

Announcements to Attentive Agents

Bolander, Thomas; van Ditmarsch, Hans; Herzig, Andreas ; Lorini, Emiliano ; Pardo, Pere ; Schwarzentruher, Francois

Published in:
Journal of Logic, Language and Information

Link to article, DOI:
[10.1007/s10849-015-9234-3](https://doi.org/10.1007/s10849-015-9234-3)

Publication date:
2016

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Bolander, T., van Ditmarsch, H., Herzig, A., Lorini, E., Pardo, P., & Schwarzentruher, F. (2016). Announcements to Attentive Agents. Journal of Logic, Language and Information, 25(1), 1-35. DOI: 10.1007/s10849-015-9234-3

DTU Library

Technical Information Center of Denmark

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

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.

Announcements to Attentive Agents

Thomas Bolander* Hans van Ditmarsch† Andreas Herzig‡
Emiliano Lorini§ Pere Pardo¶ François Schwarzentruber||

November 20, 2016

Abstract

In public announcement logic it is assumed that all agents pay attention to the announcement. Weaker observational conditions can be modelled in action model logic. In this work, we propose a version of public announcement logic wherein it is encoded in the states of the epistemic model which agents pay attention to the announcement. This logic is called attention-based announcement logic. We give an axiomatization of the logic and prove that complexity of satisfiability is the same as that of public announcement logic, and therefore lower than that of action model logic. An attention-based announcement can also be described as an action model. We extend our logic by integrating attention change. Finally, we add the notion of common belief to the language, we exploit this to formalize the concept of joint attention, that has been widely discussed in the philosophical and cognitive science literature, and we provide a corresponding axiomatization. This axiomatization also employs the auxiliary notion of attention-based relativized common belief.

1 Introduction

In public announcement logic [18] it is assumed that announcements are observed by all agents: it models the consequences of each of the agents incorporating a new formula into its set of beliefs. It does not model why the agent should wish to incorporate that new information. These epistemic actions are called *public announcements*, as they have the legal connotation that the agents cannot from then on be excused not to have heard what was said. Once the government has announced a new election, they cannot be held liable when you forget to vote on election day. You were supposed to know.

*DTU Compute, Technical University of Denmark

†LORIA, CNRS – University of Lorraine, France

‡University of Toulouse, CNRS, IRIT, France

§University of Toulouse, CNRS, IRIT, France & IAST, Toulouse, France

¶LORIA, CNRS – University of Lorraine, France

||ENS Rennes, IRISA, France

In this work we take one step back from that point of view. When an announcement is made, it may well be that some agents were not paying attention and therefore did not hear it. Also, there may be uncertainty among the agents about who is paying attention and who not, and therefore, who heard the message and who not. In our work we model restricted attention and uncertainty about it.

Additional to the usual set of propositional variables we add designated variables for each agent, that express that the agent is paying attention. A given state of a Kripke model therefore contains information about which agents are paying attention and which agents are not paying attention. This determines the meaning of what we call *attention-based announcements*. A special case is that of introspective agents that know whether they are paying attention.

An announcement by an outsider (an agent who is not explicitly modelled in the logical system) that is public for a subset of all agents is modelled in [9, 5, 8] as a *private announcement* to that subset of agents. The agents' attention configuration behind such announcements can be modelled in our logic by a particular formula built from our attention variables. However, our logic generalizes [9], because the 'attention level' of a given agent can vary between the states. Our logic can also be seen as a fragment of action model logic [5].

We show by way of a tableau calculus that the complexity of satisfiability in our logic remains in the same range as that of public announcement logic, viz. PSPACE [17]. This contrasts with the higher complexity of action model logic, that is NEXPTIME [2]. As the action models corresponding to attention-based announcements can be quite large, we consider that this is indeed a valuable result.

We also add further dynamics to our logic, namely change of attention. This is an elementary further addition to the logical framework and this logic also has a complete axiomatization. With this addition we can embed Gerbrandy's believed public announcement logic into attention-based announcement logic.

In attention-based announcement logic we can formalize a concept that has been widely discussed in the philosophical and in the cognitive science literature, namely *joint attention* [19, 16, 7]. This concept has been shown to be crucial for explaining the genesis of common belief in a group of agents. This then sets the stage for the final part of our work, wherein we add common belief and attentive relativized common belief to the logical language, and we axiomatize the resulting logic employing the techniques of [20].

Example 1

Before we jump into the definitions, let us give a motivating example. Consider two agents Anne and Bob who are both uncertain of dinner being served soon (enough...). Now Carol says that dinner will be at 11 at night (we are in Seville, people tend to eat rather late there). Let us ignore what Carol may learn from her announcement, in other words: we do not model her knowledge or belief. If neither Anne nor Bob are paying attention, they won't gain any information at all. If Anne is paying attention but not Bob, Anne now knows that dinner is at 11 but Bob doesn't. Does Anne know that? If Anne is uncertain whether Bob was paying attention, she cannot conclude that Bob is still uncertain about

the dinner time. Let us restart the scenario. Now, prior to her announcement, Carol claps her hands. Anne and Bob are startled. They are paying attention. Now Carol says: dinner will be at 11! They now have common knowledge of the dinner time. Next restart of the scenario: instead of clapping her hands, she taps Anne on her shoulder and then says: dinner will be at 11 (and Bob does not notice that), and then taps Bob on his shoulder and does likewise. However, as Anne was already attentive, she obviously notices that Carol informs Bob. Now, Anne and Bob both know that dinner is at 11. But they have no common knowledge of that. More complex scenarios are conceivable, and, given that things may happen that you do not notice, we are rather modelling Anne’s and Bob’s beliefs than their knowledge. An important modelling contribution in our work is that under certain conditions common belief can still be obtained even when Anne and Bob are not both paying attention and common belief was so far not established. \dashv

Section 2 presents the language, structures, and semantics of attention-based announcement logic. Section 3 provides the axiomatization. Section 4 establishes the PSPACE complexity of satisfiability. Section 5 gives an embedding into action model logic. Section 6 adds operators for attention change to the attention-based announcement logic. Section 7 addresses joint attention, and, finally, Section 8 introduces different notions for common belief and provides a complete axiomatization for that addition.

2 Attention-Based Announcement Logic ABAL

Let AGT be a finite set of agents, let ATM be a countable set of propositional variables, and let $H = \{h_a \mid a \in AGT\}$ be a set of propositional variables that is disjoint from ATM . A proposition h_a (for ‘ a is hearing what is being said’ or ‘ a is listening’) expresses that agent a is paying attention and so will hear announcements.

Definition 2 (Language)

The *language* \mathcal{L} of Attention-Based Announcement Logic (ABAL) is defined as follows, where $p \in ATM$ and $a \in AGT$ (so that, implicitly, $h_a \in H$).

$$\mathcal{L} \ni \varphi ::= p \mid h_a \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid B_a\varphi \mid [\varphi]\varphi$$

For $A \subseteq AGT$, we abbreviate $\bigwedge_{a \in A} h_a$ by h_A . Let \mathcal{L}^- be the language \mathcal{L} (of multi-agent modal logic) without the inductive construct for announcement. \dashv

We write q to denote a variable that is either a $p \in ATM$ or a $h_a \in H$. If we wish to be explicit about the sets of agents and propositional variables we write $\mathcal{L}(ATM \cup H, AGT)$ instead of \mathcal{L} . Other propositional connectives (including \perp and \top), and the dual modalities, are defined as usual. Formula $B_a\varphi$ is read as ‘agent a believes that φ is true’, and formula $[\varphi]\psi$ as ‘after the public announcement of φ , ψ holds’. We write $\psi[\varphi/p]$ for *uniform substitution of φ for p in ψ* .

Definition 3 (Epistemic attention model)

An *epistemic attention model* is a triple $M = (S, R, V)$ with S a non-empty set of states, R a function assigning to each agent an accessibility relation $R_a \subseteq S \times S$ and V a function assigning to each propositional variable $q \in ATM \cup H$ the subset $V(q) \subseteq S$ where the variable is true. \dashv

Definition 4 (Attention introspection)

Given an epistemic attention model $M = (S, R, V)$, the model satisfies the property of *attention introspection* if for all $s, t \in S$, if $(s, t) \in R_a$, then $s \in V(h_a)$ iff $t \in V(h_a)$. \dashv

We will see that, when attention introspection holds, an agent knows whether she is paying attention. In this paper we focus on two classes of models: \mathcal{K}_n , where $n = |AGT|$, that is, multiple agents and no special properties of the accessibility relations, and $\mathcal{K}45_n^h$, that is, multiple agents with transitive and euclidean accessibility relations, and with attention introspection as well. We emphasize that attention introspection is a model property (it depends on the valuation of the propositional variables in H) and not a frame property, so that $\mathcal{K}45_n^h$ is indeed a model class and not a *frame* class.

Definition 5 (Semantics)

Let $M = (S, R, V)$ be an epistemic attention model and $s \in S$. We define the semantics by induction on $\varphi \in \mathcal{L}$. (To disambiguate we occasionally put parentheses around M, s .)

$$\begin{aligned}
 M, s \models q & \quad \text{iff} \quad s \in V(q) & \quad \text{for } q \in ATM \cup H \\
 M, s \models \neg\varphi & \quad \text{iff} \quad M, s \not\models \varphi \\
 M, s \models \varphi \wedge \psi & \quad \text{iff} \quad M, s \models \varphi \text{ and } M, s \models \psi \\
 M, s \models B_a\varphi & \quad \text{iff} \quad M, t \models \varphi \text{ for all } (s, t) \in R_a \\
 M, s \models [\varphi]\psi & \quad \text{iff} \quad M_\varphi, (s, h) \models \psi
 \end{aligned}$$

where $M_\varphi = (S', R', V')$ is defined as follows.

- $S' = S \times \{h, \bar{h}\}$ (where h denotes ‘heard’ copy of the model and \bar{h} ‘not heard’);
- for each agent a , $((s, i), (t, j)) \in R'_a$ if and only if $(s, t) \in R_a$ and:
 1. $i = h, j = h, (M, s) \models h_a$ and $(M, t) \models \varphi$, or
 2. $i = h, j = \bar{h}$, and $(M, s) \not\models h_a$, or
 3. $i = \bar{h}, j = \bar{h}$;
- for each $p \in ATM$, $(s, h) \in V'(p)$ iff $s \in V(p)$ and $(s, \bar{h}) \in V'(p)$ iff $s \in V(p)$. \dashv

Validity is defined as usual, as truth in all states in all models. The set of valid \mathcal{L} -formulas on the class of models \mathcal{K}_n is called **ABAL**, and the set of valid \mathcal{L} formulas on the class of models $\mathcal{K}45_n^h$ is called **ABAL^{int}**.

The truth condition for attention-based announcements $[\varphi]$ is different from that of state eliminating public announcement [18] and also different from that of arrow eliminating public announcement [9], although it comes closer to the latter in spirit: it is also arrow

eliminating. We recall the semantics of arrow eliminating (‘believed’) public announcements, for which we write $\llbracket\varphi\rrbracket$: $M, s \models \llbracket\varphi\rrbracket\psi$ iff $M|\varphi, s \models \psi$ where $M|\varphi = (S, R'', V)$ is such that for each agent a , $(s, t) \in R''_a$ iff $M, t \models \varphi$. It is well-known that state eliminating public announcement and arrow eliminating public announcement have the same update effect (they result in bisimilar models) on the part of the model where the announced formula is true [14].

Attention-based public announcements have the semantics of arrow eliminating public announcement for the agents that are paying attention, but not for the agents that are not paying attention. After the announcement of φ , the agents that are attentive only consider possible the h -copies of the states of the original model M in which φ is true. In contrast, the agents that are not attentive only consider possible \bar{h} -copies of the states of the original model M . This construction of the updated model M_φ ensures that attentive agents learn φ while inattentive agents do not learn anything, they think that nothing at all happened. In that sense, the semantics is not unlike that of the Gerbrandy-style private announcement to a subgroup [9]. This is a special case of the already mentioned arrow eliminating public announcement. An important difference with Gerbrandy is that in private announcements the informed agents know that the other agents did not notice anything at all, whereas in attention-based announcements the attentive agents may be uncertain about who is paying attention. Dually, the announcer of private announcements only addresses that subgroup, whereas an attention-based announcement is addressed to all agents. It merely may not reach all agents. A more precise relation between private announcement and attention-based announcements can be given after we also introduce attention change, in Section 6.

The agents who are not attentive may have incorrect beliefs after an attention-based announcement, even if they only had correct beliefs before. This is obvious, for example, after an announcement of p to a and b it may be that a now believes p but that b , who was not paying attention, still believes that a does not believe p . The class $\mathcal{S}5_n^h$ is therefore not preserved after attention-based announcements. Also, as usual for public announcements to agents having possibly incorrect beliefs, the class $\mathcal{KD}45_n^h$ is not preserved: an agent who believes p , will have the empty accessibility relation after incorporating the information that p is false. Seriality (D) may therefore not be preserved after an announcement. This explains the restriction of our results to the classes \mathcal{K}_n and $\mathcal{K}45_n^h$.

