A modal logic of information change*

Joeri Engelfriet† Yde Venema‡

July 20, 1998

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, or knowledge, they have about the world may change as a result of performing observations, communication with other agents, or through non-monotonic reasoning (where an agent makes certain plausible assumptions). The most basic kinds of change are increase (information is added) and decrease of knowledge (information is deleted), and in a sense all changes in information can be seen as combinations of these basic kinds: first the old information is thrown away and then the new information is added. In this paper, we will study changing knowledge, and we use modal logic as a vehicle. Our perspective is semantical: we study models in which the worlds represent information states, and in which 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*).

*Technical Report IR-438, Faculty of Mathematics and Computer Science, Vrije Universiteit Amsterdam, 1997.

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

‡ibid., `yde@cs.vu.nl`. The research of this author has been made possible by a fellowship of the Royal Netherlands Academy of Arts and Sciences.

Many further choices have to be made in formalizing these intuitions, and one of them is the nature of knowledge (or information) and 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 (so $K\varphi$ means that the agent knows – or believes – φ). This again suggests using modal logic, and we will use S5 for this purpose. Then, as our information states we will 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. As we have a modal logic (with $\Diamond_u$ and $\Diamond_d$) defined "on top of" S5, our approach falls into the category of *combining logics* (see for instance [12]).

In our "supermodels" we will group together such information states: we want to study modal models in which the states themselves are (modal) S5-models. 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 to this field. For a survey, the reader is referred to [1] or [16]; in [7], a general program of 'dynamifying logics' is sketched. Our system is closely related to the update semantics of [14], which too is concerned with information states consisting of valuations for classical propositional logic. The aim of update semantics is to study the effect of incoming information (in a discourse) on the information state of an agent. The specific piece of information that induces the change is explicitly modeled in this framework, in contrast to our proposal, in which we only look at increases (decreases) as such, abstracting from the *kind* of information leading to the change.

Change of information over time, and applications to nonmonotonic reasoning, are studied in [4]. A conservative temporal epistemic model in the logic studied there (MTEL), is in a sense a path through our supermodel, following the accessibility relation of increasing information. Alternatively, MTEL can be seen as a temporalization of the logic of "only knowing" of Halpern and Moses ([8]). This logic aims to answer questions like: "if I *only* know p, what else do I then know, and what do I not know?" For instance, if you only know p, you do not know q. A model in which the agent *only knows* φ is an S5-model satisfying $K\varphi$, with as little knowledge as possible. In our "supermodels" this would be a state in which $K\varphi$ is true, but any decrease in information leads to a state where it is false, and we can express this in our language: $K\varphi \wedge \Box_d \neg K\varphi$. We shall see that a faithful translation of their consequence relation in our logic exists.

The problem with Halpern and Moses' logic is that their consequence relation turns out to be 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 has been taken by Levesque ([10]), who introduced a modal operator $\mathcal{O}$, where $\mathcal{O}\varphi$ intuitively means that the agent only knows φ. For our purposes, his logic has at least two disadvantages. In the first place, for the axiomatization an extra operator N is needed, where $N\alpha$ means that "α at most is believed to be false". The intuitive meaning of this operator in the nested case (when it applies to a formula which itself contains N or K-operators) is difficult to grasp. In the second place, it talks only about only knowing, but does not say anything about increase or decrease in information in general.

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 [13], [9]). This preference relation is in fact the modal accessibility relation (of decreasing knowledge) in our supermodels. 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 [3]. In that paper, Boutilier gives axiomatizations which are sound with respect to certain classes of preferential models (a preferential model consists of a set of worlds, a preference relation and a mapping that assigns a propositional valuation to each world). 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 with Boutilier: 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.

In Section 2, we will formally introduce our language and the supermodel semantics we have in mind for it. Section 3 contains some technical results connected with the effect of some of our choices on the logic of a supermodel. In Section 4 we introduce a Hilbert-style proof calculus; we prove it to be sound for one of the supermodel semantics. Section 5 is devoted to proving completeness of this axiomatization. Conclusions and suggestions for further research are given in Section 6.

2 The supermodels

As we already mentioned in the introduction, in our formalization of information change we take a layered approach. On the base level, we are dealing with a propositional logic, of 'facts' if one likes, 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 top level language is then obtained by adding two more operators $\Diamond_d$ and $\Diamond_u$ to this epistemic language; we have already seen that $\Diamond_d\varphi$ ($\Diamond_u\varphi$) is to be read as "it is possible to decrease (increase) your knowledge to a state where φ holds'.

What we still have to discuss is the *exact* structure of our language. For instance, as yet we do not want to consider the agent having information about its possible updates; that is to say, we will not consider formulas in which one of the diamonds is in the scope of the knowledge operator. Also, since at the top level we are only interested in the knowledge the agent may or may not have, and not in the "real" world, there is no need to allow formulas like $\Diamond_u(p \wedge Kq)$, as p refers to an actual world. Thus, 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. A more formal definition follows now.

Definition 2.1 (syntax) *We fix a set* $\mathbb{V}$ *of propositional variables* $p_0, p_1, \ldots, q, r, \ldots$ *.* $\mathcal{L}_0$ *is the* ***base language*** *of classical propositional logic; formally,* $\mathcal{L}_0$ *denotes the set of those formulas that we can build up from* $\mathbb{V}$ *using* $\perp$, $\wedge$ *and* $\neg$*. We will employ the standard defined connectives, like* $\vee$ *and* $\rightarrow$*. The (meta-)variables* $\alpha, \beta, \gamma, \ldots$ *will be used to range over formulas in this base language.*

At the intermediate level, $\mathcal{L}_1$ *denotes the* ***epistemic language****. Formally,* $\mathcal{L}_1$ *is defined as the smallest set of formulas containing* $\mathcal{L}_0$ *and such that* $K\mu$, $\neg\mu$ *and* $\mu \wedge \nu$ *are in* $\mathcal{L}_1$ *whenever* μ *and* ν *are.*[1] *The connective* K *is called the* ***knowledge operator****. For its dual we use the abbreviation* M*; that is to say,* $M\mu$ *denotes* $\neg K \neg \mu$*. We use* $\mu, \nu, \rho, \ldots$ *to range over elements of this language.*

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* $\Diamond_u$*. We let* $\mathcal{L}_{2d}$ *denote the set of* ***downdate formulas****; that is,* $\mathcal{L}_{2d}$ *consists of all* $\mathcal{L}_2$*-formulas in which the operator* $\Diamond_u$ *does not occur. To denote* $\mathcal{L}_2$*-formulas we use* $\varphi, \psi, \chi, \ldots$

The duals of these diamonds are denoted as $\Box_d$ *and* $\Box_u$*, respectively; we use the 'only' operator* O *as the following abbreviation:*

$$\mathsf{O}\varphi \;=\; \varphi \wedge \Box_d \neg\varphi.$$

On some occasions we will have reason to study fragments of a language $\mathcal{L}_i$ *in which only a restricted set, say,* $\mathbb{X}$*, of propositional variables may occur; to denote this set of formulas we will use obvious notation like* $\mathcal{L}_i(\mathbb{X})$*.*

Even without having provided the precise definition of the semantics of these languages, we can already informally discuss the meaning of its operators. For instance, the informal reading of $K\alpha$ is that the agent knows α, and hence, of $M\alpha$, that the agent considers α to be possible. The meaning of $\Diamond_d\varphi$ will

[1] There is some ambiguity in this definition; for instance, the formula $p \wedge q$ can be read as a conjunction in $\mathcal{L}_0$ or in $\mathcal{L}_1$. In order to disambiguate the language, one could 'split up' each boolean connectives into three connectives, one for each level. We do not worry about this ambiguity since it does not lead to any semantic confusion.

be that it is possible, by decreasing the agent's knowledge, to reach a state where φ holds; and likewise for $\Diamond_u\varphi$. 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 now 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. Conversely, if the agent only knows α, and it loses knowledge, it can no longer know α. Thus we see that

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

We will show later that we can embed the approach of Halpern & Moses mentioned in the introduction, in our logic.

It may seem natural to study the top level language in which both diamonds are available, rather than the language $\mathcal{L}_{2d}$ in which only the downdate diamond can be used. However, for the definition of the 'only knowing' operator the update diamond is not needed, and in the next section we will see that the downdate fragment of $\mathcal{L}_2$ has some interesting properties.

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 valuation w, denoted by $\mathfrak{m}, w \Vdash \mu$, is recursively defined as follows:*

$\mathfrak{m}, w \Vdash \alpha$	*if*	$w \models \alpha$, *for propositional* α
$\mathfrak{m}, w \Vdash \neg\mu$	*if*	$\mathfrak{m}, w \nVdash \mu$
$\mathfrak{m}, w \Vdash \mu \wedge \nu$	*if*	$\mathfrak{m}, w \models \mu$ *and* $\mathfrak{m}, w \models \nu$
$\mathfrak{m}, w \Vdash K\mu$	*if*	$\mathfrak{m}, v \models \mu$ *for all* $v \in \mathfrak{m}$.

Truth *of an epistemic formula in a model is denoted and defined as follows: $\mathfrak{m} \Vdash \mu$ if $\mathfrak{m}, v \Vdash \mu$ for all $v \in \mathfrak{m}$.*

In other words, the kind of models for the epistemic language that we are considering are the simplest S5-models.

For a propositional formula α we thus have two notions of truth in a model: $\mathfrak{m} \models \alpha$ and $\mathfrak{m} \Vdash \alpha$. These notions are equivalent, but in the sequel we will make use of the notational distinction, writing '$\mathfrak{m} \models \alpha$' if we see $\mathfrak{m}$ as a set of valuations, and '$\mathfrak{m} \Vdash \alpha$' if we see $\mathfrak{m}$ as an information state in some supermodel.

Note that the notion $\mathfrak{m} \Vdash \alpha$ is *not* two-valued: if m contains both valuations in which α is true, and ones in which it is false, then neither $\mathfrak{m} \Vdash \alpha$, nor $\mathfrak{m} \Vdash \neg\alpha$. However, it is easy to see that for a subjective formula μ it *does* hold that either $\mathfrak{m} \Vdash \mu$ or $\mathfrak{m} \Vdash \neg\mu$.

We are now ready to introduce the main characters of this story, namely the structures that we use to model the notion of information change. 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. The intuition behind this information ordering is that one model $\mathfrak{m}$ is smaller than a model $\mathfrak{n}$ iff $\mathfrak{n}$ contains more information or knowledge than $\mathfrak{m}$. The underlying idea is that $\mathfrak{n}$ contains more information than $\mathfrak{m}$ if $\mathfrak{n}$ consists of less worlds than $\mathfrak{m}$.

Definition 2.3 *Let $\sqsubset$ denote the following* ***information ordering*** *on models:*

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

Here $\subset$ denotes strict set inclusion.

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 are now ready to give the definition of the supermodel; in fact, we will define *three* alternative options. These three different versions of the supermodel are very similar; the only difference is the number, or more precisely, the *kind* of models that we allow as information states.

Definition 2.4 *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 $\mathfrak{m} = Mod(\Delta_\mathfrak{m})$; note that $\Delta_{Mod(\Gamma)} = \{\alpha \mid \Gamma \models \alpha\}$. 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*[2] *(as an example we take $\mathcal{S}$, the definition for $\mathcal{S}^+$ and $\mathcal{S}_f$ is analogous):*

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

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 $\mathcal{S}_f$).*

[2] Recall that only subjective $\mathcal{L}_1$-formulas belong to $\mathcal{L}_2$. Given our clause for negation, this is needed to keep the definition unambiguous.

