

Towards an Axiom System for Default Logic

Gerhard Lakemeyer
Dept. of Computer Science
RWTH Aachen
52056 Aachen
Germany
gerhard@cs.rwth-aachen.de

Hector J. Levesque
Dept. of Computer Science
University of Toronto
Toronto, Ontario
Canada M5S 3A6
hector@cs.toronto.edu

Abstract

Recently, Lakemeyer and Levesque proposed a logic of only-knowing which precisely captures three forms of nonmonotonic reasoning: Moore's Autoepistemic Logic, Konolige's variant based on moderately grounded expansions, and Reiter's default logic. Defaults have a uniform representation under all three interpretations in the new logic. Moreover, the logic itself is monotonic, that is, nonmonotonic reasoning is cast in terms of validity in the classical sense. While Lakemeyer and Levesque gave a model-theoretic account of their logic, a proof-theoretic characterization remained open. This paper fills that gap for the propositional subset: a sound and complete axiom system in the new logic for all three varieties of default reasoning. We also present formal derivations for some examples of default reasoning. Finally we present evidence that it is unlikely that a complete axiom system exists in the first-order case, even when restricted to the simplest forms of default reasoning.

Introduction

Recently, Lakemeyer and Levesque (2005) proposed a logic of only-knowing called $\mathcal{O}_3\mathcal{L}$, which precisely captures three forms of nonmonotonic reasoning: Moore's Autoepistemic Logic (AEL) (Moore 1985), Konolige's variant of AEL using moderately grounded expansions (Konolige 1988), and Reiter's default logic (DL) (Reiter 1980). In $\mathcal{O}_3\mathcal{L}$, defaults have the following uniform representation in all cases:

$$K\alpha \wedge M\beta \supset \gamma. \quad (1)$$

Here K and M are modal operators in a first-order language, and the default may be read as "if α is believed and it is consistent to believe β , then conclude γ ."¹ To get the correspondence with DL, (1) is understood as the translation of the Reiter default rule $\frac{\alpha:\beta}{\gamma}$. This translation was first proposed by Konolige (1988) in his attempt to map DL into AEL. That attempt failed essentially because AEL, including Konolige's variant, assumes that M is the dual of K , that is, that $M\beta \equiv \neg K\neg\beta$ is valid, while DL requires this duality to be given up.

Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

¹In general, there may be a conjunction of $M\beta_i$ and, if α, β_i, γ contain free variables, the default is $\forall \vec{x}. K\alpha \wedge \bigwedge M\beta_i \supset \gamma$.

$\mathcal{O}_3\mathcal{L}$ deals with each form of default reasoning using a variant of *only-knowing*, a notion which was originally introduced by Levesque (1990). The idea, roughly, is this: Consider a Reiter default theory $\langle \phi, D \rangle$, where ϕ and D are finite sets of sentences (facts) and closed default rules, respectively, and let Δ be the translation of the defaults according to (1). Skeptical reasoning in DL, that is, the question as to whether a sentence α is an element of all Reiter extensions, turns into the question as to whether the sentence

$$O_r(\phi \wedge \Delta) \supset K\alpha \quad (2)$$

is valid in $\mathcal{O}_3\mathcal{L}$, that is, whether knowing α follows logically from knowing only $\phi \wedge \Delta$. By varying the meaning of O_r , the question can be modified into whether α is an element of all Moore extensions² or all Konolige extensions. Technically, this is done by considering additional only-knowing operators O_M and O_K , for Moore and Konolige, respectively.

A nice feature of $\mathcal{O}_3\mathcal{L}$ is its simple possible-world semantics using two sets of worlds (epistemic states), one for the interpretation of K and one for M . When considering the meaning of O_M and O_K , the two sets are always the same so that K and M are in fact duals as required by Moore and Konolige. It is only in the case of O_r where the two sets differ and the duality no longer holds. Note also that $\mathcal{O}_3\mathcal{L}$ itself is a *monotonic* logic. In particular, the different forms of nonmonotonic reasoning are all expressed in terms of certain *valid* sentences of the logic such as (2).

