Relevance in a Logic of "Only Knowing About" and its Axiomatization

Gerhard Lakemeyer
Institute of Computer Science III
University of Bonn
Römerstr. 164
D-53117 Bonn, Germany
gerhard@cs.uni-bonn.de

Abstract

In previous work we defined the concept of *x is all an agent knows about y* by augmenting Levesque's logic of only-knowing with appropriate modal operators. In this paper we demonstrate how various notions of relevance can be captured within such a logic of only-knowing-about. For example, we are able to formalize what it means for a sentence to be relevant to a subject matter (set of atomic propositions) or for one atomic proposition to be relevant to another relative to some background theory. Moreover, we provide a proof theory for the logic of only-knowing-about. While we have not been able to prove completeness for the whole logic, completeness is established for the class of sentences which is used to define our various notions of relevance.

Introduction

Being able to ask a knowledge base queries like *"tell me all the (relevant) information you have about aunt Leonore"* is not only of practical relevance, but its meaning is so far poorly understood. One attempt of giving precise meaning to such notions was made in [Lak92, Lak93], where we proposed a formalization of *all an agent knows about a subject matter (only-knowing-about)* for both the single and multi-agent case in a propositional logic. In particular, we developed a formal semantics based on possible-world semantics [Kri63, Hin62, HM92] and, in particular, on Levesque's formalization of only-knowing [Lev90], which can be thought of as the limit-case of only-knowing-about, where the subject matter includes everything the agent has any information about.

In this paper we demonstrate how various notions of relevance can be captured within such a logic of only-knowing-about. For example, we are able to formalize what it means for a sentence to be relevant to a subject matter (set of atomic propositions) or for one atomic proposition to be relevant to another relative to some background theory. Moreover, we provide a proof theory for the logic of only-knowing-about. While we have not been able to prove completeness for the whole logic, completeness is established for the class of sentences which is used to define our various notions of relevance.

In Section 2, we define the semantics of the only-knowing-about, which is based on the semantics presented in [Lak92]. Section 3 contains our various definitions of relevance using the logic of the previous section. Section 4 tackles the axiomatization followed by some concluding remarks.

The Semantics of Only-Knowing-About

From Only-Knowing to Only-Knowing-About

In order to provide a semantics to "all the agent knows about x is y," we start out with Levesque's logic of only-knowing [Lev90]. There an agent knows a sentence α, denoted as $\mathbf{L}\alpha$, just in case α is true in all the worlds the agent thinks possible. (Formal definitions are deferred to Section 3 below.) To define only-knowing, Levesque considers another modality $\mathbf{N}$, where $\mathbf{N}\alpha$ means that α is true in all the *impossible* (or non-accessible) worlds. While $\mathbf{L}\alpha$ is best understood as "the agent knows *at least* that α is true," $\mathbf{N}\alpha$ should be read as "the agent knows *at most* that α is false." With that only-knowing α, denoted as $\mathbf{O}\alpha$, reduces to knowing at least α and at most α, that is, $\mathbf{O}\alpha$ holds just in case both $\mathbf{L}\alpha$ and $\mathbf{N}\neg\alpha$ hold.

Let us now consider how to extend these ideas to include a subject matter. Since we confine ourselves to propositional logic, we define a *subject matter* π as a finite set of atomic propositions. For each such π we introduce a new modal operator $\mathbf{O}\langle\pi\rangle$, where $\mathbf{O}\langle\pi\rangle\alpha$ should be read as "all the agent knows about π is α." Levesque defined $\mathbf{O}$ in terms of $\mathbf{L}$ and $\mathbf{N}$ mainly to simplify the axiomatization of $\mathbf{O}$. Here we use a similar approach to capture $\mathbf{O}\langle\pi\rangle$. However, rather than introducing operators $\mathbf{L}\langle\pi\rangle$ and $\mathbf{N}\langle\pi\rangle$ for arbitrary π, we do so only in the case where the subject matter consists of a single atom. Hence for any $\pi = \{p\}$ we introduce $\mathbf{L}\langle p\rangle$ and $\mathbf{N}\langle p\rangle$, where $\mathbf{L}\langle p\rangle\alpha$ is read as "the agent knows at least α about p" and $\mathbf{N}\langle p\rangle\alpha$ as "the agent knows at most that α is false about p." $\mathbf{O}\langle\{p\}\rangle\alpha$, which we also write as $\mathbf{O}\langle p\rangle\alpha$, is then defined as $\mathbf{L}\langle p\rangle\alpha \wedge \mathbf{N}\langle p\rangle\neg\alpha$. We will see later that $\mathbf{O}\langle\pi\rangle\alpha$ for an arbitrary π reduces to an expression containing operators of the form $\mathbf{O}\langle\pi'\rangle$ only for singleton π'. Thus the operators $\mathbf{L}\langle p\rangle$ and $\mathbf{N}\langle p\rangle$ are sufficient to account for $\mathbf{O}\langle\pi\rangle$ in its full generality.[1]