Remark 2.5 In the remainder of this section we will discuss these three different options ($\mathcal{S}^+$, $\mathcal{S}$ and $\mathcal{S}_f$); for instance, we will see how this choice will affect the set of valid $\mathcal{L}_2$-sentences. The reader should note however, that in fact there are *many* parameters that we needed to fix when setting up this framework. For instance, we could have opted for an intuitionistic or partial base logic; for a language with only a finite number of basic facts; and/or an information ordering with a more complex definition. Any of such choices may influence the resulting 'logic of information change'.

Obviously, our terminology (of closed and clopen sets) has a topological origin; for the interested reader we mention that we are in fact considering the topology over the set of valuations which is induced by the *Stone* embedding of the free Boolean algebra over $\mathbb{V}$, into the Boolean power set of $\mathcal{V}$. Although this topological connection is not one of our main concerns in this paper, the following facts, which all have rather straightforward proofs in this topological context, will be put to good use later on.

Proposition 2.6

1. *The collection of clopen sets is closed under taking finite intersections and unions.*

2. *The collection of closed sets contains all singletons, and is closed under taking arbitrary intersections and finite unions.*

PROOF. The proofs of these facts are all rather straightforward, so we restrict ourselves to proving that the collection of closed models is closed under finite unions. Obviously it suffices to show that if $\mathfrak{m}$ and $\mathfrak{n}$ are closed, then so is their union. But $\mathfrak{m} \cup \mathfrak{n} = Mod(\{\alpha \vee \beta \mid \alpha \in \Delta_{\mathfrak{m}}, \beta \in \Delta_{\mathfrak{n}}\})$, as a direct proof reveals.
QED

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 (we omit its rather direct proof).

Proposition 2.7 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* $\vdash \delta_{\mathfrak{n}} \to \delta_{\mathfrak{m}}$ *and* $\not\vdash \delta_{\mathfrak{m}} \to \delta_{\mathfrak{n}}$.

It follows immediately from Proposition 2.7 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*.

Corollary 2.8 *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 $\mathfrak{m} = Mod(\alpha)$.*

This observation constitutes a difference between $\mathcal{S}$ and $\mathcal{S}_f$ that we will make good use of in the completeness proof later on, but its importance seems to be only of a technical nature. 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 $\mathcal{S}$.

Perhaps surprisingly, the 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. This matter will be dealt with in the next section.

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; recall that an ordering $<$ is *dense* if $\forall xy\,(x < y \to \exists z\,(x < z < y))$, and *discrete* if $\forall xy\,(x < y \to \exists z\,(x < z \leq y \wedge \neg\exists u\,(x < u < z)))$.

Proposition 2.9 *On the three supermodels the ordering behaves as follows:*

1. *$\sqsubset$ and $\sqsupset$ are discrete on $\mathcal{S}^+$*
2. *$\sqsubset$ is dense on $\mathcal{S}_f$ (and, therefore, so is $\sqsupset$)*
3. *$\sqsupset$ is discrete on $\mathcal{S}$, but $\sqsubset$ is not.*

In fact, there is far more to say about these orderings. For instance, we have already mentioned the connection between M_f and the free Boolean algebra over countably many generators; $\sqsubset$ is in fact the naturally induced ordering on this algebra (with the bottom element taken out). We do not pursue such investigations since they are not of interest for the main line of the paper.

PROOF. PART 1. The discreteness of $\sqsupset$ and $\sqsubset$ follows immediately from the fact that M^+ is the full power set (except for the empty set) of the set $\mathcal{V}$. Adding or taking away singleton valuations to a model provides immediate successors (predecessors, respectively).

PART 2. Consider two clopen models $\mathfrak{m}$ and $\mathfrak{n}$ such that $\mathfrak{m} \sqsupset \mathfrak{n}$. It follows from Proposition 2.7.4 that $\delta_{\mathfrak{m}} \to \delta_{\mathfrak{n}}$ and $\not\vdash \delta_{\mathfrak{n}} \to \delta_{\mathfrak{m}}$. Now let q be some propositional variable that does not occur in $\delta_{\mathfrak{m}}$ or $\delta_{\mathfrak{n}}$, and consider the formula $\alpha := \delta_{\mathfrak{m}} \vee (q \wedge \delta_{\mathfrak{n}})$. It is easy to prove that $\vdash \delta_{\mathfrak{m}} \to \alpha$ and $\vdash \alpha \to \delta_{\mathfrak{n}}$, while $\not\vdash \alpha \to \delta_{\mathfrak{m}}$ and $\not\vdash \delta_{\mathfrak{n}} \to \alpha$. But then $\mathfrak{m} \sqsupset Mod(\alpha) \sqsupset \mathfrak{n}$.

PART 3. The proof that $\sqsupset$ is discrete on $\mathcal{S}$, is similar to the proof given in part 1 — we need the fact that singletons are closed and that the collection of closed sets is closed under taking finite unions.

In order to show that $\sqsubset$ is not discrete, consider the model $\mathfrak{m} = Mod(\alpha)$ for some satisfiable propositional formula α. In order to arrive at a contradiction, assume that $\mathfrak{m}$ has an immediate successor $\mathfrak{n}$ (i.e., a closed model $\mathfrak{n}$ such that $\mathfrak{m} \sqsubset \mathfrak{n}$ while there is no model properly between $\mathfrak{m}$ and $\mathfrak{n}$). It follows from part 2 that $\mathfrak{n}$ cannot be clopen; hence, there is no finite part Δ of $\Delta_{\mathfrak{n}}$ such that $\Delta \vdash \Delta_{\mathfrak{n}}$.

Since $\mathfrak{n} \subset \mathfrak{m}$ and $\mathfrak{m} \models \alpha$, we have $\mathfrak{n} \models \alpha$ and hence, $\mathcal{S}, \mathfrak{n} \Vdash K\alpha$. By Proposition 2.7.1, we obtain $\Delta_{\mathfrak{n}} \vdash \alpha$. By compactness, there is some finite set $\Delta_0 \subset \Delta_{\mathfrak{n}}$ such that $\Delta_0 \vdash \alpha$. Now consider any finite subset Γ of $\Delta_{\mathfrak{n}}$ such that $\Gamma \vdash \Delta_0$, while $\Gamma \not\vdash \Delta_{\mathfrak{n}}$ and $\Delta_0 \not\vdash \Gamma$; such a set must exist by the finiteness of Δ_0 and the non-'finiteness' of $\Delta_{\mathfrak{n}}$ mentioned earlier on. We leave it to the reader to verify that this gives $\mathfrak{n} \sqsupset Mod(\Gamma) \sqsupset \mathfrak{m}$. QED

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 $\mathcal{S}_f \Vdash \mu \vee \Diamond_d\mu \vee \Diamond_u\mu \vee \Diamond_d\Diamond_u\mu$. The proof system we will give later on, axiomatizes validity in $\mathcal{S}_f$, which means that we have a proof system for non-validity in S5: $\not\models_{S5} \mu$ iff $\vdash \neg\mu \vee \Diamond_d\neg\mu \vee \Diamond_u\neg\mu \vee \Diamond_d\Diamond_u\neg\mu$ (where $\vdash$ is provability of the system we will give in Section 4). The restriction to subjective formulas is not severe, since a (possibly not 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, and we start with the logic of only knowing of Halpern & Moses ([8]). 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 2.10 *Let μ be an honest formula, and ν any S5-formula. Then*

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

PROOF. Suppose $\mu \mathrel{|\!\sim} \nu$. Let $\mathfrak{m} \in M_f$, and suppose that $\mathcal{S}_f, \mathfrak{m} \Vdash \mathsf{O}K\mu$. Then $\mathfrak{m}$ is a maximal model of μ: it satisfies $K\mu$, and if there were a larger model of μ, then it can be shown that there is also a larger *clopen* model of μ (we leave this to the reader), which would contradict the $\Box_d\neg K\mu$ part of $\mathsf{O}K\mu$. But then we have $\mathfrak{m} \Vdash \nu$, which implies $\mathcal{S}_f, \mathfrak{m} \Vdash K\nu$. We have proved that $\mathcal{S}_f \Vdash \mathsf{O}K\mu \to K\nu$.

For the other direction, suppose $\mathcal{S}_f \Vdash \mathsf{O}K\mu \to K\nu$. Let $\mathfrak{m}$ be the unique maximal model of μ. Then $\mathcal{S}_f, \mathfrak{m} \Vdash \mathsf{O}K\mu$: it is obvious that $\mathfrak{m} \Vdash K\mu$, and if there is a $\mathfrak{n} \in M_f$ with $\mathfrak{n} \sqsubset \mathfrak{m}$ and $\mathcal{S}_f, \mathfrak{n} \Vdash K\mu$, then $\mathfrak{n} \supset \mathfrak{m}$, contradicting the assumption that $\mathfrak{m}$ is a maximal model of μ. But then $\mathcal{S}_f, \mathfrak{m} \Vdash K\nu$, so $\mathfrak{m}$ satisfies ν, which gives us the desired conclusion, $\mu \mid\!\sim \nu$. QED

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

Proposition 2.11 *Let $\mu \in \mathcal{L}_1$, then the following statements are equivalent.*

1. *μ is honest.*
2. *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 $\mathcal{S}_f \Vdash \mathsf{O}K\mu \to \neg\varphi$.*
3. *There exists an $\alpha \in \mathcal{L}_0$ such that $\mathcal{S}_f \Vdash \mathsf{O}K\mu \leftrightarrow \mathsf{O}K\alpha$.*

Proof. First, suppose μ is honest. Then, there is a unique maximal model $\mathfrak{m}$ of μ, so we have that $\mathcal{S}_f, \mathfrak{n} \Vdash \mathsf{O}K\mu$ iff $\mathfrak{n} = \mathfrak{m}$. But this means for any φ, if $\mathcal{S}_f, \mathfrak{m} \Vdash \varphi$, then $\mathcal{S}_f \Vdash \mathsf{O}K\mu \to \varphi$, and if $\mathcal{S}_f, \mathfrak{m} \Vdash \neg\varphi$, then $\mathcal{S}_f \Vdash \mathsf{O}K\mu \to \neg\varphi$. For the other direction, it is easy to show that for any two different states $\mathfrak{m}, \mathfrak{n} \in M_f$, there is an epistemic formula φ true in $\mathfrak{m}$ but not in $\mathfrak{n}$. If there are two different states $\mathfrak{m}, \mathfrak{n}$ in M_f for which $\mathcal{S}_f, \mathfrak{m} \Vdash \mathsf{O}K\mu$ and $\mathcal{S}_f, \mathfrak{n} \Vdash \mathsf{O}K\mu$, then for an epistemic formula φ differentiating between $\mathfrak{m}$ and $\mathfrak{n}$ as described above, neither $\mathcal{S}_f \Vdash \mathsf{O}K\mu \to \varphi$, nor $\mathcal{S}_f \Vdash \mathsf{O}K\mu \to \neg\varphi$.

The third statement expresses the fact that there is exactly one state (named '$\mathsf{O}K\alpha$') in which $\mathsf{O}K\mu$ holds. Its equivalence to the first statement can be proved straightforwardly. QED

Proposition 2.10 means we can use the proof system for $\mathcal{S}_f$ of Section 4 to prove all entailments in Halpern & Moses' logic. But this is not a new accomplishment. In [10], 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 briefly review Levesque's logic.

