University of Groningen

Arrow update logic

Kooi, B.P.; Renne, B.

Published in:
The Review of Symbolic Logic

DOI:
[10.1017/S1755020311000189](https://doi.org/10.1017/S1755020311000189)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Early version, also known as pre-print

Publication date:
2011

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Kooi, B. P., & Renne, B. (2011). Arrow update logic. *The Review of Symbolic Logic*, 4(4), 536-559.
<https://doi.org/10.1017/S1755020311000189>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

ARROW UPDATE LOGIC

BARTELD KOOI AND BRYAN RENNE

Faculty of Philosophy, University of Groningen

Abstract. We present *Arrow Update Logic*, a theory of epistemic access elimination that can be used to reason about multi-agent belief change. While the belief-changing “arrow updates” of Arrow Update Logic can be transformed into equivalent belief-changing “action models” from the popular Dynamic Epistemic Logic approach, we prove that arrow updates are sometimes exponentially more succinct than action models. Further, since many examples of belief change are naturally thought of from Arrow Update Logic’s perspective of eliminating access to epistemic possibilities, Arrow Update Logic is a valuable addition to the repertoire of logics of information change. In addition to proving basic results about Arrow Update Logic, we introduce a new notion of common knowledge that generalizes both ordinary common knowledge and the “relativized” common knowledge familiar from the Dynamic Epistemic Logic literature.

§1. Introduction. Public Announcement Logic is a modal logic theory for reasoning about multi-agent belief changes brought about by completely trustworthy announcements. This logic generally comes in two flavors: Plaza’s (1989, 2007) logic, in which a public announcement of a statement eliminates all epistemic possibilities in which the statement does not hold, and Gerbrandy & Groeneveld’s (1997) logic, in which a public announcement of a statement merely eliminates access to all epistemic possibilities in which the statement does not hold. The most popular Dynamic Epistemic Logic (DEL), due to Baltag *et al.* (1998) or “BMS” (see also Baltag & Moss, 2004; van Ditmarsch *et al.*, 2007), generalizes Plaza’s logic by eliminating epistemic possibilities that do not satisfy certain preconditions. However, many examples of epistemic or doxastic update are naturally thought of in terms of merely eliminating access to such possibilities.

In this paper, we present Arrow Update Logic (AUL), a theory of epistemic access elimination that generalizes Gerbrandy and Groeneveld’s logic of public announcements. AUL is inspired by the arrow precondition language proposed by Renne *et al.* (2009, 2010); there are also connections to the work of Renardel de Lavalette (2004) and of Aucher *et al.* (2009) on accessibility relation change and to the work of van Benthem (2005) on arbitrary arrow elimination.

In the Arrow Update Logic we present in this paper, we shall restrict ourselves to very simple arrow updates in which epistemic arrows are merely eliminated; no new arrows will be created. Accordingly, while agents may have different beliefs and different ways of processing incoming information, it is common knowledge among the agents how each agent will process the incoming information. We will revisit this restriction at the end of the paper.

To motivate Arrow Update Logic, we observe that it is natural to think of certain belief changes in terms of eliminating arrows. For example, let us consider a van Ditmarsch (2000) card-playing scenario: the cards *r*, *w*, and *b* are dealt to agents 1, 2, and 3, one card per agent. A Kripke model *M* representing this situation is pictured on the left in Figure 1.

Received: August 2, 2011.

Fig. 1. Kripke models for agents 1 (solid), 2 (dashed), and 3 (dotted).

In this picture, each world is labeled by an expression $c_1c_2c_3$, which represents the deal in which agent 1 is dealt card $c_1 \in \{r, w, b\}$, agent 2 is dealt card $c_2 \in \{r, w, b\}$, and agent 3 is dealt card $c_3 \in \{r, w, b\}$. Solid arrows are reserved for agent 1, dashed arrows for agent 2, and dotted arrows for agent 3. To have an agent's arrow point between two worlds means that the agent in question cannot distinguish between the deals represented by the two worlds.

Now suppose that all the agents look at their own cards, an action we denote by **pickup**. After this action, it is common knowledge that each agent knows her own card but does not know the cards of the other agents. We represent the situation after this action using a Kripke model $M * \text{pickup}$ pictured on the right in Figure 1.

Note that a natural way to obtain $M * \text{pickup}$ from M is to simply delete arrows. To specify which arrows to delete, we can simply say which arrows are to remain, with the understanding that all other arrows are to be deleted. To say that an agent- a arrow from world w to world w' should remain, we will use a triple of the form (φ, a, ψ) , where the expression φ describes a condition to be satisfied by the source world w , the expression a names the agent, and the expression ψ describes a condition to be satisfied by the target world w' . We collect together a number of such arrow specifications in the set **pickup** defined by

$$\text{pickup} \stackrel{\text{def}}{=} \{(c_a, a, c_a) \mid c \in \{r, w, b\} \text{ and } a \in \{1, 2, 3\}\},$$

where c_a is the condition “agent a has card c .” The set **pickup** describes all arrows that connect worlds in which an agent has the same card. Intuitively, after the agents look at their cards, only the arrows described by **pickup** should remain: each agent knows her own card (and hence we should delete arrows between worlds in which her card is different because she can distinguish these worlds), but she does not know the cards of the others (and hence we should maintain arrows between worlds in which she has the same card but the others' cards may be different because she cannot distinguish these worlds). The basic idea we put forward in this paper is that a finite set of triples such as **pickup** is an arrow-changing prescription called an *arrow update* that can be reasoned about within a doxastic modal logic. We develop such a logic and call it AUL.

After presenting the language and semantics (Section §2) and the axiomatics (Section §3) of AUL, we show that arrow updates can be transformed into DEL action models

(Section §4). We work out a number of examples of arrow updates (Section §5) and then prove our key result: arrow updates are sometimes exponentially more succinct than action models (Section §6). We then outline a common knowledge extension of AUL called AUL* (Section §7), and conclude with directions for further research.

§2. Language and semantics.

DEFINITION 2.1 (Agents, Propositional Variables). *We fix a nonempty finite set $\mathcal{A}$ of agents, and we fix an at most countable set $\mathcal{P}$ of propositional variables.*

DEFINITION 2.2 (AUL Language). *Grammar $\mathcal{G}$ is defined as follows.*

$$\begin{aligned}\varphi &::= \perp \mid \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box_a\varphi \mid [U]\varphi \\ U &::= (\varphi, a, \varphi) \mid (\varphi, a, \varphi), U \\ p &\in \mathcal{P}, a \in \mathcal{A}\end{aligned}$$

We will often omit parentheses around expressions when doing so ought not cause confusion. Expressions built using φ as a start symbol in grammar $\mathcal{G}$ are called **AUL-formulas** (or just formulas); we let $\mathcal{F}$ denote the set of formulas. Given formulas φ and φ' and an agent $a \in \mathcal{A}$, the syntactic object (φ, a, φ') is called an a -arrow specification (or just an arrow specification) and is said to have source condition φ , label a , and target condition φ' . Expressions built using U as a start symbol in grammar $\mathcal{G}$ are called **AUL-arrow updates** (or just arrow updates); we let $\mathcal{U}$ denote the set of arrow updates. We identify an arrow update $U \in \mathcal{U}$ with the finite nonempty set of arrow specifications defined by writing set-formation brackets around U . Expressions written using propositional connectives other than $\wedge$ and $\neg$ are to be understood as abbreviations for formulas in the usual way. Given an agent $a \in \mathcal{A}$, an arrow update $U \in \mathcal{U}$, and a formula $\varphi \in \mathcal{F}$, the expression $\Diamond_a\varphi$ abbreviates the formula $\neg\Box_a\neg\varphi$ and the expression $\langle U \rangle\varphi$ abbreviates the formula $\neg[U]\neg\varphi$. An update modal is an expression $[U]$ for which $U \in \mathcal{U}$. To say that a formula is reduced means that the formula does not contain update modals.

The formula $\Box_a\varphi$ is assigned the informal reading “agent a believes φ .” The formula $[U]\varphi$ is assigned the informal reading “after arrow update U , formula φ is true.”

The AUL language is interpreted on Kripke models.

DEFINITION 2.3 (Kripke Model). *To say that M is a Kripke model means that M is a tuple (W^M, R^M, V^M) consisting of a nonempty set W^M of worlds in M , a function $R^M : \mathcal{A} \times W^M \rightarrow \wp(W^M)$ mapping each agent–world pair $(a, w) \in \mathcal{A} \times W^M$ to a set $R_a^M(w) \in \wp(W^M)$ of worlds in M , and a function $V^M : \mathcal{P} \rightarrow \wp(W^M)$ mapping each propositional variable $p \in \mathcal{P}$ to a set $V^M(p) \in \wp(W^M)$ of worlds in M . A pointed Kripke model is a pair (M, w) consisting of a Kripke model M and a world $w \in W^M$; the world w is called the point of (M, w) .*

Given a Kripke model M and an agent $a \in \mathcal{A}$, the function R_a^M induces the binary relation $\bar{R}_a^M \stackrel{\text{def}}{=} \{(w, w') \mid w' \in R_a^M(w)\}$ on W^M . The members of $\bar{R}_a^M$ are called a -arrows.

DEFINITION 2.4 (AUL Semantics). *Given a pointed Kripke model (M, w) and a formula φ , we write $M, w \models \varphi$ to mean that φ is true at (M, w) , and we write $M, w \not\models \varphi$ for the negation of $M, w \models \varphi$. The relation $\models$ is defined by the following induction on formula construction.*

$$\begin{array}{ll}
M, w \not\models \perp & \text{and } M, w \models \top \\
M, w \models p & \text{iff } w \in V^M(p) \text{ for } p \in \mathcal{P} \\
M, w \models \neg\phi & \text{iff } M, w \not\models \phi \\
M, w \models \phi \wedge \psi & \text{iff } M, w \models \phi \text{ and } M, w \models \psi \\
M, w \models \Box_a \phi & \text{iff } M, v \models \phi \text{ for each } v \in R_a^M(w) \\
M, w \models [U]\phi & \text{iff } (M * U), w \models \phi \\
W^{M*U} & \stackrel{\text{def}}{=} W^M \\
R_a^{M*U}(v) & \stackrel{\text{def}}{=} \{v' \in R_a^M(v) \mid \exists(\phi, a, \phi') \in U : \\
& \quad (M, v \models \phi \text{ and } M, v' \models \phi')\} \\
V^{M*U}(p) & \stackrel{\text{def}}{=} V^M(p)
\end{array}$$

To say that a formula ϕ is valid in a Kripke model M , written $M \models \phi$, means that $M, w \models \phi$ for each world $w \in W^M$. To say that a formula ϕ is valid, written $\models \phi$, means that $M \models \phi$ for each Kripke model M . The negation of $\models \phi$ is written $\not\models \phi$.

To see how the semantics works, consider our card example from the previous section. We define the set $\mathcal{A}$ of agents, the set C of cards, the set $\mathcal{P}$ of propositional variables, the Kripke model M , and the arrow update pickup as follows.

$$\begin{array}{ll}
\mathcal{A} & \stackrel{\text{def}}{=} \{1, 2, 3\} \\
C & \stackrel{\text{def}}{=} \{r, w, b\} \\
\mathcal{P} & \stackrel{\text{def}}{=} \{c_a \mid c \in C \text{ and } a \in \mathcal{A}\} \\
W^M & \stackrel{\text{def}}{=} \{w : \mathcal{A} \rightarrow C \mid w \text{ is a bijection}\} \\
R_a^M(w) & \stackrel{\text{def}}{=} W^M \\
V^M(c_a) & \stackrel{\text{def}}{=} \{w \in W^M \mid w(a) = c\} \\
\text{pickup} & \stackrel{\text{def}}{=} \{(c_a, a, c_a) \mid c \in C \text{ and } a \in \mathcal{A}\}
\end{array}$$

For convenience, we identify each function $w \in W^M$ with the expression $w(1)w(2)w(3)$. Hence $\text{rwb} \in W^M$ is the function of type $\mathcal{A} \rightarrow C$ satisfying $\text{rwb}(1) = r$, $\text{rwb}(2) = w$, $\text{rwb}(3) = b$. Model M is pictured on the left in Figure 1.

