University of Groningen

The expressivity of update logics

Kuijjer, Louwe B.

Published in:
Journal of Logic and Computation

DOI:
[10.1093/logcom/exu047](https://doi.org/10.1093/logcom/exu047)

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
Publisher's PDF, also known as Version of record

Publication date:
2015

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

Citation for published version (APA):
Kuijjer, L. B. (2015). The expressivity of update logics. *Journal of Logic and Computation*, 25(3), 719-742.
<https://doi.org/10.1093/logcom/exu047>

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.

The expressivity of update logics

LOUWE B. KUIJER, *University of Groningen, Faculty of Philosophy, Oude Boteringestraat 52, 9712 GL Groningen, The Netherlands.*
E-mail: l.b.kuijer@rug.nl

Abstract

We prove two new results about logics involving updates and common knowledge. The first result is that the logic $\mathcal{L}_{AU^*}$ using Arrow Common Knowledge is more expressive than the logic $\mathcal{L}_{AR}$ using Relativized Common Knowledge. The second result is that the logic $\mathcal{L}_{AUC}$ using Arrow Updates and normal Common Knowledge is equally expressive as $\mathcal{L}_{AU^*}$.

Together with previously known results this fully determines the expressivity landscape of all logics involving any combination of normal Common Knowledge (C), Relativized Common Knowledge (R), Arrow Common Knowledge (U^*), Public Announcements (P) and Arrow Updates (U).

Keywords: Expressivity, public announcements, arrow updates, relativized common knowledge, arrow common knowledge.

1 Introduction

In this article, we consider all logics that can be obtained by adding a combination of common knowledge, relativized common knowledge [12], arrow common knowledge [11], public announcements [4, 13] and arrow updates [11] to a basic modal logic.

Generally, we can use only one logic at a time. So if we have multiple logics we have to choose between them. As such it becomes interesting to compare them to each other. Usually every logic has its own strengths and weaknesses so we cannot conclude that one logic is simply better than another. We can however, sometimes say that one logic is better than another in one particular aspect, so with respect to some specific criterion.

There are several such criteria that can be used to compare logics. We could e.g. look at the computational complexity for one of the decision problems associated with the logic, or at the succinctness of the logic. Here we want to compare logics by another criterion, namely that of their *expressivity* (or *expressive power*). As the word suggests the expressivity of a logic is a measure of what can be expressed in the logic. So if a logic $\mathcal{L}_2$ is at least as expressive as a logic $\mathcal{L}_1$ then everything that can be expressed in $\mathcal{L}_1$ can also be expressed in $\mathcal{L}_2$.

What this means is that for every $\mathcal{L}_1$ formula φ_1 there is a $\mathcal{L}_2$ formula φ_2 with the same meaning; so for every $\mathcal{L}_1$ formula there is an equivalent $\mathcal{L}_2$ formula. The formulas of $\mathcal{L}_1$ can then be seen as abbreviations for the formulas of $\mathcal{L}_2$, so everything that can be done using $\mathcal{L}_1$ can also be done using $\mathcal{L}_2$. If we only look at what a logic can do (and not at how efficiently it does so) there is then no reason to use $\mathcal{L}_1$, since $\mathcal{L}_2$ does at least as well in every situation.

Let us consider a well-known example. Let $\mathcal{L}_{\text{prop}}$ be propositional logic with the connectives $\neg, \vee, \wedge, \rightarrow$ and $\leftrightarrow$, and let $\mathcal{L}_{\{\neg, \vee\}}$ be propositional logic with only the connectives $\neg$ and $\vee$. It is well known that for every $\mathcal{L}_{\text{prop}}$ formula there is an equivalent $\mathcal{L}_{\{\neg, \vee\}}$ formula, so $\mathcal{L}_{\{\neg, \vee\}}$ is at least as expressive as $\mathcal{L}_{\text{prop}}$.

Here, we chart the expressivity landscape of all logics under consideration, so all 32 combinations of common knowledge, relativized common knowledge, arrow common knowledge, public announcements and arrow updates. The approach we take is very similar to the one in [10], where the expressivity landscape of a different (but partially overlapping) set of logics is charted.

Vol. 25 No. 3, © The Author, 2014. Published by Oxford University Press. All rights reserved.

For Permissions, please email: journals.permissions@oup.com

Published online September 1, 2014. doi:10.1093/logcom/exu047

For many of the logics the relative expressivity is already known. There are however, two important new expressivity results introduced in this article, as well as a number of results that follow from these two results. The first is that the logic using relativized common knowledge is not as expressive as the logic using arrow common knowledge. The second is that the logic using arrow updates and normal common knowledge is equally expressive as the logic using arrow common knowledge.

1.1 Overview

In Section 2, we briefly introduce and informally discuss some properties and applications of the different operators. Then, in Section 3, we give a number of definitions that are required to compare the expressivity of the logics under consideration. In Section 4, the expressivity landscape is shown and an overview is given of both the previously known results and the new results. In Section 5, a proof is given of the first new result, that the logic using relativized common knowledge is not as expressive as the logic using arrow common knowledge. In Section 6, a proof is given of the second new result, that the logic using arrow updates and normal common knowledge is as expressive as the logic using arrow common knowledge.

2 Introducing the Operators

Multi-agent Kripke models can be used to model the information states of agents. One important property of information states is that they can change. A common way to see information change is to consider it as *changes made to the Kripke model*.¹

One important and very general approach to information change is to use *action models* (see e.g. [2–4, 8]). In a logic using action models every action $[\alpha]$ is associated with an action model M_α and performing $[\alpha]$ in a model $\mathcal{M}$ changes the model to a certain submodel of the product model $\mathcal{M} \times M_\alpha$. A notable consequence of this is that applying an action may increase the size of your model. Another very general approach to information (and factual) change is to use *Global Graph Modifiers*, see [1]. Global Graph Modifiers allow one to add worlds, add or remove arrows and change the valuations of propositional variables. Because worlds and arrows can be added the use of global graph modifiers can also increase the size of your model.

Here however we focus on a particular kind of information change, where only new information is acquired (and nothing forgotten or proven false) and the new information is made publicly available. This restricted kind of information change can be described using the general Action Models or Global Graph Modifiers, but there are simpler options. New public information can only remove access to alternatives that were previously considered possible, it can never add new alternatives. This allows us to restrict ourselves to model changing operators that go from a model to one of its submodels.

A model consists of a set of possible worlds, accessibility relations between the possible worlds and the valuation of the propositional variables on the worlds. We are modelling information change, not factual change, so the values of the propositional variables should remain unchanged. Since we want to go from a model to one of its submodels this leaves us with the choice to let the information change operator remove either worlds or accessibility arrows.

¹Another approach is to consider information change as a state transition inside a larger model. The two approaches are not fundamentally different; a change from model $\mathcal{M}_1$ to model $\mathcal{M}_2$ can be seen as a state transition in a larger model containing both $\mathcal{M}_1$ and $\mathcal{M}_2$. See [5] for a discussion of dynamic epistemic logic with the dynamic operations seen as state transitions.

The most commonly used choice is to remove certain possible worlds using *public announcements* (see e.g. [4, 8, 13]). A public announcement $[\varphi]$ removes all worlds where the formula φ does not hold from the model. A less commonly used alternative is to remove certain accessibility arrows. A very simple version of this is the variation on public announcements in [9, 10], where $[\varphi]$ does not remove the worlds where φ does not hold but merely the arrows to such worlds. Removing arrows in this way has the same result as removing the $\neg\varphi$ worlds, an inaccessible world might as well not exist.

A more powerful way to remove arrows is to use *arrow updates*, see [11]. An arrow update $[U]$ consists of a number of clauses, $U = \{(u_1, a_1, u'_1), \dots, (u_n, a_n, u'_n)\}$, where we do not require that $a_i \neq a_j$ whenever $i \neq j$. An arrow satisfies a clause (u_i, a_i, u'_i) iff it is an arrow for agent a_i and it goes from a world that satisfies the start condition u_i and to a world that satisfies the end condition u'_i .² The update removes those arrows that satisfy none of the clauses.³

Another operator that is often used in logics about information is the common knowledge operator C_B , where B is a group of agents. The formula $C_B\varphi$ holds in a world w iff φ holds in all worlds w' that are reachable from w by a ' B -path' (i.e. a sequence of arrows belonging to agents in B that connect w to a successor w_1 of w , w_1 to a successor w_2 of w_1 and so up to an arrow that connects w_n to a successor w' of w_n).

For both public announcements and arrow updates there is an associated variant of common knowledge. The common knowledge variant for public announcements is *relativized common knowledge*, defined in [12]. The formula $C_B(\varphi_1, \varphi_2)$ stands for φ_2 being common knowledge relative to φ_1 . It holds in a world w iff φ_2 holds in all worlds w' that are reachable from w by a B -path *that only consists of φ_1 worlds*.

The corresponding variant of common knowledge for arrow updates is *arrow common knowledge*, defined in [11]. The formula $\{U\}^*\varphi$ stands for φ being common knowledge relative to the arrow update U . It holds in a world w iff φ holds in all worlds w' that are reachable from w by a path *that only consists of arrows that satisfy U* .⁴

The different building blocks discussed so far can be combined in different ways. We could e.g. define a logic $\mathcal{L}_{\text{APU}^*}$ that uses basic modal logic together with public announcements and arrow common knowledge, or a logic $\mathcal{L}_{\text{AR}}$ that uses basic modal logic together with relativized common knowledge.

2.1 Dynamic and static operators

The operators under consideration here can be split into two different kinds: dynamic operators and static operators. The difference between the two kinds is that dynamic operators change the model when they are interpreted whereas, static operators do not.

²We use the slightly awkward terms 'start condition' and 'end condition', as opposed to 'precondition' and 'postcondition', to prevent confusion with the preconditions in action models.

³The operator $[U]$ has some similarities to the operator $[a_1 - (\varphi_1, \psi_1); \dots; a_n - (\varphi_n, \psi_n)]$ from [1], but there are also two differences. The first difference is that $[a_1 - (\varphi_1, \psi_1); \dots; a_n - (\varphi_n, \psi_n)]$ specifies the arrows that are to be removed, whereas $[U]$ specifies the arrows that are to be retained. The second and more important difference is that the clauses from $[U]$ are considered *simultaneously*, whereas the clauses $a_1 - (\varphi_1, \psi_1); \dots; a_n - (\varphi_n, \psi_n)$ are considered *sequentially*.

⁴The $*$ in $\{U\}^*$ is a Kleene star; the operator $\{U\}^*$ is an iterated version of $\{U\}$, which is defined by $\mathcal{M}, w \models \{U\}\varphi$ iff $\mathcal{M}, w' \models \varphi$ for all worlds that are accessible from w by an arrow satisfying a clause from U . Unfortunately $\{U\}$ is visually too similar to $\{U\}^*$ to be practical, we will write $\Box_U$ instead of $\{U\}$. We should be careful to distinguish $\{U\}^*$ from the operator $[U]^*$ (which is not used in this article) that repeatedly applies an update.

