
The Dynamics of Syntactic Knowledge

THOMAS ÅGOTNES, *Department of Informatics,
University of Bergen P.B. 7800, N-5020 Bergen, Norway.*
E-mail: agotnes@ii.uib.no

NATASHA ALECHINA, *School of Computer Science and IT,
University of Nottingham, Nottingham, UK.*
E-mail: nza@cs.nott.ac.uk

Abstract

The syntactic approach to epistemic logic avoids the logical omniscience problem by taking knowledge as primary rather than as defined in terms of possible worlds. In this study, we combine the syntactic approach with modal logic, using transition systems to model reasoning. We use two syntactic epistemic modalities: ‘knowing at least’ a set of formulae and ‘knowing at most’ a set of formulae. We are particularly interested in models restricting the set of formulae known by an agent at a point in time to be finite. The resulting systems are investigated from the point of view of axiomatization and complexity. We show how these logics can be used to formalise non-omniscient agents who know some inference rules, and study their relationship to other systems of syntactic epistemic logics, such as Ågotnes and Walicki (2004, Proc. 2nd EUMAS, pp. 1–10), Alechina *et al.* (2004, Proc. 3rd AAMAS, pp. 601–613), Duc (1997, J. Logic Comput., 7, 633–648).

Keywords: Modal logic, epistemic logic, logical omniscience.

1 Introduction

The most studied formal logics of knowledge and belief, epistemic and doxastic logics¹ [16, 25], explain knowledge and belief in terms of a possible worlds semantics [21]. An alternative approach is to take knowledge or belief as *primary*. The most straightforward way of doing this is to replace the standard Kripke structures describing the possible worlds with functions assigning truth values to formulae of the form $B\phi$, meaning that ϕ is believed, *directly*. We call this approach the *syntactic* approach, since the truth value of two epistemic formulae can differ if the formulae differ syntactically. The syntactic approach allows modelling more general assumptions about knowledge and belief, for example the lack of closure conditions or the possibility of believing two mutually inconsistent formulae at the same time. Indeed, most work using the syntactic approach have been motivated by the logical omniscience problem (LOP) [22]; that epistemic logic based on possible worlds semantics describes knowledge which is closed under logical consequence. This is a problem if we want to model the explicitly computed knowledge of resource bounded reasoners, for example an agent who chooses which action to perform next based on the formulae it has deduced. 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 [32]. See e.g. [16, 27, 31] for surveys of proposed solutions to the LOP.

¹ In this article we do not distinguish between knowledge and belief, and will use both terms for the same concept. In particular, we do not assume that something which is known must be true.

In this work we study models of the *dynamics* of syntactic knowledge, i.e. how syntactically ascribed knowledge can change over time. Such models provide a description of *reasoning* as an alternative to logical omniscience. The logics we introduce in this article do not suffer from the LOP, but are at the same time 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. Formally, the logics are obtained by combining the standard syntactic approach with modal logics. Transitions in the Kripke structures used to interpret the logics correspond to acts of reasoning or communication. Thus, the modal ‘diamond’ operator can be used in a formula such as $(Bp \wedge B(p \rightarrow q)) \rightarrow \Diamond Bq$ to express that an agent can use modus ponens: if it believes both p and $p \rightarrow q$, then it can get to believe q at the next point in time.

In addition to being interpreted on transition systems which reflect the evolution of knowledge, our logics have two more distinguishing features. The first one is the *finiteness assumption*: at any point in time, any agent knows only finitely many formulae. We argue that the finiteness assumption must hold for the explicitly computed knowledge of agents with finite memory. Second, the language of the logics contains epistemic operators Δ (‘knowing at least’) and ∇ (‘knowing at most’), which apply to finite sets of formulae. These operators were introduced in [1], where the logic of static epistemic states was axiomatized and studied in detail, and a combination of syntactic epistemic knowledge operators and *Alternating Time Temporal Logic (ATL)* [7], was introduced. In this article, we investigate intermediate systems, namely adding Δ and ∇ to standard modal logic. In the modal syntactic epistemic logics developed in this article, we can express properties such as ‘if the agent knows exactly the set of formulae T now, then after n steps it may know T closed under n applications of modus ponens, and nothing else’.

Syntactic approaches to epistemic logic can be classified into logics concerned with *static* knowledge, i.e. knowledge at a point in time, or with the *dynamics* of knowledge, i.e. with how knowledge evolves over time. Another dimension for classification is the finiteness assumption, i.e. whether it is assumed that an agent can only know finitely many formulae at the same time, or whether it may know infinitely many. The logics we present in this article are dynamic logics of finite knowledge.

Classical syntactic approaches to knowledge [11, 16, 17, 26] generally model static knowledge with no assumption about finiteness. The *logic of general awareness* [15] combines a syntactic and an epistemic approach to static knowledge, and can be used to model finite knowledge. Ågotnes and Walicki [4, 5] investigate a logic of static syntactic knowledge without and with, respectively, the finiteness assumption.

Among dynamic approaches, the *deduction model* [23] 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*) [12]. 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)* [6] 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). Duc’s Dynamic Epistemic Logic (DEL)² [10] adds an operator $\langle F_i \rangle$

to the standard epistemic language; the intuitive meaning of a formula such as $\langle F_i \rangle \phi$ is that ϕ will be true after some train of thought of agent i . The described logic can be seen as a dynamic version of the epistemic logic $S4_n$ —however, no semantics is given for the logic. Yet another dynamic approach is *Dynamic Syntactic Epistemic Logic (DSEL)* [3], 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. DSEL is based on ATL.

Our combination of standard syntactic models and simple transition systems yields an abstract, general and powerful formalism. Indeed, we show that our framework can be used to capture several aspects of reasoning and other knowledge dynamics modelled by the approaches mentioned earlier. For example, our framework gives a semantics for which Duc’s DEL is sound and complete.

The article is organized 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 axiomatizable by only the axioms of the modal system \mathbf{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 axiomatization 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 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 condition 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. In Section 5, we present some applications and examples, showing that we can capture important aspects of the dynamics of syntactic knowledge in TRL, DEL and DSEL, in our framework (including soundness and completeness of DEL). Finally, we discuss some other related work 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, such as consistency or closure under consequence, on this set are assumed. Furthermore, we assume that the agent can *act* in order to change its current state, that it 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

² Which should not be confused with recent work in modelling updates of knowledge states in the light of new or changed information, often described by the same term [33]. Duc’s logic is a logic of knowledge change as a result of *reasoning*.

we generalize this model to the multiagent case, where agents cannot 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. Thus, a model will be used to interpret an epistemic atom $B\phi$ as a primitive proposition, and not as a formula with a modal operator. 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 $V : W \rightarrow \wp(\Theta)$. The class of all models is denoted $\mathcal{M}(OL)$ (or just \mathcal{M}).

The relation R is used to model acts of reasoning taking the agent from a state where it knows some finite set of formulae to another state where it knows some other such set. Note that we do not model time explicitly; intuitively, no transitions happen unless the agent chooses to perform a reasoning action.

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 object language formulae 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.

The intended meaning of a formula of the form $\diamond\phi$ is that the agent can choose some reasoning act such that ϕ becomes true in the next state.

We use the usual derived propositional connectives, in addition to $\Box\phi$ for $\neg\diamond\neg\phi$. The intended meaning of a formula of the form $\Box\phi$ is thus that ϕ will be true after *any* reasoning act performed by the agent. $B\phi$ 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 a shorthand:

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

If X is a singleton set, we often omit brackets and use the equivalent notations $\Delta\phi$, $\Delta\{\phi\}$ and $B\phi$ interchangeably.

The interpretation of \mathcal{L} in general 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 axiomatized by the modal logic **K** (in the language \mathcal{L}). The logical system **K** consists of the axiom schemes

(**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 $\Box\phi$

In completeness proofs, we will often need to transform one (general) model into another while preserving satisfaction of certain formulae. We describe here two such transformations. First, we review the well-known concept of *bisimulation*, and, second, we describe a new and more general type of transformation which we call *Σ -bisimulation*.

A brief review of the concept of a bisimulation (see, e.g. [8] for further details):

DEFINITION 3 (Bisimulation)

Let $M = (W, R, V)$ and $M' = (W', R', V')$ be two (general) models. A relation $Z \subseteq W \times W'$ is a bisimulation between M and M' iff the following three conditions hold:

- (1) For all $w \in W, w' \in W'$: if Zww' then $V(w) = V'(w')$ (w and w' satisfy the same atoms)
- (2) For all $w \in W, w' \in W'$: if Zww' and Rwv then there is a $v' \in W'$ such that Zvv' and $R'w'v'$
- (3) For all $w \in W, w' \in W'$: if Zww' and $R'w'v'$ then there is a $v \in W$ such that Zvv' and Rwv

It is well known that bisimulation preserves satisfiability: if Z is a bisimulation between M and M' and Zww' , then $M, w \models \phi$ iff $M', w' \models \phi$ for any formula $\phi \in \mathcal{L}$.

Next, we define Σ -simulation and show a corresponding preservation result.

DEFINITION 4 (Σ -bisimulation)

Let $\Sigma \subseteq \Theta$, and let $M = (W, R, V)$ and $M' = (W', R', V')$ be two (general) models. A relation $Z \subseteq W \times W'$ is a Σ -bisimulation between M and M' iff the following condition holds, in addition to conditions (2) and (3) for standard bisimulation:

(3') For all $w \in W, w' \in W'$: if Zww' then $V(w) \cap \Sigma = V'(w') \cap \Sigma$ (w and w' satisfy the same atoms from Σ)

We write $M \xleftrightarrow{\Sigma} M'$ to denote the fact that there is a Σ -bisimulation between M and M' ; we write $w \xleftrightarrow{\Sigma} w'$ to denote the fact that there is a Σ -bisimulation Z between M and M' such that wZw' .

