2.2 Reasoning about Intentional Notions

Suppose one wishes to reason about intentional notions in a logical framework. Consider the following statement (after [73, pp210–211]):

$$\text{Janine believes Cronos is the father of Zeus.} \tag{2.1}$$

A naive attempt to translate (2.1) into first-order logic might result in the following:

$$Bel(Janine, Father(Zeus, Cronos)) \tag{2.2}$$

Unfortunately, this naive translation does not work, for at least two reasons. The first is syntactic: the second argument to the *Bel* predicate is a *formula* of first-order logic, and is not, therefore a term. So (2.2) is not a well-formed formula of classical first-order logic. The second problem is semantic, and is more serious. The constants *Zeus* and *Jupiter*, by any reasonable interpretation, denote the same individual: the supreme deity of the classical world. It is therefore acceptable to write, in first-order logic:

$$(Zeus = Jupiter). \tag{2.3}$$

Given (2.2) and (2.3), the standard rules of first-order logic would allow the derivation of the following:

$$Bel(Janine, Father(Jupiter, Cronos)) \tag{2.4}$$

But intuition rejects this derivation as invalid: believing that the father of Zeus is Cronos is *not* the same as believing that the father of Jupiter is Cronos.

So what is the problem? Why does first-order logic fail here? The problem is that the intentional notions — such as belief and desire — are *referentially opaque*, in that they set up *opaque contexts*, in which the standard substitution rules of first-order logic do not apply. In classical (propositional or first-order) logic, the denotation, or semantic value, of an expression is dependent solely on the denotations of its sub-expressions. For example, the denotation of the propositional logic formula $p \wedge q$ is a function of the truth-values of p and q. The operators of classical logic are thus said to be *truth functional*.

In contrast, intentional notions such as belief are *not* truth functional. It is surely not the case that the truth value of the sentence:

$$\text{Janine believes } p \tag{2.5}$$

is dependent solely on the truth-value of p[7]. So substituting equivalents into opaque contexts is not going to preserve meaning. This is what is meant by referential opacity. The existence of referentially opaque contexts has been known since the time of Frege. He suggested a distinction between *sense* and *reference*. In ordinary formulae, the "reference" of a term/formula (i.e., its denotation) is needed, whereas in opaque contexts, the "sense" of a formula is needed (see also [147, p3]).

Clearly, classical logics are not suitable in their standard form for reasoning about intentional notions: alternative formalisms are required. A vast enterprise has sprung up devoted to developing such formalisms.

The field of formal methods for reasoning about intentional notions is widely reckoned to have begun with the publication, in 1962, of Jaakko Hintikka's book *Knowledge and Belief: An Introduction to the Logic of the Two Notions* [87]. At that time, the subject was considered fairly esoteric, of interest to comparatively few researchers in logic and the philosophy of mind. Since then, however, it has become an important research area in its own right, with contributions from researchers in AI, formal philosophy, linguistics and economics. There is now an enormous literature on the subject, and with a major biannual international conference devoted solely to theoretical aspects of reasoning about knowledge, as well as the input from numerous other, less specialized conferences, this literature is growing ever larger.

Despite the diversity of interests and applications, the number of basic techniques in use is quite small. Recall, from the discussion above, that there are two problems to be addressed in developing a logical formalism for intentional notions: a syntactic one, and a semantic one. It follows that any formalism can be characterized in terms of two independent attributes: its *language of formulation*, and *semantic model* [100, p83].

There are two fundamental approaches to the syntactic problem. The first is to use a *modal* language, which contains non-truth-functional *modal operators*, which are applied to formulae. An alternative approach involves the use of a *meta-language*: a many-sorted first-order language containing terms which denote formulae of some other *object-language*. Intentional notions can be represented using a meta-language predicate, and given

[7] Note, however, that the sentence (2.5) is itself a proposition, in that its denotation is the value true or false.

	MODAL LANGUAGE	META-LANGUAGE
POSSIBLE WORLDS	ungrounded [87] [52] [29] [83] grounded [142] [143] [58] [105] [53]	[124]
OTHER	[100] [153]	[121] [98] [125] [35]

Table 2.1: Some Intentional Logics

whatever axiomatization is deemed appropriate. Both of these approaches have their advantages and disadvantages, and will be discussed at length in the sequel.