To define the semantics of these modalities, suppose the beliefs of the agent are given by the set of worlds M the agent thinks possible. To find out what the agent knows about π we construct a set of worlds $M|_\pi$ which, intuitively, represents what the agent knows after forgetting everything that is not relevant to π. With that the operators $\mathbf{L}\langle p\rangle$,$\mathbf{N}\langle p\rangle$, and $\mathbf{O}\langle\pi\rangle$ are interpreted just like $\mathbf{L}$,$\mathbf{N}$, and $\mathbf{O}$ except that we are using $M|_p$ and $M|_\pi$ instead of M. For example, an agent believes α about p at a set of worlds M just in case he believes α at $M|_p$.

The logic is both an extension and a simplification of the one presented in [Lak92]. It is simpler because we completely ignore nested modalities, which are not essential for our purposes. It is an extension because the operators $\mathbf{L}\langle p\rangle$ and $\mathbf{N}\langle p\rangle$ were not present in the previous version.

The Language and Other Notation

The primitives of the language are a countably infinite set $\mathcal{P}$ of atomic propositions (or atoms), the connectives $\vee$, $\neg$, and the modal operators $\mathbf{L}$,$\mathbf{N}$,$\mathbf{O}$, $\mathbf{L}\langle p\rangle$, $\mathbf{N}\langle p\rangle$, and $\mathbf{O}\langle\pi\rangle$ for every atom p and every finite set of atomic propositions π with the restriction that none of the modal operators occurs within the scope of another modal operator. Sentences are formed in the usual way from these primitives.[2]

Notation: As usual, literals are either atoms or negated atoms and clauses are disjunctions of literals. We write `false` as an abbreviation for $(p \wedge \neg p)$, where p is some atom, and `true` for $\neg$`false`. It is often convenient to identify a clause with the set of literals occurring in the clause. A clause c is contained in a clause c' ($c \subseteq c'$) if every literal in c occurs in c'. We write $c \subsetneq c'$ instead of $c \subseteq c'$ and $c' \not\subseteq c$. Given a finite set of sentences Γ, $\bigwedge_{\gamma\in\Gamma}\gamma$ denotes the conjunction of all the sentences occurring in Γ. If Γ is empty, $\bigwedge_{\gamma\in\Gamma}\gamma$ denotes `true`. Finally, a sentence is called *objective* if it contains no modal operator.

[1] The main motivation for restricting ourselves to singleton subject matters in the case of the operators $\mathbf{L}\langle p\rangle$ and $\mathbf{N}\langle p\rangle$ is to reduce the complexity of the proof theory.

[2] We will freely use other connectives like $\wedge$, $\supset$ and $\equiv$, which should be understood as syntactic abbreviations of the usual kind.

A Formal Semantics

Worlds are defined extensionally as propositional truth assignments. Hence for any given set of worlds M (the possible worlds) its complement (the impossible worlds) is always well defined.

Definition 1 (Worlds) *A world w is a function $w : \mathcal{P} \longrightarrow \{\mathbf{t},\mathbf{f}\}$.*

We begin by reviewing the semantics of Levesque's logic of only-knowing:

$$
\begin{array}{lcl}
M,w\models p & \iff & w(p)=\mathbf{t},\text{ where } p \text{ is an atom}\\
M,w\models\neg\alpha & \iff & M,w\not\models\alpha\\
M,w\models\alpha\vee\beta & \iff & M,w\models\alpha \text{ or } M,w\models\beta\\
M,w\models\mathbf{L}\alpha & \iff & \text{for all } w'\in M,\ M,w'\models\alpha\\
M,w\models\mathbf{N}\alpha & \iff & \text{for all } w'\notin M,\ M,w'\models\alpha\\
M,w\models\mathbf{O}\alpha & \iff & M,w\models\mathbf{L}\alpha \text{ and } M,w\models\mathbf{N}\neg\alpha
\end{array}
$$