It is easy to see that when $\Sigma = \Theta$, Σ -bisimulation and bisimulation coincide.

For a \mathcal{L} formula ϕ , let $Subf(\phi)$ be the set of subformulae of ϕ , where subformulae of the form $B\psi$ are treated as atomic subformulae (not parsed further). Let $At(\phi)$ be the set of epistemic atoms occurring in ϕ , namely $B\psi \in At(\phi)$ if $B\psi \in Subf(\phi)$.

The following theorem shows that Σ -bisimulation preserves satisfaction of formulae with epistemic atoms in Σ .

When $\Sigma \subseteq \Theta$, let $\mathcal{L}(\Sigma)$ denote the subset of \mathcal{L} containing formulae ϕ with $At(\phi) \subseteq \Sigma$.

THEOREM 1

Let w be a state in a (general) model M , and w' a state in a (general) model M' .

$$w \xleftrightarrow{\Sigma} w' \Rightarrow \forall \phi \in \mathcal{L}(\Sigma) (M, w \models \phi \Leftrightarrow M', w' \models \phi)$$

PROOF. Let $M = (W, R, V)$ and $M' = (W', R', V')$. If we assume that $w \xleftrightarrow{\Sigma} w'$, we can prove that the equivalence holds for any $\phi \in \mathcal{L}(\Sigma)$ by structural induction. For the base case, $\phi \in \Theta$ and $\phi \in \Sigma$. $M, w \models \phi$ iff $\phi \in V(w)$ iff $\phi \in V(w) \cap \Sigma$ iff, by (1) in Definition 4, $\phi \in V'(w') \cap \Sigma$ iff $M', w' \models \phi$. The inductive step can be shown in exactly the same way as for standard bisimulation. ■

Theorem 1 formally shows the intuitive property that satisfaction of a formula does not depend on the valuation of atoms not mentioned in the formula.

We say that a logic is weakly complete with respect to a class of models, if all formulae valid in the class are provable. A logic is strongly complete if for any set of formulae Γ , any logical consequence of Γ , i.e. any formula valid in the class consisting of all the models of Γ , is provable from Γ .

THEOREM 2

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

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

The semantics we have given for the language \mathcal{L} is not compact. A counterexample 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 with respect to \mathcal{M} .

2.2 Useful axioms

In this section, we consider imposing additional conditions on the accessibility relation in the models.

2.2.1 Unbounded reasoning

Many syntactic approaches to epistemic logic [3, 6, 12] 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

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

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

THEOREM 3

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

PROOF. Like the proof of Theorem 2, with the canonical model for **KD**. ■

2.2.2 Deterministic reasoning

The models in \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

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

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

THEOREM 4

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 result in modal logic, 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 . ■

THEOREM 5

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

PROOF. Axioms F and D are valid in \mathcal{M}^{ds} , which gives soundness. The proof of completeness is analogous to the proof of Theorem 4. ■

2.3 *Knowing at most*

In the previous section, we presented different axiomatizations 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 [4] 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 6

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 Y be the set of atoms occurring in ϕ . $M' = (W, R, V')$ where $V'(w') = V(w')$ when $w' \neq w$ and $V'(w) = V(w) \cup \{B\beta\}$ where $\beta \notin Y \cup X$. Existence of such a β is ensured by the fact that the object language is infinite and both Y and X are finite. M and M' are Y -bisimilar, so by Theorem 1, $M', w \models \phi$. However, because $B\beta \in V'(w)$ and $\beta \notin X$, $M', w \not\models \nabla X$. ■

The proof above essentially uses the fact that $\mathcal{L}^\nabla(OL)$ formulae are not preserved under Σ -bisimulation, or that their truth value depends on assignment to atoms which do not occur in the formula. Most logics satisfy the principle of locality: the truth value of a formula does not depend on the assignment to variables other than the formula's free variables. This is such an obvious property that it usually goes unremarked; however some logics do violate it. This phenomenon was investigated for predicate logics in e.g. [28]; for propositional logics, the only example we know of in addition to the logic of ∇ , is a related logic of operators 'knowing at least n formulae' [2] and the logic of *only knowing* [24].

Although formulae of $\mathcal{L}^\nabla(OL)$ are not preserved under Σ -bisimulation, they are preserved under an operation which we call epistemic filtration through Σ .

Recall the definition of $At(\phi)$ for $\phi \in \mathcal{L}(OL)$. We extend it to $\phi \in \mathcal{L}^\nabla(OL)$ as follows. For an $\mathcal{L}^\nabla(OL)$ formula ϕ , let $Subf(\phi)$ be the set of subformulae of ϕ , where subformulae of the form $B\psi$ and ∇X are treated as atomic formulae. $B\psi \in At(\phi)$ if either $B\psi \in Subf(\phi)$, or for some X such that $\psi \in X$, $\nabla X \in Subf(\phi)$.

DEFINITION 5

Let $\Sigma \subseteq \Theta$, and $M = (W, R, V)$ be a (general) model. An epistemic filtration of M through Σ is a model $M' = (W', R', V')$ where $W' = W$, $R' = R$, and for every w , $V'(w)$ is the same as $V(w)$ with respect to the atoms in Σ , and all atoms $B\psi \in V(w)$ which are not in Σ are replaced in $V'(w)$ by some fresh atom $B\alpha \notin \Sigma$.

Note that $V(w)$, for every w , can be represented as $Y \cup Z$, where $Y \subseteq \Sigma$ and $Z \cap \Sigma = \emptyset$. The operation above can be seen as replacing $V(w)$ with $V'(w) = Y$ if $Z = \emptyset$, and $V'(w) = Y \cup \{B\alpha\}$ otherwise.

As before, we refer to the sublanguage of $\mathcal{L}^\nabla(\Theta)$ which only refers to atoms from Σ as $\mathcal{L}^\nabla(\Sigma)$. In particular, $\nabla\{\phi_1, \dots, \phi_n\}$ is in $\mathcal{L}^\nabla(\Sigma)$ if $B\phi_1, \dots, B\phi_n \in \Sigma$.

THEOREM 7

$\mathcal{L}^\nabla(\Sigma)$ -formulae are preserved under epistemic filtration through Σ .

PROOF. Let $\phi \in \mathcal{L}^\nabla(\Sigma)$ and let $M = (W, R, V)$ be a (general) model. We argue by induction on the complexity of ϕ that for every $w \in W$, $M, w \models \phi$ iff $M', w \models \phi$, where M' is a result of filtrating M through Σ .

Clearly, for $B\psi \in \Sigma$, $B\psi \in V(w)$ iff $B\psi \in V'(w)$. The cases for booleans and \diamond are trivial. The only interesting case is the case of ∇ . We are going to show that for every set of epistemic atoms $X \subseteq \Sigma$ the following condition holds:

$$V(w) \subseteq X \text{ iff } V'(w) \subseteq X$$

From this it would immediately follow that if $\{B\alpha : \alpha \in X\} \subseteq \Sigma$, then $M, w \models \nabla X$ iff $M', w \models \nabla X$.

First, consider the case when $V(w) \subseteq \Sigma$. Then $V'(w) = V(w)$ and the condition above holds trivially, for any set of epistemic atoms X . In the other case, $V(w)$ contains at least one atom $B\psi \notin \Sigma$. Then $V'(w) = (V(w) \cap \Sigma) \cup \{B\alpha\}$ for some fresh $B\alpha \notin \Sigma$. For $X \subseteq \Sigma$, the condition $V(w) \subseteq X$ iff $V'(w) \subseteq X$ holds trivially, since if $V(w)$ contains a non- Σ atom, $V(w) \subseteq X$ is false, and so is $V'(w) \subseteq X$ (because $V'(w)$ contains $B\alpha$ which is not in Σ and hence not in X). ■

COROLLARY 1

Every $\mathcal{L}^\nabla(OL)$ formula which has a general model, has a model.

We now present axiomatizations of logics with the ∇ operator. The axiomatizations extend the axiomatization of the purely epistemic fragment (i.e. without the modal logic) in [4].

2.3.1 Axiomatization

The properties of ∇ are captured by the following axiom schemata [4]:

(E1) $\Delta X \rightarrow \neg \nabla Y$ when $X \not\subseteq Y$

(E2) $\nabla(Y \cup \{\gamma\}) \wedge \neg B\gamma \rightarrow \nabla Y$

(E3) $\nabla X \rightarrow \nabla Y$ when $X \subseteq Y$

E1 says essentially that $\Delta X \wedge \nabla Y$ (believing all formulae in X and at most the formulae in Y) implies $X \subseteq Y$. Note that $X \subseteq Y$ is not a formula in the logic, so instead we formulate the axiom schema which declares all instances of $\Delta X \wedge \nabla Y$ with $X \not\subseteq Y$ inconsistent. E2 allows us to narrow down the set of beliefs: if we know that the agent believes at most $\{\gamma_1, \dots, \gamma_n\}$ and does not believe γ_n , then we can conclude that it believes at most $\{\gamma_1, \dots, \gamma_{n-1}\}$. Finally, E3 simply states that ∇ is monotone.

We write \mathbf{K}_∇ for \mathbf{K} extended with E1–E3, 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 use the fact that each general model for a formula ϕ can be transformed into a model for ϕ by epistemic filtration.