As with the syntactic problem, there are two basic approaches to the semantic problem. The first, best known, and probably most widely used approach is to adopt a *possible worlds* semantics, where an agent's beliefs, knowledge, goals, etc. are characterized as a set of so-called *possible worlds*, with an *accessibility relation* holding between them. Possible worlds semantics have an associated *correspondence theory* which makes them an attractive mathematical tool to work with [28]. However, they also have many associated difficulties, notably the well-known logical omniscience problem, which implies that agents are perfect reasoners. A number of minor variations on the possible-worlds theme have been proposed, in an attempt to retain the correspondence theory, but without logical omniscience.

The commonest alternative to the possible worlds model for belief is to use a *sentential*, or *interpreted symbolic structures* approach. In this scheme, beliefs are viewed as symbolic formulae explicitly represented in a data structure associated with an agent. An agent then believes ϕ if ϕ is present in the agent's belief structure. Despite its simplicity, the sentential model works well under certain circumstances [100].

Table 2.1 characterizes a number of well known intentional logics in terms of their syntactic and semantic properties (after [100, p85]). The next part of this chapter contains detailed reviews of some of these formalisms. First, the idea of possible worlds semantics is discussed, and then a detailed analysis of normal modal logics is presented, along with some variants on the possible worlds theme. Next, some meta-language approaches are discussed, and one hybrid formalism is described. Finally, some alternative formalisms are described.

Before the detailed presentations, a note on terminology. Strictly speaking, an *epistemic logic* is a logic of knowledge, a *doxastic logic* is a logic of belief, and a *conative logic* is a logic of desires or goals. However, it is common practice to use "epistemic" as a blanket term for logics of knowledge and belief. This practice is adopted in this thesis; a distinction is only made where it is considered significant. Also, the reviews focus on knowledge/belief to the virtual exclusion of goals/desires; this is because most of the principles are the same, and little work has addressed the issue of goals (but see the comments on Cohen and Levesque's formalism, below).

2.3 Possible Worlds Semantics

The possible worlds model for epistemic logics was originally proposed by Hintikka ([87]), and is now most commonly formulated in a normal modal logic using the techniques developed by Kripke ([103])[8].

Hintikka's insight was to see that an agent's beliefs could be characterized in terms of a set of *possible worlds*, in the following way. Consider an agent playing the card game Gin Rummy[9]. In this game, the more one knows about the cards possessed by one's opponents, the better one is able to play. And yet complete knowledge of an opponent's cards is generally impossible, (if one excludes cheating). The ability to play Gin Rummy well thus depends, at least in part, on the ability to deduce what cards are held by an opponent, given the limited information available. Now suppose our agent possessed the ace of spades. Assuming the agent's sensory equipment was functioning normally, it would be rational of her to believe that she possessed this card. Now suppose she were to try to deduce what cards were held by her opponents. This could be done by first calculating all the various different ways that the cards in the pack could possibly have been distributed among the various players. (This is not being proposed as an actual card playing strategy, but for illustration!) For argument's sake, suppose that each possible configuration is described on a separate piece of paper. Once the process was complete, our agent can then begin to systematically eliminate from this large pile of paper all those configurations which are *not possible, given what she knows*. For example, any configuration in which she did not possess the ace of spades could be rejected immediately as impossible. Call each piece of paper remaining after this process a *world*. Each world represents one state of affairs considered possible, given what she knows. Hintikka coined the term *epistemic alternatives* to describe the worlds possible given one's beliefs. Something true in *all* our agent's epistemic alternatives could be said to be believed by the agent. For example, it will be true in all our agent's epistemic alternatives that she has the ace of spades.

On a first reading, this technique seems a peculiarly roundabout way of characterizing belief, but it has two advantages. First, it remains neutral on the subject of the cognitive structure of agents. It certainly doesn't posit any internalized collection of possible worlds. It is just a convenient way of characterizing belief. Second, the mathematical theory associated with the formalization of possible worlds is extremely appealing (see below).

The next step is to show how possible worlds may be incorporated into the semantic framework of a logic. This is the subject of the next section.