With this in mind, one important question has remained open, namely what a proof theory for $\mathcal{O}_3\mathcal{L}$ would be like. This paper answers that question for the propositional subset of the language. There are at least two reasons why having a proof theory for $\mathcal{O}_3\mathcal{L}$ over and above a semantic account is useful. For one, it puts nonmonotonic inference back into the realm of *proving theorems* in the classical sense rather than having to resort to meta-logical arguments involving extensions, set intersections, or fixed points. For another, axiom systems in general provide a compact representation of the valid sentences of a logic, which often sheds new light on its properties. And indeed, as we will see below, the axiomatization of O_r and hence of Reiter's DL turns out to be surprisingly simple once the axioms and inference rules for

²Historically, AEL uses the term *expansion* instead of extension. For uniformity, we call them all extensions.

O_M and O_K are in place. Thus the connections among the approaches of Reiter, Moore and Konolige are further illuminated.

Why do we limit ourselves to the propositional subset? First of all, Halpern and Lakemeyer (1995) already proved the incompleteness of Levesque's original axiomatization of only-knowing, which is a proper subset of $O_3\mathcal{L}$. And, as we will argue at the end of the paper, even when $O_3\mathcal{L}$ is restricted to the simplest forms of first-order default reasoning, it is unlikely that a complete axiom system exists. So we settle for the propositional case. For the sake of simplicity, we further limit the language so that it just covers the translations of Reiter default theories, the main concern from an AI perspective.

Before $O_3\mathcal{L}$, there have been a number of proposals showing how DL can be embedded faithfully in a modal epistemic logic such as (Lin and Shoham 1988; Lifschitz 1994; Amati et al. 1997; Denecker et al. 2003). However, in contrast to $O_3\mathcal{L}$, these approaches require either that AEL defaults be represented differently from DL defaults or a very complex semantics using certain fixed-point constructions. As already mentioned, Levesque (1990) considered a subset of $O_3\mathcal{L}$ called \mathcal{OL} , which is essentially $O_3\mathcal{L}$ minus the operators O_K and O_R . Levesque provided a sound and complete proof theory for the propositional subset of \mathcal{OL} , which we will reuse here. Other notions of only-knowing were considered in (Levesque and Lakemeyer 2001; Halpern and Lakemeyer 2001; Waaler 2004), but in terms of nonmonotonic reasoning they did not go beyond AEL. There have been proof-theoretic characterizations of DL such as (Bonatti and Olivetti 1997). However, in contrast to our work, these require nonmonotonic inference rules.

The rest of the paper is organized as follows. In the next section, we review the formal details of the logic $O_3\mathcal{L}$, its syntax and semantics. Then we turn our attention to an axiom system for the propositional subset of $O_3\mathcal{L}$, followed by some example derivations. After that we briefly examine the first order case and conclude.

Syntax and semantics

The symbols of the $O_3\mathcal{L}$ language are the usual logical connectives, quantifiers, punctuation, variables, the equality symbol, predicates (of every arity), a countably infinite set of *standard names*, and the modal operators M , K , O_M , O_K , and O_R . For simplicity, constants and function symbols are omitted.³ The *terms* of $O_3\mathcal{L}$ are the variables and standard names. The *formulas* of $O_3\mathcal{L}$ are defined by the following:

1. if t_1, \dots, t_k are terms and P is a predicate of arity k , then $P(t_1, \dots, t_k)$ is an (atomic) formula;
2. if t_1 and t_2 are terms, then $(t_1 = t_2)$ is a formula;
3. if α and β are formulas and x is any variable, then $\neg\alpha$, $(\alpha \wedge \beta)$, and $\forall x.\alpha$ are formulas, as are the *modal* formulas, $M\alpha$, $K\alpha$, $O_M\alpha$, $O_K\alpha$, and $O_R\alpha$.

³The standard names can be thought of as constants that satisfy the unique name assumption and an infinitary version of domain closure.

As usual, we treat $(\alpha \vee \beta)$, $(\alpha \supset \beta)$, $(\alpha \equiv \beta)$, and $\exists x.\alpha$ as abbreviations. The notion of a free and bound variable is defined in the usual way, and α_n^x means α with all free occurrences of x replaced by n . Similarly, we write α_γ^β to mean α with all occurrences of the subformula β replaced by γ . A formula without free variables is called a *sentence*, and a formula of the form $P(n_1, \dots, n_k)$, where the n_i are standard names, a *primitive sentence*. Formulas without modal operators are called *objective*, and those where all the predicates appear in the scope of a modal operator are called *subjective*. Formulas without O_M , O_K and O_R are called *basic*. To simplify our axiomatization, we require that the operators O_M , O_K and O_R apply only to basic sentences where K and M do not occur nested. Note that this restriction still allows us to cover translations of default theories as discussed in the introduction.