Assuming that “rwb” is the deal, let us check that agent 1 believes that she has card r after the agents pick up their cards. That is, we verify that $M, \text{rwb} \models [\text{pickup}]\Box_1 r_1$. Proceeding, it follows by the definition of $M * \text{pickup}$ that

$$R_1^{M*\text{pickup}}(\text{rwb}) = \{w \in W^M \mid w(1) = r\},$$

from which it follows that $M * \text{pickup}, \text{rwb} \models \Box_1 r_1$. (Model $M * \text{pickup}$ is pictured on the right in Figure 1.) But then it follows by the semantics that $M, \text{rwb} \models [\text{pickup}]\Box_1 r_1$, as desired.

§3. Theory. In this section, we define an axiomatic theory called AUL that is sound and complete with respect to the semantics from the previous section. The completeness proof follows the *reduction axiom method* of Dynamic Epistemic Logic (Baltag & Moss, 2004; Baltag *et al.*, 1998; Gerbrandy, 1999; Plaza, 2007; van Benthem *et al.*, 2006; van Ditmarsch *et al.*, 2007). Before we define the theory AUL, we first introduce a notion of arrow update composition.

AXIOM SCHEMES

CL. Classical Propositional Logic

BK. $\Box_a(\varphi \rightarrow \psi) \rightarrow (\Box_a\varphi \rightarrow \Box_a\psi)$ U1. $[U]q \leftrightarrow q$ for $q \in \mathcal{P} \cup \{\perp, \top\}$ U2. $[U]\neg\varphi \leftrightarrow \neg[U]\varphi$ U3. $[U](\varphi \wedge \psi) \leftrightarrow ([U]\varphi \wedge [U]\psi)$ U4. $[U]\Box_a\varphi \leftrightarrow \bigwedge_{(\psi, a, \chi) \in U} (\psi \rightarrow \Box_a(\chi \rightarrow [U]\varphi))$ U5. $[U][U']\varphi \leftrightarrow [U \circ U']\varphi$

RULES

$\frac{\varphi \rightarrow \psi \quad \varphi}{\psi}$ (MP)	$\frac{\varphi}{\Box_a\varphi}$ (BN)	$\frac{\varphi}{[U]\varphi}$ (UN)
--	--------------------------------------	-----------------------------------

Table 1. *The theory AUL*

DEFINITION 3.1 (Composition). *Let U and U' be arrow updates. The composition of U with U' , written $U \circ U'$, is the arrow update defined by setting*

$$U \circ U' \stackrel{\text{def}}{=} \{(\varphi \wedge [U]\varphi', a, \psi \wedge [U]\psi') \mid \exists(\varphi, a, \psi) \in U \text{ and } \exists(\varphi', a, \psi') \in U'\}.$$

DEFINITION 3.2 (AUL Theory). *The axiomatic theory AUL is defined in Table 1. We write $\text{AUL} \vdash \varphi$ (or sometimes just $\vdash \varphi$) to mean that the formula φ is derivable in the axiomatic theory AUL; the negation of $\text{AUL} \vdash \varphi$ is written $\text{AUL} \nvdash \varphi$ (or sometimes just $\nvdash \varphi$). To say that a set $S \subseteq \mathcal{F}$ of formulas is AUL-inconsistent (or just inconsistent) means that there is a finite subset $S' \subseteq S$ such that $\vdash \neg \bigwedge S'$, where $\bigwedge S' \stackrel{\text{def}}{=} \bigwedge_{\varphi \in S'} \varphi$ if $S' \neq \emptyset$ and $\bigwedge_{\varphi \in \emptyset} \varphi \stackrel{\text{def}}{=} \top$. To say that a set of formulas is AUL-consistent (or just consistent) means that the set is not inconsistent. Consistency or inconsistency of a formula refers to the consistency or inconsistency of the singleton set containing the formula.*

Axioms U1–U5 are called *reduction axioms*. When reading these axioms from left to right, we note that the “complexity” of the formulas to which the update modal applies decreases; in U1, the update modal is eliminated entirely. As for the intuitive meanings of the reduction axioms, U1 says that atomic facts about the world do not change due to arrow updates. U2 expresses the fact that arrow updates are functional: there is only one way that an arrow update can update a Kripke model. U3 says that arrow updates distribute over conjunction. U4 characterizes an agent’s beliefs after an arrow update in terms of her beliefs before the update: a believes φ after arrow update U if and only if $[U]\varphi$ is true in all of the worlds that can be reached by an a -arrow satisfying an a -arrow specification in U .¹ Axiom U5 says that the effect of a sequence of two arrow updates is equivalent to the effect of their composition.

THEOREM 3.3 (AUL Soundness). *$\text{AUL} \vdash \varphi$ implies $\models \varphi$ for each $\varphi \in \mathcal{F}$.*

¹ To say that an a -arrow $(w, w') \in \bar{R}_a^M$ in a Kripke model M satisfies an a -arrow specification (φ, a, φ') means that the source condition φ is true at (M, w) and the target condition φ' is true at (M, w') .

Proof. By induction on the length of derivation in AUL. We restrict our attention to the base cases for Axioms U4 and U5; the other cases are straightforward. We begin with the proof that the left-to-right implication of Axiom U4 is valid. Proceeding, choose an arbitrary pointed Kripke model (M, w) satisfying $M, w \models [U]\Box_a\varphi$ and choose an arbitrary a -arrow $(\psi, a, \chi) \in U$ such that $M, w \models \psi$. It follows by the semantics that for each $v \in R_a^{M*U}(w)$, we have $(M * U), v \models \varphi$. But the latter is equivalent to the following: for each $v \in R_a^M(w)$ for which $\exists(\psi', a, \chi') \in U$ such that $M, w \models \psi'$ and $M, v \models \chi'$, we have that $M, v \models [U]\varphi$. Therefore, if we have $M, w' \models \chi$ for a world $w' \in R_a^M(w)$, then, since $(\psi, a, \chi) \in U$ and we assumed that $M, w \models \psi$, it follows that $M, w' \models [U]\varphi$ and hence that $M, w' \models \chi \rightarrow [U]\varphi$. We have therefore shown that $M, w \models \Box_a(\chi \rightarrow [U]\varphi)$ and hence that $M, w \models \psi \rightarrow \Box_a(\chi \rightarrow [U]\varphi)$. Conclusion: the left-to-right implication of Axiom U4 is valid. Let us now argue that the right-to-left implication is valid. Choose an arbitrary pointed Kripke model (M, w) satisfying

$$M, w \models \bigwedge_{(\psi, a, \chi) \in U} (\psi \rightarrow \Box_a(\chi \rightarrow [U]\varphi)) . \quad (1)$$

To prove that $M, w \models [U]\Box_a\varphi$, it suffices for us to show that $M, v \models [U]\varphi$ for each $v \in R_a^{M*U}(w)$. So choose an arbitrary $v \in R_a^{M*U}(w)$. It follows by the definition of R_a^{M*U} that $v \in R_a^M(w)$ and $\exists(\psi, a, \chi) \in U$ such that $M, w \models \psi$ and $M, v \models \chi$. But then it follows by assumption (1) that $M, v \models [U]\varphi$, as desired. Conclusion: Axiom U4 is valid.

We now turn to the base case for Axiom U5. We will show that the axiom is sound by showing that $(M * U) * U' = M * (U \circ U')$ for any Kripke model M . It is clear that the sets of possible worlds of these models are identical and so are the valuations. So it remains to be shown that for all worlds w and agents a that $R_a^{(M*U)*U'}(w) = R_a^{M*(U \circ U')}(w)$. Take an arbitrary w and a and suppose that $v \in R_a^{(M*U)*U'}(w)$ for some v . Therefore $v \in R_a^{M*U}(w)$ and there is an arrow specification $(\varphi', a, \psi') \in U'$ such that $M * U, w \models \varphi'$ and $M * U, v \models \psi'$. So $M, w \models [U]\varphi'$ and $M, v \models [U]\psi'$. Moreover we can conclude that $v \in R_a^M(w)$ and there must also be an arrow specification $(\varphi, a, \psi) \in U$ such that $M, w \models \varphi$ and $M, v \models \psi$. Therefore $M, w \models \varphi \wedge [U]\varphi'$ and $M, v \models \psi \wedge [U]\psi'$. By Definition 3.1, we have $(\varphi \wedge [U]\varphi', a, \psi \wedge [U]\psi') \in U \circ U'$ and hence $v \in R_a^{M*(U \circ U')}(w)$. This argument also works the other way around. Therefore $R_a^{(M*U)*U'}(w) = R_a^{M*(U \circ U')}(w)$, and so $(M * U) * U' = M * (U \circ U')$.

Now choose an arbitrary pointed Kripke model (M, w) and suppose that $M, w \models [U][U']\varphi$. Hence $M * U \models [U']\varphi$ and $(M * U) * U', w \models \varphi$. Therefore $M * (U \circ U'), w \models \varphi$ and so $M, w \models [U \circ U']\varphi$. The argument also holds the other way around. Hence Axiom U5 is valid. $\square$

Our proof of the completeness of AUL generally follows the proof in Dynamic Epistemic Logic (van Ditmarsch *et al.*, 2007). In particular, we begin by defining a notion of formula and arrow update “complexity” and then prove by induction on the complexity of formulas that every formula is provably equivalent to a reduced (i.e., update modal-free) formula. This allows us to reduce the completeness of AUL to the completeness of its underlying multimodal logic K (AUL without U1–U5 and UN and applied to reduced formulas only).

DEFINITION 3.4 (AUL Complexity). We define a function $c : (\mathcal{F} \cup \mathcal{U}) \rightarrow \mathbb{N}$ as follows.

$$\begin{aligned} c(q) &\stackrel{\text{def}}{=} 1 \text{ for } q \in \mathcal{P} \cup \{\perp, \top\} \\ c(\neg\varphi) &\stackrel{\text{def}}{=} 1 + c(\varphi) \end{aligned}$$

$$\begin{aligned}
c(\varphi \wedge \psi) &\stackrel{\text{def}}{=} 1 + \max\{c(\varphi), c(\psi)\} \\
c(\Box_a \varphi) &\stackrel{\text{def}}{=} 1 + c(\varphi) \\
c([U]\varphi) &\stackrel{\text{def}}{=} (c(U) + 2) \cdot c(\varphi) \\
c(\{(\varphi, a, \varphi')\}) &\stackrel{\text{def}}{=} 1 + \max\{c(\varphi), c(\varphi')\} \\
c(\{(\varphi, a, \varphi'), U\}) &\stackrel{\text{def}}{=} 1 + \max\{c(\varphi), c(\varphi'), c(U)\}
\end{aligned}$$

For each $X \in \mathcal{F} \cup \mathcal{U}$, the number $c(X)$ is called the complexity of X .

