A Nonstandard Approach to the Logical Omniscience Problem

Ronald Fagin
Joseph Y. Halpern
Moshe Y. Vardi

IBM Almaden Research Center
San Jose, CA 95120-6099, USA
E-mail: fagin@ibm.com, halpern@ibm.com, vardi@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*.

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 valid formulas (including all tautologies of standard propositional logic) and they know all the logical consequences of their knowledge, so that if an agent knows p and if p logically implies q, then the agent also knows q.

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, Wan89]. 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? Notice that notions like "validity" and "logical consequence" (which played a prominent part in our informal description of logical omniscience) are not absolute notions; 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 here (since, in computing his knowledge, 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 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 possible worlds. 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.

Similarly to relevance logic [AB75], our starting point in choosing a nonstandard logic 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 [Bel77], 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, Bel77]) 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. Intuitively, s provides the support for positive formulas and s^* provides the support for negative formulas. (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 ψ is a logical consequence of φ 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 $\hookrightarrow$ in NPL so that, among other things, we have that ψ is a logical consequence of φ 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 then $\varphi \hookrightarrow \psi$ is valid exactly if φ entails ψ in relevance logic. However, in formulas with nested occurrences of $\hookrightarrow$, the semantics of $\hookrightarrow$ is quite different from the relevance logic notion of entailment.

When our nonstandard semantics is applied 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. For example, 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.

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 the next section, we review the standard possible-worlds approach. In Section 3, we describe our nonstandard approach to possible worlds and investigate some of its properties. In Section 4, we consider the logic NPL, which results from adding $\hookrightarrow$ to the syntax, and give the complete axiomatization for the logic of knowledge using NPL as a basis rather than propositional logic. We describe a concrete application of our approach in Section 5, and relate our results to those of Levesque and Lakemeyer in Section 6.

2 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 that describe basic facts about the domain of discourse. 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., $\pi(s) : \Phi \to \{\mathbf{true}, \mathbf{false}\}$ for each world $s \in S$), and $\mathcal{K}_i$ is a *binary relation* on S. 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 also use the connectives $\vee$ and $\Rightarrow$. They are defined as abbreviations: $\varphi \vee \psi$ for $\neg(\neg\varphi \wedge \neg\psi)$ and $\varphi \Rightarrow \psi$ for $\neg\varphi \vee \psi$. 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* in $\mathcal{M}$, and write $\mathcal{M} \models \varphi$, if it is valid in all structures of $\mathcal{M}$, and it is *satisfiable* in $\mathcal{M}$ 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 in $\mathcal{M}$) if and only if $\neg\varphi$ is not satisfiable in M (resp., not satisfiable in $\mathcal{M}$).

To get a sound and complete axiomatization for validity in $\mathcal{M}$, one starts with propositional reasoning and add to it axioms and inference rules for knowledge. By propositional reasoning we mean all sound propositional inference rules of propositional logic. An inference rule for a logic L is a statement of the form "from Σ infer σ", where $\Sigma \cup \{\sigma\}$ is a set of L-formulas. Such an inference rule is sound if for every substitution τ of L-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 in L, then $\tau[\sigma]$ is also valid in L. 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 [FHV89], 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.

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 2.1: *[Che80] K is a sound and complete axiomatization for validity 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 nonstandard logic in Section 4.

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

Theorem 2.2: *[Lad77] The problem of determining validity in $\mathcal{M}$ is PSPACE-complete.*

3 Nonstandard Possible Worlds

Although by now it is fairly well entrenched, standard propositional logic has several undesirable and counterintuitive properties. One problem is that material implication, where "$\varphi \Rightarrow \psi$" is taken to be simply an abbreviation for $\neg\varphi \vee \psi$, does not quite capture our intuition about what implication is. For example, the fact that $(p \Rightarrow q) \vee (q \Rightarrow p)$ is valid is quite counterintuitive, as p and q may be completely unrelated facts. 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. In the case we just discussed of a knowledge base, φ being true would mean that the fact φ has been put into the knowledge base. Since it is possible for both φ and $\neg\varphi$ to have been put in the knowledge base, it is possible for both φ and $\neg\varphi$ to be true. Similarly, if neither φ nor $\neg\varphi$ has been put into the knowledge base, then this would correspond to neither φ nor $\neg\varphi$ being true.

