A modal logic of information change

Joeri Engelfriet* Yde Venema†

Abstract

We study the dynamics of information change, using modal logic as a vehicle. Our semantic perspective is that of a supermodel in which a state represents some agent's information, and the accessibility relations are those of increasing and decreasing knowledge. We concentrate on two specific settings in which an information state consists of all valuations that are models for some propositional formula, or theory, respectively; treating such a set of valuations as an epistemic S5-model, allows us to interpret epistemic formulas in it in the standard fashion. For the validities of one of these two supermodels we provide a Hilbert-style derivation system; our main technical result shows this derivation system to be sound and complete.

1 Introduction

Agents in a dynamic world have to deal with changing information. The information they have about the world may change as a result of performing observations, communication with other agents, or through nonmonotonic reasoning (where an agent makes certain plausible assumptions). The most basic kinds of change are increase and decrease of knowledge, and in a sense all changes in information can be seen as combinations of these basic kinds. In this paper, we study changing knowledge, using modal logic as a vehicle. In our models, the worlds represent information states, and there is one modal accessibility relation representing increase of knowledge. This relation is used to interpret two modal operators $\Diamond_u$ and $\Diamond_d$. The formula $\Diamond_u\varphi$ informally means: "It is possible to increase your knowledge to a state where φ holds" (*update*), and $\Diamond_d\varphi$ means: "It is possible to decrease your knowledge to a state where φ holds" (*downdate*).

Many further choices have to be made in formalizing these intuitions of an information state. We will take propositional logic as the basic logic in which the information of an agent is expressed. In order to describe what the agent knows and does not know, we add a knowledge operator K. This again suggests a modal approach, and we use S5 for this purpose. As our information states we take what are probably the simplest models for S5, namely, sets of valuations. Via the standard modal semantics for epistemic logic, every such S5-model naturally determines a collection of known facts.

Our 'supermodels' group together such information states; the accessibility relation in such supermodels connects two states if the agent knows more in one state than in the other. Still many choices remain: do we use a finite (propositional) language or an infinite one, do we allow an agent to possess an inherently infinite amount of knowledge or not, are there further constraints on the accessibility relation, etc. It turns out that many of these choices really affect the logic we obtain.

There is by now extensive literature on formal models of information change, and in particular, on modal approaches. For a survey, the reader is referred to Chapter 10 in [1]. Our system is

*Faculty of Mathematics and Computer Science, Vrije Universiteit Amsterdam, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands, joeri@cs.vu.nl.

†Institute for Logic, Language and Computation, University of Amsterdam, Plantage Muidergracht 24, 1018 TV Amsterdam. E-mail: yde@wins.uva.nl. The research of this author has been made possible by a fellowship of the Royal Netherlands Academy of Arts and Sciences.

closely related to the update semantics of [8]; the main difference is that in our proposal we do not model the specific piece of information inducing the change in the agent's knowledge.

Our formalism is inspired by the temporal epistemic logic MTEL of [4] which was used to study temporal aspects of information change, and applications to nonmonotonic reasoning. A connection with the latter field lies in Halpern and Moses' logic of "only knowing", cf. [5]. This logic aims to answer questions like: "What do I know and what do I not know, if I *only* know p?" (for instance, if you only know p, you do not know q). In our 'supermodels', the agent only knows φ in an information state that satisfies $K\varphi$ in such a minimal way that any decrease in information leads to a state where $K\varphi$ no longer holds; the formula for this is $K\varphi \wedge \Box_d \neg K\varphi$. We shall see that a faithful translation of Halpern and Moses' consequence relation in our logic exists. The problem with their consequence relation is that it is hard to axiomatize directly (as yet, no-one has come up with a direct axiomatization), as is the case for many nonmonotonic logics. Via a translation into a (monotonic) logic with an axiomatization, proofs for this logic can be carried out. Such an approach was taken by Levesque ([7]), who introduced a modal operator $\mathcal{O}$, where $\mathcal{O}\varphi$ means that the agent only knows φ. We will discuss the connection between his formalism and ours at the end of section 2.