The semantics of $O_3\mathcal{L}$ builds on the semantics of \mathcal{OL} from (Levesque and Lakemeyer 2001). The starting point is the notion of a *world* (or world state) which is a function from the primitive sentences to $\{0, 1\}$. We let W be the set of all worlds. An epistemic state in \mathcal{OL} is any set of worlds. What is different in $O_3\mathcal{L}$ is that *two* epistemic states are used, one to interpret formulas with K , and one to interpret formulas with M (since, as we noted, there will be contexts where the two operators are not duals).

Let w be a world, and e_1 and e_2 be epistemic states. We can define when a basic sentence α is *true* wrt e_1 , e_2 , and w , which we write as $e_1, e_2, w \models \alpha$, as follows:

1. $e_1, e_2, w \models P(n_1, \dots, n_k)$ iff $w[P(n_1, \dots, n_k)] = 1$;
2. $e_1, e_2, w \models (n_1 = n_2)$ iff n_1 and n_2 are the same standard name;
3. $e_1, e_2, w \models \neg\alpha$ iff $e_1, e_2, w \not\models \alpha$;
4. $e_1, e_2, w \models (\alpha \wedge \beta)$ iff $e_1, e_2, w \models \alpha$ and $e_1, e_2, w \models \beta$;
5. $e_1, e_2, w \models \forall x.\alpha$ iff $e_1, e_2, w \models \alpha_n^x$ for every standard name n ;
6. $e_1, e_2, w \models K\alpha$ iff $e_1, e_2, w' \models \alpha$ for every $w' \in e_1$;
7. $e_1, e_2, w \models M\alpha$ iff $e_1, e_2, w' \models \alpha$ for some $w' \in e_2$.

Observe that when $e_1 = e_2$, K and M will behave like the usual duals. Next we define O_M to coincide with the O of (Levesque and Lakemeyer 2001):

8. $e_1, e_2, w \models O_M\alpha$ iff for every $w' \in W$, $e_1, e_2, w' \models \alpha$ iff $w' \in e_1$.

This has the effect of replacing an “if” in the clause for K by an “iff”. Finally, the definitions of O_K and O_R use this one:

9. $e_1, e_2, w \models O_K\alpha$ iff for every e' such that $e_1 \subseteq e'$, $e', e', w \models O_M\alpha$ iff $e' = e_1$;
10. $e_1, e_2, w \models O_R\alpha$ iff for every e' such that $e_1 \subseteq e'$, $e', e_2, w \models O_M\alpha$ iff $e' = e_1$;

Note that the definition of O_K and O_R differ only in one place: where O_K uses the e' for its second epistemic argument (thus keeping the two arguments identical), O_R uses the given e_2 .

To complete the specification of the logic, we define $e, w \models \alpha$ to mean $e, e, w \models \alpha$, and we say that a sentence α is *valid* (which we write as $\models \alpha$) iff $e, w \models \alpha$ for every

e and w . If α is objective, we often omit the e and write $w \models \alpha$; if α is subjective, we write $e \models \alpha$ or $e_1, e_2 \models \alpha$.

Various properties of $\mathcal{O}_3\mathcal{L}$ were discussed in (Lakemeyer and Levesque 2005) which we will not repeat here. We remark that the rule for \mathbf{O}_R was slightly different there, namely

$$10' \quad e_1, e_2, w \models \mathbf{O}_R\alpha \text{ iff for every } e' \text{ such that } e_1 \subseteq e', \\ e', e_1, w \models \mathbf{O}_M\alpha \text{ iff } e' = e_1,$$

that is, e_2 did not appear on the R.H.S. It is easy to verify that the logic remains the same given the fact that validity is defined with respect to identical e_1 and e_2 , and the restriction that \mathbf{O}_R does not occur nested.

Now we turn to the axiomatization of the propositional subset of the language.

