather with allowing "unfettered" instantiation into a default; hence as described in the previous section, similar problems arise in other approaches dealing with defaults.

The development follows [11]. [6] The language is based on that of classical first-order logic, beginning with denumerable sets of variables, constants, and predicate symbols, but not, for simplicity, function symbols; again the language is augmented with the operator $\Rightarrow$. The symbol $\forall$ is used for universal quantification, and the existential quantifier $\exists$ is introduced by definition in the usual way. Variables and constants constitute the set of *terms* in the language. Formulas are interpreted in a model $M = \langle W, R, D, V\rangle$ where W and R are as before, D is a nonempty domain of individuals, and V is a function on terms and predicate symbols where:

(1) for term t, $V(t) \in D$,
(2) for any n-place predicate symbol P, $V(P)$ is a set of $(n+1)$-tuples $\langle d_1, \ldots, d_n, w\rangle$ where $d_i \in D$ for $1 \leqslant i \leqslant n$, and $w \in W$.

Given a model $M = \langle W, R, D, V\rangle$, truth at world w is given by:

(1) $M, w \models \neg\alpha$ and $M, w \models \alpha \supset \beta$ and $M, w \models \alpha \Rightarrow \beta$ are defined as in the propositional case.
(2) For n-place predicate symbol P, terms $t_1, \ldots, t_n$, and $w \in W$,

$$M, w \models P(t_1, \ldots, t_n) \quad \text{iff} \quad \langle V(t_1), \ldots, V(t_n), w\rangle \in V(P).$$

(3) For arbitrary variable x and formula α, $M, w \models \forall x\alpha$ iff for every V' which is the same as V except possibly $V(x) \neq V'(x)$, and where $M' = \langle W, R, D, V'\rangle$, we have $M', w \models \alpha$.

Definitions of satisfiability and validity are analogous to those in the propositional case. This approach will be referred to as the "naive approach" to distinguish it from what follows. Note that, in particular, the following sentence is valid in this approach:

$$\forall x\alpha(x) \supset \alpha(t) \text{ if } t \text{ is a term free for } x \text{ in } \alpha(x). \tag{13}$$

Thus from $\forall x(Bird(x) \Rightarrow Fly(x))$ we can derive $Bird(Opus) \Rightarrow Fly(Opus)$.

The domain D is fixed across possible worlds, and so the Barcan formula and its converse (see [19], for example) are valid. Moreover the assignment of terms to domain individuals is also fixed across possible worlds, and so in a model terms denote the same domain element across possible worlds. That is, terms *rigidly designate* across possible

[6] The propositional fragment of this logic in fact was based on a world-selection function, rather than an accessibility relation. This difference, for our concerns here, is irrelevant.

worlds. This is not a limitation, in that there is no problem also incorporating nonrigid terms (in Artificial Intelligence, see, for example, [26]).

3.2. Related work

In Artificial Intelligence, most work in conditional logic has focussed on logics for default statements. The base logic of the previous section provides a "core" of default inferences (via (1)) that arguably should hold in any logic of defaults; one may add conditions to augment this basic set. Alternately, by augmenting the base approach one may obtain logics suitable for other weak conditionals, such as counterfactuals.

Used as a logic of defaults, conditional logics have an advantage over approaches such as default logic [43] and circumscription [31], in that one can reason about default conditionals, and it makes sense to talk of a set of defaults implying another, or of being unsatisfiable. Given a conditional logic of defaults, one may subsequently determine what nonmonotonic inferences ought to obtain. So in such approaches, representational issues, having to do with what it means to be a default, are separated from issues of default application. [7] The large majority of work in conditional logic has centred on propositional systems.

Conditional logic however is but one way in which formal properties of defaults may be investigated. Other approaches have been based on intuitions from probability theory [36], or on the explicit development of nonmonotonic consequence operators [22,25]. Despite the diversity of approaches and intuitions, there has been a close agreement on what constitutes a core set of inferences that ought to be common to all nonmonotonic systems. Systems such as ε-entailment [36] (or 0-entailment or p-entailment [1]), possibilistic logic [12], preferential entailment [22], and CT4 [7], among others, essentially allow the same inferences as S, and may be taken as specifying a *conservative core* [37] of default inferences. These approaches though are too weak to constitute a general theory of nonmonotonicity. For example, *irrelevant* properties are not handled. Thus, even though a bird may be concluded to fly by default, a green bird cannot be concluded to fly by default since there may be models of a theory where a green bird does not fly. So these systems, while arguably adequate for *representing* and *reasoning about* sets of defaults, do not adequately address the problem of what nonmonotonic consequences should follow from a set of defaults. Subsequent work has focussed on principled means to extend a system's basic inferences.

Again, there has been a strong convergence on means of strengthening these systems. Approaches including System **Z** and 1-entailment [38], CO^* [7], possibilistic entailment [5], and rational closure [28] all assume, in a semantic sense, that a world is as unexceptional as possible. This may be done by ranking default rules according to a notion of specificity, and then using this ranking to obtain a preference ordering on worlds (or, in some approaches, models). Thus, since there is no reason to suppose that greenness has any bearing on flight, one assumes that greenness has no effect on flight. While this assumption seems reasonable enough, its realisation in these systems is not unproblematic.

[7] This distinction is not clear-cut however. For example, (1) shows how a logic of conditionals may be used to provide a base notion of default inference.

In brief, these approaches fail to fully address issues of *(ir)relevance* and *inheritance* of properties; as well they allow unwanted *specificity* relations [38]. These approaches have been extended in various ways, including [4,16,17]. However no extension is completely satisfactory. Our concerns in this paper though are orthogonal to these, which address propositional nonmonotonic reasoning. Rather, here we look at first-order representational issues.

With respect to first-order systems, there has been little work in quantified conditional logic. Stalnaker and Thomason [47] present an analysis of one such logic, of counterfactual conditionals. Domains of individuals (i.e., the range of bound individual variables) may vary across possible worlds; a designated set of individuals $\mathcal{D}'$ serves as the *outer domain*, a set of "individuals" which exist in no possible world, while a designated world λ serves as the *absurd* world. The outer domain is used for the assignment of truth values to nonreferring terms. The absurd world serves as a technical device for handling the truth values of conditionals with impossible antecedents. The logic has properties that make it unsuitable for a logic of defaults. Foremost is the fact that if α is true at world w, then $\alpha \Rightarrow \beta$ behaves the same as $\alpha \supset \beta$ at that world. Moreover, at world w there is at most a single "least" world accessible from w in which α is true. Lastly, quantification over $\Rightarrow$ is handled roughly as outlined in the previous subsection, in that the truth of sentence $\forall x(\alpha(x) \Rightarrow \beta(x))$ at world w is determined with respect to the domain of individuals at w. As we argue subsequently, we in fact want to use domains of individuals at *other* possible worlds.

Asher and Morreau [2] describe a first-order logic for generic sentences. The propositional fragment is weaker than the base logic $\mathcal{S}$, although the authors also describe a process of "normalisation" whereby further inferences may be made. Again, the sentence

$$\forall x(\alpha(x) \Rightarrow \beta(x)) \supset (\alpha(c) \Rightarrow \beta(c))$$

is valid for any constant c. So quantification into $\Rightarrow$ is handled basically as described in the previous subsection. Similar remarks may be made concerning [28], which treats quantification in the setting of nonmonotonic consequence relations.

Bacchus et al. [6] present an approach to default reasoning based on statistical notions, wherein default reasoning is founded on a principle of *indifference* among possible worlds. Friedman et al. [13] develop a first-order conditional logic based on *plausibility measures*. A plausibility measure is a function that associates a notion of plausibility (a member in some partially ordered set) with a set of possible worlds. Schlechta [44] formalises a notion of "normally $\alpha(x)$" using restricted quantification. In this interpretation of defaults, notions of preference are not relevant in the semantics; this approach is combined with that of preferential models in [45]. The work of Lehmann and Magidor [28], Schlechta [44,45] and Friedman et al. [13] is discussed further in Section 6. For quantified modal logic, Hughes and Cresswell [19] provide a readable introduction, while Fitting [14] and Garson [15] provide comprehensive surveys.

Of other approaches to nonmonotonic reasoning, most, including default logic [43], circumscription [31,32], autoepistemic logic [33], and Theorist [39], allow quantification. The problems described in Section 2 appear in these systems, but in a less severe form. For example, the natural representation of Example 2 in default logic is:

$$\frac{El(x) \wedge Ke(y) : Likes(x, y)}{Likes(x, y)},$$
$$\frac{El(x) \wedge Ke(Fred) : \neg Likes(x, Fred)}{\neg Likes(x, Fred)},$$
$$\frac{El(Clyde) \wedge Ke(Fred) : Likes(Clyde, Fred)}{Likes(Clyde, Fred)}.$$

Given that $El(Clyde) \wedge Ke(Fred)$ is true, two extensions (or sets of default beliefs) are obtained, one in which $Likes(Clyde, Fred)$ is true and one in which $\neg Likes(Clyde, Fred)$ is true. So we do obtain default conclusions in this case, unlike the conditional analogue, which results in inconsistency. However, we would have to add a condition such as $x \neq Clyde$ in the second default to block the undesirable extension.

4. An alternative approach

4.1. Initial considerations

In this section I develop an alternative to the naive approach. Two modifications to the semantics appear to be necessary:

(1) The domain of individuals may vary across possible worlds.
(2) The conditional operator may (effectively) [8] bind variables.

Example 2 is sufficient to argue for these points, in that, for the assertion that elephants like their keepers (4) we want to have that at the least elephant-and-keeper worlds, in some fashion, "elephants like their keepers" is true. However, we exclude *Fred* from consideration, since he is an exceptional individual whom elephants do not like. Thus, we want to restrict this default to *relevant* individuals or individuals not known to be exceptional, of which *Fred* is not one. Similarly for elephants not liking Fred (5): we implicitly exclude *Clyde*, who gets along with everyone, including Fred (6).

Now, an instance of (5) is $El(Clyde) \wedge Ke(Fred) \Rightarrow \neg Likes(Clyde, Fred)$, and this conflicts with (6). So arguably we do not want to allow this instantiation; this is intuitively plausible since *Clyde* is exceptional with respect to (5). Intuitively, in (5) we want x to range over nonexceptional individuals, including those elephants that do not like *Fred*. Hence, for the expression of this conditional, we want x, one way or another, to range over individuals determined by worlds *other* than that being modelled: in this case the most normal elephant-and-keeper-Fred worlds.

So for a default such as "birds normally fly", what we really mean is "at the least exceptional worlds in which there are birds, birds fly". Essentially we look at the least worlds in which there are birds, and if all of these birds fly then "birds normally fly" is true. At these most normal worlds an individual such as *Opus* may or may not be in the domain of quantification. So for "birds normally fly" we are in fact quantifying over individuals at *other* worlds and, as we concluded above, the domain of quantification may differ at those worlds from that at the world at hand.

[8] The qualification is due to the fact that, as subsequently shown, the variable-binding operator can be defined away.

Informally, "birds normally fly" is true iff:

At the least worlds in which there are birds: for every x, if x is a bird then x flies. (14)

Syntactically one way to represent this is by:

$$Bird(x) \Rightarrow_x Fly(x) \tag{15}$$

where $\Rightarrow$ now takes a list or tuple of distinct variables, bound in its scope,[9] along with an antecedent and consequent constituting the scope of the binding.

However, assuming that we allow the domain of individuals to vary (in some sense) across possible worlds, (14) can be expressed directly, and without recourse to such a variable binding operator. If we use $\exists_c$, $\forall_c$ to indicate that we are quantifying over "actual" individuals, we can express "birds fly" as follows:

$$\exists_c x Bird(x) \Rightarrow \forall_c x (Bird(x) \supset Fly(x)).$$

That is: at the least worlds in which there are "actual" or "unexceptional" birds, for every such individual x, if x is a bird then x flies. This last formula is somewhat cumbersome, and (15) seems to more naturally express "birds fly". So in the next subsection I retain (15) as the official expression of "birds fly", with $\Rightarrow_{\bar{x}}$ introduced by definition.