The dynamic operators used here are public announcements and arrow updates, the static operators are the three types of common knowledge as well as the Boolean operators and $\Box_a$ of basic modal logic. Two of the static operators are combinations of common knowledge with a dynamic operator. It is worthwhile to spend a few moments to see what it means for a static operator to be related to a dynamic operator in such a way.

Let us start by considering a logic where we have the dynamic operators, the Boolean operators, the modal $\Box_a$ and a normal common knowledge operator C_B but not the two other common knowledge operators. Suppose that in this logic we use one of the dynamic operators, say a public announcement $[\varphi]$ in a pointed model $\mathcal{M}, w$. This announcement removes all $\neg\varphi$ worlds from $\mathcal{M}$, a process that cannot be undone. This means that in the updated model $\mathcal{M}_{[\varphi]}$ some of the information contained in the model $\mathcal{M}$ is lost. In particular, we generally cannot determine from $\mathcal{M}_{[\varphi]}, w$ whether or not $\Box_a\varphi$ held before the update, so whether or not $\mathcal{M}, w \models \Box_a\varphi$. Arrow updates destroy information in the same way.

But occasionally we want to use something similar to the dynamic modalities but that does not destroy information. This can be done by adding static operators that correspond to the combination of a dynamic operator and another operator. Such a new static operator applies its update, performs its associated operation and then un-applies the update. Or, to put it another way, it temporarily *pretends* to apply a dynamic operator.

The static connective $\Box_\varphi$ is the combination of a $\Box$ operator and a public announcement. It first applies the announcement φ , then takes a step in the updated model with $\Box$ and finally it undoes the update. So we have $\mathcal{M}, w \models \Box_\varphi\psi$ if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are accessible from w in the updated model $\mathcal{M}_{[\varphi]}$. Likewise, $\mathcal{M}, w \models \Box_U\psi$ if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are accessible from w in the updated model $\mathcal{M}_{[U]}$.

The operators $\Box_\varphi$ and $\Box_U$ do not add expressivity, however. We have $\mathcal{M}, w \models \Box_\varphi\psi$ if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are accessible from w in the updated model $\mathcal{M}_{[\varphi]}$, so if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are accessible from w in $\mathcal{M}$ that satisfy φ , so if and only if $\mathcal{M}, w \models \Box(\varphi \rightarrow \psi)$. Formulating a formula equivalent to $\Box_U\psi$ is harder but it can also be done; we have $\mathcal{M}, w \models \Box_U\psi$ if and only if $\mathcal{M}, w \models \bigwedge_{(u_1, a, u_2) \in U} (u_1 \rightarrow \Box(u_2 \rightarrow \psi))$. That the operators $\Box_\varphi$ and $\Box_U$ do not add expressivity means they would not add anything fundamentally new to the logic. This does not mean that they are useless; the operator $\Box_U$ is in fact used in several of the proofs in this article. But there is no need to take them as primitive, they can be seen as abbreviations.

Things get more complicated if we combine the dynamic operators not with $\Box$ but with common knowledge. Earlier we defined relativized common knowledge $C_B(\varphi, \psi)$ as meaning ‘ ψ holds in all worlds that are reachable by a B -path that contains only φ worlds’. Note that this is equivalent to saying that $\mathcal{M}, w \models C_B(\varphi, \psi)$ if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are reachable from w by a B -path in the updated model $\mathcal{M}_{[\varphi]}$. So relativized common knowledge is indeed the static operator corresponding to the combination of common knowledge and a public announcement.

Likewise, we defined $\{U\}^*\psi$ as meaning ‘ ψ holds in all worlds that are reachable by a path that only uses arrows that satisfy U ’. This is equivalent to saying that $\mathcal{M}, w \models \{U\}^*\psi$ if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are reachable from w by a path in the updated model $\mathcal{M}_{[U]}$. So arrow common knowledge is, as the name suggests, the static operator corresponding to the combination of common knowledge and an arrow update.

The relativized common knowledge and arrow common knowledge operators are both rather complicated, and they were introduced mainly for technical reasons. Still, as [6] points out there is an informal reading of relativized common knowledge that, although not simple, can provide some intuition behind the operator. A formula $[\varphi]C_B\psi$, which contains a public announcement

and a normal common knowledge formula, can be read as ‘if ψ is announced then it will become common knowledge (among B) that ψ is the case’. A formula $C_B(\varphi, \psi)$ on the other hand can be read as ‘if φ is announced then it will become common knowledge (among B) that ψ *used to be the case* before the announcement’. Likewise, $[U]C_B\psi$ can be read as ‘if U is announced it will become common knowledge (among B) that ψ is the case’ whereas $\{U\}^*\psi$ can be read as ‘if U is announced it will become common knowledge (among B) that ψ *used to be the case* before the announcement’.

Unlike $\Box_\varphi$ and $\Box_U$ there is no obvious way to express $C_B(\varphi, \psi)$ and $\{U\}^*\psi$ without using one of the new static operators. In fact, in [6] it was shown that the logic relativized common knowledge adds expressivity to a logic with $\wedge, \neg, \Box_a, C_B$ and $[\varphi]$ operators. So relativized common knowledge adds something fundamentally new to such a logic.

Arrow common knowledge is to common knowledge and arrow updates as relativized common knowledge is to common knowledge and public announcements. As such the result in [6] suggested that arrow common knowledge would probably add expressivity to a logic with $\wedge, \neg, \Box_a, C_B$ and $[U]$ operators. Here we prove that, surprisingly, this is not the case; for any formula using $\wedge, \neg, \Box_a, C_B, [U]$ and $\{U\}^*$ there is an equivalent formula using only $\wedge, \neg, \Box_a, C_B$ and $[U]$. This means that $\{U\}^*$ does not add anything fundamentally new to such a logic and that it could in theory be used as an abbreviation. It is not very practical to consider $\{U\}^*$ in this way, however, as the translation from a formula with $\{U\}^*$ to one without is extremely complicated and causes an enormous increase in formula size.

2.2 Public announcements and arrow updates

Public announcements are quite widely used so we assume that the reader has encountered them before. Arrow updates on the other hand are not very commonly used so it seems worthwhile to give a short introduction to arrow updates, and especially the difference between arrow updates and public announcements.

The first thing to note is that everything that can be done using public announcements can also be done using arrow updates. In other words, arrow updates are at least as expressive as public announcements. To see why this is the case consider any public announcement $[\varphi]$. This announcement removes all worlds that do not satisfy φ . With arrow updates we cannot remove any worlds, but we can do something with the same effect: we can remove all arrows to $\neg\varphi$ worlds. A world that is not reachable in any way might as well not exist at all, so this has the same effect as removing all $\neg\varphi$ worlds.

But arrow updates can also be used in situations where public announcements cannot. Let us look at a simple example, loosely based on an example given in [11] (which was in turn based on an example in [7]).

EXAMPLE 1

In a very simple card game there are two players, player a and player b . Both players are dealt a single card, face down. Player a either has the ace of spades (p) or the king of spades ($\neg p$), player b either has the ace of diamonds (q) or the king of diamonds ($\neg q$). At this point neither player knows which card either one of them holds. The situation as described so far can be modeled as shown in Figure 1a.

But then suppose that a (openly) looks at his/her card without showing it to b . This action is public, because a openly looks at his/her card. But it still creates some private information for a , namely which card he/she holds. This private information makes it impossible to model

FIGURE 1. A model $\mathcal{M}$ representing a simple game and a model $\mathcal{M}_{[U]}$ representing the same game after a has looked at his/her card. Reflexive arrows are omitted.

the event using a public announcement. We can however model it quite simply using an arrow update $[U]$.

In every p world agent a learns that p is true, so in those worlds he/she no longer holds $\neg p$ worlds possible. Likewise, in every $\neg p$ world he/she no longer holds p worlds possible. This means that a -arrows should only be retained if they go from a p world to a p world or from a $\neg p$ world to a $\neg p$ world. We can do this by including the clauses (p, a, p) and $(\neg p, a, \neg p)$ to U . Agent b on the other hand learns no new information that would allow him/her to distinguish between worlds that he/she could not previously distinguish, so all b -arrows should be retained. We can do this by including a $(\top, b, \top)$ clause in U .

In the end, this gives us the update $U = (\top, b, \top), (p, a, p), (\neg p, a, \neg p)$. And indeed, if applied to $\mathcal{M}$ this gives us the model $\mathcal{M}_{[U]}$, shown in Figure 1b, which is a faithful representation of the game after a has looked at his/her card.

The most important property of the update in the above example is that the information a learns differs per world. In p worlds a learns that he/she holds the ace, whereas in $\neg p$ worlds he/she learns that he/she holds the king. This world-dependence makes it impossible to fully eliminate either the p worlds or the $\neg p$ worlds so public announcements cannot model the new information. Arrow updates on the other hand can model the new information just fine, by removing some (but not all) arrows between p and $\neg p$ worlds.

3 Definitions

Let us now define the different logics that we want to compare. To compare the expressivity of the different kinds of updates and common knowledges it is convenient to first define a logic $\mathcal{L}_{\mathcal{T}}$ that contains all the logics we consider. We can then compare the logics as fragments of $\mathcal{L}_{\mathcal{T}}$. The advantage of doing this is that it allows us to combine formulas from the different logics. For example, $[U]C_A\varphi \leftrightarrow \{U\}^*[U]\varphi$ is only a well-formed formula if we have one logic that contains all of the connectives $[U]$, C_A , $\leftrightarrow$ and $\{U\}^*$. We do have such a logic, namely $\mathcal{L}_{\mathcal{T}}$.

Let $\mathcal{A}$ be a finite non-empty set of agents and $\mathcal{P}$ a countable set of propositional variables.

DEFINITION 1 (The language of $\mathcal{L}_{\mathcal{T}}$)

The formulas of $\mathcal{L}_{\mathcal{T}}$ are given by

$$\begin{aligned} \varphi &::= p \mid \neg\varphi \mid (\varphi \vee \psi) \mid \Box_a\varphi \mid C_B\varphi \mid C_B(\varphi, \varphi) \mid [\varphi]\varphi \mid [U]\varphi \mid \{U\}^*\varphi \\ U &::= (\varphi, a, \varphi) \mid (\varphi, a, \varphi), U \end{aligned}$$

where $p \in \mathcal{P}$, $B \subseteq \mathcal{A}$ and $a \in \mathcal{A}$. Let $\Phi_{\mathcal{T}}$ be the set of formulas of $\mathcal{L}_{\mathcal{T}}$.