In Halpern and Moses' logic the consequence relation is defined in terms of a preference relation on S5-models, which prefers models with less knowledge to models with more knowledge. Consequences of a formula φ are those formulas true in all most preferred models of φ. As such, it falls into the more general scheme of preferential logics studied in Artificial Intelligence (see [6]). In our supermodels, the preference relation is in fact the modal accessibility relation (of decreasing knowledge). The idea of studying (and axiomatizing) preferential logics by considering the preference relation as an accessibility relation in a large model, has been used in [2]. In that paper, Boutilier gives axiomatizations which are sound with respect to certain classes of preferential models. The difference with our approach is twofold: in the first place, our states are not propositional valuations, but S5-models. This reflects a difference in focus: we concentrate on the dynamics of *knowledge*. In the second place, we are interested in special kinds of preference relations, namely those that reflect an increase or decrease in knowledge; hence, our preference relation is completely determined by the states.

2 The supermodels

In our formalization of information change we take a layered approach. On the base level, we are dealing with a propositional logic in which the agent's information is expressed. On top of that we have an epistemic language; since we restrict ourselves to the single agent case in this paper, we add one single knowledge operator K to the base language; the informal reading of $K\alpha$ is that the agent knows α. The top level language is then obtained by adding two more operators $\Diamond_d$ and $\Diamond_u$ to this epistemic language; $\Diamond_d\varphi$ ($\Diamond_u\varphi$) is to be read as "it is possible to decrease (increase) your knowledge to a state where φ holds". Furthermore, we should mention that as building blocks of the top level language we take *subjective* formulas, in which every propositional formula is in the scope of a K-operator.

Definition 2.1 (syntax) *We fix a set $\mathbb{V}$ of propositional variables p_0, p_1, ..., q, r, ... $\mathcal{L}_0$ is the **base language** of classical propositional logic over this set.*

*At the intermediate level, $\mathcal{L}_1$ denotes the **epistemic language** over $\mathbb{V}$. An epistemic formula is **subjective** if every propositional variable occurs in the scope of a knowledge operator.*

*Finally, our **top language** $\mathcal{L}_2$ is defined as the set of formulas obtained by closing the set of subjective formulas under the boolean connectives and the unary modal operators ('diamonds') $\Diamond_d$ and $\Diamond_u$. $\mathcal{L}_{2d}$ consists of all $\mathcal{L}_2$-formulas in which the operator $\Diamond_u$ does not occur. We use the 'only' operator O as the following abbreviation: $\mathsf{O}\varphi = \varphi \wedge \Box_d \neg\varphi$.*

The (meta-)variables α, β, γ, ... will be used to range over formulas in the base language; for the epistemic formulas we use μ, ν, ρ, ...; and for $\mathcal{L}_2$-formulas we use φ, ψ, χ, ... Note that $\Diamond_d(p \wedge Kq)$ is *not* a well-formed $\mathcal{L}_2$-formula, since $p \wedge Kq$ is not subjective.

Let us briefly discuss the intuitive meaning of the formula $\mathsf{O}K\alpha$. It says that the agent knows α, but in a sort of maximal sense: it no longer knows α after losing any piece of knowledge. This indicates that the agent has no *extra* knowledge that it might lose, apart from α. In other words, the agent *only knows* α. Thus we see that

$\mathsf{O}K\alpha$ is our formalization of *only knowing* α.

Let us now turn to a formal definition of the semantics for these languages. First we consider $\mathcal{L}_0$ and $\mathcal{L}_1$.

Definition 2.2 *Let $\mathcal{V}$ denote the set of **valuations**, that is, mappings from $\mathbb{V}$ to $\{0,1\}$; elements of $\mathcal{V}$ will also be called **worlds**. As variables ranging over valuations we use w, v, u, ... We assume familiarity with the classical propositional truth definition; truth of a formula α under a valuation w (in a set $\mathfrak{q}$ of valuations) is denoted as $w \models \alpha$ ($\mathfrak{q} \models \alpha$, respectively). Given a set Δ of propositional formulas, define $Mod(\Delta)$ as the set of valuations w such that $w \models \Delta$.*