This gives rise to a small but important consideration regarding language of discourse. I envisage a first-order conditional theory being expressed in the language of classical first-order logic augmented solely by a variable-binding conditional operator $\Rightarrow_{\bar{x}}$. However, in explicating the semantics of this approach, I appeal to a subsuming language, $\mathcal{L}_{FC}$, that includes the "weak" quantifiers, $\exists_c$ and $\forall_c$, as well as a predicate A intended to be interpreted as "unexceptional". The operator $\Rightarrow_{\bar{x}}$ can then be introduced by definition by reference to these additions. So the role of $\exists_c$, $\forall_c$, A is to explicate the meaning of $\Rightarrow_{\bar{x}}$, but there would be no need for these symbols to appear explicitly in the statement of a conditional theory.

4.2. *The first-order conditional logic*

Let $\mathcal{L}_{FC}$ be the language of classical first-order logic, involving denumerable sets of variables, constants, and predicates symbols, the truth functional connectives $\neg$ and $\supset$ as well as the binary operator $\Rightarrow$ and a second quantifier symbol $\forall_c$ in addition to the usual $\forall$. The set of terms is comprised of the set of constants and variables. Formulas are constructed as expected, by the usual recursive definition. The connectives $\wedge$, $\vee$, $\equiv$ are introduced by definition; $\exists$ is defined as $\neg\forall\neg$ and $\exists_c$ is defined as $\neg\forall_c\neg$. In addition, we have a designated unary predicate symbol A whose extension at a world is the domain of quantification, as reflected by $\forall_c$, at that world. Sentences are interpreted in terms of a model, a *QS-model*, $M = \langle W, R, D, D', V\rangle$ where W, R, and D are as in Section 3.1. D' is a function that assigns a non-empty subset $D'(w)$ of D to each possible world w. V is a function on terms, and predicate symbols and worlds, defined as follows:

(1) For term t, $V(t) \in D$.

[9] In fact, we would have a set of operators, one for each arity. For simplicity I will treat $\Rightarrow$ as a single operator.

(2) For n-place predicate symbol P we have $V(P, w)$ is a set of n-tuples $\langle d_1, \ldots, d_n \rangle$ where each $d_i \in D$.
For predicate symbol A we have in addition that $V(A, w) = D'(w)$.

The intent is that there be a "universal" domain D, consisting of the set of all individuals. At each world w there is a nonempty set $D'(w)$ of individuals "existing" or "actual" at w. The predicate A is true of just those individuals existing at a world; that is $A(t)$ is true at w just if the individual denoted by t is a member of $D'(w)$. So $A(t)$ can be read as "t is actual"; in Section 5, I will suggest reinterpreting $A(t)$ as "t is a *relevant* or *unexceptional* individual". Terms rigidly designate the same individual across possible worlds. Hence *Opus* denotes the same individual in every possible world, although *Opus* may or may not be actual at any given world.

Crucially, the quantifier $\forall_c$ will range over the set of individuals $D'(w)$ at world w. Intuitively we are concerned with knowledge bases constructed from the language of first-order logic together with the variable-binding operator $\Rightarrow_{\vec{x}}$. As indicated previously, the formal development is eased considerably if we define $\Rightarrow_{\vec{x}}$ in terms of our quantifier $\forall_c$, ranging over "actual" individuals, along with the predicate A.

Given a model $M = \langle W, R, D, D', V \rangle$, truth at a world w is given as follows:

Definition 4.1.

(1) $M, w \models \neg\alpha$, and $M, w \models \alpha \supset \beta$, and $M, w \models \alpha \Rightarrow \beta$ are defined as in the propositional case.

(2) For n-place predicate symbol P, terms $t_1, \ldots, t_n$, and $w \in W$, we have $M, w \models P(t_1, \ldots, t_n)$ iff $\langle V(t_1), \ldots, V(t_n) \rangle \in V(P, w)$.

(3) (a) $M, w \models \forall x \alpha$ iff for every V' which is the same as V except possibly $V(x) \neq V'(x)$, and where $M' = \langle W, R, D, D', V' \rangle$, we have $M', w \models \alpha$.
(b) $M, w \models \forall_c x \alpha$ iff for every V' which is the same as V except $V'(x) \in D'(w)$ where $M' = \langle W, R, D, D', V' \rangle$, we have $M', w \models \alpha$.

Definitions of satisfiability and validity are the same as in the propositional case.

For the proof theory, we need to state that the range of the quantifier $\forall_c$ is restricted to actual individuals. So for universal instantiation, informally, we need to say that we can conclude $\alpha(t)$ from $\forall_c x \alpha(x)$ only when the individual denoted by t is actual. Conversely, for universal generalisation, we can conclude that a formula is universally true (in the sense of $\forall_c$) if it is true for all individuals that are actual. This in turn means that at a world we can have that $\forall_c x(Horse(x) \supset \neg Fly(x))$ and $Horse(Pegasus) \wedge Fly(Pegasus)$ are true, without inconsistency, provided that we also have $\neg A(Pegasus)$. At this same world we would of course have that $\neg\forall x(Horse(x) \supset \neg Fly(x))$ is true.

The resulting logic is called $\mathcal{QS}$. The axioms of $\mathcal{QS}$ consist of all instances of classical propositional tautologies, as well as the following schemata, where x is taken as ranging over the set of variables.

ID. $\alpha \Rightarrow \alpha$.

ASC. $\alpha \Rightarrow \beta \supset (\alpha \Rightarrow \gamma \supset \alpha \wedge \beta \Rightarrow \gamma)$.

CA. $(\alpha \Rightarrow \gamma \wedge \beta \Rightarrow \gamma) \supset (\alpha \vee \beta \Rightarrow \gamma)$.

Ax. $\exists x A(x)$.

UI. $\forall x\alpha(x) \supset \alpha(t)$ if t is a term free for x in $\alpha(x)$.

FUI. $\forall_c x\alpha(x) \supset (A(t) \supset \alpha(t))$ if t is a term free for x in $\alpha(x)$.

CUG. $\forall x(\alpha \Rightarrow \beta) \supset (\alpha \Rightarrow \forall x\beta)$ for x not free in α.

The rules of inference are modus ponens, the two previous rules governing the weak conditional, and rules for universal generalisation:

RCEA. From $\alpha \equiv \alpha'$ infer $(\alpha \Rightarrow \beta) \equiv (\alpha' \Rightarrow \beta)$.

RCK. From $(\beta_1 \wedge \cdots \wedge \beta_n) \supset \beta$ infer $(\alpha \Rightarrow \beta_1 \wedge \cdots \wedge \alpha \Rightarrow \beta_n) \supset \alpha \Rightarrow \beta$ for $n \geqslant 0$.

UG. From $\alpha \supset \beta$ infer $\alpha \supset \forall x\beta$ where x is not free in α.

FUG. From $\alpha \supset (A(x) \supset \beta)$ infer $\alpha \supset \forall_c x\beta$ where x is not free in α.

The axiomatisation clearly subsumes that of $\mathcal{S}$. **Ax** asserts that there is an actual individual. **FUI** is a weakened version of universal instantiation: if $\forall_c x\alpha(x)$ is true, then for arbitrary term t (with proviso) if t is actual then $\alpha(t)$ is true. Similarly we have an analogue for universal generalisation in the rule **FUG**: if it is the case that β is true whenever α is and x is actual, then (given the proviso) we can infer $\alpha \supset \forall_c x\beta$.

The nonmodal fragment of the system with just the quantifier $\forall$ clearly corresponds to classical first-order logic. We obtain the following theorems and derived rules of inference, where x is taken as ranging over the set of variables:

Theorem 4.1. *Let* $\underline{\forall} \in \{\forall, \forall_c\}$.
(1) $\alpha \supset \underline{\forall} x\alpha$ *for* x *not free in* α.
(2) *From* α *infer* $\underline{\forall} x\alpha$.
(3) $\underline{\forall} x(\alpha \supset \beta) \supset (\alpha \supset \underline{\forall} x\beta)$ *for* x *not free in* α.
(4) $\underline{\forall} x(\alpha \supset \beta) \supset (\underline{\forall} x\alpha \supset \underline{\forall} x\beta)$.
(5) *Let* c *be a constant that does not appear in* α *and* x *does not occur in* α *or* $\beta(c)$.
 (a) *From* $\alpha \supset \beta(c)$ *infer* $\alpha \supset \forall x\beta(x)$.
 (b) *From* $\alpha \supset (A(c) \supset \beta(c))$ *infer* $\alpha \supset \forall_c x\beta(x)$.
(6) $\forall x\alpha \supset \forall_c x\alpha$.

Further results of a more technical interest, involving derived rules of inference and nestings of the weak conditional, are found in Section 8 in Definition 8.1 and Theorem 8.4.

The semantic and proof-theoretic characterisations coincide:

Theorem 4.2. *$\mathcal{QS}$ is sound and complete with respect to the class of $\mathcal{QS}$-models.*

We can now introduce the "variable binding" class of operators $\Rightarrow_{\bar{x}}$ by definition:

Definition 4.2. For arbitrary tuple of variables $\vec{x}$: $\alpha \Rightarrow_{\vec{x}} \beta \stackrel{\text{def}}{=} \exists_c \vec{x} \alpha \Rightarrow \forall_c \vec{x}(\alpha \supset \beta)$.

The following results are elementary, but show that this definition captures a straightforward extension of $\Rightarrow$ from the propositional logic. In particular, it shows that we can not make arbitrary sets of defaults satisfiable just by appropriately choosing the extension of A at various worlds.

Theorem 4.3. *For any wff $\alpha \in \mathcal{L}_{FC}$, define the $\mathcal{S}$-transform $\mathcal{S}(\alpha) \in \mathcal{L}_C$ as follows:*
(1) *Replace all occurrences of $\Rightarrow_{\vec{x}}$ by $\Rightarrow$.*
(2) *Delete all quantifiers and terms.*
(3) *Replace each distinct predicate variable by a distinct propositional variable; replace A by $\top$.*
We obtain: if $\vdash_{QS} \alpha$ then $\vdash_S \mathcal{S}(\alpha)$.

The next result gives a partial converse.

Theorem 4.4. *If $\vdash_S \alpha$ and α' results from α by uniformly replacing each propositional variable of α by some wff of $\mathcal{L}_{FC}$, and replacing occurrences of $\Rightarrow$ uniformly by $\Rightarrow_{\vec{x}}$ for some tuple of distinct variables $\vec{x}$, then $\vdash_{QS} \alpha'$.*

For example, since we have

$$\vdash_S (P \Rightarrow Q \wedge P \Rightarrow R) \supset P \wedge R \Rightarrow Q,$$

then

$$\vdash_{QS} (P(x) \Rightarrow_x Q(x) \wedge P(x) \Rightarrow_x R(x)) \supset P(x) \wedge R(x) \Rightarrow_x Q(x).$$

5. Discussion

5.1. Properties of the formalism

The intent in the approach is to use the quantifier $\forall_c$ to define the weak variable-binding conditional $\Rightarrow_{\vec{x}}$. However, to begin with we can compare $\forall_c$ with $\forall$. First, we have the theorem $\forall x\alpha \supset \forall_c x\alpha$. In addition, quantified modal sentences can be expressed with various "levels" of force. "Penguins are necessarily birds", for example, can be expressed by:

$$\forall_c x \Box(\mathit{Penguin}(x) \supset \mathit{Bird}(x)) \quad \text{or} \quad \forall x \Box(\mathit{Penguin}(x) \supset \mathit{Bird}(x)).$$

For the first case, we have the reading "every actual penguin is necessarily a bird". For the second case, we have the reading "every penguin (actual or not) is necessarily a bird". In the first case we could consistently conjoin that it is possible that there is a non-bird penguin (i.e., $\Diamond \exists x(\mathit{Penguin}(x) \wedge \neg \mathit{Bird}(x))$); in the second case this would be inconsistent.

We can consistently state that horses contingently do not have wings but *Pegasus* is a winged horse:

$$\forall_c x(\mathit{Horse}(x) \supset \neg \mathit{Wings}(x)), \quad \mathit{Horse}(\mathit{Pegasus}), \quad \mathit{Wings}(\mathit{Pegasus}).$$

As a logical consequence we obtain $\neg A(Pegasus)$, that is, *Pegasus* is not actual. For "actuality", something similar in kind is found in [30], although McCarthy deals with a (classical) first-order theory with reified predicates and functions. There, for example, one may assert that *Pegasus* is a horse that does not exist, and further that (existing) horses do not have wings but *Pegasus* does. In the nonmodal fragment of $\mathcal{QS}$, if all individuals were actual then the domains of quantification for $\forall$ and $\forall_c$ would of course coincide.

