A Nonstandard Approach to the Logical Omniscience Problem*

Ronald Fagin
Joseph Y. Halpern
Moshe Y. Vardi†

IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, California 95120-6099

email: {fagin, halpern, vardi}@almaden.ibm.com

Abstract

We introduce a new approach to dealing with the well-known *logical omniscience* problem in epistemic logic. Instead of taking possible worlds where each world is a model of classical propositional logic, we take possible worlds which are models of a nonstandard propositional logic we call *NPL*, which is somewhat related to *relevance logic*. This approach gives new insights into the logic of implicit and explicit belief considered by Levesque and Lakemeyer. In particular, we show that in a precise sense agents in the structures considered by Levesque and Lakemeyer are perfect reasoners in *NPL*.

*A preliminary version of this paper appeared in Proceedings of the Third Conference on Theoretical Aspects of Reasoning about Knowledge (ed. R. J. Parikh), 1990. This version will appear in *Artificial Intelligence*.

†Current address: Dept. of Computer Science, Rice University, P.O.Box 1892, Houston, TX 77251-1892.

1 Introduction

The standard approach to modelling knowledge, which goes back to Hintikka [Hin62], is in terms of *possible worlds*. In this approach, an agent is said to know a fact φ if φ is true in all the worlds he considers possible. As has been frequently pointed out, this approach suffers from what Hintikka termed the *logical omniscience* problem [Hin75]: agents are so intelligent that they know all the logical consequences of their knowledge. Thus, if an agent knows all of the formulas in a set Σ and if Σ logically implies the formula φ, then the agent also knows φ. In particular, they know all valid formulas (including all tautologies of standard propositional logic). Furthermore, the knowledge of an agent is closed under implication: if the agent knows φ and knows $\varphi \Rightarrow \psi$, then the agent also knows ψ. The reader should note that closure under implication is a special case of logical omniscience only if $\{\varphi, \varphi \Rightarrow \psi\}$ logically implies ψ; although this logical implication holds in standard propositional logic, it does not hold in our nonstandard propositional logic NPL that we shall introduce later.

While logical omniscience is not a problem under some conditions (this is true in particular for interpretations of knowledge that are often appropriate for analyzing distributed systems [Hal87] and certain AI systems [RK86]), it is certainly not appropriate to the extent that we want to model resource-bounded agents. A number of different semantics for knowledge have been proposed to get around this problem. The one most relevant to our discussion here is what has been called the *impossible-worlds* approach. In this approach, the standard possible worlds are augmented by "impossible worlds" (or, perhaps better, *nonstandard worlds*), where the customary rules of logic do not hold [Cre72, Cre73, Lev84, Ran82, Wan90]. It is still the case that an agent knows a fact φ if φ is true in all the worlds the agent considers possible, but since the agent may in fact consider some nonstandard worlds possible, this will affect what he knows.

What about logical omniscience? Although notions like "validity" and "logical consequence" (which played a prominent part in our informal description of logical omniscience) may seem absolute, they are not; their formal definitions depend on how truth is defined and on the class of worlds being considered. Although there are nonstandard worlds in the impossible-worlds approach, validity and logical consequence are taken with respect to only the standard worlds, where all the rules of standard logic hold. For example, a formula is valid exactly if it is true in all the standard worlds in every structure. The intuition here is that the nonstandard worlds serve only as epistemic alternatives; although an agent may be muddled and may consider a nonstandard world possible, we (the logicians who get to examine the situation from the outside) know that the "real world" must obey the laws of standard logic. If we consider validity and logical implication with respect to standard worlds, then it is easy to show that logical omniscience fails in "impossible-worlds" structures: an agent does not know all valid formulas, nor does he know all the logical consequences of his knowledge here (since, in deciding what the agent knows, we must take the nonstandard worlds into account).

In this paper we consider an approach which, while somewhat related to the impossible-worlds approach, stems from a different philosophy. We consider the implications of basing a logic of knowledge on a nonstandard logic rather than on standard propositional logic. The basic motivation is the observation, implicit in [Lev84] and commented on in [FH88, Var86], that if we weaken the "logical" in "logical omniscience", then perhaps we can diminish the acuteness of the logical omniscience problem. Thus, instead of distinguishing between standard and nonstandard worlds, we take all our worlds to be models of a nonstandard logic. Some worlds in a structure may indeed be models of standard logic, but they do not have any special status for us. We consider all worlds when defining validity and logical consequence; we accept the commitment to nonstandard logic. Knowledge is still defined to be truth in all worlds the agent considers possible. It thus turns out that we still have the logical omniscience problem, but this time with respect to nonstandard logic. The hope is that the logical omniscience problem can be alleviated by appropriately choosing the nonstandard logic.

There are numerous well-known nonstandard propositional logics, including intuitionistic propositional logic [Hey56], relevance logic [AB75], and the 4-valued logic in [Bel77a, Bel77b, Dun76]. We shall give our own approach in this paper, which is closely related to relevance logic and to 4-valued logic. For each of these nonstandard logics, the starting point is the observation that there are a number of properties of implication in standard logic that seem inappropriate in certain contexts. In particular, consider a formula such as $(p \wedge \neg p) \Rightarrow q$. In standard logic this is valid; that is, from a contradiction one can deduce anything. However, consider a knowledge base into which users enter data from time to time. As Belnap points out [Bel77b], it is almost certainly the case that in a large knowledge base, there will be some inconsistencies. One can imagine that at some point a user entered the fact that Bob's salary is \$50,000, while at another point, perhaps a different user entered the fact that Bob's salary is \$60,000. Thus, in standard logic anything can be inferred from this contradiction. One solution to this problem is to replace standard worlds by worlds (called *situations* in [Lev84, Lak87], and *set-ups* in [RR72, Bel77b]) in which it is possible that a primitive proposition p is true, false, both true and false, or neither true and false. We achieve the same effect here by keeping our worlds seemingly standard and by using a device introduced in [RR72, RM73] to decouple the semantics of a formula and its negation: for every world s there is a related world s^*. A formula $\neg\varphi$ is true in s iff φ is not true in s^*. It is thus possible for both φ and $\neg\varphi$ to be true at s, and for neither to be true. (The standard worlds are now the ones where $s = s^*$; all the laws of standard propositional logic do indeed hold in such worlds.)

We call the propositional logic that results from the above semantics *nonstandard propositional logic* (NPL). Unlike standard logic, for which φ logically implies ψ exactly when $\varphi \Rightarrow \psi$ is valid, where $\varphi \Rightarrow \psi$ is defined as $\neg\varphi \vee \psi$, this is not the case in NPL. This leads us to include a connective $\hookrightarrow$ ("strong implication") in NPL so that, among other things, we have that φ logically implies ψ iff $\varphi \hookrightarrow \psi$ is valid. Of course, $\hookrightarrow$ agrees with $\Rightarrow$ on the standard worlds, but in general it is different. Given our nonstandard semantics, $\varphi \hookrightarrow \psi$ comes closer than $\varphi \Rightarrow \psi$ to capturing the idea that "if φ is true, then

ψ is true." Just as in relevance logic, formulas such as $(p \wedge \neg p) \hookrightarrow q$ are not valid, so that from a contradiction, one cannot conclude everything. In fact, we can show that if φ and ψ are *standard* propositional formulas (those formed from $\neg$ and $\wedge$, containing no occurrences of $\hookrightarrow$), then $\varphi \hookrightarrow \psi$ is valid exactly if φ entails ψ in the relevance logic **R** [RR72, RM73]. In formulas with nested occurrences of $\hookrightarrow$, however, the semantics of $\hookrightarrow$ is quite different from the relevance logic notion of entailment.

We are most interested in applying our nonstandard semantics to knowledge. It turns out that although agents in our logic are not perfect reasoners as far as standard logic goes, they *are* perfect reasoners in nonstandard logic. In particular, as we show, the complete axiomatization for the standard possible-worlds interpretation of knowledge can be converted to a complete axiomatization for the nonstandard possible-world interpretation of knowledge essentially by replacing the inference rules for standard propositional logic by inference rules for NPL. We need, however, to use $\hookrightarrow$ rather $\Rightarrow$ in formulating the axioms of knowledge. Thus, the *distribution axiom*, valid in the standard possible-worlds interpretation, says $(K_i\varphi \wedge K_i(\varphi \Rightarrow \psi)) \Rightarrow K_i\psi$. This says that an agent's knowledge is closed under logical consequence: if the agent knows φ and knows that φ implies ψ, then he also knows ψ. The analogue for this axiom holds in our nonstandard interpretation, once we replace $\Rightarrow$ by $\hookrightarrow$. This is appropriate since it is $\hookrightarrow$ that captures the intuitive notion of implication in our framework. The other basic property of knowledge (knowledge generalization) remains unchanged: if φ is valid, then so is $K_i\varphi$. That is, the agents know every valid formula (although the set of valid formulas are distinct for the standard logic and for our nonstandard logic). Thus, the basic properties of knowledge (closure under logical consequence, and knowledge of valid formulas) remain unchanged; in some sense, we have decoupled the properties of the underlying propositional logic, which change drastically, from the properties of knowledge, which remain essentially the same.

Our approach has an additional nice payoff: we show that in a certain important application we can obtain a polynomial-time algorithm for reasoning about knowledge. By contrast, under the standard approach, the complexity of such reasoning in that application is co-*NP*-complete.

It is instructive to compare our approach with that of Levesque and Lakemeyer [Lev84, Lak87]. Our semantics is essentially equivalent to theirs. But while they avoid logical omniscience by giving nonstandard worlds a secondary status and defining validity only with respect to standard worlds, we accept logical omniscience, albeit with respect to nonstandard logic. Thus, our results justify and elaborate a remark made in [FH88, Var86] that agents in Levesque's model are perfect reasoners in relevance logic.

The rest of this paper is organized as follows. In Section 2, we introduce our nonstandard propositional logic, and investigate some of its properties. In Section 3, we review the standard possible-worlds approach. In Section 4, we give our nonstandard approach to possible worlds. In Section 5, we add strong implication (the propositional connective $\hookrightarrow$) to our syntax, and thereby obtain our full nonstandard propositional logic NPL. In Section 6, we give a sound and complete axiomatization for NPL, and give a sound and complete axiomatization for the logic of knowledge using NPL as a basis rather than classical propositional logic. We also show that the validity problem for NPL is co-*NP*-complete, just as for standard propositional logic, and the validity problem for our nonstandard logic of knowledge is *PSPACE*-complete, just as for the standard logic of knowledge. In Section 7, we give the payoff we promised, of a polynomial-time algorithm for querying a knowledge base in certain natural cases. We relate our results to those in the impossible-worlds approach in Section 8. Levesque and Lakemeyer's formalism is compared with ours in Section 9. We give our conclusions in Section 10.

2 A Nonstandard Propositional Logic

Although by now it is fairly well entrenched, standard propositional logic has several undesirable and counterintuitive properties. When we are first introduced to propositional logic in school, we are perhaps somewhat uncomfortable when we learn that "$\varphi \Rightarrow \psi$" is taken to be simply an abbreviation for $\neg\varphi \vee \psi$. Why should the fact that either $\neg\varphi$ is true or ψ is true correspond to "if φ is true then ψ is true"?

Another problem with standard propositional logic is that it is fragile: a false statement implies everything. For example, the formula $(p \wedge \neg p) \Rightarrow q$ is valid, even if p and q are unrelated. As we observed in the introduction, one situation where this could be a serious problem occurs when we have a large knowledge base of many facts, obtained from multiple sources, and where a theorem prover is used to derive various conclusions from this knowledge base.

To deal with these problems, many alternatives to standard propositional logic have been proposed. We focus on one particular alternative here, and consider its consequences.

The idea is to allow formulas φ and $\neg\varphi$ to have "independent" truth values. Thus, rather than requiring that $\neg\varphi$ be true iff φ is not true, we wish instead to allow the possibility that $\neg\varphi$ can be either true or false, regardless of whether φ is true or false. Intuitively, the truth of formulas can be thought of as being determined by some knowledge base. We can think of φ being true as meaning that the fact φ has been put into a knowledge base of true formulas, and we can think of $\neg\varphi$ being true as meaning that the fact φ has been put into a knowledge base of false formulas. Since it is possible for φ to have been put in both knowledge bases, it is possible for both φ and $\neg\varphi$ to be true. Similarly, if φ had not been put into either knowledge base, then neither φ nor $\neg\varphi$ would be true.

There are several ways to capture this intuition formally (see [Dun86]). We now discuss one approach, due to [RR72, RM73]. For each world s, there is an *adjunct* world s^*, which will be used for giving semantics to negated formulas. Instead of defining $\neg\varphi$ to hold at s iff φ does not hold at s, we instead define $\neg\varphi$ to hold at s iff φ does not hold at s^*. Note that if $s = s^*$, then this gives our usual notion of negation. Very roughly, we can think of a state s is as consisting of a pair $\langle B_T, B_F\rangle$ of knowledge bases; B_T is the knowledge base of true facts, while B_F is the knowledge base of false facts. The state s^* should be thought as the adjunct pair $\langle \overline{B_F}, \overline{B_T} \rangle$, where $\overline{B_T}$ is the complement of B_T, and $\overline{B_F}$ is the complement of B_F. Continuing this intuition, to see if φ is true at s, we consult B_T; to see if $\neg\varphi$ is true at s, i.e., if φ is false at s, we consult B_F. Notice that $\varphi \in B_F$ iff $\varphi \notin \overline{B_F}$. Since $\overline{B_F}$ is the knowledge base of true facts at s^*, we have an alternate way of checking if φ is false at s: we can check if φ is not true at s^*.

Notice that under this interpretation, not only is s^* is the adjunct state of s, but s is the adjunct state of s^*; i.e., $s^{**} = s$ (where $s^{**} = (s^*)^*$). To support this intuitive view of s as a pair of knowledge bases and s^* as its adjunct, we make this a general requirement in our framework.

We define the formulas of the propositional logic by starting with a set Φ of primitive propositions that describe basic facts about the domain of discourse, and forming more complicated formulas by closing off under the Boolean connectives $\neg$ and $\wedge$. Thus, if φ and ψ are formulas, then so are $\neg\varphi$ and $\varphi \wedge \psi$. When we deal only with propositional formulas, we can identify a world with a classical truth assignment to the primitive propositions, and we can decide the truth of a propositional formula at a world s by considering only s and s^*. Thus, we define an *NPL structure* to consist of an ordered pair (s, t) of classical truth assignments to the set Φ of primitive propositions. We take * to be a function that maps a truth assignment in an NPL structure to the other truth assignment in that structure. Thus, if $S = (s, t)$, then $s^* = t$ and $t^* = s$. Truth is defined relative to a pair (S, u), where S is an NPL structure and u is one of the truth assignments in S. We define $\neg\varphi$ to be true at (S, u) if φ is not true at u^*; thus, we use the other truth assignment in order to define negation. More formally, given an NPL structure $S = (s, t)$, and $u \in \{s, t\}$, we define the semantics as follows:

- $(S, u) \models p$ iff $u(p) = \textbf{true}$ for a primitive proposition p
- $(S, u) \models \varphi \wedge \psi$ iff $(S, u) \models \varphi$ and $(S, u) \models \psi$
- $(S, u) \models \neg\varphi$ iff $(S, u^*) \not\models \varphi$

We call the logic defined so far NPL$^-$. Later, we shall add strong implication ($\hookrightarrow$) to get NPL.

Note that if $S = (s, s)$ for some truth assignment s (that is, $s = s^*$), then $(S, s) \models \neg\varphi$ iff $(S, s) \not\models \varphi$. Hence, in this case, for every propositional formula φ, we have that $(S, s) \models \varphi$ precisely if φ is true under the truth assignment s, and so we are back to standard propositional logic. Note also that in the general case, it is possible for neither φ nor $\neg\varphi$ to be true at u (if $(S, u) \not\models \varphi$ and $(S, u^*) \models \varphi$) and for both φ and $\neg\varphi$ to be true at u (if $(S, u) \models \varphi$ and $(S, u^*) \not\models \varphi$).

