

From Public Announcements to Asynchronous Announcements

Philippe Balbiani¹ and Hans van Ditmarsch² and Saúl Fernández González³

Abstract. We present a multi-agent logic of belief and announcements wherein the sending of announcements and the reception of announcements by agents are separated, thus straying from the paradigm of Public Announcement Logic (PAL). Both PAL and Asynchronous Announcement Logic (recently proposed in the literature) are special cases in our framework. We provide a history-based semantics for our ‘Partially Synchronous Announcement Logic’, proposing three different interpretations of the notion of asynchronicity. We then show that the logic of our three proposals is the same (‘PSAL’) and prove soundness and completeness for a Hilbert-style axiomatisation. Finally, we propose a notion of common belief for this framework, of which we give some validities.

1 Introduction

Let us say a new documentary series has premiered, on a topic which is very much of interest to three friends, Antía, Brais and Carmiña.

Each episode of this show is released at irregular intervals on a streaming platform, and the episodes are generally watched in order. Each episode contains some factual information p_i .

Scenario 1. Let us consider a situation in which the three friends watch the show independently and without talking to each other about it. For example, Antía and Brais watch the first three episodes religiously after they are broadcast, and Carmiña binges these three episodes after the release of the third. We can represent this situation by the sequence

$$\alpha = p_1.a.b.p_2.a.b.p_3.a.b.c.c.c.$$

What does this mean? The formulas p_1 , p_2 and p_3 take the role of the announcements, novel information to be incorporated by the three friends. As all three watch the episodes individually, they receive the announcements individually too. This is unlike in Public Announcement Logic (PAL), where the information is received simultaneously (i.e., synchronously). Reading the sequence from left to right, each subsequent a reads the next unprocessed announcement (therefore there are three a ’s), and similarly for b and c . However, the order between a , b , and c is unrestricted, and also whether an announcement is received before or after the sending of the next announcement: as we see, and as described in the scenario, c reads (‘binge-watches’) all three only after they have all been sent/broadcast.

What sort of knowledge or belief do we expect to hold afterwards? As all friends watched all shows, we have that a , b , and c all know that p_1 , p_2 , and p_3 . However, as the agents are unaware of each other

having watched the show, we do not have that they know this from one another. A fortiori, no common knowledge of any kind results from this scenario.

Scenario 2. Now let us think of a more convoluted situation: Antía and Brais watch the first episode of the show independently; having discovered that they are both fans of the show, they watch the second episode together, and they convince Carmiña to start it, who then watches on her own the first two episodes in a row. After the third episode is released, the three of them get together to watch it.

To this scenario corresponds the sequence

$$\beta = p_1.a.b.p_2.ab.c.c.p_3.abc.$$

After this sequence we should clearly wish that a , b , c have common knowledge of p_3 , as they watched that episode together. We also wish to conclude that a and b have common knowledge of p_2 , as they watched that together. But do we also wish to conclude that from watching the show p_3 with Antía and Brais, Carmiña learns anything about Antía and Brais watching p_2 together before? That depends on our interpretation of such sequences, what *view* each agent has on such *histories*. We will provide several intuitive solutions to address that. Such things have of course been widely discussed in the temporal epistemic (and distributed computing) literature [5, 15, 7, 19, 17].

Scenario 3. A situation more in line with Public Announcement Logic which involves these friends and this show is one in which all the friends watch every episode together immediately after each broadcast, represented by the sequence

$$\gamma = p_1.abc.p_2.abc.p_3.abc.$$

We now have that it is common knowledge to all three afterwards that p_1 , p_2 , and p_3 , as is to be expected.

Asynchrony. What all scenarios have in common is indeed a notion of *public* announcement, in the sense of broadcast. Also, although we were not explicit about this, we always assume that the broadcast information is true when sent. So these are *truthful* announcements. However, the framework we are hinting at here differs from the paradigm of Public Announcement Logic in that the reception of these messages may be *asynchronous*. It may even be asynchronous in any way, for any subgroup, allowing for complex interactions where different intersecting subgroups acquire knowledge independently: the abode of distributed computing [13, 4, 12].

Asynchronous Announcement Logic. The example represented by history α , i.e., a case in which the epistemic agents receive the announcements individually, was proposed and studied in [11, 2]. Let us very briefly present the framework introduced in [2].

¹ IRIT, Université de Toulouse.

² LORIA, Université de Lorraine.

³ IRIT, Université de Toulouse. Email: saul.fgonzalez@irit.fr

The authors present the logic AAL of publicly sent (broadcast) but individually received announcements. Apart from an epistemic modality for each agent, the logic contains dynamic modalities for (sending of) public announcements and for individual reception of announcements. The logic is interpreted on Kripke models with arbitrary accessibility relations, where the equivalence relations as usual in Public Announcement Logic are a special case. Like in the previous scenarios, histories such as α above play an important role. To interpret epistemic modalities asynchronously, in addition to the epistemic accessibility relation between worlds, an agent-dependent notion of accessibility between histories is also given: the *view* relation. It determines what histories β , from the perspective of agent a , are accessible from a given history α (denoted $\alpha \triangleright_a \beta$). And finally, to guarantee that announcements are true when made, a notion of *executability* of a history in a state is needed. Then, agent a believes/knows a formula ϕ at world w given history α , if and only if ϕ holds in all worlds t accessible from w , for all histories β accessible from α , on condition that β is executable in world t (denoted $t \triangleright \beta$). This will be a special case of the semantics we later provide in the paper.

A complete axiomatization is given for the class of models with arbitrary accessibility relations, and also for the special case with equivalence relations. Still, the epistemic notion does not correspond to *knowledge*, but to *belief*. This is a consequence of the modelling decision that the agents do not reason over all possible histories, but only over histories containing as many announcements as they themselves have received. That means that in a situation where agents a and b are commonly uncertain about a propositional variable p , after p has been announced and received by a but not yet by b , agent a correctly believes that p , but agent b may incorrectly believe that agent a is ignorant about p . However, all *positive* knowledge (no negations before epistemic modalities; corresponding to the universal fragment in first-order logic translation) is correct. A detailed discussion about such differences between knowledge and belief is part of the investigation in [2]. It should be noted that the methods and techniques employed are roughly based on the discussion of inconsistent cuts in [16] and the history-based semantics of [17, 21].

Our results. The present paper is concerned with the example represented by β : what if several agents group together to receive announcements? We provide a language with epistemic, announcement, and reception modalities (for arbitrary subgroups), and a logical semantics with several intuitive variations involving the ‘view’ of different agents on previously received messages. Then, we provide a complete axiomatization of the logic for the class of epistemic models with empty histories. We address how the logic relates to the logics PAL [18] and AAL [2]. We also add a notion of common belief and prove a number of obvious validities for this notion (although it is not obvious that they hold in our asynchronous interpretation). Although these notions of common belief (or common knowledge) are therefore truly asynchronous notions, and different from the usual notion, they are also different from, for example, the concurrent common knowledge (or belief) of [16]: we can also obtain that, but by the usual iteration of agents continuing to send each other (iterated) shared knowledge. Examples are provided to all results. Although in our paper we restrict ourselves to the three given scenarios, with our results we can also model, for example, the rich variety of scenarios for the Muddy Children Problem presented in [14], and also in general address complex interactions between subgroup common knowledge for different intersecting subgroups that would be hard to achieve (if achievable at all) in Dynamic Epistemic Logic in general,

even in the presence of semi-public or private announcements [3, 22].

Outline. This paper is structured as follows: in Sections 2, 3 and 4 we present the technicalities of Partially Synchronised Announcement Logic (PSAL). In Sections 5 and 6 we show how we can emulate AAL and PAL, respectively, within our framework. In Section 7 we give a proposal for a notion of Common Belief in this framework, of which we provide some validities. We conclude in Section 8.

2 Syntax and histories

Throughout this paper, let P be a countable set of propositional variables and let A be a (finite, nonempty) set of agents. Let $\mathcal{G} = \mathcal{P}(A) \setminus \{\emptyset\}$.

