

How Arbitrary Are Arbitrary Public Announcements?

Louwe Bouke Kuijer

University of Groningen, The Netherlands

Abstract. Public announcements are used in dynamic epistemic logic to model certain kinds of information change. A formula $\langle\psi\rangle\varphi$ represents the statement that after ψ is publicly announced φ will be the case.

Sometimes we want to reason about whether it is possible for φ to become true after some announcement. In order to do this an arbitrary public announcement operator \diamond can be added to an epistemic logic with public announcements. Ideally a formula $\diamond\varphi$ would hold if and only if there is a formula ψ such that $\langle\psi\rangle\varphi$. However, in order to avoid circularity the \diamond operator can only quantify over those ψ that are \diamond -free. So $\diamond\varphi$ holds if and only if there is a \diamond -free ψ such that $\langle\psi\rangle\varphi$.

As a result it does not follow immediately from the definition that $\langle\psi\rangle\varphi$ implies $\diamond\varphi$ if ψ contains a \diamond . But the implication may still hold in some cases. In this paper I show that on finite models $\langle\psi\rangle\varphi$ implies $\diamond\varphi$ for every ψ , and that on finitely branching models $\langle\psi\rangle\varphi$ implies $\diamond\varphi$ for every ψ if φ is \diamond -free. Finally I also show that there are φ and ψ such that $\langle\psi\rangle\varphi$ does not imply $\diamond\varphi$ even on a finitely branching model.

1 Introduction

In epistemic logic we can reason about basic facts (represented by propositional variables) and about knowledge of different agents (represented by one operator K_a per agent). A commonly used example in epistemic logic is that of a simple card game. Suppose two agents a and b are playing a game where they each hold one card, and they know their own card but not the other's card. Then if a holds a queen (and we use the propositional variable q to represent this basic fact) the formulas (i) $K_a q$, (ii) $q \wedge \neg K_b q$ and (iii) $K_a \neg K_b q$ represent the (true) statements that (i) a knows that she holds a queen, (ii) a holds a queen but b does not know this, and (iii) a knows that b does not know that she holds a queen.

In such a basic epistemic logic we cannot however express information change. For example, we cannot reason about what would happen if a were to show her card to b in basic epistemic logic. If we want to reason about information change we need to use a *dynamic* epistemic logic. There are many different kinds of dynamic epistemic logic, see for example [1] for an overview. One of the most common ways to turn a (static) epistemic logic into a dynamic epistemic logic is to add *public announcements* [2,3] to the logic. A public announcement is a binary operator of the form $\langle\psi\rangle\varphi$. The formula $\langle\psi\rangle\varphi$ is true if φ will hold after ψ is announced *truthfully* and *publicly*.

Using public announcements we can reason about what would happen if a were to show her card to b ; the showing of a card can be considered an announcement of the card that a holds. The statement that after a shows her card b knows what card a holds is therefore represented by the (true) formula $\langle q \rangle K_b q$. One thing to note about the formula $\langle q \rangle K_b q$ is that after q is announced agent b knows that q , so the announcing of q is a way for b to get to know q .

However, not all formulas can be learned in such a way. Consider the formula $q \wedge \neg K_b q$, representing a holding a queen and b not knowing this. This formula was introduced in [4] as a formula that can never be known by b even if it is true. Since $q \wedge \neg K_b q$ can never be known by b there is also no announcement such that b will know $q \wedge \neg K_b q$ after the announcement. So not only is it impossible for b to get to know the truth of $q \wedge \neg K_b q$ by announcing $q \wedge \neg K_b q$, there is no formula ψ such that $\langle \psi \rangle K_b (q \wedge \neg K_b q)$.

This last property, whether for a given φ there exists a ψ such that $\langle \psi \rangle K_b \varphi$ requires us to quantify over all formulas. We can of course do this quantification meta-logically, but epistemic logic with public announcements does not allow us to perform this quantification inside the logic. This is unfortunate, as this means we cannot use public announcements to reason about whether it is possible to get to know something. A solution proposed in [5,6] is to add one more operator \diamond , representing *arbitrary public announcements*.

Such arbitrary public announcements can be useful when considering problems of knowability, but also in more practical scenarios such as in cryptography where it is important to know whether it is possible to make a public statement such that agent b learns the content p of a message but another agent e does not, so whether $\diamond(K_b p \wedge \neg K_e p)$.

We would like to define \diamond in such a way that $\diamond\varphi$ holds if and only if there is an announcement ψ such that $\langle \psi \rangle \varphi$ holds. There is a technical problem with this kind of definition, however. If we allow the announcement ψ to be *any* formula the evaluation of $\diamond\varphi$ would become circular. After all, in order to know whether $\diamond\varphi$ holds we would have to check whether $\langle \diamond\varphi \rangle \varphi$ holds. But in order to know whether $\langle \diamond\varphi \rangle \varphi$ holds we would among other things have to know whether $\diamond\varphi$ is a truthful announcement, so whether $\diamond\varphi$ holds.

This circularity is removed in [5,6] by restricting ψ to formulas that do not themselves contain \diamond operators. So $\diamond\varphi$ holds if and only if there is a \diamond -free formula ψ such that $\langle \psi \rangle \varphi$. Unfortunately this means that the announcements in an arbitrary announcement operator are not in fact entirely arbitrary. But while the definition of \diamond cannot allow completely arbitrary announcements it might be possible to get entirely arbitrary announcements as an “emergent property”. Suppose that whenever there is a ψ containing \diamond such that $\langle \psi \rangle \varphi$ there would always also be a ψ' that is \diamond -free such that $\langle \psi' \rangle \varphi$. Then $\langle \psi \rangle \varphi$ would imply $\diamond\varphi$, even if ψ happens to contain a \diamond .

A different way of phrasing this is to ask whether $\langle \psi \rangle \varphi \rightarrow \diamond\varphi$ is valid for every ψ . It was shown in [5] that the implication is valid if there is only a single agent. In this paper I show that if there are multiple agents the validity of the implication depends on the class of models we use to evaluate the logic on and

on φ . If we use only finite models then $\langle \psi \rangle \varphi \rightarrow \diamond \varphi$ is valid. If we allow finitely branching infinite models then $\langle \psi \rangle \varphi \rightarrow \diamond \varphi$ is valid for every ψ and every \diamond -free φ . But if we allow models that are infinitely branching or if we do not restrict to \diamond -free φ then there are φ and ψ such that $\langle \psi \rangle \varphi \rightarrow \diamond \varphi$ is not valid.

In Section 2 I give some definitions needed to formulate and prove the results. Then in Section 3 I show that for finite models $\langle \psi \rangle \varphi \rightarrow \diamond \varphi$ is valid. In Section 4.1 I prove that for finitely branching models $\langle \psi \rangle \varphi \rightarrow \diamond \varphi$ is valid if φ is \diamond -free. In Section 4.2 I construct ψ and \diamond -free φ such that $\langle \psi \rangle \varphi \rightarrow \diamond \varphi$ is not valid on infinitely branching models. Finally, in Section 4.3 I construct φ and ψ containing \diamond such that $\langle \psi \rangle \varphi \rightarrow \diamond \varphi$ is not valid on finitely branching models.