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 will transform M^c to a model $M \in \mathcal{M}(OL)$). The results extend those of [5], 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 6

Let $\hat{p} \notin OL$. Define the 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 $\forall \phi \in \mathcal{L}^\nabla(OL) (\phi \in v \Rightarrow \Diamond \phi \in w)$
- $V^c : W^c \rightarrow \wp(\Theta(OL \cup \{\hat{p}\}))$ is defined by cases. If w contains an occurrence of ∇X for some $X \in \wp^{fin}(OL)$, then $V^c(w) = \{B\alpha : B\alpha \in w\}$. Otherwise $V^c(w) = \{B\alpha : B\alpha \in w\} \cup \{B\hat{p}\}$.

The reason for adding $B\hat{p}$ in the last clause is to provide a ‘witness’ for failure of ∇X for any $X \in \wp^{fin}(OL)$: $\hat{p} \notin X$ but $\hat{p} \in \overline{V^c(w)}$.

THEOREM 8

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

Before we prove Theorem 8, 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 $\nabla X' \in w$. For the direction to the right, let $\nabla X \in w$ and let $\alpha \in Bel^w$. If it were the case that $\alpha \notin X$, then, since $B\alpha \in w$, $\neg \nabla X \in w$ by **E1** which contradicts the consistency of w . Thus, $\alpha \in X$. For the direction to the left, let $Bel^w \subseteq X$. If $X' \subseteq X$, then $\nabla X \in w$ by **E3** and we are done, so assume that $X' \not\subseteq X$. Let $\alpha_1, \dots, \alpha_k$ be the formulae in X' but not in X , i.e. $X' \setminus X = \{\alpha_1, \dots, \alpha_k\}$ (X', X are finite). Since $Bel^w \subseteq X$ and $\alpha_i \notin X$, $\alpha_i \notin Bel^w$ for any $i \in [1, k]$. Thus, for each i , $B\alpha_i \notin w$, and by maximality $\neg B\alpha_i \in w$. By (k applications of) **E2**, $\nabla(X' \setminus \{\alpha_1, \dots, \alpha_k\}) \in w$, and since $X' \setminus \{\alpha_1, \dots, \alpha_k\} \subseteq X$, $\nabla X \in w$ by **E3**. ■

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 8 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 . We will only do the case for ∇ , as the rest is standard:

$\phi = \nabla X$: We have two cases.

First, assume that for some X' , $\nabla X' \in w$. Then by the definition of V^c , $\overline{V^c(w)} = Bel^w$. By Lemma 1, $\nabla X \in w$ iff $Bel^w \subseteq X$, which is equivalent to $\nabla X \in w$ iff $\overline{V^c(w)} \subseteq X$, hence $\nabla X \in w$ iff $M^c, w \models \nabla X$.

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 $\overline{V^c(w)} \not\subseteq X$ which is the same as $M^c, w \not\models \nabla X$. ■

Theorem 8 follows from Lemma 3 and the standard Lindenbaum argument.

By Theorem 7, since ϕ is satisfied in a general model $M^c = (W^c, R^c, V^c)$, it also has a model M^f which is a result of filtrating M^c through $At(\phi)$.

THEOREM 9

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. This follows immediately from Theorem 7: we take an epistemic filtration of M through $At(\phi)$, choosing some fresh formula $\alpha \in OL$ to obtain an epistemic atom $B\alpha \notin At(\phi)$ to replace all non- $At(\phi)$ atoms in V^f , including \hat{p} . ■

THEOREM 10

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

PROOF. Soundness follows from the easily seen fact that E1-E3 are valid—in addition to validity of \mathbf{K} and the fact that MP and Gen preserve validity.

Completeness follows from Theorem 8 and Theorem 9. ■

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

Formulae 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

Formulae 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 [8]; 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 formulae of depth less or equal to n .

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

The 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 have checked all formulae in X and the counter is equal to 0, then the formula ∇X is true, else it is false (there are formulae 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 $\text{NPSpace}=\text{PSPACE}$, 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 $\mathbf{K}\nabla$ produces logics which are sound and weakly complete with respect to the corresponding classes of models.

THEOREM 11

$\mathbf{KF}\nabla$ 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 $\mathbf{KF}\nabla$ -consistent formula ϕ . By Theorem 8, 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 formulae, and identical worlds can be glued together, yielding a bisimilar model). By Theorem 9, ϕ has a model in \mathcal{M}^d . ■

THEOREM 12

$\mathbf{KD}\nabla$ is sound and complete with respect to \mathcal{M}^s .

PROOF. Straightforward from Theorem 8 and Theorem 9. ■

THEOREM 13

$\mathbf{KDF}\nabla$ is sound and complete with respect to \mathcal{M}^{ds} .

PROOF. Straightforward from Theorem 8 and Theorem 9. ■

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

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

THEOREM 14

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

PROOF. Straightforward from Theorem 8 and Theorem 9. ■

3 The multiagent 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 multiagent case where there are n agents $\Lambda = \{1, \dots, n\}$. In this case both syntax and semantics are defined over $\Theta_n = \{B_i\phi : \phi \in OL, i \in \Lambda\}$ in place of Θ .

DEFINITION 7

A model in the multiagent 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 $V : W \rightarrow \wp^{fin}(\Theta_n)$. The class of all multiagent 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 no longer corresponds 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 multiagent system. Note that the transitions available to agent i depend the state of the whole system.

The logical language for the case of n agents, denoted \mathcal{L}_n^∇ , is defined from Θ_n in the same way \mathcal{L}^∇ is defined from Θ , with the operators \diamond and ∇ replaced by indexed operators \diamond_i and ∇_i for each agent i ($1 \leq i \leq n$). The interpretation in multiagent models is straightforward, with R_i interpreting \diamond_i and s_i interpreting B_i and ∇_i . 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. In details: if it is the case that, right now, agent 1 knows at most p , i.e. it knows either only p or nothing at all, and agent 2 knows at most q , then agent 1 can perform an action (a communication action) such that after the action it holds that agent 2 at most knows p and q . In other words, after the action agent 2 might still know q , if it already knew it before the action, and might know p if it was communicated by agent 1, but since it was assumed that agent 2 has no other ways of acquiring new beliefs, it can not know anything else at this point. Further examples involving the knowledge of rules are discussed in Section 4.

Axiomatization of the multiagent case is obtained by indexing modalities and epistemic operators by agents; for example, the epistemic axioms become

$$(E1) \Delta_i X \rightarrow \neg \nabla_i Y \text{ when } X \not\subseteq Y$$

$$(E2) \nabla_i(Y \cup \{\gamma\}) \wedge \neg B_i \gamma \rightarrow \nabla_i Y$$

$$(E3) \nabla_i X \rightarrow \nabla_i Y \text{ when } X \subseteq Y$$

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 corresponding results hold also for the multiagent case.

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

In this section, we consider multimodal 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:

$$\Delta_i MP : \forall \phi, \psi [\phi, \phi \rightarrow \psi \in s_i \Rightarrow \exists s' (R_i(s, s') \wedge \psi \in s'_i)]$$

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

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

This definition of 'knowing at most MP' can be seen as 'the only possible updates are MP updates', and, again, several other definitions are possible.

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

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

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

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

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

In this section, we show some completeness results. One concerns axiomatizing monotonic reasoners:

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

This condition says that after any transition (by any agent *j*), agent *i*'s belief state does not decrease. Syntactically, this corresponds to the following axiom schema:

(**M_i**) $(\Delta_i X \rightarrow \Box_j \Delta_i X) \wedge (\Diamond_j \nabla_i X \rightarrow \nabla_i X)$ (for all *j*).

THEOREM 15

KM_i ∇_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 6, where \mathbf{M}_i holds in all worlds. This forces epistemic states of i in the general canonical model to 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 let $Aware_i(\phi) = \{\psi : B_i \psi \in At(\phi)\}$. Now consider an epistemic filtration M^f of M^c through $At(\phi)$. Let us denote the epistemic state of i in state s as s'_i in the new model. For all states s where s'_i is defined as $s_i \cap Aware_i(\phi)$, $R'_j(s', s'')$ implies $s'_i \subseteq s''_i$, because $s_i \subseteq s'_i$. Suppose s'_i is defined as $(s_i \cap Aware_i(\phi)) \cup \{\alpha_i\}$, that is, a special formula α_i is added to s'_i . Mon_i would be violated if for some s' , $R'_j(s', s'')$ and $\alpha_i \notin s''_i$. 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. ■

Note that we also can axiomatize Mon_i for 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 relations, which we will call addition conditions. Those conditions correspond to ‘knowing at least’ rules of the following form: if agent i believes formulae ϕ_1, \dots, ϕ_n , then agent i can reach a state where agent j (possibly $j = i$) believes formulae 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 16

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 axiomatizable by adding to \mathbf{K}_n axiom schemes 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 2 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. ■

THEOREM 17

\mathbf{K}_n together with an axiom schema

$$(A1) \Delta_1\{\phi, \phi \rightarrow \psi\} \rightarrow \Diamond_1 \Delta_1\{\psi\}$$

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

PROOF. The theorem follows from Theorem 16. ■

It is much harder to axiomatize knowing ‘at most’ a rule. To get a feeling for the reason why, consider axiomatizing 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 $\phi \rightarrow \psi$, ϕ which are in the agent’s state. However, there are infinitely many formulae ϕ ; expressing this would require quantification over formulae:

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

A partial solution—which still requires only the language with ‘at least’ only—is to restrict our attention to rules which have a ‘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 18

The class of models satisfying $\nabla_i \wedge_{int}$ is axiomatizable by adding to \mathbf{K}_n the following axiom schemes:

- (C1) $\diamond_i \Delta_i (\phi \wedge \psi) \rightarrow [\Delta_i (\phi \wedge \psi) \vee (\Delta_i \phi \wedge \Delta_i \psi)]$
- (C2) $\diamond_i \Delta_i X \rightarrow \Delta_i X \vee \bigvee_{\phi \wedge \psi \in X} \Delta_i \{X \setminus \{\phi \wedge \psi\}\}$
- (C3) $\diamond_i \Delta_j X \rightarrow \Delta_j X$ where $j \neq i$
- (C4) $\Delta_j X \rightarrow \square_i \Delta_j X$

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 to 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. ■

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

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

THEOREM 19

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

- (B1) $\nabla_i X \rightarrow \square_i \bigvee_{\phi, \phi \rightarrow \psi \in X} \nabla_i (X \cup \{\psi\})$
- (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 $j \neq i$

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

$$\mathbf{C5} \ \Diamond_i \nabla_j X \rightarrow \nabla_j X$$

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 formulae X in the next state, then it either already believes all formulae 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 Definition 6. We are going to transform it into a model for a consistent formula ϕ' which 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 formulae 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 $At(\phi')$, 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 . ■

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 Applications and examples

As mentioned in the introduction, dynamic aspects of syntactic knowledge have been studied before. Here, we apply our framework to reason about examples of properties from the literature involving reasoning and communication.

5.1 Timed reasoning logic

TRL was introduced in [6]. TRL is a family of logics parameterized by a set of agents A and a rule system (set of inference rules and associated rule application strategy) for each agent. Here, we show that we can apply our framework to reason about concepts from TRL.

Each agent $i \in A$ has a local state which is a finite set of formulae 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 formulae 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 formulae 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 formulae in the language \mathcal{L}_t^i to another finite set of formulae in the language \mathcal{L}_{t+1}^i (an agent's computation of the next state).

DEFINITION 8

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 $\langle obs, \{inf_i : i \in A\}, \{m_t^i : i \in A, t \in \mathbb{N}\} \rangle$ where obs is a function which maps a pair (i, t) to a finite set of formulae in \mathcal{L}_t^i , inf_i is a function from finite sets of formulae in \mathcal{L}_t^i to finite sets of formulae in \mathcal{L}_{t+1}^i , and each m_t^i is a finite set of formulae 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 a logical consequence of a set of labelled formulae Γ , $\Gamma \models (i, t) : \phi$, iff for every model M , if $M \models \Gamma$ then $M \models (i, t) : \phi$.

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 formulae derived as a result). The name of the logic reflects its similarity with *step logic* [13].

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 formulae Γ using R ($\Gamma \vdash_R (i, t) : \phi$) if there is a sequence of labelled formulae $(i_1, t_1) : \phi_1, \dots, (i_n, t_n) : \phi_n$ such that:

- (1) Each formula in the sequence is either a member of Γ , or is obtained from the previous formulae by one of the inference rules in R ; 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 in R of the form

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

inf_i in M satisfies the property

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

(2) For each communication rule in R 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 [6]:

THEOREM 20

Given a set of TRL(STEP) rules R , for any finite set of labelled formulae 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 and $\models_{\mathcal{R}}$ is the consequence relation restricted to the models in \mathcal{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 $m_0, m_1, \dots, m_t, \dots$ where the set of epistemic states associated with each m_t is (m_t^1, \dots, m_t^n) . Each pair m_{t-1}, m_t is connected by a successor relation $R = R_1 = \dots = R_n$ which is deterministic and serial. Note that this class of models can be axiomatized by adding to **KDF_n** a set of axiom schemes

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

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

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

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

THEOREM 21

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

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

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

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

an axiom schema

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

PROOF. The proof is very similar to the proof of Theorem 16, 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 . ■

5.2 Dynamic epistemic logic

The logic we have presented can be seen as a generalization of DEL [10]. Duc [10] defines the logical language L_{DE} by extending the standard epistemic language 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 [20]: 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 inside the modal operators—there is only one modal operator $\langle F_i \rangle$ for each agent and no PDL-like composite operators such as $\langle F_i * \rangle$ or $\langle F_i \cup F_j \rangle$.

Formally, L_{DE} is defined over a number of agents n and a set of primitive propositions Φ as follows. L_E is the standard epistemic language over n and Φ (Duc uses K_i instead of our B_i for belief/knowledge). L_E and L_{DE} are the least sets such that:

$$\begin{array}{ll} \Phi \subseteq L_E & L_E \subseteq L_{DE} \\ \text{if } 1 \leq i \leq n \text{ and } \phi \in L_E, \text{ then } K_i \phi \in L_E & \text{if } \phi \in L_{DE}, \text{ then } \neg \phi \in L_{DE} \\ \text{if } \phi \in L_E, \text{ then } \neg \phi \in L_E & \text{if } \phi, \psi \in L_{DE}, \text{ then } (\phi \rightarrow \psi) \in L_{DE} \\ \text{if } \phi, \psi \in L_E \text{ then } (\phi \rightarrow \psi) \in L_E & \text{if } 1 \leq i \leq n \text{ and } \phi \in L_{DE}, \text{ then } \langle F_i \rangle \phi \in L_{DE} \end{array}$$

The operator $[F_i]$ is defined as a dual to $\langle F_i \rangle$ in the usual way: $[F_i] \phi \equiv \neg \langle F_i \rangle \neg \phi$. DEL is presented via a logical system $DES4_n$ over the language L_{DE} , intended to be a ‘dynamic’ version of $S4_n$. In the definition of the system, a sublanguage $L_E^+ \subseteq L_E$ of *persistent formulae* is used. L_E^+ contains all the formulae of L_E without any occurrences of the knowledge operators K_i (called *objective formulae*) and is closed under the following conditions:

$$\begin{array}{l} \text{If } \phi, \psi \in L_E^+, \text{ then } (\phi \wedge \psi) \in L_E^+ \\ \text{If } \phi, \psi \in L_E^+, \text{ then } (\phi \vee \psi) \in L_E^+ \\ \text{If } \phi \in L_E^+, \text{ then } K_i \phi \in L_E^+ \end{array}$$

$DES4_n$ 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:

$$\begin{array}{l}
R1. \quad \frac{\vdash_{DES4_n} \phi, \vdash_{DES4_n} \phi \rightarrow \psi}{\vdash_{DES4_n} \psi} \\
R2. \quad \frac{\vdash_{DES4_n} \phi}{\vdash_{DES4_n} [F_i]\phi}
\end{array}$$

PC1-PC3 and R1 axiomatize 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.³ The main motivation behind DEL is to describe agents who are non-omniscient but nevertheless rational and non-ignorant. Examples of $DES4_n$ theorems intended to illustrate this are:

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

We now show that we can successfully apply our framework to the mentioned open problem of DEL: to provide a semantics for which the logic is sound and complete. Our logic can be seen as a generalization 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

³ In [9], he describes semantics and proves completeness for a simpler system, although with respect to what we call general models, where an agent may know infinitely many formulae.

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. The assumptions 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} satisfying additional conditions. We can work in the language \mathcal{L} (Section 2.1), as we do not need the ∇ operator. We can view L_{DE} formulae as \mathcal{L} formulae by:

- (1) Taking the object language $OL = L_E$.
- (2) Making the following very minor modification of \mathcal{L} , due to the fact that we for reasons of simplicity up until now have not included primitive propositions directly as well formed formulae (as defined in Section 2.1, while Bp is a well formed formula, p is not). Add the primitive DEL propositions Φ as primitive propositions in \mathcal{L} , and change the definition of V in the models as follows to interpret also Φ : $V : W \rightarrow \mathcal{P}^{fin}(\Theta \cup \Phi)$ (still with the proviso that for every world w and agent i the set $\{B_i\alpha : B_i\alpha \in V(w)\}$ is finite (otherwise it is a general model), and we still keep the notation of $\overline{V(w)}$ for the set $\{\alpha : B_i\alpha \in V(w)\}$). The truth definition for Φ is standard: $M, w \models p \Leftrightarrow p \in V(w)$.
- (3) Reading a K_i operator as B_i .
- (4) Reading a $\langle F_i \rangle$ operator as \diamond_i .

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.

Let Ax_i be the set of formulae DE4-DE7 prescribe that agent i must know in some accessible state:

$$Ax_i = \left\{ \begin{array}{l} \phi \rightarrow (\psi \rightarrow \phi), \\ (\phi \rightarrow (\psi \rightarrow \gamma)) \rightarrow ((\phi \rightarrow \psi) \rightarrow (\phi \rightarrow \gamma)), \\ (\neg\psi \rightarrow \neg\phi) \rightarrow (\phi \rightarrow \psi), \\ K_i\phi \rightarrow \phi \end{array} \right\}$$

for all $\phi, \psi, \gamma \in L_{DE}$.

