

Reduction axioms for epistemic actions

Barteld Kooi* and Johan van Benthem†

Abstract: Current dynamic epistemic logics often become cumbersome and opaque when common knowledge is added for groups of agents. We propose new versions that extend the underlying static epistemic languages in such a way that completeness proofs for the full dynamic systems can be obtained by perspicuous reduction axioms.

Keywords: dynamic logic, epistemic logic.

1 Introduction

Epistemic logic typically deals with what agents consider possible given their current information. This includes knowledge about facts, but also *higher-order information* about information that other agents have. A prime example is *common knowledge*. A formula φ is common knowledge if everybody knows φ , everybody knows that everybody knows φ , and so on.

The aim of dynamic epistemic logics is to analyze changes in basic and higher-order information. Completeness proofs for such logics are either easy, or hard. For instance, the logic of public announcements without common knowledge has an easy completeness proof due to axioms such as $[\varphi]\Box_i\psi \leftrightarrow (\varphi \rightarrow \Box_i[\varphi]\psi)$. We call these reduction axioms, because the announcement operator is “pushed through” the epistemic operators. The completeness proof works by way of a translation that follows the reduction axioms. Formulas with announcements are translated to provably equivalent ones without announcements. Then completeness follows from the known completeness of the epistemic base logic. This approach also is taken in [2] and [1] for more general epistemic actions.

Completeness proofs for dynamic epistemic logics with common knowledge are hard. Reduction axioms are not available, as the logic with epistemic actions is more expressive than the logic without them [1]. In this paper we extend the base language with static operators in such a way that reduction axioms do work. Section 2 does this for public announcement logic, Section 3 for general epistemic actions. Section 4 draws conclusions and indicates directions for further research. We see our proposal as more than a technical trick for smoothening completeness proofs. It also addresses a significant issue of independent interest: what is the best epistemic language for describing information models of a group of agents?

2 Public announcement logic

Section 2.1 is an introduction to public announcement logic. In Section 2.2 we give a new logic of *relativized common knowledge*. It ties in closely to the idea of viewing updates as a kind of relativization, first introduced in [6]. This logic is expressive enough to allow a reduction axiom for common knowledge. A proof system is defined in Section 2.3, and shown to be complete in Section 2.4. The system is extended with reduction axioms for public announcements in Section 2.5.

*Department of Philosophy, University of Groningen, A-weg 30, 9718 CW Groningen, The Netherlands, barteld@philos.rug.nl

†ILLC, University of Amsterdam, Plantage Muidergracht 24, 1018 TV Amsterdam, The Netherlands & Philosophy Department, Stanford University, johan@science.uva.nl

2.1 Language and semantics

Public announcement logic (PAL) was first developed by Plaza [5]. A public announcement is an epistemic action where all agents commonly know that they learn a certain formula. This is modeled by a modal operator $[\varphi]$. A formula of the form $[\varphi]\psi$ is read as “ ψ holds after the announcement of φ ”. This language \mathcal{L}_{PAL} is interpreted in models for epistemic logic.

Definition 1 (Epistemic models) Let a finite set of propositional variables P and a finite set of agents N be given. An epistemic model is a triple $M = (W, R, V)$ such that $W \neq \emptyset$ is a set of possible worlds, $R : N \rightarrow \wp(W \times W)$ assigns an accessibility relation to each agent, and $V : P \rightarrow \wp(W)$ assigns a set of worlds to each propositional variable. \square

In epistemic logic R is usually restricted to equivalence relations. In this paper we treat the general modal case. The semantics are defined with respect to models with a distinguished ‘actual world’: (M, w) .

Definition 2 (Semantics of PAL) Let a model (M, w) with $M = (W, R, V)$ be given. Let $i \in N$, $B \subseteq N$, and $\varphi, \psi \in \mathcal{L}_{\text{PAL}}$. For atomic propositions, negations, and conjunctions we take the usual definition.

$$\begin{aligned} (M, w) \models \Box_i \varphi & \text{ iff } (M, v) \models \varphi \text{ for all } v \text{ such that } (w, v) \in R(i) \\ (M, w) \models C_B \varphi & \text{ iff } (M, v) \models \varphi \text{ for all } v \text{ such that } (w, v) \in R(B)^* \\ (M, w) \models [\varphi]\psi & \text{ iff } (M, w) \models \varphi \text{ implies } (M|\varphi, w) \models \psi \end{aligned}$$

where $R(B) = \bigcup_{i \in B} R(i)$, and $R(B)^*$ is its reflexive transitive closure. The updated model $M|\varphi = (W', R', V')$ is defined by restricting M to those worlds where φ holds. Let $\llbracket \varphi \rrbracket = \{v \in W \mid (M, v) \models \varphi\}$. Now $W' = \llbracket \varphi \rrbracket$, $R'(i) = R(i) \cap \llbracket \varphi \rrbracket^2$, and $V'(p) = V(p) \cap \llbracket \varphi \rrbracket$. \square

A completeness proof with reduction axioms is impossible for this logic.

2.2 Relativized common knowledge

For public announcement logic there is no reduction axiom for formulas of the form $[\varphi]C_B\psi$, given the results in [1]. However the semantic intuition is clear. If φ is true in the old model, then every B -path in the new model ends in a ψ world. This implies that in the old model every B -path that consists exclusively of φ -worlds ends in a $[\varphi]\psi$ world. To facilitate this, we introduce a new operator $C_B(\varphi, \psi)$, which expresses that every B -path which consists exclusively of φ -worlds ends in a ψ world. We call this operator *relativized common knowledge*. The crucial clause in the semantics of the logic of relativized common knowledge, RCL, is: $(M, w) \models C_B(\varphi, \psi)$ iff

$$(M, v) \models \psi \text{ for all } v \text{ such that } (w, v) \in (R(B) \cap \llbracket \varphi \rrbracket^2)^*$$

where $(R(B) \cap \llbracket \varphi \rrbracket^2)^*$ is the reflexive transitive closure of $R(B) \cap \llbracket \varphi \rrbracket^2$. The semantics of the other operators is standard. Ordinary common knowledge can be defined with the new notion: $C_B \varphi \equiv C_B(\top, \varphi)$. The new operator is like the “until” of temporal logic. A temporal sentence “ φ until ψ ” is true iff there is some point in the future where ψ holds and φ is true up to that point.

2.3 Proof system

Relativized common knowledge still resembles common knowledge, and so we need just a slight adaptation of the usual axioms¹.