2 Definitions

Let us start by defining *arbitrary public announcement logic* $\mathcal{L}_{\text{APAL}}$ and the \diamond -free fragment *public announcement logic* \mathcal{L}_{PAL} of $\mathcal{L}_{\text{APAL}}$. Let us fix a countably infinite set \mathcal{P} of propositional variables and a finite set \mathcal{A} of agents. The language of $\mathcal{L}_{\text{APAL}}$ is then defined as follows.

Definition 1. *The formulas of $\mathcal{L}_{\text{APAL}}$ are given by*

$$\varphi ::= p \mid \neg \varphi \mid (\varphi \vee \varphi) \mid K_a \varphi \mid \langle \psi \rangle \varphi \mid \diamond \varphi$$

where p ranges over \mathcal{P} and a ranges over \mathcal{A} .

Definition 2. *The logic \mathcal{L}_{PAL} is the \diamond -free fragment of $\mathcal{L}_{\text{APAL}}$.*

Parentheses are omitted where this should not cause confusion and \wedge , \rightarrow , \leftrightarrow , \vee and \bigwedge are used in the usual way as abbreviations. Furthermore, \hat{K}_a , $[\varphi]$ and \square are used as abbreviations for $\neg K_a \neg$, $\neg \langle \varphi \rangle \neg$ and $\neg \diamond \neg$ respectively. Integer superscripts are used to indicate multiple copies of an operator, so K_a^3 stands for $K_a K_a K_a$. Finally, if B is a set of agents then K_B stands for $\bigwedge_{a \in B} K_a$ and \hat{K}_B for $\bigvee_{a \in B} K_a$.

The intended reading of the non-boolean operators is as follows:

- $K_a \varphi$ is read as “agent a knows that φ ”,
- $\langle \psi \rangle \varphi$ is read as “after it is publicly announced that ψ is the case φ holds”,
- $\diamond \varphi$ is read as “there is a \diamond -free announcement ψ such that $\langle \psi \rangle \varphi$ holds”.

Since $\mathcal{L}_{\text{APAL}}$ and \mathcal{L}_{PAL} are epistemic logics they are usually considered over the class of S5 models. We will follow this tradition, but it should be noted that none of the proofs in this paper depend on the special properties of S5 models. So all the results presented here also hold over the class of K models.

Definition 3. *A model \mathcal{M} is a triple (W, R, v) where W is a set of worlds, $R : \mathcal{A} \rightarrow \wp(W \times W)$ assigns to each agent an equivalence relation on W and $v : \mathcal{P} \rightarrow \wp(W)$ is a valuation function that assigns an extension to each propositional variable.*

A model $\mathcal{M} = (W, R, v)$ is said to be finitely branching if for each $w \in W$ and each $a \in \mathcal{A}$ the set $\{w' \mid (w, w') \in R(a)\}$ is finite. A model $\mathcal{M} = (W, R, v)$ is said to be finite if W is a finite set.

The semantics for most operators of $\mathcal{L}_{\text{APAL}}$ are as usual. For the only unusual operator \diamond it should be noted that it quantifies not over the formulas of $\mathcal{L}_{\text{APAL}}$ but over the formulas of \mathcal{L}_{PAL} .

Definition 4. Given a model $\mathcal{M} = (W, R, v)$, a world w of \mathcal{M} and φ, ψ formulas of $\mathcal{L}_{\text{APAL}}$ the satisfaction relation \models is given by

$$\begin{aligned} \mathcal{M}, w &\models p && \Leftrightarrow w \in v(p) \\ \mathcal{M}, w &\models \neg\varphi && \Leftrightarrow \mathcal{M}, w \not\models \varphi \\ \mathcal{M}, w &\models \varphi \vee \psi && \Leftrightarrow \mathcal{M}, w \models \varphi \text{ or } \mathcal{M}, w \models \psi \\ \mathcal{M}, w &\models K_a\varphi && \Leftrightarrow \mathcal{M}, w' \models \varphi \text{ for all } w' \in W \text{ such that } (w, w') \in R(a) \\ \mathcal{M}, w &\models \langle \varphi \rangle \psi && \Leftrightarrow \mathcal{M}, w \models \varphi \text{ and } \mathcal{M}_\varphi, w \models \psi \\ \mathcal{M}, w &\models \diamond\varphi && \Leftrightarrow \text{there is a } \mathcal{L}_{\text{PAL}} \text{ formula } \psi \text{ such that } \mathcal{M}, w \models \langle \psi \rangle \varphi \end{aligned}$$

with $\mathcal{M}_\varphi = (W_\varphi, R_\varphi, v_\varphi)$ where $W_\varphi = \{w \in W \mid \mathcal{M}, w \models \varphi\}$ and R_φ and v_φ are the restrictions of R and v to W_φ .

We write $\mathcal{M} \models \varphi$ if $\mathcal{M}, w \models \varphi$ for every $w \in W$ and $\models \varphi$ if $\mathcal{M} \models \varphi$ for every model \mathcal{M} . Furthermore, we write $\models_{\text{br}} \varphi$ if $\mathcal{M} \models \varphi$ for every finitely branching model \mathcal{M} and $\models_{\text{fin}} \varphi$ if $\mathcal{M} \models \varphi$ for every finite model \mathcal{M} .

3 APAL on Finite Models

With the definitions out of the way I can show that $\models_{\text{fin}} \langle \psi \rangle \varphi \rightarrow \diamond\varphi$ for all $\mathcal{L}_{\text{APAL}}$ formulas ψ . This is not a very surprising result; in a finite model we can replace any \diamond operator by the announcement of a disjunction of \mathcal{L}_{PAL} formulas, one for each world where the \diamond is replaced by the ‘‘chosen announcement’’ for that world.

Lemma 1. Fix a finite model $\mathcal{M} = (W, R, v)$ and a $\mathcal{L}_{\text{APAL}}$ formula φ . Then there is a \mathcal{L}_{PAL} formula ψ such that $\mathcal{M} \models \varphi \leftrightarrow \psi$.

Proof. By induction on the construction of φ . The lemma trivially holds if φ is atomic. Suppose then as induction hypothesis that φ is not atomic, and that the lemma holds for all finite models and all subformulas of φ .

The formula φ is not atomic, so it is of one of the following forms:

1. $\varphi = \neg\varphi'$,
2. $\varphi = \varphi' \vee \varphi''$,
3. $\varphi = K_a\varphi'$,
4. $\varphi = \langle \varphi'' \rangle \varphi'$ or
5. $\varphi = \diamond\varphi'$.