An Axiom System

We begin with an axiomatic characterization of \mathbf{O}_M , which was originally proposed in (Levesque 1990). For this purpose, it is convenient to consider \mathbf{O}_M not as a primitive notion but to define it in terms of \mathbf{K} and yet another modal operator \mathbf{N} . One way to read $\mathbf{O}_M\alpha$ is to say that α is believed and nothing more, whereas $\mathbf{K}\alpha$ says that α is believed, and perhaps more. In other words, $\mathbf{K}\alpha$ means that α *at least* is believed to be true. A natural dual to this is to say that α *at most* is believed to be false, which we write $\mathbf{N}\alpha$. The idea is that $\mathbf{O}_M\alpha$ would then be *definable* as $(\mathbf{K}\alpha \wedge \mathbf{N}\neg\alpha)$, that is, at least α is believed and at most α is believed. So, *exactly* α is believed. In other words, we are taking \mathbf{K} to specify a lower bound on what is believed (since there may be other beliefs) and \mathbf{N} to specify an upper bound on beliefs (since there may be fewer beliefs).

These bounds can be seen most clearly when talking about objective sentences. Given an epistemic state as specified by a set of world states e , to say that $\mathbf{K}\phi$ is true wrt e is to say that e is a subset of the states where ϕ is true. By symmetry then, $\mathbf{N}\neg\phi$ will be true when the set of states satisfying ϕ are a subset of e . The fact that e must contain all of these states means that nothing else can be believed that would eliminate any of them. This is the sense in which no more than ϕ is known. Finally, as before, $\mathbf{O}_M\phi$ is true iff both conditions hold and the two sets coincide.

This leads us to the precise definition of $\mathbf{N}\alpha$:

$$e_1, e_2, w \models \mathbf{N}\alpha \text{ iff} \\ \text{for every } w', \text{ if } e_1, e_2, w' \models \alpha \text{ then } w' \in e_1.$$

On closer inspection, it turns out that \mathbf{N} behaves just like an ordinary belief operator (like \mathbf{K}). This is most clearly seen by rephrasing very slightly the definition of \mathbf{N} . Let $\overline{e_1}$ stand for the set of worlds not in e_1 .

$$e_1, e_2, w \models \mathbf{N}\alpha \text{ iff for every } w' \in \overline{e_1}, e_1, e_2, w' \models \alpha.$$

We are now ready to present an axiomatization:

Axioms:

1. The axioms of propositional logic
2. The axioms for \mathbf{K} , \mathbf{N} , and \mathbf{M} :
 - Let L stand for both \mathbf{K} and \mathbf{N} .
 - (a) $L\alpha$, where α is an instance of an axiom (1)
 - (b) $L(\alpha \supset \beta) \supset L\alpha \supset L\beta$

(c) $\sigma \supset L\sigma$, where σ is subjective

(d) Axiom for \mathbf{M} : $\mathbf{M}\alpha \equiv \neg\mathbf{K}\neg\alpha$

3. The axioms for \mathbf{O}_M :

(a) The definition of \mathbf{O}_M : $\mathbf{O}_M\alpha \equiv (\mathbf{K}\alpha \wedge \mathbf{N}\neg\alpha)$.

(b) The \mathbf{N} vs. \mathbf{K} axiom: $(\mathbf{N}\phi \supset \neg\mathbf{K}\phi)$, where ϕ is any objective sentence such that $\not\models \phi$.

4. The axioms for \mathbf{O}_R :

(a) $\mathbf{O}_R\alpha \equiv \mathbf{O}_K\alpha$, provided α has no \mathbf{M} operators

(b) $\mathbf{M}\phi \supset (\mathbf{O}_R\alpha \equiv \mathbf{O}_R\alpha_{true}^{M\phi})$

(c) $\neg\mathbf{M}\phi \supset (\mathbf{O}_R\alpha \equiv \mathbf{O}_R\alpha_{false}^{M\phi})$

5. The axiom for \mathbf{O}_K : $\mathbf{O}_K\alpha \supset \mathbf{O}_M\alpha$

Inference Rules:

1. Modus Ponens: From α and $\alpha \supset \beta$, derive β .

2. Rules for \mathbf{O}_K : (Let ϕ_i and ψ be objective.)

(a) From $\mathbf{O}_M\phi_1 \supset \mathbf{O}_M\alpha$,
 $\mathbf{O}_M\phi_2 \supset \mathbf{O}_M\alpha$,
 $\mathbf{O}_M\phi_1 \supset \mathbf{K}\phi_2$,
 and $\mathbf{O}_M\phi_2 \supset \neg\mathbf{K}\phi_1$,
 derive $\mathbf{K}\phi_1 \supset \neg\mathbf{O}_K\alpha$.

