

A TRACTABLE KNOWLEDGE REPRESENTATION SERVICE WITH FULL INTROSPECTION

Gerhard Lakemeyer
Hector J. Levesque¹

Department of Computer Science
University of Toronto
Toronto, Ontario
Canada, M5S 1A4

gerhard@toronto.csnet
hector@toronto.csnet

Abstract

A Knowledge Representation service for a knowledge-based system (or agent) can be viewed as providing, at the very least, two operations that (a) give precise information about what is and is not believed (ASK) and (b) add new facts to the knowledge base when they become available (TELL). An appropriate model of belief for such operations should support the notion that *only* certain facts are believed, in particular those that have been added to a knowledge base via TELL. For logically omniscient and fully introspective agents, models of this kind lead to intractable ASK and TELL operations. In this paper, we show that tractability can be retained by giving up logical omniscience, but without sacrificing full introspection. This is done within the framework of a propositional logic of belief. In particular, the logic allows us to express that *only* a sentence (or finite set of them) is believed. We show that the validity of certain classes of sentences involving belief can be decided efficiently. These results are then applied to the specification of efficient TELL and ASK operations.

¹Fellow of The Canadian Institute for Advanced Research

1 Introduction

A Knowledge Representation service for a knowledge-based system (or agent) can be viewed as providing, at the very least, two operations that (a) give precise information about what is and is not believed (ASK) and (b) add new facts to the knowledge base when they become available (TELL) [8]. The meaning of both operations rests squarely with the model of belief adopted. What properties should it have?

Clearly, as a minimal requirement, all of the facts that have been added, which we take to be sentences in some logical language, and which make up a knowledge base, should be believed. In a very intuitive sense, this is also *all* that should be believed.

For example, assume the system has been told a single fact ‘John likes Mary or Sue’ (*). When we ask it whether it believes that John likes Mary we would expect it to answer NO, since intuitively there seems to be nothing that could justify such a belief. What we are saying is that any belief of the system should be somehow derivable from (*). ‘John likes Mary’ certainly isn’t and should therefore not be believed. Notice, however, that this argument relies on the assumption that (*) is all that is believed. In particular, the fact that ‘John likes Mary’ is not believed is *not* a logical consequence of believing (*).

The example suggests that under the assumption of believing only what is in the knowledge base or, as we will say, *only-believing* it, the knowledge base alone completely determines what the agent believes and does not believe (the agent’s epistemic state). This would then give a very intuitive interpretation to the operations ASK and TELL. For example, asking a query α reduces, at least conceptually, to simply testing whether α is believed at the current epistemic state.

Despite the intuition, a precise characterization of what it means for an agent to only-believe a sentence (or set of them) is far from obvious. For perfect reasoners with full introspection, this question has been studied in some detail [12,13,3,8,10,11]. In all cases, the underlying model of belief is possible-world semantics. The basic idea in possible-world semantics is that there are different ways the world can be (called possible worlds), which determine the truth or falsity of sentences in a complete and coherent fashion. An agent may not know exactly the way the world really is, but only that it lies within some range of possibilities. These possible worlds are then said to be accessible to the agent, and an agent at a world w is said to believe a sentence α , if it is true at all worlds accessible from w . To model a perfect and fully introspective reasoner, it suffices to have a set of worlds that are globally accessible from every world. In other words, the agent imagines the same set of worlds in every world. The corresponding belief logic is called weak S5 [4].

As discussed in [10], the idea of only-believing a sentence α can be modelled by a set of worlds that maximizes ignorance, while still supporting the truth of α in all worlds. Maximizing ignorance essentially means making the set of accessible worlds as large as possible, because the more worlds one thinks possible, the more uncertain one is about the way the world really is.

If we call a sentence *objective* when it does not mention belief, then only-believing objective sentences turns out to determine such sets uniquely. Starting with a knowledge base of this kind, ASK and TELL turn out to be definable themselves in purely objective terms [8]. In particular,

adding a new sentence will always result in a knowledge base representable by an objective sentence, even if the new assertion mentions beliefs.

Unfortunately, the assumption of a perfect reasoner (or *logical omniscience* as it is sometimes called) leads to specifications of TELL and ASK that are not effectively computable for first-order logic, and intractable in a propositional setting. This makes the claim of a *specification* for a KR service somewhat less convincing. Weaker logics of belief have recently been studied such as [9] or [2], which defeat logical omniscience in one way or another. To distinguish these approaches, will use the term *implicit* for beliefs of perfect reasoners (e.g. closed under logical implication) and *explicit* for beliefs under limited reasoning. In addition, we are concerned with only-believing and will talk about *implicitly only-believing* (in the sense of [10]) and *explicitly only-believing*, which is the new concept being formalized here.

In this paper, we focus on a logic of explicit belief with full introspection, which serves as a basis for the specification of *tractable* ASK and TELL operations. The notion of explicitly only-believing is the basis for the two operations and allows us to discuss the main complexity results as properties of the logic. The major result we show (theorem 11) is that, once we eliminate logical omniscience, it is possible to handle full introspection in a tractable way.

The paper is organized as follows. We lay out the semantics of the logic of explicit belief and briefly discuss its properties. In the following section we address computational properties of belief with an emphasis on deriving beliefs that follow from explicitly only-believing a knowledge base. We then show how the ASK and TELL operations are specified and derive their complexity from results of the previous section.

Due to space constraints we include mainly proof sketches and omit some altogether. A detailed treatment is in preparation [7].

2 The language \mathcal{L} and its semantics

The language \mathcal{L} we consider consists of sentences over a countably infinite set of propositional letters (or atoms) \mathcal{P} together with the logical connectives \neg and \vee and two modal operators **B** and **O**. $\mathbf{B}\alpha$ is read as “the agent explicitly believes α ” and $\mathbf{O}\alpha$ as “the agent explicitly only-believes α ” or “ α is all that is explicitly² believed”.³ It is also convenient to include two special atoms **true** and **false** with their obvious intended meaning. A sentence is called *objective* if it does not contain any **B** or **O**, *basic* if it does not contain any **O**, and *subjective* if every atom other than **true** or **false** occurs within the scope of a **B** or **O**.

The semantics is essentially an extension of weak **S5** models. An agent’s beliefs are modelled by a set of states of affairs which is globally accessible. One key difference to standard possible-world semantics is the use of *situations*, as described in [9], instead of worlds. Situations, in contrast to worlds, need no longer be complete and coherent. That is, instead of every proposition being either true or false, a proposition may have neither its truth nor its falsity supported or, in the other extreme, it may have both true and false support. This style of

²We usually omit the term “explicit” from now on. Whenever we mention belief, we mean explicit belief unless otherwise noted.

