

**REPORTS
IN
INFORMATICS**

ISSN 0333-3590

The Dynamics of Syntactic Knowledge

Thomas Ågotnes and Natasha Alechina

REPORT NO 304

September 2005



Department of Informatics
UNIVERSITY OF BERGEN
Bergen, Norway

This report has URL

<http://www.ii.uib.no/publikasjoner/texrap/pdf/2005-304.pdf>

Reports in Informatics from Department of Informatics, University of Bergen, Norway, is available
at <http://www.ii.uib.no/publikasjoner/texrap/>.

Requests for paper copies of this report can be sent to:

Department of Informatics, University of Bergen, Høyteknologisenteret,
P.O. Box 7800, N-5020 Bergen, Norway

The Dynamics of Syntactic Knowledge

Thomas Ågotnes^{*} Natasha Alechina[†]

Abstract

We add two syntactic epistemic modalities ‘knowing at least’ and ‘knowing at most’, to standard systems of modal logic, and investigate the resulting systems from the point of view of axiomatisation and complexity. We show how these logics can be used to formalise non-omniscient agents which know some inference rules, and study their relationship to other systems of syntactic epistemic logics, such as [2, 5, 8].

Keywords: Modal logic, epistemic logic, logical omniscience.

1 Introduction

In this paper, we pursue a standard line of research in epistemic logic: we model a development of a multi-agent system in which the knowledge or beliefs of the agents changes over time. Transitions in the Kripke structures used to interpret the logic correspond to acts of reasoning or communication. The distinguishing feature of the logics we are considering is the syntactic nature of the agents’ beliefs, which makes the agents non logically omniscient, and the use of two epistemic operators Δ (‘knowing at least’) and ∇ (‘knowing at most’), which apply to finite sets of formulas. These operators were introduced in [1], where the logic of static epistemic states was axiomatised and studied in detail, and a combination of syntactic epistemic knowledge operators and ATL was introduced. In this paper, we investigate intermediate systems, namely adding Δ and ∇ to standard modal logic.

The motivation behind this research is the desire to model non-omniscient reasoners, whose set of beliefs is not closed under logical consequence. Development of epistemic logics which do not suffer from the problem of logical omniscience is perceived to be an important challenge in formal specification and verification of agent systems [21]. At the same time, the epistemic logics should be able to express the property that although the agent’s knowledge is not closed under logical consequence, the agent may know some inference rules and may be able to derive consequences from its beliefs eventually. In modal syntactic epistemic logics developed in this paper, we can express properties such as ‘if the agent knows exactly the set of formulas T now, then after n steps it may know T closed under n applications of modus ponens, *and nothing else*’.

The paper is organised as follows. In section 2, we introduce the language of modal syntactic epistemic logic and corresponding transition structures. We briefly consider the logic without the ∇ operator, and show that it is axiomatisable by just the axioms of the modal system K . We consider the properties of accessibility relation which we may find useful to model time, for example seriality and determinism, and show that adding the corresponding modal axioms causes the logic to be complete with respect to the expected kind of structures. Then we introduce the ‘at most’ operator ∇ , show that it is not definable using ‘at least’, and provide a complete and sound axiomatisation for arbitrary transition systems and for transition systems satisfying various restrictions on the accessibility relation, such as seriality, functionality, and transitivity. We show that the complexity of the

^{*}Department of Informatics, University of Bergen PB. 7800, N-5020 Bergen, Norway. Email agotnes@ii.uib.no

[†]School of Computer Science and IT, University of Nottingham, Nottingham, UK. Email nza@cs.nott.ac.uk

model-checking problem and satisfiability-checking problem for the basic modal logic \mathbf{K} with ∇ is the same as for \mathbf{K} . In section 4 we introduce a new kind of conditions on accessibility relation, corresponding to the agents knowing some inference rules and/or being able to communicate with each other. We prove several completeness results for logics with rules. Finally, we compare our approach to some related work on non-omniscient reasoning (providing semantics for Duc’s logic [8] and proving its soundness and completeness), and conclude.

2 Syntactic Knowledge

We assume that an agent can have different internal states at different times, and that in each state we can identify a finite set of formulae the agent believes — for example the formulae stored in its knowledge base. Apart from finiteness, no restrictions, e.g. consistency or closure under consequence, on this set is assumed. Furthermore we assume that the agent can *act* in order to change its current state, that he may have several alternative ways to act available, and that the available actions are a function of the current state. Such a transition system between states with syntactically ascribed beliefs can be seen as a general model of reasoning, including non monotonic reasoning, belief revision, etc. In later sections we generalise this model to the multi agent case, where agents can not only reason to change their own beliefs but also communicate to change other agents beliefs.

Formally, we assume that the agent represents its beliefs in an arbitrary object language OL . Let $\Theta(OL) = \{B\phi : \phi \in OL\}$ (or just Θ when OL is clear from context) be the set of *epistemic atoms*. A model is a relational structure with a valuation of the epistemic atoms in each state, with the restriction that only finitely many epistemic atoms can be true in each state. We write $\wp(X)$ for the powerset of a set X , and $\wp^{fin}(X)$ for the set of all finite subsets of X .

Definition 1 *A model is a tuple $M = (W, R, V)$, where W is a non-empty set of states, R a binary relation over W , and V a function $V : W \rightarrow \wp^{fin}(\Theta)$. The class of all models is denoted $\mathcal{M}(OL)$ (or just \mathcal{M}).*

A *general model* is a model without the finiteness condition:

Definition 2 *A general model is a tuple $M = (W, R, V)$ where W is a non-empty set of states, R a binary relation over W , and V a function $V : W \rightarrow \wp(\Theta)$. The class of all general models is denoted $\mathcal{M}^{gen}(OL)$ (or just \mathcal{M}^{gen}).*

We are not interested in general models as such, but they are often useful as an intermediate stage in constructing a proper model.

Often, it is convenient to be able to refer to the set of formulas an agent believes in a given state. We write $\overline{V}(w)$ for the set $\{\phi : B\phi \in V(w)\}$, and call $\overline{V}(w)$ the agent’s *epistemic state* in state w .

2.1 Knowing At Least

\mathcal{M} can be used to interpret the language of propositional modal logic, with epistemic atoms as primitive propositions, in the usual way.

The language $\mathcal{L}(OL)$ (or just \mathcal{L}) is the least language such that:

- $\Theta(OL)$ are formulae
- If ϕ, ψ are formulae, then $\phi \wedge \psi$ is a formula
- If ϕ is a formula, then $\neg\phi$ is a formula
- If ϕ is a formula, then $\diamond\phi$ is a formula

For the sake of brevity, we do not introduce non-epistemic primitive propositions, however their introduction would not require any non-trivial changes to the subsequent results.

We use the usual derived propositional connectives, in addition to $\Box\phi$ for $\neg\Diamond\neg\phi$. $B\phi$ means that the agent knows *at least* ϕ — he knows ϕ and may or may not know something else. When X is a finite set of object formulae, we write ΔX — read “the agent knows at least X ” — as shorthand for $\bigwedge_{\alpha \in X} B\alpha$, formally:

$$\Delta\{\alpha_1, \dots, \alpha_n\} \stackrel{df}{=} B\alpha_1 \wedge \dots \wedge B\alpha_n.$$

If X is a singleton set, we often omit brackets and write $\Delta\phi$ for $\Delta\{\phi\}$.

The interpretation of \mathcal{L} in \mathcal{M} models is defined as usual in modal logic. When $M = (W, R, V)$ and $w \in W$:

$$\begin{aligned} M, w \models B\alpha &\Leftrightarrow \alpha \in \overline{V(w)} \\ M, w \models \Diamond\phi &\Leftrightarrow \exists_{(w, w') \in R} M, w' \models \phi \\ M, w \models \neg\phi &\Leftrightarrow M, w \not\models \phi \\ M, w \models \phi \wedge \psi &\Leftrightarrow M, w \models \phi \text{ and } M, w \models \psi \end{aligned}$$

Some examples of formulae and informal interpretations are:

- $\Delta\{p, p \rightarrow q\} \rightarrow \Diamond \Delta q$: the agent can reason with modus ponens from $p, p \rightarrow q$
- $\Diamond \Delta q \rightarrow \Delta\{p, p \rightarrow q\}$: the agent can only infer q by using modus ponens from the premises p and $p \rightarrow q$
- $\Delta p \rightarrow \Diamond \neg \Delta p$: the agent may forget p in the next state
- $\Delta p \rightarrow \Diamond \Delta p$: the agent may remember p in the next state
- $\Delta p \rightarrow \Box \Delta p$: the agent must remember p in the next state

The class of models \mathcal{M} is completely axiomatised by the modal logic \mathbf{K} (in the language \mathcal{L}). The logical system \mathbf{K} consists of the axiom schemas

(Prop) ϕ , when ϕ is a substitution instance of a propositional tautology

(K) $\Box(\phi \rightarrow \psi) \rightarrow (\Box\phi \rightarrow \Box\psi)$

and the rules

(Modus Ponens) From $\phi, \phi \rightarrow \psi$ prove ψ

(Gen) From ϕ , prove $\Box\phi$

Theorem 1 \mathbf{K} is sound and weakly complete with respect to \mathcal{M} .

Proof. It suffices to show that any \mathbf{K} -consistent formula ϕ is satisfied in \mathcal{M} . Let $M^{\mathbf{K}} = (W^{\mathbf{K}}, R^{\mathbf{K}}, V^{\mathbf{K}})$ be the canonical model for \mathbf{K} (this is a general model, and not necessarily a model in \mathcal{M} , since it may be that $V(w)$ is infinite for some state w). ϕ is true in at least one of the states in $M^{\mathbf{K}}$. Let $M^f = (W^f, R^f, V^f)$ be as follows: $W^f = W^{\mathbf{K}}$, $R^f = R^{\mathbf{K}}$ and for every $w \in W^f$, $V^f(w) = V^{\mathbf{K}}(w) \cap At(\phi)$, where $At(\phi)$ is the set of epistemic atoms occurring in ϕ . Clearly, for every world w , $M^{\mathbf{K}}, w \models \phi$ iff $M^f, w \models \phi$. $V^f(w)$ is finite for each w , since there are only finitely many epistemic atoms in ϕ . Thus, $M^f \in \mathcal{M}$ and ϕ is satisfied in M^f . \square

The semantics we have given for the language \mathcal{L} is not compact. A counter example to compactness is the theory $\Gamma = \{B\phi : \phi \in OL\}$. Γ is not satisfiable in \mathcal{M} , but each of its finite subsets is. Thus, \mathbf{K} is not *strongly* complete wrt. \mathcal{M} .