Our language will be \mathcal{L}_{PSAL} , defined as:

$$\phi ::= p \mid \top \mid (\phi \wedge \psi) \mid \neg \phi \mid \hat{B}_a \phi \mid \langle G \rangle \phi \mid \langle \phi \rangle \phi,$$

where $p \in P$ and $G \in \mathcal{G}$. The Boolean connectives \perp , \vee , \rightarrow , \leftrightarrow are defined by the usual abbreviations. We define dual modalities $B_a \phi = \neg \hat{B}_a \neg \phi$, $[\psi] \phi = \neg \langle \psi \rangle \neg \phi$ and $[G] \phi = \neg \langle G \rangle \neg \phi$.

We will consider words of the shape of β in our introductory example: finite sequences made out of formulas and subsets of A , i.e. words $\alpha \in (\mathcal{L}_{PSAL} \cup \mathcal{G})^*$. For cleanliness in presentation, when writing down these histories explicitly we will separate announcements and readings with dots and write down the agents in a group G as a concatenation rather than as a set (as β in the introduction). For each such word, the formula $\langle \alpha \rangle \phi$ represents an abbreviation of the sequence of announcement and reading modalities corresponding to the announcements and readings which appear in α , defined recursively as follows:

$$\langle \epsilon \rangle \phi = \phi; \langle \alpha.\psi \rangle \phi = \langle \alpha \rangle \langle \psi \rangle \phi; \langle \alpha.G \rangle \phi = \langle \alpha \rangle \langle G \rangle \phi,$$

where ϵ is the empty word. Every formula in \mathcal{L}_{PSAL} is thus of the form $\langle \alpha \rangle \phi$ for some $\alpha \in (\mathcal{L}_{PSAL} \cup \mathcal{G})^*$.

Let us order the elements appearing in α by \sqsubset_α and let us use $\alpha \upharpoonright_A$ and $\alpha \upharpoonright_!$ to denote the projection of α to, respectively, \mathcal{G} and \mathcal{L}_{PSAL} . We use $|\alpha|_!$ to denote the length of $\alpha \upharpoonright_!$, i.e., the number of announcements occurring in α .

Whether such a word constitutes a history, and crucially, whether two such histories are in some form of epistemic accessibility relation for agent a , will depend on the interpretation we are trying to make. Interestingly, there are (at least) three intuitively appealing interpretations which, surprisingly, result in almost identical logics. Let us rely on our TV show analogy from the introduction in order to present them here.

First interpretation. There is common knowledge among the group of friends that they all care greatly about the narrative and the continuity of the plot. That is, Antía knows that, if Brais is watching the eighth episode with her, this means that at some point in the past he has watched, in order, the previous seven episodes. Antía does not know, however, when or with whom did Brais enjoy the preceding instalments of the show.

In this situation, the number of announcements an agent $a \in A$ reads in a word α , denoted by $|\alpha|_a$, corresponds with the number of times this agent is included in the groups occurring in α . The chain of announcements read by the agent, $\alpha \upharpoonright_a$, will be the first $|\alpha|_a$ announcements in $\alpha \upharpoonright_!$. That is, if $\alpha \upharpoonright_! = \phi_1 \dots \phi_n$, we have

$$|\alpha|_a = |\{G \in \alpha \upharpoonright_A : a \in G\}|; \alpha \upharpoonright_a = \phi_1 \dots \phi_{|\alpha|_a}.$$

We will use $\alpha|_a$ to denote the set $\{G \in \alpha|_A : a \in G\}$ linearly ordered by \sqsubset_α .

In order for a word to be a *history*, we need that an agent can, at any point along the history, read no more announcements than those which have been made and, as explained above, we also need to require that if two agents $a, b \in G \in \alpha$ read an announcement together, then they have both read the same amount of announcements before. That is, a word $\alpha \in (\mathcal{L}_{PSAL} \cup \mathcal{G})^*$ is a history iff, for every prefix $\beta.G \sqsubseteq \alpha$ and every $a, b \in G$, we have that $|\beta.G|_a = |\beta.G|_b \leq |\beta.G|_!$.

Two histories α and β are in the view relation \triangleright_a whenever, from the perspective of a , the same announcements have been made and have been received by the same groups including a and, in β , no further announcements have been made than those a knows about. In other words,

$$\alpha \triangleright_a \beta \text{ iff } \begin{cases} \alpha|_a = \beta|_a \text{ and} \\ \alpha|_!a = \beta|_!a = \beta|_!. \end{cases}$$

Example 1 Let $\alpha_1 = p.\neg B_ap.ab.c.B_bB_ap.bc.a$. The formulas p and $\neg B_ap$ are announced in succession, after which the agents read the first announcement p together. After this occurs, the formula B_bB_ap is now true, and it is announced. Afterwards, b and c read the announcement $\neg B_ap$ together, and then a reads it alone. (Note that they are all aware this announcement is false by the time they read it.)

Second interpretation. The friends do care about the continuity of the show but they sometimes have to skip an episode due to their frantic lifestyles. When a group of friends meets to watch an episode, they will explain to each other the plot of the previous episodes and, to make sure they avoid spoilers in future social situations, they will share all they know about who watched which episode with whom. They do not wish to skip ahead, so they watch the $n + 1$ th episode, where n is the latest episode any member of the group has watched. In our case, if Antía and Brais are meeting to watch the eight episode, it means that at least one of them (let us say Brais) knows what happened in the seventh.

Here, when the agents in a group G communicate they tell each other not only with whom they have read the previous announcements and what this announcements were, but also they communicate with each other information that, during those prior readings, was communicated to them. The way to express that is by considering a partial order \leq_α^* on $\alpha|_A$, which we can define as the reflexive and transitive closure of $\bigcup_{a \in A} \xrightarrow{a}$, where $G \xrightarrow{a} G'$ whenever $a \in G$ and G' is the next group in which a appears, i.e., $G \xrightarrow{a} G'$ iff (i.) $G \sqsubset_\alpha G'$, (ii.) $a \in G \cap G'$, and (iii.) for all G'' such that $G \sqsubset_\alpha G'' \sqsubset_\alpha G'$, we have $a \notin G''$. Under this definition, $G \leq_\alpha^* G'$ whenever there is a chain $G = G_0, G_1, \dots, G_n = G'$ and agents $a_1, \dots, a_n \in A$ such that $a_i \in G_{i-1}$ and G_i is the next element of $\alpha|_A$ including a_i .

We define $\text{last}_a \alpha$ as the last G occurring in α including a , i.e.

$$\text{last}_a \alpha = \max_{\sqsubset_\alpha} \{G \in \alpha|_A : a \in G\}.$$

The communications that agent a is aware of are exactly those that were discussed by all agents in the last reading that a was involved in. These are, in turn, precisely those groups that can be traced back from $\text{last}_a \alpha$ via the relation $\bigcup_{b \in A} \xrightarrow{b}$. In other words,

$$\alpha|_a := \downarrow_{\leq_\alpha^*} \text{last}_a \alpha = \{G \in \alpha|_A : G \leq_\alpha^* \text{last}_a \alpha\}, \leq_\alpha^*.$$

Note that $\downarrow_{\leq_\alpha^*} G$ is a partially ordered set representing all the communications the agents in G are aware of. But the agents in G are reading an announcement too, let us call it ϕ_n , the n th one occurring in α . Now, this n is one plus the highest number of announcements any agent in G has read before, all the way back to the moment the agents in G read their first announcement, ϕ_1 . We see that this number n corresponds to the length of the longest chain in $\downarrow_{\leq_\alpha^*} G$, or rather the *height* of G in the poset. For $G \in \alpha|_A$,

$$h(G) = \max\{n : \exists G_1, \dots, G_n \in \alpha|_A (G_i <_\alpha^* G_{i+1} \ \& \ G_n = G)\}.$$

The number of announcements read by a is therefore $|\alpha|_a = h(\text{last}_a \alpha)$, and the announcements read by a under this interpretation are precisely the first $|\alpha|_a$ announcements occurring in $\alpha|_!$, i.e. $\alpha|_!a = \phi_1 \dots \phi_{|\alpha|_a}$.