³We will also use $\alpha \wedge \beta$, $\alpha \supset \beta$, and $\alpha \equiv \beta$ as syntactic abbreviations.

semantics is strongly related to a four-valued semantics for *tautological entailment*, a form of *relevance logic* [1].

Another difference between possible-world semantics and this logic is that for an agent to believe a sentence α , it is not sufficient for α to have true support in all situations the agent thinks possible. In addition, we require the agent to be “aware” of all atomic propositions occurring in α . We say that an agent is aware of an atomic proposition p , if every situation the agent thinks possible either supports the truth or the falsity of p . This notion of awareness was also suggested by Fagin and Halpern [2] in their discussion of [9]. However, later in their paper, when they present their own logic of awareness, they use a separate awareness function, which leads to somewhat different properties.

Finally, the interpretation of only-believing is quite similar to the treatment in [11]. Intuitively, $\mathbf{O}\alpha$ is true if the set of situations the agent thinks possible is as large as it can be. Some care must be taken so that this set also guarantees the awareness of α in the sense mentioned above.

More formally, a situation s is a tuple $\langle T, F \rangle$, where both T and F are subsets of \mathcal{P} with **true** always being a member of T and **false** a member of F . Worlds are, of course, those situations, where T and F partition \mathcal{P} , that is, every proposition is either true or false, but not both.

We now consider what it means for the truth or falsity of sentences to be supported. Because of subjective sentences to be supported, we need to specify the set of situations M that the agent thinks are possible and which we assume to be non-empty. For objective sentences, we need a single situation s to say how the world really is. We can now define what it means for M and s to support the truth (\models_T) or falsity (\models_F) of sentences in \mathcal{L} . Atoms and the propositional connectives are treated as in classical logic except that the rules are a bit more long-winded, because we need to talk about true and false support separately:

1. $M, s \models_T p \iff p \in T$ with $p \in \mathcal{P}$ and $s = \langle T, F \rangle$
 $M, s \models_F p \iff p \in F$
2. $M, s \models_T \neg\alpha \iff M, s \models_F \alpha$
 $M, s \models_F \neg\alpha \iff M, s \models_T \alpha$
3. $M, s \models_T \alpha \vee \beta \iff M, s \models_T \alpha$ or $M, s \models_T \beta$
 $M, s \models_F \alpha \vee \beta \iff M, s \models_F \alpha$ and $M, s \models_F \beta$

As we have said earlier, believing a sentence α means not only that α has true support at all situations the agent imagines, but also that the agent is aware of all atomic propositions p occurring in α . Being aware of p simply says that either p or $\neg p$ has true support at all accessible situations. If we let $det(\alpha)$ denote the conjunction of $(p \vee \neg p)$ over all atoms p occurring in α , we obtain the following semantic rule for **B**:

4. $M, s \models_T \mathbf{B}\alpha \iff$ for all t , if $t \in M$ then $M, t \models_T (\alpha \wedge det(\alpha))$
 $M, s \models_F \mathbf{B}\alpha \iff M, s \not\models_T \mathbf{B}\alpha$

Finally, as for only-believing, the intuition is that for $\mathbf{O}\alpha$ to have true support α should be believed and the set of possible situations accessible to the agent should be as large as it can

be. Both constraints seem to be achievable by simply replacing the “then” in rule 4 by an “iff” so that we get:

$$\begin{aligned} 5. \quad M, s \models_{\mathbf{T}} \mathbf{O}\alpha &\iff \text{for all } t, t \in M \text{ iff } M, t \models_{\mathbf{T}} (\alpha \wedge \mathit{det}(\alpha)) \\ M, s \models_{\mathbf{F}} \mathbf{O}\alpha &\iff M, s \not\models_{\mathbf{T}} \mathbf{O}\alpha \end{aligned}$$

However, this is not quite right. The problem is that there are far more sets of situations than there are distinct sets of basic sentences that can be believed.⁴ This has the effect that there are sets of situations that agree exactly on their basic beliefs but differ in what they only-believe. For example, consider the set $M_1 = \{s \mid s \models_{\mathbf{T}} p\}$. One can easily verify that M_1 only-believes p . Now remove any situation s^* from M_1 and call the resulting set M_2 . It is not hard to show that M_1 and M_2 agree on all basic beliefs, but, according to rule 5, M_2 does not only-believe p , since M_2 is missing s^* , which supports the truth of p . To remedy the situation, we have to make sure that only-believing does not distinguish between sets like M_1 and M_2 . One way of achieving this is to consider the largest superset M^+ of a set M that has exactly the same beliefs as M .

The right choice⁵ turns out to be $M^+ = \{s \mid \text{for all objective } \alpha, M, s \models_{\mathbf{T}} \mathbf{B}\alpha \supset \alpha\}$.⁶ The correct version of rule 5 is now as follows:

$$\begin{aligned} 5. \quad M, s \models_{\mathbf{T}} \mathbf{O}\alpha &\iff \text{for all } t, t \in M^+ \text{ iff } M^+, t \models_{\mathbf{T}} (\alpha \wedge \mathit{det}(\alpha)) \\ M, s \models_{\mathbf{F}} \mathbf{O}\alpha &\iff M, s \not\models_{\mathbf{T}} \mathbf{O}\alpha \end{aligned}$$

This completes the semantics of \mathcal{L} . The notions of *truth*, *satisfiability*, and *validity* are defined with respect to worlds and sets of situations: we say that a sentence α is true at a world w and a set of situations M if $M, w \models_{\mathbf{T}} \alpha$, otherwise α is false; α is satisfiable if it is true for some w and M ; α is valid ($\models \alpha$) if α is true at all worlds and sets of situations M . Since the truth of subjective sentences depends only on M , we use $M \models_{\mathbf{T}} \sigma$ as shorthand for $M, s \models_{\mathbf{T}} \sigma$ for any situation s . Similarly, for objective sentences α we write $s \models_{\mathbf{T}} \alpha$ instead of $M, s \models_{\mathbf{T}} \alpha$. Finally, $\mathbf{A}\alpha$ (read “the agent is aware of α ”) is used as an abbreviation for $\mathbf{B}\mathit{det}(\alpha)$.

3 Properties of the Logic

It is easy to see that outside a belief context, the logic behaves just like any ordinary propositional logic. For example, all substitution instances of propositional tautologies are valid. In the rest of this section, we explore the properties of $\mathbf{B}\alpha$ and $\mathbf{O}\alpha$ before turning to the computational aspects of the logic.