By the induction hypothesis there is a \mathcal{L}_{PAL} formula ψ' such that $\mathcal{M} \models \varphi' \leftrightarrow \psi'$ and, if applicable, a \mathcal{L}_{PAL} formula ψ'' such that $\mathcal{M} \models \varphi'' \leftrightarrow \psi''$. So if we take ψ to be $\neg\psi'$, $\psi' \vee \psi''$ or $K_a\psi'$ then we have $\mathcal{M} \models \varphi \leftrightarrow \psi$ in the first, second or third case respectively.

Let us then consider fourth case. By the induction hypothesis there are \mathcal{L}_{PAL} formulas ψ'' such that $\mathcal{M} \models \varphi'' \leftrightarrow \psi''$ and ψ' such that $\mathcal{M}_{\varphi''} \models \varphi' \leftrightarrow \psi'$. This implies that $\mathcal{M} \models \varphi \leftrightarrow \langle \psi'' \rangle \psi'$.

Let us then consider the fifth case, $\varphi = \Diamond\varphi'$. Let W' be the extension of φ , so $W' := \{w \in W \mid \mathcal{M}, w \models \Diamond\varphi'\}$. For each $w_i \in W'$ we have $\mathcal{M}, w_i \models \Diamond\varphi'$, so there is a \mathcal{L}_{PAL} formula φ''_i such that $\mathcal{M}, w_i \models \langle\varphi''_i\rangle\varphi'$. By the induction hypothesis there is a \mathcal{L}_{PAL} formula ψ'_i such that $\mathcal{M}_{\varphi''_i} \models \psi'_i \leftrightarrow \varphi'$. We therefore have $\mathcal{M} \models \langle\varphi''_i\rangle\psi'_i \leftrightarrow \langle\varphi''_i\rangle\varphi'$.

Now let $\psi := \bigvee_{w_i \in W'} \langle\varphi''_i\rangle\psi'_i$. This is a \mathcal{L}_{PAL} formula, since all its subformulas are \mathcal{L}_{PAL} formulas and W' is a finite set. Furthermore, for each $w_i \in W'$ we have $\mathcal{M}, w_i \models \psi$.

Suppose now towards a contradiction that for some $w' \in W \setminus W'$ we have $\mathcal{M}, w' \models \psi$. Then one of the disjuncts of ψ holds in w' , so for some $w_i \in W'$ we have $\mathcal{M}, w' \models \langle\varphi''_i\rangle\psi'_i$. Then we also have $\mathcal{M}, w' \models \langle\varphi''_i\rangle\varphi'$, since $\mathcal{M} \models \langle\varphi''_i\rangle\psi'_i \leftrightarrow \langle\varphi''_i\rangle\varphi'$. But φ''_i is a \mathcal{L}_{PAL} formula so this implies that $\mathcal{M}, w' \models \Diamond\varphi'$. This contradicts w' being an element of $W \setminus W'$, so we must have $\mathcal{M}, w' \not\models \psi$.

This shows that $\mathcal{M} \models \varphi \leftrightarrow \psi$, which completes the induction step and thereby the proof. \square

It now follows immediately that $\models_{\text{fin}} \langle\psi\rangle\varphi \rightarrow \Diamond\varphi$.

Theorem 1. *For every $\mathcal{L}_{\text{APAL}}$ formulas φ and χ we have $\models_{\text{fin}} \langle\varphi\rangle\chi \rightarrow \Diamond\chi$.*

Proof. Fix any $\mathcal{L}_{\text{APAL}}$ formulas φ and χ , and any finite model \mathcal{M} . Then by Lemma 1 there is a \mathcal{L}_{PAL} formula ψ such that $\mathcal{M} \models \varphi \leftrightarrow \psi$. This implies that $\mathcal{M} \models \langle\varphi\rangle\chi \leftrightarrow \langle\psi\rangle\chi$. But ψ is a \mathcal{L}_{PAL} formula so $\models \langle\psi\rangle\chi \rightarrow \Diamond\chi$ and therefore $\mathcal{M} \models \langle\varphi\rangle\chi \rightarrow \Diamond\chi$. Since this is true for any finite model \mathcal{M} this implies that $\models_{\text{fin}} \langle\varphi\rangle\chi \rightarrow \Diamond\chi$. \square

4 APAL on Infinite Models

On infinite models we cannot use the method that worked for finite models, $\bigvee_{w_i \in W'} \langle\varphi''_i\rangle\psi'_i$ is in general not a formula on infinite models since W' may be infinite. Here I show that no other method can work; there are infinite models \mathcal{M} , worlds w of \mathcal{M} and $\mathcal{L}_{\text{APAL}}$ formulas φ and ψ such that $\mathcal{M}, w \models \langle\psi\rangle\varphi \wedge \neg\Diamond\varphi$.

Like the result for the finite case this should not surprise us. What is somewhat surprising however is that the result extends to finitely branching models; there are φ and ψ such that $\not\models_{\text{br}} \langle\psi\rangle\varphi \rightarrow \Diamond\varphi$. To see why it is unexpected that $\not\models_{\text{br}} \langle\psi\rangle\varphi \rightarrow \Diamond\varphi$ consider the following. Fix any finitely branching model \mathcal{M} and any world w of \mathcal{M} . We cannot guarantee the existence of a \mathcal{L}_{PAL} formula ψ' such that $\mathcal{M} \models \psi \leftrightarrow \psi'$, but since \mathcal{M} is finitely branching we can for any $n \in \mathbb{N}$ guarantee the existence of a \mathcal{L}_{PAL} formula ψ'' such that $\mathcal{M}, w' \models \psi \leftrightarrow \psi''$ for every world w' that is reachable within n steps from w .

The language of $\mathcal{L}_{\text{APAL}}$ does not contain common knowledge, so it would at first glance seem like such a ψ'' that is equivalent to ψ up to a given distance might be sufficient to make φ have the same value after both announcements. If φ does not contain any \Diamond operators then this does indeed work, for any $\mathcal{L}_{\text{APAL}}$ formula ψ and any \mathcal{L}_{PAL} formula φ we have $\models_{\text{br}} \langle\psi\rangle\varphi \rightarrow \Diamond\varphi$. But a \Diamond operator (or more precisely: a \square operator) can make a formula depend on worlds that are

arbitrarily far away in such a way that in certain models no finite approximation ψ'' of ψ will suffice.

I first show that for \diamond -free φ we have $\models_{\text{br}} \langle \psi \rangle \varphi \rightarrow \diamond \varphi$, then that there are ψ and \diamond -free φ such that $\not\models \langle \psi \rangle \varphi \rightarrow \diamond \varphi$ and finally that for some φ that do contain \diamond we have $\not\models_{\text{br}} \langle \psi \rangle \varphi \rightarrow \diamond \varphi$. This order of proofs is chosen for reasons of clarity of exposition; the proof that $\not\models_{\text{br}} \langle \psi \rangle \varphi \rightarrow \diamond \varphi$ uses more complicated variants on some of the same techniques that are used in the proof of $\not\models \langle \psi \rangle \varphi \rightarrow \diamond \varphi$.