*A **model** is any non-empty subset of $\mathcal{V}$; the set of all models is denoted by M^+. Later on, when we will view models as constituents of bigger entities, we will also use the term **information state** for a model.*

***Truth** of an epistemic formula μ in a model $\mathfrak{m}$ at a world w, denoted by $\mathfrak{m}, w \Vdash \mu$, is defined by a standard recursive definition. For instance, the clause for K is that $\mathfrak{m}, w \Vdash K\mu$ iff $\mathfrak{m}, v \Vdash \mu$ for all $v \in \mathfrak{m}$. An epistemic formula is **true in a model** $\mathfrak{m}$ if it is true at every world in that model.*

In other words, the kind of models for the epistemic language that we are considering are the simplest S5-models. As we mentioned in the introduction, our basic idea is to gather various models into one 'supermodel' which also imposes an information ordering on the models. Let $\sqsubset$ denote the following **information ordering** on models:

$$\mathfrak{m} \sqsubset \mathfrak{n} \text{ iff } \mathfrak{n} \subset \mathfrak{m}.$$

Here $\subset$ denotes strict set inclusion. The underlying idea of this definition is that $\mathfrak{n}$ contains more information than $\mathfrak{m}$ if $\mathfrak{n}$ consists of less worlds than $\mathfrak{m}$. Now the relation $\sqsubset$ is indeed an information ordering: if $\mathfrak{m} \sqsubset \mathfrak{n}$ then at $\mathfrak{n}$ the agent possesses at least as much information as at $\mathfrak{m}$, in the sense that $\mathfrak{n} \Vdash K\alpha$ whenever $\mathfrak{m} \Vdash K\alpha$. We will now give *three* alternative options for the definition of the supermodel.

Definition 2.3 *A set of valuations $\mathfrak{m}$ is called **closed** if $\mathfrak{m} = Mod(\Gamma)$ for some set Γ of propositional formulas, **clopen** if it is of the form $Mod(\gamma)$ for some propositional formula γ. The sets of closed and clopen models are denoted by M and M_f, respectively. Finally, the **supermodels** $\mathcal{S}^+$, $\mathcal{S}$ and $\mathcal{S}_f$ are defined by: $\mathcal{S}^+ = (M^+, \sqsubset)$, $\mathcal{S} = (M, \sqsubset)$ and $\mathcal{S}_f = (M_f, \sqsubset)$, respectively.*

*Given a closed model $\mathfrak{m}$, we let $\Delta_{\mathfrak{m}}$ denote the **diagram** of $\mathfrak{m}$, that is, the set of classical formulas holding at $\mathfrak{m}$ — this gives $\mathfrak{m} = Mod(\Delta_{\mathfrak{m}})$. For a clopen $\mathfrak{m}$, $\delta_{\mathfrak{m}}$ denotes some (canonically chosen) formula such that $\mathfrak{m} = Mod(\delta_{\mathfrak{m}})$.*

*Now given these models, we define the notion of **truth** of an $\mathcal{L}_2$-formula at an information state as follows (as an example we take $\mathcal{S}$)*

$$\begin{array}{lll}
\mathcal{S}, \mathfrak{m} \Vdash \mu & \textit{if} & \mathfrak{m} \Vdash \mu \\
\mathcal{S}, \mathfrak{m} \Vdash \neg\varphi & \textit{if} & \mathcal{S}, \mathfrak{m} \nVdash \varphi \\
\mathcal{S}, \mathfrak{m} \Vdash \varphi \wedge \psi & \textit{if} & \mathcal{S}, \mathfrak{m} \Vdash \varphi \textit{ and } \mathcal{S}, \mathfrak{m} \Vdash \psi \\
\mathcal{S}, \mathfrak{m} \Vdash \Diamond_d\varphi & \textit{if} & \mathcal{S}, \mathfrak{n} \Vdash \varphi \textit{ for some } \mathfrak{n} \in \mathcal{S} \textit{ with } \mathfrak{n} \sqsubset \mathfrak{m} \\
\mathcal{S}, \mathfrak{m} \Vdash \Diamond_u\varphi & \textit{if} & \mathcal{S}, \mathfrak{n} \Vdash \varphi \textit{ for some } \mathfrak{n} \in \mathcal{S} \textit{ with } \mathfrak{m} \sqsubset \mathfrak{n}.
\end{array}$$