3.1 Explicit Belief

Awareness, which is really a special kind of belief, has the property that it is impossible for an agent to believe not to be aware of something ($\models \neg \mathbf{B}\neg \mathbf{A}\alpha$).⁷ While this is similar for Fagin

⁴See [8] for more on this with respect to possible-world semantics.

⁵The proof that this set is the correct choice is presented in the next section.

⁶Note that even though M^+ is defined in terms of $\models_{\mathbf{T}}$, there is no danger of circularity, because the interpretation of $\mathbf{B}\alpha \supset \alpha$ involves only rules 1–4.

⁷This suggests that awareness should be understood as “thinking about” rather than “having knowledge of”.

and Halpern's concept in their logic of awareness [2], there is also an important difference. For example, $\mathbf{B}(p \vee q) \supset (\mathbf{A}p \wedge \mathbf{A}q)$ is valid in our logic but not in theirs, where an agent need only be aware of either p or q . Another way of viewing our version of awareness is that it fixes a vocabulary or a subject matter, which restricts the scope of what beliefs can be about.

Logical omniscience is overcome in many ways. The only valid belief is $\mathbf{B}\text{true}$. Any other belief can fail because the agent may not be aware of it. For example: $\neg\mathbf{B}(p \vee \neg p)$ and $\neg\mathbf{B}(\mathbf{B}p \vee \neg\mathbf{B}p)$ are satisfiable. Due to the existence of incoherent situations, beliefs about the world share many properties with the explicit beliefs of [9]. For instance, they are not closed under modus ponens and inconsistent beliefs can be held without believing everything. (This is not the case for meta-beliefs, however!) Differences between our version of explicit belief and that of [9] are due to our concept of awareness. For example, simple disjunctive weakening like $\mathbf{B}p \supset \mathbf{B}(p \vee q)$ is no longer valid in our logic. Instead we have $\models (\mathbf{B}p \wedge \mathbf{A}q) \supset \mathbf{B}(p \vee q)$. Similarly $(\mathbf{B}\alpha \vee \mathbf{B}\beta) \supset \mathbf{B}(\alpha \vee \beta)$ is not valid, whereas $(\mathbf{A}\alpha \wedge \mathbf{A}\beta) \supset ((\mathbf{B}\alpha \vee \mathbf{B}\beta) \supset \mathbf{B}(\alpha \vee \beta))$ is. In short, the subject matter must be compatible before a belief can follow from another belief. If that is the case, however, an agent can do better than in [9]. In particular, the agent is able to figure out all those tautologies that have to do with the application domain, or, more formally, $\models \mathbf{B}\alpha \supset \mathbf{B}(\text{taut}_\alpha)$, where taut_α is any tautology made up of atoms mentioned in α .

As to the introspective capabilities, if an agent is aware of a *subjective* sentence σ , then σ is true exactly when it is believed: $\models \mathbf{A}\sigma \supset (\sigma \equiv \mathbf{B}\sigma)$. This property indicates that introspection in this logic is very similar to introspection in the logic weak S5 with consistency. As a consequence we get positive introspection ($\models \mathbf{B}\alpha \supset \mathbf{B}\mathbf{B}\alpha$), negative introspection with the proviso of awareness ($\models \mathbf{A}\alpha \supset (\neg\mathbf{B}\alpha \supset \mathbf{B}\neg\mathbf{B}\alpha)$), and consistency of beliefs in subjective sentences ($\models \mathbf{B}\sigma \supset \neg\mathbf{B}\neg\sigma$).

3.2 Only-Believing

The beliefs that follow from explicitly only-believing a sentence are, not surprisingly, in general much more constrained than what follows from implicitly only-believing. For example, if p is an atom, then $\mathbf{O}p$ is satisfied by $M = \{s \mid s \models_{\mathbf{T}} p\}$. In this case, the agent is just aware of p and nothing is believed about sentences containing atoms other than p . In fact, $\mathbf{O}p \supset \neg\mathbf{B}\neg\mathbf{B}q$ is valid, which is not the case under implicitly only-believing. What does follow from $\mathbf{O}p$ are sentences like $\mathbf{B}\mathbf{B}p$, $\mathbf{B}\neg\mathbf{B}\neg p$, $\mathbf{B}\neg\neg\mathbf{B}p$, and $\mathbf{B}(\mathbf{B}\neg\neg p \wedge \mathbf{O}p)$.

3.2.1 Maximal Sets

Before we go on, we should take a careful look at how the semantics of $\mathbf{O}\alpha$ has been defined. In particular, we must convince ourselves that the choice of M^+ (*a*) solves the problem we discovered with the first attempt to define the meaning of \mathbf{O} and (*b*) preserves the beliefs of the original M . Concerning (*a*) we will show that for any two sets M_1 and M_2 that agree on all objective beliefs, $M_1^+ = M_2^+$. In a sense, M^+ is a unique representative of all sets that have the same beliefs as M .

Theorem 1 *Two sets of situations M_1 and M_2 agree on all their objective beliefs iff $M_1^+ = M_2^+$.*

Proof : Let $M_1^+ = M_2^+$ and assume wlog that for some objective α , $M_1 \models_T \mathbf{B}\alpha$ and $M_2 \not\models_T \mathbf{B}\alpha$. Then for some $s \in M_2$, $s \not\models_T \alpha$. Therefore $M_1, s \not\models_T \mathbf{B}\alpha \supset \alpha$ and thus $s \notin M_1^+$. However, it is easy to show that $M_2 \subseteq M_2^+$, which implies that $s \in M_1^+$, a contradiction.

On the other hand, let M_1 and M_2 agree on all objective beliefs. Let $s \in M_1^+$. Then for all objective α , $M_1, s \models_T \mathbf{B}\alpha \supset \alpha$, i.e., if $M_1 \models_T \mathbf{B}\alpha$ then $s \models_T \alpha$. But $M_1 \models_T \mathbf{B}\alpha$ iff $M_2 \models_T \mathbf{B}\alpha$ by assumption. Therefore, $M_2, s \models_T \mathbf{B}\alpha \supset \alpha$ for all objective α . Thus $s \in M_2^+$, which proves $M_1^+ \subseteq M_2^+$. By the same argument, one can show $M_2^+ \subseteq M_1^+$ and the theorem follows. ■

Lemma 3.1 M and M^+ agree on all objective beliefs.

Proof : Easy exercise. ■

We call a set M *maximal* if $M = M^+$. It follows immediately that M^+ itself is maximal. To prove that M and M^+ agree not just on objective beliefs but on all beliefs, we first show how the true and false support of any sentence always reduces to the support of an objective sentence.