Concerning properties of the conditional logic, we have first that the following sentences are jointly consistent:

$$Bird(x) \Rightarrow_x Fly(x), \quad Bird(Opus), \quad Bird(Opus) \Rightarrow \neg Fly(Opus).$$

That is, birds normally fly, but *Opus* is a bird that normally does not fly. In the naive approach this of course would be inconsistent. Further, given our birds-fly default, we can consistently assert that *Opus* necessarily does not fly:

$$\Box(Bird(Opus) \wedge \neg Fly(Opus)).$$

This is no problem since at the most normal worlds in which there are (actual) birds, *Opus* may not be actual. On the other hand, we can not consistently assert that birds normally fly, but necessarily there is an actual nonflying bird. That is, the following sentences are jointly inconsistent:

$$Bird(x) \Rightarrow_x Fly(x), \quad \Box\exists_c x(Bird(x) \wedge \neg Fly(x)).$$

This is a desirable result, since the first formula states that at the most normal bird worlds, birds fly, whereas the second formula states that at every accessible world (including the most normal bird worlds) there is an actual nonflying bird. Further, given that we accept that normally birds fly, we can not consistently assert that necessarily all birds (actual or not) do not fly nor that birds normally do not fly.

The problems described in Section 2 are satisfactorily resolved. The first example is phrased as follows:

$$\begin{aligned}&Bird(x) \Rightarrow_x Fly(x),\\ &Penguin(x) \Rightarrow_x \neg Fly(x),\\ &\forall x \Box(Penguin(x) \supset Bird(x)).\end{aligned} \tag{16}$$

These sentences are jointly consistent, and have $Bird(x) \Rightarrow_x \neg Penguin(x)$ as a logical consequence. However, given just these sentences, we cannot conclude that $Penguin(Opus) \Rightarrow \neg Fly(Opus)$ is true, as we could in the naive approach. Thus, in contrast with the earlier approach, we do not have the result that at the least worlds where *Penguin*(*Opus*) is true, $\neg Fly(Opus)$ is true. What we *do* have is that at the least worlds in which there are penguins, penguins do not fly; Opus may or may not be among the individuals existing at such worlds. Now, it seems quite reasonable that we might want to conclude that *Opus* is an actual (or: unexceptional) penguin at such worlds, and so does not fly. However, such an assumption of *unexceptionality* is a nonmonotonic assumption, outside the present framework. I return to this theme briefly in Section 5.2, where I outline how such an assumption can be effected.

The second example involves "specificity" arising from the nesting of quantifiers. This example would now be expressed as follows:

$$(El(x) \wedge Ke(y)) \Rightarrow_{xy} Likes(x, y),$$
$$(El(x) \wedge Ke(Fred)) \Rightarrow_x \neg Likes(x, Fred),$$
$$(El(Clyde) \wedge Ke(Fred)) \Rightarrow Likes(Clyde, Fred).$$

Given that $El(Clyde) \wedge Ke(Fred)$ is contingently true at a world w, the above sentences are now simultaneously satisfiable. At the least exceptional $El(x) \wedge Ke(y)$ worlds, *Fred* would not be actual; at these worlds all elephants would indeed like their keepers. Among the least exceptional $\exists_c(El(x) \wedge Ke(Fred))$ worlds we would have that elephants do not like Fred, as desired. If instead we wished to assert that elephants normally like their keepers, but there is some keeper that is not normally liked by elephants, we could represent this as follows:

$$(El(x) \wedge Ke(y)) \Rightarrow_{xy} Likes(x, y), \ \exists y((El(x) \wedge Ke(y)) \Rightarrow_x \neg Likes(x, y)).$$

The third example deals with finite and infinite versions of the lottery paradox. While the lottery paradox seems largely bound with notions of default inference, and the present paper deals with representational issues, nonetheless we can still say something about the lottery paradox in the present framework. In the finite case, knowledge about bird species may be expressed either by (8) or (9). In (8), for a given world w, a bird is contingently one of some number of species. At other worlds such an individual need not be of that species, nor even a bird at all. Arguably such a notion is overly weak. As discussed further below, we well might want to say that if an individual is a bird then it is necessarily a bird. Hence we would adopt (9), and assert that a bird is necessarily one of some number of species. In the present framework this presents no difficulty: we can consistently assert, for example, that, along with the usual bird properties, every species is exceptional in some fashion, and that every individual that is a bird is (necessarily) of one of these species, and, perhaps, is necessarily a bird.

The "standard" (infinite) version of the lottery paradox is different. While $\forall_c x(\top \Rightarrow \neg Winner(x))$ no longer entails $\top \Rightarrow \forall_c x \neg Winner(x)$, both formulations (11) and (12) are still unsatisfiable. However, this is as things should be: what this shows is not that the present approach is inadequate, but rather that the standard lottery paradox is not plausibly represented using notions of "normality" that are not probabilistically based. In stating that "normally one does not win the lottery", the sense of "normal" is that of statistical likelihood. However, in the approach at hand, the notion of "normally" bears only a tenuous relation with statistical notions.

To further argue this last point, take a relatively extreme case: we may have a conditional $\alpha \Rightarrow_x \beta$ where in the domain being modelled $\alpha(t) \wedge \beta(t)$ is true of *no* individual. Is it reasonable to allow this situation? The answer is "yes", as two examples (from [41,42]) illustrate. Consider first the statement "lemons are yellow", $Lemon(x) \Rightarrow_x Yellow(x)$. It is not inconceivable that some new disease may come along, in which *all* lemons are affected. If this disease turns lemons blue, say, then we would still hold $Lemon(x) \Rightarrow_x Yellow(x)$ even though we contingently would have that $\forall x(Lemon(x) \supset \neg Yellow(x))$. In this case, it would make perfect sense to say that lemons are normally yellow even though presently none are, due to the presence of the disease. The second example is from physics (and assumes for sake of argument a Newtonian universe). We know that "planets move in ellipses" is true. The only difficulty is that, in point of fact, no planet ever has, nor ever will, move in an ellipse: the gravitational pull of other planets, moons, stars, etc. ensures that such idealised motion would not occur. Nonetheless, one still holds to the assertion that, in ideal circumstances (i.e., factoring out exceptional conditions), a planet would indeed move in an ellipse.

So essentially the notion of "normally" formalised here corresponds to that found in (naive, commonsense) scientific theories; and a conditional $\alpha \Rightarrow_x \beta$ corresponds to a law in such a theory. This stance further justifies the lack of "full" instantiation for general defaults. Given that one accepts "lemons are normally yellow", it seems clear that "individual l, which is a lemon, is normally yellow" should not be a logical consequence. On the other hand, this last assertion is a reasonable nonmonotonic instantiation. In the next subsection I briefly consider how such an assumption can be effected.

5.2. Minimizing exceptions

Consider an earlier example, that birds normally fly but *Opus* does not:

$$Bird(x) \Rightarrow_x Fly(x), \quad Bird(Opus) \Rightarrow \neg Fly(Opus).$$

A logical consequence is that *Opus* is not actual at the most normal bird worlds:

$$Bird(x) \Rightarrow_x \neg A(Opus) \quad \text{or equally} \quad \exists_c x Bird(x) \Rightarrow \neg A(Opus).$$

Now, the predicate A is simply a predicate with a particular fixed semantic interpretation, and we are free to informally interpret it as we wish. Instead of "exists" or "actual", an alternative reading for A is as "relevant" or "unexceptional". Which is to say, $\alpha(x) \Rightarrow_x \beta(x)$ is true just when in the least $\exists_c x\alpha$ worlds, for every actual individual, if α is true of that individual then so is β. But this is perhaps better read as saying that this is true for every "relevant" or "unexceptional" individual at such worlds. Thus the predicate A can be taken as denoting such relevant (or unexceptional) individuals.

If we are given only the assertion $Bird(x) \Rightarrow_x Fly(x)$ then, lacking information to the contrary, it seems that we should be able to conclude $Bird(Opus) \Rightarrow Fly(Opus)$. That is, lacking information to the contrary, we would want to assume that an individual (here *Opus*) is unexceptional wherever possible. This can be carried out straightforwardly in terms of preference among models. Let $P_1, \ldots, P_n$ be a tuple of predicate symbols, denoted by $\vec{P}$. The intent is to choose those models with maximal D', where the extension of the predicates in $\vec{P}$ is allowed to vary.

Assume that $\mathcal{T}$ is a consistent theory of $\mathcal{L}_{FC}$, and let $M_1 = \langle W, R, D, D'_1, V_1\rangle$ and $M_2 = \langle W, R, D, D'_2, V_2\rangle$ be two models of $\mathcal{T}$.

Write $M_1 \leqslant_{\vec{P}} M_2$ if

(1) $V_1(t) = V_2(t)$ for every term t.

(2) $V_1(Q, w) = V_2(Q, w)$ for predicate symbol Q different from A and not in $\vec{P}$, and for every world w.

(3) $D'_2(w) \subseteq D'_1(w)$ for every world w.

The $\leqslant_{\vec{P}}$-minimal models of $\mathcal{T}$ are those models with maximal D', or maximal sets of actual (or unexceptional) individuals. We can write $\mathcal{T} \models_{\leqslant_{\vec{P}}} \alpha$ just when every $\leqslant_{\vec{P}}$-minimal model of $\mathcal{T}$ is also a model of α. Clearly in our motivating example we now obtain

$$\{Bird(x) \Rightarrow_x Fly(x)\} \models_{\leqslant_{Fly}} Bird(Opus) \Rightarrow Fly(Opus).$$

Essentially we assume that A holds of as many individuals as consistently possible. Maximising the extension of A at a world is the same as minimising $\neg A$; if we read $A(t)$ as "t is unexceptional" then $\neg A(t)$ of course is "t is exceptional". So our minimisation of abnormalities looks very much like variable circumscription [32], but in a modal context. Presumably, the minimization policy would be such that the predicate symbols in $\vec{P}$ are chosen from among the predicates appearing in the consequents of a weak conditional.

Note though that there are differences with variable circumscription. In variable circumscription, there are various "Ab" predicates corresponding to conditions with respect to which an individual may be exceptional. In contrast we minimize just one predicate, A, corresponding to the set of "unexceptional" individuals at a world. In addition, an individual is exceptional or not with respect to each world. The minimization here is intended solely to provide a means of allowing instantiation "by default" for a general weak conditional. This minimization is clearly orthogonal to others, such as that employed in the rational closure of a theory [28], where the minimization is essentially carried out on the rank of worlds. Consequently, other issues concerning nonmonotonic reasoning with weak conditionals, such as dealing with irrelevant properties, are not addressed here. Presumably such issues may be handled as in the propositional case using, for example, schemes analogous to those such as rational closure.

6. Comparison with related work

As a quantificational modal logic, the present approach notably has world-relative domains and rigidly designating terms. The notion of nonexisting terms has been discussed in Artificial Intelligence in for example [30]; as well there are echoes of free logic [49] in the nonmodal fragment of the system. Hirst [21] discusses issues concerning existence and nonexistence of individuals in Artificial Intelligence that pertain to choices made here. With respect to conditional logics and default reasoning, there are three approaches in Artificial Intelligence that may be compared with the present.

Kraus et al. [22] and Lehmann and Magidor [29] investigate (propositional) nonmonotonic consequence operators. An operator $\mid\sim$ is specified such that $\alpha \mid\sim \beta$ has the intended reading that if α is true then β nonmonotonically (or normally or plausibly) follows as a consequence. Relations among these operators are given using Gentzen-style rules. Notably the symbol $\mid\sim$ is not an element in the language, but represents a conditional assertion between formulas. So analogous to the axiom **ASC** is the relation of **Cautious Monotony**:

$$\frac{\alpha \mid\sim \beta, \quad \alpha \mid\sim \gamma}{\alpha \wedge \beta \mid\sim \gamma}.$$

There are analogous rules corresponding to the axioms and rules of conditional logics. The system P of *preferential entailment* corresponds to the system $\mathcal{S}$.[10] In [28] an extension to the first-order case is given; the following rules are proposed:

$$\frac{\alpha \mid\sim \beta}{\exists x \alpha \mid\sim \exists x \beta}, \tag{17}$$

[10] See [8] for connections between conditional logics and nonmonotonic consequence operators.

$$\frac{\exists x\alpha \mathrel{|\!\sim} \beta}{\alpha \mathrel{|\!\sim} \beta} \quad (x \text{ not free in } \beta). \tag{18}$$

One representation of (17) in the present approach is

From $\alpha \Rightarrow \beta$ infer $\exists_c x\alpha \Rightarrow \exists_c x\beta$.