*A formula φ is **valid** in $\mathcal{S}$, denoted as $\mathcal{S} \Vdash \varphi$, if $\mathcal{S}, \mathfrak{m} \Vdash \varphi$ for all $\mathfrak{m} \in M$ (and analogously for $\mathcal{S}^+$ and $\mathcal{S}_f$).*

We believe the supermodels $\mathcal{S}$ and $\mathcal{S}_f$ to have some advantages over $\mathcal{S}^+$. The main one is that in $\mathcal{S}^+$, the fact that $\mathfrak{m} \sqsubset \mathfrak{n}$ not necessarily implies that the agent has *strictly* more knowledge in $\mathfrak{n}$ than in $\mathfrak{m}$. Consider for instance the case where $\mathfrak{n} = \mathcal{V}$ and $\mathfrak{m} = \mathcal{V} \setminus \{w\}$ for some valuation w. It is not difficult to prove that for all propositional formulas α, $\mathfrak{m} \models \alpha$ iff $\mathfrak{n} \models \alpha$ (using the fact that $\mathbb{V}$ is infinite). This gives that $\mathfrak{m} \Vdash \mu$ iff $\mathfrak{n} \Vdash \mu$ for all epistemic formulas μ. But $\mathfrak{m}$ is properly included in $\mathfrak{n}$! This problem cannot occur with closed sets: it is rather easy to show that a model $\mathfrak{m}$ is closed if and only if it contains all valuations w such that $w \models \{\alpha \mid \mathfrak{m} \models \alpha\}$. In fact, both $\mathcal{S}_f$ and $\mathcal{S}$ behave nicely in this respect, as the following Proposition shows.

Proposition 2.4 *1. If $\mathfrak{m}$ is closed, then $\mathfrak{m} \Vdash K\alpha$ iff $\Delta_{\mathfrak{m}} \vdash \alpha$.*

2. If $\mathfrak{m}$ is clopen, then $\mathfrak{m} \Vdash K\alpha$ iff $\delta_{\mathfrak{m}} \vdash \alpha$.

3. For closed models $\mathfrak{m}$ and $\mathfrak{n}$, $\mathfrak{m} \sqsubset \mathfrak{n}$ iff $\Delta_{\mathfrak{m}} \subset \Delta_{\mathfrak{n}}$.

4. For clopen models $\mathfrak{m}$ and $\mathfrak{n}$, $\mathfrak{m} \sqsubset \mathfrak{n}$ iff $\vdash \delta_{\mathfrak{n}} \rightarrow \delta_{\mathfrak{m}}$ and $\not\vdash \delta_{\mathfrak{m}} \rightarrow \delta_{\mathfrak{n}}$.

It follows immediately from Proposition 2.4 that a clopen model $\mathfrak{m}$ is the *only* state in the clopen supermodel where the formula $\mathsf{O}K\delta_{\mathfrak{m}}$ holds. From this perspective we can say that every state of the clopen model has a *name*. More precisely, we can prove that for any propositional formula α and $\mathfrak{m} \in M_f$ it holds that $\mathcal{S}_f, \mathfrak{m} \Vdash \mathsf{O}K\alpha$ iff $\mathfrak{m} = Mod(\alpha)$. The motivation for choosing either $\mathcal{S}$ or $\mathcal{S}_f$ will come from the intuitions concerning knowledge that one wants to model. It might be more realistic to allow only information states in which the agent has a (unbounded) *finite* amount of knowledge; in that case, $\mathcal{S}_f$ seems to be the natural choice. If one finds it more natural to allow the agent to possess an (inherently) infinite amount of knowledge, one should obviously opt for $\mathcal{S}$.