4.1 Validity of $\langle \psi \rangle \varphi \rightarrow \diamond \varphi$ for \diamond -free φ

Before proving that $\models_{\text{br}} \langle \psi \rangle \varphi \rightarrow \diamond \varphi$ we need one auxiliary lemma.

Lemma 2. *Let \mathcal{M} be any finitely branching model and w_1, w_2 two worlds of \mathcal{M} . Then there is a $\mathcal{L}_{\text{APAL}}$ formula that distinguishes between \mathcal{M}, w_1 and \mathcal{M}, w_2 if and only if there is a \mathcal{L}_{PAL} formula that distinguishes between them.*

Proof. If there is a \mathcal{L}_{PAL} formula ψ' that distinguishes between two worlds then there is also a $\mathcal{L}_{\text{APAL}}$ formula ψ that distinguishes between the two worlds, namely $\psi = \psi'$. Left to show is that if $\mathcal{L}_{\text{APAL}}$ can distinguish between two worlds then so can \mathcal{L}_{PAL} .

The formulas of $\mathcal{L}_{\text{APAL}}$ are invariant under bisimulation (see [6]), so if a $\mathcal{L}_{\text{APAL}}$ formula distinguishes between \mathcal{M}, w_1 and \mathcal{M}, w_2 then \mathcal{M}, w_1 and \mathcal{M}, w_2 are not bisimilar. On finitely branching models worlds are bisimilar if and only if they are indistinguishable by basic modal logic (see for example [1]). So since \mathcal{M}, w_1 and \mathcal{M}, w_2 are not bisimilar they can be distinguished by a \mathcal{L}_{PAL} formula. \square

Lemma 2 also holds for models that are not finitely branching, but that requires a more complicated proof and we only need the result for finitely branching models.¹

Lemma 3. *Let ψ be any $\mathcal{L}_{\text{APAL}}$ formula and let φ be any \mathcal{L}_{PAL} formula. Then $\models_{\text{br}} \langle \psi \rangle \varphi \rightarrow \diamond \varphi$.*

Proof. Fix any finitely branching model \mathcal{M} and any world w of \mathcal{M} . It was shown in [2] that every \mathcal{L}_{PAL} formula is equivalent to a \mathcal{L}_{PAL} formula that does not contain any public announcements. Let φ' be the announcement-free formula equivalent to φ , and let n be the maximum nesting depth of K operators in φ' . Then the truth of φ' —and therefore also φ —on \mathcal{M}, w does not depend on changes to worlds that are not reachable from w in at most n steps.

Let W' be the set of worlds that are reachable from w in at most n steps, and let $W_1 := \{w' \in W' \mid \mathcal{M}, w' \models \psi\}$ and $W_2 := W' \setminus W_1$. Then for each $w_i \in W_1$ and $w_j \in W_2$ the formula ψ distinguishes \mathcal{M}, w_i from \mathcal{M}, w_j , so by Lemma 2 there is also a \mathcal{L}_{PAL} formula that distinguishes the two worlds. Let

¹ For an idea of why Lemma 2 also holds for infinitely branching models consider the case where $\mathcal{M}, w \models \diamond \varphi$ and $\mathcal{M}, w' \not\models \diamond \varphi$. Then there is a ψ such that $\mathcal{M}, w \models \langle \psi \rangle \varphi$ and in particular $\mathcal{M}, w' \not\models \langle \psi \rangle \varphi$ so the formula $\langle \psi \rangle \varphi$ distinguishes the two worlds as well. This can be extended to any formula containing a \diamond operator.

$\psi'_{i,j}$ be this distinguishing \mathcal{L}_{PAL} formula and assume without loss of generality that $\mathcal{M}, w_i \models \psi'_{i,j}$ and $\mathcal{M}, w_j \not\models \psi'_{i,j}$.

For $w_i \in W_1$ let $\psi'_i := \bigwedge_{w_j \in W_2} \psi'_{i,j}$. Then $\mathcal{M}, w_i \models \psi'_i$ and $\mathcal{M}, w_j \not\models \psi'_i$ for each $w_j \in W_2$. Finally, let $\psi' := \bigvee_{w_i \in W_1} \psi'_i$. This ψ' satisfies $\mathcal{M}, w_i \models \psi'$ for each $w_i \in W_1$ and $\mathcal{M}, w_j \not\models \psi'$ for each $w_j \in W_2$.

As such, the models \mathcal{M}_ψ and $\mathcal{M}_{\psi'}$ only differ in worlds that are not reachable from w within n steps, so $\mathcal{M}, w \models \langle \psi \rangle \varphi \leftrightarrow \langle \psi' \rangle \varphi$. Because ψ' is a \mathcal{L}_{PAL} formula this implies that $\mathcal{M}, w \models \langle \psi \rangle \varphi \rightarrow \diamond \varphi$. The model \mathcal{M} and world w were chosen as any finitely branching model and any world of that model, so we have $\models_{\text{br}} \langle \psi \rangle \varphi \rightarrow \diamond \varphi$. \square

4.2 Invalidity of $\langle \psi \rangle \varphi \rightarrow \diamond \varphi$ on Infinitely Branching Models

If we do not restrict ourselves to finite or finitely branching models there are φ and ψ such that $\langle \psi \rangle \varphi \rightarrow \diamond \varphi$ is not valid. Let

$$\begin{aligned} \varphi_1 &:= \hat{K}_{cp} \wedge K_c(r \rightarrow \hat{K}_d \neg r) \wedge K_c((p \wedge \neg r) \rightarrow \hat{K}_e r), \\ \varphi_2 &:= K_c(\neg q \rightarrow (\hat{K}_f(\neg \hat{K}_c q \wedge K_a p) \wedge \hat{K}_f \neg K_a p)), \\ \varphi &:= \varphi_1 \wedge \varphi_2 \\ \psi &:= p \vee q \vee K_a \neg \diamond (\hat{K}_b K_a p \wedge \hat{K}_b \neg K_a p). \end{aligned}$$

Furthermore, let \mathcal{M} be the model shown in Figure 1 and let \mathcal{M}_n for $n \in \mathbb{N}$ be the submodels indicated in Figure 1.