This indeed is a derived rule in $\mathcal{QS}$, but it is not particularly interesting, if only because *theorems* of the form $\alpha \Rightarrow \beta$ are not interesting. The alternative representation:

$$\alpha(x) \Rightarrow \beta(x) \supset \exists_c x\alpha(x) \Rightarrow \exists_c x\beta(x)$$

is not a theorem in the present approach, since the truth values of the two conditionals may be determined with respect to differing worlds (i.e., there is no connection between the least $\alpha(x)$ worlds and the least $\exists_c x\alpha(x)$ worlds). The following formula, which is of similar form and arguably captures the intent of (17), is valid in $\mathcal{QS}$:

$$\alpha \Rightarrow_x \beta \supset \exists_c x\alpha \Rightarrow \exists_c x\beta.$$

Thus, for example, if birds normally fly, then in the least worlds in which there is a bird, some individual flies. As well,

$$\alpha \Rightarrow \beta \supset \exists_c x(\alpha \Rightarrow \beta) \quad \text{for } x \text{ not free in } \alpha, \beta$$

is valid in our approach. This last formula is not expressible using nonmonotonic consequence operators, since $\exists x(\alpha \mathrel{|\!\sim} \beta)$ is not a well-formed syntactic object.

For (18) there are two natural representations as rules of inference, depending on how we take the scope of the quantifier:

From $\exists_c x(\alpha \Rightarrow \beta)$ infer $\alpha \Rightarrow \beta$,

From $(\exists_c x\alpha) \Rightarrow \beta$ infer $\alpha \Rightarrow \beta$ for x not free in α, β.

Both are derived rules in $\mathcal{QS}$; only the second has a meaningful representation using nonmonotonic consequence operators. For the analogous formulas, $\exists_c x(\alpha \Rightarrow \beta) \supset \alpha \Rightarrow \beta$ (for x not free in α, β) is valid, while $(\exists_c x\alpha) \Rightarrow \beta \supset \alpha \Rightarrow \beta$ is not. A representation arguably closer to the intent of Lehmann and Magidor [28], and valid in $\mathcal{QS}$, is the following:

$$(\exists_c x\alpha) \Rightarrow \beta \supset \alpha \Rightarrow_x \beta \quad \text{for } x \text{ not free in } \beta.$$

Finally Lehmann and Magidor [28] reject the following rule:

$$\text{From } \alpha(x) \Rightarrow \beta(x) \text{ infer } \alpha(t) \Rightarrow \beta(t) \tag{19}$$

where t is free for x in $\alpha(x)$. The reasons given are essentially those that we gave for the elephant-keeper example in (4)–(6) in the naive approach: that unrestricted substitution leads to inconsistency in "reasonable" situations. In $\mathcal{QS}$, using **FUI** and **RCK**, we obtain the theorem (assuming t is free for x), expressing an acceptable version of (19):

$$[\alpha(x) \Rightarrow_x \beta(x)] \supset [\exists_c x\alpha(x) \Rightarrow ((A(t) \wedge \alpha(t)) \supset \beta(t))]. \tag{20}$$

Thus, for example, if birds normally fly, then at the least bird worlds, if Opus is actual (i.e., unexceptional) and a bird then Opus flies. One might want to assume that Opus is in fact unexceptional, a nonmonotonic assumption described in the previous section.

$\mathcal{QS}$ allows more distinctions than [28], and is more broadly applicable to a wider range of situations. It is unclear how the conditional $\alpha \Rightarrow_{\bar{x}} \beta$ would be translated into a system of nonmonotonic consequence operators (since in the obvious translation the image of $\Rightarrow_{\bar{x}}$ would bind variables across the $\mid\!\sim$ operator, yielding an ill-formed expression); most likely one would need to use the more cumbersome expression given in Definition 4.2. However $\mathcal{QS}$ is a logic of first-order conditionals, and so does not address notions of nonmonotonicity; the natural next step of course is to investigate nonmonotonic reasoning in this framework. Nonmonotonic reasoning is something that Lehmann and Magidor [28] address, where it is conjectured that a finite first-order knowledge base as given therein has a rational closure.

Schlechta [44] formalises a notion of "normally $\alpha(x)$" using restricted quantification. $\nabla x\alpha(x)$ is true if α holds on an "important" subset of the domain. If D is the domain of individuals, then $\mathcal{N}(D)$ consists of a family of important subsets of D such that the intersection of any two elements of this family is nonempty. The nonempty intersection condition rules out a default and its negation both being true. $\nabla x\alpha(x)$ is true just if there is a member D' of $\mathcal{N}(D)$ for which α is true for every element of D'. Our $\alpha \Rightarrow_{\bar{x}} \beta$ would be expressed $\nabla x\alpha(x) : \beta(x)$, with the interpretation that all x's that satisfy α normally satisfy β. In the resultant logic, relations among quantifiers include

$$\nabla x\alpha(x) \supset \neg\nabla\neg x\alpha(x) \text{ as well as } \forall x\alpha \supset \nabla x\alpha(x) \text{ and } \nabla x\alpha \supset \exists x\alpha(x).$$

The notion of "normally" then corresponds roughly to that of "actual" or "unexceptional" in $\mathcal{QS}$, and Schlechta's ∇ corresponds roughly to our $\forall_c$, although it is in some respects weaker. (For example, $\nabla x(\alpha \supset \beta) \supset (\nabla x\alpha \supset \nabla x\beta)$—our Theorem 4.1(4)—is not valid in Schlechta's approach, although $\forall x(\alpha \supset \beta) \supset (\nabla x\alpha \supset \nabla x\beta)$ is.) Assuming such a correspondence, in $\mathcal{QS}$ there is only one "important" set per world, as opposed to multiple sets in Schlechta's approach; however this set varies from world to world. Possible worlds are not used in Schlechta's approach, and so notions of normalcy, defined in terms of preferences among worlds, are not available. On the other hand, Schlechta is concerned with interpreting defaults in the sense given, say, in [43], where, arguably such notions of preference are not relevant in the semantics.

In [45] this approach is joined with a second "inference-greedy" step wherein as many instances of the defaults are made true as possible. This is accomplished by defining a preference relation on the models of a default theory, where roughly M is preferred to M' if M satisfies all positive instances of the defaults that M' does, and for no default is there a negative instance which holds in M but not in M', and moreover there is some default with a positive instance in M which is not positive in M'. This step then is analogous in intent to that outlined in Section 5.2. However whereas Schlechta compares models (corresponding to our possible worlds) by considering the positive and negative instances of all defaults, I maximize solely on the predicate A, but across all possible worlds, in comparing models. In the present approach I retain a notion of specificity intrinsic in the ordering on possible worlds whereas Schlechta does not. Hence, using the definition given in (1) for default inference, in Example 1 we would obtain only that a penguin flies while this would not be the case in [45].

Lastly, Friedman et al. [13] develop a first-order conditional logic based on *plausibility measures*. A plausibility measure is a function that associates a notion of plausibility (a member in some partially ordered set) with a set of possible worlds. For sets of worlds A and B, $Pl(A) \leqslant Pl(B)$ means that A is no more plausible than B. It is stipulated that:

A1. If $A \subseteq B$ then $Pl(A) \leqslant Pl(B)$.

Two further conditions are required to obtain the properties of system $\mathcal{S}$. From this, a first-order logic is defined. As with $\mathcal{QS}$, the approach is intended to deal with representational issues, and does not address issues of nonmonotonic inference. However there are notable differences between the approaches.

First, Friedman et al. [13] are generally guided by intuitions deriving from statistics and direct inference. So "birds typically fly" is foremost for them a statistical statement. As a consequence, a specific concern is to be able to adequately address the lottery paradox. In contrast, in $\mathcal{QS}$ the central concern is to address first-order issues in conditional logics in general. I have used examples concerning defaults throughout this paper, but need not have done so; rather, the approach is *generally* applicable to subjunctive conditionals, and so applies to counterfactual statements, hypotheticals, etc. Note too that the lemons-are-yellow and planets-move-in-ellipses examples of Section 5.1 illustrate that, from a representational point of view, a statistical approach is not appropriate for these notions of normality. Thus the approaches are intended for different applications.

Insofar as technical differences are concerned, Friedman et al. [13] make primary use of nonrigid designators, so a term can denote different individuals in different worlds. But (see Section 2) this seems curious, in that if we are talking about Opus' (likely) flying ability, and this notion depends on other possible worlds, then it seems that the term *Opus* should pick out the same individual across possible worlds. [11] So for our elephant/keeper example in (4)–(6), in [13] consistency can be maintained only by having *Clyde* and *Fred* designate different individuals at different worlds.

This has further ramifications, in that where we have **FUI**, Friedman et al. [13] have:

F1. $\forall x \alpha(x) \supset \alpha(t)$

where t is a term free for x in $\alpha(x)$, and if α is a formula that has occurrences of $\Rightarrow$ then the only terms substitutable for x in α are other variables.

In contrast we have limited substitution; for example, in (20), we have substitution into a modal context. We also have as an instance of **FUI** that for constant c:

$$\forall_C x(\alpha(x) \Rightarrow \beta(x)) \supset (A(c) \supset (\alpha(c) \Rightarrow \beta(c)))$$

which provides an example of, in contradistinction to [13], substitution of a constant into a modal context.

7. Conclusion

This paper has addressed issues in first-order conditional logics, and has developed a general framework within which to address quantificational concerns in conditional

[11] Rigid terms may in fact be incorporated in [13]; the point is that the truth of weak conditionals in their approach is based on nonrigid terms.

logics. The approach applies generally to conditional logics intended to represent defaults, counterfactuals, and other subjunctive conditionals. Interestingly, the classic treatment of counterfactuals, and arguably conditional logics in general [27], begins with the (counterfactual) example "if kangaroos had no tails, they would topple over". This statement is clearly and easily representable in the first-order analogue to Lewis' "official" counterfactual logic **VC** by means of the present approach. And equally clearly we could allow for a particularly coordinated individual *Hoppy* where it is true that "if *Hoppy* had no tail then he would not topple over".

The approach is most notably characterised by its adopting world-relative domains. As well, default (and other subjunctive) assertions are expressed by means of a variable-binding conditional. This "variable-binding" conditional is in turn defined in terms of the "standard" weak conditional $\Rightarrow$, as well as with a "weak" quantifier $\forall_c$, and a designated predicate A that picks out actual or unexceptional individuals at each world. The approach does not incorporate all of the formal machinery that it might in order to highlight relevant issues and to make the presentation clear. The introduction of functions, equality, and nonrigid terms would present little difficulty.

The focus in this paper is on addressing representational issues; clearly the major topic for future work is to address default inference in a first-order setting. One obvious strategy would be to extend the present approach by effectively welding on the machinery of some extant approach. The difficulty with such a strategy is that no current approach is as yet unproblematic. That being said, probably the easiest approach to incorporate with the present framework would be (an analogue of) rational closure [29]. The central intuition in rational closure is that one assumes that worlds (or their equivalents) are as normal as possible; so one attempts to minimise worlds in a ranking, based on an initial set of defaults. Very roughly, one can conclude that a green bird flies because there is nothing blocking the assumption that green-bird worlds are among the most normal bird worlds.

In the present framework there is also a second minimisation, described in Section 5.2, where we assume that individuals at a worlds are unexceptional, whenever possible. This allows the conclusion that Opus normally flies, given that birds fly along with the fact that we have no reason to believe that Opus is not exceptional. This is effected via an assumption that A holds of as many individuals as consistently possible.

8. Proofs of theorems

$\vdash$ will mean $\vdash_{\mathcal{QS}}$ throughout Section 8.

Proof of Theorem 4.1. I just give proofs for $\forall_c$; $\forall$ follows in each case analogously.