Levesque's logic has a Kripke semantics. The models he considers are closed sets of valuations, so they are just the elements of our M. A modal operator B has the same semantics as our K operator, i.e. $W, w \Vdash B\alpha$ iff $\forall w' \in W : W, w' \Vdash \alpha$. The difference with our approach is that the current world (w) need not be an element of W, and that W may be empty. This means that the B operator has the K45 (or 'weak S5') axioms, instead of the S5 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"). The formal semantics of the N operator is given by the clause

$$\begin{array}{lll} W, w \Vdash N\alpha & \text{iff} & \forall w' : (W, w' \nVdash \alpha \Rightarrow w' \in W) \\ & \text{iff} & \forall w' \notin W : W, w' \Vdash \alpha \end{array}$$

Finally, $\mathcal{O}\alpha$ is defined as $B\alpha \wedge N\neg\alpha$.

The B operator corresponds to our K operator, but what about the N operator, can we express it in our logic? First of all, we will see that nesting of modal operators inside the scope of N can be avoided. Every formula in Levesque's logic is equivalent to a conjunction of disjunctions of the form $\alpha \vee \beta$, where α is propositional and β is a disjunction of formulas of the form $N\varphi, \neg N\varphi, B\varphi$ and $\neg B\varphi$. The N operator is a normal operator, so $N(\varphi \wedge \psi) \equiv N\varphi \wedge N\psi$. But it is also straightforward to prove that for a disjunction $\alpha \vee \beta$ as described above, $N(\alpha \vee \beta) \equiv N\alpha \vee N\beta$ (this is also the case for the S5 operator K). Furthermore, the N distributes over the disjunction of β, and can then be eliminated, since $NN\varphi \equiv N\varphi$, $N\neg N\varphi \equiv \neg N\varphi$, $NB\varphi \equiv B\varphi$, and $N\neg B\varphi \equiv \neg B\varphi$. This means that the only remaining question is whether we can express a formula $N\alpha$, where α is propositional.

Suppose we have a closed state W, then it is easy to see that $W, w \Vdash N\alpha$ is independent of w. Thus it holds that

$$W \Vdash N\alpha \text{ iff } Mod(\neg\alpha) \subseteq W \text{ iff } \mathcal{S}, W \Vdash \mathsf{O}K\neg\alpha \vee \Diamond_u(\mathsf{O}K\neg\alpha)$$

So we can indeed express the N operator. Finally, we can see that

$$\mathcal{O}\alpha \equiv \mathsf{O}K\alpha.$$

For, $\mathcal{O}\alpha \equiv B\alpha \wedge N\neg\alpha \equiv K\alpha \wedge (\mathsf{O}K\neg\neg\alpha \vee \Diamond_u(\mathsf{O}K\neg\neg\alpha) \equiv \mathsf{O}K\alpha \vee (\Diamond_u(K\alpha \wedge \Box_d\neg K\alpha) \wedge K\alpha) \equiv \mathsf{O}K\alpha$.

3 Equivalence of downdate formulas

This section is devoted to the proof of the following proposition stating that the difference between the models $\mathcal{S}$ and $\mathcal{S}_f$ is not reflected in the logic if we restrict ourselves to downdate-formulas.

Proposition 3.1 *Let φ be a formula in $\mathcal{L}_{2d}$. Then*

1. *For any clopen state $\mathfrak{n}$:*

$$\mathcal{S}, \mathfrak{n} \Vdash \varphi \textit{ iff } \mathcal{S}_f, \mathfrak{n} \Vdash \varphi.$$

2.

$$\mathcal{S} \Vdash \varphi \textit{ iff } \mathcal{S}_f \Vdash \varphi.$$

The key step of the proof of Proposition 3.1 concerns a lemma stating that a certain rather naturally defined relation on closed models is some sort of bisimulation between $\mathcal{S}$ and $\mathcal{S}_f$. This lemma seems to be of independent interest; it, and some related results, are also used later on, and therefore these results are stated and proved separately. First we introduce some useful concepts, including the definition of the relation mentioned above.

Definition 3.2 *Let $\mathbb{X} \subseteq \mathbb{V}$ be a set of propositional variables. We say that two valuations w and v* ***agree on*** *$\mathbb{X}$, notation: $w \sim_{\mathbb{X}} v$, if for all $p \in \mathbb{X}$, $w(p) = v(p)$. The equivalence class of w under this relation is denoted by $[w]_{\mathbb{X}}$.*

A set $\mathfrak{q}$ of valuations is called $\mathbb{X}$***-empty*** *on an $\sim_{\mathbb{X}}$-equivalence class W if $\mathfrak{q} \cap W = \varnothing$,* $\mathbb{X}$***-full*** *on W if $W \subseteq \mathfrak{q}$. Two sets of valuations $\mathfrak{q}$ and $\mathfrak{y}$ are* ***equally*** $\mathbb{X}$***-full on*** *W if they are either both $\mathbb{X}$-empty or both not $\mathbb{X}$-empty, and also either both $\mathbb{X}$-full or both not $\mathbb{X}$-full. Finally, we say that two models $\mathfrak{m}$ and $\mathfrak{n}$ are* $\mathbb{X}$***-alike****, notation: $\mathfrak{m} \equiv_{\mathbb{X}} \mathfrak{n}$, if they are equally $\mathbb{X}$-full for each $\sim_{\mathbb{X}}$-equivalence class W.*

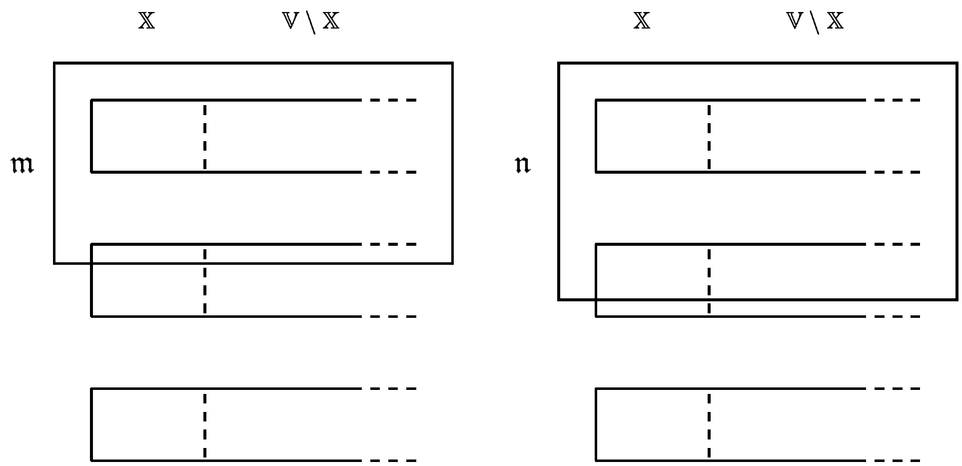

Figure 1: Two $\mathbb{X}$-alike models

Figure 1 gives a graphic representation of two $\mathbb{X}$-alike models. The half-open rectangles represent equivalence classes. Each of the valuations in such a class has the same (initial) $\mathbb{X}$-part, but differs on the $\mathbb{V} \setminus \mathbb{X}$-part. Both $\mathfrak{m}$ and $\mathfrak{n}$ are $\mathbb{X}$-full on the first (top) class, and they are both $\mathbb{X}$-empty on the third class. On the second class, they are both neither $\mathbb{X}$-full nor $\mathbb{X}$-empty. Observe that $\mathfrak{m}$ and $\mathfrak{n}$ do not (necessarily) contain the same valuations from the second class (as long as they contain some, but not all of them).

For *finite* $\mathbb{F}$, there are only finitely many equivalence classes under $\sim_{\mathbb{F}}$, each of which corresponds to a formula in $\mathcal{L}_0(\mathbb{F})$.

Lemma 3.3 *A model $\mathfrak{m}$ is clopen iff for some finite set $\mathbb{F}$ of variables, $\mathfrak{m}$ is a finite union of $\sim_{\mathbb{F}}$-equivalence classes.*

Proof. Assume that $\mathbb{F} \subset \mathbb{V}$ is finite. With each valuation w on $\mathbb{F}$ we will associate an $\mathcal{L}_0(\mathbb{F})$-formula θ_w as follows:

$$\theta_w := \bigwedge_{p \in \mathbb{F}, w(p)=1} p \wedge \bigwedge_{p \in \mathbb{F}, w(p)=0} \neg p. \tag{1}$$

We leave it to the reader to verify that for each equivalence class $[w]_{\mathbb{F}}$,

$$[w]_{\mathbb{F}} = Mod(\theta_w),$$

showing that indeed, $[w]_{\mathbb{F}}$ is clopen. This shows the direction from right to left. The other direction is left to the reader. QED

As we mentioned before, for any finite $\mathbb{F}$, $\equiv_{\mathbb{F}}$ is some sort of bisimulation between $\mathcal{S}$ and $\mathcal{S}_f$. First we show that this holds on $\mathcal{S}$.

Lemma 3.4 *For every finite set $\mathbb{F}$, the relation $\equiv_{\mathbb{F}}$ is a $\sqsupset$-bisimulation on $\mathcal{S}$. That is, for all closed information states $\mathfrak{m}$, $\mathfrak{n}$ and $\mathfrak{m}'$:*

$$\mathfrak{m}' \equiv_{\mathbb{F}} \mathfrak{m} \sqsupset \mathfrak{n} \textit{ only if for some closed state } \mathfrak{n}'\textit{: } \mathfrak{m}' \sqsupset \mathfrak{n}' \equiv_{\mathbb{F}} \mathfrak{n}.$$

PROOF. Let $\mathfrak{m}$, $\mathfrak{m}'$ and $\mathfrak{n}$ be closed information states such that $\mathfrak{m}' \equiv_{\mathbb{F}} \mathfrak{m} \sqsupset \mathfrak{n}$. We have to define a closed model $\mathfrak{n}'$ such that $\mathfrak{m}' \sqsupset \mathfrak{n}' \equiv_{\mathbb{F}} \mathfrak{n}$. We will treat each $\sim_{\mathbb{F}}$-equivalence class separately; that is, for each equivalence class W, we will define a set $\mathfrak{n}'_W$ which is to be the 'W-part of $\mathfrak{n}'$'; finally, $\mathfrak{n}'$ is defined as the union of the (finitely many) $\mathfrak{n}'_W$. Hence, let W be some equivalence class under $\sim_{\mathbb{F}}$. Distinguish the following cases:

1. If $\mathfrak{n}$ is $\mathbb{F}$-empty on W, define $\mathfrak{n}'_W := \varnothing$.
2. If $\mathfrak{n}$ is $\mathbb{F}$-full on W, define $\mathfrak{n}'_W := W$.
3. Otherwise, make a further case distinction:
   (a) If $W \cap \mathfrak{m} = W \cap \mathfrak{n}$ (that is, if $\mathfrak{n}$ is not bigger than $\mathfrak{m}$ on W), put $\mathfrak{n}'_W := \mathfrak{m}' \cap W$.
   (b) Finally, the case where $W \cap \mathfrak{m} \subset W \cap \mathfrak{n}$ (that is, where $\mathfrak{n}$ *is* bigger than $\mathfrak{m}$ on W), is the most interesting one. It follows from our case and subcase assumptions that $\mathfrak{m}$ is not full on W; but then by $\mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{m}'$, the same applies to $\mathfrak{m}'$. Hence, it is possible to find a valuation w' in $W \setminus \mathfrak{m}'$. Define $\mathfrak{n}'_W := (W \cap \mathfrak{m}') \cup \{w'\}$.