This approach is equivalent to Belnap's 4-valued logic [Bel77a, Bel77b], in which he has four truth values: **True**, **False**, **Both**, and **None**. Belnap's approach avoids the use of the * to define negation. The reason we make use of * is so that we can treat negation in a uniform manner. For example, later on we shall extend to an epistemic logic, and the use of * decouples the semantics of $K_i\varphi$ and $\neg K_i\varphi$. By contrast, in order to extend Levesque's propositional logic in [Lev84] to an epistemic logic where the semantics of $K_i\varphi$ and $\neg K_i\varphi$ are decoupled, Lakemeyer [Lak87] finds it necessary to introduce two possibility relations, $\mathcal{K}_i^+$ and $\mathcal{K}_i^-$. As we shall discuss in Section 9, the truth of a formula $K_i\varphi$ is determined by the possibility relation $\mathcal{K}_i^+$, while the truth of $\neg K_i\varphi$ is determined by the possibility relation $\mathcal{K}_i^-$. By using *, we need only one possibility relation $\mathcal{K}_i$ for agent i, not two. Furthermore, when we add a new connective to the language, as we do later when we add strong implication ($\hookrightarrow$), it may not be clear how to define the negation (for a formula $\neg(\varphi_1 \hookrightarrow \varphi_2)$) in a natural manner that decouples its semantics from that of $\varphi_1 \hookrightarrow \varphi_2$. This is done automatically for us by the use of *.

Just as in standard propositional logic, we take $\varphi_1 \vee \varphi_2$ to be an abbreviation for $\neg(\neg\varphi_1 \wedge \neg\varphi_2)$, and $\varphi_1 \Rightarrow \varphi_2$ to be an abbreviation for $\neg\varphi_1 \vee \varphi_2$. Since the semantics of negation is now nonstandard, it is not *a priori* clear how the propositional connectives behave in our nonstandard semantics. For example, while $\varphi_1 \wedge \varphi_2$ holds by definition precisely when φ_1 and φ_2 both hold, it is not clear that $\varphi_1 \vee \varphi_2$ holds precisely when at least one of φ_1 or φ_2 holds. It is even less clear how negation will interact in our nonstandard semantics with conjunction and disjunction.

The next proposition shows that even though we have decoupled the semantics for φ and $\neg\varphi$, the propositional connectives $\neg$, $\wedge$, and $\vee$ still behave and interact in a fairly standard way.

Proposition 2.1:

1. $(S,u) \models \neg\neg\varphi$ *iff* $(S,u) \models \varphi$.
2. $(S,u) \models \varphi_1 \vee \varphi_2$ *iff* $(S,u) \models \varphi_1$ *or* $(S,u) \models \varphi_2$.
3. $(S,u) \models \neg(\varphi_1 \wedge \varphi_2)$ *iff* $(S,u) \models \neg\varphi_1 \vee \neg\varphi_2$.
4. $(S,u) \models \neg(\varphi_1 \vee \varphi_2)$ *iff* $(S,u) \models \neg\varphi_1 \wedge \neg\varphi_2$.
5. $(S,u) \models \varphi \wedge (\psi_1 \vee \psi_2)$ *iff* $(S,u) \models (\varphi \wedge \psi_1)$ *or* $(S,u) \models (\varphi \wedge \psi_2)$.
6. $(S,u) \models \varphi \vee (\psi_1 \wedge \psi_2)$ *iff* $(S,u) \models (\varphi \vee \psi_1)$ *and* $(S,u) \models (\varphi \vee \psi_2)$.

We prove only (1) and (2), since the proofs of the rest are similar.

$$\begin{array}{rll} (S,u) \models \neg\neg\varphi & \text{iff} & (S,u^*) \not\models \neg\varphi \\ & \text{iff} & (S,u^{**}) \models \varphi \\ & \text{iff} & (S,u) \models \varphi. \end{array}$$

As for (2),

$$\begin{aligned}
(S,u) \models \varphi_1 \vee \varphi_2 \quad & \text{iff} \quad (S,u) \models \neg(\neg\varphi_1 \wedge \neg\varphi_2) \\
& \text{iff} \quad (S,u^*) \not\models \neg\varphi_1 \wedge \neg\varphi_2 \\
& \text{iff} \quad (S,u^*) \not\models \neg\varphi_1 \text{ or } (S,u^*) \not\models \neg\varphi_2 \\
& \text{iff} \quad (S,u^{**}) \models \varphi_1 \text{ or } (S,u^{**}) \models \varphi_2 \\
& \text{iff} \quad (S,u) \models \varphi_1 \text{ or } (S,u) \models \varphi_2. \ \blacksquare
\end{aligned}$$

In contrast to the behavior of $\neg$, $\wedge$, and $\vee$, the connective $\Rightarrow$ behaves rather peculiarly, since $(S,u) \models \varphi_1 \Rightarrow \varphi_2$ holds precisely when $(S,u^*) \models \varphi_1$ implies that $(S,u) \models \varphi_2$. We will come back to the issue of the definition of implication later.

Validity and logical implication are defined in the usual way: φ is valid if it holds at every (S,u), and φ logically implies ψ if $(S,u) \models \varphi$ implies $(S,u) \models \psi$ for every (S,u). What are the valid formulas? The formula $(p \wedge \neg p) \Rightarrow q$, which wreaked havoc in deriving consequences from a knowledge base, is no longer valid. What about even simpler tautologies of standard propositional logic, such as $\neg p \vee p$? This formula, too, is not valid. How about $p \Rightarrow p$? It is not valid either, since $p \Rightarrow p$ is just an abbreviation for $\neg p \vee p$, which, as we just said, is not valid. In fact, no formula is valid!

Theorem 2.2: *No formula of NPL$^-$ is valid.*

Proof: This follows from a stronger result (Theorem 4.2) that we shall prove in Section 4. $\blacksquare$

Thus, the validity problem is *very* easy: the answer is always "No, the formula is not valid!" Thus, the notion of validity is trivially uninteresting here. In contrast, there are many nontrivial logical implications; for example, as we see from Proposition 2.1, $\neg\neg\varphi$ logically implies φ, and $\neg(\varphi_1 \wedge \varphi_2)$ logically implies $\neg\varphi_1 \vee \neg\varphi_2$.

The reader may be puzzled why Proposition 2.1 does not provide us some tautologies. For example, Proposition 2.1 tells us that $\neg\neg\varphi$ logically implies φ. Doesn't this mean that $\neg\neg\varphi \Rightarrow \varphi$ is a tautology? This does not follow. In classical propositional logic, φ logically implies ψ iff the formula $\varphi \Rightarrow \psi$ is valid. This is not the case in NPL. For example, φ logically implies φ, yet $\varphi \Rightarrow \varphi$ (i.e., $\neg\varphi \vee \varphi$) is not valid in NPL. In Section 5, we define a new connective that allows us to express logical implication *in the language*, just as $\Rightarrow$ does for classical logic. We close this section by characterizing the complexity of deciding logical implication in NPL$^-$.

Theorem 2.3: *The logical implication problem in NPL$^-$ is co-NP-complete.*

The proof of this theorem will appear in Section 8, when we have developed some more machinery. This theorem says that logical implication in NPL$^-$ is as hard as logical implication in standard propositional logic, that is, co-*NP*-complete. We shall see in Theorem 4.3 that a similar phenomenon takes place for knowledge formulas.

3 Standard Possible Worlds

We review in this section the standard possible-worlds approach to knowledge. The intuitive idea behind the possible-worlds model is that besides the true state of affairs, there are a number of other possible states of affairs or "worlds". Given his current information, an agent may not be able to tell which of a number of possible worlds describes the actual state of affairs. An agent is then said to *know* a fact φ if φ is true at all the worlds he considers possible (given his current information).

The notion of possible worlds is formalized by means of Kripke structures. Suppose that we have n agents, named $1, \ldots, n$, and a set Φ of primitive propositions. A *standard Kripke structure* M for n agents over Φ is a tuple $(S, \pi, \mathcal{K}_1, \ldots, \mathcal{K}_n)$, where S is a set of *worlds*, π associates with each world in S a truth assignment to the primitive propositions of Φ (i.e., $\pi(s) : \Phi \rightarrow \{\mathbf{true}, \mathbf{false}\}$ for each world $s \in S$), and $\mathcal{K}_i$ is a *binary relation* on S, called a *possibility* relation. We refer to standard Kripke structures as *standard structures* or simply as *structures*.

Intuitively, the truth assignment $\pi(s)$ tells us whether p is true or false in a world w. The binary relation $\mathcal{K}_i$ is intended to capture the possibility relation according to agent i: $(s,t) \in \mathcal{K}_i$ if agent i considers world t possible, given his information in world s. The class of all structures for n agents over Φ is denoted by $\mathcal{M}_n^\Phi$. Usually, neither n nor Φ are relevant to our discusion, so we typically write $\mathcal{M}$ instead of $\mathcal{M}_n^\Phi$.

We define the formulas of the logic by starting with the primitive propositions in Φ, and form more complicated formulas by closing off under Boolean connectives $\neg$ and $\wedge$ and the modalities $K_1, \ldots, K_n$. Thus, if φ and ψ are formulas, then so are $\neg\varphi$, $\varphi \wedge \psi$, and $K_i\varphi$, for $i = 1, \ldots, n$. We define the connectives $\vee$ and $\Rightarrow$ to be abbreviations as before. The class of all formulas for n agents over Φ is denoted by $\mathcal{L}_n^\Phi$. Again, when n and Φ are not relevant to our discussion, we write $\mathcal{L}$ instead of $\mathcal{L}_n^\Phi$. We refer to $\mathcal{L}$-formulas as *standard formulas*.

We are now ready to assign truth values to formulas. A formula will be true or false at a world in a structure. We define the notion $(M,s) \models \varphi$, which can be read as "*φ is true at (M,s)*" or "*φ holds at (M,s)*" or "*(M,s) satisfies φ*", by induction on the structure of φ.

$(M,s) \models p$ (for a primitive proposition $p \in \Phi$) iff $\pi(s)(p) = \mathbf{true}$
$(M,s) \models \neg\varphi$ iff $(M,s) \not\models \varphi$
$(M,s) \models \varphi \wedge \psi$ iff $(M,s) \models \varphi$ and $(M,s) \models \psi$
$(M,s) \models K_i\varphi$ iff $(M,t) \models \varphi$ for all t such that $(s,t) \in \mathcal{K}_i$.

The first three clauses in this definition correspond to the standard clauses in the definition of truth for propositional logic. The last clause captures the intuition that agent i knows φ in world s of structure M exactly if φ is true at all worlds that i considers possible in s.

Given a structure $M = (S, \pi, \mathcal{K}_1, \ldots, \mathcal{K}_n)$, we say that φ is *valid in M*, and write $M \models \varphi$, if $(M,s) \models \varphi$ for every world s in S, and say that φ is *satisfiable in M* if

$(M, s) \models \varphi$ for some world s in S. We say that φ is *valid* with respect to $\mathcal{M}$, and write $\mathcal{M} \models \varphi$, if it is valid with respect to all structures of $\mathcal{M}$, and it is *satisfiable* with respect to $\mathcal{M}$ if it is satisfiable in some structure in $\mathcal{M}$. It is easy to check that a formula φ is valid in M (resp., valid with respect to $\mathcal{M}$) if and only if $\neg\varphi$ is not satisfiable in M (resp., not satisfiable with respect to $\mathcal{M}$).

To get a sound and complete axiomatization, one starts with propositional reasoning and adds to it axioms and inference rules for knowledge. By propositional reasoning we mean all substitution instances of sound propositional inference rules of propositional logic. An inference rule is a statement of the form "from Σ infer σ", where $\Sigma \cup \{\sigma\}$ is a set of formulas. (See [FHV92] for a discussion of inference rules.) Such an inference rule is sound if for every substitution τ of formulas $\varphi_1, \ldots, \varphi_k$ for the primitive propositions $p_1, \ldots, p_k$ in Σ and σ, if all the formulas in $\tau[\Sigma]$ are valid, then $\tau[\sigma]$ is also valid. Modus ponens ("from φ and $\varphi \Rightarrow \psi$ infer ψ") is an example of a sound propositional inference rule. Of course, if σ is a valid propositional formula, then "from $\emptyset$ infer σ" is a sound propositional inference rule. It is easy to show that "from Σ infer σ" is a sound propositional inference rule iff σ is a propositional consequence of Σ [FHV92], which explains why the notion of inference is often confused with the notion of consequence. As we shall see later, the two notions do not coincide in our nonstandard propositional logic NPL.

Consider the following axiom system K, which in addition to propositional reasoning consists of one axiom and one rule of inference given below:

A1. $(K_i\varphi \wedge K_i(\varphi \Rightarrow \psi)) \Rightarrow K_i\psi$ (Distribution Axiom)

PR. All sound inference rules of propositional logic

R1. From φ infer $K_i\varphi$ (Knowledge Generalization)

One should view the axioms and inference rules above as *schemes*, i.e., K actually consists of all $\mathcal{L}$-instances of the above axioms and inference rules.

Theorem 3.1: [Che80] *K is a sound and complete axiomatization for validity of $\mathcal{L}$-formulas in $\mathcal{M}$.*

We note that PR can be replaced by any complete axiomatization of standard propositional logic that includes modus ponens as an inference rule, which is the usual way that K is presented (cf. [Che80]). We chose to present K in this unusual way in anticipation of our treatment of NPL in Section 5.

Finally, instead of trying to prove validity, one may wish to check validity algorithmically.

Theorem 3.2: [Lad77] *The problem of determining validity of $\mathcal{L}$-formulas in $\mathcal{M}$ is PSPACE-complete.*

4 Nonstandard Possible Worlds

Our main goal in this paper is to help alleviate logical omniscience by defining Kripke structures that are based on a nonstandard propositional logic, rather than basing them on classical propositional logic. We shall base our nonstandard Kripke structures on our nonstandard propositional logic; in particular, we make use of the * operator of Routley and Meter [RR72, RM73].

A *nonstandard Kripke structure* is a tuple $(S, \pi, \mathcal{K}_1, ..., \mathcal{K}_n, ^*)$, where $(S, \pi, \mathcal{K}_1, ..., \mathcal{K}_n)$ is a (Kripke) structure, and where * is a unary function with domain and range the set S of worlds (where we write s^* for the result of applying the function * to the world s) such that $s^{**} = s$ for each $s \in S$. We refer to nonstandard Kripke structures as *nonstandard structures*. We call them nonstandard, since we think of a world where φ and $\neg\varphi$ are both true or both false as nonstandard. We denote the class of nonstandard structures for n agents over Φ by $\mathcal{NM}_n^\Phi$. As before, when n and Φ are not relevant to our discussion, we write $\mathcal{NM}$ instead of $\mathcal{NM}_n^\Phi$.

The definition of $\models$ for the language $\mathcal{L}$ for nonstandard structures is the same as for standard structures, except for the clause for negation:

$(M, s) \models \neg\varphi$ iff $(M, s^*) \not\models \varphi$.

In particular, the clause for K_i does not change at all:

$(M, s) \models K_i\varphi$ iff $(M, t) \models \varphi$ for all t such that $(s, t) \in \mathcal{K}_i$.

Our semantics is closely related to that of Levesque [Lev84] and Lakemeyer [Lak87]. We discuss their approach in Section 9. Unlike our approach, in their approach it is necessary to introduce two $\mathcal{K}_i$ relations for each agent i, to deal separately with the truth of formulas of the form $K_i\varphi$ and the truth of formulas of the form $\neg K_i\varphi$.

Similarly to before, we say that φ is *valid* with respect to $\mathcal{NM}$, and write $\mathcal{NM} \models \varphi$, if $(M, s) \models \varphi$ for every nonstandard structure M and every state s of M.