2.2 Useful axioms

In this section, we consider imposing additional conditions on the accessibility relation in the models, to make it easier to compare our logic to other approaches.

2.2.1 Unbounded Reasoning

Many syntactic approaches to epistemic logic [10, 5, 2] are based on the view that reasoning does not have an end point, but goes on indefinitely. This explains logical non-omniscience without sacrificing rationality: an agent can eventually get to know *any* particular fact it is able to compute, but can never get to know all of them at the same time. In the models \mathcal{M} , the assumption that an agent should be able to do any reasoning at all in a given state of the system is not made. In this section, we restrict the logic by adding this assumption.

Semantically, it corresponds to requiring that the accessibility relation is serial. A *serial* model is a model (W, R, V) where the accessibility relation is serial, i.e. where for each world $w \in W$ there exists a $u \in W$ such that Rwu . The class of all serial models is denoted \mathcal{M}^s .

Proof-theoretically, the assumption of unbounded reasoning corresponds to adding the axiom schema

$$(D) \quad \Box\phi \rightarrow \Diamond\phi$$

The modal system **KD** is **K** extended with the D axiom.

Theorem 2 *KD is sound and complete with respect to \mathcal{M}^s .*

Proof. Like the proof of Theorem 1, with the canonical model for **KD**. \square

2.2.2 Deterministic Reasoning

The models \mathcal{M} are models of nondeterministic reasoning, in the sense that an agent may have several possible transitions from one state. In this section we explore the special case when reasoning is deterministic, i.e. when there is at most (or *exactly*, in the case of unbounded reasoning)) one possible next state for each state. Formally, a *deterministic model* is one in which the accessibility relation is a (partial) function. The set of all deterministic models is

$$\mathcal{M}^d = \{(W, R, V) \in \mathcal{M} : Rwu \text{ and } Rvw \Rightarrow u = v\}$$

and the class of all deterministic serial models is $\mathcal{M}^{ds} = \mathcal{M}^d \cap \mathcal{M}^s$.

Proof-theoretically, we add the axiom schema

$$(F) \quad \Diamond\phi \rightarrow \Box\phi$$

The modal systems **KF** and **KDF** are **K** and **KD** extended with the F axiom, respectively.

Theorem 3 *KF is sound and complete with respect to \mathcal{M}^d .*

Proof. The axiom *F* is valid in \mathcal{M}^d , which suffices to show soundness. For completeness, let ϕ be a **KF** consistent \mathcal{L} formula. By the standard modal logic result, there is a general model M with a deterministic accessibility relation, where ϕ is satisfied. M can be transformed into a deterministic model M^f , where all epistemic states are finite, by setting $V^f(w) = V(w) \cap At(\phi)$ for every w . M^f satisfies ϕ and belongs to \mathcal{M}^d . \square

Theorem 4 *KDF is sound and complete with respect to \mathcal{M}^{ds} .*

Proof. Axioms *F* and *D* are valid in \mathcal{M}_n^{ds} , which gives soundness. The proof of completeness is analogous to the proof of Theorem 3. \square

2.3 Knowing At Most

In the previous section we presented different axiomatisations of “knowing at least” a finite set of formulae in basic modal logic. The Δ operator was defined by the B operator. In this section, we add a dual operator ∇ from [3] to the language. The intended meaning of ∇X is “the agent knows at most the finite set X ”.

The language $\mathcal{L}^\nabla(OL)$ (or just \mathcal{L}^∇) is the language $\mathcal{L}(OL)$ with the following additional clause in its definition: ∇X is a formula for every finite set $X \in \wp^{fin}(OL)$ of object formulae. We use the same derived connectives as in $\mathcal{L}(OL)$, in addition to $\boxtimes X$ for $\Delta X \wedge \nabla X$. The intended meaning of $\boxtimes X$ is “the agent knows exactly the finite set X ”. An example of an \mathcal{L}^∇ formula is:

$$\nabla(\{\phi, \neg\phi\} \cup X) \rightarrow \Box(\nabla(\{\phi\} \cup X) \vee \nabla(\{\neg\phi\} \cup X))$$

(which is true if the agent resolves all inconsistencies in the next state), or

$$\boxtimes(\{\phi, \phi \rightarrow \psi\} \cup X) \rightarrow \Diamond \boxtimes(\{\phi, \phi \rightarrow \psi, \psi\} \cup X)$$

(which is true if the agent can apply the rule modus ponens).

The interpretation of the language $\mathcal{L}^\nabla(OL)$ in $\mathcal{M}(OL)$ is defined as for the language \mathcal{L} , with the following definition for the new clause:

$$M, w \models \nabla X \Leftrightarrow \overline{V(w)} \subseteq X$$

It is easy to see that

$$M, w \models \Delta X \Leftrightarrow \overline{V(w)} \supseteq X$$

$$M, w \models \boxtimes X \Leftrightarrow \overline{V(w)} = X$$

Theorem 5 *The operator ∇ is not definable in $\mathcal{L}(OL)$ when the object language is infinite.*

Proof. Let X be a finite set of object language formulae. We show that for any \mathcal{L} formula ϕ , there is a model M' and a state w such that $M', w \not\models \phi \leftrightarrow \nabla X$. Let ϕ be an \mathcal{L} formula, and let $M, w \models \phi$ where $M = (W, R, V)$ (if ϕ is unsatisfiable we are done, since every formula ∇X is satisfiable). Let $Subf(\phi)$ be the set of subformulas of ϕ , including ϕ itself, (where subformulas of the form $B\alpha$ are treated as atoms). Let $Aware(\phi) = \{\alpha : B\alpha \in Subf(\phi)\}$. Let $M' = (W, R, V')$ where $V'(w') = V(w')$ when $w' \neq w$ and $V'(w) = V(w) \cup \{B\beta\}$ where $\beta \notin Aware(\phi) \cup X$. Existence of such a β is ensured by the fact that the object language is infinite and both $Aware(\phi)$ and X are finite. It can easily be shown by structural induction over ψ , that for any $\psi \in Subf(\phi)$,

$$\forall w' \in W (M, w' \models \psi \Leftrightarrow M', w' \models \psi).$$

Thus, $M', w \models \phi$. However, because $B\beta \in V'(w)$ and $\beta \notin X$, $M', w \not\models \nabla X$. \square

We now present axiomatisations of resulting logics. The axiomatisations extend the axiomatisation of the purely epistemic fragment (i.e. without the modal logic) in [3].

2.3.1 Axiomatisation

The properties of ∇ are captured by the following axioms [3]:

(E1) $\nabla X \wedge \nabla Y \rightarrow \nabla(X \cap Y)$

(E2) $\Delta X \rightarrow \neg \nabla Y$ when $X \not\subseteq Y$

(E3) $\nabla(Y \cup \{\gamma\}) \wedge \neg B\gamma \rightarrow \nabla Y$

(E4) $\nabla X \rightarrow \nabla Y$ when $X \subseteq Y$

We write \mathbf{K}_∇ for \mathbf{K} extended with E1-E4, and similarly for other systems and axioms. A logical system is defined relative to a logical language, and whenever ∇ is involved we implicitly take the language to be $\mathcal{L}^\nabla(OL)$.

Before proving weak completeness for \mathbf{K}_∇ with respect to \mathcal{M} , we are going to show an auxiliary result: that every \mathbf{K}_∇ -consistent formula is satisfied in a *general* model (that is, a model with possibly infinite epistemic states). Then we will show that each general model for a formula ϕ can be transformed into a model for ϕ by an ‘epistemic filtration’ technique.

The general model $M^c = (W^c, R^c, V^c)$ we construct is “almost” a canonical model for \mathbf{K}_∇ : W^c is the set of $\mathcal{L}^\nabla(OL)$ -maximal \mathbf{K}_∇ -consistent sets, R^c is the canonical relation, but V^c is *not* the canonical valuation. Particularly, M^c is constructed over an extended object language $OL \cup \{\hat{p}\}$, where \hat{p} is some primitive proposition not in the original object language OL . Thus, $V^c : W^c \rightarrow \wp(\Theta(OL \cup \{\hat{p}\}))$, and $M^c \in \mathcal{M}^{gen}(OL \cup \{\hat{p}\})$ (the epistemic filtration technique transforms M^c to a model $M \in \mathcal{M}(OL)$). The results extend those of [4], who use an auxiliary proposition \hat{p} to get satisfiability in the non modal case.

A notation which will be useful is $Bel^w = \{\alpha : B\alpha \in w\}$ for a set $w \subseteq \mathcal{L}^\nabla(OL)$.

Definition 3 Let $\hat{p} \notin OL$. Define general canonical model $M^c = (W^c, R^c, V^c) \in \mathcal{M}^{gen}(OL \cup \{\hat{p}\})$ as follows:

- W^c is the set of all $\mathcal{L}^\nabla(OL)$ -maximal \mathbf{K}_∇ -consistent subsets of $\mathcal{L}^\nabla(OL)$
- $(w, v) \in R^c$ iff $\forall \phi \in \mathcal{L}^\nabla(OL) (\phi \in v \Rightarrow \diamond\phi \in w)$
- $V^c : W^c \rightarrow \wp(\Theta(OL \cup \{\hat{p}\}))$:

$$B\alpha \in V^c(w) \Leftrightarrow B\alpha \in w \text{ or } (\alpha = \hat{p}, Bel^w \text{ is finite, and } \forall X \in \wp^{fin}(OL) \nabla X \notin w)$$

Theorem 6 Every \mathbf{K}_∇ -consistent set of $\mathcal{L}^\nabla(OL)$ formulas is satisfied in the general canonical model M^c defined in Definition 3.

Before we prove Theorem 6, we need the following intermediate result.