Finally, define

$$\mathfrak{n}' := \bigcup_{w \in \mathcal{V}} \mathfrak{n}'_{[w]_{\mathbb{F}}}$$

This is a *finite* union since there are only finitely many equivalence classes.

We will prove the following five claims concerning this $\mathfrak{n}'$.

Claim 1 $\mathfrak{n}'$ is closed.

PROOF OF CLAIM We already saw in Lemma 3.3, that each $\sim_{\mathbb{F}}$-equivalence class is clopen. But then an inspection of the definition of $\mathfrak{n}'_W$, taken together with the fact that the collection of closed sets of valuations is closed under taking intersections, reveals that for each equivalence class W, $\mathfrak{n}'_W$ is closed. Hence, $\mathfrak{n}'$ itself, being a finite union of closed sets, must be closed. ◀

Claim 2 For each equivalence class W, $\mathfrak{n}' \cap W = \mathfrak{n}'_W$.

PROOF OF CLAIM It is obvious from the definition that $\mathfrak{n}'_W \subseteq W$ for each W. From this and the definition of $\mathfrak{n}'$, the claim is immediate. ◀

In the remainder of the proof, claim 2 will be used without notice.

Claim 3 For each equivalence class W, $(\mathfrak{m} \cap W \;\subset\; \mathfrak{n} \cap W)$ if and only if $(\mathfrak{m}' \cap W \;\subset\; \mathfrak{n}' \cap W)$.

PROOF OF CLAIM Let W be an $\sim_{\mathbb{F}}$-equivalence class. The equivalence stated in the claim is proved according to the case distinction used in the definition of $\mathfrak{n}'_W$. Note that by the assumption of the Lemma we have that $W \subseteq \mathfrak{m}$ iff $W \subseteq \mathfrak{m}'$ and that $W \cap \mathfrak{m} = \varnothing$ iff $W \cap \mathfrak{m}' = \varnothing$.

In case 1, we cannot have $\mathfrak{m} \cap W \subset \mathfrak{n} \cap W$. But since $\mathfrak{n}'_W = \varnothing$, neither can it be the case that $\mathfrak{m}' \cap W \subset \mathfrak{n}' \cap W$.

Now suppose that we are in case 2. Since $\mathfrak{n}'_W$ is defined as W, we have $\mathfrak{m}' \cap W \subset \mathfrak{n}'_W$ iff $\mathfrak{m}' \cap W \subset W$. Since $\mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{m}'$, we have $\mathfrak{m}' \cap W \subset W$ iff $\mathfrak{m} \cap W \subset W$. Finally, our case assumption $(\mathfrak{n} \cap W = W)$ tells us that $\mathfrak{m} \cap W \subset W$ iff $\mathfrak{m} \cap W \subset \mathfrak{n} \cap W$. Taking all this together, we have that indeed $(\mathfrak{m}' \cap W \;\subset\; \mathfrak{n}' \cap W)$ if and only if $(\mathfrak{m} \cap W \;\subset\; \mathfrak{n} \cap W)$.

Finally, the definition in case 3 is *tailored* towards making this claim hold. ◀

Claim 4 $\mathfrak{m}' \sqsupset \mathfrak{n}'$.

PROOF OF CLAIM It is is rather straightforward to check, by an inspection of the definition, that $\mathfrak{m}' \subseteq \mathfrak{n}'$. In order to check that the inclusion is *strict*, observe that $\mathfrak{m} \subset \mathfrak{n}$ iff $\mathfrak{m} \cap W \subset \mathfrak{n} \cap W$ for some equivalence class W, and likewise for $\mathfrak{m}'$ and $\mathfrak{n}'$. But then the Claim is immediate by Claim 3. ◀

Claim 5 $\mathfrak{n} \equiv_{\mathbb{F}} \mathfrak{n}'$.

PROOF OF CLAIM We will prove that for each equivalence class W, $\mathfrak{n}$ is full (empty) on W iff $\mathfrak{n}'$ is full (empty, respectively) on W. For the cases 1, 2 and 3a, this is immediate by definition and by Claim 2.

The interesting situation occurs in case 3b; in this case $\mathfrak{n}$ is neither full nor empty on W, so we have to prove that the same applies to $\mathfrak{n}'$.

It is immediate by the definition that $\mathfrak{n}'$ is not empty on W. Now, in order to derive a contradiction, assume that $\mathfrak{n}'$ is full on W. From the definition of $\mathfrak{n}'$ in case 3b, we may infer that $W \cap \mathfrak{m}' = W \setminus \{w'\}$ for some valuation w'. But this contradicts the fact that $\mathfrak{m}'$, and hence, $W \cap \mathfrak{m}'$ is closed. ◀

It follows from Claim 5 that $\mathfrak{n}'$ is not empty, so by Claim 1, it is a closed model. Finally, it follows immediately from Claim 4 and 5 that $\mathfrak{n}'$ has the required properties. QED

Lemma 3.5 *For every finite set $\mathbb{F}$, the relation $\equiv_{\mathbb{F}} \cap (M \times M_f)$ is a $\sqsupset$-bisimulation between $\mathcal{S}$ and $\mathcal{S}_f$. That is, for all closed $\mathfrak{m}$ and clopen $\mathfrak{m}'$ satisfying $\mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{m}'$:*

1. *For all closed $\mathfrak{n}$ such that $\mathfrak{m} \sqsupset \mathfrak{n}$ there is a clopen $\mathfrak{n}'$ such that $\mathfrak{m}' \sqsupset \mathfrak{n}' \equiv_{\mathbb{F}} \mathfrak{n}$.*

2. *For all clopen $\mathfrak{n}'$ such that $\mathfrak{m}' \sqsupset \mathfrak{n}'$ there is a closed $\mathfrak{n}$ such that $\mathfrak{m} \sqsupset \mathfrak{n} \equiv_{\mathbb{F}} \mathfrak{n}'$.*

Proof. By an argument which is similar to, but slightly more sophisticated than the one used in the proof of Lemma 3.4. To start with, note that part 2 of the Lemma is an immediate consequence of Lemma 3.4; hence, we may confine ourselves to the proof of part 1.

Assume that $\mathfrak{m}$ and $\mathfrak{n}$ are closed models such that $\mathfrak{n} \sqsubset \mathfrak{m}$; that $\mathfrak{m}'$ is clopen and that $\mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{m}'$. It follows from Lemma 3.4 that there is a *closed* $\mathfrak{n}'$ such that $\mathfrak{m}' \sqsupset \mathfrak{n}' \equiv_{\mathbb{F}} \mathfrak{n}$, but we need a *clopen* $\mathfrak{n}'$ with this property — this is why we have to take a bit more care.

As before, we define sets $\mathfrak{n}'_W$ for each $\sim_{\mathbb{F}}$-equivalence class W, and $\mathfrak{n}'$ is to be the finite union of these sets. The case distinction is the same as before, and so is the definition of $\mathfrak{n}'_W$ in the cases 1, 2 and 3a. Note that in each of these cases, $\mathfrak{n}'_W$ is clopen.

The difference lies in the definition of $\mathfrak{n}'_W$ in case 3b; recall that this is the case where $\mathfrak{n}$ is neither full nor empty on W, while $\mathfrak{m} \cap W$ is *properly* contained in $\mathfrak{n} \cap W$. Take some propositional variable $q \notin \mathbb{F}$ which does *not* occur in $\delta_{\mathfrak{m}}$, and define

$$\mathfrak{n}'_W := W \cap \mathit{Mod}(\delta_{\mathfrak{m}} \vee q).$$

Then $\mathfrak{n}'_W$ is clopen, being the intersection of two clopen sets. Hence $\mathfrak{n}'_W$ *properly* contains $W \cap \mathfrak{m} = W \cap \mathit{Mod}(\delta_{\mathfrak{m}})$ since q does not occur in $\delta_{\mathfrak{m}}$ and $q \notin \mathbb{F}$. Apart from this point, the proof is completely analogous to the proof of Lemma 3.4. QED

Proof of Proposition 3.1. Fix a formula φ, and let $\mathbb{F}$ be the set of propositional variables occurring in φ. Recall that thus, $\mathcal{L}_{2d}(\mathbb{F})$ denotes the set of formulas in which only propositional variables may occur that occur in φ.

We will show that φ is satisfiable in $\mathcal{S}$ iff it is satisfiable in $\mathcal{S}_f$. First we prove the following claim, according to which $\mathbb{F}$-alike states agree on the truth of $\mathbb{F}$-formulas. Note that part 1 of the Proposition is immediate by this claim and the observation that $\mathfrak{n} \equiv_{\mathbb{F}} \mathfrak{n}$ for *any* state $\mathfrak{n}$.

Claim 1 Let ψ be a formula in $\mathcal{L}_{2d}(\mathbb{F})$. Then for all closed $\mathfrak{m}$ and clopen $\mathfrak{m}'$ such that $\mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{m}'$:

$$\mathcal{S}, \mathfrak{m} \Vdash \psi \text{ iff } \mathcal{S}_f, \mathfrak{m}' \Vdash \psi.$$

Proof of Claim First we prove a preliminary fact concerning the truth of epistemic formulas. Let μ be some formula in $\mathcal{L}_1(\mathbb{F})$, then

$$\begin{array}{l}\text{if } \mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{m}', \text{ and } w \in \mathfrak{m},\ w' \in \mathfrak{m}' \text{ are such that } w \sim_{\mathbb{F}} w', \\ \text{then } \mathfrak{m}, w \Vdash \mu \text{ iff } \mathfrak{m}', w' \Vdash \mu.\end{array} \tag{2}$$

We prove (2) by induction on the complexity of μ.

For the atomic case, where μ is some classical formula, we need the following result which holds for any $\alpha \in \mathcal{L}_0(\mathbb{F})$:

$$\text{if } w \sim_{\mathbb{F}} w', \text{ then } w \models \alpha \text{ iff } w' \models \alpha. \tag{3}$$

The proof of (3), which goes by a straightforward induction on α, is left to the reader. The base case of (2) is an immediate consequence of (3).

For the induction step of (2), we omit the trivial boolean cases, and concentrate on the case where $\mu \equiv K\nu$.

Assume that $\mathfrak{m}, w \Vdash K\nu$, and let $w' \in \mathfrak{m}'$ be such that $w \sim_{\mathbb{F}} w'$. In order to show that $\mathfrak{m}', w' \Vdash K\nu$, let v' be an arbitrary valuation in $\mathfrak{m}'$. We have to show that $\mathfrak{m}, v' \Vdash \nu$.

It follows rather easily from the definition of $\equiv_{\mathbb{F}}$ that for all closed $\mathfrak{n}$, $\mathfrak{n}'$:

$$\mathfrak{n} \equiv_{\mathbb{F}} \mathfrak{n}' \text{ only if for all } u \in \mathfrak{n} \text{ there is a } u' \in \mathfrak{n}' \text{ such that } u \sim_{\mathbb{F}} u'. \tag{4}$$

Hence, there must be a $v \in \mathfrak{m}$ such that $v \sim_{\mathbb{F}} v'$. By the clause for K in the truth definition, we obtain $\mathfrak{m}, v \Vdash \nu$, so by the inductive hypothesis, $\mathfrak{m}', v' \Vdash \nu$, as required. The other direction, i.e., the proof that $\mathfrak{m}', w' \Vdash K\nu$ only if $\mathfrak{m}, w \Vdash K\nu$, is of course completely symmetric. This proves (2).