Proposition 6

When attention introspection holds, an agent knows whether she is paying attention: $B_a h_a \vee B_a \neg h_a$ is valid in ABAL. ⊢

Proof Both $h_a \rightarrow B_a h_a$ and $\neg h_a \rightarrow B_a \neg h_a$ are valid. □

Our definition of attention introspection matches that of *awareness introspection* in the economics literature [12]. Just as it is counterintuitive that an agent is aware (of worlds, of formulas, ...) but believes that she is not aware, it is counterintuitive that an agent is attentive (of possibly incoming new information) but believes that she is not attentive. It seems less counterintuitive that an agent is unattentive but believes that she is attentive.

all propositional tautologies	$B_a\varphi \rightarrow B_aB_a\varphi$	(*)
$B_a(\varphi \rightarrow \psi) \rightarrow (B_a\varphi \rightarrow B_a\psi)$	$\neg B_a\varphi \rightarrow B_a\neg B_a\varphi$	(*)
$[\varphi]q \leftrightarrow q$	$h_a \rightarrow B_a h_a$	(**)
$[\varphi]\neg\psi \leftrightarrow \neg[\varphi]\psi$	$\neg h_a \rightarrow B_a\neg h_a$	(**)
$[\varphi](\psi \wedge \chi) \leftrightarrow ([\varphi]\psi \wedge [\varphi]\chi)$	From φ infer $B_a\varphi$	
$[\varphi]B_a\psi \leftrightarrow ((h_a \rightarrow B_a(\varphi \rightarrow [\varphi]\psi)) \wedge (\neg h_a \rightarrow B_a\psi))$	From φ and $\varphi \rightarrow \psi$ infer ψ	
	From $\varphi \leftrightarrow \psi$ infer $\chi[\varphi/p] \leftrightarrow \chi[\psi/p]$	

Table 1: The axiomatizations for ABAL (minus the (*) and (**) axioms) and ABAL^{int} (all)

Proposition 7 (Preservation of attention introspection)

If M satisfies attention introspection then M_φ satisfies attention introspection. \dashv

Proof Obvious. \square

We finish this section with an example illustrating the semantics.

Example 8

Ann (a) and Bill (b) have lunch in Excelsior and they both ruminate about snowfall this afternoon (p) — a regular occurrence in Nancy, many times of the year. In fact Bill has seen the weather report and knows whether it will snow, while Ann does not. However, Bill never knows whether Ann is attentive, whereas Ann knows that Bill is attentive — let us not delve into the reasons for this credulous but justified stance. This situation is depicted as model M in Fig. 1. Note that both agents know whether they are attentive: we have attention introspection. Cath comes along and says she just read the weather report: it will snow. Cath’s announcement is modelled as an announcement of p by an outsider. This results in the model transition depicted in Figure 1. Any of the four leftmost points in M_p can be the actual state. For example, if Ann and Bill pay attention and p is true (the bottom-left state in M) then after the announcement of p (the bottom-left state in M_p) Bill remains uncertain if Ann now knows that p , as he considers it possible that she was not paying attention (top-left state), in which case she would have remained uncertain about p (the two a -arrows pointing to the right part of M_p , that copies the original model M). \dashv

3 Axiomatization

Definition 9

The axiomatization of ABAL consists of all the derivation rules and the non-*-ed axioms of Table 1. The axiomatization of ABAL^{int} consists of ABAL plus the *-ed axioms. \dashv

The crucial axiom in the axiomatization is that for belief after attention-based announcement:

$$[\varphi]B_a\psi \leftrightarrow ((h_a \rightarrow B_a(\varphi \rightarrow [\varphi]\psi)) \wedge (\neg h_a \rightarrow B_a\psi))$$




Figure 1: Example 8 illustrates an attention-based announcement. The left grey square is the h part of the model M , the consequences of b paying attention. The right grey square is the \bar{h} part of the model M , the (absence of) consequences of a not paying attention.

It says that the belief consequences of an attention-based announcement are either, if the agent pays attention, what the agent believes to be the consequences of the announcement in case it is true, or else, if the agent does not pay attention, what the belief consequences were before the announcement. In other words, for the latter, an agent not hearing the announcement does not change her beliefs.

The axioms * formalize that agents have introspective beliefs. They are not uncertain about what they believe. But their beliefs may still be incorrect. The axioms ** formalize attention introspection. The agents know whether they are paying attention.

The axiomatization follows the pattern of arrow eliminating (‘believed’) public announcements, not that of state eliminating (‘true’) public announcements. Believed public announcements of φ have precondition \top , not precondition φ . Therefore, there is no relativization to φ in the axioms. For example, we have axioms $[\varphi]\neg\psi \leftrightarrow \neg[\varphi]\psi$ and $[\varphi]p \leftrightarrow p$ instead of $[\varphi]\neg\psi \leftrightarrow (\varphi \rightarrow \neg[\varphi]\psi)$ and $[\varphi]p \leftrightarrow (\varphi \rightarrow p)$. The reduction axiom for belief after attention-based announcement is reminiscent of that of belief after arrow eliminating public announcement: $[\varphi]B_a\psi \leftrightarrow B_a(\varphi \rightarrow [\varphi]\psi)$.

Theorem 10

The axiomatization of ABAL is sound and complete for the class of \mathcal{K}_n models, and is equally expressive as the base modal logic \mathcal{K}_n . \dashv

Proof

Soundness

We only show the validity of $[\varphi]B_a\psi \leftrightarrow ((h_a \rightarrow B_a(\varphi \rightarrow [\varphi]\psi)) \wedge (\neg h_a \rightarrow B_a\psi))$. Let a pointed epistemic attention model (M, s) be given.

\Rightarrow

First assume $M, s \models h_a$, and let $(s, t) \in R_a$ such that $M, t \models \varphi$. In that case, $(t, h) \in M_\varphi$. Now using the assumption $M, s \models [\varphi]B_a\psi$, it follows that $M_\varphi, (s, h) \models B_a\psi$, and with $((s, h), (t, h)) \in R_a^\varphi$ (because $(s, t) \in R_a$, $M, s \models h_a$, and $M, t \models \varphi$, as spelled out in Definition 5) we obtain the required $M_\varphi, (t, h) \models \psi$.

Now assume $M, s \models \neg h_a$, and let $(s, t) \in R_a$. Then $((s, h), (t, \bar{h})) \in R_a^\varphi$, and with $M, s \models [\varphi]B_a\psi$ again we get $M_\varphi, (s, h) \models B_a\psi$, and therefore $M_\varphi, (t, \bar{h}) \models \psi$. One then uses that $M_\varphi, (t, \bar{h}) \models \psi$ iff $M, t \models \psi$ (which is easily obtained).

\Leftarrow

The other direction proceeds fairly similarly.

Completeness and expressivity

The standard reduction argument applies: all axioms for the consequences of announcements push the announcement operator deeper into the formula on the right hand side, until one finally arrives at an announcement before a propositional variable, $[\varphi]q$, which is equivalent to q (where $q \in ATM \cup H$). By an inside-out reduction strategy, wherein we also Therefore the logic is as expressive as the base modal logic — that is complete. \square

Theorem 11

The axiomatization ABAL^{int} is sound and complete for the class $\mathcal{K}45_n^h$. \dashv

Proof The usual introspection axioms $B_a\varphi \rightarrow B_aB_a\varphi$ and $\neg B_a\varphi \rightarrow B_a\neg B_a\varphi$ are valid on $\mathcal{K}45_n^h$, and correspond to frame properties. The axioms $h_a \rightarrow B_a h_a$ and $\neg h_a \rightarrow B_a\neg h_a$ are also valid on the class $\mathcal{K}45_n^h$. They do not correspond to frame properties (as they depend on the valuation of propositional variables h_a), but they enforce in the canonical model (for the base model logic) a requirement that maximal consistent sets always contain both h_a and $B_a h_a$, or $\neg h_a$ and $B_a\neg h_a$. \square

Soundness and completeness of the logics also follow indirectly from the action model modelling of attention-based announcements in Section 5.

We observe that $[\varphi][\psi]\chi \leftrightarrow [\varphi \wedge [\varphi]\psi]\chi$ is not an axiom. It is well-known that, in order to demonstrate completeness of an axiomatization in dynamic epistemic logic by way of reduction axioms, one has either needs the derivation rule of substitution of equivalents (From $\varphi \leftrightarrow \psi$ infer $\chi[\varphi/p] \leftrightarrow \chi[\psi/p]$), or one needs an axiom on the composition of dynamic modalities (as above, for attentive announcements); see [25] for an in-depth discussion.

As we frequently use this composition property of attentive announcements, for example, in the tableaux construction and in the common belief section, it may be useful to point out why it is valid.

Proposition 12

$$\models [\varphi][\psi]\chi \leftrightarrow [\varphi \wedge [\varphi]\psi]\chi \quad \dashv$$

Proof Let (M, s) be a pointed epistemic attention model, and $\varphi, \psi, \chi \in \mathcal{L}$. We wish to show that $M, s \models [\varphi][\psi]\chi$ iff $M, s \models [\varphi \wedge [\varphi]\psi]\chi$. Using the semantics of attention-based announcement, this is equivalent to $(M_\varphi)_\psi, ((s, h), h) \models \chi$ iff $M_{\varphi \wedge [\varphi]\psi}, (s, h) \models \chi$. We now observe that $((M_\varphi)_\psi, ((s, h), h)) \Leftrightarrow (M_{\varphi \wedge [\varphi]\psi}, (s, h))$ by way of the bisimulation \mathfrak{R} defined as (using infix notation): for all $s \in S$, $((s, h), h)\mathfrak{R}(s, h)$ and $((s, \bar{h}), \bar{h})\mathfrak{R}(s, \bar{h})$. Note that (i) the domain of $(M_\varphi)_\psi$ is twice the size of the domain of $M_{\varphi \wedge [\varphi]\psi}$, and that (ii) we do not define the bisimulation on all elements of the domain of $(M_\varphi)_\psi$, namely not on those of shape $((s, h), \bar{h})$ or shape $((s, \bar{h}), h)$. We only wish to demonstrate bisimilarity of the *pointed* models $((M_\varphi)_\psi, ((s, h), h))$ and $(M_{\varphi \wedge [\varphi]\psi}, (s, h))$, and from points $((s, h), h)$ only those of shape $((t, h), h)$ and shape $((t, \bar{h}), \bar{h})$ are accessible (namely, the first for attentive agents, and the second for non-attentive agents). (So, notwithstanding the domains of different size, the bisimulation \mathfrak{R} is a bijection from the part of the domain of $(M_\varphi)_\psi$ on which it is defined, to the domain of $M_{\varphi \wedge [\varphi]\psi}$.) Then, having established bisimilarity, we use that bisimilarity implies logical equivalence for attention-based announcement logic (which can be proved either by induction on the structure of formulas, with non-trivial case $[\chi_1]\chi_2$, or by referring to the action model embedding in Section 5). From the logical equivalence of $(M_\varphi)_\psi, ((s, h), h)$ and $M_{\varphi \wedge [\varphi]\psi}, (s, h)$ then follows, as required, that $(M_\varphi)_\psi, ((s, h), h) \models \chi$ iff $M_{\varphi \wedge [\varphi]\psi}, (s, h) \models \chi$. \square

4 Satisfiability

We now focus on the ABAL satisfiability problem on the class of all \mathcal{K}_n frames. The satisfiability problem of a formula in action model logic plus the union operator over

actions is NEXPTIME-complete [2]. In Section 5 we will see that the logic ABAL is a fragment of action model logic. So the satisfiability problem of ABAL is decidable. In the underlying section we will show that, however, its complexity remains that of public announcement logic: PSPACE.

It is difficult to turn the tableau method of [2] into a PSPACE procedure, because each node in the tableau may contain an exponential amount of information in the length of the input formula. This is, because the action model corresponding to an attention-based announcement is exponential in the size of the number of agents (see the next Section 5). Surprisingly, we can adapt the tableau method of [2] (for action model logic) so that the amount of information in a node is polynomial in the size of the input. We now proceed by first defining the usual tableau terminology, and then prove PSPACE-completeness of ABAL in Proposition 15.

Let \mathfrak{Lab} be a countable set of labels designed to represent states of the epistemic attention model (M, s) . Our tableau method manipulates terms that we call *tableau terms*. They are of the following kind:

- $(\sigma \Sigma \varphi)$ where $\sigma \in \mathfrak{Lab}$ is a symbol (that represents a state in the initial model), Σ is a sequence of formulas (where $[]$ denotes the empty list, and where $\Sigma::\varphi$ denotes concatenation of list Σ with formula φ), and $\varphi \in \mathcal{L}$. This means that φ is true after the announcement of the formulas of the sequence Σ (in that order) in the state denoted by σ ;
- $(\sigma \Sigma \checkmark)$ means that in σ all the announcements in the sequence Σ are true when they are sequentially announced; e.g. $(\sigma [\varphi_1, \varphi_2] \checkmark)$ means that $\varphi_1 \wedge [\varphi_1]\varphi_2$ is true in σ ;
- $(\sigma \Sigma \times)$ means that in σ some announcement in the sequence Σ is not true when announced; e.g. $(\sigma [\varphi_1, \varphi_2] \times)$ means that $\varphi_1 \wedge [\varphi_1]\varphi_2$ is false in σ ;
- $(\sigma R_a \sigma_1)$ means that the state denoted by σ is linked by R_a to the state denoted by σ_1 ;
- \perp denotes an inconsistency.

A *tableau rule* is represented by a *numerator* \mathcal{N} above a line and a finite list of *denominators* $\mathcal{D}_1, \dots, \mathcal{D}_k$ below this line, separated by vertical bars:

$$\frac{\mathcal{N}}{\mathcal{D}_1 \mid \dots \mid \mathcal{D}_k}$$

The numerator and the denominators are finite sets of tableau terms.

A *tableau tree* is a finite tree with a set of tableau terms at each node. A rule with numerator \mathcal{N} is *applicable* to a node carrying a set Γ , if Γ contains an instance of \mathcal{N} . If no rule is applicable, Γ is said to be *saturated*. We call a node σ an *end node*, if the set of formulas Γ it carries is saturated or if $\perp \in \Gamma$. The tableau tree is extended as follows:

1. Choose a leaf node n carrying Γ where n is not an end node, and choose a rule ρ applicable to n .
2. (a) If ρ has only one denominator, add the appropriate instantiation to Γ .
 (b) If ρ has k denominators with $k > 1$, create k successor nodes for n , where each successor i carries the union of Γ with an appropriate instantiation of denominator \mathcal{D}_i .

A branch in a tableau tree is a path from the root to an end node. A branch is *closed* if its end node contains \perp , otherwise it is *open*. A tableau tree is *closed* if all its branches are closed, otherwise it is *open*. The *tableau tree for a formula* $\varphi \in \mathcal{L}$ is the tableau tree obtained from the root $\{(\sigma_0 \sqcup \varphi)\}$ when all leaves are end nodes.