(b) From $\mathbf{O}_M\alpha \supset \mathbf{O}_M\psi \vee \bigvee \mathbf{O}_M\phi_i$,
 $\mathbf{O}_M\psi \supset \mathbf{O}_M\alpha$,
 and $\mathbf{O}_M\psi \supset \bigwedge \neg\mathbf{K}\phi_i$,
 derive $\mathbf{O}_M\psi \supset \mathbf{O}_K\alpha$.

Let us review this proof theory. Axiom (1) and Rule (1) give us classical propositional logic. When we add Axioms (2), we get that \mathbf{K} and \mathbf{N} behave exactly like operators in the modal logic $K45$ (Chellas 1980). Moreover, because of (2c), they are mutually introspective. Axiom (2d) makes \mathbf{M} the dual of \mathbf{K} but does not mean that \mathbf{M} can be replaced everywhere by $\neg\mathbf{K}\neg$. In particular, the replacement will not be sanctioned within \mathbf{O}_R .

When we add the Axioms (3), we get that \mathbf{O}_M behaves just like the \mathbf{O} operator in (Levesque 1990). Note that Axiom (3b) appeals to satisfiability in classical propositional logic. This could be axiomatized separately, but we do not do so here. This dependence on satisfiability means that the set of instances of this axiom in a first-order setting would not be recursively enumerable. This is unfortunately how it must be, however, since the valid sentences are not recursively enumerable either. This is the price we pay for trying to formalize a notion of default reasoning that is “consistency-based” like Reiter’s DL.

So far, the axioms and rules are all in (Levesque 1990) (modulo \mathbf{M} , which was not used there). The new material concerns \mathbf{O}_K and \mathbf{O}_R . Turning to \mathbf{O}_R first, there are three axioms which taken together characterize \mathbf{O}_R in terms of \mathbf{O}_K in an inductive fashion. Axiom (4a) says that \mathbf{O}_R reduces to \mathbf{O}_K if \mathbf{M} is not mentioned anywhere. Axiom (4b) and (4c) show how $\mathbf{O}_R\alpha$ reduces to a case with fewer occurrences of \mathbf{M} . In essence, this allows us to remove subformulas $\mathbf{M}\phi$ from α one by one until they are all gone, and then applying Axiom (4a) to obtain the reduction to \mathbf{O}_K . The Axioms (4) show very clearly the difference between Konolige and Reiter: while Konolige extensions are the result of a global minimization policy with respect to all Moore extensions

(as we will see below), Reiter extensions are obtained by first fixing the assignment to the M -sub-formulas and then applying Konolige-style minimization.

Finally, we get to O_K , which requires Axiom (5) and Rules (2a) and (2b). To understand what these are doing, it is useful to recall a result from (Levesque and Lakemeyer 2001):

Theorem 1: [Levesque and Lakemeyer] *Let α be a basic sentence without quantifiers. Then there is a set of objective sentences $\Phi = \{\phi_1, \dots, \phi_n\}$ such that*

$$\models O_M\alpha \equiv (O_M\phi_1 \vee \dots \vee O_M\phi_n).$$

In other words, the set Φ characterizes precisely the Moore extensions of α . Konolige extensions correspond to those $\phi_i \in \Phi$, which have the additional property (*) that there is no $\phi_j \in \Phi$ such that ϕ_i classically entails ϕ_j but not vice versa. With this in mind, Axiom (5) makes sure that all Konolige extensions are Moore extensions, Rule (2a) eliminates those that violate (*), and Rule (2b) ensures that those that satisfy (*) are included.

Now we turn to the main technical results of the paper, the soundness and completeness of the axioms. To prove soundness, we need the following lemma:

Lemma 1: *For any epistemic states e and e' ,*

1. *if $e \models M\phi$ then $e', e \models O_M\alpha$ iff $e', e \models O_{M\alpha_{true}}^{M\phi}$;*
2. *if $e \models \neg M\phi$ then $e', e \models O_M\alpha$ iff $e', e \models O_{M\alpha_{false}}^{M\phi}$.*

Proof: To prove (1), let $e \models M\phi$. Then a simple induction establishes that for any e' and w , $e', e, w \models \alpha$ iff $e', e, w \models \alpha_{true}^{M\phi}$, using the assumption that $e \models M\phi$. The lemma then follows from the semantic definition of O_M . The proof of (2) is completely analogous. ■

Theorem 2: *The axiom system presented above is sound for propositional $O_3\mathcal{L}$.*

Proof: The proof is not very difficult. Here we show the soundness of the Axioms (4). The soundness of Axiom (4a) was already proved in Theorem 5 of (Lakemeyer and Levesque 2005). To show that Axiom (4b) is valid, let $e \models M\phi$. Then $e \models O_R\alpha$ iff $e, e \models O_M\alpha$ and for all $e' \supseteq e$, $e', e \not\models O_M\alpha$ iff $e, e \models O_{M\alpha_{true}}^{M\phi}$ and for all $e' \supseteq e$, $e', e \not\models O_{M\alpha_{true}}^{M\phi}$ (by Lemma 1, Part 1) iff $e \models O_R\alpha_{true}^{M\phi}$.