Definition 1 Let α be a sentence and M be a set of situations.

$$\begin{aligned} \|\alpha\|_M &= \alpha, \text{ for objective } \alpha \\ \|\neg\alpha\|_M &= \neg \|\alpha\|_M \\ \|\alpha \vee \beta\|_M &= \|\alpha\|_M \vee \|\beta\|_M \\ \|\mathbf{B}\alpha\|_M &= \text{RES}_{\mathbf{B}}[M, \|\alpha\|_M \wedge \text{det}(\alpha)] \\ \|\mathbf{O}\alpha\|_M &= \text{RES}_{\mathbf{O}}[M, \|\alpha\|_M \wedge \text{det}(\alpha)] \end{aligned}$$

For objective α

$$\begin{aligned} \text{RES}_{\mathbf{B}}[M, \alpha] &= \begin{cases} \text{true} & \text{if } M \models_T \mathbf{B}\alpha \\ \text{false} & \text{otherwise} \end{cases} \\ \text{RES}_{\mathbf{O}}[M, \alpha] &= \begin{cases} \text{true} & \text{if } M \models_T \mathbf{O}\alpha \\ \text{false} & \text{otherwise} \end{cases} \end{aligned}$$

Theorem 2 For any set of situations M , situation s and sentence α ,

1. $M, s \models_T \alpha$ iff $M, s \models_T \|\alpha\|_M$ and $M, s \models_F \alpha$ iff $M, s \models_F \|\alpha\|_M$.
2. $\|\alpha\|_M = \|\alpha\|_{M^+}$

Proof : By induction on the length of α .

Both statements obviously hold for objective α . Similarly for $\neg\alpha$ and $\alpha \vee \beta$.

For the case $\mathbf{B}\alpha$, $M, s \models_T \mathbf{B}\alpha$ iff $\forall s' \in M, M, s' \models_T \alpha \wedge \text{det}(\alpha)$ iff $\forall s' \in M, M, s' \models_T \|\alpha\|_M \wedge \text{det}(\alpha)$ (by induction hypothesis) iff

$M, s \models_T \mathbf{B} \|\alpha \wedge \text{det}(\alpha)\|_M$ [since $\|\alpha \wedge \text{det}(\alpha)\|_M = \|\alpha\|_M \wedge \text{det}(\alpha)$ ($\text{det}(\alpha)$ is objective) and $\text{det}(\|\alpha\|_M \wedge \text{det}(\alpha)) = \text{det}(\alpha)$] iff

$\text{RES}_{\mathbf{B}}[M, \|\alpha \wedge \text{det}(\alpha)\|_M] = \text{true}$ iff

$\text{RES}_{\mathbf{B}}[M, \|\alpha\|_M \wedge \text{det}(\alpha)] = \text{true}$ iff $M, s \models_T \text{RES}_{\mathbf{B}}[M, \|\alpha\|_M \wedge \text{det}(\alpha)]$ iff $M, s \models_T \|\mathbf{B}\alpha\|_M$.

Furthermore, by induction hypothesis, $\|\alpha\|_M = \|\alpha\|_{M^+}$ and since M and M^+ agree on all objective beliefs, $\text{RES}_{\mathbf{B}}[M, \|\alpha\|_M \wedge \text{det}(\alpha)] = \text{RES}_{\mathbf{B}}[M^+, \|\alpha\|_{M^+} \wedge \text{det}(\alpha)]$ and therefore $\|\mathbf{B}\alpha\|_M = \|\mathbf{B}\alpha\|_{M^+}$.

$M, s \models_{\tau} \mathbf{O}\alpha$ iff $(\forall s', s' \in M^+ \text{ iff } M^+, s' \models_{\tau} \alpha \wedge \det(\alpha))$ iff $(\forall s', s' \in M^+ \text{ iff } M^+, s' \models_{\tau} \|\alpha\|_{M^+} \wedge \det(\alpha))$ (by induction hypothesis) iff $M^+, s \models_{\tau} \mathbf{O} \|\alpha \wedge \det(\alpha)\|_{M^+}$ iff $M, s \models_{\tau} \mathbf{O} \|\alpha \wedge \det(\alpha)\|_{M^+}$ (since M and M^+ obviously agree on only-believing due to the fact that $M^+ = M^{++}$) iff $\text{RES}_{\mathbf{O}}[M, \|\alpha \wedge \det(\alpha)\|_{M^+}] = \text{true}$ iff $\text{RES}_{\mathbf{O}}[M, \|\alpha\|_M \wedge \det(\alpha)] = \text{true}$ (since $\det(\alpha)$ is objective and $\|\alpha\|_M = \|\alpha\|_{M^+}$ by ind. hyp.) iff $M, s \models_{\tau} \text{RES}_{\mathbf{O}}[M, \|\alpha\|_M \wedge \det(\alpha)]$ iff $M, s \models_{\tau} \|\mathbf{O}\alpha\|_M$. $\|\mathbf{O}\alpha\|_M = \|\mathbf{O}\alpha\|_{M^+}$ is proved just as in the previous case. ■

Theorem 3 M and M^+ agree on all their beliefs, that is, for any sentence α and set of situations M , $M \models_{\tau} \mathbf{B}\alpha$ iff $M^+ \models_{\tau} \mathbf{B}\alpha$.

Proof : Follows directly from lemma 3.1 theorem 2 ■

3.2.2 Stable Sets

For implicit belief, i.e. the model of belief of an ideal introspective agent, one can give a very precise characterization of the set of sentences believed by such agents. Assume, for the moment, that we have a propositional language with a single modal operator \mathbf{L} , where $\mathbf{L}\alpha$ is read as “ α is implicitly believed”. The original definition of Stalnaker is that a set of sentences Γ is *stable* iff

1. Γ is closed under logical implication.
2. If $\alpha \in \Gamma$, then $\mathbf{L}\alpha \in \Gamma$.
3. If $\alpha \notin \Gamma$, then $\neg\mathbf{L}\alpha \in \Gamma$.

Based on this concept, Moore [12] defined a *stable expansion* Γ of a set A of sentences as a least fixed-point that is closed under logical implication and contains $A \cup \{\mathbf{L}\alpha \mid \alpha \in \Gamma\} \cup \{\neg\mathbf{L}\alpha \mid \alpha \notin \Gamma\}$. In [11] it is shown that for the logic of implicitly only-believing, a sentence α is only-believed iff the set of sentences believed is a stable expansion of $\{\alpha\}$.