Now the proof of Claim 1 is by a straightforward induction on the complexity of ψ. The base case, where ψ is a subjective formula, follows from (2) and (4). the induction step is completely trivial for each of the boolean connectives, and the case where ψ is of the form $\Diamond_d \chi$ follows by Lemma 3.5 and the induction hypothesis. ◀

Now we turn to the proof of part 2 of the Proposition. First assume that φ is satisfiable in $\mathcal{S}_f$, say $\mathcal{S}_f, \mathfrak{m} \Vdash \varphi$. It is obvious that $\equiv_{\mathbb{F}}$ is reflexive, and since clopen sets are closed, it follows from Claim 1 that $\mathcal{S}, \mathfrak{m} \Vdash \varphi$. Hence, φ is satisfiable in $\mathcal{S}$.

For the other direction, we need the following claim.

Claim 2 For every closed $\mathfrak{m}$ there is a clopen $\mathfrak{n}$ such that $\mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{n}$.

Proof of Claim We will transform $\mathfrak{m}$ into a clopen $\mathfrak{n}$ such that $\mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{n}$. As in the proof of Lemma 3.4, we will treat each $\sim_{\mathbb{F}}$-equivalence class separately. Let p be some propositional variable not in $\mathbb{F}$.

Define, for an arbitrary $\sim_{\mathbb{F}}$-equivalence class W:

$$\mathfrak{n}_W := \begin{cases} \varnothing & \text{if } \mathfrak{m} \text{ is empty on } W, \\ W & \text{if } \mathfrak{m} \text{ is full on } W, \\ W \cap Mod(p) & \text{otherwise,} \end{cases}$$

and put

$$\mathfrak{n} := \bigcup_{w \in \mathcal{V}} \mathfrak{n}_{[w]_{\mathbb{F}}}$$

We leave it to the reader to verify that $\mathfrak{n}_W$ is clopen for each W (use the fact that W itself is clopen), and that $\mathfrak{n} \cap W = \mathfrak{n}_W$. Now assume, for arbitrary W, that $\mathfrak{m}$ is full on W. Then by the definition of $\mathfrak{n}_W$, $\mathfrak{n}_W = W$; in other words, $\mathfrak{n}$ is full on W. For the other direction, assume that $\mathfrak{n}$ is full on W; this gives $\mathfrak{n}_W = W$, which can only happen if $\mathfrak{m}$ is full on W. We leave it to the reader to verify that $\mathfrak{m}$ is empty on W iff $\mathfrak{n}$ is empty on W. But by definition this means that $\mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{n}$. ◀

In order to finish the proof of part 2 of the Proposition, assume that φ is satisfied in $\mathcal{S}$, say, at the closed information state $\mathfrak{m}$. By Claim 2, there is a clopen model $\mathfrak{m}'$ such that $\mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{m}'$. It follows from Claim 1 that $\mathcal{S}_f, \mathfrak{m}' \Vdash \varphi$. Hence, φ is satisfiable in $\mathcal{S}$ only if it is satisfiable in $\mathcal{S}_f$. QED

We have shown that the difference between the supermodels $\mathcal{S}_f$ and $\mathcal{S}$ is not reflected in the downdate fragment of our logic. It is reflected, however, in the full language $\mathcal{L}_2$. Consider, for example, the formula $\Diamond_u \top$ (where $\top$ is $\neg\bot$) of $\mathcal{L}_2$, expressing the fact that we can always increase our knowledge. It is not valid in $\mathcal{S}$: take any valuation w, and consider the model $\{w\}$. As singletons are closed, we have $\{w\} \in M$ and $\mathcal{S}, \{w\} \nVdash \Diamond_u \top$. On the other hand, it is valid in $\mathcal{S}_f$. Let $\mathfrak{m} \in M_f$ be arbitrary, and suppose $\mathfrak{m} = Mod(\delta_{\mathfrak{m}})$. Define $\mathfrak{n} = Mod(\delta_{\mathfrak{m}} \wedge p)$, where p is a propositional variable not occurring in $\delta_{\mathfrak{m}}$. Since $\delta_{\mathfrak{m}}$ must be consistent ($\mathfrak{m} \neq \varnothing$), $\delta_{\mathfrak{m}} \wedge p$ is consistent so $\mathfrak{n} \neq \varnothing$. Thus we have that $\mathfrak{n} \in M_f$ and $\mathfrak{n} \sqsupset \mathfrak{m}$, from which it follows that $\mathcal{S}_f, \mathfrak{m} \Vdash \Diamond_u \top$.

From these observations it follows that $\equiv_{\mathbb{F}}$ does not imply modal equivalence between closed states in $\mathcal{S}$ and clopen ones in $\mathcal{S}_f$. But if we confine ourselves to $\mathcal{S}_f$, the relation $\equiv_{\mathbb{F}}$ *is* a bisimulation with respect to *both* $\sqsupset$ and $\sqsubset$.

Proposition 3.6 *For every finite set of variables $\mathbb{F}$, the relation $\equiv_{\mathbb{F}} \cap (M_f \times M_f)$ is a $\sqsupset, \sqsubset$-bisimulation on $\mathcal{S}_f$. That is, for all clopen $\mathfrak{m}, \mathfrak{m}'$ satisfying $\mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{m}'$:*

1. *For all clopen $\mathfrak{n}$ with $\mathfrak{m} \sqsupset \mathfrak{n}$ there is a clopen $\mathfrak{n}'$ such that $\mathfrak{m}' \sqsupset \mathfrak{n}' \equiv_{\mathbb{F}} \mathfrak{n}$.*
2. *For all clopen $\mathfrak{n}$ with $\mathfrak{m} \sqsubset \mathfrak{n}$ there is a clopen $\mathfrak{n}'$ such that $\mathfrak{m}' \sqsubset \mathfrak{n}' \equiv_{\mathbb{F}} \mathfrak{n}$.*

Proof. Part 1 follows from part 1 of Lemma 3.5. For part 2, we have a similar proof of which we only give a sketch here.

Let $\mathfrak{m}, \mathfrak{m}'$ and $\mathfrak{n}$ be as in the formulation of part 2 of this Proposition. As before, we find $\mathfrak{n}'$ as the finite union of sets $\mathfrak{n}'_W$ which are defined as follows.

$$\mathfrak{n}'_W \;:=\; \begin{cases} \varnothing & \text{if } \mathfrak{n} \text{ is empty on } W, \\ W & \text{if } \mathfrak{n} \text{ is full on } W, \\ \mathfrak{m}' \cap W & \text{if } W \cap \mathfrak{m} = W \cap \mathfrak{n}, \\ Mod(\delta_{\mathfrak{m}} \wedge q) \cap W & \text{if } W \cap \mathfrak{n} \subset W \cap \mathfrak{m} \\ & \quad \text{(where } q \notin \mathbb{F} \text{ does not occur in } \delta_{\mathfrak{m}}). \end{cases}$$

We leave it to the reader to verify that this definition indeed gives the required $\mathfrak{n}'$. QED

As a corollary, we have the following.

Proposition 3.7 *Let $\mathbb{F}$ be a finite set of variables, and $\mathfrak{m}$ and $\mathfrak{m}'$ clopen states such that $\mathfrak{m} \equiv_{\mathbb{F}} \mathfrak{m}'$. Then for every $\mathcal{L}_2(\mathbb{F})$-formula ψ:*

$$\mathcal{S}_f, \mathfrak{m} \Vdash \psi \iff \mathcal{S}_f, \mathfrak{m}' \Vdash \psi.$$

Proof. Completely analogous to the proof of Claim 1 in the proof of Proposition 3.1. QED

As a final remark, remember that the ordering $\sqsupset$ is dense on $\mathcal{S}_f$. Hence, we have that $\mathcal{S}_f$ validates the density axiom: $\mathcal{S}_f \Vdash \Diamond_d\varphi \rightarrow \Diamond_d\Diamond_d\varphi$ for any φ. From Proposition 3.1 it follows that $\mathcal{S} \Vdash \Diamond_d\varphi \rightarrow \Diamond_d\Diamond_d\varphi$, even though $\sqsupset$ is not dense on $\mathcal{S}$.

4 A proof system

In this section, we will present a Hilbert-style proof system for validities of $\mathcal{S}_f$. Since the logic is built 'on top of' S5, the proof system contains S5-axioms. Furthermore, there are axioms about the properties of the accessibility relation $\sqsubset$. As this relation is completely fixed (our class of models contains just one element, $\mathcal{S}_f$), there are axioms describing the fact that $\sqsubset$ is the superset relation on S5-models.

Before we present the axioms and rules, we first need to distinguish a special class of formulas. Remember that $M\mu$ is an abbreviation of $\neg K \neg \mu$.

Definition 4.1 *Define the class of* downward persistent *formulas,* DP*, 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.

These formulas only express ignorance of the agent, and it should be expected that whenever $\mathfrak{m} \Vdash \mu$ for a downward persistent formula, it also holds that $\mathfrak{n} \Vdash \mu$ if $\mathfrak{n} \sqsubset \mathfrak{m}$, since the state $\mathfrak{n}$ holds less information than $\mathfrak{m}$. This is indeed the case, and in [4] it is proved that these formulas are the only subjective formulas, up to S5-equivalence, for which this is always the case.

Proposition 4.2 *Let μ be a subjective S5-formula, then the following are equivalent.*

1. *For all $\mathfrak{m}, \mathfrak{n} \in \mathcal{S}_f$:*
$$\mathfrak{n} \sqsubset \mathfrak{m} \ \&\ \mathfrak{m} \Vdash \mu \Rightarrow \mathfrak{n} \Vdash \mu$$
2. *There is a formula $\nu \in \mathsf{DP}$ which is S5-equivalent to μ.*

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 4.3 *The proof system* IC *(for* information change*) consists of the following axioms:*

CT	*all instances of classical tautologies*	
DB	$K(\mu \to \nu) \to (K\mu \to K\nu)$	
	$\Box_d(\varphi \to \psi) \to (\Box_d\varphi \to \Box_d\psi)$	
	$\Box_u(\varphi \to \psi) \to (\Box_u\varphi \to \Box_u\psi)$	
$A1$	$K\mu \to \mu$	
$A2$	$K\mu \to KK\mu$	
$A3$	$\neg K\mu \to K\neg K\mu$	
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	$\Diamond_d \Diamond_d \varphi \to \Diamond_d \varphi$	

and the following derivation rules:

MP *Modus Ponens (for both sorts)*

$$\frac{\varphi \to \psi \qquad \varphi}{\psi}$$

N *Necessitation for K, $\Box_d$ and $\Box_u$*

$$\frac{\mu}{K\mu} \quad \frac{\varphi}{\Box_d \varphi} \quad \frac{\varphi}{\Box_u \varphi}$$

OE 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.

The condition of **OE** may require some explanation: although there are infinitely many formulas in $\mathcal{L}_0(\mathbb{V}_\varphi \cup \{p\}\}$, one needs to prove $\mathsf{O}K\alpha \to \varphi$ only for mutually inequivalent formulas α. There are only finitely many of these.