2.3.1 Normal Modal Logics

Epistemic logics are usually formulated as *normal modal logics* using the semantics developed by Kripke [103]. Before moving on to explicitly epistemic logics, this section describes normal modal logics in general.

Modal logics were originally developed by philosophers interested in the distinction between *necessary* truths and mere *contingent* truths. Intuitively, a necessary truth is something that is true *because it could not have been otherwise*, whereas a contingent truth is something that could, plausibly have been otherwise. For example, it is a fact that as I write, the Conservative Party of Great Britain hold a majority of twenty-one seats in the House of Commons. But although this is true, it is *not* a necessary truth; it could quite easily have turned out that the Labour Party won a majority at the last general election. This fact is thus only a contingent truth.

Contrast this with the following statement: *the square root of 2 is not a rational number*. There seems no earthly way that this could be anything *but* true, (given the standard reading of the sentence). This latter fact is an example of a necessary truth. Necessary truth is usually defined as something true in *all possible worlds*. It is actually quite difficult to think of any necessary truths other than mathematical laws.

To illustrate the principles of modal epistemic logics, a normal propositional modal logic is defined.

Syntax and Semantics

This logic is essentially classical propositional logic, extended by the addition of two operators: "$\Box$" (necessarily), and "$\Diamond$" (possibly). First, its syntax.

Definition 1 *Let Prop* $= \{p, q, \ldots\}$ *be a countable set of atomic propositions. The syntax of normal propositional modal logic is defined by the following rules:*

1. *If* $p \in$ *Prop then* p *is a formula.*

[8]In Hintikka's original work, he used a technique based on "model sets", which is equivalent to Kripke's formalism, though less elegant. See [88, Appendix Five, pp351–352] for a comparison and discussion of the two techniques.

[9]This example was adapted from [81].

$$
\begin{array}{llll}
\langle M, w \rangle & \models & \mathsf{true} & \\
\langle M, w \rangle & \models & p & \text{where } p \in \mathit{Prop}, \text{ iff } p \in \pi(w) \\
\langle M, w \rangle & \models & \neg\phi & \text{iff } \langle M, w \rangle \not\models \phi \\
\langle M, w \rangle & \models & \phi \vee \psi & \text{iff } \langle M, w \rangle \models \phi \text{ or } \langle M, w \rangle \models \psi \\
\langle M, w \rangle & \models & \Box\phi & \text{iff } \forall w' \in W \cdot \text{ if } (w, w') \in R \text{ then } \langle M, w' \rangle \models \phi \\
\langle M, w \rangle & \models & \Diamond\phi & \text{iff } \exists w' \in W \cdot (w, w') \in R \text{ and } \langle M, w' \rangle \models \phi
\end{array}
$$

Figure 2.1: The Semantics of Normal Modal Logic

2. *If ϕ, ψ are formulae, then so are:*
   $\mathsf{true} \quad \neg\phi \quad \phi \vee \psi$

3. *If ϕ is a formula then so are:*
   $\Box\phi \quad \Diamond\phi$

The operators "$\neg$" (not) and "$\vee$" (or) have their standard meaning; true is a logical constant, (sometimes called *verum*), that is always true. The remaining connectives of propositional logic can be defined as abbreviations in the usual way. The formula $\Box\phi$ is read: "necessarily ϕ", and the formula $\Diamond\phi$ is read: "possibly ϕ". Now to the semantics of the language.

Normal modal logics are concerned with truth at worlds; models for such logics therefore contain a set of worlds, W, and a binary relation, R, on W, saying which worlds are considered possible relative to other worlds. Additionally, a valuation function π is required, saying what propositions are true at each world.

Definition 2 *A model for a normal propositional modal logic is a triple $\langle W, R, \pi \rangle$, where W is a non-empty set of worlds, $R \subseteq W \times W$, and*

$$\pi : W \to \mathit{powerset}\ \mathit{Prop}$$

is a valuation function, which says for each world $w \in W$ which atomic propositions are true in w. An alternative, equivalent technique would have been to define π as follows:

$$\pi : W \times \mathit{Prop} \to \{\mathsf{true}, \mathsf{false}\}$$

though the rules defining the semantics of the language would then have to be changed slightly.