There are several ways to capture this intuition formally (see [Dun77]). We now discuss one approach, due to [RR72, RM73]. For each world s, there is an associated 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. We are interested in the case where $\neg\neg\varphi$ has the same truth value as of φ. To do this, we require that $s^{**} = s$ (where $s^{**} = (s^*)^*$), for each world s.

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 $\mathcal{NM}_n^\Phi$ (or by $\mathcal{NM}$ when Φ and n are clear from the context).

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 $(s, t) \in \mathcal{K}_i$.

Recall that $\varphi \vee \psi$ stands for $\neg(\neg\varphi \wedge \neg\psi)$. It can be shown that $\vee$ still behaves as disjunction, i.e., $(M, s) \models \varphi \vee \psi$ iff $(M, s) \models \varphi$ or $(M, s) \models \psi$. We still take $\varphi \Rightarrow \psi$ to be an abbreviation for $\neg\varphi \vee \psi$, but now $\Rightarrow$ does *not* behave like material implication, due to the nonstandard semantics we have given negation.

Our semantics is closely related to that of Levesque [Lev84] and Lakemeyer [Lak87]. In their semantics, they have *situations* rather than worlds. In a given situation, a primitive proposition can be either true, false, both, or neither. This gives them a way to decouple the semantics of p and $\neg p$ for a primitive proposition p. In order to decouple the semantics of $K_i\varphi$ and $\neg K_i\varphi$, Lakemeyer introduces two possibility relations, $\mathcal{K}_i^+$ and $\mathcal{K}_i^-$. There are also two versions of $\models$, denoted $\models_T$ and $\models_F$, where $\models_T$ means "supports the truth of" and $\models_F$ means "supports the falsity of".[1] We call the structures introduced by Levesque and Lakemeyer *LL structures*. Although, superficially, our semantics seems quite different from the Levesque-Lakemeyer semantics, in fact the two approaches are equivalent in the following sense. For each nonstandard structure M and world s in M, we can find an LL structure M' and world s' in M' such that for each $\mathcal{L}$-formula φ, we have that $(M, s) \models \varphi$ iff $(M', s') \models_T \varphi$ and

[1] 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 focus here on explicit knowledge, since this is the type that avoids logical omniscience. The reader who is familiar with Levesque and Lakemeyer's work should read all our references to knowledge as "explicit knowledge".

$(M,s) \models \neg\varphi$ iff $(M',s') \models_F \varphi$. Conversely, for each LL structure M and world s in M, we can find a nonstandard structure M' and world s' in M' such that for each $\mathcal{L}$-formula φ, we have $(M,s) \models_T \varphi$ iff $(M',s') \models \varphi$, and $(M,s) \models_F \varphi$ iff $(M',s') \models \neg\varphi$. Details will be given in the full paper.

We return now to examine in detail our nonstandard semantics. Note that it is possible for neither φ nor $\neg\varphi$ to be true at world s (if $(M,s) \not\models \varphi$ and $(M,s^*) \models \varphi$) and for both φ and $\neg\varphi$ to be true at world s (if $(M,s) \models \varphi$ and $(M,s^*) \not\models \varphi$). 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. A standard world s is both complete and coherent: for each formula φ exactly one of φ or $\neg\varphi$ is true at s.

Validity and logical implication for $\mathcal{NM}$ are defined in the usual way: φ is valid in $\mathcal{NM}$ if it holds in every world of every structure of $\mathcal{NM}$, and φ logically implies ψ in $\mathcal{NM}$ if $(M,s) \models \varphi$ implies $(M,s) \models \psi$ for all nonstandard structures M and worlds s in M. There are many nontrivial logical implications in $\mathcal{NM}$; for example, $\neg\neg\varphi$ logically implies φ and $\varphi \wedge \psi$ logically implies φ. What are the valid formulas in $\mathcal{NM}$? It is easy to verify that certain tautologies of standard propositional logic are not valid. For example, the formula $(p \Rightarrow q) \vee (q \Rightarrow p)$, whose validity in standard propositional logic disturbed us, is not valid anymore. The formula $(p \wedge \neg p) \Rightarrow q$, which wreaked havoc in deriving consequences from a knowledge base, is not valid either. What about even simpler tautologies of standard propositional logic, such as $p \vee \neg p$? 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 with respect to nonstandard structures! Even more, there is a single counterexample that simultaneously shows that no formula is valid!