Lemma 1 Let w be a $\mathcal{L}^\nabla(OL)$ -maximal \mathbf{K}_∇ -consistent set. If there is an X' such that $\nabla X' \in w$, then for every $X \in \wp^{fin}(OL)$

$$\nabla X \in w \Leftrightarrow Bel^w \subseteq X$$

Proof. Let $\nabla X' \in w$, and let $U = \cap\{Y : \nabla Y \in w\}$. For the direction to the right, let $\nabla X \in w$. By (a finite number of applications of) **E1**, $\nabla U \in w$. By **E2**, it must be the case that $Bel^w \subseteq U$, and thus $Bel^w \subseteq X$. For the direction to the left, let $Bel^w \subseteq X$. Since $\nabla X' \in w$, $Bel^w \not\subseteq X'$ by **E2**. If $X' \subseteq X$, then $\nabla X \in w$ by **E4** and we are done, so assume that $X' \not\subseteq X$. Let $X' \setminus X = \{\alpha_1, \dots, \alpha_k\}$ (X', X are finite). $\alpha_i \notin Bel^w$, and by maximality $\neg\alpha_i \in Bel^w$ for each $i \in [1, k]$. By (k applications of) **E3**, $\nabla X \in w$. \square

The Lindenbaum lemma which states that we can extend every \mathbf{K}_∇ -consistent set to a \mathbf{K}_∇ -consistent and $\mathcal{L}^\nabla(OL)$ -maximal set, holds by the usual proof and will be used without explicit reference.

The existence lemma also has a standard proof:

Lemma 2 (Existence Lemma for M^c) For any $w \in W^c$, if $\diamond\phi \in w$ there is a $v \in W^c$ such that $(w, v) \in R^c$ and $\phi \in v$.

The final piece we need for the proof of Theorem 6 is the truth lemma:

Lemma 3 (Truth Lemma for M^c) For each $w \in W^c$ and $\phi \in \mathcal{L}^\nabla(OL)$

$$M^c, w \models \phi \Leftrightarrow \phi \in w$$

Proof. The proof is by induction on the length of ϕ .

$\phi = B\alpha$: For the direction to the right, let $B\alpha \in V^c(w)$. Since $\phi \in \mathcal{L}^\nabla(OL)$, $\alpha \neq \hat{p}$; thus $B\alpha \in w$. For the direction to the left, if $B\alpha \in w$ then $B\alpha \in V^c(w)$.

$\phi = \nabla X$: We have two cases. First, assume that for some X' , $\nabla X' \in w$. Then $\overline{V(w)} = Bel^w$, and thus $\overline{V^c(w)} \subseteq X$ iff $Bel^w \subseteq X$ iff, by Lemma 1, $\nabla X \in w$. Second, assume that $\forall X', \nabla X' \notin w$, in which case $B\hat{p} \in V^c(w)$. We must show that $M^c, w \not\models \nabla X$. Since $X \in \wp^{fin}(OL)$ ($\phi \in \mathcal{L}^\nabla(OL)$), $\hat{p} \notin X$ and thus $\overline{V(w)} \not\subseteq X$ which is the same as $M^c, w \not\models \nabla X$.

$\phi = \diamond\psi$: For the direction to the right, assume that there is a $(w, w') \in R^c$ such that for any $\chi \in \mathcal{L}^\nabla(OL)$ ($\chi \in w' \Rightarrow \diamond\chi \in w$) and $M^c, w' \models \psi$. By the induction hypothesis, $\psi \in w'$, and thus $\diamond\psi \in w$. For the direction to the left, let $\diamond\psi \in w$. By the existence lemma, there is a $(w, w') \in R^c$ such that $\psi \in w'$, and by the induction hypothesis $M^c, w' \models \psi$. Thus, $M, w \models \diamond\psi$.

Negation and conjunction: as usual. □

Theorem 6 follows from Lemma 3 and the standard Lindenbaum argument.

The next theorem shows how to produce a proper model given a general model, using *epistemic filtration*.

Theorem 7 *If a formula $\phi \in \mathcal{L}^\nabla(OL)$ is satisfiable in a general model $M = (W, R, V) \in \mathcal{M}^{gen}(OL \cup \{\hat{p}\})$, then it is satisfiable in a model $M^f = (W^f, R^f, V^f) \in \mathcal{M}(OL)$, where $W^f = W$ and $R^f = R$.*

Proof. We are going to build a model M^f with the same W and R but possibly different V^f , such that for every state w in M and for every subformula ψ of ϕ ,

$$M, w \models \psi \Leftrightarrow M^f, w \models \psi$$

Let $Subf(\phi)$ be the set of subformulas of ϕ , including ϕ itself, (where subformulas of the form $B\alpha$ and ∇X are treated as atoms). Let $Aware(\phi) = \{\alpha : B\alpha \in Subf(\phi) \text{ or } \alpha \in X \text{ for some } X \text{ such that } \nabla X \in Subf(\phi)\}$.

An obvious attempt at a solution to modifying V would be to set $\overline{V^f(w)} = \overline{V(w)} \cap Aware(\phi)$. However, ϕ may be $\neg\nabla\{\phi_1, \dots, \phi_k\}$ and $Aware(\phi) = \{\phi_1, \dots, \phi_k\}$, in which case ϕ would become false in M^f . This suggests adding an extra formula to each epistemic state. However, in other cases (as when for example ϕ is $\nabla\{\phi_1, \dots, \phi_k\}$), adding an extra formula will make ϕ false. So we have to modify the epistemic states on a case-by-case basis. For each w , let $\overline{V^f(w)} = \overline{V(w)} \cap Aware(\phi)$ if for some subformula of ϕ of the form ∇X is true at M, w , and $(\overline{V(w)} \cap Aware(\phi)) \cup \{\alpha\}$ otherwise, where α is a new formula not occurring in ϕ : $\alpha \in OL$, $\alpha \notin Aware(\phi)$.

Let us prove that the truth value of subformulas of ϕ of the form ΔX and ∇X has not changed in M^f . We reason by cases:

- if ΔX is true at M, w , this means that $X \subseteq \overline{V(w)}$. Since ΔX is a subformula of ϕ , $X \subseteq Aware(\phi)$. So, $X \subseteq \overline{V(w)} \cap Aware(\phi)$, and ΔX is true at M^f, w .
- if ΔX is false, $X \not\subseteq \overline{V(w)}$, hence $X \not\subseteq \overline{V(w)} \cap Aware(\phi)$. Since $\alpha \notin X$, $X \not\subseteq (\overline{V(w)} \cap Aware(\phi)) \cup \{\alpha\}$. This means that $X \not\subseteq \overline{V^f(w)}$, hence the formula is still false.
- if ∇X is true, then $\overline{V^f(w)} = \overline{V(w)} \cap Aware(\phi)$. We know that $X \subseteq Aware(\phi)$, $\overline{V(w)} \subseteq X$, so $\overline{V^f(w)} \subseteq X$, and the formula is true at M^f, w .

- if ∇X is false, $\overline{V^f(w)} = (\overline{V(w)} \cap \text{Aware}(\phi)) \cup \{\alpha\}$, and $\alpha \notin X$. So $\overline{V^f(w)} \not\subseteq X$, and the formula is false at M^f, w .

Now we have defined V^f and shown that the truth values of all epistemic subformulas of ϕ are unchanged in M^f in any state. For all other subformulas of ϕ , the proof is by an easy induction.

Note that for every w , $V^f(w) \subseteq \Theta(OL)$ (the filtration removes the proposition \hat{p}) and that $V^f(w)$ is finite. \square

Theorem 8 $\mathbf{K}\nabla$ is sound and weakly complete with respect to \mathcal{M} .

Proof. Soundness follows from the easily seen fact that E1-E4 are valid — in addition to validity of \mathbf{K} and the fact that MP and Gen preserve validity.

Completeness follows from Theorem 6 and Theorem 7. \square

Before we proceed to investigate the complexity of $\mathbf{K}\nabla$, it is useful to state explicitly some obvious facts concerning the logic.

Fact 1 Formulas of $\mathbf{K}\nabla$ are preserved under bisimulation.

The proof of this fact is totally standard.

The modal degree of a $\mathbf{K}\nabla$ formula is the (greatest) depth of nesting of modal operators \diamond in the formula.

Fact 2 Formulas of $\mathbf{K}\nabla$ of modal degree less or equal to n are preserved under n -bisimulation.

The notion of n -bisimulation is defined for example in [7]; intuitively, it means that the similarity between models can be maintained not indefinitely, as in the case of bisimulation, but for at least n steps.

Also, by the standard modal filtration argument,

Fact 3 Each satisfiable formula of $\mathbf{K}\nabla$ is satisfiable in a finite model.

From the facts above, it follows that

Fact 4 Each satisfiable formula of $\mathbf{K}\nabla$ of modal degree less or equal to n has a finite tree model of depth n .

Given a finite model M and a state w satisfying a formula ϕ of modal degree n , we can unravel M into a tree model with the root w , and cut it off at depth n . In the resulting model M' , the root is n -bisimilar to w in M , hence satisfies the same formulas of depth less or equal to n .

2.4 Complexity of $\mathbf{K}\nabla$

Model-checking complexity of $\mathbf{K}\nabla$ is the same as the model-checking complexity of \mathbf{K} . Checking whether a subformula of the form ∇X is true in an epistemic state s can be done in time linear in the size of X , provided that it is possible to check in constant time whether a formula is in the epistemic state or not, and provided we keep a record of the size of each epistemic state. To verify whether ∇X is true, given an epistemic state s , set up a counter initially equal to the size of s , and check for every formula in X whether it is in s , and if it is, decrement the counter. If we checked all formulas in X and the counter is equal to 0, then the formula ∇X is true, else it is false (there are formulas in s which are not in X).

To see that the satisfiability problem for $\mathbf{K}\nabla$ is PSPACE-complete, observe first that it is PSPACE-hard since $\mathbf{K}\nabla$ includes \mathbf{K} ; on the other hand, $\mathbf{K}\nabla$ has models of polynomial depth, and the usual NPSpace algorithm can be used to guess branches of the model, which can be written using polynomial space. Since NPSpace=PSPACE, the satisfiability problem for $\mathbf{K}\nabla$ is in PSPACE.

2.5 Additional axioms

We take assumptions about unbounded and deterministic reasoning into account — now for the language \mathcal{L}^∇ . It is straightforward to show that adding the standard modal axioms such as F , D and 4 to \mathbf{KF}^∇ produces logics which are sound and weakly complete with respect to the corresponding classes of models.

Theorem 9 \mathbf{KF}^∇ is sound and weakly complete with respect to \mathcal{M}^d .

Proof. Soundness follows from the fact that the F axiom is valid in \mathcal{M}^d .