The soundness of Axiom (4c) is established in a completely symmetric way, using Part 2 of Lemma 1 instead of Part 1. ■

The proof of completeness is much more challenging.

Theorem 3: *The axiom system presented above is complete for propositional $O_3\mathcal{L}$.*

Proof: (Sketch) The proof is adapted from the completeness proof for the logic $O\mathcal{L}$ in (Levesque and Lakemeyer 2001), which first appeared in (Halpern and Lakemeyer 1995). The idea is roughly this: First, we introduce a slight variant of the semantics of $O_3\mathcal{L}$, where N is not interpreted with respect to \bar{e} but with respect to a set e^N , which together with e covers all worlds, yet may overlap with e .

Then we show that (a) this semantics satisfies all the axioms and (b) a sentence is satisfiable in $O_3\mathcal{L}$ iff it is satisfiable under the new semantics (using a construction similar to the proof of Theorem 10.3.5 in (Levesque and Lakemeyer 2001)). The completeness proof is carried out with respect to the new semantics using a Henkin-style argument over maximal consistent sets (similar to the proof of Theorem 10.3.8 in (Levesque and Lakemeyer 2001)). ■

Some sample derivations

Before we turn to derivations of some nonmonotonic-reasoning examples, recall that $O_K\alpha \supset O_M\alpha$ is an axiom. Its counterpart, using O_R instead of O_K , is provable too, and will be used later.

Theorem 4: $O_R\alpha \supset O_M\alpha$ is a theorem.

This follows from Theorem 3 of (Lakemeyer and Levesque 2005) and the completeness theorem above.

Example 1: Tweety and Chilly.

The objective facts are that Tweety and Chilly are birds, Chilly does not fly (self-propelled through the air, that is), but Tweety's flying status is left unspecified (somewhat surprisingly, after all this time). The default information is that birds typically fly, which we need to instantiate for the two individuals. Thus we have:

$$\begin{aligned} \text{KB} &= \{ \text{Bird}(tw), \text{Bird}(ch), \neg\text{Fly}(ch) \} \\ \Delta &= \left\{ \begin{array}{l} \text{KBird}(tw) \wedge \text{MFLy}(tw) \supset \text{Fly}(tw) \\ \text{KBird}(ch) \wedge \text{MFLy}(ch) \supset \text{Fly}(ch) \end{array} \right\} \end{aligned}$$

We want to consider what logically follows from only knowing KB and Δ . We will show *formal derivations* in our axiom system that prove that $\text{Fly}(tw)$ is an element of every Moore, Konolige, and Reiter extension.

First, Moore. To derive $\text{K}\neg\text{Fly}(ch)$ from $O_M(\text{KB} \wedge \Delta)$ is easy: we use Axiom (3a) to get $\text{K}(\text{KB} \wedge \Delta)$, then apply ordinary modal $K45$ reasoning to do the rest (ignoring Δ). For Tweety, to derive $\text{KFLy}(tw)$, we again use (3a) to get $\text{K}(\text{KB} \wedge \Delta)$, then apply ordinary modal $K45$ to get to $(\text{KFLy}(tw) \vee \text{K}\neg\text{Fly}(tw))$ (*), but this is not enough. To go further, we apply Axiom (3a) to get $\text{N}\neg(\text{KB} \wedge \Delta)$, and then ordinary modal $K45$ reasoning to get $\text{N}(\text{KB} \supset \neg\text{Fly}(tw))$, using the fact that $\text{NK}\neg\text{Fly}(ch)$ is derivable by mutual introspection. Since $\text{KB} \not\models \neg\text{Fly}(tw)$, we get to use the Axiom (3b), to derive $\neg\text{K}(\text{KB} \supset \neg\text{Fly}(tw))$, and then, using modal $K45$ reasoning, $\neg\text{K}\neg\text{Fly}(tw)$. Combining this with (*), we get the desired conclusion.