We want to show that $\mathcal{M}, w \not\models \langle \psi \rangle \varphi \rightarrow \diamond \varphi$. This requires us to show that $\mathcal{M}, w \models \langle \psi \rangle \varphi$ and that $\mathcal{M}, w \not\models \diamond \varphi$. In order to prove that $\mathcal{M}, w \not\models \diamond \varphi$ we have to demonstrate that if $\mathcal{M}, w \models \langle \psi' \rangle \varphi$ then ψ' contains a \diamond operator. The subformula φ_1 is constructed in such a way that if $\mathcal{M}, w \models \langle \psi' \rangle \varphi$ then the update $\langle \psi' \rangle$ retains an infinite number of worlds. The subformula φ_2 guarantees that if $\mathcal{M}, w \models \langle \psi' \rangle \varphi$ and $\langle \psi' \rangle$ retains an infinite number of worlds then ψ' must perform an infinite number of different updates, which cannot be done without a \diamond operator. But before looking at the details of the proof that $\mathcal{M}, w \not\models \diamond \varphi$ let us start by proving the simpler part of the statement, namely that $\mathcal{M}, w \models \langle \psi \rangle \varphi$.

Lemma 4. *We have $\mathcal{M}, w \models \langle \psi \rangle \varphi$.*

Proof. To show is that $\mathcal{M}_\psi \models \varphi$, so let us look at which worlds are retained by $\langle \psi \rangle$. The disjuncts p and q of ψ guarantee that any world in the leftmost three columns is retained.

The worlds in the fourth column from the left satisfy neither p nor q though, so they are retained only if they satisfy $K_a \neg \diamond (\hat{K}_b K_a p \wedge \hat{K}_b \neg K_a p)$. These worlds themselves always satisfy $\neg \diamond (\hat{K}_b K_a p \wedge \hat{K}_b \neg K_a p)$; there is no update that would let them satisfy $\hat{K}_b K_a p$ because every b -reachable world satisfies $\neg p$.

So the worlds in the fourth column are retained if and only if the p world to the left of them (which they are a -connected with) satisfies $\neg \diamond (\hat{K}_b K_a p \wedge \hat{K}_b \neg K_a p)$.

Now we reach the difference between the rows of a submodel \mathcal{M}_n . Consider the p world in the bottom row of \mathcal{M}_n for any n . The only world b -reachable from

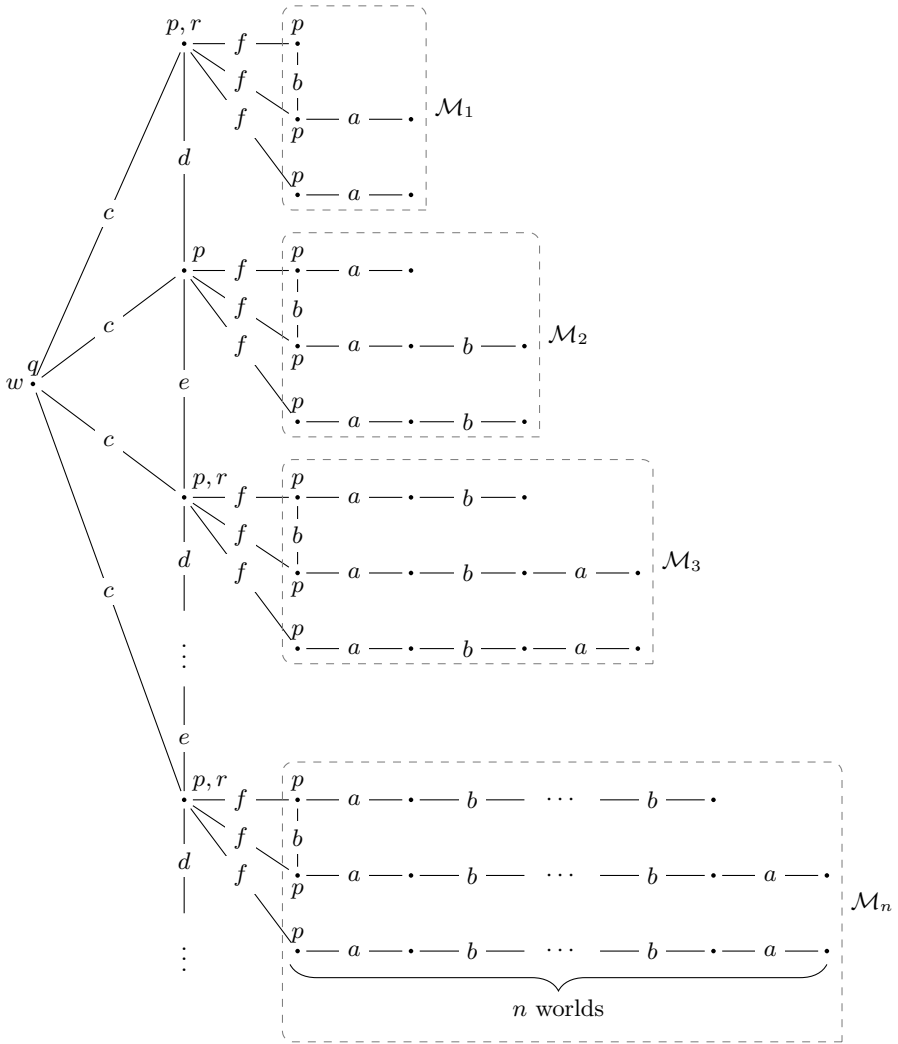


Fig. 1. The model \mathcal{M} . Some accessibility arrows are not drawn.

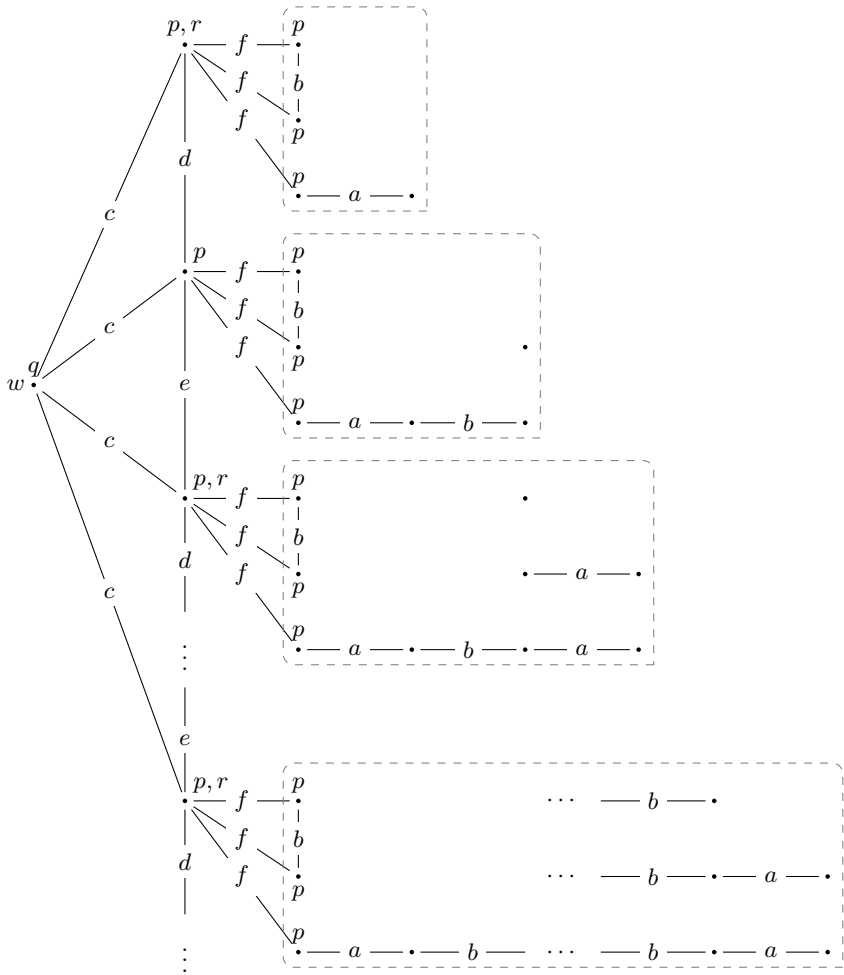


Fig. 2. The model \mathcal{M}_ψ . Some accessibility arrows are not drawn.