As we noted earlier, it is possible for neither φ nor $\neg\varphi$ to be true at world s, and for both φ and $\neg\varphi$ to be true at world s. Let us refer to a world where neither φ nor $\neg\varphi$ is true as *incomplete* (with respect to φ); otherwise, s is *complete*. The intuition behind an incomplete world is that there is not enough information to determine whether φ is true or whether $\neg\varphi$ is true. What about a world where both φ and $\neg\varphi$ are true? We call such a world *incoherent* (with respect to φ); otherwise, s is *coherent*. The intuition behind an incoherent world is that it is overdetermined: it might correspond to a situation where several people have provided mutually inconsistent information. A world s is *standard* if $s = s^*$. Note that for a standard world, the definition of the semantics of negation is equivalent to the standard definition. In particular, a standard world s is both complete and coherent: for each formula φ exactly one of φ or $\neg\varphi$ is true at s.

Remark 4.1: If we consider a fixed structure, it is possible for a world to be both complete and coherent without being standard. Nevertheless, there is an important sense in which this can be viewed as "accidental", and that the only worlds that can be complete and coherent are those that are standard. To understand this, we must work at the level of *frames* [Gol92, HC84] rather than structures. Essentially, a frame is a structure without the truth assignment π. Thus, in our present context, we define a *(nonstandard) frame* F to be a tuple $(S, \mathcal{K}_1, \ldots, \mathcal{K}_n, ^*)$, where S is a set of worlds, $\mathcal{K}_1, \ldots, \mathcal{K}_n$ are binary relations on S, and * is a unary function with domain and range the set S of worlds, such that $s^{**} = s$. We say that the nonstandard structure $(S, \pi, \mathcal{K}_1, \ldots, \mathcal{K}_n, ^*)$ is *based on* the frame $(S, \mathcal{K}_1, \ldots, \mathcal{K}_n, ^*)$. We say that a world s is *complete* (resp. *coherent*) *with respect to the frame* F if s is complete (resp. coherent) with respect to every structure based on F; the world s is *standard with respect to* F exactly if $s^* = s$. It is now easy to see that if s is complete and coherent with respect to a frame F if and only if s is standard in F. ∎

What are the properties of knowledge in nonstandard structures? One way to characterize the formal properties of a semantic model is to consider all the validities under that semantics. In our case, we should consider the formulas valid in $\mathcal{NM}$. Theorem 2.2 tells us that no formula of NPL$^-$ is valid. It turns out that even though we have enlarged the language to include knowledge modalities, it is still the case that no formula (of $\mathcal{L}$) is valid. Even more, there is a single counterexample that simultaneously shows that no formula is valid!

Theorem 4.2: *There is no formula of $\mathcal{L}$ that is valid with respect to $\mathcal{NM}$. In fact, there is a nonstandard structure M and a world s of M such that every formula of $\mathcal{L}$ is false at s, and a world t of M such that every formula of $\mathcal{L}$ is true at t.*

Proof: Let $M = (S, \pi, \mathcal{K}_1, ..., \mathcal{K}_n, ^*)$ be a special nonstandard structure, defined as follows. Let S contain only two worlds s and t, where $t = s^*$ (and so $s = t^*$). Define π by letting $\pi(s)$ be the truth assignment where $\pi(s)(p) = \mathbf{false}$ for every primitive proposition p, and letting $\pi(t)$ be the truth assignment where $\pi(t)(p) = \mathbf{true}$ for every primitive proposition p. Define each $\mathcal{K}_i$ to be $\{(s,s),(t,t)\}$. By a straightforward induction on formulas, it follows that for every formula φ of $\mathcal{L}$, we have $(M,s) \not\models \varphi$ and $(M,t) \models \varphi$. In particular, every formula of $\mathcal{L}$ is false at s, and every formula of $\mathcal{L}$ is true at t. Since every formula of $\mathcal{L}$ is false at s, no formula of $\mathcal{L}$ is valid with respect to $\mathcal{NM}$. ∎

It follows from Theorem 4.2 that we cannot use validities to characterize the properties of knowledge in nonstandard structures, since there are no validities! We will come back to this point later.

As we noted in the introduction, our basic motivation is the observation that if we weaken the "logical" in "logical omniscience", then perhaps we can diminish the acuteness of the logical omniscience problem. Logical implication is indeed weaker in nonstandard structures than in standard structures, as we now show. If φ logically implies ψ in nonstandard structures, then φ logically implies ψ in standard structures, since standard structures can be viewed as a special case of nonstandard structures. However, the converse is false, since, for example, $\{\varphi, \varphi \Rightarrow \psi\}$ logically implies ψ in standard structures but not in nonstandard structures.

Nevertheless, logical omniscience did not go away! If an agent knows all of the formulas in a set Σ, and if Σ logically implies the formula φ, then the agent also knows φ. Because, as we just showed, we have weakened the notion of logical implication, the problem of logical omniscience is not as acute as it was in the standard approach. For example, knowledge of valid formulas, which is one form of omniscience, is completely innocuous here, since there are no valid formulas. Also, an agent's knowledge need not be closed under implication; an agent may know φ and $\varphi \Rightarrow \psi$ without knowing ψ, since, as we noted above, φ and $\varphi \Rightarrow \psi$ do not logically imply ψ with respect to nonstandard structures.

We saw that the problem of determining validity is easy (since the answer is always "No"). Validity is a special case of logical implication: a formula is valid iff it is a logical consequence of the empty set. Unfortunately, logical implication is not that easy to determine.

Theorem 4.3: *The logical implication problem for $\mathcal{L}$-formulas in nonstandard structures is PSPACE-complete.*

As with Theorem 2.3, the proof of this theorem will appear in Section 8, when we have developed some more machinery.

Theorem 4.3 asserts that nonstandard logical implication for knowledge formulas (i.e., $\mathcal{L}$-formulas) is as hard as standard logical implication for knowledge formulas, that is, *PSPACE*-complete. This is analogous to Theorem 2.3, where the same phenomenon takes place for propositional formulas.

We saw in Theorem 4.2 that there are no valid formulas. In particular, we cannot capture properties of knowledge by considering all of the formulas that are valid, since there are none. By contrast, Theorem 4.3 tells us that the structure of logical implication is quite rich (since the logical implication problem is *PSPACE*-complete). In classical logic, we can capture logical implication in the language by using $\Rightarrow$: thus, φ logically implies ψ precisely if the formula $\varphi \Rightarrow \psi$ is valid. In the next section, we enrich our language by adding a new propositional connective $\hookrightarrow$, with which it is possible to express logical implication in the language.

5 Strong Implication

In Section 2 we introduced a nonstandard propositional logic, motivated by our discomfort with certain classic tautologies, such as $(p \land \neg p) \Rightarrow q$, and—lo and behold!—under this semantics these formulas are no longer valid. Unfortunately, the bad news is that other formulas, such as $\varphi \Rightarrow \varphi$, that blatantly seem as if they should be valid, are not valid either under this approach. In fact, no formula is valid in the nonstandard approach! It seems that we have thrown out the baby with the bath water. In particular, we could not characterize the properties of knowledge in the nonstandard approach by considering validities, because there are no validities.

To get better insight into this problem, let us look more closely at why the formula $\varphi \Rightarrow \varphi$ is not valid. Our intuition about implication tells us that $\varphi_1 \Rightarrow \varphi_2$ should say "if φ_1 is true, then φ_2 is true". However, $\varphi_1 \Rightarrow \varphi_2$ is defined to be $\neg\varphi_1 \vee \varphi_2$. In standard propositional logic, this is the same as "if φ_1 is true, then φ_2 is true". However, in nonstandard structures, these are not equivalent. Thus, the problem is not with our semantics, but rather with the definition of $\Rightarrow$. This motivates the definition of a new propositional connective $\hookrightarrow$, which we call *strong implication*, where $\varphi_1 \hookrightarrow \varphi_2$ is defined to be true if whenever φ_1 is true, then φ_2 is true. Formally, in the pure propositional case where $S = (s, t)$ is an NPL structure and $u \in \{s, t\}$, we define

$$(S, u) \models \varphi_1 \hookrightarrow \varphi_2 \text{ iff } (S, u) \models \varphi_2 \text{ whenever } (S, u) \models \varphi_1.$$

That is, $(S, u) \models \varphi_1 \hookrightarrow \varphi_2$ iff either $(S, u) \not\models \varphi_1$ or $(S, u) \models \varphi_2$. Similarly, if M is a nonstandard structure and s is a world of M, then

$$(M, s) \models \varphi_1 \hookrightarrow \varphi_2 \text{ iff (if } (M, s) \models \varphi_1\text{, then } (M, s) \models \varphi_2).$$

Equivalently, $(M, s) \models \varphi_1 \hookrightarrow \varphi_2$ iff either $(M, s) \not\models \varphi_1$ or $(M, s) \models \varphi_2$.

In the pure propositional case, we refer to this logic as *nonstandard propositional logic*, or *NPL*. In the case of knowledge formulas, we denote by $\mathcal{L}_n^{\Phi,\hookrightarrow}$, or $\mathcal{L}^{\hookrightarrow}$ for short, the set of formulas obtained by modifying the definition of $\mathcal{L}_n^{\Phi}$ by adding $\hookrightarrow$ as a new propositional connective.

Strong implication is indeed a new connective, that is, it cannot be defined using $\neg$ and $\wedge$. For, there are no valid formulas using only $\neg$ and $\wedge$, whereas by using $\hookrightarrow$, there are validities: $\varphi \hookrightarrow \varphi$ is an example, as is $\varphi_1 \hookrightarrow (\varphi_1 \vee \varphi_2)$.

The next proposition shows a sense in which strong implication is indeed stronger than implication.

Proposition 5.1: *Let φ_1 and φ_2 be formulas in $\mathcal{L}$. If $\varphi_1 \hookrightarrow \varphi_2$ is valid with respect to nonstandard Kripke structures, then $\varphi_1 \Rightarrow \varphi_2$ is valid with respect to standard Kripke structures. However, the converse is false.*

Proof: Assume that $\varphi_1 \hookrightarrow \varphi_2$ is valid with respect to nonstandard Kripke structures. As we remarked after the proof of Theorem 4.2, a standard Kripke structure can be viewed as a special case of a nonstandard Kripke structure. Hence, $\varphi_1 \hookrightarrow \varphi_2$ is valid with respect to standard Kripke structures. In a standard Kripke structure, $\varphi_1 \hookrightarrow \varphi_2$ is equivalent to $\varphi_1 \Rightarrow \varphi_2$. So $\varphi_1 \Rightarrow \varphi_2$ is valid with respect to standard Kripke structures.

The converse is false, since the formula $(p \wedge \neg p) \Rightarrow q$ is valid in standard propositional logic, whereas the formula $(p \wedge \neg p) \hookrightarrow q$ is not valid in NPL. ∎

As we promised earlier, we can now express logical implication in $\mathcal{L}^{\hookrightarrow}$, using $\hookrightarrow$, just as we can express logical implication in standard structures, using $\Rightarrow$. The following proposition is almost immediate.

Proposition 5.2: *Let φ_1 and φ_2 be formulas in $\mathcal{L}^{\hookrightarrow}$. Then φ_1 logically implies φ_2 in nonstandard structures iff $\varphi_1 \hookrightarrow \varphi_2$ is valid with respect to nonstandard structures.*

The connective $\hookrightarrow$ is somewhat related to the connective $\rightarrow$ of relevance logic, which is meant to capture the notion of *relevant entailment*. A formula of the form $\varphi_1 \rightarrow \varphi_2$, where φ_1 and φ_2 are standard propositional formulas, is called a *first-degree entailment*. (See [Dun86] for an axiomatization of first-degree entailments.) It is not hard to show that if φ_1 and φ_2 are standard propositional formulas (and so have no occurrence of $\hookrightarrow$), then $\varphi_1 \rightarrow \varphi_2$ is a theorem of the relevance logic **R** [RR72, RM73] exactly if $\varphi_1 \hookrightarrow \varphi_2$ is valid in NPL (or equivalently, φ_1 logically implies φ_2 in NPL$^-$). So $\varphi_1 \rightarrow \varphi_2$ can be viewed as saying that $\varphi_1 \hookrightarrow \varphi_2$ is valid. In formulas with nested occurrences of $\hookrightarrow$, however, the semantics of $\hookrightarrow$ is quite different from that of relevant entailment. In particular, while $p \hookrightarrow (q \hookrightarrow p)$ is valid in NPL, the analogous formula $p \rightarrow (q \rightarrow p)$ is not a theorem of relevance logic [AB75].

With $\hookrightarrow$, we greatly increase the expressive power of our language. For example, in $\mathcal{L}$ (the language without $\hookrightarrow$), we cannot say that a formula φ is false. That is, there is no formula ψ such that $(M,t) \models \psi$ iff $(M,t) \not\models \varphi$. For suppose that there were such a formula ψ. Let M and t be as in Theorem 4.2. Then $(M,t) \not\models \psi$ and $(M,t) \not\models \varphi$, a contradiction. What about the formula $\neg\varphi$? This formula says that $\neg\varphi$ is true, but does not say that φ is false. However, once we move to $\mathcal{L}^{\hookrightarrow}$, it is possible to say that a formula is false, as we shall see in the next proposition. In order to state this and other results, it turns out to be convenient to have an abbreviation for the proposition *false* (which is false at every world). The way we abbreviate *false* depends on the context. When dealing with the standard semantics in the language $\mathcal{L}$, we take *true* to be an abbreviation for some fixed standard tautology such as $p \Rightarrow p$. When dealing with the nonstandard semantics in the language $\mathcal{L}^{\hookrightarrow}$, we take *true* to be an abbreviation for some fixed nonstandard tautology such as $p \hookrightarrow p$. In both cases, we abbreviate $\neg$*true* by *false*. In fact, it will be convenient to think of $true$ and $false$ as constants in the language (rather than as abbreviations) with the obvious semantics. The next proposition, which shows how to say that a formula is false, is straightforward.

Proposition 5.3: *Let M be a nonstandard structure, and let s be a world of M. Then $(M,s) \not\models \varphi$ iff $(M,s) \models \varphi \hookrightarrow$ false.*

This proposition enables us to embed standard propositional logic into NPL, by replacing $\neg\varphi$ by $\varphi \hookrightarrow$ *false*. We shall make use of this technique in the next section, when we give a sound and complete axiomatization for NPL, and analyze the complexity of the validity problem.

6 Axiomatizations and Complexity

In this section, we provide sound and complete axiomatizations for our nonstandard propositional logic NPL, and for our nonstandard epistemic logic, and prove their correctness. We also show that the validity problem for NPL is co-*NP*-complete, just as for standard propositional logic, and the validity problem for our nonstandard logic of knowledge is *PSPACE*-complete, just as for the standard logic of knowledge.

6.1 A sound and complete axiomatization for NPL

In this subsection we give an axiomatization for NPL and prove that it is sound and complete. We also show that the validity problem is co-*NP*-complete, just as for propositional logic, and discuss an interesting new inference rule. For the purposes of this subsection only, it is convenient to enrich our standard propositional language so that $\Rightarrow$ and *false* are first-class objects, and not just abbreviations. Thus, let $\mathcal{L}_1$ contain all formulas built up out of *false* and the primitive propositions in Φ, by closing off under the Boolean connectives $\neg$, $\wedge$, and $\Rightarrow$. Let $\mathcal{L}_1^+$ be the negation-free formulas in $\mathcal{L}_1$ (those built up out of *false* and the primitive propositions in Φ, by closing off under the Boolean connectives $\wedge$ and $\Rightarrow$). We define $\mathcal{L}_2$ and $\mathcal{L}_2^+$ identically, but using $\hookrightarrow$ instead of $\Rightarrow$.