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 proof, theorem, consistency and maximal consistent set are standard:

Definition 4.4 *A **derivation** in IC is a finite sequence $\varphi_1, \ldots, \varphi_n$ of formulas such that for every $i \in \{1, \ldots, n\}$, φ_i is either an axiom, or the result of applying a rule to a subset of the formulas $\{\varphi_1, \ldots, \varphi_{i-1}\}$. A formula φ is a **theorem** of IC, denoted $\vdash_{\mathsf{IC}} \varphi$, if there is a derivation $\varphi_1, \ldots, \varphi_n$ such that $\varphi_n = \varphi$. A formula φ is **provable** in IC **from** a set of formulas A, denoted $A \vdash_{\mathsf{IC}} \varphi$, if there are $\varphi_1, \ldots \varphi_n \in A$ such that $\vdash_{\mathsf{IC}} (\varphi_1 \wedge \ldots \wedge \varphi_n) \to \varphi$. A set of formulas A is **consistent** in IC if $A \not\vdash_{\mathsf{IC}} \bot$. A formula φ is consistent if $\{\varphi\}$ is. A set of formulas A is a **maximally consistent set** (MCS) if it is consistent, and any proper superset is inconsistent.*

The system IC axiomatizes validity in $\mathcal{S}_f$.

Theorem 4.5 (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 and their soundness. The completeness of IC is the subject of the next section.

The axioms CT ('classical tautologies'), DB ('distribution') for K, $A1$ through $A3$ and the rules **N** and **MP** form a standard proof system for S5, sound and complete with respect to the semantics of Definition 2.2 (see [11] for a proof). As we are only considering subjective formulas, one could replace

$A1$ by the axiom $\neg K\bot$ (the resulting system is often called KD45, which is also sound and complete with respect to our semantics – with the restriction to subjective formulas –). The DB axioms and the rules Modus Ponens and Necessitation are sound in any Kripke semantics. The CV axioms ('converse') express the fact that $\sqsubset$ and $\sqsupset$ are each other's converses. The soundness of the axiom DP ('downward persistence') forms part of Proposition 4.2. Axiom 4 is sound when the accessibility relation is transitive, which $\sqsubset$ obviously is. The other axioms and rules are much less standard, and we will treat them in more detail.

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. An important instance of this rule is the following theorem (taking $\alpha = \top$): $K\beta \to \Diamond_d(\neg K\beta)$ provided that $\not\vdash_{CL} \beta$: the agent may forget any non-tautological information.

For a proof of the soundness of this rule, suppose $\mathfrak{m} \Vdash K\alpha \wedge K\beta$. Define $\mathfrak{n} = Mod(\alpha)$, then we have $\mathfrak{m} \subseteq Mod(\alpha \wedge \beta) \subset Mod(\alpha)$, where the last inclusion is strict since $\not\vdash_{CL} \alpha \to \beta$. Then $\mathfrak{n} \sqsubset \mathfrak{m}$ and $\mathfrak{n} \Vdash K\alpha \wedge \neg K\beta$ so $\mathfrak{m} \Vdash \Diamond_d(K\alpha \wedge \neg K\beta)$.

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. So suppose $\mathfrak{m} \Vdash K\alpha \wedge \Diamond_d K\alpha$. Then $\mathfrak{m} \subset Mod(\alpha)$, and by Corollary 2.8 $Mod(\alpha) \Vdash \mathsf{O}K\alpha$ whence $\mathfrak{m} \Vdash \Diamond_d \mathsf{O}K\alpha$.

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 $\not\vdash_{CL} \beta \to \alpha$), then it may perform an update to a state where all it knows is α. As for the soundness, suppose $\mathfrak{m} \Vdash \mathsf{O}K\beta \wedge \neg K\alpha$. Then by Corollary 2.8 it follows that $\mathfrak{m} = Mod(\beta)$. As $\vdash_{CL} \alpha \to \beta$, we have $Mod(\alpha) \subset Mod(\beta)$; this inclusion is strict since $\mathfrak{m} \Vdash \neg K\alpha$. But since $Mod(\alpha) \Vdash \mathsf{O}K\alpha$, we have $\mathfrak{m} \Vdash \Diamond_u \mathsf{O}K\alpha$.

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. The soundness of this rule requires a bit more of explanation.

Proposition 4.6 *The rule* ***OE*** *is sound: if* $\mathcal{S}_f \Vdash \mathsf{O}K\alpha \to \varphi$ *for all non-equivalent formulas* α *in* $\mathcal{L}_0(\mathbb{V}_\varphi \cup \{p\})$ *(with* p *not occurring in* φ*), then* $\mathcal{S}_f \Vdash \varphi$.

Proof. We will prove the contraposition. Suppose $\mathcal{S}_f \not\Vdash \varphi$, then there is an $\mathfrak{m} \in M_f$ such that $\mathcal{S}_f, \mathfrak{m} \Vdash \neg\varphi$. Let $\mathbb{F} = \mathbb{V}_\varphi$. As in the second claim of the proof of Proposition 3.1, we define, for any $\sim_\mathbb{F}$-equivalence class W:

$$\mathfrak{n}_W := \begin{cases} \varnothing & \text{if } \mathfrak{m} \cap W = \varnothing \\ W & \text{if } W \subset \mathfrak{m} \\ W \cap Mod(p) & \text{otherwise,} \end{cases}$$

and put

$$\mathfrak{n} := \bigcup_{w \in \mathcal{V}} \mathfrak{n}_{[w]_\mathbb{F}}$$

Now set $\alpha = \bigvee\{\theta_w \mid [w]_{\mathbb{F}} \subseteq \mathfrak{m}\} \vee \bigvee\{\theta_w \wedge p \mid \mathfrak{m}$ neither $\mathbb{F}$-full nor $\mathbb{F}$-empty on $[w]_{\mathbb{F}}\}$. It is straightforward to prove that $\mathfrak{n} = Mod(\alpha)$, which implies that $\mathcal{S}_f, \mathfrak{n} \Vdash \mathsf{O}K\alpha$, with $\alpha \in \mathcal{L}_0(\mathbb{V}_\varphi \cup \{p\})$. As $\mathfrak{n} \equiv_{\mathbb{F}} \mathfrak{m}$, using Proposition 3.7 we have that $\mathcal{S}_f, \mathfrak{n} \Vdash \neg\varphi$. But this means that $\mathcal{S}_f \nVdash \mathsf{O}K\alpha \to \varphi$. QED

As an example of the use of this proof system, we will sketch the derivation of the formula $\Diamond_u \top$ we discussed earlier. Using **OE**, it is sufficient to prove the following four formulas:

$$\mathsf{O}K\bot \to \Diamond_u\top \quad \mathsf{O}K\top \to \Diamond_u\top \quad \mathsf{O}Kp \to \Diamond_u\top \quad \mathsf{O}K\neg p \to \Diamond_u\top$$

The first of these is easy, since $\vdash_{\mathsf{IC}} \mathsf{O}K\bot \to K\bot \to \bot \to \Diamond_u\top$.

For the second, $\vdash_{\mathsf{IC}} \mathsf{O}K\top \to \neg Kp$: by axiom SF, we have $\vdash_{\mathsf{IC}} K\top \wedge Kp \to \Diamond_d(K\top \wedge \neg Kp)$ (as $\nvdash_{CL} \top \to p$). Also, $\vdash_{\mathsf{IC}} \Diamond_d(K\top \wedge \neg Kp) \to \Diamond_d K\top$. Obviously, $\Diamond_d K\top$ contradicts $\mathsf{O}K\top$. Furthermore, by OU we may conclude $\vdash_{\mathsf{IC}} \mathsf{O}K\top \wedge \neg Kp \to \Diamond_u \mathsf{O}Kp$ (as $\vdash_{CL} p \to \top$). Combining these derivations, we get $\vdash_{\mathsf{IC}} \mathsf{O}K\top \to (\mathsf{O}K\top \wedge \neg Kp) \to \Diamond_u\mathsf{O}Kp \to \Diamond_u\top$.

The third and fourth formula have the same proof modulo exchanging $\neg p$ for p, so we will only treat the third. As for this formula, we want to find a formula we do not know when we only know p, and we can take $p \wedge q$ for this purpose. Then $\vdash_{\mathsf{IC}} K(p \wedge q) \to (Kp \wedge Kq) \to \Diamond_d(Kp \wedge \neg Kq) \to \Diamond_d Kp \to \neg\mathsf{O}Kp$, where the use of axiom SF needs $\nvdash_{CL} p \to q$. This means that $\vdash_{\mathsf{IC}} \mathsf{O}Kp \to \neg K(p \wedge q)$. Then we can use axiom OU: $\vdash_{\mathsf{IC}} \mathsf{O}Kp \wedge \neg K(p \wedge q) \to \Diamond_u\mathsf{O}K(p \wedge q)$ (where $\vdash_{CL} (p \wedge q) \to p$). Concluding, this means that $\vdash_{\mathsf{IC}} \mathsf{O}Kp \to (\mathsf{O}Kp \wedge \neg K(p \wedge q)) \to \Diamond_u\mathsf{O}K(p \wedge q) \to \Diamond_u\top$.

Rule **OE** now allows us to conclude that $\vdash_{\mathsf{IC}} \Diamond_u\top$.

5 Completeness

This section is devoted to the completeness proof for IC with respect to $\mathcal{S}_f$; that is to say, we will prove that all $\mathcal{S}_f$-validities are derivable in IC.

As usual, we will prove our completeness result via contraposition, showing that any IC-consistent formula ξ can be satisfied in $\mathcal{S}_f$. Usually in modal completeness proofs, the key idea is to build a satisfying model in which the states are MCSs and the accessibility relation between MCSs is defined according to the internal structure of these sets of formulas. Our set-up is basically the same, with some twists that we will discuss further on.

We will first take care of the epistemic part of the language.

Definition 5.1 *For any* MCS Φ, *we let* Φ_K *denote the set of epistemic formulas in* Φ*; formally,*

$$\Phi_K := \Phi \cap \mathcal{L}_1.$$

The following result will simplify some of the proofs.

Lemma 5.2 *Every S5-formula is (provably) equivalent to a disjunction* $\bigvee_{i=1}^{n} \mu_i$*, where each* μ_i *is of the form*

$$K\alpha \wedge \neg K\beta_1 \wedge \ldots \wedge \neg K\beta_k \wedge \gamma$$

where each of the formulas α, $\beta_1, \ldots, \beta_k$ *and* γ *is propositional. In case* μ *is subjective, each conjunct can be taken in such a way that no* γ *is present.*

Proof. For the first part of the lemma, see [11], which contains a procedure for rewriting a formula into such a normal form. By inspection of this procedure, it is straightforward to see that the normal form of a subjective formula contains no purely propositional parts. QED

Lemma 5.3 *For each* MCS Φ *there is a unique closed model* $\mathfrak{m}$ *such that for all subjective epistemic formulas* μ*:*

$$\mathfrak{m} \Vdash \mu \iff \mu \in \Phi_K. \tag{5}$$

Proof. Define $\mathfrak{m} := \mathit{Mod}(\{\alpha \in \mathcal{L}_0 \mid K\alpha \in \Phi\})$, which is closed by definition. The set of formulas true in $\mathfrak{m}$ on the one hand, and the set of epistemic formulas in Φ on the other hand, are closed under S5-provability. For the former, this follows from soundness of the S5-axioms. For the latter, MCSs are closed under provability of $\vdash_{\mathsf{IC}}$ (which includes axioms for S5). Furthermore, as is always the case for MCSs, we have $\mu \wedge \nu \in \Phi$ iff $\mu \in \Phi$ and $\nu \in \Phi$, $\mu \vee \nu \in \Phi$ iff $\mu \in \Phi$ or $\nu \in \Phi$, and $\neg\mu \in \Phi$ iff $\mu \notin \Phi$. Using Lemma 5.2, it is sufficient to prove the statement of the lemma for formulas of the form $K\alpha$ with $\alpha \in \mathcal{L}_0$. The right to left direction is true by definition of $\mathfrak{m}$. For the other direction, suppose $\mathfrak{m} \Vdash K\alpha$, then $\{\beta \in \mathcal{L}_0 \mid K\beta \in \Phi\} \vdash_{CL} \alpha$. This in turn implies that $K\alpha \in \Phi$.