For completeness, consider a \mathbf{KF}^∇ -consistent formula ϕ . By Theorem 6, it is satisfied in a general canonical model. By the standard modal argument, the presence of the F axiom allows to transform this model into a deterministic model (F forces all accessible worlds to satisfy the same formulas, and identical worlds can be glued together, yielding a bisimilar model). By Theorem 7, ϕ has a model in \mathcal{M}^d . \square

Theorem 10 \mathbf{KD}^∇ is sound and complete with respect to \mathcal{M}^s .

Proof. Straightforward from Theorem 6 and Theorem 7. \square

Theorem 11 \mathbf{KDF}^∇ is sound and complete with respect to \mathcal{M}^{ds} .

Proof. Straightforward from Theorem 6 and Theorem 7. \square

Finally, if we would like to consider \diamond as a temporal ‘at some point in the future’ operator (which is useful in the subsequent comparison section), we need to consider transitive models. Let \mathcal{M}^t be the class of all transitive models.

Syntactically, this corresponds to the axiom schema

$$(4) \quad \diamond\diamond\phi \rightarrow \diamond\phi$$

Theorem 12 $\mathbf{K4}^\nabla$ is sound and complete with respect to \mathcal{M}^t .

Proof. Straightforward from Theorem 6 and Theorem 7. \square

3 The Multi-Agent Case

We have up to now considered only the single agent case, for simplicity. Henceforth we will be more general and prove results also for the multi-agent case where there are n agents $\Sigma = \{1, \dots, n\}$. In this case both syntax and semantics are defined over $\Theta_n = \{B_i\phi : \phi \in OL, i \in \Sigma\}$ in place of Θ .

Definition 4 A model in the multi-agent case is a tuple $M = (W, R_1, \dots, R_n, V)$ where W is a non-empty set of states, R_i a binary relation over W , and V a function $V : W \rightarrow \wp^{fin}(\Theta)$. The class of all multi-agent models is denoted \mathcal{M}_n , and similarly for other model classes (\mathcal{M}_n^s for serial models, etc.).

A tuple (s_1, \dots, s_n) of finite epistemic states ($s_i \subset OL$) is associated with each state $s \in W$: $s_i = \{\phi : B_i\phi \in V(s)\}$.

Different from the single agent case is that a state w does not now correspond to the state of a single agent, but to a global state of a system consisting of several agents, each with its own epistemic state. Intuitively, the fact that $(w, w') \in R_i$ means that agent i can change the system from state w to state w' . Thus, agent i can potentially affect the epistemic states

of other agents in addition to his own. In section 4 we use this fact to model both reasoning and communication in a multi-agent system. Note that the transitions available to agent i depends on the state of the whole system.

An example of an \mathcal{L}^{∇_2} formula is:

$$\nabla_1 p \wedge \nabla_2 q \rightarrow \diamond_1 \nabla_2 \{p, q\}$$

This formula may be true of a system where agent 1 may communicate its beliefs to agent 2, and agent 2 has no other way of acquiring new beliefs. Further examples involving the knowledge of rules are discussed in section 4.

Axiomatisation of the multi-agent case is obtained by indexing modalities and epistemic operators by agents; for example, the epistemic axioms become

$$(E1) \nabla_i X \wedge \nabla_i Y \rightarrow \nabla_i (X \cap Y)$$

$$(E2) \Delta_i X \rightarrow \neg \nabla_i Y \text{ when } X \not\subseteq Y$$

$$(E3) \nabla_i (Y \cup \{\gamma\}) \wedge \neg B_i \gamma \rightarrow \nabla_i Y$$

$$(E4) \nabla_i X \rightarrow \nabla_i Y \text{ when } X \subseteq Y$$

and we can express that agent i can always make a transition by adding an axiom schema

$$(D) \Box_i \phi \rightarrow \diamond_i \phi$$

All the preceding proofs can easily be modified to show that the results hold also for the multi-agent case.

4 Examples of complete logics which capture knowledge of inference or communication rules

In this section, we consider multi-modal logics defined by conditions on the accessibility relations, which involve epistemic states. Intuitively, in these logics, transitions between states correspond to an epistemic action performed by one of the agents. We find it useful to distinguish two kinds of epistemic actions. The first kind is an internal action by the agent, corresponding to applying an inference rule (or an internal state update rule, in general). This kind of action only affects the agent's state. The second kind of action is communicating something to other agents. This is seen as broadcast communication, and as a result not only the agent's own state, but also the states of other agents may be updated.

An internal inference rule could be, for example, modus ponens. There are a number of possibilities for expressing the fact that an agent is capable of reasoning by modus ponens. For example, we can say that an agent *can* make a transition using MP:

$$\Delta_i MP \quad \phi, \phi \rightarrow \psi \in s_i \Rightarrow \exists s' (R_i(s, s') \wedge \psi \in s'_i)$$

Another possibility is requiring that MP is *the only* way to update the agent's state, or that the agent knows *at most* MP:

$$\nabla_i MP \quad R_i(s, s') \Rightarrow \exists \phi, \phi \rightarrow \psi \in s_i (s'_i = s_i \cup \{\psi\} \wedge \bigwedge_{j \neq i} s'_j = s_j)$$

A communication rule could, for example, allow agent i to tell agent j one of the formulas it believes:

$$\Delta_i Comm_{ij} \quad \phi \in s_i \Rightarrow \exists s' (R_i(s, s') \wedge \phi \in s'_j)$$

As above, we can also require that this is the only epistemic action agent i could perform:

$$\nabla_i Comm_{ij} \quad R_i(s, s') \Rightarrow \exists \phi \in s_i (s'_j = s_j \cup \{\phi\} \wedge \bigwedge_{k \neq j} s'_k = s_k)$$

Any number of similar conditions on the accessibility relations can be considered.

In this section, we show some completeness results which are useful for comparisons with related work. One concerns axiomatising monotonic reasoners:

$$Mon_i \quad \forall j (R_j(s, s') \Rightarrow s_i \subseteq s'_i)$$

This condition says that after any transition, agent i does not loose any beliefs. Syntactically, this corresponds to the following axiom schema :

$$\mathbf{M}_i \quad (\Delta_i X \rightarrow \Box_j \Delta_i X) \wedge (\Diamond_j \nabla_i X \rightarrow \nabla_i X) \text{ (for all } j \text{)}.$$

Theorem 13 $\mathbf{KM}_i \nabla_n$ is sound and weakly complete with respect to models satisfying Mon_i .

Proof. Soundness is obvious; for completeness, consider a general model M^c as in Definition 3, where \mathbf{M}_i holds in all worlds. This forces epistemic states of i in the general canonical model satisfy the inclusion Mon_i : whenever s' is reachable from s , $s_i \subseteq s'_i$ (first conjunct of \mathbf{M}_i). Consider a $\mathbf{KM}_i \nabla_n$ -consistent formula ϕ , and define $Aware_i(\phi)$ as $\{\psi : \psi \in X \text{ for some } X \text{ with } \Delta_i X \text{ or } \nabla_i X \in Subf(\phi)\}$. Now consider an epistemic filtration M^f of M^c as described in the proof of Theorem 7. Let us denote the epistemic state of i in state s as s_i^f in the new model. For all states s where s_i^f is defined as $s_i \cap Aware_i(\phi)$, $R_j^f(s^f, s'^f)$ implies $s_i^f \subseteq s'^f_i$, because $s_i \subseteq s'_i$. Suppose s_i^f is defined as $(s_i \cap Aware_i(\phi)) \cup \{\alpha_i\}$, that is, a special formula α_i is added to s_i^f . Mon_i would be violated if for some s' , $R_j^f(s^f, s'^f)$ and $\alpha_i \notin s'^f_i$. This is only possible if for some $\nabla_i X$, $\nabla_i X \in s'$ and $\nabla_i X \notin s$. But this is forbidden by the second conjunct of \mathbf{M}_i . Hence, ϕ has a model where Mon_i holds. \square

Note that we also can axiomatise Mon_i in the language without ∇ - then we only need the first conjunct of \mathbf{M}_i .

Now consider the following natural class of conditions on the accessibility relation, which we will call *addition conditions*. Those conditions correspond to ‘knowing at least’ rules of the following form: if agent i believes formulas ϕ_1, \dots, ϕ_n , then agent i can reach a state where agent j (possibly $j = i$) believes formulas ψ_1, \dots, ψ_k :

$$\phi_1, \dots, \phi_n \in s_i \Rightarrow \exists s' (R_i(s, s') \wedge \psi_1, \dots, \psi_k \in s'_j)$$

Examples of such addition conditions are $\Delta_i MP$ and $\Delta_i Comm_{ij}$ above.

Theorem 14 Any set of addition conditions of the form

$$\phi_1, \dots, \phi_n \in s_i \Rightarrow \exists s' (R_i(s, s') \wedge \psi_1, \dots, \psi_k \in s'_j)$$

is axiomatisable by adding to \mathbf{K}_n axiom schemas of the form

$$\Delta_i \{\phi_1, \dots, \phi_n\} \rightarrow \Diamond_i \Delta_j \{\psi_1, \dots, \psi_k\}$$

Proof. Soundness is straightforward.

For completeness, consider a general canonical model where the axioms hold. In the general canonical model, if $\Delta_i \{\phi_1, \dots, \phi_n\} \in s$, there is an R_i -accessible state s' with $\Delta_j \{\psi_1, \dots, \psi_k\} \in s'$, and the addition condition holds.

Now we need to produce a model for a consistent formula ϕ with finite epistemic states, where the semantic condition still holds. We modify the proof of Theorem 1 as follows. Take a world s which satisfies ϕ , and unravel the model M so that s is the root of the model. Now intersect the epistemic state of all agents j in s with $Aware_j(\phi)$ as before; however for states s' reachable from s in k steps, intersect s'_j with $Aware_j(\phi)$ closed under k applications of the addition condition. For example, if $i = j$ and the condition is $\Delta_i MP$, then $Aware_i(\phi)$ would be closed under k applications of modus ponens. \square

Theorem 15 \mathbf{K}_n together with an axiom schema

$$\mathbf{A1} \quad \Delta_1\{\phi, \phi \rightarrow \psi\} \rightarrow \Diamond_1 \Delta_1 \{\psi\}$$

is sound and weakly complete with respect to models satisfying $\Delta_1 MP$.