We use $\wedge, \vee, \bigwedge, \rightarrow, \leftrightarrow, \top, \perp$ and $\Diamond_a$ in the usual way as abbreviations, omit parenthesis where this should not cause confusion and write a for $\{a\}$. We also slightly abuse notation by identifying an update $U = (u_1, a_1, u'_1), \dots, (u_k, a_k, u'_k)$ with the set $U = \{(u_1, a_1, u'_1), \dots, (u_k, a_k, u'_k)\}$. Furthermore, if $B \subseteq \mathcal{A}$ we write $\Box_B\varphi$ for $\bigwedge_{a \in B} \Box_a\varphi$ and $(\varphi_1, B, \varphi_2)$ for $\{(\varphi_1, a, \varphi_2) \mid a \in B\}$. Finally, we write $\Box$ for $\Box_{\mathcal{A}}$.

The models for $\mathcal{L}_{\mathcal{T}}$ are the standard Kripke models. It should be noted that although we speak of (common) *knowledge* we do not assume any of the frame conditions usually associated with epistemic logic.

DEFINITION 2 (The models of $\mathcal{L}_{\mathcal{T}}$)

An $\mathcal{L}_{\mathcal{T}}$ model $\mathcal{M}$ is a triple $\mathcal{M} = (W, R, v)$ where W is a set of worlds, $R: \mathcal{A} \rightarrow \wp(W \times W)$ assigns to each agent an accessibility relation on W and $v: \mathcal{P} \rightarrow \wp(W)$ is a valuation that assigns to each propositional variable a subset of the worlds.

We say that w is a world of $\mathcal{M} = (W, R, v)$ iff $w \in W$. We can now define the semantics of $\mathcal{L}_{\mathcal{T}}$.

DEFINITION 3 (The semantics of $\mathcal{L}_{\mathcal{T}}$)

Given an $\mathcal{L}_{\mathcal{T}}$ model $\mathcal{M} = (W, R, v)$, a world w of $\mathcal{M}$ and φ, ψ formulas of $\mathcal{L}_{\mathcal{T}}$ define the *satisfaction relation* $\models$ by

$$\begin{aligned} \mathcal{M}, w &\models p && \text{if } w \in v(p), \\ \mathcal{M}, w &\models \neg\varphi && \text{if } \mathcal{M}, w \not\models \varphi, \\ \mathcal{M}, w &\models \varphi \vee \psi && \text{if } \mathcal{M}, w \models \varphi \text{ or } \mathcal{M}, w \models \psi, \\ \mathcal{M}, w &\models \Box_a\varphi && \text{if } \mathcal{M}, w' \models \varphi \text{ for all } w' \text{ such that } (w, w') \in R(a), \\ \mathcal{M}, w &\models [\psi]\varphi && \text{if } \mathcal{M}, w \models \psi \text{ implies } \mathcal{M}_{[\psi]}, w \models \varphi, \\ \mathcal{M}, w &\models [U]\varphi && \text{if } \mathcal{M}_{[U]}, w \models \varphi, \\ \mathcal{M}, w &\models C_B\varphi && \text{if } \mathcal{M}, w' \models \varphi \text{ for all } w' \text{ such that } (w, w') \in R(B)^*, \\ \mathcal{M}, w &\models C_B(\psi, \varphi) && \text{if } \mathcal{M}, w' \models \varphi \text{ for all } w' \text{ such that } (w, w') \in R_{[\psi]}(B)^*, \\ \mathcal{M}, w &\models \{U\}^*\varphi && \text{if } \mathcal{M}, w' \models \varphi \text{ for all } w' \text{ such that } (w, w') \in R_{[U]}^*(a), \end{aligned}$$

where

- $W_{[\varphi]} = \{w \in W \mid \mathcal{M}, w \models \varphi\}$,
- $R_{[\varphi]}(a) = R(a) \cap (W_{[\varphi]} \times W_{[\varphi]})$ for $a \in \mathcal{A}$,
- $v_{[\varphi]}(p) = v(p) \cap W_{[\varphi]}$ for $p \in \mathcal{P}$,
- $\mathcal{M}_{[\varphi]} = (W_{[\varphi]}, R_{[\varphi]}, v_{[\varphi]})$,
- $R_{[U]}(a) = \{(w_1, w_2) \in R(a) \mid \exists (u, a, u') \in U : \mathcal{M}, w_1 \models u \text{ and } \mathcal{M}, w_2 \models u'\}$ for $a \in \mathcal{A}$,
- $\mathcal{M}_{[U]} = (W, R_{[U]}, v)$,
- $R(B)^*$ is the reflexive transitive closure of $\bigcup_{a \in B} R(a)$,
- $R_{[\varphi]}(B)^*$ is the reflexive transitive closure of $\bigcup_{a \in B} R_{[\varphi]}(a)$,
- $R_{[U]}^*(a)$ is the reflexive transitive closure of $\bigcup_{a \in \mathcal{A}} R_{[U]}(a)$.

We write $\mathcal{M} \models \varphi$ if $\mathcal{M}, w \models \varphi$ for every world w of $\mathcal{M}$ and $\models \varphi$ if $\mathcal{M} \models \varphi$ for every model $\mathcal{M}$.

Most of the semantics are as usual, although there are two things worth pointing out. The first is that the common knowledge operators take the *reflexive* transitive closure of the relevant relation. This is not very unusual, but neither is taking the transitive closure instead. The second is that a public announcement formula $[\psi]\varphi$ is automatically true in every world $\neg\psi$ world. Again, this is not unusual but there are other options. The results presented in this article also hold for the alternative semantics, only some very minor changes in the proofs would be required.

What we are really interested in is not $\mathcal{L}_{\mathcal{T}}$ but certain fragments of it. We define these fragments as in [10].

DEFINITION 4 (Fragments of $\mathcal{L}_{\mathcal{T}}$)

Let

- A, representing ‘agents’, stand for the connectives $\neg, \vee$ and $\Box_a$
- C, representing ‘common knowledge’, stand for the connective C_B
- R, representing ‘relativised common knowledge’, stand for the connective $C_B(\psi, \varphi)$
- P, representing ‘public announcement’, stand for the connective $[\varphi]$
- U, representing ‘arrow updates’, stand for the connective $[U]$
- U*, representing ‘arrow common knowledge’, stand for the connective $\{U\}^*$.

The logic $\mathcal{L}_X$ for a finite string X is the logic $\mathcal{L}_{\mathcal{T}}$ with the language restricted to only those connectives that belong to a letter in X. Let Φ_X be the set of formula of $\mathcal{L}_X$.

So e.g. the logic $\mathcal{L}_{AUC}$ is the logic using the connectives $\neg, \vee, \Box_a, C_B$ and $[U]$. We also sometimes denote the logic $\mathcal{L}_X$ by the string X, so AUC is the logic $\mathcal{L}_{AUC}$. We can easily define the relative expressivity of such fragments.

We write $\models$ for the satisfaction relation of the fragments as well as for the satisfaction relation of $\mathcal{L}_{\mathcal{T}}$. There is no risk of confusion as the different satisfaction relations coincide whenever multiple ones are defined.

DEFINITION 5

Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be fragments of $\mathcal{L}_{\mathcal{T}}$. Then $\mathcal{L}_2$ is *at least as expressive as* $\mathcal{L}_1$, denoted $\mathcal{L}_1 \preceq \mathcal{L}_2$, if for each $\mathcal{L}_1$ formula φ_1 there is an $\mathcal{L}_2$ formula φ_2 such that

$$\models \varphi_1 \leftrightarrow \varphi_2.$$

We say that $\mathcal{L}_2$ is *more expressive than* $\mathcal{L}_1$, denoted $\mathcal{L}_1 \prec \mathcal{L}_2$, if $\mathcal{L}_1 \preceq \mathcal{L}_2$ and $\mathcal{L}_2 \not\preceq \mathcal{L}_1$. We say that $\mathcal{L}_2$ and $\mathcal{L}_1$ are *equally expressive*, denoted $\mathcal{L}_1 \equiv \mathcal{L}_2$, if $\mathcal{L}_1 \preceq \mathcal{L}_2$ and $\mathcal{L}_2 \preceq \mathcal{L}_1$.

We can coherently write $\models \varphi_1 \leftrightarrow \varphi_2$ even though φ_1 and φ_2 are in different logics because both logics are fragments of $\mathcal{L}_{\mathcal{T}}$. Note that the relation $\preceq$ is reflexive and transitive. The relation $\equiv$ inherits the reflexivity and transitivity of $\preceq$ and is also symmetric, so it is an equivalence relation.

It will also be useful to define the depth of a formula.

DEFINITION 6

For $\varphi \in \Phi_{\mathcal{T}}$ define the *depth* $d(\varphi)$ of φ recursively by

- $d(p) = 0$ for $p \in \mathcal{P}$,
- $d(\neg\varphi_1) = d(\varphi_1)$,
- $d(\varphi_1 \vee \varphi_2) = \max(d(\varphi_1), d(\varphi_2))$,
- $d(\Box_a\varphi_1) = d(C_B\varphi_1) = d(\varphi_1) + 1$,
- $d(C_B(\varphi_1, \varphi_2)) = d([\varphi_1]\varphi_2) = \max(d(\varphi_1), d(\varphi_2)) + 1$,

- $d([U]\varphi_1) = d(\{U\}^*\varphi_1) = \max(d(\varphi_1), d(U)) + 1$,
- $d(U) = \max_{(u,a,u') \in U} (d(u), d(u') + 1)$.

We say that φ is of *pure depth* n if $d(\varphi) = n$ and there is no subformula φ' of φ such that $d(\varphi') = n$.

The only clause that may be somewhat surprising is that of $d(U)$. The reason for adding an extra $+1$ to the depth of end conditions is that they are evaluated in the next world, and thus reach one world further than a start condition of the same depth.

The concept of pure depth is useful to restrict the number of possibilities for the form of a formula; an AC formula of pure depth 1 e.g. must be either of the form $\Box_a \varphi'$ or of the form $C_B \varphi'$. The formulas of depth n are the Boolean combinations of the formulas of pure depth at most n .

Finally, it is useful to have a flexible definition of a path.

DEFINITION 7

Given an $\mathcal{L}_{\mathcal{T}}$ model $\mathcal{M} = (W, R, \nu)$ and two worlds w_1 and w_n of $\mathcal{M}$ a path π from w_1 to w_n is an ordered set of triples

$$\pi = ((w_1, a_1, w_2), (w_2, a_2, w_3), \dots, (w_{n-1}, a_{n-1}, w_n))$$

where $n \in \mathbb{N}$, and $a_i \in \mathcal{A}$ and $(w_i, w_{i+1}) \in R(a_i)$ for $1 \leq i \leq n-1$.