For a given set of worlds M and a subject matter that consists of a single atom p, we define $M|_p$ as the set of all worlds that satisfy precisely the known objective sentences that are about p. Since we are dealing with propositional beliefs only, we can, first of all, confine ourselves to clauses instead of arbitrary objective sentences. The idea to get at only those beliefs that are about p is to consider the smallest clauses which are believed and which mention at least one of the atoms in p. $M|_p$ is then simply the set of all worlds that satisfy all of these clauses. The generalization of this idea from a subject matter with one atom to an arbitrary subject matter π is straightforward. $M|_\pi$ is simply the intersection of all the $M|_p$, where p occurs in π, that is $M|_\pi$ believes precisely those minimal clauses c that mention some atom included in π. Formally:

Definition 2 *Let M be a set of worlds and π a subject matter.*

1. *A clause c is called M-**minimal** iff $M\models\mathbf{L}c$ and for all clauses $c' \subsetneq c$, $M\not\models\mathbf{L}c'$.*
2. *A clause c is called M-p-**minimal** iff c is M-minimal and, in addition, c mentions p, that is, c contains either p or $\neg p$.*
3. *$M|_p = \{w \mid w\models c$ for all M-p-minimal clauses $c\}$.*
4. *$M|_\pi = \bigcap_{p\in\pi} M|_p$.*[3]

By restricting ourselves to M-p-minimal clauses, we rule out clauses that mention the subject matter but do not really tell us anything about it. For example, let the subject matter be p and assume all we know is q, that is, $M = \{w \mid w\models q\}$. Then we certainly also know $(p \vee q)$, which is not M-minimal because q is known

[3] In [Lak92], $M|_\pi$ was defined in a slightly different way. However, it is easily seen that the two definitions coincide.

as well. While $(p \vee q)$ mentions the subject matter p, it does so, in a sense, only accidentally, since it does not convey us what is really known about p, namely nothing. The only M-minimal clause mentioning p is $(p \vee \neg p)$, which gives us the right information.

Given these definitions of what it means to forget irrelevant things, we obtain the following semantic rules for knowing and only-knowing about.

$$
\begin{array}{lcl}
M, w \models \mathbf{L}\langle p\rangle\alpha & \iff & M|_p, w \models \mathbf{L}\alpha \\
M, w \models \mathbf{N}\langle p\rangle\alpha & \iff & M|_p, w \models \mathbf{N}\alpha \\
M, w \models \mathbf{O}\langle \pi\rangle\alpha & \iff & M|_\pi, w \models \mathbf{L}\alpha \wedge \mathbf{N}\neg\alpha
\end{array}
$$

Note: In the original definition of $\mathbf{O}\langle\pi\rangle$ in [Lak92], $M, w \models \mathbf{O}\langle\pi\rangle\alpha$ was defined as $M|_\pi, w \models \mathbf{O}\alpha$ and $M, w \models \mathbf{L}\alpha$. The restriction "$M, w \models \mathbf{L}\alpha$" was necessary to prevent unintuitive properties in the case of nested beliefs. In the unnested case, as in this paper, these problems do not occur and we get by with the simpler definition. Notice also that for singleton subject matters, $M, w \models \mathbf{O}\langle p\rangle\alpha$ iff $M, w \models \mathbf{L}\langle p\rangle\alpha \wedge \mathbf{N}\langle p\rangle\neg\alpha$.

As pointed out in [Lev90], there are sets of worlds, which, according to the above definition, believe precisely the same objective sentences, yet disagree on what they only-believe. To avoid this anomaly, Levesque introduced the notion of a *maximal* set of worlds M^+, which is a unique representative of all sets of worlds with the same objective beliefs as M^+.

Definition 3 (Levesque) *Maximal Sets*
Let M be any a set of worlds and let $M^+ = \{w \mid M, w \models \mathbf{L}\alpha \supset \alpha$ for all objective $\alpha\}$. M is called maximal iff $M = M^+$.

Note that the sets $M|_p$ and $M|_\pi$ introduced earlier are themselves maximal sets.