Proof. The theorem follows from Theorem 14. \square

It is much harder to axiomatise knowing ‘at most’ a rule. To get a feeling for the reason why, consider axiomatising knowing at most modus ponens. We want to say something like: if a formula ψ is in the agent’s state after a transition, then it has either been there before the transition, or has been added as a result of applying modus ponens to some $\phi \rightarrow \psi, \phi$ which are in the agent’s state. However, there are infinitely many formulas ϕ ; expressing this would require quantification over formulas:

$$\Diamond_i \Delta_i \psi \rightarrow \Delta_i \psi \vee \exists \phi (\Delta_i \{\phi, \phi \rightarrow \psi\})$$

A partial solution - which still requires only the language with ‘at least’ only - is to restrict our attention to rules which have ‘subformula property’, such as the rule of conjunction introduction:

$$\nabla_i \wedge_{int} R_i(s, s') \Rightarrow \exists \phi, \psi \in s_i (s'_i = s_i \cup \{\phi \wedge \psi\} \wedge \bigwedge_{j \neq i} s'_j = s_j)$$

Theorem 16 The class of models satisfying $\nabla_i \wedge_{int}$ is axiomatisable by adding to \mathbf{K}_n the following axiom schemas:

$$\mathbf{C1} \quad \Diamond_i \Delta_i (\phi \wedge \psi) \rightarrow [\Delta_i(\phi \wedge \psi) \vee (\Delta_i \phi \wedge \Delta_i \psi)]$$

$$\mathbf{C2} \quad \Diamond_i \Delta_i X \rightarrow \Delta_i X \vee \bigvee_{\phi \wedge \psi \in X} \Delta_i \{X \setminus \{\phi \wedge \psi\}\}$$

$$\mathbf{C3} \quad \Diamond_i \Delta_j X \rightarrow \Delta_j X \text{ where } j \neq i$$

$$\mathbf{C4} \quad \Delta_j X \rightarrow \Box_i \Delta_j X$$

Proof. Note that **C1** corresponds specifically to knowing the rule of conjunction introduction, whereas **C2 – C4** in some form are necessary for any ‘knowing at most’ a rule and state that at most one formula is added after any transition, and the states of other agents are not changed after the transition.

Soundness is straightforward. For completeness, consider a formula ϕ consistent with the axioms. Build a general canonical model for ϕ , with potentially infinite epistemic states. The axioms **C3** and **C4** will force the epistemic states of agents other than i stay unchanged along the accessibility relation R_i . **C2** makes sure that if $R_i(s, s')$, then at most one new formula is in s'_i compared to s_i , and this formula is of the form $\psi_1 \wedge \psi_2$. Finally, **C1** makes sure that if a formula of the form $\psi_1 \wedge \psi_2$ is added to s'_i , then $\psi_1, \psi_2 \in s_i$.

Using this general canonical model, we build a model for ϕ with finite epistemic states as before, intersecting each epistemic state with $Aware_i(\phi)$. This means that the resulting model contains only finitely many non-trivial transitions for conjunction introduction, but this does not violate the $\nabla_i \wedge_{int}$ condition. \square

Completeness results for ‘knowing at most’ other inference rules with subformula property can be proved analogously.

In general, however, we can axiomatise ‘knowing at most’ a rule, when the rule does not have subformula property, using ‘at most’ operator rather than quantification over formulas.

Theorem 17 The class of models satisfying $\nabla_i MP$ is axiomatisable by adding to \mathbf{K}_{∇_n} the following axiom schemas:

$$\mathbf{B1} \quad \nabla_i X \rightarrow \Box_i \bigvee_{\phi, \phi \rightarrow \psi \in X} \nabla_i (X \cup \{\psi\})$$

B2 $\diamond_i \Delta_i X \rightarrow \Delta_i X \vee \bigvee_{\psi \in X} \Delta_i(X \setminus \{\psi\})$

C3 $\diamond_i \Delta_j X \rightarrow \Delta_j X$ where $j \neq i$

C4 $\Delta_j X \rightarrow \Box_i \Delta_j X$

C5 $\diamond_i \nabla_j X \rightarrow \nabla_j X$

Proof. For soundness, observe that B1 states that agent i 's epistemic state grows by at most one formula derived by MP at each step. B2 states that if agent i believes a set of formulas X in the next state, then it either already believes all formulas in X in the current state, or it believes exactly one formula less. C3 as before says that agents other than i do not acquire new beliefs along R_i , and C4 and C5 assert that all agents' epistemic states grow monotonically along R_i .

For completeness, consider a general model M^c as in the Definition 3. We need to show that the model satisfies condition $\nabla_i MP$. Assume that $R_i(s, s')$ holds. Axioms C4 and C5 make sure that for all $j \neq i$, $s'_j = s_j$. We need to show that $s'_i = s_i \cup \{\psi\}$ for some $\phi, \phi \rightarrow \psi \in s'_i$. Axiom C4 guarantees that $s_i \subseteq s'_i$. Axiom B2 makes sure that there is at most one extra formula ψ in s'_i compared to s_i ; now we know that $s'_i = s_i$ or $s'_i = s_i \cup \{\psi\}$ for some $\psi \notin s_i$.

If for some X , $\nabla_i X \in s$, then by axiom B1, in all R_i -accessible worlds the epistemic state of i can have at most one extra ψ , and it has to be derived by MP. Note that $\nabla_i X \rightarrow \Box_i \perp$ follows from B1 if there are no formulas of the form $\phi, \phi \rightarrow \psi \in X$; in this case there are no R_i accessible worlds.

The world s may however contain no formula of the form $\nabla_i X$. In this case, we cannot guarantee that if a formula was added, then it was added by MP. We fix those worlds, together with the finiteness requirement for epistemic states, by epistemic filtration through $Aware(\phi)$, closed under applications of MP. Now if a world s does not contain $\nabla_i X$, s_i contains a special formula α_i ; in the accessible world s' , s'_i also contains α_i . If there are no $\phi, \phi \rightarrow \psi$ in s_i such that ψ is the formula justifying the R_i transition, we add $\alpha_i \rightarrow \psi$ to s_i . \square

Note that the notions of “knowing at least” and “knowing at most” a rule we have discussed in this section are not the only possible ones. For example, interpretations of knowing modus ponens different from the ones used in $\Delta_i MP$ or $\nabla_i MP$ can be proposed. In Section 5.3 we briefly discuss other versions.

5 Comparisons with related work

Much of the work on syntactic knowledge has been motivated by the logical omniscience problem [16]; see e.g. [19, 20, 14] for surveys.

Syntactic approaches to epistemic logic can be classified into logics about *static* knowledge, i.e. about knowledge at a point in time, or about the *dynamics* of knowledge, i.e. about how knowledge evolves over time. Another dimension for classification is whether it is assumed that an agent can only know finitely many formulae at the same time, or whether he may know infinitely many. The approach in the current paper is a dynamic logic of finite knowledge, and in this section we compare it to similar approaches from the literature.

Classical syntactic approaches to knowledge [9, 18, 15, 14] generally model static knowledge with no assumption about finiteness. The *logic of general awareness* [13] combines a syntactic and an epistemic approach to static knowledge, and can be used to model finite knowledge. [4] and [3] investigate a logic of static syntactic knowledge without and with, respectively, the finiteness assumption.

Among dynamic approaches, the *deduction model* [17] assumes that agents' knowledge is closed under, possibly incomplete, deduction rules, and knowledge is not required to be

finite. Another very general dynamic approach is *active logics* (formerly *step logics*) [10]. An active logic consists of a formal language and inference rules and models the evolution of a belief set as the rules are applied. Each step in a derivation is assigned a moment of time. Active logics can be seen as describing transition systems over discrete linear time, but as far as we are aware this has not been explicitly stated or used by the authors. *Timed reasoning logics (TRL)* [5] are in a sense a version of active logics which makes this intuition precise, and where the applicability of inference rules always depends only on the current belief set (in active logics, it may depend on the derivation history). We compare our logic to TRL rather than to active logics because of greater similarities between the two.

5.1 Timed Reasoning Logic

Timed Reasoning Logic (TRL) was introduced in [5]. TRL is a family of logics parametrised by a set of agents A and a rule system (set of inference rules and associated rule application strategy) for each agent.

Each agent $i \in A$ has a local state which is a finite set of formulas in some logical language (propositional, predicate, modal, etc.). Different agents may use different languages at different points in time: at time t , agent i speaks the language \mathcal{L}_t^i . The local state of agent i at time t , m_t^i , is a finite set $\{\phi_1, \dots, \phi_n\}$ of formulas of the agent's language at time t , \mathcal{L}_t^i . This set may be empty or inconsistent.

Each agent has some rules to produce a new state given its current state and any new beliefs obtained by observation. Each model is equipped with a function obs , which takes a step t and an agent i as arguments and returns a finite set of formulas in the agent's language at that step (observed facts). This set is added to the agent's state at the same step (observations are instantaneous). Each agent has an associated function inf_i , which maps a finite set of formulas in the language \mathcal{L}_t^i to another finite set of formulas in the language \mathcal{L}_{t+1}^i (agent's computation of the next state).

Definition 5 Let A be a set of agents and $\{\mathcal{L}_t^i : i \in A, t \in \mathbb{N}\}$ a set of agent languages. A TRL model M is a tuple $\langle obs, inf_i, \{m_t^i : i \in A, t \in \mathbb{N}\} \rangle$ where obs is a function which maps a pair (i, t) to a finite set of formulas in \mathcal{L}_t^i , inf_i is a function from finite sets of formulas in \mathcal{L}_t^i to finite sets of formulas in \mathcal{L}_{t+1}^i , and each m_t^i is a finite set of formulas in \mathcal{L}_t^i such that $m_{t+1}^i = inf_i(m_t^i) \cup obs(i, t + 1)$.

It is assumed that there is a definition of a well formed formula associated with each of the agent's languages \mathcal{L}_t^i . If i is an agent, t is a moment of time, and ϕ a well-formed formula of the language \mathcal{L}_t^i , then $(i, t) : \phi$ is a well-formed labelled formula of TRL.

A labelled formula $(i, t) : \phi$ is true in a model, $M \models (i, t) : \phi$, iff $\phi \in m_t^i$ (the state indexed by (i, t) in M contains ϕ). A labelled formula $(i, t) : \phi$ is satisfiable, $\models (i, t) : \phi$, iff there exists a model M , $M \models (i, t) : \phi$.

One of the versions of TRL, TRL(STEP), models agents which at each step apply all their inference rules to all their beliefs (but not to formulas derived as a result). The name of the logic reflects its similarity with *step logic* [11].