Definition 3 (Proof system for RCL) The proof system for RCL contains the following axioms and rules:

Taut	all instantiations of propositional tautologies	
Dist	$\Box_i(\varphi \rightarrow \psi) \rightarrow (\Box_i\varphi \rightarrow \Box_i\psi)$	(distribution)
Dist	$C_B(\varphi, \psi \rightarrow \chi) \rightarrow (C_B(\varphi, \psi) \rightarrow C_B(\varphi, \chi))$	(distribution)
Mix	$C_B(\varphi, \psi) \leftrightarrow (\varphi \rightarrow (\psi \wedge E_B(\varphi \rightarrow C_B(\varphi, \psi))))$	(mix)
Ind	$((\varphi \rightarrow \psi) \wedge C_B(\varphi, \psi \rightarrow E_B(\varphi \rightarrow \psi))) \rightarrow C_B(\varphi, \psi)$	(induction)
MP	$\frac{\varphi \quad \varphi \rightarrow \psi}{\psi}$	(modus ponens)
Nec	$\frac{\varphi}{\Box_i\varphi}$	(necessitation)
Nec	$\frac{\varphi}{C_B(\psi, \varphi)}$	(necessitation)

In the mix axiom and the induction axiom $E_B\varphi$ is an abbreviation of $\bigwedge_{i \in B} \Box_i\varphi$ (*everybody* knows φ). A proof consists of a sequence of formulas such that each is either an instance of an axiom, or it can be obtained from formulas that appear earlier in the sequence by applying a rule. If there is a proof of φ , we write $\vdash \varphi$. \square

The soundness of the proof system can easily be shown by induction on the length of proofs, and we do not provide it explicitly.

2.4 Completeness for the static language

To prove completeness for our extended static language, we follow [4]. We take maximally consistent sets with respect to finite fragments of the language that form a canonical model for that fragment. In particular, for any given formula φ we work with a finite fragment called the *closure* of φ .

Definition 4 (Closure) The *closure* of φ is the minimal set Φ such that 1. $\varphi \in \Phi$, 2. Φ is closed under taking subformulas, 3. If $\psi \in \Phi$ and ψ is not a negation, then $\neg\psi \in \Phi$, 4. If $C_B(\psi, \chi) \in \Phi$, then $\Box_i(\psi \rightarrow C_B(\psi, \chi)) \in \Phi$ for all $i \in B$. \square

Definition 5 (Canonical model) The *canonical model* M_φ for φ is the triple $(W_\varphi, R_\varphi, V_\varphi)$ where $W_\varphi = \{\Gamma \subseteq \Phi \mid \Gamma \text{ is maximally consistent in } \Phi\}$; $(\Gamma, \Delta) \in R_\varphi(i)$ iff $\psi \in \Delta$ for all ψ with $\Box_i\psi \in \Gamma$; and $V_\varphi(p) = \{\Gamma \mid p \in \Gamma\}$. \square

Next, we show that a formula in such a finite set is true in the canonical model where that set is taken to be a world, and vice versa.

Lemma 1 (Truth Lemma) For all $\psi \in \Phi$, $\psi \in \Gamma$ iff $(M_\varphi, \Gamma) \models \psi$. \square

Proof (A sketch:) By induction on ψ . The cases for propositional variables, negations, conjunction, and individual epistemic operators are straightforward. Therefore we focus on the case for relativized common knowledge.

¹It is also helpful to write $C_B(\varphi, \psi)$ as a sentence in PDL: $[?\varphi; (\bigcup_{i \in B} i; ?\varphi)^*]\psi$. Our proof system below essentially follows the usual PDL-axioms for this formula.

From left to right. Suppose $C_B(\psi, \chi) \in \Gamma$. If $\psi \notin \Gamma$, then by the induction hypothesis $(M_\varphi, \Gamma) \not\models \psi$, and by the semantics $(M_\varphi, \Gamma) \models C_B(\psi, \chi)$.

Otherwise, if $\psi \in \Gamma$, take any $\Delta \in W_\varphi$ such that $(\Gamma, \Delta) \in (R(B) \cap \llbracket \psi \rrbracket^2)^*$. We have to show that $\Delta \models \chi$, but we can show something stronger, namely that $\Delta \models \chi$ and $C_B(\psi, \chi) \in \Delta$. This is done by induction on the length of the path from Γ to Δ and the **Mix** axiom. We omit the details.

From right to left. Let $(M_\varphi, \Gamma) \models C_B(\psi, \chi)$. Now consider the set $\Lambda = \{\delta_\Delta \mid (\Gamma, \Delta) \in (R(B) \cap \llbracket \psi \rrbracket^2)^*\}$, where δ_Δ is $\bigwedge \{\psi \mid \psi \in \Delta\}$. Let $\delta_\Lambda = \bigvee_{\Delta \in \Lambda} \delta_\Delta$. We can show that $\vdash \delta_\Lambda \rightarrow E_B(\psi \rightarrow \delta_\Lambda)$. By necessitation then $\vdash C_B(\psi, \delta_\Lambda \rightarrow E_B(\psi \rightarrow \delta_\Lambda))$. Applying the induction axiom, we get $\vdash (\psi \rightarrow \delta_\Lambda) \rightarrow C_B(\psi, \delta_\Lambda)$. Since $\vdash \delta_\Lambda \rightarrow \chi$, we also get $\vdash (\psi \rightarrow \delta_\Lambda) \rightarrow C_B(\psi, \chi)$. Now $\delta_\Gamma \rightarrow \delta_\Lambda$, and hence $\vdash \delta_\Gamma \rightarrow (\psi \rightarrow \delta_\Lambda)$. Therefore $C_B(\psi, \chi) \in \Gamma$. \square

This argument is an easy adaptation of the usual completeness proof for common knowledge, reinforcing our idea that our language extension is a natural one: since existing arguments yield more than is usually realized.

Theorem 1 (Completeness for RCL) If $\models \varphi$, then $\vdash \varphi$. \square

Proof Let $\not\models \varphi$, i.e. $\neg\varphi$ is consistent. One easily finds a maximally consistent set Γ in the closure of $\neg\varphi$ with $\neg\varphi \in \Gamma$, as only finitely many formulas matter. By the Truth Lemma, $(M_{\neg\varphi}, \Gamma) \models \neg\varphi$, i.e., $(M_{\neg\varphi}, \Gamma) \not\models \varphi$. \square

2.5 Reduction axioms

Next, let RCL^+ be the epistemic dynamic logic with both relativized common knowledge and public announcements. Its semantics combines those for **PAL** and **RCL**. RCL^+ is no more expressive than **RCL** by a direct translation.