Logical implication and validity are defined with respect to worlds and maximal sets of worlds only: A set of sentences Γ *logically implies* a sentence α ($\Gamma \models \alpha$) iff for all worlds w and for all maximal sets of worlds M, if $M, w \models \gamma$ for all $\gamma \in \Gamma$, then $M, w \models \alpha$. α is *valid* ($\models \alpha$) iff $\{\} \models \alpha$. α is satisfiable iff $\neg\alpha$ is not valid.

Expressing $\mathbf{O}\langle p_1, \ldots, p_n\rangle$ in terms of $\mathbf{O}\langle p_1\rangle$, $\ldots$ $\mathbf{O}\langle p_n\rangle$

While the subject may contain an arbitrary number of atoms, it can be shown that $\mathbf{O}\langle\pi\rangle\alpha$ reduces to a sentence where the subject matter in each occurrence of only-knowing-about is a singleton set. To obtain this result, we first need to define the notion of *prime implicates*.[4]

Definition 4 (Prime Implicates) *Let α be an objective sentence. A clause c is called a prime implicate of α iff*

1. $\models \alpha \supset c$ and
2. for all $c' \subsetneq c$, $\not\models \alpha \supset c'$.

Let p be an atom. Then $\mathcal{P}(\alpha, p) = \{c \mid c$ is a prime implicate of α mentioning $p\}$ and $\mathcal{P}(\alpha) = \bigcup_{p \in \alpha} \mathcal{P}(\alpha, p)$.

For example, if $\alpha = (p \vee q) \wedge (p \vee r \vee s) \wedge \neg s$, then $\mathcal{P}(\alpha, p) = \{(p \vee q), (p \vee r)\}$. In cases like $\alpha = q$, where p is not contained in α at all, $\mathcal{P}(\alpha, p) = \{(p \vee \neg p)\}$.

It is easy to see that for any given α and p, $\mathcal{P}(\alpha, p)$ and $\mathcal{P}(\alpha)$ are finite assuming we identify, as is customary, clauses with sets of literals and hence eliminate redundancies.

Theorem 1

$$\models \mathbf{O}\langle\pi\rangle\alpha \equiv (\bigwedge_{i=1}^{n} \mathbf{O}\langle p_i\rangle\alpha_i \wedge [(\bigwedge_{i=1}^{n} \alpha_i) \equiv \alpha]),$$

where $\pi = \{p_1, p_2, \ldots, p_n\}$ and $\alpha_i = (\bigwedge_{c \in \mathcal{P}(\alpha, p_i)} c)$ for $1 \le i \le n$.

What is known about p relative to a sentence (background theory) α has a simple characterization in terms of prime implicates of α.

Theorem 2 *$\models \mathbf{O}\alpha \supset \mathbf{O}\langle p\rangle\beta$ iff $\models \beta \equiv \bigwedge_{\gamma \in \mathcal{P}(\alpha, p)} \gamma$.*

From these theorems we obtain the general case for arbitrary π immediately:

Corollary 3 *$\models \mathbf{O}\alpha \supset \mathbf{O}\langle\pi\rangle\beta$ iff $\models \beta \equiv \bigwedge_{p \in \pi} \bigwedge_{\gamma \in \mathcal{P}(\alpha, p)} \gamma$.*

Incidentally, this result proves that Lin and Reiter's notion of *remembering* [LR94] coincides with our concept of only-knowing-about relative to a theory α.

Shades of Relevance

The logic of only-knowing-about provides a natural way of characterizing various notions of (logical) relevance. We begin by defining what it means for a sentence to be relevant to some subject matter. The intuition behind α being relevant to π is that α must contain non-trivial information about π. Our logic allows us to express this directly.

Definition 5 *An objective sentence α is relevant to a subject matter π iff $\models \mathbf{O}\alpha \supset \neg\mathbf{O}\langle\pi\rangle(p \vee \neg p)$.*

Example 1 *While $\neg p$ and $(p \supset q) \wedge (q \supset r)$ are relevant to p, $(q \supset r)$ and $p \supset (q \supset p)$ are not.*

Lemma 4 *α is relevant to π iff there is some $\gamma \in \mathcal{P}(\alpha)$ such that $\not\models \gamma$ and γ mentions some $p \in \pi$.*

[4] In [Lak92] we pointed out the close connection between only-knowing-about and deKleer's ATMS[deK86], where prime implicates play a central role as well [RdK87].