THEOREM 3.5 (AUL Reduction). *For each $\varphi \in \mathcal{F}$, there is a reduced $r(\varphi) \in \mathcal{F}$ such that $\text{AUL} \vdash \varphi \leftrightarrow r(\varphi)$.*

Proof. By induction on the complexity of formulas, one can show that the equations in Table 2 define a function $r : \mathcal{F} \rightarrow \mathcal{F}$. To show this, one of the key results that one must argue is that the equations in Table 2 are *complexity respecting*, by which we mean that r is applied on the left side of an equation to a syntactic object whose complexity is strictly greater than any syntactic object on the right side of the equation to which r is applied. Arguing the equations are complexity respecting is a tedious exercise in reasoning with inequalities; we only prove that $c([U][U']\varphi) > c([U \circ U']\varphi)$. We have the following.

$$\begin{aligned}
&c(U \circ U') \\
&= 1 + \max_{(\varphi, a, \psi) \in U, (\varphi', a, \psi') \in U'} \{c(\varphi \wedge [U]\varphi'), c(\psi \wedge [U]\psi')\} \\
&= 2 + \max_{(\varphi, a, \psi) \in U, (\varphi', a, \psi') \in U'} \{c(\varphi), c(\psi), (c(U) + 2) \cdot c(\varphi'), (c(U) + 2) \cdot c(\psi')\} \\
&\leq 2 + \max\{c(U), (c(U) + 2) \cdot c(U')\} \\
&= 2 + c(U) \cdot c(U') + 2 \cdot c(U')
\end{aligned}$$

Hence $c([U \circ U']\varphi) \leq 4 \cdot c(\varphi) + c(U) \cdot c(U') \cdot c(\varphi) + 2 \cdot c(U') \cdot c(\varphi)$. Further, it is easy to see that $c([U][U']\varphi) = c(U) \cdot c(U') \cdot c(\varphi) + 2 \cdot c(U) \cdot c(\varphi) + 2 \cdot c(U') \cdot c(\varphi) + 4 \cdot c(\varphi)$. It follows that $c([U][U']\varphi) > c([U \circ U']\varphi)$. To prove the statement of the theorem, one then argues by induction on the complexity k of formulas that for each formula θ having $c(\theta) \leq k$, it follows that $\vdash \theta \leftrightarrow r(\theta)$. The argument is broken up into a number of cases, one for each of the forms to which r is applied on the left side of an equation in Table 2.

$$\begin{aligned}
r(q) &\stackrel{\text{def}}{=} q \text{ for } q \in \mathcal{P} \cup \{\perp, \top\} \\
r(\neg\varphi) &\stackrel{\text{def}}{=} \neg r(\varphi) \\
r(\varphi \wedge \psi) &\stackrel{\text{def}}{=} r(\varphi) \wedge r(\psi) \\
r(\Box_a \varphi) &\stackrel{\text{def}}{=} \Box_a r(\varphi) \\
r([U]q) &\stackrel{\text{def}}{=} q \text{ for } q \in \mathcal{P} \cup \{\perp, \top\} \\
r([U]\neg\varphi) &\stackrel{\text{def}}{=} \neg r([U]\varphi) \\
r([U](\varphi \wedge \psi)) &\stackrel{\text{def}}{=} r([U]\varphi) \wedge r([U]\psi) \\
r([U]\Box_a \varphi) &\stackrel{\text{def}}{=} \bigwedge_{(\psi, a, \chi) \in U} r(\psi) \rightarrow \Box_a (r(\chi) \rightarrow r([U]\varphi)) \\
r([U][U']\varphi) &\stackrel{\text{def}}{=} r([U \circ U']\varphi)
\end{aligned}$$

Table 2. Definition of $r : \mathcal{F} \rightarrow \mathcal{F}$

Let us check the induction case for the form $[U]\Box_a\varphi$. Proceeding, we have

$$\vdash (\bigwedge_{(\psi,a,\chi)\in U} \psi \rightarrow \Box_a(\chi \rightarrow [U]\varphi)) \leftrightarrow (\bigwedge_{(\psi,a,\chi)\in U} r(\psi) \rightarrow \Box_a(r(\chi) \rightarrow r([U]\varphi)))$$

by the induction hypothesis—which is applicable because the equations in Table 2 are complexity respecting—and modal reasoning. By Axiom U4, the definition of r , and propositional logic, it follows that $\vdash [U]\Box_a\varphi \leftrightarrow r([U]\Box_a\varphi)$. The other cases are handled similarly. $\square$

THEOREM 3.6 (AUL Completeness). $\models \varphi$ implies $\mathbf{AUL} \vdash \varphi$ for each $\varphi \in \mathcal{F}$.

Proof. Suppose we are given an arbitrary $\varphi \in \mathcal{F}$ satisfying $\models \varphi$. By Soundness (Theorem 3.3) and Reduction (Theorem 3.5), it follows that $\models r(\varphi)$. Since $r(\varphi)$ is reduced, we have $\mathbf{K} \vdash r(\varphi)$ by the completeness of the multi-modal logic $\mathbf{K}$ (Blackburn *et al.*, 2001) and hence $\mathbf{AUL} \vdash r(\varphi)$ because $\mathbf{AUL}$ extends $\mathbf{K}$. Applying Reduction once more, it follows by propositional reasoning that $\mathbf{AUL} \vdash \varphi$. $\square$

§4. Arrow updates and action models. In this section, we show every arrow update has the “same effect” as a certain *action model* of Dynamic Epistemic Logic. Devised by Baltag *et al.* (1998), an action model is a Kripke model-like object that describes the agents’ beliefs about incoming information. Instead of a set of possible worlds, an action model A has a set E^A of possible “events,” each of which comes with an assertion $\text{pre}^A(e)$, called the *precondition*, that might be made. But the agents are uncertain as to which assertion will in fact be made: given an agent a and an event e , there is a set $R_a^A(e)$ of events whose preconditions the agent thinks might have been asserted whenever the precondition $\text{pre}^A(e)$ of event e is in fact asserted. In this way, an action model represents a set of possible events, each of which conveys the information content of a certain assertion, but the agents are uncertain as to which event is taking place.

DEFINITION 4.1 (Action Model). Let F be a set of formulas. To say that A is an F -action model means that A is a tuple (E^A, R^A, pre^A) consisting of a nonempty finite set E^A whose members will be called events in A , a function $R^A : \mathcal{A} \times E^A \rightarrow \wp(E^A)$ mapping each agent–event pair $(a, e) \in \mathcal{A} \times E^A$ to a set $R_a^A(e) \in \wp(E^A)$ of events in A , and a function $\text{pre}^A : E^A \rightarrow F$ mapping each event $e \in E^A$ to a formula $\text{pre}^A(e) \in F$ called the precondition of e . A multi-pointed F -action model is a pair (A, E) consisting of an action model and a set $E \subseteq E^A$ of events; each $e \in E$ is called a point of (A, E) . If F' is a set of formulas, then we write $\mathfrak{A}(F')$ to denote the set of multi-pointed F' -action models.

The points E of a multi-pointed action model (A, E) represent the set of possible events that may be executed according to a nondeterministic choice. The language of DEL is obtained from that of multi-modal logic by adding multi-pointed action modals as modalities.

DEFINITION 4.2 (DEL Language). We define grammar $\mathcal{G}_0$ as follows.

$$\begin{aligned} \varphi &::= \perp \mid \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box_a\varphi \\ p &\in \mathcal{P}, a \in \mathcal{A} \end{aligned}$$

We define $\mathcal{F}_0$ to be the set of expressions built using φ as a start symbol in grammar $\mathcal{G}_0$. Then, whenever $\mathcal{F}_k$ is defined, we define grammar $\mathcal{G}_{k+1}$ as follows.

$$\begin{aligned} \varphi &::= \psi \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box_a\varphi \mid [A, E]\varphi \\ \psi &\in \mathcal{F}_k, a \in \mathcal{A}, (A, E) \in \mathfrak{A}(\mathcal{F}_k) \end{aligned}$$

We define $\mathcal{F}_{k+1}$ to be the set of expressions built using φ as a start symbol in grammar $\mathcal{G}_{k+1}$. Finally, we define the set $\mathcal{F}_{\text{DEL}} \stackrel{\text{def}}{=} \bigcup_{k \in \mathbb{N}} \mathcal{F}_k$ whose members will be called DEL-formulas. We adopt conventions and terminology similar to those established for the AUL language (Definition 2.2). To say that a DEL-formula φ is reduced means $\varphi \in \mathcal{F}_0$. An action model is an $\mathcal{F}_{\text{DEL}}$ -action model. A multi-pointed action model is a multi-pointed $\mathcal{F}_{\text{DEL}}$ -action model.

The language of DEL is also interpreted on Kripke models. The core semantic definition is the so-called *product update* (Baltag & Moss, 2004; Baltag *et al.*, 1998), which determines the truth of a formula $[A, E]\varphi$ at a pointed Kripke model (M, w) by constructing a new Kripke model $M[A]$.

DEFINITION 4.3 (DEL Semantics). *Given a pointed Kripke model (M, w) and a DEL-formula φ , we write $M, w \models \varphi$ to mean that φ is true at (M, w) , and we write $M, w \not\models \varphi$ for the negation of $M, w \models \varphi$. The relation $\models$ is defined by the following induction on DEL-formula construction.*

$M, w \not\models \perp$	and	$M, w \models \top$
$M, w \models p$	iff	$w \in V^M(p)$ for $p \in \mathcal{P}$
$M, w \models \neg\varphi$	iff	$M, w \not\models \varphi$
$M, w \models \varphi \wedge \psi$	iff	$M, w \models \varphi$ and $M, w \models \psi$
$M, w \models \Box_a\varphi$	iff	$M, v \models \varphi$ for each $v \in R_a^M(w)$
$M, w \models [A, E]\varphi$	iff	$M, w \models \text{pre}^A(e)$ implies $M[A], (w, e) \models \varphi$ for all $e \in E$
$W^{M[A]}$	$\stackrel{\text{def}}{=}$	$\{(v, f) \in W^M \times E^A \mid M, v \models \text{pre}^A(f)\}$
$R_a^{M[A]}((v, f))$	$\stackrel{\text{def}}{=}$	$\{(v', f') \in W^{M[A]} \mid v' \in R_a^M(v) \text{ \& } f' \in R_a^A(f)\}$
$V^{M[A]}(p)$	$\stackrel{\text{def}}{=}$	$\{(v, f) \in W^{M[A]} \mid v \in V^M(p)\}$

Validity for DEL-formulas is defined as for AUL-formulas (Definition 2.4).

By Theorem 3.5 and by a similar theorem for DEL (Baltag & Moss, 2004; Baltag *et al.*, 1998; van Ditmarsch *et al.*, 2007), AUL and DEL can each be reduced to the underlying multi-modal logic K, by which we mean that each formula from either system is semantically equivalent to a K-formula. (The K-formulas are just the reduced formulas—the members of $\mathcal{F}_0$.) This tells us that each of AUL and DEL may be viewed as a repackaging of K. However, things are not so simple: the translation from DEL to K will bring about an exponential blow-up of certain formulas (Lutz, 2006), and it would be surprising if the same were not the case for AUL (though whether it is indeed the case is an open question we leave for future work).

While the common K-link provides one connection between AUL and DEL, we now wish to study another. In particular, we wish to understand the relationship between AUL's arrow updates and DEL's action models. As a first step in understanding this relationship, we will show that every arrow update can be transformed into an action model that has the “same effect.” This relates arrow updates to action models, telling us that every arrow update is expressible by an action model. An open question for future work is to characterize the reverse expressive relationship: which action models are expressible using arrow updates?

Our sense of an action model having the “same effect” as an arrow update is given by the following notion of *update equivalence*, which is a simple adaptation of a related

notion of *action equivalence* from Dynamic Epistemic Logic (van Eijck *et al.*, 2008) to the framework of this paper.

DEFINITION 4.4 (Update Equivalence; adapted from van Eijck *et al.*, 2008). A u-modal for an action model A is any expression of the form $[A, E]$ for some $E \subseteq E^A$, a u-modal for a multi-pointed action model (A', E') is the expression $[A', E']$, and a u-modal for an arrow update U is the expression $[U]$. For each $i \in \{1, 2\}$, let X_i be either an action model A_i , a multi-pointed action model (A_i, E_i) or an arrow update U_i . To say that X_1 and X_2 are update equivalent means that there is a u-modal $[X_1]$ for X_1 and there is a u-modal $[X_2]$ for X_2 such that, for each pointed Kripke model (M, w) and each reduced formula $\varphi \in \mathcal{F}_0$, we have $M, w \models [X_1]\varphi$ if and only if $M, w \models [X_2]\varphi$.

We now prove that every arrow update U can be transformed into an update equivalent action model $A[U]$ defined as follows.