The semantics of the language are given via the satisfaction relation, "$\models$", which holds between pairs of the form $\langle M, w \rangle$, (where M is a model, and w is a reference world), and formulae of the language. The semantic rules defining this relation are given in Figure 2.1. The definition of satisfaction for atomic propositions thus captures the idea of truth in the "current" world, (which appears on the left of "$\models$"). The semantic rules for "true", "$\neg$", and "$\vee$", are standard. The rule for "$\Box$" captures the idea of truth in all accessible worlds, and the rule for "$\Diamond$" captures the idea of truth in at least one possible world.

Note that the two modal operators are *duals* of each other, in the sense that the universal and existential quantifiers of first-order logic are duals:

$$\Box\phi \Leftrightarrow \neg\Diamond\neg\phi.$$

It would thus have been possible to take either one as primitive, and introduce the other as a derived operator.

Correspondence Theory

To understand the extraordinary properties of this simple logic, it is first necessary to introduce *validity* and *satisfiability*. A formula is *satisfiable* if it is satisfied for some model/world pair, and *unsatisfiable* otherwise. A formula is *true in a model* if it is satisfied for every world in the model, and *valid in a class of models* if it true in every model in the class. Finally, a formula is valid *simpliciter* if it is true in the class of all models. If ϕ is valid, we write $\models \phi$.

The two basic properties of this logic are as follows. First, the following axiom schema is valid.

$$\models \Box(\phi \Rightarrow \psi) \Rightarrow (\Box\phi \Rightarrow \Box\psi)$$

NAME	AXIOM	CONDITION ON R	FIRST-ORDER CHARACTERIZATION
T	$\Box\phi \Rightarrow \phi$	Reflexive	$\forall w \in W \cdot (w, w) \in R$
D	$\Box\phi \Rightarrow \Diamond\phi$	Serial	$\forall w \in W \cdot \exists w' \in W \cdot (w, w') \in R$
4	$\Box\phi \Rightarrow \Box\Box\phi$	Transitive	$\forall w, w', w'' \in W \cdot (w, w') \in R \wedge (w', w'') \in R \Rightarrow (w, w'') \in R$
5	$\Diamond\phi \Rightarrow \Box\Diamond\phi$	Euclidean	$\forall w, w', w'' \in W \cdot (w, w') \in R \wedge (w, w'') \in R \Rightarrow (w', w'') \in R$

Table 2.2: Some Correspondence Theory

This axiom is called K, in honour of Kripke. The second property is as follows.

If $\models \phi$ then $\models \Box\phi$

Proofs of these properties are trivial, and are left as an exercise for the reader. Now, since K is valid, it will be a theorem of any complete axiomatization of normal modal logic. Similarly, the second property will appear as a rule of inference in any axiomatization of normal modal logic; it is generally called the *necessitation* rule. These two properties turn out to be the most problematic features of normal modal logics when they are used as logics of knowledge/belief (this point will be examined later).

The most intriguing properties of normal modal logics follow from the properties of the accessibility relation, R, in models. To illustrate these properties, consider the following axiom schema.

$$\Box\phi \Rightarrow \phi$$

It turns out that this axiom is *characteristic* of the class of models with a *reflexive* accessibility relation. (By characteristic, we mean that it is true in all and only those models in the class.) There are a host of axioms which correspond to certain properties of R: the study of the way that properties of R correspond to axioms is called *correspondence theory*. In Table 2.2, we list some axioms along with their characteristic property on R, and a first-order formula describing the property. Note that the table only lists those axioms of specific interest to this thesis; (see [28] for others). The names of axioms follow historical tradition.

The results of correspondence theory make it straightforward to derive completeness results for a range of simple normal modal logics. These results provide a useful point of comparison for normal modal logics, and account in a large part for the popularity of this style of semantics.

A *system of logic* can be thought of as a set of formulae valid in some class of models; a member of the set is called a *theorem* of the logic (if ϕ is a theorem, this is usually denoted by $\vdash \phi$). The notation $K\Sigma_1 \ldots \Sigma_n$ is often used to denote the smallest normal modal logic containing axioms $\Sigma_1, \ldots, \Sigma_n$ (recall that any normal modal logic will contain K; cf. [78, p25]).