Let $B \subseteq \mathcal{A}$, φ a formula and U an update. The path π is a B -path if $a_i \in B$ for $1 \leq i \leq n-1$, a φ -path if $\mathcal{M}, w_i \models \varphi$ for $1 \leq i \leq n$ and a U -path if for $1 \leq i \leq n-1$ there is a clause $(u, a, u') \in U$ such that $\mathcal{M}, w_i \models u$, $a = a_i$ and $\mathcal{M}, w_{i+1} \models u'$. Conditions can be combined, π is an $(X_1, \dots, X_k)$ -path if it is an X_j -path for all $1 \leq j \leq k$.

4 The logics under consideration

Using different combinations of A, C, R, P, U and U^* we could define $2^6 = 64$ different fragments of $\mathcal{L}_{\mathcal{T}}$. Not all these fragments are interesting, however.

In this article, we do not consider logics that do not have the A connectives. This is not because logics without some or all of the A connectives are guaranteed to be uninteresting; logics with one of the common knowledge operators but not the $\Box_a$ operator are e.g. comparable to temporal logics with ‘future’ but not ‘next’. But such logics without A are outside the scope of this article and left for further work.

This leaves us with 32 fragments. There are however a few easy reductions of some connectives to other ones that allow us to reduce that number further.

LEMMA 1

For any $\mathcal{L}_{\mathcal{T}}$ formulas φ, ψ and any $B \subseteq \mathcal{A}$ we have $\models [\psi]\varphi \leftrightarrow (\psi \rightarrow [(\psi, \mathcal{A}, \psi)]\varphi)$, $\models C_B \varphi \leftrightarrow C_B(\top, \varphi)$ and $\models C_B(\psi, \varphi) \leftrightarrow \{(\psi, B, \psi)\}^* \varphi$.

The proof should be immediately clear and is left to the reader. Lemma 1 allows us to restrict ourselves to fragments having at most one of the update connectives $[\varphi]$ or $[U]$ and at most one of the common knowledge connectives C_B , $C_B(\cdot, \cdot)$ and $\{U\}^*$; if more than one of these connectives occurs in a logic only the ‘strongest’ one is relevant.⁵

This leaves 12 logics that can be ordered two dimensionally, with the update connective (if any) on one axis and the common knowledge connective (if any) on the other. The logics and their relative expressivities are shown in Figure 2.

⁵Note that among other things this implies that $\mathcal{L}_{\mathcal{T}} \equiv \mathcal{L}_{AUU^*}$.

FIGURE 2. The landscape of logics using basic modal logic (A) and a combination of public announcements (P), arrow updates (U), common knowledge (C), relativized common knowledge (R) and arrow update common knowledge (U*). Arrows $X \longrightarrow Y$ indicate that $X < Y$. Double arrows $X \longleftrightarrow Y$ indicate that $X \equiv Y$. Dashed gray arrows indicate previously established results, solid black arrows indicate new results. For arrows that are part gray and dashed and part black and solid the result in one direction was previously established but the result in the other direction is new. Boundaries around nodes indicate equivalence classes of logics that are equally expressive. For reasons of clarity not all arrows are drawn, but the omitted arrows all follow from the drawn ones by transitivity.

Note that although not all arrows are drawn the arrows in Figure 2 are sufficient to know the relative expressivity of any of the logics by transitivity and reflexivity of $\leq$.

4.1 Overview of previously known results

The arrows in Figure 2 that are drawn dashed and in gray were previously known. That $A \equiv AP$ was shown in [13]. That $AC < APC$ was shown in [4]. That $AR \equiv APR$ was shown in [12]. In [6] it was shown that $APC < AR$, which implies that $APC < APR$. Finally, in [11] it was shown that $AU \equiv AP$, $AUU^* \equiv AU^*$ and $AR \leq AU^*$. The remaining dashed arrows in Figure 2 are either trivial or follow from other dashed arrows by transitivity.

4.2 New expressivity results

The arrows that are drawn solid and in black in Figure 2 are new results. There are ten such new results, each corresponding to one half of an arrow. They all follow by transitivity from two results, however.

The first result is that $AU^* \not\leq AR$. This result is proven in Section 5. This result is not very surprising, it was already predicted in [11]. The second result is that $AU^* \leq AUC$. This result is proven in Section 6. Unlike the previous result this result is rather surprising, considering that the difference between AUC and AU^* is very similar to the difference between APC and AR and we have $AR \not\leq APC$.

The proof that $AU^* \leq AUC$ is very complicated and technical. We, therefore, only give an overview of the proof in this article itself and provide the details as supplementary data. But before looking at even an overview of the proof it might be worthwhile to consider what the difference between AUC and APC is that causes APC to be less expressive than AR whereas AUC is as expressive as AU^* . This difference is the ability to *store information*.

Recall that the informal readings of the operators $C_B(\varphi, \psi)$ and $\{U\}^*\psi$ look to the past. The formula $C_B(\varphi, \psi)$ holds if, after announcing $[\varphi]$, it will be common knowledge that ψ *used to* hold before the announcement. Likewise, $\{U\}^*\psi$ holds if, after announcing $[U]$, it will be common knowledge that ψ *used to* hold before the announcement. So if we want to simulate $C_B(\varphi, \psi)$ in APC or $\{U\}^*\psi$ in AUC we have to find a way to store information about the current model in such a way that it is not destroyed by the update. In APC we cannot do this, so AR is more expressive than APC. But arrow updates are more powerful than public announcements, and this extra power allows AUC to store just enough information to simulate $\{U\}^*\psi$. The ways in which AUC stores information are many and varied, but let us briefly consider one of them.

Suppose we want to apply an update $U = (u_1, a_1, u'_1), \dots, (u_n, a_n, u'_n)$ but first store information about which worlds satisfy a given formula φ . Then instead of updating with $[U]$ we could update with a $[U'] = [(u_1 \wedge \neg\varphi, a_1, u'_1), \dots, (u_n \wedge \neg\varphi, a_n, u'_n)]$. After this update every world that *used to* satisfy φ now satisfies $\Box\perp$, allowing us to recognize it. Of course there are many complications to this method. In particular, there might be worlds that satisfy $\neg\varphi \wedge [U]\Box\perp$. Such worlds will satisfy $\neg\varphi \wedge [U']\Box\perp$, so they will be false positives. But, through a lot of technical work, we can exclude that possibility. The important thing to note is that there is no similar way to store information using APC: there is no public announcement $[\alpha]$ that guarantees that every φ world satisfies $\neg[\alpha]\perp \wedge [\alpha]\Box\perp$ because there is no formula α that identifies the worlds that are accessible from a φ world (as opposed to the $\Diamond\varphi$ worlds from which a φ world is accessible).

5 AU* is more expressive than AR

The proof we give of the fact that AU* is more expressive than AR is very similar to the usual proofs of such results. We want to show that $AU^* \not\leq AR$, so that there are AU* formulas for which there is no equivalent AR formula. The most straightforward way to do this is to construct an AU* formula ξ and show that there is no AR formula equivalent to it.

At this point we would like to proceed by constructing a model $\mathcal{M}_c$ with two points w^+ and w^- such that ξ distinguishes between $\mathcal{M}_c, w^+$ and $\mathcal{M}_c, w^-$ even though there is no AR formula that distinguishes between the two worlds. This would be sufficient to show that there is no AR formula equivalent to ξ . Unfortunately this is too hard, we cannot find such a model. So we do something slightly more complicated.

Instead of creating a model where ξ distinguishes between two particular worlds w^+ and w^- we construct a model, where ξ distinguishes between two worlds u_{2i} and u_{2i+1} for every $i \in \mathbb{N}$. For each i there will be some AR formula φ_i that distinguishes between u_{2i} and u_{2i+1} , but we will show that for such φ_i it must hold that $d(\varphi_i) > i$. Every AR formula has a fixed and finite depth, so while for each i there is an AR formula φ_i that distinguishes u_{2i} from u_{2i+1} there is no AR formula φ_∞ that distinguishes u_{2i} from u_{2i+1} for every $i \in \mathbb{N}$. This shows that there is no AR formula equivalent to the AU* formula ξ .

Now let us construct the model $\mathcal{M}_c$ and the formula ξ . Let $\mathcal{M}_c$ be the model shown in Figure 3. Furthermore, let $\xi = \{(\top, a, \top), (p, c, \neg p \wedge \neg q), (q, b, \neg p \wedge \neg q)\}^*(\Diamond_a(p \vee q) \vee \Diamond_b(p \vee q))$. Note that the accessibility relations on $\mathcal{M}_c$ are reflexive, transitive and Euclidean so $\mathcal{M}_c$ is an S5 model.

LEMMA 2

For every $i \in \mathbb{N}$ we have $\mathcal{M}_c, u_{2i} \models \xi$ and $\mathcal{M}_c, u_{2i+1} \not\models \xi$.

PROOF. First let us look at the subformula $\Diamond_a(p \vee q) \vee \Diamond_b(p \vee q)$ of ξ . Every world other than u_0 and w_0 satisfies $\Diamond_a(p \vee q)$. Furthermore, u_0 satisfies $\Diamond_b(p \vee q)$. So w_0 is the only world in $\mathcal{M}_c$ that does

FIGURE 3. The model $\mathcal{M}_c$. This is an S5 models so every world has a reflexive arrow for every agent, but these are not drawn for reasons of clarity. For every $n \in \mathbb{N}$ the worlds s_n and u_{n+1} are connected by a , as are t_n and w_{n+1} .

FIGURE 4. The arrows in $\mathcal{M}_c$ that satisfy $\{(\top, a, \top), (p, c, \neg p \wedge \neg q), (q, b, \neg p \wedge \neg q)\}$. Reflexive arrows are omitted.

not satisfy $\Diamond_a(p \vee q) \vee \Diamond_b(p \vee q)$. As a result, any world in $\mathcal{M}_c$ satisfies ξ if and only if w_0 is not reachable from that world by a $\{(\top, a, \top), (p, c, \neg p \wedge \neg q), (q, b, \neg p \wedge \neg q)\}$ -path.

Now let us look at the set $\{(\top, a, \top), (p, c, \neg p \wedge \neg q), (q, b, \neg p \wedge \neg q)\}$ of clauses. The clause $(\top, a, \top)$ is satisfied by all a arrows. Most b - and c -arrows do not satisfy any of the clauses, however; only the b -arrows from t_i to u_i and the c -arrows from s_i to w_i satisfy $(q, b, \neg p \wedge \neg q)$ and $(p, c, \neg p \wedge \neg q)$, respectively. The arrows in $\mathcal{M}_c$ that satisfy any of the clauses are, therefore, as shown in Figure 4. The world w_0 is reachable from u_j by a $\{(\top, a, \top), (p, c, \neg p \wedge \neg q), (q, b, \neg p \wedge \neg q)\}$ -path if and only if j is odd. This implies that $\mathcal{M}_c, u_{2i} \models \xi$ and $\mathcal{M}_c, u_{2i+1} \not\models \xi$ for every $i \in \mathbb{N}$. ■

Now we should show that no AR formula φ of depth $n := d(\varphi)$ can distinguish between u_{2i} and u_{2i+1} for any $i \geq n$. To make the induction step in the proof proceed smoothly it is convenient to prove a slightly stronger lemma.