Finally, we believe it is simply very *interesting* to see how these choices affect the properties of the models, and in particular, the properties of the induced logics. As an example, we consider the nature of the ordering relation; one can show that $\sqsubset$ and $\sqsupset$ are *discrete* on $\mathcal{S}^+$ and *dense* on $\mathcal{S}_f$, while only $\sqsupset$ is discrete on $\mathcal{S}$, and $\sqsubset$ is neither dense nor discrete. Perhaps surprisingly, this difference between the models $\mathcal{S}$ and $\mathcal{S}_f$ is not reflected in the logic, at least, not if we restrict ourselves to downdate-formulas. Note that this implies that the behaviour of *only knowing* does not depend on a choice between $\mathcal{S}$ and $\mathcal{S}_f$ as our supermodel.

Proposition 2.5 *Let φ be a formula in $\mathcal{L}_{2d}$. Then for any clopen state $\mathfrak{n}$: $\mathcal{S}, \mathfrak{n} \Vdash \varphi$ iff $\mathcal{S}_f, \mathfrak{n} \Vdash \varphi$, and $\mathcal{S} \Vdash \varphi$ iff $\mathcal{S}_f \Vdash \varphi$.*

In the introduction, we mentioned the fact that all changes in information can be seen as a combination of decrease (throw away the old information) and increase (add the new information). In our supermodels (taking $\mathcal{S}_f$ as an example), this is indeed the case: if μ is a subjective epistemic ($\mathcal{L}_1$) formula that is S5-satisfiable, then from any state we can reach a state where μ is the case by (possibly) performing a downdate, (possibly) followed by an update. The reader can check that in fact, any subjective epistemic formula μ is S5-sastisfiable iff $\mathcal{S}_f \Vdash \mu \vee \Diamond_d\mu \vee \Diamond_u\mu \vee \Diamond_d\Diamond_u\mu$. The proof system we will present in the next section axiomatizes validity in $\mathcal{S}_f$; by our previous observation then, we have a proof system for non-validity in S5. (The restriction to subjective formulas is not severe, since a non-subjective formula μ is S5-satisfiable if and only if $M\mu$ is satisfiable.)

In the remainder of this section, we will briefly consider the relation between our approach and two others. We start with Halpern & Moses' logic of 'only knowing' ([5]).

Let μ be an S5-formula. A model $\mathfrak{m}$ is a *maximal model* of μ if $\mathfrak{m} \Vdash \mu$ and there exists no model $\mathfrak{n}$ with $\mathfrak{m} \subset \mathfrak{n}$ and $\mathfrak{n} \Vdash \mu$. A formula μ is called *honest* if it possesses a unique maximal

model. For an honest μ, define $\mu \mathrel{|\!\sim} \nu$ if ν is true in the unique maximal model of μ. We have the following result.

Proposition 2.6 *Let μ be an honest formula, and ν any S5-formula. Then*

$$\mu \mathrel{|\!\sim} \nu \text{ iff } \mathcal{S}_f \Vdash \mathsf{O}K\mu \to K\nu$$

Also, we can *characterize* honesty in $\mathcal{S}_f$.

Proposition 2.7 *Let $\mu \in \mathcal{L}_1$, then μ is honest iff for any $\varphi \in \mathcal{L}_2$ (or even $\varphi \in \mathcal{L}_1$), either $\mathcal{S}_f \Vdash \mathsf{O}K\mu \to \varphi$ or $\mathcal{S}_f \Vdash \mathsf{O}K\mu \to \neg\varphi$, iff there exists an $\alpha \in \mathcal{L}_0$ such that $\mathcal{S}_f \Vdash \mathsf{O}K\mu \leftrightarrow \mathsf{O}K\alpha$.*

Proposition 2.6 means we can use the proof system for $\mathcal{S}_f$ of Section 3 to prove all entailments in Halpern & Moses' logic. But this is not a new accomplishment. In [7], Levesque introduces a modal logic with an operator $\mathcal{O}$, where $\mathcal{O}\alpha$ means the agent only knows α. An axiom system for this logic is given, which can be used to prove entailments in the logic of Halpern and Moses. We will now briefly review Levesque's logic.