For the axioms T, D, 4, and 5, it would seem that there ought to be sixteen distinct systems of logic (since $2^4 = 16$). However, some of these systems turn out to be equivalent (in that they contain the same theorems), and as a result there are only eleven distinct systems. The relationships between these systems are described in Figure 2.2 (after [100, p99], and [28, p132]). In this diagram, an arc from A to B means that B is a strict superset of A: every theorem of A is a theorem of B, but not *vice versa*; $A = B$ means that A and B contain precisely the same theorems. Because some modal systems are so widely used, they have been given names:

KT	is known as	T
$KT4$	is known as	S4
$KD45$	is known as	weak-S5
$KT5$	is known as	S5

Normal Modal Logics as Epistemic Logics

To use the logic developed above as an epistemic logic, the formula $\Box\phi$ is read as: "it is known that ϕ". The worlds in the model are interpreted as epistemic alternatives, the accessibility relation defines what the alternatives

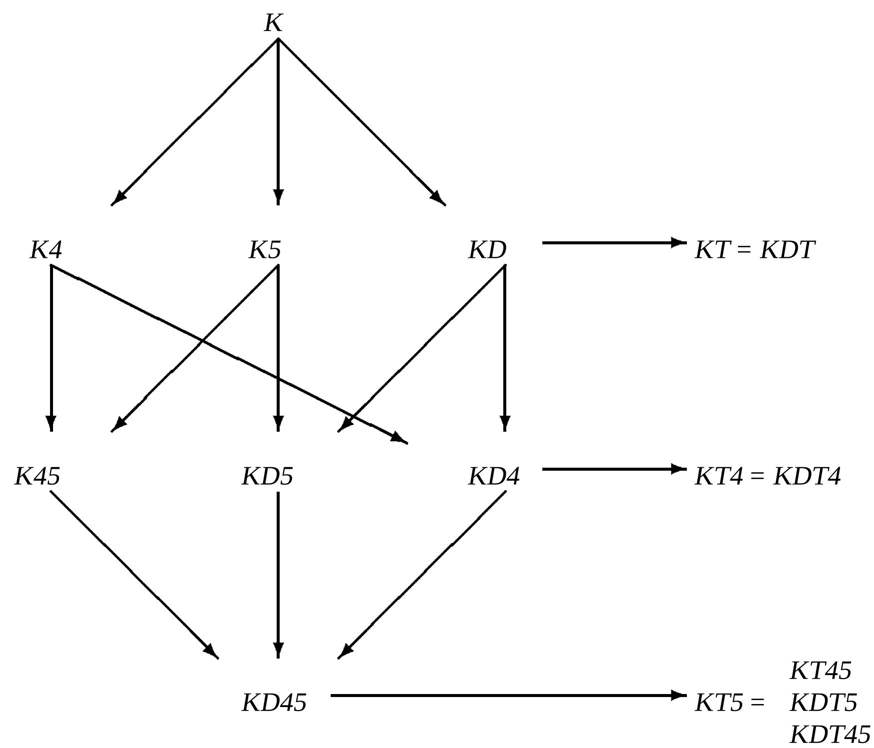

Figure 2.2: The Modal Systems Based on Axioms *T*, *D*, *4* and *5*

are from any given world. The logic deals with the knowledge of a single agent. To deal with multi-agent knowledge, one adds to a model structure an indexed set of accessibility relations, one for each agent. A model is then a structure:

$$\langle W, R_1, \ldots, R_n, \pi \rangle$$

where R_i is the knowledge accessibility relation of agent i. The simple language defined above is extended by replacing the single modal operator "$\Box$" by an indexed set of unary modal operators $\{\mathsf{K}_i\}$, where $i \in \{1, \ldots, n\}$. The formula $\mathsf{K}_i\phi$ is read: "i knows that ϕ". The semantic rule for "$\Box$" is replaced by the following rule:

$$\langle M, w \rangle \models \mathsf{K}_i\phi \quad \text{iff } \forall w' \in W \cdot \text{ if } (w, w') \in R_i \text{ then } \langle M, w' \rangle \models \phi$$

Each operator K_i thus has exactly the same properties as "$\Box$". Corresponding to each of the modal systems Σ, above, a corresponding system Σ_n is defined, for the multi-agent logic. Thus K_n is the smallest multi-agent epistemic logic and $\mathrm{S5}_n$ is the largest.