Let $\mathcal{M}^{DEL} \subset \mathcal{M}$ (respectively, $\mathcal{M}^{gen,DEL} \subset \mathcal{M}^{gen}$) be the class of (general) models $M = (W, R_1, \dots, R_n, V)$ satisfying the following conditions for each agent i and each state $w \in W$:

- D1** R_i is transitive
- D2** If $R_i(w, v)$, $K_i\phi \in V(w)$ and $\phi \in \mathcal{L}_E^+$, then $K_i\phi \in V(v)$ (Monotonicity)

D3 For every ϕ , if $K_i\phi \in V(w)$, then $M, w \models \phi$ (Knowledge)

D4 If $K_i\phi, K_i(\phi \rightarrow \psi) \in V(w)$, then $\exists v(R_i(w, v) \ \& \ K_i\psi \in V(v))$

D5 For every $\tau \in Ax_i$, $\exists v(R_i(w, v) \ \& \ K_i\tau \in V(v))$

D6 If $K_i\phi \in V(w)$ and $\phi \in \mathcal{L}_E^+$, then $\exists v(R_i(w, v) \ \& \ K_iK_i\phi \in V(v))$.

As an intermediate result, we first give soundness and completeness of $DES4_n$ with respect to the class of general models (that is, models where the epistemic states do not have to be finite), satisfying conditions **D1-D6**, $\mathcal{M}^{gen,DEL}$. A similar result was proved in [9]. We then prove the main, and more difficult, result: completeness with respect to \mathcal{M}^{DEL} .

THEOREM 22

$DES4_n$ is sound and strongly complete with respect to $\mathcal{M}^{gen,DEL}$.

PROOF. Appendix A. ■

THEOREM 23

$DES4_n$ is sound and weakly complete with respect to \mathcal{M}^{DEL} .

PROOF. Appendix B. ■

5.3 *Dynamic syntactic epistemic logic*

DSEL [3] is, as mentioned in the introduction, 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, based on ATL. In Section 4, we studied how to express reasoning rules, such as modus ponens, in our framework, by using the operators Δ, ∇, \diamond , etc. One of the distinguishing features of DSEL is that it has syntactic objects standing for rules, allowing referring to rules directly in the logical language. The logical language has operators similar to Δ and ∇ which take sets of rules, rather than sets of formulae, as arguments, in order to express properties of reasoning similar to Δ_iMP and ∇_iMP from Section 4. However, the exact notions of knowing at least and knowing at most a set of rules are different from the ones we discussed in Section 4—exactly how will be clear momentarily—and we shall in this section apply our framework to characterise these notions. There is no complete axiomatization of DSEL, and we saw in Section 4 that complete axiomatization of concepts related to ‘knowing at most’ a rule is difficult, so we limit the discussion here to defining the corresponding model classes with our language rather than investigating completeness.

In order to avoid too many technical details, we do not present the formal syntax of DSEL rules here. It suffices here to know that there exist a set RUL of rules (these are syntactic objects that can be parts of formulae, rather than semantic objects), and an *interpretation function* $\llbracket \cdot \rrbracket$ interpreting each set of rules $R \subseteq RUL$ as a relation over finite sets of object language formulae:

$$\llbracket R \rrbracket \subseteq \wp^{fin}(OL) \times \wp^{fin}(OL).$$

Intuitively, when $(X, Y) \in \llbracket R \rrbracket$, X and Y are a sets of formulae corresponding to the antecedent and consequent, respectively, of some rule in R . An important difference between DSEL rules and the type of rules we discussed in Section 4 is that the former describe *complete* epistemic states—they say that the rule (or set of rules) R can take you from knowing exactly the

formulae in a set X to knowing exactly the formulae in a set Y ($(X, Y) \in \llbracket R \rrbracket$), while the latter describe what we called addition constraints.

The following are the DSEL variants of knowing at least/at most a (set of) rule(s) R (the semantic framework of DSEL is different from ours, so the following are reasonable re-interpretations of the DSEL concepts in our semantics). We say that an agent who knows exactly formulae X knows *at least* rules R for reasoning (in the DSEL sense) 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 for reasoning if for every world he can access where he knows exactly C , $(X, C) \in \llbracket R \rrbracket$. Formally, let $M = (W, R_1, \dots, R_n, V)$ be a model and R a set of rules. In a state $w \in W$:

$$\begin{aligned} \text{ATLEAST} - R & \quad \overline{(V_i(w), s')} \in \llbracket R \rrbracket \Rightarrow \exists_{(w, w') \in R_i} \overline{V_i(w')} = s' \\ & \quad \text{Agent } i \text{ knows at least the rule(s) } R \\ \text{ATMOST} - R & \quad \overline{(V_i(w), s')} \in \llbracket R \rrbracket \Leftarrow \exists_{(w, w') \in R_i} \overline{V_i(w')} = s' \\ & \quad \text{Agent } i \text{ knows at most the rule(s) } R \\ \text{EXACTLY} - R & \quad \overline{(V_i(w), s')} \in \llbracket R \rrbracket \Leftrightarrow \exists_{(w, w') \in R_i} \overline{V_i(w')} = s' \\ & \quad \text{Agent } i \text{ knows exactly the rule(s) } R \end{aligned}$$

In DSEL, these three semantic properties are expressible directly by expressions $\tilde{\Delta}_i R$, $\tilde{\nabla}_i R$ and $\tilde{\boxtimes}_i R$, respectively, where $\tilde{\Delta}_i$, $\tilde{\nabla}_i$, $\tilde{\boxtimes}_i$ are special operators that work on rules. The question is whether such properties can be expressed indirectly in \mathcal{L}_n^∇ , without an explicit syntax for rules and operators that work on rules. It is easy to see that none of the three properties given above are expressible by a single \mathcal{L}_n^∇ formula. We instead give characterizations by translating a set of rules to a set of \mathcal{L}_n^∇ formulae.

DEFINITION 9

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

$$\begin{aligned} f(R) & = \{\tilde{\boxtimes}_i X \rightarrow \diamond_i \tilde{\boxtimes}_i Y : (X, Y) \in \llbracket R \rrbracket\} \\ g(R) & = \{\tilde{\boxtimes}_i X \rightarrow \square_i \neg \tilde{\boxtimes}_i Y : (X, Y) \notin \llbracket R \rrbracket\} \end{aligned}$$

THEOREM 24

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 $\overline{(V_i(w), s')} \in \llbracket R \rrbracket$. $M, w \models \tilde{\boxtimes}_i \overline{V_i(w)}$, so $M, w \models \diamond_i \tilde{\boxtimes}_i s'$, i.e. there exists $(w, w') \in R_i$ such that $M, w \models \tilde{\boxtimes}_i s'$ and thus $s' = \overline{V_i(w')}$.
 \Leftarrow) Let ATLEAST-R hold in w , let $(X, Y) \in \llbracket R \rrbracket$ and let $M, w \models \tilde{\boxtimes}_i X$. Since $\overline{(V_i(w), Y)} \in \llbracket R \rrbracket$, there exists $(w, w') \in R_i$ such that $Y = \overline{V_i(w')}$. Thus, $M, w \models \diamond_i \tilde{\boxtimes}_i Y$.
- (2) \Rightarrow) Let $M, w \models g(R)$ and $(w, w') \in R_i$ and $s' = \overline{V_i(w')}$. $M, w \models \diamond_i \tilde{\boxtimes}_i s'$. Assume that $\overline{(V_i(w), s')} \notin \llbracket R \rrbracket$. $M, w \models \tilde{\boxtimes}_i \overline{V_i(w)}$, so $M, w \models \square_i \neg \tilde{\boxtimes}_i s'$, which is a contradiction. Thus, $\overline{(V_i(w), s')} \in \llbracket R \rrbracket$.

\Leftarrow) 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 $\forall_{(w, w') \in R_i} Y \notin \overline{V_i(w')}$. 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 (ATOMOST-R) holds in every state.

5.3.1 Communication

In DSEL, reasoning is seen as a generalization of communication: reasoning is ‘communication to oneself’. Thus, rules can be used to reason about both concepts.

However, there is a difference in the intuitive interpretation of knowing a reasoning rule and knowing a communication rule. In the former case, formalized earlier, 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:

$$\begin{aligned} \text{ATLEAST} - R' & \quad (\overline{V_i(w)}, s') \in \llbracket R \rrbracket \Rightarrow \exists_{(w, w') \in R_i} s' \subseteq \overline{V_j(w')} \\ & \quad \text{Agent } i \text{ knows at least the rule(s) } R \text{ for communication to } j \\ \text{ATOMOST} - R' & \quad (\overline{V_i(w)}, s') \in \llbracket R \rrbracket \Leftarrow \exists_{(w, w') \in R_i} s' \subseteq \overline{V_j(w')} \\ & \quad \text{Agent } i \text{ knows at most the rule(s) } R \text{ for communication to } j \\ \text{EXACTLY} - R' & \quad (\overline{V_i(w)}, s') \in \llbracket R \rrbracket \Leftrightarrow \exists_{(w, w') \in R_i} s' \subseteq \overline{V_j(w')} \\ & \quad \text{Agent } i \text{ knows exactly the rule(s) } R \text{ for communication to } j \end{aligned}$$

A corresponding translation is:

DEFINITION 10

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

$$\begin{aligned} f'(R) & = \{\boxtimes_i X \rightarrow \diamond_i \Delta_j Y : (X, Y) \in \llbracket R \rrbracket\} \\ g'(R) & = \{\boxtimes_i X \rightarrow \square_i \neg \Delta_j Y : (X, Y) \notin \llbracket R \rrbracket\} \end{aligned}$$

THEOREM 25

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) $M, w \models g'(R) \Leftrightarrow \text{ATOMOST} - R'$ 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 24. ■

6 Related work

In the introduction we discussed some classes of logics which use the syntactic approach. In the previous section, we established certain relationships between our framework and three other logics describing the dynamics of syntactically ascribed knowledge. It was shown that we can completely axiomatize an interesting timed reasoning logic [6] which was in turn inspired by

step logics [13]. We also saw that our framework can also be seen as a generalization of DEL [10]. When it comes to DSEL [1], in Section 5 we only applied our framework to express knowledge of reasoning and communication rules. DSEL does, however, also have temporal operators from ATL which cannot be expressed in our framework, so we cannot completely describe DSEL. Aside from that, there are other interesting differences in the semantic assumptions of DSEL as well. 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.