LEMMA 3

Take any $n \in \mathbb{N}$, any $i \geq n$, any $x \in \{s, t, u, w\}$ and any $\varphi \in \Phi_{\text{AR}}$ such that $d(\varphi) \leq n$. Then $\mathcal{M}_c, x_{2i} \models \varphi$ if and only if $\mathcal{M}_c, x_{2i+1} \models \varphi$.

PROOF. By induction on n . As base case suppose that $n=0$. Then φ is a Boolean combination of propositional variables. The worlds x_{2i} and x_{2i+1} have the same values for all propositional variables so this φ cannot distinguish between them.

Suppose then as induction hypothesis that $n > 0$ and that the lemma holds for all $n' < n$. The proof for $x=s$ is completely analogous to the proof for $x=t$ and the proof for $x=u$ is likewise completely analogous to the proof for $x=w$. We therefore omit the proofs for $x \in \{t, w\}$ and only show that the lemma holds for $x \in \{s, u\}$.

Suppose towards a contradiction that φ distinguishes between x_{2i} and x_{2i+1} . If a Boolean combination of formulas distinguishes between two worlds then so does at least one of the combined formulas. We can therefore assume without loss of generality that φ is of pure depth. Furthermore, if

FIGURE 5. An illustration of why $C_B(\varphi_1, \varphi_2)$ cannot distinguish between u_j and u_{j+1} if w_j and s_j do not satisfy φ_1 . Worlds not satisfying φ_1 are crossed out.

$d(\varphi) < n$ it follows immediately from the induction hypothesis that φ does not distinguish between x_{2i} and x_{2i+1} so we can also assume that $d(\varphi) = n$.

So φ is of pure depth $n > 0$, which implies that it must be of the form $\Box_\alpha \varphi_1$ or $C_B(\varphi_1, \varphi_2)$ for some $\alpha \in \mathcal{A}$, $B \subseteq \mathcal{A}$ and $\varphi_1, \varphi_2 \in \Phi_{AR}$ such that $d(\varphi_1), d(\varphi_2) < n$.

First, let us suppose that $\varphi = \Box_\alpha \varphi_1$. We have $2i \geq 2n > 0$ so for every arrow from x_{2i} to y_k (with $y \in \{s, t, u, w\}$ and $2i+1 \geq k \geq 2i-1$) there is a corresponding arrow from x_{2i+1} to y_{k+1} and vice versa. This implies that in order for φ to distinguish between x_{2i} and x_{2i+1} it is necessary that φ_1 distinguishes between some y_k and y_{k+1} with $y \in \{s, t, u, w\}$ and $2i+1 \geq k \geq 2i-1$. But $k \geq 2i-1 \geq 2n-1 > 2(n-1) \geq 2d(\varphi_1)$ so φ distinguishing between y_k and y_{k+1} would contradict the induction hypothesis. The formula φ therefore cannot distinguish between x_{2i} and x_{2i+1} .

Let us then suppose that $\varphi = C_B(\varphi_1, \varphi_2)$. If B does not contain both a and at least one of b, c then (B, φ) -paths cannot take us far in this model: if $a \notin B$ then at most the worlds s_{2i}, u_{2i}, t_{2i} and w_{2i} are B -reachable from x_{2i} and if $b, c \notin B$ then at most the worlds s_{2i}, u_{2i+1} are B -reachable from s_{2i}, t_{2i} and w_{2i+1} from t_{2i}, u_{2i} and s_{2i-1} from u_{2i} or w_{2i} and t_{2i-1} from w_{2i} . In each case there are counterparts to these worlds reachable from x_{2i+1} and by the induction hypothesis φ_2 cannot distinguish between these counterpart worlds. So $C_B(\varphi_1, \varphi_2)$ cannot distinguish between x_{2i} and x_{2i+1} .

So to distinguish between x_{2i} and x_{2i+1} the set B must contain both a and at least one of b, c . If x_{2i} and x_{2i+1} are reachable from each other by a (B, φ_1) -path then, independent of φ_2 , we have $\mathcal{M}_c, x_{2i} \models C_B(\varphi_1, \varphi_2)$ if and only if $\mathcal{M}_c, x_{2i+1} \models C_B(\varphi_1, \varphi_2)$. So to distinguish the worlds there must be at least one $\neg\varphi_1$ world on every B -path from x_{2i} to x_{2i+1} . Furthermore, this $\neg\varphi_1$ world cannot be either x_{2i} or x_{2i+1} , since then $C_B(\varphi_1, \varphi_2)$ would reduce to φ_2 on both worlds and φ_2 cannot distinguish between them.

The model $\mathcal{M}_c$ is constructed in such a way that if every B -path from x_{2i} to x_{2i+1} contains at least one $\neg\varphi_1$ world then so does every B -path from either of those worlds to any world further than two steps away. Exactly which worlds must satisfy $\neg\varphi_1$ depends on x and B though.

Suppose $x = u$ and $B \cap \{a, b, c\} = \{a, b, c\}$. Then in order for every B -path between u_{2i} and u_{2i+1} to contain at least one $\neg\varphi_1$ world we have to have $\mathcal{M}_c, s_j \not\models \varphi_1$ and either $\mathcal{M}_c, t_j \not\models \varphi_1$ or $\mathcal{M}_c, w_{2i+1} \not\models \varphi_1$. But, since φ_1 is of depth $\leq n-1$ it follows from the induction hypothesis that it cannot distinguish between $s_{2i+2}, s_{2i+1}, s_{2i}$ and s_{2i-1} , between $w_{2i+2}, w_{2i+1}, w_{2i}$ and w_{2i-1} or between $t_{2i+2}, t_{2i+1}, t_{2i}$ and t_{2i-1} . But then every (B, φ_1) -path from u_{2i} can contain at most the worlds u_{2i} and t_{2i} . By the induction hypothesis none of these worlds can be distinguished from their counterpart by φ_2 , so $C_B(\varphi_1, \varphi_2)$ cannot distinguish between u_{2i} and u_{2i+1} . See Figure 5 for an illustration of the (B, φ_1) -paths if w_{2i} and s_{2i} do not satisfy φ_1 .

The situations for the other options for x and B are similar. The following table shows which worlds have to satisfy $\neg\varphi_1$ and which worlds (B, φ_1) -paths from x_{2i} can contain, at the most. Note that only maximal (with respect to set inclusion) sets are given. For example, one of the options for

$\neg\varphi_1$ worlds in the case $x=u$ and $\{a, b, c\}$ is s_k and t_k (for all $2i+2 \geq k \geq 2i-1$) in which case paths from x_{2i} can contain only the world x_{2i} itself. But the other option is for s_k and w_k to satisfy $\neg\varphi_1$ in which case paths from x_{2i} could contain both u_{2i} and t_{2i} .

x	$B \cap \{a, b, c\}$	Must satisfy $\neg\varphi_1$ ($\forall k \in \{2i-1, 2i, 2i+1, 2i+2\}$)	Paths from x_{2i} can contain at most
u	$\{a, b, c\}$	s_k and t_k or w_k	$\{u_{2i}, t_{2i}\}$
u	$\{a, b\}$	s_k	$\{u_{2i}, t_{2i}, w_{2i+1}\}$
u	$\{a, c\}$	s_k, t_k or w_k	$\{u_{2i}, s_{2i-1}, t_{2i-1}\}$ or $\{u_{2i}, s_{2i-1}, w_{2i-1}\}$
s	$\{a, b, c\}$	u_k and w_k or t_k	$\{s_{2i}, t_{2i}\}$ or $\{s_{2i}, w_{2i}\}$
s	$\{a, b\}$	u_k	$\{s_{2i}, t_{2i}, w_{2i+1}\}$
s	$\{a, c\}$	t_k or w_k	$\{s_{2i}, u_{2i+1}, t_{2i}\}$

In every one of these cases it follows from the induction hypothesis that φ_2 cannot distinguish between the path from x_{2i} and the path from x_{2i+1} . So $C_B(\varphi_1, \varphi_2)$ does not distinguish between x_{2i} and x_{2i+1} .

For every possible form of φ we have now shown that φ does not distinguish between x_{2i} and x_{2i+1} , contradicting the assumption that φ does distinguish between them and thereby completing the proof. ■

Recall that we did not assume any of the frame conditions often used in epistemic logic, so we use any K-model as opposed to only KD45-, S4- or S5-models. The model $\mathcal{M}_c$ that we used in the proof of Lemma 3 is an S5 model (and therefore also a KD45- and S4-model) though. As a result the proof does not depend on the fact that we use K-models, allowing us to conclude that AU* is more expressive than AR not only over K but also over KD45, S4 and S5.

6 AUC is equally expressive as AU*

In this section, we show that $AU^* \leq AUC$. Unfortunately the proof is very long and technical. We therefore give only an overview of the proof here and leave the full proof as Appendices A, B and C in the supplementary data.

6.1 Notation

Due to the technical nature of the proof even the overview is made easier by introducing some more notation. First, let us define some abbreviations regarding $\Box$ and $\Diamond$.

DEFINITION 8

For any $\varphi \in \Phi_{\mathcal{T}}$, $B \subseteq \mathcal{A}$ and U an arrow update let

- $\Box_B \varphi$ stand for $\varphi \wedge \Box_B \varphi$ and $\Diamond_B \varphi$ stand for $\varphi \vee \Diamond_B \varphi$,
- $\Box_U \varphi$ stand for $\bigwedge_{(u_1, a, u_2) \in U} (u_1 \rightarrow \Box_a(u_2 \rightarrow \varphi))$,
- $\Diamond_U \varphi$ stand for $\bigvee_{(u_1, a, u_2) \in U} (u_1 \wedge \Diamond_a(u_2 \wedge \varphi))$.

The formulas $\Box_U \varphi$ and $\Diamond_U \varphi$ thus state that φ holds in every/at least one U -successor.⁶

It is also convenient to be able to specify certain arrows that are *not* to be retained. We do this by overlining the clauses that specify arrows that should be removed. Let

⁶Note that $\Box_U$ is the static operator associated with $\Box$ and $[U]$ and that $\Diamond_U$ is the static operator associated with $\Diamond$ and $[U]$, as discussed in Section 2.1.