A word, then, is a history whenever every agent, at any stage, cannot receive more announcements than those which have been sent: α is a history iff, for all $a \in A$, and for all $\beta \sqsubseteq \alpha$, we have that $|\beta|_a \leq |\beta|_!$.

A history β is an epistemic alternative to α , from the perspective of agent a , if the poset of communications a is aware of coincides for both histories, and moreover if the announcements read by a are the same in both histories and β has no further announcements. That is,

$$\alpha \triangleright_a \beta \text{ iff } \begin{cases} \alpha|_a = \beta|_a \text{ and} \\ \alpha|_!a = \beta|_!a = \beta|_!. \end{cases}$$

Example 2 Let $\alpha_2 = p.\neg B_ap.ab.B_bB_ap.bc.B_cp.ca$. Here, p and $\neg B_ap$ are announced, after which a and b together read p . Then, B_bB_ap is announced and b and c read together the second announcement $\neg B_ap$. Now c is also aware of p because b has communicated to c the first announcement. After this, B_cp is announced, and finally, c having read the second announcement and a having so far read only the first, a and c read the third one, B_bB_ap , after which both a and c are aware of the first three announcements. Note that α_2 is not a valid history under the first interpretation.

Third interpretation. The friends do not care too much about the continuity of the show; whenever two friends meet, they watch the next episode to the latest one either of them has watched, but they do not feel the need to explain to each other previous plot lines or to even mention in which context they have enjoyed the previous episodes. It could be that Antía and Brais are watching together the third episode, but it is the first episode Antía sees: Brais watched the second episode with Carmina, who watched the first one by herself.

For this we define \leq_α^* , $\text{last}_a \alpha$ and $h(G)$ as above. Under this interpretation, a only reads as many announcements as times it occurs in α . We then define $\alpha|_a$ to be the set $\{G \in \alpha : a \in G\}$ linearly ordered by \sqsubset_α , and $|\alpha|_a$ to be the cardinality of this set, just as in the first interpretation.

The difference this time, however, is that a does not read the first $|\alpha|_a$ announcements. Let α_k^a denote the k -th group in which a reads an announcement, i.e., the k -th element of $\alpha|_a$. The k -th announcement read by a corresponds to one plus the latest announcement number any member of α_k^a has read, i.e., the height of α_k^a . Therefore the announcements read by a are $\alpha|_!a = \phi_{n_1^a} \dots \phi_{n_{|\alpha|_a}^a}$, where $n_k^a = h(\alpha_k^a)$.

A history is therefore a word α where, at any stage, the last announcement read by an agent corresponds to a number not higher than the number of announcements made. In other words, α is a history if, for all $a \in A$ and all $\beta \sqsubseteq \alpha$, $h(\text{last}_a \beta) \leq |\beta|_!$.

And, as before, we define the relation \triangleright_a as:

$$\alpha \triangleright_a \beta \text{ iff } \begin{cases} \alpha \upharpoonright_a = \beta \upharpoonright_a \text{ and} \\ \alpha \upharpoonright_{!a} = \beta \upharpoonright_{!a} = \beta \upharpoonright_{!}. \end{cases}$$

Example 3 Let $\alpha_3 = p.\neg B_a p.ab.B_b B_a p.bc.\neg B_c p.ca$. As in previous examples, p and $\neg B_a p$ are announced and then a and b read p . $B_b B_a p$ is then announced, after which b and c read together $\neg B_a p$, since b has already read the first announcement. Unlike the previous example, c never gets to read the first announcement, and in this example is not aware that p is true; $\neg B_c p$ is announced and c and a read together the third announcement, $B_b B_a p$. This history, while valid, would never be executable under the second interpretation because $\neg B_c p$ would always be false after the second reading takes place (see Section 3 for formal details on executability).

While each of these interpretations gives rise to different notions of executability (\bowtie) and indistinguishability (\triangleright_a), and thus to different semantics (see Section 3), we have that, curiously enough, the logic of validities for each interpretation will be the same (or rather, will have the same shape: see Section 4 for details).

Unless stated otherwise, the results in the remainder of this paper are valid for all three interpretations. Let \mathcal{H} be the set of histories (under one's preferred interpretation). Note that, given a history α , the set $\{\beta \in \mathcal{H} : \alpha \triangleright_a \beta\}$ is finite for any of the definitions.

3 Semantics

We read formulas of \mathcal{L}_{PSAL} on models of the form (W, R, V) , where W is a nonempty set of worlds, $R = \{R_a\}_{a \in A}$ is a family of accessibility relations on $W \times W$ and $V : \text{Prop} \rightarrow \mathcal{P}(W)$ is a valuation. We evaluate them with respect to pairs (w, α) where w is a world and α is a history such that α is executable in w , represented by $w \bowtie \alpha$. In order to define the executability relation \bowtie and the satisfaction relation \models we shall first introduce a well-founded partial order \ll between pairs (α, ϕ) .

Definition 4 We define $\deg \phi$ and $\|\phi\|$ recursively:

$$\begin{aligned} \deg p &= 0 & \|p\| &= 2 \\ \deg \top &= 0 & \|\top\| &= 1 \\ \deg(\neg\phi) &= \deg \phi & \|\neg\phi\| &= \|\phi\| + 1 \\ \deg(\phi \wedge \psi) &= \max\{\deg \phi, \deg \psi\} & \|\phi \wedge \psi\| &= \|\phi\| + \|\psi\| \\ \deg(\langle G \rangle \phi) &= \deg \phi & \|\langle G \rangle \phi\| &= \|\phi\| + 2 \\ \deg(\langle \phi \rangle \psi) &= \deg \phi + \deg \psi & \|\langle \phi \rangle \psi\| &= 2\|\phi\| + \|\psi\| \\ \deg(\hat{B}_a \phi) &= \deg \phi + 1 & \|\hat{B}_a \phi\| &= \|\phi\| + 1. \end{aligned}$$

For a word α , we set $\deg \alpha := \sum \{\deg \psi : \psi \text{ occurs in } \alpha\}$ and

$$\|\epsilon\| = 0, \|\alpha.G\| = \|\alpha\| + 1, \|\alpha.\psi\| = \|\alpha\| + \|\psi\|.$$

Finally, for pairs (α, ϕ) we set:

$$\deg(\alpha, \phi) = \deg \alpha + \deg \phi \text{ and } \|(\alpha, \phi)\| = \|\alpha\| + \|\phi\|,$$

and we define a well-founded order \ll as a lexicographical ordering on these quantities, i.e. $(\alpha, \phi) \ll (\beta, \psi)$ iff

$$\begin{cases} \deg(\alpha, \phi) < \deg(\beta, \psi), \text{ or} \\ \deg(\alpha, \phi) = \deg(\beta, \psi) \ \& \ \|(\alpha, \phi)\| < \|(\beta, \psi)\|. \end{cases}$$

Definition 5 (Semantics of PSAL) Given a pair (w, α) , where w is a world and α is a history, we define $w \bowtie \alpha$ and $w, \alpha \models \phi$ by double \ll -recursion on (α, ϕ) as it appears in Table 1.