Another related strand of work which we did not analyse earlier is the work by Sierra *et al.* [30] on architectures of multiunit knowledge bases. Their work has different motivations, but quite similar outcome to our approach to epistemic logic.

Sierra *et al.* propose to use a version of dynamic logic, which they call Descriptive Dynamic Logic (DDL), to describe a system of communicating ‘units’, each with a different initial logical language and theory and its own internal inference rules; in addition there are ‘bridge rules’ which allow units to exchange information.

In each state, units are assigned finite sets of formulae in their languages; we would call those formulae beliefs, Sierra *et al.* call them theories of the units. A significant restriction made in [30] is that not just the theories of units, but also the sets of all possible formulae in the units’ languages are finite.

Internal actions of units in DDL correspond to ‘deduction steps’ in the internal logic of the unit, that is, all possible instances of applications of inference rules (denoted as $[\Gamma \vdash_{kk} \phi]$, where k is the name of the unit, $\Gamma \cup \{\phi\}$ is a set of formulae in k ’s language, and k has an inference rule by which Γ derives ϕ). Inter-unit actions correspond to bridge rules, so that if there is a bridge rule by which Γ in the theory of k implies ϕ in the theory of l , then this instance of its application corresponds to an operator $[\Gamma \vdash_{kl} \phi]$.

DDL is given a complete axiomatization parametrized by unit’s languages and inference rules. This results relies on having the unit’s languages finite and cannot be used to axiomatize the logics considered in this article, as far as we can see.

The work on logics of awareness [14, 15] and unawareness [19] can arguably also be classified as a syntactic approach to belief, since a syntactic awareness filter is superimposed on the agent’s beliefs which are defined using the possible worlds approach. However, the main difference between our approach and the logics of awareness is that the latter abstract from dynamics of beliefs and do not model the evolution of the awareness set as a result of the agent’s inference or computation. The same applies to logics of algorithmic knowledge [18, 16, 29], where the agent’s beliefs are modelled as closed under the agent’s ability to compute consequences (which is elegantly modelled in [29] as a set of rewrite rules).

7 Conclusions

We combined the syntactic approach to epistemic logic with modal logics in order to model transitions between states of knowledge. The two epistemic operators Δ and ∇ , were used to express (finite) lower and upper bounds on knowledge, respectively. We were mostly interested in models where agents are restricted to only knowing finitely many formulae at each point in time. We studied axiomatization and complexity of the resulting logics. In particular, and perhaps most interesting from a logical point of view, we demonstrated techniques for dealing with the finiteness assumption and the ∇ operator, separately and combined.

We have thus 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 these logics 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 axiomatize several interesting classes of transition systems in them. One of the classes provides a semantics for the logic introduced in [10]. 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

We would like to thank the anonymous referees for their useful comments and suggestions.

The work presented in this article 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 the Royal Society for sponsoring related work on logics for resource-bounded agents, and the Isaac Newton Institute for hospitality during work on the final version of the article.

References

- [1] T. Ågotnes. *A Logic of Finite Syntactic Epistemic States*. PhD thesis, Department of Informatics, University of Bergen, Norway, 2004.
- [2] T. Ågotnes and N. Alechina. Knowing minimum/maximum n formulae. In *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI 2006)*, in press.
- [3] T. Ågotnes. and M. Walicki. Syntactic knowledge: A logic of reasoning, communication and cooperation. In *Proceedings of the 2nd European Workshop on Multiagent Systems (EUMAS)*, C. Ghidini, P. Giorgini and W. van der Hoek, eds, pp. 1–10, Barcelona, Spain, December 2004.
- [4] T. Ågotnes and M. Walicki. Complete axiomatizations of finite syntactic epistemic states. In *Declarative Agent Languages and Technologies III: Third International Workshop, DALT 2005, Utrecht, The Netherlands, July 25, 2005, Selected and Revised Papers, volume 3904 of Lecture Notes in Computer Science (LNCS)*, M. Baldoni, U. Endriss, A. Omicini and P. Torroni, eds, pp. 33–50, Springer, Berlin/Heidelberg, 2006.
- [5] T. Ågotnes and M. Walicki. Strongly complete axiomatizations of “knowing at most” in syntactic structures, In *CLIMA VI, volume 3900 of Lecture Notes in Computer Science (LNCS)*, F. Toni and P. Torroni, eds, pp. 57–76, London, UK, Springer, Berlin/Heidelberg, June 2006.
- [6] N. Alechina, B. Logan and M. Whitsey. A complete and decidable logic for resource-bounded agents. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*. pp. 606–613, ACM Press, New York, 2004.
- [7] R. Alur, T. A. Henzinger and O. Kupferman. Alternating-time temporal logic. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, pp. 100–109, Florida, Miami Beach, IEEE, New York, 20–22 October 1997.

- [8] P. Blackburn, M. deRijke and Y. Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge, 2001.
- [9] H. N. Duc. Logical omniscience vs. logical ignorance on a dilemma of epistemic logic. In *Progress in Artificial Intelligence, 7th Portuguese Conference on Artificial Intelligence, EPIA '95, Funchal, Madeira Island, Portugal, October 3-6, 1995, Proceedings*, volume 990 of *Lecture Notes in Computer Science*, C. A. Pinto-Ferreira and N. J. Mamede, eds, pp. 237–248, Springer, Berlin, 1995.
- [10] H. N. Duc. Reasoning about rational, but not logically omniscient, agents. *Journal of Logic and Computation*, **7**, 633–648, 1997.
- [11] R. A. Eberle. A logic of believing, knowing and inferring. *Synthese*, **26**, 356–382, 1974.
- [12] J. Elgot-Drapkin, S. Kraus, M. Miller, M. Nirkhe and D. Perlis. Active logics: A unified formal approach to episodic reasoning. *Technical Report CS-TR-4072 and UMIACS-TR-99-65*, University of Maryland, 1999.
- [13] J. J. Elgot-Drapkin and D. Perlis. Reasoning situated in time I: Basic concepts. *Journal of Experimental and Theoretical Artificial Intelligence*, **2**, 75–98, 1990.
- [14] R. Fagin and J. Y. Halpern. Belief, awareness and limited reasoning. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pp. 491–501, Los Angeles, CA, Morgan Kaufmann, Los Altos, California, 1985.
- [15] R. Fagin and J. Y. Halpern. Belief, awareness and limited reasoning. *Artificial Intelligence*, **34**, 39–76, 1988. A preliminary version appeared in [14].
- [16] R. Fagin, J. Y. Halpern, Y. Moses and M. Y. Vardi. *Reasoning About Knowledge*. The MIT Press, Cambridge, MA, 1995.
- [17] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, **37**, 549–587, 1990.
- [18] J. Y. Halpern, Y. Moses and M. Y. Vardi. Algorithmic knowledge. In *Proceedings of the 5th Conference on Theoretical Aspects of Reasoning about Knowledge, Pacific Grove, CA, USA, March 1994*, R. Fagin, ed., pp. 255–266, Morgan Kaufmann, Los Altos, California, 1994.
- [19] J. Y. Halpern and L. C. Rêgo. Reasoning about knowledge of unawareness, In *Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning*, P. Doherty, J. Mylopoulos and C. A. Welty, eds, 6–13, Lake District of the United Kingdom, AAAI Press, Menlo Park, California, June 2–5, 2006.
- [20] D. Harel, D. Kozen and J. Tiuryn. *Dynamic Logic*. The MIT Press, Cambridge, MA, 2000.
- [21] J. Hintikka. *Knowledge and Belief*. Cornell University Press, Ithaca, NY, 1962.
- [22] J. Hintikka. Impossible possible worlds vindicated. *Journal of Philosophical Logic*, **4**, 475–484, 1975.
- [23] K. Konolige. *A Deduction Model of Belief*, Morgan Kaufmann, Los Altos, California, 1986.
- [24] H. J. Levesque. All I know: a study in autoepistemic logic. *Artificial Intelligence*, **42**, 263–309, 1990.
- [25] J. -J. C. Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science*. Cambridge University Press, Cambridge, UK, 1995.
- [26] R. C. Moore and G. G. Hendrix. Computational models of beliefs and the semantics of belief sentences. Technical Note 187, SRI International, Menlo Park, CA, 1979.
- [27] A. Moreno. Avoiding logical omniscience and perfect reasoning: a survey. *AI Communications*, **11**, 101–122, 1998.

- [28] I. Németi. Fine-structure analysis of first-order logic. In *Arrow logic and multi-modal logic*, M. Masuch and L. Pólos, eds, pp. 221–247, CSLI Publications, Stanford, 1996.
- [29] R. Pucella Deductive algorithmic knowledge. In *proceeding of the AI&M 1-2004, Eighth International Symposium on Artificial Intelligence and Mathematics, January 4-6, 2004*. Fort Lauderdale, Florida, USA, 2004.
- [30] C. Sierra, L. Godo, R. López deMántaras and M. Manzano. Descriptive dynamic logic and its application to reflective architectures. *Future Gener. Comput. Syst.*, **12**, 157–171, 1996.
- [31] K. M. Sim. Epistemic logic and logical omniscience: A survey. *International Journal of Intelligent Systems*, **12**, 57–81, 1997.
- [32] W. van der Hoek and M. Wooldridge. Towards a logic of rational agency. *Logic Journal of the IGPL*, **11**, 135–159, 2003.
- [33] H. P. van Ditmarsch, W. van der Hoek and B. P. Kooi. Dynamic epistemic logic. Springer, Berlin, in press.