$U_1 = \{(u_1, a_1, u'_1), \dots, (u_k, a_k, u'_k)\}$ and $U_2 = \overline{\{(u_{k+1}, a_{k+1}, u'_{k+1}), \dots, (u_{k+l}, a_{k+l}, u'_{k+l})\}}$. An arrow is retained by the update $U = U_1 \cup U_2$ iff it satisfies at least one of the clauses of U_1 and none of the clauses of U_2 . We can define updates with overlined clauses as abbreviations of updates without overlined clauses. The trick is to note that an arrow from w_1 to w_2 for agent a satisfies none of the clauses in U_2 if and only if for every clause $(u, a, u') \in U_2$ either $\mathcal{M}, w_1 \not\models u$ or $\mathcal{M}, w_2 \not\models u'$. In such a case we can partition U_2 into U' and $U_2 \setminus U'$, where w_1 does not satisfy the start conditions of U' and w_2 does not satisfy the end conditions of $U_2 \setminus U'$. So, for any $U' \subseteq U_2$, if $(u_i, a_i, u'_i) \in U_1$ then $\left(u_i \wedge \bigwedge_{(u_j, a_j, u'_j) \in U'} \neg u_j, a_i, u'_i \wedge \bigwedge_{(u_j, a_j, u'_j) \in U_2 \setminus U'} \neg u'_j \right)$ satisfies a clause from U_1 and none of the clauses from U_2 .

DEFINITION 9

For $U_1 = \{(u_1, a_1, u'_1), \dots, (u_k, a_k, u'_k)\}$ and $U_2 = \overline{\{(u_{k+1}, a_{k+1}, u'_{k+1}), \dots, (u_{k+l}, a_{k+l}, u'_{k+l})\}}$ let $U_1 \cup U_2$ stand for

$$\left\{ \left(u_i \wedge \bigwedge_{(u_j, a_j, u'_j) \in U'} \neg u_j, a_i, u'_i \wedge \bigwedge_{(u_j, a_j, u'_j) \in U_2 \setminus U'} \neg u'_j \right) \mid (u_i, a_i, u'_i) \in U_1, U' \in \wp(U_2) \right\}.$$

We use $\overline{U}$ as shorthand for $\{\top, \mathcal{A}, \top\} \cup \{\overline{(u, a_i, u')} \mid (u, a_i, u') \in U\}$. The formulas $\Box_{\overline{U}}\varphi$ and $\Diamond_{\overline{U}}\varphi$ thus state that φ holds in every/at least one world that is a successor but not a U -successor.

We also need notation for two more concepts about formulas.

DEFINITION 10

For $\varphi \in \Phi_{\mathcal{T}}$ let $\text{Pvar}(\varphi)$ be the set of propositional variables that occur in φ .

DEFINITION 11

For $Q \subseteq \mathcal{P}$ and $n \in \mathbb{N}$ let $\Phi_Q^n := \{\varphi \in \Phi_{\text{AUC}} \mid d(\varphi) \leq n \text{ and } \text{Pvar}(\varphi) \subseteq Q\}$.

The main use of Φ_Q^n will be in conjunctions $\bigwedge_{\varphi \in \Phi_Q^n} \psi_\varphi$. Strictly speaking this is of course not a formula, as it contains an infinite number of conjuncts. However, if Q is finite—as it will be when we use it—the set Φ_Q^n contains only a finite number of *mutually non-equivalent* formulas. We can then consider $\bigwedge_{\varphi \in \Phi_Q^n} \psi_\varphi$ to be a conjunction over some maximal choice of non-equivalent formulas in Φ_Q^n .

6.2 Variable use

In the proof that $\text{AU}^* \leq \text{AUC}$ a large number of formulas are defined. Although the names given to the formulas do not, strictly speaking, matter there is a pattern in the naming, and knowing this pattern may aid in understanding the proof. The proof finds a AUC formula α that is equivalent to $\{U\}^*\varphi$ by using a case distinction.

A formula δ_i is a AU^* formula corresponding to case i . A formula β_i is an AUC formula that is both necessary and sufficient for being in case i , so $\models \delta_i \leftrightarrow \beta_i$. A formula α_i finally is a AUC formula that is equivalent to $\{U\}^*\varphi$ given that we are in case i , so $\models (\delta_i \wedge \{U\}^*\varphi) \leftrightarrow (\beta_i \wedge \alpha_i)$.⁷

⁷As the choice of the symbols α, β and δ suggests there are also formulas γ_i . These only occur in the detailed proof in the supplementary data though, where γ_i is a necessary condition for being in case i , so $\models \delta_i \rightarrow \gamma_i$.

FIGURE 6. An example of using $\neg p$ worlds as witnesses. The formula φ holds everywhere except where noted otherwise.

6.3 The main strategy

Fix any arrow update U containing only AUC formulas and any AUC formula φ and let $\chi := \{U\}^* \varphi$. If we can find a AUC formula α such that $\models \alpha \leftrightarrow \chi$ that would suffice to show that AUC is at least as expressive as AU*.

What we need then is a strategy to find such α . This poses two challenges. First, given any pointed model $\mathcal{M}, w$ we must identify the worlds that are U -reachable from w . Secondly, we must check whether φ holds in all of those worlds.

The most straightforward method to identify the U -reachable worlds is to update with $[U]$; the U -reachable worlds in $\mathcal{M}$ are exactly the reachable worlds in $\mathcal{M}_{[U]}$. Unfortunately the update $[U]$ may destroy information, so given a world w' of $\mathcal{M}$ it may be impossible to determine from $\mathcal{M}_{[U]}$ whether $\mathcal{M}, w' \models \varphi$. By using this simple method to solve the first problem we would make it impossible to solve the second problem.

So to solve both problems we need to update with a different arrow update U' . This U' will be very similar to U so the reachable worlds in $\mathcal{M}_{[U']}$ are mostly the U -reachable worlds in $\mathcal{M}$. But in addition to most arrows from U the update U' will retain just enough structure to create *witnesses* for the existence of certain worlds in $\mathcal{M}$. The question then is what worlds we want to create witnesses for and how we want to use them. Again there is a straightforward choice, namely to create witnesses for $\neg\varphi$ worlds. But, again, this straightforward choice runs into trouble. So instead we create witnesses for worlds that are ‘on the boundary’ of the U -reachable area, so those worlds reached by a U -arrow from which a $\bar{U}$ -arrow departs. So we take U' in such a way that in $\mathcal{M}_{[U']}$ every maximal path ends in a witness world.

We then make one final change. Let $U'' = U' \cup \{(\neg\varphi, \bar{\mathcal{A}}, \top)\}$. This change cuts all paths that contain a $\neg\varphi$ world. So in $\mathcal{M}_{[U'']}$ maximal paths end in a witness world if and only if they do not contain a world that was a $\neg\varphi$ world in $\mathcal{M}$. So $\mathcal{M}, w \models \{U\}^* \varphi$ if and only if every maximal path from $\mathcal{M}_{[U'']}$ ends in a witness world. There are of course several complications, but those can be dealt with. Let us look at a very simple example.

EXAMPLE 2

Suppose $U = (p, \bar{\mathcal{A}}, p)$. Then we can create witnesses for U -reachable worlds from which a $\bar{U}$ -arrow departs by retaining arrows from p worlds to $\neg p$ worlds, so by taking $U' = (p, \bar{\mathcal{A}}, p), (p, \bar{\mathcal{A}}, \neg p)$. We then have $U'' = (p, \bar{\mathcal{A}}, p), (p, \bar{\mathcal{A}}, \neg p), (\neg\varphi, \bar{\mathcal{A}}, \top)$. Consider the pointed models $\mathcal{M}, w$ and $\mathcal{M}', w'$ as shown in Figure 6. We want to check whether every path from w (resp. w') in $\mathcal{M}_{[U'']}$ (resp. $\mathcal{M}'_{[U'']}$)

ends in a witness world, so we check for $C_A(\Box \perp \rightarrow \neg p)$. We have $\mathcal{M}_{[U'']}, w \not\models C_A(\Box \perp \rightarrow \neg p)$ and $\mathcal{M}'_{[U'']}, w \models C_A(\Box \perp \rightarrow \neg p)$ so $\mathcal{M}, w \not\models \{U\}^* \varphi$ and $\mathcal{M}', w' \models \{U\}^* \varphi$.

But even with an example as simple as $U = (p, \mathcal{A}, p)$ complications can occur if we look at different models. For example, the method as described above will not work if the origin world w satisfies either $\neg p$ or $C_A p$. Both these cases can be easily dealt with, however, by checking for them before applying $[U'']$ and treating them separately.

A more interesting kind of complication is if $\mathcal{M}, w \not\models C_A p$ but there are some branches that are reachable from w that do satisfy $C_A p$. If we simply apply $[U'']$ these branches will end in a p world, and, therefore, look just like branches that were cut short due to the presence of a $\neg \varphi$ world. So we risk getting false positives for the detection of $\neg \varphi$ worlds. The solution to this complication is to cut off all $C_A p$ branches unless they contain a $\neg \varphi$ world. This means we have to modify U'' to $U''' = U'' \cup \{(\overline{\top}, \mathcal{A}, C_A(p \wedge \varphi))\}$.

These are all the complications we can encounter for this simple U , however. The formula $\{(p, \mathcal{A}, p)\}^* \varphi$ is equivalent to

$$\begin{aligned} & \varphi \wedge (C_A p \rightarrow C_A \varphi) \wedge (\neg C_A p \rightarrow \\ & [(p, \mathcal{A}, p), (p, \mathcal{A}, \neg p), (\neg \varphi, \mathcal{A}, \top), (\top, \mathcal{A}, C_A(p \wedge \varphi))] C_A(\Box \perp \rightarrow \neg p)), \end{aligned}$$

where the first two conjuncts take care of the two degenerate cases and the third uses witnesses as described above.

6.4 Creating witnesses

The main strategy requires us to create witnesses for worlds that are on the boundary of a U -area. We leave the details of how to do this to the supplementary data, but here we do present a global overview of why it is always possible to create such a witness.

The important realization is that whenever there is a world reached by a U -arrow and from which a $\overline{U}$ -arrow departs then there must be a *simple* difference between two *nearby* objects. We call this difference between the two objects a *boundary condition*.⁸

Here *nearby* means ‘reachable in at most $d(\{U\}^* \varphi)$ steps’. What it means for a difference to be *simple* is a little more complicated. Let us take a closer look at the situation. We have a world, call it w_2 , that is reached by a U -arrow and from which a $\overline{U}$ arrow departs. Let w_1 be the source of the U -arrow and w_3 the destination of the $\overline{U}$ -arrow.

Let us focus on the arrows from w_1 to w_2 and from w_2 to w_3 for a moment. The arrow from w_1 to w_2 is a U -arrow so there is a clause $(\psi_1, a, \psi_2) \in U$ that the arrow satisfies. The arrow from w_2 to w_3 is not a U -arrow so in particular it does not satisfy (ψ_1, a, ψ_2) . Then there are three possibilities: the first possibility is that the arrow from w_2 to w_3 is not an a -arrow. The second possibility is that $\mathcal{M}, w_2 \not\models \psi_1$. The third possibility is that $\mathcal{M}, w_3 \not\models \psi_2$. For the latter two possibilities see Figure 7.