this world is itself, so there is no update that can make the world satisfy $\hat{K}_b K_{ap} \wedge \hat{K}_b \neg K_{ap}$. So the p world in the bottom row satisfies $\neg \diamond (\hat{K}_b K_{ap} \wedge \hat{K}_b \neg K_{ap})$.

Now consider one of the p worlds in the top two rows of \mathcal{M}_n . These two worlds can be distinguished from each other because their “tails” are of different lengths. This allows us to create an update χ_n that removes the $\neg p$ world adjacent to the top p world but not the one adjacent to the second row p world. The formula $\chi_n := \neg p \rightarrow \hat{K}_{\{a,b\}}^{n-1} K_{\{a,b\}}^n \neg p$ for example does this.

The specific formula χ_n that works for a submodel \mathcal{M}_n depends on n , but in every case it is a PAL formula so for every n the top two p worlds of \mathcal{M}_n satisfy $\diamond (\hat{K}_b K_{ap} \wedge \hat{K}_b \neg K_{ap})$.

This means that the worlds in the fourth column are retained by $\langle \psi \rangle$ if and only if they are in the third row of any submodel \mathcal{M}_n . The model \mathcal{M}_ψ is therefore as shown in Figure 2. It is straightforward to verify that w satisfies φ in that model. \square

Now to show that there is no PAL formula ψ' that satisfies $\mathcal{M}, w \models \langle \psi' \rangle \varphi$. Recall that the two parts of φ have different purposes. The part φ_1 guarantees that ψ' retains an infinite number of worlds while φ_2 guarantees that ψ' performs an infinite number of different updates, which cannot be done without using a \diamond operator.

Lemma 5. *For every \mathcal{L}_{PAL} formula ψ' we have $\mathcal{M}, w \not\models \langle \psi' \rangle \varphi$.*

Proof. Suppose towards a contradiction that there is a \mathcal{L}_{PAL} formula ψ' such that $\mathcal{M}, w \models \langle \psi' \rangle \varphi$. Then we have $\mathcal{M}, w \models \langle \psi' \rangle \varphi_1$ and $\mathcal{M}, w \models \langle \psi' \rangle \varphi_2$.

Consider $\mathcal{M}, w \models \langle \psi' \rangle \varphi_1$. The conjunct $\hat{K}_c p$ guarantees that $\langle \psi' \rangle$ retains at least one of the p worlds that are accessible from w , so at least one of the worlds in the second column.

The worlds in the second column alternate between r and $\neg r$, and the arrows between those worlds alternate between d and e . As a result the conjunct $K_c (r \rightarrow \hat{K}_d \neg r)$ implies that if ψ' retains an r world in the second column then it also retains the $\neg r$ world below it. Likewise, the conjunct $K_c ((p \wedge \neg r) \rightarrow \hat{K}_e r)$ implies that if ψ' retains a $\neg r$ world in the second column then it also retains the r world below it.

So the three conjuncts of φ_1 together imply that ψ' retains at least one of the worlds in the second column as well as all worlds below it.

Consider then $\mathcal{M}, w \models \langle \psi' \rangle \varphi_2$. The formula φ_2 says something about all c -reachable worlds that do not satisfy q , so all worlds in the second column (that are retained by $\langle \psi' \rangle$). Of these worlds it says that they can reach two worlds by using f , one world satisfying $\neg \hat{K}_c q \wedge K_{ap}$ and one satisfying $\neg K_{ap}$.

The worlds in the first two columns all satisfy $\hat{K}_c q$ and K_{ap} so these two f -reachable worlds must be in the third column. If the n -th world of the second column is retained by $\langle \psi' \rangle$ there must therefore be two p worlds retained in \mathcal{M}_n . Furthermore, one of those worlds in \mathcal{M}_n must be adjacent to a $\neg p$ world that is retained while the other must not be adjacent to a retained $\neg p$ world.

One of the $\neg p$ worlds in the second column of \mathcal{M}_n (so the fourth column of \mathcal{M}) must be retained and one must not be retained, so in particular ψ' must

distinguish between two of those worlds. But the only way to distinguish between those worlds is to use the fact that one “tail” is shorter than the others, and doing this requires a formula with K -depth at least $2n - 2$.

The K -depth of ψ' is fixed and finite, so there is some $N \in \mathbb{N}$ such that for every $n \geq N$ the formula ψ' cannot distinguish between the worlds in the second column of \mathcal{M}_n . Putting all of the above together, we get that ψ' :

- must retain all worlds in the second column below a certain point,
- must distinguish between two worlds in the second column of \mathcal{M}_n if the n -th world of the second column is retained and
- cannot distinguish between the worlds in the second column of \mathcal{M}_n for all n greater than some number N .

This is a contradiction, so our initial assumption that such a ψ' exists must be false, which proves the lemma. \square

The theorem now follows easily.

Theorem 2. *There are a \mathcal{L}_{PAL} formula φ and a $\mathcal{L}_{\text{APAL}}$ formula ψ such that $\not\models \langle \psi \rangle \varphi \rightarrow \diamond \varphi$.*

Proof. Let \mathcal{M} , w , φ and ψ be as defined above. Then $\mathcal{M}, w \models \langle \psi \rangle \varphi$ by Lemma 4. Furthermore, by Lemma 5 we know that $\mathcal{M}, w \not\models \langle \psi' \rangle \varphi$ for every \mathcal{L}_{PAL} formula ψ' so we have $\mathcal{M}, w \not\models \diamond \varphi$. This implies that $\mathcal{M}, w \not\models \langle \psi \rangle \varphi \rightarrow \diamond \varphi$ and so that $\not\models \langle \psi \rangle \varphi \rightarrow \diamond \varphi$. \square

4.3 Invalidity of $\langle \psi \rangle \varphi \rightarrow \diamond \varphi$ on Finitely Branching Models

Now to show that $\not\models_{\text{br}} \langle \psi \rangle \varphi \rightarrow \diamond \varphi$. The method used to show this is very similar to the method used to show that $\not\models \langle \psi \rangle \varphi \rightarrow \diamond \varphi$. We use φ to force ψ to retain an infinite number of worlds in a pointed model (\mathcal{N}, w) . Additionally we force ψ to distinguish between infinitely many pairs of worlds, and we let the difference between the two worlds in a pair get further and further away.