Theorem 3.1: *There is no formula of $\mathcal{L}$ that is valid in nonstandard structures. 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 $\pi(s)(p) =$ **false** for every primitive proposition p, and letting $\pi(t)$ be the truth assignment where $\pi(t)(p) =$ **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 nonstandard structures. ∎

It follows from Theorem 3.1 that the validity problem with respect to nonstandard structures is *very* easy: the answer is always "No, the formula is not valid!" The reader may be puzzled why there are no valid formulas. For example, $\neg\neg\varphi$ logically implies φ, as noted earlier. Doesn't this mean that $\neg\neg\varphi \Rightarrow \varphi$ is valid? This does not follow. With *standard* structures, φ logically implies ψ iff the formula $\varphi \Rightarrow \psi$ is valid. This is not the case for *nonstandard* structures. For example, φ logically implies φ, yet $\varphi \Rightarrow \varphi$ is not valid with respect to nonstandard structures. In the next section, we define a new connective that allows us to express logical implication *in the language*, just as $\Rightarrow$ does for standard structures.

What about logical omniscience? It 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, however, 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 material implication; an agent may know φ and $\varphi \Rightarrow \psi$ without knowing ψ, since φ and $\varphi \Rightarrow \psi$ do not logically imply ψ in $\mathcal{NM}$.

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 3.2: *The logical implication for propositional $\mathcal{L}$-formulas in nonstandard structures is co-NP-complete, and the logical implication for $\mathcal{L}$-formulas in nonstandard structures is PSPACE-complete.*

Theorem 3.2 asserts that nonstandard logical implication is as hard as standard validity; that is, it is co-*NP*-complete for propositional formulas and *PSPACE*-complete for knowledge formulas (i.e., $\mathcal{L}$-formulas).

4 Strong Implication

Certain classic tautologies, such as $(p \Rightarrow q) \vee (q \Rightarrow p)$ made us uncomfortable. In the previous section, we introduced nonstandard structures and—lo and behold!—under this approach, these formulas are no longer valid. However, 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 (in fact, no formula is valid). It seems that we have thrown out the baby with the bath water.

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 $\neg\varphi_1 \vee \varphi_2$, which says that if $\neg\varphi_1$ is false, then φ_2 is true. In standard propositional logic, these are the same, since $\neg\varphi_1$ is false in standard logic iff φ_1 is true. However, in nonstandard structures, these are not equivalent. So let us define 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,

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

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

We denote by $\mathcal{L}_n^{\Phi,\hookrightarrow}$, or $\mathcal{L}^{\hookrightarrow}$ for short, the set of formulas obtained by replacing $\Rightarrow$ by $\hookrightarrow$ in formulas of $\mathcal{L}_n^{\Phi}$. We call the propositional fragment of $\mathcal{L}^{\hookrightarrow}$ and its interpretation by nonstandard structures *nonstandard propositional logic* (NPL).

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 $\varphi_1 \hookrightarrow (\varphi_1 \vee \varphi_2)$.

Strong implication is indeed stronger than implication, in the sense that if φ_1 and φ_2 are standard formulas, and 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, 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 nonstandard propositional logic. (We note also that the analogue to the distressing propositional tautology $(p \Rightarrow q) \vee (q \Rightarrow p)$, namely $(p \hookrightarrow q) \vee (q \hookrightarrow p)$, is not valid in nonstandard propositional logic.)

As we promised in the previous section, we can now express logical implication in $\mathcal{L}^{\hookrightarrow}$, using $\hookrightarrow$, just as we can express logical implication in standard structures, using $\Rightarrow$.

Proposition 4.1: *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 in 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*. In particular, 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][2] exactly if $\varphi_1 \hookrightarrow \varphi_2$ is valid in $\mathcal{NM}$ (or equivalently, φ_1 logically implies φ_2 in $\mathcal{NM}$). However, in formulas with nested occurrences of $\hookrightarrow$, the semantics of $\hookrightarrow$ is quite different from that of relevant entailment. In particular, while $p \hookrightarrow (q \hookrightarrow p)$ is valid in $\mathcal{NM}$, the analogous formula $p \rightarrow (q \rightarrow p)$ is not a theorem of relevance logic.

In $\mathcal{L}$, we cannot say that a formula φ is false. That is, there is no formula ψ such that $(M,s) \models \psi$ iff $(M,s) \not\models \varphi$. This is because no formula is true at the world t of Theorem 3.1, and so no $\psi \in \mathcal{L}$ can do the job, for any formula $\varphi \in \mathcal{L}$. 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 what follows, we add to $\mathcal{L}$ and $\mathcal{L}^{\hookrightarrow}$ the abbreviations $true$ and $false$. In $\mathcal{L}$, we take *true* to be an abbreviation for some fixed standard tautology such as $p \Rightarrow p$, while in $\mathcal{L}^{\hookrightarrow}$, 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.