$w \bowtie \epsilon$	always;
$w \bowtie \alpha.\phi$	iff $w \bowtie \alpha$ and $w, \alpha \models \phi$;
$w \bowtie \alpha.G$	iff $w \bowtie \alpha$;
$w, \alpha \models p$	iff $w \in V(p)$;
$w, \alpha \models \top$	always;
$w, \alpha \models \neg\phi$	iff $w, \alpha \not\models \phi$;
$w, \alpha \models \phi_1 \wedge \phi_2$	iff $w, \alpha \models \phi_i, i = 1, 2$;
$w, \alpha \models \langle G \rangle \phi$	iff $\alpha.G$ is a history and $w, \alpha.G \models \phi$;
$w, \alpha \models \langle \psi \rangle \phi$	iff $w, \alpha \models \psi$ and $w, \alpha.\psi \models \phi$;
$w, \alpha \models \hat{B}_a \phi$	iff $t, \beta \models \phi$ for some $(t, \beta) \in W \times \mathcal{H}$ such that $R_a w t, \alpha \triangleright_a \beta, t \bowtie \beta$.

Table 1: Semantics of PSAL

The semantics of belief warrants some discussion. Note that the relation \triangleright_a , for any of the interpretations, is not reflexive (it is however postreflexive, in the sense that $\alpha \triangleright_a \beta$ implies $\beta \triangleright_a \beta$). For this reason, it is not the case that $w, \alpha \models B_a \phi$ implies $w, \alpha \models \phi$. Our modality is not factual and this is the reason we favour a doxastic interpretation of it over than an epistemic one.

We make the assumption that an agent forms her beliefs based on announcements she has so far received, ignoring possible future announcements (indeed, note that $B_a[G] \perp$ is true whenever $a \in G$: an agent never believes there are unread announcements).

Let us see an example before moving on:

Example 6 Let us see why α_3 from Example 3 can never be executable under our second interpretation. Let (W, R, V) be a model and $w \in W$. Consider the prefix $\beta = p.\neg B_a p.ab.B_b B_a p.bc$. We will see that, if $w \bowtie \beta$, then $w, \beta \not\models \neg B_c p$, which entails $w \not\bowtie \alpha_3$. Indeed, suppose $\beta \triangleright_c \gamma$. We have that $\beta \upharpoonright_c = \downarrow_{\leq \beta}^* \text{last}_c \beta = \langle \{ab, bc\}, \leq_{\beta}^* \rangle$, and $\beta \upharpoonright_{!c} = p.\neg B_a p$. By our definition of \triangleright_a according to the second interpretation, γ can only be $p.\neg B_a p.ab.bc$ or $p.ab.\neg B_a p.bc$. Now, suppose $R_c w t$ and $t \bowtie \gamma$ (for either of these γ 's). In particular, this means the prefix $p \sqsubseteq \gamma$ is executable at t , i.e., $t, \epsilon \models p$, i.e., $t \in V(p)$. Thus $t, \gamma \models p$ for every pair (t, γ) with $R_c w t, \beta \triangleright_c \gamma, t \bowtie \gamma$, and therefore $w, \beta \models B_c p$.

4 The logic PSAL

We will say that a formula ϕ is ϵ -valid if, for every model (W, R, V) and every $w \in W$, it is the case that $w, \epsilon \models \phi$, and ϕ is $*$ -valid if, for every model (W, R, V) and $w \in W$, and for every history α such that $w \bowtie \alpha$, it is the case that $w, \alpha \models \phi$.

In the remainder of this section we will be concerned with ϵ -validities.

Now, let \mathcal{L}_{EL} be the language of the $\langle \phi \rangle$ and $\langle G \rangle$ -free fragment of the logic, i.e. $\phi ::= p \mid \top \mid \neg\phi \mid (\phi \wedge \phi) \mid \hat{B}_a \phi$.

Lemma 7 Given a model (W, R, V) and $w \in W$, for any formula $\phi \in \mathcal{L}_{EL}$ we have that $w, \epsilon \models \phi$ in the sense of PSAL if and only if $w \models \phi$ in the sense of the regular Kripke semantics.

In particular, $\phi \in \mathcal{L}_{EL}$ is ϵ -valid if and only if it is valid on Kripke models if and only if ϕ is a theorem of the minimal modal logic K.

Axioms of the logic of ϵ -validities. Let us give a sound and complete axiomatisation of the set of ϵ -valid formulas.

The following lemmas regarding histories will be useful. They can both be easily proven by induction on the length of α .

Lemma 8 If α is a history, w is a world in a model, and $w \bowtie \alpha$, then for every prefix $\beta \sqsubseteq \alpha$, we have that β is a history and $w \bowtie \beta$.

Lemma 9 For any model (W, R, V) and any pair $(w, \beta) \in W \times \mathcal{H}$ with $w \bowtie \beta$ we have: $w, \beta \models \langle \alpha \rangle \phi$ if and only if (i) the concatenation $\beta\alpha$ is a history, (ii) $w \bowtie \beta\alpha$, and (iii) $w, \beta\alpha \models \phi$.

The axioms and rules of the logic PSAL are displayed in Table 2.

- i. All the axioms and rules of the minimal modal logic K for each of the B_a modalities;
- ii. the following reduction axioms (where $\alpha \in (\mathcal{L}_{PSAL} \cup \mathcal{G})^*$):
 - $(R_{T1}) \quad \langle \alpha.G \rangle \top \leftrightarrow \langle \alpha \rangle \top$ if $\alpha.G$ is a history;
 - $(R_{T2}) \quad \langle \alpha.G \rangle \top \leftrightarrow \perp$ otherwise;
 - $(R_{T3}) \quad \langle \alpha.\phi \rangle \top \leftrightarrow \langle \alpha \rangle \phi$;
 - $(R_p) \quad \langle \alpha \rangle p \leftrightarrow (\langle \alpha \rangle \top \wedge p)$;
 - $(R_-) \quad \langle \alpha \rangle \neg \phi \leftrightarrow (\langle \alpha \rangle \top \wedge \neg \langle \alpha \rangle \phi)$;
 - $(R_\vee) \quad \langle \alpha \rangle (\phi \vee \psi) \leftrightarrow (\langle \alpha \rangle \phi \vee \langle \alpha \rangle \psi)$;
 - $(R_B) \quad \langle \alpha \rangle \hat{B}_a \phi \leftrightarrow (\langle \alpha \rangle \top \wedge \bigvee_{\alpha \triangleright_a \beta} \hat{B}_a \langle \beta \rangle \phi)$;
- iii. the Modus Ponens rule.

Table 2: The logic PSAL

As we remarked at the end of Section 2, given that the definition of \triangleright_a and \mathcal{H} differs for the interpretations, the big disjunct appearing in (R_B) will be different even for the same α . There is thus a slight abuse of notation in using the same acronym, PSAL, to refer to three different logics. We justify this by assuming one has fixed one's favourite interpretation.

Soundness. Validity of the rules of K is a routine check, and so is the fact that (R_p) , (R_-) and (R_\vee) are ϵ -valid. The (R_T) rules follow immediately from unpacking the semantics. Now for the other one:

Proposition 10 (R_B) is ϵ -valid.

Proof. Let (W, R, V) be a model. Suppose $w, \epsilon \models \langle \alpha \rangle \hat{B}_a \phi$. Then by Lemma 9 we have that $w \bowtie \alpha$ and $w, \alpha \models \hat{B}_a \phi$, which entails that $w, \epsilon \models \langle \alpha \rangle \top$ and that (by the semantics of B_a) there exist some $(t, \beta) \in W \times \mathcal{H}$ such that $R_a w t$, $\alpha \triangleright_a \beta$, $t \bowtie \beta$ and $t, \beta \models \phi$. But then (again by Lemma 9), we have that $t, \epsilon \models \langle \beta \rangle \phi$. Now, given the fact that $R_a w t$ plus the fact that $\epsilon \triangleright_a \gamma$ iff $\gamma = \epsilon$, the semantic definition gives us that $w, \epsilon \models \hat{B}_a \langle \beta \rangle \phi$ for some β such that $\alpha \triangleright_a \beta$ and therefore $w, \epsilon \models \bigvee_{\alpha \triangleright_a \beta} \hat{B}_a \langle \beta \rangle \phi$.