Definition 6 (Translation) The translation function t takes a formula from the language of RCL^+ and yields a formula in the language of **RCL**.

$$\begin{array}{llll}
t(p) & = & p & t([\varphi]p) & = & t(\varphi) \rightarrow p \\
t(\neg\varphi) & = & \neg t(\varphi) & t([\varphi]\neg\psi) & = & t(\varphi) \rightarrow \neg t([\varphi]\psi) \\
t(\varphi \wedge \psi) & = & t(\varphi) \wedge t(\psi) & t([\varphi](\psi \wedge \chi)) & = & t([\varphi]\psi) \wedge t([\varphi]\chi) \\
t(\Box_i\varphi) & = & \Box_i t(\varphi) & t([\varphi]\Box_i\psi) & = & t(\varphi) \rightarrow \Box_i t([\varphi]\psi) \\
t(C_B(\varphi, \psi)) & = & C_B(t(\varphi), t(\psi)) & t([\varphi]C_B(\psi, \chi)) & = & C_B(t(\varphi) \wedge t([\varphi]\psi), t([\varphi]\chi)) \\
& & & t([\varphi][\psi]\chi) & = & t([\varphi]t([\psi]\chi))
\end{array}$$

\square

Lemma 2 (Translation Correctness) For all dynamic-epistemic formulas φ and all models (M, w) , $(M, w) \models \varphi$ iff $(M, w) \models t(\varphi)$. \square

This observation underlies the soundness of the following reduction axioms, with **C-Red** the crucial reduction of relativized common knowledge.

Definition 7 (Proof system for RCL^+) The proof system for RCL^+ is that for **RCL** plus the following reduction axioms:

$$\begin{array}{ll}
\mathbf{At} & [\varphi]p \leftrightarrow (\varphi \rightarrow p) \quad (\text{atoms}) \\
\mathbf{PF} & [\varphi]\neg\psi \leftrightarrow (\varphi \rightarrow \neg[\varphi]\psi) \quad (\text{partial functionality}) \\
\mathbf{Dist} & [\varphi](\psi \wedge \chi) \leftrightarrow ([\varphi]\psi \wedge [\varphi]\chi) \quad (\text{distribution}) \\
\mathbf{KA} & [\varphi]\Box_i\psi \leftrightarrow (\varphi \rightarrow \Box_i[\varphi]\psi) \quad (\text{knowledge-announcement}) \\
\mathbf{C-Red} & [\varphi]C_B(\psi, \chi) \leftrightarrow C_B(\varphi \wedge [\varphi]\psi, [\varphi]\chi) \quad (\text{common reduction})
\end{array}$$

as well as an inference rule of necessitation for all announcement modalities. \square

The formulas on the left of these equivalences are of the form $[\varphi]\psi$. In **At** the announcement operator no longer occurs on the right-hand side. In the other reduction axioms formulas within the scope of an announcement are of higher complexity on the left than on the right.

Theorem 2 (Completeness for RCL⁺) If $\models \varphi$, then $\vdash \varphi$. □

Proof The proof system for RCL is complete (Theorem 1), and every formula in $\mathcal{L}_{\text{RCL}^+}$ is provably equivalent to one in \mathcal{L}_{RCL} . □

2.6 Model comparison games

The notion of relativized common knowledge is of independent interest, just as irreducibly binary general quantifiers (such as *Most A are B*) lead to natural completions of logics with only unary quantifiers. We provide some more information through characteristic games.

Definition 8 (Relativized common knowledge game) Let two models $M = (W, R, V)$ and $M' = (W', R', V')$ be given. Starting from each $w \in W$ and $w' \in W'$, the *n-round relativized common knowledge game* between Spoiler and Duplicator is given as follows. In each round Spoiler can initiate one of two scenarios:

\square_i -**move** Spoiler chooses a point x in one model which is an i -successor of the current w or w' , and Duplicator responds with a matching successor y in the other model. Play continues with the new link x, y .

C_B -**move** Spoiler chooses a B -path $x_0 \dots x_n$ in either of the models with x_0 the current w or w' . Duplicator responds with a B -path $y_0 \dots y_m$ in the other model, with $y_0 = w'$. Then Spoiler can (a) make the end points x_n, y_m the output of this round, or (b) he can choose a world z on the M' -path, and Duplicator must respond by choosing a matching world u on the M -path, and z, u becomes the output. □

The game continues with the new output states. If these differ in their atomic properties, Spoiler wins – otherwise, a player loses whenever he cannot perform a move while it is his turn. If Spoiler has not won after all n rounds, Duplicator wins the whole game.

Definition 9 (Modal depth) The *modal depth* of a formula is defined by: $d(\perp) = d(p) = 1$, $d(\neg\varphi) = d(\varphi)$, $d(\varphi \wedge \psi) = \max(d(\varphi), d(\psi))$, $d(\square_i\varphi) = d(\varphi) + 1$, $d(C_B(\varphi, \psi)) = \max(d(\varphi), d(\psi)) + 1$ If two models (M, w) and (M', w') have the same theory up to depth n , we write $(M, w) \equiv_n (M', w')$. □

The following result holds for most logical languages.

Lemma 3 (Propositional finiteness) For every n , up to depth n , there are only finitely many different propositions up to logical equivalence. □

Theorem 3 (Adequacy) Duplicator has a winning strategy for the n -round game from (M, w) , (M', w') iff $(M, w) \equiv_n (M', w')$. □

Proof The argument is by induction on n . The base case is obvious, and all inductive cases are also standard in modal logic, except that for relativized common

knowledge. As usual, perspicuity is increased somewhat by using the dual existential modality $\hat{C}_B(\varphi, \psi)$.

First, suppose that Duplicator has a winning strategy in games of length $n + 1$, and that $(M, w) \models \hat{C}_B(\varphi, \psi)$, a formula of depth $n + 1$. By the truth definition, there is a finite sequence of φ -worlds in M starting from w , and ending in a world v where ψ holds. Suppose this sequence is picked by Spoiler as his opening move. Duplicator's winning strategy produces a matching sequence in M' , starting at w' and ending in some world v' . Suppose now that Spoiler outputs the end points v, v' as the new link. Duplicator's winning strategy then works for the n -round game starting from M, v, M', v' , and by the inductive hypothesis, $(M', v') \models \psi$. Suppose next that Spoiler chooses any s' on the finite M' -sequence. Then Duplicator's winning strategy yields a matching point s on Spoiler's initial sequence, and again Duplicator has an n -round winning strategy left for $(M, s), (M', s')$. Once more by the inductive hypothesis, then, $(M, s) \models \varphi$. Thus, $(M', w') \models \hat{C}_B(\varphi, \psi)$.