If the first possibility holds the nearby objects that differ are arrows, and the simple difference between them is that they belong to different agents. So the first boundary condition is that there are nearby arrows belonging to different agents. Let us suppose then that there are no two nearby arrows that belong to different agents. So we are working in a single-agent part of the model.

⁸Note that boundary conditions are not called so because they only occur on the boundary. Instead they are called so because they have to occur on every boundary.

FIGURE 7. Two possibilities for a $\overline{U}$ -arrow following a U -arrow.

Let us take a closer look at the second and third possibilities. In both cases there is a U -reachable world w_i satisfying $\psi_i \wedge \Diamond \neg \psi_i$ with $i \in \{1, 2\}$. This difference between the world w_i satisfying ψ_i and the world w_{i+1} satisfying $\neg \psi_i$ has to ‘come from somewhere’.

One possible cause for the difference between w_i and w_{i+1} is the existence of two nearby worlds w' and w'' and a propositional variable p such that $\mathcal{M}, w' \models p$ and $\mathcal{M}, w'' \not\models p$. This is the second kind of boundary condition; a difference in the value of a propositional variable in two worlds.

Suppose then that we are in a situation where there are no nearby arrows for different agents and no nearby difference in propositional variables. Then the only possible cause for the difference between w_i and w'_i is the existence of one or more worlds satisfying $\Box \perp$ (and some satisfying $\Diamond \top$). This is the third kind of simple difference, some worlds having a successor and other worlds having none.

For technical reasons it is convenient to split this simple difference up into two boundary conditions. The first boundary condition is when there is a nearby world satisfying $\Diamond \psi' \wedge \Diamond \neg \psi'$, with ψ' a formula of depth lower than ψ_i . We are still in the situation where there are no different agents and no different propositional variables, so either the ψ' or the $\neg \psi'$ branch must contain a nearby $\Box \perp$ world. But it is not necessary that both branches contain such a world. The other boundary condition is when w_i is near a dead end. In that case $\mathcal{M}, w_i \models \Diamond^k \Box \perp \wedge \Box^{k+1} \perp$ for some $k \leq n$.

So the four kinds of boundary condition are:

- (1) Two arrows that belong to different agents.
- (2) Two worlds that have a different value for some propositional variable.
- (3) A world satisfying $\Diamond \psi' \wedge \Diamond \neg \psi'$ for some ψ' .
- (4) The world w_i being one step further away from a dead end than the world w_{i+1} .

In addition to these four conditions there are also two ‘degenerate boundary conditions’. These represent situations where there is no boundary at all.

- (5) There are reachable U -arrows but no reachable $\overline{U}$ -arrows.
- (6) There are no departing U -arrows at all.

These are the only possibilities. This should be intuitively clear, but we also provide a full proof in Section C in the supplementary data.

For each of the first four boundary conditions we can create a witness, and in the last two we do not need a witness because $\{U\}^* \varphi$ reduces to $C_A \varphi$ or φ respectively. The witness for the first

FIGURE 8. The U -reachable part of an example model $\mathcal{M}_t$ used to illustrate the method of working backward through the cases. The dots at the end of each branch represent the fact that the model continues, but with non- U arrows.

boundary condition is simply the arrow that belongs to a different agent. The witness for the second boundary condition is a world with a different value for the variable. The witness for the fourth boundary condition is the dead end itself. The only difficult boundary condition is the third, where we have $\Diamond\psi' \wedge \Diamond\neg\psi'$. In that case we cut off either the branch starting at the ψ' successor or the branch starting at the $\neg\psi'$ successor. This results in a world satisfying $\Diamond\Box\perp \wedge \Diamond\Box\top$, which we use as witness.

6.5 The case distinction

We have six different boundary conditions (two of which are degenerate). Unfortunately each of those is solved in a different way. This can be problematic, because there may be different branches with different boundary conditions and we can only solve one of them at a time. The solution is to make a case distinction that allows us to solve the different conditions one at a time.

We have our four types of boundary condition and two types of degenerate boundary condition. Based on these types we want to make a case distinction. To every world in any model we assign one of six cases. Case number i is associated with boundary condition type i , but not in an entirely straightforward way: a world is in case i if and only if there is at least one U -reachable world satisfying boundary condition type i and no U -reachable world satisfying boundary condition type $j < i$. Do recall that boundary conditions may ‘coincidentally’ hold in worlds that are not on the boundary of the U -area. These ‘coincidental’ boundary conditions do count for which case we are in.

Defining the cases in this way gives us three important properties for our case distinction. First, the cases are mutually exclusive, since case i requires a lack of worlds with condition $j < i$. Secondly, the cases are exhaustive, there is always an accessible branch with one of the (possibly degenerate) boundary conditions. Finally, if a world w is in case i and w' is U -reachable from w then w' is either in case i or in a case $j > i$. After all, if a world with boundary condition $k < i$ would be U -reachable from w' then the same branch would be U -reachable from w contradicting the assumption that w is in case i .

These three properties allow us to solve all cases by ‘working backwards’. We first find a formula α_{i+1} that is equivalent to $\{U\}^*\varphi$ on worlds in case $i+1$. Then we use the fact that we have already solved case $i+1$ to solve case i . The process is best explained with the help of a series of example figures, so consider Figures 8–11.

FIGURE 9. The model $\mathcal{M}_t[U_3'']$. Worlds not in case 3 are grayed out. A witness for condition 3 is represented by c_3 .

EXAMPLE 3

First, consider the model $\mathcal{M}_t$ as shown in Figure 8. If a world satisfies boundary condition i then the world is labelled b_i . There could be boundary conditions of many different types but to keep the example simple we consider an example where only boundary conditions of types 1, 2 and 3 occur.

Note that some worlds satisfy multiple boundary conditions. In particular, every world that is the endpoint of a branch, and therefore, satisfies $\Box U \perp$, satisfies condition 6 in addition to any other conditions it may satisfy. This does not matter though; if a world w satisfies conditions i and $j > i$ then any world that can reach w is either in case i or in some case $k < i$. Also note that there is one ‘coincidental’ boundary condition of type 1 in the lowest branch of the model. Some of the worlds are labelled w_j for some $1 \leq j \leq 5$. We will be taking a closer look at those worlds. The worlds w_1 and w_2 are in case 3, the world w_3 is in case 2 and the worlds w_4 and w_5 are in case 1.

To determine on which worlds in this model the formula $\{U\}^*\varphi$ holds we start by considering those worlds that are in case 3. There are no boundary conditions of type 4, 5 or 6 in this model so we know how to solve case 3; we use an update $[U_3'']$ to create witnesses for worlds satisfying boundary condition 3 and additionally remove arrows to $\neg\varphi$ worlds. See Figure 9. We can then see that every path from w_1 in $\mathcal{M}_t[U_3'']$ ends in a witness world, so $\mathcal{M}_t, w_1 \models \{U\}^*\varphi$. There is, however, a path from w_2 in $\mathcal{M}_t[U_3'']$ that does not end in a witness world so $\mathcal{M}_t, w_2 \not\models \{U\}^*\varphi$.

Now we can use what we have learned about case 3 to solve case 2. Just like in case 3 we use an update U_2'' that creates witnesses for boundary condition 2 and that removes arrows to $\neg\varphi$ worlds. But we add one additional clause. Recall that we use β_3 for the formula that identifies the case 3 worlds and α_3 for the formula that is equivalent to $\{U\}^*\varphi$ in case 3 worlds. This means that the clause $(\top, \mathcal{A}, \alpha_3 \wedge \beta_3)$ removes arrows to case 3 worlds if and only if they satisfy $\{U\}^*\varphi$.⁹ The resulting model $\mathcal{M}_t[U_2'' \cup (\top, \mathcal{A}, \alpha_3 \wedge \beta_3)]$ is shown in Figure 10. Note that all paths from w_3 end in a witness world, so $\mathcal{M}_t, w_3 \models \{U\}^*\varphi$.

Finally, we can use both previous cases to solve case 1. Again we start with an update U_1'' to create witnesses and remove arrows to $\neg\varphi$ worlds. But now we add two extra clauses, $(\top, \mathcal{A}, \alpha_2 \wedge \beta_2)$ and $(\top, \mathcal{A}, \alpha_3 \wedge \beta_3)$. See Figure 11 for the resulting model $\mathcal{M}_t[U_3'' \cup (\top, \mathcal{A}, \alpha_2 \wedge \beta_2) \cup (\top, \mathcal{A}, \alpha_3 \wedge \beta_3)]$. We have $\mathcal{M}_t, w_2 \not\models \alpha_3$, so the arrow to w_2 is not removed by the update. As a result there is a path from w_5 that does not end in a witness world, so $\mathcal{M}_t, w_5 \not\models \{U\}^*\varphi$. Note that, unlike the arrow to w_2 , the arrow to

⁹If there were branches in cases 4, 5 or 6 we would have to add additional clauses $(\top, \mathcal{A}, \alpha_4 \wedge \beta_4)$, $(\top, \mathcal{A}, \alpha_5 \wedge \beta_5)$ and $(\top, \mathcal{A}, \alpha_6 \wedge \beta_6)$ as well.

FIGURE 10. The model $\mathcal{M}_{t[U_2'' \cup (\overline{T}, \mathcal{A}, \alpha_3 \wedge \beta_3)]}$. Worlds not in case 2 are grayed out. A witness for condition 2 is represented by c_2 .

FIGURE 11. The model $\mathcal{M}_{t[U_3'' \cup (\overline{T}, \mathcal{A}, \alpha_2 \wedge \beta_2) \cup (\overline{T}, \mathcal{A}, \alpha_3 \wedge \beta_3)]}$. Worlds not in case 1 are grayed out. A witness for condition 1 is represented by c_1 .

w_3 is removed, because $\mathcal{M}_t, w_3 \models \alpha_2 \wedge \beta_2$. As a result every path from w_4 ends in a witness world, so $\mathcal{M}_t, w_4 \models \{U\}^* \varphi$.

6.6 Further complications

The method detailed above allows us to deal with the most important complication, namely that there could be different branches with different boundary conditions. Unfortunately there are several remaining complications. Notable examples include the fact that the $\overline{U}$ -arrows retained to create witnesses might connect one U -area to another, and the fact that for one of the witness types it is impossible to tell whether a given path ends in (as opposed to contains) a witness world. We also need a large number of subcases in addition to the six main cases.

These complications only arise in the detailed proof, however, so the ways to deal with these complications are also given there.

6.7 Formulas representing the cases

Above the different cases were described informally. But before proving that $\text{AU}^* \leq \text{AUC}$ we should find formal descriptions of the cases. Which case we are in is determined by the existence or non-existence of certain kinds of worlds within a certain distance of a U -reachable world. This kind of

condition is easily phrased in AU^* but not in AUC so let us first describe the conditions in English and AU^* .