Conversely, if $w, \epsilon \models \langle \alpha \rangle \top \wedge \bigvee_{\alpha \triangleright_a \beta} \hat{B}_a \langle \beta \rangle \phi$, the first conjunct gives us that $w \bowtie \alpha$ and the second gives us that there is some t with $R_a w t$ such that $t, \epsilon \models \langle \beta \rangle \phi$ for some β with $\alpha \triangleright_a \beta$. Lemma 9 then gives us that $t \bowtie \beta$ and $t, \beta \models \phi$, for some (t, β) with $R_a w t$, $t \bowtie \beta$ and $\alpha \triangleright_a \beta$, which means that $w, \alpha \models \hat{B}_a \phi$ and therefore (again by Lemma 9) $w, \epsilon \models \langle \alpha \rangle \hat{B}_a \phi$. ■

Completeness. The previous logic is complete with respect to our models in any of the interpretations. The proof of this fact is virtually identical to the completeness proof in [2]. Let us briefly sketch this here:

- i. We show by \ll -induction that, for every formula $\phi \in \mathcal{L}_{PSAL}$ there exists an announcement-free formula ψ such that $\phi \leftrightarrow \psi$ is provable in PSAL. In particular, $\psi = s_A \phi$, where the translation s_A is defined by \ll -induction in Table 3.
- ii. Since $\phi \leftrightarrow s_A \phi$ is ϵ -valid and $s_A \phi \in \mathcal{L}_{EL}$, it follows that $w \models s_A \phi$ with the usual Kripke semantics if and only if $w, \epsilon \models \phi$. Thus ϕ is a theorem of PSAL if and only if $s_A \phi$ is a theorem of K. Completeness follows.

Therefore, we have:

$s_A \langle \epsilon \rangle \top$	$= \top$
$s_A \langle \epsilon \rangle p$	$= p$
$s_A \langle \epsilon \rangle (\psi_1 \wedge \psi_2)$	$= s_A \langle \epsilon \rangle \psi_1 \wedge s_A \langle \epsilon \rangle \psi_2$
$s_A \langle \epsilon \rangle \neg \psi$	$= \neg s_A \langle \epsilon \rangle \psi$
$s_A \langle \epsilon \rangle \hat{B}_a \psi$	$= \hat{B}_a s_A \langle \epsilon \rangle \psi$
$s_A \langle \alpha'.G \rangle \top$	$= \begin{cases} s_A \langle \alpha' \rangle \top, & \alpha'.G \text{ is a history} \\ \perp & \text{otherwise} \end{cases}$
$s_A \langle \alpha' \rangle \phi$	$= s_A \langle \alpha' \rangle \phi$
$s_A \langle \alpha \rangle p$	$= s_A \langle \alpha \rangle \top \wedge p$
$s_A \langle \alpha \rangle \neg \psi$	$= s_A \langle \alpha \rangle \top \wedge \neg s_A \langle \alpha \rangle \psi$
$s_A \langle \alpha \rangle (\psi_1 \wedge \psi_2)$	$= \bigwedge_{i=1,2} s_A \langle \alpha \rangle \psi_i$
$s_A \langle \alpha \rangle \hat{B}_a \psi$	$= s_A \langle \alpha \rangle \top \wedge \bigvee_{\alpha \triangleright_a \beta} \hat{B}_a s_A \langle \beta \rangle \psi$

Table 3: The map $s_A : \mathcal{L}_{PSAL} \rightarrow \mathcal{L}_{EL}$, defined by \ll -recursion on (α, ϕ) . In the last four rows of the table we assume that $\alpha \neq \epsilon$.

Theorem 11 PSAL is a sound and complete axiomatisation of the logic of ϵ -validities for Partially Synchronised Announcements.

Let us see in the following sections how the framework presented so far can be seen as a generalisation of both Asynchronous Announcement Logic and Public Announcement Logic.

5 PSAL generalises AAL

Let us delve with some more detail into the framework of Asynchronous Announcement Logic introduced in [2].

The language \mathcal{L}_{AA} is defined as

$$\phi ::= p \mid \top \mid \neg \phi \mid (\phi \wedge \phi) \mid \langle a \rangle \phi \mid \hat{B}_a \phi,$$

with $p \in \mathbf{P}$ and $a \in A$. A *history* in the context of AAL is a word $\alpha \in (\mathcal{L}_{AA} \cup A)^*$ with the property that, for all prefixes $\beta \sqsubseteq \alpha$, and for all $a \in A$, it holds that $|\beta|_a \leq |\beta|_!$, where $|\beta|_!$ is the number of announcements appearing in β and $|\beta|_a$ is the number of its a 's. (Informally this means: when an agent is making her n th reading, there must be at least n announcements which have been broadcast.)

$\alpha|_!$ refers to the projection of α to \mathcal{L}_{AA} (i.e., the formulas occurring in α) and $\alpha|_{!a}$ is the prefix of length $|\alpha|_a$ of $\alpha|_!$ (i.e., the formulas a reads). In this context, we define $\alpha \triangleright_a \beta$ iff $\alpha|_{!a} = \beta|_{!a} = \beta|_!$.

The semantics of AAL is very similar to that in Table 1, with the following crucial changes:

$w \bowtie \alpha.a$	iff $w \bowtie \alpha$;
$w, \alpha \models \langle a \rangle \phi$	iff $\alpha.a$ is a history (i.e. $ \alpha _a < \alpha _!$) and $w, \alpha.a \models \phi$;
$w, \alpha \models \hat{B}_a \phi$	iff $t, \beta \models \phi$ for some (t, β) s.t. $R_a w t$, $\alpha \triangleright_a \beta$, $t \bowtie \beta$.

Likewise the logic AAL is very similar to the logic PSAL as it appears in Table 2, with the following changes in the reduction axioms:

(R'_{T1})	$\langle \alpha.a \rangle \top \leftrightarrow \langle \alpha \rangle \top$ if $ \alpha _a < \alpha _!$;
(R'_{T2})	$\langle \alpha.a \rangle \top \leftrightarrow \perp$ otherwise;
(R'_B)	$\langle \alpha \rangle \hat{B}_a \phi \leftrightarrow (\langle \alpha \rangle \top \wedge \bigvee_{\alpha \triangleright_a \beta} \hat{B}_a \langle \beta \rangle \phi)$.

PSAL generalises Asynchronous Announcement Logic in a very straightforward way: we can say that a formula in the language of AAL is simply a formula in the language of PSAL in which all group readings are singletons. Or, looking at it in the other direction, we can claim that the fragment of PSAL in which all groups are singletons is precisely AAL.

We can see as well that PSAL generalises Public Announcement Logic. We expand on this in the next section.

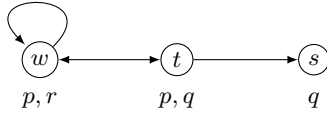
6 PSAL generalises PAL

Let us emulate Public Announcement Logic in our framework.

Let us consider as an example the following formula in the language of PAL: $\phi = \langle p \rangle (B_a p \wedge B_b p \wedge B_c p)$ (“after p is announced, all three agents believe (know) p ”). There is an implicit synchronicity in Public Announcement Logic: an announcement of p comes equipped with a simultaneous reading of this message by all agents (plus common knowledge that all agents have received the message). We can emulate this within PSAL by simply having the whole set of agents, A , read the announcement immediately after it is produced. It is not hard to see that the formula $\phi' = \langle p.A \rangle (B_a p \wedge B_b p \wedge B_c p)$ will be true at a pair (w, ϵ) if and only if ϕ is true at w in the sense of PAL.

This becomes slightly more complicated if we are dealing with formulas which have successive announcements. One might be tempted to translate a formula of the shape $\langle \phi \rangle \langle \psi \rangle \theta$ into $\langle \phi.A \rangle \langle \psi.A \rangle \theta$. This is not the right translation, as the following example illustrates:

Example 12 Consider the following one-agent model:



and let $\phi = \langle p \rangle \langle \hat{B}_a q \rangle B_a r$ and $\psi = \langle \langle p \rangle \hat{B}_a q \rangle B_a r$. Their respective translations are

$$\phi' = \langle p.a.\hat{B}_a q.a \rangle B_a r, \quad \psi' = \langle \langle p.a \rangle \hat{B}_a q.a \rangle B_a r.$$

However, while $\phi \leftrightarrow \psi$ is a theorem of PAL, we can see that $w, \epsilon \models \phi'$ whereas $w, \epsilon \not\models \psi'$.