The tableau rules are depicted in Figure 2. They contain the classical propositional rules (\wedge), ($\neg\neg$) and a rule ($\neg\wedge$) with two denominators that handles disjunctions. The rules (\leftarrow_q) and ($\leftarrow_{\neg q}$) (for $q \in ATM \cup H$) express that valuations are not changed by announcements.

Depending on the value of h_a , there are two versions of the rules B_a and $\neg B_a$. Let us explain the rule ($\neg B_a^h$), for when agent a is attentive: if after the announcements in Σ the formula $\neg B_a \varphi$ is true in σ , then agent a can reach from σ a possible state denoted σ_{new} , such that all the announcements in Σ are true when announced, and such that afterwards $\neg\varphi$ holds. On the contrary, in the rule ($\neg B_a^{\bar{h}}$), agent a does not consider announcements in σ_{new} . This is why we impose $\neg\varphi$ to be true in σ_{new} after absence of announcement.

The rule (*hear*) is a cut-rule for the atomic propositions h_a in the root. It is required to guarantee the application of rules (B_a) and ($\neg B_a$). The rule (*hear*) is non-analytic: the formulas in its denominator are not subformulas of the input formula.

The rules (\checkmark), (\times), (*clash* $_{\checkmark, \times}$) and (\sqcup_{\times}) deal with the truth of the sequence Σ .

Example 13

Figure 3 shows a tableau for formula $[p]B_a q \wedge \neg(\neg(h_a \wedge \neg B_a(\neg(p \wedge \neg q))) \wedge \neg(\neg h_a \wedge \neg B_a q))$. As all branches are closed, we conclude that the formula is unsatisfiable.

Theorem 14 (Soundness and completeness of the tableau method)

A \mathcal{L} -formula φ is satisfiable iff there exists a tableau for φ with an open branch. ⊣

Proof

\Rightarrow

Suppose that formula φ is satisfiable. Then there exists a pointed model (M, s) where $M = (S, R, V)$ such that $M, s \models \varphi$. We use the model M and the updated models from M to construct an open branch in the tableau tree for φ . In the following, we index a model M with a list of announced formulas Σ as follows: we inductively define M_Σ by $M_\emptyset = M$ and $M_{\Sigma::\psi} = (M_\Sigma)_\psi$. We also write (s, \vec{h}) for $(\dots (s, h), \dots h)$.

Now we construct step by step a branch and in parallel we construct a (partial) mapping $f : \mathcal{L}\mathbf{ab} \rightarrow S$ such that, if Γ is the set of tableau terms carried by the last node in the already constructed branch, then:

$$\begin{array}{c}
\frac{(\sigma \Sigma \varphi \wedge \psi)}{(\sigma \Sigma \varphi)} (\wedge) \\
(\sigma \Sigma \psi)
\end{array}
\qquad
\frac{(\sigma \Sigma \neg\neg\varphi)}{(\sigma \Sigma \varphi)} (\neg\neg)$$

$$\frac{(\sigma \Sigma \neg(\varphi \wedge \psi))}{(\sigma \Sigma \neg\varphi) \mid (\sigma \Sigma \neg\psi)} (\neg\wedge)
\qquad
\frac{(\sigma \Sigma q)(\sigma \Sigma \neg q)}{\perp} (\perp)$$

$$\frac{(\sigma \Sigma [\varphi]\psi)}{(\sigma \Sigma::\varphi \psi)} ([\varphi])
\qquad
\frac{(\sigma \Sigma \neg[\varphi]\psi)}{(\sigma \Sigma::\varphi \neg\psi)} (\neg[\varphi])$$

$$\frac{(\sigma \Sigma q)}{(\sigma \parallel q)} (\leftarrow q)
\qquad
\frac{(\sigma \Sigma \neg q)}{(\sigma \parallel \neg q)} (\leftarrow_{\neg q})$$

$$\frac{(\sigma \Sigma::\varphi \checkmark)}{(\sigma \Sigma \varphi)} (\checkmark)
\qquad
\frac{(\sigma \Sigma::\varphi \times)}{(\sigma \Sigma \checkmark) \mid (\sigma \Sigma \times)} (\times)$$

$$\frac{(\sigma \Sigma \times)(\sigma \Sigma \checkmark)}{\perp} (clash_{\checkmark, \times})
\qquad
\frac{(\sigma \parallel \times)}{\perp} (\parallel \times)$$

$$\frac{(\sigma \Sigma B_a \varphi)(\sigma \parallel h_a)}{(\sigma R_a \sigma_1)} (B_a^h)
\qquad
\frac{(\sigma \Sigma B_a \varphi)(\sigma \parallel \neg h_a)}{(\sigma R_a \sigma_1)} (B_a^{\bar{h}})$$

$$\frac{(\sigma \parallel h_a)(\sigma \Sigma \neg B_a \varphi)}{(\sigma R_a \sigma_{\text{new}})} (\neg B_a^h)
\qquad
\frac{(\sigma \parallel \neg h_a)(\sigma \Sigma \neg B_a \varphi)}{(\sigma R_a \sigma_{\text{new}})} (\neg B_a^{\bar{h}})$$

$$\frac{}{(\sigma \parallel h_a) \mid (\sigma \parallel \neg h_a)} (hear)$$

Figure 2: **Tableau rules for ABAL**




Figure 3: Example of a tableau tree

- for all labels σ appearing in Γ , $f(\sigma)$ is defined;
- if $(\sigma \Sigma \varphi) \in \Gamma$ then $M_\Sigma, (f(\sigma), \vec{h}) \models \varphi$;
- if $(\sigma \Sigma \checkmark) \in \Gamma$ then if $\Sigma = [\psi_1, \dots, \psi_n]$, we have $M, f(\sigma) \models \psi_1$, $M_{\psi_1}, (f(\sigma), h) \models \psi_2$, \dots , $M_{[\psi_1, \dots, \psi_{n-1}]}, (f(\sigma), \vec{h}) \models \psi_n$;
- if $(\sigma \Sigma \times) \in \Gamma$ then if $\Sigma = [\psi_1, \dots, \psi_n]$, we have that either $M, f(\sigma) \not\models \psi_1$ or $M_{\psi_1}, (f(\sigma), h) \not\models \psi_2, \dots$, or $M_{[\psi_1, \dots, \psi_{n-1}]}, (f(\sigma), \vec{h}) \not\models \psi_n$;
- if $(\sigma R_a \sigma_1) \in \Gamma$ then $(f(\sigma), f(\sigma_1)) \in R_a$.

We start the construction with the root of the tableau tree carrying $\Gamma_0 = \{(\sigma_0 \sqcap \varphi)\}$ and $f(\sigma_0) = s$ and for all $\sigma \neq \sigma_0$, $f(\sigma)$ is undefined.

We then assume by inductive hypothesis (IH) an already constructed branch and mapping f with the above properties, and demonstrate how this branch can be extended and this f can be expanded, for the example of the tableau rule $(\neg B_a^h)$. We recall the rule:

$$\frac{(\sigma \sqcap h_a)(\sigma \Sigma \neg B_a \varphi)}{\begin{array}{c} (\sigma R_a \sigma_{\text{new}}) \\ (\sigma_{\text{new}} \Sigma \checkmark) \\ (\sigma_{\text{new}} \Sigma \neg \varphi) \end{array}} (\neg B_a^h)$$

Let Γ be the set of tableau terms that we reach with the already constructed branch. In order to apply rule $(\neg B_a^h)$ we may suppose (IH) that $(\sigma \sqcap h_a), (\sigma \Sigma \neg B_a \varphi) \in \Gamma$ and that $f(\sigma)$ is defined such that $M, f(\sigma) \models h_a$ and $M_\Sigma, (f(\sigma), \vec{h}) \models \neg B_a \varphi$. According to the truth conditions, this means that there exists an a -successor of $(f(\sigma), \vec{h})$ in M_Σ satisfying $\neg \varphi$. As $(f(\sigma), \vec{h}) \models h_a$, this successor is of the form (t, \vec{h}) . We define $f(\sigma_{\text{new}}) = t$ and we extend Γ to

$$\Gamma \cup \{(\sigma R_a \sigma_{\text{new}}), (\sigma_{\text{new}} \Sigma \checkmark), (\sigma_{\text{new}} \Sigma \neg \varphi)\}.$$

The proof is similar for other tableaux rules.

Observe that in this process of branch extension we can never match the numerator of a rule that has an inconsistency in its denominator. For instance, take the rule (\perp) and suppose that wish to extend a branch with $(\sigma \Sigma q), (\sigma \Sigma \neg q) \in \Gamma$. By induction we already have that $M_\Sigma, f(\sigma) \models q$ and $M_\Sigma, f(\sigma) \models \neg q$, which is in contradiction to our inductive assumption of satisfiability.

We conclude that there exists a tableau for φ with an open branch.

◀

Given an open branch, we construct a model M where states are the nodes σ , relations are inferred from terms of the form $(\sigma R_a \sigma_1)$, and valuations are inferred from terms of the form $(\sigma \sqcap p)$ and $(\sigma \sqcap \neg p)$. We then prove by induction over the complexity of Σ and ψ that, if $(\sigma \Sigma \checkmark)$ appears on the branch, then $(\sigma \Sigma \psi)$ is in the branch iff $M_\Sigma, (\sigma, \vec{h}) \models \psi$.

For instance, suppose $(\sigma \Sigma \neg B_a \varphi)$ and $(\sigma \sqcap h_a)$ appear in the branch. Then rule $(\neg B_a^h)$ has been applied. Therefore, we find terms $(\sigma R_a \sigma_{\text{new}}), (\sigma_{\text{new}} \Sigma \checkmark), (\sigma_{\text{new}} \Sigma \neg \varphi)$ on the

branch. By induction hypothesis, $M_\Sigma, (\sigma_{\text{new}}, \vec{h})$ is well-defined and $M_\Sigma, (\sigma_{\text{new}}, \vec{h}) \models \neg\varphi$. By the construction of M , $(\sigma, \sigma_{\text{new}}) \in R_a$, so that $M_\Sigma, (\sigma, \vec{h}) \models \neg B_a\varphi$.

We conclude that, when $(\sigma \sqcup \varphi)$ appears on the branch, then $M, \sigma \models \varphi$, i.e., φ is satisfiable. \square

Proposition 15

Satisfiability of \mathcal{L} formulas in the class of \mathcal{K}_n models is PSPACE-complete. \dashv

Proof As explained in [11], a tableau method leads to a procedure using a polynomial amount of space in the input formula length, if we can apply the rules by only keeping in memory the content of a branch. In our case the argument is essentially the same (see also [3]): we only keep in memory the information concerning the current node and its path to the root node, in order to be able to backtrack. We moreover restrict the applicability of the (*hear*) rule to the atomic propositions h_a for which a occurs in the input formula. The satisfiability problem is PSPACE-hard, because we can reduce polynomially the satisfiability problem for \mathbf{K}_n (the multi-agent version of the minimal modal logic \mathbf{K}). \square

We leave open the complexity of satisfiability of the logic $\mathbf{ABAL}^{\text{int}}$, that is interpreted on the class of $\mathcal{K}45_n^h$ models. As the complexity of satisfiability of the underlying epistemic logic $\mathbf{K}45_n$ (for $n \geq 2$) is PSPACE-complete, we conjecture that the complexity of $\mathbf{ABAL}^{\text{int}}$ is also PSPACE-complete. Proving this seemed a highly technical exercise of limited additional value to our current undertaking.

This ends the core presentation of the logic \mathbf{ABAL} . In the remainder of the paper we model attention-based announcements as an *action model* (Section 5), we model *attention change* (Section 6), and the concept of *joint attention* and the related issue of *common belief* (Section 7 and Section 8).

5 Action model for attention-based announcement

Every attention-based announcement is definable as an action model. Whether an announcement φ is heard in a given state, depends on the value of h_a for every agent a in that state. The agents who hear the announcement retain all arrows pointing to states where φ holds and delete all arrows pointing to states where φ does not hold, and that is independent of the truth of φ in the actual state; whereas the agents who do not hear the announcement think that nothing has happened, i.e., also independent of the truth of φ they think that the trivial action with precondition \top happened. We can define such an action model economically, in the sense of producing a resulting model with a minimal duplication of states into bisimilar states. Still, it is a fairly big model, exponential in the number of agents.

In Definition 20 we give a translation from attention-based announcement logic \mathbf{ABAL} into the action model logic of [5]. Definition 19 spells out the inductive clause for attention-based announcement of that translation, wherein the action model for attention-based

announcement is constructed. Prior to that, we first define the relevant technical machinery for action models.

An *action model* is a structure like epistemic model but with a precondition function instead of a valuation function. We recall (def. 2) that \mathcal{L}^- is the language \mathcal{L} without the inductive construct for announcement.

Definition 16 (Action model)

An *action model* $\mathbf{M} = (\mathbf{S}, \mathbf{R}, \mathbf{pre})$ consists of a *domain* \mathbf{S} of *actions*, an *accessibility function* $\mathbf{R} : \mathbf{AGT} \rightarrow \mathcal{P}(\mathbf{S} \times \mathbf{S})$, where each \mathbf{R}_a is an accessibility relation, and a *precondition function* $\mathbf{pre} : \mathbf{S} \rightarrow \mathcal{L}^-$. A pointed (and possibly multi-pointed) action model (\mathbf{M}, \mathbf{s}) , where $\mathbf{s} \in \mathbf{S}$, is an *epistemic action*. ⊣

Performing an epistemic action in an epistemic state means computing what is known as their restricted modal *product*. This product models the new state of information.