While the previous definition, in a sense, only requires part of the sentence to be about π, we can be even more restrictive and require that everything α tells us is about π in a relevant way.[5]

Definition 6 *An objective sentence α is strictly relevant to a subject matter π iff $\not\models\alpha$ and $\mathbf{O}\langle\pi\rangle\alpha$ is satisfiable.*

Example 2 *$(p \supset q) \wedge (q \supset r)$ is not strictly relevant to p because $(q \supset r)$ is not about p. However, $(p \equiv q) \wedge (q \supset r)$ is strictly relevant to p. This time $(q \supset r)$ is recognized as being about p since p and q are assumed to be equivalent.*

Lemma 5 *Let α be an objective sentence such that $\not\models\alpha$. Then the following statements are equivalent.*

1. *α is strictly relevant to π.*
2. *$\models\mathbf{O}\alpha \supset \mathbf{O}\langle\pi\rangle\alpha$.*
3. *$\models\alpha \equiv \bigwedge_{p\in\pi}\bigwedge_{\gamma\in\mathcal{P}(\alpha,p)}$.*

Next we would like to express that an atom p contributes to what is known about another atom q in a relevant way.

Definition 7 *Let p and q be atoms, α and β objective sentences such that $\models\mathbf{O}\alpha \supset \mathbf{O}\langle q\rangle\beta$. p is relevant to q with respect to α iff $p = q$ or β is relevant to p.*

In other words, p is relevant to q if whatever is known about q contains some non-trivial information about p.

Example 3 *Let $\alpha = (p \vee q) \wedge (q \vee r)$. Since $\models\mathbf{O}\alpha \supset \mathbf{O}\langle q\rangle\alpha$ and $\models\mathbf{O}\alpha \supset \mathbf{O}\langle p\rangle(p\vee q)$, we obtain immediately that p is relevant to q. Similarly, q is relevant to r. However, p is not relevant to r, since $\models\mathbf{O}\alpha \supset \mathbf{O}\langle r\rangle(q \vee r)$ and $\models\mathbf{O}(q \vee r) \supset \mathbf{O}\langle p\rangle(p \vee \neg p)$.*

Lemma 6 *p is relevant to q with respect to α iff there is a $\gamma \in \mathcal{P}(\alpha)$ such that γ mentions both p and q.*

This relevance relation between atoms is obviously reflexive by definition. While symmetry is not obvious from the definition, it nevertheless follows immediately from Lemma 6, that is, if p is relevant to q with respect to α, then q is relevant to p. Note, however, that transitivity does not hold. Example 3 provides a counterexample.

While the previous definition requires p and q only to be weakly connected to each other, the following, and last, definition forces this connection to be much stronger. In particular, we require that whatever is known about p is also known about q.

Definition 8 *q subsumes p with respect to α ($p \prec_\alpha q$) iff $\models\mathbf{O}\alpha \supset (\mathbf{L}\langle p\rangle\beta \supset \mathbf{L}\langle q\rangle\beta)$ for all objective β. p and q are equivalent with respect to α ($p \approx_\alpha q$) iff $p \prec_\alpha q$ and $q \prec_\alpha p$.*

[5]This definition was first introduced in [Lak93].

It is easy to see that $\models\alpha \supset (p \equiv q)$ implies $p \approx_\alpha q$. Note, however, that the converse does not hold. For example, $p \approx_\alpha q$ holds even for $\alpha = (p \supset q)$.

Lemma 7 *$p \prec_\alpha q$ iff $\models\mathcal{P}(\alpha, q) \supset \mathcal{P}(\alpha, p)$.*

In all cases of our various definitions of relevance we were able to also offer alternative characterizations in terms of simple properties of certain prime implicates. In the next section, we present an axiomatization of only-knowing-about, which will give us yet another characterization and serves to better understand what we are actually doing.

A Proof Theory

Since our logic reduces to Levesque's logic of only-knowing when restricted to sentences mentioning at most the modal operators **L**,**N**, and **O**, Levesque's proof theory is part of ours as well:[6]

Axioms for L,N, and O (Levesque)

- **A1** Axioms of classical propositional logic.
- **A2** $\mathbf{L}(\alpha \supset \beta) \supset (\mathbf{L}\alpha \supset \mathbf{L}\beta)$.
- **A3** $\mathbf{N}(\alpha \supset \beta) \supset (\mathbf{N}\alpha \supset \mathbf{N}\beta)$.
- **A4** $\mathbf{L}\alpha \supset \neg\mathbf{N}\alpha$ for every falsifiable objective α.
- **A5** $\mathbf{O}\alpha \equiv (\mathbf{L}\alpha \wedge \mathbf{N}\neg\alpha)$.