(1) $\vdash \alpha \supset (A(x) \supset \alpha)$ is an instance of a propositional tautology. We can choose variable x so that it does not occur in α. **FUG** gives $\vdash \alpha \supset \forall_c x\alpha$.

(2) Given $\vdash \alpha$, from propositional logic we obtain $\vdash A(x) \supset \alpha$. Applying **FUG** gives $\vdash \forall_c x\alpha$.

(3) An instance of **FUI** is $\vdash \forall_c x(\alpha \supset \beta) \supset (A(x) \supset (\alpha \supset \beta(x)))$. By propositional logic, $\vdash (\alpha \wedge \forall_c x(\alpha \supset \beta)) \supset (A(x) \supset \beta(x))$ and by **FUG** $\vdash (\alpha \wedge \forall_c x(\alpha \supset \beta)) \supset$

$\forall_c x\beta(x)$, where x is not free in α. Rearranging terms we obtain $\vdash \forall_c x(\alpha \supset \beta) \supset (\alpha \supset \forall_c x\beta)$.

(4) Two combined instances of **FUI** give $\vdash (\forall_c x(\alpha \supset \beta) \wedge \forall_c x\alpha) \supset (A(x) \supset ((\alpha(x) \supset \beta(x)) \wedge \alpha(x)))$ from which by propositional logic we obtain $\vdash (\forall_c x(\alpha \supset \beta) \wedge \forall_c x\alpha) \supset (A(x) \supset \beta(x))$. An application of **FUG** and a rearranging of terms yields the result.

(5) Let $\psi_0, \ldots, \psi_n$ be a proof of $\alpha \supset (A(c) \supset \beta(c))$ where variable x does not appear in the proof. (This can always be obtained by uniformly renaming variables.) For each ψ_i, $0 \leqslant i \leqslant n$, we obtain the formula ψ_i' by uniformly substituting variable x for constant c. It is straightforward to argue that $\psi_0', \ldots, \psi_n'$ is a proof of $\alpha \supset (A(x) \supset \beta(x))$: if ψ_i is an axiom then so is ψ_i'. If ψ_i results from ψ_j and $\psi_j \supset \psi_i$ by modus ponens then ψ_i' results from ψ_j' and $\psi_j' \supset \psi_i'$ by modus ponens. Similar remarks apply to the other rules of inference. Since x does not appear in α, an application of **FUG** to ψ_n' for variable x yields $\alpha \supset \forall_c x\beta(x)$.

(6) $\vdash \forall x\alpha \supset \alpha(x)$ via **UI**. Hence $\vdash \forall x\alpha \supset (A(x) \supset \alpha(x))$ by propositional logic; hence $\vdash \forall x\alpha \supset \forall_c x\alpha$ by **FUG**. □

Proof of Theorem 4.3. The proof follows by induction on a proof of α in $\mathcal{QS}$. The image of every axiom of $\mathcal{QS}$ is valid in $\mathcal{S}$, and the image of the rules of inference are easily seen to preserve validity. We obtain by induction that $\mathcal{S}(\alpha)$ is valid in $\mathcal{S}$, hence $\vdash_{\mathcal{S}} \mathcal{S}(\alpha)$. □

Proof of Theorem 4.4. The proof follows by a straightforward induction on a proof in the propositional logic: the image of propositional axioms under the replacement are theorems, and the rules **RCEA** and **RCEC** and modus ponens preserve theoremhood. □

Theorem 8.1 (Bound alphabetic variants). *Call $\alpha(x)$ and $\alpha(y)$* similar *if x is free in $\alpha(x)$ precisely where y is free in $\alpha(y)$. If $\alpha(x)$ and $\alpha(y)$ are similar then*

(1) $\vdash \forall x\alpha(x) \equiv \forall y\alpha(y)$;

(2) $\vdash \forall_c x\alpha(x) \equiv \forall_c y\alpha(y)$.

Proof. Both results are straightforward; for the second, by **FUI**, $\vdash \forall_c x\alpha(x) \supset (A(y) \supset \alpha(y))$ since y is free for x in $\alpha(x)$. Rearranging terms we obtain $\vdash A(y) \supset (\forall_c x\alpha(x) \supset \alpha(y))$ and from **FUG** we get $\vdash \forall_c y(\forall_c x\alpha(x) \supset \alpha(y))$. By Theorem 4.1(3) we get $\vdash \forall_c x\alpha(x) \supset \forall_c y\alpha(y)$. The same argument gives $\vdash \forall_c y\alpha(y) \supset \forall_c x\alpha(x)$, from which we obtain the equivalence. □

Theorem 8.2 (Soundness of $\mathcal{QS}$). *If $\vdash \alpha$ then $\models \alpha$.*

Proof. Soundness is proved in the general case by an inductive argument on the length of a proof. The only interesting cases are **CUG**, **FUI** and **FUG**.

(**CUG**) Assume that in $M = \langle W, R, D, D', V\rangle$ we have $M, w \models \forall x(\alpha \Rightarrow \beta)$ where x is not free in α. Now $M, w \models \forall x(\alpha \Rightarrow \beta)$ iff for every V' like V, except possibly $V(x) \neq V'(x)$, and where $M' = \langle W, R, D, D', V'\rangle$, we have $M', w \models \alpha \Rightarrow \beta$. By definition $M', w \models \alpha \Rightarrow \beta$ iff for every $w_1 \in min(w, \|\alpha\|^{M'})$ we have $w_1 \in \|\beta\|^{M'}$. Since x is not free in α, α must have the same value at w_1 in M' as it does in M. So for every V' the same as V, except possibly $V(x) \neq V'(x)$, we have $w_1 \in \|\beta\|^{M'}$, or by the definition of V for $\forall$, we have $w_1 \in \|\forall x\beta\|^M$. That is, for every $w_1 \in min(w, \|\alpha\|^M)$ we have $w_1 \in \|\forall x\beta\|^M$. So $M, w \models \alpha \Rightarrow \forall x\beta$.

(FUI) Assume that in $M = \langle W, R, D, D', V\rangle$ we have $M, w \models \forall_c x\alpha(x)$. So for every V' the same as V except that $V'(x) \in D'(w)$ where $M' = \langle W, R, D, D', V'\rangle$ we have $M', w \models \alpha(x)$. Assume that V assigns to t some member of $D'(w)$ at w. Then among the various V' there must be one that assigns to x the same member of $D'(w)$ as V assigns to t at w. So for this V' we have $V'(x) = V(t)$. Since, $V(t) \in D'(w)$ we have $M, w \models A(t)$. Since $M', w \models \alpha(x)$ and $V'(x) = V(t)$, so $M, w \models \alpha(t)$, whence $M, w \models A(t) \supset \alpha(t)$. Alternatively, assume that V assigns to t some member of D not in $D'(w)$. Since $V(t) \notin D'(w)$ then $M, w \not\models A(t)$, so $M, w \models \neg A(t)$, so $M, w \models A(t) \supset \alpha(t)$.

(FUG) Assume that $M, w \models \alpha$ for model M and world w (otherwise $M, w \models \neg\alpha$, and so $M, w \models \alpha \supset \forall_c x\beta$). Since $\alpha \supset (A(x) \supset \beta)$ is valid by assumption, so $M, w \models A(x) \supset \beta$. For any model $M' = \langle W, R, D, D', V'\rangle$ which is the same as M, except V' assigns x to a member of $D'(w)$, we have $M', w \models A(x) \supset \beta$ (since x does not occur in α and so $M', w \models \alpha$, hence the result by modus ponens). So for every such V' we have $M', w \models A(x) \supset \beta$. Since V' assigns x to a member of $D'(w)$, so $M', w \models A(x)$ and so (via modus ponens) we obtain $M', w \models \beta$. But by the definition of $\models$ this means $M, w \models \forall_c x\beta$. So for any model and any world in which $M, w \models \alpha$ where x is not free in α we have $M, w \models \forall_c x\beta$, whence $M, w \models \alpha \supset \forall_c x\beta$ for every model and world, assuming x is not free in α. □

The following results are straightforward but tedious. Where no confusion arises, I have on occasion combined steps in a proof.

Theorem 8.3.

(1) **CSO**: $((\alpha \Rightarrow \beta) \wedge (\beta \Rightarrow \alpha)) \supset ((\alpha \Rightarrow \gamma) \equiv (\beta \Rightarrow \gamma))$.

(2) $((\alpha \Rightarrow \gamma) \wedge (\beta \Rightarrow \delta)) \supset (\alpha \vee \beta) \Rightarrow (\gamma \vee \delta)$.

(3) $((\alpha \vee \beta \Rightarrow \alpha) \wedge (\beta \vee \gamma \Rightarrow \beta)) \supset \alpha \vee \gamma \Rightarrow \alpha$.

Proof.

(1) [9, B3].

(2)
(1) $\alpha \Rightarrow \gamma \supset \alpha \Rightarrow \gamma \vee \delta$ — Theorem, from RCK;
(2) $\beta \Rightarrow \delta \supset \beta \Rightarrow \gamma \vee \delta$ — Theorem, from RCK;
(3) $(\alpha \Rightarrow \gamma \wedge \beta \Rightarrow \delta) \supset \alpha \vee \beta \Rightarrow \gamma \vee \delta$ — (1), (2), CA, PC.

(3)
(1) $((\alpha' \vee \beta' \Rightarrow \alpha') \wedge (\alpha' \Rightarrow \alpha' \vee \beta')) \supset ((\alpha' \Rightarrow \gamma') \equiv (\alpha' \vee \beta' \Rightarrow \gamma'))$ — CSO;
(2) $\alpha' \vee \beta' \Rightarrow \alpha' \supset ((\alpha' \Rightarrow \gamma') \supset (\alpha' \vee \beta' \Rightarrow \gamma'))$ — (1), Theorem: $\alpha \Rightarrow \alpha \vee \beta$, PC;
(3) $((\alpha \vee \beta \vee \gamma \Rightarrow \alpha \vee \beta) \wedge (\alpha \vee \beta \Rightarrow \alpha)) \supset (\alpha \vee \beta \vee \gamma \Rightarrow \alpha)$ — Instance of 2;
(4) $((\alpha \vee \beta \Rightarrow \alpha) \wedge (\beta \vee \gamma \Rightarrow \beta)) \supset \alpha \vee \beta \vee \gamma \Rightarrow \alpha \vee \beta$ — Theorem 8.3.2;
(5) $((\alpha \vee \beta \Rightarrow \alpha) \wedge (\beta \vee \gamma \Rightarrow \beta)) \supset$
$(\alpha \vee \beta \vee \gamma \Rightarrow \alpha)$ — (4), (3), Transitivity of $\supset$, PC;
(6) $((\alpha \vee \beta \vee \gamma \Rightarrow \alpha \vee \gamma) \wedge (\alpha \vee \beta \vee \gamma \Rightarrow \alpha)) \supset \alpha \vee \gamma \Rightarrow \alpha$ — ASC;
(7) $(\alpha \vee \beta \vee \gamma \Rightarrow \alpha) \supset (\alpha \vee \beta \vee \gamma \Rightarrow \alpha \vee \gamma)$ — Theorem, from RCK;
(8) $(\alpha \vee \beta \vee \gamma \Rightarrow \alpha) \supset \alpha \vee \gamma \Rightarrow \alpha$ — (6), (7), PC;
(9) $((\alpha \vee \beta \Rightarrow \alpha) \wedge (\beta \vee \gamma \Rightarrow \beta)) \supset \alpha \vee \gamma \Rightarrow \alpha$ — (5), (8), PC. □

The completeness proof is a bit awkward since we have two quantifiers to deal with. $\forall$ is handled basically as in [11,20]; $\forall_c$ is handled in an extension of the technique given therein. In particular, where [20] make use of the *$\forall$-property*, I define a related notion called here the *$\forall_c$-property*. For the proof we need to show (see Lemma 8.3) that a set of formulas satisfying certain conditions can be extended to a maximal consistent set satisfying both the $\forall$-property and the $\forall_c$-property.

Definition 8.1. An $\mathcal{I}$-form is a formula having the following shape:

(1) A formula of the nonmodal first-order logic is an $\mathcal{I}$-form.
(2) If β is an $\mathcal{I}$-form according to steps (1) and (2) only then $\alpha \Rightarrow \beta$ is an $\mathcal{I}$-form.
(3) If β is an $\mathcal{I}$-form then $\alpha \supset \beta$ is an $\mathcal{I}$-form.