The syntax of TRL(STEP) rules is as follows:

$$\frac{(i_1, t) : \phi_1, \dots, (i_n, t) : \phi_n}{(i, t + 1) : \psi}$$

Here, t is a universally quantified variable over time points, and i_1, \dots, i_n, i are fixed labels corresponding to names of agents.

Let R be a set of TRL(STEP) inference rules. A labelled formula $(i, t) : \phi$ is *derivable* from a set of labelled formulas Γ using R ($\Gamma \vdash_R (i, t) : \phi$) if there is a sequence of labelled formulas $(i_1, t_1) : \phi_1, \dots, (i_n, t_n) : \phi_n$ such that:

1. each formula in the sequence is either a member of Γ , or is obtained from Γ by one of the inference rules in L ; and
2. the last labelled formula in the sequence is $(i, t) : \phi$, namely $(i_n, t_n) : \phi_n = (i, t) : \phi$.

There are two kinds of TRL(STEP) rules. The first kind of rule involves just one agent and corresponds to this agent's internal inference rules (inf_i function). These rules are called *internal rules*.

The second kind of rule involves several agents and corresponds to exchange of information between agents, which is modelled using the obs function. These rules are called *communication rules*. Communication rules have the form:

$$\frac{(i, t) : \phi}{(j, t + 1) : \psi}$$

A TRL model M conforms to a set of TRL(STEP) rules R if

1. For every internal rule of the form

$$\frac{(i, t) : \phi_1, \dots, (i, t) : \phi_n}{(i, t + 1) : \psi}$$

inf_i in M satisfies the property

$$\phi_1, \dots, \phi_n \in m_t^i \implies \psi \in inf_i(m_t^i)$$

in other words, inf_i is computed using all, and only, the rules in R_{inf} .

2. for each communication rule of the form

$$\frac{(i, t) : \phi}{(j, t + 1) : \psi}$$

obs in M satisfies the property

$$\phi \in m_t^i \implies \psi \in obs(j, t + 1)$$

The following theorem was proved in [5]:

Theorem 18 *Given a set of TRL(STEP) rules R , for any finite set of labelled formulas Γ and a labelled formula ϕ , $\Gamma \vdash_R \phi$ iff $\Gamma \models_{\mathcal{R}} \phi$, where \mathcal{R} is the set of models conforming to R .*

There is a straightforward correspondence between a TRL model (conforming to a set of rules) and a model in \mathcal{M}^{ds} (satisfying a corresponding set of addition conditions). Indeed, a TRL model can be represented as consisting of a sequence of states

$$m_0 = \langle m_0^1, \dots, m_0^n \rangle, m_1 = \langle m_1^1, \dots, m_1^n \rangle, \dots, m_t = \langle m_t^1, \dots, m_t^n \rangle, \dots$$

connected by a successor relation $R = R_1 = \dots = R_n$ which is deterministic and serial. Note that this class of models can be axiomatised by adding to **KDF_n** a set of axiom schemas

$$\diamond_i \phi \leftrightarrow \diamond_j \phi$$

for every pair $i, j \in A$. Then for convenience we can use an unindexed modality defined as

$$\diamond \phi \stackrel{df}{=} \diamond_1 \phi$$

For each rule of the form

$$\frac{(i, t) : \phi_1, \dots, (i, t) : \phi_n}{(j, t + 1) : \psi}$$

the TRL model conforms to, there is an addition condition

$$\phi_1, \dots, \phi_n \in s_i \Rightarrow \exists s' (R_i(s, s') \wedge \psi \in s'_j)$$

in the corresponding model in \mathcal{M}^{ds} .

Theorem 19 Any class of TRL models conforming to a set of TRL(STEP) rules \mathbf{R} can be axiomatised by adding to \mathbf{KDF}_n the set of axiom schemas

$$\mathbf{A}_{ij} \quad \diamond_i \phi \leftrightarrow \diamond_j \phi$$

and, for each rule $r \in \mathbf{R}$ of the form

$$\frac{(i, t) : \phi_1, \dots, (i, t) : \phi_n}{(j, t + 1) : \psi}$$

an axiom schema

$$\mathbf{A}_r \quad \Delta_i \{\phi_1, \dots, \phi_n\} \rightarrow \diamond_i \Delta_j \psi$$

Proof. The proof is very similar to the proof of Theorem 14, but in the completeness proof the general canonical model has deterministic serial accessibility relations R_i with the additional property that $R_i = R_j$ for all i, j . \square

5.2 Dynamic Epistemic Logic

The logic we have presented can be seen as a generalisation of Dynamic Epistemic Logic (DEL) [8]. Duc defines the logical language L_{DE} by extending the modal epistemic language \mathcal{L}_n with one modal operator $\langle F_i \rangle$ for each agent i . The intended interpretation of the formula $\langle F_i \rangle \phi$ is that ϕ is true after agent i has performed some train of thought (some sequence of reasoning steps). While the name given to the logic may suggest otherwise, DEL seems to be closer related to temporal logic than to dynamic logic: the intended interpretation of the $\langle F_i \rangle$ operator is like the interpretation of the “future” operator from temporal logic, and the logical language allows no composition of actions (there is only one modal operator for each agent).

Formally, L_{DE} is defined over a number of agents n and a set of primitive propositions Φ as follows. $\mathcal{L}_n(\Phi)$ is the modal epistemic language over n and Θ .

$$\mathcal{L}_n(\Phi) \subseteq L_{DE}$$

$$\text{If } \phi \in L_{DE}, \text{ then } \neg \phi \in L_{DE}$$

$$\text{If } \phi, \psi \in L_{DE}, \text{ then } (\phi \rightarrow \psi) \in L_{DE}$$

$$\text{If } \phi \in L_{DE}, \text{ then } \langle F_i \rangle \phi \in L_{DE}$$

The operator $[F_i]$ is defined as a dual to $\langle F_i \rangle$ in the usual way: $[F_i] \phi \equiv \neg [F_i] \neg \phi$. DEL is presented as via a logical system DES_n^+ over the language L_{DE} , intended to be a “dynamic” version of S_n^+ . In the definition of the system, a sublanguage $L_E^+ \subseteq \mathcal{L}_n(\Phi)$ of *persistent formulae* is used. L_E^+ contains all the formulae of $\mathcal{L}_n(\Phi)$ without any occurrences of the knowledge operators K_i (called *objective formulae*) and is closed under the following conditions:

$$\text{If } \phi, \psi \in L_E^+, \text{ then } (\phi \wedge \psi) \in L_E^+$$

$$\text{If } \phi, \psi \in L_E^+, \text{ then } (\phi \vee \psi) \in L_E^+$$

$$\text{If } \phi \in L_E^+, \text{ then } K_i \phi \in L_E^+$$

DES_{4n} has the following axiom schemata:

$$\begin{array}{ll}
PC1. & \phi \rightarrow (\psi \rightarrow \phi) \\
PC2. & (\phi \rightarrow (\psi \rightarrow \gamma)) \rightarrow ((\phi \rightarrow \psi) \rightarrow (\phi \rightarrow \gamma)) \\
PC3. & (\neg\psi \rightarrow \neg\phi) \rightarrow (\phi \rightarrow \psi) \\
TL1. & [F_i](\phi \rightarrow \psi) \rightarrow ([F_i]\phi \rightarrow [F_i]\psi) \\
TL2. & [F_i]\phi \rightarrow [F_i][F_i]\phi \\
DE1. & K_i\phi \wedge K_i(\phi \rightarrow \psi) \rightarrow \langle F_i \rangle K_i\psi \\
DE2. & K_i\phi \rightarrow \phi \\
DE3. & K_i\phi \rightarrow [F_i]K_i\phi, \text{ if } \phi \in L_E^+ \\
DE4. & \langle F_i \rangle K_i(\phi \rightarrow (\psi \rightarrow \phi)) \\
DE5. & \langle F_i \rangle K_i((\phi \rightarrow (\psi \rightarrow \gamma)) \rightarrow ((\phi \rightarrow \psi) \rightarrow (\phi \rightarrow \gamma))) \\
DE6. & \langle F_i \rangle K_i((\neg\psi \rightarrow \neg\phi) \rightarrow (\phi \rightarrow \psi)) \\
DE7. & \langle F_i \rangle K_i(K_i\phi \rightarrow \phi) \\
DE8. & K_i\phi \rightarrow \langle F_i \rangle K_i K_i\phi, \text{ if } \phi \in L_E^+
\end{array}$$

and the following inference rules:

$$\begin{array}{l}
R1. \quad \frac{\vdash_{DES_{4n}} \phi, \vdash_{DES_{4n}} \phi \rightarrow \psi}{\vdash_{DES_{4n}} \psi} \\
R2. \quad \frac{\vdash_{DES_{4n}} \phi}{\vdash_{DES_{4n}} [F_i]\phi}
\end{array}$$

PC1-PC3 and R1 axiomatise propositional logic, TL1-TL2 and R2 the temporal operator $\langle F_i \rangle$ and DE1-DE8 the epistemic operator K_i and the interaction between the temporal and epistemic operators.

Duc does not define a formal semantics for DEL. The main motivation behind the logic is to describe agents who are non-omniscient but nevertheless rational and non-ignorant. Examples of DES_{4n} theorems intended to illustrate this are:

$$\begin{array}{l}
K_i(\phi \wedge \psi) \rightarrow \langle F_i \rangle K_i\phi \\
(K_i\phi \wedge K_i\psi) \rightarrow \langle F_i \rangle K_i(\phi \wedge \psi) \\
K_i\neg\neg\phi \rightarrow \langle F_i \rangle K_i\phi
\end{array}$$

Our logic can be seen as a generalisation of DEL along two dimensions: the temporal and the epistemic.

In the temporal dimension, the two logics have the same set of modalities, written \diamond_i in our logic and $\langle F_i \rangle$ in DEL for each agent i . In our basic system, \diamond_i has the temporal interpretation “some next state”, while as mentioned above the intended interpretation of $\langle F_i \rangle$ is “some future state”. As we have seen in Section 2.5, we can easily extend our basic system to give the latter interpretation of \diamond_i : by using transitive frames and extending the logical systems with the 4 axiom.