Conversely, suppose that $(M, w) \equiv_{n+1} (M', w')$. A winning strategy for Duplicator in the $(n + 1)$ -round game can be described as follows. If Spoiler makes an opening move of type $[\square_i\text{-move}]$, then the usual modal argument works. Next, suppose that Spoiler opens with a finite sequence in one of the models, say M , without loss of generality. By the Finiteness Lemma, we know that there is only a finite number of complete descriptions of points up to logical depth n , and each point s in the sequence satisfies one of these: say $\Delta(s, n)$. In particular, the end point v satisfies $\Delta(v, n)$. Let $\Delta(n)$ be the disjunction of all formulas $\Delta(s, n)$ occurring on the path. Then, the initial world w satisfies the following formula of modal depth $n + 1$: $\hat{C}_B(\Delta(n), \Delta(v, n))$. By our assumption, we also have $(M', w') \models \hat{C}_B(\Delta(n), \Delta(v, n))$. But any sequence witnessing this by the truth definition is a response that Duplicator can use for her winning strategy. Whatever Spoiler does in the rest of this round, Duplicator always has a matching point that is n -equivalent in the language. \square

Thus, games for \mathcal{L}_{RCL} are straightforward. But it is also of interest to look at the extended dynamic language $\mathcal{L}_{\text{RCL}^+}$ with announcement modalities. Here, the shift modality passing to definable submodels requires a new type of move, where players can decide to change the current model. The following description of what happens is 'modular': a model changing move can be added to model comparison games for ordinary epistemic logic (perhaps with common knowledge), or for our relativized common knowledge game. By way of explanation: we let Spoiler propose a model shift. Players first discuss the 'quality' of that shift, and Duplicator can win if it is deficient; otherwise, the shift really takes place, and play continues within the new models. This involves a somewhat unusual sequential composition of games, but perhaps one of independent interest.

Definition 10 (Public announcement move) Let the setting be the same as for the n -round game in Definition 9.

$[\varphi]$ -move Spoiler chooses a number $r < n$, and sets $S \subseteq W$ and $S' \subseteq W'$, with the current $w \in S$ and likewise $w' \in S'$. *Stage 1:* Duplicator chooses states s in $S \cup S'$, \bar{s} in $\bar{S} \cup \bar{S}'$. Then Spoiler and Duplicator play the r -round game for these worlds. If Duplicator wins this subgame, she wins the n -round game. *Stage 2:* Otherwise, the game continues in the relativized models $M|S, w$ and $M'|S', w'$ over $n - r$ rounds. \square

The definition of depth is easily extended to formulas $[\varphi]\psi$ as $d([\varphi]\psi) = d(\varphi) + d(\psi)$. For the sake of illustration, assume that the new move has been added to the

relativized common knowledge game.

Theorem 4 (Adequacy) Duplicator has a winning strategy for the n -round game on (M, w) and (M', w') iff $(M, w) \equiv_n (M', w')$ in $\mathcal{L}_{\text{RCL}^+}$. \square

Proof We only discuss the inductive case demonstrating the match between announcement modalities and model-changing steps.

First, let Duplicator have a winning strategy in the n -round game between M, w and M', w' . Suppose that $(M, w) \models \langle \varphi \rangle \psi$, with total modal depth n . Consider the case where Spoiler chooses $r = d(\varphi)$ and sets S, S' equal to the extensions of φ in the two models. Then Spoiler has a winning strategy in all r -round games in Stage 1, exploiting the $\varphi, \neg\varphi$ -difference between whatever points Duplicator chooses: by the inductive hypothesis for $r = d(\varphi)$. Suppose that Spoiler plays such a strategy. In that case, the $n - r$ -round subgame for the relativized submodels is reached, and Duplicator must have a winning strategy there. But that means, now by the inductive hypothesis for $n - r$, that $(M'|_{\varphi}, w') \models \psi$, and hence $(M', w') \models \langle \varphi \rangle \psi$.

Next, suppose that $(M, w), (M', w')$ are equivalent up to depth n . We need to describe Duplicator's winning strategy. Consider any opening choice of r, S, S' made by Spoiler. *Case 1:* Duplicator can choose two points s, \bar{s} in Stage 1 giving her a winning strategy in the initial r -round game. Then we are done. *Case 2:* Duplicator has no such winning strategy, which means that Spoiler has one – or equivalently by the inductive hypothesis, there is some formula of depth r distinguishing s from \bar{s} . In that case, we can find a formula A defining both set S in M and S' in M' . To find this, consider any point x in S . Using the preceding observation, we can find a formula δ_x of depth n which holds at x but at no world in $M - S$ or $M' - S'$. (Note that there can be infinitely many worlds involved in the comparison, but finitely many difference formulas will suffice by the Finiteness Lemma, which also holds for this extended language.) Let Δ_S be the disjunction of all these δ_x . A formula Δ'_S is found likewise, and we let A be the *disjunction of Δ'_S and Δ_S* . It is easy to see that this formula of depth r defines S in M and S' in M' . Now we use the given language equivalence between M, w and M', w' with respect to all depth n -formulas $\langle A \rangle \psi$ where ψ runs over all formulas of depth $n - r$. We can conclude that $M|A, w$ and $M'|A, w'$ are equivalent up to depth $n - r$, and hence Duplicator has a winning strategy for the remaining game, by the inductive hypothesis. \square

Finally, our two games must be related, since $\mathcal{L}_{\text{RCL}^+}$ has the same expressive power as \mathcal{L}_{RCL} . This means that players who can win one of our games should also be able to win the other, given suitable game lengths. An explicit description of the relevant strategy conversion is beyond the scope of this paper.

2.7 Complexity

Update logics are about processes that manipulate information, and hence they raise natural questions of complexity. In particular, all of the usual complexity questions concerning a logical system make sense:

Model checking: When is a formula true in a model, i.e., when do agents know given propositions in an information state, or when do specific epistemic actions in the model produce specified effects?

Satisfiability testing: When does a formula have a model, or more generally: when can we find an informational setting realizing given epistemic specifications? Or, in terms of *validity*: e.g., when will a given epistemic action always produce some global specified effect?

Model comparison: When do two given models satisfy the same formulas, or equivalently, when can Duplicator win any game over them - or: when do two representations describe the same information about some group of agents?

Now technically, the translation of definition 6 combined with known algorithms for model checking, satisfiability, validity, or model comparison for epistemic logic yield similar algorithms for public announcement logic. But, worst case, the length of the translation of a formula is exponential in the length of the formula. E.g., the translation of φ occurs three times in that of $[\varphi]C_B(\psi, \chi)$, and hence a direct complexity analysis is worth-while.