As a tool in developing an axiomatization for NPL, and motivated by Proposition 5.3, we explore the relationship between the standard and nonstandard semantics. This will make it possible to use (in part) the standard axiomatization. If $\varphi \in \mathcal{L}_1$, then we define the formula $\varphi^{nst} \in \mathcal{L}_2^+$ by recursively replacing in φ all subformulas of the form $\neg\psi$ by $\psi \hookrightarrow false$ and all occurrences of $\Rightarrow$ by $\hookrightarrow$ (the superscript ^{nst} stands for *n*on*st*andard). Note that φ^{nst} is negation-free. We also define what is essentially the inverse transformation: if $\varphi \in \mathcal{L}_2^+$, let $\varphi^{st} \in \mathcal{L}_1^+$ be the result of replacing in φ all occurrences of $\hookrightarrow$ by $\Rightarrow$. It is easy to see that the transformations ^{nst} and st are inverses when restricted to negation-free formulas. In particular:

Lemma 6.1: *If $\varphi \in \mathcal{L}_2^+$, then $(\varphi^{st})^{nst} = \varphi$.*

If s is a truth assignment, and $\varphi \in \mathcal{L}_1$, then we write $s \models \varphi$ if φ is true under the truth assignment s.

Proposition 6.2: *Assume that $S = (s,t)$ is an NPL structure, $u \in \{s,t\}$, and $\varphi \in \mathcal{L}_1$. Then $(S,u) \models \varphi^{nst}$ iff $u \models \varphi$.*

Proof: We prove this proposition by induction on the structure of φ. The result is immediate if φ is *false*, a primitive proposition, or of the form $\varphi_1 \wedge \varphi_2$. If φ is $\neg\psi$, then

$$\begin{aligned}(S,u) \models \varphi^{nst} \quad &\text{iff} \quad (S,u) \models \psi^{nst} \hookrightarrow false \\ &\text{iff} \quad (S,u) \not\models \psi^{nst}\end{aligned}$$

$$
\begin{array}{ll}
\text{iff} & \text{(by induction hypothesis) } u \not\models \psi \\
\text{iff} & u \models \varphi.
\end{array}
$$

If φ is $\psi_1 \Rightarrow \psi_2$, then

$$
\begin{array}{lll}
(S,u) \models \varphi^{nst} & \text{iff} & (S,u) \models (\psi_1)^{nst} \hookrightarrow (\psi_2)^{nst} \\
& \text{iff} & (S,u) \not\models (\psi_1)^{nst} \text{ or } (S,u) \models (\psi_2)^{nst} \\
& \text{iff} & \text{(by induction hypothesis) } u \not\models \psi_1 \text{ or } u \models \psi_2 \\
& \text{iff} & u \models \varphi. \blacksquare
\end{array}
$$

Proposition 6.2 tells us that ^{nst} gives an embedding of standard propositional logic into NPL. The following corollary is immediate. If $\varphi \in \mathcal{L}_1$, then when we say that φ is a standard propositional tautology, we mean that $s \models \varphi$ for every truth assignment s.

Corollary 6.3: *Assume that $\varphi \in \mathcal{L}_1$. Then φ^{nst} is valid in NPL iff φ is a standard propositional tautology.*

In particular, it follows from Corollary 6.3 that the validity problem for NPL is at least as hard as that of propositional logic, namely, co-*NP*-complete. In fact, this is precisely the complexity.

Theorem 6.4: *The validity problem for NPL formulas is co-NP-complete.*

Proof: The lower bound is immediate from Corollary 6.3. The upper bound follows from the fact that to determine if an NPL formula σ is not valid, we can simply guess an NPL structure $S = (s,t)$ and $u \in \{s,t\}$, and verify that $(S,u) \not\models \sigma$. $\blacksquare$

Another connection between standard propositional logic and NPL is due to the fact that negated propositions in NPL behave in some sense as "independent" propositions. We say that a formula $\varphi \in \mathcal{L}_2$ is *pseudo-positive* if $\neg$ occurs in φ only immediately in front of a primitive proposition. For example, the formula $p \wedge \neg p$ is pseudo-positive, while $\neg(p \wedge q)$ is not. If φ is a pseudo-positive formula, then φ^+ is obtained from φ by replacing every occurrence $\neg p$ of a negated proposition by a new proposition $\bar{p}$. Note that φ^+ is a negation-free formula.

Proposition 6.5: *Let φ be a pseudo-positive formula. Then φ is valid in NPL iff φ^+ is valid in NPL.*

Proof: We shall prove the "only if" direction, since the proof of the converse is very similar. Assume that φ is valid in NPL. Let Φ be the set of primitive propositions (so that in particular, every primitive proposition that appears in φ is in Φ), and let $\Phi' = \Phi \cup \{\bar{p} \,|\, p \in \Phi\}$. Let s and t be arbitrary truth assignments over Φ', and let

$S = (s,t)$. Take $u \in \{s,t\}$. To show that φ^+ is valid in NPL, we must show that $(S,u) \models \varphi^+$. Let s',t' be truth assignments over Φ defined by letting $s'(p) = \mathbf{true}$ iff $u(p) = \mathbf{true}$, and $t'(p) = \mathbf{true}$ iff $u(\bar{p}) = \mathbf{false}$. Let $S' = (s',t')$. Assume that $p \in \Phi$. It is easy to see that $(S',s') \models p$ iff $(S,u) \models p$, and $(S',s') \models \neg p$ iff $(S,u) \models \bar{p}$. A straightforward induction on the structure of formulas (where we take advantage of the fact that φ is pseudo-positive) then shows that $(S',s') \models \varphi$ iff $(S,u) \models \varphi^+$. But φ is valid, so $(S',s') \models \varphi$. Hence, $(S,u) \models \varphi^+$, as desired. ∎

Corollary 6.6: *Let φ be a pseudo-positive formula. Then φ is valid in NPL iff $(\varphi^+)^{st}$ is a standard propositional tautology.*

Proof: By Proposition 6.5, φ is valid in NPL iff φ^+ is valid in NPL. By Lemma 6.1, $\varphi^+ = ((\varphi^+)^{st})^{nst}$. By Corollary 6.3, $((\varphi^+)^{st})^{nst}$ is valid in NPL iff $(\varphi^+)^{st}$ is a standard propositional tautology. The result follows. ∎

We can use Corollary 6.6 to obtain an axiomatization of NPL. To prove that a propositional formula φ in $\mathcal{L}^{\hookrightarrow}$ is valid, we first drive negations down until they apply only to primitive propositions, by applying the equivalences given by the next lemma.

Lemma 6.7:

1. *$\neg\neg\varphi$ is logically equivalent to φ.*
2. *$\neg(\varphi \hookrightarrow \psi)$ is logically equivalent to $((\neg\psi \hookrightarrow \neg\varphi) \hookrightarrow false)$.*
3. *$\neg(\varphi \wedge \psi)$ is logically equivalent to $(\neg\varphi \hookrightarrow false) \hookrightarrow \neg\psi$.*

Proof: (1) is simply Proposition 2.1(1). We now show (2).

$$\begin{array}{rll} (S,u) \models \neg(\varphi \hookrightarrow \psi) & \text{iff} & (S,u^*) \not\models (\varphi \hookrightarrow \psi) \\ & \text{iff} & (S,u^*) \models \varphi \text{ and } (S,u^*) \not\models \psi \\ & \text{iff} & (S,u) \not\models \neg\varphi \text{ and } (S,u) \models \neg\psi \\ & \text{iff} & (S,u) \not\models (\neg\psi \hookrightarrow \neg\varphi) \\ & \text{iff} & (S,u) \models ((\neg\psi \hookrightarrow \neg\varphi) \hookrightarrow false). \end{array}$$

As for (3),

$$\begin{array}{rll} (S,u) \models \neg(\varphi \wedge \psi) & \text{iff} & (S,u^*) \not\models \varphi \wedge \psi \\ & \text{iff} & (S,u^*) \not\models \varphi \text{ or } (S,u^*) \not\models \psi \\ & \text{iff} & (S,u) \models \neg\varphi \text{ or } (S,u) \models \neg\psi \\ & \text{iff} & (S,u) \not\models (\neg\varphi \hookrightarrow false) \text{ or } (S,u) \models \neg\psi \\ & \text{iff} & (S,u) \models (\neg\varphi \hookrightarrow false) \hookrightarrow \neg\psi. \end{array}$$ ∎

Consider the following axiom system N, where $\psi_1 \rightleftharpoons \psi_2$ is an abbreviation of $(\psi_1 \hookrightarrow \psi_2) \wedge (\psi_2 \hookrightarrow \psi_1)$:

PL. All substitution instances of formulas φ^{nst}, where $\varphi \in \mathcal{L}_1$ is a standard propositional tautology

NPL1. $(\neg\neg\varphi) \rightleftharpoons \varphi$

NPL2. $\neg(\varphi \hookrightarrow \psi) \rightleftharpoons ((\neg\psi \hookrightarrow \neg\varphi) \hookrightarrow false)$

NPL3. $\neg(\varphi \wedge \psi) \rightleftharpoons ((\neg\varphi \hookrightarrow false) \hookrightarrow \neg\psi)$

R0. From φ and $\varphi \hookrightarrow \psi$ infer ψ (modus ponens)

Note that rule R0 is different from the standard modus ponens, in that $\hookrightarrow$ is used instead of $\Rightarrow$. When necessary, we shall refer to "from φ and $\varphi \Rightarrow \psi$ infer ψ" as *standard modus ponens*, and rule R0 as *nonstandard modus ponens*.

An example shows the importance of considering substitution instances in PL. The formula $\neg q \hookrightarrow \neg q$ is *not* of the form φ^{nst}, since every formula φ^{nst} is negation-free. However, $\neg q \hookrightarrow \neg q$ is an instance of PL, since it is the result of substituting $\neg q$ for p in the formula $p \hookrightarrow p$, which is $(p \Rightarrow p)^{nst}$.

Note that NPL1–NPL3 correspond to Lemma 6.7. As we noted, they are useful in driving negations down.

Remark 6.8: PL can be replaced by the nonstandard version of any complete axiomatization of standard propositional logic for the language $\mathcal{L}_1$. That is, assume that standard modus ponens along with axioms S1, ..., Sk give a sound and complete axiomatization of standard propositional logic for the language $\mathcal{L}_1$. We can replace PL by S1^{nst}, ..., Sk^{nst}, and get an equivalent axiomatization. This is because on the one hand, S1^{nst}, ..., Sk^{nst} are special cases of PL, so the new axiomatization is no stronger. On the other hand, let $\varphi \in \mathcal{L}_1$ be a standard propositional tautology, so that φ^{nst} is an an instance of PL. By completeness, there is a proof $\varphi_1, \ldots, \varphi_m$ of φ using S1, ..., Sk and standard modus ponens, where φ_m is φ, and where each φ_i is either an axiom (an instance of one of S1, ..., Sk) or the result of applying standard modus ponens to earlier formulas in the proof. We now show by induction on i that each φ_i^{nst} (and in particular, φ^{nst}) is provable from S1^{nst}, ..., Sk^{nst} along with nonstandard modus ponens. This is immediate if φ_i is an instance of one of S1, ..., Sk. Assume now that φ_i is the result of applying standard modus ponens to earlier formulas φ_j and $\varphi_j \Rightarrow \varphi_i$ in the proof. By induction assumption, φ_j^{nst} and $(\varphi_j \Rightarrow \varphi_i)^{nst}$ (that is, $\varphi_j^{nst} \hookrightarrow \varphi_i^{nst}$) are provable from S$1^{nst}$, ..., S$k^{nst}$ along with nonstandard modus ponens. By one more application of nonstandard modus ponens, it follows that φ_i is similarly provable. ∎

Theorem 6.9: *N is a sound and complete axiomatization for NPL.*

Proof: See Appendix A. ∎

The only inference rule in our axiom system N that we just asserted to be sound and complete is modus ponens. We now introduce a new propositional inference rule that we shall show is also sound. Of course, we do not need it for completeness (since it is not a rule of N). However, it will be useful in the next subsection, when we give a complete axiomatization for our nonstandard logic of knowledge. The new rule, which we call *negation replacement*, is:

From $\varphi \hookrightarrow false$ infer $\neg\varphi$.

Lemma 6.10: *The negation replacement rule is sound for NPL.*

Proof: Assume that $\varphi \hookrightarrow false$ is valid. Assume that $S = (s,t)$ is an arbitrary NPL structure, and $u \in \{s,t\}$. Then $(S,u^*) \models (\varphi \hookrightarrow false)$, since $\varphi \hookrightarrow false$ is valid. So $(S,u^*) \not\models \varphi$, that is $(S,u) \models \neg\varphi$. So $\neg\varphi$ is valid. ∎

We remark that by a similar argument, the converse rule "from $\neg\varphi$ infer $\varphi \hookrightarrow false$" is also sound.

For both standard propositional logic and NPL, if Σ logically implies σ, then "from Σ infer σ" is a sound inference rule. As we noted earlier, the converse is true for standard propositional logic, but not for NPL in general. For example, even though the negation replacement rule "from $\varphi \hookrightarrow false$ infer $\neg\varphi$" is sound, $\varphi \hookrightarrow false$ does not logically imply $\neg\varphi$ (since $(S,u) \models (\varphi \hookrightarrow false)$ precisely if $(S,u) \not\models \varphi$, which is not the same as $(S,u) \models \neg\varphi$). Nevertheless, it can shown that testing soundness of nonstandard inference rules has the same computational complexity as testing logical implication in NPL; they are both co-*NP*-complete [FHV92].

6.2 A sound and complete axiomatization for the logic of knowledge

In this subsection, we give a sound and complete axiomatization for our nonstandard logic of knowledge. We also show that a natural modification (where the only propositional inference rule is modus ponens) does *not* provide a sound and complete axiomatization. Finally, we show that the complexity of the validity problem is *PSPACE*-complete, just as for the standard case K.

The axiomatization that we shall show is sound and complete is obtained by modifying the axiom system K by (a) replacing propositional reasoning by nonstandard propositional reasoning, and (b) replacing standard implication ($\Rightarrow$) in the other axioms and rules by strong implication ($\hookrightarrow$). Thus, we obtain the axiom system, which we denote by $K^{\hookrightarrow}$, which consists of all instances (for the language $\mathcal{L}^{\hookrightarrow}$) of the axiom scheme and rules of inference given below:

A1$^{\hookrightarrow}$. $(K_i\varphi \wedge K_i(\varphi \hookrightarrow \psi)) \hookrightarrow K_i\psi$ (Distribution Axiom)

NPR. All sound inference rules of NPL

R1. From φ infer $K_i\varphi$ (Knowledge Generalization)

Thus, one can say that in our approach agents are "nonstandardly" logically omniscient.

We shall actually show that the result of replacing NPR in $\mathrm{K}^{\hookrightarrow}$ by modus ponens and negation replacement, along with all sound axioms of NPL, is complete. It follows easily that NPR can be replaced by any complete axiomatization of NPL that includes modus ponens and negation replacement as inference rules.

In the rest of this section, when we say simply that a formula is *provable*, we mean provable in $\mathrm{K}^{\hookrightarrow}$. We say that a formula φ is *consistent* if $(\varphi \hookrightarrow \mathit{false})$ is not provable. A finite set $\{\varphi_1, \ldots, \varphi_k\}$ of formulas is said to be consistent exactly if $\varphi_1 \wedge \ldots \wedge \varphi_k$ is consistent, and an infinite set of formulas is said to be consistent exactly if all of its finite subsets are consistent. (We do not worry about which parenthesization of $\varphi_1 \wedge \ldots \wedge \varphi_k$ to use, since they are all provably equivalent by NPR.)

Before we prove completeness of $\mathrm{K}^{\hookrightarrow}$, we need to prove some lemmas.

Lemma 6.11: *Let V be a consistent set of formulas of $\mathcal{L}^{\hookrightarrow}$. Then V is a maximal consistent set iff for each formula φ of $\mathcal{L}^{\hookrightarrow}$, either φ or $(\varphi \hookrightarrow \mathit{false})$ is in V.*

Proof: We first prove the "if" direction. Let φ be a formula of $\mathcal{L}^{\hookrightarrow}$ that is not in V. By assumption, $(\varphi \hookrightarrow \mathit{false}) \in V$. But

$$(\varphi \wedge (\varphi \hookrightarrow \mathit{false})) \hookrightarrow \mathit{false}$$

is an instance of NPR. Hence, $V \cup \{\varphi\}$ is inconsistent. So V is maximal.