Proposition 4.2: *Let M be a nonstandard structure, and let s be a world of M. Then $(M,s) \not\models \varphi$ iff $(M,s) \models \varphi \hookrightarrow false$.*

A close examination of all the constructs in our logic shows that in fact the only nonstandard connective is $\neg$; all other connectives "behave" standardly. We now formalize this observation by considering certain transformations on formulas and structures.

Let M be a nonstandard structure. We define M^{st}, the *standardization* of M, to be the structure obtained by replacing the $*$ of M by the the identity function. Note that if M is

[2] 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 [Dun77] for an axiomatization of first-degree entailments.

standard then $M^{st} = M$. Let φ be a standard formula. We define a nonstandard formula φ^{nst} by recursively replacing in φ all subformulas of the form $\neg\psi$ by $\psi \hookrightarrow false$ and all occurences of $\Rightarrow$ by $\hookrightarrow$. Note that φ^{nst} is negation free. We also define what is essentially the inverse transformation on negation-free nonstandard formulas. Let φ be a nonstandard negation-free formula. We define a standard formula φ^{st} by replacing in φ all occurences of $\hookrightarrow$ by $\Rightarrow$. Notice that the transformations ^{nst} and st are inverses when restricted to negation-free formulas.

Proposition 4.3: *Let M be a nonstandard structure, let s be a world of M, and let φ be a standard formula. Then $(M, s) \models \varphi^{nst}$ iff $(M^{st}, s) \models \varphi$.*

Corollary 4.4: *Let φ be a standard formula. Then φ is valid in standard structures iff φ^{nst} is valid in nonstandard structures.*

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 φ 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 $\neg(p \vee \neg p)$ 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 4.5: *Let φ be a pseudo-positive formula. Then φ is valid in nonstandard structures iff φ^{+} is valid in nonstandard structures.*

Corollary 4.6: *Let φ be a pseudo-positive formula. Then φ is valid in nonstandard structures iff $(\varphi^{+})^{st}$ is valid in standard structures.*

We can use the above facts to obtain an axiomatization of NPL. To prove that a propositional formula ψ_1 in $\mathcal{L}^{\hookrightarrow}$ is valid, we first drive negations down until they apply only to primitive propositions, by applying the following equivalences: (a) $\neg\neg\varphi$ is equivalent to φ, (b) $\neg(\varphi \hookrightarrow \psi)$ is equivalent to $((\neg\psi \hookrightarrow \neg\varphi) \hookrightarrow false)$, and (c) $\neg(\varphi \wedge \psi)$ is equivalent to $(\neg\varphi \hookrightarrow false) \hookrightarrow \neg\psi$. This gives us a pseudo-positive formula ψ_2 equivalent to ψ_1. By Corollary 4.6, it then suffices to prove that $(\psi_2^{+})^{st}$ is valid in standard structures.

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 formulas φ^{nst}, where φ is a valid formula of standard propositional logic

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)

Again, one should view the axioms and inference rules above as *schemes*, i.e., N actually consists of all propositional $\mathcal{L}^{\hookrightarrow}$-instances of the above axioms and inference rules.

Theorem 4.7: *N is a sound and complete axiomatization for NPL.*

We note that PL can be replaced by the nonstandard version of any complete axiomatization of standard propositional logic (i.e, by applying the ^{nst} operator to a complete axiomatization of standard propositional logic).

What is a sound and complete axiomatization for the full nonstandard logic? Interestingly, it 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)

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, consider the rule "from $\neg\varphi$ infer $\varphi \hookrightarrow \mathit{false}$", which we call *negation replacement.* It is not hard to verify that for any nonstandard formula φ, if $\neg\varphi$ is valid in nonstandard structures, then $\varphi \hookrightarrow \mathit{false}$ is also valid in nonstandard structures. Thus, negation replacement is a sound NPL inference rule. On the other hand, $\varphi \hookrightarrow \mathit{false}$ is clearly not a logical consequence of $\neg\varphi$ in nonstandard structures. 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 [FHV89].

Theorem 4.8: $K^{\hookrightarrow}$ *is a sound and complete axiomatization with respect to* $\mathcal{NM}$ *for formulas in the language* $\mathcal{L}^{\hookrightarrow}$.