Appendix

A Proof of Theorem 22

First, let us consider soundness. We want to prove that for every set of formulae Γ and formula ϕ_0 , if $\Gamma \vdash_{DES4_n} \phi_0$, then $\Gamma \models_{gen,DEL} \phi_0$, where $\Gamma \models_{gen,DEL} \phi_0$ denotes logical consequence with respect to $\mathcal{M}^{gen,DEL}$ ($M, w \models \Gamma$ implies that $M, w \models \phi_0$ for every $M \in \mathcal{M}^{gen,DEL}$). The proof is by induction on the length of the derivation of ϕ_0 from Γ . Clearly, the inference rules preserve validity. We need to show that every instance of an axiom schema of $DES4_n$ is valid on the class $\mathcal{M}^{gen,DEL}$. The axiom schemes $PC1 - PC3$ are valid classical tautologies. It is well known that $TL1$ is valid in all Kripke models and that $TL2$ is valid in all models with a transitive accessibility relation. $DE1$ is valid because of condition **D4**: if some state w satisfies $K_i\phi$ and $K_i(\phi \rightarrow \psi)$, then by **D4**, w has a successor which satisfies $K_i\psi$, so w satisfies $\langle F_i \rangle K_i\psi$. $DE2$ is valid because of the condition **D3**: if $K_i\phi$ is true in w , then ϕ has to be true in w . $DE3$ is valid due to the Monotonicity condition **D2**. $DE4 - DE7$ are valid because of **D5**. Finally, $DE8$ is valid because of **D6**.

The proof of completeness proceeds in a standard way. Suppose $\Gamma_0 \not\vdash_{DES4_n} \phi_0$; we show that then there exists a model and a state where Γ_0 is satisfied and ϕ_0 is not, so $\Gamma_0 \not\models_{gen,DEL} \phi_0$. In other words, we show how to construct a satisfying model for $\Gamma_0 \cup \{\neg\phi_0\}$ provided this set is $DES4_n$ -consistent. For convenience, we will refer to this set as Γ . Any $DES4_n$ -consistent set can be extended to a maximal consistent set in a standard way. Let $M^c = (W^c, R_1^c, \dots, R_n^c, V^c)$ be the canonical model for $DES4_n$. Recall that M^c is a general model; we must show that $M^c \in \mathcal{M}^{gen,DEL}$. It is straightforward to show that each R_i is transitive, because each state in the model contains $TL2$. Now we have to show that conditions **D2-D6** hold. Suppose for some w , $K_i\phi$ and $K_i(\phi \rightarrow \psi) \in V^c(w)$. This means, by the definition of V^c , that $K_i\phi, K_i(\phi \rightarrow \psi) \in w$. Since w is maximal, it also contains $K_i\phi \wedge K_i(\phi \rightarrow \psi)$ and $K_i\phi \wedge K_i(\phi \rightarrow \psi) \rightarrow \langle F_i \rangle K_i\psi$ ($DE1$). Since it is closed under inference, it also contains $\langle F_i \rangle K_i\psi$. By the truth lemma, $\langle F_i \rangle K_i\psi$ is true in w . Hence w has a successor v where $K_i\psi$ is true, which by the truth lemma implies $K_i\psi \in V^c(v)$. Analogously, we can show that **D3-D6** hold. This completes the argument that $M^c \in \mathcal{M}^{gen,DEL}$; since Γ is consistent, it is contained in one of the states in M^c , and by the truth lemma, is satisfied there.

B Proof of Theorem 23

The proof of soundness is identical to the proof in Theorem 22.

For completeness, assume that ϕ_0 is a $DES4_n$ -consistent formula. We will construct a model satisfying ϕ_0 , in several stages. First we will construct a general model M^c satisfying ϕ_0 . Then we will transform this model into a proper model with finite epistemic states. An overview of the transformations starting with M^c and ending with a model M with finite epistemic states satisfying ϕ_0 is as follows:

$$M^c \xrightarrow{\text{generated submodel}} M^g \xrightarrow{\text{unravelling}} M^t \xrightarrow{\text{transitive closure}} M^t \xrightarrow{\text{restriction of valuations}} M$$

Let $M^c = (W^c, R_1^c, \dots, R_n^c, V^c)$ be as in the proof of Theorem 22 (the canonical model for $DES4_n$). We showed that M^c satisfies **D1–D6**, but it may have infinite epistemic states and thus not be in the class \mathcal{M}^{DEL} . At least one of the states in M^c satisfies ϕ_0 ; let us call that state w_0 . We will now unravel M^c around w_0 , and then take the transitive closure; and we will end up with a transitive tree (general) model M^t . The unravelling around w_0 is defined as follows, cf. [8], pp. 220–221 for further discussion. Let $M^g = (W^g, R_1^g, \dots, R_n^g, V^g)$ be the (general) submodel of M^c generated by w_0 : it is the smallest submodel of M^c such that $w_0 \in W^g$ and $u \in W^g$ whenever $w \in W^g$ and $R_i^c(w, u)$ for some i . The *unravelling* of the general model M^g around $w_0 \in W^g$ is the general model $M^t = (W^t, R_1^t, \dots, R_n^t, V^t)$ such that:

- W^t is the set of all finite sequences (w_0, w_1, \dots, w_n) , $n \geq 0$, such that $w_1, \dots, w_n \in W^g$ and $R_{i_1}^g(w_0, w_1), R_{i_2}^g(w_1, w_2), \dots, R_{i_n}^g(w_{n-1}, w_n)$ for some i_1, \dots, i_n .
- $R_i^t(w_1, \dots, w_k)(v_1, \dots, v_l)$ iff $(w_1, \dots, w_k) = (w_1, \dots, v_{l-1})$ and $R_i^g(v_{l-1}, v_l)$
- $V^t((w_0, w_1, \dots, w_n)) = V^g(w_n)$

The result of unravelling around w_0 is an intransitive tree where w_0 is the root. The *transitive unravelling* $M^t = (W^t, R_1^t, \dots, R_n^t, V^t)$ is obtained by taking the transitive closure of each of the relations in the unravelling M^t . Note the following properties of the transitive unravelling:

$$R_i^t((w_0, \dots, w_k), (v_0, \dots, v_l)) \Rightarrow R_i^c(w_k, v_l) \quad (\text{A.1})$$

$$R_i^t(r_1, r_2) \Rightarrow r_1 \text{ is a prefix of } r_2 \quad (\text{A.2})$$

We argue that (A.1) and (A.2) hold. Let $r_1 = (w_0, \dots, w_k)$ and $r_2 = (v_0, \dots, v_l)$. If $R_i^t(r_1, r_2)$, then there exist a path $s_0, \dots, s_j \in W^t$, $j \geq 0$, of R_i^t steps between the two sequences; i.e., such that $R_i^t(r_1, s_0), \dots, R_i^t(s_j, r_2)$. In turn, this means that there exist $u_0, \dots, u_j \in W^g$ such that

$$\begin{array}{ll} s_0 = (w_0, \dots, w_k, u_0) & R_i^g(w_k, u_0) \\ \vdots & \\ s_j = (w_0, \dots, w_k, u_0, \dots, u_j) & R_i^g(u_{j-1}, u_j) \\ (v_0, \dots, v_l) = (w_0, \dots, w_k, u_0, \dots, u_j, v_l) & R_i^g(u_j, v_l) \end{array}$$

The previous equation shows that (A.2) holds (r_1 is a *proper* prefix of r_2 since R_i^t , and thus R_i^t , is irreflexive). Since R_i^c is transitive and includes R_i^g , it also follows that $R_i^c(w_k, v_l)$ and thus

(A.1) holds. We now show that M^t and M^c are bisimilar. Let $Z \subseteq W^t \times W^c$ be defined as follows:

$$(w_0, \dots, w_n)Zw_n$$

We show that Z is a bisimulation between M^t and M^c , by the three required conditions (Definition 3):

- (1) $V^t(w_0, \dots, w_k) = V^t(w_0, \dots, w_k) = V^g(w_k) = V^c(w_k)$
- (2) Immediate by (1).
- (3) We must show that if $(w_0, \dots, w_k) \in W^t$ and $R_i^c(w_k, v)$ for some $v \in W^c$, then there exists a sequence $(v_0, \dots, v_l) \in W^t$ with $v_l = v$ such that $R_i^t((w_0, \dots, w_k), (v_0, \dots, v_l))$. This holds immediately by taking $(v_0, \dots, v_l) = (w_0, \dots, w_k, v)$: $w_k \in W^g$ and thus $v \in W^g$ and thus $R_i^g(w_k, v)$, and $R_i^t((w_0, \dots, w_k)(w_0, \dots, w_k, v))$.