So an $\mathcal{I}$-form has shape $\alpha_0 \supset (\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \beta)\ldots)$ where $n \geqslant 0$. Since α_0 may be a theorem this means that informally any of the $\supset$, $\Rightarrow$ may be missing. An $\mathcal{I}$-form where β is the "innermost" formula, as given in step (1) of Definition 8.1, will be denoted $\mathcal{I}(\beta)$. For a given $\mathcal{I}$-form $\mathcal{I}(\beta)$, $\mathcal{I}(\gamma)$ will denote the $\mathcal{I}$-form obtained by replacing β by γ.

The following theorem generalizes **RCK**, **CUG**, **FUI** and **FUG**.

Theorem 8.4. *Let $\mathcal{I}$-form $\mathcal{I}(\beta)$ be $\alpha_0 \supset (\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \beta)\ldots)$.*

(1) *If $\vdash \beta \supset \gamma$ then $\vdash \mathcal{I}(\beta) \supset \mathcal{I}(\gamma)$.*
(2) *$\vdash \forall x \mathcal{I}(\beta) \supset \mathcal{I}(\forall x \beta)$ for x not free in α_i, $0 \leqslant i \leqslant n$.*
(3) *$\vdash \mathcal{I}(\forall_c x \beta(x)) \supset \mathcal{I}(A(t) \supset \beta(t))$ for t free for x in β.*
(4) *From $\vdash \mathcal{I}(A(x) \supset \beta)$ infer $\vdash \mathcal{I}(\forall_c x \beta)$ for x not free in α_i, $0 \leqslant i \leqslant n$.*

Proof.

(1) If $\vdash \beta \supset \gamma$ then **RCK** yields $\vdash (\alpha_n \Rightarrow \beta) \supset (\alpha_n \Rightarrow \gamma)$. Repeated application of **RCK** gives $\vdash (\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \beta)\ldots) \supset (\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \gamma)\ldots)$. By classical logic we get $\vdash (\alpha_0 \supset (\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \beta)\ldots)) \supset (\alpha_0 \supset (\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \gamma)\ldots))$.
(2) Assume that x is not free in the various α_i's. Proof is by induction on the depth of nesting of $\Rightarrow$ in the construction of an $\mathcal{I}$-form with **CUG** providing the base case. For the general case, an instance of **CUG** gives $\forall x(\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \beta)\ldots) \supset \alpha_1 \Rightarrow (\forall x \alpha_2 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \beta)\ldots)$. By the induction hypothesis, $\vdash (\forall x \alpha_2 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \beta)\ldots) \supset (\alpha_2 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \forall x \beta)\ldots)$. An application of **RCK** yields $\vdash (\alpha_1 \Rightarrow (\forall x \alpha_2 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \beta)\cdots)) \supset (\alpha_1 \Rightarrow (\alpha_2 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \forall x \beta)\ldots))$ from which we conclude $(\forall x(\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \beta)\ldots)) \supset (\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \forall x \beta)\ldots))$ and so by classical logic $(\forall x \alpha_0 \supset (\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \beta)\ldots)) \supset (\alpha_0 \supset (\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow \forall x \beta)\ldots))$.
(3) **FUI** is $\vdash \forall_c x \beta(x) \supset (A(t) \supset \beta(t))$ for t free for x in β. The result follows from Theorem 8.4.1.
(4) From $\vdash \mathcal{I}(A(x) \supset \beta)$ we infer $\vdash \forall x \mathcal{I}(A(x) \supset \beta)$ via Theorem 4.1(2) and $\vdash \mathcal{I}(\forall x(A(x) \supset \beta))$ via Theorem 8.4(2) (assuming that x is not free in the various α_i's). From Theorem 4.1(6) we have $\vdash \forall x \beta \supset \forall_c x \beta$ from which we obtain $\vdash \forall x(A(x) \supset \beta) \supset \forall_c x \beta$. Theorem 8.4(1) gives $\vdash \mathcal{I}(\forall x(A(x) \supset \beta)) \supset \mathcal{I}(\forall_c x \beta)$ and since we have $\vdash \mathcal{I}(\forall x(A(x) \supset \beta))$ we obtain $\vdash \mathcal{I}(\forall_c x \beta)$ $\square$

Definition 8.2.

(1) A set of wffs Λ has the $\forall$*-property* iff for every formula α and variable x,

$$\text{if } \Lambda \vdash \alpha(t) \text{ for every term } t \text{ then } \Lambda \vdash \forall x \alpha(x).$$

(2) A set of wffs Λ has the $\forall_c$*-property* iff for every formula α and variable x,

$$\text{if } \Lambda \vdash \mathcal{I}(A(t) \supset \alpha(t)) \text{ for every term } t \text{ then } \Lambda \vdash \mathcal{I}(\forall_c x \alpha(x)).$$

(3) A set of wffs with both the $\forall$-property and the $\forall_c$-property will be said to have the $\forall,\forall_c$*-property*.

Lemma 8.1.

(1) *$\emptyset$ has the $\forall$-property.*
(2) *$\emptyset$ has the $\forall_c$-property.*

Proof. I just prove the second part; the first part is a simplification of this case. Assume that $\vdash \mathcal{I}(A(t) \supset \beta(t))$ for every term t. If $\mathcal{I}(A(t) \supset \beta(t))$ is $\alpha_0 \supset (\alpha_1 \Rightarrow \cdots \Rightarrow (\alpha_n \Rightarrow (A(t) \supset \beta(t)))\ldots)$ then we have $\vdash \mathcal{I}(A(t) \supset \beta(t))$ for t a variable not appearing in α_i, $0 \leqslant i \leqslant n$. From Theorem 8.4(4) we obtain $\vdash \mathcal{I}(\forall_c x \beta(x))$. That the general result obtains for every variable follows by Theorem 8.1 and Theorem 8.4(1). □

Lemma 8.2.

(1) *If Λ has the $\forall$-property then $\Lambda \cup \{\beta\}$ has the $\forall$-property.*
(2) *If Λ has the $\forall_c$-property then $\Lambda \cup \{\beta\}$ has the $\forall_c$-property.*

Proof.

(1) Assume that Λ has the $\forall$-property and that $\Lambda \cup \{\beta\} \vdash \alpha(t)$ for every term t. So $\Lambda \vdash \beta \supset \alpha(t)$. Since Λ has the $\forall$-property we obtain $\Lambda \vdash \forall x(\beta \supset \alpha(x))$ for any variable x. We can choose x so that it does not occur in β; hence (via Theorem 4.1(3)) $\Lambda \vdash \beta \supset \forall x \alpha(x)$, so $\Lambda \cup \{\beta\} \vdash \forall x \alpha(x)$ and $\Lambda \cup \{\beta\}$ has the $\forall$-property.
(2) Assume that Λ has the $\forall_c$-property and that $\Lambda \cup \{\beta\} \vdash \mathcal{I}(A(t) \supset \alpha(t))$ for every term t. So $\Lambda \vdash \beta \supset \mathcal{I}(A(t) \supset \alpha(t))$. So by propositional logic $\beta \supset \mathcal{I}(A(t) \supset \alpha(t))$ is an $\mathcal{I}$-form and since Λ has the $\forall_c$-property we obtain $\Lambda \vdash \beta \supset \mathcal{I}(\forall_c x \alpha(x))$. Consequently $\Lambda \cup \{\beta\} \vdash \mathcal{I}(\forall_c x \alpha(x))$ and $\Lambda \cup \{\beta\}$ has the $\forall_c$-property. □

Lemma 8.3. *Assume that Λ has the $\forall,\forall_c$-property and β is a wff such that $\Lambda \cup \{\beta\}$ is consistent. Then there exists a maximal consistent set Γ of wffs such that $\Lambda \cup \{\beta\} \subseteq \Gamma$ and Γ has the $\forall,\forall_c$-property.*

Proof. Suppose that the wffs of $\mathcal{L}_{FC}$ are arranged in some determinate order $\delta_1, \delta_2, \ldots,$ as are the terms. Define a sequence of wffs $\gamma_0, \gamma_1, \ldots,$ by:

(1) γ_0 is $\Lambda \cup \{\beta\}$.
(2) (a) If δ_i is $\neg\forall x \alpha(x)$ and if $\gamma_i \cup \{\neg\forall x \alpha(x)\}$ is consistent then

$$\gamma_{i+1} = \gamma_i \cup \{\neg\forall x \alpha(x)\} \cup \{\neg\alpha(t)\}$$

for some term t that leaves γ_{i+1} consistent.

Such a term is guaranteed to exist: we have that $\gamma_i \cup \{\neg\forall x\alpha(x)\}$ is consistent. If $\gamma_i \cup \{\neg\forall x\alpha(x)\} \cup \{\neg\alpha(t)\}$ were inconsistent for every term t then we would have $\gamma_i \cup \{\neg\forall x\alpha(x)\} \vdash \alpha(t)$ for every t. But $\gamma_i \cup \{\neg\forall x\alpha(x)\}$ is of the form $\Lambda \cup \Psi$ for finite Ψ. Since Λ has the $\forall$-property, so by Lemma 8.2(1) and propositional logic does $\Lambda \cup \Psi$, or $\gamma_i \cup \{\neg\forall x\alpha(x)\}$. But this means that if $\gamma_i \cup \{\neg\forall x\alpha(x)\} \vdash \alpha(t)$ for every t then $\gamma_i \cup \{\neg\forall x\alpha(x)\} \vdash \forall x\alpha(x)$; so $\gamma_i \cup \{\neg\forall x\alpha(x)\}$ is inconsistent, contradiction. So there is always a term t that leaves γ_{i+1} consistent.

(b) Otherwise if δ_i is $\neg\mathcal{I}(\forall_c x\alpha(x))$ and if $\gamma_i \cup \{\neg\mathcal{I}(\forall_c x\alpha(x))\}$ is consistent then

$$\gamma_{i+1} = \gamma_i \cup \{\neg\mathcal{I}(\forall_c x\alpha(x))\} \cup \{\neg\mathcal{I}(A(t) \supset \alpha(t))\}$$

for some term t that leaves γ_{i+1} consistent.
Such a term is guaranteed to exist: we have that $\gamma_i \cup \{\neg\mathcal{I}(\forall_c x\alpha(x))\}$ is consistent. If $\gamma_i \cup \{\neg\mathcal{I}(\forall_c x\alpha(x))\} \cup \{\neg\mathcal{I}(A(t) \supset \alpha(t))\}$ were inconsistent for every term t then we would have $\gamma_i \cup \{\neg\mathcal{I}(\forall_c x\alpha(x))\} \vdash \mathcal{I}(A(t) \supset \alpha(t))$ for every t. But $\gamma_i \cup \{\neg\mathcal{I}(\forall_c x\alpha(x))\}$ is of the form $\Lambda \cup \Psi$ for finite Ψ. Since Λ has the $\forall_c$-property, so by Lemma 8.2(2) does $\Lambda \cup \Psi$, or $\gamma_i \cup \{\neg\mathcal{I}(\forall_c x\alpha(x))\}$. But this means that if $\gamma_i \cup \{\neg\mathcal{I}(\forall_c x\alpha(x))\} \vdash \mathcal{I}(A(t) \supset \alpha(t))$ for every t then $\gamma_i \cup \{\neg\mathcal{I}(\forall_c x\alpha(x))\} \vdash \mathcal{I}(\forall_c x\alpha(x))$; so $\gamma_i \cup \{\neg\mathcal{I}(\forall_c x\alpha(x))\}$ is inconsistent, contradiction. So there is always a term t that leaves γ_{i+1} consistent.

(c) Otherwise if δ_i is not of the form $\neg\forall x\alpha(x)$ or $\neg\mathcal{I}(\forall_c x\alpha(x))$ and $\gamma_i \cup \{\delta_i\}$ is consistent, then $\gamma_{i+1} = \gamma_i \cup \{\delta_i\}$.

(d) Otherwise $\gamma_{i+1} = \gamma_i$.

So our sequence $\gamma_0, \gamma_1, \ldots$ is well-defined.