We will now develop analogous results for explicitly only-believing. Before we can define the appropriate concept of stability, we need a form of *implication* that reflects the limited reasoning capabilities under explicit belief.

Definition 2 Let Γ be a set of sentences and α a sentence in \mathcal{L} . Γ **strongly implies** α ($\Gamma \succ \alpha$) iff

1. All atoms of α occur in some sentence of Γ .
For all M and s ,
2. If $M, s \models_{\tau} \gamma$ for all $\gamma \in \Gamma$, then $M, s \models_{\tau} \alpha$.
3. If $M, s \models_{\mathbf{F}} \alpha$, then $M, s \models_{\mathbf{F}} \gamma$ for some $\gamma \in \Gamma$.

If Γ is a singleton set $\{\gamma\}$, we also write $\gamma \succ \alpha$ instead of $\{\gamma\} \succ \alpha$.

Two sentences α and β are **strongly equivalent** ($\alpha \approx \beta$) iff $\alpha \succ \beta$ and $\beta \succ \alpha$.

To get better intuitions about the power of strong implication and equivalence, we list a few key properties:

Lemma 3.2 *If $\alpha \succ \beta$ then $\models \alpha \supset \beta$.*

Proof : This follows directly from the fact that the set of worlds is a subset of the set of situations. ■

Lemma 3.3 *Substitutivity*

Let α , β , and $\gamma \in \mathcal{L}$ and let $\alpha^{\beta/\gamma}$ be a formula obtained from α by replacing 0 or more occurrences of β by γ . If $\beta \approx \gamma$, then $\alpha \approx \alpha^{\beta/\gamma}$.

Proof : By a simple induction argument on the structure of α using the definition of \approx . ■

Definition 3 *A sentence α is in conjunctive normal form (CNF), if α is a conjunction of disjunctions, where every disjunct is either a literal (an atom or its negation) or of the form $\mathbf{B}\beta$, $\mathbf{O}\beta$, $\neg\mathbf{B}\beta$, or $\neg\mathbf{O}\beta$ such that β itself is in CNF.*

Lemma 3.4 *For any sentence α there is an α_{CNF} in CNF s.t. $\alpha \approx \alpha_{\text{CNF}}$.*

Proof : Tedious but easy. Essentially, one needs to show that the usual transformation rules (associativity, DeMorgan's law etc.) carry over to situations. Also important is the fact that all of these rules preserve the atoms occurring in a sentence. ■

The previous two lemmas essentially express that sentences and their CNF-transforms carry the same meaning in our logic. Formulas in CNF will become important when we talk about complexity issues.

We can now define what it means for a set of sentences to be stable under explicit belief. Throughout the rest of this section we will restrict ourselves to basic sentences (no Os) only.

Definition 4 *A set of basic sentences Γ is a stable set iff*

1. *If $\Gamma \succ \alpha$ then $\alpha \in \Gamma$.*
2. *If $\alpha \in \Gamma$ then $\text{det}(\alpha) \in \Gamma$.*
3. *If $\alpha \in \Gamma$ then $\mathbf{B}\alpha \in \Gamma$.*
4. *If $\alpha \notin \Gamma$ and $\text{det}(\alpha) \in \Gamma$ then $\neg\mathbf{B}\alpha \in \Gamma$.*

When interpreting Γ as the beliefs of an agent, (1.) tells us that the agent can only perform limited reasoning. (2.) relates to our concept of awareness while (3.) and (4.) express that the agent is fully introspective on matters she is aware of.

Lemma 3.5 *Every stable set is uniquely determined by its objective subset.*

Proof : The proof is a slight variation of the one in [3] for the logic S5. Instead of their lemma 1 (part a), for example, we have:

(let Γ be stable) $\mathbf{B}\alpha \vee \beta \in \Gamma$ if $\text{det}(\mathbf{B}\alpha \vee \beta) \in \Gamma$ and $(\alpha \in \Gamma \text{ or } \beta \in \Gamma)$.

Using facts like these, a simple induction argument on the depth of the nesting of Bs will show that two stable sets that agree on all objective sentences must be identical. ■

A set of basic sentences Γ is called a *basic belief set* if there is an M s.t. for all basic sentences α , $\alpha \in \Gamma$ iff $M \models_{\text{T}} \mathbf{B}\alpha$.

Theorem 4 *A set of basic sentences Γ is stable iff Γ is a basic belief set.*

This corresponds to lemma 2 and proposition 3 of Halpern and Moses [3] in the case of implicit belief. (The same result for implicit belief was obtained apparently independently by R. Moore, M. Fitting, and J. van Benthem. Theorem 2 in [10] generalizes their result to the first-order case.)

Definition 5 *Γ is a stable expansion of a set A of basic sentences iff Γ is a least fixed-point that is closed under strong implication and contains*

$$A \cup \{det(\alpha) | \alpha \in A\} \cup \{B\alpha | \alpha \in \Gamma\} \cup \{\neg B\alpha | \alpha \notin \Gamma, det(\alpha) \in \Gamma\}.$$

Theorem 5 *For any basic α and set of situations M , $M \models_{\tau} O\alpha$ iff the basic belief set is a stable expansion of $\{\alpha\}$.*

3.2.3 Determinate Sentences

Of particular interest to us are sentences that, when only-believed, uniquely determine an epistemic state, i.e., sentences that can truly be taken as representations of an agent's beliefs. Similar to [11], we call a sentence α *determinate* iff there is a unique maximal set M such that $M \models_{\tau} O\alpha$. The following two theorems tell us that (a) the determinate sentences are the only ones that, when only-believed, uniquely determine an epistemic state, and (b) epistemic states that are captured by determinate sentences can always be expressed in purely objective terms.

Theorem 6 *α is determinate iff for all $\beta \in \mathcal{L}$, exactly one of*

$$O\alpha \supset B\beta \text{ or } O\alpha \supset \neg B\beta \text{ is valid.}$$

Proof : If α is determinate, then there is a unique maximal M_{\max} s.t. $M_{\max} \models_{\tau} O\alpha$. Then for every M s.t. $M \models_{\tau} O\alpha$, $M^+ = M_{\max}$. By theorem 3, $M \models_{\tau} B\beta$ iff $M^+ \models_{\tau} B\beta$ for all $\beta \in \mathcal{L}$ and the only-if part of the theorem follows.