We now prove the "only if" direction. Assume that V is a maximal consistent set. Assume that neither φ nor $(\varphi \hookrightarrow \mathit{false})$ is in V. Since $\varphi \notin V$, it follows by maximality of V that there is a finite conjunction ψ of certain members of V such that $\psi \wedge \varphi$ is inconsistent, that is, $((\psi \wedge \varphi) \hookrightarrow \mathit{false})$ is provable. Similarly, there is a finite conjunction ψ' of certain members of V such that $(\psi' \wedge (\varphi \hookrightarrow \mathit{false}) \hookrightarrow \mathit{false})$ is provable. Now the formula

$$((\psi \wedge \varphi) \hookrightarrow \mathit{false}) \hookrightarrow ((\psi' \wedge (\varphi \hookrightarrow \mathit{false}) \hookrightarrow \mathit{false}) \hookrightarrow ((\psi \wedge \psi') \hookrightarrow \mathit{false})). \qquad (1)$$

is an instance of NPR, since it is a substitution instance of

$$(\neg(p \wedge q) \Rightarrow (\neg(p' \wedge \neg q) \Rightarrow \neg(p \wedge p')))^{nst}.$$

By (1) and two applications of modus ponens, we see that $\psi \wedge \psi'$ is inconsistent. This contradicts consistency of V. ∎

Lemma 6.12: *Let V be a maximal consistent set.*

1. *If φ is provable, then $\varphi \in V$.*

2. If $\varphi \in V$ and $(\varphi \hookrightarrow \psi) \in V$, then $\psi \in V$.

3. If $\varphi \wedge \psi \in V$, then $\varphi \in V$ and $\psi \in V$.

Proof: Assume that φ is provable. If $\varphi \notin V$, then by Lemma 6.11, $(\varphi \hookrightarrow \mathit{false}) \in V$. Now

$$\varphi \hookrightarrow ((\varphi \hookrightarrow \mathit{false}) \hookrightarrow \mathit{false}) \tag{2}$$

is an instance of NPR. Since φ is provable, by (2) and modus ponens, so is $((\varphi \hookrightarrow \mathit{false}) \hookrightarrow \mathit{false})$. Since $(\varphi \hookrightarrow \mathit{false}) \in V$, it follows that V is inconsistent. So $\varphi \in V$.

Assume now that $\varphi \in V$ and $(\varphi \hookrightarrow \psi) \in V$. If $\psi \notin V$, then by Lemma 6.11, $(\psi \hookrightarrow \mathit{false}) \in V$. But

$$(\varphi \wedge (\varphi \hookrightarrow \psi) \wedge (\psi \hookrightarrow \mathit{false})) \hookrightarrow \mathit{false})$$

is an instance of NPR. So V is inconsistent, a contradiction.

Assume now that $\varphi \wedge \psi \in V$, but that $\varphi \notin V$ or $\psi \notin V$. Say for the sake of definiteness that $\varphi \notin V$. By Lemma 6.11, $(\varphi \hookrightarrow \mathit{false}) \in V$. Now

$$((\varphi \hookrightarrow \mathit{false}) \wedge (\varphi \wedge \psi)) \hookrightarrow \mathit{false}))$$

is an instance of NPR. That is, the set $\{\varphi \wedge \psi, (\varphi \hookrightarrow \mathit{false})\}$ is inconsistent, so V is inconsistent, a contradiction. ∎

We are now ready to state completeness of $\mathrm{K}^{\hookrightarrow}$.

Theorem 6.13: $\mathrm{K}^{\hookrightarrow}$ *is a sound and complete axiomatization with respect to $\mathcal{NM}$ for formulas in the language $\mathcal{L}^{\hookrightarrow}$.*

Proof: See Appendix A. ∎

Remark 6.14: We can, of course, replace NPR by those propositional axioms and rules that are actually used in the proof of Theorem 6.13 (including those used in the proofs of lemmas). The propositional rules that were used are modus ponens and negation replacement. ∎

When we presented the axiom system K we remarked that PR can be replaced by any complete axiomatization of standard propositional logic that includes modus ponens as an inference rule. Surprisingly, this is not the case here, as the next theorem shows.

Theorem 6.15: *The result of replacing NPR by all substitution instances of valid formulas of NPL, with modus ponens as the sole propositional inference rule, is not a complete axiomatization with respect to $\mathcal{NM}$ for formulas in the language $\mathcal{L}^{\hookrightarrow}$.*

Proof: Let A be the axiom system described in the statement of the theorem. Let γ be the formula $(\neg K_1 true) \hookrightarrow \mathit{false}$. We leave to the reader the straightforward verification that γ is valid. However, we now show that γ is not provable in A. For the purposes of this proof only, we shall treat not only the primitive propositions, but also all formulas of the form $K_i\psi$, where $\psi \in \mathcal{L}^{\hookrightarrow}$, as if they were primitive propositions.

Let us call this enlarged set of primitive propositions Φ'. Similarly to the proof of Theorem 4.2, let s be the truth assignment where $s(p) = \mathbf{true}$ for every $p \in \Phi'$, and let t be the truth assignment where $t(p) = \mathbf{false}$ for every $p \in \Phi'$. Let S be the NPL structure (s, t), and let T be the set of all formulas φ such that $(S, s) \models \varphi$. We now show

1. Every formula provable in A is in T.

2. γ is not in T.

Of course, this is sufficient to show that γ is not provable in A, as desired.

We first show that every formula provable in A is in T. Let $\varphi_1, \ldots, \varphi_m$ be a proof in A. We shall show, by induction on j, that each φ_j is in T, that is, $(S, s) \models \varphi_j$.

1. If φ_j is an instance $(K_i\varphi \wedge K_i(\varphi \hookrightarrow \psi)) \hookrightarrow K_i\psi$ of the distribution axiom, then $(S, s) \models \varphi_j$, since $(S, s) \models K_i\psi$ (because $K_i\psi \in \Phi'$).

2. If φ_j is a substitution instance of a valid formula of NPL, then $(S, s) \models \varphi_j$, because (S, s) is an NPL structure.

3. If φ_j is proven from an earlier φ_k by knowledge generalization, then φ_j is of the form $K_i\psi$, and so $(S, s) \models \varphi_j$ by construction (because $\varphi_j \in \Phi'$).

4. If φ_j follows from earlier formulas φ_k and φ_l (where φ_l is $\varphi_k \hookrightarrow \varphi_j$) by modus ponens, then by induction assumption $(S, s) \models \varphi_k$ and $(S, s) \models \varphi_k \hookrightarrow \varphi_j$. Therefore, once again, $(S, s) \models \varphi_j$.

We close by showing that γ is not in T. By construction, $(S, t) \not\models K_1 true$, because $K_1 true \in \Phi'$. Therefore, $(S, s) \models \neg K_1 true$, so $(S, s) \not\models (\neg K_1 true) \hookrightarrow \mathit{false}$, that is, $(S, s) \not\models \gamma$. So $\gamma \notin T$, as desired. ∎

It follows immediately from Theorem 6.15 that if NPR is replaced by a complete axiomatization of NPL with modus ponens as the sole propositional inference rule (such as the system N of the previous subsection), then the result is not a complete axiomatization for our nonstandard logic of knowledge. However, the proof of Theorem 6.13 shows that NPR can be replaced by any complete axiomatization of NPL that includes modus ponens and negation replacement as inference rules, and still maintain completeness.

Just as we can embed standard propositional logic into NPL by using ^{nst} (see Proposition 6.2), we can similarly embed standard epistemic logic into our nonstandard epistemic logic. It then follows, as with the propositional case (Theorem 6.4), that the complexity of the validity problem for standard epistemic logic is a lower bound on the complexity of the validity problem in our nonstandard epistemic logic. The corresponding upper bound can be proved by well-known techniques [Lad77]. Thus, the complexity of the validity problem is *PSPACE*-complete, just as for the standard case K.

Theorem 6.16: *The validity problem for $\mathcal{L}^{\hookrightarrow}$-formulas with respect to $\mathcal{NM}$ is PSPACE-complete.*

7 A Payoff: Querying Knowledge Bases

As we have observed, logical omniscience still holds in the nonstandard approach, though in a weakened form. We also observed that the complexity of reasoning about knowledge has not improved. Thus, the gain from our nonstandard approach seems quite modest. We now show an additional nice payoff for our approach: we show that in a certain important application we can obtain a polynomial-time algorithm for reasoning about knowledge.

The application is one we have alluded to earlier, where there is a (finite) knowledge base of facts. Thus, the knowledge base can be viewed as a formula κ. A query to the knowledge base is another formula φ. There are two ways to interpret such a query. First, we can ask whether φ is a consequence of κ. Second, we can ask whether knowledge of φ follows from knowledge of κ. Fortunately, these are equivalent questions, as we now see.

Proposition 7.1: *Let φ_1 and φ_2 be $\mathcal{L}^{\hookrightarrow}$-formulas. Then φ_1 logically implies φ_2 with respect to $\mathcal{NM}$ iff $K_i\varphi_1$ logically implies $K_i\varphi_2$ with respect to $\mathcal{NM}$.*

Proof: It is easy to see that if φ_1 logically implies φ_2 with respect to $\mathcal{NM}$, then $K_i\varphi_1$ logically implies $K_i\varphi_2$ with respect to $\mathcal{NM}$. We now show the converse.

Assume that φ_1 does not logically imply φ_2 with respect to $\mathcal{NM}$. Let $M = (S, \pi, \mathcal{K}_1, \ldots, \mathcal{K}_n, {}^*)$ be a nonstandard structure and u a world of M such that $(M, u) \models \varphi_1$ and $(M, u) \not\models \varphi_2$. Define a new nonstandard structure $M' = (S', \pi', \mathcal{K}'_1, \ldots, \mathcal{K}'_n, {}^\dagger)$ with one additional world $t \notin S$ by letting (a) $S' = S \cup \{t\}$, (b) $\pi'(s) = \pi(s)$ for $s \in S$, and $\pi'(t)$ be arbitrary, (c) $\mathcal{K}'_j = \mathcal{K}_j$ for $j \neq i$, and $\mathcal{K}'_i = \mathcal{K}_i \cup \{(t, u)\}$, and (d) $s^\dagger = s^*$ for $s \in S$, and $t^\dagger = t$. It is straightforward to see that since $(M, u) \models \varphi_1$ and $(M, u) \not\models \varphi_2$, also $(M', u) \models \varphi_1$ and $(M', u) \not\models \varphi_2$. But then $(M', t) \models K_i\varphi_1$ and $(M', t) \not\models K_i\varphi_2$, and hence $K_i\varphi_1$ does not logically imply $K_i\varphi_2$ with respect to $\mathcal{NM}$. ∎

We focus here on the simple case where both the knowledge base and the query are standardly propositional (i.e., no $\hookrightarrow$). We know that in the standard approach determining whether κ logically implies φ is co-NP-complete. Is the problem of determining the consequences of a knowledge base in the nonstandard approach (i.e., determining whether κ logically implies φ, or equivalently, by Proposition 7.1, whether $K_i\kappa$ logically implies $K_i\varphi$) any easier? Unfortunately, the answer to this question is negative (since if

φ is *false*, then the problem is the same as deciding whether $\neg\kappa$ is a tautology of NPL, which is co-*NP*-hard by Theorem 6.4.) There is, however, an interesting special case where using the nonstandard semantics does make the problem easier.

Define a *literal* to be a primitive proposition p or its negation $\neg p$, and define a *clause* to be a disjunction of literals. For example, a typical clause is $p \vee \neg q \vee r$. We can consider a traditional database as being a collection of atomic facts, which can be thought of as primitive propositions. It is often an implicit assumption that if an atomic fact does not appear in a database, then its negation can be considered to be in the database (this assumption is called the *closed world assumption* [Rei78]). We can imagine a database that explicitly contains not only atomic facts but also negations of atomic facts. This would correspond to a database of literals. More generally yet, we could consider a database (or knowledge base) of clauses, that is, disjunctions of literals. In fact, there are many applications in which the knowledge base consists of a finite collection of clauses. Thus, κ (which represents the knowledge base) is a conjunction of clauses. A formula (such as κ) that is a conjunction of clauses is said to be in *conjunctive normal form* (or *CNF*).

Hence, we can think of the knowledge base κ as being a formula in CNF. What about the query φ? Every standard propositional formula is equivalent to a formula in CNF (this is true even in our nonstandard semantics, because of Proposition 2.1). Thus, we will assume that the query φ has been transformed to CNF. (Note that we assumed that the knowledge base is given in CNF, while the query has to be transformed to CNF. The reason for this distinction is the fact that the transformation to CNF may involve an exponential blow-up. Consequently, while we might be reasonable to apply it to the query, it is not reasonable to apply it to the knowledge base, which is typically orders of magnitude larger than the query.)

Let us now reconsider the query evaluation problem, where both the knowledge base and the query are in CNF. The next proposition tells us that under the standard semantics, the problem is no easier than the general problem of logical implication in propositional logic, that is, co-*NP*-complete.

Proposition 7.2: *The problem of deciding whether κ logically implies φ in standard propositional logic, for CNF formulas κ and φ, is co-NP-complete.*

Proof: Let κ be an arbitrary CNF formula, and let p be a primitive proposition that does not appear in κ. Now κ logically implies p in standard propositional logic iff κ is unsatisfiable in standard propositional logic. This is because if $\kappa \Rightarrow p$ is valid, then so is $\kappa \Rightarrow \neg p$, and hence $\kappa \Rightarrow (p \wedge \neg p)$. This is sufficient to prove the proposition, since the problem of determining nonsatisfiability of a CNF formula is co-*NP*-complete. ∎

By contrast, the problem is feasible under the nonstandard semantics. Before we show this, we need a little more machinery.

Let us say that clause α_1 *includes* clause α_2 if every literal that is a disjunct of α_1 is a disjunct of α_2. For example, the clause $p \vee \neg q \vee \neg r$ includes the clause $p \vee \neg q$. The next theorem characterizes when κ logically implies φ in NPL, for CNF formulas κ and φ.

Theorem 7.3: *Let κ and φ be propositional formulas in CNF. Then κ logically implies φ in NPL iff every clause of φ includes a clause of κ.*

Proof: The "if" direction, which is fairly straightforward, is left to the reader. We now prove the other direction. Assume that some clause α of φ includes no clause of κ. We need only show that there is an NPL structure $S = (s,t)$ such that $(S,s) \models \kappa$ but $(S,s) \not\models \varphi$. Define $s(p) =$ **false** iff p is a disjunct of α, and $t(p) =$ **true** iff $\neg p$ is a disjunct of α, for each primitive proposition p. We now show that $(S,s) \not\models \alpha'$, for each disjunct α' of α. If α' is a primitive proposition p, then $s(p) =$ **false**, so $(S,s) \not\models \alpha'$; if α' is $\neg p$, where p is a primitive proposition, then $t(p) =$ **true**, so $(S,t) \models p$, so again $(S,s) \not\models \alpha'$. Hence, $(S,s) \not\models \alpha$, so $(S,s) \not\models \varphi$. However, $(S,s) \models \kappa$, since every conjunct κ' of κ has a disjunct κ'' where $(S,s) \models \kappa''$ (otherwise, α would include κ'). ∎

An example where Theorem 7.3 would be false in standard propositional logic occurs when κ is $q \vee \neg q$ and φ is $p \vee \neg p$. Then κ logically implies φ in standard propositional logic, but the single clause $p \vee \neg p$ of φ does not include the single clause of κ. Note that κ does not logically imply φ in NPL.

It is clear that Theorem 7.3 gives us a polynomial-time decision procedure for deciding whether one CNF formula implies another in the nonstandard approach.