Lemma 4 Deciding whether a model (M, w) satisfies a formula $\varphi \in \mathcal{L}_{\text{RCL}}$ is computable in polynomial time in the length of φ and the size of M . \square

Proof Let $\|M\|$ be the cardinality of the set of worlds of M plus that of the accessibility relation of M , and $|\varphi|$ the length of φ . A formula φ has at most $|\varphi|$ subformulas. We make a list $\varphi_0, \dots, \varphi_n$ (where $\varphi_n = \varphi$) of these such that for all formulas ψ , their subformulas occur earlier. We process the list by successively labeling the worlds in the model where the list formulas φ_i hold. The two crucial cases are as follows. For $\Box_i \psi$, we check for a given world in the model whether ψ holds in every accessible world. Since we have already labeled the states where the subformula ψ holds, this can be done in $\|M\|$ steps. For $C_B(\psi, \chi)$ we proceed as follows. First label all those worlds with $\neg\psi$ as worlds where $C_B(\psi, \chi)$ holds, and label worlds where ψ and $\neg\chi$ hold as worlds where $C_B(\psi, \chi)$ fails. Then iterate the following step until the labeled set does not grow anymore: pick an unlabeled world that can reach a world labeled with $\neg C_B(\psi, \chi)$ in a single i -step (for any $i \in B$) and also label it as a world where $C_B(\psi, \chi)$ fails. Each round takes at most $\|M\|$ steps for checking accessibilities, and the total set of labelled worlds can grow at most $\|M\|$ steps. When the set stops growing, all still unlabelled worlds are labelled with $C_B(\psi, \chi)$. By induction of formula complexity, this algorithm can be proved correct. So, the complexity of model checking for RCL is in time $\mathcal{O}(|\varphi| \times \|M\|^2)$. \square

This algorithm does not suffice for the case with public announcements. The truth values of φ and ψ in the given model do not fix that of $[\varphi]\psi$. We must also know the value of ψ in the model restricted to φ worlds.

Lemma 5 Deciding whether a model (M, w) satisfies a formula $\varphi \in \mathcal{L}_{\text{RCL}+}$ is computable in polynomial time in the length of φ and the size of M . \square

Proof Again there are at most $|\varphi|$ subformulas of φ . Now we make a binary tree of these formulas which splits with formulas of the form $[\psi]\chi$. On the left subtree all subformulas of ψ occur, on the right all those of χ . This tree can be constructed in time $\mathcal{O}(|\varphi|)$. Labeling the model is done by processing this tree from bottom to top from left to right. The only new case is when we encounter a formula $[\psi]\chi$. In that case we first label those worlds where ψ does not hold as worlds where $[\psi]\chi$ holds, then we process the right subtree under $[\psi]\chi$ where we restrict the model to ψ worlds. After this process we label those worlds that were labeled with χ as worlds where $[\psi]\chi$ holds and the remaining as worlds where it does not hold. We can see by induction on formula complexity that this algorithm is correct.

Also by induction on φ , this algorithm takes time $\mathcal{O}(|\varphi| \times \|M\|^2)$. The only difficult step is labeling the model with $[\psi]\chi$. By the induction hypothesis, restricting the model to ψ takes time $\mathcal{O}(|\psi| \times \|M\|^2)$. We simply remove (temporarily) all

worlds labelled $\neg\varphi$ and all arrows pointing to such worlds. Again by the induction hypothesis, checking χ in this new model takes $\mathcal{O}(|\psi| \times \|M\|^2)$ steps. The rest of the process takes $\|M\|$ steps. So, this step takes over-all time $\mathcal{O}(|\psi|\chi| \times \|M\|^2)$. \square

Moving on from model checking, the satisfiability and the validity problem of epistemic logic with common knowledge are both EXPTIME-complete. In fact, this is true for almost any logic that contains a transitive closure modality. Satisfiability and validity for PDL are also EXPTIME-complete. Now there is a linear time translation of the language of RCL to that of PDL. Therefore the satisfiability and validity problems for RCL are also EXPTIME-complete. For $\mathcal{L}_{\text{RCL}^+}$ and even PAL, however, the complexity of satisfiability and validity are still unknown.

Finally, the complexity of model comparison for finite models is the same as that for ordinary epistemic logic, viz. PTIME. The reason is that even basic modal equivalence on finite models implies the existence of a bisimulation, while all our extend languages are bisimulation-invariant.

2.8 Other logics with relativization

Languages with relativizations are very common in logic. Indeed, closure under relativization is sometimes stated as a defining condition on logics in abstract model theory. Basic modal or first-order logic as they stand are closed under relativizations $[A]\varphi$, often written $(\varphi)^A$. And the same is true for logics with fixed-point constructions, like PDL (cf. [6]) or the modal μ -calculus. E.g., computing a relativized least fixed-point $[A]\mu p.\varphi(p)$ yields the same result as $\mu p.\varphi(p) \wedge A$. Relativization looks much like *restricting quantifiers*, as in the earlier-mentioned shift from unary “Most objects are ψ ” to binary “Most φ are ψ ”. By ‘Conservativity’ for generalized quantifiers, “Most φ are ψ ” is equivalent to “Most φ are $\psi \wedge \varphi$ ”. But note that the latter principle does not relativize ψ to evaluation wholly inside the φ -area! Thus, the expressive power of the two sorts of extension is not evidently the same. A similar issue arises in our setting. We defined \mathcal{L}_{RCL} as an extension with binary common knowledge in the second sense. We have shown how this allows us to define all relativizations and all ordinary common knowledge operators, i.e., the language PAL. But the converse is still open.

Question: Do \mathcal{L}_{RCL} and \mathcal{L}_{PAL} have the same expressive power?

If the answer to this question is negative, we would have two competing relativization-closed versions of epistemic logic with common knowledge, even though \mathcal{L}_{RCL} seems the more elegant one of the two.

3 Logic of epistemic actions

Our proposed methodology for epistemic logic with announcements also works more generally. In this section, we make the same move in the general dynamic logic of epistemic actions, which also lacks a reduction axiom for common knowledge. In Section 3.1 we introduce the logic of epistemic action LEA. In Section 3.2 we briefly introduce a variant of PDL, called automata PDL, which is our technical tool for creating a suitably enriched base language for LEA that allows for perspicuous reduction axioms for common knowledge. These axioms are introduced in Section 3.3.

3.1 Language and semantics

Dynamic *actions* with epistemic aspects, such as communication or other information-bearing events, are quite similar to static epistemic *situations*. In [1] this analogy is used as the engine for general update of epistemic models under epistemic actions. In particular, individual actions come with *preconditions* holding only at those worlds where they can be performed.

Definition 11 (Action models) An action model for a finite set of agents N with a language \mathcal{L} is a triple $\mathbf{A} = (W, R, \text{pre})$ where $W \neq \emptyset$ is a set of actions; $R : N \rightarrow \wp(W \times W)$ assigns an accessibility relation to each agent, and $\text{pre} : W \rightarrow \mathcal{L}$ assigns a precondition in \mathcal{L} to every action. A pair (\mathbf{A}, w) is an action model with a distinguished actual action $w \in W$. \square

Here \mathcal{L} can be any language that can be interpreted in the models of definition 1. The effect of executing an action is modeled by the following *product construction*.