Definition 17 (Update of an epistemic attention model with an action model)

Given epistemic attention model $M = (S, R, V)$ and action model $\mathbf{M} = (\mathbf{S}, \mathbf{R}, \mathbf{pre})$, their *update* (product) $M \otimes \mathbf{M}$ is the epistemic attention model $M \otimes \mathbf{M} = (S', R', V')$ such that

$$\begin{aligned} S' &= \{(t, \mathbf{t}) \mid M, t \models \mathbf{pre}(\mathbf{t})\} \\ ((t, \mathbf{t}), (t', \mathbf{t}')) \in R'_a &\text{ iff } (t, t') \in R_a \text{ and } (\mathbf{t}, \mathbf{t}') \in \mathbf{R}_a \\ (t, \mathbf{t}) \in V'(q) &\text{ iff } t \in V(q) \quad \text{for all } q \in \mathbf{ATM} \cup H \end{aligned}$$

If $s \in S$, $\mathbf{s} \in \mathbf{S}$, and $M, s \models \mathbf{pre}(\mathbf{s})$, then $(M \otimes \mathbf{M}, (s, \mathbf{s}))$ is the update of pointed model (M, s) with epistemic action (\mathbf{M}, \mathbf{s}) . ⊣

The domain of $M \otimes \mathbf{M}$ is the product of the domains of M and \mathbf{M} , but restricted to state/action pairs (t, \mathbf{t}) such that $M, t \models \mathbf{pre}(\mathbf{t})$, i.e., such that the action can be executed in that state. An agent cannot distinguish pair (t, \mathbf{t}) from pair (t', \mathbf{t}') in the next epistemic state if she cannot distinguish states t and t' in the initial epistemic state and also cannot distinguish action \mathbf{t} (that is executed in t) from \mathbf{t}' (that is executed in t').

In the logical language we can associate a dynamic operator with the execution of an epistemic action, similar to the dynamic operator for a truthful public announcement in the state elimination setting: $[\mathbf{M}, \mathbf{s}]\varphi$ means that after every execution of epistemic action (\mathbf{M}, \mathbf{s}) , φ is true. The model \mathbf{M} occurring in modalities $[\mathbf{M}, \mathbf{s}]$ needs to have a finite domain (for linguistic reasons). The semantics of this modality is then as follows.

Definition 18 (Action model logic)

Let \mathcal{L}^\otimes be the logical language with the inductive construct $[\mathbf{M}, \mathbf{s}]\varphi$ instead of $[\varphi]\varphi$. Then:

$$M, s \models [\mathbf{M}, \mathbf{s}]\varphi \quad \text{iff} \quad M, s \models \mathbf{pre}(\mathbf{s}) \text{ implies } (M \otimes \mathbf{M}), (s, \mathbf{s}) \models \varphi \quad \text{⊣}$$

Let $S' \subseteq S$. By abbreviation we define $[\mathbf{M}, S']\varphi$ as $\bigwedge_{s \in S'} [\mathbf{M}, \mathbf{s}]\varphi$ for the multi-pointed action model (\mathbf{M}, S') , and $[\mathbf{M}]\varphi$ as $\bigwedge_{s \in S} [\mathbf{M}, \mathbf{s}]\varphi$.

Definition 19 (Action model \mathbf{A}_φ)

Given a formula $\varphi \in \mathcal{L}^\otimes$, the action model $\mathbf{A}_\varphi = (\mathbf{S}, \mathbf{R}, \mathbf{pre})$ is defined as:

- $S = \{(i, J) \mid i \in \{0, 1\} \text{ and } J \subseteq AGT\} \cup \{s_\top\}$;

- R maps each agent $a \in AGT$ to

$$R_a = \{((i, J), (1, K)) \mid i \in \{0, 1\}, J, K \subseteq AGT \text{ and } a \in J\} \cup \{((i, J), s_\top) \mid J \subseteq AGT \text{ and } a \notin J\} \cup \{(s_\top, s_\top)\};$$

- $\text{pre} : S \rightarrow \mathcal{L}$ is defined as follows, given $J \subseteq AGT$:

- $\text{pre}((0, J)) = \neg\varphi \wedge \bigwedge_{a \in J} h_a \wedge \bigwedge_{a \notin J} \neg h_a$;
- $\text{pre}((1, J)) = \varphi \wedge \bigwedge_{a \in J} h_a \wedge \bigwedge_{a \notin J} \neg h_a$;
- $\text{pre}(s_\top) = \top$;

⊢

Definition 20 (From attention-based announcements to action models)

We define a recursive translation $\text{tr} : \mathcal{L} \rightarrow \mathcal{L}^\otimes$. All clauses except the one for announcement are trivial. In the clause for announcement, $S^- = S \setminus \{s_\top\}$.

$$\begin{array}{lll} \text{tr}(q) = q & \text{tr}(\neg\varphi) = \neg\text{tr}(\varphi) & \text{tr}(\varphi \wedge \psi) = \text{tr}(\varphi) \wedge \text{tr}(\psi) \\ \text{tr}(B_a\varphi) = B_a(\text{tr}(\varphi)) & \text{tr}([\psi]\varphi) = [A_{\text{tr}(\psi)}, S^-]\text{tr}(\varphi) & \end{array} \quad \dashv$$

Definition 21 (Action model for attention-based announcement)

The multi-pointed action model $(A_{\text{tr}(\varphi)}, S^-)$ is the *action model for the attention-based announcement of φ* . ⊢

Given $|AGT| = n$ agents, the action model for the attention-based announcement of φ consists of $2^{n+1}+1$ actions and has 2^{n+1} designated points. (The set of designated points is the second argument in a pair that constitutes a multi-pointed action model.) The precondition for any of these points lists whether the announced φ is true or false and which agents are attentive and which agents not. Moreover, there is a ‘nothing happens’ action s_\top with precondition \top . This action cannot be a designated point: we assume that the announcement is in fact made. We recall that in our approach, announcements can be made independently from the truth of the announced formula. Therefore, φ may be true but it may also be false. However, as the announcement is believed (to be true), an attentive agent believes that any action with precondition entailing φ may be the actual action. (‘Any’ action: including those wherein the agent is not paying attention. This seems inconsistent with the requirement of attention introspection. But in fact it does not matter, as in that case the agent only considers states possible in the epistemic attention model that satisfy awareness introspection, a property that is preserved after announcement — Prop. 7.) But an inattentive agent believes that the action with precondition \top is the actual action: she believes that no announcement was made.

This action model matches nicely with the semantics of $[\varphi]\psi$ which, given a model M , produces a model M_φ twice the size of M , namely consisting of a unchanged copy of M plus copy M' of M wherein the attentive agents restrict their accessibility to the φ states.




Figure 4: The action model \mathbf{A}_p corresponding to an attention-based announcement p to two agents a and b . The model consists of nine actions, represented by their preconditions only. The six arrows pointing to the dashed box point to all four actions in the box.

As all the 2^{n+1} different preconditions of actions (i, J) in the action model are exclusive, the product of that entire part of the action model produces a model of the same size as M , but with merely some removed arrows, whereas the single action s_{\top} with precondition \top produces the other, trivial copy.

Example 22

For the example of two agents a and b and the announcement p , the action model for attention-based announcements is depicted in Figure 4. When p is false, h_a is true, and h_b is false, then agent a hears the announcement p and believes it to be true, therefore she believes the real action to be the one where p is true — regardless of the values of h_a and h_b in states wherein it can be executed. So that makes for *four* arrows. On the other hand, agent b does not hear the announcement and believes that nothing at all happens: therefore there is a single arrow to the alternative with precondition \top , and with reflexive arrows for a and b : b believes, incorrectly, that both agents believe that nothing happened. \dashv

Proposition 23

Let M be an epistemic attention model, and $\varphi \in \mathcal{L}$. Then $M_\varphi \Leftrightarrow M \otimes \mathbf{A}_{\text{tr}\varphi}$. \dashv

Proof The bisimulation relation \mathfrak{R} consists of the following pairs, for all s in M ; and where $J = \{a \in AGT \mid M, s \models h_a\}$:

$$\begin{aligned} & ((s, h) , (s, (0, J))) \\ & ((s, h) , (s, (1, J))) \\ & ((s, \bar{h}) , (s, s^\top)) \end{aligned}$$

□

Proposition 24

Let φ in \mathcal{L} , and let (M, s) be a pointed model. Then $M, s \models \varphi$ iff $M, s \models \text{tr}(\varphi)$. \dashv

Proof The proof is by induction on the structure of φ . In the proof we essentially use that *bisimilarity implies logical equivalence* (*) in either semantics. The non-trivial case of the inductive proof is $[\psi]\varphi$. Other cases are left to the reader. Below, let $J = \{a \in AGT \mid M, s \models h_a\}$. Observe that any $\mathbf{t} \in \mathbf{S}^-$ has shape (i, J') , for $i = 0, 1$ and $J' \subseteq AGT$. The only pairs (i, J') satisfying the precondition in state s are $(i, J') = (0, J)$ and $(i, J') = (1, J)$. This is used in (**).

$$\begin{aligned}
& M, s \models [\psi]\varphi \\
\Leftrightarrow & M_\psi, (s, h) \models \varphi \\
\Leftrightarrow & M \otimes \mathbf{A}_{\text{tr}(\psi)}, (s, (i, J)) \models \varphi, \text{ for } i = 0, 1 && \text{(by Prop. 23, and (*))} \\
\Leftrightarrow & M \otimes \mathbf{A}_{\text{tr}(\psi)}, (s, (i, J)) \models \text{tr}(\varphi), \text{ for } i = 0, 1 && \text{(by induction hypothesis)} \\
\Leftrightarrow & \text{for all } \mathbf{t} \in \mathbf{S}^- : M, s \models \text{pre}(\mathbf{t}) \text{ implies } M \otimes \mathbf{A}_{\text{tr}(\psi)}, (s, \mathbf{t}) \models \text{tr}(\varphi) && \text{(**)} \\
\Leftrightarrow & M, s \models [\mathbf{A}_{\text{tr}(\psi)}, \mathbf{S}^-]\text{tr}(\varphi) \\
\Leftrightarrow & M, s \models \text{tr}([\psi]\varphi)
\end{aligned}$$

□

6 Attention change

Example 25

Consider agents a and b that are both not paying attention, and that are also both uncertain whether the other is paying attention, and that are uncertain about p . (Cf. Example 1.)

A private way to make someone listen to you is to tap on her shoulder before you speak. This only makes that person attentive and not the other agents. So, if I tap on a 's shoulder and then say p , only a and not b will hear it. I can then tap on agent b 's shoulder and say p again, then in the resulting situation a and b both know p , but b does not know that: he still considers it possible that a was not paying attention. On the other hand, a knows that b is paying attention, because she was already paying attention when b was tapped on the shoulder. (Here, we assume that agents that are paying attention hear what is being said, but also observe other public actions.) An epistemic attention model with the same information content (i.e., bisimilar) would have resulted if I had first tapped a and then b , and only then had said p (i.e., once only).

If I first tap b 's shoulder and then a 's shoulder, then again both a and b will now hear the subsequent announcement of p . In this case we have the dual situation where a considers it possible that b is not paying attention when the announcement was made. So the result is again that both know that p , but now a does not know that.

In contrast, if you clap your hands, this is a public way to get attention. After clapping hands, everybody is paying attention, and everybody knows that everybody is paying

attention, etc. After clapping hands in the example, a and b are both attentive but now also know that the other is attentive (and so on). When p is now announced, again they both know that p afterwards, but they also know that they know that. In fact, after clapping hands they have *joint attention*, such that after the announcement of p they have *common knowledge* (correct common belief) of p . Joint attention and common belief will be treated in the next sections. \dashv

Drawing inspiration from [22, 20], we model such attention change by an assignment. Given a set of agents $A \subseteq AGT$, we distinguish the assignment $+A$ that makes all agents $a \in A$ pay attention and hear subsequent announcements, from an assignment $-A$ that makes all h_a false. Such an assignment $+A$, for $A = \{a_1, \dots, a_n\}$, is merely a convenient shorthand for a simultaneous assignment $h_{a_1} := \top, \dots, h_{a_n} := \top$.

Definition 26 (Attention assignment)

To the inductive definition of the language \mathcal{L} (Def. 2) we add clauses $[+A]\varphi$ and $[-A]\varphi$, for $A \subseteq AGT$. The resulting language is called \mathcal{L}^+ . \dashv

For $[+\{a\}]\varphi$ we write $[+a]\varphi$, and for $[-\{a\}]\varphi$ we write $[-a]\varphi$.

Definition 27 (Semantics of attention assignment)

Let $M = (S, R, V)$ and $\psi \in \mathcal{L}$ be given, and $A \subseteq AGT$. Then

$$\begin{aligned} M, s \models [+A]\psi &\text{ iff } M_{+A}, (s, h) \models \psi \\ M, s \models [-A]\psi &\text{ iff } M_{-A}, (s, h) \models \psi \end{aligned}$$

where $M_{+A} = (S', R', V')$ is defined as follows.

1. $S' = S \times \{h, \bar{h}\}$;
2. if $a \in A$ and $s, t \in S$ then $((s, i), (t, j)) \in R'_a$ iff $(s, t) \in R_a$ and
 - (a) $i=h$ and $j=h$; or
 - (b) $i=\bar{h}$ and $j=\bar{h}$;
3. if $a \notin A$ and $s, t \in S$ then $((s, i), (t, j)) \in R'_a$ iff $(s, t) \in R_a$ and
 - (a) $i=h, j=h$, and $(M, s) \models h_a$; or
 - (b) $i=h, j=\bar{h}$, and $(M, s) \not\models h_a$; or
 - (c) $i=\bar{h}$ and $j=\bar{h}$;
4. for $p \in ATM$ and $s \in S$: $(s, h) \in V'(p)$ iff $s \in V(p)$, and $(s, \bar{h}) \in V'(p)$ iff $s \in V(p)$;
5. if $a \in A$ then
 - (a) $(s, h) \in V'(h_a)$, and
 - (b) $(s, \bar{h}) \in V'(h_a)$ iff $s \in V(h_a)$.