Theorem 7.4: *There is a polynomial-time decision procedure for deciding whether κ logically implies φ in NPL (or $K_i\kappa$ logically implies $K_i\varphi$ with respect to $\mathcal{NM}$), for CNF formulas κ and φ.*

Theorems 7.3 and 7.4 yield an efficient algorithm for the evaluation of a CNF query ψ with respect to a CNF knowledge base κ: answer "Yes" if κ logically implies ψ in NPL. By Theorem 7.4, logical implication of CNF formulas in NPL can be checked in polynomial time, and Theorem 7.3 implies that any positive answer we obtain from testing logical implication between CNF fomulas in the nonstandard semantics will provide us with a correct positive answer for standard semantics as well. This means that even if we are ultimately interested only in conclusions that are derivable from standard reasoning, we can safely use the positive conclusions we obtain using nonstandard reasoning. Thus, the nonstandard approach yields a feasible query-answering algorithm for knowledge bases. Notice that the algorithm need not be correct with respect to negative answers. It is possible that κ does not logically imply ψ in NPL even though κ logically implies ψ with respect to standard propositional logic.

Theorem 7.4 was essentially proved in [Lev84]. The precise relationship to Levesque's results will be clarified in Section 9. Levesque's result (like Theorem 7.4) applies only to propositional formulas κ. Lakemeyer [Lak87] extended it to modal formulas for the single-agent case. He defined the class of *extended-conjunctive-normal-form* (ECNF) formulas and showed that Theorem 7.4 holds also for ECNF formulas [Lak87]. Thus, his result shows that under the nonstandard semantics, there are nontrivial tractable fragments of the language that include modal formulas. Interestingly, a 4-valued semantics was also used in a different context in order to deal with computational complexity; Patel-Schneider defined a 4-valued terminological logic with tractable subsumption [PS89].

8 Standard-World Validity

Logical omniscience arises from considering knowledge as truth in all possible worlds. In the approach of this paper, we modify logical omniscience by changing the notion of truth. In this section, we consider the *impossible-worlds* approach, where we modify logical omniscience by changing the notion of possible world. The idea is to augment the possible worlds by impossible worlds, where the customary rules of logic do not hold. Even though these worlds are logically impossible, the agents nevertheless may consider them possible. Unlike our approach, where nonstandard worlds are considered just as realistic as standard worlds, under the impossible-worlds approach the impossible worlds are a figment of the agents' imagination; they serve only as epistemic alternatives.

Since agents consider the impossible worlds when computing their knowledge, logical omniscience need not hold. For example, suppose that an agent knows all formulas in Σ, and Σ logically implies φ. Since the agent knows all formulas in Σ, all formulas in Σ must hold in all the worlds that the agent considers possible. However, even though Σ logically implies φ, it can happen that φ does not hold at one of the impossible worlds the agent considers possible, and so the agent may not know φ. The key point here is that logical implication is determined by us, rational logicians for whom impossible worlds are indeed impossible. We do not consider impossible worlds when determining logical implication.

There are various impossible-worlds approaches (see, for example, [Ran82] and [Wan90]), depending on how we choose the possible and impossible worlds. In what follows, we shall take the possible worlds to be the standard worlds, and the impossible worlds to be the nonstandard worlds.

The difference between our approach and the impossible-worlds approach is that in our approach the distinction between standard and nonstandard worlds does not play any role. In the impossible-worlds approach, however, the standard worlds (those where $s = s^*$) have a special status. Intuitively, although an agent (who is not a perfect reasoner) might consider nonstandard worlds possible (where, for example, $p \wedge \neg p$ or $K_i p \wedge \neg K_i p$ holds), we as logicians do not consider such worlds possible; surely in the real world a formula is either true or false, but not both.

This distinction plays an important role in the way validity and logical implication are defined. In the impossible-worlds approach we consider nonstandard worlds to be "impossible", and thus consider a formula φ to be valid if it is true at all of the "possible" worlds, that is, at all of the standard worlds. Formally, define a formula of $\mathcal{L}$ to be *standard-world valid* if it is true at every standard world of every nonstandard structure. The definition for *standard-world logical implication* is analogous.

The reader may recall that, under the nonstandard semantics, $\Rightarrow$ behaves peculiarly. In particular, $\Rightarrow$ does not capture the notion of logical implication. In fact, that was one of the motivations to the introduction of strong implication. At standard worlds, however, $\Rightarrow$ and $\hookrightarrow$ coincide, that is, $\varphi_1 \Rightarrow \varphi_2$ holds at a standard world precisely if $\varphi_1 \hookrightarrow \varphi_2$ holds. It follows that even though $\Rightarrow$ does not capture logical implication, it does capture standard-world logical implication. The following analogue to Proposition 5.2 is immediate.

Proposition 8.1: *Let φ_1 and φ_2 be formulas in $\mathcal{L}$. Then φ_1 standard-world logically implies φ_2 iff $\varphi_1 \Rightarrow \varphi_2$ is standard-world valid.*

The main feature of the impossible-worlds approach is the fact that knowledge is computed over all worlds, while logical implication is evaluated only over standard worlds. As a result we avoid logical omniscience. For example, an agent does not necessarily know valid formulas of *standard* propositional logic. Specifically, although the classical tautology $p \vee \neg p$ is standard-world valid, an agent may not know this formula at a standard world s, since the agent might consider an incomplete world possible.

Let φ be a formula that contains precisely the primitive propositions $p_1, \ldots, p_k$. Define Complete(φ) to be the formula

$$(p_1 \vee \neg p_1) \wedge \ldots \wedge (p_k \vee \neg p_k).$$

Thus, Complete(φ) is true at a world s precisely if s is complete as far as all the primitive propositions in φ are concerned. In particular, if φ is propositional, and if Complete(φ) is true at s, then it follows by a simple induction on formulas that s is complete with respect to φ.

Let φ be a tautology of standard propositional logic. Clearly φ is true at every world s that is complete and coherent (with respect to all of the primitive propositions in φ). The next proposition implies that if we assume only that s is complete, then this is still enough to guarantee that φ be true at s.

Proposition 8.2: *Let φ be a standard propositional formula. Then φ is a tautology of standard propositional logic iff Complete(φ) logically implies φ in NPL.*

Proof: Assume first that Complete(φ) logically implies φ in NPL. To show that φ is a tautology of standard propositional logic, we need only show that φ is true at every world that is complete and coherent. But this is the case, since if s is complete, then by assumption φ is true at s.

Assume now that φ is a tautology of standard propositional logic, and $(S, u) \models$ Complete(φ). Let Ψ be the set of primitive propositions that appear in φ. Thus, $(S, u) \models$

$p \vee \neg p$ for each $p \in \Psi$. Hence, either $(S, u) \models p$ or $(S, u) \models \neg p$, for each $p \in \Psi$. Define the truth assignment v by letting $v(p) = \mathbf{true}$ if $(S, u) \models p$ and $v(p) = \mathbf{false}$ otherwise. Note in particular that if $v(p) = \mathbf{false}$, then $(S, u) \models \neg p$. By a straightforward induction on formulas, we can show that for each propositional formula ψ all of whose primitive propositions are in Ψ, we have

1. If ψ is true under v, then $(S, u) \models \psi$.
2. If ψ is false under v, then $(S, u) \models \neg\psi$.

Now φ is true under v, since φ is a tautology of standard propositional logic. So from what we just showed, it follows that $(S, u) \models \varphi$. Hence, Complete(φ) logically implies φ in NPL. ∎

From Theorem 8.2 we obtain immediately the proof we promised of part of Theorem 2.3, that the logical implication problem in NPL$^-$ is co-*NP*-complete. The proof of Theorem 4.3 (that the logical implication problem for $\mathcal{L}$-formulas in nonstandard structures is *PSPACE*-complete) follows from a generalization of Proposition 8.2. Define $E\varphi$ ("everyone knows φ") to be $K_1\varphi \wedge \ldots \wedge K_n\varphi$, where the agents are $1, \ldots, n$. Define $E^0\varphi$ to be φ, and inductively define $E^{r+1}\varphi$ to be $EE^r\varphi$. We now define the *depth* of a formula φ, denoted depth(φ), as follows: depth(p) $= 0$ if p is a primitive proposition; depth($\neg\varphi$) $=$ depth(φ); depth($\varphi_1 \wedge \varphi_2$) $= \max$ (depth(φ_1), depth(φ_2)); and depth($K_i\varphi$) $=$ depth(φ) $+ 1$.

Proposition 8.3: *Assume $\varphi \in \mathcal{L}$ has depth d. Then φ is valid with respect to standard structures iff Complete(φ) $\wedge$ E(Complete(φ)) $\wedge \ldots \wedge$ E^d(Complete(φ)) logically implies φ in nonstandard structures.*

Proof: The proof is a fairly straightforward generalization of that of Proposition 8.2. The details are omitted. ∎

Ladner [Lad77] showed that the *PSPACE* lower bound for validity in standard structures (Theorem 3.2) holds even when there is only one agent. If we replace E in Proposition 8.3 by K_1, then it follows from Proposition 8.3 that there is a polynomial reduction of validity in standard structures with one agent to the logical implication problem for $\mathcal{L}$-formulas in nonstandard structures with one agent. The *PSPACE* lower bound in Theorem 4.3 now follows (even when there is only one agent). The upper bound follows from Theorem 6.16.

If φ is a tautology of standard propositional logic, then an agent need not know φ, even at a standard world, since φ may be false at an incomplete world that the agent considers possible. The next theorem says that if the agent knows that the world is complete, then he must know the tautology φ. This theorem follows from results in [FH88].

Theorem 8.4: *Let φ be a tautology of standard propositional logic. Then $K_i(Complete(\varphi)) \Rightarrow K_i\varphi$ is standard-world valid.*

Proof: By Proposition 8.2, Complete(φ) logically implies φ in NPL. Hence, by Proposition 7.1, K_i(Complete(φ)) logically implies $K_i\varphi$ with respect to $\mathcal{NM}$. It follows by Proposition 8.1 that $K_i(\text{Complete}(\varphi)) \Rightarrow K_i\varphi$ is standard-world valid. ∎

Another form of logical omniscience that fails under the impossible-worlds approach is closure under implication: it is easy to see that the formula $(K_i\varphi \wedge K_i(\varphi \Rightarrow \psi)) \Rightarrow K_i\psi$ is not standard-world valid. This lack of closure results from considering incoherent worlds possible: indeed, it is not hard to see that $(K_i\varphi \wedge K_i(\varphi \Rightarrow \psi)) \Rightarrow K_i(\psi \vee (\varphi \wedge \neg\varphi))$ is standard-world valid. That is, if an agent knows that φ holds and also knows that $\varphi \Rightarrow \psi$ holds, then he knows that either ψ holds or the world is incoherent. If the agent knows that the world is coherent, then his knowledge is closed under logical implication. We now formalize this observation.

Let φ be a formula that contains precisely the primitive propositions $p_1, \ldots, p_k$. Define Coherent(φ) to be the formula

$$((p_1 \wedge \neg p_1) \hookrightarrow \mathit{false}) \wedge \ldots \wedge ((p_k \wedge \neg p_k) \hookrightarrow \mathit{false}).$$

Thus, Coherent(φ) is true at a world s precisely if s is coherent as far as all the primitive propositions in φ are concerned. In particular, if Coherent(φ) holds at s, then s is coherent with respect to φ.[1]

The next theorem says that knowledge of coherence implies that knowledge is closed under implication.

Theorem 8.5: *Let φ and ψ be standard propositional formulas. Then $(K_i(Coherent(\varphi)) \wedge K_i\varphi \wedge K_i(\varphi \Rightarrow \psi)) \Rightarrow K_i\psi$ is standard-world valid.*

Proof: Denote $K_i(\text{Coherent}(\varphi)) \wedge K_i\varphi \wedge K_i(\varphi \Rightarrow \psi)$ by τ. By Proposition 8.1, it is sufficient to show that τ standard-world logically implies $K_i\psi$. We shall show the stronger fact that τ logically implies $K_i\psi$ with respect to $\mathcal{NM}$. Let $M = (S, \pi, \mathcal{K}_1, \ldots, \mathcal{K}_n, ^*)$ be a nonstandard structure, and s a world of M. Assume that τ is true at s and that $(s,t) \in \mathcal{K}_i$, so Coherent(φ) is true at t. By a straightforward induction on formulas, we can show that for every propositional formula γ all of whose primitive propositions are contained in φ, it is not the case that both γ and $\neg\gamma$ are true at t. Now φ and $\varphi \Rightarrow \psi$ are both true at t, since $K_i\varphi$ and $K_i(\varphi \Rightarrow \psi)$ are true at s. Since φ is true at t, it follows from what we just showed that $\neg\varphi$ is not true at t. Since $\varphi \Rightarrow \psi$ is an abbreviation for $\neg\varphi \vee \psi$, it follows that ψ is true at t. Hence, $K_i\psi$ is true at s. ∎

[1]Note that Coherent(φ) is not definable in $\mathcal{L}$ but only in $\mathcal{L}^{\hookrightarrow}$. This is because if there were a formula in $\mathcal{L}$ that says that at most one of p or $\neg p$ is true, then φ would have to be false at the state t of Theorem 4.2, since both p and $\neg p$ are true at t. However, φ (along with every formula of $\mathcal{L}$) is true at t.

Theorems 8.4 and 8.5 explain why agents are not logically omniscient: "logically" is defined here with respect to standard worlds, but the agents may consider nonstandard worlds possible. If an agent considers only standard worlds possible, so that we have the antecedents $K_i(\text{Complete}(\varphi))$ and $K_i(\text{Coherent}(\varphi))$ of Theorems 8.4 and 8.5, then this agent is logically omniscient (more accurately, he knows every tautology of standard propositional logic and his knowledge is closed under implication).

We conclude the discussion of the impossible-world approach by reconsidering the knowledge base situation discussed earlier, where the knowledge base is described by a formula κ and the query is described by a formula φ. We saw earlier (Proposition 7.1) that in the nonstandard approach, φ is a consequence of κ precisely when knowledge of φ is a consequence of knowledge of κ.

The situation is different under the impossible-worlds approach. On one hand, implication of knowledge coincides in both approaches.

Proposition 8.6: *Let φ_1 and φ_2 be $\mathcal{L}^{\hookrightarrow}$-formulas. Then $K_i\varphi_1$ standard-world logically implies $K_i\varphi_2$ iff $K_i\varphi_1$ logically implies $K_i\varphi_2$ in nonstandard structures.*

Proof: The proof, which is very similar to that of Proposition 7.1, is left to the reader. ∎

On the other hand, the two interpretations of query evaluation differ in the impossible-worlds approach. In contrast to Proposition 7.1, it is possible to find φ_1 and φ_2 in $\mathcal{L}$ such that φ_1 standard-world logically implies φ_2, but $K_i\varphi_1$ does not standard-world logically imply $K_i\varphi_2$ (let φ_1 be $p \wedge \neg p$, and let φ_2 be q). The reason for this failure is that φ_1 standard-world logically implying φ_2 deals with logical implication in standard worlds, whereas $K_i\varphi_1$ standard-world logically implying $K_i\varphi_2$ deals with logical implication in worlds agents consider possible, which includes nonstandard worlds.

The difference between the two interpretations of query evaluation in the standard approach can have a significant computational impact. Consider the situation where both κ and φ are CNF propositional formulas. Theorem 7.4 and Proposition 8.6 tell us that testing whether $K_i\kappa$ standard-world logically implies $K_i\varphi$ can be done in polynomial time. However, in this case, testing whether κ standard-world logically implies φ is co-*NP*-complete, according to Proposition 7.2.

9 Levesque and Lakemeyer's formalism

In this section, we relate our results to those of Levesque [Lev84] and Lakemeyer [Lak87]. First, we relate our syntax and semantics to theirs.

Levesque and Lakemeyer also attempt to decouple the semantics of a formula from that of its negation, but their approach is different from ours. We briefly discuss the details, and then present their formal sematics.