Definition 12 (Execution) Given a static epistemic model (M, w) and an action model (\mathbf{A}, w) with $(M, w) \models \text{pre}(w)$, we say that the result of executing (\mathbf{A}, w) in (M, w) is the static model $(M \cdot \mathbf{A}, (w, w)) = ((W', R', V'), (w, w))$ where $W' = \{(v, v) \mid (M, v) \models \text{pre}(v)\}$, $R'(i) = \{((v, v), (u, u)) \mid (v, u) \in R(i) \text{ and } (v, u) \in R(i)\}$, and $V'(u, v) = V(u)$. \square

Definitions 11 and 12 provide a semantics for the logic of epistemic action LEA of [1]. The basic epistemic language \mathcal{L}_{LEA} is extended with dynamic modalities $[\mathbf{A}, w]\varphi$, where a \mathbf{A} is any *finite action model* for \mathcal{L}_{LEA} . These say that “every execution of (\mathbf{A}, w) yields a model where φ holds”:

$$(M, w) \models [\mathbf{A}, w]\varphi \quad \text{iff} \quad (M, w) \models \text{pre}(w) \text{ implies that } (M \cdot \mathbf{A}, (w, w)) \models \varphi$$

[1] presents a proof system for this logic with a complicated completeness proof, and without reduction axioms for common knowledge (which were already lacking for public announcement actions). So, we must extend this language to get reduction axioms after all. Again, the semantic intuition about the crucial case $(M, w) \models [\mathbf{A}, w]C_B\varphi$ is clear. It says that, if there is a B -path w_0, \dots, w_n (with $w_0 = w$) in the static model and a matching B -path w_0, \dots, w_n (with $w_0 = w$) in the action model with $(M, w_i) \models \text{pre}(w_i)$ for all $i \leq n$, then $(M, w_n) \models \varphi$. To express all this in the initial static model, it turns out convenient to choose a representation of complex epistemic assertions that meshes well with action models. Now, the relevant finite paths in static models involve strings of agent accessibility steps and tests on formulas, as programs in dynamic logic are associated with regular string languages. These are the sort of object that can be recognized by a *finite automaton*. But action models resemble finite automata, too, with regular languages of accessibility transitions and tests for preconditions. All this leads us to *automata PDL* a variant of PDL where finite automata tag modalities, rather than programs.

3.2 Automata PDL

The system for APDL presented here is taken from [3, Section 10.3], where relevant basic references can be found for what follows. Here, in our epistemic perspective, atomic programs will be viewed as agents.

Definition 13 (Finite automata) Let Σ be an alphabet. A *finite automaton* for Σ is a quadruple $\mathcal{A} = (S, I, F, \delta)$, where S is a finite set of states; $I, F \subseteq S$ are the

initial state and the final state of the automaton respectively; and $\delta \subseteq S \times \Sigma \times S$ is the set of transitions. \square

Intuitively, if $(s_0, \sigma, s_1) \in \delta$, then s_1 is reachable from s_0 by symbol σ . A path from the initial state to the final state yields a string, which is said to be accepted by the automaton. Ordinary finite automata theory allows more than one final state, but our set-up loses no generality (cf. [3]).

Definition 14 (Acceptance) A string $\sigma_1 \dots \sigma_n \in \Sigma^*$ is *accepted by* $\mathcal{A} = (S, I, F, \delta)$ iff there exists a sequence of states s_0, \dots, s_n such that for all $i < n$ it holds that $(s_i, \sigma_i, s_{i+1}) \in \delta$, where $s_0 = I$ and $s_n = F$. \square

In APDL, automata feature as modal operators. Their alphabet consists of the atomic programs together with tests on formulas of the language itself.

Definition 15 (Language of APDL) Let a set of atomic programs Π be given. The language $\mathcal{L}_{\text{APDL}}$ is given by the following BNF:

$$\varphi ::= p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid [i]\varphi \mid [?\varphi_1]\varphi_2 \mid [\mathcal{A}]\varphi$$

where $i \in N$ and \mathcal{A} is an automaton over $\Pi \cup \{?\varphi \mid \varphi \in \mathcal{L}_{\text{APDL}}\}$. \square

A formula of the form $[\mathcal{A}]\varphi$ should be read as “ φ holds after every execution of a string accepted by \mathcal{A} ”.

Definition 16 (Semantics of APDL) Let (M, w) be any model with $M = (W, R, V)$. Let $i \in \Pi$, and $\varphi, \psi \in \mathcal{L}_{\text{APDL}}$. For atomic propositions, negations, and conjunctions we take the usual definition. Next, we set

$$\begin{aligned} (M, w) \models [i]\varphi & \text{ iff } (M, v) \models \varphi \text{ for all } v \text{ such that } (w, v) \in R(i) \\ (M, w) \models [?\varphi]\psi & \text{ iff } (M, v) \models \varphi \text{ implies that } (M, v) \models \psi \\ (M, w) \models [\mathcal{A}]\varphi & \text{ iff } (M, v) \models \varphi \text{ for all } v \text{ and } \vec{\sigma} \text{ such that } (w, v) \in \llbracket \vec{\sigma} \rrbracket \\ & \text{ and } \vec{\sigma} \text{ is accepted by } \mathcal{A} \end{aligned}$$

$$\begin{aligned} \llbracket [i; \vec{\sigma}] \rrbracket & = R(i) \circ \llbracket \vec{\sigma} \rrbracket \\ \llbracket [?\varphi; \vec{\sigma}] \rrbracket & = \{(w, w) \mid (M, w) \models \varphi\} \circ \llbracket \vec{\sigma} \rrbracket \\ \llbracket [\varepsilon] \rrbracket & = \{(w, w) \mid w \in W\} \end{aligned}$$

where ε is the empty string. \square

APDL and PDL have the same expressive power [3], but the former system offers a more convenient ‘intensional’ description of actions, which we will exploit in our account of reduction axioms. In particular, given the earlier-mentioned connection between epistemic logic and PDL, one can also translate epistemic logic into APDL. Agents’ accessibility relations become atomic programs, and e.g., common knowledge among group B involves a single-state automaton $\mathcal{A}_{C_B} = (\{0\}, 0, 0, \delta)$ with $\delta = \{0\} \times B \times \{0\}$. Henceforth, we will think of APDL in this epistemic guise.

3.3 Reduction axioms

[3] has a sound and complete proof system for APDL by itself. Now we show that adding epistemic actions to this static language does not increase expressive power, while we can also find a complete proof system with reduction axioms.



Figure 1: A three-state epistemic action model and its six-node automaton. The middle nodes are only reachable by executing the appropriate preconditions, which label the arrows. Solid lines are transitions for i , dashed lines for j .