In the epistemic dimension, while the agents described by DEL are not logically omniscient in the sense that there are no conditions on what an agent *must* know, there are nevertheless quite strong assumptions about the agents’ reasoning mechanisms. Particularly, an agent:

- can reason perfectly in propositional logic (DE4-DE6, R1)
- can never get to know anything false or inconsistent (and knows this) (DE2, DE7)
- always reasons monotonically. Only persistent formulae are included in this definition of monotone reasoning, however. Particularly, temporal indexicals (e.g. $K_i\langle F_i \rangle\phi$) and negative epistemics (e.g. $\neg K_i\phi$) are not persistent. Monotonicity is a quite strong assumption, and is incompatible with e.g. belief revision or “forgetting”.

- can do positive introspection (DE8)

Our logic is a general framework which does not make any of these assumptions. They do, however, define a special class of our models.

In the remainder of this section, we provide a complete and sound semantics for Duc's logic, based on the models in $\mathcal{M}(OL)$ satisfying additional conditions. However, we need to assume that \mathcal{L} contains all atoms from OL . This adds non-epistemic primitive propositions to the language; we assume that they are interpreted using V in the usual way.

In what follows, we write $K_i\phi$ for $B_i\phi$ and $\langle F_i \rangle$ for \diamond_i , to be consistent with Duc's notation.

Consider syntactic epistemic models in \mathcal{M} which satisfy the following conditions, for each agent i :

D1 R_i is reflexive

D2 R_i is transitive

D3 Monotonicity for epistemic atoms: $R_i(w, v) \Rightarrow \{K_i\phi : K_i\phi \in V(w)\} \subseteq \{K_i\phi : K_i\phi \in V(v)\}$, for $\phi \in L_E^+$

D4 Knowledge: $K_i\phi \in V(w)$ implies that ϕ true in w

D5 DE1: if $K_i\phi, K_i(\phi \rightarrow \psi) \in V(w)$ then $\exists v(R_i(w, v) \& K_i\psi \in V(v))$

D6 DE4-7: $\exists v(R_i(w, v) \& K_i\tau \in V(v))$, where $\tau \in Ax$ (Ax being the list of all relevant axiom instances)

D7 DE8: if $K_i\phi \in V(w)$ then $\exists v(R_i(w, v) \& K_iK_i\phi \in V(v))$, for $\phi \in L_E^+$.

Let us call the class of models satisfying **D1 - D7**, \mathcal{M}^{DEL} .

Theorem 20 DES_{4n} is sound and weakly complete with respect to \mathcal{M}^{DEL} .

Proof. It is easy to see that DES_{4n} is sound with respect to \mathcal{M}^{DEL} . To prove completeness, assume that ϕ is a DES_{4n} -consistent formula. We will construct a model satisfying ϕ , in two stages. First we will construct a model M for ϕ which satisfies **D1-D4**. Then we will transform this model into a model which satisfies **D5-D7** as well.

We construct M as follows. Build a canonical model M^c for ϕ . Note that M^c satisfies **D1-D7**, but it may have infinite epistemic states. Applying filtration will remedy this problem, but will cause **D5-D7** to be violated. Nevertheless, we filtrate M^c using transitive filtration through the following set $Closure(\phi)$:

- $Subf(\phi) \subseteq Closure(\phi)$
- If $\tau \in Closure(\phi)$ then $\langle F_i \rangle K_i\tau$ for each $\tau \in Ax$.
- If $K_i\phi, K_i(\phi \rightarrow \psi) \in Closure(\phi)$, then $K_i\psi \in Closure(\phi)$

We unravel the resulting model into a transitive tree to obtain the model M , where the root satisfies ϕ . Note that M satisfies **D1-D4**, but properties like **D5-D7** hold only for axioms instances etc. which are in $Closure(\phi)$.

However, it is easy to prove that adding a formula which is not in $Closure(\phi)$ to the epistemic state of any world in M is not going to change the truth value of ϕ or any of its subformulas. Adding such extra properties may violate the Knowledge and Monotonicity conditions, of course. In particular, we should only add a formula ψ to the epistemic state of any world, if ψ is true in that world.

Consider the set of formulas $\Gamma = \{\psi : K_i\psi \text{ is true at the root of } M\}$. Note that the formulas in Γ are true at the root, by DE2. By DE3, the corresponding beliefs persist in all

successor epistemic states, so Γ is true in all successors. Also, all logical consequences of Γ are true in perpetuity.

To satisfy **D5-D7**, each world w should have a set of successors which between them house all of the following formulas (while still having finite epistemic states):

$$\{K_i\psi : K_i\phi, K_i(\phi \rightarrow \psi) \in w\} \cup \{K_i\tau : \tau \in Ax\} \cup \{K_iK_i\phi : K_i\phi \in w\}$$

For every w , the set above is a subset of the set $KCons = \{K_i\phi : \phi \text{ is a tautology, or a consequence of } \Gamma, \text{ or the result of putting } K_i^n \text{ before a formula of the first two kinds}\}$.

Note that adding any formula from $KCons$ to any epistemic state will not violate the Knowledge condition.

Each world in M has countably infinitely many children, and we can use them to house the formulas from $KCons$.

Enumerate $KCons$ in any way as ϕ_1, ϕ_2, \dots . Add ϕ_1 to the epistemic states of worlds on level 1 in M , ϕ_1 and ϕ_2 to level 2, and in general, $\{\phi_1, \dots, \phi_n\}$ to the epistemic states of the worlds at level n . Call the resulting model M' . Epistemic states at each level are still finite, and any formula in $KCons$ is included in the epistemic state of some successor, Knowledge and Monotonicity are not violated, and ϕ is still satisfied at the root of M . \square

5.3 Dynamic Syntactic Epistemic Logic

Dynamic Syntactic Epistemic Logic (DSEL) [2] is a logic describing how finite syntactic epistemic states can evolve in a branching-time future, and how coalitions of agents can cooperate strategically to reach certain epistemic states. It is based on *Alternating-time Temporal Logic (ATL)* [6]. The logical language contains three fragments.

First, epistemic operators $\Delta_i X$ and $\nabla_i X$, meaning that the agent knows at least and at most, respectively, the finite set of formulae X .

The second fragment is temporal operators taken from ATL. Formulae $\langle\langle G \rangle\rangle \bigcirc \phi$, $\langle\langle G \rangle\rangle \mathcal{F}\phi$ and $\langle\langle G \rangle\rangle \square \phi$ means that group of agents G can cooperate to ensure that ϕ is true in the next state, in some future state, and in all future states, respectively, no matter what the agents outside the coalition G do.

The third fragment is *rules* and *rule operators*. The language has an explicit syntax for rules, and is parameterised by a set $V_F = \{a, b, \dots\}$ of formal variables for *formulae* and a set $V_T = \{t, u, \dots\}$ of formal variables for terms. For example,

$$R_1 = \frac{t \sqcup \{a, a \rightarrow b\}}{t \sqcup \{b\}}$$

expresses the rule *modus ponens*. RUL denotes the set of all rules. The *interpretation* of a set of rules R is a relation between finite sets of *OL* formulae, and is defined as follows (we leave most technical details, including those of the formal syntax of rules; below we only assume that a set of rules gives rise to a relation over finite sets):

$$\llbracket R \rrbracket = \left\{ (A, C) : \begin{array}{l} A \text{ and } C \text{ are the results of simultaneous uniform substitution} \\ \text{of variables } t \in V_T \text{ with a finite sets of } OL \text{ formulae, and} \\ \text{variables } a \in V_F \text{ with single } OL \text{ formulae, of the antecedent} \\ \text{and consequent, respectively, of a rule in } R \end{array} \right\}$$

There are two rule operators, for knowing a (set of) rule(s) *at least* and *at most*. Knowing at least a set of rules means that if the agent's current state matches the antecedent of a rule in the set, then he is able to get to know all the formulas in the consequent of the rule in the next state. The agents also communicate by sending formulae to each other, and all agents act simultaneously (unlike in the logic described in this paper), so an agent's next epistemic state may include more than just the formulas in the consequent of the rule. For

example, $\widetilde{\Delta}_i\{R_1\}$ express the fact that agent i can reason with *modus ponens*. The dual operator $\widetilde{\nabla}_i$ expresses an upper bound on what an agent can do: an agent knows *at most* a rule if all he can do is described by the rule. For example,

$$\widetilde{\nabla}_i\left\{\frac{t}{t \sqcup u}\right\}$$

says that agent i *must* reason monotonically.

The semantic assumptions in DSEL and in this paper are quite different. In particular, in DSEL it is assumed that the next state of the system is a function of choices/actions made by *all* agents simultaneously, while we assume that the next state is a function of a choice/action made by a single agent. Because of this, and the fact that the temporal fragment in DSEL is much more expressive than ours, we limit the discussion here to an interpretation of the DSEL versions of “knowing at least” and “knowing at most” a rule in our logic rather than comparing the complete logics in detail. The interpretations of these concepts are slightly different from the ones we used in Section 4. Also, there is no complete axiomatisation of DSEL, and we saw in Section 4 that complete axiomatisation of concepts related to “knowing at most” a rule is difficult, and we limit the discussion here to defining the corresponding model classes with our language rather than investigating completeness.

A possible interpretation in our logic is as follows. We say that an agent knows *at least* rules R (in the DSEL sense) when he knows exactly formulae X if he can access a world where he knows exactly C for each $(X, C) \in \llbracket R \rrbracket$. Dually, he knows *at most* R when he knows exactly formulae X if for every world he can access where he knows exactly C there exists $(X, C) \in \llbracket R \rrbracket$. Formally, let $M = (W, R_1, \dots, R_n, V)$ be a model and R a set of rules. In a state $w \in W$:

$$\text{ATLEAST-R } (\overline{V_i(w)}, s') \in \llbracket R \rrbracket \Rightarrow \exists_{(w, w') \in R_i} \overline{V_j(w')} = s'$$

$$\text{ATMOST-R } (\overline{V_i(w)}, s') \in \llbracket R \rrbracket \Leftarrow \exists_{(w, w') \in R_i} \overline{V_j(w')} = s'$$

$$\text{EXACTLY-R } (\overline{V_i(w)}, s') \in \llbracket R \rrbracket \Leftrightarrow \exists_{(w, w') \in R_i} \overline{V_j(w')} = s'$$

It is easy to see that any of the three properties given above is not expressible by a single \mathcal{L}^∇ formula. We instead give characterisations by translating a set of rules to a set of \mathcal{L}^∇ formulae.