6. if $a \notin A$ then

- (a) $(s, h) \in V'(h_a)$ iff $s \in V(h_a)$, and
- (b) $(s, \bar{h}) \in V'(h_a)$ iff $s \in V(h_a)$;

The definition of M_{-A} is similar. We call the resulting logics ABAL_{ch} and $\text{ABAL}_{\text{ch}}^{\text{int}}$. ⊢

In the case of the singleton attention assignment $+a$, agent a will now pay attention in the h -copy of the initial model and may or may not be paying attention in the \bar{h} -copy (that copies the prior information state). If another agent b was already paying attention he will now know that a is now paying attention (clause 6a above, for agents not paying attention, wherein arrows point to other h -worlds); else his knowledge of a 's attention span is as before (clause 6b above). The only factual information change takes place in the h -copy, and only for h_a . This is the part $(s, h) \in V'(h_a)$ in clause 5a, i.e., h_a is now true for all states s in the h -copy.

Proposition 28

Attention assignment preserves attention introspection. ⊢

Proof Obvious. □

The order matters in successive attention change, and this is different again from simultaneous attention change. Recalling the introductory Example 25 in this section, we have that neither $[+a][+b]\varphi \leftrightarrow [+b][+a]\varphi$ nor $[+ab]\varphi \leftrightarrow [+b][+a]\varphi$ are valid.

Just as for attention-based announcements, an attention assignment corresponds to an action model that is a function of who is paying attention and the assignment. The logic to which attention assignment has been added can therefore also easily be axiomatized. This inspired the reduction axioms for $[+A]$ and $[-A]$ assignments in the axiomatization in Table 2.

Just as in the axiomatization for ABAL , we do not have interaction axioms between different dynamic modalities, but instead we use the rule of substitution of equivalents. Otherwise, in order to obtain completeness by rewriting, we would have needed no less than 8 additional reduction axioms, namely for $[+A][+A']\varphi$, $[+A][-A']\varphi$, $[-A][+A']\varphi$, $[-A][-A']\varphi$, $[+A][\psi]\varphi$, $[-A][\psi]\varphi$, $[\psi][+A']\varphi$, $[\psi][-A']\varphi$. And it is not even clear if any of those are straightforwardly formulated with $[+A]$, $[-A]$, and $[\psi]$ modalities (see above for some invalid principles). This motivated us to prefer the version of the axiomatization with substitution of equivalents.

Definition 29 (Axiomatizations ABAL_{ch} and $\text{ABAL}_{\text{ch}}^{\text{int}}$)

The axiomatization for ABAL_{ch} consists of the rules and axioms of ABAL plus the axioms in Table 2. Similarly, we get the axiomatization for $\text{ABAL}_{\text{ch}}^{\text{int}}$ by adding these axioms to those for ABAL^{int} . ⊢

Theorem 30

ABAL_{ch} and $\text{ABAL}_{\text{ch}}^{\text{int}}$ are sound and complete. ⊢

Proof Straightforward. □

$[+A]p$	\leftrightarrow	p	if $p \in ATM$
$[-A]p$	\leftrightarrow	p	if $p \in ATM$
$[+A]h_a$	\leftrightarrow	$\begin{cases} \top & \text{if } a \in A \\ h_a & \text{if } a \notin A \end{cases}$	
$[-A]h_a$	\leftrightarrow	$\begin{cases} \perp & \text{if } a \in A \\ h_a & \text{if } a \notin A \end{cases}$	
$[+A]\neg\varphi$	\leftrightarrow	$\neg[+A]\varphi$	
$[-A]\neg\varphi$	\leftrightarrow	$\neg[-A]\varphi$	
$[+A](\varphi \wedge \psi)$	\leftrightarrow	$[+A]\varphi \wedge [+A]\psi$	
$[-A](\varphi \wedge \psi)$	\leftrightarrow	$[-A]\varphi \wedge [-A]\psi$	
$[+A]B_a\varphi$	\leftrightarrow	$\begin{cases} B_a[+A]\varphi & \text{if } a \in A \\ (h_a \rightarrow B_a[+A]\varphi) \wedge (\neg h_a \rightarrow B_a\varphi) & \text{if } a \notin A \end{cases}$	
$[-A]B_a\varphi$	\leftrightarrow	$\begin{cases} B_a[-A]\varphi & \text{if } a \in A \\ (h_a \rightarrow B_a[-A]\varphi) \wedge (\neg h_a \rightarrow B_a\varphi) & \text{if } a \notin A \end{cases}$	

Table 2: The axiomatizations $ABAL_{ch}$ and $ABAL_{ch}^{int}$.

Embedding private announcement and public announcement We now apply the operators for attention-based announcement and change of attention to provide an embedding of Gerbrandy’s already mentioned logic of private announcements (or *conscious updates*) [9] and of (believed) public announcement logic into ABAL. Gerbrandy’s logic has constructs with dynamic modalities $\llbracket\varphi\rrbracket_A\psi$, for $A \subseteq AGT$, expressing that ψ is the case after the members of group A learn that φ , where the agents outside A believe that nothing happens. The arrow eliminating public announcement $\llbracket\varphi\rrbracket$ is the special case of the ‘private’ announcement to all agents. The crucial reduction axiom is

$$\llbracket\varphi\rrbracket_A B_a\psi \leftrightarrow \begin{cases} B_a\psi & \text{if } a \notin A \\ B_a(\varphi \rightarrow \llbracket\varphi\rrbracket_A\psi) & \text{if } a \in A \end{cases}$$

The semantics of $\llbracket\varphi\rrbracket_A\psi$ is given in terms of non-wellfounded sets; this is somewhat cumbersome, but fortunately we have an alternative action model modelling at our disposal. The most suggestive way to model private announcements (namely suggesting action models for attention-based announcements) is as the three-action action model in Figure 5 (see also [23, Section 6.9]).

Definition 31 (Embedding of private announcements)

Let \mathcal{L}^G (G for Gerbrandy) be the language \mathcal{L} of ABAL with the construct $[\varphi]\varphi$ replaced by $\llbracket\varphi\rrbracket_A\varphi$ (for all $A \subseteq AGT$). Define the embedding tr of $\mathcal{L}^G(ATM, AGT)$ into $\mathcal{L}(ATM \cup H, AGT)$ as follows : $\text{tr}(p) = p$, $\text{tr}(\neg\varphi) = \neg\text{tr}(\varphi)$, $\text{tr}(\varphi \wedge \psi) = \text{tr}(\varphi) \wedge \text{tr}(\psi)$, $\text{tr}(B_a\varphi) = B_a\text{tr}(\varphi)$, and $\text{tr}(\llbracket\varphi\rrbracket_A\psi) = [+A][-(AGT \setminus A)][\text{tr}(\varphi)]\text{tr}(\psi)$. ⊣




Figure 5: Action model for a private announcement of φ to agents $a \in A$, while agents $b \in AGT \setminus A$ think that nothing happened. The actions are labelled with their preconditions.

Proposition 32

Let (M, s) be a pointed epistemic attention model (for $ATM \cup H$ and AGT), and $\varphi \in \mathcal{L}^G(ATM)$. Then $(M, s) \models \varphi$ iff $(M, s) \models \text{tr}(\varphi)$. \dashv

Proof After the agents in A have been made attentive, the other agents are made inattentive. Therefore, the attentive agents in A notice this. We now have the setting of Gerbrandy, wherein the agents in A (correctly) assume that the other agents do not receive the private information φ . Instead of a detailed proof involving non-wellfounded set semantics, observe that Figure 5 results if in Figure 4 we only consider the alternatives with preconditions $\neg\varphi \wedge h_a \wedge \neg h_b$ and $\varphi \wedge h_a \wedge \neg h_b$ (and \top). \square

Corollary 33

Let $\varphi, \psi \in \mathcal{L}^-$, then $\llbracket \varphi \rrbracket_A \psi$ is equivalent to $[+A][-(AGT \setminus A)]\llbracket \varphi \rrbracket \psi$, and (the case of arrow eliminating public announcement) $\llbracket \varphi \rrbracket \psi$ is equivalent to $[+AGT]\llbracket \varphi \rrbracket \psi$. \dashv

We emphasize that $\text{tr} : \mathcal{L}^G(ATM, AGT) \rightarrow \mathcal{L}(ATM \cup H, AGT)$ is more an embedding than a translation, as the sets of atomic propositions in the respective logics are different. More precisely, extending the input of tr to $\mathcal{L}^G(ATM \cup H, AGT)$ makes it incorrect. For example, if one were to do so, consider two agents a, b and one variable p . Then $\text{tr}(h_a \wedge \llbracket p \rrbracket_b h_a) = h_a \wedge [+b][-a][p]h_a$. But the former is now no longer equivalent to the latter: consider a model satisfying h_a . After private announcement of p to b , h_a remains true. But after a is made inattentive, h_a is false.

Still, in the setting of Proposition 32, (M, s) can be a model for atoms $ATM \cup H$ and agents AGT , even with $\varphi \in \mathcal{L}^G(ATM, AGT)$. All values for variables $h_a \in H$ in any state of M are ‘don’t care’ values, as they are reset on the entire model by the $+A$ and $-A$ assignments.

7 Joint attention

An interesting justification of attention-based announcement logic is that it allows us to formalize a concept that has been widely discussed in the philosophical and in the cognitive science literature, namely the concept of *joint attention* or *joint attentional state* [19, 16, 7].

Let $A \subseteq AGT$. The idea is that the agents in A have *joint attention* (or are in a joint attentional state) if and only if they are looking at the source of information *together*, that is to say, every agent in A is looking at the source of information, every agent in A believes that every agent in A is looking at the source of information, and so on. More concisely, the agents in A are in a joint attentional state if and only if each of them is looking at the source of information and focusing his attention on it and they have common belief that each of them is looking at the source of information and focusing his attention on it.

We will only introduce common belief in the next section, but we may as well already informally assume a common belief operator C_A for any subgroup A of the set AGT of all agents, such that $C_A\varphi$ stands for ‘the agents in group A commonly believe φ ’, and which is interpreted in the usual way by the transitive closure of the union of all accessibility relations R_a for the agents in A .

Definition 34 (Joint attention)

Given $M = (S, R, V)$, the model M satisfies *joint attention* for set $A \subseteq AGT$ of agents if $V(h_A) = S$. The formula JointAtt_A is defined as $h_A \wedge C_A h_A$. ⊣

Clearly, a model satisfying joint attention for A makes JointAtt_A true, and if JointAtt_A is valid on a model then it satisfies joint attention. We further have (as a generalization of Prop. 7) that

Proposition 35 (Preservation of joint attention)

If M satisfies joint attention for $A \subseteq AGT$, then M_φ satisfies joint attention for A . ⊣

We also have that, modulo an otherwise trivial translation (compare to Corollary 33):

Proposition 36

Let M satisfy joint attention for AGT , and s in the domain of M . Then $M, s \models [\varphi]\psi$ iff $M, s \models \llbracket \varphi \rrbracket \psi$. ⊣

Attention introspection for all agents $a \in A$, with corresponding axiom $\bigwedge_{a \in A} (h_a \rightarrow B_a h_a)$, is insufficient to obtain joint attention for A : it may be that h_A is true while some $a \in A$ does not believe that h_A .

As pointed out by [16, 7], joint attention explains the genesis of common beliefs in the context of social interaction. Such genesis is often considered as related to public events in the sense that a common belief is either a consequence of an event whose occurrence is so evident (viz. public) that agents cannot but recognize it — as when, during a soccer match, players mutually believe that they are playing soccer — or the product of a communication process — as when the referee publicly announces that one player is expelled. From there on each player believes that each other player believes and so on that one of them has been expelled. Intuitively, an event is considered public if its occurrence is epistemically accessible by everybody such that it becomes common belief between them. But what are the intuitive conditions that make an event public? What are the reasons to believe that an occurring event is commonly believed? In a normal situation (what is announced is true, there is no noise in the communication channel, etc.) looking at the source of information

and having a common belief that everyone is looking at the source of information (i.e., being in a joint attentional state) provide a sufficient condition for the formation of a common belief. Let us (avoiding Moorean phenomena) restrict ourselves to propositional variables. We then have that

$$\models \text{JointAtt}_A \rightarrow [p]C_{AP}$$

but, as we already know, we also have

$$\not\models [p]C_{AP}$$

This now begs the question how we can obtain common knowledge for A other than when all agents in A are attentive? This is indeed possible, even when some in A are attentive and others in A are inattentive. The next section provides a full answer to that question.

8 Common belief

8.1 Attentive relativized common belief

A well-known issue with the logic of public announcements is that the addition of common belief (common knowledge) makes the axiomatization more involved, because the interaction between announcement and common belief cannot be described in an axiom, but only in a derivation rule [5]. The elimination of announcements by reduction no longer works (and indeed, the logic of public announcements with common belief is more expressive than the logic of public announcements without common belief). A conceptual and technical innovation addressing this issue is the notion of relativized common belief: if one not merely considers what is true along all accessibility paths labelled with the believing agents, but only considers such paths that satisfy a relativization condition, then there is a reduction axiom again, and public announcements can be eliminated from the logic of public announcement with relativized common belief [20]. Of course, this is at the prize of higher expressivity of that announcement-free logic.

Following the previous section on joint attention, it is clear that a proper treatment of common belief would make the attention-based announcement logic more valuable to analyze phenomena in cognitive science, philosophy, and multi-agent systems. Semantically, this does not pose any problem. But the question then is how to axiomatize such a logic. We were not successful with the two obvious ways out of this predicament: finding an axiomatization with only the addition of common belief (analogous to [5]), or finding an axiomatization with only the addition of relativized common belief (analogous to [20]). On first sight, such roads seem well-trod, as their methods not merely apply to public announcement logic but also to action model logic, and in Section 5 we modelled attention-based announcement as an action model. So, it seemed that we merely needed to plug in that action model, and either compute the derivation rule (for common belief after attention-based announcement), or the reduction (for relativized common belief after attention-based announcement), where the latter would be clearly preferable over the former (as it is then typically easier to show completeness).