On the other hand, for any M_1 and M_2 , let $M_1 \models_{\tau} O\alpha$ and $M_2 \models_{\tau} O\alpha$. By the right hand side of the theorem, M_1 and M_2 must agree on all objective beliefs and therefore, by theorem 1, $M_1^+ = M_2^+$. Thus there is a unique maximal set M_{\max} s.t. $M_{\max} \models_{\tau} O\alpha$. Hence α is determinate. ■

Theorem 7 *If α is determinate, then there is an objective α' s.t. $\models O\alpha \equiv O\alpha'$.*

Proof : Let M be the unique maximal set s.t. $M \models_{\tau} O\alpha$. By theorem 2, $M \models_{\tau} O\alpha$ iff $M \models_{\tau} \|\alpha\|_M$ iff $M \models_{\tau} O(\|\alpha\|_M \wedge det(\alpha))$. Let $\alpha' = \|\alpha\|_M \wedge det(\alpha)$ and the theorem follows. ■

Lemma 3.6 *Let α be objective and let $\mathcal{R}[\alpha] = \{s | s \models_{\tau} \alpha \wedge det(\alpha)\}$. Then α is determinate and $\mathcal{R}[\alpha]$ is the unique maximal set s.t. $\mathcal{R}[\alpha] \models_{\tau} O\alpha$.*

Proof : Easy exercise. ■

Finally, if α and β are objective, sentences of the form $O\alpha \supset O\beta$ and $O\alpha \supset B\beta$ can be reduced to sentences involving Bs only, a much more familiar territory.

Theorem 8 *If α and β are objective, then*

1. $\models \mathbf{O}\alpha \supset \mathbf{O}\beta \iff \models \mathbf{O}\alpha \supset \mathbf{B}\beta$ and $\models \mathbf{O}\beta \supset \mathbf{B}\alpha$
2. $\models \mathbf{O}\alpha \supset \mathbf{B}\beta \iff \models \mathbf{B}\alpha \supset \mathbf{B}\beta$

Proof : 1. Since all objective sentences are determinate, $\models \mathbf{O}\alpha \supset \mathbf{O}\beta$ means that there is a unique maximal set M that only-believes both α and β , i.e. $\mathcal{R}[\alpha] = \mathcal{R}[\beta]$. Since only-believing implies believing, the only-if direction follows. For the if-part of the proof, notice that $\mathbf{O}\alpha \supset \mathbf{B}\beta$ means that $\mathcal{R}[\alpha] \subseteq \mathcal{R}[\beta]$ and $\mathbf{O}\beta \supset \mathbf{B}\alpha$ that $\mathcal{R}[\beta] \subseteq \mathcal{R}[\alpha]$. Hence they are equal.

2. The if-part is immediate. On the other hand, any set M that believes α is a subset of $\mathcal{R}[\alpha]$ and therefore believes β . ■

4 Computational Properties

In general, deciding validity is intractable in this logic, since it subsumes propositional logic. Concerning beliefs implying other beliefs, however, there are interesting tractable cases.

An important one is the decision problem for the validity of $\mathbf{BKB} \supset \mathbf{B}\alpha$, where both \mathbf{KB} and α are objective sentences in CNF. Note that by lemma 3.4, every sentence β has a strongly equivalent CNF-transform β_{CNF} so that $\mathbf{B}\beta \equiv \mathbf{B}\beta_{\text{CNF}}$ by lemma 3.3. We therefore do not lose any expressiveness if we restrict our attention to CNF-formulas.

Theorem 9 *Let $\mathbf{KB} = \bigwedge \mathbf{KB}_i$ and $\alpha = \bigwedge \alpha_j$ be objective sentences in CNF.*

Then $\models \mathbf{BKB} \supset \mathbf{B}\alpha$ iff

1. *Every atom in α occurs in \mathbf{KB} (i.e. $\models \mathbf{BKB} \supset \mathbf{A}\alpha$).*
2. *For all α_j , either α_j is a tautology or there is a \mathbf{KB}_i such that every literal in \mathbf{KB}_i is contained in α_j .⁸*

Testing whether every atom in α occurs in \mathbf{KB} is clearly polynomial in \mathbf{KB} and α . Also, since α is in CNF, checking whether a clause α_j is a tautology is trivial. Finally, testing whether a clause is contained in another is obviously tractable. So, letting $|\alpha|$ denote the length of α , we have:

Corollary 4.1 *If \mathbf{KB} and α are as above, the validity of $\mathbf{BKB} \supset \mathbf{B}\alpha$ can be determined in time $O(|\mathbf{KB}| \times |\alpha|)$.*

Unfortunately, even a slight generalization, where α is still in CNF but is allowed to contain \mathbf{O} s and \mathbf{B} s, leads to intractability. In order to see why, let α and β be in CNF with $\models \mathbf{A}\alpha \supset \mathbf{A}\beta$ and $\mathbf{KB} = \text{det}(\alpha)$. Then it is easy to show that $\mathbf{BKB} \supset \mathbf{B}(\neg \mathbf{B}\alpha \vee \mathbf{B}\beta)$ ⁹ iff $\models \mathbf{B}\alpha \supset \mathbf{B}\beta$. Furthermore one can show that, if α and β are subjective, one can simulate all of propositional reasoning. For example, $\mathbf{B}(\mathbf{B}p \wedge (\neg \mathbf{B}p \vee \mathbf{B}q)) \supset \mathbf{B}\mathbf{B}q$ is valid and embodies *modus ponens*. (See [6] for

⁸This last part corresponds exactly to the way the validity of $\mathbf{BKB} \supset \mathbf{B}\alpha$ is decided in the logic of explicit belief in [9].

⁹Note that the right hand side is in CNF.

logics where this simulation fails and tractability for the general case can be achieved by trading off introspection.)

Intuitively, deciding $\models \mathbf{BKB} \supset \mathbf{B}\alpha$ is hard in general, because \mathbf{BKB} really means “at least KB is believed”, which leaves open too many possible epistemic states that must be taken into consideration. On the other hand, \mathbf{OKB} nails down a unique epistemic state. The advantage is that in order to determine whether $\mathbf{B}\alpha$ follows from \mathbf{OKB} , any belief mentioned in α can actually be replaced by its *truth value* in a kind of preprocessing. In fact, this procedure is nothing but the evaluation of $\|\alpha\|_M$ of definition 1 with $M = \mathcal{R}[\mathbf{KB}]$. For example, to determine the validity of $\mathbf{O}(p \vee q) \supset \mathbf{B}(p \wedge \neg \mathbf{B}q)$, one first constructs $\|p \wedge \neg \mathbf{B}q\|_{\mathcal{R}[\mathbf{KB}]}$ ($\mathbf{KB} = p \vee q$), which is the same as $p \wedge \neg \|\mathbf{B}q\|_{\mathcal{R}[\mathbf{KB}]}$. $\|\mathbf{B}q\|_{\mathcal{R}[\mathbf{KB}]} = \mathbf{false}$ since $\mathcal{R}[\mathbf{KB}] \not\models_{\mathbf{T}} \mathbf{B}(q \wedge \mathit{det}(q))$. Therefore $\|p \wedge \neg \mathbf{B}q\|_{\mathcal{R}[\mathbf{KB}]} = p \wedge \neg \mathbf{false}$, which reduces the original question to testing for the validity of $\mathbf{OKB} \supset \mathbf{B}(p \wedge \neg \mathbf{false})$ or, equivalently, $\mathbf{BKB} \supset \mathbf{B}(p \wedge \neg \mathbf{false})$. by theorem 9.