DEFINITION 4.5 (Maximal U -Consistency). Given $U \in \mathcal{U}$, define $\Phi(U) \subseteq \mathcal{F}$ by setting

$$\Phi(U) \stackrel{\text{def}}{=} \{\varphi \in \mathcal{F} \mid \exists(\varphi, a, \varphi') \in U\} \cup \{\varphi' \in \mathcal{F} \mid \exists(\varphi, a, \varphi') \in U\},$$

and define $\Phi^\pm(U) \subseteq \mathcal{F}$ by setting $\Phi^\pm(U) \stackrel{\text{def}}{=} \Phi(U) \cup \{\neg\varphi \mid \varphi \in \Phi(U)\}$. A U -set is a subset $\Gamma \subseteq \Phi^\pm(U)$. To say that Γ is maximal U -consistent means that Γ is a consistent U -set and any U -set $\Gamma' \supset \Gamma$ is inconsistent.

DEFINITION 4.6 (Action Model $A[U]$). Given an arrow update $U \in \mathcal{U}$, define the $\mathcal{F}$ -action model $A[U]$ as follows.

$$\begin{aligned} E^{A[U]} &\stackrel{\text{def}}{=} \{\Gamma \mid \Gamma \text{ is maximal } U\text{-consistent}\} \\ R_a^{A[U]}(\Gamma) &\stackrel{\text{def}}{=} \{\Gamma' \in E^{A[U]} \mid \exists(\varphi, a, \varphi') \in U : (\varphi \in \Gamma \text{ and } \varphi' \in \Gamma')\} \\ \text{pre}^{A[U]}(\Gamma) &\stackrel{\text{def}}{=} r(\bigwedge \Gamma) \end{aligned}$$

THEOREM 4.7. For each $U \in \mathcal{U}$, the action model $A[U]$ is update equivalent to U .

Proof. We begin by proving that for each Kripke model M and arrow update $U \in \mathcal{U}$, the Kripke models $M * U$ and $M[A[U]]$ are isomorphic.² Proceeding, fix a Kripke model M and an arrow update $U \in \mathcal{U}$. We first show that for each $v \in W^M$, there is a unique $\Gamma_v \in E^{A[U]}$ such that $M, v \models r(\bigwedge \Gamma_v)$. Proceeding, choose an arbitrary $v \in W^M$ and let $\langle \varphi_i \rangle_{i=0}^n$ be an enumeration of $\Phi(U)$. For each $i \in \mathbb{N}$ with $i \leq n$, define the formula ψ_i^v by setting

$$\psi_i^v \stackrel{\text{def}}{=} \begin{cases} \varphi_i & \text{if } M, v \models \varphi_i, \\ \neg\varphi_i & \text{if } M, v \not\models \varphi_i. \end{cases}$$

Defining $\Gamma_v \stackrel{\text{def}}{=} \{\psi_i^v \mid i \in \mathbb{N} \text{ \& } i \leq n\}$, it follows by construction, AUL Reduction (Theorem 3.5), and AUL Soundness (Theorem 3.3) that Γ_v is maximal U -consistent and that $M, v \models r(\bigwedge \Gamma_v)$. Further, if $\Gamma \in E^{A[U]}$ satisfies $\Gamma \neq \Gamma_v$, then it follows by maximal

² To say that Kripke models M and M' are isomorphic means that there exists an isomorphism between M and M' . To say that f is an isomorphism between Kripke models M and M' means that $f : W^M \rightarrow W^{M'}$ is a bijection satisfying the property that for each $w \in W^M$, we have that $v \in R_a^M(w)$ if and only if $f(v) \in R_a^{M'}(f(w))$ and that $w \in V^M(p)$ if and only if $f(w) \in V^{M'}(p)$ for each $p \in \mathcal{P}$.

U -consistency that there is a $j \in \mathbb{N}$ with $j \leq n$ such that either $\varphi_j \in \Gamma$ and $\neg\varphi_j \in \Gamma_v$ or else $\neg\varphi_j \in \Gamma$ and $\varphi_j \in \Gamma_v$, but each of these possibilities implies $M, v \not\models r(\bigwedge \Gamma)$. We conclude that Γ_v is the unique maximal U -consistent U -set satisfying $M, v \models r(\bigwedge \Gamma_v)$.

To prove the statement of the theorem, define the function $f : W^{M*U} \rightarrow W^{M[A[U]]}$ by setting $f(v) \stackrel{\text{def}}{=} (v, \Gamma_v)$. It follows by what we showed in the previous paragraph that f is a bijection; in particular, we see that $(v, \Gamma) \in W^{M[A[U]]}$ implies $M, v \models r(\bigwedge \Gamma)$, which implies that $\Gamma = \Gamma_v$. Let us now show that we have $v' \in R_a^{M*U}(v)$ if and only if $(v', \Gamma_{v'}) \in R_a^{M[A[U]]}((v, \Gamma_v))$. Proceeding, $v' \in R_a^{M*U}(v)$ means that $v' \in R_a^M(v)$ and $\exists(\varphi, a, \varphi') \in U$ such that $M, v \models \varphi$ and $M, v' \models \varphi'$, which is equivalent to the statement that $v' \in R_a^M(v)$ and $\exists(\varphi, a, \varphi') \in U$ such that $\varphi \in \Gamma_v$ and $\varphi' \in \Gamma_{v'}$. But the latter is equivalent to the statement that $v' \in R_a^M(v)$ and $\Gamma_{v'} \in R_a^{A[U]}(\Gamma_v)$. Since $M, v \models \text{pre}^{A[U]}(\Gamma_v)$ and $M, v' \models \text{pre}^{A[U]}(\Gamma_{v'})$, the statement at the end of the previous sentence is equivalent to $(v', \Gamma_{v'}) \in R_a^{M[A[U]]}((v, \Gamma_v))$. Hence we have shown that we have $v' \in R_a^{M*U}(v)$ if and only if $(v', \Gamma_{v'}) \in R_a^{M[A[U]]}((v, \Gamma_v))$. Finally, we observe that for each $p \in \mathcal{P}$, we have $v \in V^{M*U}(p)$ if and only if $v \in V^M(p)$ if and only if $(v, \Gamma_v) \in V^{M[A[U]]}(p)$. Conclusion: f is an isomorphism between the Kripke models $M * U$ and $M[A[U]]$.

So to prove that $A[U]$ and U are update equivalent, choose an arbitrary world $w \in W^M$ and an arbitrary reduced formula $\varphi \in \mathcal{F}_0$. We have $M, w \models [U]\varphi$ if and only if $M * U, w \models \varphi$. But $f(w) = (\Gamma_w, w)$ and hence $M * U, w \models \varphi$ is equivalent to $M[A[U]], (\Gamma_w, w) \models \varphi$ because $\varphi \in \mathcal{F}_0$ and modal truth is invariant under bisimulation (Blackburn *et al.*, 2001) (of which isomorphism is a special case). But we have $M[A[U]], (\Gamma_w, w) \models \varphi$ if and only if $M, w \models [A[U], \{\Gamma_w\}]\varphi$ if and only if $M, w \models [A[U], E^{A[U]}]\varphi$ by what was shown above. Conclusion: $A[U]$ and U are update equivalent. $\square$

We note that the construction from Definition 4.6 yields an action model $A[U]$ that is exponentially larger than the equivalent arrow update U . In Section §6, we will see that sometimes it is not possible to avoid this exponential blow-up. Therefore, while the belief changes described by arrow updates are also describable using action models, arrow updates are sometimes exponentially more succinct than action models in describing belief changes.

§5. Examples. In this section, we show how public announcements are captured in AUL and give an example of a belief change operation whose arrow update implementation is obvious but whose action model implementation is difficult to discern.

5.1. Public announcements. A public announcement of an assertion φ is generally realized in one of two ways: Plaza's (2007) operation $\mathbf{p}(\varphi)$ that deletes all $\neg\varphi$ -worlds (and thereby arrows that target such worlds) and Gerbrandy & Groeneveld's (1997) operation $\mathbf{g}(\varphi)$ that deletes all arrows whose target is a $\neg\varphi$ -world but does not delete any worlds. In AUL the operation $\mathbf{g}(\varphi)$ is implementable using an arrow update we call $\mathbf{APUB}(\varphi)$.

DEFINITION 5.1 (Public Announcements). Let M be a Kripke model and φ be a formula. We define the operations $\mathbf{p}(\varphi) : M \mapsto M[\varphi]_{\mathbf{p}}$ and $\mathbf{g}(\varphi) : M \mapsto M[\varphi]_{\mathbf{g}}$ as follows.

$$\begin{array}{llll} W^{M[\varphi]_{\mathbf{p}}} & \stackrel{\text{def}}{=} & \{w \in W^M \mid M, w \models \varphi\} & W^{M[\varphi]_{\mathbf{g}}} & \stackrel{\text{def}}{=} & W^M \\ R_a^{M[\varphi]_{\mathbf{p}}}(w) & \stackrel{\text{def}}{=} & R_a^M(w) \cap W^{M[\varphi]_{\mathbf{p}}} & R_a^{M[\varphi]_{\mathbf{g}}}(w) & \stackrel{\text{def}}{=} & R_a^M(w) \cap W^{M[\varphi]_{\mathbf{p}}} \\ V^{M[\varphi]_{\mathbf{p}}}(p) & \stackrel{\text{def}}{=} & V^M(p) \cap W^{M[\varphi]_{\mathbf{p}}} & V^{M[\varphi]_{\mathbf{g}}} & \stackrel{\text{def}}{=} & V^M \end{array}$$

Fig. 2. The action model $A[\text{APUB}(\varphi)]$, where neither φ nor $\neg\varphi$ is inconsistent.

DEFINITION 5.2 (AUL Public Announcement). *Given a formula φ , the AUL public announcement of φ is the arrow update $\text{APUB}(\varphi)$ defined by setting $\text{APUB}(\varphi) \stackrel{\text{def}}{=} \{(\top, a, \varphi) \mid a \in \mathcal{A}\}$.*

If neither φ nor $\neg\varphi$ is inconsistent, then $A[\text{APUB}(\varphi)]$ (Definition 4.6) is pictured as in Figure 2. If $\neg\varphi$ is inconsistent, then the picture for $A[\text{APUB}(\varphi)]$ is obtained from that in Figure 2 by deleting node $\{\top, \neg\varphi\}$ along with the right-to-left $\mathcal{A}$ -arrow. If φ is inconsistent, then the picture for $A[\text{APUB}(\varphi)]$ is obtained from that in Figure 2 by deleting node $\{\top, \varphi\}$ along with both $\mathcal{A}$ -arrows.

It is easy to see that for each Kripke model M and each formula φ , we have $M[\varphi]_g = M[\text{APUB}(\varphi)]$. Hence $\text{APUB}(\varphi)$ implements the Gerbrandy–Groeneveld operation $g(\varphi)$. It will therefore be convenient to establish the following notational abbreviation.

DEFINITION 5.3. $[\varphi]_g\psi$ abbreviates $[\text{APUB}(\varphi)]\psi$ for formulas φ and ψ .

To describe Plaza’s announcements using Gerbrandy and Groeneveld’s announcements (and hence using arrow updates), let $\mathcal{F}_{\text{PUB}}$ be the set of all formulas that can be formed after we extend the AUL-formula formation grammar $\mathcal{G}$ (Definition 2.2) by adding the following formula formation rule: from formulas φ and ψ , form the formula $[\varphi]_p\psi$. The definition of the semantical relation $\models$ for $\mathcal{F}_{\text{PUB}}$ -formulas is obtained by extending the semantics for AUL-formulas (Definition 2.4) by adding the following inductive clause (Plaza, 2007; van Ditmarsch *et al.*, 2007).

$$M, w \models [\varphi]_p\psi \quad \text{iff} \quad (M, w \models \varphi \text{ implies } M[\varphi]_p, w \models \psi) \quad (2)$$

Note that in the above equivalence, the expression $[\varphi]_p\psi$ on the left-hand side is a formula and the expression $M[\varphi]_p$ on the right-hand side is the Kripke model defined in Definition 5.1.³ It is not difficult to see that $\models [\varphi]_p\psi \leftrightarrow (\varphi \rightarrow [\varphi]_g\psi)$. Hence both Plaza’s and Gerbrandy and Groeneveld’s announcements can be expressed using arrow updates.