Inference Rules

- **MP** From α and $\alpha \supset \beta$ infer β
- **Nec** From α infer $\mathbf{L}\alpha$ and $\mathbf{N}\alpha$.

Note that axiom **A3** and the rule of necessitation tell us that **N** is just an ordinary belief operator (such as **L**). The fact that **L** and **N** are defined over complementary sets of worlds is reflected by axiom **A4**.

Let p be an arbitrary atom. In the following we confine ourselves to singleton subject matters (p), which in light of Theorem 1 is sufficient to fully account for $\mathbf{O}\langle\pi\rangle$ in the general case.

The next four axioms and the inference rule **Nec**$^+$ completely mirror **A2**–**A5** and **Nec**. In other words, $\mathbf{L}\langle p\rangle$,$\mathbf{N}\langle p\rangle$, and $\mathbf{O}\langle p\rangle$, taken by themselves, behave just like **L**, **N**, and **O**.

- **A6** $\mathbf{L}\langle p\rangle(\alpha \supset \beta) \supset (\mathbf{L}\langle p\rangle\alpha \supset \mathbf{L}\langle p\rangle\beta)$.
- **A7** $\mathbf{N}\langle p\rangle(\alpha \supset \beta) \supset (\mathbf{N}\langle p\rangle\alpha \supset \mathbf{N}\langle p\rangle\beta)$.
- **A8** $\mathbf{L}\langle p\rangle\alpha \supset \neg\mathbf{N}\langle p\rangle\alpha$ for every falsifiable obj. α.
- **A9** $\mathbf{O}\langle p\rangle\alpha \equiv \mathbf{L}\langle p\rangle\alpha \wedge \mathbf{N}\langle p\rangle\neg\alpha$.

Nec$^+$ From α infer $\mathbf{L}\langle p\rangle\alpha$ and $\mathbf{N}\langle p\rangle\alpha$ for obj. α.

The following axioms, which are probably the most interesting, deal with the connection between **L**,**N**,**O** and their versions restricted to a subject matter. Let $Min(c) = \bigwedge_{c'\subsetneq c}\neg\mathbf{L}c'$ for any clause c.

[6]Levesque's original axiomatization is more complicated since he treats nested beliefs as well as first-order sentences.

A10 $\mathbf{L}\langle p\rangle\alpha \supset \mathbf{L}\alpha$

A11 $\mathbf{N}\alpha \supset \mathbf{N}\langle p\rangle\alpha$

A12 $\mathbf{L}\langle p\rangle l \wedge \neg\mathbf{Lfalse} \supset \mathbf{N}\langle p\rangle\neg l$, where $l = p$ or $l = \neg p$.

A13 $\mathbf{L}c \wedge Min(c) \supset \mathbf{L}\langle p\rangle c$, where c is a clause that mentions p.

A14 $\mathbf{N}\langle p\rangle\neg\alpha \wedge \mathbf{L}c \supset \mathbf{N}\langle p\rangle\neg\alpha'$, where
$$\alpha' = \bigwedge_{\gamma\in\mathcal{P}(\alpha)-\gamma^*} \gamma$$
for any $\gamma^* \in \mathcal{P}(\alpha)$ such that $c \subseteq \gamma^*$ if c does not mention p, and $c \subsetneq \gamma^*$ if c mentions p.