Let $\Gamma = \bigcup_{i=0}^{\infty}(\Lambda \cup \gamma_i)$. Γ is consistent (otherwise there is some k such that $\Lambda \cup \gamma_k$ is inconsistent, contradiction), maximal, and (via an induction using Lemma 8.2) has the $\forall,\forall_c$-property. □

Definition 8.3. For $\Gamma \subseteq \mathcal{L}_{FC}$, $\Gamma^-(\alpha) = \{\gamma \mid \alpha \Rightarrow \gamma \in \Gamma\}$.

Lemma 8.4. *If Γ is a consistent set of formulas containing $\neg(\alpha \Rightarrow \beta)$ then $\Gamma^-(\alpha) \cup \{\neg\beta\}$ is consistent.*

Proof. Assume $\Gamma^-(\alpha) \cup \{\neg\beta\}$ is inconsistent. So there exists a finite subset $\{\gamma_1, \ldots, \gamma_n\}$ of $\Gamma^-(\alpha)$ where $\{\gamma_1, \ldots, \gamma_n\} \vdash \beta$ or $\vdash \left(\bigwedge_{i=1}^{n} \gamma_i\right) \supset \beta$. From **RCK** we obtain $\vdash (\alpha \Rightarrow \gamma_1) \wedge \cdots \wedge (\alpha \Rightarrow \gamma_n) \supset \alpha \Rightarrow \beta$. Now $\{\alpha \Rightarrow \gamma_1, \ldots, \alpha \Rightarrow \gamma_n\} \subseteq \Gamma$ hence $\Gamma \vdash \alpha \Rightarrow \beta$. But $\neg(\alpha \Rightarrow \beta) \in \Gamma$ so Γ is inconsistent, contradiction. □

Lemma 8.5.

(1) *If Γ is a maximal consistent set of formulas with the $\forall$-property then $\Gamma^-(\alpha)$ has the $\forall$-property.*

(2) *If Γ is a maximal consistent set of formulas with the $\forall_c$-property then $\Gamma^-(\alpha)$ has the $\forall_c$-property.*

Proof.

(1) Assume that $\Gamma^-(\alpha) \vdash \gamma(t)$ for every term t. Then, for every term t there are $\gamma_1, \ldots, \gamma_n \in \Gamma^-(\alpha)$ such that $\vdash (\gamma_1 \wedge \cdots \wedge \gamma_n) \supset \gamma(t)$. So $\vdash (\alpha \Rightarrow \gamma_1 \wedge \cdots \wedge \alpha \Rightarrow$

$\gamma_n) \supset \alpha \Rightarrow \gamma(t)$ by **RCK**. But $\alpha \Rightarrow \gamma_1, \ldots, \alpha \Rightarrow \gamma_n \in \Gamma$, since each $\gamma_i \in \Gamma^-(\alpha)$. So, $\alpha \Rightarrow \gamma(t) \in \Gamma$, or $\Gamma \vdash \alpha \Rightarrow \gamma(t)$ by Lindenbaum's Lemma. Since this holds for every t, and Γ has the $\forall$-property, $\Gamma \vdash \forall x(\alpha \Rightarrow \gamma)$ for a choice of x not free in α. By **CUG**, $\Gamma \vdash \alpha \Rightarrow \forall x\gamma$. So $\forall x\gamma \in \Gamma^-(\alpha)$ and so $\Gamma^-(\alpha) \vdash \forall x\gamma$, and so (appealing to Theorems 8.1 and 8.4(1)) $\Gamma^-(\alpha) \vdash \forall y\gamma$ for any variable, and so $\Gamma^-(\alpha)$ has the $\forall$-property.

(2) Assume that $\Gamma^-(\alpha) \vdash \mathcal{I}(A(t) \supset \gamma(t))$ for every term t. So for each term t there are $\{\alpha \Rightarrow \gamma_1, \ldots, \alpha \Rightarrow \gamma_n\} \subseteq \Gamma$ such that $\{\gamma_1, \ldots, \gamma_n\} \vdash \mathcal{I}(A(t) \supset \gamma(t))$. So $\vdash \bigwedge_{i=1}^n \gamma_i \supset \mathcal{I}(A(t) \supset \gamma(t))$. From **RCK**, $\vdash (\alpha \Rightarrow \gamma_1 \wedge \cdots \wedge \alpha \Rightarrow \gamma_n) \supset \alpha \Rightarrow \mathcal{I}(A(t) \supset \gamma(t))$. Since $\{\alpha \Rightarrow \gamma_1, \ldots, \alpha \Rightarrow \gamma_n\} \subseteq \Gamma$, we obtain $\Gamma \vdash \alpha \Rightarrow \mathcal{I}(A(t) \supset \gamma(t))$ for every term t. Since Γ has the $\forall_c$-property, so $\Gamma \vdash \alpha \Rightarrow \mathcal{I}(\forall_c x\gamma(x))$. Since Γ is maximal consistent, $\alpha \Rightarrow \mathcal{I}(\forall_c x\gamma(x)) \in \Gamma$. Thus $\mathcal{I}(\forall_c x\gamma(x)) \in \Gamma^-(\alpha)$ and so $\Gamma^-(\alpha) \vdash \mathcal{I}(\forall_c x\gamma(x))$ and $\Gamma^-(\alpha)$ has the $\forall_c$-property. □

Theorem 8.5. *Let Γ be a maximal consistent set of formulas with the $\forall$,$\forall_c$-property that contains a formula $\neg(\alpha \Rightarrow \beta)$. Then there is a maximal consistent set of formulas that includes $\Gamma^-(\alpha) \cup \{\neg\beta\}$ and has the $\forall$,$\forall_c$-property.*

Proof. By Lemma 8.4, $\Gamma^-(\alpha) \cup \{\neg\beta\}$ is consistent. By Lemma 8.5, $\Gamma^-(\alpha)$ has the $\forall$,$\forall_c$-property. Lemma 8.3 gives the desired result. □

For the following definitions let $w, w_1, \ldots$ be maximal consistent sets of formulas, and let W be the set of maximal consistent sets of formulas. Anticipating Theorem 8.7, I refer to $w, w_1, \ldots$ as *worlds*. The next definition gives a proof-theoretic analogue to $\|\cdot\|^M$.

Definition 8.4. $|\alpha|_w = \{w \in W \mid \alpha \in w\}$.

The following definition (again, defined in terms of maximal consistent sets of sentences) is designed to coincide with the similarly-named function, defined on sets of possible worlds. Use of the term "*min*" anticipates Theorem 8.8.

Definition 8.5. $min(w, |\alpha|_w) = \{w_1 \in W \mid \{\beta \mid \alpha \Rightarrow \beta \in w\} \subseteq w_1\}$.

Definition 8.6. The *canonical model* is given by the following. $M = \langle W, R, D, D', V\rangle$ where:

(1) $W = \{w \mid w$ is a maximal consistent set of formulas with the $\forall$,$\forall_c$-property$\}$.
(2) For $w, w_1, w_2 \in W$, $R_w w_1 w_2$ iff
   (a) there are γ, δ such that $w_1 \in min(w, |\gamma|_w)$ and $w_2 \in min(w, |\delta|_w)$ and
   (b) for every α, β such that $w_1 \in min(w, |\alpha|_w)$ and $w_2 \in min(w, |\beta|_w)$ we have $\alpha \vee \beta \Rightarrow \beta \in w$.
(3) D is the set of terms.
(4) $D'(w) = \{t \mid A(t) \in w\}$ for term t.
(5) V is given by: for term t, $V(t) = t$. For n-place predicate symbol P, terms $t_1, \ldots, t_n$, and $w \in W$, $\langle t_1, \ldots, t_n\rangle \in V(P, w)$ iff $P(t_1, \ldots, t_n) \in w$.

Theorem 8.6.

(1) $\forall w \in W$ *and* $\forall w_1 \in W_w$, *we have* $R_w w_1 w_1$.

(2) $\forall w \in W$ *and* $\forall w_1, w_2, w_3 \in W_w$, *we have that if* $R_w w_1 w_2$ *and* $R_w w_2 w_3$ *then* $R_w w_1 w_3$.

Proof.

(1) Assume $w_1 \in W_w$. For any α where $w_1 \in min(w, |\alpha|_w)$ we have $\alpha \vee \alpha \Rightarrow \alpha \in w$, so $R_w w_1 w_1$.

(2) Assume that $R_w w_1 w_2$ and $R_w w_2 w_3$. First, our construction guarantees that there is a formula α such that $w_1 \in min(w, |\alpha|_w)$; similar comments apply to w_2 and w_3. Since $R_w w_1 w_2$, for every α, β where $w_1 \in min(w, |\alpha|_w)$ and $w_2 \in min(w, |\beta|_w)$ we have $\alpha \vee \beta \Rightarrow \beta \in w$. Similarly, since $R_w w_2 w_3$, if $w_2 \in min(w, |\beta|_w)$ and $w_3 \in min(w, |\gamma|_w)$ we have $\beta \vee \gamma \Rightarrow \gamma \in w$. By Theorem 8.3.3, we have that $((\alpha \vee \beta \Rightarrow \beta) \wedge (\beta \vee \gamma \Rightarrow \gamma)) \supset (\alpha \vee \gamma \Rightarrow \gamma) \in w$ and since w is deductively closed, so $\alpha \vee \gamma \Rightarrow \gamma \in w$. But this means that for every α such that $w_1 \in min(w, |\alpha|_w)$ and for every γ such that $w_3 \in min(w, |\gamma|_w)$ we have $\alpha \vee \gamma \Rightarrow \gamma \in w$, whence $R_w w_1 w_3$. □

Theorem 8.7. *The canonical model is, in fact, a QS-model.*

Proof. Straightforward. $W \neq \emptyset$ since, if we let $Th(\top)$ be the set of theorems, then $Th(\top)$ has the $\forall,\forall_c$-property by Lemma 8.1. From Lemma 8.3, $Th(\top) \cup \{\alpha\}$, for arbitrary consistent α, can be extended to a maximal consistent set with the $\forall,\forall_c$-property. R was shown to have the appropriate properties in Theorem 8.6. That $D'(w) \neq \emptyset$ for every $w \in W$ follows from the fact that (via **Ax**) $\exists x A(x) \in w$: since $\exists x A(x) \in w$, so $\neg\forall\neg x A(x) \in w$; so in Lemma 8.3, in the enumeration of formulas in the construction of w, where δ_i is $\neg\forall\neg x A(x)$ we add $\neg\neg A(t)$ to w, hence $A(t)$, for some term t, from which by the definition of the canonical model $t \in D'(w)$, so $D'(w) \neq \emptyset$. □

Theorem 8.8. *Let* $M = \langle W, R, D, D', V\rangle$ *be the canonical model. For any* $\alpha \in \mathcal{L}_{FC}$ *and* $w \in W$ *we have* $M, w \models \alpha$ *iff* $\alpha \in w$.

Proof. Note that the last condition in the theorem can be expressed as $\|\alpha\|^M = |\alpha|_w$.

The proof is by induction on the composition of a formula.

(1) If α is an atomic formula $P(t_1, \ldots, t_n)$, then for $w \in W$ we have $M, w \models P(t_1, \ldots, t_n)$ iff $\langle V(t_1), \ldots, V(t_n)\rangle \in V(P, w)$ iff $\langle t_1, \ldots, t_n\rangle \in V(P, w)$ iff $P(t_1, \ldots, t_n) \in w$.

(2) The proof for $\neg$ and $\supset$ is straightforward.

(3) For $\Rightarrow$, assume first that $\alpha \Rightarrow \beta \in w$. By the induction hypothesis $|\alpha|_w = \|\alpha\|^M$, so $min(w, |\alpha|_w) = min(w, \|\alpha\|^M)$. If $w_1 \in min(w, |\alpha|_w)$ then by Definition 8.5 we have $\beta \in w_1$. By the induction hypothesis $M, w_1 \models \beta$ and since this holds for all such w_1, by the definition of truth of $\Rightarrow$ we obtain $M, w \models \alpha \Rightarrow \beta$.