We can put all this in a natural-deduction style proof annotated with justifications for each step. Here PL and $K45$ refer respectively to reasoning using the propositional logic and $K45$ (for K and N , including mutual introspection).

Proof: To prove that $(O_M(\text{KB} \wedge \Delta) \supset \text{KFLy}(tw))$ is a theorem, we proceed as follows:

1. $O_M(\text{KB} \wedge \Delta)$ Assumption.
2. $\text{K}(\text{KB} \wedge \Delta)$ 1;defn. of O_M .
3. $\text{K}\neg\text{K}\neg\text{Fly}(tw) \supset \text{KFLy}(tw)$ 2; $K45$.
4. $\neg\text{K}\neg\text{Fly}(tw) \supset \text{KFLy}(tw)$ 3; $K45$.

5. $K\text{-Fly}(ch)$ 2;K45.
6. $NK\text{-Fly}(ch)$ 5;K45 (mutual introspection).
7. $N\neg(KB \wedge \Delta)$ 1;defn. of O_M .
8. $N(KB \supset \neg\text{Fly}(tw))$ 6,7;K45.
9. $\neg K(KB \supset \neg\text{Fly}(tw))$ 8;N vs. K.
10. $\neg K\neg\text{Fly}(tw)$ 9;K45.
11. $K\text{Fly}(tw)$ 4,10;K45. ■

Konolige is next. To derive $K\text{-Fly}(ch)$ and $K\text{Fly}(tw)$ from $O_K(KB \wedge \Delta)$ is now easy: we use the previous derivation together with Axiom (5): $O_K\alpha \supset O_M\alpha$.

Finally, we turn to Reiter and DL. To derive $K\text{-Fly}(ch)$ and $K\text{Fly}(tw)$ from $O_R(KB \wedge \Delta)$, we proceed as with Konolige, but this time using Theorem 4. Strictly speaking, then, a derivation of $K\text{Fly}(tw)$ would need the above derivation and another generic one corresponding to Theorem 4.

Example 2: $O_K(Kp \supset p) \supset \neg Kp$ is a theorem.

The sentence $\alpha = (Kp \supset p)$ is perhaps the simplest that demonstrates the difference between Moore and Konolige. While α has two Moore extensions, one containing p and another containing only tautologies, Konolige rules out the former, and so in every Konolige extension, p is not known.⁴

Proof:

1. $O_Mp \supset O_M\alpha$ \mathcal{OL} .
2. $O_M\text{true} \supset O_M\alpha$ \mathcal{OL} .
3. $O_Mp \supset K\text{true}$ \mathcal{OL} .
4. $O_M\text{true} \supset \neg Kp$ \mathcal{OL} .
5. $O_Mp \supset \neg O_K\alpha$ 1,2,3,4;Rule 2a.
6. $O_M\alpha \supset (O_M\text{true} \vee O_Mp)$ \mathcal{OL} .
7. $O_M\text{true} \supset O_R\alpha$ 6,2,4;Rule 2b.
8. $O_R\alpha \supset O_M\alpha$ Ax. 5.
9. $O_R\alpha \supset (O_M\text{true} \vee O_Mp)$ 8,6;PL.
10. $O_R\alpha \supset O_M\text{true}$ 9,5;PL.
11. $O_R\alpha \supset \neg Kp$ 10,4;PL. ■

As a corollary we obtain that $O_R\alpha \supset \neg Kp$ is a theorem as well because of Axiom (4a).

Example 3: $\neg O_R[(Kp \supset p) \wedge (M\neg p \supset p)]$ is a theorem.

The sentence $\alpha = (Kp \supset p) \wedge (M\neg p \supset p)$ has been used to demonstrate the key difference between AEL and DL having to do with the duality or non-duality of K and M . It is easy to see that both $O_M\alpha$ and $O_R\alpha$ are logically equivalent to O_Mp because $M\neg p$ can be replaced by $\neg Kp$ there. In other words, there is a unique Moore and Konolige extension containing p . There is no Reiter extension, however, and so $\neg O_R\alpha$ is provable using our axiom system.