Unfortunately, forcing ψ to retain an infinite number of worlds is much more complicated in a finitely branching frame, so we need more complex formulas and models. Let \mathcal{N} be the model shown in Figure 3 and let

$$\begin{aligned} \psi &:= (\neg p \wedge \hat{K}_b(p \wedge \hat{K}_a(q \vee r))) \rightarrow \diamond(\hat{K}_a K_b p \wedge \hat{K}_a(p \wedge \neg K_b p)), \\ \varphi_1 &:= (q \vee r) \rightarrow (\hat{K}_a K_b p \wedge \hat{K}_a(p \wedge \neg K_b p)), \\ \varphi_2 &:= (q \rightarrow \neg \hat{K}_c \hat{K}_a \hat{K}_b \hat{K}_c r) \wedge (r \rightarrow \neg \hat{K}_c \hat{K}_a \hat{K}_b \hat{K}_c q), \\ \varphi &:= \langle \varphi_1 \rangle (\hat{K}_a \hat{K}_b \hat{K}_c q \wedge \hat{K}_a \hat{K}_b \hat{K}_c r \wedge \langle \varphi_2 \rangle) \square \neg (\hat{K}_a K_b p \wedge \hat{K}_a(p \wedge \neg K_b p)). \end{aligned}$$

Note the recurring a -triangles with two p worlds in the model and the recurring subformula $\hat{K}_a K_b p \wedge \hat{K}_a(p \wedge \neg K_b p)$. These subformulas have the property that they hold in the $\neg p$ world of such a triangle if and only if for one of the p worlds in the triangle a b -reachable $\neg p$ world is retained but for the other it is not.




Fig. 3. The model \mathcal{N} . Reflexive arrows are not drawn. The submodels \mathcal{N}_n^x for $n \in \mathbb{N}$, $x \in \{q, r\}$ are shown in Figure 4.




Fig. 4. The submodel \mathcal{N}_n^x for $x \in \{q, r\}$ and $n \in \mathbb{N}_{>0}$. The origin world that connects it to \mathcal{N} is the world satisfying x .

Lemma 6. *We have $\mathcal{N}, w \models \langle \psi \rangle \varphi$.*

Proof. Let us consider the update $\langle \psi \rangle$. It places the conditions on $\neg p \wedge \hat{K}_b(p \wedge \hat{K}_a(q \vee r))$ worlds that they must satisfy $\diamond(\hat{K}_a K_b p \wedge \hat{K}_a(p \wedge \neg K_b p))$. The $\neg p \wedge \hat{K}_b(p \wedge \hat{K}_a(q \vee r))$ worlds are exactly those that are in the third line from the bottom in \mathcal{N}_n^x submodels. Furthermore, of the two such worlds in a submodel

\mathcal{N}_n^x the left one satisfies $\diamond(\hat{K}_a K_b p \wedge \hat{K}_a(p \wedge \neg K_b p))$, and the right one does not.² The updated submodel $\mathcal{N}_{n\psi}^x$ is therefore as shown in Figure 5. (The worlds of \mathcal{N} that are not in one of the submodels \mathcal{N}_n^x are all retained by the update so nothing changes there.)




Fig. 5. The submodel $\mathcal{N}_{n\psi}^x$ for $x \in \{q, r\}$ and $n \in \mathbb{N}_{>0}$

After the update $\langle\psi\rangle$ the formula $\hat{K}_a K_b p \wedge \hat{K}_a(p \wedge \neg K_b p)$ therefore holds in the origin world of each submodel \mathcal{N}_n^x . Since q and r only hold in the origin worlds of these submodels the update $\langle\varphi_1\rangle = \langle(q \vee r) \rightarrow (\hat{K}_a K_b p \wedge \hat{K}_a(p \wedge \neg K_b p))\rangle$ does nothing if executed immediately after $\langle\psi\rangle$. We therefore have $\mathcal{N}, w \models \langle\psi\rangle \langle\varphi_1\rangle (\hat{K}_a \hat{K}_b \hat{K}_c q \wedge \hat{K}_a \hat{K}_b \hat{K}_c r)$.

Finally consider the third update $\langle\varphi_2\rangle$. It places conditions on $q \vee r$ worlds; q worlds must satisfy $\neg \hat{K}_c \hat{K}_a \hat{K}_b \hat{K}_c r$ and r worlds must satisfy $\neg \hat{K}_c \hat{K}_a \hat{K}_b \hat{K}_c q$. After the other updates there are no q or r worlds that satisfy this condition.

As such the result of applying the three updates $\langle\psi\rangle \langle\varphi_1\rangle \langle\varphi_2\rangle$ removes the origin worlds of all \mathcal{N}_n^x submodels. In the resulting model the two p worlds that are a -reachable from w are indistinguishable, so $\mathcal{N}, w \models \langle\psi\rangle \langle\varphi_1\rangle \langle\varphi_2\rangle \Box \neg(\hat{K}_a K_b p \wedge \hat{K}_a(p \wedge \neg K_b p))$. Together with the previous result $\mathcal{N}, w \models \langle\psi\rangle \langle\varphi_1\rangle (\hat{K}_a \hat{K}_b \hat{K}_c q \wedge \hat{K}_a \hat{K}_b \hat{K}_c r)$ this shows that $\mathcal{N}, w \models \langle\psi\rangle \varphi$. \square

² Announcements that make $\hat{K}_a K_b p \wedge \hat{K}_a(p \wedge \neg K_b p)$ true in the leftmost world in the third row do so by removing one of the $\neg p$ worlds in the fifth row but not the other. This can be done because there are formulas that differentiate between a “tail” of $2n$ worlds and a “tail” of $2n - 1$ worlds, as in the infinitely branching case.

Lemma 7. *For every \mathcal{L}_{PAL} formula ψ' we have $\mathcal{N}, w \not\models \langle \psi' \rangle \varphi$.*

Proof. Suppose towards a contradiction that there is a \mathcal{L}_{PAL} formula ψ' such that $\mathcal{N}, w \models \langle \psi' \rangle \varphi$. Then after the updates $\langle \psi' \rangle \langle \varphi_1 \rangle$ the formula $\hat{K}_a \hat{K}_b \hat{K}_c q \wedge \hat{K}_a \hat{K}_b \hat{K}_c r$ must hold in w . The origin worlds of \mathcal{N}_1^q and \mathcal{N}_1^r and the paths to those worlds must therefore be retained by $\langle \psi' \rangle \langle \varphi_1 \rangle$.