Define a *nonstandard truth assignment* to be a function that assigns to each literal a truth value. (Recall that a literal is either a primitive proposition p or its negation $\neg p$). Thus, although an ordinary truth assignment assigns a truth value to each primitive proposition p, a nonstandard truth assignment assigns a truth value to both p and $\neg p$, for each primitive proposition p. Under a given nonstandard truth assignment, it is possible that both p and $\neg p$ can be assigned the value **true**, or that both can be assigned **false**, or that one can be assigned **true** and the other **false**. This decouples the semantics of p and $\neg p$. As we shall show below, it is quite straightforward to decouple the semantics of a conjunction from its negation, once we have already done so for each of its conjuncts. Levesque and Lakemeyer do not have $\hookrightarrow$ in the language, so there is no need for them to decouple the semantics of $\varphi \hookrightarrow \psi$ from $\neg(\varphi \hookrightarrow \psi)$. Finally, in order to decouple the semantics of $K_i\varphi$ and $\neg K_i\varphi$, Lakemeyer introduces two possibility relations, $\mathcal{K}_i^+$ and $\mathcal{K}_i^-$.

A *Levesque-Lakemeyer structure* (or *LL structure* for short) is a tuple

$$M = (S, \pi, \mathcal{K}_1^+, \ldots, \mathcal{K}_n^+, \mathcal{K}_1^-, \ldots, \mathcal{K}_n^-),$$

where S is a set of worlds, $\pi(s)$ is a nonstandard truth assignment for each world $s \in S$, and each $\mathcal{K}_i^+$ and K_i^- is a binary relation on S. To define the semantics, Levesque introduces two "support relations" $\models_T$ and $\models_F$. Intuitively, $(M,s)\models_T \varphi$ (where T stands for "true") means that the truth of φ is supported at (M,s), while $(M,s)\models_F \varphi$ (where F stands for "false") means that the falsity of φ is supported at (M,s). We say that $(M,s) \models \varphi$ if $(M,s)\models_T \varphi$. The semantics is as follows:

$(M,s)\models_T p$ (for a primitive proposition p) iff $\pi(s)(p) = \mathbf{true}$

$(M,s)\models_F p$ (for a primitive proposition p) iff $\pi(s)(\neg p) = \mathbf{true}$

$(M,s)\models_T \neg\varphi$ iff $(M,s)\models_F \varphi$

$(M,s)\models_F \neg\varphi$ iff $(M,s)\models_T \varphi$

$(M,s)\models_T \varphi_1 \wedge \varphi_2$ iff $(M,s)\models_T \varphi_1$ and $(M,s)\models_T \varphi_2$

$(M,s)\models_F \varphi_1 \wedge \varphi_2$ iff $(M,s)\models_F \varphi_1$ or $(M,s)\models_F \varphi_2$

$(M,s)\models_T K_i\varphi$ iff $(M,t)\models_T \varphi$ for all t such that $(s,t) \in \mathcal{K}_i^+$

$(M,s)\models_F K_i\varphi$ iff $(M,t)\not\models_T \varphi$ for some t such that $(s,t) \in \mathcal{K}_i^-$.

We also remark that Levesque and Lakemeyer have two different flavors of knowledge in their papers: explicit knowledge and implicit knowledge. (Actually, they talk about belief rather than knowledge, but the distinction is irrelevant to our discussion here.) We consider here only their notion of explicit knowledge, since this is the type that avoids logical omniscience.

Although, superficially, our semantics seems quite different from the Levesque-Lakemeyer semantics, it is straightforward to show that in fact, the two approaches are equivalent in the sense of the following proposition.

Proposition 9.1: *For each nonstandard structure M and world s in M, there is an LL structure M' and world s' in M' such that for each $\mathcal{L}$-formula φ,*

$$(M,s) \models \varphi \textit{ iff } (M',s') \models_T \varphi \tag{3}$$

$$(M,s) \models \neg\varphi \textit{ iff } (M',s') \models_F \varphi. \tag{4}$$

Conversely, for each LL structure M' and world s' in M', there is a nonstandard structure M and world s in M such that (3) and (4) hold for each $\mathcal{L}$-formula φ.

Proof: Given a nonstandard structure $M = (S, \pi, \mathcal{K}_1, ..., \mathcal{K}_n, ^*)$, define an LL structure $M = (S, \pi', \mathcal{K}_1^+, \ldots, \mathcal{K}_n^+, \mathcal{K}_1^-, \ldots, \mathcal{K}_n^-)$ with the same set S of worlds, where for each state s and primitive proposition p, we have

$$\pi'(s)(p) = \pi(s)(p) \tag{5}$$

$$\pi'(s)(\neg p) = \textbf{true} \text{ iff } \pi(s^*)(p) = \textbf{false}, \tag{6}$$

and where

$$\mathcal{K}_i^+ = \mathcal{K}_i$$

$$(s,t) \in \mathcal{K}_i^- \text{ iff } (s^*, t) \in \mathcal{K}_i.$$

It is easy to show by induction on the structure of φ that for every world s,

$$(M,s) \models \varphi \text{ iff } (M',s) \models_T \varphi \tag{7}$$

$$(M,s) \models \neg\varphi \text{ iff } (M',s) \models_F \varphi. \tag{8}$$

For the converse, let $M' = (S, \pi', \mathcal{K}_1^+, \ldots, \mathcal{K}_n^+, \mathcal{K}_1^-, \ldots, \mathcal{K}_n^-)$ be an LL structure. We define a nonstandard structure $M = (S \cup S^*, \pi, \mathcal{K}_1, ..., \mathcal{K}_n, ^*)$ by letting s^* be a new world for each $s \in S$, letting $S^* = \{s^* \mid s \in S\}$, defining $\pi(s)$ and $\pi(s^*)$ for $s \in S$ so that (5) and (6) hold for every primitive proposition p, and defining $\mathcal{K}_i$ to consist precisely of

1. all (s,t) such that $s \in S$, $t \in S$, and $(s,t) \in \mathcal{K}_i^+$, and
2. all (s^*,t) such that $s^* \in S^*$, $t \in S$, and $(s,t) \in \mathcal{K}_i^-$.

By the identical argument to before, (7) and (8) hold for every formula φ and every state $s \in S$. ∎

Remark 9.2: We note that there is an equivalent semantics to that of Levesque and Lakemeyer that avoids the use of two satisfaction relations $\models_T$ and $\models_F$. That is, we can define a notion $\models'$ of satisfaction directly such that if M is an LL structure, s is a world of M, and φ is an $\mathcal{L}$ formula, then $(M,s) \models \varphi$ iff $(M,s) \models' \varphi$. Rather than defining $\neg\varphi$ to be true iff φ is not true (as we do with standard structures), and rather than giving a uniform definition of when $\neg\varphi$ is true (as we do with nonstandard structures, using *), we instead define separately what it means for φ to be true and what it means for $\neg\varphi$ to be true, for each type of formula φ (that is, for primitive propositions, and formulas of the form $\varphi_1 \wedge \varphi_2$, $\neg\varphi$, and $K_i\varphi$). This way we can make the truth of φ and $\neg\varphi$ independent. The definition is as follows (where we write $\models$ for $\models'$, for readability):

$(M, s) \models p$ (for a primitive proposition p) iff $\pi(s)(p) = \mathbf{true}$

$(M, s) \models \varphi_1 \wedge \varphi_2$ iff $(M, s) \models \varphi_1$ and $(M, s) \models \varphi_2$

$(M, s) \models K_i\varphi$ iff $(M, t) \models \varphi$ for all t such that $(s, t) \in \mathcal{K}_i^+$

$(M, s) \models \neg p$ (for a primitive proposition p) iff $\pi(s)(\neg p) = \mathbf{true}$

$(M, s) \models \neg(\varphi_1 \wedge \varphi_2)$ iff $(M, s) \models \neg\varphi_1$ or $(M, s) \models \neg\varphi_2$

$(M, s) \models \neg\neg\varphi$ iff $(M, s) \models \varphi$

$(M, s) \models \neg K_i\varphi$ iff $(M, t) \not\models \varphi$ for some t such that $(s, t) \in \mathcal{K}_i^-$.

The proof of the equivalence to Levesque and Lakemeyer's semantics is straightforward, and is left to the reader. ∎

Remark 9.3: While Proposition 9.1 shows the equivalence of our approach to the Levesque-Lakemeyer approach, a difference between the two aproaches emerges when we try to model agents with certain attributes. It is well known that in the standard possible-world approach, agents' attributes can often be captured by imposing certain restrictions on the possibility relations. For example, positive introspection of agent i (i.e., $K_i\varphi \Rightarrow K_iK_i\varphi$) is captured by requiring $\mathcal{K}_i$ to be transitive (cf. [HM92]). The same holds in our nonstandard approach here, i.e., positive introspection of agent i is captured by requiring $\mathcal{K}_i$ to be transitive. On the other hand, in the Levesque-Lakemeyer approach it suffices to impose transitivity on $\mathcal{K}_i^+$. Now the properties of knowledge differ in our approach and in the Levesque-Lakemeyer approach. For example, in our approach $K_i\neg K_iK_i\varphi \Rightarrow K_i\neg K_i\varphi$ becomes valid, but this is not the case in the Levesque-Lakemeyer approach. Thus, the Levesque-Lakemeyer approach allows an extra degree of freedom in modeling agents. ∎

Levesque and Lakemeyer use standard-world validity as their notion of validity. Thus, their notion of logical implication is standard-world logical implication, so as in Proposition 8.1, $\Rightarrow$ can be used to express logical implication in the language. Therefore, unlike us, they do not enlarge their language to include $\hookrightarrow$. They obtain a completeness result (with some restrictions on the allowable formulas). Because they use standard-world validity, their axiomatization contains all standard tautologies. However (as is the point with impossible-worlds approaches), agents need not know all standard tautologies. Thus, for example, $p \Rightarrow p$ is valid for them, since they are considering standard world validity, but $K_i(p \Rightarrow p)$ is not, since agent i may consider a nonstandard world possible where $p \Rightarrow p$ does not hold.

Levesque [Lev84] proves that there is a polynomial-time decision procedure for deciding whether $K_i\kappa$ logically implies $K_i\varphi$, for CNF formulas κ and φ (this is the decision procedure described in Theorem 7.3). The existence of this polynomial-time decision procedure is analogous to part of Theorem 7.4. The other part of Theorem 7.4 (that there is a polynomial-time decision procedure for deciding whether κ logically implies φ, for CNF formulas κ and φ) is false in Levesque's context, since he is considering standard-world logical implication. In particular, the analogue of Proposition 7.1 does not hold for him. We originally obtained Theorem 7.4 by using Levesque's result, along with Proposition 8.6 (and Proposition 7.1).

10 Conclusions

We have investigated a new approach to dealing with the well-known logical omniscience problem in epistemic logics. The idea is to base the epistemic logic on a nonstandard logic, in the hope that by taking an appropriate nonstandard logic, we can lessen the logical omniscience problem.

The nonstandard propositional logic we use is NPL, which we introduce in this paper. NPL has a number of attractive features, including a clean semantics and an elegant complete axiomatization. In addition, there is a tractable (polynomial-time) decision procedure for evaluating a natural class of knowledge base queries. Thus, there is a sense in which the logical omniscience problem is not as acute when considering an epistemic logic based on NPL. Our approach is closely related to that of Levesque and Lakemeyer. Indeed, we feel that thinking in terms of NPL sheds new light on their results.

There is, of course, nothing special about the role of NPL in our approach. We could just as well considered epistemic logics based on other nonstandard logics. Perhaps by considering other logics we can obtain other desirable properties. We leave consideration of this point to future research.

Acknowledgements. We are grateful to the referees for their helpful comments and to Michael Dunn for clarifying the history of 4-valued logic.

A Completeness Proofs

Proof of Theorem 6.9: The axiom scheme PL is sound by Corollary 6.3. The soundness of NPL1, NPL2, and NPL3 follow by Lemma 6.7 and by Proposition 5.2 (which says that logical implication is expressed by $\hookrightarrow$).

We now prove completeness. Assume that φ is valid in NPL. Let $\varphi_1, \ldots, \varphi_k$ be a sequence of formulas obtained by driving negations down (using Lemma 6.7), working from the outside in, until the negations apply only to primitive propositions. Thus,

$\varphi_1 = \varphi$,

φ_k is pseudo-positive, and

φ_{j+1} is obtained from φ_j by driving down a negation that is as high in the parse tree as possible, for $j = 1, \ldots, k-1$.

Now each φ_{j+1} is obtained from φ_j by replacing a formula by a (provably) equivalent formula. Since $\varphi_1 = \varphi$ is valid, it is easy to see that each φ_j is valid. Our goal is to show that each φ_j is provable in our axiom system N, so that in particular, φ_1 (that is, φ) is provable. Since φ_k is pseudo-positive, it is fairly straightforward to use Corollary 6.6 to show that φ_k is provable (we give the demonstration below). We would then like to conclude the proof by showing that if φ_{j+1} is provable, then so is φ_j, since after all, the only difference between φ_j and φ_{j+1} is that some subformula γ of φ_j is replaced by a provably equivalent formula γ' to obtain φ_{j+1}. It then follows easily (as shown below) that we would be done if we showed that the formula

$$(\gamma \rightleftharpoons \gamma') \hookrightarrow (\varphi_{j+1} \hookrightarrow \varphi_j) \tag{9}$$

were provable. We shall show below that in fact, (9) is an instance of PL. This is not obvious, since there may be negations in (9).

We now give the details of the proof. We show by backwards induction on j (for $j = k, k-1, \ldots, 1$) that each φ_j is provable. As we observed above, each φ_j is valid, so in particular, φ_k is valid. Since φ_k is also pseudo-positive, it follows from Corollary 6.6 that $((\varphi_k)^+)^{st}$ is a standard propositional tautology. So $(((\varphi_k)^+)^{st})^{nst}$ is an instance of axiom scheme PL. By Lemma 6.1, $(((\varphi_k)^+)^{st})^{nst} = (\varphi_k)^+$. Thus, $(\varphi_k)^+$ is an instance of axiom scheme PL. A substitution instance of $(\varphi_k)^+$, and hence an instance of PL, is obtained by replacing every occurrence of $\bar{p}$ by $\neg p$ (where $\bar{p}$ is the new primitive proposition that replaces every occurrence $\neg p$ in φ when we form $(\varphi_k)^+$ from φ_k). Therefore, φ_k is simply an instance of PL, and so is of course provable. This takes care of the base case $j = k$ of the induction.

Assume inductively that φ_{j+1} is provable. Now φ_{j+1} is obtained from φ_j by replacing some (negated) subformula γ of φ_j by another formula γ'. Let p_γ and $p_{\gamma'}$ be new primitive propositions. Let ψ (resp. ψ') be the result of replacing this occurrence of γ (resp. γ') in φ_j (resp. φ_{j+1}) by p_γ (resp. $p_{\gamma'}$). So ψ and ψ' are identical, except that the unique occurrence of p_γ in ψ is replaced by $p_{\gamma'}$ in ψ'. If a negation appears in ψ (and hence in ψ'), then let μ be a negated subformula of ψ that appears as high as possible in the parse tree of ψ. Since γ is a negated subformula of φ_j that appears as high as possible in the parse tree of φ_j, it is not hard to see that p_γ is not a subformula of μ. Hence μ is a subformula of φ_{j+1}. Replace μ in ψ (resp. ψ') by a new primitive proposition p_μ. Continue this process until all negations are replaced. Call the final result β (resp. β'). Note that β and β' are negation-free, and are identical, except that the unique occurrence of p_γ in β is replaced by $p_{\gamma'}$ in β'. The formula

$$(p_\gamma \rightleftharpoons p_{\gamma'}) \hookrightarrow (\beta' \hookrightarrow \beta)$$

is an instance of PL. By construction, φ_j (resp. φ_{j+1}) is a substitution instance of β (resp. β'). Hence, the formula (9) above is an instance of PL. Now $(\gamma \rightleftharpoons \gamma')$ is an instance of one of the axiom schemes NPL1, NPL2, or NPL3, and so is provable. So by modus ponens, $\varphi_{j+1} \hookrightarrow \varphi_j$ is provable. Since by induction assumption φ_{j+1} is provable, it follows by modus ponens that so is φ_j, as desired. This completes the induction step. ∎

Proof of Theorem 6.13: Soundness is easy to verify. We give a Makinson-style [Mak66] proof of completeness, which follows the same general lines as the proof of completeness of K for standard structures that is given by Halpern and Moses [HM92]. In order to prove completeness, we must show that every formula in $\mathcal{L}^{\hookrightarrow}$ that is valid with respect to $\mathcal{NM}$ is provable. We now show that it suffices to prove:

Every consistent formula in $\mathcal{L}^{\hookrightarrow}$ is satisfiable with respect to $\mathcal{NM}$. $(*)$

For suppose we can prove $(*)$, and φ is a valid formula in $\mathcal{L}^{\hookrightarrow}$. If φ is not provable, then neither is $((\varphi \hookrightarrow \mathit{false}) \hookrightarrow \mathit{false})$, since

$$(((\varphi \hookrightarrow \mathit{false}) \hookrightarrow \mathit{false}) \hookrightarrow \varphi)$$

is an instance of NPR. So by definition, $(\varphi \hookrightarrow \mathit{false})$ is consistent. It follows from $(*)$ that $\varphi \hookrightarrow \mathit{false}$ is satisfiable with respect to $\mathcal{NM}$, contradicting the validity of φ with respect to $\mathcal{NM}$.