Proof: Let $\gamma = M\neg p$ and $\bar{\gamma} = \neg M\neg p$. Also let

$$\alpha^\gamma = \alpha_{\text{true}}^{M\neg p} = (Kp \supset p) \wedge (\text{true} \supset p) \text{ and}$$

$$\alpha^{\bar{\gamma}} = \alpha_{\text{false}}^{M\neg p} = (Kp \supset p) \wedge (\text{false} \supset p).$$

We proceed by showing that $\gamma \supset \neg O_R\alpha$ and $\bar{\gamma} \supset \neg O_R\alpha$ are both theorems, from which the result follows by propositional reasoning.

⁴To keep these derivations short, we use \mathcal{OL} as a justification for steps that only involve properties of O_M .

1. $\neg O_Mp \supset \neg O_M\alpha^\gamma$ \mathcal{OL} .
2. $\neg O_M\alpha^\gamma \supset \neg O_K\alpha^\gamma$ Ax. 5.
3. $\neg O_Mp \supset \neg O_K\alpha^\gamma$ 1,2;PL.
4. $\gamma \supset \neg Kp$ Ax. 2d;K45.
5. $\neg Kp \supset \neg O_Mp$ \mathcal{OL} .
6. $\gamma \supset \neg O_Mp$ 4,5;K45.
7. $\neg O_K\alpha^\gamma \supset \neg O_R\alpha^\gamma$ Ax. 4a;PL.
8. $\gamma \supset (O_R\alpha \equiv O_R\alpha^\gamma)$ Ax. 4b.
9. $\gamma \supset \neg O_R\alpha$ 6,3,7,8;PL.
10. $O_Mp \supset O_M\alpha^{\bar{\gamma}}$ \mathcal{OL} .
11. $O_M\text{true} \supset O_M\alpha^{\bar{\gamma}}$ \mathcal{OL} .
12. $O_Mp \supset K\text{true}$ \mathcal{OL} .
13. $O_M\text{true} \supset \neg Kp$ \mathcal{OL} .
14. $Kp \supset \neg O_K\alpha^{\bar{\gamma}}$ 10,11,12,13;Rule 2a.
15. $\bar{\gamma} \supset Kp$ Ax. 2d;K45.
16. $\neg O_K\alpha^{\bar{\gamma}} \supset \neg O_R\alpha^{\bar{\gamma}}$ Ax. 4a;PL.
17. $\bar{\gamma} \supset (O_R\alpha \equiv O_R\alpha^{\bar{\gamma}})$ Ax. 4c.
18. $\bar{\gamma} \supset \neg O_R\alpha$ 15,14,16,17;PL.
19. $\neg O_R\alpha$ 9,18;PL. ■

First-order defaults

The work of Moore and Konolige is purely propositional, and so the proof theory of $\mathcal{O}_3\mathcal{L}$ as it stands correctly captures those systems. Reiter's DL, on the other hand, is first-order, although it uses defaults in a limited way. For example, using an open default such as

$$\frac{\text{Bird}(x) : \text{Fly}(x)}{\text{Fly}(x)}$$

one can conclude by default that a bird will fly, but one does not conclude by default that all birds fly. So the variable x in the default is actually just a place holder for arbitrary ground terms. In this sense, Reiter does not really have quantified defaults, but rather sets of ground defaults, which again the proof theory of $\mathcal{O}_3\mathcal{L}$ can capture.

However, as noted in (Lakemeyer and Levesque 2005), we do have the additional possibility in $\mathcal{O}_3\mathcal{L}$ of actual quantified defaults, which raises the question as to whether we can provide an axiom system for them as well. As already noted, this appears to be unlikely, even if we restrict ourselves to Levesque's logic \mathcal{OL} , which is $\mathcal{O}_3\mathcal{L}$ without O_K and O_R . Initially, Levesque proposed an axiomatization, which consisted of our Axioms (1), (2) and (3), with (1) replaced by the axioms of first-order logic (plus some to deal with standard names) and the additional Barcan formula $\forall x.L\alpha \supset L\forall x.\alpha$ where L can be either K or N . Later, Halpern and Lakemeyer (1995) proved that these axioms were incomplete by showing that the sentence $N\zeta \supset \neg K\zeta$ is not derivable for

$$\zeta = [\exists x.P(x) \wedge \neg KP(x)] \vee [\exists x.\neg P(x) \wedge KP(x)].$$

The reason is, roughly, because the N vs K axiom (3b) is too weak when restricted to objective sentences.⁵

⁵We remark that the incompleteness hinges on the ability to quantify into a modal context as in the example. If we give up quantifying-in, first-order completeness is easy to establish, but we would not be able to handle quantified defaults.