Since M^c and M^t are bisimilar, for every state w reachable from w_0 , $M^c, w \models \phi$ iff $M^t, (s, w) \models \phi$, for every formula ϕ and sequence $s \in W^t$ (we use (s, w) to denote the sequence resulting from concatenating the element w to the sequence s). This includes the special case when s is empty, $w = w_0$, and $\phi = \phi_0$. Since for every state (s, w) in M^t there is a bisimilar state w in M^c , we can also show that all of the conditions **D1–D6** are satisfied in M^t . **D1** (transitivity) holds immediately for M^t . **D2** (monotonicity) holds because if $K_i\phi \in V^t((s, w))$, and $\phi \in \mathcal{L}_E^+$, and $R_i^t((s, w), (s', v))$, then $K_i\phi \in V^c(w)$ (by construction of V^t), $R^c(w, v)$ (by (1)), and $K_i\phi \in V^c(v)$ (because M^c satisfies **D2**). By the bisimulation between (s', v) and v , $K_i\phi \in V^t((s', v))$. Similarly for **D3**: suppose $K_i\phi \in V^t((s, w))$. Then $K_i\phi \in V^c(w)$. Then, since M^c satisfies **D3**, $M^c, w \models \phi$. By the bisimulation, $M^t, (s, w) \models \phi$. **D4–D6** hold because for every $(s, w) \in W^t$, if (s, w) satisfies a certain formula, then it has a bisimilar state w in W^c , which satisfies the same formula and has a successor v which satisfies another required formula; by the bisimulation, (s, w) then also has a successor (s, w, v) which satisfies the same formula.

So, the root w_0 of M^t satisfies ϕ_0 ; conditions **D1–D6** continue to hold. Now we are going to transform M^t into a proper model, by intersecting the epistemic state in all states in M^t with a finite set of formulae; we need to do this in such a way that ϕ_0 is still satisfied at the root and conditions **D1–D6** continue to hold. More precisely, the epistemic states at level k in the tree model M^t (where w_0 is at level 0, and its one step successors at level 1, etc.) are going to be intersected with a finite set of formulae L_k , defined inductively as follows. Note that there are infinitely many levels in the tree due to seriality imposed by *DE4–DE7*.

Let henceforth $At(\phi)$ denote the set of primitive propositions and epistemic atoms which occur outside the scope of any epistemic operators in a formula ϕ . For example, $At(((F_i)K_iK_i(p \wedge q) \rightarrow p) \wedge p) = \{K_iK_i(p \wedge q), p\}$ —this set does *not* contain $K_i(p \wedge q)$ nor q since both the epistemic atom $K_i(p \wedge q)$ and the primitive proposition q only occur within the scope of an epistemic operator. In the following definition, given a set of epistemic atoms $X \subseteq \Theta$, $Cl(X)$ denotes the closure of X under nested epistemic atoms: $Cl(X)$ is the least set such that (i) $X \subseteq Cl(X)$ and (ii) if $K_i\phi \in Cl(X)$ then $At(\phi) \subseteq Cl(X)$.

$$L_0 = Cl(At(\phi_0))$$

$L_k = Cl(Inf(L_{k-1} \cup \{K_1\tau_{k-1}^1, \dots, K_n\tau_{k-1}^n\}))$, where, for each agent i , $\tau_1^i, \tau_2^i, \dots$ is some enumeration of the (countable) set Ax_i , and $Inf(X)$, for some set of epistemic atoms X , is the set containing X and formulae derived from X by a single application of the following rules:

$$\frac{\frac{K_i\phi, K_i(\phi \rightarrow \psi)}{K_i\psi} \{K_i\phi\}}{K_iK_i\phi}$$

Note that if X is finite, so is $Cl(Inf(X))$, and thus that L_k is finite for each k . Consider the resulting model $M = (W, R_1, \dots, R_n, V)$: $W = W^t$, $R_i = R_i^t$ and

$$V(w_0, \dots, w_k) = V^t(w_0, \dots, w_k) \cap L_k$$

We need to prove the following statements about M :

Truth $M, w_0 \models \phi_0$

Finiteness Epistemic states in M are finite.

D1–D6 Conditions **D1–D6** are satisfied.

Here is the proof for each of these statements.

Truth M and M^t are $At(\phi_0)$ -bisimilar (take Z to be the identity relation on W^t).

Theorem 1 says that formulae only containing atoms from $At(\phi_0)$ are preserved under $At(\phi_0)$ -bisimulation. Since M^t, w_0 and M, w_0 agree on $At(\phi_0)$ and $M^t, w_0 \models \phi_0$, it follows that $M, w_0 \models \phi_0$.

Finiteness $At(\phi_0)$ is finite, and each epistemic state is a subset of a finite set L_k .

D1 the accessibility relation in M is transitive, because it is transitive in M^t .

D2 Let $R_i(w, v)$, $\phi \in \mathcal{L}_E^+$, $K_i\phi \in V(w)$. We need to show that $K_i\phi \in V(v)$. We have that $R_i^t(w, v)$, so by (A.2) the sequence w is a prefix of the sequence v ; say $w = (v_0, \dots, v_k)$ and $v = (v_0, \dots, v_l)$ with $k \leq l$. $K_i\phi \in V^t(w)$, and since Monotonicity holds in M^t , $K_i\phi \in V^t(v)$. We also have that $K_i\phi \in L_k$, and since $L_k \subseteq L_l$, $K_i\phi \in L_l$. Thus, $K_i\phi \in V(v) = V^t(v) \cap L_l$.

D3 Let $K_i\phi \in V(w) = V^t(w) \cap L_k$, where $w = w_0, \dots, w_k$. We must show that $M, w \models \phi$. We first point out that

$$Cl(\{\phi'\}) \subseteq L_k \Rightarrow (M, w \models \phi' \Leftrightarrow M^t, w \models \phi')$$

holds for all ϕ' by induction over the structure of ϕ' . We have that $M^t, w \models K_i\phi$ and thus $M^t, w \models \phi$ by Knowledge for M^t . We also have that $K_i\phi \in L_k$ and thus that $Cl(\{\phi\}) \subseteq L_k$. It follows that $M, w \models \phi$.

D4 Let $K_i\phi, K_i(\phi \rightarrow \psi) \in V(w) = V^t(w) \cap L_k$, where $w = (w_0, \dots, w_k)$. By **D4** for M^t , there is a $v = (v_0, \dots, v_l)$ such that $R_i^t(w, v)$ and $K_i\psi \in V^t(v)$. By (2) $l \geq k$, and since $K_i\phi, K_i(\phi \rightarrow \psi) \in L_k$ it follows that $K_i\psi \in L_l$ by construction of L_{k+1} . Thus, $K_i\psi \in V(v) = V^t(v) \cap L_l$. Since $R_i^t(w, v)$, $R_i(w, v)$.

D5 Let $\tau \in Ax_i$ and $w = (w_0, \dots, w_k) \in W$. We must show that there is a $v = (v_0, \dots, v_l) \in W$ such that $R_i(w, v)$ and $K_i\tau \in V(v) = V^l(v) \cap L_l$. τ is one of the elements in the enumeration of Ax_i in the definition of L_k , say $\tau = \tau_j^i$. We first show that for any $u \in W$:

$$u = (u_0, \dots, u_m), m \geq j \Rightarrow \exists u' \in W \begin{cases} R_i(u, u') \\ K_i\tau \in V(u') \end{cases} \quad (A.3)$$

By **D5** for M^c , there is a $u_{m+1} \in W^c$ such that $R_i^c(u_m, u_{m+1})$ and $K_i\tau \in V^c(u_{m+1})$. Let $u' = (u_0, \dots, u_m, u_{m+1})$. $R_i^g(u_m, u_{m+1})$ ($u_m \in W^g$), so $R_i^l(u, u')$ and thus $R_i^l(u, u')$ and $R_i(u, u')$. $K_i\tau \in V^g(u_{m+1})$, so $K_i\tau \in V^l(u') = V^l(u')$. $K_i\tau = K_i\tau_j^i \in L_j$, and, since $m \geq j$, $L_j \subseteq L_{m+1}$, we have that $K_i\tau_j^i \in V^l(u') \cap L_{m+1} = V(u')$ and (A.3) holds. Now in the case that $k \geq j$, we are done by (A.3). Let $k < j$. R_i^g is serial (**D5** holds for M^g), so there exist $w_{k+1}, \dots, w_j \in W^g$ such that $R_i^g(w_k, w_{k+1}), \dots, R_i^g(w_{j-1}, w_j)$. Thus, $R_i^l(w, (w_0, \dots, w_k, w_{k+1})), \dots, R_i^l((w_0, \dots, w_{j-1}), (w_0, \dots, w_{j-1}, w_j))$. By transitivity, $R_i^l(w, (w_0, \dots, w_j))$. By (A.3) there is a u' such that $R_i((w_0, \dots, w_j), u')$ and $K_i\tau \in V(u')$, and by transitivity again we get that $R_i(w, u')$ and we are done.

D6 Let $w = (w_0, \dots, w_k) \in W$, $\phi \in \mathcal{L}_E^+$ and $K_i\phi \in V(w) = V^l(w) \cap L_k$. $K_i\phi \in V^l(w) = V^l(w)$. Since M' satisfies **D6**, there is a $v \in W'$ such that $R_i^l(w, v)$ and $K_iK_i\phi \in V^l(v)$. It must be the case that $v = (w_0, \dots, w_k, w_{k+1})$ for some $w_{k+1} \in W'$. It follows that $R_i^l(w, v)$ and $K_iK_i\phi \in V^l(v)$. Again, it follows that $R_i(w, v)$. By construction of L_{k+1} , since $K_i\phi \in L_k$, we have that $K_iK_i\phi \in L_{k+1}$. Thus, $K_iK_i\phi \in V(v) = V^l(v) \cap L_{k+1}$.

Received 26 September 2005