5.2. Cautious update and lying. We begin with a retelling of the well-known *surprise exam paradox* (Clark, 2007).

A teacher tells her students that there will be a surprise exam on a week-day next week. The students reason as follows: “The exam cannot take place on Friday because we would know on Thursday night that the exam would take place on Friday and it would not then be a surprise. Therefore Thursday is the last possible day for the exam. But then the exam cannot take place on Thursday because we would know on Wednesday night that the exam would take place on Thursday and it would not then be a surprise. Therefore Wednesday is the last possible day for the exam.” Continuing their reasoning in this way, the students rule out all week-

³ A Kripke model must have a *nonempty* set of worlds. Therefore, properly speaking, $M[\varphi]_p$ is a Kripke model if and only if $W^{M[\varphi]_p} \neq \emptyset$. Note that the truth of the antecedent $M, w \models \varphi$ of the right-hand side of (2) guarantees this property (and hence guarantees that the structure $M[\varphi]_p$ is in fact a Kripke model).

days, concluding that the exam cannot take place at all. The teacher was lying! But then, much to the students' surprise, the exam takes place on Tuesday.

Gerbrandy (2007) analyzed this paradox in Dynamic Epistemic Logic, showing that the statement the teacher makes may be true. Here our focus is on another aspect of the paradox: what happens to the students' beliefs when they conclude that the teacher is lying?

A student who believes the teacher is lying ought not change her beliefs upon hearing what the teacher says; otherwise, if the student does not think the teacher is lying, then it is reasonable to accept what the teacher says. Put another way: if the student believes the teacher's statement to be false, then the student should ignore the statement and leave her beliefs unchanged; otherwise, if the student believes that the teacher's statement might be true, then the student should trust the teacher and update her beliefs by accepting what the teacher has said. However, in a multi-agent setting, things are a bit more complicated: a student must not only decide whether to process the information in terms of her own beliefs but must also take into account the fact that it is common knowledge that the other students must make their own decisions based on their respective beliefs. This leads to a cautious yet eager belief change policy that we call *cautious updating*. Steiner (2006) first studied cautious updating as a model-changing operation in its own right. However, cautious updates are in fact a special case of arrow updates.

DEFINITION 5.4 (Cautious Update). *Given a formula φ , the cautious update with φ is the arrow update $\text{CU}(\varphi)$ defined by setting*

$$\text{CU}(\varphi) \stackrel{\text{def}}{=} \{(\Diamond_a \varphi, a, \varphi) \mid a \in \mathcal{A}\} \cup \{(\Box_a \neg \varphi, a, \top) \mid a \in \mathcal{A}\}.$$

The idea is that when an agent a believes at a world w that φ is false, the a -arrows leaving w should remain—thereby leaving her beliefs about φ unchanged at w —because she rejects information she believes to be false. But if an agent believes at w that φ might be true, then all a -arrows from w to a non- φ world should be deleted—thereby causing her to accept φ at w —because she accepts information that she believes might be true.

While it is easy to construct an arrow update for cautious updating—indeed, the arrow update $\text{CU}(\varphi)$ comes quite naturally from the intuitive description of what the cautious update with φ ought to do—it is difficult to construct a multi-agent action model for cautious updating. Perhaps illustrative of this difficulty, we see that the action model $A[\text{CU}(\varphi)]$ (Definition 4.6) for the two-agent case $\mathcal{A} = \{1, 2\}$ with neither φ nor $\neg\varphi$ inconsistent is already quite complicated: see Figure 3, in which we have simplified the labeling of nodes by omitting occurrences of formulas that are AUL-provably equivalent to other formulas already in the set.

We note that the action model in Figure 3 contains “spurious arrows”: if M is a Kripke model, $\Box_a \neg \varphi$ is a conjunct of a precondition of an event $e \in E^{A[\text{CU}(\varphi)]}$, and φ is a conjunct of a precondition of another event $e' \in E^{A[\text{CU}(\varphi)]}$, then there will never be an a -arrow from a world in $M[A[\text{CU}(\varphi)]]$ of the form (w, e) to a world in $M[A[\text{CU}(\varphi)]]$ of the form (w', e') because $M, w \models \text{pre}^{A[\text{CU}(\varphi)]}(e)$ implies $M, w \models \Box_a \neg \varphi$, which implies $M, v \not\models \varphi$ for each $v \in R_a^M(w)$ and hence $w' \notin R_a^M(w)$. Therefore such “spurious arrows” may be deleted, yielding a “reduced” action model $A'[\text{CU}(\varphi)]$ that is update equivalent to $A[\text{CU}(\varphi)]$.

Cautious updating is a natural way to react to information that comes from a generally trustworthy but occasionally faulty source: accept the information if you believe it might

Fig. 3. The action model $A[\text{CU}(\varphi)]$ for agents 1 (solid) and 2 (dashed), where neither φ nor $\neg\varphi$ is inconsistent; node labels have been simplified.

be true and reject the information if you believe it is false. In particular, if the source lies by asserting a false statement φ , then cautious updating will not lead astray those agents who initially believe that φ is false: these agents will simply ignore the assertion all together. We contrast this with the approach of van Ditmarsch *et al.* (2010), who propose that agents respond to a lie that asserts a false statement φ by executing Gerbrandy and Groeneveld’s public announcement operation $\mathbf{g}(\varphi)$: delete all arrows whose target is a $\neg\varphi$ -world, thereby accepting the false information φ without regard to one’s prior beliefs about φ . One consequence of this approach is that an agent who initially believes that φ is false will respond to the lie asserting φ by coming to believe everything (including inconsistent statements), thereby trivializing the agent’s beliefs. Cautious updating avoids this outcome and is therefore a better policy for responding to information coming from a trustworthy but occasionally faulty source.⁴

The number of events in the arrow update $\text{CU}(\varphi)$ grows linearly in the cardinality of $\mathcal{A}$, whereas the action model $A[\text{CU}(\varphi)]$ grows exponentially in the cardinality of $\mathcal{A}$. In the next section, we will show that the difference in the rates of growth between arrow updates and their equivalent Definition 4.6 action models is sometimes unavoidable.

§6. Succinctness. Though arrow updates are no more expressive than action models (Theorem 4.7), we will see in this section that *arrow updates are sometimes exponentially more succinct than action models*. That is, we will show that there is a sequence of arrow updates such that k -th arrow update is of size $\Theta(k)$ but the smallest equivalent action model

⁴ Recent work on Dynamic Epistemic Logic frameworks for Belief Revision (see, e.g., Baltag & Smets, 2007; van Benthem, 2004) would allow us to specify more subtle policies wherein an agent tentatively accepts information she believes might be possible but will later give up this belief if further reliable information contradicts it. However, since our focus here is on Kripke model operations—as opposed to operations on the “plausibility models” of the DEL Belief Revision literature—the study of connections with and applications to these more flexible belief revision frameworks will be left for future work.

is of size $2^{\Theta(k)}$. In order to prove this claim, we will first define the notions of the length for formulas and of size for arrow updates and for action models.

DEFINITION 6.1 (Length). We define the function $\text{len} : \mathcal{F} \cup \mathcal{U} \rightarrow \mathbb{N}$ by letting $\text{len}(\varphi)$ be the number of occurrences of each symbol appearing in the expression φ ; this includes parentheses, brackets, and commas and, for each $a \in \mathcal{A}$ and each $p \in \mathcal{P}$, counts each of the expressions “ $\Box_a$ ”, “ a ”, and “ p ” as a single symbol. For each $X \in \mathcal{F} \cup \mathcal{U}$, we call $\text{len}(X)$ the length of X .

DEFINITION 6.2 (Size). Given an action model A and an arrow update U , we define the following natural-number quantities.

$$\begin{aligned}
 n(A) &\stackrel{\text{def}}{=} |E^A| && \text{(number of events in } A) \\
 m(A) &\stackrel{\text{def}}{=} \sum_{a \in \mathcal{A}} |\bar{R}_a^A| && \text{(number of arrows in } A) \\
 p(A) &\stackrel{\text{def}}{=} \sum_{e \in E^A} \text{len}(\text{pre}^A(e)) && \text{(length of preconditions in } A) \\
 s(A) &\stackrel{\text{def}}{=} m(A) + p(A) && \text{(size of } A) \\
 s((\varphi, a, \varphi')) &\stackrel{\text{def}}{=} \text{len}(\varphi) + \text{len}(\varphi') && \text{(size of } (\varphi, a, \varphi')) \\
 s((\varphi, a, \varphi'), U) &\stackrel{\text{def}}{=} s((\varphi, a, \varphi')) + s(U) && \text{(size of } (\varphi, a, \varphi'), U)
 \end{aligned}$$

We have $p(A) \geq n(A)$ because $\text{len}(\text{pre}^A(e)) \geq 1$ for each $e \in E^A$. Hence $s(A) \geq n(A)$.

We now define a sequence $\{U_k\}_{k \in \mathbb{N}}$ of arrow updates such that U_k has size $\Theta(k)$ and the Theorem 4.7 equivalent action model $A_k \stackrel{\text{def}}{=} A[U_k]$ has $2^{\Theta(k)}$ events.

DEFINITION 6.3 (U_k, A_k). Fix an agent $a \in \mathcal{A}$ and assume that $\mathcal{P} \stackrel{\text{def}}{=} \{p_i \mid i \in \mathbb{N}\}$. For each $k \in \mathbb{N}$, we define the arrow update U_k by setting $U_k \stackrel{\text{def}}{=} \{(\neg p_i, a, p_i) \mid i \in \mathbb{N} \text{ and } i \leq k\}$, and we define the action model A_k by setting $A_k \stackrel{\text{def}}{=} A[U_k]$ according to Definition 4.6. Writing $\bar{p}_i$ as an abbreviation for $\neg p_i$ and $x_0 x_1 x_2 \dots x_n$ as an abbreviation for $\bigwedge_{i=0}^n x_i$, the action models A_0 , A_1 , and A_2 are pictured in Figure 4; note that we use dotted arrows simply to improve the visual presentation (every arrow in A_k is an a -arrow).

LEMMA 6.4 (Update Size). For each $k \in \mathbb{N}$, we have $s(U_k) = 3k+3$ and $n(A_k) = 2^{k+1}$.

Proof. For each $i \in \mathbb{N}$ with $i \leq k$, the triple $(\neg p_i, a, p_i) \in U$ contributes $\text{len}(\neg p_i) + \text{len}(p_i) = 3$ to the size. Since $k+1$ pairwise distinct triples of this form occur in U , it follows that $s(U_k) = 3k+3$. As for the size of A_k , we have $\Phi^\pm(U_k) = \{p_i, \neg p_i, \neg\neg p_i \mid$

Fig. 4. Action models A_0 (left), A_1 (middle), and A_2 (right) from Definition 6.3; some arrows are dotted for visual clarity.

$i \in \mathbb{N}$ and $i \leq k$. Hence for each $i \in \mathbb{N}$ with $i \leq k$, a maximal U_k -consistent set has exactly one of the two sets $\{p_i, \neg\neg p_i\}$ and $\{\neg p_i\}$ as a subset. Since a maximal U_k -consistent set is a union of $k + 1$ pairwise disjoint sets, each of which has one of these two distinct forms, it follows that $n(A_k) = 2^{k+1}$. $\square$

It will be our task to show that an action model having fewer than $n(A_k)$ states cannot be update equivalent to U_k . For convenience, we first define an auxiliary notion of modal equivalence between pointed Kripke models called $\Box_a$ -equivalence.

DEFINITION 6.5. *Given $a \in \mathcal{A}$, to say that pointed Kripke models (M, w) and (M', w') are $\Box_a$ -equivalent means that for each reduced formula φ , we have $M, w \models \Box_a \varphi$ if and only if $M', w' \models \Box_a \varphi$.*

We now prove that any action modal having fewer than $n(A_k) = 2^{\Theta(k)} = 2^{\Theta(s(U_k))}$ states is not update equivalent to U_k . Since the size of an action model is no less than the number of states in the action model, it therefore follows that no action model whose size is less than $2^{\Theta(s(U_k))}$ is update equivalent to U_k . Hence arrow updates are sometimes exponentially more succinct than action models.