Indeed, let $\alpha := p.a.\hat{B}_a q.a$ and $\beta := p.\hat{B}_a q.a.a$. Note that $t \bowtie \beta$, because $t, \epsilon \models p$ and $t, p \models \hat{B}_a q$, this last one on account that Rt s, $p \triangleright_a \epsilon$ and $s, \epsilon \models q$. However, $t, \beta \not\models r$. Thus there exist t, β with Rwt , $\alpha \triangleright_a \beta$, $t \bowtie \beta$ and $t, \beta \not\models r$ and therefore $w, \alpha \not\models B_a r$, which entails $w, \epsilon \not\models \langle p.a.\hat{B}_a q.a \rangle B_a r$.

On the other hand, let $\alpha' = \langle p.a \rangle \hat{B}_a q.a$. Note that, if $\alpha' \triangleright_a \beta'$, then necessarily $\beta' = \alpha'$, and note moreover that $t \not\bowtie \alpha'$, for $t, p.a \not\models \hat{B}_a q$ (given that the only successor of t executable in $p.a$ is w and $w, p.a \not\models q$). Therefore the only pair (x, γ) such that Rwx , $\alpha' \triangleright_a \gamma$ and $\gamma \bowtie \alpha'$ is (w, α') itself, and since $w, \alpha' \models r$, we get $w, \alpha' \models B_a r$ and thus $w, \epsilon \models \langle \langle p.a \rangle \hat{B}_a q.a \rangle B_a r$.

In this example, ψ' (instead of ϕ') seems like the right translation of ϕ . The translation we need is a bit more complicated and it is defined below:

Definition 13 Let $\tau : \mathcal{L}_{PAL} \rightarrow \mathcal{L}_{PSAL}$ be defined, by recursion on the length of ϕ , as follows:

$$\begin{array}{ll} \tau \top &= \top; & \tau \langle \phi \rangle p &= \langle \tau \phi.A \rangle p; \\ \tau p &= p; & \tau \langle \phi \rangle (\psi \wedge \chi) &= \tau \langle \phi \rangle \psi \wedge \tau \langle \phi \rangle \chi; \\ \tau \neg \phi &= \neg \tau \phi; & \tau \langle \phi \rangle \neg \psi &= \langle \tau \phi.A \rangle \neg \tau \psi; \\ \tau \langle \phi \wedge \psi \rangle &= \tau \phi \wedge \tau \psi; & \tau \langle \phi \rangle \hat{B}_a \psi &= \langle \tau \phi.A \rangle \hat{B}_a \tau \psi; \\ \tau \hat{B}_a \phi &= \hat{B}_a \tau \phi; & \tau \langle \phi \rangle \langle \psi \rangle \chi &= \langle \tau \langle \phi \rangle \psi.A \rangle \tau \chi. \end{array}$$

The following holds for all three interpretations.

Theorem 14 For every model (W, R, V) , for all $w \in W$ and all $\phi \in \mathcal{L}_{PAL}$, we have that $w \models \phi$ in the sense of PAL if and only if $w, \epsilon \models \tau \phi$ in the sense of PSAL.

In order to prove this, let us first recover the translation we “ruled out” in the previous example, namely $t : \mathcal{L}_{PAL} \rightarrow \mathcal{L}_{PSAL}$, defined as: $t \top = \top$, $t p = p$, $t(\neg \phi) = \neg t \phi$, $t(\phi \wedge \psi) = t \phi \wedge t \psi$, $t(B_a \phi) = B_a t \phi$, $t(\langle \phi \rangle \psi) = \langle t \phi.A \rangle t \psi$. We have:

Lemma 15 If ϕ is an announcement-free formula in the language of PAL, then it holds that $\phi = t \phi$ and $w \models_{PAL} \phi$ iff $w, \epsilon \models_{PSAL} \phi$.

Proof. By induction on announcement-free ϕ . It is trivial for $\phi = \top$ and $\phi = p$, and the induction steps for disjunction and negation are straightforward. If $\phi = \hat{B}_a \psi$ for some announcement-free ψ satisfying the induction hypothesis we have that $w \models_{PAL} \hat{B}_a \psi$ if and only if $t \models_{PAL} \psi$ for some t such that $R_a w t$ if and only if (by induction hypothesis) $t, \epsilon \models_{PSAL} \psi$ for some t with $R_a w t$ if and only if (given that $\epsilon \triangleright_a \beta$ implies $\beta = \epsilon$) $w, \epsilon \models_{PSAL} \hat{B}_a \psi$. ■

For the following result, we will say that a formula $\phi \in \mathcal{L}_{PAL}$ is in *standard form* whenever it does not contain any subformulas of the form $\langle \psi_1 \rangle \langle \psi_2 \rangle \chi$. Every formula in PAL is equivalent to a formula in standard form, which we can obtain by using recursively the equivalence $\langle \psi_1 \rangle \langle \psi_2 \rangle \phi \leftrightarrow \langle \langle \psi_1 \rangle \psi_2 \rangle \phi$.

We will moreover use the fact that every formula in the language of PSAL or PAL is equivalent to some announcement-free formula via a translation that we can define by \ll -recursion:

Lemma 16 i. The map $s_A : \mathcal{L}_{PSAL} \rightarrow \mathcal{L}_{EL}$ defined in Table 3 satisfies: for every formula ϕ and every model (W, R, V) we have that $s_A \phi$ is announcement-free and $w, \epsilon \models \phi$ iff $w, \epsilon \models s_A \phi$.
ii. There exists a translation map $s_P : \mathcal{L}_{PAL} \rightarrow \mathcal{L}_{EL}$ such that for every formula ϕ in standard form and every model (W, R_a, V) we have that $s_P \phi$ is announcement-free and $w \models \phi$ iff $w \models s_P \phi$. This map is defined by \ll -recursion⁴ in Table 4.

$s_P \top$	$= \top$
$s_P p$	$= p$
$s_P (\psi_1 \wedge \psi_2)$	$= s_P \psi_1 \wedge s_P \psi_2$
$s_P \neg \psi$	$= \neg s_P \psi$
$s_P \hat{B}_a \psi$	$= \hat{B}_a s_P \psi$
$s_P (\langle \phi \rangle \top)$	$= s_P \phi$
$s_P (\langle \phi \rangle p)$	$= s_P \phi \wedge p$
$s_P (\langle \phi \rangle \neg \psi)$	$= s_P \phi \wedge \neg s_P (\langle \phi \rangle \psi)$
$s_P (\langle \phi \rangle (\psi_1 \wedge \psi_2))$	$= \bigwedge_{i=1,2} s_P (\langle \phi \rangle \psi_i)$
$s_P (\langle \phi \rangle \hat{B}_a \psi)$	$= s_P \phi \wedge \hat{B}_a s_P (\langle \phi \rangle \psi)$

Table 4: The map s_P defined by \ll -recursion.

With this:

Lemma 17 For any formula ϕ in the language of PAL in standard form, $s_P \phi = s_A t \phi$.

Proof. By \ll -induction. Trivial for the cases in which $\phi = \top, p, \psi_1 \wedge \psi_2, \hat{B}_a \psi$.

Now suppose $\phi = \langle \chi \rangle \top$. Then $s_A t \phi = s_A \langle t \chi.A \rangle \top = s_A \langle t \chi \rangle \top = s_A t \chi = s_P \chi = s_P \langle \chi \rangle \top$, the fourth equality given by the induction hypothesis and the rest directly by Table 4.