But after those two updates it must also hold in w that $\langle \varphi_2 \rangle \Box \neg (\hat{K}_a K_b p \wedge \hat{K}_a (p \wedge \neg K_b p))$, so after the update $\langle \varphi_2 \rangle$ the two worlds that are b -reachable from the p worlds that are a -reachable from w must be indistinguishable. In particular this means that neither the origin world of \mathcal{N}_1^q nor that of \mathcal{N}_1^r may be reachable, as otherwise $\hat{K}_c q$ or $\hat{K}_c r$ would distinguish the worlds.

Since the update $\langle \varphi_2 \rangle$ only removes $q \vee r$ worlds this implies that the origin worlds of \mathcal{N}_1^q and \mathcal{N}_1^r must satisfy $\neg \varphi_2$ after the first two updates. But then $\hat{K}_c \hat{K}_a \hat{K}_b \hat{K}_c r$ must hold in the origin of \mathcal{N}_1^q and $\hat{K}_c \hat{K}_a \hat{K}_b \hat{K}_c q$ in the origin of \mathcal{N}_1^r .

But then the origins of \mathcal{N}_2^q and \mathcal{N}_2^r must be reachable after the first two updates. But these two origin worlds must also be removed by $\langle \varphi_2 \rangle$ as otherwise $\hat{K}_a \hat{K}_b \hat{K}_c q$ would distinguish the two worlds that must be indistinguishable. But then the origins of \mathcal{N}_3^q and \mathcal{N}_3^r must be retained. Repeating the argument shows that if the origins of \mathcal{N}_n^q and \mathcal{N}_n^r remain reachable then so do those of \mathcal{N}_{n+1}^q and \mathcal{N}_{n+1}^r . Therefore, the updates $\langle \psi' \rangle \langle \varphi_1 \rangle$ must leave the origin of every \mathcal{N}_n^x submodel reachable.

But then consider the update $\langle \varphi_1 \rangle$. This update retains the origin of a \mathcal{N}_n^x submodel if and only if it satisfies $\hat{K}_a K_b p \wedge \hat{K}_a (p \wedge \neg K_b p)$. This implies that for each $n \in \mathbb{N}$ and $x \in \{q, r\}$ the update $\langle \psi' \rangle$ must retain one of the worlds on the third row of the submodel but not the other. However, in \mathcal{N}_n^x these worlds are indistinguishable up to depth $2n$, so a \mathcal{L}_{PAL} formula must contain at least $2n + 1$ iterations of a K -operator to distinguish them. There is therefore no single formula in \mathcal{L}_{PAL} that distinguishes the two worlds for every submodel. This contradicts the assumption that such a ψ' exists. \square

The theorem now follows easily.

Theorem 3. *There are $\mathcal{L}_{\text{APAL}}$ formulas φ, ψ such that $\not\models_{\text{br}} \langle \psi \rangle \varphi \rightarrow \Diamond \varphi$.*

Proof. For the $\mathcal{L}_{\text{APAL}}$ formulas φ, ψ , finitely branching model \mathcal{N} and world w of \mathcal{N} as defined above we have $\mathcal{N}, w \models \langle \psi \rangle \varphi$ by Lemma 6 and $\mathcal{N}, w \not\models \langle \psi' \rangle \varphi$ for every \mathcal{L}_{PAL} formula ψ' by Lemma 7. This implies that $\mathcal{N}, w \not\models \langle \psi \rangle \varphi \rightarrow \Diamond \varphi$ so $\not\models_{\text{br}} \langle \psi \rangle \varphi \rightarrow \Diamond \varphi$. \square

5 Conclusion and Further Research

I showed that for any $\mathcal{L}_{\text{APAL}}$ formula φ and ψ we have $\models_{\text{fin}} \langle \psi \rangle \varphi \rightarrow \Diamond \varphi$ and that for any \mathcal{L}_{PAL} formula φ and any $\mathcal{L}_{\text{APAL}}$ formula ψ we also have $\models_{\text{br}} \langle \psi \rangle \varphi \rightarrow \Diamond \varphi$. Additionally, I showed that there are $\mathcal{L}_{\text{APAL}}$ formulas φ and ψ such that $\not\models_{\text{br}} \langle \psi \rangle \varphi \rightarrow \Diamond \varphi$ and that there are a \mathcal{L}_{PAL} formula φ and a $\mathcal{L}_{\text{APAL}}$ formula ψ such that $\not\models \langle \psi \rangle \varphi \rightarrow \Diamond \varphi$.

The operator \diamond therefore only represents a truly arbitrary public announcement on finite models. There are scenarios that can be modeled in finite models and where arbitrary public announcements are useful, such as the cryptography example mentioned in the introduction. The message p for which we want to know whether $\diamond(K_b p \wedge \neg K_e p)$ is generally taken from a finite set of possible messages which allows for a finite model to be used.

However, not all interesting scenarios allow for finite modeling, so it seems like an interesting topic for further research whether semantics for a different arbitrary public announcement operator \blacklozenge can be found such that for any $\mathcal{L}_{\text{PAL}+\blacklozenge}$ formulas φ, ψ we have $\models \langle \psi \rangle \varphi \rightarrow \blacklozenge \varphi$. One possibility that might work is an infinite hierarchy of \diamond_i operators, where each \diamond_i quantifies over all formulas that use only \diamond_j with $j < i$. I conjecture that if we then define $\blacklozenge \varphi$ as $\bigvee_{i \in \mathbb{N}} \diamond_i \varphi$ we have $\models \langle \psi \rangle \varphi \rightarrow \blacklozenge \varphi$.

References

1. van Ditmarsch, H., van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic. Springer (2007)
2. Plaza, J.: Logics of public communication. In: Emrich, M., Phifer, M., Hadzikadic, M., Ras, Z. (eds.) Proceedings of the Fourth International Symposium on Methodologies for Intelligent Systems, Poster Session Program, pp. 201–216. Oak Ridge National Laboratory, ORNL/DSRD-24 (1989); Reprinted as [7]
3. Baltag, A., Moss, L., Solecki, S.: The logic of public announcements, common knowledge, and private suspicions. In: Gilboa, I. (ed.) Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge, pp. 43–56. Morgan Kaufmann Publishers Inc. (1998)
4. Fitch, F.: A logical analysis of some value concepts. The Journal of Symbolic Logic 28, 135–142 (1963)
5. Balbiani, P., Baltag, A., van Ditmarsch, H., Herzig, A., Hoshi, T., de Lima, T.: What can we achieve by arbitrary announcements? A dynamic take on Fitch’s knowability. In: Samet, D. (ed.) Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge, pp. 42–51 (2007)
6. Balbiani, P., Baltag, A., van Ditmarsch, H., Herzig, A., Hoshi, T., de Lima, T.: ‘Knowable’ as ‘Known after an Announcement’. The Review of Symbolic Logic 1, 305–334 (2008)
7. Plaza, J.: Logics of public communication. Synthese 158, 165–179 (2007)