Definition 17 (Language of APDL⁺) Let a finite set of propositional variables P and a finite set of atomic programs Π be given. The language $\mathcal{L}_{\text{APDL}^+}$ is given by the following BNF:

$$\varphi ::= p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid [i]\varphi \mid [?\varphi_1]\varphi_2 \mid [\mathcal{A}]\varphi \mid [\mathbf{A}, \mathbf{w}]\varphi$$

where $p \in P$, $i \in \Pi$, \mathcal{A} is an automaton over $N \cup \{?\varphi \mid \varphi \in \mathcal{L}_{\text{APDL}^+}\}$ and (\mathbf{A}, \mathbf{w}) is an action model for $\mathcal{L}_{\text{APDL}^+}$. \square

To translate formulas of the form $[\mathbf{A}, \mathbf{w}]\varphi$ to the language without epistemic actions, we can use reduction axioms from the logic of epistemic actions without common knowledge, but now there is an extra case, namely for sentences of the form $[\mathbf{A}, \mathbf{w}][\mathcal{A}]\varphi$, with complex epistemic postconditions $[\mathcal{A}]\varphi$ such as common knowledge statements. Our idea is to merge the two modalities using a *product* of two automata: one being the epistemic $[\mathcal{A}]$, and one for the action model $[\mathbf{A}, \mathbf{w}]$.

As to the latter, we construct automata $\mathcal{A}_{(\mathbf{A}, \mathbf{w}, \mathbf{v})}$ for each reachable action world \mathbf{v} , making paths in $\mathcal{A}_{(\mathbf{A}, \mathbf{w}, \mathbf{v})}$ and (\mathbf{A}, \mathbf{w}) correspond. By way of explanation, action worlds play two roles: their preconditions determine whether the action can be executed, but they also serve as epistemic alternatives for agents. A path in our automaton takes both roles into account by having two copies of each world. Agents can reach only one sort of copy, from which the other is accessible by executing the right precondition. From the latter, agents can again reach other states.

Definition 18 (Automata for action models) Let $\mathbf{A} = (\mathbf{W}, \mathbf{R}, \text{pre})$ be an action model, with worlds w, v . The automaton $\mathcal{A}_{(\mathbf{A}, \mathbf{w}, \mathbf{v})}$ is the four-tuple (S, I, F, δ) where $S = \mathbf{W} \times \{0, 1\}$, $I = (w, 0)$, $F = (v, 1)$, and $\delta = \{((u, 0), ?\text{pre}(u), (u, 1)) \mid u \in \mathbf{W}\} \cup \{((u, 1), i, (t, 0)) \mid (u, t) \in \mathbf{R}(i)\}$ for all agents i . \square

For instance, consider an epistemic action model for two agents i and j where nothing happens or agent i is informed about the truth of p . Agent i knows exactly what is going on, but j does not know what is going on at all. This action model is shown on the left in Figure 1. Its automaton is shown on the right.

Now we must combine the two sorts of automata. In the simple case of epistemic common knowledge, the automaton for the action model itself would be virtually what we want, but in general, the desired paths, and hence the combined automaton

must be restricted. For this we use a kind of multiplication, which is actually not unlike product update.

Definition 19 Let an action model A , two action worlds $w, v \in W$, and an automaton \mathcal{A} be given. Then the automaton $\mathcal{A}_{(A,w,v)} \otimes \mathcal{A} = (S', I', F', \delta')$ is defined as follows: $S' = W \times \{0, 1\} \times S$, $I' = ((w, 0, I), F' = ((v, 1, F)$ and

$$\begin{aligned} \delta' = & \{((u, 1, s_j), i, ((t, 0, s_k) \mid (u, t) \in R(i) \text{ and } (s_j, i, s_k) \in \delta)\} \cup \\ & \{((u, 0, s), ?\varphi, ((u, 1, s) \mid \text{pre}(u) = \varphi)\} \cup \\ & \{((u, 1, s_i), ?\langle A, u \rangle \varphi, ((u, 1, s_j) \mid (s_i, ?\varphi, s_j) \in \delta)\} \end{aligned}$$

As a special case, multiplication with the epistemic automaton for common knowledge yields nothing new, i.e. $\mathcal{A}_{A,w,v} \otimes \mathcal{A}_{C_N} = \mathcal{A}_{A,w,v}$. Now we can translate the language of APDL⁺ to the language of APDL.

Definition 20 (Translation) The translation map t takes a formula or automaton from the language of APDL⁺ and yields a formula of APDL:

$$\begin{aligned} t(p) &= p \\ t(\neg\varphi) &= \neg t(\varphi) \\ t(\varphi \wedge \psi) &= t(\varphi) \wedge t(\psi) \\ t([i]\varphi) &= [i]t(\varphi) \\ t([?\varphi]\psi) &= t(? \varphi) \rightarrow t(\psi) \\ t([\mathcal{A}]\varphi) &= [t(\mathcal{A})]t(\varphi) \\ t([A, w]p) &= t(\text{pre}(w)) \rightarrow p \\ t([A, w]\neg\varphi) &= t(\text{pre}(w)) \rightarrow \neg t([A, w]\varphi) \\ t([A, w](\varphi \wedge \psi)) &= t([A, w]\varphi) \wedge t([A, w]\psi) \\ t([A, w][i]\varphi) &= t(\text{pre}(w)) \rightarrow \bigwedge_{(w,v) \in R(i)} [i]t([A, v]\varphi) \\ t([A, w][?\varphi]\psi) &= t([A, w](\varphi \rightarrow \psi)) \\ t([A, w][\mathcal{A}]\varphi) &= \bigwedge_{v \in W} t([\mathcal{A}_{(A,w,v)}] \otimes \mathcal{A})[A, v]\varphi \\ t([A, w][A', w']\varphi) &= t([A, w]t([A', w']\varphi)) \end{aligned}$$

where $t(\mathcal{A})$ is the automaton where every occurrence of a formula as a condition has been replaced by its t -translation. \square

Every formula is provably equivalent to its translation. This is essentially the soundness of the following proof system.

Definition 21 (Proof system for APDL⁺) The proof system for APDL⁺ consists of all the axioms and rules of APDL plus the following axioms:

$$\begin{aligned} \mathbf{At} \quad [A, w]p &\leftrightarrow (\text{pre}(w) \rightarrow p) && (\text{atoms}) \\ \mathbf{PF} \quad [A, w]\neg\psi &\leftrightarrow (\text{pre}(w) \rightarrow \neg[A, w]\psi) && (\text{partial functionality}) \\ \mathbf{Dist} \quad [A, w](\psi \wedge \chi) &\leftrightarrow ([A, w]\psi \wedge [A, w]\chi) && (\text{distribution}) \\ \mathbf{KA} \quad [A, w][i]\varphi &\leftrightarrow (\text{pre}(w) \rightarrow \bigwedge_{(w,v) \in R(i)} [i][A, v]\varphi) && (\text{knowledge-action}) \\ \mathbf{Red} \quad [A, w][\mathcal{A}]\varphi &\leftrightarrow \bigwedge_{v \in W} [\mathcal{A}_{(A,w,v)}] \otimes \mathcal{A}[A, v]\varphi && (\text{reduction axiom}) \end{aligned}$$

plus necessitation for action model modalities. \square

The difficult case for the soundness of these axioms is the reduction axiom.