A15 $\mathbf{O}\langle\pi\rangle\alpha \equiv (\bigwedge_{i=1}^{n} \mathbf{O}\langle p_i\rangle\alpha_i \wedge [(\bigwedge_{i=1}^{n} \alpha_i) \equiv \alpha])$, where where $\pi = \{p_1, p_2, \ldots, p_n\}$ and
$$\alpha_i = (\bigwedge_{c\in\mathcal{P}(\alpha, p_i)} c).$$

A10 and **A11** simply account for the fact that for any set of worlds M, $M \subseteq M|_p$ holds. **A12** says, in essence, that if you have complete information about p, that is, you know either it or its negation, then you cannot know any more about it (unless you are inconsistent). **A13** accounts for our semantic construction, where minimal clauses that mention p are among the beliefs about p. **A15** is best illustrated by example. One instance is $\mathbf{N}\langle p\rangle\neg[q \wedge r] \wedge \mathbf{L}q \supset \mathbf{N}\langle p\rangle\neg r$. $\mathbf{N}\langle p\rangle\neg[q \wedge r]$ says that any clause believed about p contains either q or r. However, if q is believed, then any clause containing q is irrelevant to p. Hence every clause believed about p must mention r denoted by $\mathbf{N}\langle p\rangle\neg r$. Axiom **A14**, finally, provides the reduction from $\mathbf{O}\langle\pi\rangle$ for arbitrary π to singleton $\mathbf{O}\langle p_i\rangle$'s as in Theorem 1.

Theorem 8 (Soundness)
For all α, if $\vdash \alpha$ then $\models\alpha$.[7]

So far we have not obtained a completeness proof for the axiomatization. In fact, as discussed below, there is good reason to believe that it is incomplete. However, the proof theory is complete for the class of sentences that we used in the previous section on relevance.

Theorem 9 (Partial Completeness) *The axiomatization is complete for sentences of the form $\mathbf{O}\alpha \supset \beta$, that is, if $\models\mathbf{O}\alpha \supset \beta$ then $\vdash \mathbf{O}\alpha \supset \beta$.*

With that, we get a syntactic characterization of relevance (Definitions 5–8). Simply replace $\models$ by $\vdash$ everywhere.[8]

[7] We assume the usual notion of provability ($\vdash$).

[8] In the case of strict relevance (Definition 6), use the alternative characterization of Lemma 5, that is, α is strictly relevant to π iff $\models\mathbf{O}\alpha \supset \mathbf{O}\langle\pi\rangle\alpha$.

Conclusion

In this paper we showed how to define various forms of relevance within a logic of only-knowing-about. We showed that the semantic definitions can be captured both in terms of prime implicates and in proof theoretic terms of the underlying logic.

An important open question is whether the proof theory is complete for the whole language. While I have no formal proof yet, I conjecture that it is indeed incomplete. For example, the sentence $\mathbf{N}\langle p\rangle\neg(q \vee r \vee s) \supset \neg\mathbf{L}\langle p\rangle(q \vee r \vee s \vee t)$ is valid for any atoms q, r, s, and t distinct from p, which seems impossible to account for with the axioms we have so far.

Besides finding a complete axiomatization, another interesting question is how to extend our results to the first-order case. Lin and Reiter [LR94] have a metalogical formalization of *remembering* in the first-order case. We have seen that remembering is expressible within our logic when restricted to the propositional case. Whether similar parallels exist in first-order logic remains to be seen.

References

[deK86] de Kleer, J., An assumption-based TMS. *Artificial Intelligence*, **28**(2), 1986, pp. 127–162.

[HM92] Halpern, J. Y. and Moses, Y. O., A Guide to completeness and complexity for modal logics of knowledge and belief, *Artificial Intelligence*, **54**, pp. 319–379, 1992. *Proc. of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA, 1985, pp. 480–490.

[Hin62] Hintikka, J., *Knowledge and Belief: An Introduction to the Logic of the Two Notions.* Cornell University Press, 1962.

[Kri63] Kripke, S. A., Semantical considerations on modal logic. *Acta Philosophica Fennica* **16**, pp. 83–94, 1963.

[Lak92] Lakemeyer, G., All You ever Wanted to Know about Tweety, in *Proc. of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR'92)*, Morgan Kaufmann, San Mateo, CA, 1992, pp. 639–648.

[Lak93] Lakemeyer, G., All They Know About, *Proc. of the 11th National Conference on Artificial Intelligence (AAAI-93)*, MIT Press, Cambridge, MA, pp. 662–667.

[Lev90] Levesque, H. J., All I Know: A Study in Autoepistemic Logic. *Artificial Intelligence*, North Holland, **42**, 1990, pp. 263–309.

[LR94] Lin, Fangzhen and Reiter, R., Forget It!, these proceedings.

[RdK87] Reiter, R. und de Kleer, J., Foundations of assumption-based truth maintenance systems: preliminary report. *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, Seattle, WA, 1987, pp. 183–188.