Theorem 10 *If KB is objective and α any sentence,*
 $\models \mathbf{OKB} \supset \mathbf{B}\alpha \iff \models \mathbf{OKB} \supset \mathbf{B} \|\alpha \wedge \mathit{det}(\alpha)\|_{\mathcal{R}[\mathbf{KB}]}$.

Proof : By theorem 2, $\mathcal{R}[\mathbf{KB}] \models_{\mathbf{T}} \mathbf{B}\alpha$ iff $\mathcal{R}[\mathbf{KB}] \models_{\mathbf{T}} \|\mathbf{B}\alpha\|_{\mathcal{R}[\mathbf{KB}]}$ iff
 $\mathbf{RES}_{\mathbf{B}}[\mathcal{R}[\mathbf{KB}], \|\alpha \wedge \mathit{det}(\alpha)\|_{\mathcal{R}[\mathbf{KB}]}] = \mathbf{true}$ iff $\mathcal{R}[\mathbf{KB}] \models_{\mathbf{T}} \mathbf{B} \|\alpha \wedge \mathit{det}(\alpha)\|_{\mathcal{R}[\mathbf{KB}]}$. ■

Lemma 4.2 *For an objective KB and arbitrary α , both in CNF, $\|\alpha\|_{\mathcal{R}[\mathbf{KB}]}$ is computable in time $O(|\mathbf{KB}| \times |\alpha|)$.*

Proof : From the way recursion is used in defining $\|\alpha\|_{\mathcal{R}[\mathbf{KB}]}$ and the fact that $|\mathit{det}(\alpha)|$ is proportional to $|\alpha|$, it is easy to show that the theorem holds if $\|\mathbf{B}\beta\|_{\mathcal{R}[\mathbf{KB}]}$ and $\|\mathbf{O}\beta\|_{\mathcal{R}[\mathbf{KB}]}$ can be determined fast. First, note that if β is CNF, then $\|\beta\|_{\mathcal{R}[\mathbf{KB}]} \wedge \mathit{det}(\beta)$ is in CNF. Now assume we have evaluated $\|\beta\|_{\mathcal{R}[\mathbf{KB}]}$ and let $\gamma = \|\beta\|_{\mathcal{R}[\mathbf{KB}]} \wedge \mathit{det}(\beta)$. $\mathbf{RES}_{\mathbf{B}}[\mathcal{R}[\mathbf{KB}], \gamma] = \mathbf{true}$ iff $\mathcal{R}[\mathbf{KB}] \models_{\mathbf{T}} \mathbf{B}\gamma$ iff $\models \mathbf{OKB} \supset \mathbf{B}\gamma$ iff $\models \mathbf{BKB} \supset \mathbf{B}\gamma$, which can be done efficiently by corollary 4.1 for objective KB and γ in CNF. The argument for $\|\mathbf{O}\beta\|_{\mathcal{R}[\mathbf{KB}]}$ is similar. ■

From this it is clear that deciding $\mathbf{OKB} \supset \mathbf{B}\alpha$ for an arbitrary α in CNF is no harder than for an objective α in CNF.

Theorem 11 $\models \mathbf{OKB} \supset \mathbf{B}\alpha$ can be decided in time $O(|\mathbf{KB}| \times |\alpha|)$ for any objective KB and arbitrary α , both in CNF.

5 ASK and TELL

We begin by giving a very general specification of ASK and TELL in terms of epistemic states, modelled by sets of situations, and sentences in the logic.

Definition 6

$$\begin{aligned} \mathbf{ASK}[M, \alpha] &= \begin{cases} \mathbf{YES} & \text{if } M \models_{\mathbf{T}} \mathbf{B}\alpha \\ \mathbf{NO} & \text{otherwise} \end{cases} \\ \mathbf{TELL}[M, \alpha] &= M \cap \{s \mid M, s \models_{\mathbf{T}} \alpha \wedge \mathit{det}(\alpha)\} \end{aligned}$$

ASK inquires whether a sentence is believed. Since every sentence is either believed or not believed, there will always be a YES/NO answer. Note that both $\text{ASK}[\mathcal{R}[M, \alpha]]$ and $\text{ASK}[\mathcal{R}[M, \neg\alpha]]$ may return YES, since the system may have inconsistent beliefs about the world. However, this is not the case if α is subjective.

Adding a new assertion means intersecting the current epistemic state with the one where α is believed (which includes awareness of all propositions in α). This approach makes sure that only a minimal amount of knowledge is added, which is consistent with the nature of only-believing, where a maximum of ignorance is desired. Notice also that the set that is intersected with M has itself a reference to M . This is because the query may contain references to beliefs, which are meant to be those *currently* in effect, i.e., with respect to M .

Since we want to do knowledge representation, we now consider objective sentences KB, which uniquely represent what an agent only-believes (KB is determinate). For those sentences, it is possible to define the ASK operation without explicit reference to a particular set of situations.

Theorem 12 *Let α be an arbitrary and KB an objective sentence.*

$$\text{ASK}[\mathcal{R}[\text{KB}], \alpha] = \begin{cases} \text{YES} & \text{if } \models_{\text{OKB}} \supset \mathbf{B}\alpha \\ \text{NO} & \text{otherwise} \end{cases}$$

Proof : A simple consequence of definition 6 and the fact that any objective KB is determinate. ■

The significance of the above results about ASK so far hinges on the assumption that there is an objective sentence KB representing the epistemic state. Since we want TELL to maintain the epistemic state, it needs to be shown that TELL indeed preserves representability. The following theorem does just that by guaranteeing that any belief can be “added” to an epistemic state by conjoining an objective sentence to the one representing it.