Lemma 6 $[A, w][\mathcal{A}]\varphi$ is equivalent to $\bigwedge_{v \in W} [\mathcal{A}_{(A,w,v)}] \otimes \mathcal{A}[A, v]\varphi$ \square

Proof (Sketch only:) From left to right. By contraposition. Suppose $(M, w) \models \bigvee_{v \in W} \langle \mathcal{A}_{(A, w, v)} \otimes \mathcal{A} \rangle \langle A, v \rangle \varphi$. Let $\langle \mathcal{A}_{(A, w, v)} \otimes \mathcal{A} \rangle \langle A, v \rangle \varphi$ be a disjunct that makes the formula true. Therefore there is a string $\vec{\sigma} = \sigma_1 \dots \sigma_n$ accepted by $\mathcal{A}_{(A, w, v)} \otimes \mathcal{A}$ such that some path w_0, \dots, w_n runs from w to v with $(w, v) \in \llbracket \vec{\sigma} \rrbracket$ and $(M, v) \models \langle A, v \rangle \varphi$. Let $((w, 0, I), \dots, (v, 1, F))$ be the path in $\mathcal{A}_{(A, w, v)} \otimes \mathcal{A}$ that yields $\vec{\sigma}$. We can now make two strings: a string $\sigma_{\mathcal{A}_{(A, w, v)}}$ that is accepted by $\mathcal{A}_{(A, w, v)}$ and a string $\sigma'_{\mathcal{A}_{(A)}}$ that is accepted by \mathcal{A} . These can be “read” from the path $((w, 0, I), \dots, (v, 1, F))$. In particular, $(w, v) \in \llbracket \sigma_{\mathcal{A}_{(A, w, v)}} \rrbracket$, because of the path w_0, \dots, w_n . In the product update model $M \cdot A$ this path becomes $((w_0, w_0), \dots, (w_n, w_n))$. Along this path $\sigma_{\mathcal{A}}$ can be executed. Therefore $(M, w) \models \langle A, w \rangle \langle \mathcal{A} \rangle \varphi$.

From right to left. Suppose $(M, w) \models \langle A, w \rangle \langle \mathcal{A} \rangle \varphi$. Then $(M \cdot A, (w, w)) \models \langle \mathcal{A} \rangle \varphi$. Therefore some string $\vec{\sigma} = \sigma_1 \dots \sigma_n$ is accepted by \mathcal{A} such that $((w, w), (v, v)) \in \llbracket \vec{\sigma} \rrbracket_{M \cdot A}$ and $(M \cdot A, (v, v)) \models \varphi$. Therefore there is a path w_0, \dots, w_n from w to v and a path w_0, \dots, w_n from w to v such that $(M, w_i) \models \text{pre}(w_i)$. Let $\vec{\sigma}' = \sigma'_1 \dots \sigma'_n$ be the string obtained by replacing every ψ in $\vec{\sigma}$ with $\langle A, u \rangle \psi$. Then $(w, v) \in \llbracket \vec{\sigma}' \rrbracket_M$. Let $\vec{\sigma}''$ be the string obtained from $\vec{\sigma}'$ by prefixing it with $?\text{pre}(w)$, while every σ'_i of the form i is replaced by $i; ?\text{pre}(w_i)$. It can be shown that $\vec{\sigma}''$ is accepted by $\mathcal{A}_{(A, w, v)} \otimes \mathcal{A}$ and that $(w, v) \in \llbracket \vec{\sigma}'' \rrbracket$. Therefore $(M, w) \models \bigvee_{v \in W} \langle \mathcal{A}_{(A, w, v)} \otimes \mathcal{A} \rangle \langle A, v \rangle \varphi$. \square

Theorem 5 (Completeness) If $\models \varphi$, then $\vdash \varphi$. \square

Proof The proof system for APDL is complete and every formula in $\mathcal{L}_{\text{APDL}^+}$ is provably equivalent to a formula in $\mathcal{L}_{\text{APDL}}$. \square

4 Conclusion and further research

Dynamic-epistemic logics provide excellent means for studying exchange of factual and higher-order information. In this many-agent setting, common knowledge is an essential concept. We have presented two extended languages for dynamic-epistemic logic that admit explicit action/common knowledge reduction axioms: one (PAL) for public announcement only, and one (APDL) for general action update. These systems make proof and complexity analysis more perspicuous than earlier attempts in the literature. Still, PAL and epistemic APDL are just two extremes on a spectrum, and many further natural update logics may lie in between. For instance, we found an open question of expressiveness of \mathcal{L}_{RCL} versus PAL.

In addition, our methods raise new model-theoretic questions about languages that admit of ‘update closure’ by reduction axioms and translation procedures. Examples are temporal **UNTIL** logic, fragments of the μ -calculus, or even first-order logic and its fixed-point extensions. Also, our automata seem a natural setting for generalizing other open questions, known mainly so far for public announcement (cf. [7]). These include extensions of our current concerns, such as axiomatizing the *schematic validities* of update languages, closed under arbitrary substitutions for proposition letters— an issue which is not solved by our reduction axioms alone.

References

- [1] A. Baltag, L. S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. Technical Report SEN-R9922, CWI, Amsterdam, 1999.

- [2] J. Gerbrandy. *Bisimulations on Planet Kripke*. ILLC Dissertation Series, Amsterdam, 1999.
- [3] D. Harel, D. Kozen, and J. Tiuryn. *Dynamic Logic*. Foundations of Computing. MIT Press, Cambridge, Massachusetts, 2000.
- [4] D. Kozen and R. Parikh. An elementary proof of the completeness of PDL. *Theoretical Computer Science*, 14:113–118, 1981.
- [5] J. A. Plaza. Logics of public communications. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z. W. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216, 1989.
- [6] J. van Benthem. Update as relativization. ILLC, University of Amsterdam, 1999, manuscript.
- [7] J. van Benthem. One is a lonely number: on the logic of communication. Technical Report PP-2002-27, ILLC, Amsterdam, 2002. To appear in P. Köpke, ed. *Colloquium Logicum, Münster, 2001*, AMS Publications, Providence.