The Kripke models he considers are closed sets of valuations, so they are just the elements of our M. A modal operator B has almost the same semantics as our K operator, the difference being that B is a *belief* operator satisfying only the $K45$ axioms. There is a second modal operator, N, where $N\alpha$ intuitively means that "α *at most* is believed to be false" (dual to the intuition that $B\alpha$ means that "α is *at least* believed to be true"). Finally, $\mathcal{O}\alpha$ is defined as $B\alpha \wedge N\neg\alpha$.

The differences with our formalism are twofold. First, Levesque's language is static since it cannot express any *change* in the agent's knowledge. And second, Levesque needs an extra operator N for his axiomatization, the meaning of which in nested cases is not very intuitive. Nevertheless, we can show that in our formalism, we can express Levesque's operator $\mathcal{O}\alpha$ by our formula $\mathsf{O}K\alpha$.

3 A proof system

In this section, we will present a Hilbert-style proof system for validities of $\mathcal{S}_f$. Before we present the axioms and rules, we first need to distinguish a special class of formulas. Define the class DP of *downward persistent* formulas as follows:

$$\mathsf{DP} ::= M(\alpha) \mid \mathsf{DP} \vee \mathsf{DP} \mid \mathsf{DP} \wedge \mathsf{DP} \mid M(\mathsf{DP})$$

where α is any propositional formula, and $M\mu$ is an abbreviation of $\neg K\neg\mu$. The name 'downward persistent' is explained by a result in [4] stating these formulas are the only subjective formulas (up to S5-equivalence) for which it holds that $\mathfrak{n} \sqsubset \mathfrak{m}$ & $\mathfrak{m} \Vdash \mu \Rightarrow \mathfrak{n} \Vdash \mu$.

We are now ready to give the proof system. The expression $\vdash_{CL} \alpha$ denotes the fact that α is provable in classical propositional logic.

Definition 3.1 *The proof system* IC *(for* information change*) consists of the following axioms, besides the classical tautologies and the modal distribution axioms:*

$A1$	$K\mu \to \mu$	
$A2$	$K\mu \to KK\mu$	
$A3$	$\neg K\mu \to K\neg K\mu$	
CV	$\varphi \to \Box_u\Diamond_d\varphi$	
	$\varphi \to \Box_d\Diamond_u\varphi$	
DP	$\mu \to \Box_d\mu$	*whenever μ is in* DP
SF	$(K\alpha \wedge K\beta) \to \Diamond_d(K\alpha \wedge \neg K\beta)$	*provided that* $\nvdash_{CL} \alpha \to \beta$
OD	$(K\alpha \wedge \Diamond_d K\alpha) \to \Diamond_d \mathsf{O}K\alpha$	
OU	$(\mathsf{O}K\beta \wedge \neg K\alpha) \to \Diamond_u \mathsf{O}K\alpha$	*whenever* $\vdash_{CL} \alpha \to \beta$
4	$\Diamond_d\Diamond_d\varphi \to \Diamond_d\varphi$	

Its derivation rules are Modus Ponens, Necessitation for all boxes, and the rule OE of O*-Elimination:*

$$\frac{\{\mathsf{O}K\alpha \to \varphi \mid \alpha \in \mathcal{L}_0(\mathbb{V}_\varphi \cup \{p\})\}}{\varphi}$$

Here α ranges over the finite set of classical formulas that can be built using the propositional variables in φ and one new letter p, modulo equivalence.

The axioms SF and OU depend on propositional provability. Since this is decidable, the set of axioms of IC is recursive.

Given the system IC, the notions of derivation, proof, theorem, consistency and the like are standard. The system IC axiomatizes validity in $\mathcal{S}_f$:

Theorem 3.2 (Soundness and Completeness) *For every $\mathcal{L}_2$-formula φ:*

$$\mathcal{S}_f \Vdash \varphi \iff \vdash_{\mathsf{IC}} \varphi.$$

We will proceed with an informal discussion of the axioms and rules.