Conversely, assume that $\alpha \Rightarrow \beta \notin w$. So $\neg(\alpha \Rightarrow \beta) \in w$, since w is maximal consistent. By Theorem 8.5, $w^-(\alpha) \cup \{\neg\beta\}$ has a maximal consistent extension w_1 that has the $\forall,\forall_c$-property. By Definition 8.5, $w_1 \in min(w, |\alpha|_w)$. Since $\neg\beta \in w_1$ so

$\beta \notin w_1$. So $w_1 \notin |\beta|_w$, so by the induction hypothesis $M, w_1 \not\models \beta$, so $M, w \not\models \alpha \Rightarrow \beta$.

(4) For $\forall_c$, assume first that $\forall_c x\alpha(x) \in w$. We need to show that $M, w \models \forall_c x\alpha$, that is, for every V' which is the same as V except $V'(x) \in D'(w)$ where $M' = \langle W, R, D, D', V'\rangle$, that $M', w \models \alpha$. So consider any V' like V but where $V'(x) \in D'(w)$. There are two cases:
   (a) $V'(x) = t$ for some term t free for x in α. Since $\forall_c x\alpha(x) \supset (A(t) \supset \alpha(t)) \in w$ for t free for x in α (**FUI**), we have by modus ponens that $A(t) \supset \alpha(t) \in w$. By the induction hypothesis, $M, w \models A(t) \supset \alpha(t)$ for t free for x in α, and so $M', w \models A(x) \supset \alpha(x)$. Since $M', w \models A(x)$ so by modus ponens $M', w \models \alpha(x)$.
   (b) $V'(x) = y$ for some variable y for which x is not free in α. Consider a bound alphabetic variant α' of α such that $\forall_c y$ does not occur in α'. We can safely replace free occurrences of x in α' by y, since no quantifier containing y occurs in α'. Since α and α' are bound alphabetic variants, by Theorem 8.1 we have $\vdash \forall_c x\alpha \equiv \forall_c x\alpha'$, and since $\forall_c x\alpha \in w$, so $\forall_c x\alpha' \in w$. Now $\forall_c x\alpha'(x) \supset (A(y) \supset \alpha'(y)) \in w$ so by modus ponens, $A(y) \supset \alpha'(y) \in w$. By the induction hypothesis, we have $M, w \models A(y) \supset \alpha'(y)$, and so by the same argument as in the first part, $M', w \models A(x) \supset \alpha'(x)$. Since α' is a bound alphabetical variant of α, so $M', w \models A(x) \supset \alpha(x)$. Since $M', w \models A(x)$ so by modus ponens $M', w \models \alpha(x)$.

   So we have shown that for every V' the same as V except $V'(x) \in D'(w)$ where $M' = \langle W, R, D, D', V'\rangle$, that $M', w \models \alpha$. Thus, by the definition of V, we have $M, w \models \forall_c x\alpha$.
   Conversely, assume that $\forall_c x\alpha \notin w$. Since w has the $\forall_c$-property, there must be some t such that $A(t) \supset \alpha(t) \notin w$. By the induction hypothesis, $M, w \not\models A(t) \supset \alpha(t)$. If V' is exactly like V, except that $V'(x) = V(t)$, then $M', w \not\models A(x) \supset \alpha(x)$, and so $M, w \not\models \forall_c x\alpha$.
   This completes the case for $\forall_c$.

(5) $\forall$ follows essentially as a simpler version of $\forall_c$. □

Theorem 8.9 (Completeness of $\mathcal{QS}$). *If* $\models \alpha$ *then* $\vdash \alpha$.

Proof. Let α be a consistent formula, and let $Th(\top)$ be the set of theorems. $Th(\top)$ has the $\forall$,$\forall_c$-property by Lemma 8.1. From Lemma 8.3, $Th(\top) \cup \{\alpha\}$ can be extended to a maximal consistent set with the $\forall$,$\forall_c$-property. So, by Theorem 8.8, every consistent wff is true in some world in the canonical model. Similarly, every non-theorem is false in some world in the canonical model. Since the canonical model is a $\mathcal{QS}$-model (Theorem 8.7), we obtain the desired result. □

Acknowledgement

I thank the reviewers for their detailed and helpful comments. This research was funded by the Natural Science and Engineering Research Council of Canada grant A0884.

References

[1] E. Adams, The Logic of Conditionals, D. Reidel, Dordrecht, Netherlands, 1975.

[2] N. Asher, M. Morreau, Commonsense entailment: a model theory of nonmonotonic reasoning, in: Proceedings International Joint Conference on Artificial Intelligence (IJCAI-91), Sydney, Australia, August 1991, pp. 387–392.

[3] F. Bacchus, Representing and Reasoning with Probabilistic Knowledge, MIT Press, Cambridge, MA, 1990.

[4] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, H. Prade, Inconsistency management and prioritized syntax-based entailment, in: Proceedings International Joint Conference on Artificial Intelligence (IJCAI-93), Chambéry, France, 1993, pp. 640–645.

[5] S. Benferhat, D. Dubois, H. Prade, Representing default rules in possibilistic logic, in: Proceedings 3rd International Conference on the Principles of Knowledge Representation and Reasoning, Cambridge, MA, October 1992, pp. 673–684.

[6] F. Bacchus, A. Grove, J. Halpern, D. Koller, Statistical foundations for default reasoning, in: Proceedings International Joint Conference on Artificial Intelligence (IJCAI-93), Chambéry, France, 1993, pp. 563–569.

[7] C. Boutilier, Conditional logics for default reasoning and belief revision, Ph.D. Thesis, Department of Computer Science, University of Toronto, Toronto, Ontario, 1992.

[8] C. Boutilier, Conditional logics of normality: a modal approach, Artificial Intelligence 68 (1) (1994) 87–154.

[9] J.P. Burgess, Quick completeness proofs for some logics of conditionals, Notre Dame Journal of Formal Logic 22 (1) (1981) 76–84.

[10] B.F. Chellas, Modal Logic, Cambridge University Press, Cambridge, 1980.

[11] J.P. Delgrande, A first-order conditional logic for prototypical properties, Artificial Intelligence 33 (1) (1987) 105–130.

[12] D. Dubois, H. Prade, Possibilistic logic, in: D.M. Gabbay, C.J. Hogger, J.A. Robinson (Eds.), Nonmonotonic Reasoning and Uncertain Reasoning, Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3, Clarendon Press, Oxford, 1994, pp. 439–513.

[13] N. Friedman, J. Halpern, D. Koller, First-order conditional logic revisited, in: Proceedings AAAI National Conference on Artificial Intelligence (AAAI-96), Portland, OR, August 1996, pp. 1305–1323.

[14] M. Fitting, Basic modal logic, in: D.M. Gabbay, C.J. Hogger, J.A. Robinson (Eds.), Logical Foundations, Handbook of Logic in Artificial Intelligence and Logic Programming, Clarendon Press, Oxford, 1993, pp. 365–448.

[15] J.W. Garson, Quantification in modal logic, in: D. Gabbay, F. Guenthner (Eds.), Handbook of Philosophical Logic, Vol. 2, D. Reidel, Dordrecht, Netherlands, 1984, pp. 249–307.

[16] M. Goldszmidt, P. Morris, J. Pearl, A maximum entropy approach to nonmonotonic reasoning, in: Proceedings AAAI National Conference on Artificial Intelligence (AAAI-90), Boston, MA, 1990.

[17] M. Goldszmidt, J. Pearl, System-Z+: a formalism for reasoning with variable-strength defaults, in: Proceedings AAAI National Conference on Artificial Intelligence (AAAI-91), Anaheim, CA, 1991, pp. 399–404.

[18] H. Geffner, J. Pearl, Conditional entailment: bridging two approaches to default reasoning, Artificial Intelligence 53 (2–3) (1992) 209–244.

[19] G.E. Hughes, M.J. Cresswell, An Introduction to Modal Logic, Methuen and Co., London, 1968.

[20] G.E. Hughes, M.J. Cresswell, A Companion to Modal Logic, Methuen and Co., London, 1984.

[21] G. Hirst, Existence assumptions in knowledge representation, Artificial Intelligence 49 (1–3) (1991) 199–242.

[22] S. Kraus, D. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, Artificial Intelligence 44 (1–2) (1990) 167–207.

[23] H.E. Kyburg, Probability and the Logic of Rational Belief, Wesleyan University Press, Middletown, CT, 1961.

[24] P. Lamarre, S4 as the conditional logic of nonmonotonicity, in: Proceedings 2nd International Conference on the Principles of Knowledge Representation and Reasoning, Cambridge, MA, April 1991, pp. 357–367.

[25] D. Lehmann, What does a conditional knowledge base entail?, in: R. Brachman, H. Levesque, R. Reiter (Eds.), Proceedings 1st International Conference on the Principles of Knowledge Representation and Reasoning, Toronto, Ontario, May 1989, pp. 212–222.

[26] H.J. Levesque, Foundations of a functional approach to knowledge representation, Artificial Intelligence 23 (1984) 155–212.
[27] D. Lewis, Counterfactuals, Harvard University Press, Cambridge, MA, 1973.
[28] D. Lehmann, M. Magidor, Preferential logics: the predicate calculus case (extended abstract), in: Proceedings 3rd Conference on Theoretical Aspects of Reasoning about Knowledge, Pacific Grove, CA, 1990, pp. 57–72.
[29] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, Artificial Intelligence 55 (1) (1992) 1–60.
[30] J. McCarthy, First-order theories of individual concepts and propositions, in: D. Michie (Ed.), Machine Intelligence 9, Edinburgh University Press, Edinburgh, UK, 1979, pp. 129–147.
[31] J. McCarthy, Circumscription—a form of nonmonotonic reasoning, Artificial Intelligence 13 (1980) 27–39.
[32] J. McCarthy, Applications of circumscription to formalizing common-sense knowledge, Artificial Intelligence 28 (1986) 89–116.
[33] R.C. Moore, Semantical considerations on nonmonotonic logic, Artificial Intelligence 25 (1985) 75–94.
[34] W. Nejdl, The P-systems: a systematic classification of logics of nonmonotonicity, in: Proceedings AAAI National Conference on Artificial Intelligence (AAAI-91), Anaheim, CA, 1991, pp. 366–372.
[35] D. Nute, Topics in Conditional Logic, Philosophical Studies Series in Philosophy, Vol. 20, D. Reidel, Dordrecht, Netherlands, 1980.
[36] J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, Morgan Kaufmann, San Mateo, CA, 1988.
[37] J. Pearl, Probabilistic semantics for nonmonotonic reasoning: a survey, in: Proceedings 1st International Conference on the Principles of Knowledge Representation and Reasoning, Toronto, Ontario, May 1989, pp. 505–516.
[38] J. Pearl, System Z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning, in: Proceedings 3rd Conference on Theoretical Aspects of Reasoning about Knowledge, Pacific Grove, CA, 1990, pp. 121–135.
[39] D.L. Poole, A logical framework for default reasoning, Artificial Intelligence 36 (1) (1988) 27–48.
[40] D.L. Poole, The effect of knowledge on belief: conditioning, specificity, and the lottery paradox in default reasoning, Artificial Intelligence 49 (1–3) (1991) 281–307.
[41] H. Putnam, Is semantics possible?, in: Mind, Language and Reality: Philosophical Papers, Vol. II, Cambridge University Press, 1975, pp. 139–152.
[42] H. Putnam, The 'corroboration' of theories, in: Mind, Language and Reality: Philosophical Papers, Vol. I, 2nd ed., Cambridge University Press, 1979, pp. 250–270.
[43] R. Reiter, A logic for default reasoning, Artificial Intelligence 13 (1980) 81–132.
[44] K. Schlechta, Defaults as generalized quantifiers, J. Logic Comput. 5 (4) (1995) 473–494.
[45] K. Schlechta, A two-stage approach to first-order default reasoning, Fundamenta Informaticae 28 (3–4) (1996) 377–402.
[46] D. Scott, Advice on modal logic, in: K. Lambert (Ed.), Philosophical Problems in Logic: Some Recent Developments, D. Reidel, Dordrecht, Netherlands, 1970, pp. 143–173.
[47] R.C. Stalnaker, R.H. Thomason, A semantic analysis of conditional logic, Theoria 36 (1970) 23–42.
[48] R.C. Stalnaker, A theory of conditionals, in: N. Rescher (Ed.), Studies in Logical Theory, Basil Blackwell, Oxford, 1968, pp. 98–112.
[49] B.C. van Fraassen, The completeness of free logic, Z. Math. Logik Grundlagen Math. 12 (1966) 219–234.