THEOREM 6.6 (Update Succinctness). *Assume that $\mathcal{P} \stackrel{\text{def}}{=} \{p_i \mid i \in \mathbb{N}\}$. For each $k \in \mathbb{N}$, no action model A having $n(A) < n(A_k) = 2^{k+1} = 2^{\Theta(s(U_k))}$ is update equivalent to U_k .*

Proof. Fix $k \in \mathbb{N}$ and $a \in \mathcal{A}$ such that a appears in U_k . Let C_k denote the k -dimensional epistemic hypercube; that is, defining $\bar{k} \stackrel{\text{def}}{=} \{i \in \mathbb{N} \mid i \leq k\}$, we let C_k be the Kripke model defined as follows.

$$\begin{aligned} W^{C_k} &\stackrel{\text{def}}{=} \wp(\bar{k}) \\ R_b^{C_k}(v) &\stackrel{\text{def}}{=} W^{C_k} \text{ for each } b \in \mathcal{A} \\ V^{C_k}(p_i) &\stackrel{\text{def}}{=} \{v \in W^{C_k} \mid i \in v\} \end{aligned}$$

Note that every world $w \in W^{C_k}$ has a unique event in $e_w \in E^{A_k}$ such that $C_k, w \models \text{pre}^{A_k}(e_w)$.

Suppose that w_1 and w_2 are two different worlds in C_k . Therefore there is an $i \in \bar{k}$ such that, without loss of generality, $w_1 \in V^{C_k}(p_i)$ and $w_2 \notin V^{C_k}(p_i)$. Hence $(w_1, e_{w_1}) \in R_a^{C_k[A_k]}((w_2, e_{w_2}))$ because $(\neg p_i, a, p_i) \in U_k$. So

$$C_k, w_2 \not\models [A_k, E^{A_k}] \Box_a \neg \text{pre}^{A_k}(e_{w_1}) .$$

Note that A_k is irreflexive, since if $e_w \in R_a^{A_k}(e_w)$, then there must be $(\neg p_j, a, p_j) \in U_k$ such that $\{\neg p_j, p_j\} \subseteq e_w$, contradicting the U_k -consistency of e_w . Therefore $(w_1, e_{w_1}) \notin R_a^{C_k[A_k]}((w_1, e_{w_1}))$ and hence $C_k, w_1 \models [A_k, E^{A_k}] \Box_a \neg \text{pre}^{A_k}(e_{w_1})$. So $(C_k[A_k], (w_1, e_{w_1}))$ and $(C_k[A_k], (w_2, e_{w_2}))$ are not $\Box_a$ -equivalent.

Now let (A, E) be a multi-pointed action model having $n(A) < n(A_k)$. We shall prove that (A, E) and U_k are not update equivalent. Proceeding, define the functions $\bar{E} : W^{C_k} \rightarrow \wp(E)$ and $\bar{W} : E \rightarrow \wp(W^{C_k})$ as follows.

$$\begin{aligned} \bar{E}(w) &\stackrel{\text{def}}{=} \{e \in E \mid (w, e) \in W^{C_k[A]}\} \\ \bar{W}(e) &\stackrel{\text{def}}{=} \{w \in W^{C_k} \mid (w, e) \in W^{C_k[A]}\} \end{aligned}$$

If $|\bar{E}(w)| = 0$ for some $w \in W^{C_k}$, then (A, E) and U_k are not update equivalent because $C_k, w \models [A, E] \perp$ and $C_k, w \not\models [U_k] \perp$.

So let us assume that $|\bar{E}(w)| > 0$ for each $w \in W^{C_k}$. If there is a $w \in W^{C_k}$, an $e_1 \in \bar{E}(w)$, an $e_2 \in \bar{E}(w)$, and a reduced formula $\Box_a \varphi$ such that $C_k[A], (w, e_1) \models \Box_a \varphi$ and $C_k[A], (w, e_2) \not\models \Box_a \varphi$, then it follows that $C_k, w \not\models [A, E]\Box_a \varphi$ and $C_k, w \models [U_k]\Box_a \varphi$. But since we have either $C_k, w \models [U_k]\Box_a \varphi$ or $C_k, w \models [U_k]\neg\Box_a \varphi$, it follows that (A, E) and U_k are not update equivalent.

So let us assume that for each $w \in W^{C_k}$, we have not only that $|\bar{E}(w)| > 0$ but also that $(C_k[A], (w, e_1))$ and $(C_k[A], (w, e_2))$ are $\Box_a$ -equivalent for each $e_1 \in \bar{E}(w)$ and each $e_2 \in \bar{E}(w)$. Choosing an event $\tilde{e}_w \in \bar{E}(w)$ for each $w \in W^{C_k}$, it follows from the assumption that $C_k, w \models [A, E]\Box_a \chi$ if and only if $C_k, w \models [A, \{\tilde{e}_w\}]\Box_a \chi$ if and only if $C_k[A], (w, \tilde{e}_w) \models \Box_a \chi$ for each reduced formula $\Box_a \chi$. Since $n(A) < n(A_k) = 2^{k+1} = |W^{C_k}|$, there is an $\hat{e} \in E^A$ such that $|\bar{W}(\hat{e})| \geq 2$; that is, there is a $w_1 \in \bar{W}(\hat{e})$ and a $w_2 \in \bar{W}(\hat{e})$ with $w_1 \neq w_2$. For each $v \in \bar{W}(\hat{e})$ and each $v' \in \bar{W}(\hat{e})$, it follows that $(w, e) \in R_a^{C_k[A]}((v, \hat{e}))$ if and only if $w \in R_a^{C_k}(v)$ and $e \in R_a^A(\hat{e})$ if and only if $e \in R_a^A(\hat{e})$ if and only if $w \in R_a^{C_k}(v')$ and $e \in R_a^A(\hat{e})$ if and only if $(w, e) \in R_a^{C_k[A]}((v', \hat{e}))$. It follows that $(C_k[A], (w_1, \hat{e}))$ and $(C_k[A], (w_2, \hat{e}))$ are $\Box_a$ -equivalent. But we argued above that $(C_k[A_k], (w_1, e_{w_1}))$ and $(C_k[A_k], (w_2, e_{w_2}))$ are not $\Box_a$ -equivalent, and hence there is a reduced formula $\Box_a \chi$ such that, without loss of generality, $C_k[A_k], (w_1, e_{w_1}) \models \Box_a \chi$ and $C_k[A_k], (w_2, e_{w_2}) \not\models \Box_a \chi$.⁵ Further, either $C_k[A], (w, \hat{e}) \models \Box_a \chi$ for each $w \in \{w_1, w_2\}$ or $C_k[A], (w, \hat{e}) \not\models \Box_a \chi$ for each $w \in \{w_1, w_2\}$. If we have that $C_k[A], (w, \hat{e}) \models \Box_a \chi$ for each $w \in \{w_1, w_2\}$, then we have $C_k, w_2 \not\models [A_k, E^{A_k}]\Box_a \chi$ and $C_k, w_2 \models [A, \{\hat{e}\}]\Box_a \chi$, from which it follows by Theorem 4.7 and what was shown above that $C_k, w_2 \models [U_k]\Box_a \chi$ and $C_k, w_2 \models [A, E]\Box_a \chi$ and hence that (A, E) and U_k are not update equivalent. On the other hand, if we have that $C_k[A], (w, \hat{e}) \not\models \Box_a \chi$ for each $w \in \{w_1, w_2\}$, then we have $C_k, w_1 \models [A_k, E^{A_k}]\Box_a \chi$ and $C_k, w_1 \not\models [A, \{\hat{e}\}]\Box_a \chi$, from which it follows by Theorem 4.7 and what was shown above that $C_k, w_1 \models [U_k]\Box_a \chi$ and $C_k, w_1 \not\models [A, E]\Box_a \chi$ and hence that (A, E) and U_k are not update equivalent. $\square$

§7. Common knowledge. Let $\mathcal{F}_0^C$ be the extension of the basic multi-modal language $\mathcal{F}_0$ obtained by adding ordinary common knowledge operators C_G for each group $G \subseteq \mathcal{A}$ of agents. Extending the language $\mathcal{F}_0^C$ by adding public announcement operators such as Plaza's operator $[\varphi]_p$ (Section §5) strictly increases language expressivity (Baltag *et al.*, 2005; Renne, 2008), which makes it impossible to prove completeness via the reduction axiom method of Dynamic Epistemic Logic (Baltag & Moss, 2004; Baltag *et al.*, 1998; Gerbrandy, 2007; Gerbrandy & Groeneveld, 1997; Plaza, 2007; van Benthem *et al.*, 2006; van Ditmarsch *et al.*, 2007). Since AUL can express public announcements (Section §5), adding AUL's arrow update modals $[U]$ to the language $\mathcal{F}_0^C$ also strictly increases language expressivity.

In the case of adding public announcement operators to $\mathcal{F}_0^C$, one way to avoid the expressive increase is to begin with a base language $\mathcal{F}_0^{C+}$ built by extending $\mathcal{F}_0$ with a more expressive common knowledge operators C_G^+ such as the operators of *relativized common knowledge* (Kooi & van Benthem, 2004; van Benthem *et al.*, 2006). We wish to do something similar for the case where we add arrow update modals $[U]$: beginning with a base language $\mathcal{F}_0^*$ obtained by extending $\mathcal{F}_0$ with new common knowledge operators

⁵ The argument for the case where $C_k[A_k], (w_1, e_{w_1}) \not\models \Box_a \chi$ and $C_k[A_k], (w_2, e_{w_2}) \models \Box_a \chi$ is obtained by interchanging w_1 and w_2 in the remainder of the proof.

$\{U\}^*$ —here $\{U\}^*$ is viewed not as a belief change operator but instead as a certain kind of common knowledge operator that we define below—we will see that the further addition of update modals $[U]$ does not increase language expressivity. For brevity, we will approach this study by beginning with the full language $\mathcal{F}^*$ obtained by extending $\mathcal{F}_0$ with both common knowledge operators $\{U\}^*$ and belief-changing update modals $[U]$. After defining a semantics for this language, we will develop a sound and complete axiomatics for the validities of the language with respect to the semantics. On the way to proving completeness, we will prove (in Theorem 7.8) that the full language $\mathcal{F}^*$ and the fragment $\mathcal{F}_0^*$ obtained by omitting belief-changing update modals $[U]$ are equally expressive.

DEFINITION 7.1 (AUL* Language). *The grammar $\mathcal{G}^*$ is defined as follows.*

$$\begin{aligned}\varphi &::= \perp \mid \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box_a\varphi \mid [U]\varphi \mid \{U\}^*\varphi \\ U &::= (\varphi, a, \varphi) \mid (\varphi, a, \varphi), U \\ p &\in \mathcal{P}, a \in \mathcal{A}\end{aligned}$$

Expressions built using φ as a start symbol in grammar $\mathcal{G}^*$ are called **AUL***-formulas; we let $\mathcal{F}^*$ denote the set of **AUL***-formulas. Expressions built using U as a start symbol in grammar $\mathcal{G}^*$ are called **AUL***-arrow updates; we let $\mathcal{U}^*$ denote the set of **AUL***-arrow updates. We adopt conventions and terminology similar to those established for the **AUL** language (Definition 2.2).

DEFINITION 7.2 (AUL* Semantics). *We extend the **AUL** semantics from Definition 2.4 as follows.*

$$\begin{aligned}M, w \models \{U\}^*\varphi &\text{ iff } M, v \models \varphi \text{ for each } v \in (R_U^M)^*(w) \\ R_U^M(v) &\stackrel{\text{def}}{=} \{v' \in W^M \mid \exists(\varphi, a, \varphi') \in U : \\ &\quad (M, v \models \varphi, v' \in R_a^M(v), \text{ and } M, v' \models \varphi')\}\end{aligned}$$

Here $(R_U^M)^*$ denotes the reflexive–transitive closure case of R_U^M . $\mathcal{F}^*$ -validity is defined as is $\mathcal{F}$ -validity (Definition 2.4). Given an **AUL***-arrow update $U \in \mathcal{U}^*$ and a Kripke model M , a U -path in M is a finite nonempty sequence $\langle w_i \rangle_{i=0}^n$ of worlds in M satisfying the property that $w_{i+1} \in R_U^M(w_i)$ for each $i \in \mathbb{N}$ with $i < n$. We say that the U -path $\langle w_i \rangle_{i=0}^n$ in M begins at w_0 , ends at w_n , and is from w_0 to w_n . A step in a U -path $\langle w_i \rangle_{i=0}^n$ is a subsequence $\langle w_j, w_{j+1} \rangle$ obtained by choosing some $j \in \mathbb{N}$ satisfying $0 \leq j < n$.