This method turned out to be not as straightforward as we thought, because the action model for attention-based announcement is rather large (we recall that it is exponential in the number of agents, see Section 5). The already not so intuitive derivation rule for common belief after the action model modality of [5] would then become really complex. However, the more obvious relativized common belief axiom was also evasive, as the role of transformation matrices in the rewriting techniques of [20] proved a stumbling block as well for our ‘very large’ action models. Something comes out, but for it to make sense, you need a computer.

Motivated by the relativized common belief by van Benthem et al. in [20], we here provide an adaptation of that notion that also takes into account whether agents are paying attention. With that adaptation, we regain, we think, some elegance. The interaction between attention-based announcement and common belief is then a different case.

Definition 37 (Languages with common belief)

The language \mathcal{L}_C of ABAL with *common belief* is obtained by adding to the logical language \mathcal{L} (Def. 2) inductive constructs $C_A\varphi$ and $\mathbf{C}_A^\varphi\varphi$, for any $A \subseteq AGT$. Without announcements we get the language \mathcal{L}_C^- . \dashv

For $C_A\varphi$, read ‘the agents in A commonly believe φ ’. For $\mathbf{C}_A^\chi\varphi$, read ‘relative to χ , the attentive agents in A commonly believe φ ’.

Prior to defining the semantics, let R^+ be the transitive closure of a relation R . Further, given a model (S, R, V) , let $\mathbf{R}_a^\chi = \{(s, t) \in R_a \mid M, s \models h_a \text{ and } M, t \models \chi\}$. Now define $\mathbf{R}_A^\chi = \bigcup_{a \in A} \mathbf{R}_a^\chi$.

For more notational flexibility in this section, instead of $(x, y) \in R$ we also write $y \in xR$ or xRy , and if M_φ is a model resulting from attention-based announcement of φ , we assume it to have shape $M_\varphi = (\varphi S, \varphi R, \varphi V)$. This unorthodox notation is an advantage in constructs like ${}_\varphi\mathbf{R}_a^\chi$; which stands for the ‘attentive’ relativized accessibility relation that is based on the accessibility relation ${}_\varphi R_a$ in the model M_φ .

Definition 38 (Semantics of common belief)

$$M, s \models \mathbf{C}_A^\chi\varphi \text{ iff } M, t \models \varphi \text{ for all } (s, t) \in (\mathbf{R}_A^\chi)^+ \quad \dashv$$

Given a group of agents $A \subseteq AGT$, *common belief* $C_A\psi$ is the fixpoint of *shared belief* $E_A\psi = \bigwedge_{a \in A} B_a\psi$ and we recall from [20] that, given a formula χ , *relativized common belief* $\mathbf{C}_A^\chi\varphi$ is the fixpoint of *relativized shared belief* $E_A(\chi \rightarrow \psi)$ (i.e., $\bigwedge_{a \in A} B_a(\chi \rightarrow \psi)$). In the same way, we can define *attentive relativized shared belief* by abbreviation as

$$\mathbf{E}_A^\chi\varphi = \bigwedge_{a \in A} (h_a \rightarrow B_a(\chi \rightarrow \varphi))$$

and then we can see *attentive relativized common belief* among the agents in A as the fixpoint of \mathbf{E}_A^χ :

$$\mathbf{C}_A^\chi\varphi = \mathbf{E}_A^\chi\varphi \wedge \mathbf{E}_A^\chi\mathbf{E}_A^\chi\varphi \wedge \dots$$

The semantics of attentive relativized shared belief therefore is: $M, s \models \mathbf{E}_A^\chi\varphi$ iff $M, t \models \varphi$ for all $(s, t) \in \mathbf{R}_A^\chi$.




Figure 6: Visualization of attentive relativized common belief

Example 39 (Attentive common belief)

Figure 6 illustrates the semantics of attentive relativized common belief. The formula $\mathbf{C}_{ab}^p q$ (short for $\mathbf{C}_{\{a,b\}}^p q$) is true in the root of the model. We can verify this by observing that q is true in all grey nodes. Each arrow in one of the grey paths satisfies that the agent who labels the link is attentive at the start of the arrow and that the relativization condition p is true at the end of the arrow. \dashv

Let us now give an impression of the consequences of this semantics. Clearly we do not have that $[p]C_{AGT}p$ is valid for any atomic proposition p , because agents not paying attention will not have heard the announcement of p and may still be uncertain about p . But we have that $[p]\mathbf{C}_{AGT}^\top p$ is valid. For this, see Example 40. Example 41, building on Example 8, demonstrates where $[p]C_{AGT}p$ can be false and $[p]\mathbf{C}_{AGT}^\top p$ true. The difference is, because the semantics of \mathbf{C}_{AGT}^\top takes attention into account, while that of C_{AGT} does not.

Example 40

We show that $\models [p]\mathbf{C}_{AGT}^\top p$. Let (M, s) be an epistemic attention model. Note that $M, s \models [p]\mathbf{C}_{AGT}^\top p$ iff $M_p, (s, h) \models \mathbf{C}_{AGT}^\top p$. We show the latter, i.e., for any $(t, i) \in (s, h)({}_p\mathbf{R}_{AGT}^\top)^+$: $M_p, (t, i) \models p$. This follows from the stronger claim that for all $n \in \mathbb{N}$:

$$\text{For all } (t, i) \in (s, h)({}_p\mathbf{R}_{AGT}^\top)^n: M_p, (t, i) \models p \text{ and } i = h.$$

Proof by induction on n : Let $(t, i) \in (s, h)({}_p\mathbf{R}_a^\top)$, for some $a \in AGT$. This implies $M, s \models h_a$ and $M, t \models p$. As $M, s \models h_a$, i must be h ; while $M, t \models p$, together with $M, t \models p$ iff

$M_p, (t, h) \models p$, delivers $M_p, (t, h) \models p$ and thus completes the proof. The inductive case is similar. \dashv

Example 41

We build on Example 8 (page 6), illustrated in Figure 1. Consider the state in model M in Figure 1 that is labeled $\{p, h_a, h_b\}$. We denote this state as s . We note that $AGT = \{a, b\}$. In M_p , resulting from the attention-based announcement that p , we have the path

$$\{p, h_a, h_b\} \longrightarrow_b \{p, \neg h_a, h_b\} \longrightarrow_a \{\neg p, \neg h_a, h_b\}$$

where the former two are in the h part of M_p and the latter is in the \bar{h} part. From this, it follows that $M_p, (s, h) \not\models C_{AGT}p$ and thus $M, s \not\models [p]C_{AGT}p$. \dashv

8.2 Axiomatizations with common belief

The axioms and rules for the logic ABALC, i.e., attention-based announcement logic with common belief and attentive relativized common belief, are those of ABAL together with the axioms of Table 3. The axiomatization for the logic KC, that is, multi-agent modal logic with attentive relativized common belief, consists of the part of ABALC that does not contain announcement modalities. In the axiomatization, we have also used the abbreviation: $\bar{\mathbf{E}}_A\psi = \bigwedge_{a \in A} (\neg h_a \rightarrow B_a\psi)$, which stands for shared knowledge among the non-attentive agents in A of ψ . (With an obvious associated accessibility relation $\bar{\mathbf{R}}_A$ such that $M, s \models \bar{\mathbf{E}}_A\psi$ iff $M, t \models \psi$ for all $(s, t) \in \bar{\mathbf{R}}_A$.) Among the axioms, we have the usual Mix and Induction axioms characterizing the fixpoint modalities \mathbf{C}_A^x and C_A . Prior to demonstrating the soundness of the axiomatization, let us explain informally the two axioms for common belief after announcement, starting with the one for attentive relativized common belief.

$$[\varphi]\mathbf{C}_A^x\psi \leftrightarrow \mathbf{C}_A^{\varphi \wedge [\varphi]\chi}[\varphi]\psi$$

The relativized common belief operator only goes through paths where agents are attentive: we call these paths *attentive paths*. (More precisely, a path s_1, \dots, s_n is *attentive* iff $M, s_i \models h_a$ for all links $(s_i, s_{i+1}) \in R_a$ in the path.) After announcing φ , the relativized common belief operator goes through all attentive paths satisfying χ in the updated models. These paths are exactly the attentive paths satisfying $\varphi \wedge [\varphi]\chi$ in the initial model.

Now consider the axiom for common belief after an announcement.

$$[\varphi]C_A\psi \leftrightarrow \mathbf{C}_A^\varphi([\varphi]\psi \wedge \bar{\mathbf{E}}_A(\psi \wedge C_A\psi)) \wedge \bar{\mathbf{E}}_A(\psi \wedge C_A\psi)$$

This describes that in order for A to obtain common belief of ψ after announcement of φ , two conditions must be satisfied. Firstly, for inattentive agents in A , ψ must already be true and have been common belief before the announcement. Secondly, the attentive agents in A must have common belief of that, and also commonly believe that after the announcement of φ , ψ is true.

So, after all, recalling Section 7, growth of common belief is possible in this logic, even when not all agents are attentive. We consider this an important result.

$\mathbf{C}_A^x \psi$	\leftrightarrow	$\mathbf{E}_A^x(\psi \wedge \mathbf{C}_A^x \psi)$	Mix \mathbf{C}
$\mathbf{C}_A^x \psi$	\leftarrow	$\mathbf{E}_A^x \psi \wedge \mathbf{C}_A^x(\psi \rightarrow \mathbf{E}_A^x \psi)$	Induction \mathbf{C}
$\mathbf{C}_A^x \psi$	\leftarrow	$\mathbf{C}_A^x \varphi \wedge \mathbf{C}_A^x(\varphi \rightarrow \psi)$	Distribution \mathbf{C}
$[\varphi] \mathbf{C}_A^x \psi$	\leftrightarrow	$\mathbf{C}_A^{\varphi \wedge [\varphi]x} [\varphi] \psi$	Reduction \mathbf{C} *
$C_A \psi$	\leftrightarrow	$E_A(\psi \rightarrow C_A \psi)$	Mix C
$C_A \psi$	\leftarrow	$E_A \psi \wedge C_A(\psi \rightarrow E_A \psi)$	Induction C
$C_A \psi$	\leftarrow	$C_A \varphi \wedge C_A(\varphi \rightarrow \psi)$	Distribution C
$[\varphi] C_A \psi$	\leftrightarrow	$\mathbf{C}_A^\varphi([\varphi] \psi \wedge \overline{\mathbf{E}}_A(\psi \wedge C_A \psi)) \wedge \overline{\mathbf{E}}_A(\psi \wedge C_A \psi)$	Reduction C *
From φ infer $\mathbf{C}_A^x \varphi$		Necessitation \mathbf{C}	
From φ infer $C_A \varphi$		Necessitation C	

Table 3: Axioms and rules involving common belief for the axiomatizations for ABALC and (without the *-ed axioms) for KC.

Lemma 42

$$\models [\varphi] \mathbf{C}_A^x \psi \leftrightarrow \mathbf{C}_A^{\varphi \wedge [\varphi]x} [\varphi] \psi. \quad \dashv$$

Proof Let (M, s) be given. We first show that:

$$(s, h)(\varphi \mathbf{R}_A^x)^+(t, h) \text{ iff } s(\mathbf{R}_A^{\varphi \wedge [\varphi]x})^+ t \quad (*)$$

For this, it suffices to show that: $(s, h)_\varphi \mathbf{R}_A^x(t, h)$ iff $s \mathbf{R}_A^{\varphi \wedge [\varphi]x} t$; and that follows from: for all $a \in A$: $(s, h)_\varphi \mathbf{R}_a^x(t, h)$ iff $s \mathbf{R}_a^{\varphi \wedge [\varphi]x} t$, shown below. The proof for the transitively closed relations then follows immediately.

$$\begin{aligned} & (s, h)(\varphi \mathbf{R}_a^x)(t, h) \\ \text{iff } & M_\varphi, (s, h) \models h_a, (s, h) R_a^\varphi(t, h), \text{ and } M_\varphi, (t, h) \models \chi \\ \text{iff } & M, s \models h_a, s R_a t, M, t \models [\varphi] \chi, \text{ and } M, t \models \varphi \\ \text{iff } & s \mathbf{R}_a^{\varphi \wedge [\varphi]x} t \end{aligned}$$

We now proceed as follows:

$$\begin{aligned} & M, s \models [\varphi] \mathbf{C}_A^x \psi \\ \text{iff } & M_\varphi, (s, h) \models \mathbf{C}_A^x \psi \\ \text{iff } & \text{for all } (t, h) \in (s, h)(\varphi \mathbf{R}_A^x)^+ : M_\varphi, (t, h) \models \psi \quad (\text{by } (*)) \\ \text{iff } & \text{for all } t \in s(\mathbf{R}_A^{\varphi \wedge [\varphi]x})^+ : M_\varphi, (t, h) \models \psi \quad (\text{by the truth condition of } [\varphi] \psi) \\ \text{iff } & \text{for all } t \in s(\mathbf{R}_A^{\varphi \wedge [\varphi]x})^+ : M, t \models [\varphi] \psi \\ \text{iff } & M, s \models \mathbf{C}_A^{\varphi \wedge [\varphi]x} [\varphi] \psi \end{aligned}$$

□

Proposition 43

$$\models [\varphi]C_A\psi \leftrightarrow \mathbf{C}_A^\varphi([\varphi]\psi \wedge \overline{\mathbf{E}}_A(\psi \wedge C_A\psi)) \wedge \overline{\mathbf{E}}_A(\psi \wedge C_A\psi) \quad \dashv$$

Proof In the proof, the symbol ‘ \cdot ’ stands for relational composition.