Definition 6 *Translations* $f : \wp(RUL) \rightarrow \wp(\mathcal{L}^\nabla_n)$ and $g : \wp(RUL) \rightarrow \wp(\mathcal{L}^\nabla_n)$ are defined as follows:

$$f(R) = \{\boxtimes_i X \rightarrow \diamond_i \boxtimes_i Y : (X, Y) \in \llbracket R \rrbracket\}$$

$$g(R) = \{\boxtimes_i X \rightarrow \square_i \neg \boxtimes_i Y : (X, Y) \notin \llbracket R \rrbracket\}$$

Theorem 21 *Let* $M = (W, R_1, \dots, R_n, V)$ *be a model and* R *a set of rules.*

1. $M, w \models f(R) \Leftrightarrow \text{ATLEAST-R holds in } w$
2. $M, w \models g(R) \Leftrightarrow \text{ATMOST-R holds in } w$
3. $M, w \models f(R) \cup g(R) \Leftrightarrow \text{EXACTLY-R holds in } w$

Proof.

1. \Rightarrow) Let $M, w \models f(R)$ and $(\overline{V_i(w)}, s') \in \llbracket R \rrbracket$. $M, w \models \boxtimes_i \overline{V_i(w)}$, so $M, w \models \diamond_i \boxtimes_i s'$, i.e. there exists $(w, w') \in R_i$ such that $M, w \models \boxtimes_i s'$ and thus $s' = \overline{V_j(w')}$.

\Leftrightarrow Let ATLEAST-R hold in w , let $(X, Y) \in \llbracket R \rrbracket$ and let $M, w \models \boxtimes_i X$. Since $(\overline{V_i(w)}, Y) \in \llbracket R \rrbracket$, there exists $(w, w') \in R_i$ such that $Y = \overline{V_j(w')}$. Thus, $M, w \models \diamond_i \boxtimes_i Y$.

2. \Rightarrow Let $M, w \models g(R)$ and $(w, w') \in R_i$ and $s' = \overline{V_j(w')}$. $M, w \models \diamond_i \boxtimes_i s'$. Assume that $(\overline{V_i(w)}, s') \notin \llbracket R \rrbracket$. $M, w \models \boxtimes_i \overline{V_i(w)}$, so $M, w \models \square_i \neg \boxtimes_i s'$, which is a contradiction. Thus, $(\overline{V_i(w)}, s') \in \llbracket R \rrbracket$.

\Leftrightarrow Let ATMOST-R hold in w , let $(X, Y) \notin \llbracket R \rrbracket$ and let $M, w \models \boxtimes_i X$. Since $(\overline{V_i(w)}, Y) \notin \llbracket R \rrbracket$, by ATMOST-R $\forall (w, w') \in R_i Y \notin \overline{V_j(w')}$. Thus, $M, w \models \square_i \neg \boxtimes_i Y$.

3. Immediate. □

It follows that the models of $f(R)$ ($g(R)$) are exactly the models where ATLEAST-R (ATMOST-R) holds in every state.

5.3.1 Communication

In DSEL, rules and rule operators are also used to express properties about *communication* in addition to properties about reasoning. For example,

$$\widetilde{\Delta}_{ij} \left\{ \frac{t \sqcup \{a\}}{\{a\}} \right\}$$

expresses that agent i can communicate to j any single formula known to him (i); while

$$\widetilde{\nabla}_{ij} \left\{ \frac{t \sqcup \{a\}}{\{a\}} \right\}$$

expresses that agent i can *only* communicate formulae known to him (i).

In our logic, where only one agent acts at any single time, there is a difference in the intuitive interpretation of knowing a reasoning rule and knowing a communication rule. In the former case, formalised above, the consequent of the rule is assumed to describe the *complete* next epistemic state of the agent. In the latter, the consequent should only be interpreted as a *part* of the next epistemic state of the agent. Formally, we can describe the fact that i knows a set of rules R for communication to j in state w of model $M = (W, R_1, \dots, R_n, V)$ as follows:

$$\text{ATLEAST-R}' \quad (\overline{V_i(w)}, s') \in \llbracket R \rrbracket \Rightarrow \exists (w, w') \in R_i s' \subseteq \overline{V_j(w')}$$

$$\text{ATMOST-R}' \quad (\overline{V_i(w)}, s') \in \llbracket R \rrbracket \Leftarrow \exists (w, w') \in R_i s' \subseteq \overline{V_j(w')}$$

$$\text{EXACTLY-R}' \quad (\overline{V_i(w)}, s') \in \llbracket R \rrbracket \Leftrightarrow \exists (w, w') \in R_i s' \subseteq \overline{V_j(w')}$$

A corresponding translation is:

Definition 7 Translations $f' : \wp(RUL) \rightarrow \wp(\mathcal{L}^{\nabla}_n)$ and $g' : \wp(RUL) \rightarrow \wp(\mathcal{L}^{\nabla}_n)$ are defined as follows:

$$f'(R) = \{\boxtimes_i X \rightarrow \diamond_i \Delta_j Y : (X, Y) \in \llbracket R \rrbracket\}$$

$$g'(R) = \{\boxtimes_i X \rightarrow \square_i \neg \Delta_j Y : (X, Y) \notin \llbracket R \rrbracket\}$$

Theorem 22 Let $M = (W, R_1, \dots, R_n, V)$ be a model and R a set of rules.

1. $M, w \models f'(R) \Leftrightarrow \text{ATLEAST-R}'$ holds in w

2. $M, w \models g'(R) \Leftrightarrow \text{ATMOST-}R'$ holds in w
3. $M, w \models f'(R) \cup g'(R) \Leftrightarrow \text{EXACTLY-}R'$ holds in w

Proof. Similar to the proof of Theorem 21. □

6 Conclusions

In this paper, we have shown how a modal epistemic logic may be used to formalise rational, but non-omniscient agents. This problem is a challenge to the standard epistemic modal logics, where knowledge is closed under logical consequence. This makes them unsuitable for specifying and verifying properties of non-omniscient agents, which may be able to reliably apply a set of inference rules. We have investigated the formal properties of such logics, and shown how to axiomatise several interesting classes of transition systems in them. One of the classes provides a semantics for the logic introduced in [8]. We are optimistic that this type of logics can be used to formulate properties of agents in standard verification tasks, and plan to investigate their use in model-checking resource-bounded reasoners in further work.

Acknowledgements The work presented in this paper was done while Natasha Alechina visited the Department of Informatics, University of Bergen. We are grateful to Michal Walicki, Marc Bezem and the NFR project Modularity in large software and information systems (nr. 146967/431) for enabling this collaboration. Natasha Alechina would like to thank Royal Society for sponsoring related work on logics for resource-bounded agents.

References

- [1] Thomas Ågotnes. *A logic of Finite Syntactic Epistemic States*. Ph.D. thesis, Department of Informatics, University of Bergen, Norway, 2004. 1
- [2] Thomas Ågotnes and Michal Walicki. Syntactic knowledge: A logic of reasoning, communication and cooperation. In *Proceedings of the Second European Workshop on Multi-Agent Systems (EUMAS)*, Barcelona, Spain, December 2004. (document), 2.2.1, 5.3
- [3] Thomas Ågotnes and Michal Walicki. Complete axiomatizations of finite syntactic epistemic states. In *Proceedings of the 3rd International Workshop on Declarative Agent Languages and Technologies (DALT 2005)*, Utrecht, the Netherlands, July 2005. 2.3, 2.3, 2.3.1, 5
- [4] Thomas Ågotnes and Michal Walicki. Strongly complete axiomatizations of "knowing at most" in standard syntactic assignments. In Francesca Toni and Paolo Torroni, editors, *Pre-proceedings of the 6th International Workshop on Computational Logic in Multi-agent Systems (CLIMA VI)*, London, UK, June 2005. 2.3.1, 5
- [5] Natasha Alechina, Brian Logan, and Mark Whitsey. A complete and decidable logic for resource-bounded agents. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2004)*, pages 606–613. ACM Press, Jul 2004. (document), 2.2.1, 5, 5.1, 5.1
- [6] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. In *38th Annual Symposium on Foundations of Computer Science*, pages 100–109, Miami Beach, Florida, 20–22 October 1997. IEEE. 5.3

- [7] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge University Press, 2001. 2.3.1
- [8] H. N. Duc. Reasoning about rational, but not logically omniscient, agents. *Journal of Logic and Computation*, 7(5):633–648, 1997. (document), 1, 5.2, 6
- [9] R. A. Eberle. A logic of believing, knowing and inferring. *Synthese*, 26:356–382, 1974. 5
- [10] Jennifer Elgot-Drapkin, Sarit Kraus, Michael Miller, Madhura Nirkhe, and Donald Perlis. Active logics: A unified formal approach to episodic reasoning. Techn. Rep. CS-TR-4072, 1999. 2.2.1, 5
- [11] Jennifer J. Elgot-Drapkin and Donald Perlis. Reasoning situated in time I: Basic concepts. *Journal of Experimental and Theoretical Artificial Intelligence*, 2:75–98, 1990. 5.1
- [12] Ronald Fagin and Joseph Y. Halpern. Belief, awareness and limited reasoning. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 491–501, Los Angeles, CA, 1985. 13
- [13] Ronald Fagin and Joseph Y. Halpern. Belief, awareness and limited reasoning. *Artificial Intelligence*, 34:39–76, 1988. A preliminary version appeared in [12]. 5
- [14] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. The MIT Press, Cambridge, Massachusetts, 1995. 5
- [15] Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, July 1990. 5
- [16] J. Hintikka. Impossible possible worlds vindicated. *Journal of Philosophical Logic*, 4:475–484, 1975. 5
- [17] Kurt Konolige. *A Deduction Model of Belief and its Logics*. PhD thesis, Stanford University, 1984. 5
- [18] R. C. Moore and G. Hendrix. Computational models of beliefs and the semantics of belief sentences. Technical Note 187, SRI International, Menlo Park, CA, 1979. 5
- [19] Antonio Moreno. Avoiding logical omniscience and perfect reasoning: a survey. *AI Communications*, 11:101–122, 1998. 5
- [20] Kwang Mong Sim. Epistemic logic and logical omniscience: A survey. *International Journal of Intelligent Systems*, 12:57–81, 1997. 5
- [21] W. van der Hoek and M. Wooldridge. Towards a logic of rational agency. *Logic Journal of the IGPL*, 11(2):135–159, 2003. 1