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; if we replace NPR by all valid formulas of NPL with modus ponens as the sole propositional inference rule, then the resulting system would *not* be complete. It can shown, however, that NPR can be replaced by any complete axiomatization of NPL that includes modus ponens and negation replacement as inference rules. We discuss the details in the full paper.

The reader should note the similarity between the axiom system K for knowledge in standard Kripke structures and the nonstandard system $K^{\hookrightarrow}$. The latter is obtained from the former by replacing the inference rules for standard propositional logic by inference rules for nonstandard propositional logic and by replacing $\Rightarrow$ by $\hookrightarrow$ in the distribution axiom. Thus, one can say that in our approach agents are "nonstandardly" logically omniscient.

Since $\hookrightarrow$ can capture logical implication it is easy to see that our lower bound results for logical implication in the language $\mathcal{L}$ from Section 3 translate immediately to results on validity for the language $\mathcal{L}^{\hookrightarrow}$. We can show that these bounds are tight.

Theorem 4.9: *The validity problem for propositional* $\mathcal{L}^{\hookrightarrow}$*-formulas in* $\mathcal{NM}$ *is co-NP-complete and the validity problem for* $\mathcal{L}^{\hookrightarrow}$*-formulas in* $\mathcal{NM}$ *is PSPACE-complete.*

5 A Concrete Example

An interesting application of our approach is in the situation 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. Secondly, we can ask whether knowledge of φ follows from knowledge of κ. Fortunately, these are equivalent questions, as we now see.

Proposition 5.1: *Let φ_1 and φ_2 be $\mathcal{L}$-formulas. Then φ_1 logically implies φ_2 in $\mathcal{NM}$ iff $K_i\varphi_1$ logically implies $K_i\varphi_2$ in $\mathcal{NM}$.*

The problem of determining the consequences of a knowledge base (whether κ logically implies φ, or equivalently, by Proposition 5.1, whether $K_i\kappa$ logically implies $K_i\varphi$) is co-*NP*-complete, by Theorem 3.2, even when the database is propositional. However, there is an interesting special case where the problem is not hard.

Define a *clause* to be a disjunction of literals. For example, a typical clause is $p \vee \neg q \vee r$. Suppose that the knowledge base consists of a finite collection of clauses. Thus, κ 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*). Every standard propositional formula is equivalent to a formula in CNF (this is true even in our nonstandard semantics).

We now consider the question of whether κ logically implies another clause φ. In standard propositional logic, this problem is no easier than the general problem of logical implication in propositional logic, that is, co-*NP*-complete. By contrast, there is a polynomial-time decision procedure for deciding whether κ logically implies φ in nonstandard propositional logic. In fact, even when φ is a CNF formula (rather than just a clause), there is a polynomial-time decision procedure for deciding whether κ logically implies φ in nonstandard propositional logic. In particular, the task of computing whether a set of clauses logically implies another clause (and whether an agent's knowledge of a set of clauses logically implies his knowledge of another clause) is feasible.

Theorem 5.2: *There is a polynomial-time decision procedure for deciding whether κ logically implies φ in nonstandard propositional logic (or $K_i\kappa$ logically implies $K_i\varphi$ with respect to nonstandard structures), for CNF formulas κ and φ.*

Theorem 5.2 follows from results in [Lev84]. The precise relationship to Levesque's results will be clarified in the next section.

6 Standard-World Validity

Recall that a world s of a nonstandard structure $M = (S, \pi, \mathcal{K}_1, ..., \mathcal{K}_n, ^*)$ is *standard* if $s = s^*$. In a standard world, negation behaves classically, because at a standard world s, we have $(M,s) \models \neg\varphi$ iff $(M,s) \not\models \varphi$. As mentioned in the introduction, in the impossible worlds approach there is a distinction between standard and nonstandard worlds [Cre72, Cre73, Lev84, Ran82]. According to this approach, although an agent might consider a nonstandard world possible, the real world must be standard. Consequently, validity and logical implication are defined with respect to 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.

At standard worlds, implication ($\Rightarrow$) behaves as it does in standard Kripke structures: that is, $\varphi_1 \Rightarrow \varphi_2$ holds at a standard world precisely if it is the case that if φ_1 holds, then φ_2 holds. We now have the following analogue to Proposition 4.1.