The semantics of ordinary common knowledge C_G among a nonempty group $G \subseteq \mathcal{A}$ is given by defining $M, w \models C_G\varphi$ to mean that $M, v \models \varphi$ for each $v \in (R_a^M)^*(w)$ (Fagin *et al.*, 1995). Since we have that

$$\models C_G\varphi \leftrightarrow \{(\top, a, \top) \mid a \in G\}^*\varphi,$$

it follows that ordinary common knowledge is expressible in the language of **AUL***.

Let M be a Kripke model. Given a world $w \in W^M$, a formula φ , and a nonempty group $G \subseteq \mathcal{A}$, a φ - G -path from w is a finite nonempty sequence $\langle w_i \rangle_{i=0}^n$ of worlds in M such that $w_0 \stackrel{\text{def}}{=} w$, that $n \geq 1$, that $w_{i+1} \in \bigcup_{a \in G} R_a^M(w_i)$ for each $i < n$, and that $M, w_i \models \varphi$ for each $i \leq n$. A φ - G -path $\langle w_i \rangle_{i=0}^n$ is said to *end on* w_n . The semantics of relativized common knowledge $C_G(\varphi, \psi)$ among a nonempty group $G \subseteq \mathcal{A}$ is given by defining $M, w \models C_G(\varphi, \psi)$ to mean that $M, v \models \psi$ for each $v \in W^M$ such that there is a φ - G -path

from w that ends at v (Kooi & van Benthem, 2004; van Benthem *et al.*, 2006). Since we have that

$$\models C_G(\varphi, \psi) \leftrightarrow \bigwedge_{a \in G} \Box_a(\varphi \rightarrow \{(\top, a, \varphi) \mid a \in G\}^* \psi) ,$$

it follows that relativized common knowledge is expressible in the language of $\mathbf{AUL}^*$.

DEFINITION 7.3 ($\{U\}\psi$, $U(\psi/\psi', a, \chi/\chi')$). *Let U be an $\mathbf{AUL}^*$ -arrow update and $a \in \mathcal{A}$ be an agent. Let ψ , ψ' , χ , and χ' be $\mathbf{AUL}^*$ -formulas. We define the expressions $\{U\}\psi$ and $U(\psi/\psi', a, \chi/\chi')$ as follows.*

- $\{U\}\psi$ is the $\mathbf{AUL}^*$ -formula defined by setting

$$\{U\}\psi \stackrel{\text{def}}{=} \bigwedge_{(\varphi, b, \varphi') \in U} (\varphi \rightarrow \Box_b(\varphi' \rightarrow \psi)) .$$

- $U(\psi/\psi', a, \chi/\chi')$ is the $\mathbf{AUL}^*$ -arrow update defined by setting

$$U(\psi/\psi', a, \chi/\chi') \stackrel{\text{def}}{=} (U - \{(\psi, a, \chi)\}) \cup \{(\psi', a, \chi')\} .$$

DEFINITION 7.4 ($\mathbf{AUL}^*$ Theory). *The axiomatic theory $\mathbf{AUL}^*$ is defined in Table 3. We write $\mathbf{AUL}^* \vdash \varphi$ (or sometimes $\vdash^* \varphi$) to mean that the $\mathbf{AUL}^*$ -formula φ is derivable in the axiomatic theory $\mathbf{AUL}^*$; the negation of $\mathbf{AUL}^* \vdash \varphi$ is written $\mathbf{AUL}^* \not\vdash \varphi$ (or sometimes $\not\vdash^* \varphi$).*

LEMMA 7.5. *Let M be a Kripke model and $\langle w_i \rangle_{i=0}^n$ be a finite nonempty sequence of worlds in M . Then $\langle w_i \rangle_{i=0}^n$ is a U' -path in $M * U$ if and only if $\langle w_i \rangle_{i=0}^n$ is a $(U \circ U')$ -path in M .*

Proof. Assume that $\langle w_i \rangle_{i=0}^n$ is a U' -path in $M * U$. This means that for each step $\langle w_i, w_{i+1} \rangle$ in $\langle w_i \rangle_{i=0}^n$, there exists a $(\varphi_i, a_i, \varphi'_i) \in U'$ such that $M * U, w_i \models \varphi_i$ (equivalently: $M, w_i \models [U]\varphi_i$), that $w_{i+1} \in R_{a_i}^{M * U}(w_i)$, and that $M * U, w_{i+1} \models \varphi'_i$ (equivalently: $M, w_{i+1} \models [U]\varphi'_i$). But to have $w_{i+1} \in R_{a_i}^{M * U}(w_i)$ means that there exists a $(\psi_i, a_i, \psi'_i) \in$

AXIOM SCHEMES

AUL. Arrow Update Logic **AUL**

$$\text{U6. } [U]\{U'\}^* \varphi \leftrightarrow \{U \circ U'\}^* [U]\varphi$$

$$\text{UK. } \{U\}^*(\varphi \rightarrow \psi) \rightarrow (\{U\}^* \varphi \rightarrow \{U\}^* \psi)$$

$$\text{MIX. } \{U\}^* \varphi \leftrightarrow (\varphi \wedge \{U\}\{U\}^* \varphi)$$

$$\text{IND. } \{U\}^*(\varphi \rightarrow \{U\}\varphi) \rightarrow (\varphi \rightarrow \{U\}^* \varphi)$$

RULES

$$\frac{\varphi \rightarrow \psi \quad \varphi}{\psi} \text{ (MP)} \quad \frac{\varphi}{\Box_a \varphi} \text{ (BN)} \quad \frac{\varphi}{[U]\varphi} \text{ (UN)}$$

$$\frac{\varphi}{\{U\}^* \varphi} \text{ (*N)} \quad \frac{\psi \rightarrow \psi' \quad \chi \rightarrow \chi'}{\{U(\psi/\psi', a, \chi/\chi')\}^* \varphi \rightarrow \{U\}^* \varphi} \text{ (*W)}$$

Table 3. The theory $\mathbf{AUL}^*$

U such that $M, w_i \models \psi_i$, that $w_{i+1} \in R_{a_i}^M(w_i)$, and that $M, w_{i+1} \models \psi'_i$. Hence for each step $\langle w_i, w_{i+1} \rangle$ in $\langle w_i \rangle_{i=0}^n$, there exists a $(\psi_i, a_i, \psi'_i) \in U$ and a $(\phi_i, a_i, \phi'_i) \in U$ such that $M, w_i \models \psi_i \wedge [U]\phi_i$, that $w_{i+1} \in R_{a_i}^M(w_i)$, and that $M, w_{i+1} \models \psi'_i \wedge [U]\phi'_i$. But by the definition of $U \circ U'$ (Definition 3.1), this is what it means to say that $\langle w_i \rangle_{i=0}^n$ is a $(U \circ U')$ -path in M . $\square$

THEOREM 7.6 (AUL* Soundness). $\text{AUL}^* \vdash \varphi$ implies $\models \varphi$ for each $\varphi \in \mathcal{F}^*$.

Proof. By induction on the length of derivation in AUL^* . We restrict our attention to the base case for Axiom U6 and the induction case for Rule *W. The other cases are straightforward adaptations of the arguments for ordinary common knowledge (Fagin *et al.*, 1995).

Proceeding, we prove that Axiom U6 is valid. For the left-to-right direction, assume that $M, w \models [U]\{U'\}^*\varphi$ and that $\langle w_i \rangle_{i=0}^n$ is a $(U \circ U')$ -path in M that begins at w . It follows by Lemma 7.5 that $\langle w_i \rangle_{i=0}^n$ is a U' -path in $M * U$ that begins at w and therefore that $M * U, w_n \models \varphi$. Hence $M, w_n \models [U]\varphi$ for each $(U \circ U')$ -path $\langle w_i \rangle_{i=0}^n$ in M that begins at w . But this is what it means to have $M, w \models \{U \circ U'\}^*[U]\varphi$. For the right-to-left direction, assume that $M, w \models \{U \circ U'\}^*[U]\varphi$ and that $\langle w_i \rangle_{i=0}^n$ is a U' -path in $M * U$ that begins at w . It follows by Lemma 7.5 that $\langle w_i \rangle_{i=0}^n$ is a $(U \circ U')$ -path in M that begins at w and therefore that $M, w \models [U]\varphi$. Hence $M * U, w_n \models \varphi$ for each U' -path $\langle w_i \rangle_{i=0}^n$ in $M * U$ that begins at w . But this is what it means to have $M, w \models [U]\{U'\}^*\varphi$. This completes the proof that Axiom U6 is valid.

Let us now show that Rule *W is validity preserving. Proceeding, assume that $\models \psi \rightarrow \psi'$ and $\models \chi \rightarrow \chi'$. We wish to argue that $\models \{U(\psi/\psi', a, \chi/\chi')\}^*\varphi \rightarrow \{U\}^*\varphi$. So let (M, w) be an arbitrary pointed Kripke model satisfying the property that $M, w \models \{U(\psi/\psi', a, \chi/\chi')\}^*\varphi$. It follows that for each $U(\psi/\psi', a, \chi/\chi')$ -path $\langle w_i \rangle_{i=0}^n$ in M that begins at w , we have that $M, w_n \models \varphi$. But since $\models \psi \rightarrow \psi'$ and $\models \chi \rightarrow \chi'$, it follows that every U -path in M is also a $U(\psi/\psi', a, \chi/\chi')$ -path in M . Hence $M, w \models \{U\}^*\varphi$. Since M was chosen arbitrarily, it follows that Rule *W is validity preserving. $\square$

DEFINITION 7.7 (AUL* Complexity). We define $c : (\mathcal{F}^* \cup \mathcal{U}^*) \rightarrow \mathbb{N}$ by adding the following to the equations in Definition 3.4.

$$c(\{U\}^*\varphi) \stackrel{\text{def}}{=} 2 + c(U) + c(\varphi)$$

For each $X \in \mathcal{F}^* \cup \mathcal{U}^*$, the number $c(X)$ is called the complexity of X .

THEOREM 7.8 (AUL* Reduction). For each $\varphi \in \mathcal{F}^*$, there is a reduced $\varphi^\circ \in \mathcal{F}^*$ such that $\text{AUL}^* \vdash \varphi \leftrightarrow \varphi^\circ$.

Proof. We add the following equations to those in Table 2.

$$\begin{aligned} r(\{U\}^*\varphi) &\stackrel{\text{def}}{=} \{r(U)\}^*r(\varphi) \\ r([U]\{U'\}^*\varphi) &\stackrel{\text{def}}{=} \{r(U \circ U')\}^*r([U]\varphi) \\ r((\varphi, a, \varphi')) &\stackrel{\text{def}}{=} (r(\varphi), a, r(\varphi')) \\ r((\varphi, a, \varphi'), U) &\stackrel{\text{def}}{=} (r(\varphi), a, r(\varphi')), r(U) \end{aligned}$$

As in the proof of AUL Reduction (Theorem 3.5), we then argue that the full collection of equations defines a function $r : \mathcal{F}^* \rightarrow \mathcal{F}^*$. To see that these equations are complexity

respecting, we note that we have $c(U \circ U') \leq 2 + c(U) \cdot c(U') + 2 \cdot c(U')$ by what was shown in the proof of AUL Reduction (Theorem 3.5). Further, since we have that $c([U]\varphi) = c(U) \cdot c(\varphi) + 2 \cdot c(\varphi)$ and that $c([U]\{U'\}^*\varphi) = c(U) \cdot 2 + c(U) \cdot c(U') + c(U) \cdot c(\varphi) + 4 + 2 \cdot c(U') + 2 \cdot c(\varphi)$, it follows that $c([U]\{U'\}^*\varphi) > c(U \circ U')$ and $c([U]\{U'\}^*\varphi) > c([U]\varphi)$. The equations are therefore complexity respecting. So to prove the statement of the theorem, one then argues by induction on the complexity k of AUL*-formulas that for each formula θ having $c(\theta) \leq k$, it follows that $\vdash^* \theta \leftrightarrow r(\theta)$. The argument is broken up into a number of cases, one for each of the forms of an AUL*-formula to which r is applied on the left side of an equation from either Table 2 or from the above-listed additions to Table 2. Let us check the induction cases for AUL*-formulas of the form $\{U\}^*\varphi$ and for AUL*-formulas of the form $[U]\{U'\}^*\varphi$. The other cases are straightforward.