$$\begin{aligned} & M, s \models [\varphi]C_A\psi \\ \text{iff} & M_\varphi, (s, h) \models C_A\psi \\ \text{iff} & \text{for all } (t, i) \in (s, h)(\varphi R_A)^+ : M_\varphi, (t, i) \models \psi \\ \text{iff} & \text{for all } (t, h) \in (s, h)(\varphi \mathbf{R}_A^\top)^+ : M_\varphi, (t, h) \models \psi \text{ and} \\ & \text{for all } (t, \bar{h}) \in (s, h)((\varphi \mathbf{R}_A^\top)^* \cdot \varphi \overline{\mathbf{R}}_A \cdot (\varphi R_A)^*) : M_\varphi, (t, \bar{h}) \models \psi \\ \text{iff} & \text{for all } t \in s(\mathbf{R}_A^\varphi)^+ : M_\varphi, (t, h) \models \psi \text{ and} \\ & \text{for all } t \in s((\mathbf{R}_A^\varphi)^* \cdot \overline{\mathbf{R}}_A \cdot (R_A)^*) : M_\varphi, (t, \bar{h}) \models \psi \quad (\text{by } (*) \text{ in Lemma 42}) \\ \text{iff} & \text{for all } t \in s(\mathbf{R}_A^\varphi)^+ : M, t \models [\varphi]\psi \text{ and} \\ & \text{for all } t \in s((\mathbf{R}_A^\varphi)^* \cdot \overline{\mathbf{R}}_A \cdot (R_A)^*) : M, t \models \psi \\ \text{iff} & M, s \models \mathbf{C}_A^\varphi[\varphi]\psi \text{ and } M, s \models \overline{\mathbf{E}}_A(\psi \wedge C_A\psi) \wedge \mathbf{C}_A^\varphi \overline{\mathbf{E}}_A(\psi \wedge C_A\psi) \\ \text{iff} & M, s \models \mathbf{C}_A^\varphi([\varphi]\psi \wedge \overline{\mathbf{E}}_A(\psi \wedge C_A\psi)) \wedge \overline{\mathbf{E}}_A(\psi \wedge C_A\psi) \end{aligned}$$

□

Proposition 44 (Soundness)

The axiomatizations for ABALC and KC are sound. \dashv

Proof The only non-trivial cases are the axioms ‘Reduction C ’ (Lemma 42) and ‘Reduction C' ’ (Lemma 43). □

8.3 Completeness

The proof of completeness follows that of [20]. We first prove the completeness of the logic KC (i.e. without announcements). The completeness of the logic ABALC then simply follows from a translation of its language \mathcal{L}_C into the language \mathcal{L}_C^- of KC. The translation uses the reduction axioms for all possible inductive cases after an announcement.

Definition 45 (Closure)

The *closure* of φ is the minimal set Φ such that

1. $\varphi \in \Phi$,
2. Φ is closed under taking subformulas,
3. If $\psi \in \Phi$ and ψ is not a negation, then $\neg\psi \in \Phi$,
4. If $C_A\psi \in \Phi$, then $E_A(\psi \wedge C_A\psi) \in \Phi$,
5. If $\mathbf{C}_A^\chi\psi \in \Phi$, then $\mathbf{E}_A^\chi(\psi \wedge \mathbf{C}_A^\chi\psi) \in \Phi$. \dashv

The closure of φ is a finite set. As a consequence the canonical model (defined next) is also finite.

Definition 46 (Canonical model)

Let Φ be the closure of φ . The *canonical model* M^c for φ is the triple (S^c, R^c, V^c) where

$$\begin{aligned} S^c &= \{\Gamma \subseteq \Phi \mid \Gamma \text{ is maximally consistent in } \Phi\} \\ R_a^c &= \{(\Gamma, \Delta) \mid \psi \in \Delta \text{ for all } \psi \text{ with } B_a\psi \in \Gamma\} \\ V^c(p) &= \{\Gamma \mid p \in \Gamma\} \end{aligned} \quad \dashv$$

Lemma 47 (Truth Lemma)

Let M^c be the canonical model for $\varphi \in \mathcal{L}$, and let $\Gamma \in S^c$. For all $\psi \in \Phi$: $\psi \in \Gamma$ iff $M^c, \Gamma \models \psi$. \(\dashv\)

Proof The proof of the Truth Lemma is found in the Appendix. \(\square\)

Proposition 48 (Completeness of KC)

$\models \varphi$ implies $\vdash \varphi$. \(\dashv\)

Proof Let $\not\vdash \varphi$. Then $\neg\varphi$ is consistent. A maximal consistent set Γ , a subset of the closure of $\neg\varphi$, exists with $\neg\varphi \in \Gamma$. By the Truth lemma (Lemma 47), $M^c, \Gamma \models \neg\varphi$, so we conclude $\not\models \varphi$. \(\square\)

Theorem 49 (Soundness and Completeness of KC)

$\models \varphi$ iff $\vdash \varphi$. \(\dashv\)

Proof From Proposition 44 (Soundness) and Proposition 48 (Completeness). \(\square\)

We proceed by demonstrating completeness of ABALC, by employing a translation.

Definition 50 (Translation)

The translation function $\text{tr} : \mathcal{L}_C \rightarrow \mathcal{L}_C^-$ is defined as follows:

$$\begin{aligned} \text{tr}(p) &= p & \text{tr}([\varphi]p) &= p \\ \text{tr}(\neg\varphi) &= \neg\text{tr}(\varphi) & \text{tr}([\varphi]\neg\psi) &= \text{tr}(\neg[\varphi]\psi) \\ \text{tr}(\varphi \wedge \varphi') &= \text{tr}(\varphi) \wedge \text{tr}(\varphi') & \text{tr}([\varphi](\psi \wedge \psi')) &= \text{tr}([\varphi]\psi \wedge [\varphi]\psi') \\ \text{tr}(B_a\varphi) &= B_a\text{tr}(\varphi) & \text{tr}([\varphi]B_a\psi) &= \text{tr}((h_a \rightarrow B_a(\varphi \rightarrow [\varphi]\psi)) \wedge \\ & & & \quad (\neg h_a \rightarrow B_a\psi)) \\ \text{tr}(C_A\varphi) &= C_A\text{tr}(\varphi) & \text{tr}([\varphi]C_A\psi) &= \text{tr}(C_A^\varphi([\varphi]\psi \wedge \overline{E}_A(\psi \wedge C_A\psi)) \wedge \\ & & & \quad \overline{E}_A(\psi \wedge C_A\psi)) \\ \text{tr}(C_A^x\varphi) &= C_A^{\text{tr}(x)}\text{tr}(\varphi) & \text{tr}([\varphi]C_A^x\psi) &= \text{tr}(C_A^{\varphi \wedge [\varphi]x}[\varphi]\psi) \\ & & \text{tr}([\varphi][\varphi']\psi) &= \text{tr}([\varphi \wedge [\varphi]'\psi]) \end{aligned} \quad \dashv$$

Theorem 51 (Soundness and Completeness of ABALC)

$\models \varphi$ iff $\vdash \varphi$. \(\dashv\)

Proof The soundness follows, again, from Proposition 44. The completeness of the proof system for the logic ABALC without announcement in the language follows from the observation that every formula containing announcements is provably equivalent to its translation obtained by applying the reduction axioms (Def. 50). \(\square\)

9 Related work and conclusions

Our proposal is related to several other logics in dynamic epistemic logic, such as to arrow update logic [15], wherein a simple dynamic operator can have a large action model equivalent; and to action language approaches [1]. There are further relations to logics with dedicated dynamics, such as the framework reasoning about speech in [10], reasoning about perception as investigated in [24], reasoning about perceptual beliefs [13], and reasoning about visually oriented agents [4]. Paying attention to announcements would normally appear to mean that the agents are *listening* to these announcements, but of course it equally applies to other public events than public announcements and to other senses than hearing. If the event is a light flash, then paying attention means having your eyes open (or facing the light). If the event is inhaling the weak but ever present smell of flowering orange trees in Seville, then paying attention means inhaling. (Joint attention seems difficult to obtain for smell.) And so on.

Relations to other logics of announcement have been discussed in detail in previous sections [18, 9], as has the relation to relativized common knowledge (belief) [20]. However, in those cases the relation was more motivational; these were the sources of our ideas.

This work can be seen as part of a larger research program which consists of developing computationally interesting dynamic epistemic logics wherein one can reason about properties of perception and communication. Directions of future work are manifold. On a similar setting we intend to model trust-based announcement logic. We also intend to model announcements to agents that pay attention to other agents making announcements (as in [6]), instead of to the outsider who is making the announcement and who is not modelled as an agent.

Acknowledgements

Hans van Ditmarsch is also affiliated to IMSc, Chennai, as research associate. We acknowledge support from European Research Council grant EPS 313360. Andreas Herzig and Emiliano Lorini acknowledge support of the the ANR project EmoTES.

This work is based on a conference publication [21]. All parts have been thoroughly revised. The section on attention change has been expanded considerably. A completely novel addition is the section on common belief and the axiomatization for that extension.

References

- [1] G. Aucher. BMS revisited. In *Proc. of 12th TARK*, pages 24–33, 2009.
- [2] G. Aucher and F. Schwarzentruher. On the complexity of dynamic epistemic logic. In *Proc. of 14th TARK*, 2013.
- [3] P. Balbiani, H. van Ditmarsch, A. Herzig, and T. de Lima. Tableaux for public announcement logics. *Journal of Logic and Computation*, 20(1):55–76, 2010.

- [4] P. Balbiani, O. Gasquet, and F. Schwarzentruher. Agents that look at one another. *Logic Journal of IGPL*, 2012.
- [5] A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In *Proc. of 7th TARK*, pages 43–56. Morgan Kaufmann, 1998.
- [6] T. Bolander. Seeing is believing: Formalising false-belief tasks in dynamic epistemic logic. In A. Herzig and E. Lorini, editors, *Proc. of the European Conference on Social Intelligence*, pages 87–107, 2014. CEUR Workshop Proceedings no. 0074-1283-4.
- [7] H.H. Clark and C.R. Marshall. Definite reference and mutual knowledge. In *Elements of discourse understanding*, pages 10–63. Cambridge University Press, 1981.
- [8] T. French, J. Hales, and E. Tay. A composable language for action models. In R. Goré, B.P. Kooi, and A. Kurucz, editors, *Advances in Modal Logic 10*, pages 197–216. College Publications, 2014.
- [9] J.D. Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, University of Amsterdam, 1999. ILLC Dissertation Series DS-1999-01.
- [10] N. Guiraud, A. Herzig, and E. Lorini. Speech acts as announcements. *Logic and the Simulation of Interaction and Reasoning*, page 13, 2009.
- [11] J.Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.
- [12] A. Heifetz, M. Meier, and B.C. Schipper. Interactive unawareness. *Journal of Economic Theory*, 130:78–94, 2006.
- [13] A. Herzig and E. Lorini. A modal logic of perceptual belief. In F. Lihoreau and M. Rebuschi, editors, *Epistemology, Context and Formalism*, Synthese Library. Springer, 2013.
- [14] B. Kooi. Expressivity and completeness for public update logics via reduction axioms. *Journal of Applied Non-Classical Logics*, 17(2):231–254, 2007.
- [15] B. Kooi and B. Renne. Generalized arrow update logic. In *Proc. of 13th TARK*, pages 205–211, 2011. Poster presentation.
- [16] E. Lorini, L. Tummolini, and A. Herzig. Establishing mutual beliefs by joint attention: towards a formal model of public events. In *Proc. of CogSci*, pages 1325–1330. Lawrence Erlbaum Associates, 2005.
- [17] C. Lutz. Complexity and succinctness of public announcement logic. In *Proc. of the 5th AAMAS*, pages 137–144, 2006.

- [18] J.A. Plaza. Logics of public communications. In *Proc. of the 4th ISMIS*, pages 201–216. Oak Ridge National Laboratory, 1989.
- [19] M. Tommasello. Joint attention as social cognition. In C. Moore and P. Dunham, editors, *Joint attention: its origins and role in development*, pages 103–130. Lawrence Erlbaum Associates, 1995.
- [20] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.
- [21] H. van Ditmarsch, A. Herzig, E. Lorini, and F. Schwarzentruher. Listen to me! public announcements to agents that pay attention – or not. In *Proc. of 4th LORI*, pages 96–109. Springer, 2013. LNCS 8196.
- [22] H. van Ditmarsch, W. van der Hoek, and B. Kooi. Dynamic epistemic logic with assignment. In *Proc. of 4th AAMAS*, pages 141–148. ACM, 2005.
- [23] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2007.
- [24] J. van Eijck. Perception and change in update logic. *Games, Actions and Social Software*, pages 119–140, 2012.
- [25] Y. Wang and Q. Cao. On axiomatizations of public announcement logic. *Synthese*, 190(1 supp.):103–134, 2013.

Appendix: Proof of Truth Lemma 47

Truth Lemma 47

Let M^c be the canonical model for $\varphi \in \mathcal{L}$, and let $\Gamma \in S^c$. For all $\psi \in \Phi$: $\psi \in \Gamma$ iff $M^c, \Gamma \models \psi$. ⊖

The proof is by induction on the structure of formulas $\psi \in \Phi$, namely, on the property

$$\mathcal{P}(\psi_0) : \quad \text{for all } \Delta \in S^c: M^c, \Delta \models \psi \text{ iff } \psi \in \Delta.$$

Throughout this proof, we write R_a instead of R_a^c for the accessibility relation in the canonical model. The cases for atomic propositions, negations, conjunction and belief are as in [20]. We only prove the case for attentive relativized common belief, and the similar case for common belief. We further define for an arbitrary (finite) set of formulas Γ (such as a maximally consistent set in the domain of M^c), $\delta_\Gamma = \bigwedge_{\gamma \in \Gamma} \gamma$, and for any collection Σ of such finite sets we define $\delta_\Sigma = \bigvee_{\Gamma \in \Sigma} \delta_\Gamma$. The notation $\vdash \varphi$ stands for derivability of φ in the axiomatization.