Proposition 6.1: *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.*

What about logical omniscience? Although the classical tautology $\varphi \vee \neg\varphi$ 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. So agents do not necessarily know all standard-world valid formulas. The reason for this lack of knowledge is the inability of the agent to distinguish between complete and incomplete worlds.

Let φ be a propositional 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, at least as far as the primitive propositions in φ are concerned. If φ is a standard propositional tautology, then knowledge of Complete(φ) implies knowledge of φ. The next theorem follows from the results in [FH88].

Theorem 6.2: *Let φ be a tautology of standard propositional logic. Then $K_i(Complete(\varphi))$ logically implies $K_i\varphi$ in nonstandard structures.*

A similar phenomenon occurs with regard to closure of knowledge under material implication. 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 the inability to distinguish between coherent and incoherent worlds; indeed, it is shown in [FH88] 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 or she knows that either ψ holds or the world is incoherent.

Let φ be a propositional 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, at least as far as the primitive propositions in φ are concerned. (Note that Coherent(φ) is not definable in $\mathcal{L}$ but only in $\mathcal{L}^{\hookrightarrow}$.) Knowledge of coherence implies closure of knowledge under implication.

Theorem 6.3: *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.*

Propositions 6.2 and and 6.3 explain why the agents in Levesque's model [Lev84] are not logically omniscient: "logically" is defined there with respect to standard worlds, but the agents cannot distinguish standard from nonstandard worlds. If an agent's knowledge includes the distinction between standard and nonstandard worlds, i.e., we have the antecedents $K_i(Complete(\varphi))$ and $K_i(Coherent(\varphi))$ of Theorems 6.2 and 6.3, then this agent is logically omniscient.

Let us reconsider 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 5.1) that in the nonstandard approach φ is a consequence of κ precisely when knowledge of φ is a consequence of knowledge of κ. Furthermore, implication of knowledge coincides in the standard and nonstandard approaches.

Proposition 6.4: *Let φ_1 and φ_2 be $\mathcal{L}$-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.*

On the other hand, the two interpretations of query evaluation differ in the standard approach. In contrast to Proposition 5.1, 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 $K_i\varphi_2$. 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. In this case, testing whether κ standard-world logically implies φ is co-NP-complete, while testing whether $K_i\kappa$ standard-world logically implies $K_i\varphi$ can be done in polynomial time by Theorem 5.2 and Proposition 6.4. (In fact, Levesque proved the latter result in [Lev84], from which we obtained Theorem 5.2 using Proposition 6.4.)

References

[AB75] A. Anderson and N. D. Belnap, *Entailment: the logic of relevance and necessity*, Princeton University Press, 1975.

[Bel77] N. D. Belnap, A useful four-valued logic, *Modern Uses of Multiple-Valued Logic* (G. Epstein and J. M. Dunn, eds.), Reidel, 1977, pp. 5–37.

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

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

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

[Dun77] J. M. Dunn, Relevance logic and entailment, *Handbook of Philosophical Logic, Vol. III* (D. Gabbay and F. Guenthner, eds.), Reidel, 1977, pp. 117–224.

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

[FHV89] R. Fagin, J. Y. Halpern, and M. Y. Vardi, What is an inference rule?, 1989.

[Hal87] J. Y. Halpern, Using reasoning about knowledge to analyze distributed systems, *Annual Review of Computer Science, Vol. 2* (J. Traub et al., ed.), Annual Reviews Inc., 1987, pp. 37–68.

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

[Hin75] J. Hintikka, Impossible worlds vindicated, *Journal of Philosophy* **4**, 1975, pp. 475–484.

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

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

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

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

[RK86] S. J. Rosenschein and L. P. Kaelbling, The synthesis of digital machines with provable epistemic properties, *Theoretical Aspects of Reasoning about Knowledge: Proceedings of the 1986 Conference* (J. Y. Halpern, ed.), Morgan Kaufmann, 1986, pp. 83–97.

[RM73] R. Routley and R. K. Meyer, The semantics of entailment, I, *Truth, Syntax, and Semantics* (H. Leblanc, ed.), North-Holland, 1973, pp. 194–243.

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

[Var86] M. Y. Vardi, On epistemic logic and logical omniscience, *Theoretical Aspects of Reasoning about Knowledge: Proceedings of the 1986 Conference* (J. Y. Halpern, ed.), Morgan Kaufmann, 1986, pp. 293–305.

[Wan89] H. Wansing, A general possible worlds framework for reasoning about knowledge and belief, 1989. Unpublished manuscript.