Proceeding, we prove that $\vdash^* \{U\}^*\varphi \leftrightarrow r(\{U\}^*\varphi)$. First, we have $\vdash^* \varphi \leftrightarrow r(\varphi)$ by the induction hypothesis (which may be applied by the fact that $r : \mathcal{F}^* \rightarrow \mathcal{F}^*$ is complexity respecting). Applying modal reasoning, it follows that $\vdash^* \{U\}^*\varphi \leftrightarrow \{U\}^*r(\varphi)$. But since we have that $\vdash^* \{U\}^*r(\varphi) \leftrightarrow \{r(U)\}^*r(\varphi)$ by repeated use of Rule *W and the induction hypothesis, the result follows.

Finally, we prove that $\vdash^* [U]\{U'\}^*\varphi \leftrightarrow r([U]\{U'\}^*\varphi)$. First, we have that $\vdash^* [U]\{U'\}^*\varphi \leftrightarrow \{U \circ U'\}^*[U]\varphi$ by Axiom U6. Second, it follows by the induction hypothesis (which is applicable because $r : \mathcal{F}^* \rightarrow \mathcal{F}^*$ is complexity respecting) that $\vdash^* [U]\varphi \leftrightarrow r([U]\varphi)$ and hence that $\vdash^* \{U \circ U'\}^*[U]\varphi \leftrightarrow \{U \circ U'\}^*r([U]\varphi)$ by modal reasoning. Hence we have shown that $\vdash^* [U]\{U'\}^*\varphi \leftrightarrow \{U \circ U'\}^*r([U]\varphi)$. But since we have that $\vdash^* \{U \circ U'\}^*r([U]\varphi) \leftrightarrow \{r(U \circ U')\}^*r([U]\varphi)$ by repeated use of Rule *W and the induction hypothesis, the result follows. $\square$

THEOREM 7.9 (AUL* Completeness). $\models \varphi$ implies $\text{AUL}^* \vdash \varphi$ for each $\varphi \in \mathcal{F}^*$.

Proof. As in the proof of AUL Completeness (Theorem 3.6), though making use of AUL* Soundness (Theorem 7.6) and AUL* Reduction (Theorem 7.8). Note that completeness of the reduced fragment of AUL* follows by a standard Fischer–Ladner-style finite model construction (van Ditmarsch *et al.*, 2007, sec. 7.3).⁶ $\square$

§8. Conclusion. We presented the theory AUL of *Arrow Update Logic*, a multi-agent doxastic theory for reasoning about belief changes brought about by epistemic arrow-eliminating operations called *arrow updates*. We provided a sound and complete axiomatization of AUL along with two examples of AUL-describable belief changes: public announcements (both Plaza’s and Gerbrandy and Groeneveld’s) and cautious updating. We proved that every belief-changing operation expressible using an arrow update U is expressible using the action model $A[U]$ of Dynamic Epistemic Logic. However, using action models can be expensive: we proved that arrow updates are sometimes exponentially more succinct than action models in expressing belief changes. These results give us half of the story about the relationship between arrow updates and action models: arrow updates can be expressed using action models but it is sometimes exponentially more

⁶ The *reduced fragment* of AUL* is the theory AUL_0^* obtained from AUL* by eliminating all axioms and rules that contain nonreduced formulas. The language of AUL_0^* consists of the set of reduced AUL*-formulas.

expensive to do so. What remains for future work is to determine the other half of the story: which action models are expressible using arrow updates, and are action models sometimes exponentially more succinct than arrow updates?

Toward the end of this paper, we presented the theory AUL^* of AUL with common knowledge operators $\{U\}^*$. We proved that the language $\mathcal{F}_0^*$ obtained by adding these common knowledge operators to the basic multi-modal language $\mathcal{F}_0$ is just as expressive as the language $\mathcal{F}^*$ obtained by further extending $\mathcal{F}_0^*$ by adding arrow update modals $[U]$ (Theorem 7.8). While we showed that our common knowledge operators $\{U\}^*$ can express both ordinary and relativized common knowledge, it is an open question whether this expressive relationship holds the other way around (though we suspect that in each case it does not).

While the action models of DEL generalize Plaza's world-eliminating public announcements, the arrow updates of AUL generalize Gerbrandy and Groeneveld's arrow-eliminating public announcements, though the AUL generalization comes with an assumption that the way each of the agents responds to incoming information is common knowledge among the agents. In Kooi & Renne (2011), we study the theory of *Generalized Arrow Update Logic* (GAUL), which drops this assumption. We achieve this by expanding the arrow specifications (φ, a, φ') of AUL in a way that allows us to say that an a -arrow elimination happens privately for some subgroup of agents. The resulting *generalized arrow updates* of GAUL therefore have a more direct connection with Dynamic Epistemic Logic's action models: each generalized arrow update is update equivalent to an action model and each action model is update equivalent to a generalized arrow update. But mutual expressivity does not come for free. In particular, generalized arrow updates are at worst polyexponentially less succinct than action models, though this improves to being at worst less succinct if the action models have purely epistemic preconditions (i.e., preconditions not containing action models) or if we allow generalized arrow updates to have target conditions in the language of Dynamic Epistemic Logic. Further, we showed that generalized arrow updates are sometimes exponentially more succinct than action models. See Kooi & Renne (2011) for details. Though GAUL allows us to drop the common knowledge assumption present in AUL, we still have not identified the exact update-expressive relationship between action models and AUL arrow updates, an important question for future research.

§9. Acknowledgments. The authors thank Andreas Herzig, Ben Rodenhäuser, Hans van Ditmarsch, Guillaume Aucher, and an anonymous referee for their helpful feedback on the text.

§10. Dedication. The authors dedicate this work to the memory of Horacio Arl6-Costa, who the authors were honored to have as their editor.

BIBLIOGRAPHY

- Aucher, G., Balbiani, P., Fariñas del Cerro, L., & Herzig, A. (2009). Global and local graph modifiers. *Electronic Notes in Theoretical Computer Science*, **231**, 293–307.
- Baltag, A., & Moss, L. S. (2004). Logics for epistemic programs. *Synthese*, **139**(2), 165–224.
- Baltag, A., Moss, L. S., & Solecki, S. (1998). The logic of public announcements, common knowledge, and private suspicions. In Gilboa, I., editor. *Proceedings of the 7th*

- Conference on Theoretical Aspects of Rationality and Knowledge (TARK VII)*. Evanston, IL, pp. 43–56.
- Baltag, A., Moss, L. S., & Solecki, S. (2005). Logics for epistemic actions: Completeness, decidability, expressivity. Manuscript.
- Baltag, A., & Smets, S. (2007). A qualitative theory of dynamic interactive belief revision. In Bonanno, G., van der Hoek, W., and Wooldridge, M., editors. *Selected Papers from LOFT'06, Texts in Logic and Games*. Amsterdam: Amsterdam University Press.
- Blackburn, P., de Rijke, M., & Venema, Y. (2001). *Modal Logic*. Cambridge University Press.
- Clark, M. (2007). *Paradoxes from A to Z* (second edition). London: Routledge.
- Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. Y. (1995). *Reasoning about Knowledge*. Cambridge: The MIT Press.
- Gerbrandy, J. (1999). Bisimulations on planet Kripke. PhD Thesis, University of Amsterdam. ILLC Dissertation Series DS-1999-01.
- Gerbrandy, J. (2007). The surprise examination in dynamic epistemic logic. *Synthese*, **155**(1), 21–33.
- Gerbrandy, J., & Groeneveld, W. (1997). Reasoning about information change. *Journal of Logic, Language and Information*, **6**(2), 147–169.
- Kooi, B., & Renne, B. (2011). Generalized arrow update logic. In Apt, K. R., editor. *Proceedings of the 13th Conference of Theoretical Aspects of Rationality and Knowledge (TARK XIII)*. Groningen, The Netherlands: ACM, pp. 205–211.
- Kooi, B., & van Benthem, J. (2004). Reduction axioms for epistemic actions. In Schmidt, R., Pratt-Hartmann, I., Reynolds, M., and Wansing, H., editors. *AiML-2004: Advances in Modal Logic (Preliminary Proceedings)*. Manchester, UK, pp. 197–211.
- Lutz, C. (2006). Complexity and succinctness of public announcement logic. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-06)*. Hakodate, Japan: Association for Computing Machinery (ACM), pp. 137–144.
- Plaza, J. (1989). Logics of public communications. In Emrich, M. L., Pfeifer, M. S., Hadzikadic, M., and Ras, Z. W., editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems (ISMIS 1989): Poster Session Program*. Charlotte, NC: Oak Ridge National Laboratory ORNL/DSRD-24, pp. 201–216.
- Plaza, J. (2007). Logics of public communications. *Synthese*, **158**(2), 165–179.
- Renardel de Lavalette, G. R. (2004). Changing modalities. *Journal of Logic and Computation*, **14**(2), 215–249.
- Renne, B. (2008). Public and private communication are different: Results on relative expressivity. *Synthese*, **165**(2), 225–245.
- Renne, B., Sack, J., & Yap, A. (2009). Dynamic epistemic temporal logic. In He, X., Horty, J., & Pacuit, E., editors. *Logic, Rationality, and Interaction, Second International Workshop, LORI 2009, Chongqing, China, October 8–11, 2009, Proceedings, Volume 5834/2009 of Lecture Notes in Computer Science*. Berlin/Heidelberg, Germany: Springer, pp. 263–277.
- Renne, B., Sack, J., & Yap, A. (2010). Dynamic epistemic temporal logic. Extended Manuscript.
- Steiner, D. (2006). A system for consistency preserving belief change. In Artemov, S., and Parikh, R., editors. *Proceedings of the Workshop on Rationality and Knowledge, 18th European Summer School in Logic, Language, and Information (ESSLLI)*. Málaga, Spain, pp. 133–144.

- van Benthem, J. (2004). Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, **14**(2), 129–155.
- van Benthem, J. (2005). An essay on sabotage and obstruction. In Hutter, D., editor. *Mechanizing Mathematical Reasoning, Essays in Honor of Jörg Siekmann on the Occasion of his 69th Birthday*. Cambridge, Massachusetts: Springer. pp. 268–276.
- van Benthem, J., van Eijck, J., & Kooi, B. (2006). Logics of communication and change. *Information and Computation*, **204**(11), 1620–1662.
- van Ditmarsch, H. (2000). Knowledge games. PhD Thesis, University of Groningen. ILLC Dissertation Series DS-2000-06.
- van Ditmarsch, H., van der Hoek, W., & Kooi, B. (2007). *Dynamic Epistemic Logic*, Vol. 337 of *Synthese Library*. Dordrecht, The Netherlands: Springer.
- van Ditmarsch, H., van Eijck, J., Sietsma, F., & Wang, Y. (2010). On the logic of lying. Manuscript.
- van Eijck, J., Ruan, J., & Sadzik, T. (2008). Action emulation. Manuscript.

FACULTY OF PHILOSOPHY
UNIVERSITY OF GRONINGEN
OUDE BOTERINGESTRAAT 52
9712 GL GRONINGEN
THE NETHERLANDS
E-mail: b.p.Kooi@rug.nl, bryan@renne.org