The next step is to consider how well normal modal logic serves as a logic of knowledge/belief. Consider first the necessitation rule and axiom *K*, since any normal modal system is committed to these.

The necessitation rule tells us that an agent knows all valid formulae. Amongst other things, this means an agent knows all propositional tautologies. Since there are an infinite number of these, an agent will have an infinite number of items of knowledge: immediately, one is faced with a counter intuitive property of the knowledge operator.

Now consider the axiom *K*, which says that an agent's knowledge is closed under implication. Suppose ϕ is a logical consequence of the set $\Phi = \{\phi_1, \ldots, \phi_n\}$, then in every world where all of Φ are true, ϕ must also be true, and hence

$$\phi_1 \wedge \cdots \wedge \phi_n \Rightarrow \phi$$

must be valid. By necessitation, this formula will also be believed. Since an agent's beliefs are closed under implication, whenever it believes each of Φ, it must also believe ϕ. Hence an agent's knowledge is closed under logical consequence. This also seems counter intuitive. For example, suppose, like every good logician, our agent knows Peano's axioms. It may well be that Fermat's last theorem follows from Peano's axioms — although, despite strenuous efforts, nobody has so far managed to prove it. Yet if our agent's beliefs are closed under logical consequence, then our agent must know it. So consequential closure, implied by necessitation and the K axiom, seems an overstrong property for resource bounded reasoners.

The Logical Omniscience Problem

These two problems — that of knowing all valid formulae, and that of knowledge/belief being closed under logical consequence — together constitute the famous *logical omniscience* problem. This problem has some damaging corollaries.

The first concerns consistency. Human believers are rarely consistent in the logical sense of the word; they will often have beliefs ϕ and ψ, where $\phi \vdash \neg \psi$, without being aware of the implicit inconsistency. However, the ideal reasoners implied by possible worlds semantics cannot have such inconsistent beliefs without believing *every* formula of the logical language (because the consequential closure of an inconsistent set of formulae is the set of all formulae). Konolige has argued that logical consistency is much too strong a property for resource bounded reasoners: he argues that a lesser property, that of being *non-contradictory* is the most one can reasonably demand [100]. Non-contradiction means that an agent would not simultaneously believe ϕ and $\neg \phi$, although the agent might have logically inconsistent beliefs.

The second corollary is more subtle. Consider the following propositions (this example is from [100, p88]):

1. Hamlet's favourite colour is black.

2. Hamlet's favourite colour is black *and* every planar map can be four coloured.

The second conjunct of (2) is valid, and will thus be believed. This means that (1) and (2) are logically equivalent; (2) is true just when (1) is. Since agents are ideal reasoners, they will believe that the two propositions are logically equivalent. This is yet another counter intuitive property implied by possible worlds semantics, as: "equivalent propositions are *not* equivalent as beliefs" [100, p88]. Yet this is just what possible worlds semantics implies. It has been suggested that propositions are thus too *coarse grained* to serve as the objects of belief in this way.

The logical omniscience problem is a serious one. In the words of Levesque:

> "Any one of these [problems] might cause one to reject a possible-world formalization as unintuitive at best and completely unrealistic at worst". [112]

Axioms for Knowledge and Belief

We now consider the appropriateness of the axioms D_n, T_n, 4_n, and 5_n for logics of knowledge/belief.

The axiom D_n says that an agent's beliefs are non-contradictory; it can be re-written in the following form:

$$\mathsf{K}_i\phi \;\Rightarrow\; \neg\,\mathsf{K}_i\neg\,\phi$$

which is read: "if i knows ϕ, then i doesn't know $\neg\,\phi$". This axiom seems a reasonable property of knowledge/belief.

The axiom T_n is often called the *knowledge* axiom, since it says that what is known is true. It is usually accepted as the axiom that distinguishes knowledge from belief: it seems reasonable that one could believe something that is false, but one would hesitate to say that one could *know* something false. Knowledge is thus often defined as true belief: i knows ϕ if i believes ϕ and ϕ is true. So defined, knowledge satisfies T_n.