(Case $\mathbf{C}_A^X \psi$) Suppose $\mathcal{P}(\psi)$, $\mathcal{P}(\chi)$, and $\mathcal{P}(h_a)$ for all $a \in A$. We prove that $\mathcal{P}(\mathbf{C}_A^X \psi)$.

⇒

Suppose $\mathbf{C}_A^x\psi \in \Gamma$, and let us prove that $(M^c, \Gamma) \models \mathbf{C}_A^x\psi$. Let $\Delta \in S^c$ be arbitrary with $(\Gamma, \Delta) \in (\mathbf{R}_A^x)^+$, where \mathbf{R}_A^x is the interpretation of the operator \mathbf{E}_A^x . (Note that if no such Δ exists, we are done.) We must show that $(M^c, \Delta) \models \psi$. This is done by induction on the length of the paths from Γ to Δ . The claim is shown together with $\mathbf{C}_A^x\psi \in \Delta$. So let:

$$\mathcal{Q}(n): \quad \text{for all } \Delta \in S^c, (\Gamma, \Delta) \in (\mathbf{R}_A^x)^n: M^c, \Delta \models \psi \text{ and } \mathbf{C}_A^x\psi \in \Delta.$$

(Base Case) From the assumption $(\Gamma, \Delta) \in \mathbf{R}_A^x$, we have $(\Gamma, \Delta) \in \mathbf{R}_a^x$ for some $a \in A$. From this we get that $M^c, \Gamma \models h_a$ and that $M^c, \Delta \models \chi$, so

$$(1) \quad h_a \in \Gamma \text{ and } \chi \in \Delta \quad \text{by } \mathcal{P}(h_a) \text{ and } \mathcal{P}(\chi)$$

Also, $\mathbf{C}_A^x\psi \in \Gamma$ implies by way of the (Mix \mathbf{C}) axiom that for any $a \in A$:

$$(2) \quad \vdash \delta_\Gamma \rightarrow (h_a \rightarrow B_a(\chi \rightarrow (\psi \wedge \mathbf{C}_A^x\psi)))$$

From (1), it follows that $\vdash \delta_\Gamma \rightarrow h_a$, which combined with (2) gives

$$(3) \quad \vdash \delta_\Gamma \rightarrow B_a(\chi \rightarrow (\psi \wedge \mathbf{C}_A^x\psi))$$

The latter implies that $B_a(\chi \rightarrow (\psi \wedge \mathbf{C}_A^x\psi)) \in \Gamma$, by the maximal consistency of Γ . Hence, $\chi \rightarrow (\psi \wedge \mathbf{C}_A^x\psi) \in \Delta$ by the assumption $(\Gamma, \Delta) \in \mathbf{R}_a^x$. Combining this with $\chi \in \Delta$ (1), we conclude that

$$(4) \quad \psi \in \Delta \text{ and } \mathbf{C}_A^x\psi \in \Delta$$

by the maximal consistency of Δ . Finally, by $\mathcal{P}(\psi)$, we conclude that

$$(5) \quad M^c, \Delta \models \psi \text{ and } \mathbf{C}_A^x\psi \in \Delta$$

So, $\mathcal{Q}(1)$ is proved.

(Ind. Case) Assume $\mathcal{Q}(n)$ and let us prove $\mathcal{Q}(n+1)$. For an arbitrary $(n+1)$ -length path from Γ to Δ , let Θ be the n -th element in this path, i.e. $(\Gamma, \Theta) \in (\mathbf{R}_A^x)^n$ and $(\Theta, \Delta) \in \mathbf{R}_A^x$. Then there is an $a \in A$ such that $h_a \in \Theta$ and $\chi \in \Delta$. Also, by $\mathcal{Q}(n)$, we have $\mathbf{C}_A^x\psi \in \Theta$. Applying the same reasoning as in the Base Case, we conclude that $M^c, \Delta \models \psi$ and $\mathbf{C}_A^x\psi \in \Delta$. So, $\mathcal{Q}(n+1)$ is proved.

⇐

Suppose that $M^c, \Gamma \models \mathbf{C}_A^x\psi$. We show that $\mathbf{C}_A^x\psi \in \Gamma$. Let $\mathcal{S} = \{\Delta \mid (\Gamma, \Delta) \in (\mathbf{R}_A^x)^+\}$. The set \mathcal{S} is the set of successors reachable from Γ by the relation corresponding to attentive relativized common belief. The formula $\delta_{\mathcal{S}}$ (we recall that it is defined as $\bigvee_{\Delta \in \mathcal{S}} \delta_\Delta$) is a syntactic description of the set of worlds \mathcal{S} — recall that \mathcal{S} is finite, so $\delta_{\mathcal{S}}$ is a formula.

We now prove the following three claims:

$$1. \quad \vdash \delta_{\mathcal{S}} \rightarrow \mathbf{E}_A^x\delta_{\mathcal{S}}$$

Claim 1

$$2. \vdash \delta_\Gamma \rightarrow \mathbf{E}_A^x \delta_\mathcal{S}$$

Claim 2

$$3. \vdash \delta_\mathcal{S} \rightarrow \psi$$

Claim 3

Claim 1 says that we remain in \mathcal{S} by taking an \mathbf{R}_a^x -transition. Claim 2 says that by taking one \mathbf{R}_a^x -transition from Γ , we end up in \mathcal{S} . Claim 3 says that ψ holds in all $(\mathbf{R}_A^x)^+$ -successors.

Proof of Claim 1:

First, we prove that $(\Delta, \Delta') \notin R_a$ implies $\vdash \delta_\Delta \rightarrow B_a \neg \delta_{\Delta'}$. Indeed, $(\Delta, \Delta') \notin R_a$ iff there is a formula φ such that $B_a \varphi \in \Delta$ and $\varphi \notin \Delta'$. Thus, $\vdash \delta_\Delta \rightarrow B_a \varphi$ and $\vdash \varphi \rightarrow \neg \delta_{\Delta'}$. This implies that $\vdash \delta_\Delta \rightarrow B_a \neg \delta_{\Delta'}$.

Let now $a \in A$, $\Delta \in \mathcal{S}$, and $\Delta' \notin \mathcal{S}$. By the definition of \mathcal{S} we have that $(\Delta, \Delta') \notin \mathbf{R}_a^x$. We then proceed as follows:

$$\begin{array}{ll}
(\Delta, \Delta') \notin \mathbf{R}_a^x & \\
\text{iff } M^c, \Delta \not\models h_a, \text{ or } (\Delta, \Delta') \notin R_a, \text{ or } M^c, \Delta' \not\models \chi & \\
\text{iff } h_a \notin \Delta, \text{ or } (\Delta, \Delta') \notin R_a \text{ or } \chi \notin \Delta' & \text{by } \mathcal{P}(h_a), \mathcal{P}(\chi) \\
\text{implies } \vdash \delta_\Delta \rightarrow \neg h_a, \text{ or } \vdash \delta_\Delta \rightarrow B_a \neg \delta_{\Delta'} \text{ or } \vdash \chi \rightarrow \neg \delta_{\Delta'} & \text{see above} \\
\text{implies } \vdash \delta_\Delta \rightarrow \neg h_a, \text{ or } \vdash \delta_\Delta \rightarrow B_a \neg \delta_{\Delta'}, \text{ or } \vdash B_a(\chi \rightarrow \neg \delta_{\Delta'}) & \\
\text{implies } \vdash \delta_\Delta \rightarrow \neg h_a, \text{ or } \vdash \delta_\Delta \rightarrow B_a(\chi \rightarrow \neg \delta_{\Delta'}) & \\
\text{implies } \vdash \delta_\Delta \rightarrow (h_a \rightarrow B_a(\chi \rightarrow \neg \delta_{\Delta'})) &
\end{array}$$

As all the quantifiers range over finite sets, we also have (wherein ‘big’ disjunctions and ‘big’ conjunctions bind stronger than all other operators):

$$\begin{aligned}
& \vdash \bigwedge_{\Delta \in \mathcal{S}} \bigwedge_{\Delta' \in \mathcal{S}} \bigwedge_{a \in A} (\delta_\Delta \rightarrow (h_a \rightarrow B_a(\chi \rightarrow \neg \delta_{\Delta'}))) \\
& \vdash \bigvee_{\Delta \in \mathcal{S}} \delta_\Delta \rightarrow \bigwedge_{a \in A} (h_a \rightarrow B_a(\chi \rightarrow \bigwedge_{\Delta' \notin \mathcal{S}} \neg \delta_{\Delta'}))
\end{aligned}$$

We now use that $\vdash \bigvee_{\Delta \in S^c} \delta_\Delta$; as S^c is the domain of *all* maximal consistent sets in Φ , then if this disjunction were not a theorem, its negation would be consistent, from which we can easily obtain a contradiction. We then write this disjunction as $\vdash \bigvee_{\Delta' \notin \mathcal{S}} \delta_{\Delta'} \vee \bigvee_{\Delta \in \mathcal{S}} \delta_\Delta$, in other words, as $\vdash \bigwedge_{\Delta' \notin \mathcal{S}} \neg \delta_{\Delta'} \rightarrow \bigvee_{\Delta \in \mathcal{S}} \delta_\Delta$. We therefore can conclude:

$$\vdash \bigvee_{\Delta \in \mathcal{S}} \delta_\Delta \rightarrow \bigwedge_{a \in A} (h_a \rightarrow B_a(\chi \rightarrow \bigvee_{\Delta \in \mathcal{S}} \delta_\Delta))$$

in other words,

$$\vdash \delta_\mathcal{S} \rightarrow \mathbf{E}_A^x \delta_\mathcal{S}$$

Hence the claim is proven.

Proof of Claim 2:

It is omitted. It is similar to the proof of Claim 1, except that we replace the arbitrary $\Delta \in \mathcal{S}$ by the fixed set Γ .

Proof of Claim 3:

As $M^c, \Gamma \models \mathbf{C}_A^x \psi$, we have, for all $\Delta \in \mathcal{S}$, that $M^c, \Delta \models \psi$. By $\mathcal{P}(\psi)$, we have $\psi \in \Delta$. Therefore $\vdash \delta_\Delta \rightarrow \psi$ for all $\Delta \in \mathcal{S}$. So $\vdash \delta_{\mathcal{S}} \rightarrow \psi$.

This completes the proof of all three claims.

Having thus proved these claims, we now proceed as follows.

- (i) $\vdash \mathbf{C}_A^x(\delta_{\mathcal{S}} \rightarrow \mathbf{E}_A^x \delta_{\mathcal{S}})$ by applying (Necessitation **C**) on Claim 1
- (ii) $\vdash \mathbf{E}_A^x \delta_{\mathcal{S}} \wedge \mathbf{C}_A^x(\delta_{\mathcal{S}} \rightarrow \mathbf{E}_A^x \delta_{\mathcal{S}}) \rightarrow \mathbf{C}_A^x \delta_{\mathcal{S}}$ instance of (Induction **C**)
- (iii) $\vdash \mathbf{E}_A^x \delta_{\mathcal{S}} \rightarrow \mathbf{C}_A^x \delta_{\mathcal{S}}$ by propositional reasoning on (i) and (ii)
- (iv) $\vdash \delta_\Gamma \rightarrow \mathbf{C}_A^x \delta_{\mathcal{S}}$ by Claim 2 and (iii)
- (v) $\vdash \mathbf{C}_A^x \delta_{\mathcal{S}} \rightarrow \mathbf{C}_A^x \psi$ by (Necessitation **C**) and (Distribution **C**) on Claim 3
- (vi) $\vdash \delta_\Gamma \rightarrow \mathbf{C}_A^x \psi$ by propositional reasoning on (iv) and (v)

From (vi), the maximal consistency of Γ in Φ , and $\mathbf{C}_A^x \psi \in \Phi$ by assumption, we conclude that $\mathbf{C}_A^x \psi \in \Gamma$.

This closes the Case $\mathbf{C}_A^x \psi$ of the Truth Lemma. We continue with the Case $C_A \psi$.

(Case $C_A \psi$) Suppose that $\mathcal{P}(\psi)$, and let us prove $\mathcal{P}(C_A \psi)$. For this case, the proof is simpler than that of $\mathbf{C}_A^x \psi$: we replace the relation \mathbf{R}_A^x by R_A (i.e., R_A^c) everywhere, e.g., in the reformulation of $\mathcal{Q}(n)$ that now becomes:

$$\mathcal{Q}'(n): \quad \text{for all } \Delta \in S^c, (\Gamma, \Delta) \in R_A^n: M^c, \Delta \models \psi \text{ and } C_A \psi \in \Delta.$$

\Rightarrow

From $M^c, \Gamma \models C_A \psi$ we directly obtain, by the (Mix **C**) axiom:

$$(3') \quad \vdash \delta_\Gamma \rightarrow B_a(\psi \wedge C_A \psi)$$

so that, as before, for any $(\Gamma, \Delta) \in R_a, \psi \wedge C_A \psi \in \Delta$, thus, using $\mathcal{P}(\psi)$, $M^c, \Delta \models \psi$ and $C_A \psi \in \Delta$: $\mathcal{Q}'(1)$ is true. The inductive case for $\mathcal{Q}'(n) \Rightarrow \mathcal{Q}'(n+1)$ is again straightforward.

\Leftarrow

For this part of the proof, we redefine the set \mathcal{S} as $\mathcal{S} = \{\Delta \mid (\Gamma, \Delta) \in (R_A^c)^+\}$, and replace Claim 1 and Claim 2 by, respectively, $\vdash \delta_{\mathcal{S}} \rightarrow E_A \delta_{\mathcal{S}}$ and $\vdash \delta_\Gamma \rightarrow E_A \delta_{\mathcal{S}}$, whereas Claim 3 is as before. The proofs of these claims are then analogous to the previous cases, wherein one should replace \mathbf{C}_A^x by C_A and abstract from (remove occurrences of) χ and the atomic propositions h_a . For the remainder of this part of the proof, the steps (i)-(vi) are also similar, but now we use (Necessitation **C**), (Induction **C**) and (Distribution **C**).