If $\phi = \langle \chi \rangle p$, then $s_A t \phi = s_A \langle t \chi.A \rangle p = s_A \langle t \chi.A \rangle \top \wedge p$. The first conjunct of this last formula equals $s_P \chi$ by the previous step, therefore $s_A t \langle \chi \rangle p = s_P \chi \wedge p = s_P \langle \chi \rangle p$.

If $\phi = \langle \chi \rangle \neg \psi$ we have that $s_A t \phi = s_A \langle t \chi.A \rangle \top \wedge \neg s_A \langle t \chi.A \rangle \psi$. The first conjunct is again equal to $s_P \chi$ (as before) and the second equals $\neg s_A t \langle \chi \rangle \psi$ which, by induction hypothesis, equals $s_P \langle \chi \rangle \psi$. Therefore $s_A t \phi = s_P \chi \wedge \neg s_P \langle \chi \rangle \psi = s_P \phi$.

If $\phi = \langle \chi \rangle \hat{B}_a \psi$, then $s_A t \phi = s_A \langle t \chi.A \rangle \hat{B}_a t \psi$. Now note that if $t \chi.A \triangleright_a \beta$, then $\beta = t \chi.A$ and therefore $s_A t \phi = s_A \langle t \chi.A \rangle \top \wedge$

⁴ Slight abuse of notation: for formulas in PAL, $\phi \ll \psi$ iff $(\deg \phi < \deg \psi)$ or $(\deg \phi = \deg \psi \text{ and } \|\phi\| < \|\psi\|)$, where \deg and $\|\cdot\|$ are as in Def. 4.

$\hat{B}_a s_A \langle t \chi . A \rangle t \psi$. By induction hypothesis the first conjunct equals $s_P \chi$ and the second one equals $\hat{B}_a s_P (\langle \chi \rangle \psi)$, thus their conjunction equals $s_P \phi$. ■

And now:

Corollary 18 *If $\phi \in \mathcal{L}_{PAL}$ is in standard form, then $w \models_{PAL} \phi$ iff $w, \epsilon \models_{PSAL} t \phi$.*

Proof. $w \models_{PAL} \phi$ iff (by Lemma 16) $w \models_{PAL} s_P \phi$ iff (by Lemma 15 and the fact that $s_P \phi$ is announcement-free) $w, \epsilon \models_{PSAL} s_P \phi$ iff (by the previous Lemma) $w, \epsilon \models s_A t \phi$ iff $w, \epsilon \models t \phi$. ■

Now, let sf be the translation of formulas in PAL to their standard forms, by applying recursively the equivalences $\langle \psi_1 \rangle \langle \psi_2 \rangle \phi \leftrightarrow \langle \langle \psi_1 \rangle \psi_2 \rangle \phi$ and $\langle T \rangle \phi \leftrightarrow \phi$. More explicitly, $sf \top = \top$, $sf p = p$, $sf \neg \phi = \neg sf \phi$, $sf \hat{B}_a \phi = \hat{B}_a sf \phi$, $sf \langle \phi \rangle \psi = \langle sf \phi \rangle sf \psi$ (if ψ is not of the form $\langle \chi_1 \rangle \chi_2$), $sf \langle \phi \rangle \langle \psi \rangle \chi = \langle \langle sf \phi \rangle sf \psi \rangle sf \chi$.

Note that $sf \phi$ is always in standard form and that $\models_{PAL} \phi \leftrightarrow sf \phi$. Therefore:

Proof of Thm. 14 Simply note that $\tau = t sf$. We have: $w \models \phi$ iff $w \models sf \phi$ (by the previous corollary) $w, \epsilon \models t sf \phi$. ■

7 Common belief

As we mentioned above, a notion of common belief (or knowledge) does not make much sense in the setting of AAL, where messages are received individually by the agents and thus an agent can never be certain that others have received the same messages she has. Going back to our second example from the introduction, here we have that Antía and Brais have watched the second installment of the documentary together; therefore one would expect that not only do they both believe/know p_2 , but they both believe the other to believe it, and they believe the other believes they themselves believe it, etc.

Let $G \in \mathcal{P}(A) \setminus \{\emptyset\}$ be a group of agents. Let us propose a notion of common belief of ϕ by all the agents in G , $CB_G \phi$.

Given worlds $w, t \in W$, histories α and β , and an agent $a \in A$, let $(w, \alpha) \mathbf{R}_a (t, \beta)$ iff $R_a w t$, $\alpha \triangleright_a \beta$ and $t \triangleright \beta$. We can read $CB_G \phi$ in terms of the transitive closure of the union of these \mathbf{R}_a 's, $\mathbf{R}_G = (\bigcup_{a \in G} \mathbf{R}_a)^T$, so that $w, \alpha \models CB_G \phi$ iff $(w, \alpha) \mathbf{R}_G (t, \beta)$ implies $(t, \beta) \models \phi$. Equivalently,

Definition 19 $w, \alpha \models CB_G \phi$ if and only if for all $a_1, \dots, a_n \in G$ and for all chains

$$(w, \alpha) \xrightarrow{\mathbf{R}_{a_1}} (t_1, \beta_1) \xrightarrow{\mathbf{R}_{a_2}} \dots \xrightarrow{\mathbf{R}_{a_n}} (t_n, \beta_n),$$

it is the case that $t_n, \beta_n \models \phi$.

Two *-validities of Common Belief. We define the abbreviation “every agent in G believes that ϕ ” as $E_G \phi = \bigwedge_{a \in G} B_a \phi$. We have:

Theorem 20 *The following two principles are *-valid in all three interpretations:*

- (Fix) $CB_G \phi \rightarrow E_G (\phi \wedge CB_G \phi)$;
- (Ind) $CB_G (\phi \rightarrow E_G \phi) \rightarrow (E_G \phi \rightarrow CB_G \phi)$.

Proof. (Fix). Suppose $w, \alpha \models CB_G \phi$, take $a \in G$ and consider (t, β) such that $(w, \alpha) \mathbf{R}_a (t, \beta)$. Then $(w, \alpha) \mathbf{R}_G (t, \beta)$ (and thus $t, \beta \models \phi$) and, if $(t, \beta) \mathbf{R}_G (s, \gamma)$ we have that $(w, \alpha) \mathbf{R}_G (s, \gamma)$ and thus $s, \gamma \models \phi$, which entails $t, \beta \models CB_G \phi$. Therefore, for every $a \in G$ it holds that $w, \alpha \models B_a (\phi \wedge CB_G \phi)$ and thus $w, \alpha \models E_G (\phi \wedge CB_G \phi)$.

(Ind). Suppose $w, \alpha \models CB_G (\phi \rightarrow E_G \phi) \wedge E_G \phi$ and consider a chain

$$(w, \alpha) = (t_0, \beta_0) \xrightarrow{\mathbf{R}_{a_1}} (t_1, \beta_1) \xrightarrow{\mathbf{R}_{a_2}} \dots \xrightarrow{\mathbf{R}_{a_n}} (t_n, \beta_n),$$

with $a_1, \dots, a_n \in G$. Note that $n > 0$ and it is easy to prove by induction on n that every element in the chain satisfies $t_k, \beta_k \models E_G \phi$. In particular, $t_{n-1}, \beta_{n-1} \models B_{a_n} \phi$ and thus $(t_n, \beta_n) \models \phi$. ■

Let us now finish with an example.

Example 21 Let $\alpha_2 = p. \neg B_a p. ab. B_b B_a p. bc. B_c p. ca$, as in Example 2, and let $W = \{w, t\}$, $R_a = R_b = R_c = W^2$, $V(p) = \{w\}$.