Axiom 4_n is called the *positive introspection axiom*. Introspection is the process of examining one's own beliefs, and is discussed in detail in [100, Chapter 5]. The positive introspection axiom says that an agent knows what it knows. Similarly, axiom 5_n is the *negative introspection axiom*, which says that an agent is aware of what it doesn't know. Positive and negative introspection together imply an agent has perfect knowledge about what it does and doesn't know (cf. [100, Equation (5.11), p79]). Whether or not the two types of introspection are appropriate properties for knowledge/belief is the subject of some debate. However, it is generally accepted that positive introspection is a less demanding property than negative introspection, and is thus a more reasonable property for resource bounded reasoners.

Given the comments above, the modal system $S5_n$ is often chosen as a logic of *knowledge*, and weak-$S5_n$ is often chosen as a logic of *belief*.

Computational Aspects

Before leaving this basic logic, it is worth commenting on its computational/proof theoretic properties. Halpern and Moses have established the following ([83]):

1. The provability problem for each of the key systems K_n, T_n, $S4_n$, weak-$S5_n$, and $S5_n$ *is* decidable. Halpern and Moses sketch some tableaux decision procedures for these logics.

2. The satisfiability and validity problems for K_n, T_n, $S4_n$, (where $n \geq 1$), and $S5_n$, weak-$S5_n$ (where $n \geq 2$) are PSPACE complete.

The first result is encouraging, as it holds out at least some hope of automation. Unfortunately, the second result is extremely discouraging: in simple terms, it means that in the worst case, automation of these logics is not a practical proposition.

Discussion

To sum up, the basic possible worlds approach described above has the following disadvantages as a multi-agent epistemic logic:

- agents believe all valid formulae;
- agents beliefs are closed under logical consequence;
- equivalent propositions are identical beliefs;
- if agents are inconsistent, then they believe everything;
- in the worst case, automation is not feasible.

To which many people would add the following:

- "[T]he ontology of possible worlds and accessibility relations …is frankly mysterious to most practically minded people, and in particular has nothing to say about agent architecture". [147]

Despite these serious disadvantages, possible worlds are still the semantics of choice for many researchers, and a number of variations on the basic possible worlds theme have been proposed to get around some of the difficulties. The following sections examine various topics associated with possible worlds semantics.

2.3.2 Common and Distributed Knowledge

In addition to reasoning about what one agent knows or believes, it is often useful to be able to reason about "cultural" knowledge: the things that everyone knows, and that everyone knows that everyone knows, etc. This kind of knowledge is called *common knowledge*. The famous "wisest man" puzzle — a classic problem in epistemic reasoning — is an example of the kind of problem that is efficiently dealt with via reasoning about common knowledge (see, e.g., [100, p58] for a statement of the wisest man problem)[10].

The starting point for common knowledge is to develop an operator for things that "everyone knows". A unary modal operator EK is added to the modal language discussed above; the formulae $\mathsf{EK}\phi$ is read: "everyone knows ϕ". It can be defined as an abbreviation:

$$\mathsf{EK}\phi \triangleq \mathsf{K}_1\phi \wedge \cdots \wedge \mathsf{K}_n\phi$$

or it can be given its own semantic rule:

$$\langle M, w\rangle \models \mathsf{EK}\phi \quad \text{iff } \langle M, w\rangle \models \mathsf{K}_i\phi \text{ for all } i \in \{1, \ldots, n\}$$

The EK operator does not satisfactorily capture the idea of common knowledge. For this, another derived operator CK is required; CK is defined, ultimately, in terms of EK. It is first necessary to introduce the derived operator EK^k; the formula $\mathsf{EK}^k\phi$ is read: "everyone knows ϕ to degree k". It is defined as follows:

$$\begin{array}{lcl} \mathsf{EK}^1\phi & \triangleq & \mathsf{EK}\phi \\ \mathsf{EK}^{k+1}\phi & \triangleq & \mathsf{EK}(\mathsf{EK}^k\phi) \end{array}$$

[10] The discussion that follows was adapted and expanded from [73, Chapter 9] and [83].