The uniqueness of $\mathfrak{m}$ follows from the fact that a closed model is uniquely determined by the formulas $K\alpha$ it validates (Proposition 2.7), which, in this case, are those contained in Φ. QED

Definition 5.4 *Given an* MCS Φ, *we let* $\mathfrak{p}_\Phi$ *denote the unique model satisfying (5).*

Now we know that MCSs determine information states, it is good to see how the information ordering between these information states can be traced back to the content of the MCSs involved.

Definition 5.5 *Define the binary relation* R *on* MCS*s as follows:*

$$R\Phi\Psi \iff \Diamond_u\Psi \subseteq \Phi \text{ and } \{K\alpha \in \Phi_K\} \subset \{K\alpha \in \Psi_K\}.$$

Here $\Diamond_u\Psi$ *is defined as the set* $\{\Diamond_u\varphi \mid \varphi \in \Psi\}$.

The relation R is *not* the same as the 'canonical' accessibility relation between MCSs known from standard modal completeness proofs — that relation only uses the clause '$\Diamond_u\Psi \subseteq \Phi$'. We have to add the second clause in order to ensure irreflexivity.

Lemma 5.6 *For any pair of* MCS*s* Φ *and* Ψ: $R\Phi\Psi$ *only if* $\mathfrak{p}_\Phi \sqsubset \mathfrak{p}_\Psi$.

PROOF. Suppose $R\Phi\Psi$, then by definition we have $K\alpha \in \Phi$ only if $K\alpha \in \Psi$, for all α. From the proof of Lemma 5.3, we know that $\mathfrak{p}_\Phi = Mod(\{\alpha \in \mathcal{L}_0 \mid K\alpha \in \Phi\}) = Mod(\{\alpha \in \mathcal{L}_0 \mid K\alpha \in \Phi_K\})$. It follows that $\mathfrak{p}_\Psi \subset \mathfrak{p}_\Phi$, or $\mathfrak{p}_\Phi \sqsubset \mathfrak{p}_\Psi$. QED

Since we want to prove completeness for $\mathcal{S}_f$, we are more interested in clopen models than in arbitrary closed ones. Fortunately, there is a simple criterion on MCSs ensuring clopenness of the induced model.

Definition 5.7 *An* MCS Φ *is called* ***witnessing*** *if there is an* $\mathcal{L}_0$*-formula* α *such that* $\mathsf{O}K\alpha \in \Phi$.

Lemma 5.8 *An* MCS Φ *is witnessing only if* $\mathfrak{p}_\Phi$ *is clopen.*

PROOF. Assume that the formula $\mathsf{O}K\alpha$ belongs to Φ. We will show that

$$\mathfrak{p}_\Phi = Mod(\alpha).$$

For the inclusion from left to right, let $w \in \mathcal{V}$ be a member of $\mathfrak{p}_\Phi$. Since $K\alpha \in \Phi$, we have $w \models \alpha$ by definition of $\mathfrak{p}_\Phi$, so $w \in Mod(\alpha)$.

For the other direction, suppose that w is a valuation such that $w \models \alpha$ while $w \notin \mathfrak{p}_\Phi$. From the latter fact we infer that $w \models \neg\beta$ for some β with $K\beta \in \Phi$. But then $w \models \alpha$ implies that $\not\vdash_{CL} \alpha \to \beta$. By axiom SF this gives

$$\vdash_{\mathsf{IC}} (K\alpha \wedge K\beta) \to \Diamond_d(K\alpha \wedge \neg K\beta).$$

But this would imply that $\Diamond_d K\alpha \in \Phi$, which clearly gives a contradiction with the assumption that $\mathsf{O}K\alpha \in \Phi$. QED

Lemma 5.9 *There is a set* W *of witnessing* MCS*s satisfying the following conditions:*

1. *for each consistent formula φ there is a $\Phi \in W$ with $\varphi \in \Phi$,*

2. *for each $\Phi \in W$ and each formula $\Diamond_u \varphi \in \Phi$, there is a $\Psi \in W$ such that $R\Phi\Psi$ and $\varphi \in \Psi$,*

3. *for each $\Phi \in W$ and each formula $\Diamond_d \varphi \in \Phi$, there is a $\Psi \in W$ such that $R\Psi\Phi$ and $\varphi \in \Psi$.*

PROOF. The proof of this lemma is rather involved, and since it is a special case of a more general method, we will only give a brief outline here, referring the reader to [15] or [5] for more details.

For a fixed consistent formula φ, we want to find a set W_φ of witnessing MCSs satisfying the conditions 2 and 3 and such that $\varphi \in \Phi$ for some $\Phi \in W_\varphi$. Then the union $\bigcup\{W_\varphi \mid \varphi \text{ consistent }\}$ is easily shown to meet the requirements of the Lemma.

The basic idea of the proof is to define W_φ step by step, in a sort of parallel Lindenbaum construction on graphs. During the construction we are dealing with finite approximations of W_φ; at each stage, one of the shortcomings of the current approximations is taken care of; this can be done in such a way that the limit of the construction has no shortcomings at all.

Let us give a bit more detail; a finite approximation of W_φ consists of a finite, directed graph together with a labeling which assigns a finite set of formulas to each vertex of the graph. We associate an $\mathcal{L}_2$-formula with each of these finite labeled graphs, and require that this corresponding formula is consistent for each of the approximations. The first graph has no edges, and just one point of which the label set is the singleton $\{\varphi\}$. The construction is such that the graph is growing in the sense that edges may be added to the graph, and formulas to the label sets. All this is done to ensure that in the limit we are dealing with a possibly infinite labeled graph meeting the requirements that (1) the label set of each point is a MCS, (2) each label set contains a witness and (3) if a formula of the form $\Diamond_u \psi$ ($\Diamond_d \psi$) belongs to the label set of some vertex, then there is an edge leading from this vertex to another one (or conversely, in the case of $\Diamond_d \psi$) containing ψ in its label set. Finally, W_φ is defined as the range of this infinite labeling function. QED

Now we are ready for the final part of the completeness proof. In it, we will heavily use the fact that formulas of the form $\mathsf{O}K\alpha$ *name* states in the clopen model; that is $Mod(\alpha)$ is the *unique* clopen state where $\mathsf{O}K\alpha$ holds. Another use of this fact is that for all formulas ψ

$$\mathcal{S}_f, \mathfrak{m} \Vdash \psi \text{ iff } \mathcal{S}_f \Vdash \mathsf{O}K\delta_{\mathfrak{m}} \rightarrow \psi, \tag{6}$$

since $\mathfrak{m}$ is the unique state where the formula $\mathsf{O}K\delta_{\mathfrak{m}}$ holds. In the theory of modal logic, formalisms that have the expressive power to name points, often display some special behaviour — like the fact that special rules are needed in their axiomatization. For details concerning *modal logic with names* we refer to [2, 6].

The usual procedure in a modal completeness proof is a lemma stating that for any formula φ and any MCS Φ, φ is true at the possible world determined by

Φ if and only if φ is a member of Φ. Such a *truth lemma* is proved by induction on φ. Here, we will use a similar lemma, but in order to prove it, we need to strengthen the inductive hypothesis. In particular, in order to prove that $\mathcal{S}_f, \mathfrak{p}_\Phi \Vdash \varphi$ only if $\varphi \in \Phi$ (for every $\Phi \in W$), we need that (6) holds for all formulas ψ that are 'simpler' than φ. Now since the formula $\mathsf{O}K\delta_{\mathfrak{m}} \to \psi$ need not be a subformula of φ, we need a more complex inductive hypothesis. This explains the following definition and the precise set-up of the completeness proof where completeness itself *and* the truth lemma are proved in a *simultaneous* induction.

Definition 5.10 *We define the* **c-depth** *of a formula as follows:*

$$\begin{array}{rcl} \mathsf{c}(\mu) & = & 0 \\ \mathsf{c}(\neg\varphi) & = & \mathsf{c}(\varphi) \\ \mathsf{c}(\varphi \wedge \psi) & = & \max(\mathsf{c}(\varphi), \mathsf{c}(\psi)) \\ \mathsf{c}(\Diamond_d \varphi) & = & \mathsf{c}(\varphi) + 1 \\ \mathsf{c}(\Diamond_u \varphi) & = & \mathsf{c}(\varphi) + 2 \end{array}$$

PROOF OF COMPLETENESS. Fix a set W of witnessing MCSs satisfying the conditions of Lemma 5.9. We will prove that for any $\mathcal{L}_2$-formula φ, the following two claims hold

$$\text{for every } \Phi \in W\text{: } \varphi \in \Phi \iff \mathcal{S}_f, \mathfrak{p}_\Phi \Vdash \varphi, \tag{7}$$

and

$$\varphi \text{ is consistent only if } \varphi \text{ is satisfiable in } \mathcal{S}_f. \tag{8}$$

The proof of (7) and (8) is by induction on the c-depth of φ. Note that for any φ, (8) follows from (7) and part 1 of Lemma 5.9. Hence, we will confine ourselves to proving (7); we will *use* the inductive hypothesis of (8), however.

CASE $\mathsf{c}(\varphi) = 0$ In this case we are dealing with a subjective epistemic formula. Hence, part 7 follows from Lemma 5.3.

CASE $\mathsf{c}(\varphi) = 1$ Now φ must be a boolean combination of epistemic formulas and formulas of the form $\Diamond_d \mu$, with μ epistemic. We leave it to the reader to verify that the boolean connectives do not cause any problem. That is, we only treat the case where φ is of the form $\Diamond_d \mu$.

First assume that $\Diamond_d \mu \in \Phi$. The conditions of Lemma 5.9 imply that there is a $\Psi \in W$ with $R\Psi\Phi$ and $\mu \in \Psi$. By Lemma 5.6 and the inductive hypothesis, this implies $\mathfrak{p}_\Phi \sqsupset \mathfrak{p}_\Psi$ and $\mathcal{S}_f, \mathfrak{p}_\Psi \Vdash \mu$. But then by the truth definition, we have that $\mathcal{S}_f, \mathfrak{p}_\Phi \Vdash \Diamond_d \mu$.