Recall that $\chi = \{U\}^* \varphi$, where φ is a AUC formula and U contains only AUC formulas. Let $n := d(\chi)$. The nearby difference must then be within distance n of a U -reachable world. For technical reasons we check for some differences up to a distance of a multiple of n though. This gives us the following cases.

- (1) We are in the first case if $\Diamond_U \top$ and there are at least two agents a_1 and a_2 that have an arrow within distance $3n$ of a U -reachable world. The AU^* representation of this case is

$$\delta_1 := \Diamond_U \top \wedge \bigvee_{a_1 \neq a_2 \in \mathcal{A}} (\neg\{U\}^* \Box^{3n} \Box_{a_1} \perp \wedge \neg\{U\}^* \Box^{3n} \Box_{a_2} \perp).$$

- (2) We are in the second case if $\Diamond_U \top$, we are not in the first case and there is a propositional variable $p \in \text{Pvar}(\chi)$ such that both p and $\neg p$ hold within distance $3n$ of a U -reachable world. The AU^* representation of this case is

$$\delta_2 := \Diamond_U \top \wedge \neg \delta_1 \wedge \bigvee_{p \in \text{Pvar}(\chi)} (\neg\{U\}^* \Box^{3n} p \wedge \neg\{U\}^* \Box^{3n} \neg p).$$

- (3) We are in the third case if $\Diamond_U \top$, we are not in the first or second case and there is a world w_2 within distance n of a U -reachable world such that the successors of w_2 are distinguishable by a AUC formula of depth at most $2n$ using only the propositional variables in χ . The AU^* representation of this case is

$$\delta_3 := \Diamond_U \top \wedge \neg \delta_1 \wedge \neg \delta_2 \wedge \bigvee_{\psi \in \Phi_{\text{Pvar}(\chi)}^{2n}} \neg\{U\}^* \neg \Diamond^n (\Diamond \psi \wedge \Diamond \neg \psi)$$

- (4) We are in the fourth case if $\Diamond_U \top$, we are not in one of the previous cases and there is a U -reachable world, where $\Diamond_{\overline{U}} \top$ holds. The AU^* representation of this case is

$$\delta_4 := \Diamond_U \top \wedge \neg \delta_1 \wedge \neg \delta_2 \wedge \neg \delta_3 \wedge \neg\{U\}^* \neg \Diamond_{\overline{U}} \top$$

- (5) We are in the fifth case if $\Diamond_U \top$, we are not in one of the previous cases and there is no U -reachable world where $\Diamond_{\overline{U}} \top$ holds. The AU^* representation of this case is

$$\delta_5 := \Diamond_U \top \wedge \neg \delta_1 \wedge \neg \delta_2 \wedge \neg \delta_3 \wedge \neg \delta_4 \wedge \{U\}^* \neg \Diamond_{\overline{U}} \top$$

- (6) We are in the sixth case if $\neg \Diamond_U \top$. The AU^* representation of this case is

$$\delta_6 := \neg \Diamond_U \top.$$

To construct the AUC formula α that is equivalent to the AU^* formula χ we first have to find AUC formulas $\beta_1, \dots, \beta_6$ such that $\models \beta_i \leftrightarrow \delta_i$ for $i \in \{1, \dots, 6\}$. Then we find α_6 such that $\models (\delta_6 \wedge \chi) \leftrightarrow (\beta_6 \wedge \alpha_6)$, use this α_6 to find α_5 such that $\models (\delta_5 \wedge \chi) \leftrightarrow (\beta_5 \wedge \alpha_5)$ and so on until we have $\alpha_1, \dots, \alpha_6$ such that

$$\alpha = (\beta_1 \rightarrow \alpha_1) \wedge (\beta_2 \rightarrow \alpha_2) \wedge (\beta_3 \rightarrow \alpha_3) \wedge (\beta_4 \rightarrow \alpha_4) \wedge (\beta_5 \rightarrow \alpha_5) \wedge (\beta_6 \rightarrow \alpha_6).$$

6.8 $AU^* \preceq AUC$

LEMMA 4

There are AUC formula $\beta_1, \dots, \beta_6$ such that $\models \beta_i \leftrightarrow \delta_i$ for all $1 \leq i \leq 6$.

The proof of this is included as Section A in the supplementary data.

LEMMA 5

There are AUC formulas $\alpha_6, \dots, \alpha_1$ such that $\models (\delta_i \wedge \chi) \leftrightarrow (\delta_i \wedge \alpha_i)$ for all $6 \geq i \geq 1$.

The proof of this is included as Section B in the supplementary data.

LEMMA 6

There is an AUC formula α such that $\models \chi \leftrightarrow \alpha$.

PROOF. The cases $\delta_1, \dots, \delta_6$ are exhaustive, so $\models \chi \leftrightarrow ((\delta_1 \rightarrow \chi) \wedge \dots \wedge (\delta_6 \rightarrow \chi))$. Then by Lemma 5 there are AUC formulas $\alpha_1, \dots, \alpha_6$ such that $\models \chi \leftrightarrow ((\delta_1 \rightarrow \alpha_1) \wedge \dots \wedge (\delta_6 \rightarrow \alpha_6))$. Furthermore, by Lemma 4 there are AUC formulas $\beta_1, \dots, \beta_6$ such that $\models \chi \leftrightarrow ((\beta_1 \rightarrow \alpha_1) \wedge \dots \wedge (\beta_6 \rightarrow \alpha_6))$. This proves the lemma with $\alpha = (\beta_1 \rightarrow \alpha_1) \wedge \dots \wedge (\beta_6 \rightarrow \alpha_6)$. ■

The formula χ was taken $\chi = \{U\}^* \varphi$ with any update U containing only AUC formulas and any AUC formula φ . Lemma 6 therefore allows us to eliminate all occurrences of operators $\{U\}^*$ in any AU^* formula by first eliminating the innermost occurrences and working outward. We therefore have the following corollary.

THEOREM 1

$AUC \equiv AU^*$

PROOF. It was shown in [11] that $AUU^* \equiv AU^*$, which together with Lemma 1 shows that $AUC \preceq AU^*$. Furthermore, it follows from Lemma 6 that $AU^* \preceq AUC$. We therefore have $AUC \equiv AU^*$. ■

7 Conclusion

We have demonstrated two new expressivity results about logics using arrow updates. The first result is that the logic $\mathcal{L}_{AU^*}$ using arrow common knowledge is strictly more expressive than the logic $\mathcal{L}_{AR}$ using relativized common knowledge. This result is not surprising and had in fact been predicted in [11], where AU^* was introduced.

The second result is that the logic $\mathcal{L}_{AUC}$ using arrow updates and common knowledge is as expressive as $\mathcal{L}_{AU^*}$. This result is rather surprising considering that the logic $\mathcal{L}_{AR}$ was shown in [6] to be strictly more expressive than the logic $\mathcal{L}_{APC}$ using public announcements and common knowledge, and the difference between $\mathcal{L}_{AUC}$ and $\mathcal{L}_{AU^*}$ is comparable to the difference between $\mathcal{L}_{APC}$ and $\mathcal{L}_{AR}$.

These two new results together with results from [13], [4], [12], [6] and [11] fully determine the expressivity landscape of all logics using any combination of common knowledge, relativized common knowledge, public announcements, arrow updates and arrow common knowledge. This expressivity landscape is shown in Figure 2.

One interesting property of this landscape is that there are no logics in it with incomparable expressivity; if $\mathcal{L}_X$ and $\mathcal{L}_Y$ are logics of the kind under consideration then either $\mathcal{L}_X \preceq \mathcal{L}_Y$ or $\mathcal{L}_Y \preceq \mathcal{L}_X$.

A remaining open question is the succinctness of the different logics. In particular, the translation from $\mathcal{L}_{AU^*}$ to $\mathcal{L}_{AUC}$ demonstrated in this article has an extremely high growth in formula size. Whether this is necessarily so or there is an efficient translation is not currently known.

Acknowledgements

I would like to thank Barteld Kooi, Allard Tamminga and two anonymous referees for carefully reading an earlier version of this article (including the highly technical supplementary data) and for providing very helpful feedback.

Supplementary Data

Supplementary Data are available at *LOGCOM* online.

References

- [1] G. Aucher, P. Balbiani, L. Fariñas del Cerro, and A. Herzig. Global and local graph modifiers. In *Proceedings of the 5th Workshop on Methods for Modalities (M4M5 2007)*, C. Areces and S. Demri, eds, pp. 293–307. École Normale Supérieure de Cachan, Cachan, France, 2009.
- [2] A. Baltag. A logic of epistemic actions. In *Proceedings of ESSLLI'99 Workshop on Foundations and Applications of Collective Agent Based Systems*, W. van der Hoek, J.-J.C. Meyer, and C. Witteveen, eds, Utrecht University, Utrecht, The Netherlands, 1999.
- [3] A. Baltag and L. S. Moss. Logics for epistemic programs. *Synthese*, **139**, 165–224, 2004.
- [4] A. Baltag, L. S. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. Technical Report TR534, Indiana University Bloomington, 1999.
- [5] J. van Benthem, J. Gerbrandy, T. Hoshi, and E. Pacuit. Merging frameworks for interaction. *Journal of Philosophical Logic*, **38**, 491–526, 2009.
- [6] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, **204**, 1620–1662, 2006.
- [7] H. van Ditmarsch. *Knowledge Games*. PhD thesis, University of Groningen, The Netherlands, 2000.
- [8] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, 2004.
- [9] J. Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, University of Amsterdam, The Netherlands, 1999.
- [10] B. Kooi. Expressivity and completeness for public update logics via reduction axioms. *Journal of Applied Non-Classical Logics*, **17**, 231–253, 2007.
- [11] B. Kooi and B. Renne. Arrow update logic. *The Review of Symbolic Logic*, **4**, 536–559, 2011.
- [12] B. Kooi and J. van Benthem. Reduction axioms for epistemic actions. In *AiML-2004: Advances in Modal Logic*, volume UMCS-04-9-1 of *Technical Report Series*, R. Schmidt, I. Pratt-Hartmann, M. R., and H. Wansing, eds, pp. 197–211. University of Manchester, 2004.
- [13] J. Plaza. Logics of public communication. In *Proceedings of the Fourth International Symposium on Methodologies for Intelligent Systems, Poster Session Program*, M.L. Emrich, M.S. Phifer, M. Hadzikadic, and Z.W. Ras, eds, pp. 201–216. Oak Ridge National Laboratory, ORNL/DSRD-24, 1989. Reprinted as [14].
- [14] J. Plaza. Logics of public communication. *Synthese*, **158**, 165–179, 2007.

Received 29 June 2014