Theorem 13 *Let α be an arbitrary and KB an objective sentence.*

$$\text{TELL}[\mathcal{R}[\text{KB}], \alpha] = \mathcal{R}[\text{KB} \wedge \|\alpha\|_{\mathcal{R}[\text{KB}]} \wedge \text{det}(\alpha)]$$

Proof : By definition, $\text{TELL}[\mathcal{R}[\text{KB}], \alpha] = \mathcal{R}[\text{KB}] \cap \{s \mid \mathcal{R}[\text{KB}], s \models_{\text{T}} \alpha \wedge \text{det}(\alpha)\}$. Using the fact that $\mathcal{R}[\gamma \wedge \delta] = \mathcal{R}[\gamma] \cap \mathcal{R}[\delta]$ for any objective γ and δ , all we need to show is that

$$\mathcal{R}[\|\alpha\|_{\mathcal{R}[\text{KB}]} \wedge \text{det}(\alpha)] = \{s \mid \mathcal{R}[\text{KB}], s \models_{\text{T}} \alpha \wedge \text{det}(\alpha)\}.$$

$$\text{But } \mathcal{R}[\|\alpha\|_{\mathcal{R}[\text{KB}]} \wedge \text{det}(\alpha)] = \{s \mid s \models_{\text{T}} \|\alpha\|_{\mathcal{R}[\text{KB}]} \wedge \text{det}(\alpha)\} = \{s \mid \mathcal{R}[\text{KB}], s \models_{\text{T}} \|\alpha\|_{\mathcal{R}[\text{KB}]} \wedge \text{det}(\alpha)\}.$$

By theorem 2, $\mathcal{R}[\text{KB}], s \models_{\text{T}} \alpha \wedge \text{det}(\alpha)$ iff $\mathcal{R}[\text{KB}], s \models_{\text{T}} \|\alpha \wedge \text{det}(\alpha)\|_{\mathcal{R}[\text{KB}]}$.

Since $\|\alpha \wedge \text{det}(\alpha)\|_{\mathcal{R}[\text{KB}]} = \|\alpha\|_{\mathcal{R}[\text{KB}]} \wedge \text{det}(\alpha)$,

$$\mathcal{R}[\|\alpha\|_{\mathcal{R}[\text{KB}]} \wedge \text{det}(\alpha)] = \{s \mid \mathcal{R}[\text{KB}], s \models_{\text{T}} \alpha \wedge \text{det}(\alpha)\} \text{ and we are done. } \blacksquare$$

The representability results about ASK and TELL ensure that these operations can indeed be viewed as a complete specification of the KR service of an agent. At the beginning, when the agent knows nothing, the KB is the sentence true, which is considered an objective determinate sentence. The set of situations the agent thinks possible is the set of all situations. Using TELL

as the only means to acquire knowledge, the agent's KB is then always guaranteed to be an objective sentence. Finally, for sentences in conjunctive normal form TELL and ASK are also efficient:

Theorem 14 *Let KB be an objective and α an arbitrary sentence, both in CNF.*

$\text{ASK}[\mathcal{R}[\text{KB}], \alpha]$ is computable in time $O(|\text{KB}| \times |\alpha|)$.

$\text{TELL}[\mathcal{R}[\text{KB}], \alpha]$ is computable in time $O(|\text{KB}| \times |\alpha|)$

Proof : The tractability result follows immediately from theorem 4.2 and corollary 11. ■

Converting a sentence α into CNF can result in an exponential blow-up. However, keeping a KB in CNF is not unreasonable, since TELL guarantees that only the new assertion needs to be converted into CNF with the rest of the KB left untouched. Similarly, in case of ASK, only α needs to be converted.

6 Summary

We have presented a logic of explicit belief with full introspection. Besides a conventional belief operator **B**, the logic provides an operator **O**, which captures the intuition of only-believing, which is the appropriate concept for KR purposes. Since deciding whether only-believing an objective KB in CNF implies believing an arbitrary sentence in CNF is tractable, efficient KR service routines TELL and ASK can be specified.

We are currently investigating to what extent these results carry over to a much more expressive first-order logic of explicit belief with rigid as well as non-rigid designators [5,7].

Acknowledgements

We would like to thank Jim des Rivières for his comments on an earlier draft of this paper. This research has been supported in part by a Government of Canada Award.

References

- [1] Belnap, N. D., A Useful Four-Valued Logic, in G. Epstein and J. M. Dunn (eds.), *Modern Uses of Multiple-Valued Logic*, Reidel, 1977.
- [2] Fagin, R. and Halpern, J. Y., Belief, Awareness, and Limited Reasoning: Preliminary Report, in *Proc. of the Ninth International Joint Conference on Artificial Intelligence*, August 1985, pp. 491-501.
- [3] Halpern, J. Y. and Moses, Y. O., Towards a Theory of Knowledge and Ignorance: Preliminary Report, in *The Non-Monotonic Reasoning Workshop*, New Paltz, NY, 1984, pp. 125-143.
- [4] Halpern, J. Y. and Moses, Y. O., A Guide to the Modal Logics of Knowledge and Belief, in *Proc. of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA, August 1985, pp. 480-490.

- [5] Lakemeyer, G., Steps Towards a First-Order Logic of Explicit and Implicit Belief, in *Proc. of the Conference on Theoretical Aspects of Reasoning about Knowledge*, Asilomar, California, 1986, pp. 325–340.
- [6] Lakemeyer, G., Tractable Meta-Reasoning in Propositional Logics of Belief, in *Proc. of the Tenth International Joint Conference on Artificial Intelligence*, Milano, Italy, 1987.
- [7] Lakemeyer, G., Ph.D. thesis, forthcoming.
- [8] Levesque, H. J., Foundations of a Functional Approach to Knowledge Representation, *Artificial Intelligence*, Vol. 23, 1984, pp. 155-212.
- [9] Levesque, H. J., A Logic of Implicit and Explicit Belief, Tech. Rep. No. 32, Fairchild Lab. for AI Research, Palo Alto, 1984.
- [10] Levesque, H. J., All I Know: An Abridged Report, in *Proc. of the Sixth National Conference of the American Association for Artificial Intelligence*, Seattle, Washington, 1987, pp. 426–431.
- [11] Levesque, H. J., All I Know: A Study in Autoepistemic Logic, submitted to: *Artificial Intelligence*, 1987.
- [12] Moore, R., Semantical Considerations on Nonmonotonic Logic, in *Proc. of the Eighth International Joint Conference on Artificial Intelligence*, Karlsruhe, FRG, 1983, pp. 272–279.
- [13] Moore, R., Possible World Semantics for Autoepistemic Logic, in *The Non-Monotonic Reasoning Workshop*, New Paltz, NY, 1984, pp. 344–354.