Now, we can easily see that, after an execution of α_2 at w , both $\{b, c\}$ and $\{a, c\}$ have common knowledge of the fact that $B_c p$. Indeed, for the former, we consider any chain $\alpha_2 \triangleright_{x_1} \beta_1 \dots \triangleright_{x_n} \beta_n$ and it is straightforward that any element of this chain will contain at least the two first announcements, p and $\neg B_a p$, and the readings ab and bc . Therefore, $t \not\models \beta_i$ for $i = 1, \dots, n$ and we have to evaluate $B_c p$ on (w, β_n) . But again, if $\beta_n \triangleright_c \gamma$, then γ will contain (at least) the two first announcements and the two first readings, thus will only be executable in w . And since $w, \gamma \models p$ for any such γ , we have that $(w, \beta_n) \models B_c p$ for any such chain, and thus $w, \alpha_2 \models CB_{bc} B_c p$. (We reason similarly to see that $w, \alpha_2 \models CB_{ac} B_c p$.)

Let us see, however, that this fact is not common knowledge between a and b : let $\beta_1 = p. \neg B_a p. ab. bc$, $\beta_2 = p. ab$. Note that $\alpha_2 \triangleright_b \beta_1 \triangleright_a \beta_2$ and note that all these histories are executable on w . However, we have that $w, \beta_2 \not\models B_c p$, for we have that $R_c w t$, $\beta_2 \triangleright_c \epsilon$, $t \models \epsilon$ and $t, \epsilon \not\models p$. Therefore, $w, \alpha_2 \not\models CB_{ab} B_c p$.

8 Conclusion

We have introduced Partially Synchronised Announcement Logic, a framework which allows us to model communicative situations involving truthful announcements which are publicly sent yet received by different groups of agents at different times; we have given three intuitive interpretations of the ‘view’ of an agent, and provided the sound and complete logic of each of them (for the class of models with empty histories). We have as well given a proposal for common belief in this framework.

Our framework for partial synchronization may be of interest to model various multi-agent systems and protocols wherein agents or groups of agents send and receive messages, as in distributed computing. The typical way to simulate, in a dynamic epistemic logic, that an agent a sends a message (with content) ϕ is as an announcement of $B_a \phi$ by the environment. For such applications the belief operator functions as an acknowledgement of receipt; for example: a sends p to b (announcement $B_a p$) and, after b eventually receives this announcement, b acknowledges this by sending $B_a p$ to a (announcement $B_b B_a p$). In this way, we can model as diverse systems as: the internet protocol TCP guaranteeing correctness of initial sequences of packages [8, 20] (an example of individual reception), gossip protocols wherein agents inform each other in peer-to-peer telephone calls in the setting with rounds of calls [10, 1] (an example of full synchronization for all agents after simultaneous partial synchronization for subsets of size two – namely the two agents involved in a call), and immediate snapshots in distributed computing, involving schedules consisting of concurrency classes (an example of a partition of a set of agents into subsets of arbitrary size, namely those agents involved in joint read/write actions) [9, 6].

Some interesting research directions are yet to be explored. For example, studying the logic of *-validities of PSAL, as opposed of that of ϵ -validities, seems to be a very relevant way to move forward.

In further research we wish to find a suitable semantics for asynchronous knowledge. Unlike belief, knowledge should be correct: *what you know is true*. A knowledge semantics requires a view relation that is an equivalence relation, instead of our \triangleright_a relation for belief, which is not reflexive. It should be noted that even for the asynchronous belief semantics some beliefs are correct: assuming models where all relations are equivalences, if the believed formula is in the “positive fragment” of the language (no negations before epistemic modalities), it is correct (and could be said to constitute knowledge), as reported in [2].

In [2] it is proven that the model-checking complexity of AAL is in PSPACE. We conjecture that an analogous result will be true of PSAL.

But perhaps the most interesting direction for future work is finding a sound and complete axiomatisation for PSAL with common belief.

Acknowledgements. The authors are very thankful to the reviewers for their helpful comments and suggestions. Their careful reviews have contributed to increasing the quality and readability of this paper.

Hans van Ditmarsch is also affiliated as research associate to IMSc, Chennai, India.

REFERENCES

- [1] M. Attamah, H. van Ditmarsch, D. Grossi, and W. van der Hoek, ‘Knowledge and gossip’, in *Proc. of 21st ECAI*, pp. 21–26. IOS Press, (2014).
- [2] P. Balbiani, H. van Ditmarsch, and S. Fernández González, ‘Asynchronous announcements’, *CoRR*, **abs/1705.03392**, (2019).
- [3] A. Baltag, L.S. Moss, and S. Solecki, ‘The logic of public announcements, common knowledge, and private suspicions’, in *Proc. of 7th TARK*, pp. 43–56. Morgan Kaufmann, (1998).
- [4] I. Ben-Zvi and Y. Moses, ‘Beyond lamport’s *Happened-before*: On time bounds and the ordering of events in distributed systems’, *Journal of the ACM*, **61**(2), 13:1–13:26, (2014).
- [5] K.M. Chandy and J. Misra, ‘How processes learn’, in *Proc. of the 4th PODC*, pp. 204–214, (1985).
- [6] E. Goubault, J. Ledent, and S. Rajsbaum, ‘A simplicial complex model for dynamic epistemic logic to study distributed task computability’, in *Proc. of 9th GandALF*, volume 277 of *EPTCS*, pp. 73–87, (2018).
- [7] J.Y. Halpern and Y. Moses, ‘Knowledge and common knowledge in a distributed environment’, *Journal of the ACM*, **37**(3), 549–587, (1990).
- [8] J.Y. Halpern and L.D. Zuck, ‘A little knowledge goes a long way: Knowledge-based derivations and correctness proofs for a family of protocols’, *J. ACM*, **39**(3), 449–478, (1992).
- [9] M. Herlihy, D. Kozlov, and S. Rajsbaum, *Distributed Computing Through Combinatorial Topology*, Morgan Kaufmann, San Francisco, CA, USA, 2013.
- [10] A.-M. Kermarrec and M. van Steen, ‘Gossiping in distributed systems’, *SIGOPS Oper. Syst. Rev.*, **41**(5), 2–7, (2007).
- [11] S. Knight, B. Maubert, and F. Schwarzentruher, ‘Reasoning about knowledge and messages in asynchronous multi-agent systems’, *Mathematical Structures in Computer Science*, **29**(1), 127–168, (2019).
- [12] A.D. Kshemkalyani and M. Singhal, *Distributed Computing: Principles, Algorithms, and Systems*, Cambridge University Press, New York, NY, USA, 2008.
- [13] L. Lamport, ‘Time, clocks, and the ordering of events in a distributed system’, *Communications of the ACM*, **21**(7), 558–565, (1978).
- [14] Y.O. Moses, D. Dolev, and J.Y. Halpern, ‘Cheating husbands and other stories: a case study in knowledge, action, and communication’, *Distributed Computing*, **1**(3), 167–176, (1986).
- [15] M. Mukund and M.A. Sohoni, ‘Keeping track of the latest gossip in a distributed system’, *Distributed Computing*, **10**(3), 137–148, (1997).
- [16] P. Panangaden and K. Taylor, ‘Concurrent common knowledge: Defining agreement for asynchronous systems’, *Distributed Computing*, **6**, 73–93, (1992).
- [17] R. Parikh and R. Ramanujam, ‘A knowledge based semantics of messages’, *Journal of Logic, Language and Information*, **12**, 453–467, (2003).
- [18] J.A. Plaza, ‘Logics of public communications’, in *Proc. of the 4th IS-MIS*, pp. 201–216. Oak Ridge National Laboratory, (1989).
- [19] R. Ramanujam, ‘View-based explicit knowledge’, *Ann. Pure Appl. Logic*, **96**(1-3), (1999).
- [20] F. Stulp and R. Verbrugge, ‘A knowledge-based algorithm for the internet transmission control protocol (TCP)’, *Bulletin of Economic Research*, **54**(1), 69–94, (2002).
- [21] J. van Benthem, J.D. Gerbrandy, T. Hoshi, and E. Pacuit, ‘Merging frameworks for interaction’, *Journal of Philosophical Logic*, **38**, 491–526, (2009).
- [22] H. van Ditmarsch, W. van der Hoek, and B. Kooi, *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*, Springer, 2008.