For the other direction, without loss of generality (see Lemma 5.2) we may assume that μ is a disjunction of formulas of the following form: $K\alpha \wedge \neg K\beta_1 \wedge \ldots \wedge \neg K\beta_n$, and since $\Diamond_d$ distributes over disjunctions, we may actually assume that

$$\mu \text{ is of the form } K\alpha \wedge \neg K\beta_1 \wedge \ldots \wedge \neg K\beta_n.$$

Now assume that $\mathfrak{p}_\Phi \Vdash \Diamond_d \mu$. Then there is a clopen $\mathfrak{m} \supset \mathfrak{p}_\Phi$ such that $\mathcal{S}_f, \mathfrak{m} \Vdash K\alpha \wedge \bigwedge_i \neg K\beta_i$. This implies that $\mathfrak{p}_\Phi \Vdash K\alpha$ and for some γ, $\mathfrak{p}_\Phi \Vdash K\gamma$ while

$\mathfrak{m} \not\Vdash K\gamma$. From $\mathfrak{p}_\Phi \Vdash K\alpha \wedge K\gamma$ we may infer (by the previous case) that $K\alpha \wedge K\gamma \in \Phi$. Also, from $\mathfrak{m} \Vdash K\alpha \wedge \neg K\gamma \wedge \bigwedge_i \neg K\beta_i$ it follows that $\not\vdash_{CL} \alpha \to \gamma$ and $\not\vdash_{CL} \alpha \to \beta_i$ (all i).

Claim 1 If $\not\vdash_{CL} \alpha \to \gamma$ and $\not\vdash_{CL} \alpha \to \beta_i$ (all i), then

$$\vdash_{\mathsf{IC}} (K\alpha \wedge K\gamma) \to \Diamond_d(K\alpha \wedge \neg K\gamma \wedge \bigwedge_{0<i\leq n} \neg K\beta_i).$$

Proof of Claim We prove this claim by induction on n. In the base step ($n = 0$), it is immediate that $\not\vdash_{CL} \alpha \to \gamma$ implies $\vdash_{\mathsf{IC}} (K\alpha \wedge K\gamma) \to \Diamond_d(K\alpha \wedge \neg K\gamma)$: we are simply dealing with an instance of axiom SF.

For the induction step, assume that α, γ and $\beta_1, \ldots, \beta_n$ are such that $\not\vdash_{CL} \alpha \to \gamma$ and $\not\vdash_{CL} \alpha \to \beta_i$ (all i). For succinctness, abbreviate $\nu := \neg K\gamma \wedge \bigwedge_{0<i<n} \neg K\beta_i$; then our inductive hypothesis is that

$$\vdash_{\mathsf{IC}} (K\alpha \wedge K\gamma) \to \Diamond_d(K\alpha \wedge \nu),$$

from which it is easily inferred that

$$\vdash_{\mathsf{IC}} (K\alpha \wedge K\gamma) \to (\Diamond_d(K\alpha \wedge \nu \wedge K\beta_n) \vee \Diamond_d(K\alpha \wedge \nu \wedge \neg K\beta_n)). \tag{9}$$

Also, axiom SF and the fact that $\not\vdash_{CL} \alpha \to \beta_n$ ensure that

$$\vdash_{\mathsf{IC}} (K\alpha \wedge K\beta_n) \to (\Diamond_d(K\alpha \wedge \neg K\beta_n),$$

while axiom DP gives

$$\vdash_{\mathsf{IC}} \nu \to \Box_d \nu.$$

Bringing these two observations together we readily obtain

$$\vdash_{\mathsf{IC}} (K\alpha \wedge \nu \wedge K\beta_n) \to \Diamond_d(K\alpha \wedge \nu \wedge \neg K\beta_n),$$

so using the transitivity axiom, we find

$$\vdash_{\mathsf{IC}} \Diamond_d(K\alpha \wedge \nu \wedge K\beta_n) \to \Diamond_d(K\alpha \wedge \nu \wedge \neg K\beta_n).$$

But then (9) gives

$$\vdash_{\mathsf{IC}} (K\alpha \wedge K\gamma) \to \Diamond_d(K\alpha \wedge \nu \wedge \neg K\beta_n)),$$

which is precisely what we were after. ◀

From the claim and the fact that $K\alpha \wedge K\gamma \in \Phi$ it follows that the formula $\Diamond_d(K\alpha \wedge \neg K\gamma \wedge \neg K\beta_1 \wedge \ldots \wedge K\beta_n)$ and hence $\Diamond_d(K\alpha \wedge \neg K\beta_1 \wedge \ldots \wedge K\beta_n)$ belongs to Φ. But the latter formula is simply the formula φ, so we have proved (7) for the case $\mathsf{c}(\varphi) = 1$.

Case $\mathsf{c}(\varphi) > 1$ Again, we proceed by an induction on the structure of φ and again, we leave the boolean cases to the reader. This leaves the following two subcases.

SUBCASE φ is of the form $\Diamond_d\psi$, where we may apply the inductive hypothesis to ψ since $\mathsf{c}(\psi) < \mathsf{c}(\varphi)$.

We omit the proof that $\Diamond_d\psi \in \Phi$ implies $\mathcal{S}_f, \mathfrak{p}_\Phi \Vdash \Diamond_d\psi$, since it is similar to proofs we have seen before.

For the other direction, assume that $\mathcal{S}_f, \mathfrak{p}_\Phi \Vdash \Diamond_d\psi$. By definition this means that $\mathcal{S}_f, \mathfrak{m} \Vdash \psi$ for some clopen $\mathfrak{m}$ with $\mathfrak{p}_\Phi \subset \mathfrak{m}$. From $\mathcal{S}_f, \mathfrak{m} \Vdash \psi$ it follows that $\mathcal{S}_f \Vdash \mathsf{O}K\delta_\mathfrak{m} \to \psi$; here we use the fact that $\mathfrak{m}$ is the *only* state where $\mathsf{O}K\delta_\mathfrak{m}$ holds. By the inductive hypothesis (note: part 8!) we may infer that $\vdash_{\mathsf{IC}} \mathsf{O}K\delta_\mathfrak{m} \to \psi$. From this it is immediate that

$$\vdash_{\mathsf{IC}} \Diamond_d\mathsf{O}K\delta_\mathfrak{m} \to \Diamond_d\psi. \tag{10}$$

We now claim that

$$\Diamond_d\mathsf{O}K\delta_\mathfrak{m} \in \Phi. \tag{11}$$

For, since $\mathfrak{p}_\Phi \subset \mathfrak{m}$, we have $\mathfrak{p}_\Phi \Vdash K\delta_\mathfrak{m} \wedge \Diamond_d K\delta_\mathfrak{m}$. Hence, by the inductive hypothesis (part 7), the formula $K\delta_\mathfrak{m} \wedge \Diamond_d K\delta_\mathfrak{m}$ belongs to Φ. But then by axiom OD, $\Diamond_d\mathsf{O}K\delta_\mathfrak{m} \in \Phi$. This proves (11).

But from (10) and (11) it is immediate that $\Diamond_d\psi \in \Phi$.

SUBCASE φ is of the form $\Diamond_u\psi$, where we may apply the inductive hypothesis to ψ since $\mathsf{c}(\psi) < \mathsf{c}(\varphi)$.

The direction from left to right, being as in the previous subcase, is omitted. For the other direction, assume that $\mathcal{S}_f, \mathfrak{p}_\Phi \models \Diamond_u\psi$. By definition this means that $\mathcal{S}_f, \mathfrak{m} \Vdash \psi$ for some $\mathfrak{m}$ with $\mathfrak{m} \subset \mathfrak{p}_\Phi$. From $\mathcal{S}_f, \mathfrak{m} \Vdash \psi$ it follows that $\mathcal{S}_f \Vdash \mathsf{O}K\delta_\mathfrak{m} \to \psi$, from which it is immediate that

$$\vdash_{\mathsf{IC}} \Diamond_u\mathsf{O}K\delta_\mathfrak{m} \to \Diamond_u\psi. \tag{12}$$

This is all analogous to the previous subcase, and so is the next claim:

$$\Diamond_u\mathsf{O}K\delta_\mathfrak{m} \in \Phi, \tag{13}$$

but the remainder of the proof is different. Let γ be the formula such that $\mathsf{O}K\gamma \in \Phi$. Then by the inductive hypothesis, $\mathcal{S}_f, \mathfrak{p}_\Phi \Vdash \mathsf{O}K\gamma$. Hence, it follows from $\mathfrak{m} \subset \mathfrak{p}_\Phi$ that $\mathfrak{p}_\Phi \Vdash \mathsf{O}K\gamma \wedge \neg K\delta_\mathfrak{m}$, and also that $\vdash_{CL} \delta_\mathfrak{m} \to \gamma$ (this is by Proposition 2.7). But from this and axiom OU it follows that

$$\vdash_{\mathsf{IC}} (\mathsf{O}K\gamma \wedge \neg K\delta_\mathfrak{m}) \to \Diamond_u\mathsf{O}K\delta_\mathfrak{m}.$$

But we may use the induction hypothesis to show that $\mathsf{O}K\gamma \wedge \neg K\delta_\mathfrak{m}$ belongs to Φ, so it follows from the properties of MCSs that $\Diamond_u\mathsf{O}K\delta_\mathfrak{m} \in \Phi$. This proves (13).

And as in the previous subcase, from (12) and (13) it is immediate that $\Diamond_u\psi \in \Phi$. This finishes the proof for the case where $\mathsf{c}(\varphi) > 1$.

Finally, it is easy to see that the completeness of IC follows immediately from (8). QED

6 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. One of these choices is whether we opt for closed or finite-knowledge states.

Our approach, in which a modal logic (of increasing and decreasing information) is placed 'on top of' another modal logic (S5), fits in the recent trend of 'combining logics'. 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 ([8]) 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 [3]). There are strong connections between our system and the one of Levesque ([10]), a logic which also embeds the logic of only knowing.

For one particular kind of super model we have defined a proof system, and this system was proved to be sound and complete. In order to prove this result we used special techniques from modal logic, including non-standard derivation rules and the ability of our language to name points.

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, *Exploring Logical Dynamics*, CSLI Publications, Stanford, 1996.

[2] P. Blackburn, "Nominal tense logic", *Notre Dame Journal of Formal Logic* **14**:56–83, 1993.

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

[4] J. Engelfriet, "Minimal temporal epistemic logic", pp. 233–259 in [12].

[5] D.M. Gabbay, "An irreflexivity lemma with applications to axiomatizations of conditions on linear frames", in: U. Mönnich (ed.), *Aspects of Philosophical Logic*, Reidel, Dordrecht, 1981, pp. 67–89

[6] G. Gargov and V. Goranko, "Modal logic with names", *Journal of Philosophical Logic* **22**:607–636, 1993.

[7] J. Jaspars and E. Krahmer, "A programme of modal unification of dynamic theories", in: P. Dekker and M. Stokhof (eds.), *Proceedings of the Tenth Amsterdam Colloquium*, ILLC, Amsterdam, 1996, pp. 425–444.

[8] 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.

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

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

[11] J.-J.Ch. Meyer and W. van der Hoek, *Epistemic Logic for Computer Science and Artificial Intelligence*, Cambridge Tracts in Theoretical Computer Science, vol. 41, Cambridge University Press, Cambridge, 1995.

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

[13] Y. Shoham, "A semantical approach to nonmonotonic logics", in: J. McDermott (ed.), *Proceedings 10th IJCAI*, Kaufmann, 1987, pp. 388–392.

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

[15] Y. Venema, "Derivation rules as anti-axioms in modal logic", *Journal of Symbolic Logic* **58**:1003–1034, 1993.

[16] Y. Venema (ed.), "Modal Logic and Dynamic Semantics", special issue of the *Journal of Logic, Language and Information*, **6**, 1997.