The axioms $A1$ through $A3$ are standard axioms for S5. As we are only considering subjective formulas, one could replace $A1$ by the axiom $\neg K\bot$ (the resulting system is often called KD45). The CV axioms ('converse') express the fact that $\sqsubset$ and $\sqsupset$ are each other's converses. The axiom DP ('downward persistence') expresses that possibility is preserved when the agent's information decreases. The axiom SF ('selective forgetting') states that when the agent knows something (β), then it can perform a downdate to forget it. Any knowledge it previously held (α) can be retained, provided it does not imply β. The axiom OD ('O-down') expresses the fact that when $K\alpha$ holds in a state, it is either a maximal model of $K\alpha$ ($\mathsf{O}K\alpha$ holds in it), or such a state can be reached by a downdate. Now let us consider axiom OU ('O-up'). If $\vdash_{CL} \alpha \to \beta$, then α contains more information than β. If an agent only knows β (and not α; this is the case whenever $\nvdash_{CL} \beta \to \alpha$), then it may perform an update to a state where all it knows is α. The last rule, **OE** is perhaps the most complicated. It states that in order to prove a formula φ, it is sufficient to prove that φ is true in a number of named states. These states differ in the knowledge they have about the propositional variables mentioned in φ, but also with respect to any extra knowledge. This 'extra' knowledge only requires mentioning one propositional variable not occurring in φ.

4 Conclusions

We studied the dynamics of information change by proposing a modal logic of increasing and decreasing information; this logic is the logic of a specific 'super model' in which the states themselves are models of an epistemic language. In defining this particular set-up there were a lot of different choices to be made, in many different aspects. A few of these choices and the influence they have on the emerging set of validities, have been discussed in some detail.

Our approach, in which a modal logic ($\Diamond_u, \Diamond_d$) is placed 'on top of' another modal logic (S5), fits in the recent trend of 'combining logics' (see [3]). Combinations of logics are often (almost) orthogonal, in the sense that there is limited interaction between the two logics (this is the case in for instance [4]). In our logic, however, the two logics are very strongly tied; in fact, our logic is based on a single modal model ($\mathcal{S}_f$), in which the accessibility relation of increasing information is completely determined by the states, which are themselves S5-models.

It was shown that the logic of only knowing of Halpern and Moses [5] can be embedded in our logic, which means we can use our proof system to derive validities of their logic. The preference ordering of Halpern and Moses is the modal accessibility relation in our logic (a similar idea is used in [2]). There are strong connections between our system and the one of Levesque [7], a logic which also embeds the logic of only knowing.

For one particular kind of super model we have defined a proof system which is sound and complete.

One of our future interests concerns the logic MTEL [4], which can be seen as a temporalization of Halpern and Moses' logic. Entailment of both default logic and autoepistemic logic can be embedded in MTEL. A proof system for MTEL would thus give a system in which both derivations for default logic and for autoepistemic logic could be carried out. We hope to be able to apply similar methods (viewing the preference relation of MTEL as a modal accessibility relation) to arrive at such a proof system.

References

[1] J. van Benthem & A. ter Meulen, *Handbook of Logic and Linguistics*, North-Holland, Amsterdam, 1997.

[2] C. Boutilier, "Conditional logics of normality: A modal approach", *Artificial Intelligence* **68**:87–154, 1994.

[3] M. de Rijke and P. Blackburn (eds.), "Combining Logics", special issue of the *Notre Dame Journal of Formal Logic*, **37**(2), 1996.

[4] J. Engelfriet, "Minimal temporal epistemic logic", pp. 233–259 in [3].

[5] J.Y. Halpern and Y. Moses, "Towards a theory of knowledge and ignorance: Preliminary report", in: K.R. Apt (ed.), *Logics and Models of Concurrent Systems*, NATO ASI Series **F13**, Springer-Verlag, 1985, pp. 459–476.

[6] S. Kraus, D. Lehmann and M. Magidor, "Nonmonotonic reasoning, preferential models and cumulative logics", *Artificial Intelligence* **44**:167–207, 1990.

[7] H.J. Levesque, "All I know: A study in autoepistemic logic", *Artificial Intelligence* **42**:263–309, 1990.

[8] F. Veltman, "Defaults in update semantics", *Journal of Philosophical Logic* **25**:221–261, 1996.