Following the Makinson-style approach, we construct a canonical structure $M^c \in \mathcal{NM}$, which has a world s_V corresponding to every maximal consistent set V. We then show

$$(M^c, s_V) \models \varphi \text{ iff } \varphi \in V,$$

i.e., the worlds in M^c contains as elements precisely the formulas that they satisfy. Since it is easy to show that every consistent formula in $\mathcal{L}^{\hookrightarrow}$ belongs to some maximally consistent set, this is sufficient to prove $(*)$.

Given a maximal consistent set V of formulas, define $V^* = \{\varphi \in \mathcal{L}^{\hookrightarrow} \mid \neg\varphi \notin V\}$. We now show that V^* is a maximal consistent set, and that $V^{**} = V$.

V^* is consistent: If not, then there are $\varphi_1, \ldots, \varphi_k$ in V^* such that $((\varphi_1 \wedge \ldots \wedge \varphi_k) \hookrightarrow \mathit{false})$ is provable. By negation replacement, $\neg(\varphi_1 \wedge \ldots \wedge \varphi_k)$ is provable. Now the formula

$$\neg(\varphi_1 \wedge \ldots \wedge \varphi_k) \hookrightarrow (((\neg\varphi_1 \hookrightarrow false) \wedge \ldots \wedge (\neg\varphi_k \hookrightarrow false)) \hookrightarrow false)$$

is an instance of NPR, and hence provable. By modus ponens,

$$((\neg\varphi_1 \hookrightarrow false) \wedge \ldots \wedge (\neg\varphi_k \hookrightarrow false)) \hookrightarrow false$$

is provable, and hence by Lemma 6.12(1) is in V. So by consistency of V,

$$(\neg\varphi_1 \hookrightarrow false) \wedge \ldots \wedge (\neg\varphi_k \hookrightarrow false)$$

is not in V. Therefore, by Lemma 6.12(3), some $\neg\varphi_i \hookrightarrow false$ is not in V. Hence, by Lemma 6.11, $\neg\varphi_i$ is in V. Therefore, φ_i is not in V^*, a contradiction.

V^* is maximal consistent: By Lemma 6.11, we need only show that either φ or $(\varphi \hookrightarrow \mathit{false})$ is in V^*, for each formula φ of $\mathcal{L}^{\hookrightarrow}$. Assume that $\varphi \notin V^*$ and $(\varphi \hookrightarrow \mathit{false}) \notin V^*$. So by definition of V^*, it follows that the formulas $\neg\varphi$ and $\neg(\varphi \hookrightarrow \mathit{false})$ are each in V. But

$$((\neg p) \wedge \neg(p \hookrightarrow \mathit{false})) \hookrightarrow \mathit{false} \tag{10}$$

is easily seen to be valid. Hence, $((\neg\varphi) \wedge \neg(\varphi \hookrightarrow \mathit{false})) \hookrightarrow \mathit{false}$ is an instance of NPR. This shows that V is inconsistent, a contradiction.

$V^{**} = V$: We have $\varphi \in V^{**}$ iff $\neg\varphi \notin V^*$ iff $\neg\neg\varphi \in V$ iff $\varphi \in V$, where the last step uses Lemma 6.12(2) and the fact that $\varphi \hookrightarrow \neg\neg\varphi$ and $\neg\neg\varphi \hookrightarrow \varphi$ are both in V by Lemma 6.12(1).

Given a set V of formulas, define $V/K_i = \{\varphi : K_i\varphi \in V\}$. Let $M^c = (S, \pi, \mathcal{K}_1, \ldots, \mathcal{K}_n, {}^*)$, where

$$\begin{array}{rcl} S & = & \{s_V : V \text{ is a maximal } \mathrm{K}^{\hookrightarrow} \text{ consistent set}\} \\ \pi(s_V)(p) & = & \begin{cases} \mathbf{true} & \text{if } p \in V \\ \mathbf{false} & \text{if } p \notin V \end{cases} \\ \mathcal{K}_i & = & \{(s_V, s_W) : V/K_i \subseteq W\} \\ (s_V)^* & = & s_{V^*}. \end{array}$$

We show by induction on the structure of φ that for all V we have $(M^c, s_V) \models \varphi$ iff $\varphi \in V$. If φ is a primitive proposition p, then this is immediate from the definition of $\pi(s_V)$ above.

Assume that φ is a conjunction $\varphi_1 \wedge \varphi_2$. If $(M^c, s_V) \models \varphi$, then $(M^c, s_V) \models \varphi_1$ and $(M^c, s_V) \models \varphi_2$, so by induction assumption $\varphi_1 \in V$ and $\varphi_2 \in V$. Since $\varphi_1 \hookrightarrow (\varphi_2 \hookrightarrow (\varphi_1 \wedge \varphi_2))$ is an instance of NPR, it follows by Lemma 6.12(2) applied twice that $(\varphi_1 \wedge \varphi_2) \in V$. Conversely, if $(\varphi_1 \wedge \varphi_2) \in V$, then so are φ_1 and φ_2, because of the following instances of NPR and Lemma 6.12(2):

$$(\varphi_1 \wedge \varphi_2) \hookrightarrow \varphi_1$$

$$(\varphi_1 \wedge \varphi_2) \hookrightarrow \varphi_2$$

By induction assumption, $(M^c, s_V) \models \varphi_1$ and $(M^c, s_V) \models \varphi_2$, so $(M^c, s_V) \models (\varphi_1 \wedge \varphi_2)$.

Assume now that φ is of the form $\varphi_1 \hookrightarrow \varphi_2$. If $(M^c, s_V) \models \varphi$, then either $(M^c, s_V) \not\models \varphi_1$ or $(M^c, s_V) \models \varphi_2$. If $(M^c, s_V) \not\models \varphi_1$, then by induction assumption $\varphi_1 \notin V$, so by Lemma 6.11, $(\varphi_1 \hookrightarrow \mathit{false}) \in V$, so $\varphi_1 \hookrightarrow \varphi_2 \in V$ because of Lemma 6.12(2) and the fact that

$$(\varphi_1 \hookrightarrow \mathit{false}) \hookrightarrow (\varphi_1 \hookrightarrow \varphi_2).$$

is an instance of NPR. If $(M^c, s_V) \models \varphi_2$, then by induction assumption $\varphi_2 \in V$, so $\varphi_1 \hookrightarrow \varphi_2 \in V$, because of Lemma 6.12(2) and the fact that $\varphi_2 \hookrightarrow (\varphi_1 \hookrightarrow \varphi_2)$ is an instance of NPR.

Conversely, assume that $(\varphi_1 \hookrightarrow \varphi_2) \in V$. Then we cannot have both $\varphi_1 \in V$ and $(\varphi_2 \hookrightarrow \mathit{false}) \in V$, because the following instance of NPR would tell us that V is inconsistent:

$$(\varphi_1 \wedge (\varphi_2 \hookrightarrow \mathit{false}) \wedge (\varphi_1 \hookrightarrow \varphi_2)) \hookrightarrow \mathit{false}.$$

If $\varphi_1 \notin V$, then by induction assumption $(M^c, s_V) \not\models \varphi_1$, so $(M^c, s_V) \models \varphi$. If $(\varphi_2 \hookrightarrow \mathit{false}) \notin V$, then by induction assumption $(M^c, s_V) \not\models (\varphi_2 \hookrightarrow \mathit{false})$, so $(M^c, s_V) \models \varphi$.

Assume that φ is of the form $\neg\psi$. Then $(M^c, s_V) \models \neg\psi$ iff $(M^c, (s_V)^*) \not\models \psi$ iff $(M^c, s_{V^*}) \not\models \psi$ (since $(s_V)^* = s_{V^*}$) iff $\psi \notin V^*$ (by induction hypothesis) iff $\neg\psi \in V$.

Finally, assume that φ is of the form $K_i\psi$. Assume first that $\varphi \in V$. Then $\psi \in V/K_i$. So if $(s_V, s_W) \in \mathcal{K}_i$, then it follows by definition of $\mathcal{K}_i$ that $\psi \in W$, and so by induction hypothesis, $(M^c, s_W) \models \psi$. Therefore, $(M^c, s_W) \models \varphi$, as desired. For the other direction, assume that $(M^c, s_V) \models K_i\psi$. It follows that the set $(V/K_i) \cup \{\psi \hookrightarrow \mathit{false}\}$ is not consistent. For suppose not. Then it would have a maximal consistent extension W, and, by construction, we would have $(s_V, s_W) \in \mathcal{K}_i$. By the induction hypothesis we would have $(M^c, s_W) \models \psi \hookrightarrow \mathit{false}$, that is, $(M^c, s_W) \not\models \psi$, and so $(M^c, s_V) \not\models K_i\psi$, contradicting our original assumption. Since $(V/K_i) \cup \{\psi \hookrightarrow \mathit{false}\}$ is not consistent, there must be some finite subset, say $\{\varphi_1, \ldots, \varphi_k, \psi \hookrightarrow \mathit{false}\}$, which is not consistent. That is,

$$((\varphi_1 \wedge \ldots \wedge \varphi_k \wedge (\psi \hookrightarrow \mathit{false})) \hookrightarrow \mathit{false})$$

is provable. But

$$((\varphi_1 \wedge \ldots \wedge \varphi_k \wedge (\psi \hookrightarrow \mathit{false})) \hookrightarrow \mathit{false}) \hookrightarrow (\varphi_1 \hookrightarrow (\varphi_2 \hookrightarrow (\ldots (\varphi_k \hookrightarrow \psi) \ldots)$$

is an instance of NPR. Hence, by modus ponens,

$$\varphi_1 \hookrightarrow (\varphi_2 \hookrightarrow (\ldots (\varphi_k \hookrightarrow \psi) \ldots)$$

is provable. By the knowledge generalization rule,

$$K_i(\varphi_1 \hookrightarrow (\varphi_2 \hookrightarrow (\ldots (\varphi_k \hookrightarrow \psi) \ldots)$$

is provable. By Lemma 6.12(1), this formula is in V. Since $\varphi_1, \ldots, \varphi_k$ are all in $V/\mathcal{K}_i$, it follows that $K_i\varphi_1, \ldots, K_i\varphi_k$ are all in V. By repeated applications of the distribution axiom and by Lemma 6.12(1) and Lemma 6.12(2), it is easy to see that $K_i\psi \in V$, as desired. ∎

References

[AB75] A. Anderson and N. D. Belnap. *Entailment: The Logic of Relevance and Necessity.* Princeton University Press, Princeton, N.J., 1975.

[Bel77a] N. D. Belnap. How a computer should think. In *Contemporary Aspects of Philosophy*, pages 30–56. Oriel Press, 1977.

[Bel77b] N. D. Belnap. A useful four-valued logic. In G. Epstein and J. M. Dunn, editors, *Modern Uses of Multiple-Valued Logic*, pages 5–37. Reidel, Dordrecht, Netherlands, 1977.

[Che80] B. F. Chellas. *Modal Logic*. Cambridge University Press, Cambridge, U.K., 1980.

[Cre72] M. J. Cresswell. Intensional logics and logical truth. *Journal of Philosophical Logic*, 1:2–15, 1972.

[Cre73] M. J. Cresswell. *Logics and Languages*. Methuen and Co., London, 1973.

[Dun76] J. M. Dunn. Intuitive semantics for first-degree entailment and "coupled trees". *Philosophical Studies*, 29:149–168, 1976.

[Dun86] J. M. Dunn. Relevance logic and entailment. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, Vol. III*, pages 117–224. Reidel, Dordrecht, Netherlands, 1986.

[FH88] R. Fagin and J. Y. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34:39–76, 1988.

[FHV92] R. Fagin, J. Y. Halpern, and M. Y. Vardi. What is an inference rule? *Journal of Symbolic Logic*, 57(3):1018–1045, 1992.

[Gol92] R. Goldblatt. *Logics of Time and Computation*. CSLI Lecture Notes Number 7. Center for Studies in Language and Information, Stanford University, 2nd edition, 1992.

[Hal87] J. Y. Halpern. Using reasoning about knowledge to analyze distributed systems. In J. F. Traub, B. J. Grosz, B. W. Lampson, and N. J. Nilsson, editors, *Annual Review of Computer Science, Vol. 2*, pages 37–68. Annual Reviews Inc., Palo Alto, Calif., 1987.

[HC84] G. E. Hughes and M. J. Cresswell. *A Companion to Modal Logic*. Methuen, London, 1984.

[Hey56] A. Heyting. *Intuitionism: An Introduction*. North-Holland, Amsterdam, 1956.

[Hin62] J. Hintikka. *Knowledge and Belief*. Cornell University Press, Ithaca, N.Y., 1962.

[Hin75] J. Hintikka. Impossible possible worlds vindicated. *Journal of Philosophical Logic*, 4:475–484, 1975.

[HM92] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.

[Lad77] R. E. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal on Computing*, 6(3):467–480, 1977.

[Lak87] G. Lakemeyer. Tractable meta-reasoning in propositional logics of belief. In *Proc. Tenth International Joint Conference on Artificial Intelligence (IJCAI '87)*, pages 402–408, 1987.

[Lev84] H. J. Levesque. A logic of implicit and explicit belief. In *Proc. National Conference on Artificial Intelligence (AAAI '84)*, pages 198–202, 1984.

[Mak66] D. Makinson. On some completeness theorems in modal logic. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 12:379–384, 1966.

[PS89] P. F. Patel-Schneider. A four-valued semantics for terminological logics. *Artificial Intelligence*, 38:319–351, 1989.

[Ran82] V. Rantala. Impossible worlds semantics and logical omniscience. *Acta Philosophica Fennica*, 35:18–24, 1982.

[Rei78] R. Reiter. On closed world data bases. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 55–76. Plenum Press, New York, 1978.

[RK86] S. J. Rosenschein and L. P. Kaelbling. The synthesis of digital machines with provable epistemic properties. In J. Y. Halpern, editor, *Theoretical Aspects of Reasoning about Knowledge: Proc. 1986 Conference*, pages 83–97. Morgan Kaufmann, San Francisco, Calif., 1986.

[RM73] R. Routley and R. K. Meyer. The semantics of entailment, I. In H. Leblanc, editor, *Truth, Syntax, and Semantics*, pages 194–243. North-Holland, Amsterdam, 1973.

[RR72] R. Routley and V. Routley. Semantics of first degree entailment. *Noûs*, 6:335–359, 1972.

[Var86] M. Y. Vardi. On epistemic logic and logical omniscience. In J. Y. Halpern, editor, *Theoretical Aspects of Reasoning about Knowledge: Proc. 1986 Conference*, pages 293–305. Morgan Kaufmann, San Francisco, Calif., 1986.

[Wan90] H. Wansing. A general possible worlds framework for reasoning about knowledge and belief. *Studia Logica*, 49(4):523–539, 1990.